Internet Engineering Task Force
Internet-Draft
Expires: January 10, 2005

J. Palet
M. Diaz
C. Olvera
A. Vives
Consulintel
E. Fleischman
Boeing
D. Lanciani
July 12, 2004

Analysis of IPv6 Multihoming Scenarios
draft-palet-multi6-scenarios-00.txt

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of section 3 of RFC 3667. By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with RFC 3668.

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.

This Internet-Draft will expire on January 10, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

Multihoming seems to be one of the key pieces for the deployment of IPv6 in the enterprise scenario, but is becoming a frequent

requirement in all kinds of networks. In addition, other factors including the deployment of broadband networks, the increase in the variety of access technologies, the increase in the demand for resilience/redundancy, etc., in non-enterprise environments, for example in SOHO/home, necessarily imply the increase of IPv6 multihoming demand for a number of scenarios, which are described in this memo.

Table of Contents

1. Introduction

The term multihoming refers to the practice of having a network connected to more than one ISP (connectivity provider, transit provider, upstream provider, etc.).

A device can also be multihomed (e.g. host-centric multihoming), when it has more than one interface, and each of the interfaces is attached to different networks (may be within a multihomed network).

In addition to that, in IPv6, each interface can have multiple addresses, which also means than even with a single interface, a host can be multihomed.

Multihoming provides certain degree of resilience/redundancy against failures (link, hardware, protocols, others) and also enables features such as load balancing. Moreover, multihoming can be used in order to differentiate traffic based on policy, for non-technical reasons, such as cost associated with different flows, time of the day, etc. For highly distributed enterprises, it can also occur as an aid to address that enterprise's geographical distribution, and as a traffic engineering mechanism to improve local performance such as latency and hop count reductions for real time protocols.

The scope of this document is to provide a brief description of the motivations for multihoming and an analysis of different scenarios where IPv6 multihoming could be relevant.

2. Multihoming Motivations

Considering the goals for requiring multihoming described in [1], as well as other frequent issues, we can classify the motivations as both technical and non-technical.

2.1 Technical motivations

The technical motivations are intended to provide resilience for the multihomed networks.

1. Redundancy: In order to protect the network against failures. These failures could be either in the multihomed network itself, or in the upstream (e.g. provider) network. The failures can be of very different types, including and not limited to: Link failures (physical, logical, configuration, ...), hardware failures (interfaces, routers, other equipment), protocol failures (routing protocol, others).

2. Load sharing/balancing: In order to distribute the inbound and/or

outbound traffic (e.g. between multiple providers).

3. Performance: As a traffic engineering mechanism to dampen the affects of upstream hub congestion through the use of "complementary" providers (e.g., in case the local Internet Exchanges are congested), avoiding for example, the traffic being forwarded thru longer routes than logically required.

Some of these technical motivations could be partially solved by multi-attaching the network to the same ISP (i.e. load balancing), either in the same point-of-presence (POP) or via geographically separated POPs. Resilience can also be increased by using different routers, instead of just different interfaces. But in any case, load balancing alone doesn't provide a complete protection from failures in the upstream provider. For this reason, many sites may choose to attach to multiple ISPs (i.e. multihoming) so that a failure within the ISP will not necessarily isolate that site from the Internet as a whole.

2.2 Non-technical motivations

A network can require multihoming for non-technical reasons. Financial considerations may influence the deployment, such as different cost or fees associated with with different traffic flows (with different priorities), or differing costs at different times of the day use, destination networks, accumulated bandwidth used, etc.

Political motivations may exist such as the desire to be able to quickly change the provider, without renumbering. This motivation could stem from a wide variety of considerations including service charges, improved SLA, ISP bankruptcy, etc.

Another non-technical motivation, which occurs very frequently in some scenarios (e.g. research and educations networks), is acceptable usage policy, where commodity traffic is not accepted, in general.

3. Multihoming Scenarios

Different networking scenarios provide different multihoming cases or scenarios, which may have different peculiarities and requirements. They are listed below with no specific order or priority.

3.1 Internet Service Provider (ISP)

An ISP is naturally multihomed when connected to two or more upstream providers. Actually this is a very common case, especially for medium and large ISPs.

In this scenario, the ISP will usually have its own address pool, in a provider independent fashion, allocated from a Regional Internet Registry (RIR). A multihomed ISP uses routing protocols to advertise its address space to several upstream providers.

Some small ISPs only use one upstream provider, so in this case, multihoming will not apply.

The National Research and Education Networks (NRENs) can be considered, from the multihoming perspective, as ISPs, because they are usually connected to several upstream providers, but they also have their own addressing space (allocated from a RIR).

3.2 IX

Can an IX be multihomed, for example if they are layer 3 transit providers, distributed IXs ???.

TBD.

3.3 Enterprise

The size of any given enterprise network is a function of the size of the corporation it supports. Enterprise networks can therefore range from being small to being enormous. Larger enterprises may link together multiple buildings within a campus, multiple campuses within a region, multiple regions within a country, as well as have sites of various sizes within multiple countries. A few corporations even have corporate sites located within every country of the world. These sites may themselves be linked together via public (e.g., ISP) or private (e.g., dark fiber, satellite communications, modem connections over telephone networks) means. A distinction therefore needs to be made between a large corporation's private use of public (ISP) networking facilities to privately link together disparate parts of that corporation and that same corporation's public POPs to the worldwide Internet. The former is solely known (or, at least, it should be solely known) to the corporation and the supplier, the latter is advertised to the worldwide Internet infrastructure as the standard mechanism by which that corporation can be accessed. This distinction is important because the security mechanisms protecting these different uses (e.g., firewalls) are often very distinct from each other.

Medium and large corporations are frequently attached to several ISPs, often for multihoming and load balancing reasons. Highly distributed corporations often have additional motivations for connectivity to multiple ISPs stemming from traffic engineering and performance considerations, particularly if real time traffic (e.g.

VoIP) is being internally supported.

This kind of networks usually require a high resilience, because any outage, even when minimal, could mean a very high cost, that could justify the adoption of a multihoming solution even if this is expensive.

Medium and small enterprise networks could fall in the category of Enterprise or in the SOHO one, which is very dependent on the specific resilience requirements.

3.4 University/Campus

In principle, it seems that a University or Campus network fits in the same category as the enterprise network, except that its policies and security systems are very different, particularly because the university will not go out of business if the information contained within its computers becomes compromised.

Is also important to recognize one more special non-technical requirement. Often they are connected in addition to commercial ISPs to non-commercial ISPs (NRENs), which don't allow commodity traffic, and consequently the network must control traffic based on policy.

3.5 Security, Defense and Emergency

Military networks differ from university and corporate networks in that their computers and the networks themselves operate at specific classification levels. In military jargon, these are known as "Red networks". Each red network instance operates at a specific classification level (e.g., secret, sensitive-but-unclassified, etc.). Red networks operating at different classification levels may have disjoint (i.e., unrelated) address spaces from each other. In some military environments, Red networks are conveyed over physical media that can be protected. This is in contrast to Black networks, whose media cannot be similarly protected (e.g., wireless transmissions). Packets from Red networks may be conveyed over Black networks if the Red packets are first encrypted. In some cases, the encrypted Red packet may be encapsulated within an appropriate IP header of the Black network.

But from the multihoming perspective, in principle, it seems that security and military networks fit into the same category as the enterprise network (for instance, as different networks for the Red and Black ones, and both could be also multihomed).

This is also the case for civil and emergency networks, for example those used in airports, or by police/law enforcement, fireman, health

care, etc. and specially in the avenue of catastrophic events.

The importance of multihoming here is related to the need of high reliability, because the failure could potentially increase the risk for human life.

3.6 SOHO and Home

A Small Office/Home Office or Home, can fall into the category of a managed or unmanaged network. Usually there is a minimum management of the network, that could be done in-house, or as part of the service provided by the ISP, external consultants or systems integrators.

Is becoming very frequent the case where different ISPs provide service to the same SOHO/Home network, by means of different access networks (xDSL, Cable, Wireless, PLC/BPL, etc.).

A SOHO/Home network was usually a very small network with a single subnet, but this is also rapidly changing, and several subnets will be more frequently present in this scenario.

Some times this network can be part of an enterprise network, and consequently managed, usually connected via a VPN to the corporate network, so in this case, the multihoming could depend on the VPN access itself. But is also the case when different hosts or different applications need to chose either the VPN (for example email or corporate applications), or the ISP (for example regular web access); in this case the multihoming case is the same as in the enterprise network scenario (the VPN behaves as a separate ISP).

Today, the complete resilience of this network is already required, even if the solution cost is usually still too high. We can envision that the requirement for multihoming will become more frequent, specially considering the increase of tele-work and the usage of Internet for voice and video applications.

So the SOHO/Home network is positioned to enjoy the greatest benefit from multihoming. Although those networks typically do not have access to "enterprise-class" links with guaranteed uptime (or any guarantee at all) it is not uncommon for multiple providers to offer technologically diverse residential services in a single area. Multihoming to two or more such services (e.g., DSL, cable, satellite, BPL/PLC, etc.) can dramatically increase the resistance of a SOHO/Home network to external single points of failure. Resilience in those networks is becoming an important issue in the face of residential VoIP deployment. While it is unlikely that an enterprise environment will be without any conventional phone lines in the near

future, residential consumers are already experimenting with Internet-only phone service. It will become critical to deliver reliability approaching that of conventional phone lines if we are to maintain E911 service, fire alarm monitoring, and similar life safety functions at their current levels of effectiveness. While such reliability may be delivered by a single provider in select areas, multihoming should allow a larger set of consumers to set and meet their own standards through multiple attachment.

SOHO/Home networks may ultimately support life safety functions (e.g., health care, detectors, surveillance, etc.). The required reliability to monitor a smoke detector while a family sleeps is as least as great as that required by an enterprise where the most serious consequence of a network failure is likely financial.

In some cases, a host-centric multihoming approach could be sufficient for this scenario. This could be the case when single devices are directly connected to several access networks, for example for safety or security reasons.

3.7 3GPP

Is this the same as the mobility case ???.

Can a GGSN be multi-homed ???.

As IPv6 allows to have multiple addresses in each interface, the 3GPP terminal with a single interface can be multihomed using multiple IPv6 addresses with a single radio connection. Furthermore some 3GPP terminals are available with more than one interface (3GPP and WLAN) can be multihomed using one or several IP addresses in each interface.

Similarly, a GGSN can have multiple IPv6 addresses per interface and several upstream links.

3.8 Add-hoc

is this a special case ???. what about managed ad-hoc networks for the military domain ?

3.9 Mesh

Mesh networks are those that use its nodes as routers so that every node does not have to hear every other node directly (as in the classic ad-hoc case). It is more likely in this case that the entire network may touch more than one "external" provider, and consequently is multihomed.

4. Special Services and Applications within the Multihoming Scenarios

In addition to the generic scenarios, there are some special situations, services or applications, that could be parallel to any of the previously described cases, and should be considered in order to have a complete perspective for each of the above mentioned scenarios.

4.1 GRIDs and other temporary networks

An organization with his own network and addressing being connected to a bigger network, for example in a GRID situation, in a temporary or quasi-permanent basis. This can be also the example for research institutions that often connect their networks to other networks and may receive new addressing space creating a multi-homing situation.

4.2 Mobility

A mobile host or a mobile network can be simultaneously moving and still be attached to more than just one ISP. For example when connected simultaneously to a GPRS or 3GPP networks and a WISP (WLAN ISP).

Much more text needed here !.

4.3 Multihomed devices

Cellular phones with for example 3GPP and WLAN interfaces, laptops with 3GPP and Ethernet or even WLAN, all those are good examples of multihomed devices. This seems to be same as the host-centric multihoming case, but could also be related to the mobility scenario.

4.4 Renumbering

When a network is being renumbered, temporarily, during the renumbering process itself, it may become a multihomed network.

4.5 Real Time Traffic

The increase of the usage data networks and Internet for real time traffic (VoIP, video/audio streaming, videoconferencing, etc.).

For example, the impact of the penetration of VoIP for residential usage (SOHO, home), could bring situations where the conventional phone lines disappear, while several access networks connect those sites. In this case, multihoming is much more critical to maintain the level of service required for emergency (e.g. E911, medical, security, ...) and similar life facilities that today depend on the

telephone.

This is also specially relevant for rural areas, which today don't have copper connectivity, but are already attached to Internet by other means (satellite, PLC, WLAN, etc.).

Of course this could be also true for other kind of networks, like enterprise, emergency, etc.

4.6 Specific protocols/applications

Content Delivery Networks (CDNs), Internet Data Centers (IDC), DNS, DHCP, RA, ????.

TBD.

4.7 Others

any ??? TBD.

5. Security Considerations

This memo does not generate any new security considerations.

6. IANA Considerations

This document requests no action for IANA.

[[Note to RFC-editor: this section can be removed upon publication.]]

7. Acknowledgements

The authors would like to acknowledge the inputs from Jim Bound, Munechika Sumikawa, Antonio Tapiador, and the European Commission support in the co-funding of the Euro6IX project, where this work is being developed.

8 Informative References

[1] Abley, J., Black, B. and V. Gill, "Goals for IPv6 Site-Multihoming Architectures", RFC 3582, August 2003.

Authors' Addresses

Jordi Palet Martinez
Consulintel
San Jose Artesano, 1
Alcobendas - Madrid
E-28108 - Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
EMail: jordi.palet@consulintel.es

Miguel Angel Diaz Fernandez
Consulintel
San Jose Artesano, 1
Alcobendas - Madrid
E-28108 - Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
EMail: miguelangel.diaz@consulintel.es

Cesar Olvera
Consulintel
San Jose Artesano, 1
Alcobendas - Madrid
E-28108 - Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
EMail: cesar.olvera@consulintel.es

Alvaro Vives Martinez
Consulintel
San Jose Artesano, 1
Alcobendas - Madrid
E-28108 - Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
EMail: alvaro.vives@consulintel.es

Eric Fleischman
Boeing

Phone:
Fax:
EMail: eric.fleischman@boeing.com

Dan Lanciani

Phone:
Fax:
EMail: ddl@danlan.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

Acknowledgment

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