

April 2003

IPv6 Site Multihoming: Now What?

[draft-savola-multi6-nowwhat-00.txt](#)

Status of this Memo

This document is an Internet-Draft and is subject to 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/lid-abstracts.txt>.

To view the list Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

Abstract

ROUTING ARCHITECT'S WARNING: flagrant IPv4 site multihoming practices cause a significant increase the routing table size, change rates and instability, the tragedy of the commons by encouraging selfish routing practices, the exhaustion of the 16-bit AS number space, and may collapse the Internet interdomain routing architecture.

As there has been a desire to avoid similar problems as seen with IPv4, the use of similar techniques to achieve site multihoming has been prevented operationally in IPv6. However, the long effort to proceed with other IPv6 multihoming mechanisms has produced lots of heat but little light. This memo tries to list available techniques, split the organizations to different types to separately examine their multihoming needs, to look at the immediate and short-term solutions for these organization types, and to list a few immediate action items on how to proceed with IPv6 site multihoming.

Internet Draft [draft-savola-multi6-nowwhat-00.txt](#)

April 2003

Table of Contents

1.	Introduction	2
2.	Site Multihoming Mechanisms	3
3.	Classification of Sites and Motivations	4
3.1.	Classification	4
3.1.1.	Minimal	5
3.1.2.	Small	5
3.1.3.	Large	6
3.1.4.	International	6
4.	Choosing a Multihoming Mechanism	7
4.1.	Viable Multihoming Mechanisms	7
4.1.1.	Immediate Approaches	7
4.1.2.	Short-term Approaches	7
4.1.3.	Long Term Approaches	8
4.1.4.	Conclusion on Approaches	8
4.2.	Applicable Mechanisms for Site Types	8
4.2.1.	Minimal	9
4.2.2.	Small	9
4.2.3.	Large	9
4.2.4.	International	10
5.	Issues to Be Addressed	10
6.	Acknowledgements	11
7.	Security Considerations	11
8.	References	11
8.1.	Normative	11
8.2.	Informative	11
	Author's Address	13
A.	Independence as a Multihoming Reason	13
B.	Multi-connecting	13
	Intellectual Property Statement	14
	Full Copyright Statement	15

[1.](#) Introduction

Currently in IPv4, there seem to be three-four main mechanisms which are used to achieve at least some of the multihoming benefits: obtaining their own address space and AS number and advertising those, advertising more specific routes with a different path, using multi-connecting and leveraging NAT.

In IPv6, the first two of IPv4 mechanisms which are considered architecturally unscalable have been operationally prevented for now and the fourth does not exist.

IPv6 site multihoming effort has been in progress for many years already, with little light but lots of heat. This memo tries to increase the light, and show a few possible ways how to proceed with site multihoming in a relatively qualitative fashion.

First, the multihoming mechanisms which have been proposed and briefly considered are listed for reference; a more detailed description and analysis is omitted [[SMULTI](#)].

Then, a classification of sites and their multihoming motivations is presented. This is done to break the "site" and "multihoming" concepts to smaller, digestible pieces.

After that, the multihoming mechanisms are classified to immediate, short-term, and long-term approaches; long-term approaches are ignored when describing which techniques might fit for each organization type.

Last, some short term issues to be addressed and things to be done, are listed.

[V4MULTI] describes IPv4 multihoming properties. The motivations section seems to be missing an important factor, the explicit desire for operator-independence. This is described in [Appendix A](#). Similarly, multi-connecting -- which has typically been considered out of scope -- is briefly described in [Appendix B](#).

[2](#). Site Multihoming Mechanisms

In [[SMULTI](#)], quite a few different site multihoming mechanisms have been briefly described and analyzed:

- o Transport solutions, including TCP modifications [TCP+] and SCTP [[SCTP](#)],
- o Locator/Identifier separation solutions, including HIP [[HIP](#)],

- o LIN6 [[LIN6](#)] and Mobile IPv6 [MIPv6]),
- o Host-centric Multihoming [[HC](#)],
- o Multihoming at Site Exit Routers [[SITEEXIT](#)],
- o Geographic Address Allocation [[GEO](#)],
- o Provider Independent Addressing Derived from AS Numbers [[ASNPI](#)],

Savola

[Expires October 2003]

[Page 3]

Internet Draft

[draft-savola-multi6-nowwhat-00.txt](#)

April 2003

- o Multi-connecting (see [Appendix B](#) for introduction), and
- o Other methods, including: Advertising more specific routes [[SPECIFIC](#)], Multihoming with route aggregation [[RAGGR](#)], End-to-end multihoming [[END2END](#)], Traffic steering using routing header, Router renumbering [[RR](#)], and MHAP [[MHAP](#)].

A degree of familiarity with at least the most important techniques is assumed. The description and further analysis is out of scope.

[3.](#) Classification of Sites and Motivations

The different types of organizations are separated, and their possible motivations explored separately. In this way, it's possible to try to break the big "site" and "multihoming" concepts into smaller pieces.

First, a rough classification is presented, and then each of the four types are described at more length.

Analysis on which methods could suit each type best are described in the next section.

[3.1.](#) Classification

A table of multihoming motivations and types of sites is created; it is shown in the table below.

.-----.

	Independence	Redundancy	Load sharing
Minimal	no	no	no
Small	maybe	maybe	no
Large	maybe/yes	yes	maybe
International	yes	yes	yes

Note that the motivations "Performance" and "Policy" [\[V4MULTI\]](#) are not included above: in practice, they do not seem to be prime motivators, and it is difficult to gauge how desirable features they are.

The term "IP-users" is used in the classifications below. This is used to loosely refer to those people who have an operational need for Internet access to get their work done.

The organizations are classified roughly in four categories:

1. Minimal: a very small or restricted end-site, for example a typical home network, or a small office of less than 10 IP-users,
2. Small: a small-to-mid-size enterprise, for example with less than 50-150 IP-users,
3. Large: a large enterprise, for example a regional or national enterprise of typically less than 1000 IP-users, and
4. International: a large or very large enterprise, with a significant amount of international activity.

The distinction between the last two comes from the assumption that an international organization is assumed to have major activity in multiple countries or regions: in practice, one with a separate significant Internet connectivity with different ISPs in many areas. Naturally, even smaller organizations are likely to have some relatively minor international activity, but these are not likely counted as International.

Now, the reasons and motivations specific to each organization class

are described. These give a bit more elaboration on why the categories were written down as they were.

[3.1.1.](#) Minimal

Very minimal end-sites, such as typical home networks or very small enterprises, are quite small and typically do not include mission-critical activities.

Naturally, anyone would be willing to achieve multihoming benefits, but usually the associated costs, e.g. caused by obtaining physical connectivity to two ISPs, do not justify it.

In the very rare case that some multihoming is required, it can be done using the "Small" model.

[3.1.2.](#) Small

Small end-sites are typically small-to-mid-size enterprises, which operate mainly in one geographical location; branch offices are also possible, but these are typically handled using Virtual Private Networks (VPNs) or similar.

Small end-sites typically use the Internet in such a manner that they might find redundancy and independence important, depending on the

costs and complexity. However, typically an outage of some minutes or rarely even hours is not yet critical, and due to the size, operator-independence is not vital.

[3.1.3.](#) Large

Large end-sites are typically largish enterprises which operate in one or typically more geographical locations, however, mostly inside a region or a country. Relatively minor international branch offices are possible, but there are no major internal, private interconnects -- physical or link-layer connectivity that does not go through the Internet -- to the offices in other countries. Typically, if there are multiple locations inside a region, there are internal, private interconnects between the locations.

Redundancy is very important due to increased size: longer outages are unacceptable for productivity. Also, as the number of the nodes in the site is quite high, the effort of renumbering is also high, and some level of operator independence valuable but not always strictly necessary if significant enough ISPs operate in the area. Only in some relatively rare cases the network capacities or other parameters are exceeded so that some form of load-sharing is required.

3.1.4. International

International end-sites are typically large or very large enterprises which have significant employment internationally, connect to the Internet with high-speed links in each of these significant locations and may often have internal, private interconnects in place.

Typically, ISPs are not global enough to supply connectivity to all the locations, and being locked to scenic routing paths caused by only a few regional ISPs is not considered a serious option.

As the international enterprise typically exceeds the geographic and topological scope of any single ISP, and the network is so large, operator independence is considered extremely important. Redundancy is also critical. The required capacities and usage patterns may vary a lot, but often also some form of load sharing is necessary: all the traffic to the end-site cannot go through a single location, and it might not be reasonable either, due to geography.

4. Choosing a Multihoming Mechanism

In the previous section organizations were split into four main classifications; now it's possible to summarize the mechanisms and see if particular mechanisms could be made to meet the organizations' needs.

The first thing to realize is that the scope and breadth of

multihoming motivations and requirements vary a lot. It is highly unlikely that an "one size fits all" solution could be found: indeed, it seems fruitless to even try to look for one except possibly in purely research terms. Further, generic solutions are not probable to be found without larger architectural changes.

Therefore, by applying some rough classification one can hope to break the problem space into pieces and try to evaluate whether applicable solutions can be found for those pieces.

[4.1.](#) Viable Multihoming Mechanisms

First, it is important to take a look at and analyze different solutions to see which mechanisms could be viable and how they could be positioned.

A detailed analysis is omitted from this memo.

These are grouped to three categories: immediate, short-, and long-term. The definitions of the latter two are intentionally left vague, but short term solutions should not take more than 1-3 years to implement and deploy, at most.

[4.1.1.](#) Immediate Approaches

Multi-connecting seems to be the only obvious way to work around multihoming problems right now.

Host-centric multihoming and multihoming at site exit routers would also be applicable -- to an extent -- immediately, if no ISP would implement ingress filtering.

[4.1.2.](#) Short-term Approaches

Host-centric IPv6 multihoming and multihoming at site-exit routers, fully fleshed out, are both short-term solutions; both still need some work.

Provider independent addressing based on AS numbers is a short-term solution; the same applies to advertising more specific routes, if

done from specific allocations to make them distinguishable.

Parts of multihoming with route aggregation, traffic steering using routing header and router renumbering could be used in the short term, but the mechanisms themselves are not usable.

[4.1.3.](#) Long Term Approaches

All transport solutions are only generally applicable in the long term, if at all.

Identifier and locator separation solutions are long-term solutions; HIP the most credible of them all. LIN6 has fundamental issues which does not make it globally deployable. Mobile IPv6, if the address ownership and key management issue could be worked around, could be a short-term solution, a "hack". Architecturally, it seems a bit ill-fit as a long-term solution.

Geographic address allocation is a long-term solution if it's considered viable.

End-to-end multihoming is a long-term solution; it requires some major changes in thinking, and may be difficult to accept.

MHAP is a long-term solution; it moves into an area where most do not want to go architecturally, and it's likely it could not be made to work in practice.

[4.1.4.](#) Conclusion on Approaches

In the following analysis, none of the long-term approaches are seriously considered; they all need much more work to be adoptable.

Of the long-term solutions, it seems likely that at least identifier and locator solutions will prevail in some form or another. However, they do not solve the whole multihoming problem themselves.

[4.2.](#) Applicable Mechanisms for Site Types

Now, the mechanisms outlined in the multihoming analysis and considered viable, above, are applied to the rough end-site organization categories above.

[4.2.1.](#) Minimal

Minimal end-sites do not seem to require a multihoming mechanism. Of course, some would still find one preferable.

If they would, one building on "host-centric multihoming" approach would seem sufficient; this might not even be all too expensive using e.g. multiple xDSL connections.

[4.2.2.](#) Small

Small end-sites might find a multihoming mechanism desirable for redundancy and independence reasons.

The possible approaches are three-fold:

- o multi-connecting to one ISP, if redundancy is important,
- o host-centric multihoming, if independence is important, or
- o multihoming at site exit routers, if both are important.

Multi-connecting or multihoming at site exit routers provide a very extensive amount of redundancy and convergence; independence is obtained by connecting to multiple ISPs and deploying IP addresses from those, which would make the renumbering much less of a problem if it had to be done.

Solving the multihoming problem of small end-sites with an approach like "ASN-PI" or more specific routes is unscalable.

[4.2.3.](#) Large

Large end-sites typically require redundancy and possibly independence, sometimes desiring load sharing as well.

Often, the same mechanisms applied to small end-sites will also provide the sufficient level of redundancy and independence for also large end-sites.

One might be able to handle the scalability issues associated with approaches like "ASN-PI" with large end-sites, but such mechanisms should be avoided.

[4.2.4](#). International

International end-sites typically require redundancy, independence, and load sharing.

As their geographical scope typically exceed ISPs', assigning addresses from multiple ISPs, as with previous approaches, would lead to suboptimal routing.

One approach would be to depend on a change of the IP addressing and routing model: each significant international part of the organization would obtain separate IP address assignments from the local ISPs, and use an appropriate multihoming mechanism with that -- probably like with large end-sites above -- and interconnect the offices using mechanisms like VPN's; not even trying to obtain a homogeneous address space for all of the company.

This would be a "divide-and-conquer" approach to the problem: break it to smaller pieces and solve them independently.

If this approach can't be adopted, it seems the only realistic model would be to use something like "ASN-PI", more specific routes or separate address allocations, or similar.

[5](#). Issues to Be Addressed

Now, a few relatively short-term action items are presented which should be done as soon as possible.

- o Update and finish documenting IPv4 multihoming practices [\[V4MULTI\]](#),
- o Further than that, try to gain a better understanding why certain multihoming mechanisms are used with IPv4,
- o Finish documenting the IPv6 multihoming goals [\[GOALS\]](#),
- o Realize that certain problems like connection survivability may not be strict requirements in all cases of site multihoming, and

that such problems could be worked around.

- o Create a roadmap on how to proceed with multihoming solutions,
- o Work on short-term solutions, in particular, fix [[SITEEXIT](#)] and [[HC](#)] in the case that ISPs use ingress filtering and consider the case when uplink MTU is not higher than the one used within the site when using [[SITEEXIT](#)],

Savola

[Expires October 2003]

[Page 10]

Internet Draft

[draft-savola-multi6-nowwhat-00.txt](#)

April 2003

- o Work on procedures to ensure the separation -- who is "privileged" for which solution -- between different site types and multihoming methods, if different techniques are necessary, and
- o Start documenting how to do renumbering, and how to make renumbering easier; the renumbering doesn't have to work perfectly or even well, as long as it is satisfactory for some uses.

[6.](#) Acknowledgements

Most of the ideas on how to proceed originated in 2002 and in the beginning of 2003. The ideas on how to proceed were discussed prior to IETF56 in March 2003 with Tim Chown and Kurt Erik Lindqvist.

[7.](#) Security Considerations

This memo summarizes issues in IPv6 site multihoming and gives ideas on the way forward. Many multihoming techniques have security considerations which are not to be forgotten; they should be dealt with in the respective specifications and remembered when analyzing their applicability. Security is also an integral issue to consider when considering the overall site multihoming landscape. However, in itself, this memo does not present any security considerations.

[8.](#) References

[8.1.](#) Normative

- [V4MULTI] Abley, J., Black, B., and Gill, V., "IPv4 Multihoming Motivation, Practices and Limitations",

[draft-ietf-multi6-v4-multihoming-00](#), expired, Jun 2001.

[GOALS] Abley, J., Black, B., and Gill, V., "Goals for IPv6 Site-Multihoming Architectures", [draft-ietf-multi6-multihoming-requirements-05.txt](#), work-in-progress, Apr 2003.

8.2. Informative

[SMULTI] Savola, P., "Examining Site Multihoming in Finnish Networks", MSc Thesis, <http://staff.csc.fi/psavola/di.ps>, Apr 2003.

[TCP+] Tattam, P., "Preserving Active TCP sessions on Multihomed IPv6 Networks", <http://jazz-1.trumpet.com.au/ipv6-draft/preserve-tcp.txt>, Aug 2001.

Savola

[Expires October 2003]

[Page 11]

Internet Draft

[draft-savola-multi6-nowwhat-00.txt](#)

April 2003

[SCTP] Coene, L. (ed.), "Multihoming issues in the Stream Control Transmission Protocol", [draft-coene-sctp-multihome-03.txt](#), work-in-progress, Feb 2002.

[HIP] Moskowitz, R., "The Host Identity Payload Homepage", <http://homebase.htt-consult.com/HIP.html>.

[LIN6] Tereoka, F., et al., "LIN6: A Solution to Mobility and Multi-Homing in IPv6", [draft-teraoka-ipng-lin6-01.txt](#), expired, Aug 2001.

[MIPV6] Johnson, D., et al., "Mobility Support in IPv6", work-in-progress, Feb 2003.

[HC] Huitema, C., Draves, R., "Host-Centric IPv6 Multihoming", [draft-huitema-multi6-hosts-01.txt](#), expired, Jun 2002.

[SITEEXIT] Hagino, J., Snyder, H., "IPv6 Multihoming Support at Site Exit Routers", [RFC3178](#), Oct 2001.

[GEO] Hain, T., "Application and Use of the IPv6 Provider Independent Global Unicast Address Format", [draft-hain-ipv6-pi-addr-use-04.txt](#), work-in-progress, Apr 2003.

[ASNPI] Savola, P., "Multihoming Using IPv6 Addressing Derived

from AS Numbers", [draft-savola-multi6-asn-pi-00.txt](#), work-in-progress, Jan 2003.

[SPECIFIC] Lindqvist, K., "Multihoming in IPv6 by Multiple Announcements of Longer Prefixes", [draft-kurtis-multihoming-longprefix-00.txt](#), work-in-progress, Dec 2002.

[RAGGR] Yu, J., "IPv6 Multihoming with Route Aggregation", [draft-ietf-ipngwg-ipv6multihome-with-aggr-01.txt](#), expired, Aug 2000.

[END2END] Ohta, M., "The Architecture of End to End Multihoming", [draft-ohta-e2e-multihoming-03.txt](#), work-in-progress, Nov 2002.

[RR] Crawford, M., "Router Renumbering for IPv6", [RFC2894](#), Aug 2000.

[MHAP] Py, M., "Multi Homing Aliasing Protocol", [draft-py-mhap-01a.txt](#), expired, Apr 2002.

Savola

[Expires October 2003]

[Page 12]

Internet Draft [draft-savola-multi6-nowwhat-00.txt](#)

April 2003

Author's Address

Pekka Savola
CSC/FUNET
Espoo, Finland
EMail: psavola@funet.fi

[A.](#) Independence as a Multihoming Reason

This is an excerpt from [[SMULTI](#)].

Technical reasons for independence are mostly covered under redundancy; independence here focuses on economic, political and administrative perspectives.

Multihoming, especially with your own addresses can bring the organization some degree of provider independence. If the organization can just change its ISP's with minimal effort, one

undoubtedly has a better standing in the service agreement negotiations, e.g. being able to get a cheaper price and requested service agreement duration. ISP's could also just disappear, but in most cases services usually continue uninterrupted in the event of bankruptcy, mergers, etc. Also, especially some larger organizations consider it important to stand on their own, so that they can't be seen to depend on others.

An important factor for many, especially bigger sites, is also that the renumbering in the event of a network service provider change is considered too difficult and time-consuming, and it is desirable to find mechanisms which do not require that; such feature is provided by independence.

B. Multi-connecting

This is an excerpt from [[SMULTI](#)].

Multi-connecting means connecting multiple times to a single ISP. By some definition [[V4MULTI](#)], multi-connecting is not considered a multihoming mechanism. However, it still deserves a brief description.

Multi-connecting is typically achieved by having multiple site border routers and connecting each of them to separate routers at the ISP, usually in different locations.

Often, a routing protocol is run between the site and the ISP, to be able to quickly detect errors in the connectivity and to switch to alternative paths. Routing protocol used is often BGP with private

AS numbers.

Sometimes, typically when the ISP is also in charge of the management of site border routers, some other protocols such as OSPF or IS-IS are also possible -- this often requires strict filtering of advertisements at the ISP end, to prevent compromising the ISP's core routing system. In some cases, simply static routes may also be possible; typically this requires a link-layer medium which provides notifications if the end-to-end link-layer path becomes nonoperational.

Multi-connection is a flexible mechanism and relatively easy to accomplish: it requires no coordination between ISPs, no address applications for provider independent address space, no AS number applications for BGP AS numbers or the like. Yet it provides protection from the most common set of problems.

If one link fails, the changes do not need to be propagated further than the ISP; therefore, this is a very scalable approach when considering the global routing table.

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

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