ECAR

Transitioning to IPv6

October 2013

EDUCAUSE CENTER FOR ANALYSIS AND RESEARCH

EDUCAUSE

Authors and Contributors

Special thanks go to the following ECAR Campus Cyberinfrastructure Working Group (ECAR-CCI) authors and contributors to this report. For more information about the ECAR Working Groups, go to http://www.educause.edu/ecar/ecar-working-groups.

Guy T. Almes, *ECAR-CCI Chair*
Director, Academy for Advanced Telecommunications and Learning Technologies
Texas A&M University

Celeste Anderson
Director, External Networking Group Information Technology Services
University of Southern California

Michael R. Mundrane
Vice Provost and Chief Information Officer
University of Connecticut

Valerie E. Polichar
Infrastructure Analyst and Technical Projects Coordinator, Administrative Computing and Telecommunications
University of California, San Diego

We are grateful for the following consultants who provided early reviews and input to this report

R. P. Aditya
Senior Network Architect
University of Michigan, Ann Arbor

John Baird
HPCMP IPv6 Implementation Manager, High Performance Computing Program Office/Defense Research and Engineering Network
United States Department of Defense

Owen DeLong
IPv6 Evangelist and Director of Professional Services
Hurricane Electric

David Farmer
Technical Expert, Office of Information Technology
University of Minnesota

Jeff Harrington
Senior Network Engineer
NYSERNet

Michael H. Lambert
Network Engineer
Pittsburgh Supercomputing Center

Joe St Sauver, *Ph.D.*
Manager, Internet2 Nationwide Security Programs, and Manager, InCommon Certificate Service
University of Oregon

Ray Patrick Soucy
Network Engineer
University of Maine System

Alan Whinery
Information Technology Specialist, Information Technology Services, Technology Infrastructure
University of Hawaii

Table of Contents

1. Foreword

Internet Protocol version 6 (IPv6), the named successor to the current networking protocol that supports the transport of traffic across the Internet today, is a more mature standard that appropriately balances compatibility against the addition of capabilities and enhancements that are important for current and future growth of the Internet. It has a large address space that effectively deals with the address exhaustion present in the existing standard and contains improved support in a number of key areas that will better address the needs of big data and other cyberinfrastructure initiatives.

These clear benefits aside, there will be challenges at any given institution to incorporating IPv6 support within existing network infrastructures, but like it or not, IPv6 is coming. Partner institutions and researchers who use IPv6 will increasingly turn to institutions that can provide this service for collaborative research and prospective students and faculty will increasingly use its availability as a metric of a school's technological readiness. Granting agencies are starting to require IPv6, so lack of IPv6 will soon begin to hurt the institutional bottom line. It is extremely likely that IPv6 traffic is already on your network and this will impact both security and controllability.

But transitioning to IPv6 can't happen overnight. You need to start today to sell the concept to your campus administration and take appropriate steps to initiate the planning process, including assessing and addressing your institution's structural preparedness and developing technical, security, and rollout plans for IPv6. This paper will show you how.

2. Introduction

The Internet is a worldwide system of individual networks connected by standardized application of naming, addressing, and protocols. The ability for any device to transparently function both locally and globally has fostered true geographic independence and has been an underlying catalyst for significant innovation. The Internet has benefitted immeasurably from its organic evolution and the simple, open standards that facilitated its growth and robust operation, allowing it to function as a highly effective laboratory for change.

While Internet Protocol version 4 (IPv4) has been the global standard used to deliver both network and transport capabilities to individual network devices for the past 30 years, its ability to sustain and support the type of growth we have come to expect and demand from the Internet is questionable. IPv4 provides network capabilities both through *addresses* that allow individual devices to be uniquely located on the Internet and via *transport capabilities*, low-level services that network-connected devices use for communication. The address space (the total number of unique addresses possible) of the IPv4 standard is nearly exhausted, meaning that

almost all of the addresses have already been allocated for use. Thus far, this has been partially ameliorated by network address translation (NAT), which allows for a smaller number of addresses to be utilized by a larger number of devices. But while there are some limited situations where this technique might be attractive, in general NAT is a technical compromise that introduces unnecessary complexity to communication and services; can challenge security efforts; and can negatively impact performance (see Section 4.5, Security). The package of transport services integrated into IPv4, which was forward-looking at the time of its design and implementation, is no longer adequate for generally safe and effective networking on an open Internet. Like the addressing challenge, this has also been partially addressed by employing creative approaches in both hardware and software, but with a similar adverse impact on both complexity and performance.

The design process for IPv6 started in the early 1990s and was ready for wide-scale adoption by 2010. Designed and maintained by the same Internet Engineering Task Force (IETF) experts who oversee all Internet protocols, it differs from the older IPv4 protocol in a number of significant ways. Where IPv4 addresses are 32 bits, IPv6 addresses are 128 bits. Thus, while IPv4 supports approximately 4 billion unique addresses, the IPv6 address space is approximately 4 billion times 4 billion times 4 billion times the size of the IPv4 address space—more than enough to provide unique IP addresses for the over 7 billion people and their devices (phones, laptops, tablets, desktops, and, yes, even refrigerators and home appliances). This increase in usable address space is truly profound. Kim Davies and Kieran McCarthy have described the difference this way: "[I]f all the IPv4 space would fit in an iPod, then all the IPv6 space is the size of the entire Earth."[1] IPv6 also changes the way packets (small chunks of data and their associated addressing/metadata) are formatted, providing more integrated support for security and service protocols.

IPv6 has enjoyed intermittent attention among various technical communities and has continued to be part of the ongoing discussion pertaining to the operation and evolution of the Internet. While early emphasis was relatively modest, this has been changing steadily, with 2013 marking "the third straight year global IPv6 usage has doubled."[2] Historically, IPv6 was a technology that was not strictly needed, and most campuses responded with a "wait and see" attitude. This is changing rapidly, and current chatter about IPv6 and its association with campus cyberinfrastructure has taken on a more urgent tone for a number of key reasons. Campuses need to be familiar with the current state of IPv6, recent developments associated

[1] Kieren McCarthy, "IPv6, the iPod and the Earth," *ICANN Blog*, June 17, 2007, accessed September 26, 2013, http://blog.icann.org/2007/06/ipv6-the-ipod-and-the-earth/.
[2] Phil Roberts, "IPv6 Deployment Hits 2%, Keeps Growing," *Tech Matters*, Internet Society, September 24, 2013, accessed September 25, 2013, http://www.internetsociety.org/blog/2013/09/ipv6-deployment-hits-2-keeps-growing.

with the Internet and Internet-related activities, and the relationship of both to their cyberinfrastructure planning.

3. Why Now

What happens without IPv6 deployment? One common misperception is that the alternative to deploying IPv6 is simply to stay with the "tried and true" IPv4 Internet. This is not really true. The growth and success of the IPv4 Internet was based on globally routable addresses, with each host capable of communicating with every other Internet host. As IPv4 address exhaustion progresses, campuses are finding that new hosts are connected via NAT gateways, sometimes nested ("double NAT" or "carrier grade NAT [CGN]"). This raises at least two kinds of issues: (a) the users who do this will experience poorer functionality, performance, and reliability than with classical globally routable IPv4 addresses, and (b) very often these NAT segments are not manageable or even visible to central campus networking staff, resulting in operational confusion and creating a real risk that whole segments of or even entire campuses may be blocked if a single host causes a visible problem. Increasing use of NAT will make central organizations lose control and visibility of what's happening on the network. Even if a university itself has sufficient IPv4 space and does not move to NAT, the move towards NAT solutions by other organizations that are short on address space weakens security and manageability for the Internet as a whole. As more institutions support IPv6, the Internet community will increasingly use IPv4 NAT in favor of better solutions.

Ignoring IPv6 will not make it go away and is not the functional equivalent of deferring its adoption. Quite aside from reasons why you may wish to support IPv6 on your campus network is the simple fact that it is *already on* your network—whether you like it or not.[3] Easy-to-use "script kiddie" hacking packages are already available to attack your network using IPv6, and customers connecting to your network from within using "owned" machines may unwittingly bring hackers along with them. Microsoft's Active Directory currently defaults to using IPv6 for much of its internal communication. Default settings on many contemporary end-user machines will attempt to find a 6-to-4 tunnel (a technique for encapsulating IPv6 traffic on an IPv4 network) and use IPv6 whenever possible. This means that IPv6-based activity is already occurring on existing networks, invisible to campus security staff. These systems will simply make up workable IPv6 addresses for themselves when not provided an authentic address. While this is an intended and attractive feature of the IPv6 standard, it is far less desirable as an unexpected and unobserved communication path within existing IPv4 networks.

[3] Joseph St. Sauver, "If We Do IPv6, Will It Help or Hurt Our Security?" (presentation at the Internet2/ESnet Joint Techs meeting, College Station, TX, February 2, 2009), accessed September 26, 2013, http://pages.uoregon.edu/joe/ipv6-security/ipv6-security.pdf.

A thoughtful transition to IPv6 at your institution today will allow for the continued, effective use of any existing IPv4 address where truly needed, expand access to additional capabilities and resources, and enhance overall security by effectively dealing with all of the traffic that may be present.

3.1. State of the Industry

Key vendors made their larger routers and switches IPv6-compatible to varying degrees years ago, but over the past 2–5 years, functionality has been added to these products to make them fully and consistently usable in IPv6 environments. In addition, many smaller routers and switches now also support IPv6. The Wikipedia page on IPv6 compatibility in routing equipment indicates compatibility for the vast majority of commonly used routers.[4] Firewall equipment from major manufacturers such as Check Point, Palo Alto, Juniper, Cisco, and others is now IPv6 compatible by default. Further, enough time has passed that much of the older equipment and infrastructure that did not support IPv6 is likely to have already migrated out of active use.[5]

3.2. Government and Granting Agency Progress

The U.S. government required its agencies to make their networks IPv6 compatible by 2010 and to provide IPv6 on public-facing services by 2012,[6] while the Canadian government intends to complete IPv6 deployment to public-facing websites by early 2015.[7] By the end of 2014, U.S. government sites should be using native IPv6. Although compliance has thus far been less complete than hoped, granting agencies such as the National Science Foundation (NSF) and the National Institute of Health (NIH) have begun to encourage IPv6 use by building IPv6 requirements into their calls for proposals and awards and by using IPv6 capabilities in proposal evaluation.[8] This is particularly true in high-tech fields but will increasingly be the case in other fields as well. Some agencies began to require IPv6 access for funded projects' websites in 2012. Already some calls for proposals refer to IPv6 as a desirable functionality of awardee locations. NASA has already required that some subcontractors make data and

---

[4] Wikipedia contributors, "Comparison of IPv6 Support in Routers," *Wikipedia, The Free Encyclopedia*, accessed September 26, 2013, http://en.wikipedia.org/wiki/Comparison_of_IPv6_support_in_routers.

[5] "IPv6 Knowledge Base: General Information," Department of Defense High Performance Computing Modernization Program (DoD HPC), accessed September 25, 2013, http://www.hpc.mil/cms2/index.php/ipv6-knowledge-base-general-info.

[6] Carolyn Duffy Marsan, "White House Issues IPv6 Directive: Agencies Must Support Next-gen Internet on All Public-facing Web Sites by Fall 2012," *NetworkWorld*, September 28, 2010, accessed September 26, 2013, http://www.networkworld.com/news/2010/092810-white-house-ipv6-directive.html.

[7] "Government of Canada IPv6 Adoption Strategy," Treasury Board of Canada Secretariat, accessed September 26, 2013, http://www.tbs-sct.gc.ca/it-ti/ipv6/ipv6tb-eng.asp.

[8] See the NSF Program Solicitation 13-530 (April 3, 2013) at http://www.nsf.gov/pubs/2013/nsf13530/nsf13530.htm for one example where a campus cyberinfrastructure plan that includes a description of readiness for IPv6 is required by a federal solicitation.

websites IPv6-compatible.[9] The U.S. and Canadian governments are moving to IPv6, if slowly—IPv6 may eventually be required for at least some types of grant submissions.[10] Other governments, such as India's, are now taking an active interest in fostering IPv6 deployment, and areas such as Latin America are aggressively moving forward with their deployments.[11] There are also proposals that .gov domains not be renewed if organizations have not conformed to Domain Name System Security Extensions (DNSSEC) and IPv6.[12]

Thus IPv6 may be both *useful* for attracting awards and *necessary* to comply with award requirements. Likewise, as various national centers or international institutions continue their individual migrations towards IPv6 adoption, the ability for researchers to effectively access and leverage external resources and capabilities via IPv4 will be impeded. This may be a small problem today, but it will only increase and it is likely to grow at Internet speed.

3.3. Campus Cyberinfrastructure May Require IPv6 Sooner

IPv4 addresses will continue to be needed in the near future because the move to IPv6 will include many years of supporting "dual stack" implementations, where both IPv4 and IPv6 may be used on a single machine. The use of IPv6 in applications that require large address blocks will free up existing IPv4 addresses for other uses on campus during the transition period. New, very-high-capacity supercomputers, servers, and clusters are being developed that make use of tens to hundreds of smaller servers, requiring a large number of mostly-internal IP addresses. Even clustered devices can sometimes benefit from being globally visible on the network, and IPv6 would be helpful for these use cases. In addition, IPv6 may enable some types of experiments—such as those that create thousands of virtual endpoints—that are more difficult to support using constrained IPv4 space. Medical environments, led by our institutionally supported research and teaching hospitals, will be on the forefront of independently addressed, network-connected sensors. The robust technical needs for these environments and their significant legal and regulatory requirements should discourage all of us from advocating solutions that do not support the highest levels of performance and security. IPv6 will undoubtedly be the foundation upon which we build these critical-function, high-availability facilities.

---

[9] For more information, see Kevin L. Jones' presentation, "NASA IPv6 Implementation Status" (November 13, 2012), http://gogonetlive.com/pdf/3/kevin_jones.pdf (accessed September 26, 2013).

[10] "Estimating USG IPv6 & DNSSEC External Service Deployment Status," National Institute of Standards and Technology (NIST) Advanced Network Technologies Division, accessed September 26, 2013, http://fedv6-deployment.antd.nist.gov/cgi-bin/generate-gov.

[11] See the Latin American and Caribbean IPv6 Transition Portal page on "Who Is Implementing" for a list of implementers and the status of the implementations at http://portalipv6.lacnic.net/quienes-implementan/ or http://portalipv6.lacnic.net/en/who-is-implementing/.

[12] Scott Hogg, "U.S. Government Progress on IPv6 Deployment: Monitoring the Country's IPv6-enabled Government Services," Core Networking and Security blog, *NetworkWorld*, November 26, 2011, accessed September 26, 2013, http://www.networkworld.com/community/blog/us-government-progress-ipv6-deployment.

4. Making the Case

Campuses have historically been incubators for Internet advances and technologies and have generally assumed the role of promoting the health and evolution of the Internet while simultaneously encouraging innovation, both locally and nationally. In addition, our educational mission presses us to ensure we are "walking the talk" by deploying the technologies that we are preparing students to use and engineer.[13] Consistent with our historical role and in the interests of our institutions and communities, colleges and universities are natural candidates to promote the migration from IPv4 to IPv6. Since this will require broad resources over a number of years, it is essential to have buy-in from upper management.

Campus buy-in is particularly helpful when a vendor's implementation of IPv6 is not complete or consistent. A campus commitment will enable you to push back on vendors (e.g., urge departments to avoid purchase if the product is not IPv6-compatible) and select equipment with IPv6 support that meets a defined campus architecture standard. Getting buy-in from the campus procurement office will help establish a campus-wide commitment to IPv6 by requiring vendors to assert their state of IPv6 compliance as part of the campus' purchasing requirement. This holds for all campuses but may be especially valuable in distributed environments where departments and other administrative groups are empowered to make independent purchasing decisions.

4.1. Sufficient Time for Planning

Developing an address scheme and configuring wired and wireless network components to handle IPv6 are critical early phases of the transition process. Time is required to review the campus addressing structure, evaluate its strengths and flaws, and develop a scheme that fits the emerging needs of your environment. Simply mapping your IPv4 structure onto IPv6 may not be satisfactory, depending on your original IPv4 addressing structure and the requirements of network management, security architecture, and network programming. A thoughtful evaluation of what information is important for an address to encode or contain; of what political, functional, and directional structure exists on the campus; and of how addresses will be managed will result in an address scheme that will limit the need for reengineering later and will support network and security engineering and troubleshooting in an efficient, effective manner.

These steps, however, are only the beginning. Although much current hardware and software will run IPv6, most campus infrastructure will contain some devices that are not yet fully

[13] As of August 2013, classes on IPv6 were being offered at a number of community colleges, including Salt Lake Community College, City College of San Francisco, and Gulf Coast Community College, and at four-year universities such as University of Hawaii, University of Pennsylvania, and DePaul University.

compatible. Supporting IPv6 requires that software (including operating systems) and firmware running on network hardware (such as switches, routers, and security devices) be able to understand both IPv4 and IPv6. Internet-connected devices are "smart" in that they are aware of the infrastructure *and* actively participate in its operation. In the case of IPv6, the connecting infrastructure (such as routers and switches) needs to be sufficiently current to use the new standard, but so do all of the connected devices—including end-user computers. Servers must be migrated, and any software that is network-aware and may be looking at or storing an IP address (most web-based apps and many databases fall into this category) must be carefully examined and made compatible with the IPv6 address format. While web servers and client machines can be converted to dual stack with relative ease, the conversion of security devices, mail gateways/servers, and file servers is more complex. These devices and services have significantly greater reliability and performance requirements, and their migration will require more careful planning and additional time.

4.2. Financial Considerations

Most campuses are not in a position to replace all of their incompatible networking and security equipment and network-connected devices in one fell swoop. Older equipment may not be capable of running versions of software that provide the best, most consistent, support for IPv6. Institutions that start a methodical transition today can use IPv6 support as an evaluative criterion for upgrade and replacement selections. This places the IPv6 transition process within the normal refresh and upgrade cycle for the typical campus and will move the institution along its IPv6 migration path with only modest increases in cost. Without such planning, a transition will eventually be forced upon the institution, and this is much more likely to result in the need to replace existing equipment prior to the end of its useful life.

As mentioned above (see Section 3.2, Government and Granting Agency Progress), the availability of IPv6 will also increase the competitiveness of universities for federal grants and contracts. Delaying implementation will increase the cost of the transition, and the longer your campus waits, the more it will cost.

4.3. IPv6 and Collaboration

IPv6 will increasingly be needed for cross-institutional collaborative research and shared access to data repositories. Outreach to potential students and faculty—especially in countries that have already migrated to IPv6—will be improved by IPv6. With Europe, Asia, and Latin America moving towards IPv6 faster than the U.S., certain collaborations with remote faculty may be difficult or impossible from IPv4-only campuses. Indeed, some larger experiments may be IPv6 only. Until your campus supports IPv6 on its key sites, traffic from IPv6-only or IPv6-

preferred locations must travel through a 6-to-4 tunnel to reach your campus and departmental websites. This tunnel can be a bottleneck, slowing the user experience, and some links may not work. Also, if the tunnel encountered is outside the IPv6 user's institution, there is a question of privacy and the possibility for packets and information to be stolen as part of an attack.

In addition—as was noted in Section 3, Why Now—visiting scholars, researchers, and other visitors from IPv6-preferred institutions are already arriving on your campuses with IPv6-only configured laptops and mobile devices they would like to use (especially wireless) on campus.

4.4. IPv6 and Applications

Additional IPv6 address space, without the problems introduced by NAT and non-globally-routed IPv4 address space, will facilitate increasingly important services such as security cameras, keyless locks, wireless environmental sensors, and other applications. IPv6 address space could also be used right now for supercomputers, clusters, and other applications where hundreds or thousands of unique addresses are needed. This preserves outward-facing IPv4 address space and will make the transition years—during which many systems and services will require dual stacks (both IPv4 and IPv6 addresses)—easier.

4.5. Security

As noted, IPv6 is already on your campus. If you are not prepared to support it, you are currently in a risky state and have a responsibility to protect your campus. Moving toward IPv6 helps protect your campus' networked infrastructure and makes it more visible and controlled.

You may have heard that Microsoft's decision to sunset their Teredo service will reduce the risk of tunneled IPv6.[14] However, other tunneling technologies remain, invisible to campus security. Moreover, if you have customers using a tunneling technology to reach IPv6 resources, you have customers who want to use IPv6. It is far better to provide them with visible, native IPv6 service than allow shadow services to travel unseen through your network infrastructure.

Network Address Translation (NAT) has long been used as a mechanism to map private IPv4 addresses to public addresses that are then globally routable on the Internet. This often takes the form of "many-to-one," which means that multiple private IPv4 addresses share a single public IPv4 address. While it is not ideal, this does make NAT a technique for extending IPv4 address space.

However, many-to-one NAT adds operational complexity to the network. It makes identification and treatment of individual machines more difficult; it may not be practical to

[14] Jeff Burdette, "Microsoft Testing Sunset of Teredo," *IPv6 Knowledge*, July 19, 2013, accessed September 26, 2013, http://knowipv6.digitalelement.com/?p=101.

quickly and easily cordon off a compromised machine when it is just one host in a many-to-one block of hosts. Security groups may be faced with the option of cutting off the entire group of machines if they cannot quickly identify a single, specific culprit. For similar reasons, many-to-one NAT makes it hard to configure firewall exceptions for single machines, often leading to imposing restrictions or granting permissions to a large, unmonitored space. The way NAT operates, computers behind a many-to-one translation will not have the same direct access at the network level that non-translated hosts have. They are, in effect, less visible to the Internet and this represents a problem for a number of frequently used services that expect clients to be globally accessible at the network level. Services such as the Session Initiation Protocol (SIP) used to establish Voice over Internet Protocol (VoIP) communications simply will not work as well with machines that are translated using many-to-one NAT because they are designed around end-to-end connectivity that NAT breaks. Staff can spend valuable troubleshooting time in a futile attempt to shoehorn the full design capabilities of these services into the more restricted environment that many-to-one NAT provides.

The obfuscation of multiple hosts behind a single public address *does* have the side effect of making them somewhat more secure with respect to the rest of the network. This is simply a byproduct of how traffic between Internet devices works and was never the intended purpose or a specific design feature of NAT. Unfortunately, there is a wide belief that many-to-one NAT functions as a security mechanism in and of itself, and this has resulted in a tendency to ignore NAT areas. Regardless, compromised devices within a common address block can still infect other devices if not adequately monitored. This is true in general, but the increasing presence of "bring your own device" (BYOD) machines typically within effectively unmonitored NAT address blocks has only made this worse.

Use of many-to-one NAT does extend the IPv4 address space under the right conditions and it is a technique that network administrators will likely continue to use, but it brings with it a series of other operating problems and challenges that greatly diminish its efficacy. It is no panacea.

Knowing this, transitioning to properly configured IPv6 is actually a *conservative* move; ignoring IPv6, or transitioning to the increasingly normative IPv4-behind-NAT architecture, is actually more *radical and dangerous*. In the long run, IPv6 is likely more secure, definitely more sustainable, and much more organized and less messy than IPv4-behind-NAT. (See more about security and IPv6 in Section 5.3, Developing a Security Plan.)

5. Planning the Transition

Understanding why a campus should make the transition to IPv6 now is often overwhelmed by the practicalities of making this change. This section includes some recommended steps to

consider, including assessing and addressing structural preparedness issues, as well as developing technical, security, and rollout plans.

Part of any good transition plan will include ensuring adequate technical training for network, security, system-administration, and software-development staff. Because IPv6 includes protocol changes and is not a simple address translation, new approaches will be necessary for network and security engineering and system management. Software developers need to be aware of IPv6 address format and client addressing behaviors (for example, that a client may not always announce the same IPv6 address for itself) and must modify their code to avoid another compatibility problem of Y2K proportions.

It will also be necessary to identify and purchase or develop tools, both to serve technical staff in supporting IPv6 and to serve the campus community. Where possible, we have noted such tools in the sections below.

5.1. Structural Preparedness

Development of an overall IPv6 technical, security, and rollout plan represents an opportunity for institutions to rethink and modernize their network operations as a whole. Without this, it may be impractical to execute a structured rollout that lets departments migrate on their own timetable and this, in turn, can lead to a loss of early adopters and leave the project floundering. Although many institutions will remain functionally fairly distributed, it is still possible to move forward in a more coherent fashion.

As part of a structured rollout, the institution can have as its first aim to make IPv6 *possible* anywhere on campus. This is not the same as *enabling* it everywhere, but rather it is an exercise in identifying those infrastructure requirements needed to support IPv6 and enabling them at the wide area network (WAN), core, and distribution layers. These portions of the infrastructure are often centrally managed even in distributed environments and so this portion of the migration is typically manageable. Once complete, the central infrastructure will be prepared to enable individual migrations for early adopters or key locations. These early adopters are likely to be spurred to migrate by external requirements (e.g., access to a key resource, collaboration, etc.). For a broader, more general rollout, thought may be given to developing a campus project around a modest but useful service or capability.

5.2. Developing a Technical Plan

A technical plan will cover issues such as auditing existing capabilities, developing an addressing scheme, and designing network infrastructure for existing (IPv4), transitional (both

IPv4 and IPv6), and future (IPv6 with IPv4 support) states. Such a plan may need to include financial projections if not all network hardware is IPv6-ready.

The technical plan should also identify the tools needed to support infrastructure management processes. For example, can existing host registration systems handle IPv6? Will servers and clients be handled differently from how they were under IPv4? What tools/IPAM server will be used?

Technical planning will include (but not be limited to):

- Best-practice network segmentation
- Plans for WAN, core, distribution, and access
- Plans for how client machines will get their addresses
- Private address space in IPv6 vs. how virtual local area networks (VLANs) are currently used
- Use of IPv6 improved services
- Consideration of how Multiprotocol Label Switching (MPLS) and other network protocols that support mobility and redundancy may be used to aid transition
- Flexible addressing plan

Finally, the technical plan should discuss how the various aspects of supporting dual stack will be accomplished, including its support in network equipment. Eventually only IPv6 will be needed, but that point may be 10 or more years in the future, so while details of the IPv4 phase-out may not be required at this time, attention must be devoted to how IPv6 and IPv4 will coexist within the infrastructure.

5.3. Developing a Security Plan

The security plan should be coupled to the technical plan and address the education of campus IT support professionals about the different approaches to security necessary under IPv6. For example, under IPv4, system administrators often filtered access by IP address, using tools such as iptables. Under IPv6 client machines may not have fixed IP addresses, so filtering can generally be done only at the subnet level when it is done at all. The use of security filters in place of proper authorization tools must be deprecated. Financial planning should include projections for next-generation IPv6-aware firewalls, intrusion-detection systems, proxies, VPNs, DNS security tools, ACL tools, etc. Since not all existing network security hardware is likely to be IPv6 compliant at the outset, security planning should include the development of best practices for security with or without IPv6-compliant firewall hardware, including interim solutions (e.g., ACLs). If firewalls are in place, the plan should include making sure that they do

not block IPv6 by default. If they pass IPv6 transparently, the security plan should include a strategy for ensuring this does not represent a security hole.

Contemporary client operating systems typically have support for IPv6 but may not have protections enabled by default. An outreach and education exercise is strongly advised to educate both staff and the community on the need for and the approach to enabling the appropriate protections. There are potential IPv6 exploits even on traditional IPv4 networks, and it is highly recommended that client device protections be enabled before IPv6 is deployed broadly. Internet Protocol Security (IPSec) is a technology protocol suite for providing encryption features to networking. IPSec support is required for IPv6 implementations. It is somewhat better integrated in IPv6 than in IPv4, but making use of it still requires overcoming significant complexity hurdles with respect to configuration and operation. Further, it is not capable of replacing all of the security features and capabilities that are implemented at higher levels in the networking stack and that have become prominent techniques in existing IPv4 implementations. Any comprehensive security plan developed for the transition of IPv4 to IPv6 must weigh the realized security benefits of all current and proposed security elements and provide mechanisms to transition between them as needed. Security planning must ensure that transition activities—and the associated changes to existing infrastructure that they will entail—do not undermine or circumvent existing IPv4 security mechanisms during the overlap period (when both IPv4 and IPv6 are operating concurrently) or eliminate them at the end of transition without providing realistic alternatives.

Because one-to-one IP-to-hardware address mapping is not preserved in IPv6, tools needed to support the security effort include those that will map an IPv6 address and a timestamp to the hardware address or other identifier of the machine using that address at that time. It may be necessary to change logging behavior and/or make logs more centrally viewable for troubleshooting. Additional hardware and software may be needed.

5.4. Developing a Rollout Plan

The rollout plan should include a deployment schedule, ideally from low- to high-hanging fruit. Assessment of feasibility may include customer need, technical readiness, and availability of funding. Think about your migration priorities over the first few years. When looking at early targets for migration, consider, for example, whether VoIP or wireless will make use of IPv6. VoIP will often benefit from the better integrated quality of service support built into the IPv6 standard, and wireless conversion may benefit the greatest number of interested users quickly. At the same time, some applications (e.g., network management software, firewalls, IDS, and e-mail) may be more technically challenging, require more vendor maturity, necessitate community effort (e.g., e-mail blacklists), or have a low tolerance for risk (e-mail is a good

example of this). Others (UNIX servers in labs, web servers, and wireless clients, for example) may be easy to deploy on IPv6 and create value for the university quickly.

One possible order of deployment is suggested below. You should adjust this according to your own local needs and campus risk tolerance—for example, some may opt to deploy IPv6 within the campus network first, in order to develop operational familiarity and knowledge before deploying it on a service with high visibility (such as the campus website). One approach is to test out low-profile services with static "quad A" records and then, as experience is gained, enable IPv6 on higher-profile services.

Pre-Deployment:

1. Obtain IPv6 address space from ARIN (The American Registry of Internet Numbers) or from your Internet Service Provider (ISP).
2. Ensure IPv6 transit and peering are available.
3. Ensure campus network can support native IPv6 throughout.
4. Ensure that DNS supports IPv6 resource records and that DNS servers are accessible via SLAAC, DHCPv6, etc.
5. Provide training to campus network engineers.

Deployment:

1. Deploy on campus gateway and begin announcing routes.
2. Deploy dual stack on primary campus web servers (outward facing).
3. Deploy on campus backbone.
4. Deploy on campus distribution.
5. Make training and tools available to help desks, sysadmins, programmers etc.
6. Deploy on research VLANs for servers on request/by need.
7. Deploy on campus wireless.
8. Deploy on campus VoIP infrastructure.
9. Deploy on wired client VLANs with appropriate protection (may be ACLs at first).
10. Deploy on campus network services such as e-mail and Active Directory.
11. Begin to deploy routinely to new VLANs.
12. Enable rest of existing VLANs.

In addition to a deployment schedule, the rollout plan should address how, when, and with what support departments on your campus will deploy dual stack. It may include building general awareness, training for department system administrators and programmers, deploying tools to self-assess readiness, providing documentation for common server environments, and other support mechanisms. Because many specialized research devices both are network-aware and will not support IPv6 in the short term (one example is freezer monitors, critical to many

research endeavors), deployment should be encouraged but not forced at the expense of research. In many areas, however, risk is reduced by official IPv6 deployment, and this should be stressed.

Remember that help desks and desktop support staff will need tools for troubleshooting and support, especially for the wireless network. They will need a way to identify users of IPv6 addresses within existing troubleshooting tools, which may expect hardware addresses. Some of the tools deployed for the security team may be useful here as well.[15]

6. Conclusion

The current state of IPv4, with its address exhaustion and less integrated service support, threatens to sharply inhibit both the current operation of the Internet as well as its continued evolution. When we look at how the Internet can grow, there are two w—one of which relies on bandaging problems and propping up limitations under IPv4, and the other that moves in the direction of IPv6. In short, our campuses are migrating from classical IPv4 to *something*, and a transition to IPv6—with its conceptual similarity to the familiar IPv4 architecture and its plentiful globally routable addresses—is a wiser choice than a transition to a prevalently NAT-based IPv4 architecture.

In the past, there were sound reasons why campuses and other organizations delayed deployment of IPv6. Those reasons, however, are no longer valid, and the time to begin planning a migration to IPv6 is now. There are now distinct gains (financial, reputational, accessibility, etc.) to be had by making your campus reachable via IPv6. Perhaps more importantly, there are existing weaknesses (network infrastructure visibility, security, etc.) to maintaining an IPv4-only environment that are putting your campus at risk. An ordered, thoughtful approach to IPv6 planning will secure campus commitment and put your campus on the road to a less risky, methodical path to IPv6 that will allow you to continue to serve the needs of research, instruction, and campus business effectively; secure your cyberinfrastructure; attract research grants; and position your campus for the future.

We are at a crossroads. It is time for the higher education community to show leadership and take the necessary and proactive steps to bring our communities and the Internet into a healthy, high-performing future.

[15] See "Comparison of Network Monitoring Systems" for more information on systems that support IPv6 at http://en.wikipedia.org/wiki/Comparison_of_network_monitoring_systems. The types of tools that security teams use include those that use the NetFlow, SFlow, JFlow (sampling technologies), loggers like ArcSight, packet sniffers, etc.

Appendix A: Resources

Articles, Presentations and White Papers

Awduche, Daniel O. *Benefits of IPv6 for Enterprises*. Verizon: 2012. Accessed September 30, 2013. http://www.verizonenterprise.com/resources/whitepapers/wp_benefits-of-ipv6-for-enterprises_en_xg.pdf.

Broersma, Ron. "Developing an IPv6 Addressing Plan: Guidelines, Rules, Best Practice." *HPC DREN* November 16, 2011. Accessed September 27, 2013. http://www.v6.dren.net/AddressingPlans.pdf.

Curan, John, and Phillip Deneault. "Migration to IPv6—Has Tomorrow Finally Arrived?" EDUCAUSE Live! Webinar, March 9, 2011. Accessed September 30, 2013. http://www.educause.edu/library/resources/migration-ipv6—has-tomorrow-finally-arrived.

Curran, John. "IPv4 Depletion and Migration to IPv6." EDUCAUSE Live! Webinar, June 18, 2008. Accessed September 30, 2013. http://www.educause.edu/library/resources/ipv4-depletion-and-migration-ipv6.

Deneault, Phillip. "Starting Over from the Top: Campus IPv6 Deployment and Security." Presentation at the Security 2010 Conference, Atlanta, Georgia, April 13, 2010. Accessed September 30, 2013. http://www.educause.edu/events/security-professionals-conference/2010/starting-over-top-campus-ipv6-deployment-and-security.

EDUCAUSE. "7 Things You Should Know About IPv6." *7 Things You Should Know*: April 2011. Accessed September 30, 2013. http://www.educause.edu/library/resources/7-things-you-should-know-about-ipv6.

Frankel, Sheila, Richard Graveman, John Pearce, and Mark Rooks. *Guidelines for the Secure Deployment of IPv6: Recommendations of the National Institute of Standards and Technology*. Gaithersburg, MD: NIST Special Publication 800-119, December 2010. Accessed September 30, 2013. http://csrc.nist.gov/publications/nistpubs/800-119/sp800-119.pdf.

Heder, Brian. "IPv6 Transition Framework for the Enterprise." *Network World*, June 9, 2011. Accessed September 30, 2013. http://www.networkworld.com/news/2011/060911-ipv6-transition.html

Huston, Geoff. "A Primer on IPv4, IPv6 and Transition." *The ISP Column*, April 2013. Accessed September 30, 2013. http://www.potaroo.net/ispcol/2013-04/primer.html.

Kosiur, Dave. "IPv6: Years in the Making, Years Before Full Adoption Key Findings." EDUCAUSE Center for Applied Research, July 16, 2003. Accessed September 30, 2013. http://www.educause.edu/library/resources/ipv6-years-making-years-full-adoption.

Levy, Martin. "Six Benefits of IPv6." *Networking Computing*, June 8, 2011. Accessed September 30, 2013. http://www.networkcomputing.com/ipv6/six-benefits-of-ipv6/230500009.

Pittman, Elaine. "Why the Internet of Things Needs IPv6." *Government Technology*, April 18, 2013. Accessed September 30, 2013. http://www.govtech.com/policy-management/Why-the-Internet-of-Things-Needs-IPv6.html.

Roberts, Phil. "IPv6 Deployment Hits 2%, Keeps Growing." *Tech Matters* (blog), September 24, 2013. Accessed September 25, 2013. http://www.internetsociety.org/blog/2013/09/ipv6-deployment-hits-2-keeps-growing.

Soucy, Ray. "IPv6: Time to Move." *InformationWeek,* February 2010. Accessed September 30, 2013. http://soucy.org/ipv6/doc/S210210_InformationWeek_SOUCY.pdf.

St. Sauver, Joe. "MAAWG IPv6 Training for Senders and Others." Messaging Anti-Abuse Working Group, Crystal City VA, October 4–8, 2010. Accessed September 30, 2013. http://pages.uoregon.edu/joe/maawg-senders-ipv6-training/maawg-senders-ipv6-training.pdf.

Strategy and Planning Committee, Federal Chief Information Officers Council. *Planning Guide/Roadmap Toward IPv6 Adoption within the U.S. Government (Version 2.0)*. CIO Council, July 2012. Accessed September 30, 2013. http://cio.gov/wp-content/uploads/downloads/2012/09/2012_IPv6_Roadmap_FINAL_20120712.pdf.

York, Dan. "2% of All Traffic to Google Now Over IPv6! (Doubling in Past Year)." *CircleID*, September 24, 2013. Accessed September 30, 2013. http://www.circleid.com/posts/20130924_2_percent_to_google_now_over_ipv6_doubling_in_past_year/.

Deployment Case Studies

Bauer, J., and S. Huque. "IPv6 Deployment at the University of Pennsylvania." Presentation at the EDUCAUSE Mid-Atlantic Regional Conference, January 8, 2009, http://www.educause.edu/midatlantic-regional-conference/2009/ipv6-deployment-university-pennsylvania.

Benoit, S. "0 to IPv6 in 3 months." Presentation at NANOG52, June 2011, http://www.nanog.org/meetings/nanog52/presentations/Tuesday/Benoit-0toIPv6ACustomersViewV2.pdf.

Whinery, A. "Hawaii IPv6 Deployment Experiences." PTC 2010, January 17, 2010, http://www.ptc.org/ptc10/program/images/papers/slides/Slides_Alan Whinery_RT1.pdf.

For More Information

- ARIN: American Registry for Internet Numbers
  - ARIN IPv6 Wiki: Get IPv6
    https://getipv6.info
  - IPv6 Info Center
    https://www.arin.net/knowledge/ipv6_info_center.html
    This site includes a number of resources to assistance with IPv6 transition.
- Internet Society
  - IPv6
    http://www.internetsociety.org/what-we-do/internet-technology-matters/ipv6,
    Includes links to guides, publications and workgroups.
  - Deploy360 Programme
    http://www.internetsociety.org/deploy360/ipv6/
- Internet2 IPv6 Working Group
  http://ipv6.internet2.edu/
- IPv6 Forum: Web Links
  http://www.ipv6forum.com/modules.php?op=modload&name=Web_Links&file=index
  Includes information on applications, implementations, resources, deployments, and networks.
- IPv6 Knowledge Base: General Information
  http://www.hpc.mil/cms2/index.php/ipv6-knowledge-base-general-info
  This page from the U.S. Department of Defense High Performance Computing Modernization Program (DoD HPC) "provides a wide variety of information about planning for IPv6 deployment and installing, configuring, securing, and testing IPv6 products, services, and networks."
- IPv6 Testing Consortium
  https://www.iol.unh.edu/services/testing/ipv6/
  From the University of New Hampshire InterOperability Laboratory, the "IPv6 Consortium is focused on offering testing services that reduce the time to market for our participants, and accelerate the adoption of IPv6 technology."
- World IPv6 Launch
  - Measurements
    http://www.worldipv6launch.org/measurements/
    This page shows IPv6 global deployment on networks.
  - Participants
    http://www.worldipv6launch.org/participants/.

Appendix B: ECAR-CCI Working Group Members

The following were members of the EDUCAUSE Center for Analysis and Research Campus Cyberinfrastructure Working Group at the time this report was published.

David Ackerman
Associate Vice President, .edu Services
New York University

Vijay K. Agarwala
Senior Director, Research Computing and Cyberinfrastructure
The Pennsylvania State University

Guy T. Almes (Chair)
Director, Academy for Advanced Telecommunications and Learning Technologies
Texas A&M University

Celeste Anderson
Director, External Networking Group Information Technology Services
University of Southern California

Greg Anderson
Retired Senior Director
University of Chicago

Charles R. Bartel
Director, Global IT Services, Computing Services
Carnegie Mellon University

Asbed Bedrossian
Director, Enterprise Applications and Operations
University of Southern California

Rajendra Bose
Manager, Research Computing Services
Columbia University

James R. Bottum
Vice Provost and CIO
Clemson University

Perry Brunelli
Director of Network Services
University of Wisconsin–Madison

Duncan Buell
Professor, Computer Science
University of South Carolina

Alan Crosswell
Associate Vice President and Chief Technologist
Columbia University

James Cuff
Director of Research Computing and Chief Technology Architect
Harvard University

James Davis
Chief Information Officer
Iowa State University of Science and Technology

Thomas (Ted) Dodds
CIO and Vice President for Information Technologies
Cornell University

Edward J. Evans
Director of Software Services
Purdue University

Michael Fary
Enterprise Data Architect
University of Chicago

Steve Fleagle
CIO and Associate Vice President
The University of Iowa

Jill B. Gemmill
Executive Director, Cyberinfrastructure Technology Integration
Clemson University

Rick Golden
Assistant CIO
University of Nebraska

Richard Greenfield
Strategic Research Analyst
University of Alaska

Ted Hanss
Chief Information Officer, Medical School
University of Michigan–Ann Arbor

Thomas Hauser
Director of Research Computing
University of Colorado Boulder

Curtis W. Hillegas
Director of Research Computing
Princeton University

Bill Hogue
Vice President for Information Technology and CIO
University of South Carolina

Sally Jackson
Professor of Communication
University of Illinois at Urbana-Champaign

Kurt J. Jeschke
Manager, Network Engineering
The Pennsylvania State University

James A. Jokl
Associate Vice President and Chief Enterprise Architect
University of Virginia

Paul Killey
Executive Director of IT, CoE
University of Michigan–Ann Arbor

Richard Knepper
Manager, Campus Bridging and Research Infrastructure
Indiana University Bloomington

Boyd Knosp
Associate Dean for Information Technology
The University of Iowa

John E. Kolb
Vice President and CIO
Rensselaer Polytechnic Institute

Steve Krogull
Director, Systems Engineering & Operations
University of Wisconsin–Madison

Timothy Lance
President and Board Chair
NYSERNet, Inc.

Betty Leydon
Vice President for Information Technology and CIO
Princeton University

Clifford Lynch
Executive Director
Coalition for Networked Information

Donald (Rick) F. McMullen
Cyberinfrastructure Strategist
Great Plains Network

Gregory E. Monaco
Director for Research & Cyberinfrastructure Initiatives/Great Plains Network
Kansas State University

Randy Monroe
Associate Director, Network Services
Carnegie Mellon University

Michael R. Mundrane
Vice Provost and Chief Information Officer
University of Connecticut

Jon L. Oliver
Assistant Dean and Director of IT
Rutgers, The State University of New Jersey

Kim Owen
Advanced Applications Outreach
North Dakota State University

James Pepin
Chief Technology Officer
Clemson University

Valerie E. Polichar
Infrastructure Analyst and Technical Projects Coordinator
University of California, San Diego

Lynn Rohrs
Director, eSystems and Research Services
New York University

Brian Seiler
Network Analyst
University of Illinois at Urbana-Champaign

Michael K. Smeltzer
Director of Networking
University of Illinois at Urbana-Champaign

Oren Sreebny
Senior Director, Emerging Technologies and Communications
University of Chicago

David Stack
Chief Operating Officer and Deputy CIO
University of Wisconsin–Milwaukee

Charles Thompson
Assistant Dean and Chief Information Officer
University of Illinois at Urbana-Champaign

Scott Valcourt
Director, Strategic Technology
University of New Hampshire

Shelton Waggener
Senior Vice President
Internet2

David H. Walker
Consultant
Pleasanton, California

Harry Williams
Chief Technology Officer
Marist College

Tom Zeller
Senior Technology Analyst
Indiana University Bloomington