

6MAN Working Group
Internet-Draft
Intended status: Informational
Expires: May 2, 2021

G. Mishra
Verizon Inc.
A. Petrescu
CEA, LIST
N. Kottapalli
Benu Networks
D. Mudric
Ciena
D. Shyti
SFR
October 29, 2020

**SLAAC with prefixes of arbitrary length in PIO (Variable SLAAC)
draft-mishra-v6ops-variable-slaac-problem-stmt-00**

Abstract

In the past, various IPv6 addressing models have been proposed based on a subnet hierarchy embedding a 64-bit prefix. The last remnant of IPv6 classful addressing is a inflexible interface identifier boundary at /64. This document details the 64-bit boundary problem statement.

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 May 2, 2021.

Copyright Notice

Copyright (c) 2020 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

- [1. Terminology](#) [2](#)
- [2. Introduction](#) [2](#)
- [3. Problem Statement](#) [3](#)
- [4. Variable SLAAC Use Cases](#) [5](#)
 - [4.1. Permission-less Extension of the Network](#) [5](#)
 - [4.2. Private Networks](#) [6](#)
 - [4.3. Mobile IPv6](#) [6](#)
 - [4.4. Home and SOHO](#) [7](#)
 - [4.5. 3GPP V2I and V2V networking](#) [7](#)
 - [4.6. Smart Traffic Lights](#) [8](#)
 - [4.7. 6lo](#) [8](#)
 - [4.8. Large ISP's backbone POP](#) [9](#)
 - [4.9. Permission-less extension of the Network](#) [9](#)
- [5. Security Considerations](#) [9](#)
- [6. IANA Considerations](#) [9](#)
- [7. Contributors](#) [10](#)
- [8. Acknowledgements](#) [10](#)
- [9. References](#) [10](#)
 - [9.1. Normative References](#) [10](#)
 - [9.2. Informative References](#) [11](#)
- [Appendix A. ChangeLog](#) [11](#)
- [Authors' Addresses](#) [11](#)

1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2. Introduction

From the beginning, the IPv6 addressing plan was based on a 128 bit address format made up of 8 hexets which were broken down into a 64 bit four hexet prefix and 64 bit four hexet interface identifier. For example, the address 2001:db8:3:4::1 has the first 4 hexets forming the /64 prefix 2001:db8:3:4::/64, whereas the last four

hexets form an interface identifier abbreviated as `::1` (a 'hexet' is a group of max 4 hex digits between two columns, e.g. "2001" and "db8" are each a hexet). A comprehensive analysis of the 64-bit boundary is provided in [[RFC7421](#)]. The history of IPv6 Classful models proposed, and the last remnant of IPv6 Classful addressing rigid network interface identifier boundary at /64 is discussed in detail as well as the removal of the fixed position of the boundary for interface addressing in draft [[I-D.bourbaki-6man-classless-ipv6](#)].

This document discusses the reasons why the interface identifier has been fixed at 64 bits, and the problems that can be addressed by changing the GUA interface identifier from fixed 64 bit size to a variable interface identifier. This change would be consistent with static and DHCPv6 stateful IPv6 address assignment. This document tries to achieve clearing the confusion related to prefix length, and provide consistency of variable length prefix across the three IPv6 addressing strategies deployed, static, DHCPv6 and Stateless Address Autoconfiguration(SLAAC), and finally update all RFCs with the new variable SLAAC standard.

Over the years one of the merits of increasing the prefix length, and reducing the size of the interface identifier has been incorrectly stated as the possibility of IPv6 address space exhaustion could be circumvented, or that a 64 bit interface identifier is an efficient use of address space.

3. Problem Statement

This section details the problem statement as to what is broken today with fixed length Stateless Address Autoconfiguration SLAAC [[RFC4862](#)] and why it is critical to resolve this problem.

The main problem is that SLAAC RA or PD allocates a /64 by the wireless carrier 4G, 5G, 3GPP to mobile handset or hotspot, however segmentation of the /64 via SLAAC is required so that downstream interfaces can be further sub-netted. The use case section of this draft discusses this scenario as one of the use cases for shorter interface identifier, and this use case is the only one stated here in the problem statement as this is broken today with the current SLAAC specification [[RFC4862](#)], and there is not any workaround for this use case.

There are two reasons why this was not a problem in the past, but now with increased bandwidth there are more and more devices being piled onto a single handset or mobile hotspot. In the past generations of cellular systems (e.g. 2.5G aka GPRS and some 3G) the bandwidth available to the User Equipment was not enough to accommodate several applications; bandwidth available was roughly 256Kbit/s. For that

reason, users were rarely tempted to use an UE to link other devices than that UE to the Internet. However, with the arrival of 3G, 3G+ (e.g. HSDPA / HSUPA), and even more so with 4G and 4G+, the bandwidth made available to UE increased significantly; this became an average effective of 1Mbit/s and even more. With this available bandwidth, the users are more and more tempted to connect several devices to the Internet. This operation is named 'connection sharing' or 'tethering'. Another answer to this question is that IPv6 technology that is widely used to 'tether' several IP devices to a smartphone is '64share' [RFC7278](#). This technology is used for smartphones but is not so in vehicles. One of the reasons of not being used in vehicles is the lack of scalability: a /64 prefix is shared between the UE ptp link and the subnet (typically Wi-Fi), but can not be further sub-netted to other subnets in the car.

The reason why all devices in a car cannot remain on a single /64 are as follows. These devices have different link-layer technologies, and not all WiFi could be bridged into Ethernet such as to keep all devices into one /64. They could be on links that are not bridgeable: devices on 802.11-OCB cannot be bridged, devices on Bluetooth cannot be bridged, devices on 3GPP cannot be bridged, and so on. Other than the impossibility to bridge several such link-layer technologies there is also a problem of noise: in a vehicle one wants the braking pedal signal to not be disturbed by entertainment sites such as YouTube. That physical technical requirement separation of different link layer technologies segmentation on to different smaller IPv6 subnets cannot be achieved if all devices are on a single /64, or bridged. Therefore, the only possible solution to connect these disparate devices onto a 3GPP network for internet access is to keep these separate link layer technologies segmented onto separate greater than /64 prefix subnets and breaking the /64 boundary that exists today with a Variable SLAAC solution. Thus, when the 3GPP network gives a /64 to the car, and when there are unbridgeable technologies in the car (e.g. WiFi cannot be bridged to Bluetooth), then the only possibility is to divide that /64 into two /65s. One /65 would be used on the WiFi and another /65 would be used on Bluetooth. But in order for SLAAC to work with /65 then there is a need to have the shorter interface identifier of length 63. Hence the need of lengths of PIOs other than 64 (variable plen).

There are three scenarios that require SLAAC to be able to be routed between two greater than /64 prefix segments as part of the requirement for variable length SLAAC and what is broken with the current SLAAC specification defined in [\[RFC4862\]](#).

The first scenario is within a car using car manufacturer single SIM for internet access and being able to bridge(Route) other link layer devices like BT via variable slaac. In this scenario the

communication between downstream devices are all located within the car using the car manufacturer built in SIM card for in-vehicle communication. The in-vehicle scenario covers both the built-in car manufacturer SIM card scenario, or if the car manufacturer does not support built-in SIM card then a single mobile handset providing 3GPP internet access to all devices in the car.

The second scenario is V2V (vehicle to vehicle) between cars requiring SLAAC to subnet the >64 prefix so that the two cars have WiFi connectivity.

This third scenario is a uCPE(Universal Customer Premises Equipment) device is LTE 4G and Wi-Fi capable, and utilizes NFV (Network Function Virtualization) framework, providing SFC (Service Function Chaining), where one VNF (Virtual Network Function) is a CPE Layer 3 router and is the uCPE device which will receive a /64 prefix from 4G 3GPP Wireless provider and would like to be able to provide further segmentation. In order to provide further segmentation and subdivide the /64 into smaller longer prefix subnets variable SLAAC must be employed. In this example we would give 1st /66 to Wi-Fi users, 2nd /66 to Wired connected network device without security, 3rd /66 prefix to VNF firewall instance, and 4th /66 prefix VNF load balancer instance. The uCPE (Universal Customer Premises Equipment) defined in draft [[I-D.shytyi-opsawg-vysm](#)].

From a segmented bandwidth perspective while breaking up the /64 subnet into smaller subnets, there is not any impact to the user experience of the now shared bandwidth, as long as the cellular signal has adequate enough bars as far as signal strength to accommodate the now multiple devices sharing the single cellular signal. These scenarios described above are the problems that can only be solved with a variable SLAAC solution. There is no other solution or workaround for this problem.

4. Variable SLAAC Use Cases

This section describes real world use cases of variable slaac that cannot be done today and with fixed 64 bit prefix lengths.

4.1. Permission-less Extension of the Network

Permission-less extensions of the network with new links (and by implication with new routers) are not supported.

The lack of possibility to realize a permission-less extension of the network is an important problem, which appears at the edge of the network. The permission is 'granted' for end users situated at the edge of the network, and is 'granted' by advertising a prefix of

length 64 inside the PIO option in a RA typically. The end user receives this prefix, forms an address, and is able to connect to the internet. However, the end user has no permission to further extend the network. Although the device is able to form subsequent prefixes of a length of, say 65, and further advertise it down in the extension of the network, no other Host in that extension of the network is able to use that advertisement; a Host cannot form an address with a prefix length 65 by using SLAAC. The Linux error text reported in the kernel log upon reception of a plen 65 is "illegal" (or similar).

4.2. Private Networks

Private networks such as Service Provider core not accessible by customers and enterprises where all hosts are trusted are the primary use case for variable SLAAC as the shorter interface identifier does not create any security issues with not having a longer 64 bit interface identifier for privacy extensions stable interface identifier [[RFC8084](#)] due to all hosts being inherently trusted. Private internal networks such as corporate intranets traditionally have always used static IPv6 addressing for infrastructure. This manual IPv6 address assignment process for network infrastructure links can take long lead times to complete deployment. By changing the behavior of SLAAC to support variable length prefix and interface identifier allows SLAAC to be used programmatically to deploy to large scale IPv6 networks with thousands of point-to-point links. Note that network infrastructure technically does not require IPv6 addressing due to IPv6 next hop being a link local address for IGP routing protocols such as OSPF and ISIS as well as the link local address can be the peer IPv6 address for exterior gateway routing protocols such as BGP. However for hop by hop ping and traceroute capability to have IPv6 reachability at each hop for troubleshooting jitter, latency and drops it is an IPv6 recommended best practice to configure IPv6 address on all infrastructure interfaces.

4.3. Mobile IPv6

Old MIP6 (Mobile IPv6) Working Group and old Nemo Working Group's routing solution scenarios related to Mobile IPv6 ([\[RFC3775\]](#)) (note: nowadays most MIP-related activity is in DMM WG) where the mobile endpoint can now obtain from the home agent variable SLAAC address and not 64 bit prefix /64 address. This maybe useful in cases where a /64 can now be managed from an addressing perspective and subdivided into blocks for manageability of MIP6 endpoints instead of allocating a single /64 per endpoint.

4.4. Home and SOHO

Home and SOHO (Small Office and Home Office) environments where internet access uses a broadband service provider single or dual homed scenario. In those such Home networking Homenet environments where HNCP (Home Network Control Protocol [[RFC7788](#)] SADR (Source Address Dependent Routing) are deployed for automatic configuration for LAN Wi-Fi endpoint subnets can also now take advantage of variable length SLAAC in deployment scenarios. In cases where multiple routers are deployed in a home environment where routing prefix reachability needs to be advertised where Babel [[RFC6126](#)] routing protocol is utilized in those cases variable SLAAC can also be utilized to break up a /64 into multiple smaller subnets.

4.5. 3GPP V2I and V2V networking

In V2I networking (with 3GPP or with IEEE 802.11bd) the IP-OBU in the vehicle receives a /64 prefix from the cellular network (or from a IP-RSU - Road-Side Unit). This /64 prefix can be used to form one address for the egress interface of the Mobile Router (MR, which is also termed 'IP-OBU', for IP On-Board Unit, in IPWAVE WG documents such as [RFC8691](#)), but can not be used to form IP addresses for other hosts in the vehicle. In the following two paragraphs we explain this problem.

In certain 3GPP V2I networking use cases a /56 is allocated by the 3GPP infrastructure to the 4G modem of the IP-OBU in the vehicle. In such use case it is possible that the IP-OBU sub-divides the allocated /56 into multiple 'result' /64 prefixes. Such a 'result' /64 prefix could be used to form addresses for deeper subnets in the vehicle, by employing existing SLAAC and existing IPv6-over-foo specifications of Interface ID.

If in other 3GPP V2I networking use-cases the infrastructure does not allocate a /56 (or 'longer' prefix lengths such as a /57, /58.. /63) to the IP-OBU, i.e. a /64 is allocated to the IP-OBU, then the 'result' prefix obtained after a sub-divide operation can only be of length /65, or /66, or longer. A prefix of such length (longer than 64) can not be used with SLAAC and existing IP-over-foo Interface Identifiers, because the length of all Interface Identifiers in all IPv6-over-foo documents must always be 64, and the length of the IPv6 is always 128bit. The 64bit of an IID added to the 65bit (or more) of a prefix is larger than 128bit. It is for this reason that a SLAAC with other than 64bit Interface IDs (hence a 'Variable Prefix Length SLAAC') is needed.

The problem of /64 allocation to the vehicle is mostly present in V2I use-cases. In V2V use-cases this problem is less apparent but

deserves consideration. Until now there was no clearcut design and decision about the infrastructure allocating addresses to several vehicles (just to one, in V2I, see above). In some use-cases, the prefix allocated to one vehicle could be further extended by that vehicle to delegate prefixes to other vehicles nearby which might not have 3GPP connections, but only 802.11-OCB interfaces. In such cases it is again necessary that a /64 allocated by the infrastructure to the first vehicle be further sub-divided in multiple 'result' longer-than-/64 prefixes; and one of these longer-than-64 prefixes might be used for the second vehicle (instead of being used for the internal subnets of the first vehicle); this latter vehicle will need to use a form of SLAAC and IP-over-foo that are not limited by the /64 limit.

4.6. Smart Traffic Lights

Smart traffic lights are traffic lights equipped with a communication system. Smart traffic lights are deployed at intersections of roads and serve the purpose of safely arbitrating the passage of automobiles, pedestrians and cyclists. A typical smart traffic lights setting is made of several computers, included but not limited to: a traffic lights controller, a power controller and a communication gateway. More advanced smart traffic lights are equipped with more computers for radars, detection loops, lidars, V2X wireless capabilities, Wi-Fi, Bluetooth and cellular 4G or 5G. All these computers need to use IP addresses: at least one IP address per computer. Since smart traffic lights are deployed in areas where Internet might not be available by cable, fibre or other Wireless MAN technology the only way to connect all computers in the smart traffic lights setting is to employ a 4G (or 5G) gateway. This gateway obtains typically a /64 prefix from the network operator; there is a problem in subdividing that /64 prefix into smaller prefixes, because the obtained prefixes can not be used by SLAAC, because SLAAC uses Interface IDs of length 64 in practice. Even if the SLAAC specification is independent of the prefix length, the length of the Interface ID dictates the prefix length by side effect (128 minus IID length imposes the prefix length). SLAAC might work with a plen 65 by specification, but all IIDs in all IPv6-over-foo request that IIDs be 64; and the sum of IID len plus plen must be 128.

4.7. 6lo

6lo Working IPv6 over Network Constrained nodes working group use cases. Use cases for IoT devices where have limited network access requirements could now take advantage of variable SLAAC longer prefixes lengths /65-/128.

4.8. Large ISP's backbone POP

Large ISP backbone POPs such as IXPs where many carriers share the same backbone and ND cache exhaustion may occur due to /64 subnet size. One mitigation technique employed is the use of an ARP Sponge for IPv4 or Layer 2 multicast rate limiters for IPv6. In those particular cases a longer prefix static or variable SLAAC subnet could be utilized to reduce the maximum number of hosts on the subnet.

4.9. Permission-less extension of the Network

When one wants to extend the network, one typically wants to add new computers to it. Currently, there are two ways to achieve it: (1) ask the network administrator to provide addresses while also inserting a route towards the new subnet of devices and (2) use NAT. With IPv6, NAT is not desirable. In order to extend the network without asking for permission one needs to obtain addresses and to obtain that route inserted. In order to obtain addresses, one might take advantage of the /64 prefix typically advertised by the network to an edge of it. To do that, one needs to sub-divide the /64 prefix into /65 sub-prefixes (or longer, such as /66, /67, etc.) which could be further advertised in the extension of the network. For the action of inserting a route, the particular topic is outside the scope of this document.

5. Security Considerations

The administrator should be aware to maintain 64 bit interface identifier for privacy when connected directly to the internet so that entropy for optimal heuristics are maintained for security.

Variable length interface identifier shorter than 64 bits should be only used within corporate intranets and private networks where all hosts are trusted.

In all cases where the host is on a public network for privacy concerns to avoid traceability variable interface identifier MUST never be utilized.

6. IANA Considerations

IANA is requested to assign the new Router Advertisement flag defined in [Section 5](#) of this document. Bit 6 is the next available bit in this registry, IANA is requested to use this bit unless there is a reason to use another bit in this registry.

IANA is also requested to register this new flag bit in the IANA IPv6 ND Router Advertisement flags Registry [IANA-RF].

7. Contributors

Contributors.

8. Acknowledgements

Ole Troean.

9. References

9.1. Normative References

[I-D.bourbaki-6man-classless-ipv6]

Bourbaki, N., "IPv6 is Classless", [draft-bourbaki-6man-classless-ipv6-05](#) (work in progress), April 2019.

[I-D.shytyi-opsawg-vysm]

Shytyi, D., Beylier, L., and L. Iannone, "A YANG Module for uCPE management.", [draft-shytyi-opsawg-vysm-08](#) (work in progress), May 2020.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), DOI 10.17487/RFC3775, June 2004, <<https://www.rfc-editor.org/info/rfc3775>>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.

[RFC6126] Chroboczek, J., "The Babel Routing Protocol", [RFC 6126](#), DOI 10.17487/RFC6126, April 2011, <<https://www.rfc-editor.org/info/rfc6126>>.

[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>>.

- [RFC8084] Fairhurst, G., "Network Transport Circuit Breakers", [BCP 208](#), [RFC 8084](#), DOI 10.17487/RFC8084, March 2017, <<https://www.rfc-editor.org/info/rfc8084>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[9.2.](#) Informative References

- [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, <<https://www.rfc-editor.org/info/rfc7421>>.

[Appendix A.](#) ChangeLog

The changes are listed in reverse chronological order, most recent changes appearing at the top of the list.

-00: initial version.

Authors' Addresses

Gyan Mishra
Verizon Inc.

Email: gyan.s.mishra@verizon.com

Alexandre Petrescu
CEA, LIST
CEA Saclay
Gif-sur-Yvette, Ile-de-France 91190
France

Phone: +33169089223

Email: Alexandre.Petrescu@cea.fr

Naveen Kottapalli
Benu Networks
300 Concord Road
Billerica MA 01821
United States of America

Phone: +1 978 223 4700
Email: nkottapalli@benu.net

Dusan Mudric
Ciena
Canada

Phone: +1-613-670-2425
Email: dmudric@ciena.com

Dmytro Shytyi
SFR
Paris
France

Email: dmytro.shytyi@sfr.com

