

THE IMPACTS OF A NONPROFIT CANCER RESEARCH ORGANIZATION MOVING TO A PRIVATE CLOUD ENVIRONMENT

Matthew R. Kisow, D. Sc., NSABP Foundation, Inc., matthew.kisow@nsabp.org
John J. Scarpino, D. Sc., Robert Morris University, scarpino@rmu.edu

ABSTRACT

In 2006, the Information Technology (IT) department at the operations center of one of the nations oldest and most respected clinical trial cooperative groups located in Pittsburgh, Pennsylvania needed to shift from a non-redundant one-server/ one-application paradigm to a highly available and vastly redundant private cloud delivering computing as a service rather than a product. To achieve this goal they implemented a virtualization solution based on Hewlett Packard hardware, and VMWare's ESXi virtualization technology, consolidating seventy-five percent of their infrastructure, and developing a solid foundation for their private cloud.

Keywords: Private Cloud, Case Study, Quality Assurance, Quality Control

INTRODUCTION

The IT department at the operations center of the National Surgical Adjuvant Breast and Bowel Project (NSABP) Foundation was at a crossroads. With more than a fifty-year history in designing and conducting clinical trials that have dramatically changed the treatment and prevention of breast and colorectal cancer, the NSABP located in Pittsburgh, Pennsylvania is among the nations oldest and most respected clinical trial cooperative groups, [10]. In 1999, however the NSABP Foundation began conducting clinical trials for the biotechnological industry. Since the inclusion of the biotechnology industry sponsored clinical trials, the NSABP's reliance on IT resources began to grow. In 2006 after several years of continued growth combined with the operational complexities and resource utilization deficiencies of the one-server/ one-application paradigm the IT department looked toward to a high availability/ high redundancy private cloud.

This case study analyzes the NSABP information technology department's decision to implement a dynamic server infrastructure using VMWare's private cloud technologies, highlighting the technical challenges and success of this project.

Problem Statement

The operational complexities and resource utilization deficiencies of the one-server/ one-application paradigm exhausted costly primary storage resources and, produced delays in the provisioning of compute resources, but the development of a private cloud environment has streamlined the data management and storage processes, and eliminated costly delays in server provisioning.

Purpose of Study

The purpose of this case study is to provide a thorough analysis into how the information technology department managed its increase in operational demands, and increased its business continuity framework by migrating its existing one-server/ one-application infrastructure into a private cloud environment.

Research Questions

This case study sought to discover what operational efficiencies including the Quality Assurance (QA) and Quality Control (QC) benefits an organization would realize by moving to a private cloud environment. Specifically, this case study examined the following research questions:

1. What benefits does moving to a private cloud realize?
2. What is the ROI for moving to a private cloud?
3. What QA and QC benefits does moving to a private cloud realize?

Limitations

Research conducted by Merriam, and Hamel point out that case studies involve issues of reliability, validity and generalizability because of their lack of rigor in the collection, construction, and analysis of data [6, 8]. The researchers selected a case study design based on the nature of the research problem and the questions being asked. Case study was deemed the best way to answer the research questions for this study.

BACKGROUND: A REVIEW OF LITERATURE

Introduction

An analysis of resource utilization across its existing standalone servers revealed that each of the NSABP's twelve servers were using resources at an average of thirty percent or less. It was beginning to become evident that the one-server/ one-application paradigm was not making sense; of greater concern to IT department managers was their reliance on tape backup for their disaster recovery method.

A change in the competitive environment brought about by the technological shift from the one-server/ one application paradigm has fundamentally changed the role that technology plays within an organization. The success of this paradigm shift can be described by using the "Process Oriented Model of Business Value" as proposed by Mooney, Gurbaxani, and Kraemer [9]. This framework indicates that as IT continues to permeate the organization, it has a greater impact on the processes and eventually on the organization as a whole.

A Paradigm Shift

The IT department's management quickly realized that the one-server/ one-application paradigm would not meet the future needs of the organization. A review of their existing resource utilization revealed that on average the one-server/ one application paradigm was underutilizing each of their standalone servers by an average of seventy percent. The organization was losing funds each time a new server was deployed under this old paradigm. According to a 2003 Intel White Paper the underutilization of the servers CPU horsepower, memory, and I/O operations cost an organization an estimated \$60/tpm (Transactions Per Minute) [13].

For the NSABP information technology department an important driving force behind a private cloud infrastructure was cost. "We needed to look at the delivery of computing as a service rather than the one-server/ one application delivery of a product or application," stated a senior technology manager in the IT department. The senior technology manager went on to say, "Every time we added a new application or service that required isolated computing resources, we ended up spending upwards of \$5,500 on average".

The IT department was concerned over the scalability and up front, costs associated with a private cloud. One concern was their ability of making the system robust enough to withstand a systemic failure while continuing to maintain the 24x7x365 availability that their clients demanded. They were able to accomplish this by using industry standard hardware provided by an industry leader, Hewlett Packard and their line of BL680 blade servers [4, 7, 12, 15, 16]. Hewlett Packard offered the scalability and redundancy that was required to make the project a success. Hewlett Packard products have enabled the IT department to continue to scale their private cloud environment while offering a higher ROI over the one-server/ one-application paradigm [7].

A senior technology manager noted that, "since we placed our private cloud into production we have virtualized 60% of our one-server/ one-application servers and have been able to spin-up several others to meet the needs of our

changing environment.” The demand for computing resources has increased, and the IT department has been able to meet that demand with minimal cost when compared with the one-server/ one-application paradigm.

VMWare ESXi

VMWare’s ESXi environment is the only thin hypervisor architecture that does not depend on a general purpose operating system, and is purpose-built for virtualization, and virtual machine recovery [17]. The term hypervisor, or virtual machine manager (VMM), is a hardware virtualization technique that allows multiple operating systems, termed guests, and to run concurrently on a host computer.

The hypervisor, or VMM presents to the guest operating systems virtual hardware, and an operating platform that manages the execution of the guest operating systems. Multiple instances of a variety of operating systems may share the virtualized hardware resources. VMWare’s ESXi hypervisor is commonly installed on server class hardware, providing the ability to run many guest operating systems at the same time [17].

Disaster Recovery

A greater concern to IT department management, however, was the organization’s antiquated disaster recovery model, which had relied upon a de facto standard 1950’s style tape backup and recovery model. Although the NSABP could have consolidated all of their initially targeted 12 servers onto two HP BL680 blade servers, they chose to deploy a third similar server for n+1 redundancy and disaster recovery. VMWare’s ESXi environment is the only thin hypervisor VMM (Virtual Machine Manager) architecture that does not depend on a general purpose operating system, and is purpose-built for virtualization, and recovery [17]. This environment plays a vital role in this scenario since it allows image files of the virtual machines to be stored on a backup server or public cloud storage for near instantaneous failover and recovery.

The IT department wanted the flexibility of a disaster recovery model that would eventually include hosted cloud services that would meet the needs of its internal customers as well as the regulatory demands placed upon them by the Food and Drug Administration (FDA), and their biotechnology industry partners.

The organization’s disaster recovery initiative has recently been enhanced by a third party disk-to-disk backup solution, Veeam. Veeam extends the disaster recovery capabilities of ESXi increasing the recovery time of a virtual machine by up to seventy-seven percent [5]. Research conducted by Veeam, and released in a 2010 white paper that was updated in 2011 highlights industry concerns over the recoverability of virtual guests. Veeam’s research surveyed 500 CIO’s from organizations with 1,000 or more employees from around the world. The research concludes that there is a growing trend in industry to transition mission-critical applications, and services to virtual guests. Interestingly however, the research indicated that almost half (44%) of respondents in the study indicated hampered efforts towards enterprise wide virtualization, citing fears surrounding the ability to successfully back up and recover the virtual guests [3, 14].

The disaster, and recovery quality assurance (QA), and quality control (QC) validation conducted by the IT department concluded that using the Veeam recovery software has enhanced their business continuity model. Veeam has, “increased our operational flexibility, by allowing us to immediately recover a file or entire virtual guest without waiting,” stated a senior IT manager. This near instant recovery has enabled the IT department of the NSABP to increase the return on investment (ROI) of their recovery over tape backup.

ROI and Operational Flexibility

The Information Technology (IT) department of the NSABP had a problem. How could they increase and maximize their computing resources while reducing their overall Return on Investment (ROI) and increase their operational flexibility?

The overall flexibility that a private cloud environment offers to the IT department is tighter controls over its compute, and storage resources. One of the features of VMWare virtualization is High Availability (HA), and Distributive Resource Scheduling (DRS). These features ensure that resources will always be available when they are needed by high performance applications.

Under the traditional one-server/ one-application paradigm, when you run out of compute or storage resources you do not have the flexibility of adding more. Further, consider a circumstance where your one-server/ one-application system only needs these additional resources for a short period of time. Under this paradigm, there is a pool of unused compute or storage resources that cannot be used by any other server or application. This paradigm represents a little ROI, because there are compute, and storage resources available, on the other hand, they are not useable because they are dedicated to one-server or one-application.

VMWare's ESXi, however breaks this traditional paradigm by allowing IT departments to cluster and pool compute, and storage resources. Under this paradigm, compute and storage resources can be utilized by multiple systems. When a server or application requires additional resources under this paradigm, those resources are taken from the pool of available compute or storage resources. When these resources are no longer required, the private cloud reclaims them placing the resources back into the pool. This model represents a high ROI, because a pool of compute, and storage resources are available and the system requests and is subsequently granted access to that pool of resources.

Hardware

The Information Technology (IT) department of the NSABP is using Hewlett Packard (HP) enterprise class BL680 blade servers. These enterprise servers are capable of handling up to twenty or more virtual guests [7, 11, 13]. Hewlett Packard was chosen by the NSABP because of their exceptional market place value and industry benchmarks [4, 7, 12, 15, 16]. In their current configuration each enterprise class blade is configured with four, virtualization optimized Intel Xeon processors and 128 GB of RAM each capable of handling large virtual environments. VMWare's ESXi is installed on the hot-swappable RAID 1 (redundantly mirrored) configured internal 72 GB disks. Each host is configured with eight 1000Mb network connections; two of these are dedicated to the management network, while the remaining six is configured for the distributed virtual switch.

Each ESXi host is connected to shared SAN (Storage Area Network) storage through a redundant multi-pathed 8 GB Fiber Channel network. The multi-pathed Fiber Channel solution provides a redundant path configuration to each of the ESXi hosts and storage.

Even though the performance between Fiber Channel and iSCSI are the same; with iSCSI delivering a lower ROI, over its Fiber Channel counterpart in terms of overall hardware, labor costs, and a growing industry critique of Fiber Channel for its expense and complexity [1, 2], it was none-the-less chosen to interface with existing legacy systems.

The adoption of this storage hardware solution has enabled the IT department to streamline its data management, and storage processes. These efficiencies came with VMWare; using virtualization technology to give the appearance of more resource than what is actually physically available i.e. thin provisioning model. Using this approach, the IT department has dramatically reduced the amount of physical storage leading to a reduction in physical infrastructure.

Although iSCSI has a lower ROI when compared to Fiber Channel, the IT department was faced with an issue of interfacing with their legacy systems. The decision was made to continue using Fiber Channel for their primary on-site storage, however move to iSCSI for their co-location in an effort to increase their ROI.

Reduction in Physical Infrastructure

A senior technology manager in the IT department stated that, “[their] deployment of Hewlett Packard blade servers running Intel Xeon processors has enabled [them] to hit industry benchmarks on virtualization technology.” In fact depending on the type of service or application that the IT department has needed to deploy, they have seen a physical to virtual server ratio of 20:1; that means for every virtual host the IT department deploys, they are averaging twenty virtual instances [13], placing them well within the stated industry standards.

By using ESXi the IT department has been able to maximize their existing resources, and consolidate its infrastructure.

METHODOLOGY

Introduction

The purpose of this case study is to provide insight into what operational efficiencies the NSABP has realized because of moving to a private cloud environment. Specifically this case study examined the overall benefits; return on investment (ROI), quality assurance (QA), and quality control (QC) has brought to bear on the organization as a whole.

Research Design

The researchers used a case study as their method due to the intrinsic and instrumental insight that was required. The scope and inquiry into how this organization defined approached and came to resolve their issue is best described using the evaluation research method of case study. According to Merriam it is the best plan to answer research questions that do not apply to all business models [8].

Data Collection

This case study is the result conducting inquiry into how the IT department at the NSABP implemented their private cloud infrastructure. The analysis of this bounded system was selected because of the complexities associated with developing a private cloud and the IT departments associated success in reducing their overall ROI of hardware, and subsequent increase of their QA and QC measures.

Participant

The participant organization for this case study was the IT department at the operations center of the National Surgical Adjuvant Breast and Bowel Project (NSABP) Foundation. The NSABP has more than a fifty-year history in designing and conducting clinical trials that have dramatically changed the treatment and prevention of breast and colorectal cancer. Located in Pittsburgh, Pennsylvania, they are among the nations oldest and most respected clinical trial cooperative groups [10].

Data Collection Procedures

Data for this case study was collected through the extensive use of interview and document analysis. Interviews were conducted both individually, and in a group. These interviews were conducted both in person, and via e-mail from January 12, 2012 through March 10, 2012. The responses to the questions of this case study involved a detailed analysis of how the IT department of the NSABP setup and deployed their private cloud environment.

During one of the interviews conducted on February 2, 2012 with the private cloud project team and its stakeholders, a series of open-ended questions were asked. The questions and answers that are shown below in Appendix: A, reflect the IT department’s attitudes towards patch management, quality control and quality assurance, as well as the IT department’s efficiency and ROI.

ANALYSIS

The data for this case study was collected using interviews. Interviews were conducted in person and via e-mail from January 12, 2012 through March 10, 2012. During this time, the authors conducted half-hour interviews with each of the IT department staff members individually, and one half-hour interview with all IT department staff members as a group. The individual and group interviews were conducted on January 12, 2012. Follow up question and answer e-mails were conducted from January 13, 2012 through March 10, 2012.

Research Question 1 asked, What benefits does moving to a private cloud realize? The results presented in this case study is that; business continuity has been enhanced through the redundancy provided by VMWare's private cloud computing environment. Other enhancements have been the seamless integration with public cloud storage providers, and high availability with complete fail-over into their co-location for near one hundred percent (~100%) availability of their IT operations. With their private compute cloud environment, they are now able to deliver computing as a service rather than a product.

Research Question 2 asked, What is the ROI for moving to a private cloud? The IT department has established an increasing level of ROI for their private cloud infrastructure. The initial capital expenditure has continued to be offset with yearly increases of return approaching twenty-five percent (25%). It is important to note that the ROI would be almost fifteen percent (15%) higher, approaching forty percent (40%) overall had the IT department not needed to continue with its legacy fiber channel, storage area network.

Research Question 3 asked, What QA and QC benefits does moving to a private cloud realize? The IT department has very strict controls on the QA and QC over their computing environment. The ability to ensure a 24x7x365 operational computing environment and its accessibility is crucial towards the benefits of NSABP and the outcomes of their research.

CONCLUSION

The IT department was able to shift from a non-redundant one-server/ one-application paradigm to a highly available and vastly redundant private cloud delivering computing as a service rather than a product. These findings seem to support the literature and the transformational dimensions of IT as this support environment within which the "Process Oriented Model of Business Value" is implemented.

The NSABP was able to achieve this goal by implementing a virtualization solution based on Hewlett Packard hardware, multi-pathed fiber channel storage, integration into public cloud storage, and VMWare's ESXi virtualization technology. Consolidating seventy-five percent of their infrastructure, they have developed a solid foundation for their private cloud moving into the future.

REFERENCES

1. Anonymous. Iscsi San Finally Proves Worthy Alternative to Fibre Channel San for the Mainstream. Retrieved March 4, 2012, 2012, from <http://searchstorage.techtarget.com/iSCSI-SAN-finally-proves-worthy-alternative-to-Fibre-Channel-SAN-for-the-mainstream>
2. Bigelow, S. J. (2007). Iscsi Vs. Fibre Channel Explained: Iscsi Takes Its Rightful Place Beside Fibre Channel (p. 8). Tech Target.
3. Bourne, V. (2011). Virtualization Data Protection Report 2011 (p. 12): Veeam Software.
4. Bubic, A., & Tett, M. (2009). What's the Best Blade Server? Retrieved June 6, 2012, 2012, from http://www.zdnet.com.au/whats-the-best-blade-server_p8-339297483.htm
5. Choinski, V. (2011). Veeam Backup & Replication (p. 13): ESG Lab Review.
6. Hamel, J. (1993). Case Study Methods (Vol. 32). Thousand Oaks, CA: Sage.
7. HP. (2010). Why Scalable Blades: Hp Integrity Server Blades (Bl860c I2, Bl870c I2, and Bl890c I2) (p. 12).
8. Merriam, S. B. (2009). Qualitative Research. San Francisco: Jossey-Bass.
9. Mooney, J., Gurbaxani, V., & Kraemer, K. (1995). A Process Oriented Framework for Assessing the Business Value of Information Technology. Center for Research on Information Technology and Organizations.
10. NSABP. (2011). Over Fifty Years of Clinical Trial History. Retrieved February 11, 2012, 2012, from <http://www.nsabp.pitt.edu>
11. Packard, H. (2010). Best Practices for Deploying Vmware Vsphere 4.1 on Hp Proliant D1980 G7 Servers Technical White Paper (p. 13). Hewlett Packard.
12. Robb, D. (2011). Hp Proliant Servers Buyer's Guide. Retrieved June 6, 2012, 2012, from <http://www.serverwatch.com/trends/article.php/3937696/HP-ProLiant-Servers-Buyers-Guide.htm>
13. Solutions, I. (2003). Twenty-to-One Consolidation on Intel Architecture (p. 14).
14. Surksum, K. v. (2010). Veeam Releases Cios Survey Report About Data Protection in Vmware Environments. Retrieved March 10, 2012, 2012, from <http://virtualization.info/en/news/2010/11/veeam-releases-cios-survey-report-about-data-protection-in-vmware-environments.html>
15. Venezia, P. (2007). Blade Server Shoot-Out: Dell Vs. Hp Vs. Sun. Retrieved June 6, 2012, 2012, from <http://www.infoworld.com/d/hardware/blade-server-shootout-dell-vs-hp-vs-sun-658>
16. Venezia, P. (2010). Blade Server Review: Dell, Hp, Ibm Battle for the Virtual Data Center. Retrieved March 4, 2012, from <http://www.infoworld.com/d/computer-hardware/blade-server-review-dell-hp-ibm-battle-the-virtual-data-center-252>
17. VMWare. (2011). Esxi – Vmware's Next Generation Hypervisor Architecture. Retrieved March 10, 2012, from <http://communities.vmware.com/community/vmtn/server/vsphere/esxi>

APPENDIX: A

When and how do you conduct your patch management tasks? The NSABP conducts patch management procedures on the second Saturday of each month, or on an as needed basis.

Do you follow the vendor cycle and how do you test to ensure that it will work? Our environment is not large enough to constitute an expenditure that fully or partially mirrors our or production environment. Therefore, we wait approximately one month after an update, or patch is released prior to placing it into production. This method ensures that any issues encountered by early adopters at other organizations are vetted and remediate by the patch vendor.

Do you have a test environment ready to ensure it will work and not impact the production system and what is your process? Our environment is not large enough to constitute an expenditure that fully or partially mirrors our or production environment. This process is different for our software development team, however. Because of the robustness of ESXi we are able to mirror a small portion of our production environment so we can test and fully vet any development updates to our software, according to our QA and QC processes.

What is your SLA (service level agreement) if you have any? We have a 24x7x365x4 SLA for all of our hardware and cloud software. In the event of any systemic hardware failure, we are guaranteed an onsite field engineer within 4 hours of the reported outage of the hardware. For the software, we are guaranteed a call back within 4 hours.

How does your vendor support you and their issues and what is their response rate, timing? We have never experienced issues with any of our HP or VMWare support; I would rate them at 100%!

How can you show an efficiency and ROI within the process you use to install patches? We probably cannot show efficiency because we do this as an in-house process. As such it is a process one of our employees does as part of their typical duties. Industry leaders like Hewlett Packard and VMWare recommend using a proof-of-concept environment to mirror as closely as possible the production environment [11]. This best practice strategy allows for both a quality assurance (QA) and quality control (QC) environment to be at the center of any cloud upgrades.