

Lab Report

HP and Microsoft Serverquarium

Private Cloud Solution for Virtualized Workloads

By Kerry Dolan, Lab Analyst and Brian Garrett, Vice President, ESG Lab

November 2013

Contents

Introduction	3
Background.....	3
MMS 2013: HP and Microsoft Serverquarium	4
MMS 2013 Labs	5
Serverquarium Components.....	8
The Bigger Truth	12

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. 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 Microsoft.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

Introduction

Any organization interested in cloud computing expects to gain some kind of economic and agility benefits, but few truly understand the extent of what clouds can achieve. How many workloads can be consolidated on a single server? How complex can the workloads be? How many VMs can run at the same time? How quickly can they expect to change direction? While there is plenty of hype, actual details are often fuzzy. At the 2013 Microsoft Management Summit (MMS), attendees had the chance to see a cloud deployment “in the flesh,” complete with real-time stats on its accomplishments.

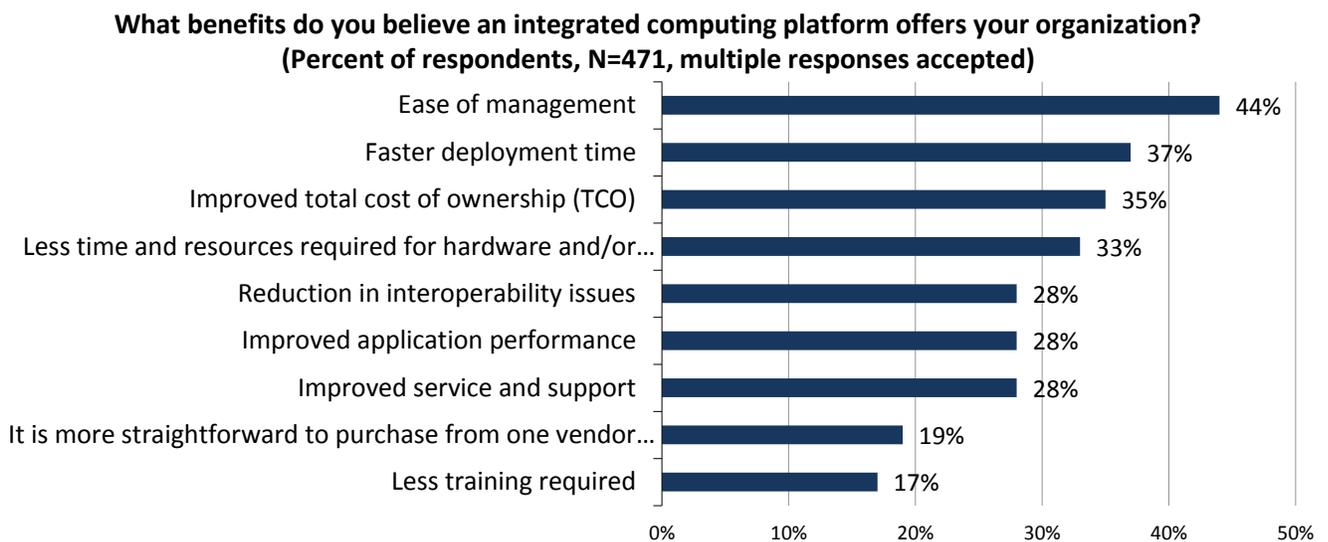
Several ESG employees attended MMS, as they have in past years. This year, once again, some of the best-attended sessions were the hands-on and instructor-led labs, where attendees can try out Microsoft management technologies on their own. And again, the private cloud infrastructure built to deliver these labs was a powerful and efficient integrated solution from HP and Microsoft. This paper looks at this “Serverquarium” infrastructure that powered the lab environments.

Background

While many organizations are enticed by the possibilities that a private cloud offers—consolidation, cost reduction, business agility—most are unsure of how to proceed, and many are simply not equipped to build what they want. Doing it yourself requires a lot of time and expense, from planning to procurement to testing, assembling, integrating, and then managing separate resource silos. It’s a hard job, and often results in inadequate resources for the workloads it supports.

Another option is an integrated computing platform, which ESG defines as a virtual computing infrastructure that integrates hardware and software components that are normally operated separately into single consumable IT systems. Integrated computing platforms can mean the difference between success and failure for a private cloud, offering relief from the tasks of acquiring and integrating individual components. Whether delivered through reference architectures or turnkey converged stacks, integrated platforms can ease the burden of deploying virtualized infrastructures and ensure successful implementation, as ESG research demonstrates. Figure 1 shows the benefits that respondents to a recent survey reported experiencing with integrated platforms, including simplified management, faster deployment time, improved TCO, reduced integration time and cost, fewer interoperability issues, and better application performance.¹

Figure 1. Top Nine Benefits of Integrated Computing Platforms



Source: Enterprise Strategy Group, 2013.

¹ Source: ESG Research Brief, [Integrated Computing Trends](#), March 2011.

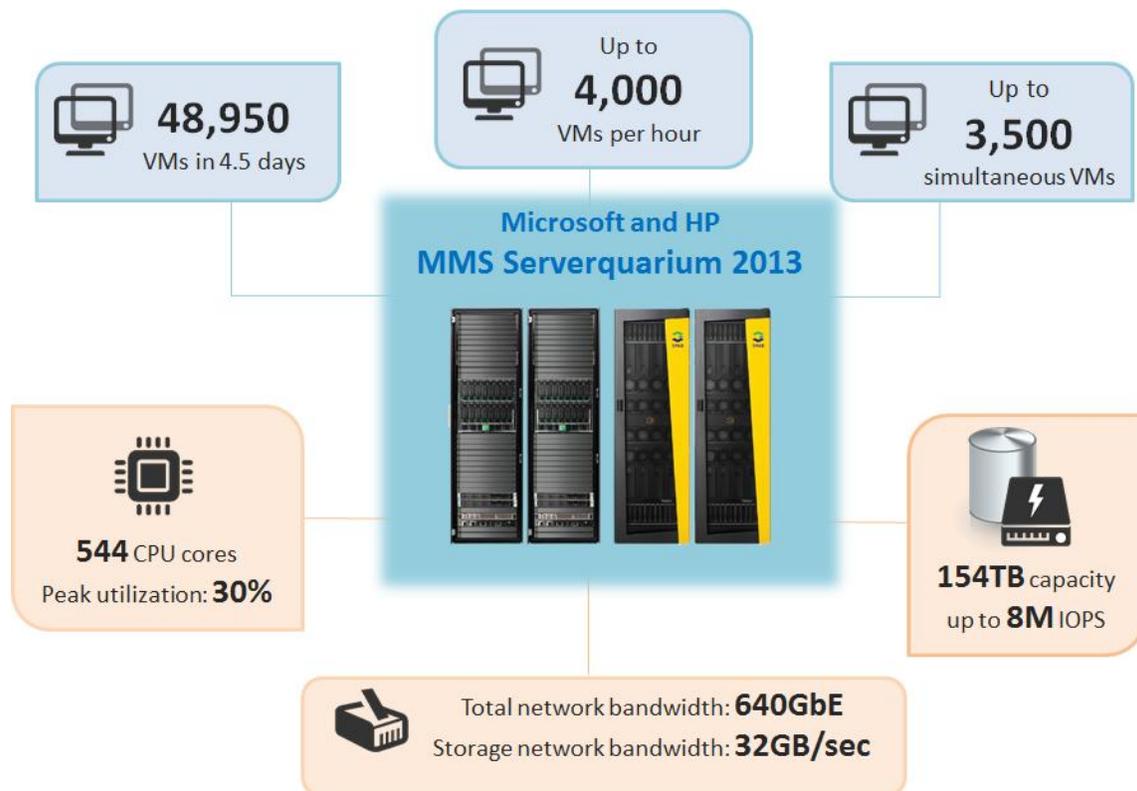
MMS 2013: HP and Microsoft Serverquarium

An integrated computing platform served as the foundation for lab delivery at MMS 2013. MMS is a technical information and training event designed to familiarize customers with the latest technologies and management solutions from Microsoft, its partners, and industry experts. The hands-on and instructor-led labs are a big part of the event, and enable attendees to get first-hand experience with the technologies.

As they have been since 2010, the labs were delivered using a 100% virtualized infrastructure built with HP ProLiant servers and HP 3PAR StoreServ storage, as well as integrated HP and Microsoft management applications. For the third year in a row, a private cloud infrastructure was built and managed by Xtreme Consulting Group and HP experts, using the latest technologies including Windows Server 2012 with Hyper-V, System Center 2012 SP1, and the HP VirtualSystem VS3 for Microsoft. (Note: The new HP Converged System 700x, engineered for Microsoft, delivers these capabilities and replaces HP VS3.) This Microsoft Fast Track certified solution is available to customers today as either a pre-configured bundle or a reference architecture; this “Serverquarium” was not a custom build designed for demonstration, but a solution you can put in your data center today.

Figure 2 shows a few highlights of the Serverquarium. The photo in the middle shows the actual infrastructure that created and delivered all the labs throughout the week—note how little floor space the four racks consume. (In fact, the entire Lab environment could have been delivered in only three racks, but for shipping insurance purposes, the compute portion was split into two racks.) This was a working private cloud on display for all to see, and like other Vegas exhibits, a Plexiglass wall and theater-style rope dividers enabled attendees to view but not touch. Real-time monitoring of the system was displayed on several large screens as well.

Figure 2. Serverquarium Highlights



Cloud Characteristics

The term “cloud” has been overused and misused, resulting in confusion. The essential components of cloud computing include self-service, on-demand provisioning using pools of centrally managed, virtualized resources that are highly scalable and elastic—easily scaling up to support growing needs, and back down again when

demand lets up. Cloud computing also enables a fee-for-use cost model that can benefit both enterprise IT departments and service providers.² The efficiencies that cloud computing offers include reduced costs of equipment, power, floor space, and management. In addition, cloud infrastructures enable greater flexibility to quickly respond to internal circumstances or external market conditions. For example, organizations may experience sudden growth, or need to temporarily scale up or scale back operations to take advantage of new opportunities or manage emerging threats. The jointly designed infrastructure powering MMS 2013 labs met all the private cloud characteristics with enterprise-class performance and resiliency while reducing complexity.

MMS 2013 Labs

ESG can attest to the popularity of the hands-on labs, as, on numerous occasions, we witnessed the sea of workstations packed with attendees trying out various Microsoft technologies. There were 650 lab stations available. Attendees would select a workstation, log into the web-based portal, and choose a lab to work on. From there, the Serverquarium would spin up the resources for that lab and present them to the user in only a few seconds. The photograph in Figure 3 gives you an idea of the scope of the user station deployment.

Figure 3. Hands-on Labs



Lab Workloads

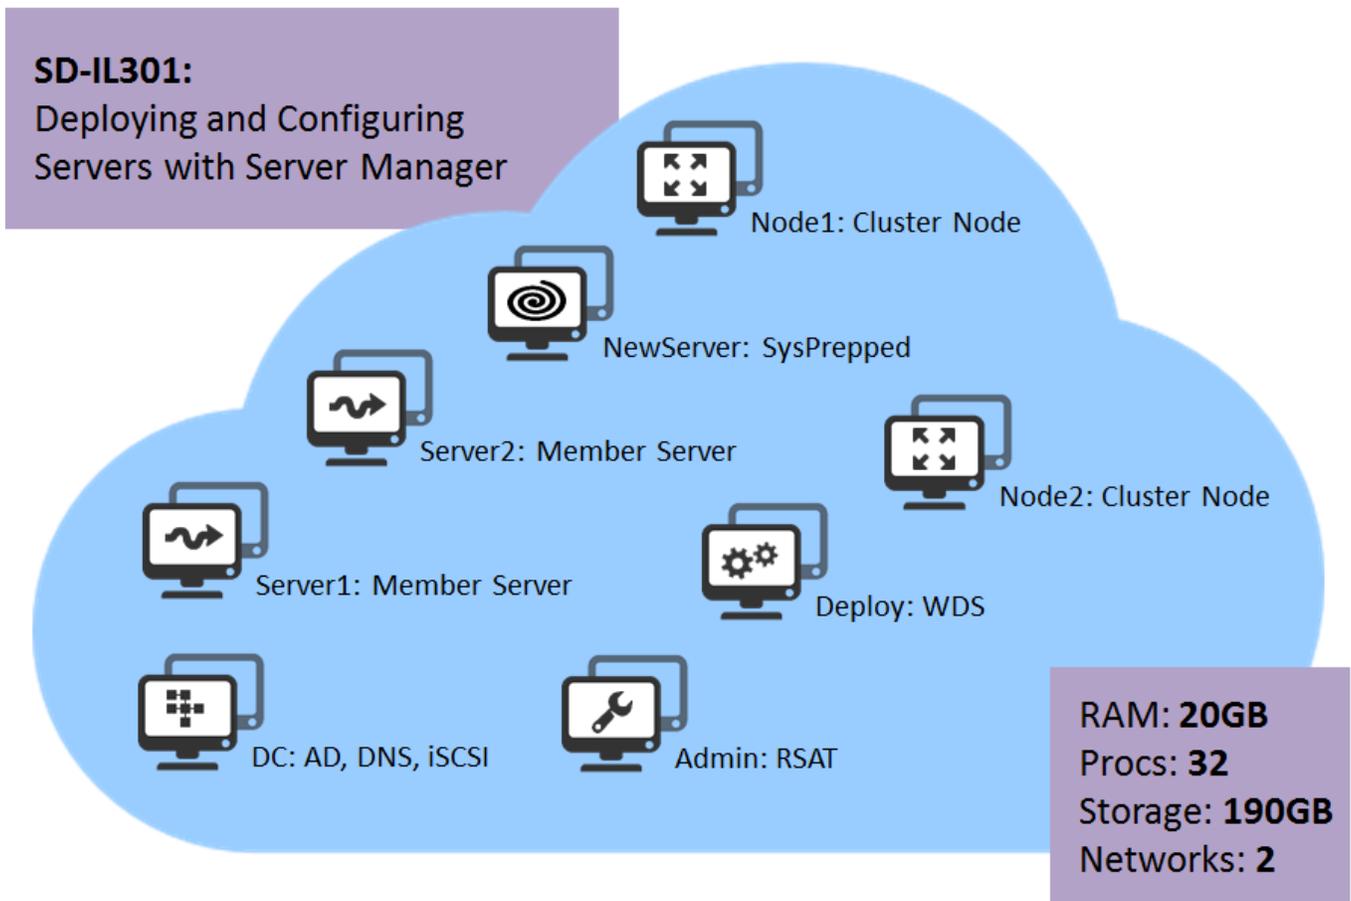
The system was designed to handle worst-case scenarios that scheduling might create. When the labs opened at 8 a.m., there was invariably a long line of attendees waiting to dive in. What if 650 people logged on simultaneously and selected the same lab? The Serverquarium had to be ready for that kind of resource drain. Similarly, all 650 seats could potentially run all different labs, a scenario that would also significantly tax any system. A virtual desktop boot storm can cause a performance slow down, and that's just pushing the *same* image out to many users at once; this lab boot storm was exponentially more complex, as the system had to deploy hundreds of unique workloads, all using between one and 10 VMs, with nothing cached.

² For more information, please refer to ESG Research Brief, [Private Cloud Attributes](#), May 2013.

It is this kind of flexibility—to offer a service catalog, and deliver whatever labs (or other workloads) are requested by any number of users—that cloud solutions make possible. In comparison, as recently as 2009 when the MMS labs were run using physical servers, this type of flexibility was not achievable. The developers had to guess which labs would be most popular and dedicate servers to them. If there were 50 servers running Lab X and 200 people arrived wanting Lab X, only the first 50 could access that lab. The remaining 150 would have to do another lab or come back at another time. This is a clear differentiation of service capabilities. Taking a cue from public cloud providers, forward-looking organizations recognize how cloud solutions like this can help them transform their infrastructures with a goal of delivering IT-as-a-service.

The hefty workloads these labs were running prove out the robustness of this HP/Microsoft cloud solution. Each lab required numerous VMs with specific amounts of memory, processing power, storage, and network resources. For example, even the fairly simple Service Delivery and Automation Lab used four VMs, 15GB of RAM, 15 virtual processors, 280GB of storage, and two virtual networks. Figure 4 depicts the Deploying and Configuring Servers with Server Manager Lab, which was more complex. For each user working on this lab, the HP/Microsoft cloud solution provisioned eight VMs, 20GB of RAM, 32 virtual processors, 190GB of 3PAR StoreServ storage, and two virtual networks. All the labs were run as secure environments, so there were up to 650 isolated Active Directory environments supporting multiple enterprise workloads at the same time.

Figure 4. Lab Workload Example



In addition to the hands-on labs, instructor-led labs were running in five sessions per day in different rooms. To deliver these labs, the Serverquarium provisioned 1,400 VMs to 400 user stations in four locations, changing them over in less than ten minutes between sessions. Jeff Woolsey, Microsoft's Principal Group Program Manager for Windows Server Virtualization, recounted a conversation he had with one of the instructors, who found his lab session ran a little bit short because he had built in time for getting user systems up and running—time that was not needed with this cloud solution.

Why This Matters

While today's enterprise workloads continue to grow in volume and complexity, IT is becoming a victim of its own success. Because IT often performs miraculous feats, like spinning up new applications as VMs and delivering access to production resources over mobile devices, users are beginning to expect that they can get any service, at any time, in an instant. This puts all the more pressure on IT to deliver on service levels that would have been impossible just a few years ago.

The MMS labs are a good example of the kinds of workloads and service levels enterprises are faced with. Organizations looking at projects such as deploying virtual desktops or delivering client virtualization to make users productive from mobile devices can take comfort that there is an answer. The joint HP/Microsoft solution proved its ability to handle enterprise-class workloads and minimize equipment and operational costs, without taking months of planning, purchasing, deployment, and integration.

Serverquarium Components

The infrastructure was based on the [HP VirtualSystem VS3 for Microsoft](#), a turnkey, [Microsoft Fast-Track-compliant](#) cloud solution built for enterprise organizations.³ Various options are available to address customer requirements, and the systems are integrated during manufacturing and delivered as a converged unit, or built with pre-installed and tuned components using a reference architecture. The VS3 includes compute, storage, and network fabric, and delivers the levels of scalability, performance, and integrated management that enterprise organizations require.⁴ The system built for MMS was completely redundant to ensure high availability.

Compute

The compute portion consisted of dual HP BladeSystem c7000 chassis, each configured with 16 BL460c Gen 8 blades in a 32-node Hyper-V cluster, plus a two-node Microsoft System Center 2012 SP1 cluster running on Hyper-V for management. Each blade included:

- 32 logical processors, provisioned using 16 cores in two sockets with Symmetric Multi-Threading enabled
- 256GB of RAM, configured with Hyper-V Dynamic Memory
- 2 local 300GB SAS disks, mirrored for OS boot per blade
- HP I/O Accelerator card, either the 768GB or 1.2TB model

The Accelerator Cards provided very high speed data transfer. This enabled the quick change over between lab sessions that required creating hundreds of VMs.

Storage

The HP 3PAR StoreServ P10800 provided a tiered SAN that was tightly integrated with Microsoft System Center 2012 with SMI-S Providers, simplifying storage management and VM provisioning. HP 3PAR Adaptive Optimization provided policy-driven automated storage tiering, keeping the hottest data on the fastest tier (usually SSD) and moving colder data to lower tiers. Adaptive Optimization provides the opportunity to mix and match storage tiers to balance capacity and IOPS, keeping costs to a minimum. Built-in thin technology provided efficient storage as well by reducing capacity requirements and by automatically and continuously freeing and reclaiming unused capacity. The storage included:

- 2 3PAR StoreServ P10800 arrays in a 4-node, scale-out configuration
- 8GB Fibre Channel SAN configured with Windows Server Multi-Path I/O
- Redundant 16-port FC switches
- Total capacity of 153.6TB:
 1. 64 x 200GB SSD
 2. 128 x 600GB 15K RPM FC disks
 3. 32 x 2TB 7200 RPM SAS disks

During a lab session, the VMs used HP 3PAR StoreServ for storage. However, the interactive user session ran on the I/O Accelerator Cards, keeping that session close to local compute and storage resources.

Networking

The network environments were created with HP VirtualConnect FlexFabric 10GbE and FCoE, which is built into the blade servers. Each node supported up to 30 virtual networks at a time. This dramatically simplified the environment and reduced the number of cables and adaptors. Windows Server 2012 NIC teaming was used to

³ The Microsoft Private Cloud Fast Track program is a joint effort between Microsoft and its hardware partners to deliver validated, pre-configured solutions that reduce complexity and risk for customers implementing private clouds.

⁴ This report describes the MMS Serverquarium as an example of a working private cloud delivering enterprise-class workloads. In addition, ESG Lab completed testing of the HP VS2 for Microsoft in 2012, and validated its fast deployment, performance scalability, and low response times for tier-1 workloads such as Exchange and SQL Server. ESG Lab has also tested most of the components of Windows Server 2012. For more information, please visit <http://www.esg-global.com/lab-reports>.

aggregate the traffic. HP Intelligent Resilient Framework (IRF) was used to flatten the network and eliminate multiple switch hops, enabling faster network performance, recovery time, and live migration between blades.

Integrated Management

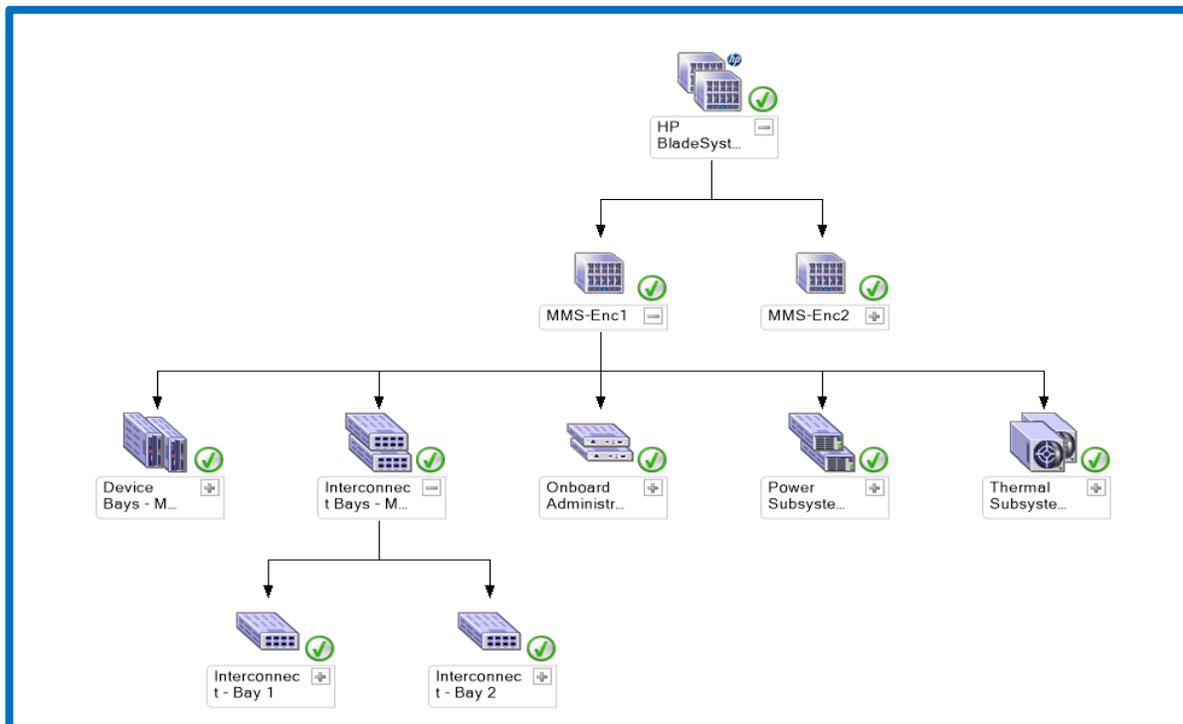
All the functionality needed to manage the unique needs of these labs is available out of the box. With the integration of HP Insight Control (a comprehensive server monitoring and management application) and SMI-S Providers for HP 3PAR StoreServ into Microsoft System Center 2012, the infrastructure was managed end to end through the System Center console. Microsoft Active Directory was used to manage availability for all locations, and Dynamic Memory ensured high efficiency for running VMs. This integration enabled full monitoring and management functionality from a central location.

System Center Components included:

- Operations Manager – to monitor health and performance
- Virtual Machine Manager – to configure and manage virtual server, network, and storage resources
- Orchestrator – for automation, provisioning, and handling portal requests
- Configuration Manager – to manage host configurations and health
- Service Manager – to automate incident management and problem resolution

HP ProLiant and HP BladeSystem Management Packs for Operations Manager were used. As a result, administrators could monitor HP BladeServers and enclosures and get alerts, see a graphical view of all nodes, and directly launch remote management applications such as iLO Advanced or SMH. As Figure 5 shows, in one screen, a topology map showed the system down to individual fans, power supplies, and interconnects.

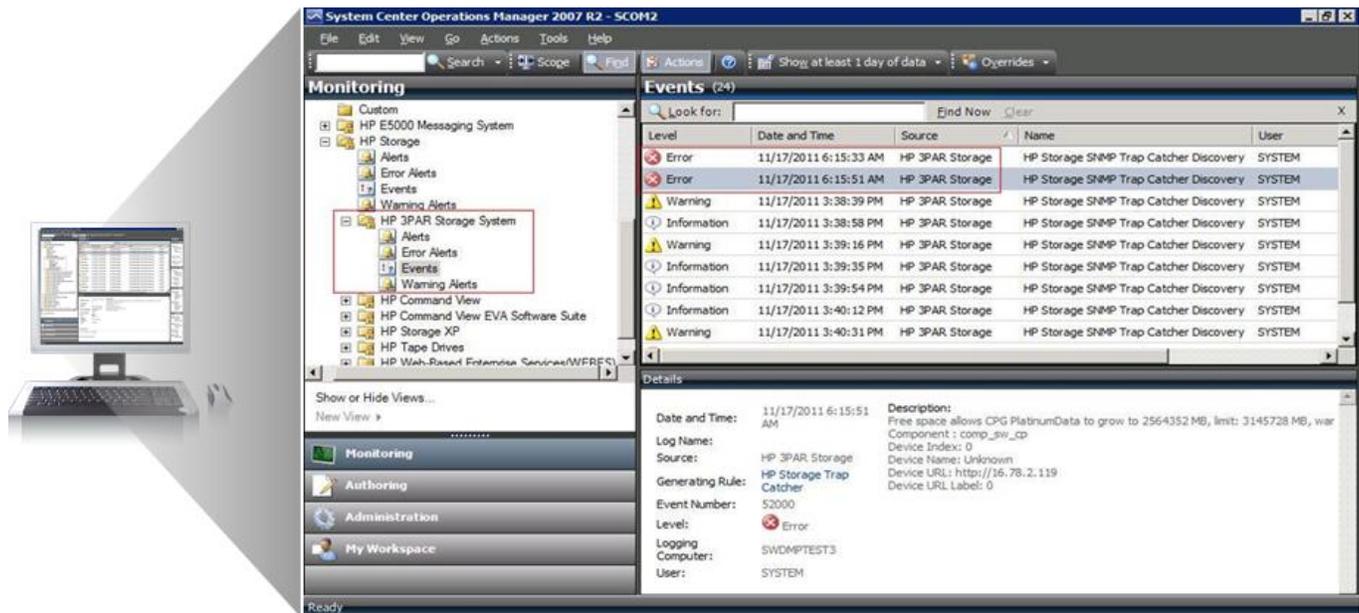
Figure 5. Topology Map



System Center Virtual Machine Manager was used to provision and manage the infrastructure, as well as to monitor and report on VMs throughout the environment. Also, a plug-in enabled administrators to view the HP FlexFabric network from within Virtual Machine Manager.

The HP 3PAR StoreServ was also integrated with System Center Virtual Machine Manager (SC-VMM) through SMI-S Providers and with System Center Operations Manager through HP Storage Management Packs. As a result, SC-VMM could automatically identify and allocate HP 3PAR StoreServ LUNs, create a class of storage, and provision the private cloud. There was a single drive failure during the MMS event. Operations Manager identified the array and drive so it could be replaced, but due to the redundant configuration the failure did not interrupt operations. Figure 6 shows the System Center Operations Manager screen displaying details about the 3PAR storage system.

Figure 6. Integrated Management



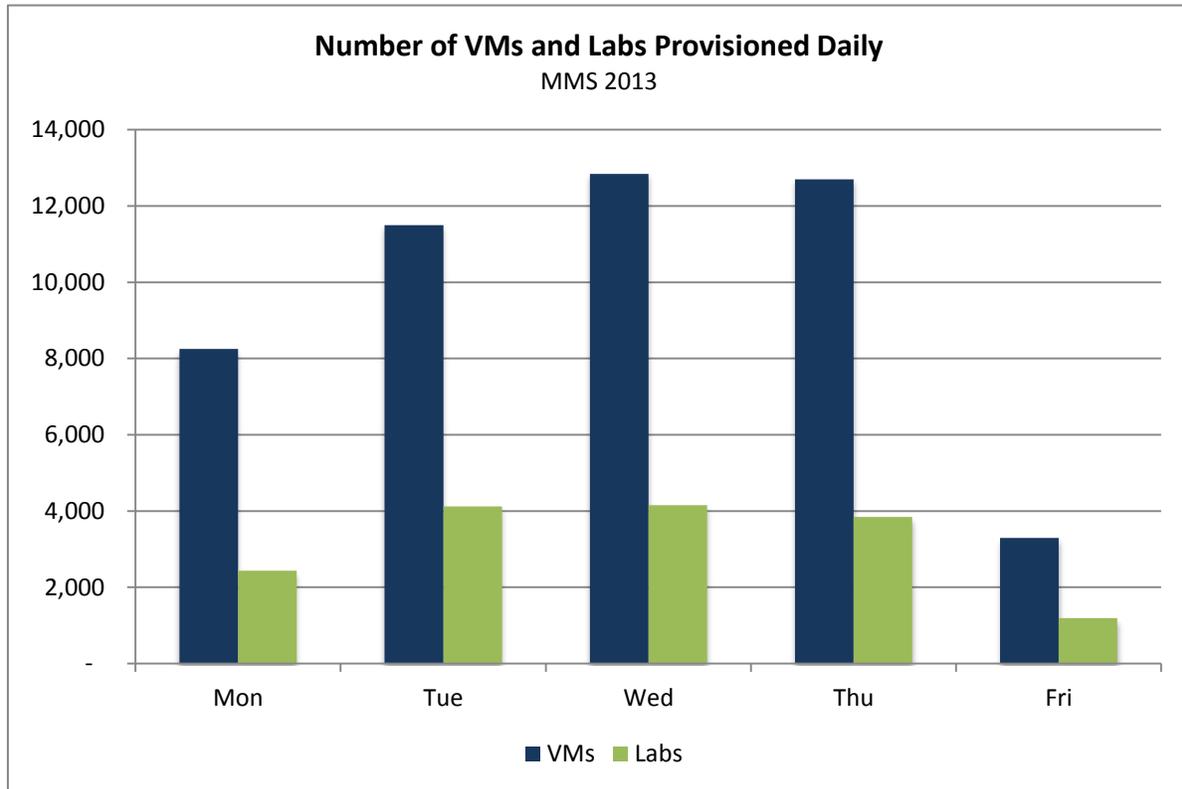
Automation

Automation was essential to delivering these labs. For example, at the end of each lab session, all snapshots of the system state had to be deleted and each workstation reset for the next session. Only a few minutes were available to complete this task; without automated management, this would not have been possible. The inclusion of PowerShell as part of Windows Server 2012 enabled developers to create workflows for Orchestrator to automate creation, deployment, and management of the HP/Microsoft cloud resources.

Selected Statistics

The Serverquarium performed flawlessly throughout the week, pushing heavy workloads with speed and dexterity—and yet, performance monitoring indicated that it was never pushed to the limit. Just the sheer number of VMs and labs run during the week is impressive. Figure 7 shows the daily volume of labs and VMs being generated through four full-day and one half-day session. The total number of VMs was more than 48,500, running more than 15,750 labs.

Figure 7. MMS VMs and Labs Provisioned Daily



Since the Serverquarium was a showcase for cloud deployments as well as a running infrastructure, the team gathered statistics on numerous factors. Among the most compelling:

- Average CPU usage remained at 15%-20%, even with thousands of VMs running at once. Peak CPU usage was only 30%. Clearly, this joint solution had a lot of head room, and could have handled even more than these complex simultaneous workloads.
- The combination of HP components, the 3PAR StoreServ P10800 (in particular due to Adaptive Optimization for automated tiering with SSD), and I/O Accelerator Cards created a system capable of generating up to eight million IOPS with performance to spare.
- Disk queues for each host remained consistently at 1.0; a peak of 1.3 was reached once, when the number of simultaneous VMs reached its peak of 3,500.

Why This Matters

Private cloud deployments often seem like an iffy proposition to customers as they try to wade through the hype and find out what is possible today. They want an idea of what a solution can do, how it will perform, and how it might serve their needs. What is often missing is the chance to see a cloud solution in action.

The Serverquarium at MMS did just that, with real enterprise-class workloads. Customers could see the racks, watch the performance monitors, and experience the service delivery just by logging onto a lab. But not just any vendor can put together this type of system. Years of research and innovation, as well as a deep understanding of customer needs, enabled HP and Microsoft to build and integrate the technologies that make this solution fast, efficient, and easy to manage.

The Bigger Truth

ESG experienced first-hand the power of the HP/Microsoft collaboration. In a small footprint, the Serverquarium spun up hundreds of unique workloads and thousands of VMs at the same time in isolated environments, changed them over in minutes, and did it all again, hour after hour. Speed, simplicity, automation, and efficiency created a balanced environment in a small footprint. While it looks a lot like magic, it came from a long-term partnership of two industry giants working together to provide an integrated system that bridges the gap between the cloud vision and its reality. Said Tom Joyce, Senior Vice President and GM, Converged Systems at Hewlett-Packard, "Simplifying transformation is the name of the game. HP and Microsoft have combined our strengths to deliver innovative solutions that help our customers get the benefits of the cloud quickly and easily."

ESG was impressed with the flexibility, scalability, and high performance this joint solution demonstrated. Changing over tier-1 workloads in minutes between lab sessions, deploying 3,500 VMs at one time, and provisioning more than 48,000 VMs in four and a half days are just a few of the notable results. This system also proves that Windows Server 2012 with Hyper-V is an enterprise-class virtualization platform that is fully ready for prime time. Combining it with the robust and full featured HP BladeSystem, 3PAR StoreServ storage, and FlexFabric solutions created an integrated cloud solution that alone could power many businesses. What was not measured was the system's ability to deliver the kind of high availability that most enterprises need, since if a lab failed, it could simply be recreated for the user. This is quite different from having your Exchange deployment fail, for example. However, the redundancy of the entire system indicates that it was up to the task.

The MMS lab infrastructure has changed dramatically over the years. As recently as 2009, the labs were run on distributed servers, with one per seat in the hall; as a result, for 650 attendees, 650 servers were needed. Virtualization and Cloud Service Design methodology have enabled this consolidation and delivered huge power savings, starting with the first highly virtualized implementation in 2010, when power consumption dropped by more than 14X. This year's complete private cloud solution also delivered more than 15,000 labs and provisioned up to 4,000 VMs per hour without ever breaking a sweat.

The best news is that Microsoft and HP have worked closely to deliver a complete private cloud infrastructure that customers can deploy today. ESG recommends that any organization looking to improve service delivery and operational efficiency take advantage of what these industry giants have done together. For many, this will offer the opportunity to be more flexible, deliver IT services faster, and minimize both infrastructure and management costs. For IT professionals who generally have to manage expectations in two directions—to users and to senior management—this solution can solve a lot of problems.



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

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