Reviewing an Academic Computing Infrastructure: Reflecting on the Past to Forge a Path into the Future

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

Abstract

Several years have passed since the Department of Mathematics and Computer Science at the University of Wisconsin-Superior revamped its computing infrastructure in 2010. As the system enters its twilight years, the author conducted a review to evaluate what worked, what didn't, reflect on lessons learned, and consider new architectures and technologies that might better suit the mission of the department. This paper begins with an overview of the current infrastructure and examines the various systems in place. Next, the paper examines the system's effectiveness considering various dimensions such as alignment with mission, classroom experiences, and adapting to changes in academic programs. Then, a technology review identifies the most prevalent challenges and considers new services that offer improved system functionality and manageability. The paper closes with thoughts and observations on the review process and efforts to secure resources for technology-based academic programs.

1 Introduction

Several years have passed since the Department of Mathematics and Computer Science at the University of Wisconsin-Superior launched an effort to redesign its computing infrastructure in 2010. As the department's computing infrastructure enters its twilight years, the author has commenced upon a review to evaluate what worked and what did not work, reflect on lessons learned, and consider new architectures and technologies that might better suit the mission of the department and the evolving needs of students and faculty who rely on this system as valuable educational resource. Although the existing infrastructure is still operational and continues to perform its intended function, the constant wear and tear of hardware and technological obsolesce necessitate the system's eventual replacement.

Performing a review, even if informal, provides a reflection point prior to a foreseen transitional period. If anything, the process of a review helps develop a vision and mental map in which to make forward progress. This is important in that the department's computing infrastructure is built around a core set of technologies that are being displaced by newer techniques that yield improved reliability and manageability without having to resort to specialized or proprietary systems. For a modest sized infrastructure, this generally entails a complete overhaul of the components that constitute the foundation of the system. Nonetheless, there are inherent limits of a review in that past performance is not necessarily an effective predictor of the future.

This paper presents the author's efforts to review the function and effectiveness of the department's computing infrastructure and to outline considerations moving forward. Discussion begins with a brief overview of the computing infrastructure, the system's intended function, and how the system supports the mission of the department. Next, the author considers various dimensions that constitute effectiveness such as alignment with mission, classroom experiences, and adapting to changes in academic programs. Finally, a technological review identifies the most prevalent challenges confronting the current infrastructure and considers new services and technologies that offer a way forward. The paper closes with commentary on the approach and inherent challenges of tailoring an academic computing infrastructure to suite a wide array of needs for a diverse set of users.

2 Computing Infrastructure Overview

The Department of Mathematics and Computer Science manages and maintains a dedicated computing environment for academic activities. Computing resources are used extensively in classes across the department's programs ranging from entry level coursework that fulfill the University's mathematics general education requirements to advanced electives in majors and minors. A wide array of applications suited for general productivity, programming, software development, analytic modeling, device emulation, network security, multimedia, and web authoring for use by faculty, instructional staff, and roughly 400 students per year.

The computing infrastructure consists of multiple labs, core systems, and application servers that 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.

The department enjoys several computer lab facilities used by students and instructional staff. The principle resources used for instruction are called the Advanced Computing Labs and contain a total of 50 computers equally divided between two separate locations. Each room contains 24 shared workstations available to all authenticated users on first-come, first-serve basis plus a dedicated teacher workstation connected to a central projection and sound system. The department also has two smaller lab facilities called the Learning Lab and Development Lab. The Learning Lab contains two workstations and is used for special projects that require high-performance hardware. The Hardware Lab is a limited access facility and is primarily used for microcontroller and embedded system projects. It contains two computers configured specifically for that purpose, one of which is connected to a microscope with a digital camera and the other to a printer used for printing circuit board layouts. The department also maintains a collection of computers for prototyping referred to as the "sandbox" that serve as a resource for academic and personal interests.

Core systems provide key services across the computing infrastructure and is comprised of a mix of physical and virtual servers. The physical infrastructure consists of the following systems:

- Two redundant domain controllers;
- A two-node failover cluster that uses a 24-port iSCSI disk array for block storage configured as a virtualization platform;
- Three separate virtualization platforms for internal operations, virtual desktop infrastructure, and MCS Club extracurricular activities;
- A network attached storage (NAS) device for file storage; and
- Two workstations that serve as administrative consoles

The core system also includes several important services that provide essential functions necessary for operations that include:

- Active Directory Domain Services—Redundant servers used for domain administration such as authentication and group policy settings
- IPv6 Gateway—A virtualized internet appliance that provides IPv6 network access using a 6in4 tunnel broker
- Edge Gateway—A virtualized internet-facing device used to provide domain name, internet proxy, and time services for internal systems
- Internet Protocol Address Management—A virtualized server that provides DHCPv4 and DHCPv6 services

- Internet Proxy Services—Two virtualized internet-facing devices used to manage web traffic to the advanced computing labs
- Management Consoles—Two dedicated computers hosting Windows Remote Server Administration Tools that allow support personnel access and manage the computing infrastructure

The computing infrastructure hosts several application servers to support operations and academic functions that include:

- Windows Update Services—A virtualized server that provides local management and deployment of Microsoft updates and patches for server and client systems
- Helpdesk Services—A virtualized server that hosts an instance of Spiceworks to manage user support tickets
- Web Servers—One virtualized internet-facing Linux server using Nginx web services to host department member web sites and another virtualized internet-facing Linux server using Apache web services to host the MCS Club website
- Academic Servers—Three separate virtualized internet-facing Linux servers hosting Apache web services customized with various applications (LAMP, WordPress, etc.) tailored for individual instructors
- User Profile Server—A virtualized file server tasked with supporting roaming profiles and redirected folders
- Shared Folders—A virtualized file server that provides central file storage for instructional materials
- Database Servers—One virtualized MS SQL server that hosts multiple database instances for instruction and another virtualized MS SQL server dedicated to the MCS Club for their annual programming contest
- MDK Professional License Key Server—A virtualized Linux server that tracks Keil Embedded Development application instances
- Sage Server—A virtualized internet-facing Linux server that host an instance of the open source SageMath application configured for multiple users
- Minecraft Server—A virtualized server managed by the MCS Club that hosts a Minecraft instance used for LAN parties

In all, the department hosts and manages about two dozen virtual instances of which nearly twenty are running fulltime in support of the department's day-to-day activities.

3 System Effectiveness

Most faculty and students using the department's computing infrastructure simply use the resource as they are and do not consider the ins and outs of its operations. The author's intent for this review is to consider how the technology impacts how the department conducts its business and those that rely on its operations. Therefore, the effectiveness of the computing infrastructure primarily considers: 1) alignment with the department mission, 2) classroom experience, and 3) adapting to changes in academic programs.

Like most academic units, the Department of Mathematics and Computer Science has a tri-partite mission that includes teaching, scholarship, and service. Of the three, teaching stands out as the largest component of the department's efforts given that the University is primarily student focused and revolves around graduating educated students with critical thinking skills. Living up to this aspect of the mission entails that faculty have a substantial teaching load that averages 12 credits per semester to sustain academic programs in mathematics, computer science, and information technology.

Computer lab use varies widely between the three programs. The mathematics program typically utilizes the department's computer facilities for the lab portion of MATH 130 Elementary Statistics, a service course that satisfies general education requirements for students enrolled in programs across campus. Two sections are normally taught each semester and are often fully enrolled. Demands on the computer lab resources are minimal and generally requires only software capable of analyzing basic statistics. Conversely, nearly all the courses in the computer science and information technology programs are held in the advanced computer labs. These classes utilize a wide range of applications that include common productivity tools, individual and team programming and database projects, network and storage intensive virtualization efforts, and multimedia loaded web applications. The computing infrastructure has thus far handled this type of use with aplomb but is now beginning to approach storage and network bandwidth limits.

Ad hoc requests to install software applications in the labs for special occasions is another area that frequently causes frustration amongst the users of the department's computing facilities. It is the author's observation that instructional staff and students have difficulty foreseeing what tools will be required until the time they need it. Some applications bypass Windows Installer (which requires administrative privileges to run) and installs directly in the user's profile. Consequently, users may experience protracted login and logoff delays since their individual profiles must be loaded and unloaded in a shared computing environment that utilizes roaming profiles. These two factors often work against the goal of the system administrators to ensure workstations maintain a uniform and stable configuration.

Adoption and migration to cloud applications by students and faculty impacts the necessity of the computing infrastructure for certain tasks and challenges local administrators to keep pace with changes. At the time the infrastructure was designed, applications the department desired needed to be hosted locally. This situation changed within a couple of years as cloud technologies made new applications available and easily accessible. For the administrator of the department's computing infrastructure, SageMath Cloud (CoCalc), GitHub, and Microsoft OneDrive exemplify this problem since they rapidly displaced traditional approaches to mathematical analysis, source management and collaboration, and file storage services.

The inability to remotely access the system also impedes student access to departmental computing resources for those not able to be physically present in the labs. There have been instances where students would like to take a class to complete a plan of study but have moved to another location that makes campus resources inaccessible. Enabling remote access in a seamless and secure way is a capability that must be designed into a system from the start. Consideration must be given to both technology and licensing issues along with the ability to justify and potentially recover costs. Based on the author's experience, configuring remote access for an individual requires a significant amount of customization that is hard to replicate in an efficient manner. Even if remote access to local computing resources are created, it often does not represent the full functionality available to local users.

Finally, as the department begins to participate in new online collaborative programs sponsored through the University of Wisconsin-Extension, the need for a local computing infrastructure is reduced since centralized computing resources are provided by the host organization. Justifying the cost of deployment and maintenance becomes more difficult when a portion of the students participate remotely and faculty resources are reallocated to programs that support dedicated systems specifically designed for the purpose, achieve economies of scale, and recover costs directly through tuition and fees.

4 Technological Review

Should a decision to replace the current computing infrastructure occur, there would be several changes that author would recommend. First, the department should employ a hybrid strategy that includes both cloud and on premises technologies that take advantage of the strengths that each approach offers. Second, partition the network to a zoned architecture that restrict portions based on function and improve security. Third, deploy a homogenous, scalable virtualization platform using commodity hardware to adapt to changes in utilization over time.

Hybrid Cloud/On Premises Strategy

The objective of a hybrid cloud/on-premise strategy is to utilize the strength of each technology to advance the mission of the department to its fullest. Multitenant cloud systems are scalable virtualization and application platforms that run in one or more datacenters that can be accessed remotely. The ability to access these computing resources remotely allow users to access the system from anywhere freeing individuals from doing their work from specific locations such as a campus computer lab. Cloud systems are particularly useful for student projects that require individual workspaces and demand high speed storage along with ample network bandwidth. Desktop virtualization is one such example where the benefits of a cloud system would greatly exceed the local resources needed to provide an equivalent service.

Cloud systems would also be well suited for persistent server instances that have been traditionally hosted on the department's computing infrastructure. Not only can cloud services expand and contract the resources needed to host server instances, but they often provide management and security tools that enable clients to optimize performance and protect their systems. There are several servers currently hosted by the department, such as the database, web, and academic servers, which could be migrated to the cloud thus freeing local computing resources for other activities.

On the other hand, there are systems that would be better located on the premise than in the cloud. For example, systems sensitive to time delays or latency such as file shares (including roaming profile and redirected folders), databases, and applications that need low-latency, high-speed network connections to local devices. Also, private systems (e.g. domain controllers, internet proxy servers, etc.) might also be better situated within a local infrastructure.

The biggest drawback to utilizing cloud systems by academic departments is converting from an ownership model to a pay-as-you-go service. A pay-as-you-go service trades upfront costs for an ongoing expense that must be absorbed by an already constrained departmental budget. In the author's experience, it is often easier to secure a periodic internal grant than to achieve even a modest increase in a department budget.

Network Architecture

The department's computing infrastructure currently resides on a single network segment where core systems, internet-facing devices, internet accessible servers, and lab computers share the same network address space. This makes it difficult to track and manage traffic between the different types of systems thus give rise to security issues. Although the author has made attempts to partition the network, the situation is an artifact from original infrastructure design and is difficult to change given its configuration within the larger campus network infrastructure. In addition, the campus network administrator maintains a gateway that manages IPv4 only traffic since the campus only uses the IPv4 protocol. Whereas, the department employs an IPv4/IPv6 dual-stack network and manages IPv6 traffic using a separate gateway in conjunction with a tunnel broker service. IPv6 traffic is encapsulated in IPv4 packets and tunneled through the IPv4 gateway managed by the IT Department of the University.

Retaining the IPv4/IPv6 dual-stack network is an important consideration in any future network architecture. As the University eventually adopts IPv6, it would be preferable to consolidate the IPv4 and IPv6 perimeter devices into a single gateway that manages traffic at the boundary of the department's network. It would also be desirable to partition the network into zones that consolidate network devices based on their security profile. This would entail acquiring a suitable Layer-3 switch and firewalls to connect and manage traffic between the various zones. A zoned architecture also enables expansion to include new zones for devices that have their own unique security profile.

Network bandwidth also needs special attention particularly amongst members of the core system and connections to the advanced computing labs. At present, links throughout the network use one or more 1 Gbps network adapters and are bonded (teamed) to provide more throughput. Bonding with load balancing will scale effectively when there are multiple sessions across a NIC team, however, a single session can only transmit and receive at the maximum rate of a single team member. In the latter case, large file transfers in excess of 300GB may take an hour or more too complete (e.g. migrating or archiving a virtual machine hard disk). Therefore, using high-speed network adapters on network links subject to bottlenecks is key to improving the overall performance of the system.

Virtualization Platform

The ability to virtualize servers and devices resides at the center of the department's computing infrastructure. Core system servers, internet-facing devices, application servers, and virtual desktop infrastructure all rely on a robust virtualization platform. Overtime, the department's computing systems hosting the virtualization platforms have reached their physical limits. In addition, piecemeal growth has led to fragmentation of the virtualization platforms that make it difficult to balance loads across the available resources.

Overprovisioning virtual machine processing cores, memory, and storage is an essential capability to maximize the utility of bare-metal hardware. There is a point, however, where the physical resources are exhausted and can no longer keep pace with the overprovisioned demand. The department's failover cluster that hosts the main virtualization platform is approaching that limit with a 2:1 ratio for processing cores and full consumption of block storage during peak usage periods. To ease the demand for physical resources, less critical virtual machines have been offloaded to standalone virtualization platforms thus losing the redundancy the failover cluster provides.

End-of-life and proprietary hardware also limit the ability to scale systems to accommodate demand. At a certain point, vendors no longer release software drivers for old servers making it difficult to install the latest operating systems that provide new functionality security features. Proprietary systems are even more difficult to manage and upgrade since maintenance and troubleshooting guides are limited or non-existent and vendor extended service agreements are often cost prohibitive for financially constrained departments.

The fragmentation of virtualization platforms is arguably the biggest challenge to ensuring a coherent virtualization architecture. Each additional platform introduces physical and logical boundaries that disrupt uniformity. The ideal virtualization platform would appear as a uniform resource where virtual machines could reside independent of the underlying hardware. In this arena, proprietary systems have given way to commodity hardware and converged architectures that scale and provide highly available solutions.

5 Closing Thoughts and Observations

Reviews are an integral part of the systems lifecycle and offer an opportunity to lookback and reflect on what worked, what didn't, and to think about how to do things differently. However, the process has many similarities to fighting the last war in that the experiences gained from designing, building, and maintaining an academic computing infrastructure may not be completely relevant for an uncertain future as academic programs evolve. Nonetheless, plans that arise from a review at least provide a baseline in which to compare alternatives as the department prepares for the eventual phase-out of a system integral to its operations and valued by students and faculty alike.

This paper highlights the most pertinent challenges facing the computing infrastructure in a discussion that attempts to balance effectiveness and technological issues. It goes without saying however that the biggest obstacles will likely stem from administrative justification, navigating budgetary constraints, and securing personnel with the skills necessary to design, deploy, and maintain the system. University of Wisconsin-Superior like many academic institutions have faced daunting economic pressures in the decade since the Great Recession and often strain to fulfill their academic mission. Yet, departments that utilize technology extensively still need the resources to fully accomplish their assigned mission and prepare students for what lies beyond.