Rebuilding an Academic Network Infrastructure Employing Virtualization and Failover Clustering

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

Abstract

Occasionally, circumstances coincide that necessitate a new vision for an academic department's technology resources. Computer labs, servers, networking, and facilities undergo a complete transformation to better serve the academic mission, engage students, spur creativity, and push forward the state of the art. The Department of Mathematics and Computer Science at the University of Wisconsin-Superior is in the midst of a transformation where a new academic facility, aging and obsolete systems, and the need to expand technological capabilities forced a rethinking of the way business is done. During January 2011, a student employed as the advanced computer lab manager proposed a new server and network infrastructure based on server virtualization and failover clustering. This one idea spurred a process to consolidate separate threads of activity into a single cohesive vision. This paper presents a narrative of the experience, changes, and lessons learned.

1 Introduction

Occasionally, circumstances coincide that necessitate a new vision for an academic department's computing resources. Computer labs, servers, networking, and facilities undergo a complete transformation to better serve the academic mission, engage students, spur creativity, and push forward the state of the art in fiscally challenging times. Creating such a vision is an incremental process punctuated by profound insights that requires an open mind and the ability to make a break from the past. The Department of Mathematics and Computer Science at the University of Wisconsin-Superior is in the midst of one such transformation where the move to a new academic facility, aging and obsolete systems, and the need to expand technological capabilities beyond the bounds of traditional academic programs forced a rethinking of the way business is done and the systems that enable learning and discovery.

Virtualization is widely used to consolidate server resources, improve hardware utilization, and centralize application management. Separating applications from the underlying hardware introduces additional degrees of freedom that enable services to operate and migrate within a static computing infrastructure. Failover clustering integrates computing and networking components into a seamless system that incorporates planned redundancy to sustain a functional computing environment that minimizes unplanned and protracted outages. These two concepts are often associated with production environments that must provide consistent service to users. However, this approach has distinct advantages for an academic environment characterized by high application turnover, ongoing development activities, and the need to scale capabilities that match fluctuating demand witnessed throughout the academic year.

This paper presents a narrative of an ongoing effort. To begin with, a brief overview of the previous system is presented that includes a description of the computing laboratory, server and network infrastructure, and processes called upon to support the various department curriculums. Next, the paper articulates the chain of events beginning with the single idea offered by a student that set off the transformation and the changes leading to a reorganization of the computing infrastructure, expanded services, and incorporation of processes utilized in the world of practice. Included is a discussion of the fundamental role students have in the design, implementation, and operation of the system. The paper closes with commentary on the approach and inherent challenges of designing and implementing an academic computing infrastructure that employs virtualization and failover clustering.

2 Previous System Overview

2.1 Computer Laboratory

The Department of Mathematics and Computer Science managed and maintained a dedicated computing laboratory for the computer science program. The Advanced

Computer Lab was located across campus in a building separate from faculty offices. As the sole departmental computing facility, the room served as classroom and development area for students and faculty. A cardkey system controlled access to the premises as well as protected computing assets. The room was available 24 hours a day and on weekends for student working on homework or class projects. Occasionally, the room was used for special activities such as departmental programming competitions and LAN parties.

The Advanced Computer Lab contained 20 computer workstations configured with dual-monitors. Workstations were replaced on a four-year schedule; however, only 25% of the workstations were replaced at a time. Although the computers were all purchased from Dell—the University's designated vendor—this led to the situation where the lab contained four different computer models all with varying capabilities and performance characteristics. The latest generation of computers contained 4GB of RAM while the remaining computers had 2GB of RAM. All but the oldest generation of computers employed 64-bit architectures.

Each workstation hosted a suite of proprietary and open source applications for instruction and development. Microsoft Windows XP was used as the base operating system and computers were joined to a private Microsoft Window 2003 domain isolated from the campus network. Software ranged from basic productivity applications to wide array of programming tools used in various courses and for personal exploration.

2.2 Server and Network Infrastructure

The server and network infrastructure used to support the advanced computer lab and provide server applications consisted of a collection of technology assembled since the turn of the millennium. The system was distributed throughout several closet spaces on multiple floors since a centralized area was not available. In addition, a dedicated multimode fiber optic linked various components of the system including the building that housed faculty offices across campus. A diagram of the system is shown in Figure 1.

The bulk of the servers and networking equipment serving the facility was located two floors above the advanced computer lab in a closet attached to a working classroom. The system consisted of five Dell PowerEdge servers (models 650, 1750, 2600, 6600, and 840), five Dell Optiplex GX620 workstations recast as dedicated application servers, and a Buffalo Teraserver Pro network attached storage (NAS) device. Three Dell PowerConnect switches (one model 5212 and three model 2508) connected the various systems.

A closet attached to the computer laboratory served as the main wiring center for the facility. The closet contained a Dell PowerConnect 5224 switch that connected the lab computers, a Dell PowerConnect 2708 switch for additional connectivity to lab devices, a second Dell PowerEdge 650 that served as the perimeter device connecting the department's private network infrastructure to the broader campus network and internet.

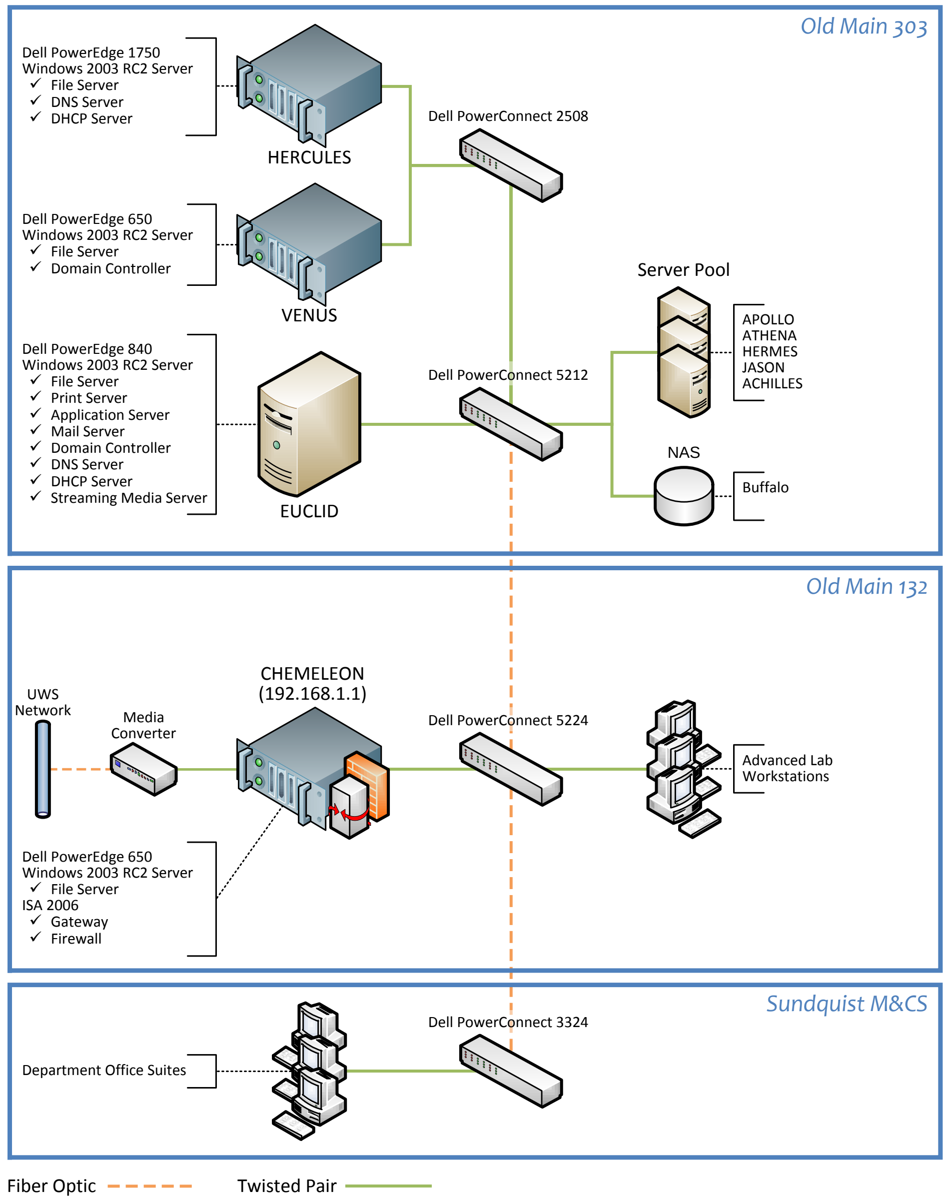

Figure 1: Diagram of the previous server and network infrastructure

The production servers employed Microsoft Windows Server 2003 as the host operating system. Most of the servers were assigned basic network support functions to include AD DS, DNS, DHCP, SQL, file, and document and print services. Linux servers with AMP applications were hosted on individual application servers and integrated into the network.

2.3 Processes and Curricular Support

The Advanced Computer Laboratory was primarily designed to support the department's computer science program. Although the department maintains five other academic programs, the lab was built with the intent to satisfy the educational needs of computer science students and faculty.

As with most computer science programs, programming is at the heart of the curriculum. The lab provided students with access to tools suited for writing, debugging, and testing software. Several integrated development environments were installed on each workstation to support procedural, object-oriented, assembly, and scripted languages; various software frameworks; and database programming. Source code management tools enabled team development processes. Since the lab was isolated from the campus, web and database servers could be deployed for application development over a network.

The department's computer science curriculum also includes courses that emphasize various technologies and professional practices embraced by the discipline. Concepts related to networking, security, and cryptography could be demonstrated in the lab to supplement class lectures or reinforce theoretical constructs. Many student projects were developed in the lab facilities simply because the tools needed to design, develop, and prototype the system were readily available and could be tailored for the specific tasks at hand.

To a lesser extent, the lab allowed certain students to work with the underlying technology needed to operate the server and network infrastructure. Students serving as lab managers routinely setup and configured servers, cabled networked devices, tested connectivity, diagnosed server and networking problems, deployed domain controllers and network applications, and maintained group policies and user accounts. This exposed students to a few of the underlying practices employed in data centers albeit at a smaller scale.

3 Transition and Transformation

3.1 The Buildup

The manner in which the department utilized the computer laboratory and the server and network infrastructure in many ways resembled a cottage industry to the main production capacity of the department. Although this configuration worked for a number of years, the need for a consolidated approach began to put pressure on the department from several different directions. However, the driving factor that led to this change was the construction of and eventual move to the new academic building—Swenson Hall.

Swenson Hall was proposed and designed to replace aging and antiquated teaching facilities and faculty offices distributed throughout Sundquist Hall, McCaskill Hall, and Old Main. The building was the third of three new facilities plus two major renovations the campus undertook since 2002. The Department of Mathematics and Computer Science was identified as a future occupant and members of the department were included in the planning process early in its development.

Eventually, the department was assigned a wing on the third floor of the complex that provided desperately needed floor space, centralized operations, and room to expand. In all, the suite of rooms included a new server room, hardware laboratory, student-teacher research room, and two advanced computer laboratories in addition to three department-priority classrooms. One of the advanced computer labs was located on the first floor of the building to provide students with 24-hour access.

Although the new facilities showed considerable promise, there was a real need to make a transition from the program-centric infrastructure to computing facilities that served the whole department. This was not clear at the time, but several indicators were present that included:

- Difficulty scaling the existing configuration and processes
- Desire to use computing technology in programs beyond computer science
- Creation of the Information Technology and Systems minor
- Silos of special interests and disconnected planning processes
- Concentrated ownership of existing facilities

In addition, the department needed adopt a broad vision of academic computing and reach out beyond department bounds to the campus community at large in need of advance computing technology.

A third factor contributed to the transition: technological and functional obsolescence. Only one of the six servers employed a 64-bit processor and had enough capacity to run a modern day operating system and applications. None of the remaining servers or workstations recast as application servers could be upgraded in a cost effective manner given their age. Also, a quarter of the computer workstations in the advanced laboratory could not be upgraded due to 32-bit architectures and another quarter were near the end of their useful life performance wise.

To compound the problem, key portions of the infrastructure were simply falling apart due to wear and tear. Essential server and network infrastructure components developed intermittent problems due to overheating since they were placed in closets that lacked proper environmental controls. Lack of maintenance, old carpets, and worn furnishing in the computer lab also contributed to equipment failures as dust and dirt entered cooling vents, clogged fans, and coated heat sinks and internal components.

Finally, management of the laboratories and network infrastructure increasingly lay in the hands of student assistants. Two factors led to this condition: First, the lab’s primary faculty architect accepted another position leaving a vacancy in their absence. Second, little documentation was available to communicate the system’s overall design and essential operational details. Fortunately, several very talented students were hired to help ensure the lab remained operational.

3.2 The Precipitating Idea

Despite all the challenges of the existing system, it was a student lab manager that proposed the initial idea to consolidate the server and network infrastructure using virtualization and failover clustering. Carl Berg was employed by the department to ensure the computing facilities remained operational and had successfully managed the lab through some of its most trying days. During mid-January of 2011, Carl proposed a design for server and network infrastructure that would transform the way the department viewed academic computing.

Before looking specifically at Carl's idea, it is important to establish the circumstances and the context in which the proposal was made. To begin with, the author had just started his term as Department Chair and the July-August 2011 move date to Swenson Hall was less than nine months away. Floor plans had been completed and construction crews were actively fleshing out the new superstructure. However, details of the computing laboratories details had yet to be worked out.

Matters were complicated since it was not just a direct move from the old facility to the new. First, the existing advanced computer laboratory (referred from here on out as the Development Computer Laboratory or Development Lab) was in need of an infrastructure upgrade. Second, the design of a second advanced computing laboratory (referred to as the Application Computer Laboratory or Application Lab) had yet to solidify. Budget shortfalls and funding uncertainties contributed to a series of back and forth discussions between the department, the Information Technology (IT) Department, and various administrative units on the design of the facility. Of course, a deadline loomed since proposals for the Classroom-Laboratory Modification Grant to fund the new system needed to be submitted by February 3, 2011.

During the period prior to Carl's proposal, the department lacked a cohesive approach to planning the new facility. Tracks had organized around each lab and were competing for the same resources. Decisions whether or not to make incremental upgrades or keep the two labs separate stemmed from past experience and different philosophies. It is important to note that this was not a contentious issue but rather a natural alignment into silos of interest among very collegial department members.

The basic premise of Carl's idea was to move away from dedicated systems serving each of the proposed new labs and adopt a centralized computing resource employing failover clustering and virtualizing the server applications. The system would employ a fault tolerant storage area network (SAN) using mirrored servers, a two-node virtual machine host cluster, and dedicated primary and secondary domain controllers. In effect, this would create generic hardware platform upon which servers could be created and deployed while retaining an individual's ability control of their own systems.

As Chair, the magic of Carl's proposal was that it created an avenue to consolidate a disjointed infrastructure and unify the department's computing facilities. Instead of a collection of piecemeal proposals advocating different approaches, a single vision for the computing facility could be put forth across all proposals. In addition, the vision included much more than the technology, but how the technology could be used to advance the department's academic mission and benefit the university.

Shortly after this insight, a meeting was scheduled with members of the IT Department on January 26, 2011, to determine the feasibility of such an option. The discussion produced a number of promising results including an alternate approach to implement the SAN using iSCSI storage array. The IT Department graciously offered to contact Dell and request a proposal for such a system in time for the classroom-laboratory modification grant proposal deadline. On January 28, 2011, a proposal arrived from Dell with a cost estimate for the system being considered.

Three classroom-laboratory modification proposals were submitted to the CIO for the committee's consideration on February 8, 2011, after a brief campus-wide deadline extension. The proposals were organized around the three major facilities: the application lab, the development lab, and the network infrastructure so requests could be distributed between two consecutive funding periods. Nonetheless, all proposals were framed as components of single academic computing vision. On May 3, 2011, the author received the results of the proposal with news that all but a modest portion of the server and network infrastructure would be funded.

3.3 The New Computing Facilities

During the proposal process, it was decided to populate the application and development computer laboratories with identical systems. Each lab received 24 student workstations plus one teacher workstation for a total of 50 workstations. This greatly simplifies management and allows additional flexibility in scheduling classes and labs between two identically configured labs. The useable workstations recovered from the old advanced computer laboratory were refurbished and upgraded for redeployment.

Each lab workstation consists of a Dell Optiplex 980 computer, a primary monitor, and a secondary monitor arranged in portrait mode as shown in Figure 2. Each computer is outfitted with an Intel i7 2600 processor, 8GB of RAM, an ATI Radeon HD 4550 video card with 512 MB of dedicated video memory, a LOM network adaptor, and a 500 GB hard drive. The primary monitor is a Dell U2410, 24-inch, 1920x1200 pixel, LCD panel while the secondary monitor is either a Dell 17-inch or 19-inch, 1280x1024 pixel, LCD panel. Workstations in each computer lab are arranged in groups of four across six rows aligned perpendicular to the teacher workstation for easy navigation. Several workstations are equipped with ADA compliant desks that can be adjusted vertically.

The server and network infrastructure provides a generic platform to virtualize and host application servers for classroom, laboratory, or scholarly activity as shown in Figure 3. To begin with, the support cluster provides necessary system services to the infrastructure and consists of two Dell PowerEdge R310 servers arranged in a failover cluster configuration. The servers are designated as integrated primary and secondary domain controllers for the system and host the AD DS and DNS roles. DHCP and iSNS are deployed as high availability applications on the mutual cluster.

Figure 2: Workstation configuration and a portion of the computer laboratory layout

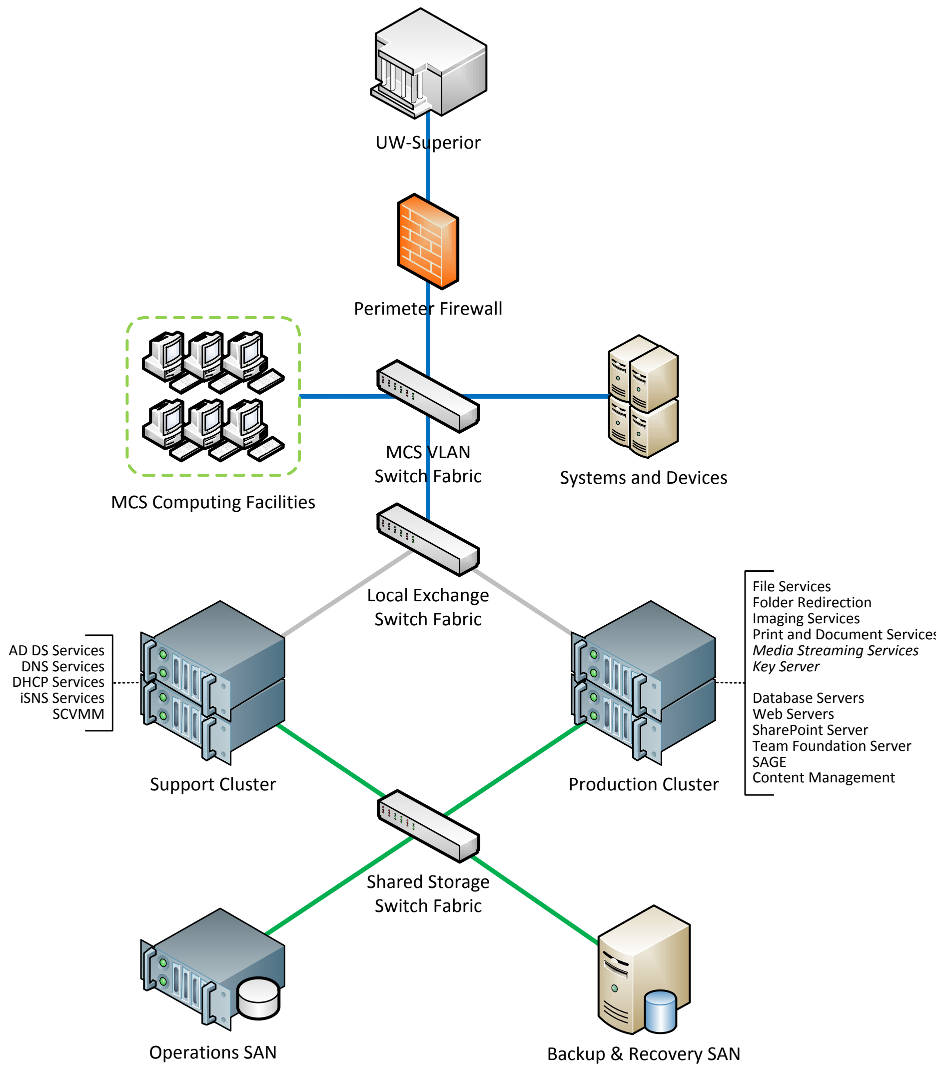

Figure 3: Diagram of the new server and network infrastructure

The production cluster hosts the virtual machines and consists of two Dell PowerEdge R610 servers arranged in a failover cluster configuration. Cluster nodes are connected to private networks that provide cluster communications and live migration for high availability applications. Both clusters are networked to the shared storage subsystem through a dedicated private network that connects a Dell PowerVault MD3220i iSCSI storage array for normal operations and a Dell PowerEdge 840 server configured as an iSCSI storage server for backup and recovery. The system is scheduled for installation beginning March 19, 2012.

The “sandbox” is a new addition to the computing infrastructure and consists of a pool of twelve personal computers, two workstations, network adaptors, hard disks, and various components suitable for experimentation and prototyping. The system is located in the server room along with the production system, but is isolated on its own private network. The sandbox allows students and faculty a place in which to explore and learn new concepts without disrupting the production environment. In addition, the sandbox also serves a mockup of the production system to test configuration changes before going live as shown in Figure 4.

Figure 4: Portion of the sandbox with a mockup of the production system

Overall, this new configuration represents a significant change from the prior system. The number of workstations increased threefold from 21 to 64 systems; whereas, the total number of servers decreased from 11 (six servers plus five PC-based application servers) down to five servers plus a dedicated storage array.

4 Implications for an Academic Department

4.1 Virtualization

Virtualization is often discussed in the context of enterprise environments that employ high availability services such as website, email, database, networking, and enterprise applications that require 24/7 availability. The benefits of virtualization are widely known and include server consolidation, improved hardware utilization, and centralized management. The question is, however, would server virtualization benefit an academic department?

Based on the experience of the author, the department has at least two distinct operating modes server applications must satisfy. The first mode is organized around the on-going academic operations of the department whose function is fairly consistent and needed in perpetuity. Whereas, the second operating mode coincides with periodicity of the academic year where server applications are deployed and utilized in synchronization with courses offered over a semester.

The use of virtualization software alleviates some of the challenges that arise when hosting server applications in a project-oriented academic environment. For instance,

- Cloning enables quick setup times using a pre-built virtual machine,
- Snapshots allow an application to reset or rewound to a known state,
- Difference or linked-clones enables multiple deployments from a single virtual machine,
- Exporting allows a virtual machine to be easily moved from one system to another, and
- Deleting the virtual machine removes all traces of the server application.

In addition, system resources can be scaled to match application demands or a virtual machine can simply be turned off when not needed freeing resources for other activities. These capabilities allow an academic department to leverage a generic computing platform in a variety of ways independent of the particular applications being hosted.

There are drawbacks, however, to incorporating a small-scale private cloud to host virtualized applications in an academic environment. First and foremost, there is a tradeoff between simplicity and complexity. A system that employs failover clustering is like an old mechanical watch—there are lots of little interconnected parts. This creates a fairly steep learning curve for both faculty and students new to these types of systems. To compound this challenge, a more formal documentation process is also needed to capture setup and configuration details so that nuances of the system can be communicated to successive system operators and managers. Regardless of the difficulties, these types of systems are common in the world of practice and both students and faculty need to be familiar with their operations.

4.2 Student Involvement

Students have always been very involved in supporting the computing facilities and infrastructure the department hosts. Historically, student assistants were employed to maintain the workstations in the previous computer laboratory and the dozen servers plus networking infrastructure located in adjoining facilities. Typical duties included managing user accounts, installing software on servers and lab workstations, troubleshooting hardware and software problems and defects, configuring and deploying systems and applications, and servicing equipment.

In this particular project, students contributed a great deal toward the design of the server and network infrastructure. As the one who proposed the original idea, Carl took a leading role in the design by prototyping the system and demonstrating proof-of-concept. In an independent study course that paralleled his lab management duties, he prepared an engineering model of the system consisting of a two computers networked in a failover cluster configuration connected to a third computer setup as a SAN using an open source application called FreeNAS. He documented the configuration and presented his findings at the end of Spring Semester 2011.

Efforts to work out the details of implementing the new system began as the design stage entered its final phase after the move to the new Swenson Hall facilities. Carl was joined by a new laboratory assistant, Aaron Walding, at the beginning of Fall Semester 2011 and together they documented and wrote a how-to manual detailing the tasks needed to configure the system. Later that same semester, Brentton Paulus joined the team to ensure continuity after Carl's graduation. All three participated in meetings between the department and the IT staff to coordinate preparations and helping layout a floor plan for the server room.

Looking toward the future, three factors will impact the activities computer laboratory assistants perform in the new computing facilities. First, the new system is much larger in scale than the prior system. Second, the system is substantially more sophisticated with the introduction of virtualization and failover clustering. Third, the computing facilities now serve more constituents and are much more visible to a broad range of students, faculty, administrators, and guests. Regardless, students will be involved in every aspect of the system since it is designed for academic purposes.

In addition, it is the author's desire to use the infrastructure as a tool to model enterprise operations and equip students with valuable professional skills. For instance, employing a configuration management process where changes are tested and evaluated in a sandbox environment first, then implemented in the production system using team deployment exposes students to accepted industry practices. Another option may include deploying a helpdesk system to track, prioritize, and manage service tickets in a manner similar to data center support operations.

5 Lessons Learned and Summary

The move to Swenson Hall triggered an avalanche of changes leading to a transformation of the department's computing facilities and infrastructure. As members of the department worked through the issues associated with the transition, a single idea put forth by a student served as a catalyst to unite a disconnected planning process and stimulated a new vision of academic computing.

The author learned three important lessons from this experience. First, circumstances occasionally require fundamental changes that lead to a departure from business as usual. Three years ago, the vision for the department's computing facilities was significantly different than what it eventually arrived at, thus highlighting the need to be flexible in thought and action. Although experience often provides the basis for decision making during complex changes, one must guard against the inertia experience creates if new opportunities and alternatives are to be pursued.

Second, take advantage of major transitions and use the change to rethink operations. Big changes—such as the move to the new academic facility—occur infrequently but often provide additional resources that otherwise would not be available. Change also serves as an opportunity to showcase new ideas and communicate a vision moving forward. The challenge for faculty and staff is overcoming the disruption change causes and seeing beyond the immediate problems.

Third, draw upon student creativity and insight and foster conditions that encourage this behavior in roles in which they serve. There are steps that can be taken to invite student participation in dimensions beyond routine activities associated with hosting a computer facility. Specifically, a department and supervisors can:

1. Listen to student ideas and create a framework that enables students to participate in the planning process
2. Incorporate unstructured time into the duties that give student assistants time to explore and investigate alternate approaches and new technologies
3. Make a pool of resources available that enable student assistants to experiment with and prototype systems

Creating the conditions where transformative ideas occur is a long-term effort that cultivates a good working relationship between students and department personnel.

This narrative captures a period of change for an academic department and its response to that challenge. Problems come at academic programs from every angle and departments that rely heavily on technological to fulfill their academic mission have additional factors to consider. This paper reminds readers to think broadly about circumstances and chart a path forward.