Network Working Group J. Abley
Internet-Draft Metromedia Fiber Network Inc.
Expires: December 24, 2001 B. Black
Layer8 Networks
V. Gill
Metromedia Fiber Network Inc.
June 25, 2001

IPv4 Multihoming Motivation, Practices and Limitations
draft-ietf-multi6-v4-multihoming-00

Abstract

Multihoming is an essential component of service for enterprises [3] which are part of the Internet. This draft describes some of the motivations, practices and limitations of multihoming as it is achieved in the IPv4 world today.

The context for this discussion is the requirements analysis for site multihoming in IPv6, which is described in a companion draft [1].

1. Introduction

Multihoming is an essential component of service for enterprises which are part of the Internet. Current IPv4 multihoming practices have been added on to the CIDR architecture [2], which assumes that routing table entries can be aggregated based upon a hierarchy of customers and service providers.

Multihoming is a mechanism by which enterprises can currently satisfy a number of high-level requirements, and is widely used in the IPv4 network today. There are some practical limitations, however, including concerns of how well (or, if) the current practice will scale as the network continues to grow.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [4].

An "enterprise" is an entity autonomously operating a network using TCP/IP and, in particular, determining the addressing plan and address assignments within that network. This is the definition of "enterprise" used in [3].

A "transit provider" is an enterprise which provides connectivity to the Internet to one or more other enterprises. The connectivity provided extends beyond the transit provider's own network.

A "multi-homed" enterprise is one with more than one transit provider. "Multihoming" is the practice of being multi-homed.

A "multi-attached" enterprise is one with more than one point of layer-3 interconnection to a single transit provider.

The term "re-homing" denotes a transition of an enterprise between two states of connectedness, due to a change in the connectivity between the enterprise and its transit providers.

3. Motivations for Multihoming

3.1 Redundancy

By multihoming, an enterprise can insulate itself from certain failure modes within one or more transit providers, as well as failures in the network providing interconnection with one or more transit providers.

Examples of failure modes from which an enterprise can obtain some degree of protection by multi-homing are:

- o Physical link failure, such as a fiber cut or router failure,
- o Logical link failure, such as a misbehaving router interface,
- o Routing protocol failure, such as a BGP peer reset,
- o Transit provider failure, such as a backbone-wide IGP failure, and
- o Exchange failure, such as a BGP reset on an inter-provider peering.

Some of these failure modes may be protected against by multi-attaching to a single transit provider, rather than multi-homing.

3.2 Load Sharing

By multihoming, an enterprise can distribute both inbound and outbound traffic between multiple transit providers.

Sometimes it is not possible to increase transit capacity to a single transit provider because that provider does not have sufficient spare capacity to sell. In this case a solution is to acquire additional transit capacity through a different provider. This scenario is common in bandwidth-starved stubs of the Internet where, for example, transit demand outpaces under-sea cable deployment.

3.3 Performance

By multihoming, an enterprise can protect itself from performance difficulties between transit providers.

For example, suppose enterprise E obtains transit from transit providers T1 and T2, and there is long-term congestion between T1 and T2. By multihoming between T1 and T2, E is able to ensure that in normal operation none of its traffic is carried over the congested interconnection T1-T2.

3.4 Policy

An enterprise may choose to load-share for a variety of policy reasons outside technical scope (e.g. cost, acceptable use conditions, etc).

For example, enterprise E homed to transit provider T1 may be able to identify a particular range of addresses within its network that correspond to non-real-time traffic (e.g. a network containing mail and Usenet/NNTP servers). It may be advantageous to shift inbound traffic destined for that range of addresses to transit-provider T2, since T2 charges less for traffic.

4. Features of IPv4 Multihoming

4.1 Simplicity

The current multihoming solution is not without complexity, but in practice it quite straightforward to deploy and maintain by virtue of the fact that it is well-known, tried and tested.

4.2 Transport-Layer Survivability

The current multihoming solution provides session survivability for transport-layer protocols; i.e. exchange of data between devices on the multi-homed enterprise network and devices elsewhere on the Internet may proceed with no greater interruption than that associated with the transient packet loss during a re-homing event.

New transport-layer sessions are able to be created following a re-homing event.

4.3 Inter-Provider Traffic Engineering

A multi-homed enterprise may influence routing decisions beyond its immediate transit providers by advertising a strategic mixture of carefully-aimed long prefixes and covering shorter-prefix routes. This precise effects of such egress policy are often difficult to predict, but an approximation of the desired objective is often easy to accomplish. This can provide a similar mechanisms to that described in Section 3.3, except that the networks whose traffic is being influenced are not transit providers of the enterprise itself.

4.4 Load Control

The current multihoming solution places control of traffic flow in the hands of the organisation responsible for the multi-homed interconnections with transit providers. A single-homed customer of a multi-homed enterprise may vary the demand for traffic that they impose on the enterprise, and may influence differential traffic load between transit providers; however, the basic mechanisms for congestion control and route propagation are in the hands of the enterprise, not the customer.

4.5 Impact on Routers

The routers at the boundary of a multi-homed enterprise are usually required to participate in BGP sessions with the interconnected routers of transit providers. Other routers within the enterprise have no special requirements beyond those of single-homed enterprises' routers.

4.6 Impact on Hosts

There are no requirements of hosts beyond those of single-homed enterprises' hosts.

4.7 Interactions between Hosts and the Routing System

There are no requirements for interaction between routers and hosts beyond those of single-homed enterprises' routers and hosts.

5. Limitations of IPv4 Multihoming

5.1 Scalability

Current IPV4 multihoming practices contribute to the significant growth currently observed in the state held in the global inter-provider routing system; this is a concern both because of the hardware requirements it imposes and also because of the impact on the stability of the routing system. This issue is discussed in great detail in [5].

6. Security Considerations

Security considerations are not discussed in this draft.

References

[1] Black, B., Gill, V. and J. Abley, "Requirements for IP Multihoming Architectures (work-in-progress)", I-D draft-ietf-multi6-v4-multihoming-01, June 2001, <http://www.automagic.org/~jabley/draft-ietf-multi6-multihoming-requirements-01.txt>.

[2] Fuller, V., Li, T., Yu, J. and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.

[3] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. and E. Lear, "Address Allocation for Private Internets", RFC 1918, February 1996.

[4] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[5] Huston, G., "Analyzing the Internet's BGP Routing Table", January 2001.

Authors' Addresses

Joe Abley
Metromedia Fiber Network Inc.
2204 Pembroke Court
Burlington, ON L7P 3X8
Canada

Phone: +1 905 319 9064
EMail: jabley@mfnx.net

Benjamin Black
Layer8 Networks

EMail: ben@layer8.net

Vijay Gill
Metromedia Fiber Network Inc.
8075 Leesburg Pike
Suite 300
Vienna, VA 22182
US

Phone: +1 410 262 0660
EMail: vgill@mfnx.net

Full Copyright Statement

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.