

IAB/IESG Recommendations on IPv6 Address Allocations to Sites

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

This document provides recommendations to the addressing registries (APNIC, ARIN and RIPE-NCC) on policies for assigning IPv6 address blocks to end sites. In particular, it recommends the assignment of /48 in the general case, /64 when it is known that one and only one subnet is needed and /128 when it is absolutely known that one and only one device is connecting.

The original recommendations were made in an IAB/IESG statement mailed to the registries on September 1, 2000. This document refines the original recommendation and documents it for the historical record.

1. Introduction

There have been many discussions between IETF and RIR experts on the topic of IPv6 address allocation policy. This memo addresses the issue of the boundary in between the public and the private topology in the Internet, that is, how much address space should an ISP allocate to homes, small and large enterprises, mobile networks and transient customers.

This document does not address the issue of the other boundaries in the public topology, that is, between the RIRs and the LIRs.

This document was developed by the IPv6 Directorate, IAB and IESG, and is a recommendation from the IAB and IESG to the RIRs.

2. Background

The technical principles that apply to address allocation seek to balance healthy conservation practices and wisdom with a certain ease of access. On one hand, when managing a potentially limited resource, one must conserve wisely to prevent exhaustion within an expected lifetime. On the other hand, the IPv6 address space is in no sense as limited a resource as the IPv4 address space, and unwarranted conservatism acts as a disincentive in a marketplace already dampened by other factors. So from a market development perspective, we would like to see it be very easy for a user or an ISP to obtain as many IPv6 addresses as they really need without a prospect of immediate renumbering or of scaling inefficiencies.

The IETF makes no comment on business issues or relationships. However, in general, we observe that technical delegation policy can have strong business impacts. A strong requirement of the address delegation plan is that it not be predicated on or unduly bias business relationships or models.

The IPv6 address, as currently defined, consists of 64 bits of "network number" and 64 bits of "host number". The technical reasons for this are several. The requirements for IPv6 agreed to in 1993 included a plan to be able to address approximately 2^{40} networks and 2^{50} hosts; the 64/64 split effectively accomplishes this. Procedures used in host address assignment, such as the router advertisement of a network's prefix to hosts [RFC2462], which in turn place a locally unique number in the host portion, depend on this split. Subnet numbers must be assumed to come from the network part. This is not to preclude routing protocols such as IS-IS level 1 (intra-area) routing, which routes individual host addresses, but says that it may not be depended upon in the world outside that zone. The 64-bit host field can also be used with EUI-64 for a flat, uniquely allocated space, and therefore it may not be globally treated as a subnetting resource. Those concerned with privacy issues linked to the presence of a globally unique identifier may note that 64 bits makes a large enough field to maintain excellent random-number-draw properties for self-configured End System Designators. That alternative construction of this 64-bit host part of an IPv6 address is documented in [RFC3041].

While the IETF has also gone to a great deal of effort to minimize the impacts of network renumbering, renumbering of IPv6 networks is neither invisible nor completely painless. Therefore, renumbering should be considered a tolerable event, but to be avoided if reasonably feasible.

In [RFC2374] and [RFC2450], the IETF's IPNG working group has recommended that the address block given to a single edge network which may be recursively subnetted be a 48-bit prefix. This gives each such network 2^{16} subnet numbers to use in routing, and a very large number of unique host numbers within each network. This is deemed to be large enough for most enterprises, and to leave plenty of room for delegation of address blocks to aggregating entities.

It is not obvious, however, that all edge networks are likely to be recursively subnetted; a single PC in a home or a telephone in a mobile cellular network, for example, may or may not interface to a subnetted local network. When a network number is delegated to a place that will not require subnetting, therefore, it might be acceptable for an ISP to give a single 64-bit prefix - perhaps shared among the dial-in connections to the same ISP router. However this decision may be taken in the knowledge that there is objectively no shortage of /48s, and the expectation that personal, home networks will become the norm. Indeed, it is widely expected that all IPv6 subscribers, whether domestic (homes), mobile (vehicles or individuals), or enterprises of any size, will eventually possess multiple always-on hosts, at least one subnet with the potential for additional subnetting, and therefore some internal routing capability. In other words the subscriber allocation unit is not always a host; it is always potentially a site. The question this memo is addressing is how much address space should be delegated to such sites.

3. Address Delegation Recommendations

The IESG and the IAB recommend the allocations for the boundary between the public and the private topology to follow those general rules:

- /48 in the general case, except for very large subscribers.
- /64 when it is known that one and only one subnet is needed by design.
- /128 when it is absolutely known that one and only one device is connecting.

In particular, we recommend:

- Home network subscribers, connecting through on-demand or always-on connections should receive a /48.
- Small and large enterprises should receive a /48.
- Very large subscribers could receive a /47 or slightly shorter prefix, or multiple /48's.

- Mobile networks, such as vehicles or mobile phones with an additional network interface (such as bluetooth or 802.11b) should receive a static /64 prefix to allow the connection of multiple devices through one subnet.
- A single PC, with no additional need to subnet, dialing-up from a hotel room may receive its /128 IPv6 address for a PPP style connection as part of a /64 prefix.

Note that there seems to be little benefit in not giving a /48 if future growth is anticipated. In the following, we give the arguments for a uniform use of /48 and then demonstrate that it is entirely compatible with responsible stewardship of the total IPv6 address space.

The arguments for the fixed boundary are:

- That only by having a provider-independent boundary can we guarantee that a change of ISP will not require a costly internal restructuring or consolidation of subnets.
- That during straightforward site renumbering from one prefix to another the whole process, including parallel running of the two prefixes, would be greatly complicated if the prefixes had different lengths (depending of course on the size and complexity of the site).
- There are various possible approaches to multihoming for IPv6 sites, including the techniques already used for IPv4 multihoming. The main open issue is finding solutions that scale massively without unduly damaging route aggregation and/or optimal route selection. Much more work remains to be done in this area, but it seems likely that several approaches will be deployed in practice, each with their own advantages and disadvantages. Some (but not all) will work better with a fixed prefix boundary. (Multihoming is discussed in more detail below.)
- To allow easy growth of the subscribers' networks without need to go back to ISPs for more space (except for that relatively small number of subscribers for which a /48 is not enough).
- To remove the burden from the ISPs and registries of judging sites' needs for address space, unless the site requests more space than a /48. This carries several advantages:
 - It may become less critical for ISPs to be able to maintain detailed knowledge of their customers' network architecture and growth plans,

- ISPs and registries may reduce the effort spent on assessing rates of address consumption, with address space ample for long-term growth plans,
 - Registry operations may be made more efficient or more focused, by reducing the urgency of tracking and assessment.
 - Address space will no longer be a precious resource for customers, removing the major incentive for subscribers to install v6/v6 NATs, which would defeat the IPv6 restoration of address transparency.
-
- To allow the site to maintain a single reverse-DNS zone covering all prefixes.
 - If and only if a site can use the same subnetting structure under each of its prefixes, then it can use the same zone file for the address-to-name mapping of all of them. And, using the conventions of [RFC2874], it can roll the reverse mapping data into the "forward" (name-keyed) zone.

Specific advantages of the fixed boundary being at /48 include

- To leave open the technical option of retro-fitting the GSE (Global, Site and End-System Designator, a.k.a., "8+8") proposal for separating locators and identifiers, which assumes a fixed boundary between global and site addressing at /48. Although the GSE technique was deferred a couple of years ago, it still has strong proponents. Also, the IRTF Namespace Research Group is actively looking into topics closely related to GSE. It is still possible that GSE or a derivative of GSE will be used with IPv6 in the future.
- Since the site-local prefix is fec0::/48, global site prefixes of /48 will allow sites to easily maintain a trivial (identity) mapping between the global topology and the site-local topology in the SLA field.
- Similarly, if the 6to4 proposal is widely deployed, migration from a 6to4 prefix, which is /48 by construction, to a native IPv6 prefix will be simplified if the native prefix is /48.

4. Conservation of Address Space

The question naturally arises whether giving a /48 to every subscriber represents a profligate waste of address space. Objective analysis shows that this is not the case. A /48 prefix under the 001 Global Unicast Address prefix contains 45 variable bits. That is, the number of available prefixes is 2 to the power 45 or about 35 trillion (35,184,372,088,832).

More precisely,

- [RFC1715] defines an "H ratio" based on experience in address space assignment in various networks. The H ratio varies between 0 and 0.3, with larger values denoting denser, more efficient assignment. Experience shows that problems start to occur when the H ratio becomes greater than 0.25. At an H ratio of 0.25, a 45 bit address space would have 178 billion (178 thousand million) identifiers.

$$H = \log_{10}(178 \cdot 10^9) / 45 = 0.25$$

This means that we feel comfortable about the prospect of allocating 178 billions /48 prefixes under that scheme before problems start to appear. To understand how big that number is, one has to compare 178 billion to 10 billion, which is the projected population on earth in year 2050 (see <http://www.census.gov/ipc/www/world.html>). These numbers give no grounds for concern provided that the ISPs, under the guidance of the RIRs, allocate /48's prudently, and that the IETF refrains from new recommendations that further reduce the remaining 45 variable bits, unless a compelling requirement emerges.

- We are highly confident in the validity of this analysis, based on experience with IPv4 and several other address spaces, and on extremely ambitious scaling goals for the Internet amounting to an 80 bit address space *per person*. Even so, being acutely aware of the history of under-estimating demand, the IETF has reserved more than 85% of the address space (i.e., the bulk of the space not under the 001 Global Unicast Address prefix). Therefore, if the analysis does one day turn out to be wrong, our successors will still have the option of imposing much more restrictive allocation policies on the remaining 85%. However, we must stress that vendors should not encode any of the boundaries discussed here either in software nor hardware. Under that assumption, should we ever have to use the remaining 85% of the address space, such a migration may not be devoid of pain, but it should be far less disruptive than deployment of a new version of IP.

To summarize, we argue that although careful stewardship of IPv6 address space is essential, this is completely compatible with the convenience and simplicity of a uniform prefix size for IPv6 sites of any size. The numbers are such that there seems to be no objective risk of running out of space, giving an unfair amount of space to early customers, or of getting back into the over-constrained IPv4

situation where address conservation and route aggregation damage each other.

5. Multihoming Issues

In the realm of multi-homed networks, the techniques used in IPv4 can all be applied, but they have known scaling problems. Specifically, if the same prefix is advertised by multiple ISPs, the routing information will grow as a function of the number of multihomed sites. To go beyond this for IPv6, we only have initial proposals on the table at this time, and active work is under way in the IETF IPNG and Multi6 working groups. Until current or new proposals become more fully developed, existing techniques known to work in IPv4 will continue to be used in IPv6.

Key characteristics of an ideal multi-homing proposal include (at minimum) that it provides routing connectivity to any multi-homed network globally, conserves address space, produces high quality routes via any of the network's providers, enables a multi-homed network to connect to multiple ISPs, does not unintentionally bias routing to use any proper subset of those networks, does not damage route aggregation, and scales to very large numbers of multi-homed networks.

One class of solutions being considered amounts to permanent parallel running of two (or more) prefixes per site. In the absence of a fixed prefix boundary, such a site might be required to have multiple different internal subnet numbering strategies, (one for each prefix length) or, if it only wanted one, be forced to use the most restrictive one as defined by the longest prefix it received from any of its ISPs. In this approach, a multi-homed network would have an address block from each of its upstream providers. Each host would either have exactly one address picked from the set of upstream providers, or one address per host from each of the upstream providers. The first case is essentially a variant on [RFC2260], with known scaling limits.

In the second case (multiple addresses per host), if two multi-homed networks communicate, having respectively M and N upstream providers, then the one initiating the connection will select one address pair from the $N \times M$ potential address pairs to connect between, and in so doing will select the providers, and therefore the applicable route, for the life of the connection. Given that each path will have a different available bit rate, loss rate, and delay, if neither host is in possession of any routing or metric information, the initiating host has only a $1/(M \times N)$ probability of selecting the optimal address pair. Work on better-than-random address selection is in progress in the IETF, but is incomplete.

The existing IPv4 Internet shows us that a network prefix which is independent of, and globally advertised to, all upstream providers permits the routing system to select a reasonably good path within the applicable policy. Present-day routing policies are not QoS policies but reachability policies, which means that they will not necessarily select the optimal delay, bit rate, or loss rate, but the route will be the best within the metrics that are in use. One may therefore conclude that this would work correctly for IPv6 networks as well, apart from scaling issues.

6. Security Considerations

This document does not have any security implications.

7. Acknowledgments

This document originated from the IETF IPv6 directorate, with much input from the IAB and IESG. The original text forming the basis of this document was contributed by Fred Baker and Brian Carpenter. Allison Mankin and Thomas Narten merged the original contributions into a single document, and Alain Durand edited the document through its final stages.

8. References

- [RFC1715] Huitema, C., "The H Ratio for Address Assignment Efficiency", RFC 1715, November 1994.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [RFC2260] Bates, T. and Y. Rekhter, "Scalable Support for Multi-homed Multi-provider Connectivity", RFC 2260, January 1998.
- [RFC2374] Hinden, R., O'Dell, M. and S. Deering, "An IPv6 Aggregatable Global Unicast Address Format", RFC 2374, July 1998.
- [RFC2450] Hinden, R., "Proposed TLA and NLA Assignment Rule", RFC 2450, December 1998.
- [RFC2462] Thompson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [RFC2874] Crawford, M. and C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering", RFC 2874, July 2000.

- [RFC3041] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.
- [MobIPv6] Johnson, D. and C. Perkins, "Mobility Support in IPv6", Work in Progress.

9. Authors Address

Internet Architecture Board

Email: iab@iab.org

Internet Engineering Steering Group

Email: iesg@ietf.org

10. Full Copyright Statement

Copyright (C) The Internet Society (2001). 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.