

## INTRODUCTION

The flexibility, and stability provided by server virtualization has allowed it to migrate from being a niche technology to now being a key component in many enterprise infrastructures. As IT responds to the ever-increasing business pressures, virtualization is now being used to ease server management and reduce overall costs – while at the same time delivering reliability, scalability and security. Microsoft's Hyper-V hypervisor provides the isolation and performance required to consolidate servers and to enable high-performance virtual environment that are key to delivering high availability to mission-critical applications. Additional use-cases for server-based virtualization include application mobility, personalization, disaster recovery, legacy support, workforce integration, and operating system migration. Added to this list is an ever increasing scenario that integrates Linux applications with the industry-leading application platform of Microsoft Windows Server. These mixed operating system workloads are of interest to organizations looking for a proven, simple way to integrate the functionality provided by the applications of one operating system with the applications, management structure and toolsets provided by another.

## HETEROGENOUS WORKLOAD VALIDATION – LAMP/ACTIVE DIRECTORY

To evaluate the feasibility and benefits of virtualizing these heterogeneous workloads, an independent team of engineers from Vision360, LLC was engaged to perform a series of validation related tasks. The purpose of this effort was to demonstrate the value of an integrated operating system environment and to compare the performance of a virtualized workload with a similar workload deployed on physical hardware.

As part of the validation Hyper-V was used to host a partition containing a SUSE Linux Enterprise Server 10 virtual machine. Installed in this virtual machine were the application components of a LAMP stack. The LAMP stack, which is a combination of Linux/Apache/MySQL/PHP (or Perl, or even Python), is one of the most wide spread application stacks utilized in enterprises that have deployed Linux. It allows webmasters to build websites and run complex applications such as, but not limited to: content management systems, webmail interface or wikis.

For many environments, the versatility, security and scalability of the LAMP stack makes it an ideal platform for delivering web-based services to both internal and external audiences. The stack's proliferation across all markets, industries and company sizes makes it an ideal candidate for this validation effort. These LAMP components were then integrated with directory, validation and authentication services provided by Microsoft Active Directory running on top of Windows Server 2008 Enterprise and delivered via a child partition created and managed by Hyper-V. After this integrated heterogeneous application stack was deployed and before starting the validation effort, each solution component was individually tested to ensure proper 'standalone' functionality.

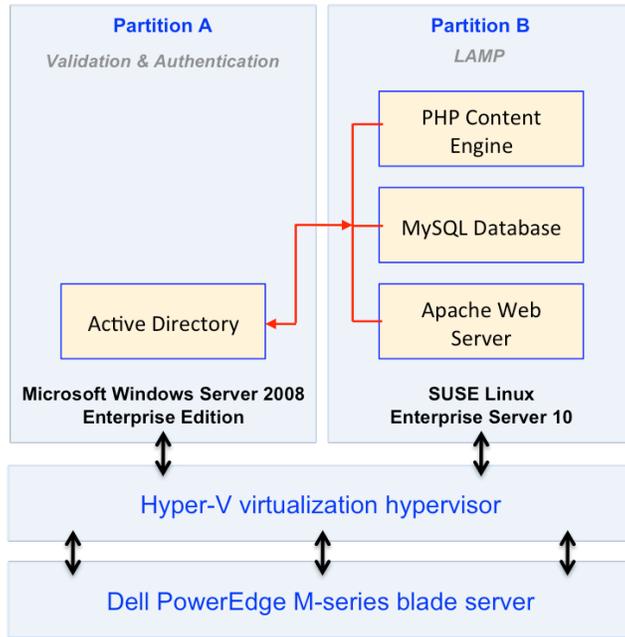


Figure 1: LAMP & Active Directory Installation

The validation efforts were conducted in the joint Microsoft-Novell Interoperability Lab located in Cambridge, MA. A series of automated installation scripts – leveraging traditional management tools - were used to ensure that the operating systems, virtualization hypervisors and all applications were uniformly installed. Two different models of Dell PowerEdge M-series Intel processor-based blades provided the required computing power. The blades had similar components, but differed in form factor. The form factor of the half-height M610 model is designed to allow enterprises achieve greater server density, while the full-height M710 model affords greater expansion opportunities.

As detailed in Figure 2, both models contained Intel Xeon processors and were configured with twenty-four gigabytes of memory, leveraged internal drives and utilized onboard gigabit Broadcom Ethernet adapters. These servers were placed in the same cabinet and connected to the same blade chassis.

	Dell PowerEdge M610 Blade Server	Dell PowerEdge M710 Blade Server
Processor	Intel Xeon E5506, 2.13Ghz 4M Cache 4.86 GT/s QPI	Intel Xeon E5506, 2.13Ghz 4M Cache 4.86 GT/s QPI
Memory	24 GB 12 x 2GB 1066Mhz UDIMMs	24 GB 12 x 2GB 800Mhz RDIMMs
Storage	Internal 250GB 7.2K RPM SATA	Internal 146GB 10K RPM SAS
Network	Broadcom 57710 Dual Port 10GbE I/O Card	Broadcom 5709 Dual Port GbE w/ TOE Card

Figure 2: The hardware configuration used in this validation environment.

Following the automated installation process described earlier, the Active Directory server was configured and with a full set of heterogeneous data types similar to those found in most large-enterprise deployments. A typical intranet application was then installed on the web, database and content-serving components of the LAMP stack, and restrictions were made as to the user's

with access authority. Authorization for users was linked to a successful validation of credentials stored in the directory server.

## HYPER-V TECHNICAL OVERVIEW

Before analyzing the results of the validation effort it is necessary to gain a basic understanding of the functionality delivered by Hyper-V. Microsoft's Hyper-V is an integrated feature of Windows Server 2008 and is a "bare-metal" hypervisor that sits directly on the compute hardware - between the physical layer and the operating system. Hyper-V provides full isolation and near-native performance this is achieved through the hypervisor with manages resource scheduling and hypercalls to the underlying processors and memory.

As shown in the high-level architecture of Hyper-V, Figure 3, each deployed hypervisor has at least one root partition, running Windows Server 2008. Hyper-V's virtualization components execute in this root partition and have direct access to the underlying computing hardware. The child partitions (VM's) are created by the root partition and they host the guest operating systems. A created VM does not have access to the physical processor; instead, it has a virtual view of the processor and runs in a guest virtual address space. The hypervisor handles the interrupt requests to the processor, and redirects them to the appropriate requesting partition using a logical interrupt controller (IC). Hyper-V can hardware accelerate the address translation between various guest virtual address-spaces by using I/O Stack memory management which is independent of the memory management hardware used by the physical CPU.

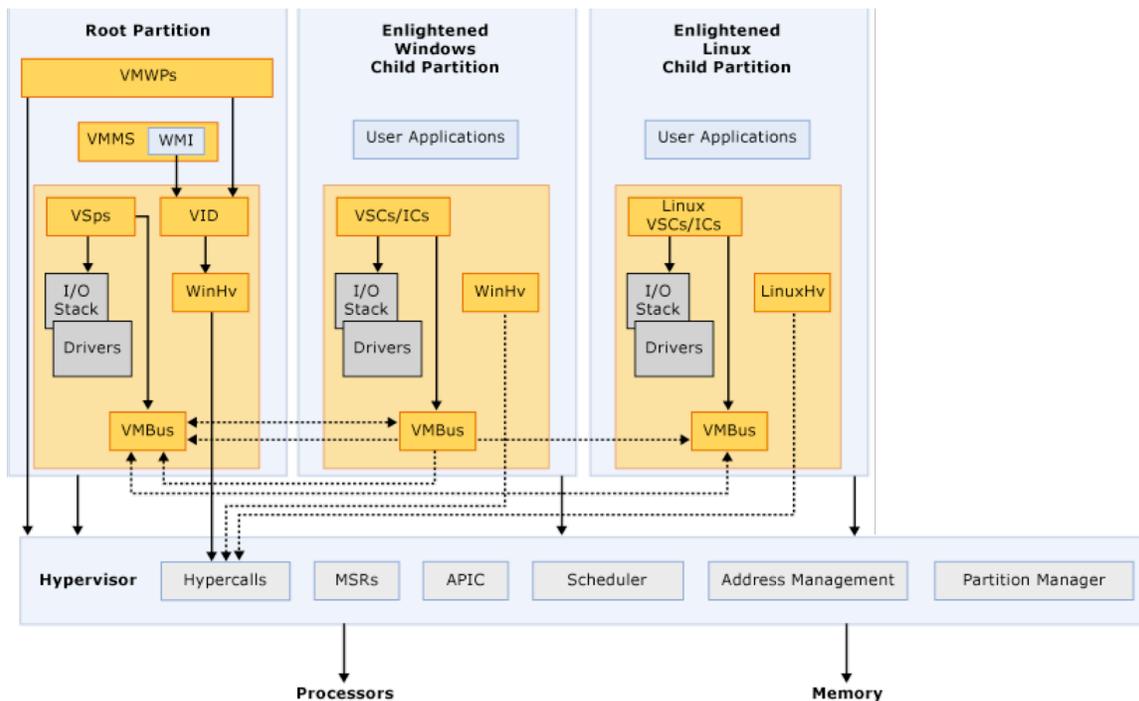


Figure 3: Hyper-V high-level architecture

As noted earlier, each individual VM's does not have direct access to hardware resources, but instead it has a virtual view of the available resources, as virtual devices. Any request to these virtual devices is redirected via the VMBus to the devices in the root partition, which will manage the requests. The VMBus also enables optimized bidirectional communication between VM partitions, eliminating the need to traverse the hypervisor and physical hardware. Additionally,

the VMBus provides full isolation ensuring that a virtual machine will only see the broadcast traffic of the other VMs connected to its virtual switch, meaning unicast traffic is fully isolated. Unicast traffic isolation is significant in that it prevents a compromised VM from easily capturing the network traffic of all other VMs on the same virtual switch.

## VALIDATION CONCLUSIONS

To determine the feasibility of deploying a virtualized heterogeneous environment, a series of 'real-life' use-cases were developed and implemented on the integrated stack environment. In addition to ensuring basic functionality, these use cases were designed to: stress the environment to evaluate relative performance, compromise one or more application components of the stack to validate isolation, and to verified vm to vm communication using the VMBus infrastructure. These use-cases were run numerous times, with varying user loads, during the validation effort and the results of each were noted. In addition, the results were compared to those done on a similar environment that was deployed in a non-virtualized manner.

Throughout the validation effort it became clear that there were no major obstacles or impediments to deploying a heterogeneous Windows/Linux solution stack on Hyper-V. In fact there were very apparent benefits of this approach, some of these included advantages related simple virtualization and the ability of hypervisors to create guest operating system partitions:

- Consolidation of numerous servers and resources on to a few number of physical resources
- The ability to deploy, manage and monitor a Linux environment using the proven Windows systems management tools familiar to most enterprises
- Full VM isolation limited the damaging effects that a failed or breached server could have on the rest of the compute environment

Additionally there were some benefits of this deployment model that specifically enhanced the performance of this heterogeneous application stack. There were demonstrated performance gains and increased security surrounding the authentication and validation process. Since all communications between VM's occur directly using the VMBus – there was a noticeable decrease in authentication response when compared to a physical server deployment as the VMBus eliminated the need for messages to connect through the physical network between two physically separated servers. Additionally, since these VM to VM communications – and the user credentials passed as part of this process - were never exposed to the Ethernet network, this scenario was deemed to have reduced the ability of a breach targeted at user login credentials or certificates. For the use cases that separated the stacks onto multiple servers - the backchannel between the blades delivered similar performance (and relative security) benefits as it provided the same category of functionality as the VMBus for traffic passing between servers.

This validation effort has proven that Hyper-V is a potentially valuable tool to consolidate multiple server roles onto a single physical machine and to effectively manage the applications of multiple different operating systems—in this case Windows Server 2008 and SUSE Linux Enterprise Server 10.