

Creating Robust IT Security and Efficiency by Reducing Infrastructure Complexity in Higher Education

GIAC (GSLC) Gold Certification

Author: Keith M. Lard, keith.lard@gmail.com
Advisor: Dr. Kees Leune

Accepted: May10th 2010

Abstract

IT systems and infrastructure have experienced a rapid change in complexity as a result of consolidating multimedia communication through one network (Internet Protocol convergence), moving from mainframes to client/server and implementation of web services. This has had a drastic effect on higher education IT resources. It has created a greater dependency on IT services as well as a disconnect between these innovations and the ability of the IT security professional to understand the entire stack well enough to secure it. The university's continued funding of antiquated infrastructure and systems has resulted in unnecessary expenses and hidden costs as it relates to security and efficiency. Reducing university system and infrastructure complexity by consolidating configuration, migration, management and maintenance of the hardware stack allows the organization to decrease the number of devices requiring setup, configuration and management while allowing for funding flexibility and realizing cost reductions. Further, system management efficiency allows the IT administrative and security professional to focus on more than just application configuration management strategies and user authorization potentially resulting in increased security. Implementation of virtualization through the Cisco UCS (Unified Computing System) coupled with Oracle's RAC (Real Application Clusters), VMware's enterprise software and SAN technology resulted in complexity and cost reductions. While specific security, efficiency and capacity goals were met, virtualization brought about new security issues.

INTRODUCTION

Recent economic conditions have created a business problem unique to higher education and its IT infrastructure. In the past ten years, IT systems and infrastructure have experienced a rapid change in complexity as a result of moving from mainframes to web services (Weinschenk, 2003). The technical landscape continues to become more complex as technology advances and application sophistication increases more rapidly, creating a greater dependency on IT services. To stay competitive and efficient, private and for-profit businesses have spent the last ten years keeping up with technology and training their staff. However, the university has been insulated in its own microcosm, having the luxury of ignoring business cycles, as the product offered has not changed drastically. Now, recent economic conditions and rapid advancement in technology have created the perfect storm within the university setting.

It has become imperative that the university finds ways to cut expenses to meet declining state budgets (CU AFR, 2009). In the private sector, businesses have focused on either cutting unprofitable products or reducing labor costs through options such as layoffs and outsourcing. While these methods work in the short term, a 2002 Price Waterhouse Coopers Study reports that financial executives admit these costs eventually “creep back into their businesses within two to three years” (p. 1). The university is faced with the issue of cutting operational expenses while trying to maintain a minimum level of infrastructure and quality of the product and services provided thus these short-sided solutions will not do. The luxury of ignoring the business cycle is gone and drastic measures are now needed to limit the number of resources that are cut.

Specific to this case study, two conclusions were drawn after primary data collection and analysis. First, vendor support contracts and maintenance renewal negotiations show that continuing to fund antiquated infrastructure and systems results in unnecessary expenses and hidden costs related to quality and efficiency. Second, skill set observation and coaching of IT personnel show a lack of training has created a disconnect between evolving technology and the ability of the IT administrative and security professional to understand the entire stack well enough to secure it. This disconnect has created an environment in which the IT administrative and security professionals are focused on application configuration management strategies and user authorizations rather than service to stakeholders and security of private data. The solution

is to reduce system and infrastructure complexity by consolidating configuration, migration, management and maintenance of the hardware stack, which will allow the university a strategic method to decrease the number of devices requiring setup, configuration and management while allowing for funding flexibility and realizing cost reductions simultaneously.

Overview

Virtualization can reduce system and infrastructure complexity. Further, virtualization can be implemented in phases allowing for capital expense spending flexibility. The purpose of virtualization is to increase efficiency and reduce costs, with each phase addressing different aspects of specific resources. Phase I should include hardware consolidation and virtualization of the web and application tiers along with commoditization of the database tier. Specific focus is on reducing hardware maintenance costs, EOSL (End Of Support Life) hardware and reducing power consumption. The main objective is to increase efficient use of CPU capacity and to create system flexibility and resiliency. Phase II should include implementation of desktop virtualization with a primary focus on change management issues and reducing the impact of human error. Further, an emphasis on training, organizational restructuring and group consolidation to support operational technologies implemented in Phase 1 should be incorporated. Phase III should include full data tier virtualization and disaster recovery focusing on full multi-site recovery in a cost effective manner.

Three Phase Virtualization

	PHASE I	PHASE II	PHASE III
TARGET RESOURCE AND OBJECTIVE	Network Infrastructure Web/Application virtualization	Human Capital Desktop virtualization And Human Capital Investment	Multi-site Infrastructure Disaster Recovery

	PHASE I	PHASE II	PHASE III
--	---------	----------	-----------

FOCUS	<ul style="list-style-type: none"> • Reduction in hardware and maintenance costs • Resolve EOSL issues • Improve system flexibility, resiliency and efficiency in power consumption and CPU capacity 	<ul style="list-style-type: none"> • Improve Efficiency in Time Management • Train/Retrain labor in new technology to reduce knowledge gap • Organizational restructuring to reduce labor costs 	<ul style="list-style-type: none"> • Fully capitalize on virtualization through integrating sites. • Reduce hardware and maintenance costs at disaster recovery site • Eliminate EOSL issues at Disaster Recovery site • Enhance disaster recovery by reducing down time and implementing full recovery
-------	---	--	---

Management will initially see the capital expense of this project as counterintuitive as the natural tendency is to cut expenses immediately. However, when taking into account an accumulation of hidden costs and unnecessary expenses, internal optimization becomes more cost competitive with alternatives such as outsourcing and can result in a more secure environment for data and improved quality of service. Therefore, one of the main hurdles to this conversion becomes a “human” rather than technical issue.

To manage both the technical and human side of the conversion requires a strong visionary leader. A strong leader will provide clear guidance for short and long-term goals and allows for both management and employees to move in the same direction. In the private sector, Fulmer (2007) found that strong leadership has been shown to be “the essence of exceptional organizational performance...[and the important traits of leaders include]...flexibility, collaboration, ability to leverage subject matter expertise and a willingness to continue their own learning journey” (pg. 12).

Leadership for planning and organizing the optimization project within the university should be no different. A broad based knowledge not only in technology (both hardware and software) but an understanding of finance, project management and even Human Resources are critical in demonstrating the value of the project investment as a long-term solution and carrying the project forward throughout the organization. Further, all these elements are necessary for gaining funding for a project of this size and nature and to keep the project on time. This often

seems like a tall order to fill. However, it is often out of crisis that a leader with “development and articulation of a leadership strategy” emerges (Fulmer, 2007 pg. 7). While managers are naturally risk-averse and strive for stability in the day-to-day tasks, a strong leader asks the tough questions, takes calculated risks and is willing to think outside the box with the best interest of the organization in mind. A leader can emerge from any level in the organization with real data and ideas on how to make the project work.

A visionary is almost a necessity in the current economic conditions. The State system begs the question of where costs should be cut and where funding can be found; nothing is off the table. However, some of the most critical cost saving measures may be less obvious. The IT professional can often lead the way in highlighting unnecessary expenditures and replacement or maintenance budget items that can be used for funding the first phase of complexity reduction and training. Further, budget forecast analysis shows that full cost recovery of the project could be realized within two to three years through operational expense savings in power, cooling, space, cabling and equipment maintenance costs. These savings can quickly offset the capital expenditures of the new hardware and software.

The importance of this must be emphasized considering the private sector has found that traditional cost-saving measures are neutralized within this time frame. Management’s natural reaction of reverting to cost-saving solutions such as outsourcing and a resistance to spend money on technical upgrades (Weinschenk, 2003) or security during budget crises has to be overcome. The hidden costs specifically related to outsourcing should be addressed, as this is often perceived as a quick, “safe” and less expensive solution and therefore implemented without analysis of the impact. Often, outsourcing leads to a less intimate process as the outsourcer does not truly understand the core business of the organization or the organization’s message gets lost in assumptions.

This is not to say that some traditional measures of cost savings should not be implemented, as eventually some will support the long-term strategy. Salary savings resulting from reduced complexity will happen and can be reinvested in the remaining staff in the form of specialized IT and security training. This allows the IT administrative and security professional to focus on more than application configuration management strategies and user authorization. Further, specific to this case study, specialized training helped to improve efficiency in service to

internal and external customers, tipping the scale in the cost/benefit analysis. Eventually, focused training and a more cohesive and efficient environment should result in cost savings and yield a better-trained work force with a more institutional security focus. The organization is no longer making short term cost savings decisions to appease the budget, but have now become better stewards of the budget as well as a better service provider and corporate neighbor to the community. All of which have become top priority goals under current constraints (UC AFR, 2009).

I. PHASE 1

1. PLANNING

The urgency in which cost saving must be found or created requires a well-orchestrated system for implementing these changes. Boyd's Law of Iteration says, "In implementing complex systems, it is better to act quickly than perfectly" (Sessions, 2008 p. 50). However, optimization requires new knowledge and a shift in the daily activity of the IT administrator and the security professional that results in the need for immediate education and training before the project begins. To fund these immediate training needs, the project should start with analysis of budget line items, compute capacity, power usage and inventory of existing infrastructure as well as review of vendor services and contracts.

Budget projections of reducing or eliminating line items such as desktop replacement, network improvements and midrange server support contracts show a possible cost-savings in the hundreds of thousands of dollars. Money for optimization can be found in budget items designated for data center improvements and plant funds. Further, now is the time to reconfigure procurement strategies and renegotiate services or contracts with vendors for support services and campus agreements as it relates to hardware and application licenses. Research shows that long-term financial sustainability and performance are greatly affected by poor procurement management within the university (UC Berkeley, 2010 pg. 80) and thus just as important as the physical hardware and software upgrades.

Budget review should also include an evaluation of the life span of infrastructure as it stands and then taken into account when looking at future costs. The university IT infrastructure

has not had the need to become faster, better or more efficient and as a result some equipment is reaching end of support. Further, IT infrastructure and the data centers are a massive drain on resources. They are in general expensive to run, requiring expensive real estate and mass power consumption (Gai, et al, 2010 p. 2). Private companies such as Google, Microsoft and Yahoo have all recognized this issue. And while the university data center may not take up the 470,000 square feet that the Microsoft Quincy center does or use the 47 MWatts of power (Gai, et al, 2010 p. 2), it should be analyzed for efficiency. Reducing the carbon footprint of the data center is environmentally friendly as well as fiscally responsible.

This reduced footprint is easily attainable when you look at the computing efficiency of the new consolidated hardware. It is well documented that most equipment in a data center often runs at five to ten percent capacity. Server virtualization can be optimized by increasing server loads from 25 to 50 percent while reducing power consumption by 14 percent (Gai, et al, 2010 p. 4). By increasing server loads, lightly loaded servers can be consolidated or eliminated altogether, resulting in less real estate needed and reduced power consumption. Finally, consolidation and elimination of lightly loaded servers reduces or eliminates physical cables resulting in improved airflow through the data center and increased cooling capacity.

Once data center, infrastructure, procurement and budget analysis is complete, optimization has won management approval and funding is found, it is time to plan for hardware, support and training for the project to begin. With a leading edge project such as this, it is often the vendors that will step up and include training and support before implementation. However, it is critical to understand the knowledge base of the team before hand and improve skills where needed. The knowledge base of the IT group now changes from network administration and desktop support to cloud computing, SaaS (Software as a Service) IaaS (Infrastructure as a Service) and PaaS (Platform as a Service).

In cloud computing, SaaS is often mistakenly defined as synonymous with “cloud”. However, SaaS is an application software service that allows customers a network-based mechanism to manage commercially available software via the Internet (CSA 2009, p. 15). It is an attractive alternative to installing and running applications on the customer’s computers as it allows remote access to applications via the Web, simplifies maintenance and support and eliminates the need for patches and upgrades. IaaS is an outsourced service that replaces the

customer's datacenter space, network equipment, hardware and software. It is a complete outsourced service of a platform virtual environment (CSA 2009, p. 16). PaaS or cloud platform services, is a complementary computing platform for cloud infrastructure and applications that reduce the complexity and costs associated with hardware and software layers (CSA 2009, p. 16).

These are all drastic changes for the university IT system and thus this paradigm shift brings with it a fear of the unknown and a need for people to take a leap of faith by abandoning legacy knowledge and dependence on how things are done. Management support and encouragement to seek and use training opportunities is crucial. Leadership and management should use incentive and reward programs to promote this new shift and publicly recognize those that are embracing what they learn. Grasping the new technology needs to come before implementation in order to succeed. It is here that gaining "followers" is more important than directing "subordinates".

To ease information overload a cohesive optimization system and process was chosen. For this case study, using Oracle RAC (Real Application Clusters) on the Cisco Unified Computing System with VMware enterprise software works to address two goals. First, this system eases the real time issue of a knowledge base shift through a cohesive and simplified deployment (Cisco, 2010 pg. 3). Second, high performance is cost effective, reducing budget concerns. Database administrators are now relieved of the time consuming task of configuring each hardware element in the stack independently (Cisco, 2010 pg. 4) by virtue of the Cisco UCS Manager. Further, the UCS Manager consistently and accurately configures servers through service-based profiles thereby eliminating human error almost completely and simplifying the scaling process. Finally, the "wire once" configuration simplifies cabling in the rack (Cisco, 2010 pg. 4) as network cabling is done once with changes implemented through software, reducing the potential error rate in the cabling process. Reduction in human error alone allows for better-managed and more effective work time resulting in higher labor efficiency.

Efficiency and performance through the Cisco UCS is also realized through the Cisco partnership with Oracle RAC. Oracle has a "mission critical" focus in their database and application technologies (Cisco, 2010 pg. 38) which are highly relevant to large corporations as well as the university. The Cisco UCS memory capability coupled with a scalable storage

system opens the performance floodgates to Oracle database environments allowing for unheard of work and performance loads (Cisco, 2010 pg. 38). Previous workload performance testing “included a realistic mix of OLTP and DSS workloads, which generated a sustained load on the eight-node Oracle RAC configuration for a period of 72 hours... far exceed[ing] the demands of typical database deployments” (Cisco, 2010 pg.38). The results showed that processors were barely at 50 percent capacity and port utilization was at 40 percent. The expansion potential of the system allows the university a cost effective way to scale up, if (or when) the need arises.

This elasticity potential needs to be highlighted. It is not uncommon for a university to have several “outposts” or campuses. The new infrastructure has the potential to ease procurement and efficiency problems across campuses. Shared services or “in-sourcing” through this project can be a better and less expensive option compared to well-known alternatives such as outsourcing. Further, with successful completion of the project, the university would be well ahead of most outsourcing companies in terms of infrastructure and capability.

Innovation is proving that outsourcing as it is currently known will eventually be irrelevant. Optimization and virtualization is the path that technology is following. Outsourcing companies will eventually have to change to cloud computing to stay competitive. This will be costly in the form of their infrastructure upgrades resulting in smaller outsourcers being swallowed by larger corporations such as Google and IBM (Overby, 2008 pg.2). These costs, both measurable and hidden, will trickle down to customers. The implications of outsourcing while on the cusp of change include higher costs for service, security risk to data and outages as well as the organization’s risk of lost revenue during outages. It is at this time the university can make the choice to be a leader in the industry or to be kept at the mercy of the outsourcing company. The organization can become world class and support their message or continue to be outdated and suffer the implications when they arise.

2. IMPLEMENTATION

2.1 Hardware

Based on original lab testing by Cisco, assumptions can be made and achievable goals set

surrounding existing hardware and the implementation of the new hardware system. Goals around initial testing and current hardware include power consumption reduction, equipment efficiency and cost reductions. Original hardware currently includes 290 servers, 24 racks and 3,500 square feet of space. Testing of original hardware proved that CPU capacity was running at the typical five to ten percent. Power consumption test results showed usage at a continuous 124 KW. Annual costs related to power, real estate and maintenance currently run at approximately \$400,000 (\$250,000 for space and maintenance, \$150,000 for power) while server maintenance runs another \$400,000. The objective is to eventually reduce equipment to 30 servers, 4 racks and 120 square feet of space while reducing costs to \$120,000 for space and power and \$40,000 for server maintenance. Further, it is assumed that power consumption will be reduced by 75% and CPU capacity increased to 50% after full implementation.

Tests were done to compare efficiency and costs after Phase I hardware implementation and configuration of 7 servers and 2 racks with 20% of virtual machines running. The test plan (Appendix A) is meant to outline the various tests and procedures that were implemented during the Accelerated Deployment (AD) of the Unified Computing System platform undertaken by the University of Colorado UIS team during September 07 to October 01, 2010 (four weeks duration). Test results showed CPU ran at 10% and power consumption at approximately 12 KW. The results were within pretest assumptions allowing for project progression.

2.2 Applications

Testing for optimization and upward capability can cement the university's destiny in the cloud-computing world as a leader. Further, test results will help lay groundwork for phase three of the project (full data tier virtualization) and possible shared services. However, the main objective of initial testing of the UCS / RAC concept should first be to simulate failures under realistic conditions, proving that the system as a whole is resilient, and that the new system does not fundamentally change the operation of legacy applications.

Application virtualization begins with the implementation of vMotion, DRS (Dynamic Resource Scheduling) and HA (High Availability) through the VMware enterprise plus software package. This allows for mobility of applications without impact of application function, as well as dynamic resource allocation and quicker failure recovery. More specifically, vMotion "allows you to:

- Perform live migrations with zero downtime, undetectable to the user.
- Continuously and automatically optimize virtual machines within resource pools.
- Perform hardware maintenance without scheduling downtime and disrupting business operations.
- Proactively move virtual machines away from failing or underperforming servers.”
(VMware.com, 2010, Products Section, vMotion)

Further, VMware DRS as a resource management mechanism offers performance, scalability and availability beyond physical infrastructure possibilities. DRS “improve service levels for all applications by continuously balanc[ing] capacity [to] ensure that each virtual machine has access to appropriate resources at any point in time” (VMware.com, 2010, Products Section, DRS). In addition, DRS will “easily deploy new capacity [and] seamlessly take advantage of the additional capacity of new servers added to a resource pool by redistributing virtual machines without system disruption” (VMware.com, 2010, Products Section, DRS). DRS will also automate server maintenance by migrating virtual machines off physical servers, allowing zero downtime for server maintenance thereby increasing productivity by allowing system administrators a mechanism to more effectively manage infrastructure. Finally, VMware HA protects against hardware and system failover by recovering all virtual machines on new hardware in the event of failover. It “monitors virtual machines to detect operating system and hardware failures, restarts virtual machines on other physical servers in the resource pool without manual intervention when server failure is detected and protects applications from operating system failures by automatically restarting virtual machines when an operating system failure is detected” (VMware.com, 2010, Products Section, HA).

Implementation of VMware built in services and initial testing of the VMware package in real time was successful in proving that legacy applications functioned without impact. In addition, DRS monitored applications and dynamically added resources to the application when specified applications reached the predefined resource limit. Also, testing of HA recovered all virtual machines from failover to new hardware within three to five minutes. The test plan developed and used outlines the various tests and procedures that were conducted during the VMware testing phase of the application installation (Appendix B). The enterprise package has proven to be resilient and thereby pushing the project to phase 3 of implementation.

2.3 RDBMS / Database (RAC)

The implementation of RAC on the RDBMS (Relational Data Base Management System) addresses three issues of particular importance. First, the non-RAC RDBMS (specific to the university setting in testing) consists of five total servers each running up to fifty databases. The database servers are isolated from each other and are not scalable beyond their current capacity. RAC RDBMS allows both the databases and database servers to be managed from a central location thereby creating flexibility not only in day-to-day database management but also scalability when the need arises. Second, additional databases require additional servers that have to be purchased. Each of these current database servers cost \$300,000-\$500,000 with an additional \$20,000 in annual maintenance. On the other hand, required servers for RAC RDBMS run approximately \$10,000 each with \$500 in annual maintenance costs thereby reducing hardware scalability costs dramatically. Finally, all five current non-RAC servers will reach EOSL (end of support life) by May 2012 (three by May of 2011 and two by May 2012) creating a reliance on third party support vendors and used hardware for mission critical applications. By replacing expensive, antiquated servers with modern, scalable equipment the university not only avoids unnecessary issues related to EOSL but also supports its mission critical applications in a more cost effective manner.

Testing of RAC RDBMS was based on the CISCO hardware failure test plan for UCS/RAC implementation. It included stressing the system with open source tools and variable manual disconnects within the system. These tests proved in all fault scenarios that overall there was no detectable disruption in the Interconnect/Ethernet or FC/SAN traffic (Appendix C). The RAC RDBMS implementation and testing showed that the system is predictable, reliable and resilient in managing the databases and database servers thereby protecting functionality of mission critical applications. With success of the final phase of implementation and testing done, the project can move to Phase II focusing on change management and desktop virtualization.

PHASE II

Keith M. Lard, keith.lard@gmail.com

1. Organizational Restructuring

With hardware and software in place and testing done, it is time to push phase two of optimization. Technically, this is the least difficult phase as the focus is on desktop virtualization. However, this can be the most difficult phase for management given the current economic conditions. Virtualization reduces desktop replacement costs, power consumption and costs related to desktop support. However, desktop virtualization equates to labor reduction and salary savings. This often means people that have been diligent workers for upward of fifteen years are now not necessarily needed. Management must now reduce their labor force. This process is more straightforward in the private sector. Unfortunately, the state system and how layoffs are handled is limited and often difficult to maneuver. Difficult or not, management must accept and being willing to commit to implementing and incorporating an organizational restructuring process in phase two for financial sustainability and growth.

Organizational restructuring has to start with an assessment of the future state of the organization and the need to redefine roles. Personal labor observations over fifteen years dictate that one desktop support person can reasonably manage 150 users. Based on this ratio, labor projections specific to this case study show that desktop support staff could be reduced by 40 percent. However, management may have to retain and retrain some of these individuals in new mobile technology as it is becoming more prevalent. The increased use of mobile computing would shift these employees to a shared helpdesk service group. Depending on the university's commitment to mobile or telecommuting, salary savings will vary. In some circumstances, salary savings may not be realized at all. This is dependent on the how the university handles state employee job regulations. If this system has been manipulated or abused, layoffs can become a shell game.

In the private sector, layoffs are simple. Notice is given and said employees are no longer employed; some may be given a severance. However, the state system by statute allows employees to retain property rights to their jobs. The implications can be dramatic. Management could be faced with positioning someone in a role that they are not qualified for due to these rights and seniority. Management could ask those close to retirement to resign but often this also means providing a severance or some other monetary incentive. On the other hand, management may be forced to lose a highly qualified employee due to the rights and

longevity of others.

While this seems an almost hopeless position for management to be in, it is manageable. However, it requires human investment, planning and patience. Phase three of optimization can be a spring board to catapult newer employees into new technology and specialized positions all but eliminating seniority and property rights to specific positions. The notion being that new technology creates new positions with specialized skills and requirements for the job. The motivation for current employees will be to embrace training, technology and apply for positions based on knowledge and merit (not politics) or risk lay-off.

This strategy will require management to do business differently. Eliminating legacy roles and creating new ones is time consuming in the state system and push back from management is difficult to overcome. The time consuming task of redefining jobs, writing position requirements, thumbing through current resumes and matching current qualifications to new positions is not what management wants to do. However, it is now necessary, as the IT administrator role has changed to be functional at multiple levels within the stack. In the university setting, the network administrator, the DBA or the desktop administrator has not cross-trained but rather stayed very specialized in tasks and training throughout the years. The depth of understanding needed to make decisions up or down the entire stack will take a strong commitment to training from not only management but also the employees themselves. Those not willing to retrain will have to be let go. The university can use this as an opportunity to engage, empower and promote those employees that have been successful in learning and collaborating during the project while eliminating the ones that have been barriers to an improved organization. The project has been an opportunity to build a cohesive team that can cross train and work up and down the stack individually as well as together.

On the other hand, the impact of retraining is a concern. Highly specialized training in virtualization results in employees being in higher demand and possibly leaving the university setting to pursue career options in the higher paying private sector. Management needs to build a strategy around retention of these individuals to avoid the revolving door scenario. Budget concerns prove difficult in monetarily rewarding those that have excelled. However, funding incentives through salary savings may be possible. Shared services, shared vision and a renewed but different sense of ownership through innovation should emerge while legacy knowledge and

experience are no longer relevant.

2. Desktop Virtualization

The second half of phase two is the actual implementation of desktop virtualization. The end result being reduced complexity and risk from remote access, reduced power consumption, and reduction in endpoint risk. A reduction in effort and risk from patching and reducing sensitive data protection efforts and risks should also be evident. This is realized through leveraging knowledge from the consolidation and virtualization put in place during phase I.

With desktop virtualization, power consumption is assumed to drop 75% based on data and test results taken from phase I. Reduced effort and risk as well as increased host security is assumed since there is no longer a requirement for the end user to interact with their local host to assure patching. Further, due to a reduction in the data footprint, there is a reasonable assumption of increased data security, as the data no longer resides at the end point desktop or laptop.

On the other hand, desktop virtualization brings on new security concerns and an urgent need to refocus the IT administrative professional's daily activities. Old infrastructure had static security solutions that are no longer adequate in the virtual environment. Further, the IT administrator can no longer choose to work in the middle ground between the customer and the technology. The middle ground, previously known as desktop support, no longer exists. Administrators now have to work directly with the customer through the helpdesk or migrate to a more technical position. Virtualization brings with it new security concerns that provide the administrator this technical option. Securing the virtual environment is more complex and must include new planning and processes to the university security policy. The specified role to manage this process is now the virtualization administrator (Haletky 2009, p. 4) and not the network administrator. All risks and threats that could potentially affect the virtual environment have not yet been identified and thus the position is an opportunity for expanded training and focus. This should be considered in the restructuring of the organization.

III. PHASE 3

The outcome of phase III is to complete full data tier virtualization and focus on enhanced disaster recovery. Full virtualization completes implementation of phase I with all 30 servers and 4 racks in service. Once hardware is in place, the VMware SRM (Site Recovery Manager) application is utilized to provide immediate and total business recoverability. Past recoverability was a “pick and choose” process. In other words, in the event of a site going down, administrators physically had to decide what systems and/or applications would be brought back up during disaster recovery and then manually bring them back up. However, SRM allows for automatic and immediate full recovery of all systems and applications all the way down to the actual desktops. Further, site equipment becomes more efficient and better utilized through SRM by managing and monitoring site changes and then duplicating changes instantly. Finally, equipment costs for the disaster recovery site are decreased and the issues related to EOSL eliminated through the final UCS/RAC/VMware implementation.

CONCLUSION

Successful implementation of the virtualization project proves that the university can reduce complexity and costs while improving system capacity, flexibility, efficiency and security. Budget and timeline reviews proved the project can be completed within budget and on schedule. Cost reductions proved to be dramatic in hardware, support and power consumption. Further, the resiliency and flexibility of the entire stack has shown promise for future growth with no impact to functionality. Issues concerning support of antiquated infrastructure have been altogether eliminated.

Difficulties and the impact of virtualization have been limited to human labor, organizational restructuring, availability of product and a shift in security risks. Engaging employees in training and targeted tasks proved to be overwhelming at times due to the vast gap in knowledge. The knowledge gap between understanding the entire stack and new technology well enough to implement each phase was reduced but not eliminated. Extra training of network and security administrators as well as DBAs, etc. through vendor assistance was required. Information retention proved difficult and remedial training is still needed. It is hoped that through remedial training and the current knowledge base, organizational restructuring will be

simplified. However, management has to be cognizant of the current state system and employment structure as it can and will play a part in the restructuring process.

On a positive note, great strides were made in improving security. Security was enhanced through simplification of account management. Clear text authentication methods moved to secure PKI mechanisms to reduce risk related to super user credentials traversing the network in the clear. A reduction in the number of policy enforcement points allowed the remaining enforcement points to become more robust. It is important to highlight that this security enhancement could have been implemented without virtualization. However, the change was less complex with virtualization due to the elimination of legacy configurations and processes. Further, much of the human error that resulted in security issues was eliminated.

It is evident that the university is playing catch up in security and more advanced security improvements are needed. Further, it becomes clear that when implementing full virtualization, trading one security risk for another is inevitable. While the university has reduced risk from human error, it has no control over the impact of vendors lagging behind the technology that has been implemented. Vendors offering security solutions in virtual environments have not yet provided options beyond the traditional tools already being used (Haletky, 2009 p. 431). The risk from human error is almost eliminated with reduced complexity and automation. However, inter-node visibility is lost within the virtualized network and not regained until the information leaves the virtual environment.

Specific to this project is the need to set up monitoring on the RAC nodes. RAC nodes communicate on a “private” interconnect via the UDP protocol. When communicating via UDP, the nodes do not necessarily need to be aware of each other and there is no mutual “handshake” that occurs. This makes it much easier to eavesdrop and thus this channel must be protected from sniffing to reduce the eavesdropping risk. Identifying and monitoring these new risks in the virtual environment now becomes a full time position and a commitment to continued learning and development. Optimization and virtualization clearly help increase efficiency while reducing complexity and operational costs. However, the trade off is in unknown security risks and committing to human capital investment.

Appendix A

Summary

This test plan is meant to outline the various tests and procedures that will be tested during the Accelerated Deployment (AD) tests of the Unified Computing System platform being undertaken by the UIS team during September 07 to October 01, 2010 (four weeks duration).

Test Bed Topology

The UCS infrastructure being used for the AD test is located at UIS.

Test bed includes Fabric Interconnects (2x6120), UCS chassis (2x5108), four B250 M2 blades and two B200 M2 blades and one B440 M1 blade in the UIS AD configuration.

Section 1 – UIS Infrastructure Setup

Assumptions:

- UCS is racked, stacked and powered on.
- Admin Access to UCS environment

Section 2 – General UCS Testing/Demonstration

This section serves to demonstrate the unique capabilities of the Unified Computing System and the stateless computing paradigm. This portion of the test plan, while supported by the UIS staff, was developed by Cisco.

Test Case Section 2.1 – Service Instantiation

Test Case # 2.1a	Title	Create Component Template Building Blocks
	Purpose	Demonstrate Hierarchical Nature of UCS Service Profile Templates
Procedure	<ol style="list-style-type: none"> 1. Create WWN Pool 2. Create MAC Address Pool 3. Create UUID Pool 4. Create Pin Groups 5. Create WWPN Pool 	

Test Case # 2.1b	Title	Create Component Templates
	Purpose	Demonstrate Hierarchical Nature of UCS Service Profile Templates

Procedure	<ol style="list-style-type: none"> 1. Create vHBA Template 2. Create vNIC Template 3. Create Server Pool 4. Create Boot Policy
------------------	--

Test Case # 2.1c	Title	Create Service Templates and Service Profiles
	Purpose	Demonstrate Hierarchical Nature of UCS Service Profile Templates
Procedure	<ol style="list-style-type: none"> 1. Create Service Profile Template 2. Using Wizard, assign the various building blocks and I/O Templates to the Service Profile 3. Once Service Profile Template is complete, use it to create a resulting service Profile or create multiple service profiles 	

Test Case Section 2.2 – Bare Metal Service Installation and Operation

Test Case # 2.2a	Title	Install Operating System
	Purpose	Show the flexibility of Service-based provisioning
Procedure	<ol style="list-style-type: none"> 1. Assign the service profile to a physical blade 2. Using the KVM Console, observe the PNU OS boot process to re-program the ‘personality’ of the physical server 3. Using the Virtual Media option through the KVM console, associate an ISO image to the ‘CD’ drive of the server 4. Progress to the main install section of VMware vSphere 4.0 to show the boot to SAN process and/or Local disk boot in action. 	

Test Case # 2.2b	Title	Demonstrate Service Profile Mobility
	Purpose	Show the benefits of stateless computing through Service Identity vs. Server Identity

Procedure	<ol style="list-style-type: none"> 1. Visually confirm ‘personality’ information on <ol style="list-style-type: none"> a. MAC Addresses b. WWN’s 2. Shutdown UIS-ESX4 Service Profile 3. Disassociate UIS-ESX4 Service Profile from server 1/1 4. Associate UIS-ESX4 Service Profile to server 2/4 5. Once booted, re-confirm UISs ‘personality’ information <ol style="list-style-type: none"> a. MAC Addresses b. WWN’s c. WWPN’s d. UUID
------------------	--

Test Case Section 2.3 – UCS Infrastructure Failure Scenarios

Test Case # 2.3a	Title	Remove a Fabric Interconnect During Operation
	Purpose	Show the full HA capabilities of the UCS platform
Procedure	<ol style="list-style-type: none"> 1. From an associated server instance, start a continuous ping test to a remote system to test IP connectivity 2. From an associated server instance, start a large file copy to the array to test Fiber Channel connectivity 3. While the above operations are running, physically remove a Fabric Interconnect to break ½ of the connectivity paths 4. Observe resulting actions 	

Test Case # 2.3b	Title	Disable a Fabric Interconnect During Operation
	Purpose	Show the full HA capabilities of the UCS platform
Procedure	<ol style="list-style-type: none"> 1. From an associated server instance, start a continuous ping test to a remote system to test IP connectivity 2. From an associated server instance, start a large file copy to the array to test Fiber Channel connectivity 3. While the above operations are running, physically turn off a Fabric Interconnect to break ½ of the connectivity paths 4. Test connectivity to the UCS Manager application 5. Observe resulting actions 	

Test Case # 2.3d	Title	Remove a power supply connection
	Purpose	Show the ability of UCS to operate on limited power supplies
Procedure	<ol style="list-style-type: none"> 1. Remove redundant power in the following scenarios: <ol style="list-style-type: none"> a. Single blade powered up b. Two blades powered up c. Multiple (and so on) 2. Monitor power usage in UCS interface for each test 	

Section 3 – Specific UCS Testing/Demonstration

Test Case Section 3.1 – UCS Use Case Testing

Test Case # 3.1a	Title	Concurrent changes to critical areas of UCS by simultaneous users
	Purpose	Demonstrate UCS behavior when multiple users make changes
Procedure	<ol style="list-style-type: none"> 1. Create two service profiles at the same time using the same MAC, UUID, WWNN, and WWPN pools (use the same template) 2. Test behavior of UCS when two admins make RBAC changes to the same user account 3. Test behavior of UCS when two admins make differing QOS changes to the same profile 4. Test behavior of other critical items 	

Test Case # 3.1b	Title	Blade Failure Simulation test
	Purpose	Build VMware infrastructure using VMware vSphere 4 operating system on multiple blades – test blade failure scenerio

Procedure	<ol style="list-style-type: none"> 1. Install vSphere 4 on multiple blades using boot from SAN 2. Install Windows 2008 R2 as a virtual machine on one of the ESX4 host 3. Install vCenter software on this Windows Virtual machine 4. Create Data Center and Cluster in vCenter. Add various vSphere4 hosts and build various virtual machines (Windows and Linux) 5. Enable HA and DRS on the cluster 6. Ping VM's IP from network for those VMs in the blade to be removed 7. Remove blade and monitor connectivity and Virtual Machines functionality
------------------	---

Test Case # 3.1c	Title	Provisioning process
	Purpose	Document provisioning process of systems under UCS
Procedure	<ol style="list-style-type: none"> 1. Document provisioning process of a system from start to finish 2. Make note of efficiencies provided by UCS solution 3. Make note of required deviations from standard operational processes needed to accommodate UCS, and any workarounds needed 4. Compare documented process to standard provisioning process – Are their efficiencies afforded on both the infrastructure AND operational sides? 	

Test Case # 3.1d	Title	Implement Microsoft clustering
	Purpose	Demonstrate Microsoft clustering on UCS
Procedure	<ol style="list-style-type: none"> 1. Install and configure Windows 2008 w 2. Install and configure SQL Server and clustering software 3. Test failover of nodes. 	

Test Case # 3.1e	Title	Local RAID Disk Relocation
	Purpose	Test movement of existing RAID set with data to different blade

Procedure	<ol style="list-style-type: none">1. Associate SP to blade using local disk policy of RAID12. Create data on local disks (install OS, create file-system, etc)3. Physically remove blade associated with SP from chassis4. Move disks from original blade into a different blade removed from chassis5. Place blade with moved disks into chassis/slot at original location of removed disks6. Accept the new blade in the slot location7. Monitor SP startup and break into LSI Configuration Utility8. Select SAS adaptor from list9. Select RAID Properties -> Manage Array10. Select Activate Array, waiting until Status is Optimal Select Synchronize Array and wait until complete11. Exit LSI utility and monitor SP startup12. Confirm data access
------------------	---

Appendix B

Summary

This test plan is meant to outline the various tests and procedures that will be tested during the VMware testing phase of the application installation. Test vMotion, VMware Dynamic Resource Scheduler, and VMware High Availability

System	Test	Objective
Cognos 8.4	Functional – usage	Observe normal system operation
InfoEd	Functional – usage	Observe normal system operation
PS HR	Functional – usage	Observe normal system operation
PS CS	Functional – usage	Observe normal system operation
PS CR	Functional – usage	Observe normal system operation
Web Servers	Functional – usage	Observe normal system operation
Open Filer	Functional – usage	Observe normal system operation
Job Invocation – UC4 to invoke SQRs, cobols, App Engine, PS agent.	Functional	See that UC4 is able to invoke these jobs

Test Case Section 1.1 – VMware vMotion

Test Case # 4.1	Title	Configure applications to use VMware vMotion
	Purpose	Demonstrate the capabilities of vMotion. Validate all hardware in the VMware “datacenter” is compatible and capable of vMotion.
Procedure	<ol style="list-style-type: none"> 1. Create application stacks with web and application tier <ol style="list-style-type: none"> a. configure and test Cognos 8.4 b. configure and test InfoEd c. configure and test PeopleSoft HR d. configure and test PeopleSoft Finance e. configure and test PeopleSoft Campus Solutions f. configure and test web servers g. configure and test Open Filer h. configure and test UC4 2. Create functional testing scenarios for each application that can be run continuously and monitored in real time 3. While testing VMotion both the web tier and the application tier and validate that the application still functions correctly 	

Test Case Section 1.2 – VMWare Dynamic Resource Scheduling

Test Case # 4.2	Title	Configure applications to use VMware DRS
	Purpose	Understand and demonstrate the capabilities of DRS.
Procedure	<ol style="list-style-type: none"> 1. Create application stacks with web and application tier <ol style="list-style-type: none"> a. configure and test Cognos 8.4 b. configure and test InfoEd c. configure and test PeopleSoft HR d. configure and test PeopleSoft Finance e. configure and test PeopleSoft Campus Solutions f. configure and test web servers g. configure and test Open Filer h. configure and test UC4 2. Create functional testing scenarios for each application that can be run continuously and monitored in real time 3. Use vMotion to move all of the machines to 2 nodes and enable DRS on each of the application servers. Validate that DRS automatically migrates web and application servers to less busy nodes. 4. Place load on several servers in a single node and validate that DRS migrates additional servers to less busy nodes. 	

Test Case Section 1.3 – VMware HA (High Availability)

Test Case # 4.3	Title	Configure applications to use VMware HA
	Purpose	Understand and demonstrate the capabilities of VMware HA.
Procedure	<ol style="list-style-type: none"> 1. Create application stacks with web and application tier <ol style="list-style-type: none"> a. configure and test Cognos 8.4 b. configure and test InfoEd c. configure and test PeopleSoft HR d. configure and test PeopleSoft Finance e. configure and test PeopleSoft Campus Solutions f. configure and test web servers g. configure and test Open Filer h. configure and test UC4 2. Simulate hardware failure and validate that virtual machines get moved and restarted on available hardware. 	

Appendix C

UCS failure test plan and results based on recommended CISCO hardware failure test plan for UCS/RAC implementation.

Test Case # 5.1	Title	Single link failure public interface
	Purpose	Understand and demonstrate traffic disruption and node eviction
Procedure	<ol style="list-style-type: none"> 1. Let the system run workload for at least 20 minutes. 2. Disconnect the first 10 GigE public network connection from the first chassis. 3. Wait for 5 minutes and reconnect this connection. 4. Wait 5 minutes after the reconnect. 5. Disconnect the first 10 GigE public network connection from the second chassis. 6. Wait for 5 minutes and reconnect this connection. 7. Wait 5 minute after the reconnect. 8. Disconnect the second 10 GigE public network connection from the first chassis. 9. Wait for 5 minutes and reconnect this connection. 10. Wait 5 minute after the reconnect. 11. Disconnect the second 10 GigE public network connection from the second chassis. 12. Wait for 5 minutes and reconnect this connection. 	
Results	<p>The single network disconnection should not cause any disruption in the Interconnect/Ethernet or FC/SAN traffic.</p> <p>No node eviction, hangs should occur.</p>	

Test Case # 5.2	Title	All links failure public interface
	Purpose	Understand and demonstrate traffic disruption and node eviction
Procedure	<ol style="list-style-type: none"> 1. Disconnect all 10 GigE network connections on the first chassis, connecting it to the public network 2. Wait for 5 minutes and reconnect the network connections. 3. Wait 5 minutes 4. Disconnect all 10 GigE network connections on the second Chassis, connecting it to the public network 5. Wait for 5 minutes and reconnect the network connections. 6. Wait 5 minutes 	

Results	<p>Single IOM Failure. Nodes should be able to continue work over the other IOM. It should not cause any disruption in the Interconnect/Ethernet or FC/SAN traffic.</p> <p>No node eviction, hangs should occur.</p>
----------------	--

Test Case # 5.3	Title	Single link failure private interface
	Purpose	Understand and demonstrate traffic disruption and node eviction
Procedure	<ol style="list-style-type: none"> 1. Let the system run workload for at least 20 minutes. 2. Disconnect the first 10 GigE private network connection from the first chassis. 3. Wait for 5 minutes and reconnect this connection. 4. Wait 5 minutes after the reconnect. 5. Disconnect the first 10 GigE private network connection from the second chassis. 6. Wait for 5 minutes and reconnect this connection. 7. Wait 5 minute after the reconnect. 8. Disconnect the second 10 GigE private network connection from the first chassis. 9. Wait for 5 minutes and reconnect this connection. 10. Wait 5 minute after the reconnect. 11. Disconnect the second 10 GigE private network connection from the second chassis. 12. Wait for 5 minutes and reconnect this connection. 13. Wait 5 minute after the reconnect. 	
Results	<p>The single network disconnection should not cause any disruption in the Interconnect/Ethernet or FC/SAN traffic.</p> <p>No node eviction, hangs should occur.</p>	

Test Case # 5.4	Title	All links failure private interface
	Purpose	Understand and demonstrate traffic disruption and node eviction
Procedure	<ol style="list-style-type: none"> 1. Disconnect all 10 GigE network connections on the first chassis, connecting it to the private network 2. Wait for 5 minutes and reconnect the network connections. 3. Wait 5 minutes 4. Disconnect all 10 GigE network connections on the second Chassis, connecting it to the private network 	

	<ol style="list-style-type: none"> 5. Wait for 5 minutes and reconnect the network connections. 6. Wait 5 minutes
Results	<p>Single IOM Failure. Nodes should be able to continue work over the other IOM. It should not cause any disruption in the Interconnect/Ethernet or FC/SAN traffic.</p> <p>No node eviction, hangs should occur.</p>

Test Case # 6.1	Title	Single Fiber Channel link failure
	Purpose	Understand and demonstrate traffic disruption and node eviction
Procedure	<ol style="list-style-type: none"> 1. Let the system run workload for at least 20 minutes. 2. Fail the first FC connection from the expansion module, connecting to the first SAN switch. 3. Wait for 5 minutes and reconnect the FC connection. 4. Wait 5 minute after reconnect. 5. Fail the second FC connection from the expansion module, connecting to the second SAN switch. 6. Wait for 5 minutes and reconnect the FC connection. 	
Results	<p>The single FC disconnection should not cause any disruption in the FC/SAN traffic.</p> <p>No delays or hangs should occur.</p> <p>Monitor the multipath software for correct path fail detection and path failover</p>	

Test Case # 6.2	Title	Double Fiber Channel link failure
	Purpose	Understand and demonstrate traffic disruption and node eviction
Procedure	<ol style="list-style-type: none"> 1. Let the system run workload for at least 20 minutes. 2. Fail both FC connections from the expansion module, connecting the SAN Switches. 3. Let the nodes reboot. 4. During reboot, reconnect the FC connections. 5. Make sure the cluster start up again, after FC reconnection. 	

Results	<p>The whole cluster is expected to reboot due to storage disconnection.</p> <p>Make sure after reconnection, node are able to start up again.</p>
----------------	--

Test Case # 7.1	Title	Both IOM's on a single UCS Chassis fail
	Purpose	Understand and demonstrate traffic disruption and node eviction
Procedure	<ol style="list-style-type: none"> 1. Disconnect all 10 GigE network connections (both public and private IOMs) on either of the chassis. 2. Wait for 10 minutes and reconnect the network connections to both the IOMs. 	
Results	<p>All nodes in the failed chassis should reboot.</p> <p>The nodes in the other chassis should continue to work.</p>	

Test Case # 8.1	Title	Fabric Interconnect failure
	Purpose	Understand and demonstrate traffic disruption and node eviction
Procedure	<ol style="list-style-type: none"> 1. Let the system run workload for at least 20 minutes. 2. Fail either of the Fabric Interconnect by powering it down. 3. Wait for 5 minutes and power up the switch. 	
Results	<p>The FI disconnection should not cause any disruption in the Ethernet or FC/SAN traffic.</p> <p>No node eviction or hangs should occur.</p>	

Test Case # 9.1	Title	Single host failure
	Purpose	Understand and demonstrate traffic disruption and node eviction

Procedure	<ol style="list-style-type: none"> 1. Let the system run workload for at least 20 minutes. 2. Fail a single blade from the first chassis by pulling it out of the chassis. 3. Wait for 5 minutes 4. Fail a single blade from the second chassis by pulling it out of the second chassis. 5. Wait for 5 minutes 6. Move all blades back in place.
Results	<p>The removed nodes should be detected in the cluster within the CSS miscount (60 seconds) time.</p> <p>Instance recovery will happen for the failed blades.</p> <p>Make sure that the failure of a single blade does not cause hangs or delays in the Ethernet or SAN network due to reconfiguration.</p> <p>Make sure that when the blades startup again they join the cluster without problems.</p>

Test Case # 9.2	Title	Multiple host failure
	Purpose	Understand and demonstrate traffic disruption and node eviction
Procedure	<ol style="list-style-type: none"> 1. Let the system run workload for at least 20 minutes. 2. Fail 4 blades simultaneously by powering down the chassis. 3. Wait 10 minutes. 4. Power up the blade chassis. 	
Results	<p>The removed nodes should be detected in the cluster within the CSS miscount (60 seconds) time.</p> <p>Instance recovery will happen for the failed blades.</p> <p>Make sure failure of blades in a single chassis does not cause hangs or delays in the Ethernet or SAN network due to reconfiguration.</p> <p>Make sure that when the blades startup again they join the cluster without problems.</p>	

Bibliography

- Cisco. (2010). *Deploying Oracle Real Application Clusters on the Cisco Unified Computing System with EMC CLARiiON Storage*. White Paper. Cisco.com. Retrieved September 1, 2010 from http://www.cisco.com/en/US/prod/collateral/ps10265/ps10280/white_paper_c11-562881.pdf.
- CSA (Cloud Security Alliance). (2009). *Security Guidance for Critical Areas of Focus in Cloud Computing V2.1*. Retrieved on November 6, 2010, from <http://www.cloudsecurityalliance.org/guidance/csaguide.v2.1.pdf>
- Fulmer, Robert M. and Bleak, Jared L. (2007). *The Leadership Advantage: How the Best Companies Are Developing Their Talent to Pave the Way for Future Success*. New York, NY: AMACOM. Retrieved June 27, 2010, from <http://www.safaribooksonline.com>
- Gai, Silvano; Salli, Tommi and Anderson, Roger (2010). *Cisco Unified Computing System (UCS) (Data Center): A Complete Reference Guide to the Cisco Data Center Virtualization Server Architecture*. Indianapolis, IN: Pearson Education, Cisco Press. Retrieved June 7, 2010, from <http://www.safaribooksonline.com>
- Haletky, Edward. (2009). *VMware vSphere™ and Virtual Infrastructure Security: Securing the Virtual Environment*. Prentice Hall. Retrieved June 12, 2010, from <http://www.safaribooksonline.com>
- Overby, Stephanie. (2010). *The End of IT Outsourcing As We Know It*. CIO.com. Retrieved August 12, 2010 from http://www.cio.com/article/603075/The_End_of_IT_Outsourcing_As_We_Know_It
- Price Coopers Waterhouse. (2002). *Myth of Corporate Cost-Cutting Revealed—Financial Executives Admit That Cost Savings Back Within Three Years*. London. Retrieved November 6, 2010, from <http://www.prpllc.com/corporatemyth.pdf>
- Sessions, Roger (2008). *Simple Architectures for Complex Enterprises*. Redmond, WA: Microsoft Press. Retrieved June 7, 2010, from <http://www.safaribooksonline.com>
- University of California, Berkeley. (2010). *Operational Excellence Phase I: Diagnostics*. Retrieved July 10, 2010, from <http://berkeley.edu/oe/phase1/>
- University of Colorado Office of University Controller. (2009). *Annual Financial Report 2009*. Retrieved September 10, 2010, from https://www.cu.edu/System_Controller/financial-rpts.html
- Weinschenk, Carl. (2003). *Hardware Today: Moving From Mainframes to Modularization, the Shrinking Server Footprint*. ServerWatch.com. Retrieved July 26, 2010, from <http://www.serverwatch.com/hreviews/article.php/2247251/Hardware-Today-Moving-From-Mainframes-to-Modularization-the-Shrinking-Server-Footprint.htm>