Deploying a Virtualized Network Infrastructure in an Academic Computing Environment

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

Abstract

During Fall Semester 2015, the author pursued a plan to partition the network infrastructure of the Department of Mathematics and Computer Science at the University of Wisconsin-Superior into a fully zoned architecture. The goal was to adopt a more rigorous security posture to protect the computing assets from external and internal threats. An innovative design was proposed to create a virtual infrastructure using network appliances configured from virtual machines and open source software. The paper provides an overview of the system and components necessary to zone an IPv4-IPv6 dual-stack network. In addition, the design, technologies, and processes used to achieve this goal; the trade-offs involved; and the difficulties encountered are presented. Finally, the paper closes with commentary on the approach and inherent challenges of securing an academic network infrastructure in a resource constrained environment that must accommodate technological uncertainties in the world of practice and academics.

1 Introduction

Traditional network infrastructures are often characterized by the physical devices that define the structure and behavior of a network and the paths packets take. Networking devices generally have specially designed hardware and firmware dedicated to the task they are assigned and fall into specific categories that include some combination of switching, routing, or filtering. Designers select network appliances based on specific functionality and performance characteristics and physically connect the device using transmission media based on metal conductor, optical link, or radio frequency. Network architectures are often illustrated schematically using diagrams that contain symbols representing the functionality of the physical devices.

In a virtual network infrastructure, connections are made through logical constructs whose instances are defined in software. There may be few, if any connections to the physical world in completely virtualized environment making it difficult to insert off-the-shelf components. Instead, networking devices are built using virtual machines running on generic computer hardware using commodity operating systems. The ability to configure a stock operating system and add functional applications greatly expands the options available to design and implement customized network appliances. This flexibility blurs boundaries typically associated with physical network devices as appliances are tailored for specific purposes. As a result, naming conventions and diagramming symbols may no longer apply or fully represent the appliances that define network structure leading to communication challenges and loss of functional identity.

During Fall Semester 2015, the author pursued a plan to partition the network infrastructure of the Department of Mathematics and Computer Science at the University of Wisconsin-Superior into a fully zoned architecture. The goal was to adopt a more rigorous security posture to protect the computing assets from external and internal threats. Discussion begins with a brief overview of the existing system and the limitation of the current implementation followed by an overview of the proposed system and components necessary to zone an IPv4-IPv6 dual-stack network. Next, the design, technologies, and processes used to achieve this goal; the trade-offs involved; and the difficulties encountered are presented. Finally, the paper closes with commentary on the approach and inherent challenges of securing an academic network infrastructure in a resource constrained environment that must accommodate technological uncertainties in the world of practice and academics.

2 Background

Deploying a virtualized network infrastructure is a significant undertaking since it entails redesigning the original architecture of the department’s network. The transition impacts a number of areas and requires coordination between several stakeholders, such as the university network administrator, department faculty and staff, and students, to ensure the computing resources remain operational.

Nonetheless, it is important to push forward on projects that protect and strengthen the computing assets the department has acquired and built up over the years. The MCS computing infrastructure exists in an ever changing environment that needs to reviewed and monitored for threats that could disrupt the computing services students and instructional personnel rely on to fulfill department's tripartite mission of teaching, scholarship, and service.

2.1 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.

The computing infrastructure consists of multiple labs, internet-facing devices, and core services that provide the backbone of the system. Lab facilities include two advanced instructional computing centers that contain a total of 50 shared computers, a hardware lab containing two dedicated computers, a learning lab that contains two high-end workstations, and a sandbox that contains assorted computers for prototyping. Internet applications are used extensively not only to access the tools necessary for class activities but also as a resource for academic and personal interests. Approximately 400 students use these facilities per year.

The department also hosts a number of internet-facing devices used for department operations as well as providing services beyond the confines of the facility. Internal services include a 6in4 tunnel broker and various software update services. External services include multiple web servers along with several application servers that host application specific tools tailored to the needs of students and faculty.

Core systems provide essential services necessary for authentication, hosting virtualized applications, manage computing assets, and deploy security policies. The physical system consists of two servers that host domain services, a two node failover cluster with an iSCSI storage array that provides a virtualization platform for application servers, a dedicated server that hosts a virtual desktop infrastructure, a dedicated virtual host for management applications, and other specialty servers that provide shared storage.

The entire computing infrastructure resides on a single network segment connected to the main campus backbone. The segment connects all the facilities used throughout the department's computing infrastructure including the server room, labs, and workspaces. The network supports IPv4 and IPv6 protocols in dual-stack configuration. Traffic utilizing the IPv4 access point must go through the campus firewall and benefits from the managed security systems the University maintains. The IPv6 access point tunnels directly through the firewall and provides an unmanaged connection to the internet.

2.2 Project Motivation

Breaches, exploits, malware, and other malicious activities have the potential to totally disrupt operations by rendering computing platforms inoperable and corrupt data. Even with a backup system in place, the potential outcome of a disruption can severely hinder normal day-to-day operations and consume the few resources available to recover from such an incident. In addition to preventative measures, containing a breach or isolating a malware attack is a prudent step toward ensuring the systems the department relies upon remain functional.

Over the past five years, the department’s computing infrastructure has steadily grown to support new services and capabilities. Resources are continuously injected to update old systems and build new ones to meet the evolving needs of students and faculty leading to a considerable investment of time and effort. Knowledge of the wide range of systems used in the department is concentrated across a few individuals that lead the charge to ensure its continued operation and effectiveness. Considerable effort is made to manage the attack surface and limit susceptibility to internal and external threats.

Allowing internet-facing systems, lab workstations, and core services to share a single segment potentially allows a breach to spread unhindered across systems. Partitioning the network into zones and managing traffic adds an additional layer of defense. Access control lists along with event logging can alert administrators to potential threats in advance of a breach or at least provide a trail for forensic analysis should a breach occur.

During the summer and fall of 2015, three new threats emerged that triggered the author to redesign the network architecture. These threats included 1) a significant increase in the number of unauthorized attempts to access particular servers in the department, 2) a new wave of ransomware attacks, and 3) student efforts to circumvent measures put in place to control internet access. The time had arrived where the current system had outgrown the original network architecture and another structure was needed to protect computing assets and contain a breach should one occur.

2.3 Considerations and Alternative Selection

It had been the author’s intention to partition the network into separate zones containing internet-facing devices, lab workstations, and core infrastructure for some time. Although thoughts had been given to this architecture, no clear path emerged how to implement this structure given the constraints of the system already in place, budget restrictions, and the time and effort needed to pursue the project. There were a number of complicating factors had to be taken into consideration before any particular alternative could be selected.

Three conditions would have to be met before the project could move forward. First, any zone would have to accommodate physical and virtual network connections distributed across multiple systems. Second, the structure would have to support high-bandwidth connections needed for desktop virtualization. Third, changes must be able to be undone with the least amount of time and effort if the system does not work as anticipated.

Several hardware and software configurations were considered but a solid design concept never surfaced. One alternative included the purchase of a layer-3 switch. Unfortunately, enterprise grade equipment is very expensive and requires significant expertise to configure making this option less favorable given the uncertainty. Another alternative included the reassignment of one of the physical servers and converting it to a router. This proposal was taken off the table because it would require substantial changes to the existing network structure making it difficult to implement. In the end, the project was put on hold until a better way could be found.

Virtualizing the network infrastructure emerged as an option during fall 2015 after the author took second look at the VLAN capabilities in Microsoft's Hyper-V virtualization platform. Partitioning the network would require two additional VLANs, one for the lab workstations and another for core services, while the existing network segment would continue to serve as the DMZ for internet-facing systems. Conversations with the campus network administrator confirmed the feasibility of adding additional VLANs to the existing infrastructure.

This approach enjoyed other benefits as well. First, no hardware would have to be purchased. Virtualizing the network infrastructure leverages the existing virtualization hardware saving money for other projects. Second, virtualization is flexible. Virtual machines and switches can be arranged in many different ways offering considerable versatility. Third, the system could be deployed incrementally and be reverted to the previous state if necessary. If the new architecture fails to perform as anticipated, then it could be abandoned without incurring significant penalties. In the end, a virtualized network infrastructure was selected and a design drawn up for deployment.

3 Design and Deployment

There is little argument that deploying a zoned, network architecture is positive step forward toward protecting department computing and data resources. Virtualizing the infrastructure, however, entails overcoming limitations associated with traditional design methods and figuring out how to take advantage of flexibilities virtualization offers. To achieve this, it is important to first consider the logical design and how it relates to its physical implementation. Next, the logical design is reconsidered in light of a virtualized infrastructure and how virtualization changes the perspective. Once complete, the actual implementation is detailed identifying the technical details of the configuration.

3.1 Logical Design and the Physical World

The basic architecture of the network structure consists of three zones labeled as DMZ, Management, and Labs shown in Figure 1. At the top-center of the diagram, a layer-3 switch defines core of the network. A layer-3 switch is routing devices that implements address lookup in hardware using application specific integrated circuits (ASICs) to substantially reduce latency. A router on the other hand, uses a routing table implemented in software and generally includes provisions to connect a variety of wide area network protocols. As the diagram shows, a layer-3 switch forwards data packets between zones while inline firewalls govern access. Data packets are directed to respective clients using layer-2 switches.

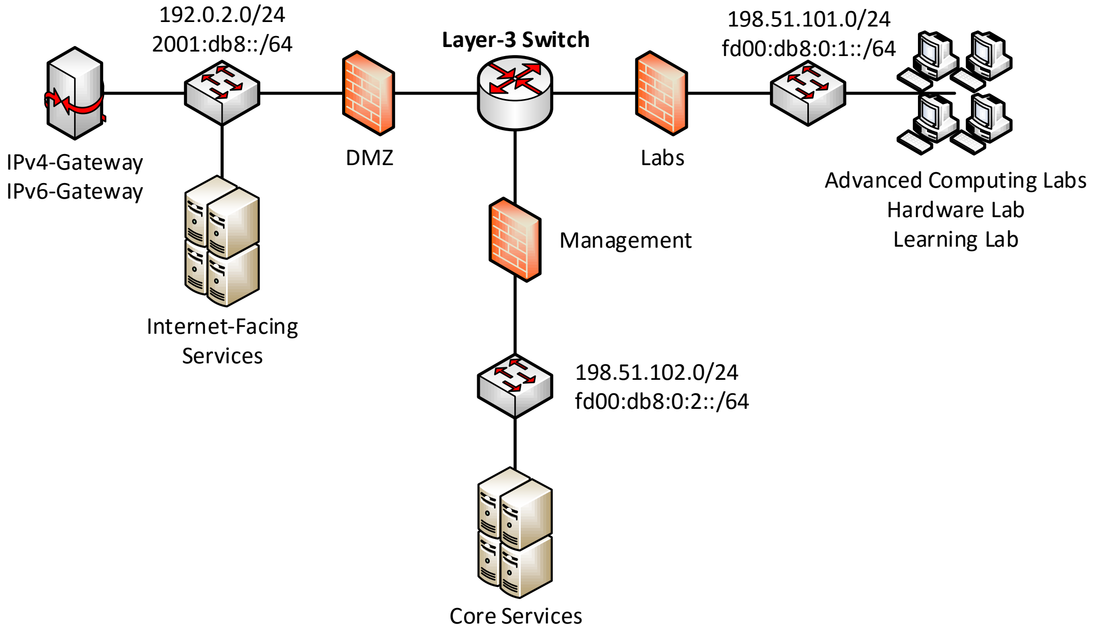

IPv4 and IPv6 documentation addresses shown as substitutes for the actual assignments.

Figure 1: Diagram of zoned network infrastructure

Each zone hosts IPv4 and IPv6 protocols arranged in a dual-stack configuration. The IPv4 network consists of three links assigned private addresses. The IPv4-Gateway serves as the default gateway for devices located in the DMZ; whereas, the router serves as the default gateway for client systems on the Labs and Management links. Network address translation (NAT) is used to connect IPv4 network links.

The IPv6 network utilizes both global and unique-local address spaces. The DMZ is mapped to a global address space while the Labs and Management links are mapped to subnets of a private address space. The IPv6-Gateway serves as the default gateway for devices located on the DMZ link; whereas, the router serves as the default gateway for client systems on the Labs and Management links. To ensure that private addresses are routed back to the router, a static route using the private address network prefix (fd00:db8::/48) is appended to the IPv6-Gateway routing table. This ensures all devices on the IPv6 network are reachable.

The diagram illustrating the network components shown in Figure 1 would accurately depict the components used in a physical network infrastructure. Each symbol, whether layer-3 switch, layer-2 switch, firewall, or gateway, would normally correspond to a particular device each maintaining a physical connection to another device using a patch cable. In this case, the logical world the network diagram illustrates coincides nicely with the physical world where off-the-shelf components could be purchased that meet the requirements of the network.

3.2 Logical Design and the Virtual World

Implementing a network infrastructure in a virtual environment replaces a number of key devices with customized virtual machines. The illustration shown in Figure 2 attempts to capture the essence of a virtual network infrastructure. The diagram illustrates the external and internal connectivity to a virtual machine host. External traffic passes from the virtual switch through a physical network interface to a layer-2 managed switch. From there, the switch directs packets to specific ports connecting physical servers or other network appliances.

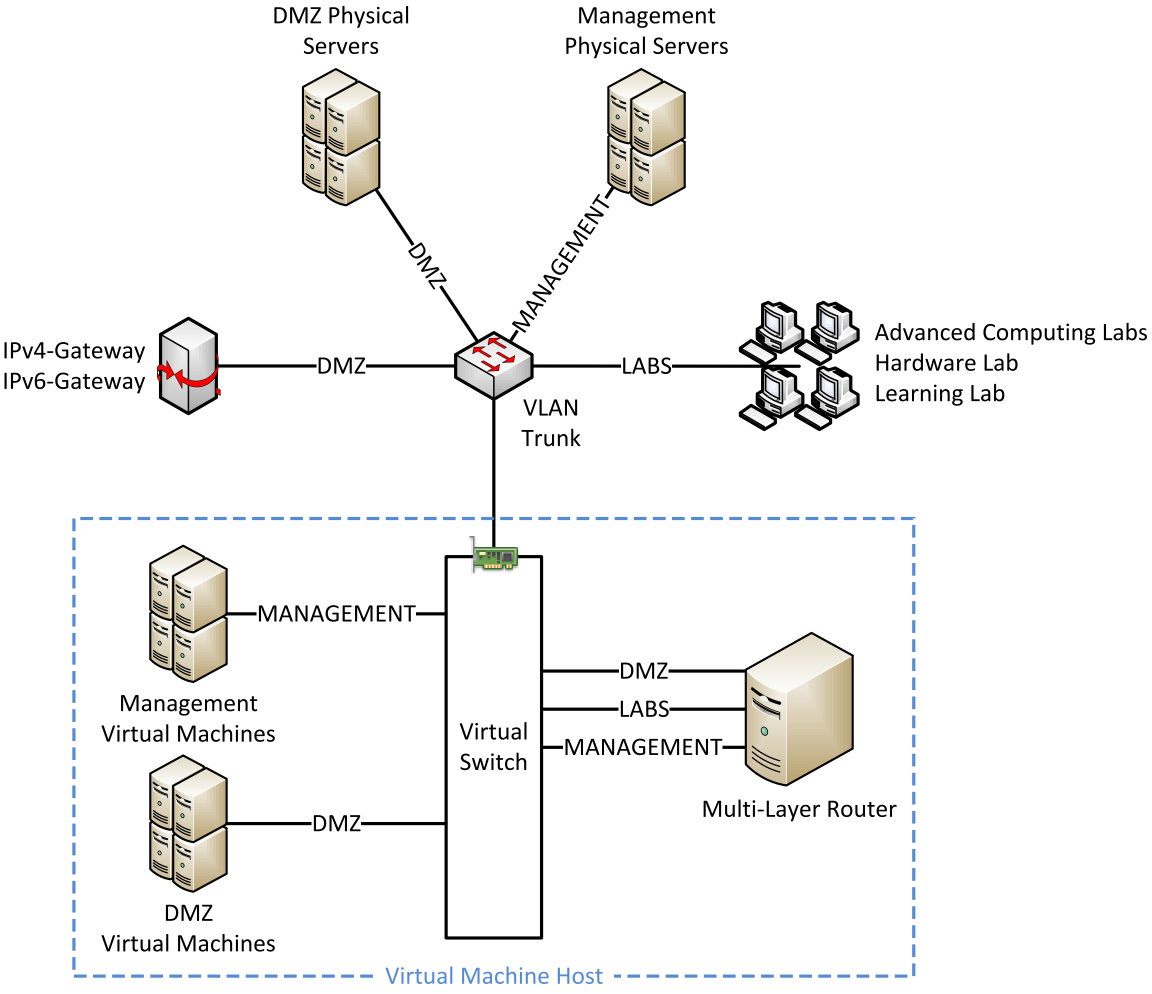

Figure 2: Diagram of virtual network infrastructure

The diagram also depicts a virtual machine host as dashed line bounding the virtual switch and assorted virtual machines. The aggregate of virtual machines shown on the left-hand side of the virtual switch represent pools of virtual machines connected to the management and DMZ zones. They have been grouped for clarity but each virtual machine is bound to a dedicated virtual network adapter that is assigned a VLAN identifier. In turn, each virtual network adapter is connected to the virtual switch.

There are two important differences worth noting between the diagram shown in Figure 2 and the diagram shown in Figure 1. First, the virtual switch plays a central role in the network. A virtual switch is software mechanism that connects virtual machines but does not provide the functionality normally associated with a physical layer-2 switch. Instead, virtual network adapters do most of the work by selecting packets for the virtual machine. In addition, the VLAN identifier is set on the virtual network adapter not the virtual switch allowing tagged packets that match the VLAN identifier through to the host. The second important difference is the representation of the multi-layer router[1] as a virtual machine. In this case, one device combines the all the functionality of the layer-3 switch and adjoining firewalls shown in Figure 1.

Creating customized network appliances using virtual machines greatly expands options for network designers but also incurs significant drawbacks. First, these customized systems do not associate with traditional network nomenclature. Every appliance is a unique device so communicating functionality is a real challenge particularly in discussions with other network professionals. Second, network appliances are limited to software implementations of common network functions. There is no provision for the integration of custom ASICs at this time thus limiting the performance of the device to the processing capabilities of the underlying virtual machine host.

3.3 Implementation

Network appliances occupy a critical role in an organization’s network infrastructure and must be specially crafted to ensure secure and reliable operation. These devices may have roles that span multiple network zones, host internet services, or manage network traffic making them primary targets for malware and malicious attacks. Best practices along with a minimalist philosophy guided the construction and implementation of the devices.

A virtual machine hosting a minimal installation of Debian 8 resides at the core of each network appliance. Debian is a well-known Linux distribution that is used as a stand-alone operating system or as a base for other popular operating systems such as Ubuntu and Linux Mint [1]. It has a reputation for stability making it a highly regarded candidate for server operating system in enterprise environments. Only two additional packages, *sudo* and *resolvconf*, were installed to round out the base configuration. A local administrative account was created and root privileges granted using *sudo*. The root account was deactivated as a security precaution.

[1] The term “multi-layer router” was co-opted by the author for lack of a better term to describe a network appliance that performs stateless and stateful packet inspection in addition to its primary routing function.

The multi-layer router required additional configuration to enable routing and zone management capabilities. Packet forwarding is enabled between interfaces by setting the following kernel variables.

```
net.ipv4.conf.all.forwarding = 1
net.ipv6.conf.all.forwarding = 1
```

Next, the packages shown in Table 1 were installed to provide the specific network services and zone management functions.

Package	Description
radvd	Provides router advertisement services. Configured only for Labs and Management zones.
isc-dhcp-relay	Provides relay service for DHCPv4 and DHCPv6 services. DHCPv6 initialization and configuration file must be created manually.
etherwake	Utility to issue wake-on-LAN packets to restore client systems from low-power states.

Table 1: Packages installed on multi-layer router

Although the network infrastructure supports IPv4 and IPv6 protocols in a dual-stack configuration, the author's preference has shifted toward IPv6 as the internet protocol of choice. The rationale for this decision is that IPv6 private networks offer a more secure alternative than IPv4 private networks that rely on network address translation (NAT). Private IPv6 networks employ unique-local addresses that are not internet routable meaning that all inbound and outbound traffic stops at the internet-facing router by default unless specific permission is granted. Private IPv4 networks using NAT allow all outbound traffic by default potentially opening a pathway to internal devices for illicit purposes such as remote control or data exfiltration.

Internal private IPv6 networks gain access to internet resources using edge devices. Edge devices serve two purposes: First, they create a point of presence for a particular internet resource within the confines of a site. Second, they serve as a control point for specific network protocols that allows administrators to track and audit inbound and outbound traffic. These devices are assigned global IPv6 addresses and reside in the DMZ. They take the form of custom designed servers or network appliances hardened to withstand attacks from the internet.

Three edge services were necessary for the department's computing infrastructure. Microsoft Active Directory and Domain Services (ADDS) requires DNS and time services. While web-based and software update applications require the service of an internet proxy. All edge services were consolidated on a single host although they could also be distributed across multiple platforms as a security measure to isolate systems. The packages shown in Table 2 were installed and configured for each edge service.

Package	Description
bind9	Provides DNS services for the site. Device configured as a primary DNS server with A and AAAA records added for selected devices in the DMZ. Private domain-integrated DNS forwards unresolved queries exclusively to this device.
ntp	Provides authenticated time services for the site. Synchronizes with Stratum-2 sources. Time source for domain controller holding the PDC emulator role.
squid3	Provides internet proxy services for http, https, and ftp. Designated proxy server for Windows Server Update Server and devices that require access to web update services. Default proxy server for private network.

Table 2: Packages installed on edge server

The access control lists that govern ingress and egress traffic for each network appliance was created and managed using FirewallBuilder. FirewallBuilder is an open source application capable of configuring cross-platform rules for multiple firewalls through a single interface [2]. Firewall rules for stateless and stateful packet filtering were compiled for Linux systems using the *iptables* and *ip6tables* commands. Unresolved packets are reported to the host's *rsyslog* service that logs the incident locally as well as sending a copy to a centralized syslog server.

3.4 Deployment Challenges

Not every activity goes as planned. Many deployment challenges can be resolved prior to rolling out a new infrastructure by prototyping the system in advance; however, some components may not be accessible or cannot be easily simulated. The department's managed switches are one such system. Configuring these switches requires coordination with the campus network administrator to ensure proper integration with the university network and cannot be accessed directly. In addition, the author did not have access to either physical switch or virtual mockup making it difficult to test settings in advance.

The biggest problem arose from the set of switches that connect the virtual hosts that provide the virtualization platform for the department. These switches require specific VLAN trunk settings that enable network traffic to pass between servers transparently. To compound the situation, settings can only be applied when classes are not in session to prevent network disruptions in the computing labs. This leaves only small windows of opportunity during the academic year and the summer months to deploy configuration changes to the managed switches. Once the switch settings are configured, the path will be open to full deployment of the new virtualized network infrastructure.

4 Closing Thoughts

The ability to virtualize a network infrastructure opens up a number of possibilities for academic departments that rely on computing systems to fulfill their academic mission. Instead of having to rely on physical network devices that are often very expensive and hard to configure, new and innovative structures can be deployed using available virtual machine technologies and open source applications. The flexibility that virtualization offers administrators of academic networks allows for setups once thought impractical or even out of reach.

For departments seeking an option to hosting their own computing infrastructure, virtual network devices open the door to cloud-based services hosted privately on campus or publicly through any number of commercial providers. An academic department could in theory utilize a cloud infrastructure to provide its entire computing and network platform without having to provide facilities for and secure the upfront cost of servers, networking gear, and other equipment.

In closing, this paper presents an innovative way to design and implement a virtualized network infrastructure. Topics include the need to protect academic computing resources from internal and external security threats, difference arising from physical and virtual implementations of a logical design, and challenges the author experienced deploying the system. Using virtual machines and open source software, the solution provides a number of capabilities that facilitate control of network traffic while enhancing the security of computing and data assets while promoting the effective use of computing technology in an academic environment.

References

[1] Debian, "About Debian," December 2015. [Online]. Available: https://www.debian.org/.

[2] FirewallBuilder, "Home," 2012. [Online]. Available: http://www.fwbuilder.org/.