SLAAC with prefixes of arbitrary length in PIO

(Variable SLAAC) – A Problem Statement

V6ops:
draft-mishra-v6ops-variable-slaac-problem-stmt-01
6man:
draft-mishra-6man-variable-slaac-01

Gyan Mishra (Verizon), speaker
A. Petrescu (CEA)
N. Kottapalli (Benu Networks)
D. Murdic (Cienna)
D. Shytyi (SFR)

IETF-109, November 2020

Extending /64

There are many potential solutions:

a1) ask the network operator for more address space.
a2) change provider
a3) introduce government regulation
b1) steal the uplink /64 (64share)
b2) steal multiple /64s from uplink
c) overlay. use e.g. LISP to tunnel across the access ISP to connect to an ISP that support multi-homing and larger address space.
d) MultiLink Subnet Routing. I.e. let a single /64 span multiple links. draft-thubert-6man-ipv6-over-wireless, draft-ietf-ipv6-multilink-subnets
e) NAT
f) P2P Ethernet. Hosts are not on the same physical link, so let's stop pretending they are. A consequence of that is that links don't need subnets. Only assign addresses to hosts. draft-troan-6man-p2p-ethernet-00
g) extend the /64 bit boundary. HNCP implementations do /80s I think (forces DHCP for address assignment)
h) Variable SLAAC (openbsd and linux implementations)
i) Mobile IP with NEMO extensions.
j) IP encapsulation and a VPN gateway

Requirements:
R-1: Permissionless. Not require an action on the network operator
R-2: Arbitrary topology
R-3: Long-lived address assignments
R-4: Support bad operational practice: flash renumbering / ephemeral addressing
I would like to suggest a requirement R-5: do not use encapsulation and do not use VPN gateways or Home Agent services. The reason is twofold: these rdv points represent additional single points of failure. The second reason comes from an observation of current work situation: very often in these electronic virtual meetings the involvement of VPN gateways induces latencies and breaks audio, or reduces its quality.

Use-case:
3GPP V2I and V2V networking

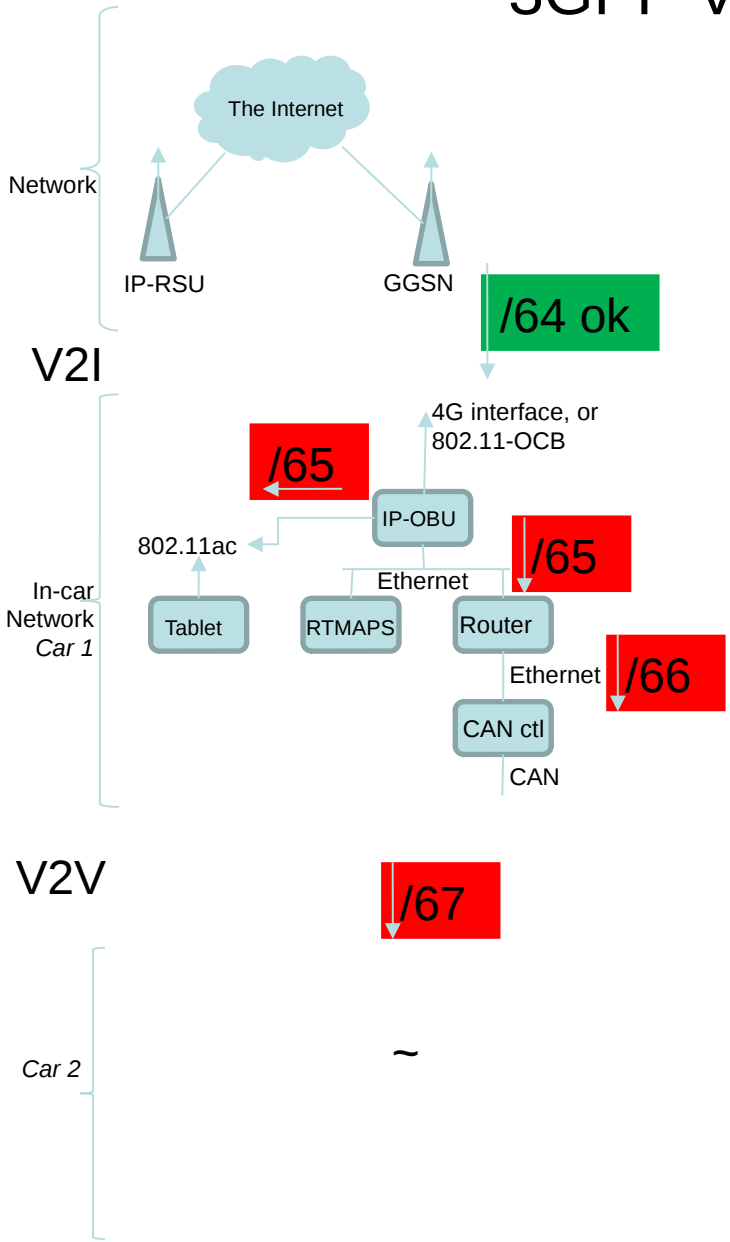

- In V2I networking the IP-OBU in the vehicle receives a /64 prefix from the cellular network. This /64 prefix can be used to form one address for the egress interface of the IP-OBU (RFC8691), but can not be used to form IP addresses for other hosts in the vehicle.
- In V2V, that /64 needs to be further extended to vehicles nearby.
- A prefix of length longer than 64 can not be used with SLAAC because the length of all Interface Identifiers must always be 64, and the length of the IPv6 is always 128bit.
- A SLAAC with other than 64bit Interface IDs is needed: a 'Variable Prefix Length SLAAC'.

64share-V2 – an interpretation and discussion

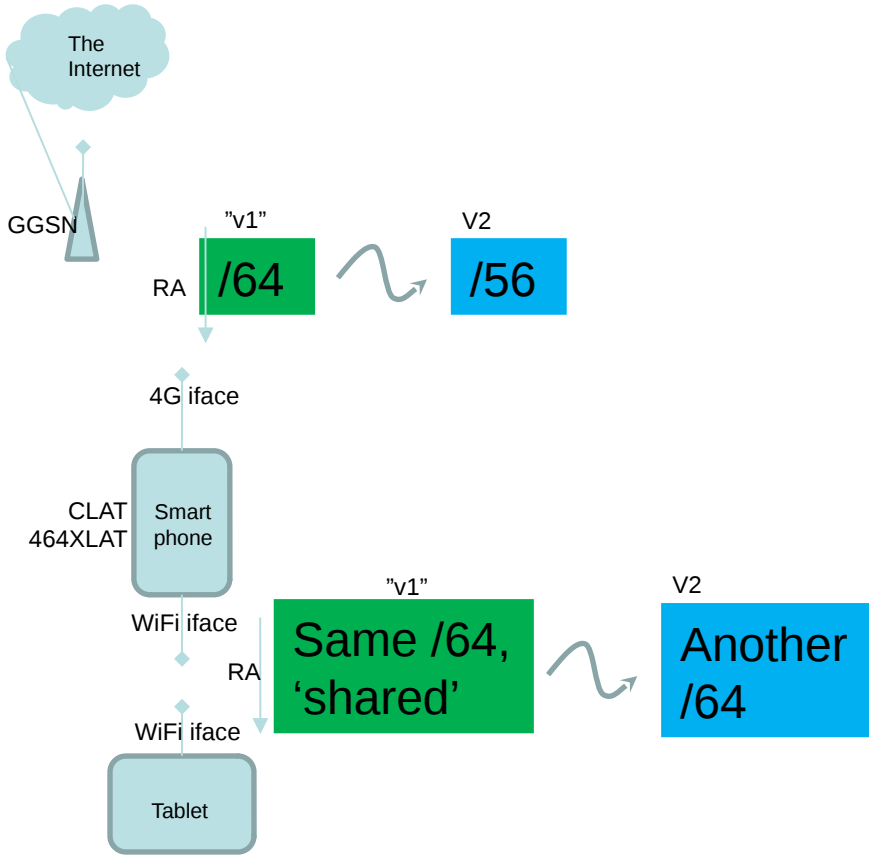

- “64share” (v1) is in RFC7278.
- 64share-V2 is in
  draft-byrne-v6ops-64sharev2-00.txt
  at https://pastebin.com/duyYRkzG
- “This memo requests the 3GPP to change this requirement to allow any prefix size less than or equal to 64 be advertised by the 3GPP gateway RA. This change allows the UE to be given a prefix such as a /56 using RA, which is consider sufficient for a home network.”
- A Question:
  - Probably the ‘A’ flag in RA is not set (not autonomous) in V2. Probably some magic would allow the smartphone to form an address even if A reset. Could same magic be able to be used in “v1”?
- Other Remarks:
  - It probably needs to allow the smartphone to form an IPv6 address starting with SLAAC from a /56, on its 4G interface. Probably needs an IID of length 72.
  - The V2 is much of a trick as “v1” is, in that it delegates a prefix to the smartphone, instead of assigning it on the link.
  - It has the advantage of not requiring new software on the GGSN (no new bits in RA, no new protocol), but just modify a configuration file in GGSN.

Internet Engineering Task Force (IETF)
Request for Comments: 6164
Category: Standards Track
ISSN: 2070-1721

M. Kohno
Juniper Networks, Keio University
B. Nitzan
Juniper Networks
R. Bush
Y. Matsuzaki
Internet Initiative Japan
L. Colitti
Google
T. Narten
IBM Corporation

April 2011

Using 127-Bit IPv6 Prefixes on Inter-Router Links

Abstract

On inter-router point-to-point links, it is useful, for security and other reasons, to use 127-bit IPv6 prefixes. Such a practice parallels the use of 31-bit prefixes in IPv4. This document specifies the motivation for, and usages of, 127-bit IPv6 prefix lengths on inter-router point-to-point links.

Status of This Memo

This is an Internet Standards Track document.

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 Internet Standards is available in Section 2 of RFC 5741.

5.2. Neighbor Cache Exhaustion Issue

As described in Section 4.3.2 of [RFC3756], the use of a 64-bit prefix length on an inter-router link that uses Neighbor Discovery (e.g., Ethernet) potentially allows for denial-of-service attacks on the routers on the link.

Consider an Ethernet link between two routers, A and B, to which a /64 subnet has been assigned. A packet sent to any address on the /64 (except the addresses of A and B) will cause the router attempting to forward it to create a new cache entry in INCOMPLETE state, send a Neighbor Solicitation message on the link, start a retransmit timer, and so on [RFC4861].

By sending a continuous stream of packets to a large number of the 2^64 - 3 unassigned addresses on the link (one for each router and one for Subnet-Router anycast), an attacker can create a large number of neighbor cache entries and cause one of the routers to send a large number of Neighbor Solicitation packets that will never receive

replies, thereby consuming large amounts of memory and processing resources. Sending the packets to one of the 2^24 addresses on the link that has the same Solicited-Node multicast address as one of the routers also causes the victim to spend large amounts of processing time discarding useless Neighbor Solicitation messages.

replies, thereby consuming large amounts of memory and processing resources. Sending the packets to one of the 2^24 addresses on the link that has the same Solicited-Node multicast address as one of the routers also causes the victim to spend large amounts of processing time discarding useless Neighbor Solicitation messages.

Careful implementation and rate-limiting can limit the impact of such an attack, but are unlikely to neutralize it completely. Rate-limiting Neighbor Solicitation messages will reduce CPU usage, and following the garbage-collection recommendations in [RFC4861] will maintain reachability, but if the link is down and neighbor cache entries have expired while the attack is ongoing, legitimate traffic (for example, BGP sessions) over the link might never be re-established, because the routers cannot resolve each others' IPv6 addresses to link-layer addresses.

This attack is not specific to point-to-point links, but is particularly harmful in the case of point-to-point backbone links, which may carry large amounts of traffic to many destinations over long distances.

While there are a number of ways to mitigate this kind of issue, assigning /127 subnets eliminates it completely.

5.3. Other Reasons

Though address space conservation considerations are less important for IPv6 than they are in IPv4, some operators prefer not to assign /64s to individual point-to-point links. Instead, they may be able to number all of their point-to-point links out of a single /64 or a small number of /64s.

Internet Engineering Task Force (IETF)
Request for Comments: 7381
Category: Informational
ISSN: 2070-1721

K. Chittimaneni
Dropbox, Inc.
T. Chown
University of Southampton
L. Howard
Time Warner Cable
V. Kuarsingh
Dyn, Inc.
Y. Pouffary
Hewlett Packard
E. Vyncke
Cisco Systems

October 2014

Enterprise IPv6 Deployment Guidelines

Abstract

Enterprise network administrators worldwide are in various stages of preparing for or deploying IPv6 into their networks. The administrators face different challenges than operators of Internet access providers and have reasons for different priorities. The overall problem for many administrators will be to offer Internet-facing services over IPv6 while continuing to support IPv4, and while introducing IPv6 access within the enterprise IT network. The overall transition will take most networks from an IPv4-only environment to a dual-stack network environment and eventually an IPv6-only operating mode. This document helps provide a framework for enterprise network architects or administrators who may be faced with many of these challenges as they consider their IPv6 support strategies.

In the data center or server room, assume a /64 per VLAN. This applies even if each individual system is on a separate VLAN. In a /48 assignment, typical for a site, there are then still 65,535 /64 blocks. Some administrators reserve a /64 but configure a small subnet, such as /112, /126, or /127, to prevent rogue devices from attaching and getting numbers; an alternative is to monitor traffic for surprising addresses or Neighbor Discovery (ND) tables for new entries. Addresses are either configured manually on the server or reserved on a DHCPv6 server, which may also synchronize forward and reverse DNS (though see [RFC6866] for considerations on static addressing). SLAAC is not recommended for servers because of the need to synchronize RA timers with DNS Times to Live (TTLs) so that the DNS entry expires at the same time as the address.

All user access networks should be a /64. Point-to-point links where NDP is not used may also utilize a /127 (see [RFC6164]).

“Race to the Bottom”

One of the reasons why 6MAN has deferred removing the 64 bit boundary is due to the ISP “race to the bottom” fear that we are at the bottom giving out /64 to mobile handsets that ISP’s will in theory do what history has shown us what was done with IPv4 where due to address depletion issues and issues with overlapping address space ISP’s made the broadband standard to dole out /32 WAN IP which is NAT port overloaded outside wan interface. NAT as well as CGNAT have solved issues with IPv4 shortage and overlapping ranges allowing overlapping ranges to co-exist with NAT as well as now ISP’s due to the risk of IPv4 address depletion made the standard a /32 WAN IP with NAT port overloaded via PAT (Port Address translation) with private 192.168.1.0/24 subnet for SOHO hosts.

IPv6 on the other hand does not have any risk of address depletion so you cannot compare what history has told us with IPv4 to IPv6 and there is no other data point to be had.

On the other side of the spectrum with 64share we are looking now at shorter prefixes and maybe an idea of creating and RFC 6177bis that allows a /48 per human per mobile device. Is that possible ???

https://www.iana.org/assignments/ipv6-address-space/ipv6-address-space.xhtml

2001::/3 GUI

IAIA RIR ISP End user Allocation

RIR allocations to Service Providers:

For the massive block allocations for ISPs is it tiered like this
/32 - General Starting point for ISP allocation
/24-26 - Medium to Large
/19-20 - Exceptions and rare

Since allocations are done blockwize and not linearly you have to account for future growth so generally over allocate for growth next 20 years. So with that the much larger allocations.

With that being said let's say for a typical large ISP /24 - that would yield 16M /48s. That still is small since most large ISPs like, we Verizon we want to scale to a billion as right now we have 150M and that's just domestic not worldwide. Safe best for large ISPs is we want to scale to the number of humans on the planet so 7 Billion is a good number. Also that does not account for broadband. So for Verizon as an example /48 per human is not possible.

I think /48 per user site is sensible for the much smaller ISPS with few million subscriber base but I think an impossibility at least for the larger Verizon size ISP’s. Looking on IAIA allocation link I don't see too many RIRs getting a /12.

https://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xhtml

IPv6 Addressing options - Static, DHCPv6, SLAAC
Typical addressing as it works today

Static IPv6 Addressing Defacto Standard by operators

- Loopback = /128
- P2P = /127 – RFC 3627
- >2 host VLSM subnet masked /64-/125

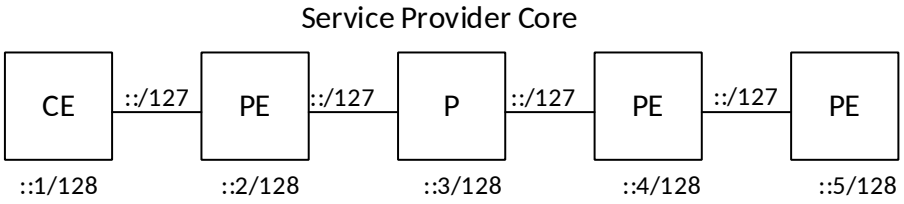

For security reasons all subnets are sized based on maximum number of hosts
For security reasons to create a ND Cache hard limit
To avoid ND Cache exhaustion.

DHCPv6 Server

/64 - Pool

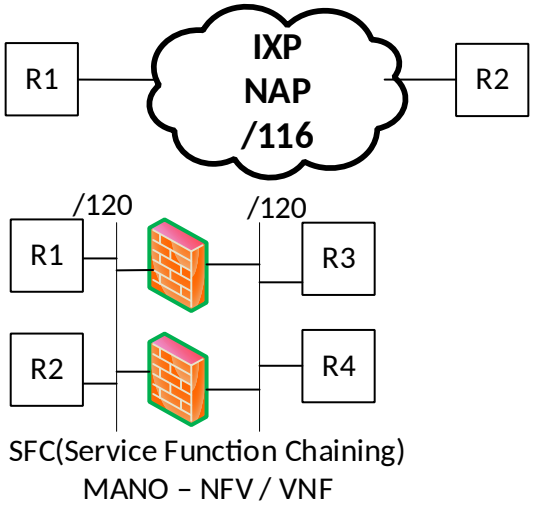

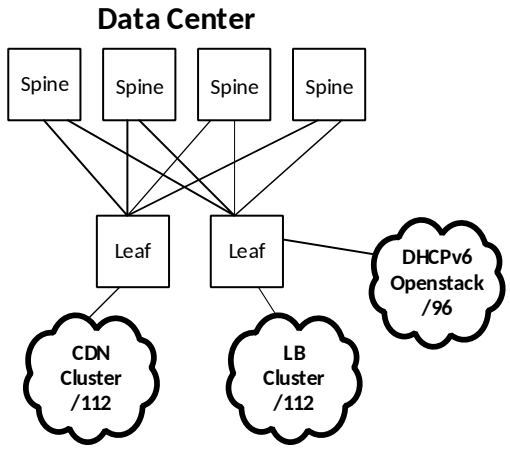

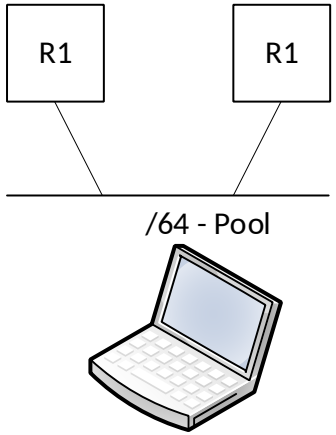

SLAAC interoperability issue with DHCPv6 & Static Problem Statement

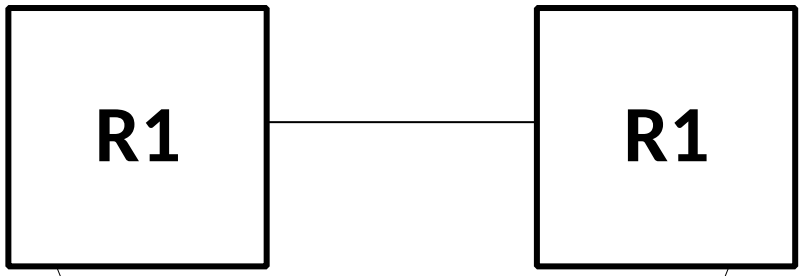

DHCPv6 Server

RFC 4291 & RFC 4862 only support /64 for SLAAC due to modified EUI64 only supporting /64 fixed boundary

Random Station ID RFC 4941 Privacy Extension & Stable IID RFC 7217

/112 Subnet

DHCPv6 /112

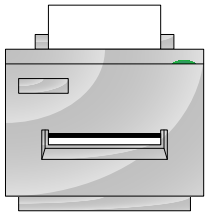

This device only supports /64

Static /112

DHCPv6 /112