June 2020
Geoff Huston

Technology Adoption in the Internet

How are new technologies adopted in the Internet? What drives adoption? What impedes adoption? These were the questions posed at a panel session at the recent EuroDiG workshop in June.

In many ways this is an uncomfortable question for the Internet, given the uncontrolled runaway success of the Internet in its first two decades. The IPv4 Internet was deployed about as quickly as capital, expertise and resources could be bought to bear on the problem, and the Internet's expansion appeared to be a case of being driven by demand pull. Whether it was bandwidth, content, or services, they were quickly saturated soon after they were deployed by an enthusiastic consumer base that appeared to have in insatiable demand. Perhaps naively, this bred a reputation for the Internet's infallibility, and there was a confident expectation that all Internet technologies would enjoy a similar enthusiastic reception.

However, this is not the case, and in this article, I'd like to look at a couple of technologies that have not been instant runaway success cases, and then look at likely reasons for this.

IPv6

Originally specified in RFC 1883, published in December 1995, following a frantic five-year developmental effort in the IETF, IPv6 did not enjoy a runaway level of success in terms of deployment. IPv6 was designed in response to the prospect of IPv4 address exhaustion, where even before the Internet took off as a consumer and enterprise product it was clear that the number of connected devices would rapidly exceed the number of unique addresses in the IPv4 protocol. The fundamental premise of IPv6 was "more addresses" and in the case of IPv6 that's just about all it did. In most respects it was a modest refinement to the IPv4 model.

Our expectations were that IPv6 deployment would be propelled by the prospect of address exhaustion in IPv4. The future risk of address exhaustion should've motivated industry actors to develop and deploy IPv6, and the risk of being caught out would ensure that the transition to IPv6 would be complete long before new handed out our last IPv4 address. After all the IT industry had been preparing itself for the Y2K date roll for more than 10 years, so it appeared that a similar abundance of caution would apply to the prospect of IPv4 address exhaustion. The address barrel would never get to empty as we would've moved on to IPv6 well before such a calamitous event could ever occur!

Obviously, we've chosen to take a different path. We really don't know how many devices are connected to the Internet today. Estimates from various industry sources range from 10 billion to 50 billion (which is an impressive level of uncertainty!). The IPv4 routing table advertises 2.8 billion addresses. That means that in one respect we have already achieved the impossible and stuffed many more devices into the Internet than we have addresses. Obviously, the combination of client/server architectures and address sharing technologies have played a big role here and we have confronted address exhaustion and worked around it.

If we look at the last nine years of deployment of IPv6 in the Internet (Figure 1) the overall picture shows an overall trend of adoption of IPv6 in the public Internet, but the picture is by no means one of rapid enthusiastic adoption. Both in 2018 and in 2020 the IPv6 deployment momentum has all but stopped, yet the overall use of IPv6 is yet to exceed one quarter of the user population.

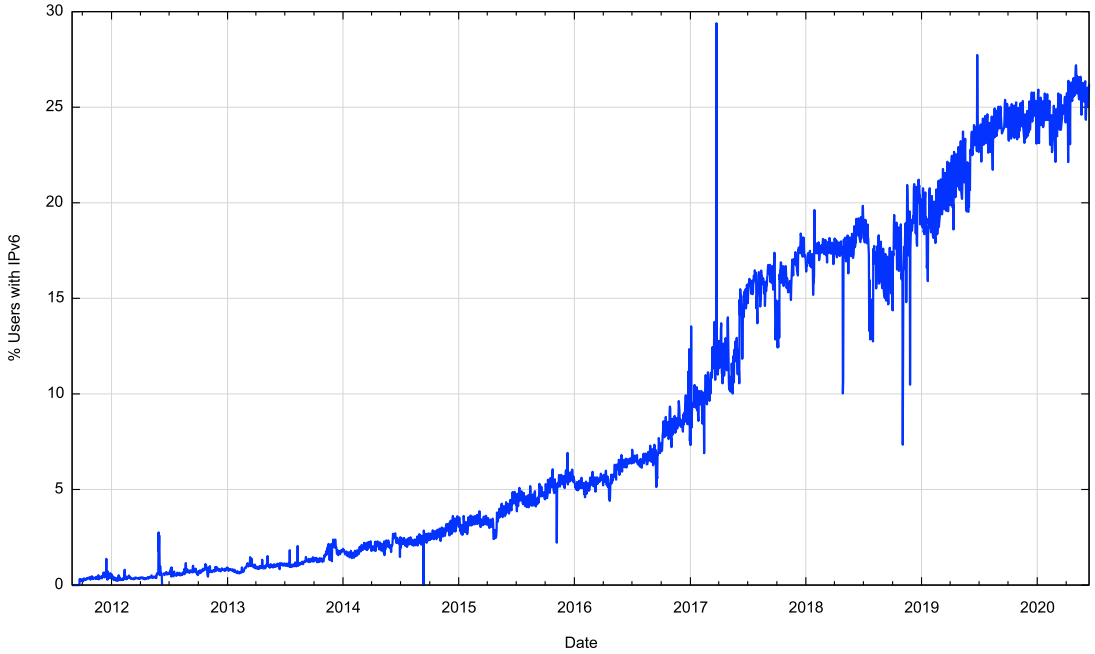

Figure 1 – Adoption of IPv6, 2012 to the present (https://stats.labs.apnic.net/ipv6/XA)

So if address exhaustion is meant to be the driver for IPv6 adoption then clearly this has not proved to be the case, and rather than being experts on IPv6 deployment, we are well and truly experts in stretching out the IPv4 address plant to outperform anything we could've envisaged in the past!

DNSSEC

Originally specified in RFC 2065, published in January 1997, the specification of a security framework for the DNS was considered to be a vital part of the larger security framework for the Internet. Given that the DNS resolution protocol operated in the clear and made extensive use of intermediate agents (recursive resolvers), it was considered essential to be able to trust that the answers provided by the DNS were genuine. Even today it is probably still an essential attribute of a trusted network, but somehow we have resigned ourselves to a DNS infrastructure that fails to achieve this and the DNS interfered with to an extent that can only be described as somewhere between prolific to universal!

DNSSEC comes in two parts: production and consumption. On the production side DNSSEC relies on DNS zones being signed and registries managing delegation signing records in a similar manner to name delegation records. DNSSEC uses the delegation hierarchy so the entire signing system "locked" into place with the signing of the root zone in July 2010. The consumption side requires clients of the DNS to request a digital signature of a DNS response, and then verify this signature by generating a validation chain of interlocking signatures back to the root key.

In theory end users should perform their own validation, as there are many risks associated with outsourcing security functions. In practice this does not happen today and DNSSEC validation is a function performed by recursive resolvers. If a validating recursive resolver cannot validate a DNS signed response it will not return the response, but instead indicate a DNS error. If a client (stub resolver) exclusively uses validating resolvers, then it will be unable to resolve DNS names where the signature cannot be validated.

We can look at the proportion of Internet users who exclusively use resolvers that perform DNSSEC validation over the past 6 1/2 years (Figure 2). Coincidentally the level of adoption of DNSSEC validation is also 25% of the user base, but the trajectory of adoption is entirely different. DNSSEC use plateaued in 2016 and fell across 2017, resuming an upward trajectory in 2018.

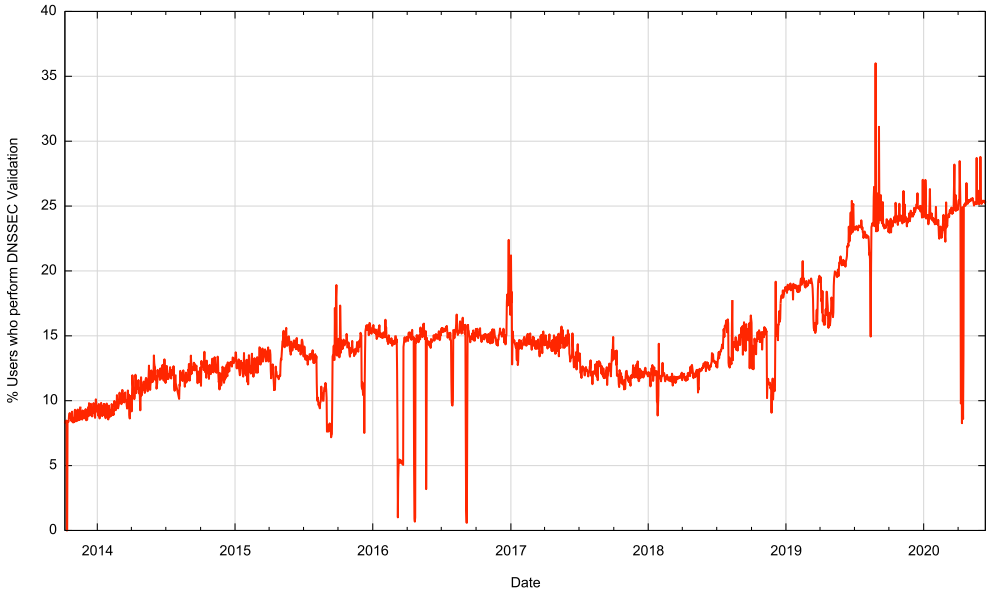

Figure 2 – Adoption of DNSSEC, 2013 to the present (https://stats.labs.apnic.net/dnssec/XA)

IPv6 and DNSSEC Adoption in Europe

It seems that while IPv6 and DNSSEC enjoy similar levels of end user adoption today, the path to get to this situation differs markedly. This leads to the suspicion that in terms of technology adoption we don't all react to the same environmental signals in the same way, and the adoption of the kinds of technologies is actually far more piecemeal than we might've suspected.

Let's compare the landscape in Europe for IPv6 adoption to that of DNSSEC adoption to illustrate this point.

IPv6 is well established in the major consumer access networks in Belgium, Germany, Greece, France and Switzerland. It is not well established in Spain, Italy, Sweden, Denmark, and Slovakia. The total European IPv6 deployment level is 20%, some 5% below the internet-wide level. (Figure 3)

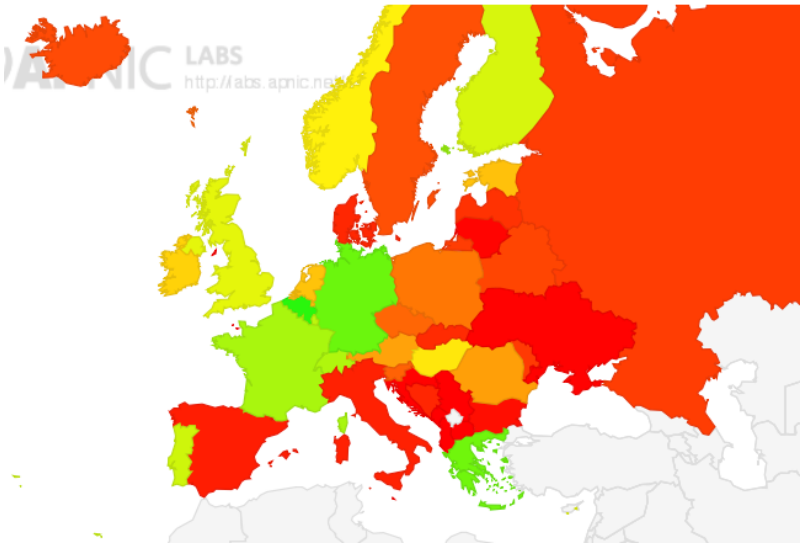

Figure 3 – IPv6 Adoption in Europe

European deployment of DNSSEC validation is different (Figure 4). This technology is extensively deployed in Sweden Norway, Denmark, Finland and Iceland. It is also deployed in Czechia and Switzerland. The European average deployment level is 30%, some 5% higher than the Internet-wide average value.

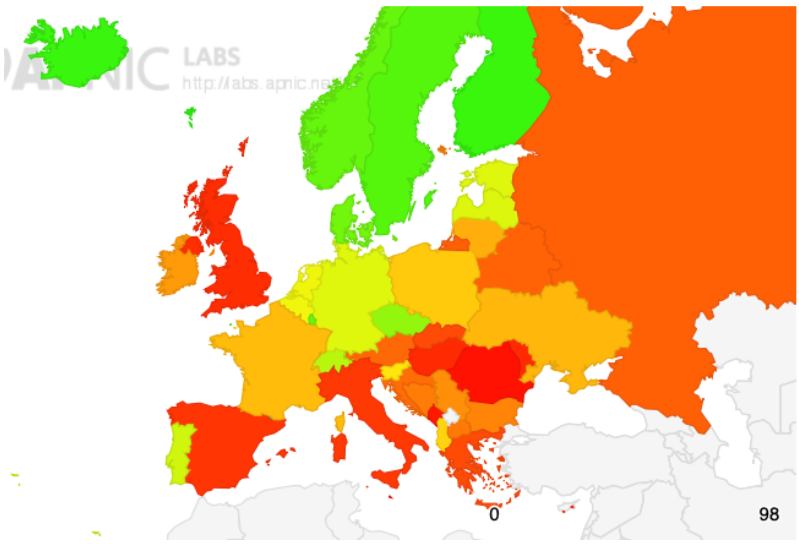

Figure 4 – DNSSEC Adoption in Europe

Why is Switzerland the only economy in both of these lists? Why so some providers choose to deploy IPv6 and others choose DNSSEC validation?

And perhaps more importantly, why is there such obvious diversity in the response to these technologies? In both cases the end goal is comprehensive deployment by all providers, yet in different locales the local network service operators appear to have made different decisions regarding technology deployment. This is in spite of the fact that there is a much greater degree of uniformity in the end user profiles, both in terms of the applications used and the platforms used.

Diversity of Adoption

This technology adoption diversity is not just endemic to Europe. We see a similar diversity in all parts of the Internet. What is going on?

Part of the issue here is that neither of these technologies gives an early adopter an obvious competitive advantage. In the case of the IPv6 transition there are two distinct phases. The first is piecemeal adoption of dual stack services, where some services and some clients can use both IPv4 and IPv6, while others still use IPv4 only. The second phase is the transition to IPv6 only services, but this phase is only viable when the extent of dual stack deployment is close to universal. At that point every major service and every major client sector can speak IPv6 and there no longer any benefit in support IPv4. But until that point is reached, we are still in a dual stack environment, and the support of IPv4 is close to mandatory. So it's the late adopters that determine the overall timetable of transition, and in the case of IPv6 transition this initial dual stack phase has been operating for almost twenty five years, and the early adopters are still waiting for the later adopters to move.

The second factor relates more to DNSSEC, and that is the that the economics of security tend not to favour early adopters. Spending resources to reduce risks in one the one hand a prudent measure, but determining how much is prudent to spend on risk mitigation depends on the quantification of the risk itself. We are notoriously bad at quantification of risk, and typically underestimate risk. It is also unclear how actually carries the risk in any case. If a resolver does not enable DNSSEC validation that that increase its liability? If it does enable DNSSEC validation has it actually prevented its clients

from being deceived? The original form of cache poising attack that DNSSEC mitigates is a relatively esoteric form of attack and this makes the quantification of risk and benefit quite uncertain.

There are a number of environmental factors that add to the uncertainty in this space.

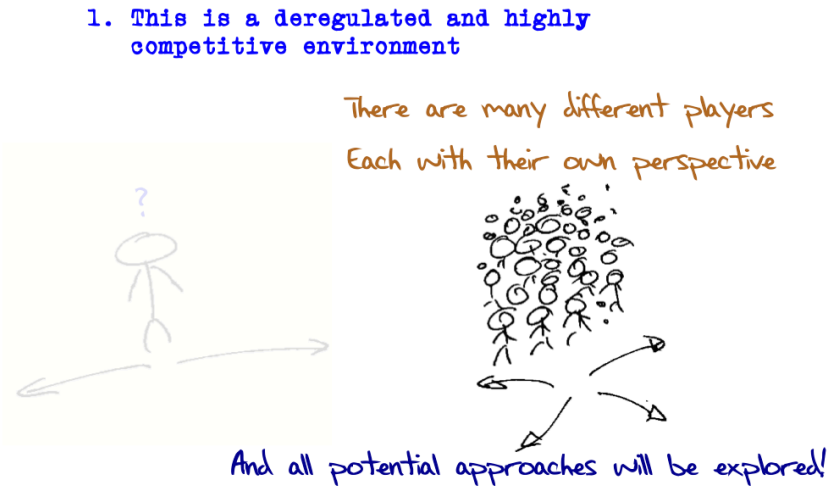

The first is that this is a largely deregulated activity driven by private capital investment in a competitive and relatively uncertain environment. There are many different actors in the Internet space. Some have short term plans to build up a business that will be purchased by a larger player, while others are determined to create a long-term market position. Some actors are now the largest enterprises on the planet, while others operate in niche market opportunities. They will all react differently to a given situation. The DNS Wars issue (https://www.potaroo.net/ispcol/2019-11/dnswars.html) is a good example of some of the drivers behind DNSSEC validation and why Google's Public DNS services make so much sense for Google's core business as a DNSSEC-validating open recursive resolver. Given such diversity in the market there is no reason to believe that all market actors will react in the same way. They evidently react quite differently.

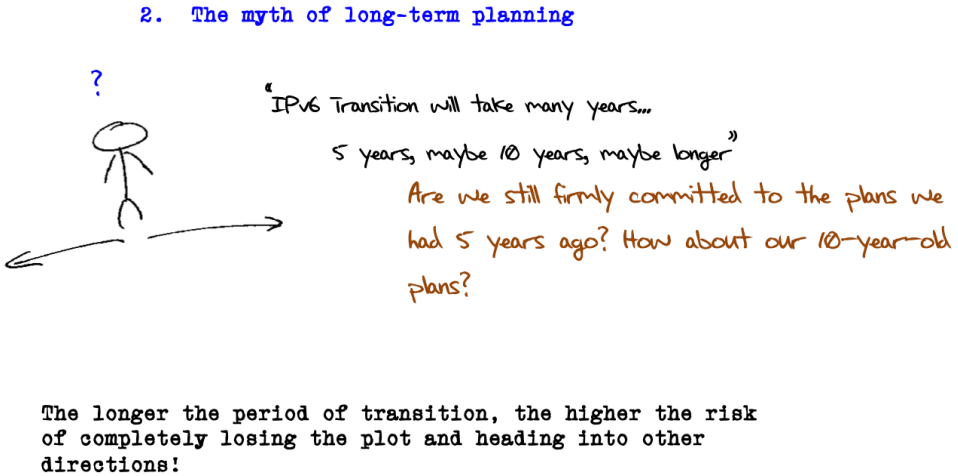

The second factor is that this is still business that is reactive rather than deliberative. The pace of technology change continues to unsettle incumbents, and this works again most forms of long-term business plans. This means that no actor has the luxury of working to long term plans. We have gone through a number of quite fundamental changes in the first two decades of this century, including mobile services, content services, social networking and streamers. It's impossible to be certain here but continued technical innovation and consequent disruption is likely. The result is that nobody is left alone to execute long term plans. When we hear the confident prediction that Ipv6 transition will

take years or even decades to complete is worth bearing in mind that none of remember the long term plans we had devices just ten years ago!

3. The Internet keeps changing

Today's Internet Architecture

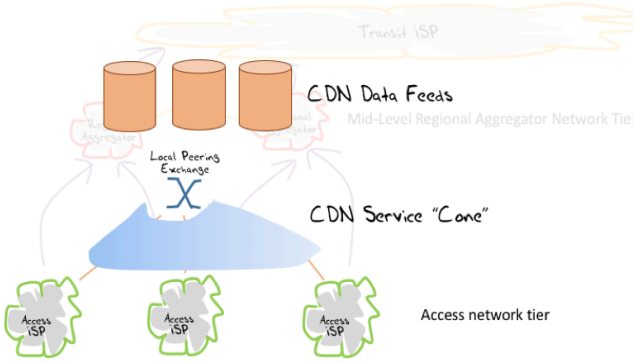

The third factor is that the Internet itself keep in changing. The original model of the internet closely resembled the model telephone network. It was a passive packet carriage platform designed to allow connected devices to communicate with each other. That's not what we have today. Today we have a content distribution network where connected clients negotiate service delivery with inbuilt service delivery platforms. Clients don't and can't talk with other clients.

What can this tell us about the dynamic of technology adoption?

Some market actors see relative advantage in early adoption. They may see cost efficiencies, or relative competitive advantage. They might perceive the technology as assisting them to defend their core activities from competitive erosion. They may see advantage in supporting services that offer the perception of enhanced utility, security and safety.

On the other hand, others may see equally persuasive reasons to wait. After there are alternate interpretations of these technologies that can sustain a case of waiting and watching. IPv6 is a 1990's technology solution to a 1980's network architectural issue. Content Feeder networks do not necessarily require persistent globally unique address schemes and on-demand addressing seems to work perfectly fine for client / server interactions. DNSSEC is a mess. If DNSSEC validation was pushed all the way to the client edge of the network we are worried that it will make the DNS tediously, unacceptably slow. The rather unfortunate fate of DANE graphically illustrates this situation. More recent examples, such as Route Origin Filtering are difficult to conceive as solutions to a real risk. Origin filtering only makes prefix hijacking only marginally more difficult for the attacker while introducing a new set of technologies with their own fragilities and operational costs.

Part of the strength of the Internet lies in the decoupled nature of the network's infrastructure, where many component service providers operate within their chosen niche of activity, and the overall orchestration of the collective efforts is left to market forces. No one is in charge. But while this is a strength it can also be a weakness, particularly in cases of cost displacement. In a centrally orchestrated environment, the costs and benefits could be directly compared, and such solutions could be deployed where it was cost-beneficial to do so. However, without such orchestration there is little in the way of incentive for individual actors to go down this path. Ideally every actor wants to retain benefit and transfer costs to others. The somewhat incoherent result is what we have today.

But that is not necessarily a poor outcome. This diversity is in and of itself a strength and the efforts of incumbents to impose some form of stasis is being constantly challenged by others who see

technology innovation as a means of leveraging relative advantage. And in all of this change does happen. Don't forget that during the period that has seen the protracted sagas of IPv6, of DNSSEC and even Routing Security absorb attention and energy we've also seen the Internet completely transform itself a number of times. We've seen the rapid rise of the mobile Internet, we've seen the rise of CDNs and content streamers, and we've seen the inexorable coming of Internet of Things (for better or worse!). None of these changes were protracted exercises in procrastination. None of these changes was even a debate within the industry. They happened because users wanted such services and were willing to pay for them. And maybe that's all that really matters!

Disclaimer

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

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net