

IPv6 Operations (v6ops)

Internet-Draft

Obsoletes: [6177](#) (if approved)

Intended status: Best Current Practice

Expires: December 29, 2018

L. Roberts

Stanford University, UIT

J. Palet Martinez

The IPv6 Company

June 27, 2018

**IPv6 Address Assignment to End-Sites**  
**draft-palet-v6ops-rfc6177-bis-00**

Abstract

The Regional Internet Registries (RIRs) policies, have different views regarding the recommendation of the prefix to be assigned to end-sites. However, all them allow up to a /48 without further justification and clearly state that the exact choice of how much address space should be assigned to end-sites is a decision of each operator. In fact, the operational community has recently developed a document with also covers this topic "Best Current Operational Practice for Operators: IPv6 prefix assignment for end-users - persistent vs non-persistent, and what size to choose".

This document reviews the architectural and operational considerations of end-site assignments, and reiterates that assignment policy and guidelines belong to the RIR community. This revision is being made to emphasize that IPv6 protocol evolution requires an ever increasing availability of subnets at the end-site and so, policy should reflect that assignment of a single subnet is no longer appropriate unless the recipient explicitly agrees to the limitations implied by such an assignment.

This document obsoletes [RFC6177](#) (IPv6 Address Assignment to End Sites).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on December 29, 2018.

## Copyright Notice

Copyright (c) 2018 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 (<https://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.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Considerations Regarding the Prefix Length . . . . .	<a href="#">4</a>
<a href="#">3.</a>	On /48 Assignments to End-Sites . . . . .	<a href="#">5</a>
<a href="#">4.</a>	Impact on IPv6 Standards . . . . .	<a href="#">7</a>
<a href="#">5.</a>	Summary . . . . .	<a href="#">7</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">9</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">9</a>
<a href="#">8.</a>	Acknowledgements . . . . .	<a href="#">9</a>
<a href="#">9.</a>	Informative References . . . . .	<a href="#">9</a>
	Authors' Addresses . . . . .	<a href="#">10</a>

## [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 (narrower) 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.

[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 those recommendations, and it triggered the development of [[RFC6177](#)]. However, some of the RIRs, have later on updated those policies, which still allow using /48, but leave the decision in the



hands of the ISP, or even, in some cases, encourage the assignment of smaller (i.e., /56) blocks to residential end-sites, while keeping /48 for business.

More recently, a Global IPv6 Deployment Survey (for residential/household services) has been conducted since May 2016, with responses from over 1.559 ISPs, from 105 different countries, which latest results have been presented in January 2018 [[IPv6-Survey](#)]. This survey is showing that 23% of the responders (generally in more advanced countries, in terms of IPv6 deployment) assign a /48, 35% assign a /56 and 33% assign a single /64.

This raises the question of a possible misinterpretation of [[RFC6177](#)] by at least 1/3rd of the operational community and consequently, the need to revisit it.

This document obsoletes [[RFC6177](#)], updating its recommendations in the following ways:

1. It is extremely discouraged 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. [[RFC3177](#)] 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. [[RFC6177](#)] moved 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.
4. This revision has been created to more clearly assert the requirement to ensure that address assignments to end-sites provide a sufficiently big number of subnets (/64 on classic networks) to each end-site, taking under consideration the end-site's future expected needs, new deployment expectations and new protocol requirements, among others. Once all these are considered, it seems unlikely that a single subnet (/64) or even a small number of them should be assigned, unless very clearly justified and agreed to by the end-site.



This document reaffirms, as [[RFC6177](#)] did, an important assumption behind [[RFC3177](#)], using however, a much stronger and clearer language:

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 many years, probably decades, rather than just months. In practice, that means that a single /64 is not a choice, even for a small residential network, which following technology trends will need a sufficiently big number of /64's. 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, beyond what has been indicated in the precedent paragraph. The exact choice of how much address space to assign end-sites is an operational issue and under that context, discussed already in [[RIPE-690](#)].

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. Considerations Regarding the Prefix Length**

[RFC7608] already discusses about the IPv6 prefix length recommendations for forwarding, and the need for routing and forwarding implementations to ensure that longest-prefix-match works on any prefix length.

Any prefix length up to /128 is treated identically by routing protocols, however for a given network, end-site, subnet or link, there always exists a Longest Acceptable Prefix (LAP), whose length is locally determined - e.g. a site or link that uses SLAAC has a LAP of /64 and will not work with a longer one.

This consideration should be noticed, across this document, in the sense that end-sites usually have subnets that use, by default, SLAAC, and consequently, the LAP is mandatorily a /64. Other technologies, may have a different LAP, which must be used accordingly.



### 3. 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.

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 a big number of subnets worth of space, by default. Hence, this document still recommends giving home sites significantly many more than a single /64.

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 [[RFC4193](#)]. 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 [[RFC3177](#)] 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.

Furthermore, to keep addressing plans usable and understandable, and to align with DNS reverse zone delegations, the size of the assigned prefix should be aligned with a nibble boundary. Each hexadecimal character in an IPv6 prefix represents one nibble, which is 4 bits. The length of a delegated prefix should therefore always be a multiple of 4, so the possible choices are /48, /52, /56 and /60.

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. While reasonable people may disagree over whether all end sites should get a /48 assignment by default, reasonable people do agree that an end-site should be able to get up to a /48 by request. It is important to stress that the strength of IPv6 is the vast size of its address space, which should allow users to easily acquire as many subnets as required for their applications, plus room to grow. Math's show that even assuming 32 billion of humans ( $2^{35}$ ), and assigning each of them 4 /48's, with a 50% routing efficiency, we can repeat that  $2^{10}$  times, so if average life span of each human is 100 years, and we don't recover back the /48's, we will be able to use IPv6 addressing space for over 100.000 years.

A large business (which may have thousands of employees) would, by default, receive the same amount of address space, for each of its end-sites, as a home user. However, it is clear that that for each corporate end-site it is perfectly feasible to justify further needs if that becomes the case, and RIR policies allow that.

Today typically, a home has already a considerable number of possible subnets (a common CE has 4 LAN ports, 2 WiFi radios which support several SSIDs each one, VoIP subnet, IPTV subnet, additional VLANs)



and if downstream routers are used, there is a need for further subnets. This means that in a short term, assigning a /60 (16 subnets), it is already a really bad decision, as it may enforce IPv6 NAT between the main CE and downstream routers.

It will become very common that homes start using technologies like HNCP [[RFC7788](#)], and this increases the need for more subnets, which means that /56 (256 subnets) may be too short also in very few years.

Finally, considerations about multiple addresses per host [[RFC7934](#)] and techniques to allow a single /64 per host/interface [[RFC7934](#)], means that we will see in the short term, many home devices that will take advantage of that, either for security reasons, or because they may need to run internally multiple virtual machines, or many other reasons, which will again, push the limit of the regular home needs, beyond the /56, and consequently, suggesting that /48 is a smarter choice.

The above-mentioned goals of [[RFC3177](#)] could be met by giving home users a default assignment of less than /48, such as a /56, as suggested in [[RFC6177](#)], however, there are new motivations and technologies, to reconsider that a /48 is a much better choice.

#### **4. Impact on IPv6 Standards**

[[RFC6177](#)], regarding [[RFC3056](#)] and [[RFC4193](#)], suggested, in order to facilitate transitioning from the address numbering scheme of those protocols, to one based on a prefix obtained from an ISP, to advise end-sites to number out of the right most bits first, using the leftmost bits only if the size of the site made that necessary.

However, that suggestion fail in several aspects, such as not enforcing actually a standard for devices to be bound on that, increasing security risk against recommendations of using sparse and smart addressing plans, and enforcing renumbering from addressing plans that have been designed before [[RFC6177](#)] was even drafted.

#### **5. Summary**

The exact choice of how much address space to assign end-sites is an issue for the operational community, discussed with much more detail in a recent document [[RIPE-690](#)].

While the recommendation to assign /48s as a default, is not a requirement of the IPv6 architecture and anything of length /64 or shorter works from a standards perspective, 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 to end-sites take into consideration the following:

- a. It should be easy and inexpensive for an end-site to obtain address space to number a sufficiently big number of subnets (i.e., a big number of /64's) and to support reasonable growth projections over long time periods (e.g., more than a decade).
- b. An end-user should be able to receive any prefix length up to /48 simply by asking. It is critical that the community shed the restrictive view of IP addresses that grew up during the end of IPv4. IPv6 addresses should be freely available, not a tiered cost structure.
- c. The default assignment size should take into consideration the likelihood that an end-site will have need for multiple subnets in the near future and many more in a medium and longer terms, avoiding the IPv4 practice of having frequent and continual justification for obtaining small amounts of additional space.
- d. Although a /64 can (in theory) address an almost unlimited number of devices, end-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, NAT, proxy or other techniques. Whether or not those techniques are an appropriate choice is an end-site matter.
- e. 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.
- f. 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.

As a consequence, it is strongly discouraged to assign to end-sites a single /64 or even a reduced number of them.

Instead, it is strongly suggested considering a /48, or alternatively, a trade-off choice is to reserve a /48 for each end-site, and actually assign them only the first /56, so in the future renumbering is not needed and either in a case by case, by demand, or across all the network, the complete /48 can be re-assigned to each end-site.



The considerations of this document are meant mainly for end-sites, regardless of being connected by cellular, wired or other technologies. They don't apply to single cellular devices such as smartphones, which typically will get a single /64 for each connection (PDP context). However, a broadband cellular router connecting an end-site falls within the scope of this document.

## **6. Security Considerations**

This document has no known security implications.

## **7. IANA Considerations**

This document has no actions for IANA.

## **8. Acknowledgements**

The authors of this document will like to acknowledge the authors of previous versions (Thomas Narten and Geoff Huston) and the inputs from .... TBD.

## **9. Informative References**

[IPv6-Survey]

Palet Martinez, J., "IPv6 Deployment Survey (Residential/Household Services)", January 2018, <<https://indico.uknof.org.uk/event/41/contribution/5/material/slides/0.pdf>>.

[RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", [RFC 3056](#), DOI 10.17487/RFC3056, February 2001, <<https://www.rfc-editor.org/info/rfc3056>>.

[RFC3177] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", [RFC 3177](#), DOI 10.17487/RFC3177, September 2001, <<https://www.rfc-editor.org/info/rfc3177>>.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/info/rfc4193>>.

[RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", [BCP 157](#), [RFC 6177](#), DOI 10.17487/RFC6177, March 2011, <<https://www.rfc-editor.org/info/rfc6177>>.





- [RFC7608] Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", [BCP 198](#), [RFC 7608](#), DOI 10.17487/RFC7608, July 2015, <<https://www.rfc-editor.org/info/rfc7608>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", [RFC 7788](#), DOI 10.17487/RFC7788, April 2016, <<https://www.rfc-editor.org/info/rfc7788>>.
- [RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", [BCP 204](#), [RFC 7934](#), DOI 10.17487/RFC7934, July 2016, <<https://www.rfc-editor.org/info/rfc7934>>.
- [RIPE-690] RIPE, "Best Current Operational Practice for Operators: IPv6 prefix assignment for end-users - persistent vs non-persistent, and what size to choose", October 2017, <<https://www.ripe.net/publications/docs/ripe-690>>.
- [Route-Scaling] "Routing and Addressing Problem Statement", February 2010, <<https://tools.ietf.org/html/draft-narten-radir-problem-statement-05>>.

#### Authors' Addresses

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

Phone: +1-650-723-3352  
EMail: [lea.roberts@stanford.edu](mailto:lea.roberts@stanford.edu)

Jordi Palet Martinez  
The IPv6 Company  
Molino de la Navata, 75  
La Navata - Galapagar, Madrid 28420  
Spain

EMail: [jordi.palet@theipv6company.com](mailto:jordi.palet@theipv6company.com)  
URI: <http://www.theipv6company.com/>

