Lab Validation Report

EMC Symmetrix VMAX and Microsoft Server Virtualization

Scalable Enterprise-Class Virtual Infrastructure

By Vinny Choinski with Brian Garrett

November 2010

Contents

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by EMC.

Introduction

While a growing number of organizations have taken advantage of the tremendous efficiency and savings that can be realized with server virtualization technology, broad-based adoption throughout the data center often stalls due to a number of challenges including scalability, reliability, and ease of deployment. This report documents how EMC VMAX storage, combined with Microsoft Cluster Shared Volumes (CSV) and Hyper-V R2 server virtualization technology, can be used to addresses those challenges as it supports up to 1,024 virtual machines running on a consolidated, enterprise-class infrastructure.

Background

Server virtualization has made giant strides over the past decade, creating heroes inside IT organizations. Accordingly, it's no surprise that interest in server virtualization remains as strong as ever. Indeed, respondents to a recent ESG survey ranked "increased use of server virtualization" as their number one IT priority over the next 12-18 months.[1] However, despite the broad success of server virtualization, nagging issues and challenges exist. As a result, a low percentage of the potential workloads that can be virtualized have been migrated to virtual machines and the consolidation ratios of virtual machines per physical server remains relatively low. Challenges that are blocking the next wave of virtualization and consolidation include concerns about scalability, performance, reliability, and ease of deployment.

A recent ESG survey of North American enterprise and larger midmarket IT professionals explored the storage challenges associated with server virtualization.[2] Given the rapid growth in the number of virtual machines being deployed, it's no surprise that scalability, performance, and the overall volume of storage capacity have been identified as key challenges.

Figure 1. Server Virtualization Storage Challenges

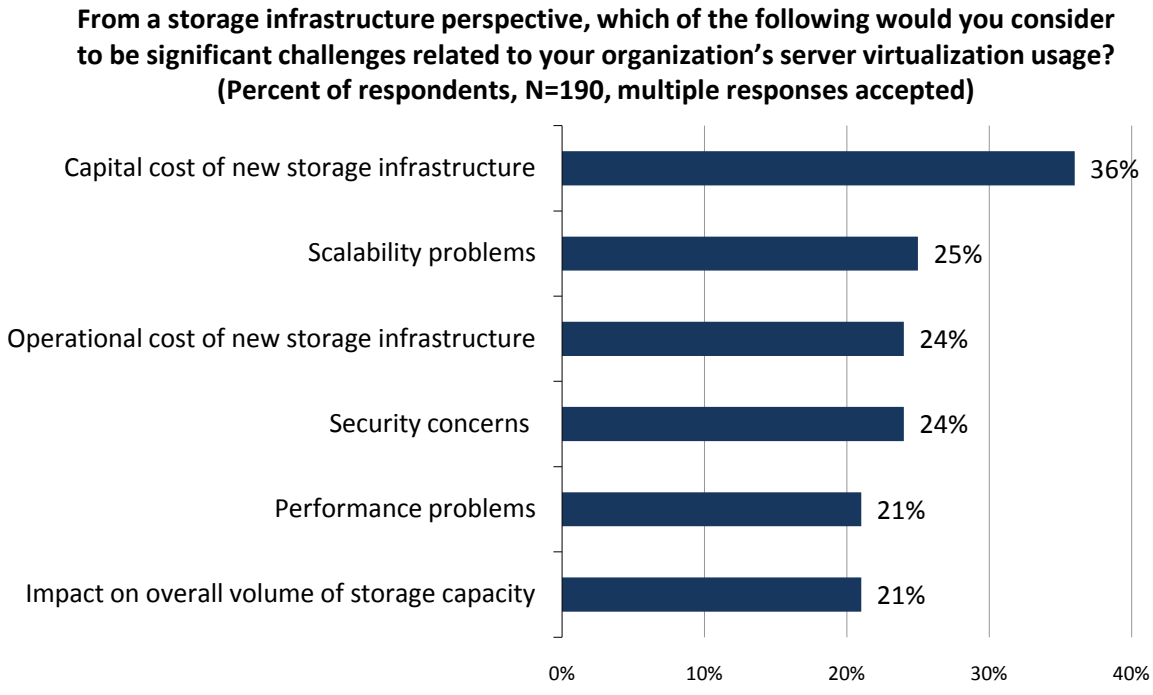

Source: Enterprise Strategy Group, 2009.

[1] Source: ESG Research Report, *2010 IT Spending Intentions Survey*, January 2010.
[2] Source: ESG Research, *2010 Server Virtualization Survey*, September 2010.

EMC Symmetrix VMAX and Microsoft Server Virtualization

This report documents testing of an EMC VMAX storage and Microsoft server virtualization solution that's designed to overcome the scalability, performance, and ease of deployment challenges associated with large-scale virtual server deployment. As shown in Figure 2, the solution is comprised of a cluster of industry-standard servers running Hyper-V R2 server virtualization software and a SAN-connected EMC Symmetrix VMAX Series disk array. Microsoft-enabled servers and EMC-enabled storage resources can be added to the solution to increase the number of virtual machines running on a consolidated virtual infrastructure.

Figure 2. Enterprise-class Virtual Machine Scalability with EMC VMAX and Microsoft Server Virtualization

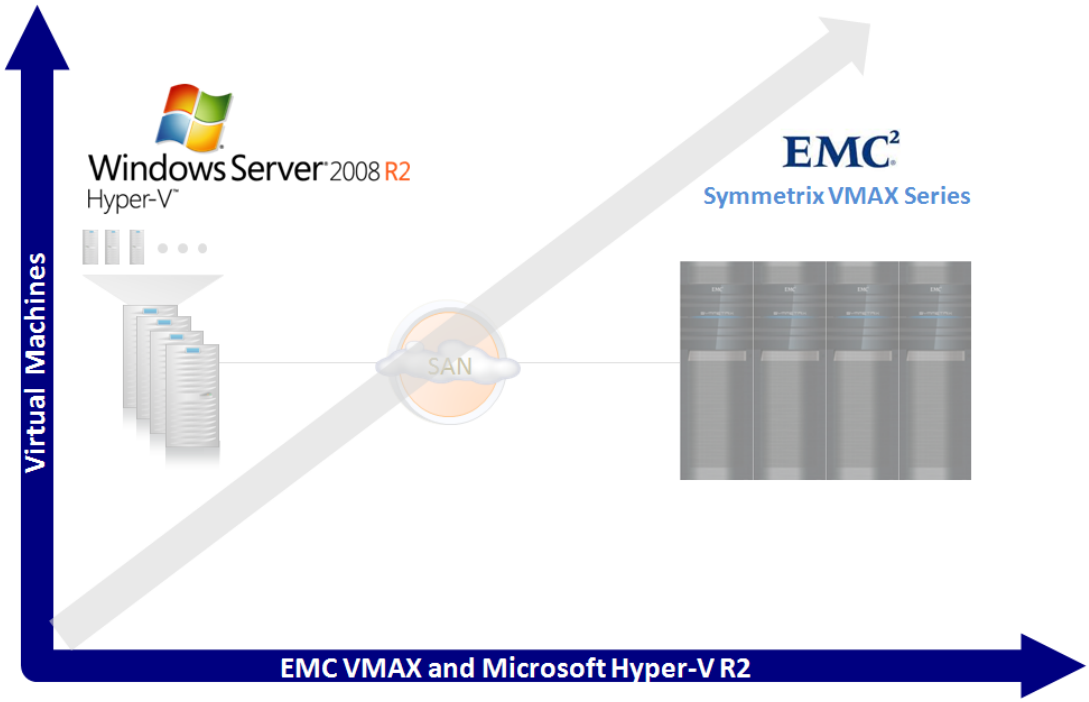

The EMC VMAX series is the latest generation of the Symmetrix enterprise-class storage product line that's ideally suited for the large virtual data center. Enginuity is the software at the core of the VMAX solution that controls and enables advanced features like seamless scale-out, easy provisioning, rapid cloning, and right-sized storage tiering. The EMC VMAX can scale to meet a wide variety of storage needs ranging from a small virtual server deployment to enterprise-class implementations of thousands of virtual machines requiring petabytes of capacity.

Hyper-V R2 is a bare metal hypervisor that enables the hosting of multiple virtual machines on the same physical server. The supported virtual machines can be a mix of almost all Microsoft (server and desktop) platforms in addition to a couple different Linux platforms. Hyper-V R2 is available in the Windows 2008 Server R2 operating system as well as Microsoft Hyper-V Server 2008 R2. With affordable licensing models, a familiar Microsoft look and feel, and a hypervisor-agnostic suite of management tools that works with physical and virtual servers (VMware included), Microsoft lets companies take advantage of existing skill sets, training programs, and certifications.

A number of enterprise-class features found in EMC VMAX storage and Microsoft virtual server technologies make this solution possible. For example, this report illustrates how EMC TimeFinder clones and Auto-provisioning Groups, combined with Microsoft System Center Virtual Machine Manager (SCVMM) templates, can be used to automate the rapid deployment of virtual machines. The report also documents how Microsoft CSV connected to a consolidated pool of EMC Symmetrix VMAX storage can be used to create a scalable platform for large application workloads.

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of VMAX storage, Microsoft CSV, and Hyper-V R2 at EMC's corporate headquarters located in Hopkinton, Massachusetts. Testing was designed to demonstrate the scalability, ease of deployment, performance, and reliability of a consolidated infrastructure with 1,024 virtual machines.

Getting Started

The ESG Lab test bed is shown in Figure 3. Hyper-V R2 was used to deploy more than a thousand virtual machines on a 16-node Microsoft Failover Cluster SAN-connected to a four-frame EMC VMAX storage system.

Figure 3. The ESG Lab Test Bed

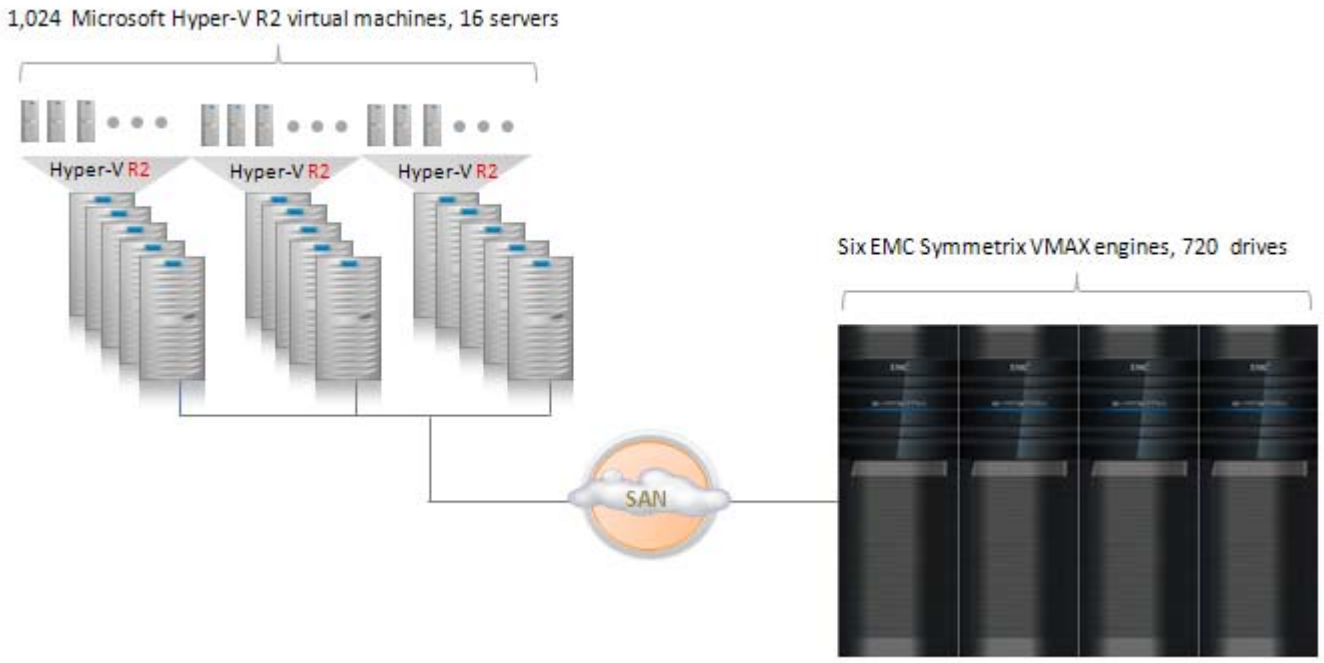

Key components of the configuration include:[3]

- Sixteen multi-core servers, each with eight processor cores and 96 GB of memory
- Windows Server 2008 Hyper-V R2
- Microsoft Cluster Shared Volumes
- A Brocade 8 Gbps SAN fabric
- A six-engine EMC Symmetrix VMAX array
- 702 450 GB 15K RPM disk drives
- EMC TimeFinder, Auto-provisioning Groups, and PowerPath software

The configuration of the Microsoft CSVs and the EMC Symmetrix VMAX solution was designed to meet the scalability requirements of a performance-intensive application workload supporting up to 1,024 virtual machines. Those requirements were met with multiple VMAX engines (six) and a large number of high performance disk drives (702).

[3] For more details, see the Appendix.

Microsoft System Center Virtual Machine Manager (SCVMM) was used to take a tour of the test bed from a virtual server infrastructure point of view. SCVMM supports a number of advanced management features including intelligent placement, physical to virtual (P2V), virtual to virtual (V2V), and centralized management and monitoring. SVCMM also includes automated provisioning tools that use templates to simplify the deployment of virtual machines.

A server-centric SCVMM view of the ESG Lab test bed is shown in Figure 4. The 16-node cluster is shown toward the left. Details about one of the physical servers in the cluster is shown toward the right (e.g., 16 Intel Xeon 2.27 GHz processor cores).

Figure 4. A Microsoft Virtual Machine Manager View of the Test Bed

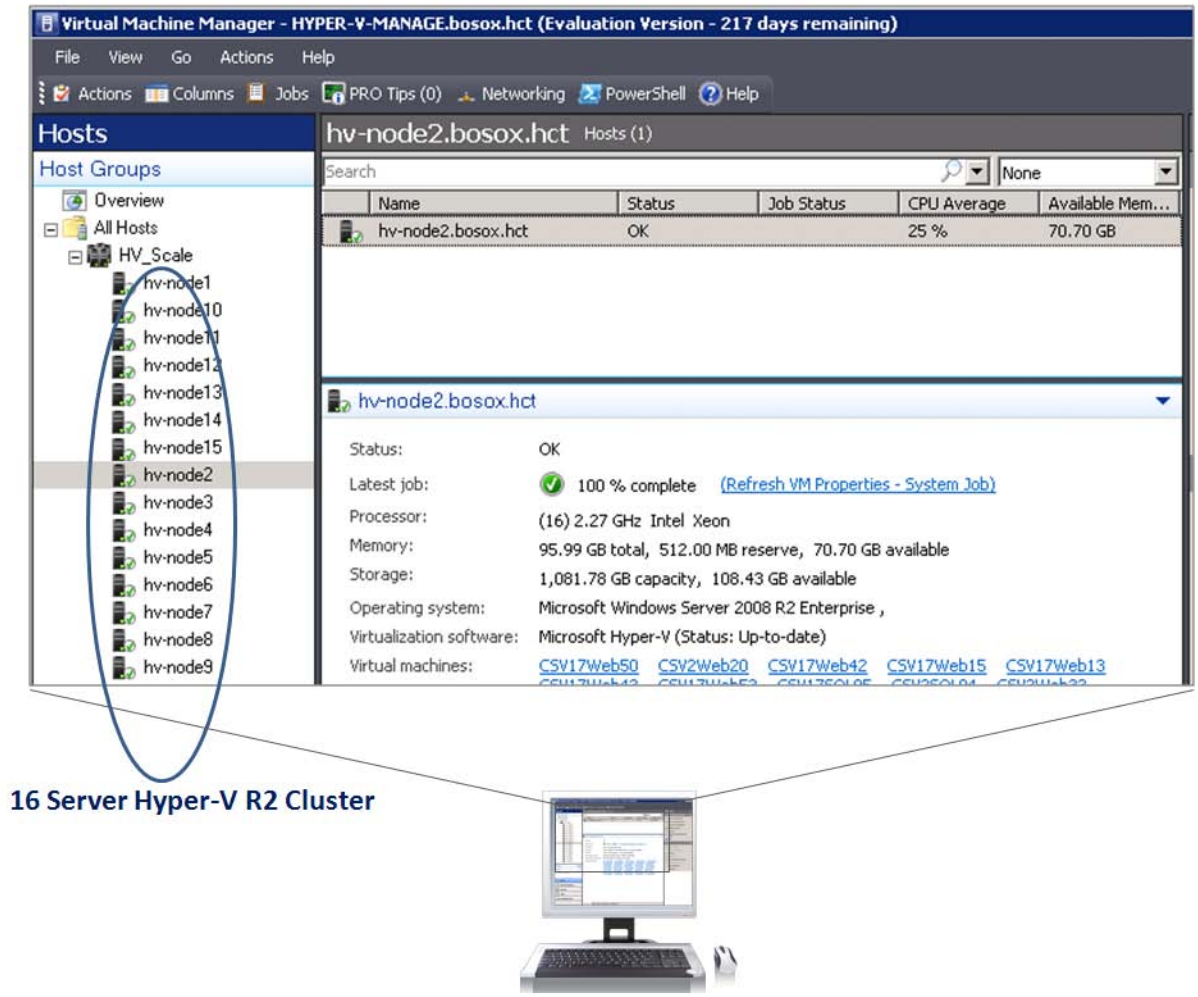

Next, ESG Lab used the Symmetrix Management Console (SMC) to examine the configuration from a storage perspective. The Symmetrix Management Console is a Web-based management interface designed to simplify storage resource management. As shown in Figure 5, 720 disk drives represented the backing storage for thin devices were used within the environment. These data devices, configured as 2-way mirrored volumes, were used to store the data associated with 1,024 virtual machines.

Figure 5. An EMC Symmetrix Management Console View of the Test Bed

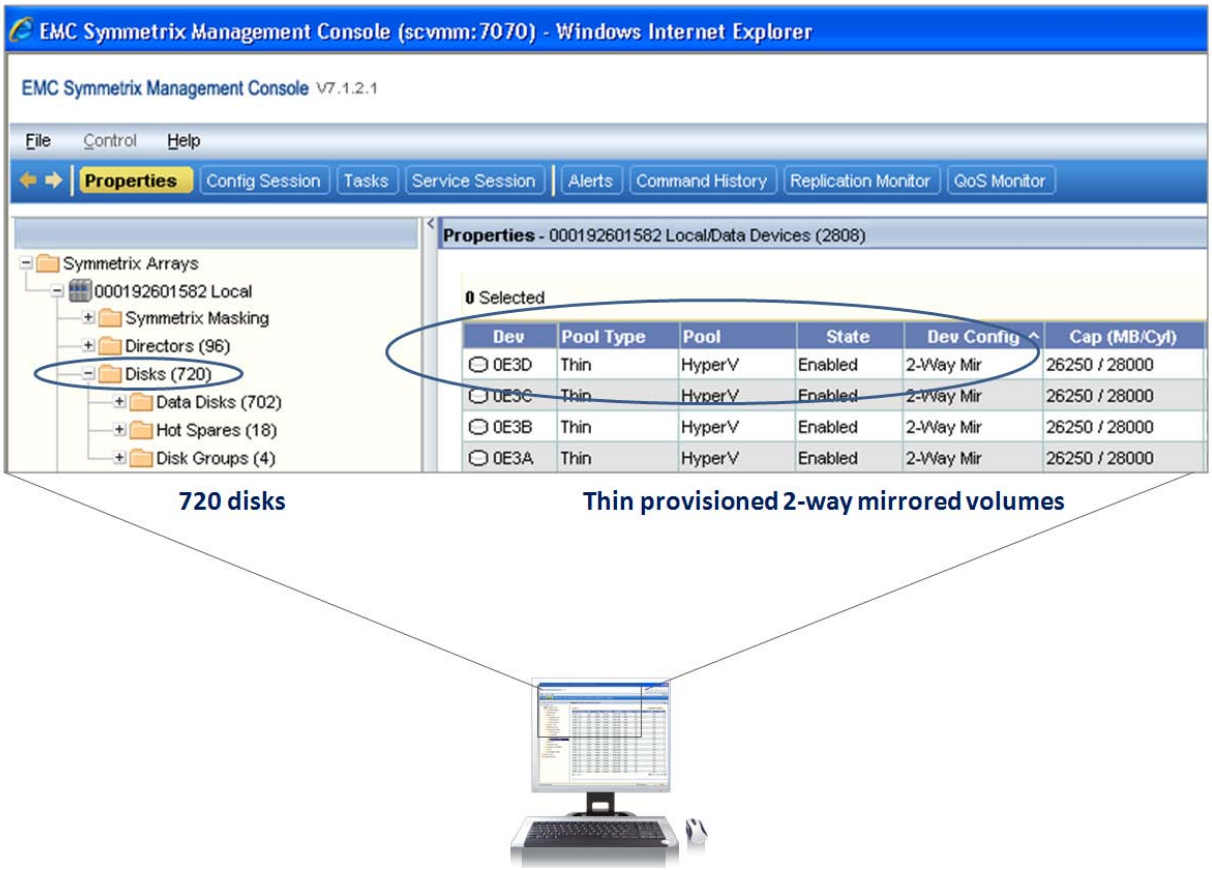

As summarized in Figure 6, EMC VMAX and Hyper-V R2 were used to create a single virtualized infrastructure that can scale to support from 64 through 1024 server-class virtual machines.

Figure 6. Scaling to 1,024 Server-class Virtual Machines with a Single Hyper-V R2 Cluster with EMC VMAX

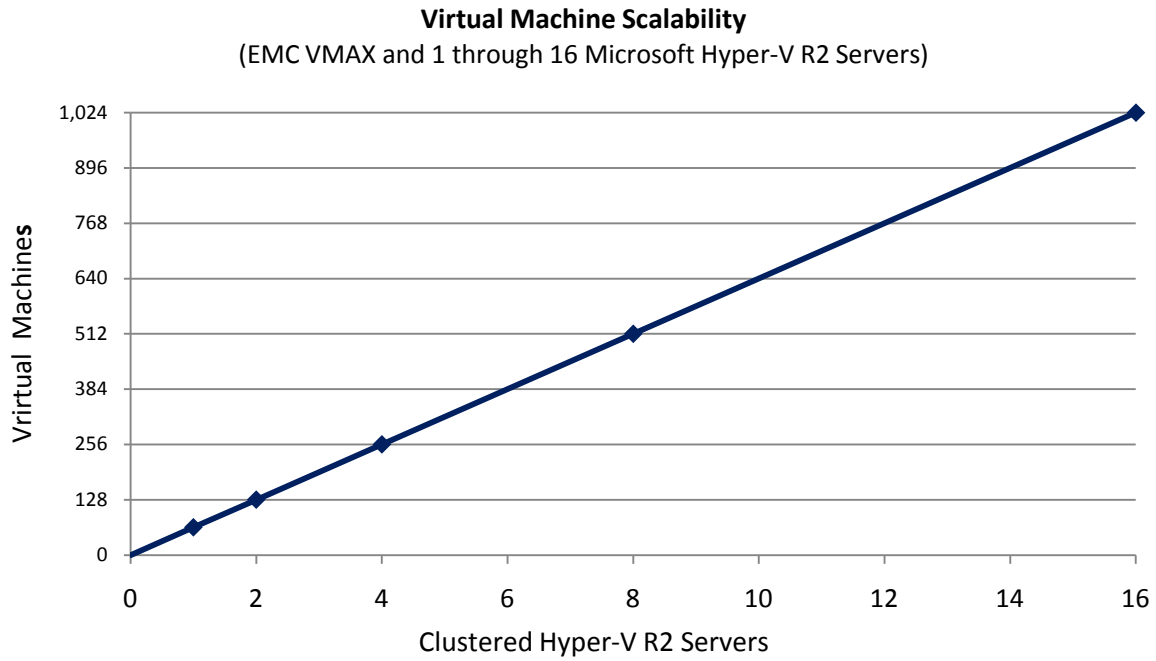

When the implementation was configured with a maximum implementation, 1024 virtual machines were distributed across 16 high performance servers. Each server was configured with two quad-core high speed Intel processors and 96 GB of RAM. The servers were connected to EMC VMAX storage through a dual fabric 8Gb Brocade SAN. Multiple paths to the VMAX storage were configured from each server for added scalability and availability. The storage was carved into three thin pools comprised of a total of 702 drives in RAID-1 mirror sets. Thin devices were used to create 24 1 TB 16-way striped data LUNs. The 24 LUNs were thin provisioned and presented to the servers as CSVs. On the volumes, 1,024 1 TB Hyper-V R2 virtual hard drives were used to host virtual machine operating system images.

The configuration used to deploy 1,024 virtual machines was analyzed by ESG Lab with a goal of estimating the server and storage horsepower needed to support smaller virtual server deployments. The results, which are summarized in Table 1, can be used to plan the resources required to get started with a server virtualization solution that can scale to meet the needs of the business. For example, the second column shows how a pair of servers attached to two VMAX engines can be used to support up to 128 server-class virtual machines. Columns to the right illustrate how servers, VMAX engines, and drives can be added to support up to 1,024 virtual machines with a consolidated Enterprise-class infrastructure.

Table 1. Scaling to 1,024 Virtual Machines in One Microsoft Hyper-V R2 Cluster with EMC VMAX

VMs	64	128	256	512	1024
Servers	2	2	4	8	16
CPU Cores	16	16	32	64	128
RAM (GB)	192	192	384	768	1,536
VMAX Frames	1	2	3	3	4
VMAX Drives	120	194	388	582	702
Engines	2	2	2	3	6
FE Ports	8	16	16	24	48
FE Bandwidth (Gbps)	32	6	64	96	192
Usable Capacity (TB)	20	40	80	160	320

Why This Matters

Managing server sprawl is a challenge faced by many mid- and large-scale IT organizations. Under-utilized servers take up space, waste resources, and are more likely to be improperly aligned with the workloads they support. Server virtualization can be used to combat server sprawl as it consolidates the number of servers in the data center.

Despite the broad success of server virtualization, a new set of scalability challenges has emerged. As a result, the consolidation ratios of virtual machines per physical server (and the savings that can be achieved) have remained relatively low. If the scalability challenges can be addressed, consolidation ratios—and savings—are bound to rise as well.

ESG Lab has confirmed that a cluster of Hyper-V R2 multi-core servers connected to an EMC VMAX with multiple storage engines can be used to meet the scalability challenges associated with wide-scale virtual server consolidation.

Rapid Deployment with Scale

EMC TimeFinder software running within an EMC VMAX disk array and Microsoft CSVs, automated with Microsoft PowerShell scripts, can be used to accelerate and simplify the large-scale deployment of virtual machines. A virtual machine template is a core concept that makes this possible; these templates contain a standardized group of hardware and software settings that can be used to repeatedly create new virtual machines with similar attributes. Templates are commonly used to automate the deployment of virtual machines. Microsoft SCVMM and PowerShell scripts can be used to automate the use of virtual machine templates.

Traditionally, templates are deployed using the TCP/IP protocol running over a local area network. Although automated, deploying a large number of VMs with this approach could take days due to the speed and bandwidth of the local area network. For a more rapid deployment approach, EMC TimeFinder can be used to create instant clones or snapshots of CSV-based templates which are then presented to physical servers. The physical servers use the snapshot-based templates to finish the process and bring the VMs online.

ESG Lab Testing

ESG Lab performed hands-on testing of TimeFinder-enabled cloning and rapid deployment of 64 new virtual machines using the 1,024 virtual machine test bed presented previously in this report. ESG Lab also observed the creation of ten virtual machines using both the traditional and TimeFinder methods. As shown in Figure 7, templates were copied over the LAN using traditional means compared to the fast and efficient TimeFinder method.

Figure 7. Rapid VM Deployment with EMC TimeFinder and Microsoft Cluster Shared Volumes

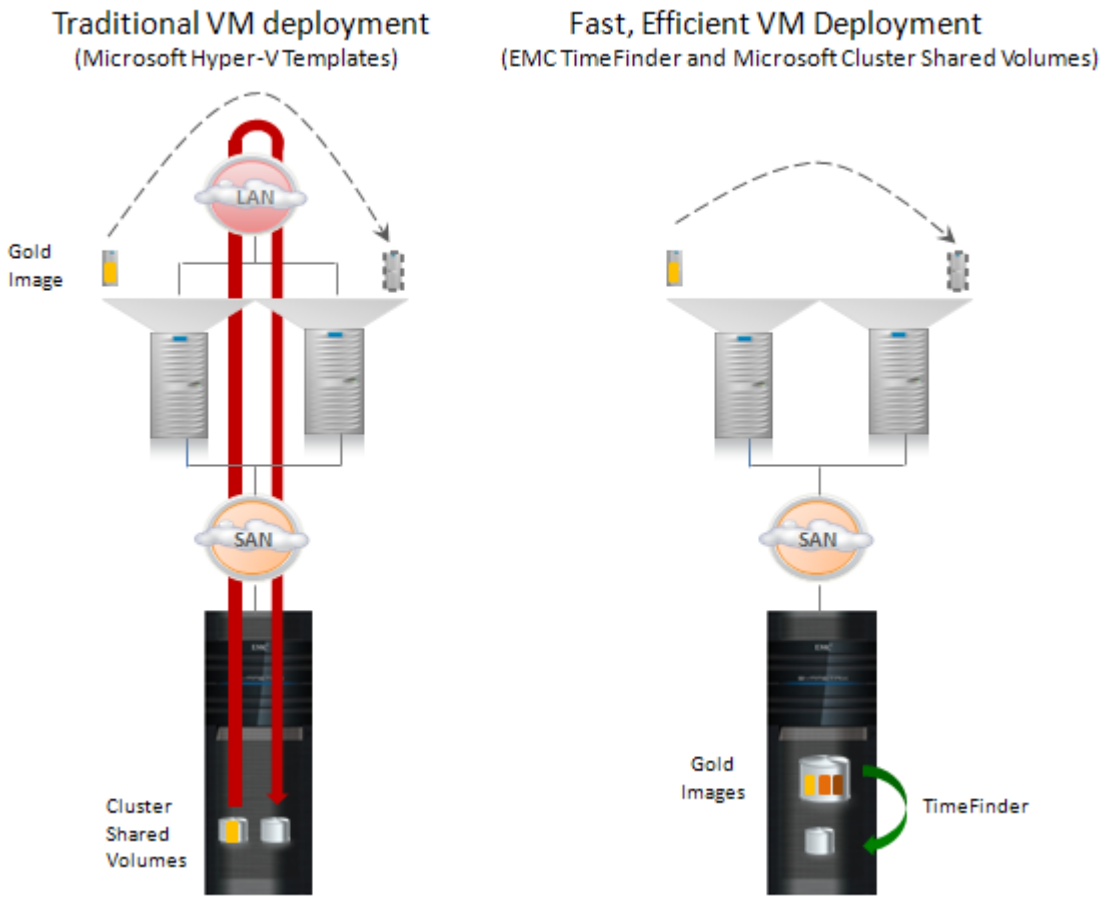

The right side of Figure 7 shows the TimeFinder approach, where multiple replicas of a Gold Master LUN were created. The replicas were then implemented as TimeFinder images. These TimeFinder images were then presented to the cluster server nodes and the LUNs were configured as CSV volumes within the cluster. At that point, the virtual machine instances were imported. As shown in Figure 8, the time to deploy virtual machines with the SAN-based approach is a fraction of what's required by the traditional network-based method.

Figure 8. Rapid VM Deployment with EMC TimeFinder and Microsoft Cluster Shared Volumes

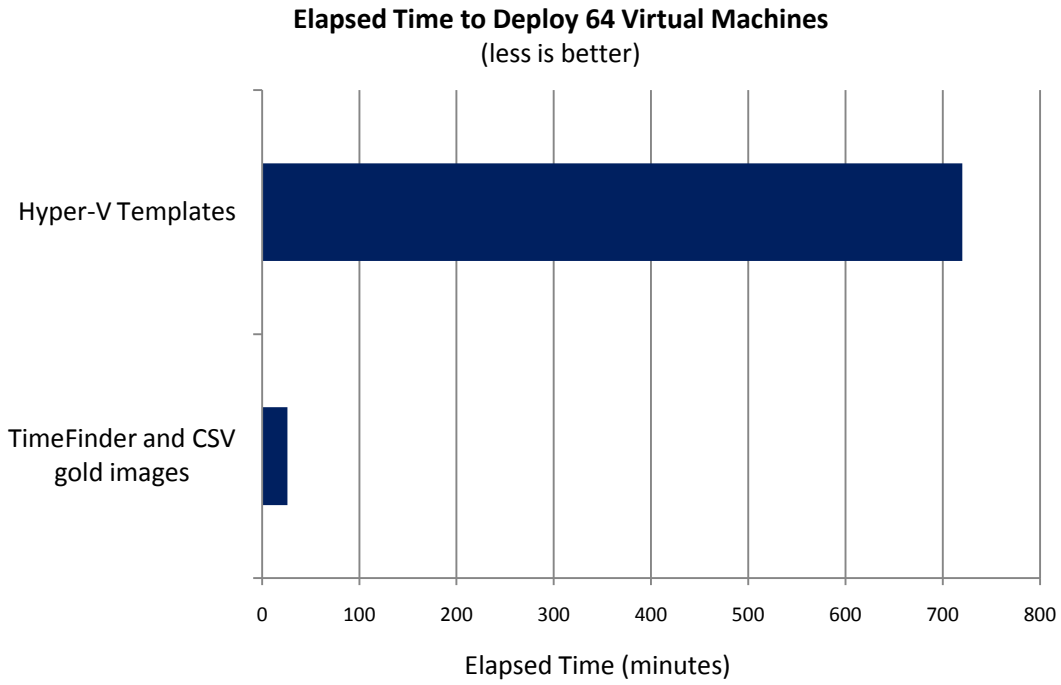

What the Numbers Mean

- Deploying 64 virtual machines using the fully automated TimeFinder method took 26 minutes during ESG Lab testing. Most of that time was spent waiting for virtual machine boot and sysprep activities.
- An estimated elapsed time of 12 hours for the traditional Hyper-V template method was projected based on the time it took to deploy ten virtual machines.
- The TimeFinder method was more than 27 times faster and generated dramatically less LAN traffic.

Why This Matters

Rapid and efficient methods are needed when deploying large numbers of virtual machines in a consolidated virtual environment. Rapid deployment of virtual machines helps the IT organization respond faster to the needs of the business. Using TimeFinder to make replicas of gold images reduces network traffic compared to traditional methods which copy them over a LAN. Less network traffic reduces the impact that large-scale VM deployment can have on production application performance.

ESG Lab confirmed that EMC Symmetrix VMAX storage with TimeFinder and Auto-provisioning Groups, combined with Hyper-V R2 and CSV, greatly simplifies and accelerates the deployment of virtual servers and storage. The ability to easily deploy large numbers VMs in minutes instead of hours (or days) increases efficiency and reduces costs.

Performance Scalability

The Hyper-V R2 and Symmetrix VMAX solution presented in this report was designed to deliver predictably fast performance for a consolidated pool of virtualized applications. The enterprise-class performance of the solution starts with a powerful cluster of industry standard multi-core servers; the processing power, memory, and internal bandwidth of the solution can be scaled to meet the needs of the business by simply adding more servers to the cluster. Next, Hyper-V R2 leverages the bare metal hypervisor acceleration technology built into the latest multi-core servers to minimize any potential overhead in the server virtualization layer. Last, but not least, the massively scalable architecture of the Symmetrix VMAX delivers predictably high levels of storage performance for a growing number of virtual machines.

ESG Lab audited a series of performance tests performed on the 1,024 virtual machine cluster. The result of an IO-intensive performance test that was run on one of the servers in the cluster is shown in Figure 9. This particular test was designed to measure the performance of a growing number of virtual machines sharing a single CSV residing on an EMC Symmetrix VMAX. The average IO response time was recorded as the number of virtual machines scaled from eight to 64. To put the response times in perspective, a Microsoft guideline for the maximum response time for Exchange database reads (20 milliseconds) is shown as a green dotted line.

Figure 9. EMC VMAX Performance Scalability with Hyper-V R2

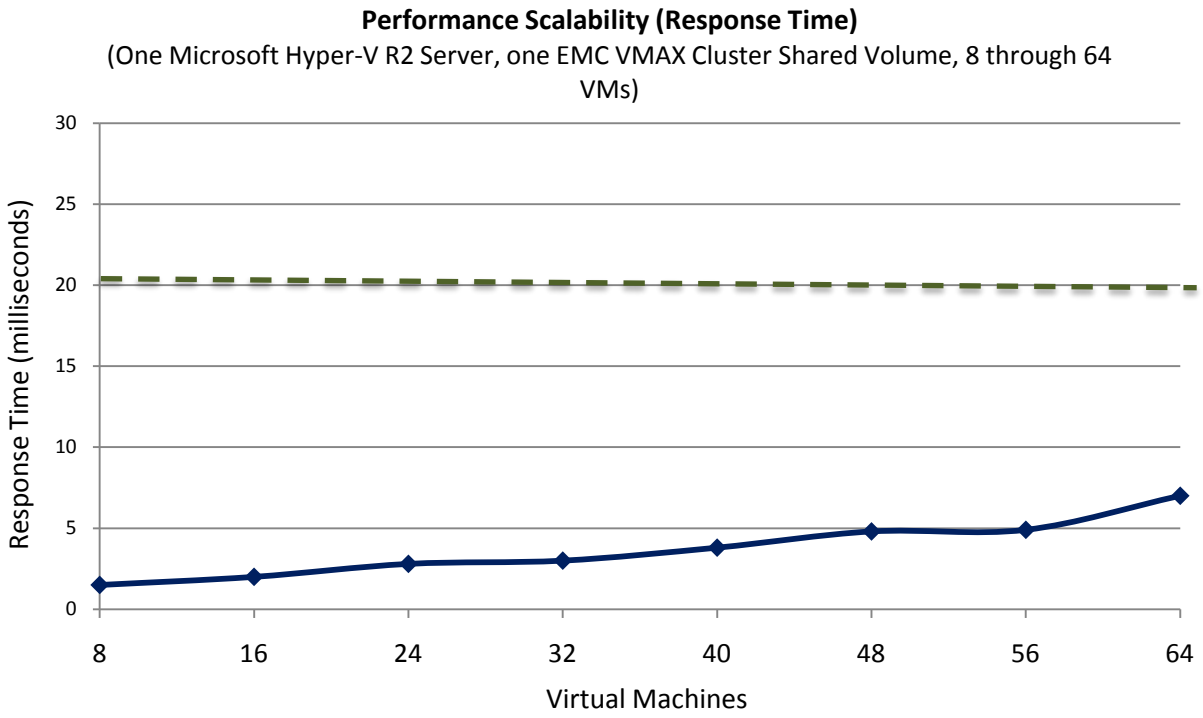

Table 2. EMC VMAX Performance Scalability with Hyper-V R2

Virtual Machines	8	16	24	32	40	48	56	64
Response Time(ms)	1.5	2.0	2.8	3.0	3.8	4.8	4.9	7.0
IOs per second	25,000	36,000	40,000	43,000	43,500	43,000	43,000	45,000

What the Numbers Mean

- Long IO response times make applications feel slow to end-users.
- Predictably low response times are needed to support a growing number of virtual machines running on a consolidated infrastructure.
- Eight virtual machines running an IO-intensive workload maintained an excellent IO response time of 1.5 milliseconds.
- Response times remained predictably low as the number of virtual machines increased.
- A maximum response time of 7 milliseconds was recorded for 64 virtual machines. Response times were under the generally accepted limits for common business applications.
- For example, the Microsoft guideline for the maximum response time for Exchange database reads is shown as a green dotted line (20 milliseconds).
- Results were recorded with a performance-intensive workload for 64 virtual machines (45,000 IOs per second) that is beyond what most IT organizations would require for a single LUN. This methodology was chosen to provide solutions-based guidelines that customers can use when configuring a scalable infrastructure based on Microsoft CSV and EMC VMAX technologies.

Why This Matters

Slow storage response times can result in loss of productivity, loss of competitiveness, and, in the worst case, loss of revenue. With more and more organizations deploying entire suites of production applications on virtual servers, predictably fast storage performance that can scale to meet the needs of the business is a critical concern. Adding storage hardware to an over-engineered solution to avoid potential performance problems increases costs as it reduces the potential savings that can be realized with consolidation.

ESG Lab has confirmed that the enterprise-class performance capabilities of two or more VMAX engines, combined with the low overhead of Hyper-V R2 and CSV, can be used to deliver the predictably fast response times needed in a consolidated environment supporting up to 1,024 virtual machines.

Flexibility and High Availability

The EMC and Microsoft solution tested by ESG Lab offers a number of Enterprise-class flexibility and high availability characteristics. The solution was architected to be highly available using a number of technologies including Microsoft Live Migration and Failover Clustering. EMC PowerPath multi-path software and redundant SAN connections were used to create a highly available connection to shared VMAX storage residing on Microsoft CSV. High availability was provided at the storage system level in the VMAX disk array with a fully redundant architecture and a rich assortment of RAID protection options. As shown in Figure 10, ESG Lab tested these flexibility and high availability capabilities with a goal of ensuring that the virtual server infrastructure remained up and available during a Live Migration, a server failure, and a hard drive failure.

Figure 10. Flexibility and Availability

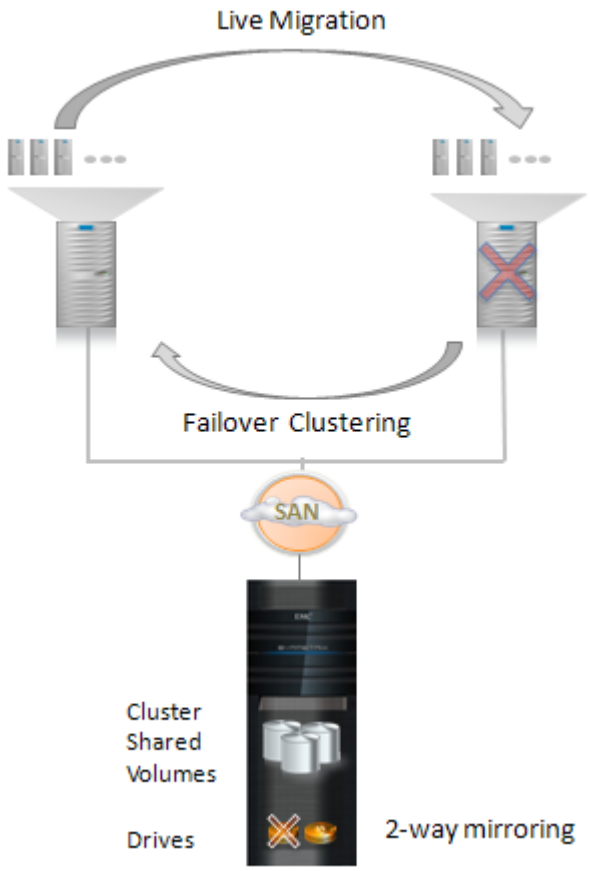

ESG Lab high availability testing was performed as an IO-intensive Iometer workload and a long running directory-level copy operation ran without error on 64 virtual machines. SCVMM was used for a Live Migration of one of the running virtual machines from one physical server to another. Removing a power cable on an active physical server triggered failover clustering of all of the running VMs to another physical server in the cluster. One of the active hard drives was removed from the VMAX to trigger an automated recovery of a 2-way mirrored volume.

Why This Matters

As more and more infrastructure is consolidated on a common platform, the impacts of failures become greater. ESG Lab has confirmed that EMC Symmetrix VMAX storage, combined with powerful Microsoft technologies (Hyper-V R2, Failover Clustering, CSV, and Live Migration), can be used to create a flexible and highly available virtual infrastructure.

ESG Lab Validation Highlights

- ☑ ESG Lab confirmed that 16 servers running Hyper-V R2 connected to a six-frame EMC Symmetrix VMAX can be used to support up to 1,024 virtual machines.
- ☑ EMC VMAX TimeFinder replicas and CSVs were used to automate the deployment of 64 Hyper-V R2 virtual machines.
- ☑ A TimeFinder-enabled deployment of 64 virtual machines completed in 26 minutes.
- ☑ The TimeFinder cloning method, which was scripted for fully automatic VM deployment using Microsoft PowerShell, was more than 27 times faster than a traditional LAN-based VM template deployment.
- ☑ An IO-intensive workload designed to measure the maximum performance capabilities of single CSV residing within an EMC VMAX achieved excellent response times (1 to 7 milliseconds) as the number of virtual machines was scaled from eight to 64.
- ☑ Live Migration was used to move a running virtual machine from one physical server to another.
- ☑ A server power failure triggered a Microsoft Cluster Failover for active virtual machines.
- ☑ Pulling an active drive triggered an automated recovery of CSVs residing on a 2-way mirrored volume in the Symmetrix VMAX.

Issues to Consider

- ☑ While not tested during this ESG Lab Validation, one of the more compelling capabilities of Microsoft's integrated suite of virtualization management tools is the ability to manage physical and virtual servers—including virtual servers running over a VMware hypervisor—from a single pane of glass.
- ☑ While Microsoft PowerShell scripts were used to rapidly deploy virtual machines using EMC Symmetrix VMAX TimeFinder technology during ESG Lab testing, this deployment method is not supported from the Microsoft SCVMM graphical user interface. EMC professional services are recommended for IT managers that lack experience with PowerShell scripting and EMC command line interfaces.
- ☑ The 1,024 virtual machine configuration documented in this report was tested in a controlled environment. Due to the many variables in each production data center environment, the number of virtual machines that can be consolidated with the test bed presented in this report may vary according to the size and resource requirements of the virtualized applications in your organization.

The Bigger Truth

Server virtualization adoption has taken off in recent years due to the fantastic savings that can be achieved as under-utilized servers are consolidated to run on a cost-efficient virtual infrastructure. Consolidating applications running on under-utilized servers onto a fewer number of servers reduced cost and complexity. Connecting growing pools of virtualized servers to a centralized SAN-attached storage infrastructure magnified the savings as it increased the flexibility, and availability, of the underlying IT infrastructure.

As the number of organizations that have embraced server virtualization continues to rise, the rate of adoption within each organization often stalls. Having realized the benefits of server virtualization for utility, test, and development applications, IT professionals are looking to extend the benefits to mission critical production applications—at scale and with high availability, predictably fast performance, and automated ease of deployment.

ESG Lab has confirmed that the combination of EMC Symmetrix VMAX and Hyper-V R2 is ready to meet the challenges of the next wave of virtual server adoption. Hyper-V R2, which is built into Windows Server 2008 R2, provides the underlying virtualization technology that enables server consolidation. EMC's flagship enterprise-class Symmetrix VMAX product line provides the scalable, reliable, and fast storage that's needed in a highly consolidated environment. With EMC TimeFinder and PowerShell scripts for rapid and efficient server provisioning and powerful Microsoft capabilities—including Live Migration, CSV, and Failover Clustering—ESG Lab has confirmed that EMC and Microsoft have created an enterprise-class platform for wide-scale virtual server deployments.

This ESG Lab project highlights the practical and compelling value of EMC ELab. EMC invests a huge amount of time, money, and effort to make sure its products work well with existing networks, servers, non-EMC storage, applications, operating systems, and server virtualization technologies—including the Hyper-V R2 solution examined in this report. EMC's ELab investments can be used to substantially lower the total cost of ownership by addressing not only the visible costs—meaning hardware, software, and services— but also the embedded, or hidden, costs of testing, interoperability, and integration.

Beyond the scalability and ease of deployment documented in this report, a couple of facts about Hyper-V should be considered—specifically related to maturity, manageability, and cost—before betting on Microsoft server virtualization technology. While it can be argued that earlier versions of Hyper-V lacked features and capabilities compared to VMware, ESG Lab feels that the Hyper-V R2 release in late 2009 has the features required for wide-scale virtual server deployment. From a manageability standpoint, it should be noted that Microsoft's growing suite of hypervisor-agnostic management tools can manage at the application level regardless of whether applications are deployed on a physical or virtual server. With Microsoft System Center Virtual Machine Manager, you can manage virtual machines running on a mix of Microsoft, VMware, or Xen hypervisors.

Last but not least, it should be noted that Microsoft Hyper-V Server is a free download, limited in terms of support for total memory, number of processors, and quick migration. Windows Server 2008 R2 Standard, Enterprise, and Datacenter editions each include a Hyper-V R2 license for a single physical server as well as licensing for one, four, and unlimited virtual machines, respectively. Compounded by the savings that can be achieved when you leverage existing investments in Microsoft training, certification, and interoperability and the field-proven reliability and efficiency of the EMC Symmetrix product line, ESG Lab feels that EMC Symmetrix VMAX and Hyper-V R2 can be used to create a cost effective and compelling platform for the next wave of scale-out virtual server deployment.

Appendix

Table 3. Test Environment Details

Servers	
16 Intel-based 2U rack mount servers	2 Quad Core L5520 CPUs @ 2.27 GHz 96 GB RAM Dual Port 8Gb FC HBA
Storage and SAN	
EMC VMAX Storage System	4 Frames 6 Engines 702 450 GB FC 15k Drives
Brocade 8Gb SAN	Dual Fabric 8Gb SAN
Software	
EMC	Enginuity microcode version 5874 EMC PowerPath Auto-provisioning Groups EMC TimeFinder Symmetrix Management Console
Microsoft	Windows 2008 R2 Hyper-V R2 Failover Clustering Cluster Shared Volumes System Center Virtual Machine Manager

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel:508.482.0188 Fax: 508.482.0218 | www.enterprisestrategygroup.com