Internet Society

The ISP Column
An occasional column on things Internet

October 2007

Geoff Huston

IPv6 Local Addresses

I suppose I'm no different to many people who have spent some time in the Internet Engineering Task Force (IETF) in making the observation that I've always been fascinated by the process of technology development.

I've often heard this process being likened to that of the development of mathematical systems, where the assumption that lurks behind the process is that every proposition is ultimately provably right or wrong, and that the exercise is one of exposing the reasoning that leads to this deterministic outcome. This is perhaps likened to an attitude that the truth is out there, somewhere, and our task is to discover it.

> A proposition that motivated Bertrand Russell and Albert Whitehead's *Principia Mathematica* attempting to demonstrate the fundamental decideability of the formal mathematical system at the start of the 20th century – a proposition that was elegantly destroyed by Kurt Gödel in 1931, which, in turn, lead to the work by Alonzo Church and Alan Turing on computability later in the 1930's.

I've heard this deterministic perspective on technology development being used to support many of the processes used in the IETF. For example, the rejection of formal voting was best expressed by the observation that voting on technology choices made about as much sense as voting on the correct value of the mathematical constants π or e! The underlying view was that technology is a highly constrained system and that technology decision points are resolveable in a deterministic manner. On the other hand I've also heard the view that technology choices are often highly subjective and that technology decisions are often made for reasons that vary between the most whimsical and the most venal.

While there are often cases where the "right" thing to do is blindingly obvious, and other cases where a few basic technology principles can lead you to a sensible outcome. Avoiding complexity and preferring minimalism whenever possible is always a good yardstick.

> This is a perspective that can be readily traced back to William of Occam's frequent reference to "Pluralitas non est ponenda sine neccesitate" in the 13th century. The term "Occam's Razor" describes this preference for simple solutions that demonstrate adequacy without excessive ornamentation or complexity.

Avoiding continual reinvention of basic tools also helps, and building upon experience is also helpful. But even so there are time when it appears to me that the decision point is one between what appears to be non-technical propositions, and at some point the decision becomes a subjective judgement.

So how does all this relate to IPv6 Site Local Addresses? Lets see.

Private Use Addresses

So-called “private use” addresses were not part of the original IPv4 address architecture. Addresses were considered to be plentiful and anyone who asked was given addresses, for no cost. Whether these addresses were to be used in the context of the nascent global Internet or used in a purely private context didn’t matter. The consistent property here was that all addresses obtained through the address distribution process were unique addresses and could be used in any context, private, semi-private or public with no risk of collision.

However this arrangement was not sustainable, for two reasons; the IPv4 address pool is of finite size, a fact that we are painfully aware of these days, and providing globally unique address space for private use was exacerbating address consumption, and, secondly, it was not possible to run an address registry at no cost, and while it was convenient that a third party was picking up the tab at the time, this was not a sustainable arrangement.

This was resolved by creating a distinction between public and private address use. Because private use contexts have no strict requirement for global uniqueness it is possible for all private use contexts to use the same addresses, leaving globally unique addresses for use in public contexts. The IETF defined the concept of private use addresses in RFC1597 in March 1994, later refined in RFC1918 in February 1996. These address blocks, 10/8, 172.16/12 and 192.168/16, are defined for private use. This was subsequently coupled with Network Address Translators and the result is that it appears that more of the Internet is now addressed using private address space than public space.

IPv6 Site Local Addresses

One view of IPv6 was that given the truly massive address span of 128 bits the entire concept of address scarcity was irrelevant, and in IPv6 there was simply no need to recycle IPv6 addresses for use in private contexts. IPv6 presented us with the ability to revive the consistency of the original IPv4 address architecture where an address had no implicit scope associated with it.

But while IPv6 are abundant, they are not necessarily readily available nor freely available The address distribution framework used today now exposes the costs of operating the registry function to the address holders and there is an associated policy framework that imposes qualifications on address holders according to the prevailing policies of address distribution. The costs and the associated policies are well aligned to the public use of addresses, where the value of utility of these addresses offsets the costs associated with the registry functions. But the same probably cannot be said with respect to private use. Why should I pay to have globally unique addresses registered simply to set up a my home network?

Given the proved utility of private use addresses in IPv4 it was not surprising that the IPv6 address architecture included a private use address pool. This address block, FEC0::/10, was intended for local private use. It was defined in many ways similar to IPv4 private space with overlapping use, no registration requirement, and no usage costs.

Problem solved – right?

Unfortunately not so!

The problem appeared to be that ‘site’ was an ill-defined concept. To quote from, RFC3879:

```
Depending on whom we ask, the definition of the site scope varies.
It may map security boundaries, reachability boundaries, routing
boundaries, QOS boundaries, administrative boundaries, funding
boundaries, some other kinds of boundaries, or a combination of
these.  It is very unclear that a single scope could satisfy all
these requirements.

…

In summary, the current concept of site is naive, and does not map
```

```
   operational requirements.
```

Oddly enough that was not considered to be such a show-stopping problem for the private use addresses of IPv4, but for IPv6 this was seen as a significant problem.

But removing private use addresses altogether from the IPv6 address architecture was not on the cards either. It was not possible to re-implement the concept of a common pool of "scope-neutral" addresses that had potential use in private or public contexts. The address registry framework was already highly attuned to supporting the public network with a strong emphasis on supporting provider-based address hierarchies that attempted to keep the inflation of the routing domain under control. In the public IPv6 address distribution framework you needed to be a service provider to qualify for an address allocation, and all the policies and costs were attuned to the characteristics of the ISP sector.

IPv6 Unique Local Addresses

What was offered to replace a single site local address prefix that every user would potentially collide with any other user in the event that these local use prefixes leaked out was a very large set of such local use address prefixes. The intent was that the collection of such local use address prefixes was large enough to meet some reasonable prediction of total demand, and use a prefix selection mechanism that allowed each intending user the ability to select a prefix that was unique, preferably without incurring the additional overheads of using an address registry to ensure that the selection is unique. Uniqueness of the local addresses would ensure that even if there was some level of intersection of locally addressed realms the result would be benign, in that there would be no consequent address collision.

The IETF's original formulation of these unique local addresses (ULAs) was a two-part approach.

The first part is that of locally selected ULAs, described in RFC4193. Here a block of addresses is reserved for local private use (FC00::/7) and one half of that block (FD00::/8) is used with a local selection method. Each intended user generates a 40 bit value at random, and appends to 0xFD to create a 48 bit local use address prefix. This self-selection mechanism certainly satisfies the criteria of being freely available, but the question of true uniqueness in this approach probably needs closer investigation.

How unique are these self-selected ULAs?

When two parties each make a random selection from a 40 bit space the number of discrete values is 1,099,511,627,776, so the chances that these 2 selection events will collide is around 1 in a trillion. This is probably acceptably low odds.

More generally, making N selections from a pool of M numbers gives a probability p that any 2 will collide as

$$p = 1 - (N! / N^{M} \times (N-M)!\ )$$

Solving this probability for the 40 bit ULA address pool gives the result that the probability of a collision will rise above 0.5 once there are more than 1.24 million random selections from the pool. In a use context strictly limited to private networks this could still be quite a comfortable outcome, but if these prefixes were ever to be used in a colliding space, such as in the reverse DNS, or in the public routed network, then any probability of address collision is not an acceptable outcome.

To address this issue of "almost but not quite unique" the second half of the local use address block, FC00::/8 was to be managed in a more conventional manner using a central address registry to ensure that all selections of address prefixes from this pool were "assuredly unique". These are "centrally assigned ULAs", or ULA-C's as distinct from the self-selected ULAs.

When this second part of the ULA approach was conveyed to the address policy forums the reaction was somewhat negative. If these address prefixes are indeed truly unique that what prevents them being treated in the same fashion as any other piece of public unicast address space? If indeed these addresses were identical to public address blocks in everything but name then it appears pretty obvious that these addresses would find their way into the public space. The problem with this possibility was that such addresses have no hierarchical structure. They are very similar to the original /24 address blocks in the IPv4 address architecture and the public use of such addresses had a similar potential to cause routing inflation in the IPv6 world.

The ULA-C proposal was withdrawn for a couple of years, but resurfaced early in 2007 as an active topic, and the discussion resumed as to whether it was truly necessary to have "assuredly unique" local uses addresses in addition to the self-selected ULAs that are already defined.

The consideration of this topic has quickly got to the essential question: Are these assuredly unique local use addresses useful or not?

ULA-C Addresses and Public Addresses

The concept of "uniqueness" is sometimes a poorly understood concept, and true uniqueness is neither easy nor cheap to maintain. Uniqueness is not a private property or an unequivocal attribute of an address. Uniqueness is a public assertion that has very overt displacement properties in both time and space. My public claim that the address I hold is "unique" displaces any claim you may want to make about the same address being unique to you. Now if I never make my claim public then you have no idea that I hold such a claim, and even if I dispute your claim, you have no basis to see whether my disputation is valid or not. So uniqueness is a public assertion relating to the association of an entity with an address. The public nature of the uniqueness assertion immediately brings into the realm of consideration the concept of validation of such assertion, and the use of trusted third parties in the form of registries. The issue of registries, registry behaviours, costs, policies and properties of a public registry are all relevant to this topic. If uniqueness is a public assertion about the properties of an association of an address and a holder then should the identity of the holder be a secret? What properties make sense in the context of uniqueness? How should such a unique local address registry be managed? Who should pay the costs of the operation of the registry? What registry behaviours are appropriate in this context? What policies are appropriate for entry to the registry?

At this point the essential difference between global unicast address space and these centrally assigned ULA's becomes fascinating. Given that there is no guarantees about local or global routability of unicast address space then what essentially is the difference between global unicast address space and ULA-C address space?

From the perspective of the address registry operator, what is the difference between ULA-C space and IPv6 Provider Independent space? There is essentially no difference in terms to registry functions, and that being the case, the cost to the address holder to maintain the registration entry in the registry that is an essential condition to ensure uniqueness is identical. The registry performs the same actions in maintaining a registry entry in both cases.

From the perspective of the end user what's the difference? It appears that ULA-C space has really very strange routing properties. In theory its not globally routeable, but given that there is no clear and coherent distinction between global and local routing then is a somewhat strange assertion. ULA-C addresses should not be used in certain contexts that we cannot coherently define. Obviously this does not appear to make a whole lot of sense! On the other hand unicast addresses have a similar amount of fuzziness - they are not assuredly routeable, whatever that may imply.

So what is the real motivation behind the ULA-C proposal? What "problems" do they solve in local contexts through the subtle distinction of being assuredly unique, as distinct from being probably unique? Are such problems of sufficient magnitude that they would justify the cost of setting up a complete address distribution framework and associated registry operation?

Some Further Questions about ULA-C's

Are ULA-C's really an effort to further dilute the concept of provider independent address space, as some sort of policy bypass operation? Is there a perceived problem with the Regional Internet Registries' IPv6 address allocation policies? Is the lack of IPv6 take up so far really a case of such restrictive address allocation policies that it has become necessary to create an alternative IPv6 address distribution channel to break the current logjam that is withholding IPv6 from its global destiny?

On the other hand, are the prevailing address distribution policies for IPv6 really very topical or useful any more? What is the objective of those policies? Are we still trying to use current address allocation policies to solve the 1999 routing explosion problem? To what extent do these address allocation policies place limitations and costs on addresses that are based on historical issues that are not an aspect of today's environment? Are ULA-C's a somewhat strange form of uniqueness-lite that is a reaction to the perception that the existing address registry function is too burdensome?

It seems to me that the entire picture behind ULA-Cs is one of confusion, mixed motives, unclear expectations and no clear and coherent concept of the problem that these local use addresses are intended to solve. This tends towards the conclusion that this is a classic case of application of the First Law of Holes (in case you haven't heard of this law, its pretty simple: if you are in one, stop digging!). Why are ULA-C's needed? What’s the true problem here? Are ULA-C's the most sensible response? Are there other potential responses? What are their relative merits and risks? Why do we need local use addresses that are assuredly unique in any case?

In trying to craft technical solutions we are often faced with the proposition that we are attempting to use technology to solve issues that are not necessarily technical in the first place. Is the entire issue about these assuredly unique local addresses really all about how good a job we’ve done so far with setting up an IPv6 address distribution framework that meets our apparent needs? Are there glaring deficiencies in the current framework of public address distribution for IPv6 that ULA-C’s can solve in a useful and productive manner?

So this leads me back to the original proposition – that technology design is often the outcome of entirely subjective decisions. The process can at times be entirely deterministic and logical while at other times we are placed into a position of having to make decisions which are very much value judgements. In the case of these centrally assigned Unique Local Addresses I suspect that there is no clear right or wrong answer but instead there are simply a collection of individual opinions. This is going to make any judgement of “consensus’, however rough, quite a tough call.

IETF Documents on IPv6 Site Local and ULA Addresses

- RFC1918 “Address Allocation for Private Internets”, Y. Rekhter el.al, February 1996.
- RFC3879, “Deprecating Site Local Addresses”, C. Huitema, B. Carpenter, September 2004.
- RFC4193, “Unique Local IPv6 Unicast Addresses”, R. Hinden, B. Haberman, October 2005.
- Internet Draft (draft-ietf-ipv6-ula-central-02.txt), “Centrally Assigned Local IPv6 Addresses”, R. Hinden, G. Huston, T. Narten, June 2007.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre, nor those of the Internet Society.

About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He 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 is currently the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of the Internet Society from 1992 until 2001.

http://www.potaroo.net