

IPv6 Multihoming Workinggroup
Internet-Draft
Expires: June 1, 2003

K. Lindqvist
Netnod
December 2002

Multihoming in IPv6 by multiple announcements of longer prefixes
draft-kurtis-multihoming-longprefix-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 June 1, 2003.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

In the current IP version 4 Internet, many organizations that are not Internet Service Providers are still getting their Internet connectivity through multiple providers by having address space that are announced via multiple paths. This is commonly referred to as multihoming. For the IP version 6 Internet, there have long been a debate as to how to achieve the same effect. This document addresses these issues by highlighting some of the current policies and how these can be used to achieve multihoming for IP version 6 connected networks.

Internet-Draft

IPv6 Multihoming

December 2002

1. Introduction

Traditionally Internet Service Providers or ISPs have had multiple connections from their own network that is servicing customers to the rest of the Internet. In order to achieve this, they have been allocated blocks of IP addresses that are unique to them. These blocks are allocated based on the need an ISP can justify to the Regional Internet Registries, or RIRs, that are handling the allocations. Many enterprises have adopted the same technique with multiple connections to the rest of the world for a number of reasons such as increased redundancy, load sharing or simply costs of connections. For this they have also applied for similar blocks and their own AS numbers just as the ISPs use today. This document shows how the organizations that today are multihoming in IPv4 can achieve the same thing with the current policies of IPv6 networks.

[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](#)].

[3](#). Network requirements

In order to make use of the methods for achieving multihoming as described in this document, the network wishing to multihome in the IPv6 Internet will need to have obtained a block of IPv6 addresses, have multiple network connections either natively or through some transition or tunneling technique. The network must also be running the Border Gateway Protocol version 4 [\[2\]](#) with the BGP Multiprotocol Extensions [\[3\]](#) on their border routers. How do this with IPv6 is described in [\[4\]](#).

[4.](#) Obtaining IPv6 addresses

Networks that are registered as Local Internet Registries, or LIRs, should apply for a IPv6 delegation from their respective RIR. Networks that are not registered as LIRs should apply for IPv6 addresses from their network service provide. In the case of already existing multiple service providers, only one of the should allocate a address block to be used for multihoming with longer prefix announcements. The allocations currently being made to LIRs are 32 bit long prefixes and allocations made to end-sites are currently 48 bit prefixes.

[5.](#) Announcing a longer prefix

[5.1](#) Description of procedure

The first step in order to achieve the multihoming solution is establish an exchange of routing information using BGP to all the upstream providers of the network. These upstreams SHOULD accept a prefix length of 48 bits from their peers, and SHOULD NOT aggregate these routes. For the provider that have allocated the addresses to the network this is especially important. The multihomed network should configure it's border routers to announce the prefix it have received from one of it's providers to all it's peers. If these providers will accept the prefix is a business relationship between the multihomed network and their upstream providers. This will have

the affect that in the global routing table at least two paths for the multihomed network will be visible. One is the aggregate route representing the block allocated to the same ISP from where the multihomed network have received their IPv6 address block, the other route is a more specific route with a path through one of the upstreams of the multihomed network. There are a number of considerations that a network that wishes to use the approach needs to make. This approach does not have a long term scaling effect and will at some point in the future be abandoned.

[5.2](#) Reasoning of model

The first advantage of this model is that it gives a jump start mechanism to the problem of multihoming that have often been quoted as one of the main reasons as to why not IPv6 is deployed in the enterprise environment. Giving the enterprises a way to handle this will prove that claim either false or right. the second advantage of this model is that it does not require the implementation of any sort to the installed base of equipment. The third advantage of this model is that this mirrors existing operating practices in the IPv4 world. This is likely to lead to easier general adoption. However, this is also the largest point of concern with this model. This model will without doubt not scale for a widespread adoption so the growth rate of the routing table and the predictions of the growth needs to be monitored carefully. This model does however provide for mechanisms to control this as described below. This might be enough of concern for some enterprises to stay way from IPv6, but it does at least give them a way to start out with the technology and mirror the their existing network functions.

[5.3](#) What this model does not give you

The reasons for going to mutlihomining is many. Mostly these are for traffic flow purposes, being either resiliency or the ability to

better control how the traffic reaches the enterprise. Another popular reason in the IPv4 Internet have also been the fact that you with a multihoming setup would get so called Provider Independent addresses. These PI addresses are allocated directly from the RIRs to the enterprise. This means that a enterprise will keep the same the addresses even through a shift of ISPs as their upstreams. This model does not address this issue, and will not allow for address

allocations outside the generally adopted allocation rules of the RIRs.

[6. Effects on the growth of the current routing table](#)

[6.1 Aggregation effects with the current allocation policies](#)

In the IPv4 Internet routing table, also known as the Default Free Zone, there is currently around 125 000 routes announced with around 14000 originating ASes, of which 1500 ASes are providing transit. In recent years the routing table for the default free zone has grown tremendously. The reasons for this growth are still unknown, but reasons might be a mix of many new providers with a lower level of routing knowledge announcing larger allocations blocks as many more specifics, and the economic upswing around Internet based services (the so called dot com bubble). More recently the growth seems to have slowed down though. This could indicate that the interest in multihoming has decreased and that operating practices have improved. With the proposal for multihoming IPv6 as described in this document, the potential for the same routing table growth explosion as we have today exists. However, given that the current allocation blocks to ISPs are blocks with 32 bit long prefixes, one of these blocks will fit the current address assignments of most ISPs. This means that the routing table will be inherently small to start with. Even if all ASes that are active today were to announce IPv6 address space, it would only be around 14 000 routes. Almost 90% less of the current routing table. This means that we have quite a lot of breathing space before we start running into the same scalability problems as today. This is due to the current allocations. Doing multihoming by announcing longer prefixes out of allocated blocks will make the routing larger but looking at the current amount of multihomed networks, this should not pose a problem.

[6.2 Scaling back an explosion in routing table growth](#)

In the case this method of achieving multihoming is in so wide use that it starts to pose a problem, or that the routing table size in combination with the increased prefix length of IPv6 starts causing a problem, this model can be backed out from by simply reducing the length of prefixes that have to be accepted. Backing out will pose the problem that sites using this model today will no longer be able to multihome. However, given what we today know of the routing table growth this should not be a problem before the time where we should have other solutions for multihoming.

7. Security considerations

Using this approach, networks need to take care not to announce their prefixes in such a way that they generate asymmetric traffic flows. Traffic following asymmetrical paths might get blocked by so called strict Reverse Path Forwarding (RPF) checks. These checks are implemented in routers and will make sure that the source address of a packet arriving on an interface also have a route to the source address through that interface. RPF checks are implemented in order to help block address spoofing. Longer prefix multihoming SHOULD NOT be used as an excuse to disable RPF checks.

[8.](#) Acknowledgments

This paper builds on a discussion at the IETF-55 in Atlanta. People who was part in the idea is Thomas Narten and Chirstian Huitema but also all the people who where are the unofficial multi6 meeting.

Internet-Draft

IPv6 Multihoming

December 2002

Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.
- [2] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4", [RFC 1771](#), March 1995.
- [3] Bates, T., Chandra, R., Katz, D. and Y. Rekhter, "", [RFC 2283](#), February 1998.
- [4] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", [RFC 1771](#), March 1995.

Author's Address

Kurt Erik Lindqvist
Netnod Internet Exchange
Bellmansgatan 30
Stockholm S-118 47
Sweden

Phone: +46 708 30 60 01
EMail: kurtis@kurtis.pp.se
URI: <http://www.kurtis.pp.se>

Full Copyright Statement

Copyright (C) The Internet Society (2002). 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 assigns.

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.