The ISP Column
A monthly column on things Internet

November 2010

Geoff Huston

Transitional Myths

I'd like to continue this mini-series of articles on IPv6. I attended the RIPE 61 meeting this month, and, not unexpectedly for a group that has some interest in IP addresses, the topic of IPv4 address exhaustion, and the related topic of the transition of the network to IPv6 has captured a lot of attention throughout the meeting. One session I found particularly interesting was one on the transition to IPv6, where folk related their experiences and perspectives on the forthcoming transition to IPv6.

I found the session interesting, as it exposed some commonly held beliefs about the transition to IPv6, so I'd like to share them here, and discuss a little about why I find them somewhat fanciful.

"We have many years for this transition"

No, I don't think we do!

The Internet is currently growing at a rate that consumes some 200 million IPv4 addresses every year, or 5% of the entire address IPv4 pool. This reflects an underlying growth of service deployment by the same order of magnitude of some hundreds of millions of new services activated per year. Throughout a dual stack transition all existing services will continue to require IPv4 addresses, and all new services will also require access to IPv4 addresses. The pool of unallocated addresses is predicted to exhaust in March 2011, and the RIRs will exhaust their local pools commencing late 2011 and through 2012. Once those pools exhaust, then all new Internet services will need access to IPv4 addresses as part of the IPv4 part of the dual stack environment, but at that point there are no more freely available addresses from the registries. Service providers have some local stocks of IPv4 addresses, but even those will not last for long.

As the network continues to grow the pressure to find the equivalent of a further 200 million or more IPv4 addresses each year will become acute, and at some point will be unsustainable. Even with the widespread use of NATs, and further incentives to recover all unused public address space, the inexorable pressure of growth will cause unsustainable pressures on the supply of addresses.

It's unlikely that we can sustain 10 more years of network growth using dual stack, so transition will need to happen faster than that. How about 5 years? Even then, at the higher level of growth forecasts, we will still need to flush out the equivalent of 1.5 billion IPv4 addresses from the existing user base to sustain a 5 year transition, and this seems to be a stretch target. A more realistic estimate of transition time, in terms of accessible IPv4 addresses from recovery operations, is in the 3 - 4 year timeframe, and no longer.

So no, we don't have many years for this transition. If we are careful, and a little bit lucky we'll have about four years.

"Its just a change of a protocol code. Users won't see any difference in the transition."

If only that were true!

In an open market environment scarcity is invariably reflected in price. For as long as this transition lasts this industry is going to have to equip new networks and new services with IPv4 addresses, and the greater the scarcity pressure on IPv4 addresses the greater the scarcity price of IPv4 addresses. Such a price escalation of an essential good is never a desirable outcome, and while there are a number of possible measures that can be taken that would be intended to mitigate, to some extent or other, the scarcity pressure and the attendant price escalation, there is still a reasonable expectation of some level of price pressure on IPv4 addresses as a direct outcome of scarcity pressure.

In addition, an ISP many not be able to rely solely on customer-owned and operated NATs to locally mask out some of the incremental costs of IPv4 address scarcity. It is likely, and increasingly so the longer the transition takes, that the ISP will also have to operate NATs. The attendant capital and operational costs of such additional network functionality will, ultimately be a cost that is borne by the service provider's customer base during the transition.

But its not just price that is impacted by this transition. The performance of the network may be impacted during the transition. Today a connection across the internet is typically made by using the DNS to translate a name to an equivalent IP address, then launching connection establishment packet (or the entire query in the case of UDP) to the address in question. But such an operation assumes a uniform single protocol. In a transition world you can no longer simply assume that everything is contactable via a single protocol, and it is necessary to extend the DNS query to two queries, one for IPv4 and one for IPv6. The client then needs to select which protocol to use if the DNS returns addresses in both protocols. And then there is the tricky issue of failover. If the initial packet fails to elicit a response within some parameter of retries and timeouts, then the client will attempt to connect using the other protocol with the same set of retries and timeouts. In a dual stack transitional world not only does failure take more time to recognise, but even partial failure make take time.

So users may see some changes in the Internet. They may be exposed to higher prices that reflect the higher costs of operating the service, and then may see some instances where the network simply starts to appear "sluggish" in response.

"NAT upon NAT upon NAT will work"

Maybe. But maybe not all of the time, and maybe not in ways that match what happens today.

The internet has been operating with a very prevalent model of a single level of address translation in the path for more than a decade now. Application designers now assume its existence, and also make some other rather critical assumptions, notably that the NAT is close to the client in a client / server world, and that there is a single NAT in the pat, and that its particular form of address translation behaviour can be determined with a number of probe tests. There is even a client-to-NAT protocol to assist certain applications to communicate port binding preferences to the local NAT. In a multi-level NAT world such assumptions do not directly translate, but its not necessarily the case that the application is aware of the added NATs in the end-to-end path.

However, its not just the added complexity of the multi-part NAT that presents challenges to applications. The NAT layering is intended to create an environment where a single IP address is dynamically shared across multiple clients, rather than being assigned to a single client at a time. Applications that make extensive use of parallelism by undertaking concurrent sessions require access to a large pool of available port addresses. Modern web browsers are a classic example of this form of behaviour. The multiple NAT model effectively shares a single address across multiple clients by using the port address, effectively placing the pool of port addresses under contention. The higher the density of port contention the greater the risk that this multiple layering of NATs has a visible impact on the operation of the application.

There is also a considerable investment in the area of logging and accountability where individual users of the network are recorded in the various log functions via their public side address. Sharing these public addresses across multiple clients at the same time, as is the intended outcome of a multi-layer NAT environment implies that the log function is now forced to record operations at the level of port usage and individual transactions. Not only does this have implications in terms of the load and volume of logged information, there is also a tangible increase in the level of potential back tracing of individual users' online activities if full port usage logging were to be instituted, with the attendant concerns that this represents an appropriate balance between accountability and traceability and personal privacy. Its also unclear whether there will be opportunity to have any such public debate on such a topic, given that the pressure to deploy multi-level NAT is already visible.

"Changing the Customer Premises Equipment (CPE) is easy"

No, not necessarily.

I think we have all seen many transition plans. Multi-level V4 NATs, NATS that perform protocol translation between IPv4 and IPv6, NATs plus tunnelling, as in Dual-Stack Lite, the IVI bi-direction mapping gateway, 6to4, 6RD, Teredo, to call up but a few of the various transitional technologies that have been proposed in recent times.

All approaches to dual stack transition necessarily make changes to some part of the network fabric, whether its changes to the end systems to include an Ipv6 protocol stack in addition to an IPv4 stack, or the addition of more NATs, or gateways into the network infrastructure. Of course, within a particular transitional model there is a selective choice as to what elements of the infrastructure are susceptible to change and what elements are resistant to change. Some models of transition, such as 6RD and Dual-Stack Lite assume that changing the CPE is easy and straightforward, or at least that such a broad set of upgrades to customer equipment is logistically and economically feasible. 6RD contains an implicit assumption that the network operator has no economic motivation to alter the network elements, and wishes to retain a single protocol infrastructure that uses IPv4.

Where the CPE is owned, operated and remotely maintained by the service provider, upgrading the image on the CPE might present fewer obstacles than upgrading others elements of the network infrastructure, such as broadband remote access servers that operate in a single protocol mode, but sweeping generalisations in this industry are unreliable. Service providers tend to operate customized cost models, and appear to be operating with specialised mixes of vendor equipment and operational support systems. For this reason operators tend to have differing perspectives on what component of their network is more malleable, and correspondingly have differing perspectives on which particular transition technology suits their particular environment.

This is a volume-based industry, where an underlying homogeneity of the deployed technology, coupled with economies of scale and precision of process are key components of reliable and cost efficient rollouts. It is somewhat unexpected to see this transition expose a relative high degree of customisation and diversity in network service environments.

"My ISP has enough IPv4 addresses to last for years, so they don't have a problem"

Well, not necessarily.

The assumption behind this statement is that everyone else is also able to persist with IPv4, and everyone you wish to reach, and every service point you wish to access, will maintain some form of connectivity in IPv4 indefinitely.

But this is not necessarily the case. At the point in time when a significant number of clients or services cannot be adequately supported on IPv4, then irrespective of how much IPv4 your ISPs,

they will need to provide their clients with IPv6 in order to reach these IPv6-only services. This is a network, and it contains network effects where the actions of others directly effects your own local actions. So if you believe that you need do nothing, and use an IPv4 service for years into the future, then this position will be inadequate at the point in time when a significant number of others encounter critical levels of scarcity such they are incapable of sustaining the Ipv4 side of a dual stack deployment, and are forced to deploy an Ipv6-only service. The greater the level of address hoarding the greater the level of pressure to deploy IPv6-only services on the part of those service providers who are badly placed in terms of access to Ipv4 addresses.

"We will always have to run IPv4 protocols"

Probably not.

Or at least not in terms and volumes that are significant to the industry over the forthcoming decades. Protocols do die. DECnet, or SNA for that matter, no longer exist as widely deployed networking protocols. In particular, networking in the public space is all about any-to-any connectivity, and to support this we need a common protocol foundation. In terms of the dynamics of transition, this is more about tipping points of the mass of the market than it is about sustained coexistence of diverse protocols. Once a new technology, or, in this case, protocol, achieves a critical level of adoption the momentum switches from resisting the change to embracing it.

The aftermath of such transitions does not leave a legacy of enduring demand for the superseded technology. As difficult as it is to foresee today, once the industry acknowledges that the new technology achieves this critical mass of adoption, the dynamics of the networking effect propels the industry into a topping point where the remainder of the transition is likely to be both inevitable and comprehensive. The likely outcome of this situation is that there is no residual significant level of demand for IPv4.

"There is a technology that will translate between IPv4 to IPv6"

Yes, but.

Such a technology effectively maps between IPv4 and IPv6 addresses. One approach, IVI provides a 1:1 mapping by embedding fields of one address in the other. Another approach, originally termed NAT-PT, uses a mapping table in a similar fashion to a conventional NAT unit. The common constraint here is that of there are no IPv4 addresses, then such a bidirectional mapping cannot be sustained in each approach Ultimately, every packet that traverses the public Internet requires public address values in the source and destination fields, and is the task of to provide a protocol bridge between Ipv4 and IPv6, then public IPv4 addresses are required to support the task.

But its not just the requirement for continued access to addresses that is the critical issue here. A reading of RFC4966, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status" should curb any untoward enthusiasm that this approach is capable of sustaining the entire load of this dual stack transition without any further implications or issues.

"We don't necessarily have to transition to IPv6. There are substitutes"

Nothing is visible from here!

If we want to continue to operate a network at the price, performance and functional flexibility that is offered by packet switched networks, then the search for alternatives to IPv6 is necessarily constrained to a set of technologies that offer approaches that are, at a suitably abstract level, isomorphic to IP. But from abstract observations to a specific protocol design is never a fast or easy process, and the lessons from the genesis of both IPv4 and IPv6 point to a period of many years of design and progressive refinement to come up with a viable approach. In our current

context any such re-design is not a viable alternative to IPv6, given the timeframe IPv4 address exhaustion. Its unlikely that such an effort would elicit a substitute to IPv6, and its more likely that such an effort may lead towards an inevitable successor to IPv6, if we dare to contemplate networking technologies further into the future.

Other approaches exist, based around application level gateways and similar forms of mapping of services from one network domain. We've been there before in the chaotic jumble of networks and services that defined much of the 1980's, and it's a past that I for one find easier to forget! Such an outcome is of considerably higher complexity, considerably less secure, harder to use, more expensive to operate and more resistant to scaling.

Like it or not the pragmatic observation of today's situation is that we don't have a viable choice here. There are no viable substitutes.

"We know what's happening"

I'm not sure that's universally true! The observations I've heard about the current situation lead me to the observation that there are many different perspectives on the situation. Each individual perspective sees the transition in terms that relate to their own circumstances and their own limitations, and a more encompassing perspective of the entire Internet and this transition is harder to assemble. So, from the perspective of the Internet as a whole, no, we are not really aware of what's happening.

"We know what we're doing"

Individually this is, hopefully true. But at the level of the entirety of the Internet, then no, we don't really have a clear perspective of this transition.

"We have a plan!"

See above.

"The Internet will be fine!"

I'm unsure about this one.

The worrying observation is that the Internet has so far thrived on diversity and competition. We've seen constant innovation and evolution on the Internet, and the entrance of new services and new service providers.

But if we rely solely on IPv4 for the future Internet, then this level of competition and diversity will be extremely challenging to sustain. If we lose that impetus of competitive pressure from innovation and creativity, then the Internet will likely stagnate under the oppression of brutal volume economics. The risks of monopoly formation under such conditions are relatively high.

There is one observation I heard at the RIPE session that I hope will be a myth, as this transition gets underway:

"The incumbents will have all the IPv4 space. Thanks for playing"

If that's not a myth, then we are going to be in serious trouble!

Disclaimer

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

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. 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 was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

www.potaroo.net