IPv6 Operations
Internet-Draft
Intended status: Best Current Practice
Expires: January 24, 2016

L. Colitti
V. Cerf
Google
S. Cheshire
D. Schinazi
Apple Inc.

July 23, 2015

# Host address availability recommendations
# draft-colitti-v6ops-host-addr-availability-01

## Abstract

This document recommends that networks provide general-purpose end hosts with multiple global addresses when they attach, and describes the benefits of and the options for doing so.

## 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 http://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 January 24, 2016.

## Copyright Notice

Table of Contents



## 1. Introduction

In most aspects, the IPv6 protocol is very similar to IPv4. This similarity can create a tendency to think of IPv6 as 128-bit IPv4, and thus lead network designers and operators to apply identical configurations and operational practices to both. This is generally a good thing because it eases the transition to IPv6 and the operation of dual-stack networks. However, in some areas it can lead to carrying over IPv4 practices that are not appropriate in IPv6 due to significant differences between the protocols.

One such area is IP adressing, particularly IP addressing of hosts. This is substantially different because unlike IPv4 addresses, IPv6 addresses are not a scarce resource. In IPv6, each link has a virtually unlimited amount of address space [RFC7421]. Thus, unlike IPv4, IPv6 networks are not forced by address availability considerations to assign only one address per host. On the other hand, assigning multiple addresses has many benefits including application functionality and simplicity, privacy, future applications, and the ability to deploy the Internet without the use

of NAT. Assigning only one IPv6 address per host negates these benefits.

This document describes the benefits of assigning multiple addresses per host and the problems with not doing so. It recommends that networks provide general-purpose end hosts with multiple global addresses when they attach, and lists current options for doing so.

### 1.1. Requirements Language

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 [RFC2119].

## 2. Common IPv6 deployment model

IPv6 is designed to support multiple addresses, including multiple global addresses, per interface ([RFC4291] section 2.1, [RFC6434] section 5.9.4). Today, many general-purpose IPv6 hosts are configured with three or more addresses per interface: a link-local address, a stable address (e.g., using EUI-64 or [RFC7217]), one or more privacy addresses [RFC4941], and possibly one or more temporary or non-temporary addresses assigned using DHCPv6 [RFC3315].

In most general-purpose IPv6 networks, including all 3GPP networks (see [RFC6459] section 5.2) and Ethernet and Wi-Fi networks using SLAAC [RFC4862], IPv6 hosts have the ability to configure additional IPv6 addresses from the link prefix(es) without explicit requests to the network.

## 3. Benefits of multiple addresses

Today, there are many host functions that require more than one IP address to be available to the host:

- Privacy addressing to prevent tracking by off-network hosts (e.g., [RFC4941]).
- Multiple processors inside the same device. For example, in many mobile devices both the application processor and baseband processor need to communicate with the network, particularly for recent technologies like ePDG.
- Extending the network (e.g., tethering).
- Running virtual machines on hosts.

- Translation-based transition technologies such as 464XLAT that provide IPv4 over IPv6. Current implementations require the availability of a dedicated IPv6 address in order to determine whether inbound packets are translated or native.

- ILA ("Identifier-locator addressing"): https://tools.ietf.org/html/draft-herbert-nvo3-ila

- Future applications (e.g., per-application IPv6 addresses, such as described in [TARP]).

Example of how the availability of multiple addresses per host has already allowed substantial deployment of new applications without explicit requests to the network are:

- 464XLAT [RFC6877]. 464XLAT is usually deployed within a particular network operator's network, but there are deployment models where the PLAT is provided as a service by a different network (e.g., [1])

- /64 sharing [RFC7278]. This was a way to provide IPv6 tethering without needing to wait for network operators to deploy DHCPv6 PD, which is only available in 3GPP release 10.

## 4. Problems with assigning a limited number of addresses per host

Assigning a limited number of addresses per host implies that functions that require multiple addresses will either be unavailable (e.g., if the network provides only one IPv6 address per host, or if the host has reached the limit of the number of addresses available), or that the functions will only be available after an explicit request to the network is granted. The necessity of explicit requests has the following drawbacks:

- Increased latency, because a provisioning operation, and possibly human intervention with an update to the service level agreement, must complete before the functionality is available.

- Uncertainty, because it is not known in advance if a particular operation function will be available.

- Complexity, because implementations need to deal with failures and somehow present them to the user. Failures may manifest as timeouts, which may be slow and frustrating to users.

- Increased load on the network's provisioning servers.

Some operators may desire to configure their networks to limit the number of IPv6 addresses per host. Reasons might include hardware limitations (e.g., TCAM or neighbour cache table size constraints), operational consistency with IPv4 (e.g., an IP address management system that only supports one address per host), or business models (e.g., a desire to charge the network's users on a per-device basis).

## 5. Overcoming limits using Network Address Translation

These limits can mostly be overcome by end hosts by using NAT, and indeed in IPv4 most of these functions are provided by using NAT on the host. Thus, the limits could be overcome in IPv6 as well by implementing NAT66 on the host.

Unfortunately NAT has well-known drawbacks. For example, it causes application complexity due to the need to implement NAT traversal. It hinders development of new applications. On mobile devices, it reduces battery life due to the necessity of frequent keepalives, particularly for UDP. Applications using UDP that need to work on most of the Internet are forced to send keepalives at least every 30 seconds [2]. For example, the QUIC protocol uses a 15-second keepalive [I-D.tsvwg-quic-protocol]. Other drawbacks are described in [RFC2993]. While IPv4 NAT is inevitable due to the limited amount of IPv4 space available, that argument does not apply to IPv6. Guidance from the IAB is that deployment of IPv6 NAT is not desirable [RFC5902].

If networks that provide limited amount of addresses become widely deployed, then the desire to overcome the problems listed in Section 4 without disabling any features may result in operating system manufacturers implementing IPv6 NAT.

This is not a desirable outcome. It is not desirable for users because they may experience application brittleness. It is likely not desirable for network operators either, as they may suffer higher support costs, and even when the decision to assign only one IPv6 address per device is dictated by the network's business model, there may be little in the way of incremental revenue, because devices can share their IPv6 address with other devices. Finally, it is not desirable for operating system manufacturers and application developers, who will have to build more complexity, lengthening development time and/or reducing the time spent on other features.

Indeed, it could be argued that the main reason for deploying IPv6, instead of continuing to scale the Internet using only IPv4 and large-scale NAT44, is because doing so can provide all the hosts on the planet with end-to-end connectivity that is limited not by technical factors but only by security policies.

## 6. Options for obtaining more than one address

Multiple IPv6 addresses can be obtained in the following ways:

- Using Stateless Address Autoconfiguration [RFC4862]. SLAAC allows hosts to create global IPv6 addresses on demand by simply forming new addresses from the global prefix assigned to the link.

- Using stateful DHCPv6 address assignment [RFC3315]. Most DHCPv6 clients only ask for one non-temporary address, but the protocol allows requesting multiple temporary and even multiple non-temporary addresses, and the server could choose to assign the client multiple addresses. It is also possible for a client to request additional addresses using a different DUID. The DHCPv6 server will decide whether to grant or reject the request based on information about the client, including its DUID, MAC address, and so on.

- DHCPv6 prefix delegation [RFC3633]. DHCPv6 PD allows the client to request and be delegated a prefix, from which it can autonomously form other addresses. The prefix can also be hierarchically delegated to downstream clients, or, if it is a /64, it be reshared with downstream clients via ND proxying [RFC4389] or /64 sharing [RFC7278].

| | SLAAC | DHCPv6 IA_NA / IA_TA | DHCPv6 PD | DHCPv4 |
|---|---|---|---|---|
| Autonomously form addresses | Yes (/64 share) | No | Yes (/64 share) | Yes (NAT44) |
| "Unlimited" endpoints | Yes* | Yes* | No | No |
| Stateful, request-based | No | Yes | Yes | Yes |
| Immune to layer 3 on-link resource exhaustion attacks | No | Yes | Yes | Yes |

[*] Subject to network limitations, e.g., ND cache entry size limits.

Table 1: Comparison of multiple address assigment options

## 7. Number of addresses required

If we itemize the use cases from section Section 3, we can estimate the number of addresses currently used in normal operations. In typical implementations, privacy addresses use up to 8 addresses (one per day). Current mobile devices may typically support 8 clients, with each one requiring one or more addresses. A client might choose to run several virtual machines. Current implementations of 464XLAT require use of a separate address. Some devices require another address for their baseband chip. Even a host performing only several of these functions simultaneously might need on the order of 20 addresses at the same time. Future applications designed to use an address per application or even per resource will require many more. These will not function on networks that enforce a hard limit on the number of addresses provided to hosts.

## 8. Recommendations

In order to avoid the problems described above, and preserve the Internet's ability to support new applications that use more than one IPv6 address, it is RECOMMENDED that IPv6 network deployments provide multiple IPv6 addresses from each prefix to general-purpose hosts when they connect to the network. To support future use cases, it is RECOMMENDED to not impose a hard limit on the size of the address pool assigned to a host. If the network requires explicit requests for address space, a /64 prefix is desirable. Using DHCPv6 IA_NA or IA_TA to request a sufficient number of addresses (e.g. 32) would accomodate current clients but sets a limit on the number of addresses available to hosts when they attach and would limit the development of future applications.

## 9. Operational considerations

### 9.1. Stateful addressing and host tracking

Some network operators - often operators of networks that provide services to third parties such as university campus networks - have made the argument that the only feasible IPv6 deployment mechanism is DHCPv6, due to the need to be able to track at all times IPv6 addresses are assigned to which hosts. (Example: [3] ). One reason frequently cited for this is protection from liability for copyright infringement or other illegal activity by maintaining persistent logs that map user IP addresses and timestamps to hardware identifiers such as MAC addresses.

It is worth noting that using DHCPv6 does not by itself ensure that hosts will actually use the addresses assigned to them by the network as opposed to using any other address on the prefix. Such guarantees

can only be provided by link-layer security mechanisms that enforce that particular IPv6 addresses are used by particular link-layer addresses (for example: SAVI [RFC7039]). If those mechanisms are available, it is possible to use them to provide tracking, instead. This form of tracking is much more reliable because it operates independently of how addresses are allocated.

Additionally, attempts to track this sort of information via DHCPv6 are likely to become decreasingly viable due to ongoing efforts to improve the privacy of DHCP: [I-D.ietf-dhc-anonymity-profile].

Many large enterprise networks, including the enterprise networks of the authors, are fully dual-stack but do not currently use or support DHCPv6.

### 9.2. Address space management

In IPv4, all but the world’s largest networks can be addressed using private space [RFC1918], and with each host receiving one IPv4 address. Many networks can be numbered in 192.168.0.0/16 which has roughly 64k addresses. In IPv6, that is equivalent to assigning one /64 per host out of a /48. Under current RIR policies, a /48 is easy to obtain for an enterprise network.

Networks that need a bigger block of private space use 10.0.0.0/8, which is is roughly 16 million addresses. In IPv6, that is equivalent to assigning a /64 per host out of a /40. Enterprises of such size can easily obtain a /40 under current RIR policies.

Currently, residential users receive one IPv4 address and a /48, /56 or /60 IPv6 prefix. While such networks do not have enough space to assign a /64 per device, today such networks almost universally use SLAAC.

Unlike IPv4 where addresses came at a premium, in all these networks, there is enough IPv6 address space to supply clients with multiple IPv6 addresses.

### 9.3. Addressing link layer scalability issues via IP routing

The number of IPv6 addresses on a link has direct impact for networking infrastructure nodes (routers, switches) and other nodes on the link. Setting aside exhaustion attacks via Layer 2 address spoofing, every (Layer 2, IP) address pair impacts networking hardware requirements in terms of memory, MLD snooping, solicited node multicast groups, etc. Many of these same impacts can be felt by neighboring hosts. Switching to a DHCPv6 PD model means there are

only forwarding decisions, with only one routing entry and one ND cache entry per device on the network.

## 10. Acknowledgements

The authors thank Dieter Siegmund, Mark Smith, Sander Steffann, James Woodyatt and Tore Anderson for their input and contributions.

## 11. IANA Considerations

This memo includes no request to IANA.

## 12. Security Considerations

None so far.

## 13. References

### 13.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

### 13.2. Informative References

[I-D.ietf-dhc-anonymity-profile]
Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity profile for DHCP clients", draft-ietf-dhc-anonymity-profile-01 (work in progress), June 2015.

[I-D.tsvwg-quic-protocol]
Jana, J. and I. Swett, "QUIC: A UDP-Based Secure and Reliable Transport for HTTP/2", draft-tsvwg-quic-protocol-01 (work in progress), July 2015.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <http://www.rfc-editor.org/info/rfc1918>.

[RFC2993] Hain, T., "Architectural Implications of NAT", RFC 2993, DOI 10.17487/RFC2993, November 2000, <http://www.rfc-editor.org/info/rfc2993>.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389, DOI 10.17487/RFC4389, April 2006, <http://www.rfc-editor.org/info/rfc4389>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.

[RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007.

[RFC5902] Thaler, D., Zhang, L., and G. Lebovitz, "IAB Thoughts on IPv6 Network Address Translation", RFC 5902, July 2010.

[RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, DOI 10.17487/RFC6434, December 2011, <http://www.rfc-editor.org/info/rfc6434>.

[RFC6459] Korhonen, J., Soininen, J., Patil, B., Savolainen, T., Bajko, G., and K. Iisakkila, "IPv6 in 3rd Generation Partnership Project (3GPP) Evolved Packet System (EPS)", RFC 6459, January 2012.

[RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, April 2013.

[RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <http://www.rfc-editor.org/info/rfc7039>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, April 2014.

[RFC7278] Byrne, C., Drown, D., and A. Vizdal, "Extending an IPv6 /64 Prefix from a Third Generation Partnership Project (3GPP) Mobile Interface to a LAN Link", RFC 7278, June 2014.

[RFC7421]  Carpenter, B., Ed., Chown, T., Gont, F., Jiang, S., Petrescu, A., and A. Yourtchenko, "Analysis of the 64-bit Boundary in IPv6 Addressing", RFC 7421, DOI 10.17487/RFC7421, January 2015, <http://www.rfc-editor.org/info/rfc7421>.

[TARP]  Gleitz, PM. and SM. Bellovin, "Transient Addressing for Related Processes: Improved Firewalling by Using IPv6 and Multiple Addresses per Host", August 2001.

## 13.3. URIs

[1] http://www.jpix.ad.jp/en/service/ipv6v4.html

[2] http://www.ietf.org/proceedings/88/slides/slides-88-tsvarea-10.pdf

[3] https://code.google.com/p/android/issues/detail?id=32621#c60

## Authors' Addresses

Lorenzo Colitti
Google
Roppongi 6-10-1
Minato, Tokyo  106-6126
JP

Email: lorenzo@google.com

Vint Cerf
Google
1600 Amphitheatre Parkway
Mountain View, CA  94043
US

Email: vint@google.com

Stuart Cheshire
Apple Inc.
1 Infinite Loop
Cupertino, CA  95014
US

Email: cheshire@apple.com

David Schinazi
Apple Inc.
1 Infinite Loop
Cupertino, CA  95014
US

Email: dschinazi@apple.com