

The ISP Column

An occasional column on things Internet

July 2005

Just how big is IPv6?

- or where did all those addresses go?

Geoff Huston

This is an article about large and small numbers and what happens to numbers when they are used in address plans for networks.

The question we are looking at here is just how long can we expect the 128 bit address set of IPv6 to last before we've run out of IPv6 addresses? And the secondary question is if we assume that we are just a little worried that we are being a little too profligate with these numbers, whether this is something we can quickly rectify without changing the basics of the address plan, or whether there is some more fundamental weakness in the way in which we've been thinking about IPv6.

And for those who want a peek at the answer at the end without wading through all the words in between, the conclusions being drawn here are that there may some level of discomfort with the overall capacity of the current IPv6 address plan, but that this can be rectified in a number of address plan parameter settings without having to redesign the entire IPv6 address plan, which is , on the whole reassuring news for those who see the Internet, and IPv6 in particular, as the underpinning of the global communications enterprise over the coming century. Now back to the details ...

Some Background to Address Architectures

To recap a little over the background here, let's go back to the early 1980's and look at the networking protocols that were prevalent at the time. At that time "computers" were those large mainframes housed in highly specialised environments, with clusters of attached data entry terminals and printing stations. The typical role of a computer network protocol was to link these satellite data entry clusters back to the mainframe mothership. Addressing plans within these networks was constrained to identifying the individual components in these computing clusters. Bandwidth was low and processing and memory capacities highly restrictive. Little wonder that computer networks were designed to suit that environment quite precisely. Per-packet overheads were strongly constrained, and expanding the packet header size to accommodate large address fields was considered an unnecessary extravagance, particularly when data communications appeared to be an extravagant and exotic

alternative to sending computer tapes through the post! So one objective of the protocol design was to limit the size of all the fields in the packet header, including the address fields, in order to improve the payload efficiency. Protocols in use at the time typically used an 8 bit or 16 bit address field, accommodating networks whose total device count was between tens and thousands of devices within each networked computing cluster. The initial address plans were 'flat' in that devices were numbered sequentially from 1. 16-bit address fields also heralded structured address plans, such as a two level plan, where the first 6 bits identified the 'area' and the next 10 bits identified the node within the area. Within these two level plans there was an associated network topology, where all nodes within any area were normally fully interconnected within the area, and all areas had a direct attachment to the 'backbone area' (often area 1).

It was this environment into which the Internet Protocol was introduced. One of the more wild-eyed radical aspects of this protocol design was the use of 32-bit address fields in the packet header. This was coupled with an approach of splitting the address into two components, a network part and a host part. The address structure of the Internet protocol was not intentionally restricted to a single computing cluster, or network, but was intended to encompass multiple networks, and span a much broader environment of interconnection. Supporting this extended address field was the establishment of an IP address registry function where anyone who wished to deploy the IP protocol on their network could apply to the IP address registry for a network prefix identifier that was unique. Once assigned in this manner, the network prefix would not be assigned to any other network in the future. This address plan did not identify any particular area as a 'backbone', and indeed made very few assumptions relating to the topology of the network. The major change in the Internet's address architecture was that from the outset there was the capability and administrative arrangements to ensure that every IP node had a unique address, whether it was connected to other Internet networks at the time or not.

At the time of the original specification of TCP/IP, the early 1980's, this choice of 32 bits, or potentially over 4 billion unique addresses, was either wild-eyed extravagance on the part of the researchers who designed this protocol or an uncanny level of prescience about the shift in the computing environment from a small number of massive lumbering mainframes into a prolific realm of chattering embedded tiny digital devices. I strongly suspect that a strong element of the latter was part of their thinking, and this remarkable choice of address design formed one of the reasons why the internet protocol was in a position to support the personal computing revolution of a decade later.

So how big is 32 bits of address space?

Each additional bit in an address field doubles the number of possible address values that are available for use. 6 bits of address space encompasses a range of values from 0 to 63 (or a bit pattern in binary notation of 000000 through to 111111). Adding an additional bit here would increase the range of possible address values from 0 through to 127. 32 bits of address space corresponds to a address value range from 0 to 2 to the power of 32 (less one), or 4,294,967,295. That's four and a half billion, or a number that is roughly two thirds of the current global human population. When the world was populated with a few hundred thousand computers this was indeed a very large number. In a world endeavouring to equip itself with ubiquitous silicon and ubiquitous communications this number is now looking like it's just too small.

IPv4 and Address Consumption

Interestingly enough, the observation that we would run out of 32 bit addresses did not take all that long to be made after the original IP address specification. By 1990 it was becoming clear that the uptake of IP on all kinds of networks was such that the Internet would run out of addresses pretty quickly. 1995 was mentioned as a likely time when the stock of useful addresses would be exhausted. Part of the problem here was the address plan for IPv4 that attempted to group networks into three sizes. There was "large", with up to 128 networks each of which had up to 16.7 million hosts (Class "A"), "medium", with up to 16,384 networks, each with up to 65,536 hosts (Class "B"), and a further 2,097,152 networks each with up to 256 hosts. None of these choices were "just right". Both the Class A and B spaces were too large, and the class C space was too small for most networks. The commentary on the imminent exhaustion of IP address space was in fact one based on the imminent exhaustion of the Class B address space and the associated very low utilization of this space because it was in fact just too large for most deployed networks.

The response to this failure in the IP address plan was in a number of directions, and the outcomes of this work included a number of refinements to the internet protocol that form a big part of the Internet today. These measures include the use of so-called "classless addresses", as a means of using address blocks as efficiently as possible, the introduction of the dynamic host configuration protocol and associated remote authentication protocols (such as RADIUS) that allow a host to use a unique address for a period of time, and then pass it back to a common pool when it is no longer required in order to allow other systems to use it, the use of various forms of network address translators to allow a small number of public addresses to be shared between a larger set of devices that are configured to use private addresses. Oh yes. And, of course IP version 6.

It was always thought that these other measures were interim steps, and the migration to of the Internet to IPv6 was inevitable. Some 15 years later these interim measures appear to have become well entrenched, and form the Internet as we know it today.

The IPv6 Address Space

Working along the lines of "if big is good, then even bigger is even better" then when it became evident in the early 1990s that 32 bits of address space was not going to be enough for sustained growth of the Internet and more address space was needed, then it should be unsurprising to observe that IPv6 was designed with 128 bit address fields. Don't forget that each additional bit doubles the total span of the address space. So adding an extra 96 bits is well into the overachiever department. Compared to the 4.4 billion addresses in IPv4, IPv6 has 3.4×10^{38} (340 undecillion) addresses, or as Wikipedia tells us,

"if the earth were made entirely out of 1 cubic millimetre grains of sand, then you could give a unique [IPv6] address to each grain in 300 million planets the size of the earth"

http://en.wikipedia.org/wiki/IP_address

Or, using a more earthly analogy:

"The optimistic estimate would allow for 3,911,873,538,269,506,102 addresses per square meter of the surface of the planet Earth."

"IP Next Generation Overview", R. Hinden, Communications of the ACM, Vol. 39, No. 6 (June 1996) pp 61 - 71, ISSN:0001-0782 <http://portal.acm.org/citation.cfm?coll=GUIDE&dl=GUIDE&id=228517>

Given this truly massive number of unique address values in IPv6 then it would be reasonable to conclude that the entire concept of running out of IPv6 addresses, ever, should be a totally alien concept. Right?

Well not quite. One should never under estimate our collective ability to make what in retrospect may be interpreted as somewhat perverse design decisions, and when faced with the seemingly infinite we have an uncanny knack of establishing just how finite we can make it!

Address Plans

In practice it is not possible to configure a distributed network to fully utilize every last address. If we look closely at the IPv4 protocol specification, not all the addresses are used in a conventional manner to support numbering each connected device. A number of address blocks have been reserved, and are not available for use in the network. These include addresses that are used for self-identification, for many-to-many communications models (multicast) and a certain amount of space that is withheld from use for some yet-to-be-defined future purpose.

A similar internal address plan has been set up in IPv6, and a number of addresses have been reserved for special purposes. The current IPv6 address plan is as follows:

Address type	Binary prefix	IPv6 notation
-----	-----	-----
Unspecified	00...0 (128 bits)	::/128
Loopback	00...1 (128 bits)	::1/128
Multicast	11111111	FF00::/8
Link-local unicast	1111111010	FE80::/10
Global unicast	(everything else)	

[RFC3513, as updated by work-in-progress: draft-ietf-ipv6-addr-arch-v4]

There have been a number of iterations of the IPv6 address plan for global unicast addresses over IPv6's extended gestation. The current address plan uses a division of the address into three components: a global network identifier, that corresponds to an address prefix announced in the public network, a subnet identifier, which is used to support internal structure within corporate or campus networks, and a device identifier part that is used to support unique identification of hosts within the local subnet.

While the current IPv4 address plan uses a variable length format for sub-fields within the address, the IPv6 address plan, by contrast has specified a fixed boundary at bit position 64, such that the device identifier part of an address is 64 bits in length, and the global identifier and local network identifier occupies the leading 64 bits. In addition, there are Internet Architecture Board guidelines that suggest that the subnet identifier should be, as far as possible, fixed at a 16 bit length [RFC 3177]. The result is that IPv6 global unicast addresses have an associated address plan that, in the general case, uses a global network identifier in a 48 bit field, a subnet identifier in a 16 bit field, and a local device identifier in a 64 bit field (Figure 1).




Figure 1 – IPv6 Address Structure (RFC 3513, RFC 3177)

It's reasonable to ask why the IPv6 design appears to have reverted to the original IPv4 model of fixed length sub-fields within the address plan. The trade-off here appears to one of simplicity versus efficiency. A fixed length address plan lends itself to simplicity in the associated procedural aspects of configuring a network and its devices with addresses. However, in adopting a 'one size fits all' approach this fixed length address plan

tends to err on the side of encompassing as many network scenarios as possible, and therefore allowing very generous sizes within the fixed size components of the address. Variable sized address plans tend to have higher procedural and operational overheads due the fact that every deployment is in effect a custom deployment, but can be adapted to meet individual requirements quite precisely. The IPv6 fixed length plan still allows customization, but the default action is the same in all cases, and, in theory, these networks can be rolled "out of the box."

Certainly one of the objectives of IPv6 is to create networking environments that can work straight out of the box, and that setting up a network should not rely on detailed configuration of each network component. This is a natural fit with the emerging concept of networking as a ubiquitous commodity. It appears that the rationale behind this design choice was that there is no shortage of address space, and therefore no reason to impose additional configuration burdens on the function of deployment of IPv6 networks in the field. For this reason IPv6 has deliberately adopted the model of using both a fixed size parameter for the device identifier, defining this field as a 64 bit value, and recommending a default size for the subnet identifier of a 16 bit field.

So does this leave us enough address space to encompass all the various visions for deployment of IPv6?

Lets now look at the demand model for IPv6 to see if we can answer this question.

The Demand Model for Global Identifiers

So within this address plan there are 48 bits for the global routing identifier. This is the number of remaining bits once the 64 bit device identifier and the 16 bit subnet identifier are stripped out of the 128 bit address field. There are 281,474,976,710,656 global identifiers in this 48 bit space.

Now if 128 bits encompasses 300 million planets the size of the earth, then one planet requires only 103 bits. By the time you get down to 48 bits then you are looking at the number of grains of sand on a large beach!

The number of grains of sand in a beach 60 km long, 30m wide and 30cm deep is some 2^{48} . This quantity, of the order of 10^{15} , also appears to be roughly the same number as the number of ants in all the world. (<http://pages.prodigy.net/jhonig/bignum/indx.html>)

The demand model for these global identifiers is effectively one of consideration of global populations over some decades to come. This includes considerations of numbers of households, numbers of workplaces, numbers of public agencies, in looking at human activities and their associated communications requirements. However we are also looking at aspects of deployment of silicon, which implies not only consideration of populations of conventional computers , but also consumer electronics, civil infrastructure elements, embedded devices and similar. The considerations of the size of these global populations is of the order of billions in each case, and in looking at a span of some five to ten decades of use this is perhaps better phrased as populations to be serviced with deployments of tens of billions in each of these segments.

We can also expect that the massive scale of deployment will also lead to further commoditization of the service provider market, so that we may expect some thousands of service enterprises each servicing tens of millions of service endpoints in markets that will be characterized by economies of volume rather than higher valued efforts of service differentiation.

The rough order of magnitude of the size of these end populations over time is one of the order of tens of billions, or even possibly low hundreds of billions. The demand population for addressed end sites is then of the order of 10^{11} to 10^{12} . This is equivalent to 2^{40} , or some 40 bits of address space if we could achieve 100% address utilization efficiency in addressing end sites.

The Routing Constraint

Network addresses have utility when they are deployed in the context of a network. For the network to be able to use these addresses then the address plan must fit into the structure of available routing technologies.

Making routing work across very large networks is a long standing issue, and our accumulated understanding of large scale routing to date is that the most effective scaling mechanism for routing is the use of aggregation of information through the imposition of hierarchies in the address plan. There have been a series of efforts to investigate future routing systems that exhibit radically different scaling properties as compared to the current capabilities of the Border Gateway Protocol, and it would be comforting to take the view that the global network will migrate to a different form of routing that has substantially improved scaling properties. However no such routing system has emerged so far from this work, and this different form of routing remains unspecified.

It may be more prudent to take the view that the changes to the inter-domain routing system will be more incremental in nature over the coming years, and that the scaling properties of the existing inter-domain routing protocol, BGP, will be a continuing factor here. The current IPv4 network carries some 160,000 entries, or of the order of 10^5 . It would be reasonable to expect that further refinements of the model and capability improvements in routing elements may lift this by some two orders of magnitude. This indicates that the constraint model of routing appears to be capable of supporting a system with the order of 10^7 entries.

The difference between these two numbers, 10^{12} and 10^7 , requires some leverage in terms of aggregation of addresses into routing entries. The tool that we have to undertake this leverage is that of hierarchies in the address space, and the associated issue is that of the level of hierarchies that need to be used within various providers' address plans. The efficiency of such address plans in terms of the ratio of total address space and the numbered end sites is a critical factor in looking at total consumption levels.

Aggregation and address hierarchies are, in general, relatively inefficient addressing plans, and in looking at total demand estimates then the expected address utilization efficiency is a factor in the overall demand estimation. It is also the case that the addressing plan has to accommodate both large and small providers.

So the next question is one of aggregation efficiency, namely, what level of efficiency can be anticipated if one were to deploy 10^{12} end sites in a network routing system capable of supporting some 10^7 routing entries?


Current IPv6 Address Allocation Policies

Before attempting to answer this question it is useful to briefly review the current address allocation policies as used for IPv6. The current structure is one where the Regional Internet Registries (RIRs) allocate address blocks to service providers. Within the terminology used by the RIRs these service providers are termed "Local Internet Registries" (LIRs).

The minimum allocation unit to LIRs is a /32. LIRs can have access to larger address blocks based on a utilization target applied to the number of end sites for which they will be providing IPv6 services. This utilisation target is based on a Host Density Ratio (more later on this ratio).

What is a /32?

The terminology used to describe address blocks is generally of the form “/n”. The meaning of this term is to describe the number of leading bits, n, that form the common prefix of the address block. In this case a /32 implies that the address block has a common 32 bit prefix, and the address block spans the remaining 96 bits. The lower 64 bits are used for the device identifier, and 16 bits are used by end sites for the subnet identifier, leaving the ISP 16 bits from which to draw identifiers for end sites, allowing for a pool of 65,536 such end site identifiers.



A /20, by comparison has 20 bits in the common prefix, allowing for 28 bits to use for end site identifiers, creating a pool of 2^{28} such identifiers, or 268,435,456 identifier values.

When the LIR has used the block in accordance with the target utilization level a further allocation is made, again with a minimum size of a /32 address block.

LIRs pass address space to end sites, or customers. The allocation policies note that where the end site is a single device with a single network interface, then the allocation is a single address, or a /128. If there is the certain knowledge that the end site will only have a requirement for a single subnet then the allocation is a /64. In all other cases the default allocation unit is a /48, allowing for a pool of 16 bits, or 65,536 values, to be used to number each subnet. Considering the range of possible subnet technologies that reach down to the level of personal networks such as Bluetooth, the anticipated general case is that each end site is assigned a /48 address block.

The RIR address policies are based on a recommendation from the Internet Architecture Board, as documented in RFC3177.

The Host Density Ratio

The above description uses the concept of a “target utilization level” as a means of determining when a block of addresses is considered to be fully utilized.

Networks are not static entities, and various parts of a network grow and shrink over time. The common practice is to divide a network's address space into continuous blocks of addresses, each of which is assigned to serve a distinct section of the network, or subnet. When a new device is added to a part of the network the

intent is that the new device is assigned an available address from the local subnet address pool, leaving the addressing of the remainder of the network unaltered. When the subnet address pool is exhausted then it is necessary to renumber the subnet into a larger address pool. Renumbering a network or even a subnet is at best an extensive and highly disruptive operation. For large networks it becomes a protracted and expensive affair that is best avoided.

For this reason a common network address plan attempts to provide each subnet with sufficient address space to number not only the current collection of attached devices, but also to encompass future expansion of the subnet over time. This implies that achieving a 100% use level of addresses is not an achievable objective. What level of utilization is achievable?

Early experience with this in the IPv4 world indicated that achieving an address utilization rate of 10%, where 10% of the address block was actually used to number devices and 90% was sitting unused in various address pools was an reasonable outcome. Subsequent refinements of the subnetting model in IPv4 with variable length subnet address blocks allowed far higher utilization rates to be achieved, and current IPv4 address distribution policies call for address utilization rates of some 80% as a threshold level that should be achieved before more address space is allocated to the service provider.

This is a relatively extreme metric, and it places a considerable burdens on local network managers to achieve such a high address utilization level. It is often the case that local managers use private addresses using a much lower address utilization level and then place network address translators on the boundary in order to meet these objectives.

With IPv6 the concept of address utilization efficiencies has been redrafted. Within end-sites each subnet has 64 bits of address pool, and no particular utilization target is specified. Even in terms of numbering of subnets there is no particular address efficiency metric, as each end site is assigned a 16 bit subnet space that they can deploy in any manner of their choosing. The only place where an efficiency metric is specified is with the ISP, and that is a metric of the efficiency of the end-site numbering within the service provider's address pool. In this case the consideration has been that a fixed threshold value imposes an unnecessary burden on service providers, and an alternative form of efficiency metric has been used.

The guiding observation in defining this host density metric is that a service provider network typically use a number of levels of hierarchy. A large service provider metric may divide the address space into regions. Within each region the address space is further divided into network access points, or POPs. Within each POP the address space may be further divided into access classes. Within each access class address pool end-site addresses are assigned. This then defines a four level internal address hierarchy. Smaller networks may have a smaller number of internal levels of hierarchy, perhaps using only one or two levels. Assuming that at any level of the hierarchy a utilization efficiency of, say 0.7 (or 70%) can be achieved, then a two level hierarchy achieves a threshold level of efficiency of 0.49 (or 0.7²) and a four level hierarchy would map to 0.24 (or 0.7⁴).

The next part of this process is to define the relative sizes of “large” and “small” networks in terms of the change in network size that corresponds to the additional of a new level of internal hierarchy. It may correspond to an increase in size of the network by a factor of, say, 4. This leads to the general observation that we are looking at a relationship of two exponential values, in which case the ratio of the log of these two values is a constant.

And this leads us to the Host Density Ratio. This ratio is expressed as:

$$HD = \frac{\log(\text{number.of.allocated.objects})}{\log(\text{pool.size})}$$

The value used in the IPv6 address allocation policies is an HD Ratio of 0.8.

The following table shows the target utilization levels for various sizes of IPv6 address blocks, where the left-most column is the threshold level of utilization according to the 0.8 HD-Ratio value.

Prefix	/48 count	end-site count
/32	65,536	7,132
/31	131,072	12,417
/30	262,144	21,619
/29	524,288	37,641
/28	1,048,576	65,536
/27	2,097,152	114,105
/26	4,194,304	198,668
/25	8,388,608	345,901
/24	16,777,216	602,249
/23	33,554,432	1,048,576
/22	67,108,864	1,825,677
/21	134,217,728	3,178,688
/20	268,435,456	5,534,417
/19	536,870,912	9,635,980
/18	1,073,741,824	16,777,216

Putting it all together

The IPv6 address plan is intentionally one that is simple and easy to use. The IPv6 address plan is intended to provide simple structures that allow low overhead deployments of small and large networks, both for the local network management or end site and for the IPv6 service provider in deploying an address plan across their network with ease of expansion while avoiding renumbering whenever possible. The IPv6 address plan is also intended to accommodate the consideration that aggregation and hierarchies of address structures are not highly efficient users of address space.

The inputs to the total consumption of address space are the factors of a 64 bit device identifier, a 16 bit subnet identifier, an HD-Ratio of 0.8 for end-site utilization, a set of global populations of network deployments and an anticipated lifetime of at least 60 years. The basic sum is an end-site population of between 50 billion and 200 billion. Applying at HD-Ratio of 0.8 to this range of values gives a total address requirement of between a /1 to a /4. That's between 1/2 and 1/16 of the total IPv6 address pool.

The corresponding 0.8 HD Ratio mapping is indicated in the following table:

Prefix	/48 count	end-site count
/17	2,147,483,648	29,210,830
/16	4,294,967,296	50,859,008
/15	8,589,934,592	88,550,677
/14	17,179,869,184	154,175,683
/13	34,359,738,368	268,435,456
/12	68,719,476,736	467,373,275
/11	137,438,953,472	813,744,135
/10	274,877,906,944	1,416,810,831
/9	549,755,813,888	2,466,810,934
/8	1,099,511,627,776	4,294,967,296
/7	2,199,023,255,552	7,477,972,398
/6	4,398,046,511,104	13,019,906,166
/5	8,796,093,022,208	22,668,973,294
/4	17,592,186,044,416	39,468,974,941
/3	35,184,372,088,832	68,719,476,736
/2	70,368,744,177,664	119,647,558,364

By comparison a similar estimate is provided in RFC3177, which provided a total end-site census of some 178 billion end-sites, and a calculation of an equivalent address requirement of a /3.

Considering that this calculation makes a number of quite sweeping assumptions there is some uncertainty associated with this number. We may need to stick with this technology for longer than 60 years. It may be that our assumptions about the ubiquity of silicon devices are inadequate, or that we may see the use of different address models, such as one-off use of addresses. These factors can be summarized as:

- Time period estimates (decades vs. centuries)
- Consumption models (recyclable vs. one-time manufacture)
- Network models (single domain vs. overlays)
- Network Service models (value-add-service vs. commodity distribution)
- Device service models (discrete devices vs. ubiquitous embedding)
- Population counts (human populations vs. device populations)
- Address Distribution models (cohesive uniform policies vs. diverse supply streams)
- Overall utilization efficiency models (aggregated commodity supply chains vs. specialized markets)

The question that arises from this is: are we comfortable with this outcome given these uncertainties over the total demand estimate?

If not, then we need to consider the various components of the IPv6 address plan and see if there are some parameter adjustments that can be made that would allow some greater margins in the total address consumption levels. The three areas of consideration are :

- the size of the interface identifier (currently set to 64 bits),
- the size of the subnet identifier (currently set to 16 bits), and
- the value of the Host Density Ratio (currently set to 0.8).

Let's look at each of these in turn to see if there is some latitude to change these settings in such a way that would provide some greater level of "comfort margin" for the total address consumption value.

The 64 bit Interface Identifier

The IPv6 address plan by divides the address into two distinct parts: a network location identification part and a device interface identification part. The dividing point is at the 64th bit position.

In the original incarnation of this approach this measure would allow each manufactured network media interface to be assigned a unique 64 bit identification code. This is functionally aligned to the operation of embedding unique media access identifiers in Ethernet interfaces devices. This interface identification code was intended to function in some sense as an endpoint identification, where, irrespective of the location of the endpoint within the network, it would maintain its unique endpoint identification. At one level this makes sense, in that wherever I may be, I'm still me. My identity is a constant. But the implication here is that the same endpoint identity values cannot be used by two or more distinct endpoints. This turns the capacity of the address space into 2^{64} possible endpoints in any one of 2^{64} network locations. That's a much smaller number than 2^{128} addresses. The benefit was an intention to provide a solution to the current semantic overloading of an IPv4 address, which encompasses elements of both location and identity. In this approach the upper half of the address is its location, while the lower half is its identity.

There are some doubts about the efficacy of this approach.

There are a number of unresolved issues here, relating to uniqueness, persistence, authenticity and privacy. In this approach uniqueness of the identity value is not assured. The IPv6 address architecture calls for this field to be based on the extended Ethernet MAC address (EUI-64). In practice this may not necessarily result in a unique identifier value, in that clashing MAC addresses continue to occur. While a MAC address had a limited context of a local subnetwork's ARP domain the consequences of a collision were contained to the subnetwork. If the interface identity is used as a global endpoint identity token then any form of collision results in an ambiguity over identity and the potential for traffic misdirection. Endpoint devices are not monolithic devices, but a system of individual components. If a device has multiple interfaces, such as a wired and wireless Ethernet port then a change of interface implies a change of MAC address which, in turn implies a change of IP address as well as a change of endpoint identity. The identity field is not protected, so that it is not directly possible to differentiate between a device that has changed its network location but retained its constant identity and an attempt by a third party to assume the identity of an end point. It has privacy implications in that remote sites can track the current network location of an individual endpoint as it changes its locator part.

If the objective of separation of location and identity requires more robust approaches that can address the above issues, then why persist with this approach of using 64 bits for a device identity? If the only residual functional objective here is the unique identification of devices in the context of individual subnets then the use of 64 bits of the address to achieve this is certainly well into a case of over-achievement.

These days the 64-bit IPv6 interface ID is an architectural boundary defined by Stateless Address Autoconfiguration [RFC2462]. This function assumes Interface identifiers are fixed at 64 bits. Changing this boundary would impact existing implementations of this function, and any transition to a different boundary would take some years. An alternative approach is to deprecate stateless autoconfiguration completely for generating interface identifiers and use the Dynamic Host Configuration Protocol (DHCP) for this function. However, client implementation of DHCP for address configuration is not mandatory in IPv6, and it is believed that a significant percentage of IPv6 implementations do not support DHCP for address configuration.

So already it appears that even prior to mass deployment IPv6 has managed to accumulate some issues of legacy here, and while a change in the length of this identifier would recover a large number of address bits, this would have some impact on existing implementations of IPv6..

There is a considerable measure of reluctance for further change here that must be acknowledged. IPv6 has had a considerable developmental lead time and there is a considerable body of opinion that its now time to cease further modifications to the protocol specification and provide the industry with a stable view of the protocol. Without this assurance of stability vendors are reluctant to commit the protocol into products, service providers are reluctant to commit to deployment programs and the protocol remains an experiment rather than a service platform for a communications network. So while this particular part of the address plan represents the greatest level of gain in terms of total address capacity, it also presents a considerable risk to the industry acceptance of IPv6, and for this reason changes in the length and structure of this part of the address plan should be contemplated only with considerable trepidation.

The Subnet Identifier

The subnet identifier part of the IPv6 address is, strictly speaking a variable length field. However, within the parameters of current address allocation policies the Regional Internet Registries assume that general case for end site assignments are /48s, and thus utilization measurements are based on an HD-ratio that counts numbers of /48 assignments. Allocating a /48 to an end sites allows each site to deploy up to 65,536 subnets. While that number may make sense for larger enterprises, it is admittedly hard to imagine a typical home network, or a personal local area network requiring this much subnet address space.

Looking back at some of the original motivations behind the /48 recommendation [RFC 3177], one overriding concern was to ensure that end sites could easily obtain sufficient address space without having to "jump through administrative hoops" to do so. As a comparison point, in IPv4 typical home users are given a single public IP address (even this is not always assured), but getting even a small number more is often a more expensive option either in terms of the effort needed to obtain additional addresses, or in the actual cost. It should be noted that the "cost" of additional addresses cannot be generally be justified by the actual cost of those addresses, but the need for additional address is sometimes used to imply a different type or "higher grade" of service, for which some ISPs charge a premium. Thus, an important goal in IPv6 was to significantly change the default end site assignment, from "a single address" to "multiple networks".

Another motivating concern was that if a site changes ISPs and subsequently renumbers, renumbering from a larger set of "subnet bits" into a smaller set is particularly painful, as it can require making changes to the network itself (e.g., collapsing links) as well as reconfiguring the network into a different prefix and associated prefix length. In contrast, renumbering a site into a subnet that has the same number of subnet bits is considered to be easier, because only the top-level bits of the common address prefix need to change. Thus, another goal of the RFC 3177 recommendation is to ensure that upon renumbering, one does not have to deal with a comprehensive reconfiguration of the local network.

These concerns were met by the /48 recommendation, but could also be realized through a more conservative approach. For example, one can imagine "classes" of users, with default sizes for each class. For example:

- A PDA device with a low bandwidth WAN connection and a personal area network (PAN) connection - a single network or /64 assignment.

The /64 assignment allows for the addressing of a number of hosts, each connected to the same PAN link as the device. This would be appropriate in deployments where the end device is not expected to provide connectivity services to a larger site, but is intended to provide connectivity for the device and a small number additional devices directly connected to the same PAN as the primary device.

- Small Office, Home Office (SOHO) - expected to have a small number of networks - a /56 assignment

This is similar to the /48 motivation, but includes the expectation that the typical small office or home environment has a limited requirement for multiple discrete subnets, and this expectation could be generally achieved within a pool for 256 discrete subnet identifiers.

- Other enterprise and organizational entities - a /48 assignment as the default

Although, as with existing allocation policies larger end site allocations are possible within this framework, according to the total end site requirement.

A change in policy (such as above) would have a significant impact on address consumption projections and the expected longevity for IPv6. For example, changing the default allocation from a /48 to /56 (for the overall majority of end sites) would result in a reduction of cumulative address consumption by some 6 or 7 bits, or around two orders of magnitude. Of course, the exact amount of change depends on the relative number of home users compared with the number of larger sites over time.

One can, of course, imagine a policy supporting the entire range of assignments between /48 and /64, depending on the size or type of end site. However, this must be balanced against the advantages of having a small number of simple policies, so that end users can easily identify which assignment size is appropriate for them, and that there is wide agreement among ISPs as to what reasonable assignment sizes are for a given customer class. Having excess flexibility in selecting an appropriate assignment size for a given customer type

can lead to different ISPs offering different assignment sizes to the same customer. This is undesirable because it may lead to a need to renumber into a smaller subnet space when switching ISPs, or may lead to ISPs attempting to differentiate their service offerings by offering the most liberal address assignment policies, defeating the purpose of having wide range of policies.

The advantage of this approach is that it does not impact on existing IPv6 protocol implementations, nor does it create a legacy or transitional impact. This sits comfortably within the realm of a change to the allocation policy parameters that allow a more precise fit of the size of the allocated address block to the nature of the intended use of the addresses without imposing a significant additional administrative overhead on service providers, vendors or end consumers.

The HD Ratio

Coming from an IPv4 deployment environment the HD-Ratio value of 0.8 represents a relatively radical change to the way in which we view end sites address allocations. For example, under the IPv4 address allocation policies a consumer market service provider with some 5 million customers would be expected to achieve an overall 80% address utilization. This would correspond to an address plan that would service this customer base from a pool of some 6.5 million /32 IPv4 addresses, or a total address allocation of a /9. A further allocation would be made only when the total addressed population exceeds 6.7 million. These days with DHCP and NATS most service providers have become accustomed to achieving even higher address utilization densities in IPv4, and it is not unusual to see such a service provider with some 5 million customers using a total address pool of a /11, or some 2 million /32 addresses.

The equivalent allocation in IPv6 would be a /20, or some 268 million /48 end site prefixes to service the same 5 million customers. And once the customer population exceeded some 5.5 million customer the allocation policies would allow for a further application of a /20, making a total of some 536 million end site addresses to draw from. This 1% utilization level of end sites addresses is well distanced from the current IPv4 allocation parameters, and the question arises as to whether this allocation policy has managed to pass across the points of sound engineering and venture into the spaces that could be associated with extravagant use.

As noted above the basic proposition behind the HD Ratio is that the number of internal levels of aggregation hierarchy within a network increases in proportion to the log of the size of the network, and that at each level in the hierarchy one can expect to achieve a fixed level of utilization efficiency. This basic proposition appears to match our understanding to the capabilities of routing and also appears to match out experience with network design, so there appears to be nothing intrinsically wrong with the capability of the HD Ratio to capture the nature of address use within deployed networks.

However, the lingering uncertainty remains that the value of 0.8 may not be the most appropriate value to capture what we would regard as reasonable engineering practice in network design, particularly with larger networks. In exploring scenarios that would result from various HD Ratio values, the first step is to look at the efficiency outcomes that would result from differing values of the HD Ratio, and map these back to the basic function of the number of internal levels of network hierarchy. Figure 2 shows the various utilization efficiency values that result from changing the HD Ratio for various sizes of address blocks.

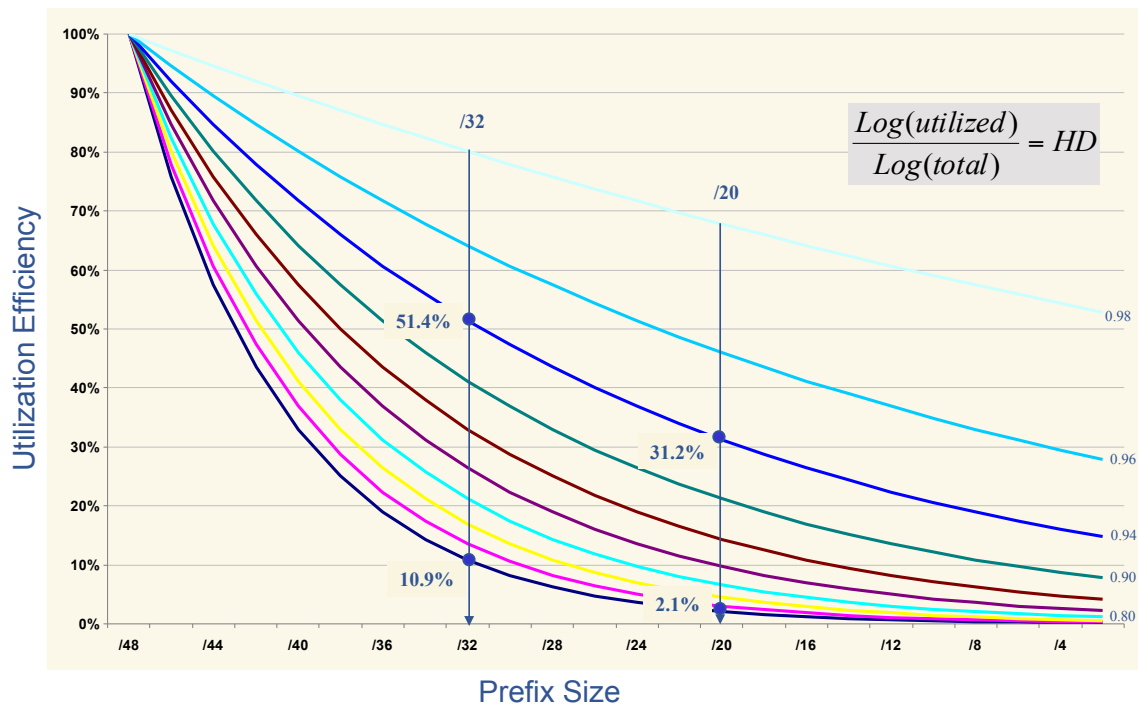


Figure 2 – HD Ratio Outcomes

The first vertical line represents the minimum allocation size of a /32. With a HD Ratio value of 0.8 a service provider can obtain a further allocation of address space once the utilization efficiency reaches 10%, or some 6,500 end sites drawn from a pool of 65,536 site identifiers. The second vertical line represents our example service provider with its 5 million customers. With an HD Ratio of 0.8 the threshold utilization efficiency level is some 2%. In terms of internal levels of network hierarchy this corresponds to 18 internal levels of hierarchy at a per-level efficiency of some 80%. Even with a per-level efficiency level of 70% this still represents 11 levels of internal hierarchy within the network.

This 0.8 value for the HD Ratio does not appear to capture reasonable engineering expectations of network design. Even the largest service provider networks do not encompass more than 5 or 6 levels of internal hierarchy and the internal routing protocols typically operate on a simple two level hierarchy.

It may be useful to consider a higher value of the HD Ratio for address allocation policies. As can be seen in Figure 2, an HD Ratio value of 0.94 would rephrase these threshold levels such that a /32 would need to be used at a level of some 50% before a further allocation is made, while the /20 allocation would need to achieve a 31% efficiency level. This latter value represents a network with 5 internal levels of hierarchy, each being utilized to an average of 80% efficiency. As an initial observation this latter value appears to represent a more realistic model of network deployment based on a competently executed network design.

Another way of looking at this data is to examine the recent past in terms of Internet business activity levels in IPv4, as expressed in address allocations, and see how this would relate to IPv6. The basic question posed in this exercise is: what would've been the total address consumption level over the past three years if we had been using IPv6 instead of IPv4? And how would this total consumption profile change if we'd been using a different value for the HD Ratio?

This simulation exercise produces some surprising outcomes. The first is that 80% of the address allocations would remain at the /32 minimum allocation size or at a /31. Varying the HD Ratio between 0.80 and 0.96 has

little impact on this outcome. So for the majority of ISPS in the last three years a change in the HD Ratio value would have no significant impact on the amount of allocated addresses that they would receive. The second outcome is that only 2% of allocations are greater than a /27. Changing the HD Ratio for these allocations would lift the average address utilization efficiency level from 4% to 25% by a change in the HD Ratio value from 0.80 to 0.94. In other words only a small number of large providers would see some change in the target efficiency levels with such a change. The third outcome is that the change in total address consumption by such a change in the HD Ratio value is a factor of 10. In other words under the current HD Ratio value of 0.8 a small fraction of the allocations (2%) is consuming over 95% of the total address pool.

So perhaps there is some benefit in reviewing this initial choice of 0.80 as an HD Ratio value. The relevant questions here is what is an appropriate HD Ratio value to use? Certainly the initial choice of 0.8 as a value was a somewhat arbitrary one, made more in an effort to define an initial set of address allocation policies than being based in a more deeply researched effort to model sound engineering design principles. In reconsidering this value it would be helpful to consider the following aspects:

- What is common practice in today's network in terms of internal architecture?
- Should we define a common 'baseline' efficiency level rather than an average attainable level? In other words, what value would be readily achievable by large and small networks without resorting to frequent network renumbering or unacceptable internal route fragmentation?
- What are the overall longer term objectives? What is the anticipated address pool lifetime of various HD Ratio values? What would be the anticipated impact on the routing space?

It would appear that some further activity is merited here to explore what value for a threshold address utilization efficiency level that represents a reasonable balance between simplicity of network deployment and the larger issues of conservatism in the impacts on the routing space and ensuring that the overall address pool can indeed accommodate extended lifetime expectations.

Putting it back together again

It appears that there are two aspects to the current address policy framework that merit further broad consideration, namely the subnet identifier size and the HD Ratio.

An additional point in the subnet allocation policy, using a /56 allocation point for SOHO end sites in addition to the current allocation points may alter the cumulative address consumption total by some 6 to 7 bits of address space, without any major impact on the engineering of end site networks, and without any significant impact on service provider procedures in address allocations to end sites. Such a measure would still preserve the essential elements of simplicity while allowing the overall majority of end sites to use an address block that is more commensurate with anticipated needs in terms of subnetting.

The HD Ratio appears to be another area of further study. Initial studies of the impacts of various HD Ratio settings indicate that if the HD Ratio setting of 0.8 implies a total consumption of a certain amount of address space, then a setting of 0.87 would imply a total consumption of $\frac{1}{2}$ of this amount and a setting of 0.94 would imply a total consumption of $\frac{1}{10}$ of this amount. In other words there is the potential to alter the cumulative consumption of address space by 3 bits.

Just these two measures would provide latitude to reduce total consumption levels by up to 10 bits, or a total consumption of between a /10 to a /17 of address space. If the initial estimates of a total consumption of a /1 to a /4 appear to represent some level of discomfort in the total capacity of IPv6 it is reasonable to estimate that a /10

to a /17 would represent a much higher level of confidence that IPv6 would be capable of meeting a much broader set of potential future scenarios for the role on the Internet across the coming century or perhaps even longer.

Public Policy and the “Fairness” factor

If the current address plan has risks of premature exhaustion, then the consequent question is whether these risks should be addresses now or later. One approach is to adopt a “wait and see” attitude, and defer consideration until more data is available.

This viewpoint is expressed in RFC3177:

*We are highly confident in the validity of this analysis, based on experience with IPv4 and several other address spaces, and on extremely ambitious scaling goals for the Internet amounting to an 80 bit address space *per person*. Even so, being acutely aware of the history of under-estimating demand, the IETF has reserved more than 85% of the address space (i.e., the bulk of the space not under the 001 Global Unicast Address prefix). **Therefore, if the analysis does one day turn out to be wrong, our successors will still have the option of imposing much more restrictive allocation policies on the remaining 85%.** However, we must stress that vendors should not encode any of the boundaries discussed here either in software nor hardware. Under that assumption, should we ever have to use the remaining 85% of the address space, such a migration may not be devoid of pain, but it should be far less disruptive than deployment of a new version of IP.*

RFC3177

An alternative way of expressing this perspective is that it appears to be premature to consider changes to the IPv6 address plan when we have so little experience with deployment of IPv6. It would appear that we are not qualified to make such decision and leave them to more qualified individuals. Who would they be? From this perspective they would be the network engineers of the future who would have had 10-20 years of IPv6 operational experience.

Lets look at this assertion in a little more detail. Now if the consumption analysis in RFC3177 is indeed flawed, and uptake is larger than has been anticipated here, then yes, there will still be large pools of unallocated address space available, and yes, it will be possible, in theory at any rate, to use a different addressing plan on this remaining space which targets a higher utilization rate. However the installed base of IPv6 will also be extremely large at this point. Indeed it will be so large that the problem of inertial mass and potential inequities in distribution structures will effectively imply that any changes will be extremely tough.

It could be argued that we have already lived through a similar transition in IPv4 in the change from class-based addressing to one of classless addressing plus Network Address Translators. The legacy of this transition is uncomfortable, with later adopters pointing to the somewhat liberal address holdings of the early adopters and asking why they have to bear the brunt of the cost and effort to achieve very high address utilization rates while the early adopters are still able to deploy relatively simple, but somewhat more extravagant addressing schemes across their networks.

When to consider such a change to the address plan is very much a public policy topic. While there is a temptation to leave well alone, from a public policy perspective we stand the risk of, yet again, visibly creating an early adopter reward and a corresponding late adopter set of barriers and penalties. I suspect that IP has already exhausted any tolerance that may have been enjoyed in the past on this type of behaviour and there is a strong impetus on the part of many developing populous economies not to see such a precise rerun of what they would term previous mistakes. This is not an abstract concept but one where we are already seeing

proposals from the ITU-T to establish an alternative address distribution system that is based around this particular concern of creating a framework that established early adopter rewards and late adopter penalties.

In other words it is possible to put forward the case that this is a premature discussion, but others, for equally valid reasons, see it as being timely, while others see this as an urgent priority. There is a case to be made that we should study the evolution of address policies in the history of IPv4 and be careful to avoid a needless repetition of earlier mistakes. It would appear to be prudent, and indeed "fairer" to plan for success rather than failure, and plan for extensive, indeed ubiquitous deployment of IPv6 for an extended period of time. In such a scenario there is little room for structural inequities in the address distribution model, and that at all times all players should be positioned evenly with respect to access to addresses. Consequently there would be little room to adjust the address plan parameters on the fly and we should exercise some care to ensure that the address plan structure we adopt at the outset has sufficient room to accommodate future requirements on a similar, if not identical, basis. From this perspective the time for consideration of the address plan and its associated parameters is now, rather than deferring the matter to some unspecified future time.

The alternative is that the installed base of IPv6 will consume very little address space in the coming decades, in which case the entire topic would be irrelevant! In other words this topic is predicated on the assumption that in some 50 or 100 years hence we will still be using IP as the base technology for a global communications enterprise.

This is a central topic to the entire consideration of IPv6 address plans. My best answer to this assumption is that I really don't know which, logically, admits the possibility of "yes, we'll still be using IP a century hence. Some technologies are "sticky" simply because they work and the cost of universal adoption of alternatives are just too high. Over a century later we still use the internal combustion engine, decades later we still use amplitude modulated radio signalling, and so on. It may well be the case that packet switching and IP is one of these "sticky" technologies, in which case the address architecture is indeed a critical issue.

Its not clear that we should be in the business of built in obsolescence, and certainly not if we can buy additional time without undue pain. We've looked at the HD ratio and the subnet boundary as potential points of variation in the IPv6 address plan that could admit more efficient utilization without substantial alteration to the overall IPv6 architecture and without undue need to alter existing equipment, software or current deployments, such as they are today. Its certainly the case that alteration of the length of the global identifier could admit vastly greater address utilization benefits but of course the question here is, simply, whether the gain is worth the pain.

However, its sensible to also note that if we think that "installed base" is an argument today in terms of the pain associated with changing the 64 bit length for the device identifier, just wait until the installed base of end sites gets to the 30 billion mark that is commensurate with a 1/4 consumption under current policies. 30 billion end sites would be a very impressively large installed base, and its inertial impetus would say to me that at that stage your wriggle room for changes in the address plan for the remaining space is pretty much a lost opportunity. So if we are having trouble now in looking at the global identifier on the basis of the inertial mass of already deployed systems and services, then you cannot also put forward the proposition that we can change things once we've deployed 30 billion end site instances of the same.

So I'm afraid that "we've still got wriggle room in the future so don't worry about it now" is not an approach that can be accepted easily - if at all. At that point the late comers will be complaining that they are exposed to tougher and more constrained policies that are deployed at a higher cost than that of the early adopters - and if all this sounds hauntingly familiar in reference to the current debates about national interests and highly populous economies and various address policy frameworks, then it should. I'm afraid that there's an increasing cynicism out there about the refrain of "don't worry we'll fix it once its visibly broken" with respect to address

policies. We should at this point be striving to instil some broad confidence in the proposition that we can provide a stable and enduring platform for the world's communications needs.

The 64/64 split is not quite in the same category here, and there is an impact on the current address architecture. Its true that the original motivations for this particular aspect of the address architecture have largely dissipated, or at least have been unable to be realized, and the residual reasons for its adoption are based more in legacy conformance than in true utility. But here its not quite so clear to me that change is necessary. Maybe it would be more practical to pursue some more conservative opportunities that represent some small scale parameter value shifts and adopt a preference to look at the HD Ratio and the Subnet identifier size points over looking at the 64 bit split point between local identification and routing identifiers.

There is some need to exercise careful judgement here. In attempting to look at measures that would ensure a prolific and valuable lifecycle for IPv6 over an extended time care needs to be exercised in ensuring that we continue to have a stable technology base in IPv6. Further changes to the IPv6 protocol at this stage would entrench attitudes that IPv6 remains a developmental exercise rather than a technology capable of sustaining a global investment of trillions of dollars over the coming decades. However, happily, there does appear to be sufficient scope to make some small parameter changes to the IPv6 address allocations policies without making any changes to the protocol itself that would ensure that even the most optimistic predictions of uptake of IPv6 across its lifetime can be readily fuelled by availability of that most essential element of networks: addresses.

References

RFC 1715 The H Ratio for Address Assignment Efficiency, C. Huitema, November 1994.

RFC 2462 IPv6 Stateless Address Autoconfiguration, S. Thomson, T. Narten, December 1998.

RFC 3177 IAB/IESG Recommendations on IPv6 Address Allocations to Sites, IAB, IESG, September 2001.

RFC 3194 The H-Density Ratio for Address Assignment Efficiency. An Update on the H Ratio, A. Durand, C. Huitema, November 2001.

RFC 3513 Internet Protocol Version 6 (IPv6) Addressing Architecture, R. Hinden, S. Deering, April 2003.

Acknowledgement

This work has been greatly facilitated and influenced by discussions held within the IAB ad-hoc IPv6 committee, and in hallway and public discussions at the ARIN XV and RIPE 50 meetings. The author is particularly indebted to Thomas Narten, particularly for the material relating to the consideration of the subnet identifier, but also for helpful discussions across the range of material considered in this document.

Disclaimer

The views expressed are the author's and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

About the Author

Geoff Huston B.Sc., M.Sc., has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and has been active in the Internet Engineering Task Force for many years.

www.potaroo.net