

Levels", [RFC 2119](#), March 1997.

IPv6 Multihoming Workinggroup
Internet-Draft
Expires: August 24, 2003

K. Lindqvist
Netnod
February 23, 2003

A road-map for multihoming in IPv6
draft-kurtis-multi6-roadmap-00

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/lid-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 August 24, 2003.

Copyright Notice

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

Abstract

IPv6 was decided as the replacement to the widely deployed IPv4 in 1992. The new protocol was designed to meet a number of criteria that at that time was seen as the failures of IPv4, such as security and limited address space.

IPv6 meet these needs, but it fails to meet other needs of the Internet of today such as scalable solutions for multihoming and portable address space for end-sites. The effects of the lack of scalable solutions to this is widely documented.

In order to solve the multihoming scalability problems, the multi6

working group was created in the IETF. This group have so far not

Lindqvist

Expires August 24, 2003

[Page 1]

Internet-Draft

Multihoming road-map for IPv6

February 2003

yet produced any documents as it has been hard to find consensus on even the most basic documents such as requirements. Multi6 have also suffered from problems reaching consensus on what type of solution will be required moving forward, and the attempt of trying to come up with a 'fit-all' solution from day one.

This document tries to outline steps that could be taken to better understand the problem and what steps could be taken as we bit by bit tries to address the problem.

1. Introduction

When one starts to look at the multihoming problem, one will find that there are number of other issues that will start to have an impact of the problem and of potential solutions.

These factors are for example portability of address space, route convergence times, size of routing tables, growth of processing power, utilization of address space etc. All these issues are linked and will have an impact on what solutions can be used.

If the Internet is to scale far further than to where we are today, and preserve the end-to-end principle, it is essential that the model of the Internet at some stage addresses all these issues. So far many of the attempts at solving the multihoming problem have tried to address all of these at once, or have effectively stopped solutions to other problems.

The idea behind this document, is to try and list the pieces and how they are linked, at the same time as try to map potential solutions models to each of the problems. Going forward, solutions can then be benchmarked and ordered. It is important to realize that the most effective way to move forward at the moment seem to be to agree on a solution at a time and bit by bit get us to the complete solution to the multihoming problem and the problems that are linked to it.

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](#) [1].

"Site", "Transit provider" and "multihomed", are defined in [draft-ietf-multi6-multihoming-requirements-03](#) [2].

"Host multihoming" denotes a single host that have routing information, and/or addresses from two or more transit providers.

[3.](#) Determining the nature, type and scale of the problem

[3.1](#) Definition of the multihoming problem

Today many sites have the requirement of connections to multiple providers. This can be for a number of reasons such as business critical applications that require redundancy; Cost benefits of being able to route different types of traffic to different providers; etc.

This in most cases creates an entry in the routing table per each upstream provider. This creates a growth of the global routing table that we with todays routing protocols can not handle. The size becomes to large to keep in routers memory and the convergence time in the case of a failure becomes to high.

A related problem leading to the same effects is the fact that many enterprise users want stable IP addresses that do not change when they change providers. Currently these addresses are routed as independent blocks and therefor also contributing to the growth of

the global routing table.

[3.2](#) IPv6 and multihoming

When the need for a new IP protocol was being discussed in the first half of the 1990s, a number of problems that needed solutions was listed in order to have them solved at the same time. The Internet back then was much more hierarchical than the Internet of today. The first plans of an address architecture reflects this with the assignments being split up into different aggregator levels.

Since then multihoming have become widespread use among the enterprises and the needs of addresses have changed dramatically. This is also the explanation as to why multihoming never made it as a requirement for IPv6.

Years later when multihoming was starting to become popular and the growth of the global routing table was starting to become visible as a problem, a number of suggestions on how to address it was made. What these would look like have varied over time. Some of them have suggested to make changes to the IPv6 address model itself, while deployment is relatively limited, some have suggested changing the way IPv6 addresses are allocated, some have suggested other methods. Common though is the understanding that the migration to IPv6 gives us the opportunity to change other characteristics of the IP protocols.

Although IPv6 in itself does not give an answer to the multihoming problem, the solution will undoubtedly be closely linked with future

IPv6 deployment.

[3.3](#) Scaling to a solution

There exists a number of proposals to a solution to the multihoming problem. All or at least most of them will mostly work, but the large question is if they will scale better than what we have today, and to an increase in complexity that we can handle.

In order to judge this, we need to come to a common model of what scaling properties the problem will have over a certain time.

[3.3.1](#) Time

Determining the time and lifetime of the solution should not be that hard. If we use previous experiences from changing the routing model of the Internet we can look at the migration to BGP. If we leave out the transition time as the Internet is much larger today, the solution have lasted for roughly fifteen years.

If we assume a somewhat similar speed of growth of the Internet a solution need to last for around fifteen years. To add to this we need the time it takes from a agreed and published solution to make general deployment. A common understanding of this time seems to be around five years.

This means that a multihoming solution needs to scale to the growth over fifteen years.

[3.3.2](#) Scale

The hardest thing to determine in building a new multihoming model is to what number it needs to scale.

Currently multihoming growth is limited by the number of Autonomous System numbers that are available. Currently there are around 60,000 that are allocated by IANA to be handed out. Of these only 20,000 have been allocated and only around 12,000 are actually visible in the global routing table. This indicates that the current need for multihoming is somewhat limited. This is not to say that it will for the future.

The current amount of AS numbers available are limited by the 32-bit field of AS-number in the BGP specification. However, there is a proposal in the IDR working group of the IETF to increase this field to 64-bits. This would increase the amount of ASes available, and therefor the potential amount of routes in the global routing table.

However, it is worth noting that the increase in multihomed networks, will most likely increase, partly because of the change in nature of business done on the Internet, partly because of new billing models, partly because an available multihoming solution will attract more people to the solution.

One effect that will play a role in the scaling of the multihoming solution will be what type of networks that will be multihoming. There is a difference in requirements if these are multinational enterprise networks with multiple exit points around the world, or if these are multihomed home networks behind DSL lines, or multihomed mobile clients.

With the current development of the Internet it is most likely that we will see all of the above on a large scale. What will limit this will mostly be economical factors, i.e. who are willing to pay for these types of services. Still, the number of multihomed sites will be very large in twenty years. It is most likely safe to make calculations based on the fact that most western Europe, the middle east, Asia and the Americas will have these needs. This includes most of the worlds population. With the enterprise market this adds up to several billion of entries.

What is worth keeping in mind though is that this will not happen immediately. Rather it will grow slow at the start limited by the above mentioned shortage of AS numbers and the cost of multiple providers. This is important as a slow start will give us time and valuable data on the adaption and usage of multihoming as a function of the network. This is needed to make sure that the solution we choose actually will last for the required twenty years to come. We will also need to use the data to build growth curves, so that we can be sure we at each stage have a workable solution.

No one knows the exact growth, and we have no data to build on, but if we go on the known parts, that we 2003 still have 16-bit AS numbers; it will take us at least 2-3 years from agreement on 32-bit AS numbers to deployment and the same time for any modifications to the routing model; it will give us tree data points

2003-2006: less than 60k routes

2006: more than 4.2M routes

2020: ~5B routes

This is a far from perfect model, but at least gives us some indication on possible scenarios.

[4.](#) The nature of solutions

There are a number of already proposed solutions that can be categorized in five groups, host, addressing, routing, protocol modifications, and combinations. Below is a short overview of them and their possible applications. This document does not make any claim on judging the suitability of any of the solutions, for doing that we would need more data on the actual scaling problem we are trying to solve.

The solutions are grouped in a "implementable time line" of "short" (less than a year from standard publication), "medium" (between one and three years from standard publication), "long" (more than three years from standard publication).

[4.1](#) Hosts

One way of achieving multihoming that exists already today is to use multiple addresses on the end hosts. The advantage from a multihoming view is that the addresses are announced as part of the aggregate routes of the upstream providers.

This solution to some extent also give stability for a service over a switch of addresses, however it does not decrease the workload of a network administrator.

In variations this will also require that certain higher level applications will know which address it should talk to.

This solution could be targeted to the home end-user suppliers or other large scale deployments done through automatic addressing.

The advantage is that this is a solution that could be deployed in the short to medium term with no or little modification to other structures.

[4.2](#) Addressing

Some of the solutions suggest changing the methods used to allocate IP addresses, and then based on this build either routing models or modifications to the IP layer to handle the aggregation of routing information.

These models build on allocating addresses after well known facts such as geographic coordinates, population or other geo data.

Common to these solutions is that they need other supporting mechanisms in order to work.

Internet-Draft

Multihoming road-map for IPv6

February 2003

These are all medium term solutions as they require a change in the allocation policy, coordination between all the Regional Internet Registries, and work on behalf of the end-user network administrators.

[4.3](#) Routing

At the heart of the multihoming problem is the fact that the current routing model does not scale in terms of size and convergence times. The natural solution would be to address the routing model and try and come up with a new routing protocol, or more effective path finding algorithms. The problem is the scale.

A solution will have to involve a new routing model, but the scale and the today known algorithms for finding paths means that we need to exclude information for the algorithm to converge.

A routing solution also means that we need to modify the installed base, which - besides time take to agree on and publish a standard - will take years due to the number of sites and nodes involved.

A modification to the routing model alone will most likely not work in the long run and it is therefor safe to assume that a solution will have to be done in combination with another method.

Given the above a routing solution will fit in the long term solution group.

[4.4](#) The protocol

One way to move forward would be to simply modify the IP protocol itself. Most of these proposals build on the notion that an address actually consists of a "locator" and "identifier" part, therefor we could also split an address into two parts.

These solutions require a modification to not only the IP layer, but also to the way layers above use the addresses.

This is a highly intrusive modification, that besides the work needed on protocol changes and new standards, also will take a long time to implement. These solutions therefor are in the long term category.

[4.5](#) Combinations

Many of the proposals made are combinations of the methods above. This is simply because many of the solutions will not work on their own alone. Especially routing solutions will need help by other solutions to abstract data to an amount that can be handled.

Lindqvist

Expires August 24, 2003

[Page 9]

Internet-Draft

Multihoming road-map for IPv6

February 2003

Where the place these solutions in implementation time is hard, as this will depend on which solutions that have been combined. These solutions are however the most likely way forward as they will provide various migration steps as well.

[4.6](#) Others

The last category is solutions that do not fit above. These are solutions that build on continuing the current model, simply to make IPv6 more attractive for the enterprise. They do not have any medium or long term validity, but instead function as a starting point that will help give data and attract enterprises and end-users to use IPv6 for multihoming.

These solutions more based on doing multihoming exactly as done today and/or modify the way allocations of either IP addresses or AS numbers are made.

These solutions are all in the short term implementation category as they do not modify anything else except policy.

[5.](#) The road ahead

All attempts at solving the multihoming problem so far have tried to address the end solution, or at least large parts of it. In order to move forward we need to accept that the road to the end solution will consist of getting there bit by bit. We need to accept that solutions we adopt in order to move forward will not be perfect, but rather far from it. This should be ok as long as it help us to a better understanding of the problem and to gain momentum.

In order to achieve the goals above, we will need to agree on a certain set of requirements that meets at least a minimum set of issues. These requirements can and should be updated as we move along the way. Trying to identify all problems from the beginning risk that we miss out on parts we don't know yet and will also make it harder to get consensus for. If we instead go by this bit by bit, by defining a small part of the problem and the agreeing on the solution, we will both learn more and move faster.

All the solutions grouped earlier have merits for one part of the problem. A end solution might consist of them all, each tailored to meet a specific goal or part of the problem, with the legacy parts still in use in parts of the address space. Migration between the various solutions and steps will be slow, but should not be an issue as long as we at the end of the time period are at a common base.

The next steps should include ordering the proposed solutions to each category, agreeing on the implementation horizon, and most important

of all, start doing multihoming with IPv6 in order to gain experience and get data to build on.

[6](#). Security considerations

This document only provides an overview of various solutions and proposes a benchmarking method for ordering solutions to the multihoming problem, therefor no security issues should arise from this document.

Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement