

Internet Engineering Task Force (IETF)  
Request for Comments: 6177  
BCP: 157  
Obsoletes: 3177  
Category: Best Current Practice  
ISSN: 2070-1721

T. Narten  
IBM  
G. Huston  
APNIC  
L. Roberts  
Stanford University  
March 2011

## IPv6 Address Assignment to End Sites

### Abstract

RFC 3177 argued that in IPv6, end sites should be assigned /48 blocks in most cases. The Regional Internet Registries (RIRs) adopted that recommendation in 2002, but began reconsidering the policy in 2005. This document obsoletes the RFC 3177 recommendations on the assignment of IPv6 address space to end sites. The exact choice of how much address space to assign end sites is an issue for the operational community. The IETF's role in this case is limited to providing guidance on IPv6 architectural and operational considerations. This document reviews the architectural and operational considerations of end site assignments as well as the motivations behind the original recommendations in RFC 3177. Moreover, this document clarifies that a one-size-fits-all recommendation of /48 is not nuanced enough for the broad range of end sites and is no longer recommended as a single default.

This document obsoletes RFC 3177.

### Status of This Memo

This memo documents an Internet Best Current Practice.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on BCPs is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6177>.

## Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

## Table of Contents

1. Introduction .....	3
2. On /48 Assignments to End Sites .....	4
3. Other RFC 3177 Considerations .....	6
4. Impact on IPv6 Standards .....	6
4.1. RFC 3056: Connection of IPv6 Domains via IPv4 Clouds .....	6
4.2. IPv6 Multicast Addressing .....	7
5. Summary .....	7
6. Security Considerations .....	8
7. Acknowledgments .....	8
8. Informative References .....	8

## 1. Introduction

There are a number of considerations that factor into address assignment policies. For example, to provide for the long-term health and scalability of the public routing infrastructure, it is important that addresses aggregate well [ROUTE-SCALING]. Likewise, giving out an excessive amount of address space could result in premature depletion of the address space. This document focuses on the (more narrow) question of what is an appropriate IPv6 address assignment size for end sites. That is, when end sites request IPv6 address space from ISPs, what is an appropriate assignment size.

RFC 3177 [RFC3177] called for a default end site IPv6 assignment size of /48. Subsequently, the Regional Internet Registries (RIRs) developed and adopted IPv6 address assignment and allocation policies consistent with the recommendations of RFC 3177 [RIR-IPV6]. In 2005, the RIRs began discussing IPv6 address assignment policy again. Since then, APNIC [APNIC-ENDSITE], ARIN [ARIN-ENDSITE], and RIPE [RIPE-ENDSITE] have revised the end site assignment policy to encourage the assignment of smaller (i.e., /56) blocks to end sites.

This document obsoletes RFC 3177, updating its recommendations in the following ways:

- 1) It is no longer recommended that /128s be given out. While there may be some cases where assigning only a single address may be justified, a site, by definition, implies multiple subnets and multiple devices.
- 2) RFC 3177 specifically recommended using prefix lengths of /48, /64, and /128. Specifying a small number of fixed boundaries has raised concerns that implementations and operational practices might become "hard-coded" to recognize only those fixed boundaries (i.e., a return to "classful addressing"). The actual intention has always been that there be no hard-coded boundaries within addresses, and that Classless Inter-Domain Routing (CIDR) continues to apply to all bits of the routing prefixes.
- 3) This document moves away from the previous recommendation that a single default assignment size (e.g., a /48) makes sense for all end sites in the general case. End sites come in different shapes and sizes, and a one-size-fits-all approach is not necessary or appropriate.

This document does, however, reaffirm an important assumption behind RFC 3177:

A key principle for address management is that end sites always be able to obtain a reasonable amount of address space for their actual and planned usage, and over time ranges specified in years rather than just months. In practice, that means at least one /64, and in most cases significantly more. One particular situation that must be avoided is having an end site feel compelled to use IPv6-to-IPv6 Network Address Translation or other burdensome address conservation techniques because it could not get sufficient address space.

This document does not make a formal recommendation on what the exact assignment size should be. The exact choice of how much address space to assign end sites is an issue for the operational community. The IETF's role in this case is limited to providing guidance on IPv6 architectural and operational considerations. This document provides input into those discussions. The focus of this document is to examine the architectural issues and some of the operational considerations relating to the size of the end site assignment.

## 2. On /48 Assignments to End Sites

Looking back at some of the original motivations behind the /48 recommendation [RFC3177], there were three main concerns. The first motivation was to ensure that end sites could easily obtain sufficient address space without having to "jump through hoops" to do so. For example, if someone felt they needed more space, just the act of asking would at some level be sufficient justification. As a comparison point, in IPv4, typical home users are given a single public IP address (though even this is not always assured), but getting any more than one address is often difficult or even impossible -- unless one is willing to pay a (significantly) increased fee for what is often considered to be a "higher grade" of service. (It should be noted that increased ISP charges to obtain a small number of additional addresses cannot usually be justified by the real per-address cost levied by RIRs, but additional addresses are frequently only available to end users as part of a different type or "higher grade" of service, for which an additional charge is levied. The point here is that the additional cost is not due to the RIR fee structures, but to business choices ISPs make.) An important goal in IPv6 is to significantly change the default and minimal end site assignment, from "a single address" to "multiple networks" and to ensure that end sites can easily obtain address space.

A second motivation behind the original /48 recommendation was to simplify the management of an end site's addressing plan in the presence of renumbering (e.g., when switching ISPs). In IPv6, a site may simultaneously use multiple prefixes, including one or more public prefixes from ISPs as well as Unique Local Addresses [ULA-ADDRESSES]. In the presence of multiple prefixes, it is significantly less complex to manage a numbering plan if the same subnet numbering plan can be used for all prefixes. That is, for a link that has (say) three different prefixes assigned to it, the subnet portion of those prefixes would be identical for all assigned addresses. In contrast, renumbering from a larger set of "subnet bits" into a smaller set is often painful, as it can require making changes to the network itself (e.g., collapsing subnets). Hence, renumbering a site into a prefix that has (at least) the same number of subnet bits is more straightforward, because only the top-level bits of the address need to change. A key goal of the recommendations in RFC 3177 is to ensure that upon renumbering, one does not have to deal with renumbering into a smaller subnet size.

It should be noted that similar arguments apply to the management of zone files in the DNS. In particular, managing the reverse (ip6.arpa) tree is simplified when all links are numbered using the same subnet ids.

A third motivation behind the /48 recommendation was to better support network growth common at many sites. In IPv4, it is usually difficult (or impossible) to obtain public address space for more than a few months worth of projected growth. Thus, even slow growth over several years can lead to the need to renumber into a larger address block. With IPv6's vast address space, end sites can easily be given more address space (compared with IPv4) to support expected growth over multi-year time periods.

While the /48 recommendation does simplify address space management for end sites, it has also been widely criticized as being wasteful. For example, a large business (which may have thousands of employees) would, by default, receive the same amount of address space as a home user, who today typically has a single (or small number of) LAN and a small number of devices (dozens or less). While it seems likely that the size of a typical home network will grow over the next few decades, it is hard to argue that home sites will make use of 65K subnets within the foreseeable future. At the same time, it might be tempting to give home sites a single /64, since that is already significantly more address space compared with today's IPv4 practice. However, this precludes the expectation that even home sites will grow to support multiple subnets going forward. Hence, it is strongly intended that even home sites be given multiple subnets

worth of space, by default. Hence, this document still recommends giving home sites significantly more than a single /64, but does not recommend that every home site be given a /48 either.

A change in policy (such as above) would have a significant impact on address consumption projections and the expected longevity for IPv6. For example, changing the default assignment from a /48 to /56 (for the vast majority of end sites, e.g., home sites) would result in a savings of up to 8 bits, reducing the "total projected address consumption" by (up to) 8 bits or two orders of magnitude. (The exact amount of savings depends on the relative number of home users compared with the number of larger sites.)

The above-mentioned goals of RFC 3177 can easily be met by giving home users a default assignment of less than /48, such as a /56.

### 3. Other RFC 3177 Considerations

RFC 3177 suggested that some multihoming approaches (e.g., Generalized Structure Element (GSE)) might benefit from having a fixed /48 boundary. This no longer appears to be a consideration.

RFC 3177 argued that having a "one-size-fits-all" default assignment size reduced the need for customers to continually or repeatedly justify the usage of existing address space in order to get "a little more". Likewise, it also reduces the need for ISPs to evaluate such requests. Given the large amount of address space in IPv6, there is plenty of space to grant end sites enough space to be consistent with reasonable growth projections over multi-year time frames. Thus, it remains highly desirable to provide end sites with enough space (on both initial and subsequent assignments) to last several years. Fortunately, this goal can be achieved in a number of ways and does not require that all end sites receive the same default size assignment.

### 4. Impact on IPv6 Standards

#### 4.1. RFC 3056: Connection of IPv6 Domains via IPv4 Clouds

RFC 3056 [RFC3056] describes a way of generating IPv6 addresses from an existing public IPv4 address. That document describes an address format in which the first 48 bits concatenate a well-known prefix with a globally unique public IPv4 address. The "SLA ID" field is assumed to be 16 bits, consistent with a 16-bit "subnet id" field. To facilitate transitioning from the address numbering scheme in RFC 3056 to one based on a prefix obtained from an ISP, an end site would be advised to number out of the right most bits first, using the leftmost bits only if the size of the site made that necessary.

Similar considerations apply to other documents that allow for a subnet id of 16 bits, including [ULA-ADDRESSES].

#### 4.2. IPv6 Multicast Addressing

Some IPv6 multicast address assignment schemes embed a unicast IPv6 prefix into the multicast address itself [RFC3306]. Such documents do not assume a particular size for the subnet id, per se, but do assume that the IPv6 prefix is a /64. Thus, the relative size of the subnet id has no direct impact on multicast address schemes.

#### 5. Summary

The exact choice of how much address space to assign end sites is an issue for the operational community. The recommendation in RFC 3177 [RFC3177] to assign /48s as a default is not a requirement of the IPv6 architecture; anything of length /64 or shorter works from a standards perspective. However, there are important operational considerations as well, some of which are important if users are to share in the key benefit of IPv6: expanding the usable address space of the Internet. The IETF recommends that any policy on IPv6 address assignment policy to end sites take into consideration the following:

- it should be easy for an end site to obtain address space to number multiple subnets (i.e., a block larger than a single /64) and to support reasonable growth projections over long time periods (e.g., a decade or more).
- the default assignment size should take into consideration the likelihood that an end site will have need for multiple subnets in the future and avoid the IPv4 practice of having frequent and continual justification for obtaining small amounts of additional space.
- Although a /64 can (in theory) address an almost unlimited number of devices, sites should be given sufficient address space to be able to lay out subnets as appropriate, and not be forced to use address conservation techniques such as using bridging. Whether or not bridging is an appropriate choice is an end site matter.
- assigning a longer prefix to an end site, compared with the existing prefixes the end site already has assigned to it, is likely to increase operational costs and complexity for the end site, with insufficient benefit to anyone.

- the operational considerations of managing and delegating the reverse DNS tree under ip6.arpa on nibble versus non-nibble boundaries should be given adequate consideration.

## 6. Security Considerations

This document has no known security implications.

## 7. Acknowledgments

This document was motivated by and benefited from numerous conversations held during the ARIN XV and RIPE 50 meetings in April-May, 2005.

## 8. Informative References

- [APNIC-ENDSITE] "prop-031: Proposal to amend APNIC IPv6 assignment and utilisation requirement policy,"  
<http://www.apnic.net/policy/proposals/prop-031>
- [ARIN-ENDSITE] "2005-8: Proposal to amend ARIN IPv6 assignment and utilisation requirement",  
[http://www.arin.net/policy/proposals/2005\\_8.html](http://www.arin.net/policy/proposals/2005_8.html)
- [RIR-IPV6] ARIN: <http://www.arin.net/policy/nrpm.html#ipv6>; RIPE Document ID: ripe-267, Date: 22 January 2003  
<http://www.ripe.net/ripe/docs/ipv6policy.html>; APNIC: <http://www.apnic.net/docs/policy/ipv6-address-policy.html>
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3306] Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6 Multicast Addresses", RFC 3306, August 2002.
- [RFC3177] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.
- [RIPE-ENDSITE] "Proposal to Amend the IPv6 Assignment and Utilisation Requirement Policy", 2005-8,  
<http://www.ripe.net/ripe/policies/proposals/2005-08>.
- [ROUTE-SCALING] "Routing and Addressing Problem Statement", Work in Progress, February 2010.

[ULA-ADDRESSES] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

#### Authors' Addresses

Thomas Narten  
IBM Corporation  
3039 Cornwallis Ave.  
PO Box 12195  
Research Triangle Park, NC 27709-2195

Phone: 919-254-7798  
EMail: narten@us.ibm.com

Geoff Huston  
APNIC

EMail: gih@apnic.net

Rosalea G Roberts  
Stanford University, Networking Systems  
P.O. Box 19131  
Stanford, CA 94309-9131

EMail: lea.roberts@stanford.edu  
Phone: +1-650-723-3352