Deploying an IPv4-IPv6 Dual Stack Network in an Academic Computing Infrastructure

Shaun M. Lynch, Ph.D.
Department of Mathematics and Computer Science
University of Wisconsin-Superior
Superior, WI 54880
slynch@uwsuper.edu

Abstract

Adoption of IPv6 has been an agonizingly slow process even though IPv4 address exhaustion constrains deployment in a time of expanding network utilization. Academics and academic programs charged to educate and train the next generation of technology professionals are acutely aware of the challenges associated with introducing new technologies into existing curriculums. IPv6 technology is one of those areas that students in technology related curriculums will need to understand going forward. In this light, the author pursued a plan to transition the computing infrastructure of the Department of Mathematics and Computer Science at the University of Wisconsin-Superior to a full-fledged IPv4-IPv6 dual-stack network. The goal was to create a network that was not only functional but capable of providing faculty and students with the tools necessary to learn and explore IPv6 technologies. This paper presents the design, technologies, and processes used to achieve this goal.

1 Introduction

The adoption of IPv6 in many regards has been an agonizingly slow process. At the start of 2015, IPv6 connectivity remains in single digits as a percentage of overall network traffic despite being available since the turn of the millennium. However, IPv4 address exhaustion is an irrefutable constraint that limits the deployment in a time of expanding network utilization due to increased use and deployment of mobile technologies and the potential explosion of connected devices and embedded systems stimulated by the Internet of Things (IoT) movement [1, 2, 3].

Academics and academic programs charged to educate and train the next generation of technology professionals are acutely aware of the challenges associated with introducing new technologies into existing curriculums. Adding new content to a curriculum generally means reducing content elsewhere. Faculty and academic programs take a calculated risk by changing programs to provide students with the skills they will need to adapt to a volatile technology landscape upon graduation. Nonetheless, maintaining curriculum so that it contains material relevant to student success is a necessary activity and requires the appropriate tools and technologies to that subject matter can be taught effectively.

IPv6 technology is one of those areas that students in technology related curriculums will need to understand going forward. Although adoption has been slow, there are no other technologies on the horizon to satisfy networking challenges expected in the future. In this light, the author pursued a plan to transition the computing infrastructure of the Department of Mathematics and Computer Science at the University of Wisconsin-Superior to a full-fledged IPv4-IPv6 dual-stack network. The goal was to create a network that was not only functional but capable of providing faculty and students with the tools necessary to learn and explore IPv6 technologies thus providing a mechanism to extend textbook concepts to the world of practice.

This paper presents the design, technologies, and processes used to achieve this goal. Discussion begins with a brief background that includes motivation for the project, computing environment, and academic value of incorporating IPv6 technology. Next, design considerations including connectivity, perimeter control, and address management are explored. Emphasis is place on viable alternatives suited to an academic setting. This is followed by an in-depth look at the deployment process where the current architecture, implementation, and system migration are presented. Particular attention is paid to the factors that influenced the selection of design alternatives. Finally, the paper closes with comments and observation about the process.

2 Background

Transitioning to IPv6 is a nontrivial challenge since academic computing infrastructures often reside within a larger campus network and must coexist without disrupting ongoing operating activities of the university. University technology operations generally

emphasize production and sustainability needed to serve the entire university community. Their operations are subject to a wide range of policy and regulatory issues that must be complied with.

Academic departments on the other hand, must equip students with knowledge and skills they will need in the future requiring a vision that looks toward the horizon. In addition, academic operations are free from many of the policy and regulatory issues enabling an environment where innovation can occur. Regardless, the challenge is being able to respect differences and make these two systems work together.

2.1 Motivation

The decision to deploy an IPv4-IPv6 dual-stack network infrastructure was driven by two motivating factors. First, implementation enables faculty and staff to gain in-depth familiarity with IPv6 networking technologies and work through the challenge of deploying a dual-stack network. It is through implementation that educators begin to understand the nuances and caveats of the technology and develop the deep-reasoned knowledge necessary to teach subject matter and guide students through the learning process. It also serves to stimulate new ideas and approaches to communicate the concepts to those engaged in the subject matter for the first time.

The second motivating factor is to create an IPv6 infrastructure that provides students with the ability to practice textbook concepts using actual hands-on projects. It is the author's experience that teaching from a book is easy. Creating meaningful projects that engage student interests and provides a physical outlet that reinforces learned concepts is hard. This process is made even more difficult in a resource constrained environment where the latest technology is often out of reach and beyond the budget of most academic institutions.

2.2 Academic Computing Environment

The department manages and maintains a dedicated computing environment used to advance the department's teaching mission. Computing resources are used extensively in classes across academic programs ranging from entry level coursework that fulfills the University's mathematics general education requirement to advanced electives that meet curricular objectives. Students, faculty, and academic staff enjoy a wide array of applications suited for general productivity, programming, software development, analytic modeling, device emulation, network security, multimedia, and web authoring.

Two advanced computing laboratories are the most visible aspect of the department's computing environment. Each advanced computing lab contains 25 computers with one set aside as a teacher workstation that drives the rooms projection and multimedia system. All workstations share a common configuration based on a core image containing Windows 7 x64 Enterprise as the operating system along with a productivity suite, integrated development tools, and a variety of utility applications. Specialty applications

are licensed and installed on all computers within a given lab to ensure consistency and allow full utilization of lab workstations.

The department also maintains a "sandbox" for special projects, experimentation, prototyping, and demonstration that includes a variety of legacy computers, one-of-a-kind systems, components, networking gear, and mobile racks. The sandbox provides a setting where students and faculty can explore without tying up shared resources needed for instruction or disrupting the production systems.

The server and network resources are the cornerstone of the department's computing environment and provide essential services to host applications, manage computing assets, and deploy security policies. The physical system consists of physical servers that host domain and internet protocol management services, a two node cluster along with an iSCSI storage array that provides a virtualization platform for a variety of application servers, a dedicated server that hosts virtual desktops, and other specialty servers that provide shared storage.

2.3 Academic Value

There is little argument that IPv6 is primarily a technical solution aimed to expand network capabilities beyond those currently provided by IPv4. Few people consider how their browsers works or how their mobile device allows people to connect, the technology just works. The value brought by enabling IPv6 technologies in an academic program is that it exposes what is hidden, transparent, and often taken for granted to students pursuing technology related curriculums.

An IPv6 enabled network offers the ability to demonstrate a wide array of technologies that simply would be out of reach using an IPv4 only network. A partial list includes:

- Transitional mechanisms such as 6in4, 6rd, NAT64, and DNS64
- IPv6 addressing such as link-local, unique local, and global addresses
- IPv6 address formats such as unicast and multicast
- Private network architectures using non-routable addresses
- Address management mechanisms such as router advertising and DHCPv6
- Stateless and stateful auto configuration
- Network protocols such ICMPv6

In combination with desktop virtualization, devices can be built using open source and proprietary operating systems and applications allowing students to directly interact with IPv6 technologies.

3 Design Considerations

Prior to deploying IPv6 in a dual-stack configuration, it is important to weigh options and the evaluate the tradeoff they may incur. Fortunately, academic computing environments

emphasize innovation over production thus tend to be more flexible in the technologies deployed. Based on the author's experience, the most important options to consider before deploying IPv6 in an academic computing infrastructure include:

- How to secure an IPv6 connection,
- How to manage the IPv6 connection at the network perimeter, and
- How to manage and assign IPv6 addresses within the department's intranet

Design alternatives are not necessarily mutually exclusive; therefore, multiple options may be available. In some cases, it is possible to create network enclaves within the academic infrastructure to test various configurations before deployment.

3.1 IPv6 Connectivity

The simplest way to attain an IPv6 connection is to use an address space provided by the host institution or internet service provider. If the host organization has not moved to an IPv6 infrastructure, then the host's perimeter and intermediate networking devices may need to be configured to allow and route IPv6 traffic to the local network. However, this approach may not always be possible and depends on the organization's network policies, intervening network configuration, and relationship with network administrators. In the absence of a native IPv6 connection, IPv6 transitional technologies are available that can provide connectivity over existing IPv4 only networks.

IPv6 transitional technologies work by tunneling IPv6 traffic between endpoints that span an IPv4 network. 6rd and 6in4 are the two most common protocols suitable for a dual-stack network. 6rd is a derivative of 6to4 (recommended for deprecation by IETF [4]) and may be offered as an optional service to ISP subscribers. In this particular scheme, the ISP operates a 6rd border relay using its own IPv6 network prefix to provide remote endpoint services. Most residential customers receive an IPv6 address with a 64 network prefix; however, non-residential service may receive an address that allows up to a 16 bit (and possibly more) subnet range.

On the other hand, 6in4 offers IPv6 connectivity without need for dedicated ISP services. Instead, an independent tunnel broker provides a remote endpoint and an IPv6 address space to the subscriber. Hurricane Electric and SixXS are two popular 6in4 tunnel brokers in the United States and provide free services to individuals experimenting with and organizations making the transition to an IPv6 infrastructure. Both service providers offer IPv6 addresses with 64 and 48 bit network prefixes.

Fortunately, IPv6 connectivity options are not mutually exclusive alternatives. Since IPv6 transitional technologies employ a virtual private network to tunnel IPv6 traffic over IPv4 networks, they can be used concurrently with native IPv6. In fact, multiple transitional technologies can be used simultaneously provided network resources are available. This feature provides considerable flexibility to academic departments and creates the ability to provide enclaves of IPv6 connectivity within an overarching IPv4-IPv6 dual-stack network infrastructure.

3.2 Perimeter Gateway

Manufactures of commercial and consumer grade routing and firewall equipment often include native and transitional IPv6 capabilities that provide gateway services for a local network. It is important to note that not all devices are IPv6 capable or support IPv6 services to the same level as traditional IPv4; therefore, it is important to review device specifications before purchasing a dedicated appliance.

In some instances, enthusiasts can upgrade IPv4 consumer grade routers with new firmware such as OpenWRT, DD-WRT, or Tomato to provide native IPv6, 6rd, and 6in4 capabilities. However, this approach can be risky in that flashing unsupported firmware often voids warrantees and may lead to an irrecoverable failure. Nonetheless, this is a viable alternative for student and faculty interested in building a home lab or small setups that need IPv6 connectivity.

Alternatively, do-it-yourself network administrators can build an IPv6 gateway device from the scratch using an old computer or virtual machine. There are a number of full-featured, web-enabled applications built upon open source operating systems can be used such as pfSense, M0n0, Untangle, Vyatta, and more. In addition, an IPv6 gateway can be built directly from open source and proprietary operating systems allowing designers to pick and choose specific applications and features. This enables the ability to create highly customized systems tailored for specific tasks.

Most gateways are built using a dual-homed platform that allows traffic to be managed between public and private networks. Dedicated routing and firewall applications generally require a minimum of two network adapters particularly if network address translation (NAT) is being employed for IPv4 to amplify the address space. However, there may be instances where an IPv4 only gateway on the perimeter cannot be replaced with a device that supports IPv6 native and transitional technologies. In this situation, it is possible to create a single-homed, intranet residing IPv4-IPv6 gateway that compliments the functionality of the IPv4 perimeter device.

A single-homed, intranet residing IPv4-IPv6 gateway takes advantage of the ability of the network interface to support both IPv4 and IPv6 protocols independently. This dual-stack capability makes a single network adapter appear as if it were two separate devices each supporting a distinct protocol. A virtual network interface using simple internet transition (SIT) provides connectivity and allows the gateway to manage traffic between IPv4 and IPv6 protocols. Therefore, all private intranet IPv6 traffic tunnels through the IPv4 network to the remote endpoint connected to the public IPv6 internet. The biggest downside to this approach is that it requires customized low-level network settings of an operating system to achieve the desired functionality.

Firewall settings for an IPv6 enabled gateway are very similar to those used in IPv4 with the notable exception of ICMPv6. ICMPv6 extends the functionality of ICMPv4; however, certain ICMPv6 packets—such as error messages—must be allowed through the firewall for IPv6 to function properly. In addition, it is essential that encapsulated

IPv6 (Protocol 41) be allowed on the IPv4 network interface if the gateway is using the 6in4 transitional technology.

3.3 Address Management

Router advertisement serves as the primary mechanism to configure IP addresses in an IPv6 infrastructure. Routers multicast ICMP router discovery messages to host devices that include the router's address (default gateway) along with various flags that indicate how the host's network interface should be configured. There are several options that range from stateless auto-configuration to stateful configuration using DHCPv6. The most common configurations include:

- Stateless – Router provides network prefix enabling host to auto-configure network interface address
- Stateless with Recursive DNS – Router provides network prefix enabling host to auto-configure network interface address and name server addresses (Note: Not all operating systems respond equally to this configuration)
- DHCPv6 Stateless – Router provides network prefix enabling host to auto-configure network interface address while DHCPv6 provides name server addresses and other options
- DHCPv6 Stateful – DHCPv6 assigns host address, name server address, and other options

Dynamic address assignment in an IPv6 environment is managed using DHCPv6. Although there similarities with DHCPv4, DHCPv6 is distinct protocol and has significant differences from its predecessor (i.e. available options, multicasting instead of broadcasting, ports used, unique identifiers, *et cetera*). Nonetheless, it should be recognized that DHCPv6 only augments the functionality of the advertising service provided by the router and not a replacement.

Domain name services in an IPv6 environment are nearly identical as those found in IPv4 infrastructure. Only when naming a reverse lookup zone is a difference is observed. The naming convention for an IPv6 reverse lookup zone uses a reverse-ordered, period-expanded, hexadecimal byte form followed by suffix *ipv6.arpa*.

4 Deployment

Fortunately for computing infrastructure administrators, deploying an IPv4-IPv6 dual-stack network is an incremental process that builds upon an existing IPv4 architecture. However, it is important to ensure the underlying system is sound before superimposing a new IPv6 layer. Guiding principles the author used during this process include:

- Partition the network into manageable zones based on physical or logical boundaries
- Articulate incoming and outgoing traffic patterns for each zone

- Ensure network zones are aligned with logical structures and policy frameworks (e.g. Active Directory and group policy)
- Hide and protect internal assets using appropriate technology and configurations
- Expose only what you have to and cordoned off those systems with a well-defined interface
- Smaller is generally more manageable than bigger
- Model and prototype offline before implementing
- Automate where appropriate
- Be consistent
- Use the principle of low-hanging fruit—start with the easiest and work up to the most difficult
- Deploy, test, repeat

Most importantly, implement better solutions as they become available! It is likely that design flaws, deficiencies, and inconstancies will be discovered as the new technology is implemented and comes online.

4.1 Current Architecture

The department's computing infrastructure consists of a core server infrastructure, two advanced computing laboratories, hardware lab and technical support workstations, and sandbox systems. The server infrastructure hosts eight physical servers, 16 virtual servers, iSCSI storage array, and network attached storage. The two advanced computing lab each contain 24 workstations and one teacher workstation. In addition, approximately a dozen workstations are distributed between the hardware lab, technical support, and sandbox systems.

These systems reside on a private network established by the University's network administrator. The department's network perimeter is controlled by a Cisco ASA 5505 that provides IPv4 NAT services to the department's intranet. The gateway lies behind the University's perimeter firewall and has access to other campus computing resources. Configuration and management of the gateway is the sole responsibility of the campus network administrator to ensure compliance with University network policies.

All department systems share an eight-bit IPv4 host address space that provides 254 usable addresses. The address space is partitioned into regions to simplify management. A portion is allocated for networking infrastructure appliances and servers that use static addresses. A dynamic portion using reserved addresses is set aside for workstations in the advanced computing lab and networked printers to ensure connectivity. The remaining address space is made available for temporary connections made by user devices, virtual machine instances hosted on local workstations, and sandbox systems.

4.2 Implementation

To begin the process, the author met with the University's network administrator to explore available options to bring IPv6 connectivity to the department's private network. At the time, native IPv6 was not available and adopting the protocol was not being considered in the near future by the University's networking staff. This left the options of using a different ISP or a tunnel broker service to provide IPv6 connectivity. Further research revealed that the current gateway device used for the department's perimeter did not have the necessary settings to create a local 6in4 tunnel endpoint potentially eliminating the option to use a tunnel broker.

The author continued to investigate various solutions and found a series of commands that would enable the gateway to pass encapsulated IPv6 traffic to a selected device residing on the private network interface. This option was well received since it allowed for continued use of the gateway along with existing settings. After working with the University's network administrator and overcoming a few technical challenges, the author opted to provide IPv6 connectivity using a tunnel broker service.

Securing a tunnel broker service was the next step in the process. Hurricane Electric and SixXS offer free tunnel broker services that would adequately meet the academic needs of the department. The author selected Hurricane Electric as the service provider for the following reasons: 1) familiarity and prior use of the service, 2) simple registration process with quick turnaround, 3) numerous configuration examples, and 4) provisions for both 64 and 48 bit network prefixes suitable for creating subnets.

The next decision revolved around whether to use a virtual machine or physical computer to host the IPv4-IPv6 gateway. In this case, the author opted to deploy the system on a physical computer for two reasons: First, installing and configuring an operating system on actual hardware allows drivers to interact with and fully utilize available hardware. This impacts both performance and the ability to troubleshoot problems should they arise. Second, isolating a prototype system with an untested configuration allows the system to be physically quarantined if needed. This is particularly important when working with a public-facing network appliance subject to wilds of the internet.

Choosing an address management scheme for the dual-stack infrastructure was the final step, however, inherent differences between IPv4 and IPv6 addressing schemes forced several tradeoffs. To begin with, the original network design placed all computing assets on a single segment behind a gateway using NAT. Therefore, all hosts are assigned a private IPv4 address that should not be routable beyond the network perimeter. However, the IPv6 address range serviced by the gateway is globally routable. Ideally, the network should be partitioned into public and private zones where globally routable and unique local addresses can be applied. Despite this limitation, the author decided to maintain a single network prefix and block access to hosts within the perimeter using firewall rules.

Another complicating factor is that the department's academic network supports both provisioned and ad hoc dynamic connections. For instance, advanced computing lab workstations must have network connectivity during class periods. Conversely,

temporary connections such as laptops or locally created virtual machines compete for a limited pool of IPv4 addresses. The reservation and leasing system in DHCPv4 provides an effective mechanism to track and prioritize network connectivity. Therefore, the author decided to align the IPv6 address management mechanism to the existing IPv4 address management scheme as closely as possible.

These conditions led a tiered address management scheme based on system type. The principal schemes include:

- Server Infrastructure—Static host, name server, and default gateway addresses using global IPv6 addresses; host automatically registers address to DNS server when available otherwise host address is manually registered with DNS.
- Advanced Computing Lab Workstations and Printers—Stateful using DHCPv6 to assign host and name server addresses using global IPv6 address; individual host addresses are reserved in DHCPv6; default gateway assigned link local address of IPv4-IPv6 gateway; DHCPv6 registers host address with DNS.
- Temporary Connections—Stateful using DHCPv6 to assign host and name server addresses using global IPv6 address; default gateway assigned link local address of IPv4-IPv6 gateway; DHCPv6 registers host address with DNS.

Where possible, a symbolic representation of IPv4 host identifier was used for the IPv6 host identifier. For example, a server with the IPv4 address 192.51.100.25/24 (192.51.100.0/24 network prefix and 25-decimal host identifier) would have an IPv6 address of 2001:db8::25 (2001:db8::/64 network prefix and 25-hexadecimal host identifier). This approach greatly simplifies host identification whether an IPv4 or IPv6 address is being referenced.

4.3 Migrating Systems to IPv6

Fortunately, all networked devices and operating systems in the department's computing infrastructure were IPv4-IPv6 dual-stack enabled. This greatly simplified the migration process once the mechanisms used to manage IPv6 addresses were deployed. However, a couple of post deployment issue emerged that escaped immediate detection.

Shortly after the deployment, the author observed inconsistencies in DNS registration. Specifically, a device would only register either the assigned IPv4 or IPv6 address, but not both. The problem occurred intermittently, was not associated to any particular device, and happened only on systems assigned addresses by DHCP. The problem was eventually traced to a setting in DHCPv4 that forced a DNS update on older client systems that do not support name mapping (Client FQDN, Option 81) [5, 6]. The setting created a timing problem that conflicted with the dynamic DNS update setting in DHCPv6. Unselecting the option and using the default dynamic DNS update settings in both DHCPv4 and DHCPv6 resolved the issue.

In another incident, the author was preparing to upgrade the core servers to Windows Server 2012 R2 and observed that static IPv6 addresses set using PowerShell would not

persist across reboots. The issue was first noticed in mockup and caused considerable consternation since the same commands worked flawlessly in Windows Server 2012. Documentation could not be found indicating a different approach was needed. The problem was eventually traced to a design change that required the DHCP setting in the network adapter interface to be disabled manually. This was in contrast to the prior operating system version that made the change automatically.

Finally, it is important to continually review network configurations to ensure they are protocol independent. For example, converting a setting that contains a physical IP address to a resolvable, fully qualified domain name (FQDN) eliminates dependence on whether IPv4 or IPv6 is being used. Sometimes these settings are not always apparent as the author recently discovered. For instance, routing web traffic through a proxy server entails setting a field called "Address" in the Internet Explorer control panel. The term *Address* suggests that a physical address is needed. After some trial and error to determine if a named resource would work, it was learned that a FQDN will satisfy address requirements for that field thus allowing protocol independence.

5 Closing Observations

Overall, deployment of the IPv4-IPv6 dual-stack network in the department's computing infrastructure went smoothly and suffered no major setbacks. It is the author's opinion that being able to practice and experiment with IPv6 technologies prior to the actual deployment contributed greatly to the success of the rollout. Using desktop virtualization software and an IPv6 connection via tunnel broker enabled the author to build and test mockups and trial configurations well in advance of the actual implementation.

The challenge of integrating meaningful and relevant IPv6 projects into existing courses is the next step now that the IPv4-IPv6 dual stack network is in place. A natural starting point is to extend existing IPv4 setups into the IPv6 domain and take advantage of the similarities between protocols. Having the ability to compare IPv4 and IPv6 protocols and deployment methods offers the ability to highlight advancements made in the IPv6 protocol and leverage IPv6 technology in the classroom.

In closing, this paper presented the process of deploying IPv6 technology in an academic computing environment to provide students and faculty with the tools to extend textbook concepts to the world of practice. Topics discussed included the challenges of deploying IPv6 technology, curricular implications, design considerations, implementation, and migration of the existing systems. A variety of design alternatives were presented to stimulate discussion and encourage academic departments to consider incorporating IPv6 technology in their curriculum.

References

[1] Czyz, J., Allman, M., Zhang, J., Lekel-Johnson, S., Osterweil, E., and Bailey, M., Measuring IPv6 Adoption (2014). ACM SIGCOMM ’14, August 17-22, Chicago, IL, USA, pp. 87-98.

[2] ICANN urges IPv6 adoption as global address shortage looms (2014). ZDNet, http://www.zdnet.com/article/icann-urges-ipv6-adoption-as-global-address-shortage-looms/.

[3] Smart cities will house 9.7 billion IoT devices by 2020: Gartner (2015). ZDNet, http://www.zdnet.com/article/smart-cities-will-house-9-7-billion-iot-devices-by-2020-gartner/.

[4] Deprecating Anycast Prefix for 6to4 Relay Routers draft-ietf-v6ops-6to4-to-historic-11 (2015). IETF, https://tools.ietf.org/html/draft-ietf-v6ops-6to4-to-historic-11.

[5] Using DNS servers with DHCP (2005). Microsoft, https://technet.microsoft.com/en-us/library/cc787034%28v=ws.10%29.aspx.

[6] The Dynamic Host Configuration Protocol (DHCP) Client Fully Qualified Domain Name (FQDN) Option—RFC 4702 (2006). IETF, https://tools.ietf.org/html/rfc4702.