

April 2024
Geoff Huston,
Nathan Ward

IPv6 Prefix Lengths

The topic of address plans for IPv6 has had a rich and varied history. From the very early concepts of “it’s just like IPv4, only with a 128-bit address field”, through the models of “Aggregation Identifiers” and the hierarchy of “Top-Level,” “Next-Level” and “Site-Level” defined in RFC 2373 from July 1998 and then the simplified adoption of a /48 Site-Level prefix in RFC 3177 from September 2001 to the address plan of RFC 6177 from March 2011, which eschews as far as possible the use of fixed boundaries in the address plan.

What we have today is a single “boundary” in the IPv6 address plan, where the low order 64 bits are locally assigned as an interface identifier, and the high order 64 bits are in essence a network identifier, where the boundary between what constitutes locally defined site networks and globally visible networks is left to each network operator.

The concept behind this 64-bit interface identifier was the idea that all hosts would maintain a constant 64-bit interface identifier irrespective of where and when the host attached to a public network. In theory a site could maintain locally connectivity based only on these 64-bit values, irrespective of the site’s external connectivity. This was supposed to aid in sites improving their resiliency through having an option for multiple external connections and allow disconnected sites to still operate.

As subsequently pointed out in RFC4941, this represented a significant issue for user privacy, allowing an external observer to correlate multiple appearances of the same mobile host on multiple access networks over time. This RFC recommended that IPv6 hosts use a random IPv6 Interface Identifier, and regularly change its value. The effective result is that for almost all IPv6 hosts (some 96.15% of IPv6 hosts as seen in APNIC Lab’s IPv6 measurement program) use these random interface identifiers in place of a static 64-bit interface identifier for public communications.

The question is: Why do we persist with this 64/64 bit boundary in the IPv6 address architecture between the network and the host identifier? Why did we not just go all the way and emulate IPv4’s address architecture and allow the network operator to select their own address length for the network?

I have no rational answer to this question. I’m left with the observation that IPv6 is not in fact a “128-bit address protocol” in the same way as IPv4 is a “32-bit address protocol”. It’s a “64-bits plus a few extra interface identifier bits” address protocol.

While the concepts of the fixed length “Top Level” and “Next Level” aggregator address fields were quickly discarded, the “Site Prefix” has persisted in the IPv6 address architecture, but instead of being a fixed 48-bit value, it’s now a variable length value, which the length determined by each individual network operator, or their IPv6 technology provider.

This leads to the question: What lengths are commonly used by network operators to assign Site Prefixes to each customer?

Unless you are located within the network and can observe the length of the IPv6 address prefix that your provider has assigned to you, this is not an easy question to answer. But if we can assemble a collection of IPv6 addresses used in the public Internet, then we can examine the address to make a reasonable estimate of the site prefix length being used.

For each IPv6 address:

- If bits 48 to 63 are all zero then assume a /48 site prefix, and the site is using subnet 0.
- If any bits between 48 and 55 are one, and bits 56 through 63 are all zero, we assume a /56 prefix.
- If any bits between 48 and 59 are one, and bits 60 through 63 all being zero, we assume a /60 prefix.
- If we see any bits between 48 and 64, we assume a /64 prefix.
- For prefix lengths 48, 56, and 60, if bit 64 is 1, we assume the address is from a second subnet within the ISP assigned prefix.

This exercise of deriving site prefix lengths from IPv6 source addresses was first conducted by Nathan Ward using a small data set of 936 IPv6 host addresses. The results of that analysis can be found at <https://github.com/nward/ipv6-prefix-lengths>. He estimated that 0.01% of these IPv6 addresses used a /48 prefix length, 21% used a /56, 22% used a /60 and 55% of these addresses used a /64 prefix length.

The obvious question is: Does this distribution of values hold when using a significantly larger collection of IPv6 source addresses?

The APNIC Labs ad-based IPv6 measurement platform, collects some 7M – 8M unique IPv6 source addresses per day. These are typically the IPv6 addresses of end user systems who received an ad impression (although there is a small level of exceptions when a VPN is in use, or when obscuring technology, such as Apple Private Data Relay service, is being used).

IPv6 Subnet Use

We’ve applied this subnet classification algorithm to these IPv6 host addresses on a day-by-day basis since the start of 2022. The result is shown in Figure 1.

The results from this exercise show that the current breakdown of subnet prefix sizes in the APNIC Labs data set is consistent with the results from the smaller data set used by Nathan Ward. The /64 subnet is the most common subnet size, used in 57% of cases in recent data. The /60 and /56 subnets are seen in 21% and 20% of cases respectively. Finally, the /48 subnet is seen in 2% of cases.

Over the past 28 months there has been some changes in this distribution, where the relative use of /64 prefixes have dropped from 72% to 57%, while the relative incidence of /56 subnets has risen by 10% over this period. There was a small relative increase in /60 subnets of 5% while the relative use of /48 subnets have remained constant.

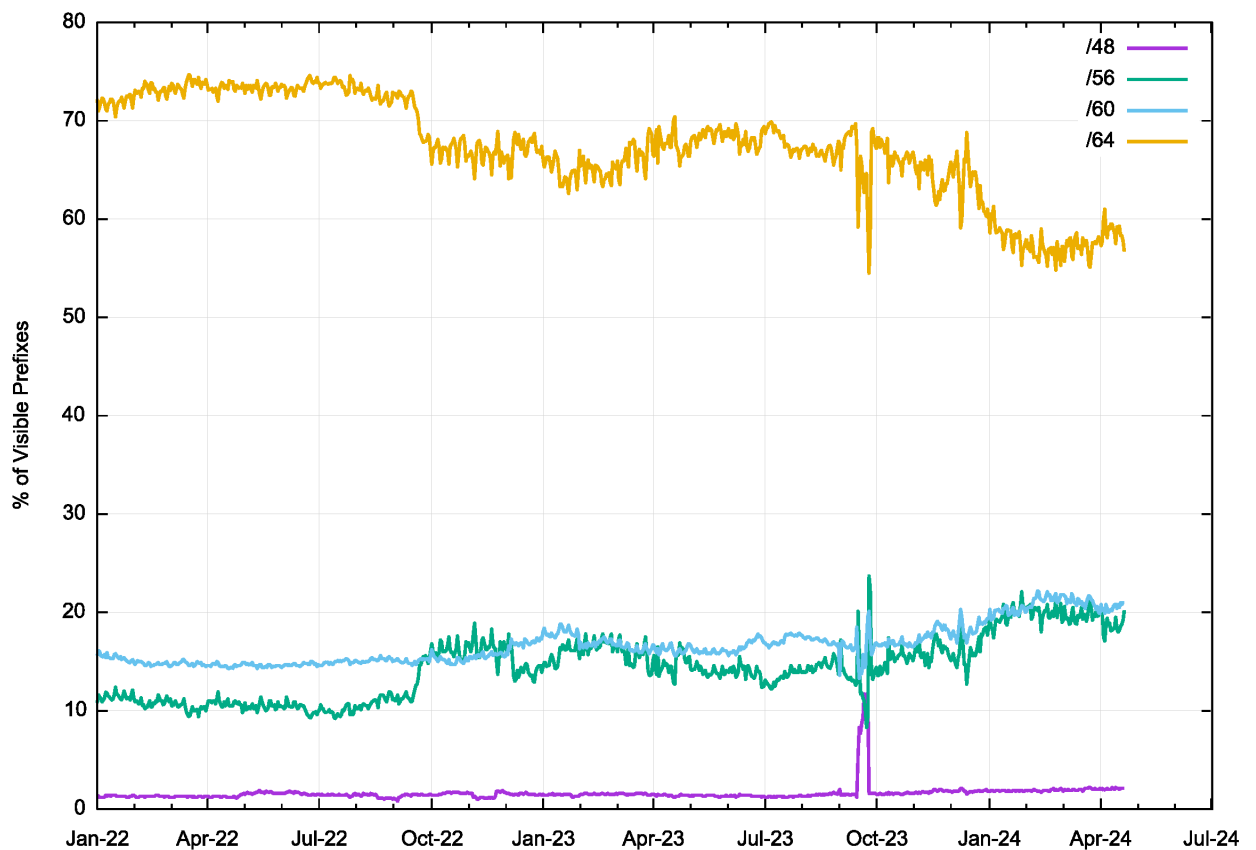


Figure 1 – IPv6 Site Prefix Length Estimation - 2022 - present

This data indicates that few ISPs assign end sites a /48 prefix. A possible explanation of this distribution of subnet sizes is that the use of a /64 prefix is prevalent in mobile services, where individual devices are assigned a /64 prefix by the mobile service provider, while the use of a /60 and /56 appears to have become the default setting in broadband deployments in this data.

It may well be the case that /48 prefixes are used more commonly in enterprise contexts, and the penetration of IPv6 into enterprise environments has been far slower than the deployment in mass market public IP services in mobile and broadband.

Internal Subnet Structure?

Prefix lengths of 63 bits or smaller allow a site to operate an internal subnet structure. The question is: How often are multiple subnets used by end sites?

Again, this is not an easy question to answer by direct measurement, as this internal structure is only visible to an observer located within the end site. However, there is a way to make a rough estimate, and that is by looking for the relative incidence of subnet “1” in visible host addresses. This assumes that most site administrators will number site-local subnets using sequential numbering of 0,1,2 and so on. The presence of subnet 1 in an address, as per the decision algorithm described above, may be an indicator of the use of multiple subnets within a site.

The relative occurrence of subnet 1 in each of the /48, /56 and /60 site prefixes are shown in Figure 2.

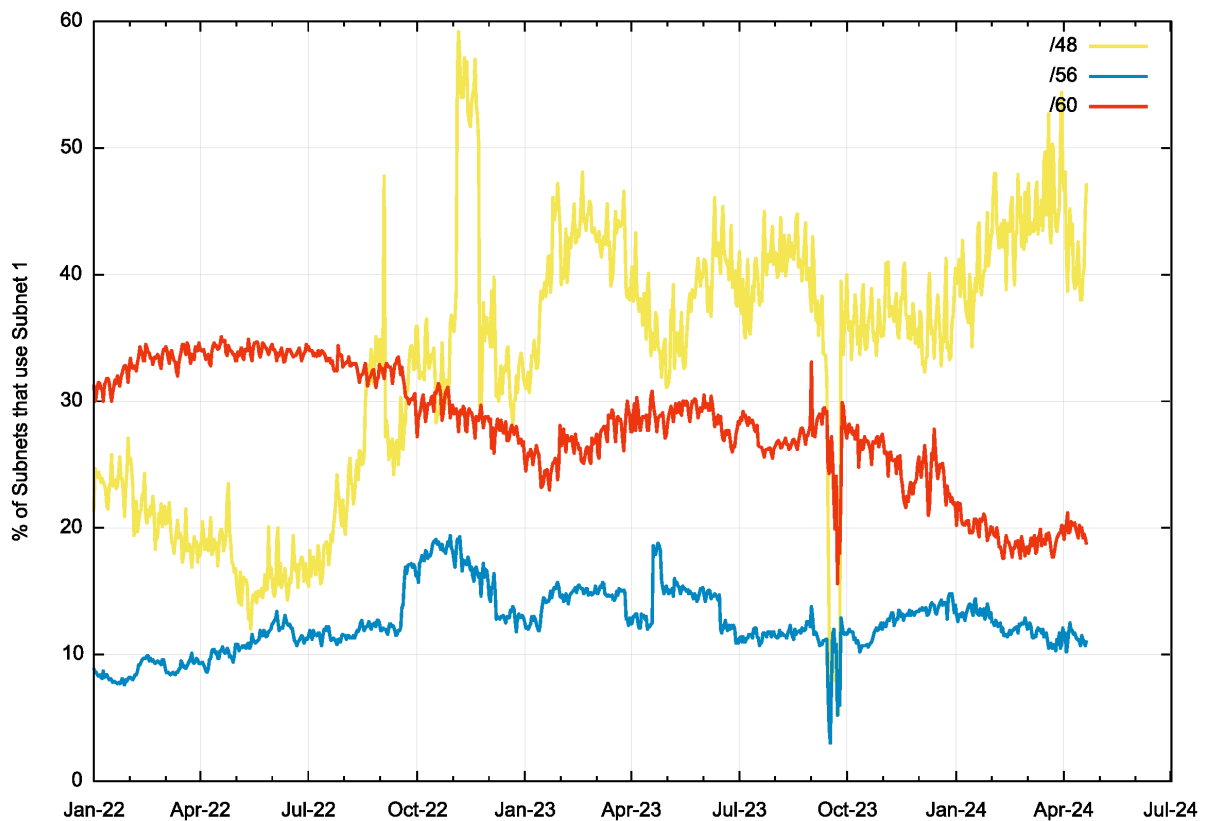


Figure 2 – Subnet 1 occurrence - 2022 - present

This data indicates that multiple subnets may be common in /48 prefixes (~45%), but less so for /56 (10%) and /60 prefixes (20%). This data is consistent with the supposition of /48 prefixes are more commonly used in enterprise scenarios, where multiple subnets are more likely to be used.

Measurement by /48s

It can be argued that the outcome shown in Figure 1, namely that /64s are the most common, is influenced by the observation that there are many more /64 prefixes than /48 prefixes. We can attempt to compensate for this by using a uniform /48 division in the IPv6 address plan. For each observed source IPv6 address, we use the classification algorithm, to derive a likely subnet size, but we then use the encompassing /48 address prefix, and assign this subnet size to the /48 subnet. When we get multiple subnets of different sizes in the same /48 well use the longer subnet (i.e. if we observe both a /48 and a /60 in a common /48 prefix then we'll use the /60 value).

The result of this analysis is shown in Figure 3.

Comparing this result to that shown in Figure 1 we see that the most common subnet size, when looking at the IPv6 space as a set of /48 prefixes, is a /56. While there are a large set of IPv6 source addresses that map to a /64 subnet, they all come from a smaller set of common /48 prefixes.

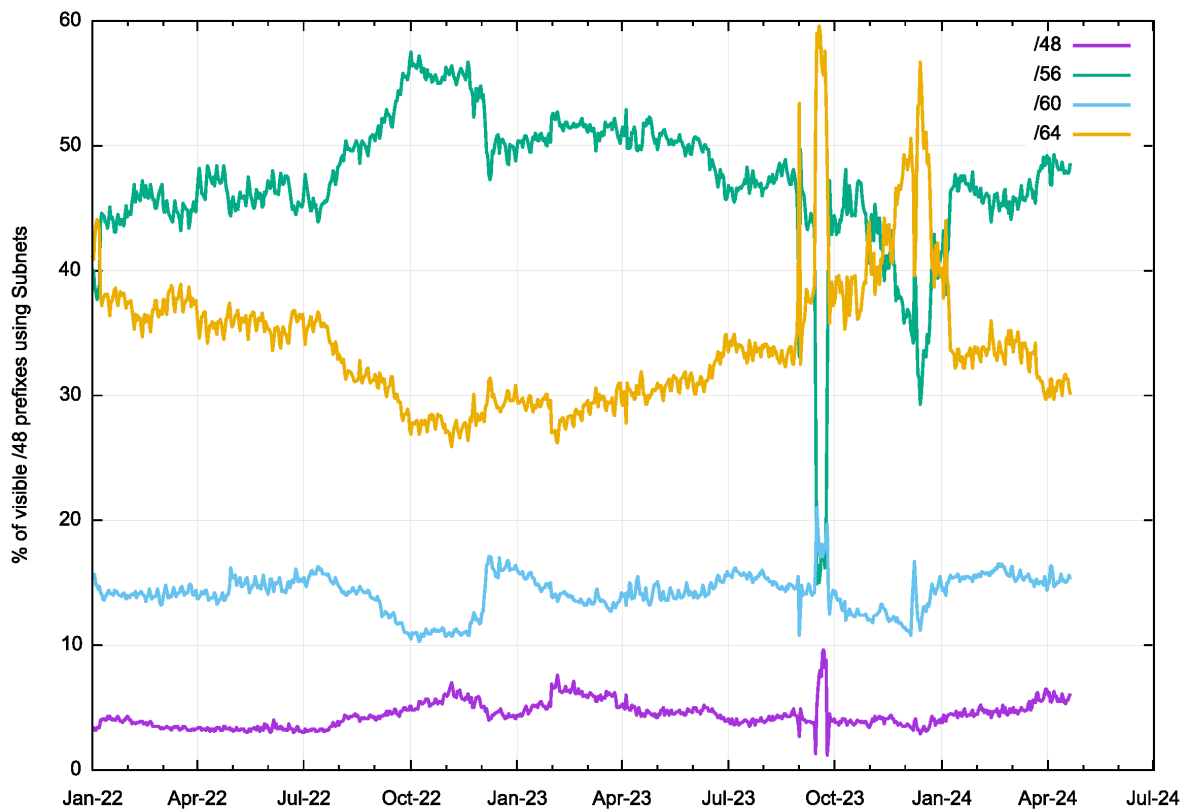


Figure 3 – Proportion of subnets per encompassing /48 prefix

RIR Allocations

It's not clear that a /48 “parent” prefix is the most appropriate one to use here. Alternatively, we could make an assumption that each network uses a uniform subnet address plan, and furthermore assume that each IPv6 address allocation from a Regional Internet Registry (RIR) corresponds to an individual network.

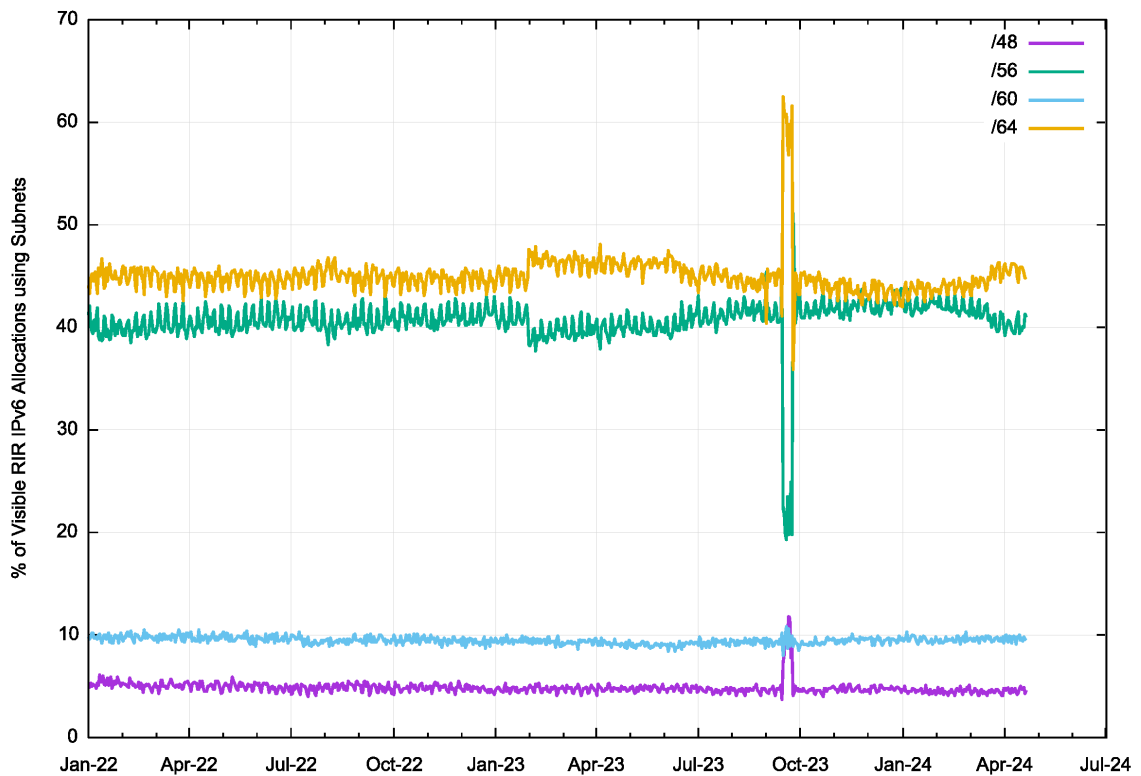


Figure 4 – Proportion of subnets per encompassing RIR address allocation

This perspective of subnet use is similar to Figure 3, where the use of /64 and /56 subnets are the most prevalent, while /60 and /48 subnets are far less common.

Anomalies

This exercise is based on the assumption that where subnetting is being used within an end site, the initial subnet, subnet 0, is the most likely to be used.

This is not necessarily the case. For example, Starlink assigns each end user a /56 (<https://starlink-enterprise-guide.readme.io/docs/dhcp-configuration>). The default subnet identifier, seen in 15,309 cases out of 16,711 in one recent day from this data set, uses subnet number 16 (hex 10) as the default on-site network, causing this algorithm to incorrectly assume that a /60 is being used for Starlink address assignments.

It appears that, at best, this approach offers an approximate view of IPv6 address assignment and subnet behaviours in the IPv6 network.

The role of Subnets in a Network Architecture

Why use subnets at all? The standard response is that “subnets make networks more efficient” by localising traffic. A subnet encompasses an internally self-connected region of a network. Traffic between the attached nodes within a subnet can be handled by the routers within that subnet. Subnets can simplify the network routing architecture, in that routers within a distinct subnet need only maintain routes for the hosts that are located within the subnet, and a “default” router pointing to the network beyond the subnet can be used for all other hosts. If all the hosts in a subnet are addressed from a common address prefix then the external network that hosts the subnet need only maintain a single route to this common address prefix, delegating the details of individual host reachability to the routers within the subnet. All of this is explained in detail in RFC 950.

The entire concept of nested hierarchies of subnets is an intrinsic part of the IPv4 architecture, particularly so when the address architecture was migrated from the old Class A, B and C network/host

fixed boundaries to a classless address architecture where every subnet was essentially defined by a common address prefix and a prefix length.

It's an interesting question to ask to what extent this address architecture is an intrinsic part of today's network designs. In the IPv4 environment the depleted IPv4 address pools have meant that for many networks at the edge the conventional subnet boundaries have been replaced by network address translation boundaries. In the IPv6 environment a subnet still appears to have some traction, but it's closer to a IPv6 Neighbor Discovery realm where the concept of "locality" is defined by virtue of a membership of a common multicast realm. In a networking world where the capabilities of layer-2 switching environments have all but supplanted the former role of direct physical connectivity the answer to a simple question of "are we connected to each other?" is sometimes deceptively complex.

As an abstraction to assist in scaling the networks, subnets still have a role to play, but as the pendulum of the technology of packet networking swings back from routing to switching, it's increasingly challenging to understand exactly what this role is!



Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.