

Unclassified

DSTI/ICCP/CISP(2014)3/FINAL

Organisation de Coopération et de Développement Économiques
Organisation for Economic Co-operation and Development

28-Nov-2014

English - Or. English

DIRECTORATE FOR SCIENCE, TECHNOLOGY AND INNOVATION
COMMITTEE ON DIGITAL ECONOMY POLICY

Working Party on Communication Infrastructures and Services Policy

The Economics of the Transition to Internet Protocol version 6 (IPv6)

S. Paltridge

JT03367357

Complete document available on OLIS in its original format

This document and any map included herein are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area.

DSTI/ICCP/CISP(2014)3/FINAL
Unclassified

English - Or. English

FOREWORD

This report was presented to the Working Party on Communication, Infrastructures and Services Policy (CISP) in June 2014 and the CISP agreed to recommend it for declassification to the Committee on Digital Economy Policy (CDEP). The CDEP Committee approved the report in October 2014. The document was prepared by Chris Forman from the Georgia Institute of Technology.

Note to Delegations

This document is also available on OECD iLibrary. Please cite as:
OECD (2014), "The Economics of Transition to Internet Protocol version 6 (IPv6)",
OECD Digital Economy Papers, No. 244, OECD Publishing.
DOI: [10.1787/5jxt46d07bhc-en](https://doi.org/10.1787/5jxt46d07bhc-en)

TABLE OF CONTENTS

FOREWORD.....	2
MAIN POINTS	4
1. INTRODUCTION	7
2. SOME FRAMEWORKS TO UNDERSTAND THE TRANSITION TO IPV6	11
2.1 Some economics of not moving to IPv6	19
3. OVERVIEW OF BENEFITS AND COSTS FOR DIFFERENT PLAYERS WITHIN THE PLATFORM.....	25
3.1 Infrastructure vendors	25
3.2 Wireline ISPs	27
3.3 Issues Specific to Mobile Providers	30
3.4 Content providers.....	32
3.5 Enterprises – Internal	33
3.6 Consumer End Users.....	34
3.7 Comments on visibility of costs and benefits	34
4. ALTERNATIVE TRANSITION SCENARIOS AND ATTEMPTS TO FACILITATE TRANSITION TO IPV6	36
4.1 Alternative transition scenarios and S-Curves	36
4.2 Alternative policy responses	39
5. CONCLUSION	41
ANNEX 1: IPV6 IN CONSUMER ELECTRONICS DEVICES AND IN THE INTERNET OF THINGS	42
BIBLIOGRAPHY	44
NOTES	52

MAIN POINTS

Internet Protocol (IP) defines the address space for the Internet. The number of addresses defined by Internet Protocol version 4 (IPv4), the version of IP that has been used since the start of the commercialisation of the Internet, is currently running out. A successor to IPv4, known as Internet Protocol version 6 (IPv6) has been available since 1998. However, the diffusion of this new protocol has been slow: by some measures the worldwide traffic over IPv6 stood at roughly 3.5% as of April 2014.

The transition to IPv6 represents an example of a diffusion of a new technology standard. Economics has developed a body of theory to identify and understand some of the key features influencing diffusion. This document describes some standard economic frameworks to understand the diffusion of new technology standards, and applies them to understand the current state of transition to IPv6. It then describes how these frameworks can be used to understand the implications of various coordination mechanisms designed to facilitate the transition to IPv6.

The report presents a canonical model of technology adoption under which users adopt a new technology when benefits exceed costs, and when the net benefits (benefits minus costs) today exceed the net benefits of adopting at any other point in the future. The document focuses particularly on the case of adoption of a new platform, in which users make investments that are specific to the platform, some of which may be long-lived. It is noted that platforms may be subject to direct network effects where the value to a user of consuming a product is increasing in the number of other users of the product. They may also be subject to indirect network effects, which occur when the value of a good depends upon the provision of a complementary good, which in turn depends on the demand for the original good. One of the features of the canonical model is that the presence of direct and indirect network effects can make it harder for a new platform standard to achieve critical mass because sellers and users are unable to coordinate their behaviour. Second, expectations of other user behaviour may be important; an actor will adopt today only if they believe that other users are likely to adopt.

The report makes the case that IPv6 represents an example of a platform. Within the context of IPv6, the sides of the platform are Internet service providers, backbone providers, device manufacturers, content providers, and so forth. The net benefits to adopting the new platform are not distributed equally across sides. For some participants, such as backbone and transit providers and manufacturers of devices such as routers, the transition to IPv6 has been relatively swift. For these participants the benefits of adoption were clear, and adoption demonstrated the technical ability of the company and fitness of its network. For others, such as many content providers for the Web, enterprises contemplating deployment of IPv6 within internal firm networks, and providers of consumer electronics equipment such as DVD or Blu-Ray players or televisions, the transition has been slower. For them the benefits have not been as clear, and many legacy devices, networks, customers, suppliers have not transitioned.

While the diffusion of IPv6 shares many features of the canonical model, it also contains elements that are not featured in many prominent models. These have shaped its diffusion. First and foremost is that most of the benefits to adopting IPv6 come from having access to an expanded IP address space. However, in most current implementations of IPv6, users run both IPv4 and IPv6 in dual stack, which continue to require some IPv4 addresses. (In this report, users refer to firms who may make decisions of whether or not to invest in IPv4 as part of using the Internet to deliver products and/or services to others. Thus, users may include Internet Service Providers, providers of economic content, and enterprises considering integrating IPv6 technology in their internal firm networks. In general, these do not include end users such as consumers.) As a result, many of the benefits of adopting IPv6 today will not be fully realised until a user has been able to turn off its support for IPv4 in the future. For example, ISPs adopting IPv6 must also

deploy converter technologies that enable incompatible networks to communicate so that their users can also access IPv4 content.¹ These converter technologies will be necessary until the vast majority of the network deploys IPv6. More broadly, in most cases users adopting dual stack IPv6 must continue to support both IPv4 and IPv6 until most Internet traffic takes place over IPv6. Thus, laggards in adoption may disproportionately affect adoption relative to the diffusion of other technologies displaying network effects, where the benefits of adopting may be increasing in the adoption behaviour of others but where the benefits of adoption do not depend upon near 100% penetration of the new technology. This feature increases the incentives to delay adoption.

Another important feature in the diffusion of IPv6 is that it cannot be implemented “out of the box” and often successful implementation requires solving new and unexpected technical challenges. As has been common in previous waves of information technology diffusion, one of the major costs of adoption are the costs to adapt the new IT to the idiosyncratic business environment in which it is used. These costs are highest in the early stages of diffusion. In the early years of its diffusion, well-developed markets for skills for implementing IPv6 did not exist, and these have slowed its diffusion. As familiarity with IPv6 increases, these barriers to its diffusion will likely fall.

A further factor which has weighed on diffusion is that over time there have been significant efforts to ensure that the old standard, IPv4, can continue to work in an environment of address depletion. For example, there is evidence that actions taken by vendors have reduced the costs to users of remaining behind NATs, an issue that is described in further detail in this document.² These efforts have been necessary given the slow diffusion of IPv6. However, they have also decreased incentives to adopt, since many users may assume they will be supported for some time to come.

Economists have noted that a focal user or application can play an important role in facilitating the transition to a new standard. Governments and large companies can sometimes play this role. Given the decentralised distribution of economic activity on the Internet, no such user can play this role during the transition to IPv6. However, there are some countries in which large ISPs have adopted IPv6 and now have a significant fraction of their traffic over the new protocol. Due in part to their unique economics, IPv6 penetration for some mobile providers is higher than for wireline. Similarly, some of the largest content providers such as Google, Facebook, and Yahoo now all support IPv6 connections. Given the high percentage of traffic over these sites, this may increase the net benefits of adopting for other players in the platform.

However, amidst these positive developments it should be noted that the typical dual stack implementation of IPv6, under which a set of IPv4 and IPv6 addresses are needed until most of the network traffic has transitioned to IPv6, may continue to weaken incentives for adoption. This may be particularly true for smaller users with low traffic growth who may have less need for new IP addresses and whose relative costs of adopting IPv6 may be high.

The diffusion of IPv6 is at an early stage, and the discussion above suggests a range of potential outcomes. One interesting question is at what point declining costs (driven by a thickening of markets for IPv6-related products and services) and increasing expected future benefits (influenced by changes in perceptions about the future path of diffusion) will engender significant adoption among smaller users who are by and large not currently contemplating adoption. In other words, in an environment with competing standards, an open question will be whether diffusion of the new technology will stall, as it has for a range of technologies, either temporarily or permanently (e.g., the Fax (Rohlfs 2001), 56K modems (Augereau, Greenstein, and Rysman 2006)).

Not transitioning to IPv6 will likely have a range of economic implications. Existing technologies that enable the sharing of limited IP addresses break some applications, and disrupt the modularity of the

Internet. This will degrade the quality of Internet service for many users, and increase the costs of developing new applications and services. The economic costs of address run-out will be asymmetric, affecting some products, services, and providers more than others. New markets have arisen to trade existing IPv4 addresses. However, as of the time of this report even in the face of address run-out activity in this market has been limited.

Governments and non-governmental institutions in the technical community such as the Internet Society have sought to facilitate the transition to IPv6. These have included efforts to begin to diffuse best practices for implementing IPv6, and publishing information on IPv6 deployments. These efforts are valuable as they help inform adopters about the potential costs and benefits of adoption, though the availability of some types of data to inform decision-makers could be improved. There have also been successful efforts at coordinating the behaviours of large ISPs and content providers by publishing data on IPv6 penetration through sources such as World IPv6 Launch. While adoption remains very low, these policies have shown some success at influencing the behaviour of lead users. One open question is the extent to which these policies will be effective at encouraging adopters by other users on the IPv6 platform.


1. INTRODUCTION

The diffusion of the commercial Internet has had a transformative effect on economies throughout the world.³ It has led to an increase in economic growth, and the transformation of markets and industries in ways that could hardly have been imagined 20 years ago when the Internet was being transformed from a predominately research network into the backbone for commerce that it has become today.⁴

Many communications on the Internet today remain based on versions of protocols that were developed many years ago. The most highly cited Internet standards remain Transmission Control Protocol (TCP) and Internet Protocol (IP), both published in the early 1980s.⁵ While there has been a burst of standardisation activity since 1990 associated with the commercialisation of the Internet and the need to develop new standards to support new types of applications and uses for the Internet these changes have continued to be built around these two protocols.⁶


Among other things, Internet Protocol includes the address space for the Internet. The design of Internet Protocol based on RFC 791, known as Internet Protocol version 4 (IPv4), uses a 32-bit address space that theoretically provided for over 4.4 billion addresses to be connected to the Internet. This was a very large address space in its time and very significant considering the much smaller address spaces of some contemporary protocols.⁷ However, due to the growth in the use of the Internet and the accompanying growth in devices requiring an IP address, this address space is currently running out (Figure 1). In February 2011 the Internet Assigned Numbers Authority (IANA) announced that the pool of available IPv4 addresses had run out. Each of the Regional Internet Registries (RIRs) are also either running out or have run out of addresses. The RIR for the Asia-Pacific Region (APNIC) has exhausted its pool of addresses as of April 2011, the RIR for Europe and the Middle East (RIPE NCC) has run out as of September 2012, and the RIR for Latin America (LACNIC) exhausted its stock of general use IPv4 addresses in June 2014. Recent estimates (as of July 2014) project that the RIR covering North America and the Caribbean Islands will exhaust their addresses for general use allocation in February 2015, and the RIR for Africa (AFRINIC) will run out in June of 2020.⁸

Figure 1. RIR IPv4 Address Run-Down



Source IPv4 Address Report, <http://www.potaroo.net/tools/ipv4/index.html>

A successor to IPv4, Internet Protocol version 6 (IPv6), has been available since 1998.⁹ While a variety of things have changed with the transition from IPv4 to IPv6, the most important is a significantly increased address space.¹⁰ The address field has increased from 32 bits to 128 bits in length, allowing for a theoretical maximum of 2^{128} addresses, offering a potential solution to the shortage of addresses described above.¹¹ While IPv6 offers a solution to the address problem described above, its diffusion has been very slow. While there are a number of ways of measuring the transition to IPv6, one way is to measure IPv6 connectivity to large web sites such as Google. Figure 2 displays some recent statistics on IPv6 diffusion from Google. Recent statistics suggest that worldwide over 3% of users access Google over IPv6 as of March 2014. This represents a significant acceleration from the close to 0% connection during the first decade of IPv6's diffusion, but remains at a low level.^{12 13 14}

Figure 2. Historical IPv6 Connectivity Based on Google Connectivity Statistics

Graph shows the percentage of users that access Google over IPv6.

Source: <https://www.google.com/intl/en/ipv6/statistics.html>.

In this document some standard economic frameworks are applied to inform understanding of the transition to IPv6. The starting point for this discussion is IPv6's incompatibility with IPv4. Thus the transition from IPv4 to IPv6 provides an example of the diffusion of a new compatibility standard that is incompatible with a previous standard. In this report, the economics literature is followed in defining compatibility standards as a set of design rules that promote interoperability (e.g., Farrell and Simcoe 2013; David and Greenstein 1990). Under this definition, the protocols that enable communication between Internet devices is an example of a compatibility standard. For brevity, this report follows convention in economics and refers simply to standards throughout. It is possible to apply lessons learned from prior economic studies of the diffusion of incompatible standards to understand the transition to IPv6. However, this report will also emphasise some of the key institutional factors that may be somewhat distinct from some prior examples and which are important to understanding the diffusion of IPv6.

This document begins in the next section by defining some standard economic frameworks that can be used to understand the diffusion of new standards and relate them to the diffusion of IPv6. These frameworks are used to understand the economic incentives of adoption for different players along the Internet value chain. The document concludes by discussing some alternative transition scenarios, and how these might be influenced by actions taken to encourage the transition to IPv6.

This document represents a preliminary investigation into the economic factors shaping the IPv6 transition. It is intentionally broad in scope, discussing many different economic actors shaping the diffusion of the protocol and many potential factors influencing diffusion. As such, the goal of this report is to motivate further discussion and further analysis, rather than to provide final answers.

Further, while the scope of the study is broad, its scope has been limited in several important ways. In particular, while this report acknowledges the considerable heterogeneity in cross-country factors influencing the diffusion of IPv6, it will not undertake a systematic study of country-level environmental factors that might facilitate or hinder IPv6 diffusion. Both of these are important features in need of future work.

2. SOME FRAMEWORKS TO UNDERSTAND THE TRANSITION TO IPV6

This section discusses some well-known economic frameworks that are known to inform the rate of transition to a new standard. The basic motivation for the transition to IPv6 is, as noted above, the depletion of IPv4 addresses. While technologies and approaches exist for mitigating the costs of address depletion by, for example, sharing IPv4 addresses among users, these approaches have been known to degrade the quality of Internet service for users and have important implications for user privacy and accountability.¹⁵ Individual users may have widely varying incentives for transitioning to IPv6 however, and this section investigates the implications of these differences. In this report, users refer to firms who may make decisions of whether or not to invest in IPv4 in the process of using the Internet to deliver products and/or services to others. Thus, users may include Internet Service Providers, providers of economic content, and enterprises considering integrating IPv6 technology in their internal firm networks. In general, the label users will not include “end users” such as consumers unless explicitly stated.

This section first presents some general frameworks and findings from canonical models of the adoption and diffusion of new technologies and how these may be affected by the presence of incompatible standards. It next details some specific ways in which the diffusion of IPv6 departs from these canonical models, and how this might influence equilibrium predictions. Finally, the document discusses some specific historical examples of the diffusion of new technologies, their similarities and differences to the IPv6 context, and in turn how they might inform understanding of the transition to IPv6.

As noted above, IPv6 is not compatible with IPv4, so there is a problem of maintaining compatibility between two different standards while users and providers of complementary hardware and software transition from IPv4 to IPv6. These different users and producers make coordinated investments around a computing *platform*. Platforms allow users and producers to share interchangeable components, so that they can coordinate their economic activity and share in the benefits of technical advance (Bresnahan and Greenstein 1999). Examples of IT-based platforms include those based around different generations of Windows software (and Intel and AMD microchips), smartphones such as the iPhone and Android, videogame platforms such as Microsoft Xbox 360 and Xbox One, as well as historical examples such as the IBM System/360 platform. Platforms have interchangeable components, often from different manufacturers, that allow users to use the same platform over time and to make long-lived investments in software and human capital with the knowledge that they will not quickly be degraded in an environment of rapid technological change.

A significant economic theory¹⁶ has developed that argues that platforms will tend to be persistent under a set of assumptions (Bresnahan and Greenstein 1999): (1) that buyers and sellers will make investments that are specific to the platform (i.e., cannot be reused across other incompatible platforms); and (2) that some of these investments will be long-lived, and so will influence the returns to adopting future technological improvements or new standards. Such standards frequently exhibit *positive feedback*: That is, the value of adopting the standard is increasing in the number of other adopters or exhibits increasing returns. This makes existing platform standards very difficult to stop, and also makes it very difficult for new standards to achieve acceptance. For these reasons, platforms tend to exhibit considerable persistence, and established platforms tend to remain dominant for long periods of time.

An important question is how incompatible platforms can achieve compatibility, or interoperate, with one another. Farrell and Simcoe (2013) describe four paths to achieving compatibility between products or

platforms. The first is decentralised choice, where the utility-maximising decisions of individual consumers lead to the establishment of a standard. The second is negotiations through voluntary Standard Setting Organizations (SSOs).¹⁷ Another is through leadership by an influential decision maker, such as a platform leader or the government (for example through (mandatory) standard setting or procurement rules). A further approach is to achieve compatibility through converters or through multi-homing (participating in multiple platforms simultaneously).

While IPv6 is the result of a formal standard-setting process, because neither the IETF nor any other authority is able to compel buyers and sellers to transition to IPv6, its transition has included elements of all four paths to compatibility. The role that these other mechanisms have played in influencing the transition to IPv6 is discussed below.¹⁸

Decentralised adoption

It is useful to first view adoption of IPv6 through the lens of decentralised adoption, as the transition to IPv6 will occur only through the uncoordinated decisions of individual users. In keeping with frameworks on technology adoption, here the focus is primarily on the behaviour of users to adopt IPv6 within their organisations, rather than the decisions of producers to incorporate IPv6 into their products. However, as is discussed in further detail in section 3, producers of complementary hardware and software will similarly trade off costs and benefits when deciding whether to incorporate IPv6 into their products. Here the document presents a standard canonical model of technology adoption and later discusses important institutional features of the case of IPv6 diffusion.

The adoption of IPv6 can be viewed through the well-known “probit model” (e.g., David 1969; Karshenas and Stoneman 1993) of technology adoption and diffusion, in which a user adopts a new technology at some time period when two conditions hold: (1) the benefits of adoption exceed the costs and (2) it is better to adopt at this time period than any other. In formal notation, the first condition means that users will adopt when $NB(x_i, t) \equiv B(x_i, t) - C(x_i, t) > 0$, where NB is equal to the net benefits of adoption, B is the gross benefits of adoption, and C is the costs of adoption. Note that each of these terms represent the flows of all future benefits of costs discounted back to time period t . x_i is often defined broadly as a set of firm-specific and external conditions that shape the costs and benefits to adoption. These could include, for example, the internal technological capabilities of users such as how much experience a user has with IPv6 and related technologies. It could also include external factors shaping adoption such as the adoption decisions of users in other parts of the platform—as might happen if the adoption decisions of content providers influence the adoption decisions of ISPs—as well as external factors such as intervention by governments or other organisations to encourage adoption of IPv6. Formally, the second condition means that a user will adopt at time t when $NB(x_i, t) > NB(x_i, t')$ for any $t' \neq t$. Adoption costs generally decline over time for new technologies; for example, in the case of IPv6, it may become easier and less costly to fix problems that commonly arise in IPv6 implementations because of increased experience within the community.¹⁹ As costs decline over time, high valuation users (those with high gross benefits) will adopt first, followed over time by successively lower valuations users (those with low gross benefits). Note that in general because B and C represent expectations of future flows, they are uncertain. Further, if the benefits of adoption depend upon the decisions of others, and expectations are that other potential users do not adopt in the future, then the expected value of B could be very small. In particular, one point that we describe in further detail below is that many of the benefits of IPv6 will be realised only when most other users have adopted IPv6 and when the focal user can in turn off support for IPv4. Thus, if users believe that outcome is unlikely, the value of B could be very low.

The approach here uses this very simple framework because it provides a straightforward way of thinking about the incentives to adopt IPv6. As discussed in further detail below, even if for some players

the costs of adoption do not seem to be especially high, the benefits are uncertain and will occur primarily in the future. Thus, for many users (and sellers) it may be best to wait.

It is useful at this point to reemphasise the role of positive feedback on influencing adoption decisions in platforms. Positive feedback is a common feature of platforms—there are increasing returns to adoption and the number of adopters needs to get over a critical point for the platform to be successful. Positive feedback means the strong get stronger and the weak get weaker (e.g., Shapiro and Varian 1999). In the setting here, positive feedback arises from demand-side economies of scale or network effects, in which the value of consuming the product for a user increases with the number of other users (e.g., Rohlfs 1974, 2001). Network effects arise in many settings, including telecommunications networks, personal computers, mobile phones and of modern social networking applications like Facebook and Twitter. Network effects are sometimes described using Metcalfe's Law, i.e. that the value of a network increases in the number of users.²⁰

Information technology goods in platforms are also sometimes characterised by indirect network effects, which occur when the value of a good or its demand depend upon the provision of some complementary good, which in turn depends upon the demand for the original good (e.g., Rysman 2009). Such indirect network effects arise in many settings, as when demand for a computer operating system depends upon the number of software titles or when the demand for a video game console depends upon the number of games. Indirect network effects can be a useful tool for understanding the diffusion of IPv6. For example, in general content providers will only adopt IPv6 when end users can access their content over an IPv6 connection, which in turn is likely to be provided only once there is IPv6 content available. Note that many large content providers such as Google, Facebook, and Yahoo currently provide content over IPv6. These content providers account for a disproportionate share of IPv6 traffic. However, the majority of content providers do not support IPv6 (this will be discussed in further detail below). As mentioned above and will be described in further detail elsewhere, the benefits of adopting dual stack IPv6 will only be realised once most participants in the platform adopt IPv6. Thus, even with the largest providers supporting IPv6 connections, indirect network effects are likely to play some role in shaping behaviour in this market.

Markets that exhibit network effects have important implications for understanding adoption of IPv6. First, they sometimes exhibit what are known as network externalities, which arise when one user's adoption affects the value that another receives from the good, but the first user does not receive compensation. (In the setting used here, one actor's adoption of IPv6 affects the value that others achieve from adopting IPv6, but they do not get paid for it.) Many of the efforts to promote IPv6 adoption have been aimed at addressing these issues. For example, efforts such as World IPv6 Launch (<http://www.worldipv6launch.org/>) and World IPv6 Day (<http://internetsociety.org/ipv6/archive-2011-world-ipv6-day>) are events that both improved coordination among users and may also shape their expectations about future adoption of IPv6.

Further, adoption of early users can influence adoption of later users. In the discussion of the probit model of adoption above, it was noted that the behaviour of adopters and non-adopters identifies the difference between high value and low value users. When network effects are important, the adoption of a sufficient critical mass of high value users is needed to start the bandwagon rolling. This has also begun to happen in the diffusion of IPv6, as large ISPs such as Comcast, Verizon Wireless, AT&T, Free, Deutsche Telekom, and KDDI all now have significant IPv6 deployments, as do some large content providers such as Google, Facebook, and Yahoo.²¹ However, so far their adoption has not given significant enough trust to move others to adopt IPv6 on an increasing and accelerating larger scale.

For many platforms the value of adopting is more or less increasing with the supply of complementary goods; for example, the value of CD players are increasing with the number of titles

(Gandal, Kende, and Rob 2000), the value of a VCR platform is increasing with the number of movies (Ohashi 2003), and the value of a video game platform is increasing with the number of games (e.g., Corts and Lederman 2009). IPv6 is different in that the value of adopting the platform only increases once the entire network migrates (or something very close to it) to the platform. Until then, users will need to support both protocols. This feature shows up in current usage statistics. While some broadband networks have deployed IPv6 widely as have some web sites, a server will record an IPv6 transaction only when all of the following occurs (OECD 2010, p.31): “(1) The client has an IPv6 stack; (2) The client’s application is configured with IPv6 support; (3) The client’s DNS configuration is able to perform an IPv6 address query; and (4) The client and server can communicate end-to-end using IPv6.” In other words, a server will record a transaction when all of the elements in the end-to-end connection are IPv6-enabled.

As suggested above, many platform markets do not share this characteristic. While it has been seen in other contexts—railroad gauges and digital television are two examples—IPv6 is unusual because the network is truly worldwide and because, at least relative to these other cases, there are few influential users to help facilitate transition to the new platform. For example, Puffert (1989) describes how the United States transition from “non-standard” (width 5’ 0”) southern railway track gauge to “standard” (4’ 8 ½”) gauge used east of the Rocky Mountains was facilitated by an active cartel organisation. The transition to digital television was also facilitated by regulators. This discontinuity of benefits when adoption approaches 100% makes some of the coordination challenges described above particularly challenging. Further, as described in further detail in section 4, it means that policies designed to improve coordination among market participants will also need to consider how to encourage adoption among lower value users. Such responses may need to be different than those targeting higher value users.

Large players imposing a standard

Another way to achieve coordination is for a large player to simply impose a standard (e.g., Farrell and Simcoe 2013). One possibility is for a platform leader to play this role.²² An example of this from earlier in the history of the Internet was the transition from NCP to IPv4 protocols.²³ The precursor to the Internet was ARPANET, an experimental packet-switched network sponsored by the Advanced Research Project Agency (ARPA) of the United States Department of Defense (DOD). ARPANET had used a protocol called Network Control Protocol (NCP), but for ARPANET to become an operational “network of networks” it needed to migrate to the then-new Internet Protocol (IPv4).

At the time ARPANET was managed by the Defense Communications Agency (DCA), which announced that it would transition to IPv4 on January 1, 1983. ARPANET was working well and many users had little incentive to switch to the new protocol. As a result, users resisted the change. The DCA undertook an extensive awareness campaign to facilitate the transition and coordinate user behaviour. NCP was even turned off twice on ARPANET during 1982, leaving non-IPv4 hosts offline. The contrast with the transition to IPv6 is instructive. The transition to IPv4 had a firm cutover date and a platform leader that could enforce the deadline (Cannon 2010). Users who failed to make the transition on January 1, 1983 were cut off from the network. Of course, no such platform leader currently exists for coordinating behaviour on the Internet and to facilitate the transition to IPv6.

Sometimes large players in a platform—customers or complementors—can play a role as a *de facto* standard setter. For example, movie studios played an important role in resolving the standards war between Blu-ray and HD-DVD (Flamm 2013). Several large players have similarly played a role in encouraging providers of complementary goods and services to transition to IPv6. One such example is Verizon Wireless, who has required any manufacturers producing devices for Verizon Wireless’s Long Term Evolution (LTE) network to support IPv6. Verizon Wireless’s LTE network is fully IPv6 enabled. As a result, IPv6 deployment at Verizon Wireless currently stands at close to 50%²⁴ and this requirement may

have increased the availability of IPv6 technology in mobile handsets, a key complementary good for the IPv6 platform. Recent (2014) analysis has shown that 40.6% of all LTE device types are IPv6-enabled.²⁵

Governments often play the “large user” role in encouraging the deployment of standards. Large buyers can require equipment that is compatible with a new standard, and also encourage the development of skills and expertise that can spillover to other organisations. Finally, governments can also develop new methods for testing the deployments of new standards.

Governments have sought to play the large user role in encouraging the transition to IPv6. As an example, the United States government has sought to play some role in encouraging IPv6 in all three of these ways. In 2003, the Department of Defense (DOD) set a 2008 deadline for deploying IPv6. The United States Office of Management and Budget (OMB) later set a similar mandate for other parts of the United States government. More recently, the United States CIO issued a mandate that by September 2012, all public content should be IPv6-enabled and by September 2014 all internal systems should be dual-stack enabled.²⁶ While these deadlines have not been met, such efforts stimulate the supply of complementary goods. Further, the experience with deployment facilitates human capital accumulation on how to deploy and use IPv6 and may have also resulted in the introduction of new product features on the vendor side, as vendors worked with the government on problems that may have risen during deployment.²⁷ Unfortunately, however, the public sector is a relatively small market participant on the Internet platform, which has limited their ability to play the “large user” role and so has limited their ability to influence the IPv6 transition. The document discusses this in further detail below.

Converters and multi-homing

Another way of achieving compatibility is to use converters, translators, and multi-homing. Converters and translators offer a means for two incompatible networks to communicate with one another, and have a long history in playing a role of maintaining compatibility between different standards, such as different word processing and video game formats. In this report the standard definition in economics of multi-homing is used: Multi-homing occurs when one set of economic agents participates in multiple platforms. Note that this is different than its use in routing, where the term has a very different meaning.

Due to the very large installed base of IPv4 deployments and the extent to which these are embedded in existing networks, converters will be needed for communication between the IPv4 and IPv6 stacks for some time. There are three approaches to obtaining compatibility between IPv4 and IPv6 deployments: dual stack, tunnelling, and translation.²⁸ In principle, when communicating with other IPv6 devices dual stack implementations act like an IPv6-only node, and when communicating with other IPv4 devices they act like an IPv4-only node. Tunnelling involves packaging an IPv6 packet so it can travel across an IPv4 network, or packaging an IPv4 packet so it can travel across an IPv6 network. For example, tunnelling can connect disparate islands of IPv6 networks that are separated by IPv4 networks. Under translation, devices translate an IPv4 packet into an IPv6 packet and vice versa.

Approaches like tunnelling and translation are intended to help ease the transition between the old and new standards. However, prior research studying earlier technologies has presented mixed results on the implications of converters for achieving compatibility. For example, David and Bunn (1988), in a study of standards competition between the AC and DC electrical supply systems, examine the impact of *one-way* converters, those that allow one of the technologies to obtain the network benefits of the other technology, but not vice-versa. They show that the introduction of a converter can “tip” the market in favour of one standard. However, models that examine the implications of two-way converters show that they can in fact hinder compatibility. Farrell and Saloner (1992) show that the existence of a symmetric two-way converter can actually lead to less compatibility between two incompatible standards than would exist in the absence of the converter. Choi (1996) explores a setting closer to the diffusion of IPv6, investigating the

implications of a two-way converter that exists to provide compatibility between an old and new standard. He shows that incompatibility between the two standards can facilitate migration to the new standard, as users of the old standard may have some fear of being stranded on an older technology. As converters reduce the risk of being orphaned on the old technology, they slow the transition to the new technology (relative to that which would have occurred in the absence of the converter). This brings to light a question that will arise elsewhere in this document; the extent to which approaches to ease the transition to IPv6 for users may in fact be slowing the transition.

Dual stack implementations are an example of multi-homing: service providers and content developers can support multiple platforms. There are many examples of platforms where agents multi-home. For example, in video game platforms such as Microsoft Xbox or Sony's PlayStation, game developers multi-home while consumers often single-home, choosing only one platform.²⁹ Some research has shown that multi-homing may inhibit the standardisation to a single platform. For example, Corts and Lederman (2009) show in the video game market that multi-homing among video game developers has led to a decline in the market concentration of the video game console market over time. The implications for the transition to IPv6 are that as certain "sides" of the Internet platform multi-home or offer converter technology (e.g., broadband providers, content providers), it may decrease the incentives for other sides of the platform to transition to IPv6.

A common problem with converters is that technically they do not always work well.³⁰ For example, the device driver software in personal computers is one of the most frequent causes of system failure (Ganapathi et al. 2006). Such problems also arise in the transition technologies for the migration from IPv4 to IPv6. For example, running a dual stack network adds additional complexity to managing the network, due both to the complexity of managing two protocols, the additional features of IPv6, and the way in which IPv6 and IPv4 interact with one another. This can have important security implications, and also add additional overhead that can have important implications for bandwidth, server capacity, and latency (e.g., InterConnect Communications 2012). Tunnelling and translation each have their own limitations. For example, some tunnelling approaches like "6to4" have security issues and cannot be reached by certain nodes on the IPv6 Internet. Translation creates a separate set of problems. For example, the Network Address Translation - Protocol Translation (NAT-PT) is an IPv6 to IPv4 translation mechanism, which allows IPv6-only devices to communicate with IPv4-only devices and vice versa. NAT-PT has a serious set of problems that are documented in RFC 4966 and which caused the IETF to deprecate its use in 2007.^{31 32}

Important factors influencing the diffusion of IPv6

In sum, the probit model provides a framework to understand the diffusion of IPv6. Under the probit model, users adopt when benefits exceed costs, and when the net benefits of adopting today are greater than the net benefits of adopting at any other time in the future. When those benefits are low or will accrue far in the future, this mutes incentives for adoption. When those benefits are uncertain, it will increase the incentives to wait until more information about the benefits of adoption can be obtained. Further, when the benefits of adoption exhibit network effects, they will be increasing in the adoption of other agents in the platform; expected future benefits will in turn depend upon the potential adopters' expectations about the adoption decisions of other users. Importantly, under such a model, it is possible to arrive at a situation where high value users (those with high benefits and/or low costs) adopt, but others choose to wait. The document returns to this point in further detail below.

The discussion in the previous section and in the paragraph above focuses upon a fairly standard model of technology adoption under network effects. Extensions to the standard model were discussed to refer to various ways in which compatibility is obtained between incompatible standards. Like all such models, they represent attempts to identify the important features shaping the behaviour of economic

agents and how these features can shape outcomes, and in general are written to represent a broad range of circumstances rather than particular cases. This document next turns to some institutional features that have shaped IPv6 diffusion.

The diffusion of IPv6 has been shaped by a range of factors. While it shares important features of the canonical model it has also been shaped by some unique features, some of which are not featured prominently in the basic model discussed above. This section highlights some of these and discuss their potential effects on IPv6 diffusion.

First and foremost, it is important to remember that the benefits to adopting IPv6 come from having access to an expanded address space. However, in most current implementations of IPv6 users run both IPv4 and IPv6 in dual stack, and continue to require some IPv4 addresses. As a result, many of the benefits of adopting IPv6 today will not be fully realised until a user has been able to turn off its support for IPv4 in the future. This in turn can occur only when the vast majority of traffic has switched to IPv6. As noted above, this is different from many other products that exhibit network effects, where the value of adoption may either be increasing continuously in the number of other adopters on the same side or other side of the platform (e.g., the value of a music player is increasing in the number of available titles) or there exists some localised network from which a network may obtain value (e.g., the value of a bank electronic payment system might be increasing in the number of other banks adopting in the same geographic area (Gowrisankaran and Stavins 2004)). This feature limits the potential range of outcomes where adopting IPv6 provides benefits to the user.

Another key feature, which is consistent with the canonical model but is particularly important to emphasise here, is that there is evidence of significant differences in the costs to implementing IPv6 across organisations. As with many initiatives that require adopting a new piece of enterprise infrastructure technology, adoption of IPv6 involves integration of new technology with new systems. Organisations must learn how to use the technology within its particular context, i.e. things do not often work “out-of-the-box.” With new technologies, there often do not exist well-developed labour markets or third party consultants to assist with the transition, and vendor products may include only incomplete support. All of these issues were present in the early stages of the transition to IPv6, and as of the writing of this article continue to exist. Users will have varying in-house capabilities to respond to these issues, and to some extent these differences will help define the margin between adopter and non-adopter. Over time, increasing adoption typically will lead to thicker markets for all of these products and services and lead to declines in the mean and variance of these costs across firms. However, as of today there remains significant uncertainty in the rate with which these declines will occur.

Moreover, there have been a range of efforts to facilitate communications over IPv4 even in an environment of public IP address exhaustion. These efforts are necessary to maintain service quality over the Internet, particularly when the percentage of end-to-end connections over IPv6 number in the single digits. However, these efforts to improve IPv4 communications have further muted the incentives for firms to invest in IPv6.

All of this points to a range of potential experiences for adopters considering adoption. A range of large service providers such as Verizon Wireless, AT&T, Comcast, Free, Deutsche Telekom, and Telenet have significant levels of penetration. Further, some large content providers such as Google, Facebook, and Yahoo have similarly supported connections over the protocol. One interesting question is at what point declining costs (driven by a thickening of markets for IPv6-related products and services) and increasing expected future benefits (influenced by changes in perceptions about the future path of diffusion) will engender significant adoption among other firms not currently contemplating adoption. In other words, in an environment with competing standards, an open question will be whether diffusion of the new technology will stall, as it did for early generations of the Fax (Rohlfs 2001) and for 56K modems

(Augereau, Greenstein, and Rysman 2006). This is important, since much of the value of adopting IPv6 will be obtained only once diffusion is widespread.

Historical examples that inform pattern of diffusion

Historical examples can provide some insights into the factors influencing the diffusion of new technologies. Below provides two such examples. To be clear, neither matches all or even most of the key features of the diffusion of IPv6 mentioned above, however each is illuminating in its own way. There is a long history of researchers investigating case studies of standards diffusion, for more examples the interested reader can refer to Rohlfs (2001), Farrell and Klemperer (2007), and the articles included in the volume edited by Greenstein and Stango (2007).

One such example was the transition from NSFNET to the commercial Internet in the mid-1990s. While it may be difficult to remember today, the rapid diffusion of the commercial Internet was unexpected. That is, within the context of our framework, the commercial Internet represented another example of a platform. It was not at all obvious to commercial participants around the time of its initial transition from NSFNET that it would transition as rapidly and widely as it did. As recounted in prior work (Greenstein 2010; Cusumano and Yoffie 1998), many large market participants did not anticipate the economic opportunity from the commercialisation of the Internet backbone, and initially made little in the way of anticipatory investments. For example, over the early 1990s IBM's enterprise division did not wish to be involved with equipment using IPv4, choosing instead to focus on proprietary equipment and protocols such as Token Ring and OS/2. Similarly, as Microsoft developed its enterprise e-mail application, later known as Exchange, it decided not to develop Exchange using design principles put forth by the Internet community. As is currently the case with IPv6, many companies during that episode lacked internal champions to promote investments in the new technology, and because the benefits for transitioning to the Internet were diffuse there were few players who had adequate incentives to nurture the transition to the new platform (Greenstein 2010; Cusumano and Yoffie 1998).

One event that helped to shape expectations about the value of the commercial Internet was the diffusion of the Netscape browser.³³ That is, the diffusion of the browser helped market participants to understand how the Internet could be used as a platform to co-ordinate economic activity (Greenstein 2010; Cusumano and Yoffie 1998). Therefore, the diffusion of the browser played a role as a focal point that acted as a coordinating device. Market participants who eventually understood its potential, quickly changed strategy in anticipation of the rapid diffusion of the commercial Internet that was to come.³⁴ Moreover, there was an explosion of entrepreneurial activity by new entrants. Thus, in this case, while there was no coordinated attempt by Netscape to shape market outcomes, the rapid diffusion of the browser—and the realisation by market participants of how it would shape outcomes—played a major role in promoting investments in the commercial Internet backbone and related technologies.

Make no mistake, the institutional features of the initial diffusion of the commercial Internet differed in critical ways from the diffusion of IPv6. The commercial Internet represented a platform that incorporated a range of technical standards, and its diffusion engendered a range of new business models and opportunities. Within the context of the framework of the probit model, the diffusion of the Netscape browser helped firms to understand that the adoption of Internet technology was associated with a range of benefits, some of which could be realised quite quickly (Forman 2005; Forman, Goldfarb, and Greenstein 2012). In contrast, diffusion of IPv6 represents a diffusion of technical standard that solves an infrastructure problem. The immediate benefits of adoption are limited, and expectations of future benefits depend heavily on the adoption decisions of a range of market actors. However, the comparison is useful, both to contrast the reasons for the differential diffusion experience, as well as to highlight the value of a focal event to igniting network effects.

Another useful point of comparison is the early diffusion of client/server (C/S) computing technology. As with the commercial Internet, C/S is not an example of a single protocol or standard, but rather a range of standards and technologies that are employed within a computing platform. The diffusion of C/S with enterprise computing environments proceeded much more slowly than was initially anticipated. Bresnahan and Greenstein (1996) describe this experience, emphasising the importance of co-invention, the process through which users make technological progress more valuable through their own invention and discovery. They describe how this process takes time, and often relies on third party markets for products and services that can be undeveloped when a new enabling innovation initially diffuses.

The diffusion of C/S provides some useful lessons about the diffusion of a new enabling technology within enterprise computing environments. First, it highlights the challenges of integrating new enabling technologies in legacy computing environments, particularly when third-party markets for products and services do not yet exist. Again, one must be very careful when drawing comparisons with IPv6. First, at least initially, applications of C/S technology did not rely heavily on inter-firm communications, and so direct network effects did not play the same role that they play in the diffusion of IPv6. Thus, the nature of the coordination problem was different.³⁵ The early focus on C/S computing was on large scale computing centres in corporate computing environments, which represents a different environment than that which is prevalent for many users examining the decision to adopt dual stack IPv6.

Moreover, co-invention emphasises identifying opportunities for creating new economic value. The early applications of IPv6 have been to solve an addressing problem. The focus has been on solving a technical problem. While some adopters have reported the ability to change and improve some of the ways in which they address their networks, this is not an example of developing valuable new economic products and services. However, the current prevalence of NATs in the Internet today has made it difficult to develop certain types of new products and services that require end-to-end connections; examples of this include peer-to-peer applications like Web RTC (Web Real-Time Communication). Further, there are some examples of how firms are experimenting with using IPv6 to develop new products and services that would not be possible in an environment where the number of IP addresses are constrained. Many of these examples are in area of the Internet of Things; some specific examples are described in Annex 1. In short, as is often case with the diffusion of a new enabling information technology, there will likely be some period of experimentation required to understand how and whether the diffusion of IPv6 will influence the development of economically valuable new products and services. As is typical, this process takes time.

2.1 Some economics of not moving to IPv6

This section discusses some of the implications of not moving to IPv6. Some of this discussion will overlap with that included below on the costs and benefits to adopting IPv6 for different players along the value chain. Here three issues are discussed that may arise and will affect multiple value chain participants if the transition to IPv6 is slow: the implications of increasing use of carrier grade NATs (CGNs), implications for the markets for IP addresses, and issues to consider in any welfare analysis of the implications for not transitioning to IPv6.

The economic implications of not transitioning to IPv6 are not distributed equally across participants in the platform. The costs of not transitioning to IPv6 will not be equally shared; indeed, some participants may actually receive some benefits in the short run. This variance in transition costs is one factor that has made the transition more difficult.

Implications of CGNs

One very obvious implication of the slow transition to IPv6 is that some way is needed to maintain communications over the Internet even in the absence of a sufficient number of IP addresses for all users.

The primary approach to doing this has been network address translation (NAT). NAT is used extensively in the IPv4 Internet to share devices in a local network. This form of NAT allows multiple devices to share a common IPv4 address, and is known as NAT44. NAT44 creates a variety of technical problems that have been documented extensively elsewhere.³⁶

CGN compounds the problems discussed in the prior paragraph. These problems have also been discussed extensively elsewhere.³⁷ Donely et al. (2013) discuss the results of a test of the effects of NAT444 and Dual-Stack Lite on the performance of several applications when CGNs are in use. The testing revealed that while some basic services such as email and web browsing worked as expected, a variety of applications, including video streaming, video gaming, and peer-to-peer file sharing were affected.

Two points about these tests are significant. First, the report notes that “performance often differs from vendor to vendor and from test environment to test environment, and the results are somewhat difficult to predict.” For example, the report records that in an initial (2010) round of testing some important nuances of peer-to-peer gaming between two Xbox 360 users: in certain testing environments peer-to-peer gaming would work, while in others it would not. In other words, by disrupting the performance of applications and making them dependent upon the particular context of lower level architectural decisions, the deployment of CGN disrupted the modular architecture of the Internet.

While it is difficult to quantify exactly what effect this will have, the benefits of modularity have been detailed extensively elsewhere. For example, Baldwin and Clark (2000) argue that by minimising dependencies across components within a system, modularity facilitates parallel design experiments. In an environment with modularity, module designers can try out a range of approaches in product design so long as they adhere to the design rules by which modules fit together. This is particularly valuable in environments undergoing rapid change, as is the case in IT markets. Multiple design experiments may allow an industry to arrive at valuable new products quickly.³⁸

More broadly, the open and modular structure of the Internet facilitates economic experiments (Greenstein 2012). Economic experiments refer to market experiments that alter knowledge about the market value of goods and services (Rosenberg 1992; Stern 2005). These are different than technical and design experiments that take place in a laboratory or design studio; they involve experiments with business models, operations, and organisational procedures that give rise to lessons about what combinations of these lead to economic value. Firms that engage in economic experiments reduce uncertainties about market value, both for themselves and other existing and potential firms in an industry (e.g., Greenstein 2007). The first twenty years of the commercial Internet have seen a rapid rate of economic experiments that have given rise to the extraordinary dynamism and economic growth surrounding the Internet platform.

The creation of new products and services through entrepreneurship requires the entrepreneur to overcome transaction costs. These can include the costs (monetary and non-monetary) of setting up procedures for delivering a new service, the direct costs of employing technical personnel to obtain the required technical knowledge, and the opportunity costs of delays and hassles for which easy secondary markets do not exist (Greenstein 2012). The structure of the Internet has lowered these costs by making the technical requirements of interconnection freely available to anyone. Firms can enter without worrying about hassles of interconnecting with the platform, for example whether the requirements of interconnection might change over time or whether there are costs of interconnection imposed by some platform leader.³⁹ Low transaction costs may facilitate entry and new innovation from small, entrepreneurial firms, for whom such costs might be a barrier to entry. By introducing uncertainty about whether applications will work in particular settings, disruption of the modular structure of the Internet could raise transaction costs. As noted above, applications may work in certain environments and not

others, and it is not clear *ex ante* when problems may arise. In an environment without modularity, potential entrants need to test whether their applications would work under a variety of circumstances. This may be feasible for large firms producing new products and services but could be more difficult for small, entrepreneurial firms. Thus, a disruption to the modular structure of the Internet could have implications for entrepreneurial innovation.

A second significant point in the Donely et al. (2013) report is the changes in the outcomes of tests over time, due to vendor efforts. For example, subsequent to the first test chronicled in the report, Microsoft released an update to Xbox in December 2011. Preliminary testing described in the report noted that it was now possible to play Xbox head-to-head behind a CGN, at least for some games. In other words, in an environment in which converters are unable to achieve complete compatibility, vendors at a different vertical layer were able to obtain compatibility through their own efforts.

These results highlight two features of the diffusion of IPv6. First, while it is impossible to say given the data that are available, as noted above it is possible that many vendors will be unable to undertake these types of efforts; in other words, smaller vendors may have fewer resources than larger ones to obtain compatibility independently.

Second, these incremental improvements in compatibility made by vendors could slow the transition to IPv6. Improvements in the old standard will reduce the incremental net benefits of adopting IPv6. In other words, by making the existing standard more valuable, they will reduce the net benefits of adopting, and so slow the diffusion of the new standard. This pattern has been seen in many prior examples in the transition from an old to a new technology. For example, Mokyr (1990, p. 90, 129) documents how improvements to waterpower slowed the adoption of the steam engine and how technological improvements to sailing slowed the transition to steam ships by several decades. Related, in resource-constrained enterprises, incremental investments in making products and services work under the old platform may delay investments in making products work under the new.

An additional issue with the increasing deployment of CGN is that, by disrupting the end-to-end IP Internet, their use makes it difficult to do “traceback” to identify the ultimate source of messages sent over the Internet. The basic issue is that because you can no longer identify a permanent public IP address with a device, it makes the task of identifying the sender much more difficult and sometimes impossible. Another concern that has sometimes been raised is that by allowing greater granularity in the traceability of individuals, the diffusion of IPv6 raises privacy concerns. This is a complicated issue: It is important to note that some IPv6 implementations use “privacy addresses” that inhibit traceback to the device. However, in general, IPv6 still makes address use less opaque than many current implementations of CGN in today’s Internet (Huston 2013a). Further, as IPv6 becomes more widely deployed, it may be used in ways that affect privacy in ways that are as yet uncertain. The tradeoffs behind the appropriate levels of privacy on the Internet are nuanced, and it is beyond the scope of this report to discuss them in detail. Here, it is simply noted that the deployment of IPv6 has implications for traceability, which will have varying implications for different segments of society.

Markets for IPv4 addresses

A second approach for dealing with the address shortage problem is address reallocation.⁴⁰ While there is a fixed supply of IP addresses, at present some addresses remain unused (i.e., are not currently being routed on the Internet). Part of this arises from the institutional details of the manner in which IP addresses have been allocated. Early in the history of the Internet, there were no formal policies for allocating addresses and IP address conservation was not a major concern. The Internet Assigned Numbers Authority (IANA) was the central body who distributed IP addresses. Requesters were not required to describe why addresses were needed or how they would be used, and large address blocks were given out

to large organisations like universities, major corporations, and some government agencies. These “legacy” address blocks obtained prior to the establishment of the RIRs are significant in size (Mueller, Kuerbis, and Asghari 2012 estimate them at 40% of IPv4 numbers). However, it remains unclear whether holders of these legacy address blocks have property rights to sell them to others.

New conservation policies for addresses were applied after the establishment of the Regional Internet Registries (RIRs) (Hubbard, Kusters, Conrad, Karrenberg, and Postel 1996). Requesters of addresses were required to demonstrate need, in particular providing the RIRs with information about how the numbers would be used (Mueller, Kuerbis, and Asghari 2012). While holders of unused addresses are supposed to return them to the RIRs, this has been done relatively infrequently. Institutions holding blocks of addresses have little incentive to return them, and the RIRs do not hold the institutional authority to compel the return of addresses (Mueller 2010; Mueller, Kuerbis, and Asghari 2012). In sum, due to the existence of large legacy address blocks and an institutional environment that has made reclamation of addresses difficult there exist significant blocks of unused addresses (Mueller 2008; Perset 2007). Over the last several years the RIRs have been engaged in a variety of activities to extend the lifetime of the IPv4 address pool, but as noted above many RIRs are now fully depleted of IPv4 addresses (for one example of how the RIRs have done this, see Nobile 2012).

Besides unused addresses, there is heterogeneity in the value of IPv4 addresses across users. As noted above, CGN does not affect all services equally. There are some services that work well behind CGN, while others (such as peer-to-peer applications) that work less well. In general, attempts to serve data behind CGN create complications. This suggests that there may be some opportunity to reallocate addresses from low value to high value uses and users.

There are costs to address reallocation. Paid transfers of addresses run the risk of breaking up large blocks of addresses into smaller blocks that are resold. This will increase the size of routing tables and so impose externalities on the rest of the Internet community (Edelman 2009; Edelman and Schwartz 2013).⁴¹ Partly as a result of these issues, the RIRs dictate that the smallest block that can be traded is 256 unique contiguous numbers, known as a /24. (Mueller, Kuerbis, and Asghari 2012).⁴²

Recent research has demonstrated a growing market for IP addresses.⁴³ Mueller, Kuerbis, and Asghari (2012) estimate that the number of IP addresses traded has grown from 11,264 in 2009 to 4,999,936 in the first half of 2012. Their paper provides other relevant statistics. In terms of number of addresses traded, North America dominates the market: 84% of all addresses traded are in the ARIN region. Further, over this period there were 22,471,424 addresses allocated by ARIN, so the number of addresses transferred over the market is 26.3% of that allocated by ARIN. This is significant because over this period, ARIN still had addresses available to allocate, and these can be acquired directly from ARIN at a much lower cost than can be obtained on an open market.

Mueller, Kuerbis, and Asghari (2012) provide suggestive evidence that some companies may be acquiring addresses through address transfers in anticipation of future use. As of the time of writing of their article, ARIN’s policy for allocating addresses from the free pool required that requesters demonstrate need for addresses within three months. In contrast, under ARIN’s 8.3 transfer policy, requesters must demonstrate need within two years.⁴⁴ Thus, it may be easier for firms to acquire addresses in anticipation of future events in the transfer market than by acquiring addresses through ARIN’s free pool. Howard (2013b) argues that many large acquirers of address blocks such as Amazon and Microsoft may have had limited IPv6 support at the time of address acquisition.

In sum, there is some evidence that the market for IPv4 addresses is growing, and some suggestive evidence that some firms may be acquiring IPv4 addresses for future use. One interesting question is; given the depletion of IP addresses, why has this market not developed further? One reason may be that until

fairly recently, pools of addresses have been available from the RIRs at lower cost. More broadly, there has been some speculation that as of the time of writing, there has been relatively little buyer interest in the address transfer market. However, if the circumstances now are such that there are buyers with higher valuations for addresses than the valuations of sellers and if the allocations one can obtain from the RIRs are limited (or, in many cases, non-existent), there remains an interesting question about why prices have not adjusted to levels to where a more active market might develop.

One potential contributing factor is the lack of transparency about the price of transactions occurring in this market. The RIRs do not publish the price of transactions on their web sites, and market participants are reluctant to reveal transacted prices. Historically, one source of highly idiosyncratic data has been from bankruptcy proceedings; in one highly publicised 2011 transaction, Microsoft paid \$11.25 per address (Edelman and Schwartz 2013). Private firms who broker transactions also provide some data. Statistics reported by one IPv4 broker (IPv4 Market Group) indicate that prices for /16 address blocks in the RIPE region were slightly over USD 10 per IP address early in 2013; prices for /16 blocks in the APNIC region were also around USD 10 per IP address in the first part of 2013.⁴⁵ However, even within markets there can be substantial variance in prices for addresses, even for blocks of the same size and from the same region, which is surprising since many blocks should approach commodity status (although for some discussion of some situations where all IP addresses may not be the same, see Huston 2013b and Huston 2013c). For example, during July 2014 one broker reported prices of US \$6.74 and US \$ 24.32 for /24 blocks of registered ARIN addresses sold at auction. Broadly speaking, as of mid-2014 the market is in an environment where market prices are unclear and where buyers and sellers may have differing beliefs about the market-clearing price today and may further have different beliefs about how prices may evolve in the future. As a result, some trades may not take place.

Edelman and Schwartz (2013) formalise this intuition. They note that the IPv4 address market incorporates aspects of private and common value auctions.⁴⁶ Bidders may have different private values of the addresses based on their individual needs. However, addresses also include a common value component because they can be resold on the IPv4 address market at a later time and buyers and sellers may have widely differing forecasts about the future demand for IPv4 addresses. In such an environment with differing valuations across market participants, there may be potential trades that do not take place even when there exists buyers with higher valuations for addresses than sellers (Maskin 1992).

A last reason may be institutional opposition to paid transfers. Paid transfers of address blocks were initially controversial (Mueller, Kuerbis, and Asghari 2012), and involved long debates about the efficacy of such a policy. It is not the purpose to review this debate here, rather the interested reader is referred to recent papers on the topic (e.g., Edelman 2008; Lehr, Vest, and Lear 2008). It is possible that this institutional opposition to paid transfers may have slowed the market's early development.

In short, while at present their activity is not high, markets for IPv4 addresses can play a role in allocating addresses from low value uses to higher value uses. This heterogeneity in circumstance highlights the co-ordination problem discussed above; in the short- to medium- run, some firms may simply have little incentive to transition to IPv6. Further, the need to acquire addresses will increase transaction costs for new entrants delivering products and services over the Internet, increasing the costs of economic experiments. In addition, a case can be made that such a shift in addresses from high value uses from low value uses may be counter to some public policy values such as universal access. In general when resources become scarce and costs get assigned to them the ability for permission-less innovation and the trial and error innovation that has characterised the Internet becomes harder and more expensive.

This could have implications for the Internet of Things (IoT) by making it more difficult and costly to roll-out these devices. It is widely expected that in the next decade billions of devices will go online. A recent OECD report estimated that a family with two teenagers could have as many as 50 devices

connected in their homes by 2022 (OECD 2013). Such numbers are plausible when industry expectations are that all LED connected light bulbs will soon be connected, in a similar vein to the Philips' Hue light system. This could result in 14 billion connected devices in households in the OECD area alone. With the rise of incomes around the world and the decrease in costs of devices, the expectation that even more devices go online is not unwarranted. In addition, businesses expect an increase in the number of connected devices from the integration into manufacturing activities to the delivery of goods and services. General Electric calls this the "Industrial Internet". The expectation here is that every system will be designed around the data needs of the user. In order to capture this data every system will be equipped with sensors, actuators and communication. As a result some estimates are that 50 billion devices will be connected by 2020-2030.

If billions of devices are to go online, the current pool of 4 billion IPv4 addresses will not be sufficient to facilitate the innovation and deployment required for economic growth. Though it may be possible for some market players to acquire IPv4 addresses for IoT applications, the scarcity of IPv4 addresses is such that it will be nigh impossible to obtain even close to sufficient address space.

There are several primary benefits in using IPv6 for IoT applications and they all demonstrate the challenges of staying with IPv4. For its part, IPv6 extends the address space as well as enabling end-to-end reachability and auto-configuration. Without a transition to IPv6 one of the main challenges faced in the use of some IoT applications is the unreliability of the underlying networks in providing critical services. This is because an open, unrestricted IP connection is needed to better enable any issue to be quickly addressed and resolved. The challenge is that in the transition to an IPv6 based Internet many players will need to use NATs to ensure connectivity to IPv4 to meet their own requirements or those of their customers.

Implications for Competition and Consumer Welfare

If the quality of Internet services received by users is degraded, this would have a direct impact on consumer welfare. Like many of the other implications of not transitioning to IPv6 that have been discussed in this section, direct measurement of these implications is difficult. The value of the Internet is hard to measure because so many of the goods and services are delivered without a direct monetary cost to consumers. Several attempts at measuring the value of the Internet to consumers have been made, however this is an extremely complicated issue and so this report only lists the different approaches with citations for those who wish to explore them further.⁴⁷ One approach has been to measure its value using the time spent outside of work at the wage rate of the consumer. The assumption here is that the marginal minute comes from work (e.g., Goolsbee and Klenow 2006; Brynjolfsson and Oh 2012). Another way is to measure the value of Internet connections; Greenstein and McDevitt (2011) study the incremental value of migrating from dial-up to broadband. A last method is to measure consumer willingness to pay for broadband using a discrete choice survey approach (Rosston, Savage, and Waldman 2010).

As noted above in the section on implications of CGNs, failure to transition to IPv6 may have important implications for some entrants seeking to deploy services over the Internet. As has been noted elsewhere, the costs of address run-out will not be born equally. Some services cannot be deployed behind NAT and so IPv4 addresses will need to be acquired for these services to be deployed. These services can include both those deployed by firms and those deployed by nongovernmental organisations that benefit society in other ways. Ultimately, use of NAT will increase entry costs for developing such services, and will lead to a wealth transfer from those who have legacy IP addresses to those who need them to deliver new products.

3. OVERVIEW OF BENEFITS AND COSTS FOR DIFFERENT PLAYERS WITHIN THE PLATFORM

The previous section described some of the high level factors influencing individual-level adoption decisions of IPv6. This section details some specific factors influencing the decisions of different economic actors in the platform. A common theme in this section is the difficulty of obtaining hard data on benefits and costs of deployment for each player. Rather than providing such estimates, here the document seeks to describe some institutional factors that influence the costs and benefits to deployment—and so, through the probit model, inform the understanding of adoption decisions—for different players in the platform.

3.1 Infrastructure vendors

3.1.1 Network gear

Support for IPv6 among network equipment such as routers for the enterprise and Internet backbone is currently excellent, with major vendors promoting their offerings.⁴⁸

The situation with customer-premises equipment (CPE) for wireline service is somewhat more complicated. In some countries like the United States, depending upon their Internet Service Provider, wireline service consumers are either required or have the option to purchase their own CPE. When this happens, lack of co-ordination between the customer and the service provider can create disruption in IPv6 service. One issue is that some consumers have purchased CPE that either does not support IPv6 or which is incompatible with the implementation used by the service provider. This creates a situation where the consumer may not receive IPv6 service even when the service provider is delivering IPv6 connectivity.⁴⁹ This issue is somewhat different than that faced by mobile providers, which is described in further detail below.

Note that this issue has the potential to influence IPv6 diffusion indirectly through its effects on ISP incentives to adopt IPv6. CPE that is incompatible with an ISP's deployment of IPv6 will reduce the potential benefits to adoption, since IPv6 service will not be delivered to the customer.⁵⁰ Further, it may increase provider costs, for example if the provider is compelled to eventually purchase new CPE for the customer.

Unfortunately, systematic data on support for IPv6 in CPE, and other devices more broadly, is difficult to obtain. Testing labs such as the University of New Hampshire's Interoperability Laboratory provide a valuable service by publishing the results of their tests of interoperability and conformance to the community, and similar tests are conducted by the National Institute of Standards and Technology through the USGv6 program. Unfortunately, such lists provide less information on products that are not IPv6 certified, and so it is difficult to get a sense, for example of, the extent of IPv6 support across different products such as CPE, routers, and hosts. Such data would be valuable to ascertain the extent to which vendor support is facilitating the migration to IPv6.⁵¹

3.1.2 General purpose computing devices

Windows operating systems have fully supported IPv6 since Windows Vista and Windows Server 2008.⁵² However, some configuration is required, which may lead to some problems in IPv6 deployment for inexperienced users.

Web browsers like Chrome, Firefox, Internet Explorer, and Safari have also supported IPv6 for some time. However, incremental improvements to browsers to improve the user experience during the transition may again be inadvertently slowing the transition from IPv4 to IPv6 (Aben 2011).

The issue here initially arose from broken IPv6 connectivity. When a dual stack user attempts to connect to a web site over IPv6 first (as in RFC 3484)⁵³, if the connection is broken over IPv6 it will have to time out before an IPv4 connection will be attempted. The timing of the wait can take tens of seconds, resulting in an unacceptable user experience (Aben 2011; Huston 2012).

The “Happy Eyeballs” algorithm addresses the above situation by taking measures so that the user does not experience significant performance degradation when one protocol is broken. One quote on the goals of the original “Happy Eyeballs” implementation is as follows: “A Happy Eyeballs algorithm has two primary goals: Provides fast connection for users, by quickly attempting to connect using IPv6 and (if that connection attempt is not quickly successful) to connect using IPv4. Avoids thrashing the network, by not (always) making simultaneous connection attempts on both IPv6 and IPv4.”⁵⁴

However, there have historically been different implementations of “Happy Eyeballs” that exhibit different behaviour across different browsers and computing platforms. In particular, some research has found that an implementation on Safari and Mac OS X and Mountain Lion would frequently fetch a dual-stacked object over IPv4 instead of IPv6 (Aben 2011). It is believed that this behaviour continues to affect the percentage of measured traffic occurring over IPv6 (Akamai 2013).

Recent statistics from Erik Nygren at Akamai bring some of these issues into sharper focus (Nygren 2013). They show the results of 262 billion requests against dual-stacked sites during June 12, 2013 (Table 1).

Table 1. Requests against dual-stacked sites observed by Akamai

Operating System	Browser	IPv6 as % of Requests
Microsoft Windows 8		4.1%
Microsoft Windows Vista		3.3%
Microsoft Windows 7		2.5%
Microsoft Windows XP		0.5%
Mac OS X 10.5 & 10.6	Chrome & Firefox	3.4%
Mac OS X 10.5 & 10.6	Safari	3.3%
Mac OS X 10.7 & 10.8	Chrome & Firefox	3.3%
Mac OS X 10.7 & 10.8	Safari	2.1%

Source: <https://blogs.akamai.com/2013/06/world-ipv6-launch-anniversary-measuring-adoption-one-year-later.html>.

There are several things to note. First, the very low percentage of IPv6 requests over the Windows XP OS, observed by Akamai, reflects the lack of IPv6 support in that operating system. Given the still-large percentage of Windows XP used in corporations (e.g., King and Yadron 2014), this will continue to act as a drag on the diffusion of IPv6 in the near future. Second, there is some evidence that the “Happy Eyeballs” implementation may be influencing IPv6 traffic, as mentioned above: in recent implementations of Mac OS X, the percentage of IPv6 requests over Safari is lower than that over Chrome and Firefox.

This echoes an issue raised elsewhere in this document; that attempts to ease the transition between IPv4 and IPv6 have increased the value of remaining on the older platform and muting user incentives to migrate to IPv6. Specifically, if clients with the capability to connect over IPv6 continue to connect over IPv4, this will send distorted signals about the number of clients that can connect over IPv6 to market participants who are contemplating IPv6 deployment. This will in turn shape their perceptions about the present and future diffusion of IPv6 in different segments of the platform. As noted above, such perceptions will influence the expected benefits of adoption, and in turn adoption decisions.

Connectivity over mobile devices is discussed in the section on mobile networks in section 3.3 below.

3.1.3 Consumer electronic devices

Obtaining systematic data on IPv6 support among consumer devices such as video game consoles and Blu-Ray players is very difficult. Even if they provided an IPv4 networking stack, the quality of that stack was limited, with many “bugs” reported, but apparently questionable action taken, according to some Internet forums devoted to such equipment. The actual use of Internet based services, such as for Blu-Ray content, was therefore very limited. Recently, however, this has been changing. Televisions and associated peripheral devices increasingly come with “connected television” functionality. Manufacturers are moving to more powerful platforms, such as Android, WebTV and implementations of Linux, which make use of more mature network stacks and as a result can more easily activate IPv6.

Whether or not manufacturers actually support IPv6 in their products may be difficult for consumers to assess. One example that displays some of this complexity is Sony. Several of Sony’s televisions support IPv6.⁵⁵ However, the PlayStation 4, which is Sony’s latest gaming console does not support IPv6 and this will not be implemented in the near future.⁵⁶ On the other hand, the competing Xbox One does support IPv6. The reason for this is that Microsoft intends to work from a more unified code bases for all devices it supports.

Like other segments of the platform, the benefits of providing support for IPv6 may be low until a large fraction of users adopt. However, unlike other hardware and device manufacturers in the ecosystem that support broadband service providers or enterprises who may be demanding or are at least having conversations about IPv6, demand for IPv6 support among end users as such is very low (though end users may demand support for applications that require or are very difficult to deliver without a public IP address, a point returned to below). Thus, relative to other features, IPv6 support may be a low priority among product development managers who are making decisions on where to invest R&D.

Through indirect network effects, lack of IPv6 support in consumer electronic devices will reduce the returns to adoption for other sides of the platform. As described in the section above on general purpose computing devices, consumer devices connecting over IPv4 will shape perceptions about the present and future diffusion of IPv6 for service providers and other segments of the platform.

3.2 Wireline ISPs

This section discusses deployment costs and benefits for ISPs, particularly from the perspective of wireline service providers; key features related to mobile deployment are discussed below.

There is significant variation in the penetration of IPv6 across wireline ISPs. Akamai has provided data, for example, on the levels of the percent of IPv6 requests from US network operators received by the Akamai network during 2013 (Nygren 2014) (Table 2).⁵⁷

Table 2. IPv6 Connectivity by Network Provider, Akamai Statistics

Network Provider	IPv6 as a % of Requests, June 2013	IPv6 as a % of Requests, Dec 2013
Google Fiber	47.5%	60.7%
Verizon Wireless	34.9%	43.4%
Hughes Network Systems	24.2%	21.9%
AT&T	8.4%	13.6%
Comcast	3.2%	13.1%
Time Warner Cable	0.3%	3.3%
T-Mobile USA	0.3%	3.2%

Source: Nygren (2014), Available at <http://www.networkcomputing.com/networking/ipv6-sweet-spots-of-adoption/a/d-id/1234689>.

The analysis here focuses on the United States network operators primarily to examine a context where there is less heterogeneity in external environment (as compared, for example, to a cross-country comparison). There are two things to note in this table. First, in general, IPv6 deployment has been growing fairly rapidly. Second, there is wide variation in IPv6 penetration across providers. This is significant, as the very simple model of adoption described above shows that this variance should be driven by differences in the perceived costs and benefits to adoption across network operators. Interviews suggest that these differences may reflect several factors. First, the costs of deploying IPv6 vary widely based upon many things such as the provider's legacy architecture and infrastructure. Second, as discussed in further detail below, the economics of IPv6 deployment for wireless providers are significantly different than for wireline service providers; this is reflected in the high percentage of requests over the Verizon Wireless network. Third, some large service providers may be seeking to influence adoption across the platform by playing the "large, influential user" role detailed in section 2 above. Fourth, as noted above, the benefits of IPv6 are uncertain, and so differences in the timing of adoption may reflect differences in expectations about diffusion patterns and the benefits of adoption conditional on different diffusion patterns. Finally, and related, some differences may reflect the presence (or lack of presence) of someone in the organisation acting as a "champion" for IPv6.

The payoffs to adopting IPv6 compared to various alternatives have probably been investigated more extensively for network operators than for any other group in the platform (e.g., Howard 2013a, 2013c; Chandler 2012, 2013).⁵⁸ Still, hard data are very difficult to obtain. As is common in any enterprise IT deployment, the returns to adopting new technologies are very uncertain. This is particularly the case for IPv6, where the payoffs for adoption depend in a very complicated way on the decisions of other players in the ecosystem. It is easier to find estimates of the factors affecting adoption costs, particularly given the availability of white papers as well as conference presentations by providers who have experience with IPv6 deployment. However, it is important to note that many published estimates are based upon a small number of limited experiences, rather than systematic surveys. Moreover, as noted above, there is a wide range of approaches to deploy IPv6, and the costs and benefits of different deployment strategies will depend in a very significant way on legacy infrastructure investments. Thus, estimating ex ante deployment costs in any particular setting will be hard.

The decision to deploy IPv6 is complicated because network operators have different options to IPv4 address depletion: deploy CGNs, purchase IPv4 addresses on the market; or deploy IPv6. These options are not mutually exclusive. Unless an operator has ample stocks of IPv4 addresses, it will need to either purchase additional IPv4 addresses or deploy CGNs to enable sharing of public IP addresses. This will be true even if the operator is deploying dual stack IPv6, as the operator will require sufficient IPv4 addresses until IPv6 diffuses sufficiently elsewhere and all forms of IPv4 support can be removed from the network.⁵⁹

Deploy CGNs.

Howard (2013a) provides a set of estimates of the payoffs to CGN deployment.⁶⁰ He estimates the capital costs of CGN to be USD 90 000 per 10 000 users, based on hardware costs of CGN, logging systems, and software development. Further, there are additional annual operating expenses of USD 10 000 per year per 10 000 users.

The additional costs to CGN deployment are related to deterioration in service quality. As noted above, certain services, such as VoIP, peer-to-peer gaming, or streaming video may not work behind CGN. Consumers who try to use these services and who experience problems may choose to leave and go to another provider. This loss of customers represents a potential cost of deploying CGN, but only if competitors are not themselves deploying CGN.⁶¹ That is, the costs of CGN will depend upon the stickiness of customers and the service qualities offered by competitors, both of which are difficult to forecast.

Purchase IP Addresses.

Another possibility for network operators who have a shortage of IPv4 addresses is to acquire them through address transfer. This was discussed in some detail above.

Deploy IPv6

In the short run deploying IPv6 is not a direct substitute to either of the other two alternatives mentioned above. Deployments like dual stack IPv6 still require IPv4 addresses. Thus, the short run payoffs to deploying IPv6 will be far lower than the long run payoffs. Further, due to the uncertainty about the timing of when other players adopt, the timing and magnitude of the payoffs are both very uncertain. The small short run benefits and uncertainty about the payoffs to IPv6 adoption in the long run are some of the key factors that characterise its diffusion.

Estimates of the costs of deploying IPv6 range widely, however some individuals interviewed for this report estimated them to be fairly small relative to total infrastructure and service budgets. One key feature of deployment costs is that they will depend significantly upon the length of the deployment. Deployments that stretch over several years will incur few hardware costs, as network operators can substitute hardware with new hardware that supports IPv6 during regular replacement cycles. In such cases, the direct costs may relate mostly to software and testing. Interviewees for this report also emphasised that the long timeframe for deployment also provided time to work with vendors and other complementors to solve problems that arose during deployment. As more ISPs and their complementors gain experience with deploying IPv6, this may become less important. In other words, deployment costs decline over time as the availability of complementary skills to deploy new technology becomes more prevalent. For example, as vendors and content providers gain more experience with deployment, costs may decline significantly over time. This has been a common experience in the diffusion of new networked platforms (e.g., Bresnahan and Greenstein 1996; Forman 2005).

Howard (2013d) estimates most of the hardware costs for an ISP to deploying IPv6 to be primarily related to updating CPE, and to be approximately USD 25 per user (again, this represents one point within a wide range).^{62 63} He finds that the training costs of deployment to be approximately \$0.15 per user. For ongoing operations costs, he further estimates development costs to be USD 6.40 per user per year, with an additional USD 0.25 to USD 1.27 per user per year in operations costs. As noted above, consumers in the United States and some other countries use their own CPE or have the option of substituting the network operator's CPE for the own devices. This can create problems as consumers can substitute IPv6-ready CPE provided by the provider with CPE that did not support the IPv6 implementation on the provider's network.

In Europe some networks have supported IPv6 for quite some time, particularly notable is the second largest broadband provider in France (Iliad-Free). Other operators have sometimes provided CPE to their customers that were not capable of supporting IPv6. An additional approach has been suggested by Deutsche Telekom. In 2013, the company presented a re-engineered network design that it calls Terastream. This network is native IPv6 only. IPv4 is only a service on the network. The Terastream design aims to simplify the operation of a network. IPv6 allows it to simplify operations under many circumstances, by for example encoding the service type in the IP-address. All voice traffic to a customer's fixed line telephone has a particular string in the IP-address, which allows the network to route it separately from other traffic. In the past VPNs and MPLS were used to achieve this benefit. Such a network could use NAT to reach legacy IPv4 devices in the home and to reach legacy IPv4 sites on the Internet.⁶⁴ In addition, such a network aims to offer the customer the possibility to multi-home, use multiple service providers and to have a fully self-configuring network. The architects argue that such a simplified network has become possible with IPv6 and would allow Deutsche Telekom to scale its network at much lower costs than traditional network designs that use IPv4 or dual-stacked IPv4 and IPv6. The network is currently being tested in Zagreb, Croatia.

In sum, there is evidence that costs to wireline network operators of deploying IPv6 may not be large. However, many of the benefits of deployment will arrive in the future and are very uncertain. Further, while there are benefits of adopting early so that deployment occurs over a period of time, the costs of adoption are declining as expertise related to diffusion of IPv6 deployment continues to become more widespread. Declining adoption costs increase the incentives to wait, particularly among smaller providers who may have less internal capabilities to deploy new technologies. This would be consistent with the traditional pattern of the probit model of diffusion described above and may be responsible for the "wait and see" attitude toward adoption taken by many smaller network operators.

3.3 Issues Specific to Mobile Providers

The economics of IPv6 deployment for providers of mobile voice and data service differ from that of wireline service in several significant ways. For one, the rate of growth in the demand for IP addresses is greater in mobile than in traditional broadband service. One reason is the rapid growth in mobile subscribers compared to the slower rate of growth in wireline service. Further, mobile phones can use multiple IP addresses at a single time. For example, some providers are deploying voice services over data networks using Voice over LTE (VoLTE), which consumes additional IP addresses. As a result, both the growth in the number of mobile devices (each of which will require a public or private IP address) and the number of IP addresses per device create increasing demand for IP addresses.

The cost structure for the investment decision for IPv6 deployment also differs from mobile providers. The architecture and infrastructure in mobile networks change more rapidly than for wireline networks. One reason for this is simply the rapid growth in mobile subscribers, which necessitate new infrastructure investments.

Another is that carriers will, in the course of transitioning from 3G to 4G, make changes to their network architectures. Some carriers have deployed IPv6 concurrently with these transitions. That is, the economics of deploying IPv6 as part of a greenfield new infrastructure deployment may be different from that of deploying it in a legacy infrastructure where investments in IPv6 represent an incremental cost that cannot be bundled with a larger investment initiative. Still, even when mobile carriers implement IPv6 as part of a migration to 4G service there remain incremental costs relative to that in a 4G service under IPv4. Some of these challenges are similar to those faced by wireline carriers. For example, routes used under IPv6 may be suboptimal. Further, some transit carriers may be unable to support increasing traffic over IPv6. Some content providers may advertise an IPv6 address but then the connection fails when a user attempts to connect.

Another issue for mobile providers transitioning to IPv6 service may be short run challenges with obtaining handsets providing IPv6 service. As noted above, support of IPv6 among recent generations of mobile devices is also robust. Some statistics have shown that 40.6% of LTE devices support IPv6 (Status of the LTE Ecosystem Report, www.gsacom.com), and that 18% of all mobile devices were potentially IPv6-compatible in 2013 (Cisco 2014). However, even among handsets that report IPv6 compatibility, there may be issues in getting them to work with a provider's network. This was particularly the case in earlier generations of the Android Operating System and iOS, which included incomplete support for IPv6. Android also created unusual challenges, as carriers would have to ensure that IPv6 would be supported across a range of hardware platforms and handset manufacturers.

Differences in how customers connect to the network will also influence the ease with which mobile networks are able to transition to IPv6. As noted above, in wireline networks customer premises equipment (CPE) is often updated infrequently, and in some cases customers may themselves purchase CPE hardware that is incompatible with the service provider's IPv6 deployment. Thus, a service provider could deploy IPv6 to the customer premises and the customer could still not achieve an IPv6 connection to his/her personal device. In mobile networks the circumstances are quite different. Customers purchase new phones every 2-3 years and, as described above, many new phones support IPv6. Some service providers compel device manufacturers to support IPv6 in order to connect to their networks. Thus, on average, the IPv6 connection rate on dual stack networks will be much higher in mobile than on wireline networks. As noted above, this may in turn increase incentives to deploy IPv6 in these networks, for reasons described above.


The economics of mobile deployment are different in another way. Several prominent mobile service providers such as T-Mobile in the United States, Orange in Poland, and Telstra in Australia have been able to deploy single stack IPv6 using the 4G4XLAT architecture.⁶⁵ Recall above the unique economics of dual stack IPv6, the most common type of IPv6 deployment to date: adopters achieve the full benefits to adoption only when they can turn off IPv4, which in turn can happen only when the vast majority of traffic on the Internet is over IPv6. Firms deploying single stack IPv6 do not have this particular problem, which can increase and reduce the future uncertainty to the benefits of adopting IPv6. As of version 4.3, which was released in 2013, 4G4XLAT, is supported by Android.

As a result of these differences, there is some evidence to suggest that mobile service provider networks may transition to IPv6 more quickly than wireline networks. For example, in July 2014 Verizon Wireless reported IPv6 deployment of 53.6% and T-Mobile reports deployment of 34.12%.⁶⁶ However, as has been the case elsewhere on the IPv6 platform, there are a range of deployment options that have varying implications for behaviour elsewhere on the platform. For example, some providers deploying IPv6 on their internal networks are using network address transition to IPv4 to connect to the public Internet. Thus, in some cases these deployments may not change the payoffs to deployment of other firms on the platform.

3.4 Content providers

According to World IPv6 Launch, 13.5% of the Alexa Top 1000 websites are currently reachable over IPv6.⁶⁷ Among the top 10 Alexa Websites, the websites for Google, Facebook, YouTube, Yahoo, and Wikipedia all support IPv6, while those for Baidu, LinkedIn, Twitter, Microsoft's Live.com and Tencent QQ all do not.⁶⁸ Data are available to illustrate the number of the top 500 Alexa websites that support IPv6 (Figure 3).

Figure 3. Percentage of IPv6 Adopters among Top 500 Websites



Source: Cisco 6Lab Statistics, available at <http://6lab.cisco.com/stats/index.php>.⁶⁹

Thus, while a significant fraction of web sites have been enabled, the majority of even the largest websites have not, and this fraction has changed little over time, even as there has been significant growth in other areas of IPv6 diffusion.

Again, deployment costs vary widely. Further, because of the importance of development and testing any deployment is likely to have a significant fixed cost component (i.e., there will be large costs that do not vary with traffic), so that larger web sites are likely to have lower per user deployment costs than smaller websites.⁷⁰ Howard (2013d) estimates that deployment costs for IPv6 for content providers to be USD 7 per user and operations costs to be USD 6.08 per user per year.⁷¹ However, these figures must again be viewed as one point estimate within a wide range of potential outcomes.

A key point underlined by stakeholders interviewed for this report is the importance of testing. Like the deployments of other IT services, adopters frequently encounter bugs in their deployments, and content providers must deploy IPv6 in production to solve many of these problems. In many cases, thousands of changes must be made.⁷² Solving these types of problems can take a long time and require sustained effort. As a result, some firms who have sought to deploy IPv6 have broken or poor implementations (e.g., Nygren 2011).

While this was true historically more than it is today, even when content deployments have been done properly some clients may experience a variety of connectivity problems to dual-stack sites.⁷³ Thus, the

existence of an IPv6 record can cause inferior response times to customers even when a site has been implemented properly. This can decrease incentives for sites to deploy content over IPv6 due to poor service quality to customers and potentially higher help desk costs, exacerbating the problem of low content availability over IPv6.

There are likely to be a range of experiences around deploying IPv6. Large content providers such as Google and Facebook will not only be able to spread their deployment costs across a larger base of economic activity, but their costs of deployment are also likely to be mitigated by the existence of internal capabilities—i.e., internal technical staffs—to undertake deployment and development efforts.⁷⁴ Some aspects of development, testing, and deployment are likely to be challenging for smaller firms with fewer internal capabilities. While such expertise can often be contracted for in the market, this may be harder to do at present in the case of IPv6 where expertise is not yet as widespread because of the formative stage of diffusion. Therefore, it is important to be careful when generalising the experiences of early adopters from other contexts. This is discussed in further detail below.

In short, while the direct costs of deploying IPv6 do not appear to be large, there are a range of circumstances—some internal to the content provider and some external—wherein deployment of IPv6 can lead to inferior service performance for customers. In the absence of a large pool of IPv6 users (and an even lower pool of users who are IPv6 only, many on IPv6 will be dual stack) it is not surprising that many content providers have chosen to defer deployment. The incentives to delay deployment will be particularly significant for content providers who already retain sufficient addresses to run their business and for smaller content providers who may lack the internal capabilities that will facilitate a smooth deployment.

3.5 Enterprises – Internal

Many enterprise networks already deploy private IPv4 numbering schemes and NATs extensively. Enterprise NATs enable devices within corporate networks to share public IP addresses.⁷⁵ The widespread deployment of private numbering and NATs within enterprise networks reduces one of the key benefits of IPv6—the need for additional addresses. Further, the costs of adopting IPv6 within enterprises may be sizeable. Prior research on the diffusion of other networked technologies has demonstrated there to be sizeable costs of integrating new technologies with legacy systems that may reduce the net benefits to adoption and slow diffusion (e.g., Bresnahan and Greenstein 1996; Forman 2005). In the case of IPv6, this can include the installation of new hardware and software as part of the upgrade, upgrading a range of legacy applications, and significant training costs (IPv6 Task Force 2006).

As noted above, new security issues may be introduced in running dual stack. Running a dual stack network increases the number of possible ways in which a network can be attacked. This occurs because introducing IPv6 adds a new network layer protocol with new transition technologies that are usually activated by default. Further, the interaction between the two protocols allows IPv4 to be attacked from IPv6 and vice-versa (InterConnect Communications 2012).

The introduction of the Internet of Things, however, in enterprises, for manufacturing, building control, track and trace and other applications is only starting and will have profound implications for such considerations. It will, when combined with cloud applications, lead to more autonomous robotic systems. It is here where the use of IPv6 is likely to be greatest. Private addressing is difficult to maintain when the use of numbers reaches tens of thousands if not millions, because companies would likely be forced to create separate networks for separate sensors and applications, where one private address might be used by multiple devices, each in a different network. This, in turn, would mean that servers interacting with those devices have to be on separate networks, creating a fragmented internal network all for the sake of conserving address space.

One critical development of the production environment of the future is the need for flexibility. Currently sensors and actuators are located on fixed locations and cannot be changed. With production being more flexible and the need for data becoming more prominent, wireless sensors and actuators become a necessity in business. Standards such as ISA100.11a support IPv6 using 6LoWPAN. This allows their users faster deployment in industrial settings, much in the same way similar standards such as Zigbee do in consumer settings. The MGM Aria hotel in Las Vegas, which has 110 thousand Zigbee sensors, exemplifies that these settings can be very large.

In sum, as a result of prior network architectural decisions, enterprises may have muted needs for new addresses. Further, many enterprises may have significant costs of integrating use of IPv6 with legacy infrastructure and applications. This may lead to muted incentives to adopt IPv6 within the enterprise though this could change with more IoT deployment

3.6 Consumer End Users

Relative to other groups in the value chain, consumer end users are less informed about the benefits and costs of deploying IPv6. In short, the average consumer likely cares little whether network services are delivered via IPv4 or IPv6, and so are unlikely to demand IPv6 service from content or service providers.

Consumer behaviour does have the potential to shape IPv6 investment indirectly, however. If, for example, consumers demand network services that perform poorly behind CGN, then consumer demand for these services could increase the benefits of IPv6 adoption among service providers as a way of reducing CGN deployments.

While there is little systematic empirical evidence to test this assertion, there is also little evidence at present that this type of consumer pressure is shaping incentives to invest in IPv6. There are likely many reasons for this, but two are mentioned here for illustrative purposes. There appears to be significant investments to increase the value of the legacy IPv4 platform, including attempts to mitigate the costs of CGN. This is a rational response by providers of applications whose service quality might be affected by increasing deployments of CGN.

Second, dual stack deployments continue to require IPv4 addresses and do not alleviate the need for CGN. Indeed, some ISPs deploying IPv6 also find the need to deploy CGN as a “bridge” to IPv6. Thus, as long as there remains significant IPv4 traffic, ISPs facing address depletion are likely to continue to see a need for CGN.

3.7 Comments on visibility of costs and benefits

As has been noted throughout this report, decisions to invest in IPv6 involve a trade-off of benefits and costs. However, this section has repeatedly mentioned how difficult it is to obtain systematic cost and benefit data. In many ways the broad contours of this problem are typical of any new technology adoption decision.

Typically, in any enterprise IT investment costs are to some degree front-loaded, and involve some tangible investments in hardware and software that are comparatively easily measured. However, there are other costs as well, including integration with existing technologies that involve the co-invention activities described in section 2. Such costs are typically high during the early stages of the diffusion of new technology, when labour markets for necessary expertise are thin and many third party providers of complementary services may also be inexperienced with the new technology. There is evidence to suggest that there are significant costs to IPv6 deployment beyond hardware and software costs, but the size of this activity varies across organisations that may be uncertain ex-ante, making it more difficult to estimate the total size of deployment costs. For example, enterprises and content providers deploying IPv6 have

reported challenges with changeovers to their addressing schemes, and that the conversion has involved changes to existing databases and applications. Existing value chain partners like ISPs or CDNs may not yet deploy IPv6. Although many products from vendors may support IPv6, they may not support some features. All of this adds uncertainty to the investment decision. Until adoption rates edge higher and markets for skills thicken, this will continue to be a factor that weighs on adoption decisions.


As is typical in any enterprise IT investment decision, benefits of IPv6 adoption occur in the future and are more uncertain. However, as is also common in a technology characterised by network effects, these benefits will depend in a complicated way on the decisions of others. As noted in the introduction, adoption of IPv6 is very low. However, as pointed out in OECD (2010), there are a wide variety of metrics of IPv6 deployment, covering things such as the number of IPv6 prefixes allocated and assigned by the RIRs, IPv6 in the global routing tables, IPv6 support among ISPs, IPv6 among content providers, and end-user connectivity. Web sites such as <http://6lab.cisco.com/stats/> and <http://www.worldipv6launch.org/measurements/> similarly provide an array of measurements. Each of these measurements plays a valuable role in describing the penetration of IPv6 in different parts of the platform, however the need for such a broad array of measurements may make it more complicated for market participants to assess the likely future path of diffusion and, as a consequence, arrive at firm projections on the likely future benefits of adoption.

4. ALTERNATIVE TRANSITION SCENARIOS AND ATTEMPTS TO FACILITATE TRANSITION TO IPV6

4.1 Alternative transition scenarios and S-Curves

One important question is how IPv6 will diffuse going forward. This is very difficult to answer because currently IPv6 is at an early stage of the diffusion process. Hall and Khan (2003) show the time paths of the diffusion of several key innovations in the 20th Century (Figure 4). The diffusion of these innovations span a range of technologies and time periods and only some represent platforms as defined above (e.g., electric service, the telephone, the PC in household), leading to a wide range of outcomes. While some historical innovations have taken a very long time to diffuse, it is also clear that some more recent innovations—not included in this graph—such as the commercial Internet and the mobile Internet have diffused much more quickly.

Figure 4. Historical Diffusion of New Technologies



Source: Hall and Khan (2002)

Many of the innovations in the graph above generate an S-shaped curve, common in the diffusion of innovations (e.g., Rogers 2003). A variety of assumptions can generate an S-shaped curve (Geroski 2000; Hall and Khan 2003). An S-curve can be generated using the probit model of diffusion described above if (1) the distribution of valuations of the innovation is approximately normal, (2) the costs of adopting decline over time, and (3) users adopt when benefits exceed costs.

These models have historically been used to forecast the future diffusion trajectories of new innovations. This has been done to forecast the future diffusion of IPv6 (e.g., Elmore, Stephens, and Camp 2008a, b; Townsley 2014; Colitti and Kline 2013).


The estimation of S-curves can be a useful tool to think about the future path of diffusion of new technologies. A challenge for the transition to IPv6 is that many of the gains will appear in the future and will further depend upon the adoption patterns of others. Thus, attempts to forecast future adoption are useful.

A careful attempt at forecasting IPv6 diffusion is beyond the scope of this document. The goal here is simply to point out a couple of important issues to think about when forecasting diffusion and highlight areas for future research.

Figure 2 illustrated that, using IPv6 connectivity among Google users, penetration of IPv6 among existing connections currently stands between 3% and 4% worldwide.⁷⁶ Because of the relatively low levels of diffusion at the moment, estimates of the shape of the S-curve will be very sensitive to changes in the path of diffusion going forward. For example, some researchers, when estimating historical S-curves, throw out data at the tails of the distribution.⁷⁷

Related, many innovations fail to diffuse widely, an empirical fact that S-curves do not allow for. To put it another way, some innovations that diffuse to those with high net benefits to adoption—lead users—do not diffuse to the entire population of adopters. These ideas were popularised by Geoffrey Moore and his book “Crossing the Chasm” (Moore 2014). Again, the probit model provides a useful way of fixing ideas. As in Rogers (2003), the distribution of valuations (net benefits) for a new innovation can be plotted assuming a normal distribution and classified into five groups: innovators, early adopters, early majority, late majority, and laggards. Innovators and early adopters are those with high valuations for the innovation, while later adopters have lower valuations. The insight of Moore is that many innovations have had difficulty transitioning from this “early market” to the “mainstream market.”

This section will close with one other issue related to aggregate S-curves. There is currently substantial cross-country variance in experience related to the diffusion of IPv6. Eric Vynke has provided a useful tool to compare cross-country IPv6 adoption using Google statistics. It demonstrates the experience of several OECD countries in different regions of the world: France, Japan, the United States and the worldwide statistics (Figure 6). Singapore is included as well for reasons that will become clearer below. While these statistics are by no means representative, they are a useful starting point to raise several issues.

Figure 5. Historical IPv6 Connectivity by Country Based on Google Connectivity Statistics

Source: <https://www.vyncke.org/ipv6status/compare.php?metric=p&countries=us,ww,fr,sg>, compilation by Eric Vyncke of Google statistics on IPv6 usage (Worldwide Google statistics were presented in Figure 2).

First, traditional lead users of IT such as the United States have seen relatively faster diffusion of IPv6. This is true even though the address depletion problem is less severe in the United States than in other regions of the world. While additional work should study this result more carefully, one interpretation can again be seen through the probit model; that is, while United States firms may have lower benefits of adopting IPv6, they may also have lower costs.

Second, it is clear that there are very different experiences across countries, and some countries are much farther along in diffusion. These differences may be caused by several factors. One is that the decisions of large service providers within a country can have significant effects on the transition experience in a particular country. For example, Japan's KDDI has one of the larger IPv6 deployments in the world. NTT was one of the first firms to bring IPv6 service to its users, and currently its NTT FLETS Hikari brings IPv6 to all of its users. However, the latter is a closed network that is not directly connected to the Internet. As a result, while in 2011 Japan had one of the highest connectivity rates for IPv6, it also had one of the highest rates of impaired IPv6 connections in the world (Colitti and Kline 2011). Through 2012 and 2013 Romania had one of the highest IPv6 connection rates in the world, due in large part to the IPv6 implementation by service provider RCS & RDS.

A second factor that may influence country-specific implementation rates is the allocation of IPv4 addresses across countries. For example, at the time of writing this report, IPv6 diffusion in South Korea is 0.01%. One reason may be that country's relatively large number of IPv4 addresses per user (approximately 2.3), which could decrease the urgency of transition. In short, there may be many country-

specific factors that could influence implementation rates. Future work should explore the source of these differences more carefully.

4.2 Alternative policy responses

The discussion above has chronicled some of the challenges of obtaining a co-ordinated transition to IPv6. The models discussed help to inform discussion about what types of actions and co-ordination mechanisms may facilitate the transition.

A variety of actors, both governmental and nongovernmental, have sought to facilitate the transition to IPv6 through various means. Many of the approaches that have been used so far, which are discussed in further detail below, can be seen as mechanisms to either increase the benefits or decrease the costs of adoption. Though lack of data makes it difficult to assign causality between particular initiatives and outcomes, many of these initiatives appear at least anecdotally to have enjoyed some success.

One challenge is that many of these initiatives are targeted at one particular user group within the platform. As noted above, the benefits and costs of adopting IPv6 may not be spread equally among different players in the platform. For example, some businesses contemplating deploying IPv6 in their e-commerce sites or in their internal firm networks may have little reason to deploy IPv6 at the moment because they have ample addresses. In contrast, other sides of the platform such as backbone and transit providers have relatively lower costs of adoption because of both their sophistication in deploying new network technologies and because they can deploy IPv6 as part of regular hardware replacement cycles. Prior work in platform theory suggests that to promote adoption of a multi-sided platform it can be effective to target the costs of adoption among the side of the platform with some of the highest adoption costs (e.g., Rochet and Tirole 2006; Eisenmann, Parker, and Van Alstyne 2006). However, some of these groups, such as content and e-commerce providers and providers of some complementary consumer product devices, may be difficult to target directly.

The technical community has been engaged in a number of very important initiatives that have been helpful in furthering IPv6 deployment. For example, the Internet Society has sponsored World IPv6 Day and World IPv6 Launch. These initiatives have served several purposes, including providing a mechanism for firms throughout the Internet ecosystem to co-ordinate their IPv6 deployments, to provide an environment for firms to test their deployments, and a forum for sharing and exchanging knowledge about the challenges of deploying IPv6. The technical community also sponsors a very wide range of conferences and meetings throughout the world that enable community members to share their knowledge about IPv6 and the business case for deployment.

As noted, these efforts have been very successful at spreading knowledge about IPv6. However, many of the successful deployments so far have been among lead user ISPs such as Comcast, Verizon Wireless, Iliad-Free, KDDI, and Time Warner Cable and among lead user content providers such as Google and Facebook. This is to be expected. Such firms, by nature of the size and technical sophistication of their workforce, are both the first to adopt IPv6 and naturally also play a role of communicating their experience with others. They also are likely to have a greater need for addresses. However, the costs and benefits of adopting IPv6 by this set of lead users may not be representative of the majority of potential adopters.⁷⁸ In other words, an open question remains on how useful this set of actions will be in promoting adoption among firms who are not lead users.

One issue is whether efforts to encourage IPv6 diffusion can shift the perceived net benefits of adoption among all of the potential players in the platform. For example, one issue that arose multiple times during interviews for this document is key decision makers among enterprise users and consumer electronics manufacturers may not be fully informed about IPv6, and so IPv6 adoption (or incorporating

IPv6 into products) may not be high on their list of priorities. An important question is whether existing efforts to encourage IPv6 diffusion will reach these decision makers. In other words, some of the key mechanisms in use in promoting IPv6 thus far may be unable to reach key bottlenecks to IPv6 deployment within the Internet ecosystem.

Governments and governmental organisations have also sought to take actions to ease the transition to IPv6. For example, as noted above, the United States government has played a role by mandating adoption of IPv6 first within the Department of Defense and later when the Office of Management and Budget set a June 2008 deadline by which federal agency network backbones must support IPv6. While many federal government agencies did not meet these goals, it is believed that the efforts have both stimulated federal government suppliers and vendors at enabling IPv6 enabled products and facilitating the delivery of IPv6 enabled services. These actions also facilitated the development of human capital related to the deployment of IPv6.

Other governments have experimented with a variety of approaches at encouraging IPv6 deployment. While a detailed survey is beyond the scope of this report, a couple of notable cases can be discussed. Sweden, for example, has sought to encourage IPv6 deployment in the governmental sector through a program of informational awareness (Hersaeus 2014). As part of the Digital Agenda for Sweden published in 2011, the government established a goal that all Swedish authorities and municipalities should be reachable over IPv6 in 2013. A wide range of practical information on how to deploy IPv6, including reference cases, fact sheets, and procurement information was provided to government agencies and municipalities (e.g., Hersaeus 2014; PTS 201). PTS further participated in a variety of public sector conferences to communicate this information. In short, as in the case of the United States government, in Sweden PTS sought to encourage public sector deployment and thereby indirectly stimulate private investment. These initiatives can be useful in stimulating human capital development on how to implement IPv6 and more broadly stimulate the expertise among complementary third party service and support firms that can facilitate IPv6 deployment. PTS has compiled statistics on deployment of IPv6 in government agencies and municipalities. While it is too early to discern the effects of these efforts on IPv6 deployment in the government, the results are instructive as they indicate the challenges that adopters with fewer IT capabilities may face when seeking to transition to IPv6; as of March 2014, 33% of Swedish government authorities have support for IPv6 over their websites, while only 13% of local municipalities (who may have fewer technical capabilities) have such support (Hersaeus 2014).

In other countries there have been collaborations between the government and the technical community that have sought to promote IPv6 awareness. For example, in Slovenia, the group go6 wrote a study describing the major issues surrounding IPv6 and its deployment (Kunc et al. 2012). Collaborations between members of go6 and RIPE also lead to an update on procurement guidelines for IPv6, the RIPE document RIPE 554 (Kao, Zorz, and Steffan 2012).

While it may not generalise to all settings, the approach used by the Infocomm Development Authority (IDA) of Singapore is significant in that it addresses multiple sides of the platform simultaneously.⁷⁹ Like other countries, they worked with government agencies on their IPv6 transition and also developed materials to help guide organisations with their IPv6 transition, including an IPv6 marketplace that aggregated the contact details of organisations that could assist with the transition. However, they also worked with some of the largest web sites and banks in Singapore to encourage their transition to IPv6, including co-funding some of their development costs. IDA also worked with local training providers, polytechnics, and universities to help local workers attain the necessary skills to implement IPv6, which lead to the training of 300 industry professionals and 6000 students. Finally, IDA worked with service providers to ensure that users are not stranded on either the IPv4 or IPv6 standard.⁸⁰ Figure 6 shows that according to Google's IPv6 connection statistics that Singapore has seen rapid growth in IPv6 diffusion since the middle of 2013.

5. CONCLUSION

This report has sought to develop some straightforward economic principles to help understand the diffusion of IPv6. After introducing some economic frameworks that have been used to understand the diffusion of other previous platforms, these frameworks were applied to understanding IPv6's diffusion. The report discussed some of the important implications for economic and social prosperity of the slow transition to IPv6. Next, a qualitative discussion of the benefits and costs of adoption of different members of the platform was included. Finally, some issues were discussed regarding the future diffusion of IPv6, and recent policy responses that have been used to accelerate IPv6 diffusion.

As noted in the introduction, this report represents only a preliminary step in understanding the economics influencing IPv6 diffusion. The findings in this report point to areas for future research. One area that has been highlighted throughout this report is the need for additional data related to IPv6 deployment and costs. By providing information on the IPv6 transition, web sites such as <http://6lab.cisco.com/stats/>, <http://www.worldipv6launch.org/measurements/>, <http://labs.apnic.net/ipv6-measurement/>, and <https://labs.ripe.net/statistics> all play a valuable role in informing adoption decisions, but even more can be done. For example, there is little hard data at present on the cost experiences of users deploying IPv6. Further, more data on IPv6 support across different types of hardware would be helpful. More broadly, most data available on IPv6 deployment experiences are based on large ISPs or content providers. That is, at present, there are little available data on the experiences of small and medium sized businesses. The focus on large users is sensible, since that is where the vast majority of activity is currently ongoing, and because such users have a disproportionate impact on the transition decisions among other players on the platform. That being said, some further dissemination of experiences from smaller users will also be valuable, given the unique economics of the benefits of IPv6 deployment.

As noted elsewhere in this report, there have been a range of experiences across countries in their transition to IPv6. Many of these cross-country differences reflect either cross-sectional differences in circumstance or institutional factors such as the existence of technical capabilities, the availability of IPv4 addresses, or other issues related to how the Internet has historically been deployed. However, other differences reflect specific investments in IPv6 that have occurred recently, such as efforts by large ISPs to deploy IPv6 or efforts by governments to encourage IPv6 deployment. It is beyond the scope of this report to investigate these issues in depth. Future work could investigate the impact these efforts have had on IPv6 deployment, as a way of informing the potential efficacy of strategies to encourage the IPv6 transition going forward.

ANNEX 1: IPV6 IN CONSUMER ELECTRONICS DEVICES AND IN THE INTERNET OF THINGS

One further factor that will increase the demand for IP addresses is the increasing use of “Smart Home” applications. In a Smart Home more and more devices will be connected to a home network and through those networks to the Internet. Many of these devices and services require a continuous connection to the Internet and are operated via cloud services. Examples of such devices are the “Nest Thermostat” and the “SmartThings” hub. These devices would be best with a clear and uninterrupted connection to the Internet and auto configuration. Instead, the suppliers currently find that each customer has only one IPv4 address and needs to use a NAT. The use of a NAT, however, can disrupt the proper functioning of these devices. As a result, despite the fact that they could benefit from IPv6, commonly used devices such as the Nest thermostat and the SmartThings hub do not support IPv6 today.

The need for IPv6 is clear when low power networks are evaluated. These devices aim to work for long periods on small cell battery charges. Many of the standard networking solutions for IPv4 and IPv6 do not work in such environments. The development of 6LowPan, however, allowed the use of IPv6 in low power environments. Adapting IPv4 appears to be substantially more difficult because it does not support auto configuration and other essential elements in a low power environment. For this reason, IPv6 was adopted by some of the largest standards consortia in the industry, such as Zigbee and Bluetooth Low Energy. Some examples of low power devices in the home that do make use of IPv6 are the Philips Hue light bulbs, that have proven to be popular and make use of 6LoWPAN and Zigbee.

IPv6 is making inroads into large scale deployments of smart energy networks, for example in the United Kingdom. Here Telefonica and Connode, will provide a network to reach smart energy meters based on IPv6 networking. Telefonica will provide the cellular network and Connode will provide a wireless mesh solution to connect devices where there is limited cellular coverage. One of the main benefits, cited by Connode in the operation of its solution, is that it allows the mesh network to work without so called concentrators, which traditionally would collect and store data, before sending it on to the main network and NAT to connect to the cellular network. Instead, each node can be addressed individually reducing the operational costs of the network.^{81 82} Given the long lifetime of a smart meter, often 30 years or more, energy companies and their suppliers are interested in assessing long-term economic costs and benefits.

Outside a residence or premises with a fixed network connection, the use of IPv6 in consumer and industrial electronics is limited by the capabilities of mobile networks. Discussions with an automobile manufacturer and one of navigation devices, undertaken for this report, indicated that a lack of support by mobile networks was the primary reason that IPv6 was not in use by them. In addition, one remarked that with careful allocation of IPv4 addresses to their active devices, it was possible to significantly reduce the need for active IPv4 addresses, because despite a few million active devices, most vehicles are not moving most of the time.

This demonstrates that management of IPv4 addresses is already being undertaken and the complexity will increase when, in the coming years, millions of new vehicles will come online. Presentations from suppliers to the automotive industry, such as Bosch, show that IPv6 is certainly on their roadmap for deployment in vehicles in the near future. In this case, the presentation cited among the largest benefits of IPv6 to be the increased address space and auto-configuration, as well as the increased support for mobility.⁸³ Given that some vehicles today have up to 70 computing devices, sensors and cameras on board, it is clear that IPv6 could be highly valuable in a vehicle.

BIBLIOGRAPHY

- Aben, Emile. 2011. Hampering Eyeballs – Observations on Two “Happy Eyeballs” Implementations. Accessed at <https://labs.ripe.net/Members/emileaben/hampered-eyeballs> on April 3, 2014.
- Akamai. 2013. The State of the Internet: 3rd Quarter, 2013 Report. White Paper. Accessed at http://www.akamai.com/dl/akamai/akamai-soti-q313.pdf?WT.mc_id=soti_Q313 on April 3, 2014.
- Aoun, C. and E. Davies. 2007. Reasons to Move the Network Address Translator-Protocol Translator (NAT-PT) to Historic Status. IETF RFC 4966, Available at <http://www.ietf.org/rfc/rfc4966.txt>.
- Augereau, Angelique, Shane Greenstein, and Marc Rysman. 2006. Coordination versus differentiation in a standards war: 56k modems. *RAND Journal of Economics* 37(4): 887-909.
- Baldwin, Carliss and Kim B. Clark. 2000. *Design Rules: The Power of Modularity, Volume 1*. Cambridge: MIT Press.
- Bresnahan, Timothy F. and Shane Greenstein. 1996. Technical progress and co-invention in computing and the uses of computers. *Brookings Papers on Economic Activity: Microeconomics* (1996): 1-83.
- Bresnahan, Timothy F. and Shane Greenstein. 1999. Technological Competition and the Structure of the Computer Industry. *The Journal of Industrial Economics* 47(1): 1-40.
- Broersma, Ron. 2012. DREN & Federal IPv6 Update. Presentation, Internet2 Joint Techs, Summer 2012, Stanford, CA.
- Brynjolfsson, Erik and Joo Hee Oh. 2012. The Attention Economy: Measuring the Value of Free Goods on the Internet. Accessed at http://conference.nber.org/confer/2012/EoDs12/Brynjolfsson_Oh.pdf on April 17, 2014.
- Cannon, Robert. 2010. Potential Impacts on Communications from IPv4 Exhaustion & IPv6 Transition. Federal Communications Commission Staff Working Paper 3.
- Ceccagnoli, Marco, Chris Forman, Peng Huang, and D.J. Wu. 2012. Co-creation of Value in a Platform Ecosystem: The Case of Enterprise Software. *MIS Quarterly* 36(1): 263-290.
- Cerf, Vinton. 1990. IAB Recommended Policy on Distributing Internet Identifier Assignment. RFC 1174.
- Ceruzzi, Paul. 2003. *A History of Modern Computing*. Cambridge: MIT Press.
- Chandler, Nav. 2012. The Business Case for Delivering IPv6 Service Now. White Paper, IDC.
- Chandler, Nav. 2013. The Business Case for IPv6 Services in Mobile Networks. White Paper, IDC.
- Chargonda, Shraavan Kumar. 2014. Promoting and Encouraging the Deployment of IPv6 in ASEAN ICT SMEs. IPv6 Conference 2014, Singapore.
- Choi, Jay Pil. 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.

- Cisco. 2014. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013-2018.
- Colitti, Lorenzo, Steinar H. Gunderson, Erik Kline, and Tiziana Refice. 2010. Evaluating IPv6 Adoption in the Internet. Passive and Active Measurement Conference 2010. Accessed at <http://static.googleusercontent.com/media/research.google.com/en/us/pubs/archive/36240.pdf> on March 25, 2014.
- Colitti, Lorenzo and Erik Kline. 2011. World IPv6 \$Next: Impact on Japan. Presentation.
- Colitti, Lorenzo and Erik Kline. 2013. IPv6 in Japan. Presentation at IPv6 Summit in Tokyo 2013.
- Corts, Kenneth S. and Mara Lederman. 2009. Software exclusivity and the scope of indirect network effects in the U.S. home video game market. *International Journal of Industrial Organization* 27: 121-136.
- Cusumano, Michael A. and Annabelle Gawer. 2002. The Elements of Platform Leadership. *MIT Sloan Management Review*. Accessed at <http://sloanreview.mit.edu/article/the-elements-of-platform-leadership/> on April 15, 2014.
- Cusumano, Michael A. and David B. Yoffie. 1998. Competing on Internet Time: Lessons from Netscape and Its Battle with Microsoft. New York: Free Press.
- David, Paul A. 1969. A contribution to the theory of diffusion. Memorandum No. 71, Stanford Center for Research in Economic Growth, Stanford University.
- David, Paul A. and Julie Ann Bunn. 1988. The Economics of Gateway Technologies and Network Evolution: Lessons from Electricity Supply History. *Information Economics and Policy* 3: 165-202.
- David, Paul A. and Shane Greenstein. 1990. The Economics of Compatibility Standards: An Introduction to Recent Research. *Economics of Innovation and New Technology* 1(1/2): 3-41.
- Deering, S. and R. Hinden. 1995. Internet Protocol, Version 6 (IPv6) Specification. IETF RFC 1883, Available at <https://tools.ietf.org/html/rfc1883>.
- Deering, S. and R. Hinden. 1998. Internet Protocol, Version 6 (IPv6) Specification. IETF RFC 2460, Available at <https://tools.ietf.org/html/rfc2460>.
- Donley, C., L. Howard, V. Kuarsingh, J. Berg, J. Doshi. 2013. Assessing the Impact of Carrier-Grade NAT on Network Applications. IETF RFC 7021, available at <http://tools.ietf.org/html/rfc7021>.
- Dranove, David, Chris Forman, Avi Goldfarb, and Shane Greenstein. Forthcoming. The Trillion Dollar Conundrum: Complementarities and Health Information Technology. *American Economic Journal: Policy*.
- Draves, R. 2003. Default Address Selection for Internet Protocol version 6 (IPv6). IETF RFC 3484, available at <http://www.ietf.org/rfc/rfc3484.txt>.
- Edelman, Benjamin. 2009. Running Out of Numbers: Scarcity of IP Addresses and What To Do About It. Harvard Business School Working Paper 09-091.
- Edelman, Benjamin and Michael Schwartz. 2013. Pricing and Efficiency in the Market for IP Addresses. Harvard Business School Working Paper 12-020.

- Eisenmann, Thomas, Geoffrey Parker, and Marshall W. Van Alstyne. 2006. Strategies for Two-Sided Markets. *Harvard Business Review OnPoint* 1463: 2-10.
- Elmore, Hillary, Brandon Stephens, and L. Jean Camp. 2008a. Diffusion and Adoption of IPv6 in the ARIN Region. Working Paper, School of Informatics, Indiana University Bloomington.
- Elmore, Hillary, Brandon Stephens, and L. Jean Camp. 2008b. Diffusion and Adoption of IPv6 in the United States. Technical Report TR661, School of Informatics, Indiana University Bloomington.
- Ericsson. 2013. IPv6: Meeting the Challenge of Connecting Everyone and Everything. Ericsson White Paper 284 23-3194.
- Evans, David S. and Andrei Hagiu. 2006. *Invisible Engines: How Software Platforms Drive Innovation and Transform Industries*. Cambridge: MIT Press.
- Farrell, Joseph and Paul Klemperer. 2007. Coordination and Lock-In.: Competition with Switching Costs and Network Effects. In *Handbook of Industrial Organization*, Volume 3, eds. M. Armstrong and R. Porter. Amsterdam: Elsevier, pp. 1967-2072.
- Farrell, Joseph and Garth Saloner. 1992. Converters, Compatibility, and the Control of Interfaces. *Journal of Industrial Economics* 40(1): 9-35.
- Farrell, Joseph and Timothy Simcoe. 2013. Four Paths to Compatibility. In *The Oxford Handbook of the Digital Economy*, eds. Martin Petiz and Joel Waldfogel, Oxford: Oxford University Press, pp. 34-58.
- Flamm, Kenneth. 2013. A Tale of Two Standards: Patent Pools and Innovation in the Optical Disk Drive Industry. NBER Working Paper 18931.
- Ford, M., M. Boucadair, A. Durand, P. Levis, and P. Roberts. 2011. Issues with IP Address Sharing. IETF RFC 6269, Available at <http://tools.ietf.org/html/rfc6269>.
- Forman, Chris. 2005. The Corporate Digital Divide: Determinants of Internet Adoption. *Management Science* 51(4): 641-654.
- Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2008. Understanding the Inputs into Innovation: Do Cities Substitute for Internal Firm Resources? *Journal of Economics and Management Strategy* 17(2): 295-316.
- Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2012. The Internet and Local Wages: A Puzzle. *American Economic Review* 102(1): 556-575.
- Forman, Chris and Anne Gron. 2011. Vertical Integration and Information Technology Investment in the Insurance Industry. *Journal of Law, Economics, and Organization* 27(1): 180-218.
- Ganapathi, A., V. Ganapathi, and D. Patterson. 2006. Windows XP Kernel Crash Analysis. *Proceedings of the 20th Large Installation System Administration Conference*, p. 149-159.
- Gandal, Neil, Michael Kende, and Rafael Rob. 2000. The Dynamics of Technological Adoption in Hardware/Software Systems: The Case of Compact Disc Players. *The RAND Journal of Economics* 31(1): 43-61.

- Gates, Bill. 1995. The Internet Tidal Wave. Internal Microsoft memo. Accessed at <http://www.justice.gov/atr/cases/exhibits/20.pdf> on March 31, 2014.
- Geroski, Paul A. 2000. Models of technology diffusion. *Research Policy* 29: 603-625.
- Goldfarb, Avi, Shane Greenstein, and Catherine Tucker. Forthcoming. Economics of Digitization. Chicago: University of Chicago Press. Preliminary Chapters Available at <http://papers.nber.org/books/gree13-1> (Accessed March 25, 2014).
- Goolsbee, Austan and Peter J. Klenow. 2006. Valuing Consumer Products by the Time Spent Using Them: An Application to the Internet. *American Economic Review* 96(2): 108-113.
- Gowrisankaran, G. and J. Stavins. 2004. Network Externalities and Technology Adoption: Lessons from Electronic Payments. *RAND Journal of Economics* 35: 260-276.
- Greenstein, Shane. 2007. Economic Experiments and Neutrality in Internet Access. NBER Working Paper 13158.
- Greenstein, Shane. 2010. Innovative Conduct in U.S. Commercial Computing and Internet Markets. In *Handbook on the Economics of Innovation*, edited by Bronwyn Hall and Nathan Rosenberg. Burlington: Academic Press. Pp. 477-538.
- Greenstein, Shane. 2012. Concentration in Internet Access and Entrepreneurial Truncation of Innovation. *Capitalism and Society*. Volume 7, Issue 1.
- Greenstein, Shane and Ryan C. McDevitt. 2011. The broadband bonus: Estimating broadband Internet's economic value. *Telecommunications Policy* 35: 617-632.
- Greenstein, Shane and Victor Stango. 2007. *Standards and Public Policy*. Cambridge: MIT Press.
- Griliches, Zvi. 1957. Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica* 25(4): 501-522.
- Grundemann, Chris. 2011. NAT444 (CGN/LSN) and What it Breaks. Blog Post. Accessed at <http://chrisgrundemann.com/index.php/2011/nat444-cgn-lsn-breaks/> on April 16, 2014.
- Hall, Bronwyn. Adoption of New Technology. NBER Working Paper 9730.
- Harley, C. 1973. On the persistence of old techniques: the case of North American wooden shipbuilding. *Journal of Economic History* 33: 372-398.
- Hersaeus, Erika. 2014. IPv6 in the Swedish Public Sector. Presentation, Netnod Spring Meeting, 2014.
- Howard, Lee. 2013a. Internet Access Pricing in a Post-IPv4 Runout World. White Paper.
- Howard, Lee. 2013b. Record of Transfers in the ARIN Region. Blog Post. Accessed at <http://www.wleecoyote.com/blog/transfers.htm> on April 2, 2014.
- Howard, Lee. 2013c. Game Theory Spreadsheet. Accessed at <http://www.asgard.org/documents.html> on April 4, 2014.

- Howard, Lee. 2013d. Total Cost of Ownership (TCO) of IPv6. INET Denver Presentation. Accessed at <https://www.youtube.com/watch?v=vXf8Zlew1j0> on August 14, 2014.
- Huang, Peng Huang, Marco Ceccagnoli, Chris Forman, and D.J. Wu. 2013. Participation in a Platform Ecosystem: Appropriability, Competition, and Access to the Installed Base. *Management Science* 59(1): 102-121.
- Hubbard, K., M. Koster, D. Conrad, D. Karrenberg, and J. Postel. 1996. Internet Registry IP Allocation Guidelines. RFC 2050.
- Huston, Geoff. 2005. IPv4 Address Utilization. White Paper.
- Huston, Geoff. 2010. Measuring More IPv6. *The ISP Column*. Accessed at <http://www.potaroo.net/ispcol/2010-04/ipv6-measure.pdf> on April 19, 2014.
- Huston, Geoff. 2012. Bemused Eyeballs: Tailoring Dual Stack Applications in a CGN Environment. White Paper. Accessed at <http://labs.apnic.net/blabs/?p=188> on April 3, 2014.
- Huston, Geoff. 2013a. IP Addresses and Traceback. *The ISP Column*. Accessed at <http://www.potaroo.net/ispcol/2013-11/traceback.html> on July 28, 2014.
- Huston, Geoff. 2013b. All IP Addresses are not the same. *The ISP Column*. Accessed at <http://www.potaroo.net/ispcol/2013-09/iptests.html> on August 2, 2014.
- Huston, Geoff. 2013c. Valuing IP addresses. *The ISP Column*. Accessed at <http://www.potaroo.net/ispcol/2013-09/valuation.html> on August 2, 2014.
- Huston, Geoff and George Michaelson. 2008. IPv6 Deployment: Just where are we? *The ISP Column*. Accessed at <http://www.potaroo.net/ispcol/2008-04/ipv6.pdf> on April 19, 2014.
- InterConnect Communications. 2012. Ofcom MC/111 Internet Protocol Version 6 Deployment Study.
- IPv6 Task Force. 2006. Technical and Economic Assessment of Internet Protocol Version 6 (IPv6). White Paper, U.S. Department of Commerce, National Institute of Standards and Technology, National Telecommunications and Information Administration.
- Karshenas, M. and P. Stoneman. 1993. Rank, stock, order and epidemic effects in the diffusion of new process technologies: An empirical model. *RAND Journal of Economics* 24: 503-528.
- Kao, Merike, Jan Zorz, and Sander Steffan. 2012. Requirements for IPv6 in ICT Equipment. RIPE Document 554.
- King, Rachel and Danny Yadron. 2014. Windows XP: Old Platforms Die Hard, Security Risks Live On. *The Wall Street Journal*. Accessed at <http://online.wsj.com/news/articles/SB10001424052702304157204579473821528178710?KEYWORDS=windows+xp&mg=reno64-wsj> on April 4, 2014.
- Kolko, Jed. 2012. Broadband and Local Growth. *Journal of Urban Economics* 71: 100-113.
- Kretschmer, Tobias and Eugenio Miravete. 2012. Competitive Pressure and the Adoption of Complementary Technologies. *American Economic Review* 104(2): 1540-1570.

- Kunc, Urban, Ivan Pepelnjak, Janez Sterle, Matjaz Straus Istenic, Andrej Kobal, Simeon Lisec, Olaf Maennel, and Jan Zorz. 2012. Study: Transition to IPv6 (Guidelines for Deliberation on the National IPv6 Strategy). Accessed at http://go6.si/docs/Study_MVZT_IPv6_en.pdf on April 9, 2014.
- Lehr, William, Tom Vest, and Elior Lear. 2008. Running on Empty: the challenge of managing Internet Addresses. Paper Presented at the 2008 TPRC Conference. Accessed at http://cfp.mit.edu/publications/CFP_Papers/Lehr%20Lear%20Vest%20TPRC08%20Internet%20Address%20Running%20on%20Empty.pdf on April 17, 2014.
- Macleod, C. 1992. Strategies for innovation: the diffusion of new technology in nineteenth century British industry. *Economic History Review* 45: 285-307.
- Maskin, Eric S. 1992. Auctions and privatization. In *Privatization*, ed. H. Siebert, Institut fur Weltwirtschaften der Universitat Kiel.
- Mawatari, M., M. Kawashima, and C. Byrne. 2013. 464XLAT: Combination of Stateful and Stateless Translation. RFC 6877. Accessed at <http://tools.ietf.org/html/rfc6877> on August 2, 2014.
- Mokyr, Joel. 1990. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford: Oxford University Press.
- Moore, Geoffrey. 2014. *Crossing the Chasm: Marketing and Selling Disruptive Products to Mainstream Customers, 3rd Edition*. Harper Business.
- Mueller, Milton L. 2008. Scarcity in IP Addresses: IPv4 Address Transfer Markets and the Regional Internet Address Registries. Internet Governance Project. Accessed at http://www.internetgovernance.org/wordpress/wp-content/uploads/IPAddress_TransferMarkets.pdf on April 17, 2014.
- Mueller, Milton L. 2010. Critical resource: An institutional economics of the Internet addressing-routing space. *Telecommunications Policy* 34(8): 405-416.
- Mueller, Milton, Brenden Kuerbis, and Hadi Asghari. 2012. Dimensioning the Elephant: An Empirical Analysis of the IPv4 Number Market. White Paper, Internet Governance Project, School of Information Studies, Syracuse University.
- Nobile, Leslie. 2012. IPv4 Countdown Plan. Presentation, ARIN 29, Vancouver, BC. Accessed at https://www.arin.net/participate/meetings/reports/ARIN_XXIX/PDF/monday/nobile_ipv4_countdown.pdf on August 2, 2014.
- Nygren, Erik. 2011. Implementing IPv6 on a Global Scale: Experiences at Akamai. USENIX LISA 2011.
- Nygren, Erik. 2013. World IPv6 Launch Anniversary: Measuring Adoption One Year Later. Blog Post. Accessed at <https://blogs.akamai.com/2013/06/world-ipv6-launch-anniversary-measuring-adoption-one-year-later.html> on April 4, 2014.
- Nygren, Erik. 2014. IPv6: Sweet Spots of Adoption. Network Computing, Accessed at <http://www.networkcomputing.com/networking/ipv6-sweet-spots-of-adoption/a/d-id/1234689?> on April 4, 2014.
- OECD. 2007. *Internet Address Space: Economic Considerations in the Management of IPv4 and in the Deployment of IPv6*. Ministerial Background Report [DSTI/ICCP\(2007\)20/FINAL](http://www.oecd.org/dataoecd/20/12/39122222.pdf).

DSTI/ICCP/CISP(2014)3/FINAL

OECD. 2010. *Internet Addressing: Measuring Deployment of IPv6*. DSTI/ICCP/CISP 17/FINAL.

OECD. 2013. *Building Blocks for Smart Networks*. [DSTI/ICCP/CISP\(2012\)3/FINAL](#).

OECD. 2014. *The Internet in Transition: The State of the Transition to IPv6 in Today's Internet*. [DSTI/ICCP/CISP\(2012\)8/FINAL](#).

Ohashi, Hiroshi. 2003. The Role of Network Effects in the US VCR Market, 1978-1986. *Journal of Economics and Management Strategy* 12(4): 447-494.

Peitz, Martin and Joel Waldfogel. 2013. *Oxford Handbook of the Digital Economy*. Oxford: Oxford University Press.

Perset, Karine. 2007. Internet Address Space: Economic Considerations in the Management of IPv4 and in the Deployment of IPv6. OECD Ministerial Background Report DSTI/ICCP.

PTS. 2011. Deploying IPv6 – Internet Protocol version 6. Report number PTS-ER-2011:18.

Puffert, Douglas. 2000. The Standardization of Track Gauge on North American Railways, 1830-1890. *The Journal of Economic History* 60(4) : 933-960.

Rochet, Jean-Charles and Jean Tirole. 2006. Two-sided markets: a progress report. *RAND Journal of Economics* 37(3): 645-667.

Rogers, Everett. 2010. *Diffusion of Innovations, 5th Edition*. New York: Free Press.

Rohlfs, Jeffrey. 2001. Bandwagon Effects in High-Technology Industries. Cambridge: MIT Press.

Rohlfs, Jeffrey. 1974. A Theory of Interdependent Demand for a Communications Service. *The Bell Journal of Economics and Management Science* 5(1): 16-37.

Rosenberg, Nathan. 1992. Economic Experiments. *Industrial and Corporate Change* 1(1): 181-203.

Rosston, Gregory, Scott Savage, and Donald Waldman. 2010. Household Demand for Broadband Internet Service. *The B.E. Journal of Economic Analysis and Policy* 10(1).

Rysman, Marc. 2009. The Economics of Two-Sided Markets. *The Journal of Economic Perspectives* 23(3): 125-143.

Saloner, Garth. 1990. Economic issues in computer interface standardization. *Economics of Innovation and New Technology* 1: 135-156.

Shapiro, Carl and Hal R. Varian. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard University Press.

Simcoe, Timothy. Forthcoming. The Endogenous Modularity of the Internet, in *The Economics of Digitization*, (eds.) Avi Goldfarb, Shane Greenstein, and Catherine Tucker. Chicago: University of Chicago Press.

Smith, Adam. 1776. *The Wealth of Nations*.

- Stern, Scott. 2005. Economic Experiments: The Role of Entrepreneurship in Economic Prosperity. In *Understanding Entrepreneurship: A Research and Policy Report*. Kauffman Foundation.
- Strang, D. and S. A. Soule. 1998. Diffusion in organizations and social movements. *Annual Review of Sociology* 24: 265-290.
- Townsley, Mark. 2014. IPv6-Centric Networking. Presentation, IPv6 World Congress, Paris France, 2014.
- Wallsten, Scott. Forthcoming. What are We Not Doing When We're Online? In *Economics of Digitization: An Agenda*, eds. Avi Goldfarb, Shane Greenstein, and Catherine Tucker.
- Wing, D. and A. Yourtchenko. 2011. Happy Eyeballs: Success with Dual-Stack Hosts. Draft IETF Specification, available at <http://tools.ietf.org/html/draft-ietf-v6ops-happy-eyeballs-05>.

NOTES

¹ The definition of converter technology commonly employed in the economics literature on standards is used. Namely, converters are technologies that reduce the costs of incompatibility between competing standards. For example, for our purposes, software that allows users to read and exchange between different file types function as converters.

² It can be noted that even network operators who invest in IPv6 will need to continue to operate IPv4, and if the network operator has insufficient addresses, will need to operate carrier grade NAT. This paragraph simply emphasizes how the existence of carrier grade NAT reduces the economic incentives to make new investments in IPv6.

³ There have been many studies that have measured the economic implications of the Internet, which are far too numerous to mention here. For recent examples of studies that have shown the economic impacts of business Internet and broadband diffusion, see Forman, Goldfarb, and Greenstein (2012) and Kolko (2012). For recent surveys, see the articles contained in Peitz and Waldfoegel (2013) and Goldfarb, Greenstein, and Tucker (Forthcoming).

⁴ For details on this transition see, for example, Ceruzzi (2003).

⁵ Based on citations from IETF and W3C publications and US Patents citations, analysed in Simcoe (Forthcoming).

⁶ Simcoe (Forthcoming) provides an overview of the time series of IETF standards and their citation patterns. According to his study, some recently developed standards are also very widely cited. For example, Session Initiation Protocol (SIP), which is used to control voice and video calls over IP networks, is the second most cited standard by IETF and W3C citations (341 times as of 2013).

⁷ Huston (2005) notes that at the time contemporary protocols were using 8- or 16-bit address fields, translating to networks of between 256 and 65,356 hosts.

⁸ Source: The IPv4 Address Report, maintained by Geoff Huston, accessed at <http://www.potaroo.net/tools/ipv4/index.html> on March 25, 2014. As noted on that page, "'Exhaustion' is defined here as the time when the pool of available addresses in each RIR reaches the threshold of no more general use allocations of IPv4 addresses." The threshold for exhaustion date differs from each RIR. Note that Figure 1 depicts the timeline for completely running out of IPv4 addresses, which differs from the concept of "exhaustion" used in the text. As a result, the dates for "running out" of IPv4 addresses differs somewhat in the text and from what is shown in Figure 1.

⁹ Here we date availability from RFC 2460 "Internet Protocol, Version 6 (IPv6) Specification" (Deering and Hinden 1998). Some RFCs defining IPv6 were released previously, starting with RFC 1883 (Deering and Hinden 1995).

¹⁰ Other changes include the management of fragmentation, the handling of the IP options field, and a simplified packet header

¹¹ However, as Geoff Huston notes, the current IPv6 address space encompasses a theoretical maximum of 282 trillion end sites (OECD 2014). This is still a very large number, but significantly smaller than the 2^{128} in the theoretical address space.

12 Statistics obtained from <https://www.google.com/intl/en/ipv6/statistics.html>. Accessed on March 25, 2014. These statistics measure the percentage of users accessing Google's services over IPv6. For discussion on methodology, see Colitti, Gunderson, Kline, and Refice (2010).

13 For discussion of other methodologies for measuring IPv6 diffusion, see OECD (2010) and the statistics described on <http://6lab.cisco.com/stats/information.php>.

14 Note that these statistics represent worldwide averages; there is considerable variability across countries.

15 For discussions of the problems with address sharing and carrier grade NATs, see for example Geoff Huston's analysis in OECD (2013) and InterConnect Communications (2012). For examples of ways in which user experience may be degraded, see Donley et al. (2013). For a description of how user privacy and accountability can be affected, see Internet Society (2013).

16 For further details, see Bresnahan and Greenstein (1999) and Greenstein (2010).

17 This report follows Farrell and Simcoe (2013) in using a broad definition of SSOs that includes globally recognised standard setters such as the ITU, private bodies that manage a standard such as the IETF, as well as smaller consortia that focus on a particular technology such as the Blu-Ray Disc Association. In some cases, it could also refer to regulatory agencies that sometimes have the power to mandate standards, such as the United States Federal Communications Commission.

18 It is not uncommon to achieve standards compatibility through multiple means. For a case study of the diffusion of 56K modem technology that included both a standards war and intervention by a SSO, see Augereau, Greenstein, and Rysman (2006).

19 For example, Nygren (2011) discusses some common issues that give rise to broken IPv6 implementations in the context of content providers. Broersma (2012) discusses some common issues that might arise when implementing IPv6 in an internal network. The document returns to this issue in greater detail below.

20 See Rohlfs (2001) for a discussion of the assumptions needed for Metcalfe's law, and why the value of a network may be increasing with the number of users but by less than its square.

21 See, for example, <http://www.worldipv6launch.org/measurements/>.

22 A platform leader is a company that drives innovation in the platform (e.g., Cusumano and Gawer 2002).

23 This episode is recounted in further detail in Cannon (2010).

24 Source: <http://www.worldipv6launch.org/measurements/>. Accessed on March 31, 2014.

25 Data source: Status of the LTE Ecosystem Report (www.gsacom.com). By comparison, a recent analysis by Cisco shows that 18% of mobile devices were potentially IPv6-capable in 2013 (Cisco 2014).

26 Source: <http://www.internet2.edu/presentations/jt2012summer/20120718-Broersma-DREN%20IPv6%20Update%20JtTechs-Jul2012.pdf>. Accessed March 31, 2014.

27 A further example comes from Australia. Internally, the Australian Government Information Management Office has been co-ordinating the transition of all government agencies' external-facing services to IPv6, which is nearing completion. As a major consumer of ICT equipment, this is expected to stimulate demand for IPv6 capable hardware and software in Australia more broadly.

28 For more complete discussions of how each of these technologies work see, for example, InterConnect Communications (2012), Ericsson (2013), and OECD (2014).

29 For an introduction to these issues, see Rysman (2009) and Evans and Hagiou (2006).

30 For further discussion of these points, see Farrell and Simcoe (2013).

31 See Aoun and Davies (2007) <http://www.ietf.org/rfc/rfc4966.txt> for further details.

32 It is important to note that translation does not eliminate the need to have access to IPv4 addresses. A difference between translation and dual stack is where these addresses are deployed. Under a dual stack scenario, IPv4 addresses are deployed either at the Network Address Translator (NAT) or at the end devices. Under translation the addresses are needed at the interface translator.

33 Note that this role of the browser as a focal application was to some extent foretold by the Mosaic browser. For more on this, see Cusumano and Yoffie (1998).

34 For an example of this, see the well-known internal memo written by then-Microsoft CEO Bill Gates, “The Internet Tidal Wave.”

35 Note that the diffusion of client/server was subject to indirect network effects, since the increasing diffusion of the platform lead to the provision of more third party products and services.

36 For a brief overview of NAT 44, some of the problems it creates, and methods used to overcome these problems see Grundemann (2011).

37 For example, see Ford et al. (2011) and Donley et al. (2013).

38 For a discussion of modularity within the context of IBM’s System/360 mainframe, see Baldwin and Clark (2000) and Bresnahan and Greenstein (1999).

39 For a counterexample where joining the platform requires explicit permission from the platform leader and where transaction costs are much higher, see Huang, Ceccagnoli, Forman, and Wu (2013) and Ceccagnoli, Forman, Huang, and Wu (2012).

40 This is a necessarily short discussion of markets for IPv4 addresses. For more information, see for example Mueller, Kuerbis, and Asghari (2012), Howard (2013b), and Edelman (2009).

41 Edelman (2009, p. 9) offers a good example of this. “For example, suppose a network needs 2^{16} addresses. If the network obtains 2^{16} contiguous addresses, others’ routers can add just a single routing entry. But if the network obtains eight non-contiguous blocks of 2^{13} , eight routing entries will be required. If a network considers only its self-interest, it will choose the eight 2^{13} blocks over the single 2^{16} any time the former costs less – imposing a negative externality through extra routing costs.”

42 For a detailed discussion of potential market rules for transfer of IP addresses, see Edelman and Schwartz (2013).

43 For some recent statistics, see <http://www.wleecoyote.com/blog/transfers.htm> as well as recent statistics at ARIN <https://www.arin.net/knowledge/statistics/transfers.html> and the other RIRs.

44 Requesters were required to demonstrate need within one year from 2009-2011

45 For more detail, see the press release at <http://www.marketwired.com/press-release/ipv4-market-group-reviews-2013-year-in-trading-predicts-trends-for-2014-1873805.htm> and for more recent data see <http://ipv4auction.com/>.

- 46 Private value auctions are auctions for which each buyer knows how much she values the good, but this valuation is known only to her. Common value auctions are those where the value of the good is the same for everyone, but each market participant has a separate signal (private information) about what the value of that good might be.
- 47 For further discussion of these issues, see Wallsten (Forthcoming).
- 48 For discussion of individual vendor IPv6 product offerings see, for example, Chandler (2012) and Ericsson (2013). While IPv6 functionality may be supported, not all networking features may be supported over IPv6 in all devices.
- 49 This lack of co-ordination between legally distinct but vertically related parties is common in new technology diffusion. For example, it is a major problem in establishing electronic links between firms. This has occurred when, for example, trying to establish electronic links between vendors and their sales force (e.g., Forman and Gron 2011) or between hospitals seeking to share electronic medical records (Dranove, Forman, Goldfarb, and Greenstein Forthcoming).
- 50 We note here, as elsewhere, that many of these benefits may occur in the future and will depend upon the adoption decisions of other agents within the platform.
- 51 For some data on the number of products that support IPv6, see the IPv6 Ready Program Approved List at <https://www.ipv6ready.org/db/index.php/public/?o=6>; the home gateway list at <https://www.iol.unh.edu/services/testing/hnc/cerouter.php>; and the USGv6 test program list at <https://www.iol.unh.edu/services/testing/ipv6/usgv6tested.php>.
- 52 Note that IPv6 was introduced in Windows XP through Service Packs, however it was not configured to be on by default.
- 53 See RFC 3483 (“Default Address Selection for Internet Protocol version 6 (IPv6)”), accessed at <http://www.ietf.org/rfc/rfc3484.txt> on April 3, 2014 (Draves 2003).
- 54 The “Happy Eyeballs” implementation was originally described in Wing and Yourtchenko (2011) at <http://tools.ietf.org/html/draft-ietf-v6ops-happy-eyeballs-05>. An updated version appears in RFC 6555.
- 55 See for example a user manual for a Sony TV. http://docs.esupport.sony.com/imanual/NA/2013/KDL32_50W650A/uc_uen/c_cntnet_ipv6.html
- 56 <https://www.mail-archive.com/ipv6-ops@lists.cluonet.de/msg00780.html> message on an IPv6 mailinglist. Similar messages in community support forums for the PlayStation.
- 57 There are other sources of IPv6 requests. In particular, <http://www.worldipv6launch.org/measurements/> includes statistics on IPv6 deployments based on statistics from a range of content providers. However, the range of network operators included is smaller, as they include only participants in World IPv6 Launch (<http://www.worldipv6launch.org/participants/?q=2>).
- 58 There are many different ways of deploying IPv6. These are discussed briefly below, but this report has largely abstracted away from these for the purposes of this report. There are many sources for further information on these differences; the interested reader is referred to, for example, Ericsson (2013) or InterConnect Communications (2012).
- 59 Note that deploying CGNs will also consume IPv4 addresses by definition. CGN technology allows customers who are not assigned a global IPv4 address to access IPv4 content, by optimising global IPv4 address usage.

60 This section provides a summary of the trade-offs. For more details, see Howard (2013a).

61 The effects of competitor behaviour on adoption decisions has a long history. For a recent example, see Kretschmer and Miravete (2012).

62 Specifically, he estimates the costs of CPE per user to be \$50 each, but estimates that only half of users will require upgrades.

63 For other estimates of the costs of IPv6 deployment, see for example IPv6 Task Force (2006) and PTS (2011).

64 TeraStream –A Simplified Delivery Model, Peter Lothberg, Deutsche Telekom
<https://ripe67.ripe.net/presentations/131-ripe2-2.pdf>

65 For further discussion of 464XLAT, see Mawatari, Kawashima, and Byrne (2013).

66 Source: <http://www.worldipv6launch.org/measurements/>. Accessed July 23, 2014.

67 Source: <http://www.worldipv6launch.org/measurements/>. Accessed April 5, 2014.

68 Source for Alexa Top 10 websites <http://www.alexa.com/topsites>. Source of IPv6 support <http://ipv6-test.com/validate.php>. Both websites accessed on April 4, 2014.

69 Data are based on IPv6 address requests made to the DNS servers to the Alexa top 500 websites. Websites are weighted based upon the number of page views.

70 A long line of studies in the technology diffusion literature have found a positive correlation between firm size and technology adoption. For a few examples, see David (1969), Karshenas and Stoneman (1993), and Forman (2005). See Geroski (2000) for discussion on this and related issues.

71 He estimates the majority of costs will come from application development and development of security appliances.

72 One example that we heard of this is that in backend databases used in production deployments of e-commerce sites, the size of an old IP address field may be insufficient to accommodate IPv6 websites. Thus, deployment of IPv6 requires complementary changes to other infrastructure and applications. As another example, Nygren (2011) describes some common problems that have historically arisen when implementing IPv6 in DNS queries.

73 Nygren lists a number of common problems, including “unreliable 6to4, the browser tries IPv6 when IPv4 would have been better, bad IPv6 connectivity, and slow fallback to IPv4 with long timeout.”

74 The existence of such capabilities have historically played an important role in the adoption of prior generations of network technologies. For examples, see Bresnahan and Greenstein (1996), Forman (2005), and Forman, Goldfarb, and Greenstein (2008).

75 Some argue that NAT also, when used in conjunction with local addressing schemes, can conceal users from unwanted communications and provide some degree of network security. In particular, because NATs prevent outside parties from initiating communications, they protect from some types of viruses and worm probes (IPv6 Task Force 2006). However, others argue this is a poor approach for providing network security (e.g., Grundemann 2013).

76 There are a range of issues around how to measure diffusion of IPv6 that, in the interests of brevity, will not be discussed here. These statistics can vary significantly based on how they are collected and the

population of users from which the statistics are drawn, among other things. For a detailed discussion of statistics measuring diffusion of IPv6, see OECD (2010), Huston and Michaelson (2008), and Huston (2010). For the present purposes, it is sufficient that IPv6 is currently at the early stages of diffusion.

77 For example, Griliches (1957), in his seminal paper on the diffusion of hybrid seed corn, estimated S-curves using observations between the 5th and 95th percentile. The reason is that such observations “are liable for very large percentage errors and would have very little weight anyway in any reasonable weighting scheme” (Griliches 1957, p. 504).

78 For another recent example of this dynamic at work within the context of healthcare IT, see Dranove, Forman, Goldfarb, and Greenstein (Forthcoming).

79 For more details on this case study, see <http://www.ida.gov.sg/Infocomm%20Landscape/Technology/IPv6> or Chargonda (2014).

80 This is the “no islanding” principle described, for example, at https://www.ida.gov.sg/~media/Files/PCDG/Consultations/20110620_NoIslandingPrinciple/NoIslandingPrinciple.pdf.

81 “UK smart meter network based on IPv6 wireless mesh tech”, August 21, 2013 <http://www.telecoms.com/173322/uk-smart-meter-network-based-on-ipv6-wireless-mesh-tech/>

82 The Connode Clients and the Mesh <http://www.connode.com/products-and-services/connode-clients-mesh/>

83 Why Bosch rigs up IPv6, <http://blog.bosch-si.com/wp-content/uploads/2011/12/20111201V05-Bosch-IPv6.pdf>