The Hidden Standards War:

Economic Factors Affecting IPv6 Deployment

Brenden Kuerbis and Milton Mueller, Internet Governance Project, Georgia Institute of Technology, School of Public Policy

Research Project # 138078
Project sponsor: Internet Corporation for Assigned Names and Numbers

Final draft (February 2019)

Preface

This report was commissioned by ICANN's Office of the Chief Technology Officer (OCTO). The motivation for the project was our shared conviction that the Internet community needs a better understanding of the economic factors affecting the transition to Internet Protocol version 6. In July 2018, ICANN's OCTO provided some funding for the project and facilitated access to data. We note that ICANN did not edit or control the content of the report; IGP's researchers were free to draw their own conclusions. IGP is fully responsible for the content of the report.

Table of Contents

The Hidden Standards War:

Economic Factors Affecting IPv6 Deployment

Introduction: Standards Migration...or Competition?

In the period between February 2011 and July 2018, the world outgrew a networking standard that underlies what we call 'the Internet.'[1] The 1981-vintage protocol that makes the internet work, known as Internet Protocol version 4 (IPv4), had no more unallocated, globally unique numbers that could be used for IP addresses. The original Internet's 32-bit address space had all been given out.

More than 20 years earlier, the Internet engineering community had anticipated this problem. They developed a new, "next generation" Internet Protocol, with a much larger address space of 128 bits. Unfortunately, their design process was unable to make this new internet protocol backwards compatible with the older version. The Internet community was then faced with the problem of "migrating" the entire Internet to the new standard, known as IP version 6 (IPv6).

Despite more than a decade of evangelism on behalf of IPv6, the Internet community has not yet succeeded in moving the world to the new networking standard. IPv6 is being deployed by some network operators, and its overall share in measurements is gradually increasing. But many network operators and enterprises are not deploying it. Only about 20% of the Internet traffic we measure comes from machines equipped to speak IPv6 to each other. Seemingly endless discoveries of things that don't work when one uses IPv6 have forced some deployments to retreat or slow down.[2] The overall status of the transition is unclear. The data we present in this report show uneven IPv6 adoption rates, and the data do not follow simple or predictable patterns.

In the meantime, technologies like Network Address Translation (NAT)[3] and secondary markets for IPv4 numbers have introduced major economic efficiencies that have extended the life of IPv4. NAT has evolved from a technology deployed at the periphery to one used throughout some operators' networks. The secondary market has incentivized the discovery of massive stocks of unused or underutilized IPv4 number blocks, which are now being traded in a vigorous market.

It is not really correct, then, to call the IPv4 - IPv6 transition a "migration," if by migration we mean an orderly and sponsored movement from one software generation to another within a defined time period. This is more like a standards *competition,* in which network operators make a choice based

---

[1] That time period begins with the final /8 allocations made by the IANA to regional Internet registries and ends with the exhaustion of AFRINIC's unallocated IPv4 address pool. Between that period RIPE, LACNIC, APNIC and ARIN exhausted their unreserved number blocks, although all hold small reserves of IPv4 numbers to facilitate interconnection between IPv4 and IPv6.

[2] In September 2017 *The Register* reported that Microsoft staff discovered that 99% of attendees were unable to connect to their corporate VPNs at a conference where the local network was running IPv6 only on WiFi. The same month Microsoft's Veronica McKillop blogged that "applications are the big unknown, not just our own but the third-party applications that often claim 'IPv6 compatible;' however when it comes to a real deployment, the experience is quite different."

[3] NAT utilizes private address spaces (RFC 1918) and interconnects them to the Internet via shared, globally routed IPv4 addresses. By sharing a few globally routable addresses, NAT economizes on IPv4 numbers.

on which standard they consider more efficient and productive. It is an economic battle between two paths of technical evolution: a legacy standard and its putative successor.

The Internet community, industry and policy makers need to understand better the dynamics of this competition, and its implications for the future of the internet. We may be headed to an Internet held together and built around two incompatible protocol suites, creating a global internet only when both of them are run in parallel or they go through some translation point.

Is this mixed-standard Internet a passing phenomenon, or will we get stuck in it? If it is only a transitional phase of a standards war and one will prevail, which one will it be? If IPv6 prevails, how long will it take us to get there? Is it possible that IPv6 actually loses the standards competition, and becomes the proverbial 'orphan' of the standards economics literature? This report tries to answer those questions.

The report examines the economic incentives affecting the competition between IPv4 and IPv6, and quantitative data about current levels and patterns of IPv6 adoption. Three research questions motivate this study:

- What economic incentives affect decisions by network operators to deploy IPv6?
- What factors can best explain the observed levels of IPv6 adoption at the country level?
- How do translation and tunneling technologies, which serve as a bridge between incompatible IPv4 and IPv6 standards, alter the economic incentives to remain with IPv4 or deploy IPv6?

We approached these questions from both a microeconomic and a macro-social perspective. In our macro-level analysis, the unit of analysis is the country and the research relates different IPv6 deployment levels to socio-economic variables, such as country-level wealth measured in Gross Domestic Product (GDP) and a measure of market concentration, Herfindahl-Hirschman Index (HHI). In our microeconomic analysis, the unit of analysis is the network operator. We examine how efficiency and cost decisions at the operator level affect IPv6 deployment decisions, including: the cost of IPv6 deployment; the cost of acquiring IPv4 addresses on the secondary market; the cost of running IPv4 NAT; the relative cost and scalability of IPv4 - IPv6 converter technologies.

1. The economics of network migration

In this section, we explain the relevant economic theories and apply them to the case of the IPv6 transition. We focus in particular on the role of *network effects* on the competition between IPv4 and IPv6, which we believe current literature has not analyzed correctly. We also highlight the important role that the value of *network growth* plays in the transition. Our analysis draws on the economics literature on network externalities, compatibility, standards competition, and converter technologies.

1.1 The network externality

Certain products and services have what are commonly called network effects. This means that the value of a product or service to one person depends on how many other people (or places, or things) use the *same* product or service. Initially modelled as "interdependent demand" (Rohlfs, 1974), the network effect arises from a demand-side economy of scope. (Economides, 1996) A network is a gigantic bundle of connections, and each connection in that bundle is, from an economic point of view, a separate product or service. A user who makes an investment in a particular technology, standard or protocol gains access to a wider scope of connections - a bigger bundle - when others join the same network service, and/or adopt compatible technology. This is a scope economy, not a scale economy, because the additional connections add *different products* to the network bundle (broaden its scope) rather than increasing the quantity of a homogeneous output (its scale). With network effects, the efficiencies of increasing scope occur on the demand side, not on the supply side. Even if the supply side costs per unit increase as a network broadens its scope, powerful network benefits on the demand side can support growth. (Mueller, 1997)

The presence of demand-side scope economies gives the market process unique characteristics. Equilibrium is highly path-dependent (Arthur, 1994); it really matters who joins a network and in what sequence. Start-up networks must reach a critical mass before the network benefits are realized, which may require subsidizing early adopters. (Rohlfs, 1974) In a competition between networks that are not connected, or are technically incompatible, the market is "tippy;" that is, once it becomes clear that a certain standard is gaining a decisive edge, actors will converge on the perceived winning standard in order to maximize their network benefits. (Farrell and Saloner, 1985; Katz & Shapiro, 1986) This produces a tendency towards winner-take all outcomes. (Shapiro and Varian, 1999)

Once a large number of actors converge on a single platform or protocol, a condition economists have called inertia or "lock-in" sets in. (Farrell and Saloner, 1986) Widespread compatibility across the entire market raises the costs of switching to a new standard that is not compatible with the old one. Few actors can afford to sacrifice the tremendous value of widespread compatibility by abandoning the established standard for a new one, and the costs of coordinating large numbers of users to change standards at the same time are very high.

The competition between IPv4 and IPv6 is strongly influenced by network effects - but mainly by inertia. The world converged on IPv4 in the mid-1990s. Now that it has created compatible data communications amongst billions of hosts and applications everywhere in the world, it is virtually impossible for any organization to abandon IPv4 altogether for a new, incompatible data communications protocol. Thus, the incompatibility between IPv4 and IPv6 has major consequences for IPv6 deployments.

1.2 Converters and gateway technologies

In industries with strong network effects, adoption decisions and competitive outcomes can be strongly affected by interconnection of competing systems, or by technologies that act as converters or bridges between incompatible technologies. There is an extensive economics literature on the impact of converters and interconnection on actor incentives and market structure (e.g., Farrell and Saloner, 1992; Choi, 1996, 1997; Seifert and Vare, 2009; Sen, Jin, Guérin, & Hosanagar, 2010). Converter technologies or interconnection can neutralize network effects among competing networks or standards.

Prior modeling and empirical work shows that converter technologies are double-edged in their effects. While they broaden compatibility and allow users of both systems or technologies to reap the same network benefits, they can also eliminate, reduce or delay an end user's incentive to migrate from an old network/standard to a newer one. (Joseph, Shetty, Shuang, & Stoica, 2007) If the alternate networks are not connected or are using incompatible technologies, then interconnection or a technology that makes them compatible eliminates the competitive advantage a network provider gains from expanding its scope. (Mueller, 1997; Hazlett, 2005)

Users can also bridge incompatible technologies by duplicate use of both at the same time, which is sometimes called "dual homing." (Parker, Van Alstyne, Choudary, & Foster, 2016) In the IPv4-IPv6 context, this is called "dual stack."

Translation or tunneling technologies can also bridge the compatibility gap between IPv4 and IPv6 without dual stack. The cost structure and economic incentive effects of these converter technologies for IPv4-IPv6 are not well understood, however. One of the contributions of this report is to explore that topic in greater depth.

1.3 The historical background of the IPv4 - IPv6 competition

The possibility of a crippling shortage of IPv4 address numbers became a concern in the IETF as early as 1990. By mid-1992 the IETF was debating the merits of various "Next Generation Internet Protocol" (IPng) designs with a bigger address space. (DeNardis, 2009) Even at that early date, the relation between IPv4 and IPng was understood by some to be a standards *competition,* one in which economic incentives would play a major role. In comments to the IETF, for example, an engineer from BBN warned that "As currently envisioned, IPng may not be ambitious enough in the delivery of new capabilities to compete against IPv4 and the inevitable arrival of network address translation devices." (Curran, RFC 1669, 1994). IBM warned that IPv4 users won't upgrade to IPng without a compelling business reason (cited in DeNardis, 2009, p. 56). Nevertheless, the IETF pushed ahead: at its July 1993 meeting, an IETF session decided that it "needed to take decisive action" to select an IPng, and "any option of letting the market decide was unacceptable." (cited in DeNardis, 2009, p. 51)

In January 1995, the description of the leading IP next generation proposal in RFC 1752, known as Simple Internet Protocol Plus (SIPP), claimed that "it can be installed as a normal software upgrade in Internet devices and is interoperable with the current IPv4." However, further study and development made it clear to the designers that there was no known method to make the differently-sized address headers in the packet interoperate.[4] As IPv6 was finalized, the goal of maintaining backwards compatibility with IPv4 was discarded.

Given the absence of backwards compatibility, early understandings of the migration process envisioned a reliance on what came to known as the *dual stack* strategy.[5] Hosts and routers would run both protocols, IPv4 and IPv6, to allow those who had not converted to IPv6 to communicate with those who had. As the new protocol gained acceptance it would become possible to turn off the older protocol. In Section 1.4 below we explain why dual stack as originally conceived was not an economically feasible migration path. But IETF's dual stack migration strategy may not have been as flawed then as it seems now. In 1994, the Internet was still relatively small and was growing exponentially. There were still multiple, competing layer 3 networking protocols in existence and a lot of change and diversity in their adoption. Many enterprises had used a dual stack method to move to IPv4 from proprietary data networking protocols such as Novell Netware. However, this had worked on a comparatively small scale. No networking protocol from the early 1990s had the large installed base and global dominance that IPv4 did from 1998 on.

As is well-understood in the economics literature, when strong network externalities are present timing is everything. Adoption decisions are path-dependent and convergence on a common protocol is not reversible as it creates inertia and lock-in. Had IPv6 been rapidly deployed from 1994 to 1996 the dual stack strategy might have worked, although as BBN and corporate users had warned, there was no real economic motivation to adopt it. So unsurprisingly, it was IPv4, not IPv6, that spread like wildfire during those years. Furthermore, the immediate threat of address runout that was supposed to motivate ISPs to switch to IPv6 had been substantially mitigated by Classless Inter-Domain Routing (CIDR) and by a move to more conservative initial number block allocation policies by the newly-formed Internet registries. By 1998, IPv4 had become too embedded for a rapid rollout of IPv6 to displace IPv4.

In these early conceptions of the transition, and in later ones (RFC 4213, 2005), configured tunneling was also proposed as a means of maintaining compatibility. Tunneling involves the use of protocols to encapsulate the payload of a packet within the header of another packet. An example would be encapsulating IPv6 packets within IPv4 headers to establish point-to-point tunnels through IPv4 routing infrastructures (or vice-versa). Inefficiency and compatibility problems with some forms of tunneling, however, seem to have prevented widespread reliance on it as a transitional mechanism. For instance, earlier tunneling protocols (e.g., Teredo and 6rd, defined in RFC 4380, RFC 5569 and RFC 5969) encapsulated IPv6 packets and transited them over IPv4 networks, but had mixed results when going through Network Address Translation (NAT) devices, which were coming into widespread use as numbers became more scarce. More importantly, just as with dual stack these protocols did not diminish demand for IPv4 addresses as networks grew. At best, tunneling was a tool to bypass a local problem, not a migration strategy. More recent approaches (e.g., Dual-Stack Lite and Lightweight 4over6, defined in RFC 6333 and RFC 7596, respectively) have combined encapsulation with carrier-grade NAT in a way that reduces demand for globally routable IPv4 addresses within operators' networks.

[4] Discussion with Alain Durand, January 2019.

[5] See, e.g., Robert E. Gilligan, Simple Internet Transition Overview, draft-gilligan-ipv6-sit-overview-01, November 1994.

1.4 Problems with dual stack migration

From an economic perspective, the flaws in the pure dual stack strategy are obvious, and fatal. Dual stack adds significant costs to network operators without any offsetting benefits. The initial decision to deploy IPv6 incurs some fixed, one time costs in infrastructure, learning and training. But the operator must continue to run IPv4, which means additional operational costs due to the need to duplicate firewall rules and other configurations. Additionally, implementers are likely to encounter unanticipated incompatibilities at the application layer, and sometimes the transport or data link layers. While additional costs are incurred, the availability of IPv6 on a network does not generate additional revenue. Nor does it provide a competitive advantage as far as we can tell with available data.

The most serious problem is that IPv4 address exhaustion is now the main reason to deploy IPv6, yet dual stacking *does not reduce the demand for IPv4 addresses*. In a pure dual stack configuration, each host and router still needs a globally routed IPv4 address in addition to an IPv6 address. So regardless of whether one deploys IPv6 or not, network growth is still constrained by the availability and cost of IPv4 addresses.[6] This means that NAT must be employed whether one adopts IPv6 or not. Only in the very final stages of the transition, when operators can begin to turn off IPv4 because nearly everyone else is already on IPv6 (which we call the *end game*), is the pressure on IPv4 address space relieved. That end game, however, comes far too late, as it assumes that almost everyone has already made the decision to adopt IPv6 and turning off IPv4 will not lose access to anything important.

Another, related problem is that network operators and content providers who run dual stack make themselves fully compatible with operators, end points and content providers who do not make the transition. While compatibility is of course a good thing, under dual stack the costs of global compatibility are absorbed entirely by those who adopt IPv6, and avoided by those who remain in the older standard. This means that adopters bear an economic penalty in the intermediate term, while organizations that have no incentive to adopt can keep using IPv4 without suffering diminished access or any additional costs. Not until the end game, when major operators can make a credible threat to turn off the older protocol, will laggards suffer any consequences.

1.5 Last mover advantage?

This last bit of analysis has led some observers to describe the IPv4 -IPv6 transition as a “last mover’s advantage.” This view implies that it is rational for network operators to wait until most other networks have converted before assuming the costs of deploying IPv6. This is a very pessimistic view, because if every operator gains an advantage from moving last, and reaps only costs and no benefits from deployment until everyone else has done so, no one will ever move. Waiting for others to go first is most likely to produce an equilibrium frozen on IPv4 forever.

Fortunately for the transition, this view is not quite right. As a factual matter, there are some first movers and numerous early movers. On the other hand, deployment decisions seem to be independent and uncoordinated, and almost randomly distributed across Autonomous System (AS)

[6] There are known methods for economizing on IPv4 addresses, such as sharing via DHCP and more importantly network address translation.

operators (see Section 2). So it does not appear as if IPv6 adoption is a coordination game, in which each networks' deployment decisions create network effects that generate pressure on other networks to follow suit.

What, then, is the incentive to deploy? Our analysis focuses on the issue of network growth. IPv4 address space exhaustion is a constraint on growth. Growth in the IPv4 internet requires purchase of increasingly expensive IPv4 number allocations in the secondary market, and ever-more intensive sharing of globally routed IPv4 numbers through Network Address Translation (see Sections 1.6 and 4.3 below). IPv6 numbers, in contrast, are abundant and inexpensive. IPv6 deployment opens the door to less constrained growth. IPv6 deployment does, however, incur significant initial and ongoing costs caused by the necessity of maintaining IPv4 compatibility. Because of that need for backwards compatibility, IPv6 deployment does *not* immediately eliminate the need for IPv4 addresses, nor does it eliminate the need to NATing (sharing) those addresses. The key drivers of the deployment game, then, is the network operator's subjective assessment of the value of network growth to its business, and the relative cost of a growth plan that involves an IPv6 deployment and one that does not. In this scenario, there is an important externality, which we explore in Section 4.3: as more of a network operators' traffic runs on the IPv6 network, it's NAT costs and its need for additional IPv4 numbers eventually starts to decline.

We formalize the cost structure in the following way:

- Let GV = the value of network growth to an individual AS operator.
- Let GC = the cost of network growth to an individual AS operator.

IPv6 deployment (or non-deployment) decisions are governed at the first order by whether GV > GC, and at the second step by whether the net benefit of growth in IPv6 (GC_6) is less than or more than the net benefit of growth in IPv4 (GC_4).

The value of growth (GV) will vary greatly across AS's. Some networks, for example, are static and have no need to grow. Others are growing rapidly. The cost of growth (GC) will also vary. Variation in GC is based on both the scope and type of the network (e.g., mobile or fixed, legacy or greenfield) and on whether growth is pursued by deploying IPv6 or not.

For a network that does *not* deploy IPv6, GC_4 is a function of

- Cost of acquiring additional IPv4 numbers (v4$)
- Cost of Extending (CE) IPv4, by operating network address translation facilities.

For a network that does decide to deploy IPv6, GC_6 is a function of:

- Initial Costs required to deploy (IC). This includes infrastructure investments, coding, learning and training. Most of these costs are one-off in a given time period, although in larger networks IC can be distributed over different sections of the network in different time periods.
- Cost of Compatibility (CC). These are the costs required to maintain compatibility with the IPv4 Internet. This includes the costs of 6 to 4 NAT translation or tunneling infrastructure, the costs of running a duplicate protocol, and the costs of discovering and fixing incompatibilities caused by the IPv6 deployment.
- Cost of acquiring IPv6 numbers (v6$). This is a negligible cost at the present time.

For IPv6 deployers, GC_6 = IC + CC + v6$. For non-deployers, GC_4 = CE + v4$. CC and CE are comparable cost terms in IPv6 and IPv4, but their relative size will vary across AS's. As v4$ costs rise, GC_4 tends to grow relative to CC. Anything that reduces IC or lowers CC relative to CE will have a major impact on IPv6 deployment incentives. Another key variable is the *traffic ratio*. In the mixed world, traffic combines two protocols, and the mix can be represented as (p)IPv6 + (p-1)IPv4 = 1, where p is a percentage of traffic. For operators who have deployed IPv6, GC_6 declines as p increases. That is, IPv6 deployment becomes more efficient as more of the network's traffic flows over its native IPv6 network. This means that the costs of a network operator are interdependent to some degree with the decisions of other network operators and end users, as others' use of IPv6 reduces a growing network's need for additional IPv4 numbers and reduces their NAT costs.

1.6 Supply and demand for IPv4 number resources

A key factor in the Growth Cost for IPv4 users (GC_4) is the price and availability of IPv4 number resources. Economic incentives affect this in two distinct ways. One is through Network Address Translation (NAT), which subdivides the global IPv4 address space into local private address spaces and interconnects them to the Internet using a smaller number of shared globally routed IPv4 addresses. A network operator can respond to more expensive IPv4 addresses or limited availability by deploying NAT. The other is secondary markets for IPv4 numbers. Number trading creates a market price for number resources and thus incentivizes conservation in use, as well as the sale/transfer of unused or underutilized IPv4 numbers to network operators who need them more.

Ever since the development and legitimization of secondary markets for IPv4 numbers, market incentives have led to the creation of an increasingly efficient process for locating unused or underutilized number blocks and for linking buyers and sellers in ways that move numbers from less valued to higher valued uses. Regional address registry policies began to legitimize secondary markets in 2008, but organized exchanges and brokerages only began to function systematically after 2011. Prior research on these markets (Mueller, Kuerbis & Asghari, 2013) has estimated that nearly 40% of the IPv4 number space was allocated inefficiently (i.e., without clear conservation criteria) to legacy holders. Thus, despite the apparent "exhaustion" of the IPv4 number space from the standpoint of IANA and the regional registries, a significant stock of untapped IPv4 number supply was known to exist. Secondary markets incentivized the discovery and exchange of these resources. By analyzing the markets for IPv4 number resources we can glean important information about both the remaining supply and intensifying demand for IPv4 numbers. This data is highly relevant to the IPv6 transition because, as indicated above, extending IPv4 is a competitive substitute for converting to IPv6. We analyze the IPv4 market in more depth in section 4.2 below.

1.7 Summary: The economic parameters of the transition

The following points summarize our economic analysis framework for the IPv4 - IPv6 transition.

1. No one uses IPv6 only. All public network operators, and nearly all private ones, must offer full compatibility with all other network operators and as many end points and applications as possible. Given that fundamental constraint, there are only three basic choices for network operators:
    a. Remain on IPv4 (do nothing)
    b. Run both IPv4 and IPv6 (implement dual stack and NAT)

c. Run native IPv6 among compatible parts of their own network with some kind of NAT technology at the boundaries to make it compatible with IPv4

2. Amongst these viable alternatives, there is no difference in the network benefits obtained; all three approaches gain access to essentially the same "Internet."
3. The costs of maintaining compatibility between the two standards are borne exclusively by networks that deploy IPv6.
4. Network growth is the critical driver of IPv6. The relative cost of network growth (GC_4 vs. GC_6) is the factor that most affects deployment decisions. This also means that networks that have no need to grow (GV = 0) have no incentive whatsoever to deploy IPv6, and can be expected to lag until the end game.
5. Decisions to deploy or not to deploy IPv6 are made independently (unilaterally) at the AS level, based on each AS's distinctive GV assessment and GC configuration. It is not a coordination game, not until the very end (an end that may never be reached).
6. Networks that deploy IPv6, particularly those who choose option (c) above, realize efficiency benefits when a greater portion of their Internet traffic can be diverted to an IPv6-only network. In this limited respect, the IPv6 deployment decisions of different network operators have some influence on each other. But one network operator's migration to IPv6 creates no variation in demand-side economies of scope, and no discernable difference in the Internet service offered.
7. As the IPv4 Internet continues to grow, and IPv4 brokerages and exchanges make unused or underutilized number blocks available, the slack in the IPv4 number space will be progressively eliminated. The resulting supply constraints on IPv4 numbers will lead to higher prices for IPv4 resources, which can narrow the cost penalty for IPv6 deployment.

The framework provides the basis for interpreting patterns of current IPv6 deployment (Section 2), explaining correlations of IPv6 deployment with macro-social variables (Section 3), and our analysis of the economics of converter technologies and IPv4 numbers (Section 4). Finally, along with the empirical evidence developed in the rest of the report, it forms the basis for our forecast of whether or not we can expect the Internet to converge on one of the competing standards, or remain in a mixed state.

2. Where are we now? The current state of IPv6 adoption

Examining the rates of IPv6 adoption, neither the evangelizers of IPv6 nor its detractors see what they expected to see. The evangelists hoped to see a standard diffusion curve, with slow beginnings steadily gaining momentum, reaching an inflection point, and taking off until 100% diffusion is reached. Critics and cynics of v6 probably hoped they would see it never take root. We are clearly somewhere in between those extremes.

2.1 Measuring IPv6 adoption

Several efforts have been made to track levels of IPv6 "adoption" globally, at the country or network level.[7] Each effort looks at adoption differently using their own methods.[8] For instance, the Asia-Pacific Regional Internet Registry (APNIC) measurement system inserts web objects into random online ads worldwide, which include domain names that trigger different IPv4 and IPv6 behaviors.[9] By observing the requests to their measurement infrastructure, APNIC calculates, for each economy and autonomous system (AS), the percentage of user populations that 1) have capability to use IPv6, and 2) prefer IPv6 connections. By design, the APNIC data collection system gets at the question: "What proportion of the Internet's users are capable of using IPv6 when offered a choice of protocols?" APNIC currently collects around 6-10 million samples per day, with measurements run since 2011 covering more than 200 countries, thereby offering comprehensive, historical, publicly archived data. Using this data, we explored both overall trends and more granular economy level experiences.

2.2 Broad overview of adoption levels worldwide

Following the June 2012 IPv6 World Launch day (an extensive industry and Internet Society-led effort to highlight the need for adopting IPv6), adoption rates overall began to increase, rising from less than 2% globally in 2012 to more than 18% in 2018. But, as noted above, the data do not indicate a technology being ignored, nor do they indicate a technology following a normal diffusion curve.

Beginning in 2016, adoption grew steeply, but leveled off in the second half of 2017. When drilling down to the regional level, we see many cases where more economically developed regions have higher adoption. In 2018, the Americas stood at 28.77%, and Africa at 1.23%. (See Section 3 for a statistical correlation between GDP and IPv6 deployment)). However, there is substantial variance between economically developed regions. E.g., in 2018, Europe was at 15%, almost half North America. Asia stood at 17.03%, Oceania at 14.49%.

---

[7] E.g., Google IPv6 statistics, Akamai IPv6 Adoption Visualization, U.S. NIST, Estimating USG IPv6 & DNSSEC External Service Deployment Status

[8] World IPv6 Launch Day Measurements

[9] APNIC Labs IPv6 Measurements

Figure 1: Global % of IPv6 capable users (Source: APNIC)

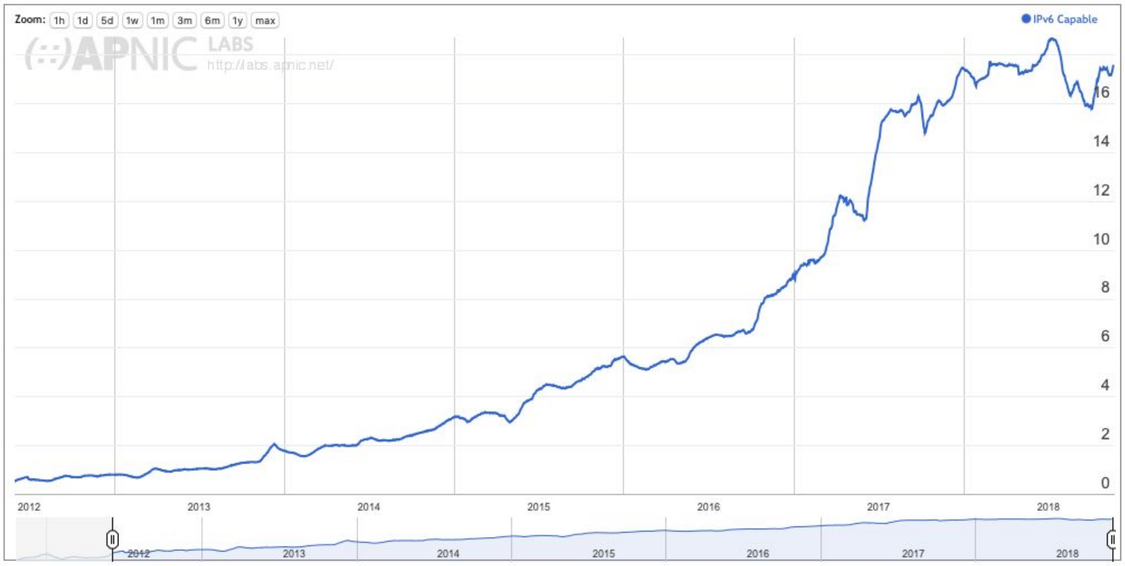

In many respects, this aggregation of all the world's networks into a single trend line is highly misleading. As Figure 2 below shows, there are huge differences in the growth pattern at the national economy level, and as our economic analysis suggests, there will be major differences in individual network operators as well.

2.3 Economy level growth trends

To better understand adoption, we collected average IPv6 capability and preference measurements from APNIC for 215 economies/countries over similar 120-day periods during 2015, 2016 and 2017. Concerned about the APNIC measurement system's reliance on a single advertising platform, we validated capability measurements against data collected by Akamai during the same 2016 period and found it to be strongly correlated, r(210) = .96, p = < .01. We also found APNIC's capability and preference measurements to be strongly correlated, r(428) = .99, p = <.01, allowing us to focus just on capability measurement as an indicator of adoption.

From this data, we identified countries with over 5% measured IPv6 capability in any one of the three years of study.[10] This left a group of only 46 economies, which we examined for year to year growth trends. Figure 2 (below) identifies the number of economies by growth trends observed. The vast majority (169 countries, or 79%) had no appreciable deployment; i.e. they remain at or below 5% during the entire period. Countries in this group were located in all regions, and included both small (e.g., Mozambique) and large (e.g., China) economies. The next largest group (26 countries, or 12%) were economies with *increasing* levels of IPv6 capability. This included economies from all regions with differing growth patterns, including:

- Consistent growth, ranging from 5-10% increase in capability/per year (such as Brazil, Finland, Hungary, Japan, Malaysia, Netherlands, Poland, Romania, United States)

[10] In our judgment, noise levels or measurement errors in the APNIC method are at least 5%, so AS and country-level measurements that are consistently below 5% are counted as non-deployers.

- Strong growth, a doubling or more of capability/per year (Argentina, Guatemala, India, Ireland, New Zealand, Saudi Arabia, Slovenia, Sri Lanka, Thailand, Trinidad and Tobago)
- Rapid growth, a greater than 10% increase in capability/per year but only over one year (Korea, Rep., Macao SAR, Mexico, Puerto Rico, Slovak Republic, Sweden, Uruguay).

Figure 2: Frequency of economy-level IPv6 capability growth trends

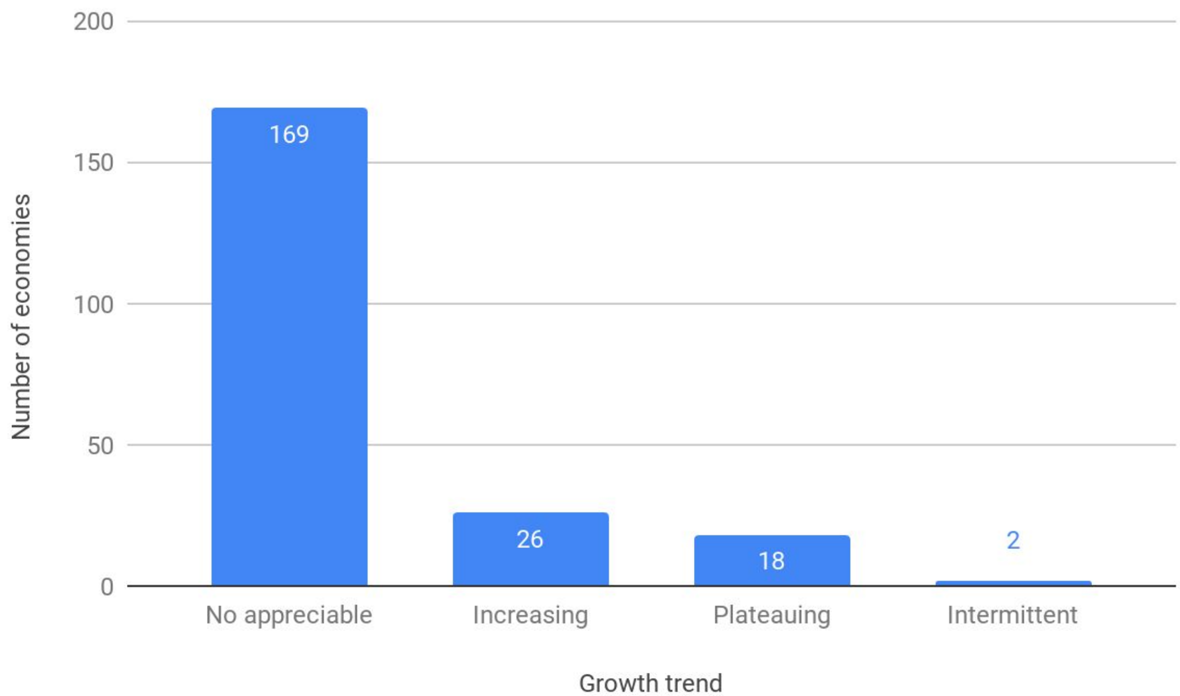

The next largest group (18 countries, or 8%) exhibited plateaus in deployment, with IPv6 capability growth stopping at levels anywhere between 8% (Austria) and 59% (Belgium). This group included numerous mature European economies, like the United Kingdom (27%), France (24%), Luxembourg (34%), Germany (43%), Belgium (59%), Norway (12%), Switzerland (32%), Estonia (17%), Czech Republic (10%), Austria (8%), Portugal (23%), and Greece (38%), as well countries in other regions including Australia (16%), Canada (21%), Vietnam (7%), Zimbabwe (13%), Ecuador (15%), and Peru (19%). The remaining two economies (Singapore, Liberia) exhibited intermittent growth. Capability fluctuated above and below 5% during the period of study.

2.4 Country-level graphs

Figure 3 (below) shows IPv6 capability for 8 countries (aggregating all ASNs operating within the country) since 2013, with data smoothed to 90 day averages. The data illustrate different growth patterns for the period from 2015-2017, from no appreciable (CN, BG), to increasing (US, NZ), plateauing (AU, BE, CZ, EC) and intermittent (LR, SG). Like the global view of IPv6 capability, none exhibit patterns associated with normal diffusion curves. Even with the data smoothed to 90 day averages, many of the growth patterns exhibit small peak(s), presumably as different ASes adopt IPv6.

Figure 3: Selected country-level IPv6 capability growth trends (Source: APNIC)

3a No appreciable movement: China[11] and Bulgaria

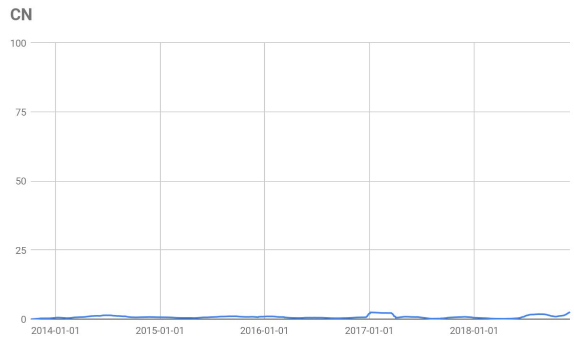

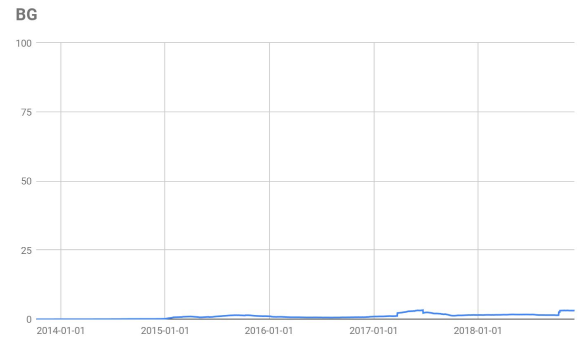

3b Increasing: United States and New Zealand

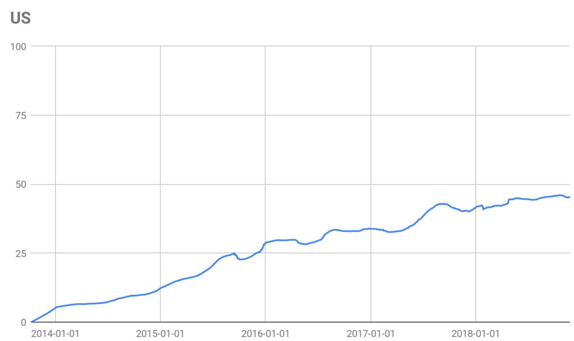

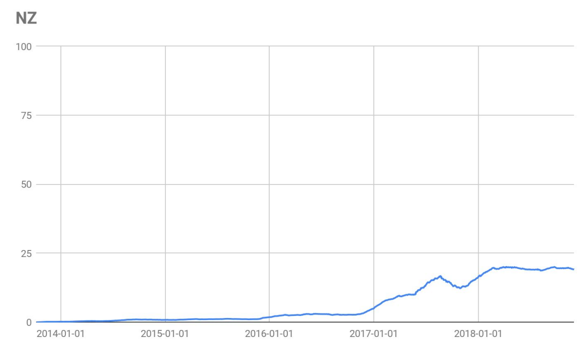

3c Plateauing: Australia and Belgium

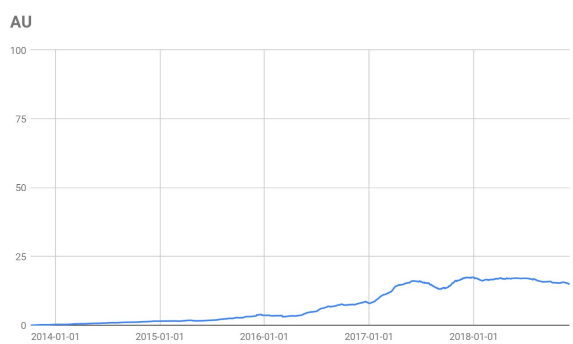

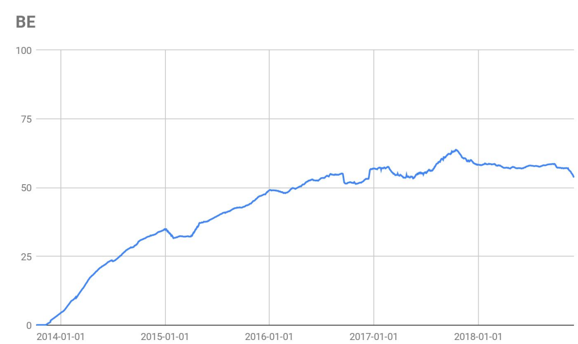

[11] Because China blocks Google the APNIC measurement method does not provide conclusive evidence of its IPv6 deployment levels. However, the existence of very low levels of network operator adoption have been supported in direct discussions with Chinese Internet technical people.

3d Plateauing or declining: Czech Republic, Ecuador

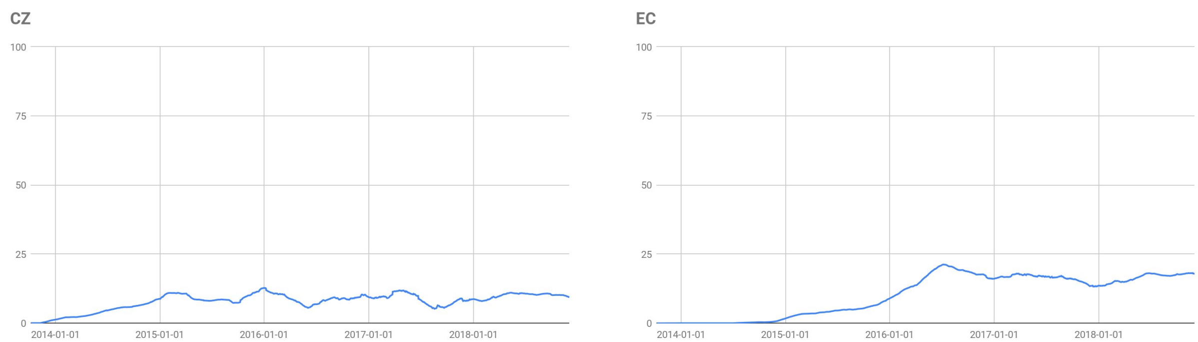

3e Intermittent: Liberia, Singapore

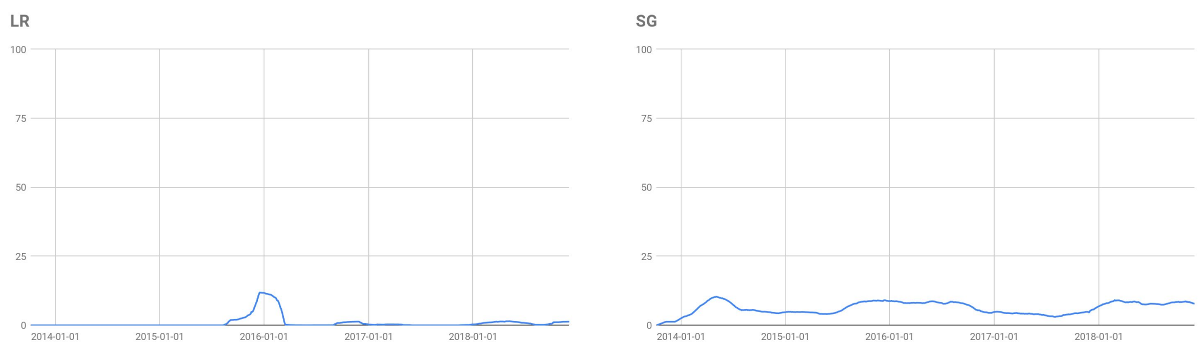

2.5 AS-level views of national markets

Percentage measures of IPv6 capability at the global or country level are just a summation of multiple AS-level deployments. Such measurements would vary depending on the deployment decisions and market share of the different operators. So the most accurate picture of IPv6 deployment comes from looking at the IPv6 adoption levels of individual AS's operating in the same market, on a disaggregated basis. The charts shown below (Figures 4a-4d) show disaggregated IPv6 deployment levels for the top 4 AS's in a national market according to samples collected using the same measurement methodology described in Section 2.4 above. Relative sample size of an AS served as a proxy for market share to determine the top four.

Our analysis of the economic factors driving deployment suggests that there should be wide variation in the deployment levels and choices of different network operators (AS's). One network operator (e.g., a newer, fast-growing wireless network) may benefit from being 75-85% IPv6 deployed while another AS operating in the same market may have no deployment at all, either because it doesn't need much growth or because it sees extending IPv4 as a better option. This prediction is confirmed by the APNIC measurement data. Of the top 4 AS's in each national market, we see wide variation in the timing of deployments from 2014 - 2018, and great variation in the level of IPv6 capability, ranging from 0% to 80% in the same market.

Graphs of IPv6 deployments of an individual AS should look like a step function, as different pieces of its network are converted over a period of one to two years. While it is theoretically possible for a company to decide to convert all of its network in a single time period, larger ISPs are unlikely to get money for a deployment that covers all of its facilities at once. They would deploy on a per market basis, as measures of IPv6 deployment reflect the conversion of different pieces of its network.

Plateaus can be explained based on this logic. Plateaus at percentage *x* at the AS level would be observed when one large AS has made a commitment to deploy IPv6 at percentage *x* of its network facilities, and has either completed that commitment, or suspended further deployments due to cost or compatibility issues. An AS may plateau at a high level (say, 80%) because it has made a financial and operational commitment to IPv6, but the cost of converting the final 20% of the infrastructure might be significantly higher than the initial 80%. Diminishing returns could occur, for example, because the revenue derived from high-density population centers would justify replacing network facilities with the newest routers, whereas the more remote, thinner-traffic segments of the network might be given the older routers which are not IPv6 capable. An AS may plateau at a lower or middle level when management has made a more tentative or experimental commitment to deploy IPv6 and has either completed that commitment, or suspended further deployments due to cost or compatibility issues.

Based on the data we have, it is difficult to attribute a clear cause to *declines* in deployment levels. In the case of T-Mobile (see Section 4.1 below), changes in the mix of end user devices reflected a real, but temporary, decline in capability due to its promotional marketing of an Apple handset that was not IPv6 capable (iPhone). The decline was reversed when Apple made its handsets IPv6 capable. Steep, intermittent increases and decreases in IPv6 capability (such as is seen in Liberia in Fig 4d below) might occur when IPv6 has been turned on and then turned off due to compatibility issues or other implementation problems. But we do not have independent confirmation of this. We are very uncertain about what is the underlying cause of gradual declines of the sort seen by AS 5610 in CZ (Fig, 4c), AS 55430 in Singapore, or AS 1221 in Australia (Fig. 4b). These could be measurement errors, changes in device mix, partial turn-offs, or a combination of all those and other, unknown factors.

Figure 4: Selected AS-level IPv6 capability growth trends (Source: APNIC)

4a - Increasing: United States and New Zealand

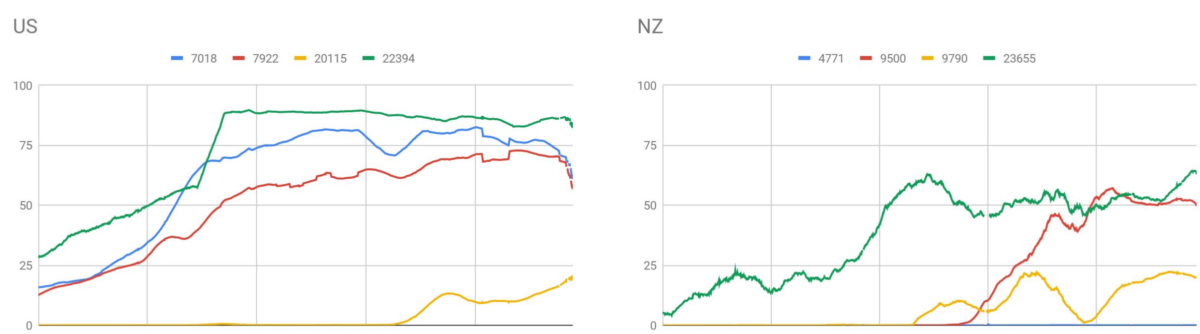

In US and NZ, what appeared to be gradual countrywide increases are revealed to be plateaus for earlier adopters, which are then joined by new deployments by other operators. Note that in New Zealand, one of the top 4 ASNs (4771) has not deployed IPv6 at all, and the earliest and most extensive deployer

(23655) is a mobile operator that has plateaued around 50%. The aggregate growth for the country comes from new deployments from two additional AS's, one of which shows an intermittent pattern.

4b - Plateauing: Australia and Belgium

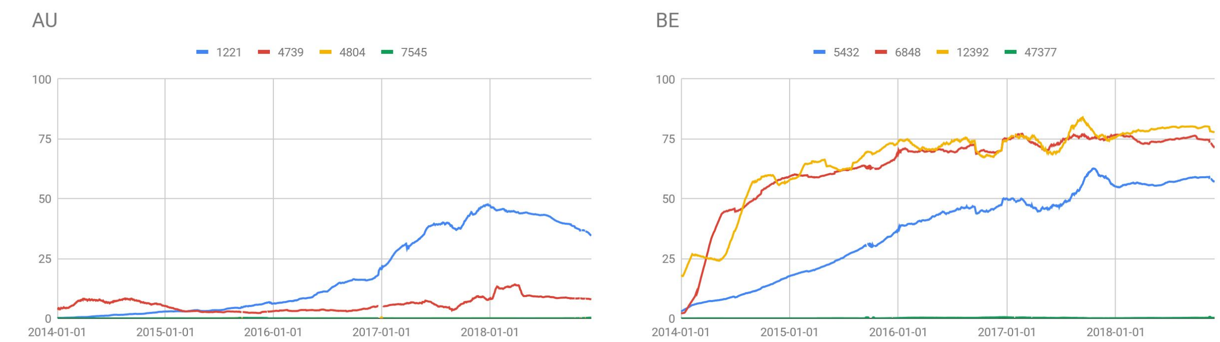

In Australia, only two of the top four AS's are deployers, and an apparent decline in the first AS's deployment level has been covered up in the aggregate statistics by a slight growth in deployment by the second AS. In Belgium, one of the top 4 AS's has no measurable deployment while the two leading AS's seem to be plateauing at 75%. The common timing of IPv6 among the top 3 ISPs in Belgium has been attributed to a coordinated "conversation with the regulator" amongst the leading ISPs; assuming that is true, however, one of the top 4 AS's still has zero deployment.

4c - Plateauing or declining: Czech Republic and Ecuador

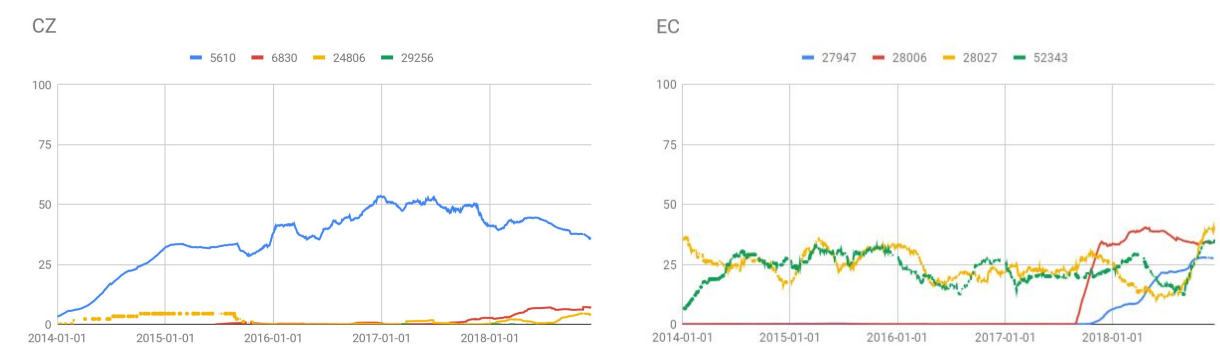

In Czech Republic, only one AS has made a major commitment to deployment and its level reached only 50% and seems to be declining. At the country level, however, the decline is mitigated somewhat by initial low-level deployments by two AS's in the last year. In Ecuador, two universities (AS 52343 and AS 28027) had implemented IPv6 in 25% of their network by 2014, although they have not demonstrated any growth beyond that level, while two carriers (AS28006 and AS27947) deployed IPv6 in the latter part of 2017 and early 2018.

4d - Intermittent: Liberia, Singapore

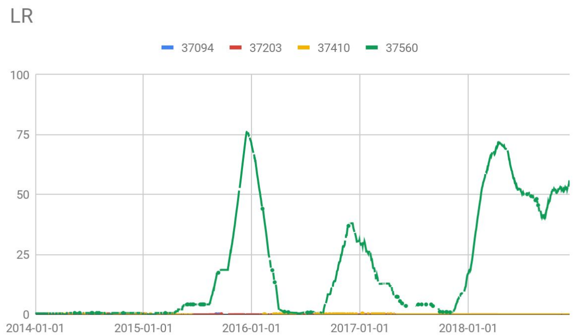

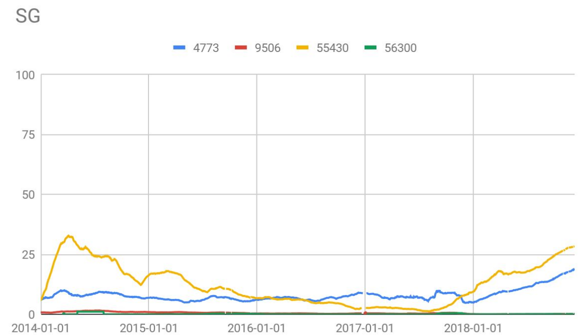

Liberia’s intermittent pattern is clearly explained by the on-off decisions of a single operator. AS 37560 has jumped several times from high levels of measured IPv6 deployment to effectively zero in 6-month periods, before leveling off around 50% in 2018. The three other top AS’s have no measurable deployment. In Singapore, only 2 of the 4 AS’s are deploying IPv6. Together, these two show an unusual U-shaped pattern, with one declining gradually from 2014 to 2017 and then going on an upward swing.

3. IPv6 adoption and macrosocial variables

Many studies have examined the technical aspects of IPv6 adoption. However, there is a dearth of work examining the relationship between IPv6 adoption and macrosocial factors such as economic and market conditions at the country level. This is somewhat surprising given the prevalence of such studies in the traditional telecommunications policy literature. For example, the bidirectional relationship between ICT infrastructure development and country-level economic performance is well studied (Pradhan et al., 2014). Market concentration is a well known factor in fixed and mobile service development and diffusion (Wallsten, 2001; Houpis, et al., 2016). The few studies that do try to look at the relationship between competition and IPv6 adoption are limited. Huston (2018) only considers certain groups of operators adopting IPv6, does not distinguish between different operator types, and does not take into account overall market structures. Given the shortcoming in the literature, we sought to better understand the relationship(s) between the kind of macrosocial variables commonly used in telecommunications policy studies and IPv6 capability rates and growth.

To examine these relationships we developed an economy-level dataset covering 2015-2017 from three main sources, including: the aforementioned APNIC IPv6 capability measurements; World Development Indicators published by the World Bank, which provided population, Gross Domestic Product (GDP) per capita (in US dollars), individuals using the Internet (as a percentage of population), fixed broadband and mobile subscriptions (per 100 people) and compounded annual growth rate; as well as the Inclusive Internet Index developed by *The Economist Intelligence Unit* for Facebook, which provided fixed broadband and wireless market concentration measurements as indicated by the Herfindahl–Hirschman Index, a traditional measure of market concentration that accounts for how operators' market shares are distributed over the number of competitors. Table 1 (below) presents the summary statistics for the dataset.

Table 3.1: Summary statistics for macrosocial variables

	N	Minimum	Maximum	Mean	Std. Deviation
Population	643	11001	1386395000	34591825.80	135054392.60
GDP per capita	581	237.44	168146.02	14092.32	19720.66
Mobile cellular subscriptions (per 100 people)	407	9.80	321.80	106.42	40.51
Mobile cellular subscriptions (per 100 people) (CAGR)	575	-45.55	72.22	2.07	9.73
Fixed broadband subscriptions (per 100 people)	398	0.00	48.38	13.38	13.77
Fixed broadband subscriptions (per 100 people) (CAGR)	561	-77.01	2091.06	21.10	123.01
Wireless market concentration (HHI)	159	1454.94	10000.00	3715.36	1184.61
Broadband market concentration (HHI)	159	1439.28	10000.00	4463.11	2051.21
IPv6 capable (% of users)	430	0.00	0.59	0.04	0.09
IPv6 capable (CAGR)	294	-1.0000	37.4404	0.71	3.21

Our initial analysis found a few significant relationships between macrosocial variables and IPv6 capability measurements when both are measured at the country level. The key findings were these:

- Per capita GDP was positively associated with IPv6 capability, r(381) = .499, p = <.01.
- Higher country-level IPv6 capability rates were correlated with lower levels of concentration in wireless and broadband markets, r(159)= -.267 p = <.01 and r(159)= -.347 p = <.01, respectively

Table 3.2: Summary statistics for macrosocial variables

	1	2	3	4	5	6	7	8	9	10
1. Population										
2. GDP per capita	-0.050									
3. Mobile cellular subscriptions (per 100 people)	-0.052	.400**								
4. Mobile cellular subscriptions (per 100 people) (CAGR)	0.033	-.092*	-0.096							
5. Fixed broadband subscriptions (per 100 people)	-0.029	.745**	.405**	-.133**						
6. Fixed broadband subscriptions (per 100 people) (CAGR)	0.007	-.093*	-.102*	-0.017	-.135**					
7. Wireless market concentration (HHI)	-0.124	-0.130	-0.220	0.145	-.240*	-0.034				
8. Broadband market concentration (HHI)	-.160*	-.371**	-.276*	0.030	-.530**	-0.015	.407**			
9. IPv6 capable (% of users)	.217**	.499**	0.128	-0.043	.448**	-0.058	-.267**	-.347**		
10. IPv6 capable (CAGR)	0.016	-0.004	.192*	-0.031	-0.023	-0.022	0.043	-0.045	.175**	
*. Correlation is significant at the 0.05 level (2-tailed).										
**. Correlation is significant at the 0.01 level (2-tailed).										

We found no relationship between fixed or mobile subscription growth rates and IPv6 capability. However, with regard to growth in IPv6 capability, we found this to be positively associated with the number of mobile cellular subscriptions, r(149) = .192 p = <.01. This is consistent with our view that deployment decisions occur at the individual AS level. The number of fixed broadband subscriptions per 100 population was positively and strongly correlated with IPv6 capability r(195) = .448, p = <.01. As Figure 6 (below) shows, this result was driven almost entirely by Western European countries, the

United States, Canada and Japan, all countries with extensive legacy fixed infrastructure. We found growth in IPv6 capability to be positively related to IPv6 capability, r(294) = .175 p = <.01, although at a low level, which is unsurprising given the step function often observed in AS adoption patterns.

Figure 5: Scatterplot of Fixed broadband subscriptions and IPv6 Capability

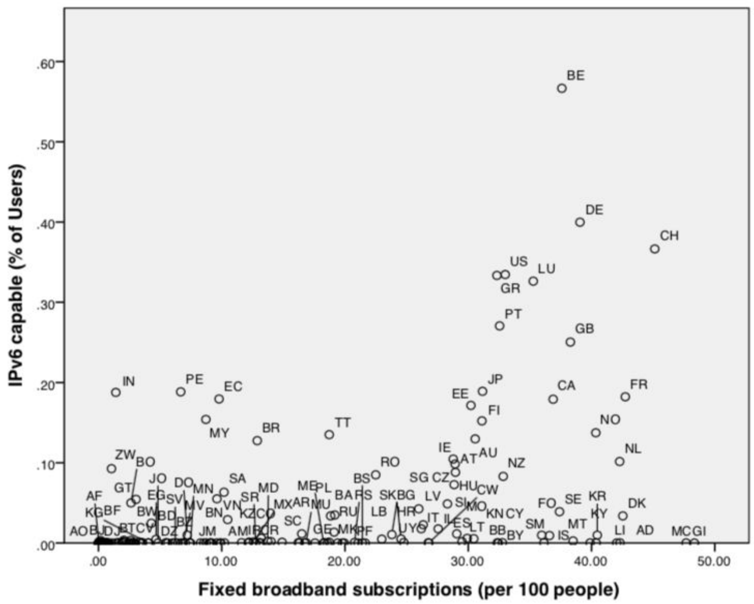

Our findings so far are consistent with our understanding of the economic incentives to adopt IPv6 described in Section 1. Higher levels of IPv6 capability are correlated with greater country-level GDP per capita because IPv6 deployment is costly. Network operators in countries with greater wealth are more likely to risk money on it. Examples would be fixed broadband networks in Western Europe, Japan, United States and Canada mentioned above. The inverse correlation between higher levels of IPv6 capability and market concentration is not as straightforward to interpret, but seems intuitive. A market with more players increases the likelihood that one of the firms will make a random decision to deploy IPv6. A less concentrated market also is more likely to permit the entry of new firms with no legacy infrastructure, for whom the cost of an IPv6 deployment is not much different from the cost of an IPv4 deployment. Higher numbers of mobile subscriptions being positively related to growth in IPv6 capability is consistent with anecdotal evidence observed of IPv6 deployment by mobile network operators.

4. IPv6 adoption and microeconomic variables

This section is based on our development of an operator cost structure model and an analysis of the market for IPv4 addresses. The model allows us to compare the way two IPv6 transition approaches (dual stack and 464XLAT) affect the demand for IPv4 numbers under varying rates of growth and varying assumptions about how much traffic on the Internet is speaking IPv6. The first section (4.1) is a case study of 464XLAT deployment by a specific operator; the second section analyzes the secondary market for IPv4 numbers; the third section isolates a few key variables to develop a model that shows how a network's IPv4 address requirements are affected by some variables.

4.1 A case study of IPv6 deployment for growth

Figure 6 (below) shows IPv6 capability growth for AS 21928 (T-Mobile USA). T-Mobile's substantial IPv6 deployment began nearly three years after IPv6 World Launch day. The handful of major content providers that launched IPv6 in 2012 accounted for 60-70% of T-Mobile's traffic. Having popular content available on IPv6-enabled networks facilitated T-Mobile's IPv6 deployment decision, but operating cost factors related to the growth of its network were more influential. Between 2012 and 2018, T-Mobile added nearly 42 million subscribers,[12] an average year to year growth rate of 13%. (T-Mobile, 2018) For T-Mobile, the business case for IPv6 was grounded partially in the need to reduce reliance on IPv4 resources, but more so in the argument that moving subscribers to IPv6 handsets and

Figure 6: AS 21928 IPv6 capability growth (Source: APNIC)

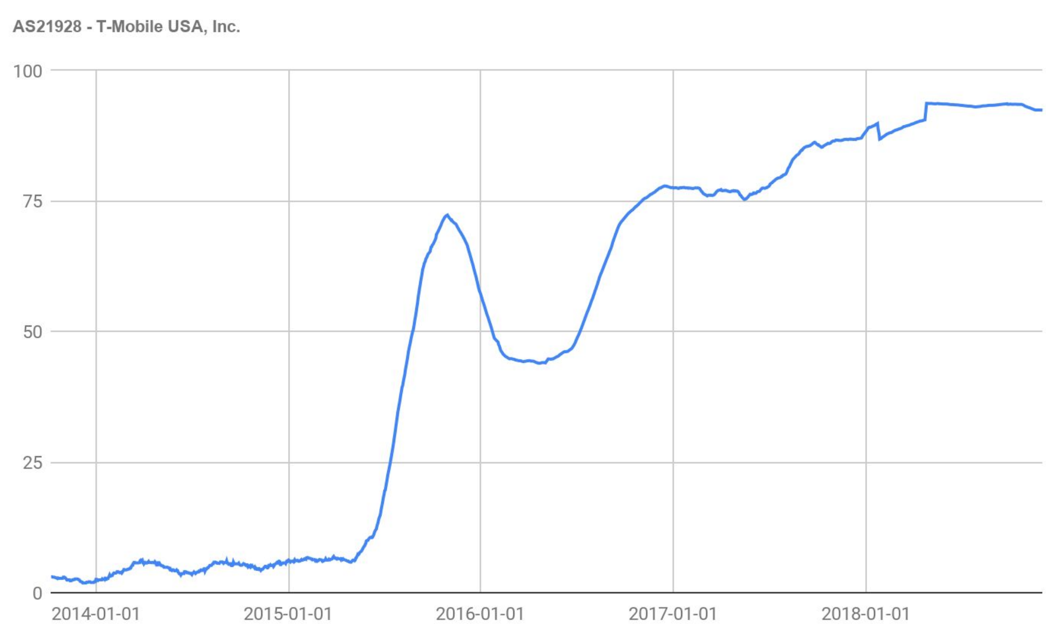

[12] Each additional subscriber is defined as a SIM card with a unique T-Mobile USA mobile identity number that generates revenue.

native IPv6 communication would decrease IPv4 Network Address Translation (NAT) costs. NAT is prevalent in all commercial mobile networks. According to our interview subject, NAT can be very challenging for operators as it has scaling, maintenance of state, and other issues that introduce additional costs. T-Mobile's modeling estimated that IPv4 NAT costs could be cut in half by switching subscribers to IPv6.

To provide compatibility between the IPv4 Internet and the IPv6 Internet, T-Mobile implemented 464XLAT, which is defined in RFC 6877 (2013). Several other mobile operators have implemented the 464XLAT approach, including BT, SK Telecom, Telus, Rogers, Orange, and Bouygues Telecom. Figure 7 (below) outlines the topography of the T-mobile 464XLAT approach.

Figure 7: 464XLAT Transition Approach

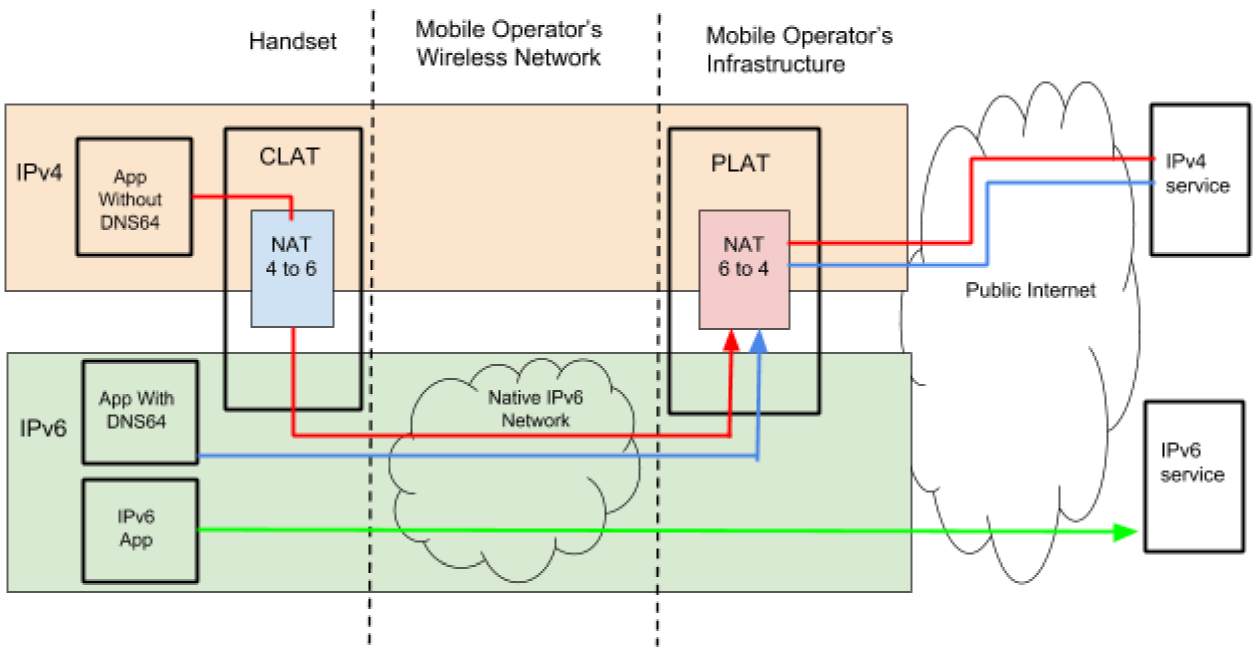

464XLAT divides the communication path into four segments: the handset, the mobile operator's wireless network, the mobile operator's internal infrastructure, and the Internet. At the core of the communication path is an IPv6-only access network (shown here as a green cloud), as opposed to two distinct networks (i.e., dual stack) or an IPv4 or IPv6 network transiting encapsulated packets (i.e., tunneling). Traffic can be categorized in three ways:

- Mobile devices and applications at the edge that support IPv6 can communicate directly with IPv6 services on the Internet without interfacing with a converter technology.
- IPv6 devices and applications communicating with IPv4 services interface with the Provider side Translator (PLAT),[13] based on the NAT64 protocol (RFC 6146), that converts (N:1) global IPv6 addresses to global IPv4 addresses, and vice versa.

[13] PLAT is supported by numerous equipment vendors and software implementations, including A10, Cisco, F5, Juniper, NEC, Huawei, Jool, Tayga, Ecdsys, Linux, OpenBSD, etc.

- IPv4 devices and applications communicating with IPv4 services interface first with the Customer side translator (CLAT),[14] based on the SIIT protocol (RFC 6145), which converts (1:1) dynamically-assigned, private IPv4 addresses to global IPv6 addresses, and vice versa. The IPv6 addressed traffic then interfaces with the aforementioned PLAT to convert it to a globally unique IPv4 address.

There are several cost advantages to the 464XLAT approach. By repurposing IPv4 NAT equipment to NAT64, T-Mobile needed no new capital expenditures to implement this approach. Operating a native IPv6 core network avoided the duplication and/or the additional complexity of dual stack or encapsulation. At the network level, an IPv6-only network is relatively cheaper to scale, thus reducing a network's capital and operating costs as it grows. While the approach continues to use converter technologies, it is no longer solely IPv4 NAT, which would perpetuate a network's demand for globally unique IPv4 addresses, albeit at some reduced level. Stateful PLAT minimizes a network's demand for globally unique IPv4 addresses (Byrne, 2014), and stateless CLAT eliminates it. Another important factor is the location of converter technology in the network's topography. The CLAT resides in the Customer Premise Equipment (in this case, the mobile handset). Because private address space is used there, the growth of subscribers is no longer coupled to the availability of globally unique IPv4 addresses.

Figure 6 shows how IPv6 capability grew in T-Mobile from 2014 to 2018. T-Mobile initially launched its IPv6 plan into production with IPv6-enabled Android devices in 2013 and 2014. T-Mobile engineers anticipated compatibility failures that their customers might experience and proactively addressed them.[15] By mid-2014, T-mobile claimed that among devices with IPv6 capabilities, over 50% of their communications were end-to-end IPv6 with no translation needed. (Bryne, 2014) Our interview subject speculated that the decline in capability observed in late 2015 was attributable to the IPv6-enabled Android device users migrating to an iPhone model promoted by T-Mobile that was IPv4-only. This trend was reversed when the next iteration of iPhones, which were IPv6-enabled, came out. Around the same time Apple started to restrict its App store to IPv6-compatible applications, leading to a noticeable peak in IPv6 capability in late 2016.[16] By 2017, T-Mobile had reached nearly 75% IPv6 capability according to APNIC and other adoption measurements.[17] With significant implementation experience and its IPv6 capability closing in on 90%, T-Mobile attempted to go "IPv6-only" for its iPhone subscribers in early 2017. After two attempts and subsequent rollbacks (presumably visible in two small dips in Figure 6) due to complaints from large customers unable to access websites, these problems were resolved. By Spring 2017 T-Mobile was IPv6 only for iPhone subscribers. (Lagerholm, 2017) Subscribers using tethering, roaming, or with IPv4 devices/applications/networks continue to use IPv6 + 464XLAT. T-Mobile has no new IPv6 implementation initiatives.

---

[14] CLAT is supported by numerous equipment vendors and software implementations in Android, Nokia, Windows, NEC, Linux, Jool, OpenWRT, Apple, etc.

[15] These included both DNS query/response and network routing errors. See Lagerholm (2017) for various explanations of how content sites T-Mobile subscribers attempted to access would fail and the solutions T-Mobile would implement.

[16] At its Worldwide Developers Conference (WWDC) 2015 Apple announced the transition to IPv6-only network services in iOS 9. Starting June 1, 2016 all apps submitted to the App Store were required to support IPv6-only networking.

[17] E.g., World IPv6 Launch, T-Mobile USA IPv6 Deployment

While different types of network operators (e.g., fixed, mobile, enterprise, content, cloud) may implement a variety of transition approaches, from this case study we can extract at least five generalizable factors that can be examined further and modeled:

- Subscriber growth
- Subscriber device IPv4/IPv6 ratio
- Converter technology
- IPv4 address price (relative to IPv6, and available IPv4 assets)
- The ratio of IPv4 and IPv6 traffic that passes through a network's gateway

The impact of these variables will be explored in part 4.3 below, after we examine the market for IPv4 addresses.

4.2 IPv4 prices and resource transfers

The supply of IPv4 numbers plays an important role in the IPv4 - IPv6 competition. The prospect of what some engineers have called "IPv4 runout" was the main reason for developing IPv6 in the first place. From an economic point of view, however, resources never just "run out;" instead, as their supply diminishes they become increasingly expensive, and consumption patterns adapt to scarcity with greater conservation and new forms of substitution.

Network operators have adapted to the tighter supply of IPv4 addresses in two ways. One is by using NAT, a conservation technique that uses a private (non-globally routable) IPv4 address space to connect local hosts, and passes traffic to the Internet by translating the many private addresses into a smaller number of globally routable IPv4 addresses. NATs are the reason why 20 billion connected devices on the modern Internet can be served by about 2 billion active IPv4 addresses. The other adaptation is a secondary market for IPv4 number blocks, which allows networks that need more IPv4 numbers to buy them from networks with an excess supply. The incentives provided by the secondary market have led to the identification of millions of unused or underutilized IPv4 numbers by brokers such as IPv4 Market Group and exchanges such as Addrex and Hilco Streambank.

Data presented below details trends in IPv4 address prices and the number of transfers in the secondary market. What little public price data exists was published by Hilco Streambank.[18] Hilco provides an online auction platform for the sale of IPv4 address blocks, including blocks registered in ARIN, RIPE, and APNIC, and ranging in size from /24 to /17.[19] This data, combined with data on the number of transfers published by the RIRs, allows us to make some important observations.

Figure 8 (below) charts the price data published by Hilco. Reviewing the 1,161 blocks transferred within that limited market, the median price per address has doubled in four years, from around $8.00 in 2014 to $17.00 in 2018. As one would expect, the range of market prices narrowed as the market matured and initial uncertainties about the value of IPv4 addresses were resolved. These price trends were corroborated in interviews with other address brokers.

Another important indicator is the number of transactions in which IPv4 blocks were transferred (Figure 9) and the total number of IPv4 addresses transferred (Figure 10).[20] Data from the Regional

[18] The relative lack of price transparency in the IPv4 address market (compared to other resource markets, e.g., spectrum) is an area that could be improved.

[19] IPV4 Auctions.com, Recently Closed Auctions

[20] Because we are focused on markets for IP addresses, we do not include IPv4 address blocks or addresses acquired through merger and acquisition. Also, Figure 9 shows the number of address blocks (or "transfer sets" as

Figure 8: Hilco Streambank IPv4 address block transfers

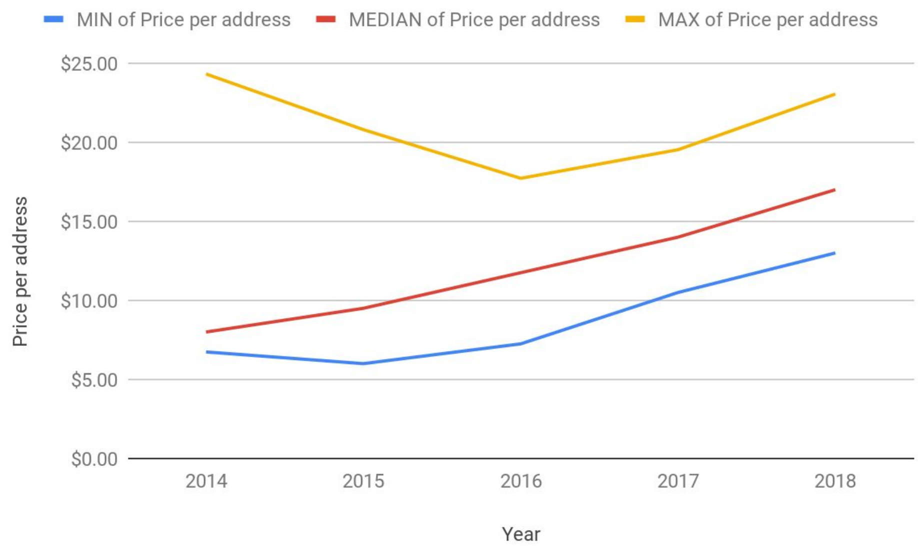

Internet Registries on the number of IPv4 address transactions shows a major increase in 2014 to 2015 as the markets developed, with smaller growth increments in 2016, 2017, and 2018. The number of transactions is still growing, however. The leveling off of growth in the amount of numbers transferred

Figure 9: Number of IPv4 address block transfers, by recipient RIR

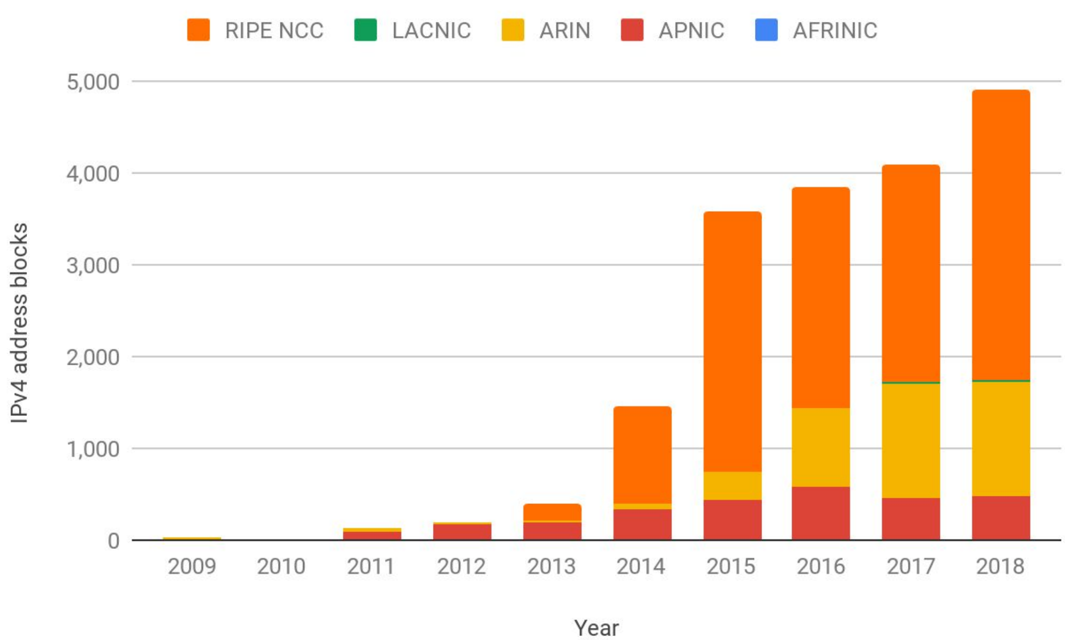

---

they are referred to in data). RIPE and APNIC record 1 or more transfer sets per transfer record (i.e., a transaction between two organizations recorded by the RIR). ARIN records one transfer set per transfer record. E.g., some RIPE transfers include 19 transfer sets, one APNIC transfer included 78 individual transfer sets, although the vast majority in both regions only record one transfer set.

Figure 10: Total IPv4 address numbers transferred, by recipient RIR

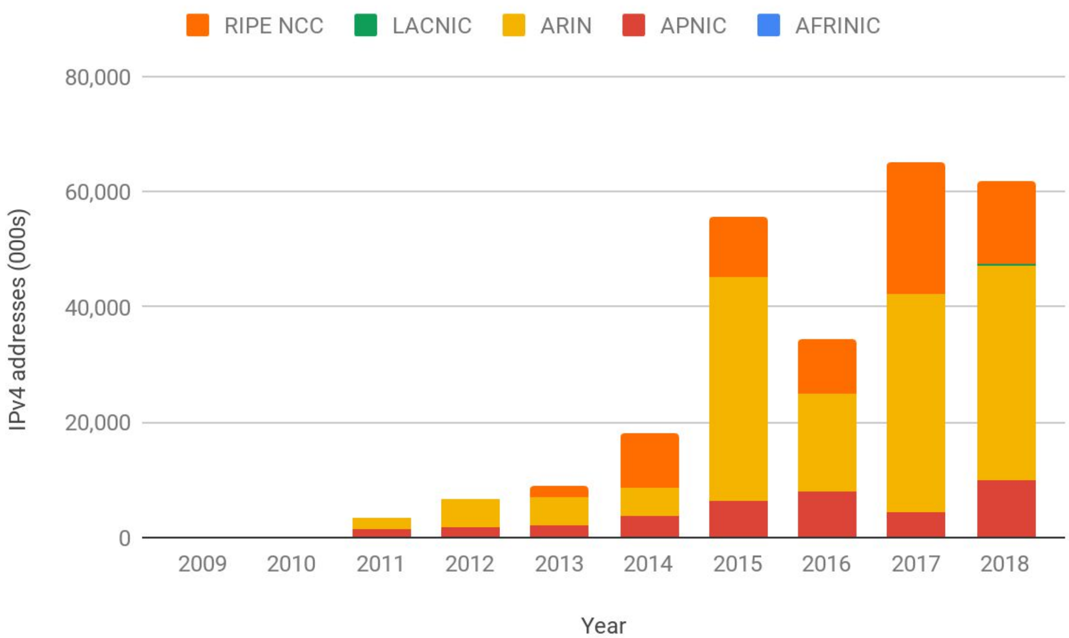

per year corresponds to steady rises in price, possibly indicating that supplies of untapped IPv4 number blocks are no longer expanding. From 2015 to 2018, an average of about 54 million numbers were transferred each year.

Who is buying IPv4 address resources? The answer, evident from Figure 11, is cloud service providers (CSPs): Amazon AWS, Microsoft Azure, Google, Alibaba, Oracle Cloud, and others. Cloud services are a rapidly growing market, currently estimated at 17% annually and worth $200 billion. (Gartner, 2018) In terms of number of customers and interconnected hosts, these cloud networks represent the largest networks ever connected to the Internet. While this market initially served only individual developers, enterprise use of them has grown tremendously over the last 3-4 years. CSPs are buying IPv4 numbers because enterprise networks are lagging in IPv6 deployment relative to public provider networks. One interviewee said only about 5-7% of enterprise network traffic is IPv6 enabled, an assertion that is corroborated by some estimates of industry IPv6 deployment.[21] Enterprises are also slow to upgrade their applications, particularly if they have current revenue from them. This drives demand for IPv4 addresses among cloud providers, as each instance of service may require a globally routed IPv4 address.

Cloud providers see clearly that there are two Internets, and they are aggressively positioning themselves to serve both of them. Many cloud providers built IPv6-enabled services to become qualified for U.S. government contracts, but the majority of their revenue is derived from services delivered over IPv4 connections. Amazon and Microsoft have tens of millions of customers, each supporting thousands

[21] E.g., as recently as December 2018, the National Institute for Standards and Technology's Advanced Network Technologies Division estimated a little more than 10% of industry unique operational web interfaces surveyed supported IPv6. See https://fedv6-deployment.antd.nist.gov/govmon.html

or millions of devices, each of which may/may not require a public facing address. Cloud service providers are positioned to feel the pain of IPv4 scarcity most acutely. So far, they have not passed on the cost to enterprise customers, but as available IPv4 resources diminish, it is reasonable to expect CSPs to pass on these costs to customers, or provide them with an incentive to deploy IPv6. However, according to one expert interviewed, some CSPs are already incentivizing customers to use IPv6 services as IPv4 addresses become more scarce and expensive.

Table 4.1: Top 10 Recipient Organizations in ARIN region of Transferred Addresses

Recipient Org	Number of Transfers	Number of Addresses	% of Total Addresses Transferred	Operator Type
Amazon Technologies Inc.	30	61,275,974	35.91%	Cloud
Microsoft Corporation	6	30,998,482	18.16%	Cloud
Charter Communications	33	8,386,444	4.91%	ISP
Amazon.com, Inc.	21	6,753,160	3.96%	eCommerce
Google LLC	3	5,243,389	3.07%	Cloud
Alibaba.com Singapore E-Commerce Private Limited	1	5,242,878	3.07%	Cloud
Frontier Communications Corporation	1	4,718,581	2.77%	ISP
Google Inc.	3	4,194,299	2.46%	Cloud
Alibaba.com LLC	7	3,014,634	1.77%	Cloud
Reliance Jio Infocomm Pte Ltd	5	2,162,672	1.27%	ISP
Google Fiber Inc.	1	2,097,151	1.23%	ISP
Oracle Public Cloud	5	1,441,771	0.84%	Cloud
VODAFONE AMERICAS INC.	2	1,118,202	0.66%	ISP
Windstream Communications LLC	1	1,048,575	0.61%	ISP

Other networks acquiring IPv4 include growing ISPs such as Charter, and Indian mobile network Reliance Jio which, while using IPv6, still have great demand for IPv4 given the scale of their networks and need to connect to IPv4 hosts. This corroborates our modeling exercise in Section 4.3 below, which shows that growing networks that deploy IPv6 may still need to increase the size of their IPv4 number holdings. While various NAT strategies and a shift in the traffic matrix to more IPv6 endpoints reduces the amount of IPv4 numbers needed, we can still expect to see intensifying use of IPv4 numbers and higher prices as they are reallocated to the most highly valued uses.

4.3 Modeling IPv4 address requirements under dual-stack and 464XLAT

In this section, we model how an operator's technical requirements for globally routable IPv4 addresses over a 15 year period are affected under two transition cases, dual-stack and 464XLAT. The model allows us to explore the way the operator's need for IPv4 addresses is affected by various subscriber growth rates and the growth of IPv6 capability in the operator's traffic matrix.[22] This provides insight into how IPv4 address scarcity might create incentives to deploy IPv6.

Under dual-stack, the operator implements a separate IPv6 network, and an IPv4 network using stateful NAT44 at the edge router, and carrier grade NAT (CGN) for IPv4 in their gateway router. Under 464XLAT, the operator implements the approach discussed in section 4.1 above, with 90% of its devices assumed to be IPv6 capable (reflecting the current state of mobile devices). Stateless NAT64 is put at the edge router converting any subscriber IPv4 sessions to IPv6, a core IPv6 network, and stateful carrier grade NAT64 at the gateway converting IPv6 sessions to public IPv4. In both cases, we assume the operator uses typical active subscriber, reserved port, and compression ratio values in configuring the CGN router.[23] The need to maintain compatibility between IPv4 and IPv6 means that networks that deploy IPv6 must continue to use IPv4 numbers as addresses. The model is intended to show how the requirements for IPv4 addresses change under various scenarios.

One simple finding is that a network operator's IPv4 requirements are not affected by which of these two IPv6 transition models, dual stack or 464XLAT, it chooses. While the IPv6 number requirements differ, the IPv4 requirements are the same in both cases. A network operator's technical requirement for globally routed IPv4 addresses is affected most strongly by the following variables:

- The baseline size of the network in terms of its subscribers/users (i.e., does it start with 1,000 users, 100,000 users, 10 million users?)
- The growth rate of subscribers/users and the specific trajectory growth takes (i.e., flat, a simple linear annual percentage growth, plateauing growth, accelerating growth or a standard s-curve growth pattern)
- The baseline IPv6 traffic ratio at the network's gateway (how much traffic is IPv6-enabled to begin with)
- The rate and pattern at which the IPv6 traffic ratio shifts (i.e., flat, a simple linear annual percentage rate, plateauing growth, accelerating growth or a s-curve growth pattern)

As we discuss the model, it is important to keep in mind that the reasonableness of assumptions regarding growth in network size and the shift in the IPv6 ratio will vary across each of the 60,000+ network operators. Some networks are large, fast-growing and embedded in environments where a large portion of the traffic is IPv6 capable. Others are small, not growing, and exchange most of their traffic with environments where IPv6 capability is low. And of course there are many points in between these extremes. Operator diversity means that any number of outcomes are possible. All decisions regarding IPv4 address acquisition and IPv6 deployment will take place at the individual operator (AS) level and will reflect this diversity in conditions and incentives. The best we can do is show various scenarios and discuss their implications for the transition.

---

[22] Our modeling only examines the number of public IP addresses required to operate the network and the scaling properties of NAT. Other factors, not considered here, will influence how an operator deploys NAT. E.g., CPU, memory, or power consumption, traffic volume (throughput) constraints, and performance factors like reliability.

[23] Specifically, we assume 80% active subscribers during peak traffic, and implementation of deterministic port sharing, with 1024 reserved ports used by popular protocols, and a compression ratio of 8 (CableLabs, 2012). It is important to note that although there are inherent tradeoffs and performance limits, NAT compression ratio can be increased which effectively reduces the number of required public IP addresses.

In general, the model indicates that the *subscriber growth rate* and the *rate at which the traffic ratio shifts towards IPv6* are the most critical factors affecting IPv4 number requirements. Rapid growth pushes IPv4 requirements upwards, shifts in the traffic ratio toward IPv6 push it downwards.

Figure 11 shows a set of assumptions involving a large mobile network with fast subscriber growth (15% per year) coupled with a slower (5%) annual shift in the traffic ratio toward IPv6. The levels of IPv6 traffic are assumed to start around 44%.[24] These assumptions produce an inflection point, after about 10 years, where the network's technical requirement for IPv4 addresses starts to decline precipitously. This happens when the traffic ratio is near 74% IPv6 and 26% IPv4. If the subscriber growth rate is lower, then the inflection point occurs earlier (shifts to the left on the X axis), at a lower IPv6/IPv4 traffic ratio. But if the user growth rate is equal to the rate at which the traffic ratio shifts towards IPv6, the network's need for IPv4 addresses starts to decline immediately.

Figure 11: Scenario 1 (mobile ISP, rapid growth)

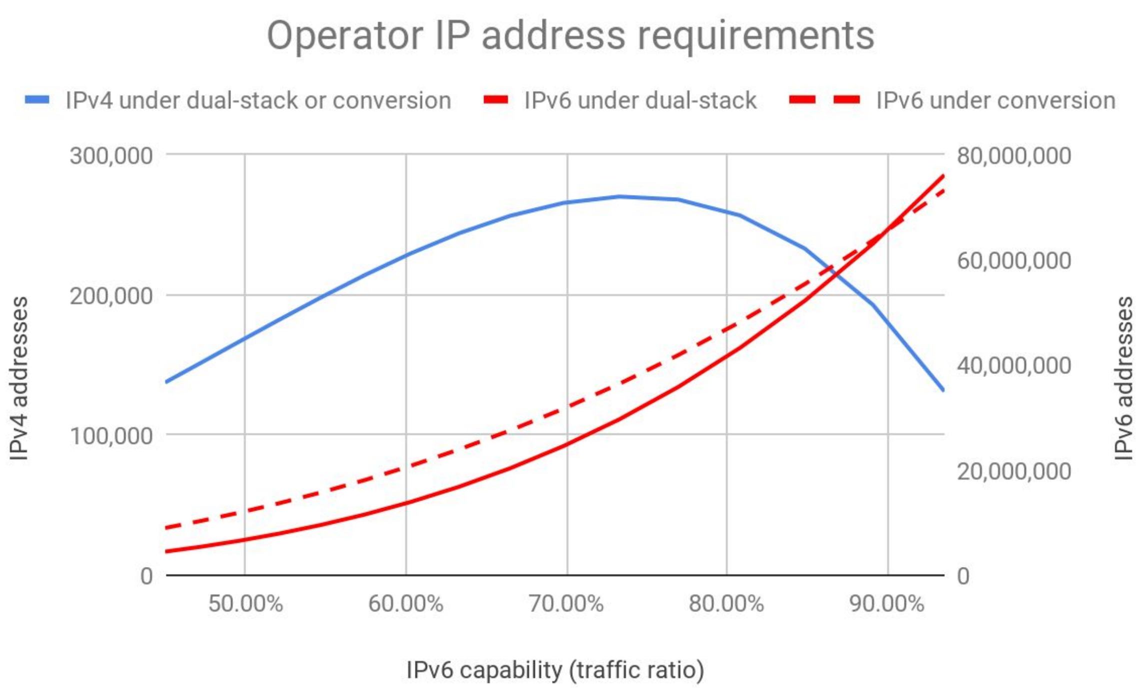

The model is also slightly sensitive to the specific pattern of growth. Growth or change in either critical variable could be linear, accelerating, plateauing or follow an s-curve pattern. For example, if the subscriber growth rate slows down and starts to plateau after 7 or 8 years and the IPv6 traffic ratio follows a more rapid S-curve shift, the operator's need for IPv4 numbers starts to decline after only 4 years and approaches 0 after 15 years. (Figure 12)

[24] These IPv6 capability assumptions and growth rates are based on mobile operators in the United States (although traffic ratios are operator specific).

Figure 12: Scenario 2 (mobile ISP, different growth pattern)

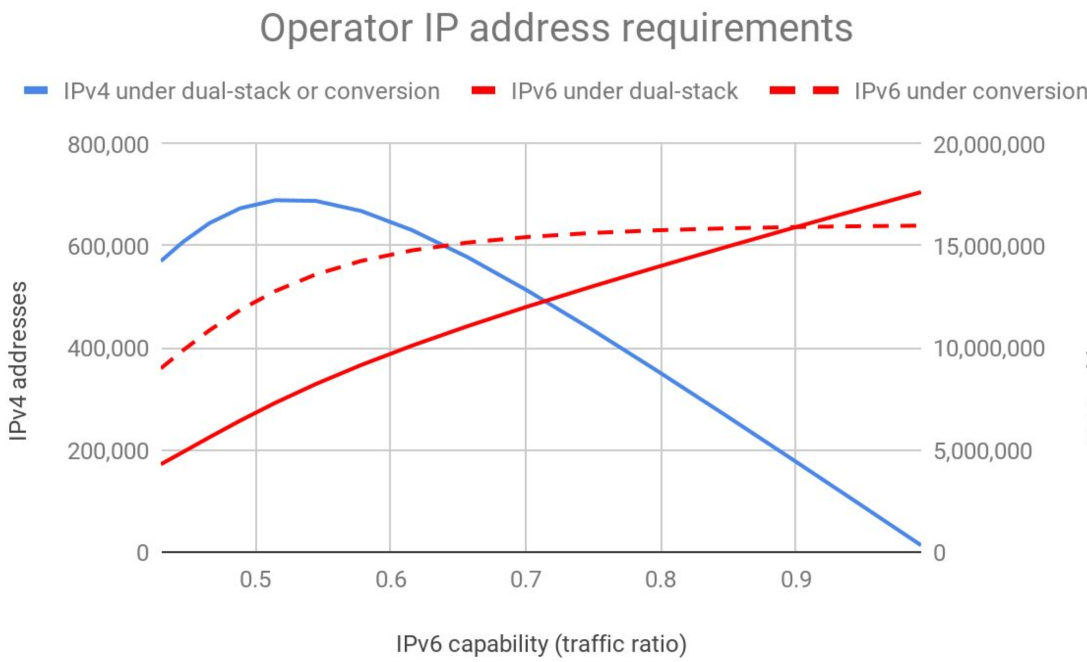

A different set of assumptions for a different type of network, however, tells a very different story. Figure 13 shows a small enterprise network that has only 10,000 users and grows only 2% per year.

Figure 13: Scenario 3 (small low growth enterprise network)

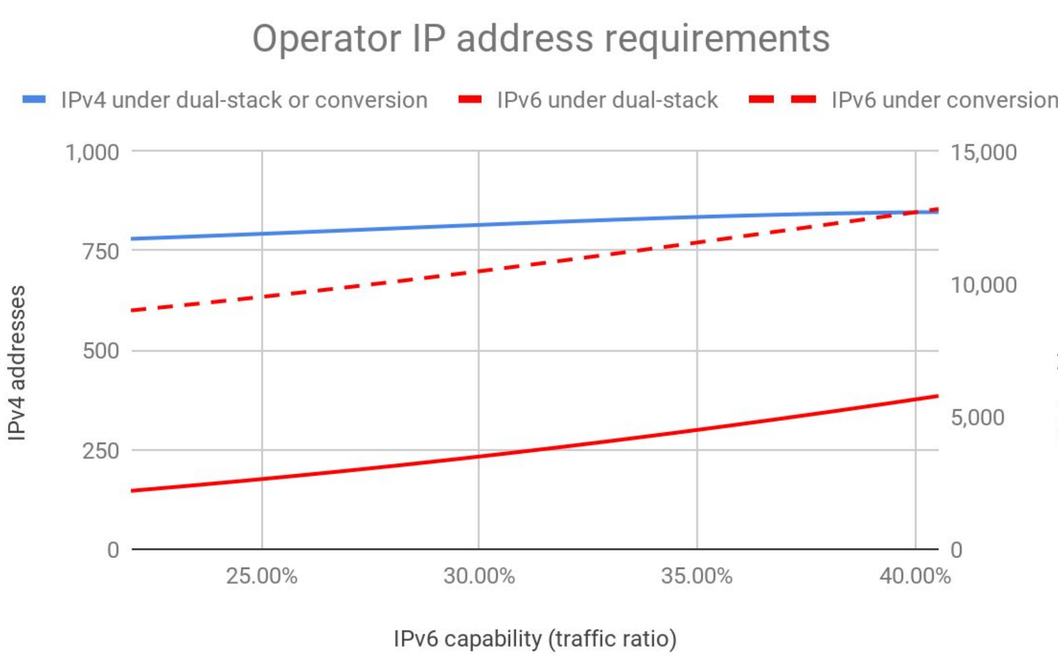

The traffic ratio shifts towards IPv6 linearly at the rate of 4% a year, but since the enterprise network is located in an older legacy software and device environment, the starting point for the IPv6 traffic ratio is set at 22% rather than 45%. In this case, implementation of IPv6 doesn't do anything for the operator. Its IPv4 number requirements increase modestly whether or not it implements IPv6, so the additional expenses for deploying it would not reduce its demand for IPv4 addresses by a worthwhile amount. While its IPv4 number requirements would also increase without IPv6, the small scale of the network would mean that the additional numbers needed (1,000 - 1500) would not be that expensive or hard to come by. At worse, the NAT compression ratios could be tweaked to stretch their address assignments further. In general, for smaller scale networks the gains from IPv6 deployment do not justify the expenses. For such networks it makes sense to continue to extend IPv4. Not until the market value of the IPv4 addresses held by the network exceeds the enterprise's costs of transitioning to IPv6 (e.g., upgrading devices/applications and running dual-stack or conversion) would there be an economic incentive to change. However, not all enterprises own their IPv4 addresses. In those cases, the pressure to shift may come from an upstream service provider.

In theory, after reaching the inflection point where their requirement for IPv4 addresses starts to decline, a network operator would be able to sell surplus IPv4 numbers into the secondary market. This release of number resources would make it easier for other operators to expand the size of their IPv4 networks. Because not all networks would reach this inflection point at the same time, it is hard to predict what kind of impact the new supply of IPv4 numbers would have on the transition. The new supply could be taken up by other expanding dual stack IPv4/IPv6 networks that had not reached the

Figure 14: Scenario 4 (cloud service provider)

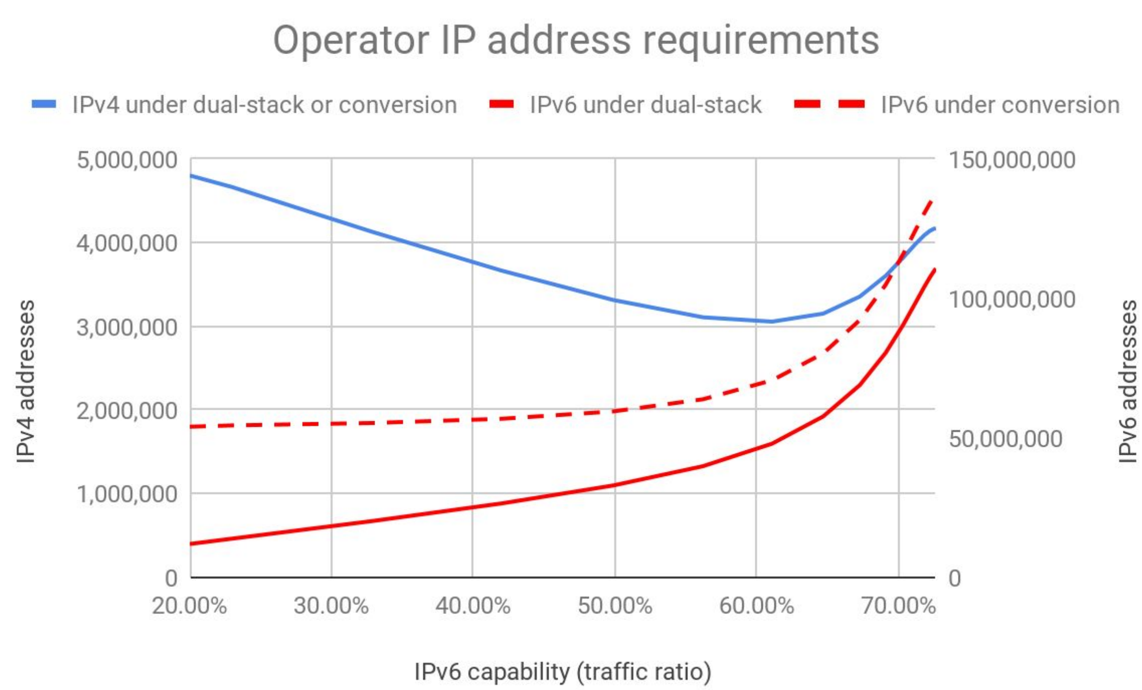

inflection point. Or it could be taken up by IPv4-only networks seeking to extend their life. If a number of large dual stack or 464XLAT networks started releasing IPv4 numbers at the same time, it could lead to major price declines in IPv4 numbers, perhaps making it easier for smaller legacy networks to continue relying on IPv4.

But it is also possible to posit scenarios where large providers might want to hang on to their IPv4 numbers. Figure 14 is based on assumptions that are intended to represent a major cloud service provider (CSP). It assumes 60 million subscribers to start, with an s-curve growth trajectory in their numbers over 15 years. The IPv6 traffic ratio shift is assumed to grow rapidly at first but then plateau around 70-73%. We have also assumed a lower starting point for the IPv6 capability traffic ratio (20%), because CSPs interact with enterprise networks likely to have lower IPv6 deployment rates. The model shows the IPv4 requirements for a CSP declining steadily for 5 years while the traffic ratio shift rate exceeds the subscriber growth rate; as the shift rate levels off, however, the requirement for IPv4 numbers starts to increase again as subscriber growth continues. Thus some operators might, despite experiencing increases in IPv6 traffic and declines in their immediate IPv4 requirements, continue to hold IPv4 resources in the face of uncertainty about future IPv6 traffic ratios and IPv4 prices. While traffic ratios are operator specific, such a scenario could occur given the numerous economies that are exhibiting plateauing growth currently.

4.4 Discussion

Instead of a smooth, linear transition from IPv4 to IPv6, we see a complicated mixed world emerging. IPv6 deployers can escape the growth constraints of IPv4 by setting up separate IPv6 networks and using dual stack/NAT or conversion approaches such as 464XLAT to remain compatible with the IPv4 Internet. But this kind of strategy only makes sense for larger-scale networks characterized by significant growth. Even for them, the escape from the need for IPv4 numbers is only partial, and full realization of its benefits may take years, because they must wait for the proportion of external traffic to shift to higher levels of IPv6. To summarize key findings of the model:

- The demand for public IPv4 addresses is essentially the same under different IPv6 transition approaches.
- It is the interaction between an operator's subscriber size/growth rate and its gateway IPv4/IPv6 traffic matrix (i.e., how many end hosts are speaking IPv6 vs. how many are speaking IPv4) that determines the number of public IPv4 addresses required.[25]
- The more quickly an operator's traffic ratio shifts towards IPv6, the more quickly they can reduce demand for public IPv4 addresses.[26]
- Ceteris paribus, the higher the growth rate in subscribers or users, the longer it will take for a network's IPv4 address requirements to begin to decline. If they are not growing rapidly, they do not have to wait as long - but the lower the growth rate, the lower the incentive to invest heavily in IPv6.

Our study did not formally model the way rising prices and diminishing supply of IPv4 address resources enter into operator's decisions. But some useful observations are possible. First, as operator's

[25] Operators' gateway IPv4/IPv6 traffic ratios vary widely depending on the mix of end hosts with which their subscribers are communicating.

[26] This explains substantial efforts (like World IPv6 Launch Day) to get the largest, most popular content providers such as Google and Facebook to enable IPv6 capability. It also suggests where future efforts might be targeted, e.g., identifying major second-tier content providers, aggregated groups of enterprises (e.g., industry sectors) or other major traffic sources, where added IPv6 capability could help drive gateway IPv6 traffic ratios.

technical requirements for IPv4 addresses start to diminish, they could begin to sell their excess supply to other network operators. This transfer of resources could help to keep the mixed world of IPv4/IPv6 in place by ameliorating some of the growth constraints of the IPv4 Internet. On the other hand, uncertainty about future needs could lead them to hold on the their IPv4 assets.

In sum, the interaction of continued Internet growth, rising prices for IPv4 addresses, the availability of transferable number resources, and the diversion of more and more traffic to native IPv6 networks are crucial to understanding the transition.

5. Summary and Conclusions

This section addresses the key question behind this report: how will all this end? Will IPv6 die out? Will the world converge on IPv6? Or will we live in a mixed world for the foreseeable future? Although we cannot be 100% certain, the most likely scenario is that we are looking at a mixed world for the next 20 years. We see a fast-growing mobile Internet deploying IPv6 and diverting greater and greater amounts of traffic to an IPv6-only network. But the incentives of large, fast-growing network operators aren't the same as those of many other networks on the Internet. IPv6 will dominate some parts of the newer, larger-scale parts of the Internet; legacy IPv4 will coexist in the long tail. Conversion technologies will interface between them for the foreseeable future, as will the use of NAT by IPv4-only networks. The most likely future equilibrium is on a mixed, IPv4/IPv6 Internet.

5.1 Summary of findings

The report examined quantitative data about current levels and patterns of IPv6 adoption, and tried to explain that data based on an analysis of the economic incentives affecting IPv6 deployers. The report characterized IPv4 and IPv6 as competing, incompatible network protocols. IPv4 is an established, legacy technology with high levels of inertia; IPv6 is a newer, incompatible networking protocol that requires additional investments to deploy and forces its deployers to absorb the costs of maintaining compatibility with the legacy protocol. Our analysis focused on the incentives of network operators, as opposed to device/equipment manufacturers or content providers.

We have shown that network operators' IPv6 deployment decisions are discrete and independent. IPv6 deployment patterns, therefore, are best tracked at the AS level. The data show that there are a few markets where one or two major AS's have converted as much as 90% of their network to IPv6 while other major AS's in the same market have no discernable deployment at all. As far as we can tell from available data, IPv6 deployers enjoy no competitive advantage when that happens, although there do seem to be cost efficiencies for large operators in expanding their network via IPv6. While the aggregate trend for IPv6 over the past 7 years is upwards, the deployment trajectory is best understood as an accumulation of discrete decisions by individual network operators to convert all or part of their networks. Consequently, it is not unusual to see plateaus in deployment, at the country level and the AS level. Plateaus occur when network operators complete a commitment to deploy IPv6 in part of their network and remain at that level. Further expansion requires another investment decision. We also see evidence of declines. We do not know for sure what explains them. Some may just be measurement errors, others may be operators turning IPv6 off.

Due to the added costs of IPv6 deployment, there is a strong positive correlation between a country's wealth (measured in per capita GDP) and country-level IPv6 deployment levels. Per capita GDP differences explain nearly half (.499) of the variation in IPv6 deployment levels across countries. The correlation is statistically strong ($<.01$). The study also found that a lower level of market concentration is correlated with higher country-level IPv6 capability rates. This was true of both wireless ($-.267$, $p = <.01$) and fixed broadband ($-.347$, $p = <.01$) service markets. We found no clear evidence indicating that IPv6 deployment creates a major competitive advantage, so the negative correlation between market concentration and IPv6 deployment probably exists for two reasons: 1) the presence of more players in a market increases the likelihood that one of them will make an arbitrary deployment decision; 2) a more open market permits the entry of new firms (such as India's Reliance Jio) with newer infrastructures, which have a more favorable cost structure for IPv6.

5.2 Good news and bad news for IPv6

There is good news and bad news for IPv6 in our findings. The good news is that IPv6 is unlikely to become an orphan. For network operators that need to grow, particularly mobile networks where the software and hardware ecosystem is mostly converted, IPv6 deployment can make economic sense. It mitigates a major constraint on growth and can provide a path out of the operational complexities and costs of large-scale NAT. Our modeling of dual stack and conversion methods shows that for fast-growing networks, an IPv6 deployment can tame - but not immediately eliminate - the requirement for more IPv4 number resources. The key variable is how quickly the traffic ratio (i.e., the portion of traffic that goes to IPv6 and IPv4 hosts) shifts toward IPv6. The rising price of IPv4 numbers provides an additional stimulus to deploy IPv6.

The bad news is that the need for deployers to maintain backwards compatibility with non-deployers eliminates many network effects that would create pressure to convert to IPv6. Additional bad news is that many enterprise networks don't need to grow much and/or may still be lodged in a slower-moving software and hardware ecosystem tied to IPv4. Another issue that emerged from our modeling exercise was that once the IPv6-only traffic ratio among IPv6 deployers reaches a certain level, their IPv4 address requirements start to decline. These operators can therefore release IPv4 address resources into the market that would alleviate shortages and facilitate continued low levels of growth for legacy IPv4 networks.

5.3 Convergence scenarios

While it is possible to posit factors that might push the entire Internet into a convergence on IPv6, it is difficult to come up with scenarios that would unambiguously lead to that outcome.

Perhaps the most likely scenario is that there will be a progressive shrinkage of the cost penalty for deploying IPv6, as legacy infrastructure depreciates and is taken out of production, and software incompatibilities are discovered and resolved. As all networking gear, devices and applications gradually become IPv6 capable, one day we all wake up and discover that IPv4 can be turned off with no effect. That could happen, but how long would it take? Any time horizon under 20 years seems unrealistic. For example, it seems to take Microsoft 12-15 years to kill off obsolete operating systems that are less embedded and dominant than the world's main networking protocol suite.[27] And Microsoft software products have a proprietary sponsor that can coordinate the withdrawal of support; IPv4 does not. Much, albeit not all, new equipment, devices and applications will work with both IP versions. Taking the final step of obsoleting IPv4 would be a globally distributed decision. As IPv4 shrinks as a portion of the world's networking operations, it is difficult to project any plausible scenarios in which major service providers would literally shut down IPv4, because their cost of maintaining residual IPv4 capabilities would be low.

---

[27] In 2016, 15% of the devices scanned in a Trend Micro study were still running Windows XP. (Huq et al, 2017) In 2018, Microsoft itself said that XP still holds a 4.5% market share of the world's operating systems.

Some observers hold hope that the growing scarcity and rising price of IPv4 addresses will push the world into convergence. Our analysis and modeling exercise provided some support for that logic, but mostly undermined it. The rising price of IPv4 numbers and the operational costs of NAT do in fact stimulate IPv6 deployment. But for static networks that already hold the IPv4 number resources they need, that is not a problem. Further, while rising prices will stimulate NAT and IPv6 deployments among fast-growing networks, the success of IPv6 deployers in reducing their own demand for IPv4 numbers could make more IPv4 numbers available for others. The eventual release of unneeded numbers by IPv6 deployers could sustain the technical requirements of other networks for some time to come.

We saw cloud providers buying up huge quantities of IPv4 numbers in order to serve as the bridge between the two Internets. Is it possible for cloud providers to reduce this cost by offering discounts or subsidies to IPv4 networks to convert? While this might create effective pressure on some organizations teetering on the brink of a deployment decision, it is difficult to see how the incentives or discounts from a single vendor could be big enough to push all marginal operators over the edge. Although one of our interviewees suggested that this was a possibility, and some CSPs do offer their services over IPv6, we did not find any evidence of differential pricing among CSPs.

Another possible convergence scenario involves the growing concentration of eyeballs on a few big IPv6-using web sites. A highly concentrated ecosystem might make it possible for the long tail of smaller IPv4-only sites to be cut off by ISPs and other intermediaries without anyone taking much notice. This would create a fragmented Internet that might incentivize IPv6 adoption by lagging networks. The problems of concentration and walled gardens, however, have very little to do with the IPv6 vs IPv4 competition. Content providers and sites that would be excluded under this scenario would be excluded more because of their small following than due to their use of IPv4. Since all major IPv6-enabled content providers already have in place a conversion or dual-stack that makes them compatible with IPv4 resources on the Internet, it is hard to see what economic gains would be realized by ending support for IPv4. The massive growth of Facebook and Google is rooted in an advertising based business model that places a premium on reaching as many people as possible. While the costs of maintaining access to the long tail of IPv4-only users are relatively small, the loss of coverage could be serious. ISPs might also experience increased technical support calls as their users encountered unexpected barriers to access.

There is, however a plausible element to this focus on concentrated providers and CSPs. As big cloud and content providers account for a bigger portion of the internet, it is possible that they could leverage that position to prod the rest of the world to move into IPv6, if there was a major economic benefit they could obtain by doing so. Market leaders frequently play a role in overcoming barriers to collective action. The problem is that neither the economic benefit they would obtain nor the method they would use to facilitate convergence have been clearly specified. Furthermore, the world is a big place and the market leaders in one part of the world (e.g., the U.S. and Europe) are not necessarily leaders in China, India, Russia, or Indonesia.

Given the vast number of countries with no discernable IPv6 deployment, their concentration in developing countries, and the presence of many enterprise networks that do not need to grow, it is difficult to envision a clean convergence on IPv6 any time in the near future. We hasten to add, however, that this conclusion is just an educated forecast, not a scientific prediction based on knowing all potentially relevant facts. There are too many variables and unknowns to be certain about the future course of standards evolution in the world for the next 20 years. Entirely new standards and technologies could come along that could disrupt the entire system, for example. Still, the wisest course of action for the global Internet technical community is to look forward to a mixed IPv4-IPv6 world for the next 20 years and plan accordingly.

References

Arthur, W. B. (1994). *Increasing Returns and Path Dependence in the Economy*. Ann Arbor: University of Michigan Press.

Byrne, C. (2014). 464XLAT: Breaking Free of IPv4. https://www.nanog.org/sites/default/files/wednesday_general_byrne_breakingfree_11.pdf

Choi, J. P. (1996). Do converters facilitate the transition to a new incompatible technology? A dynamic analysis of converters. *International Journal of Industrial Organization, 14*(6), 825-835.

Choi, J. P. (1997). The Provision of (Two-way) Converters in the Transition Process to a New Incompatible Technology. *Journal of Industrial Economics, 45*(2), 139-153.

Economides, N. (1996). The economics of networks. *International Journal of Industrial Organization, 14*(6), 673-699.

Farrell, J., & Saloner, G. (1985). Standardization, Compatibility, and Innovation. *Rand Journal of Economics, 16*, 70-83.

Farrell, J., & Saloner, G. (1986). *Competition, Compatibility and Standards: The Economics of Horses, Penguins and Lemmings*. In Gabel, ed. Product Standardization and Competitive Strategy. Amersterdam: Elsevier.

Farrell, J., & Saloner, G. (1992). Converters, Compatibility, and the Control of Interfaces. *Journal of Industrial Economics, 40*(1), 9-35.

Forman, C. (2014). *The Economics of the Transition to Internet Protocol version 6 (IPv6)*. Paris: OECD

Gartner. (2018). Gartner Forecasts Worldwide Public Cloud Revenue to Grow 17.3 Percent in 2019. Accessed February 6, 2019, from https://www.gartner.com/en/newsroom/press-releases/2018-09-12-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-17-percent-in-2019

Hazlett, T. W. (2005). *Rivalrous Telecommunications Networks With and Without Mandatory Sharing*. Retrieved from Washington D.C.:

Joseph, D., Shetty, N., Shuang, J., & Stoica, I. (2007). *Modeling the Adoption of New Network Architectures.* Paper presented at the Proceedings of the 2007 ACM CoNEXT Conference.

Katz, M., & Shapiro, C. (1986). Technology Adoption in the Presence of Network Externalities. *Journal of Political Economy, 92*(Journal Article), 822.

Lagerholm, S. (2017). Going IPv6 Only. https://www.youtube.com/watch?v=nNMNglk_CvE

Mueller, M. L. (1997). *Universal Service: Competition, Interconnection and Monopoly in the Making of the American Telephone System*. Cambridge, Mass: MIT Press.

Mueller, M., Kuerbis, B. Asghari, H. (2013) "Dimensioning the elephant: an empirical analysis of the IPv4 number market", info, Vol. 15 Issue: 6, pp.6-18.

Parker, G. G., Van Alstyne, M., Choudary, S. P., & Foster, J. (2016). *Platform revolution: How networked markets are transforming the economy and how to make them work for you New York: WW Norton.* New York: WW Norton.

Rohlfs, J. (1974). A Theory of Interdependent Demand for a Communications Service. *Bell Journal of Economics, 5*(1), 16-37.

Seifert, R. W. & Varé, A. X. (2009) Adoption of network technologies in the presence of converters, *Economics of Innovation and New Technology*, 18:1, 69-91.

Sen, S., Jin, Y., Guérin, R., & Hosanagar, K. (2010). Modeling the Dynamics of Network Technology Adoption and the Role of Converters. *IEEE/ACM Transactions on Networking, 18*(6), 1793-1805.

Shapiro, C., & Varian, H. (1999). *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard Business School Press.

Huq, N., Hilt, S., & Hellberg, N. (2017). US Cities Exposed: A Shodan-Based Security Study of Exposed Assets in the US. Tokyo: Trend Micro. https://www.trendmicro.com/content/dam/trendmicro/global/en/security-intelligence/research/reports/wp-us-cities-exposed.pdf

T-Mobile US. (2018). Total number of customers/subscribers of T-Mobile US from 1st quarter 2010 to 3rd quarter 2018 (in millions). In Statista - The Statistics Portal. Retrieved December 4, 2018, from https://www.statista.com/statistics/219577/total-customers-of-t-mobile-usa-by-quarter/.

Wu, P., Cui, Y., Wu, J., Liu, J., & Metz, C. (2013). Transition from IPv4 to IPv6: A State-of-the-Art Survey. *IEEE Communications Surveys & Tutorials, 15*(3), 1407-1423.