Microsoft | Virtualization

Implementing a High Performance, End-to-End Virtualized Solution Using Windows Server 2008 R2 Hyper-V, IBM® XIV® Storage, and Brocade Fabric Solutions

Configuration and Best Practices Guide

Authors: Eric B. Johnson[i], Taylor Brown[ii]
Editors: Bill Watson[iii], David Hartman[iv]
Date: August 2010
Rev. 2.0

Table of Contents

Executive Summary 3

Intended Audience 4

Microsoft Windows Server 2008 Hyper-V Introduction 4
- Core Hyper-V Features and Benefits 5
- New Hyper-V Features in Windows Server 2008 R2 5
- Ease of Manageability 5
- System Center Virtual Machine Manager (VMM) 6
- Dynamic and Highly Available 7
- Reduced TCO/Improved Flexibility 8

Compute, Storage and Network Hardware Overview 8
- IBM® Enterprise Servers 8
- IBM® XIV® Storage with Brocade SAN 9
- Hyper-V™ Virtual Networks 11

Layout 13
- Infrastructure Servers 13
- Virtual Machines 13
- Virtual Machine Storage 14

Installation and Configuration 15
- Hyper-V™ Host Installation 15
- Networking Configuration 16
- Storage and Host Cluster Configuration 17
- VM Guest Cluster Configuration 20

Test Objectives 21
- Enabling Cluster Shared Volumes (CSV) 21
- Virtual Machine Cluster Failovers 22
- Virtual Machine Quick Migration 22
- Virtual Machine Live Migration 23
- Hyper-V™ Hot Add Storage 23

Responding To Growth 23

Monitoring 26

Summary 28

Executive Summary

Today, virtualizing the data center is critical to the success of enterprises looking to deliver a high quality of service at a low cost to their customers. In a 24x7x365 business climate, high availability and reliability are absolute requirements that enterprise organizations must deliver to ensure their operations are available when their customers need them. These organizations look to deliver their services without adding additional expenses for personnel, infrastructure, or complex processes that take current resources away from key revenue generating applications, as well as current tactical and strategic tasks.

This white paper describes a reference architecture for a virtualized, highly available and reliable computing platform solution for enterprises using Microsoft® Windows Server® 2008 R2 with Cluster Shared Volumes including the Microsoft® live migration technology, and was created through a cooperative effort between Microsoft, IBM, and Brocade. Throughout this document the reader will find best practices and time saving steps for deploying and managing an efficient end-to-end virtualized IT environment using Microsoft Hyper-V™ and IBM XIV Storage Systems. This document showcases Microsoft's Hyper-V role on IBM's XIV Storage System connected via Brocade SAN switches and host bus adapters.

Microsoft Windows Server 2008 Hyper-V with live migration moves VM's seamlessly between server data pools in a virtualized environment. This integration enables high availability and reliability with minimal intervention by IT staff, thereby reducing risk, cost, and complexity found in traditional high availability solutions. By leveraging Microsoft Hyper-V, customers reap additional rewards including simplifying complex operational processes and a lower overall Total Cost of Ownership (TCO). By virtualizing the operating system and line of business applications, enterprise organizations are able to effectively implement green business practices by reducing data center computer hardware expenses, power consumption, cooling and space requirements, and administrative overhead.

Along with Windows Server 2008 R2, Microsoft enhanced their System Center product suite with System Center Virtual Machine Manager 2008 R2 to provide seamless management of virtualized assets. SCVMM manages both Microsoft Hyper-V and VMware virtual environments. More information about Microsoft Virtualization, System Center and Windows Server 2008 R2 Hyper-V can be found at http://www.microsoft.com/virtualization.

Virtualizing the server is just the start of providing a simple and cost effective solution. At the foundation, the storage platform also requires a virtualized approach to deliver similar value. To take advantage of virtualized enterprise storage, we opted to utilize the enterprise-proven IBM XIV Storage System. It is an industry-leading answer to powerful end-to-end virtualization demands which span the entire infrastructure platform.

The IBM XIV Storage System is a revolutionary, easily managed, high-end, open disk system, designed to meet today's highly demanding IT challenges. Its design — a virtualized grid of standard Intel®/Linux® components connected in any-to-any topology using gigabit Ethernet — represents the next generation of high-end SAN (Storage Area Networks) systems. XIV is the industry's first single-architecture storage system, providing the highest level of performance, reliability and features, along with management simplicity and exceptional TCO.

The XIV system scales non-disruptively from partial to full rack configurations (27 TB to 161TB usable), meeting growing capacity needs on-demand. Its unique virtualized grid architectural design scales throughput and IOPs proportionally along with capacity, providing ample resources needed to keep performance high. More information about the IBM XIV Storage System can be found at http://www-03.ibm.com/systems/storage/disk/xiv/index.html.

A full implementation of this solution was deployed, tested, and evaluated by Microsoft and IBM engineering staff located at Microsoft's Main Campus in Redmond, Washington. The referenced architecture used in this solution featured Microsoft Windows Server 2008 R2 Hyper-V, IBM enterprise-class System x servers and XIV Storage System as well as Brocade Communications SAN, LAN and Converged Networking technologies.

Intended Audience

IT administrators looking to gain experience with Microsoft virtualization technology are the intended audience for this document. Those new to Microsoft's Hyper-V technology will find many of the current hypervisor or virtualization concepts apply. In any event, document viewers should be comfortable with the following:

- General virtual server configuration and terminology
- General storage, SAN, Fibre Channel, Ethernet, and iSCSI knowledge

Microsoft Windows Server 2008 R2 Hyper-V Introduction

Microsoft Windows Server 2008 R2 Hyper-V builds on the architecture and functions of Windows Server 2008 Hyper-V by adding multiple new features that enhance product flexibility. The adoption of virtualization in the enterprise has increased flexibility in deployment and life cycle management of applications. IT professionals deploy and use virtualization to consolidate workloads and reduce physical server sprawl. Additionally, users can deploy virtualization with clustering technologies to provide a robust IT infrastructure with high availability and quick disaster recovery.

Windows Server 2008 R2 Hyper-V further increases flexibility with Live Migration technology. Live Migration is integrated with Windows Server 2008 R2 Hyper-V and Microsoft Hyper-V Server 2008 R2. With Hyper-V Live Migration, a running virtual machine (VM) is moved from one Hyper-V physical host to another, without any disruption or perceived loss of service. IT professionals increasingly look to Live Migration to create a dynamic and flexible IT environment that responds to emerging business needs. Live migration provides the core technology required for VM mobility and high availability for virtualized workloads during physical computer maintenance.

Core Hyper-V Features and Benefits

Windows Server 2008 R2 Hyper-V includes five core areas of improvement for creating dynamic virtual data centers:

- Increased availability for virtualized data centers
- Improved management of virtualized data centers
- Increased performance and hardware support for Hyper-V Virtual Machines
- Improved virtual networking performance
- A simplified method for physical and virtual computer deployments by using VHD files

Here is a compilation of what is at the heart of Hyper-V technology available in Windows Server 2008 and R2:

- 64-bit native hypervisor-based virtualization
- The ability to run 32-bit and 64-bit virtual machines concurrently
- Large virtual machine memory support
- Virtual machine snapshots to capture the state of a running virtual machine and revert the virtual machine to a previous state quickly and easily
- Runs on all roles of Windows Server 2008 including Server Core
- Failover cluster support for high availability virtual machines
- Move or migrate virtual machines between clustered nodes
- Microsoft Management Console (MMC) interface
- Fully Programmable WMI interface for easy scripting

New Hyper-V Features in Windows Server 2008 R2

Additionally, Microsoft has added several innovative high availability features in Windows Server 2008 R2 to enhance Hyper-V's manageability and availability, including:

- Cluster Shared Volumes (CSVs) – all nodes can read and write to CSVs at essentially the same time
- Live Migration of running virtual machines (VMs) – zero perceived downtime for VM migrations between cluster nodes
- Hot Add Storage – ability to add Virtual Hard Disks (VHDs) to running virtual machines
- Quick Storage Migration (requires VMM) – quick VM storage transfer from one host to another
- Expanded support for iSCSI SANs (Only with SCVMM)

Ease of Manageability

IT administrators working with the combination of Microsoft Hyper-V and IBM XIV Storage Systems will benefit from an easy-to-understand and simple solution which offers the benefits of:

- Server consolidation
- Application consolidation
- Lower total cost of ownership (TCO)- less maintenance, higher levels of space efficiency, and lower power consumption are just a few of the facets
- True storage virtualization – no need to manage physical-to-virtual entity mappings, RAID configurations, etc…
- Easy server cloning – SCVMM's cloning ability uses Microsoft's System Preparation Tool Sysprep - utility used to prepare Windows systems for cloning by assigning an unique Security ID (SID)

Note: For environments without SCVMM, administrators can manually copy the VHD to achieve the same cloning results. Sysprep however is still required.

The combined Microsoft and IBM solution helps reduce the risk, cost, and complexity of a customer's IT platform, so IT administrators become more efficient in producing results and can spend less time on operations, and more time on higher-value strategic initiatives.

System Center Virtual Machine Manager

To assist with ease of manageability, Microsoft introduced System Center Virtual Machine Manager 2008 R2 (SCVMM), which is an enterprise-class management software solution that enables administrators to easily and effectively manage both physical and virtual environments from a single management console. The key capabilities of VMM 2008 R2 are included in **Table 1** below:

Features	Benefits
Enterprise-class Management Suite	Manages both Hyper-V and VMware ESX virtualization environments
Intelligent VM Placement	Supports placement of virtual machines on the optimal server
System Center Operations Manager 2007 Integration	Provides proactive management of both virtual and physical environments through a single console
Performance & Resource Optimization (PRO)	System Center provides alerts to VMM remediation actions
Native P2V/V2V Migration	Native capability for Physical-to-Virtual migration and Virtual-to-Virtual migrations
Failover Cluster Integration	Supports set up of high availability clusters and manages live migration of virtual machines.
Automation	Easy automation capabilities leveraging Windows PowerShell.

Table 1: Microsoft SCVMM 2008 Features and Benefits

SCVMM helps ensure a high degree of operational continuity with its fully cluster-aware support. It can manage Hyper-V clusters as a single object for managing mission- and business-critical virtual machines. Combining the hardware redundancy of IBM servers, IBM XIV Storage, and Brocade fabric, as selected for this study, achieves a virtualized, flexibly, highly-available mission-critical solution.

Dynamic and Highly Available

Customers enjoy greater flexibility and less concern regarding their mission critical server availability that can be attained as a result of this software and hardware partnership. Additionally, the solution alleviates unplanned downtime that can have a serious impact on business productivity. To help facilitate reduced risk, Microsoft Windows 2008 R2 Hyper-V on IBM XIV Storage System provides:

- The ability to add, remove or replace physical storage devices without disruption to availability or system performance
- Storage-based backups leveraging Microsoft Hyper-V VSS-based snapshots
- Migration of virtual machines between servers
- The ability to quickly provision new virtual machines
- Sophisticated XIV grid architecture - providing HA and reliability in the face of disk and/or module failures that would normally cause other storage solutions to lose data, or degrade performance as the information on the drive is rebuilt using a hot spare
- XIV elimination of hot spots - automatically self-tunes and delivers consistently high performance with unique, automated data placement and balancing using all key system resources

Reduced TCO/Improved Flexibility

Reduced TCO and maximal return on investment (ROI) are key considerations for every IT administrator, especially in today's difficult economic climate. For the savvy, budget-conscious customer, this combined Hyper-V and XIV solution helps to deliver long term tangible returns.

Up until now, this paper has focused predominantly on the software components of this solution. Hyper-V depends on a highly available and reliable hardware platform to operate most effectively. IBM and Brocade help to meet this critical need by providing the following key storage array-based benefits:

- Reduced hardware requirements including spare parts and occupied space
- Thin Provisioning – storage and server hardware, utilizing what is needed when it is needed
- Very High Density Slower Rotation (VHDSR) drives reducing power and cost
- High Scalability
- Ability to quickly and transparently move virtual machines from one physical server to another

Compute, Storage and Network Hardware Overview

Hyper-V, IBM XIV and Brocade components operate together to form a mutually beneficial end-to-end virtualized hardware and software solution that helps optimize server and storage resources while resulting in lower customer costs. The overall tested Hyper-V reference configuration as depicted in Figure 1 below consists of three System x3950 M2 servers, two Brocade 8000 switches and one XIV Storage System.

IBM Enterprise Servers

The reference server configurations below use the IBM System x3950 M2. Engineered with the needs of enterprise organizations in mind, the IBM x3950 M2 is scalable, efficient and highly reliable. This 4th generation X-Architecture® enterprise server combines 64-bit performance in a balanced design. The x3950 M2 can help organizations meet business demands with confidence and since many organizations require servers that expand as business grows, the x3950 M2 provides the flexibility to run more applications on the same piece of hardware. The features of the x3950 M2 deliver a scale-up platform that is optimized for large enterprise applications and server consolidation workloads via large-scale Hyper-V deployments.

As required by Microsoft Hyper-V, the servers' x64 processors are enabled in the BIOS for hardware-assisted virtualization and data execution prevention (DEP).

Individual server components are configured as follows:

Server 1 - IBM x3950 M2 (7233)

- 4 – Intel 2.66GHz CPUs
- 64GB RAM
- Brocade 825 Dual Port 8GB HBA
- Brocade 1020 Dual Port 10GB Converged Enhanced Ethernet Adapter

Server 2 - IBM x3950 M2 (7233)

- 4 – Intel 2.66GHz CPUs
- 64GB RAM
- Brocade 825 Dual Port 8GB HBA
- Brocade 1020 Dual Port 10GB Converged Enhanced Ethernet Adapter

Server 3 - IBM x3950 M2 (7141)

- 4 – Intel 2.4GHz CPUs
- 64GB RAM
- Brocade 825 Dual Port 8GB HBA
- Brocade 1020 Dual Port 10GB Converged Enhanced Ethernet Adapter

IBM XIV Storage with Brocade SAN

Another critical part of this Hyper-V reference solution is the IBM XIV Storage System. IBM's market success with the IBM XIV can be attributed to many industry leading capabilities, the most notable for customers being the simplicity and ease-of-use found in the IBM XIV Management Console. Storage configuration and management of the XIV is arguably the easiest yet most expansive among enterprise-class storage systems. An Adobe Flash video presentation is available, which covers the major tasks involved in XIV storage administration.

For those who enjoy the efficiency and convenience of running scripts and command line interfaces (CLI), XIV offers a comprehensive platform via the XIVCLI. View the XIVXCLI Utility User Manual for detailed information.

IBM XIV applies automated data distribution, unique caching, and other innovations to eliminate I/O hotspots and deliver single-tier storage featuring high-end performance, reliability, availability, and scalability, with an exceptionally low TCO. The XIV architecture can scale performance with capacity via sophisticated engineering algorithms. Thus no manual tuning is required enabling capacity growth without additional complexity.

IBM XIV comes in easily scalable partial-to-full racks ranging from 27 to 161TB of usable capacity. Also, IBM XIV offers an innovative approach to software licensing – every feature and function that is traditionally a chargeable item is included at no extra charge, including:

- Snapshots
- Thin provisioning
- Replication
- Advanced management, monitoring and alerting
- Cloning
- Full support of Microsoft technologies including Geocluster support, VSS and MPIO

For the network, the Brocade 8000 switch was utilized. The Brocade 8000 switch is a top-of-rack link layer (Layer 2) Converged Enhanced Ethernet (CEE)/Fibre Channel over Ethernet (FCoE) switch. This multiprotocol switch has 24 10 gigabit Ethernet ports and 8 8Gbps Fibre Channel ports. The Ethernet ports, which are architected in a cut through and non-blocking architecture, support Ethernet, Converged Enhanced Ethernet (CEE), Fibre Channel, and Fibre Channel over Ethernet (FCoE) protocols.

Figure 2 below depicts the physical architecture of the tested configuration:

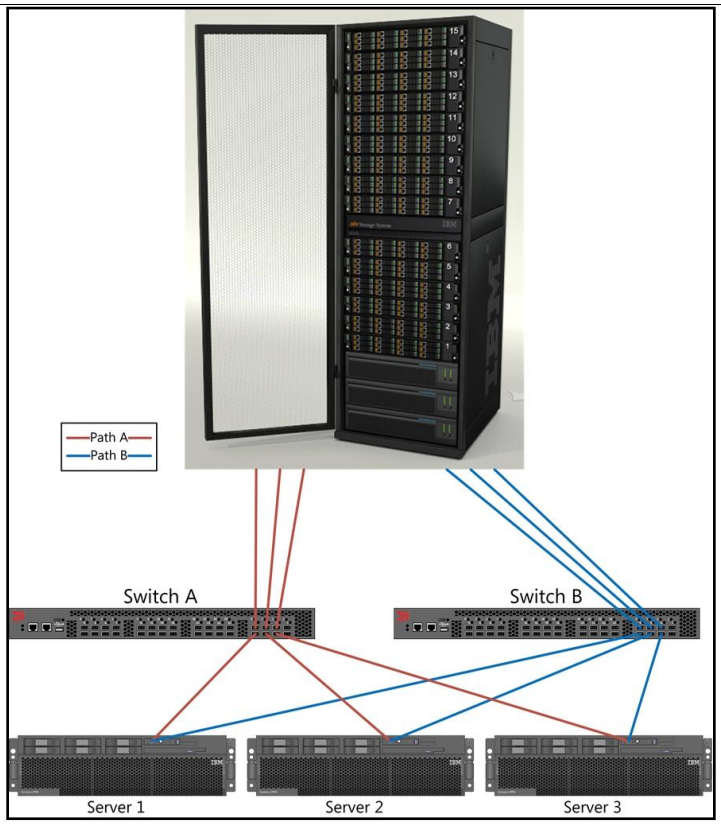

Figure 2: Reference platform for the Hyper-V end-to-end virtualization solution.

Individual storage and SAN components are configured as follows:

IBM XIV A14/2810 Storage System

- Version 10.1.0.a
- 79TB of capacity
- 6 – 4GB Fiber Channel connections active
- 6 – 1GB iSCSI connections active

Brocade Fiber Channel Fabric

- 2 - Brocade 8000 switches (24 – 10GbE ports, 8 – 8Gbps Fiber Channel ports)
- IBM XIV Storage connected via three 4Gbps FC paths to each switch
- Servers connected via one 8Gbps Fiber Channel path to each switch

Hyper-V Virtual Networking

Microsoft Hyper-V is flexible with regards to the allowed networking topologies. In the testing, all three types of virtual networks are implemented:

- **External Virtual Networks** – allows virtual machine communication with external servers as well as the parent operating system
- **Internal Virtual Networks** – allows communication between virtual machines located on the same physical server and parent operating system
- **Private Virtual Networks** – provides isolated communication to keep traffic separate from the external and host or parent traffic

For detailed Hyper-V network configuration information and guidelines, see the following Microsoft TechNet articles:

Hyper-V: Configuring Virtual Networks

Hyper-V: Live Migration Network Configuration Guide

The tested Hyper-V reference architecture uses the configuration as depicted in **Figure 3** below, and was comprised of the following:

- 1GB Dedicated management operating system connection per server
- 2 – 10GB Converged Enhanced Ethernet connections per server

    1. VLAN 10 – corporate (public) network
    2. VLAN 100 – cluster private network
    3. VLAN 101 – Live Migration network
    4. VLAN 102 – Storage (iSCSI network)

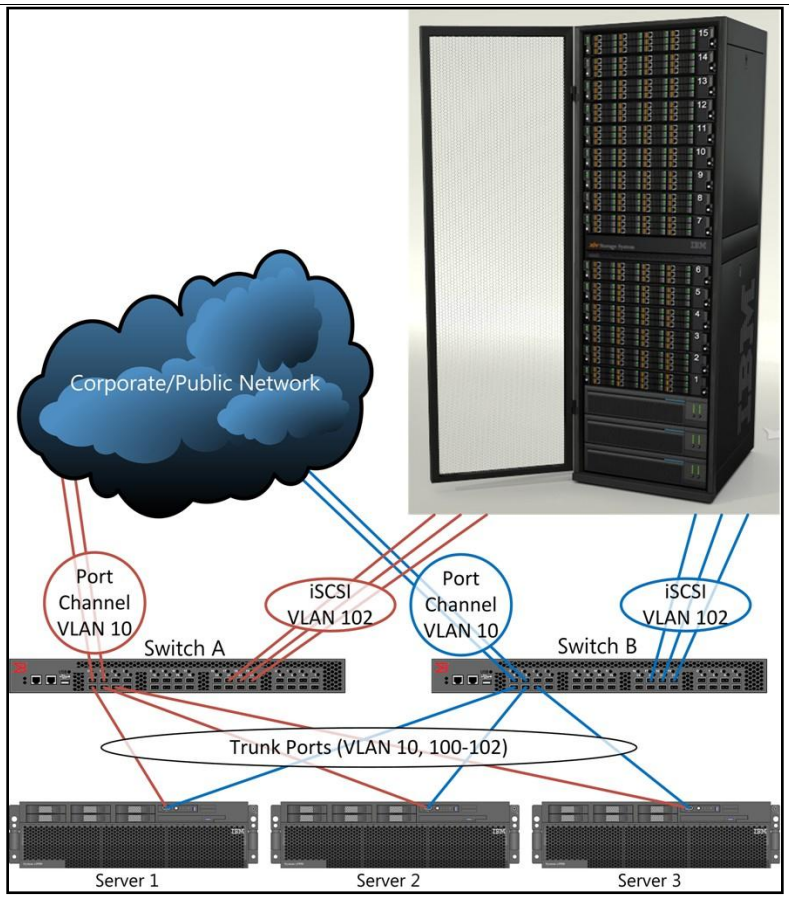

Figure 3 – Hyper-V network diagram

Layout

The layout consists of core infrastructure servers created as virtual machines. It also consists of the VM storage allocation information. Common configuration items such as detailed descriptions for setting up Microsoft Windows Server 2008 R2 and Failover Clustering, or general XIV administration and provisioning are beyond the scope of this paper. Links to vital or useful information will continue to be provided as needed throughout this document.

Note: The IBM XIV Storage System offers both iSCSI and Fibre Channel (FC) host connectivity. The XIV Device Specific Module (DSM) for Multipath Input/Output (MPIO) software is included in the IBM XIV Host Attachment Kit for Windows. In this project, the Hyper-V host FC disks are used for virtual machine VHD's. The VHD's are used to install the operating system. There is also a single FC pass-through disk which is presented to the SCVMM guest. Guest Failover Cluster virtual machines use iSCSI for all of their cluster physical disk resources.

Infrastructure Servers

Infrastructure servers in any hypervisor environment which take advantage of virtualization are commonly available from both host and guest partitions. Hyper-V hosts are configured with pre-existing infrastructure server roles such as DNS and Active Directory. Failover Clustering requires these roles to be available to both host & guest VM's.

The tested reference architecture used a virtualized infrastructure instead of the traditional physical infrastructure. The redundant infrastructure servers were installed on different physical hosts to make the testing much easier and convenient.

Virtual Machines

The following configurations used in the tests represent a diverse set of infrastructure and application virtual machines created on different hosts to promote high availability. These VM choices also emphasize ideal examples for implementing Hyper-V solutions and are further depicted in **Figure 4** below:

Domain Controller 1 - Windows Server 2008 R2 – Active Directory, DNS

Domain Controller 2 - Windows Server 2008 R2 – Active Directory, DNS

File Server 1 (clustered) - Windows Server 2008 R2 – File Services, DFS, and Clustering

File Server 2 (clustered) - Windows Server 2008 R2 – File Services, DFS, and Clustering

SQL Server 1 (clustered) - Windows Server 2008 R2 – Microsoft SQL Server® 2008 SP1

SQL Server 2 (clustered) - Windows Server 2008 R2 – Microsoft SQL Server 2008 SP1

SQL Reporting Services - Windows Server 2008 R2 – Microsoft SQL Server 2008 SP1 Reporting Services

Monitoring - Windows Server 2008 R2 – System Center Operations Manager 2008 R2

Virtualization Management - Windows Server 2008 R2 – System Center Virtual Machine Manager 2008 R2

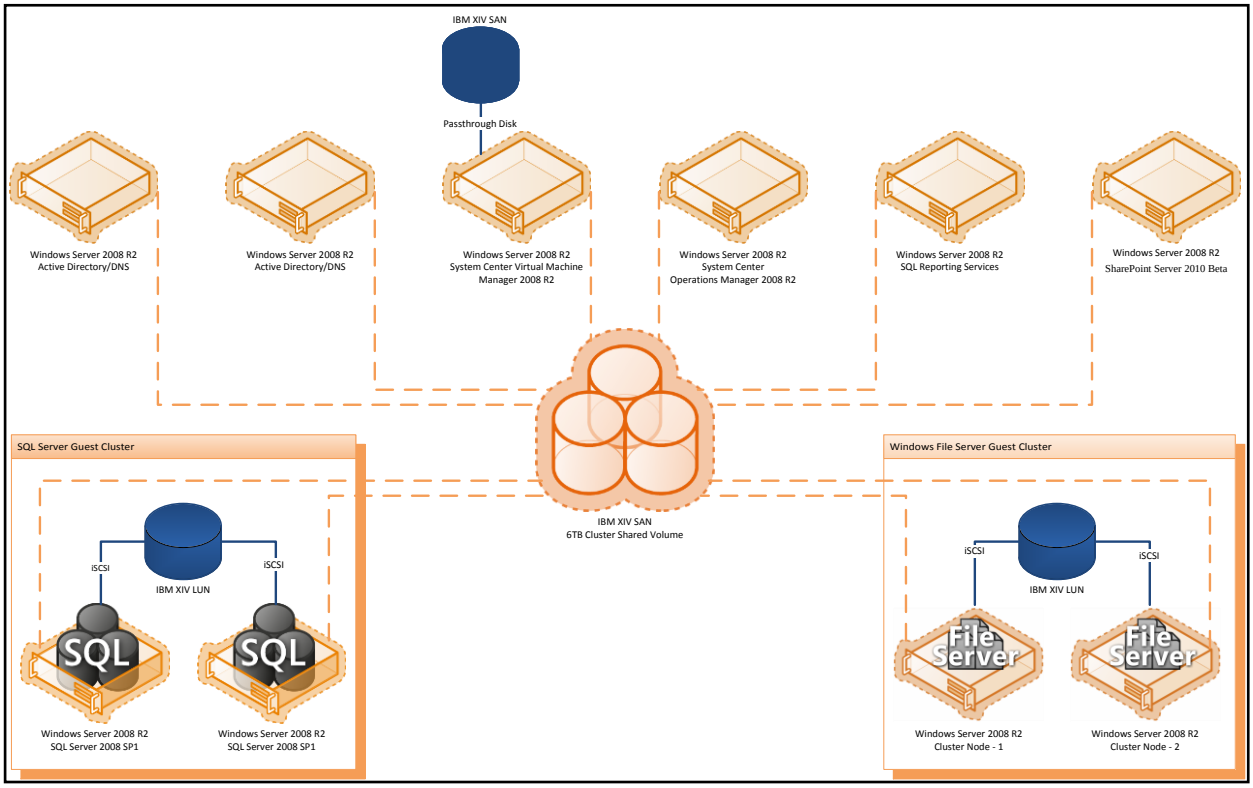

Figure 4: Virtual Machine configuration

Virtual Machine Storage

All of the VM's are configured to boot using a 60GB fixed virtual hard disk (VHD). It resides on a 6TB Cluster Shared Volume (CSV) which is attached to a two node Windows Server 2008 R2 Failover Cluster via fiber channel. The operating system is initially installed from media. Then the Sysprep tool is utilized for additional virtual machines created by copying the virtual hard disk from the initial installation. This virtual hard disk (VHD) is later imported into the System Center Virtual Machine Manager Library.

In addition to the operating system virtual hard disks, the following storage is utilized:

- The guest clustered file server utilizes a 3TB iSCSI data volume and a 17GB iSCSI cluster quorum volume
- The guest SQL Server 2008 cluster utilizes an 800GB iSCSI database volume, a 250GB iSCSI transaction log volume, and a 17GB iSCSI quorum volume
- System Center Virtual Machine Manager utilizes a 12TB pass-through volume connected to the cluster nodes with fiber channel as a primary library store

For further Microsoft Hyper-V detailed storage planning and configuration considerations, see the following TechNet articles:

Planning for Disks and Storage

Configuring Disks and Storage

Installation and Configuration

There are several considerations for performing installations and configurations in a Hyper-V environment. Configure the storage and hosts (physical servers) first to determine the design for the guests (virtual machines).

On the IBM XIV Storage System:

- Provision or verify availability of adequate volumes to be used for Hyper-V VHD's
- Setup iSCSI target Ethernet ports
- Create iSCSI based host objects and map volumes to be used for VHD's

Hyper-V Host Installation

You should complete basic server and storage configuration considerations or prerequisites prior to installing Hyper-V. The bulleted steps below outline the basic ordered installation process.

- Install the latest IBM server firmware
- Install Windows Server 2008 R2 using the local server disks
- Configure for DHCP or static IP use
- Rename the servers using a practical nomenclature
- Join the servers to the domain
- Complete the Windows updates as required
- Install the latest IBM Server drivers and agents
- Assign the servers static IP addresses for each interface
- Install latest HBA firmware, drivers, and multipath software
- Install the Failover Clustering feature
- Create the new cluster following standard best practices
- Add available storage as Cluster Shared Volumes
- Install the Hyper-V role
- Create Virtual Machines on CSV volumes and load the OS
- Make the Virtual Machines highly available

Note: The latest IBM firmware, drivers and agents are available at Fix Central

One should understand the host installation at this stage. However, the following additional references provide further clarification.

Using Hyper-V and Failover Clustering

Hyper-V: Using Live Migration and Cluster Shared Volumes in Windows Server 2008 R2

Networking Configuration

The following are the test network configuration steps.

1. Create and Configure VLAN's on all Nodes
2. Configure IP addresses, DNS names on via the cluster and Live Migration Interfaces

The following Microsoft TechNet articles cover networking best practices in greater depth.

Checklist: Setting up Failover Cluster Nodes, Networks and Storage

Hyper-V: Live Migration Network Configuration Guide

Storage and Host Cluster Configuration

The IBM XIV Storage System provides the storage foundation for the Hyper-V cluster hosts and guest VM's The host and guest failover cluster nodes use XIV storage for shared disks or physical disk resources. The guest virtual machines use it for their system disks. Everything else relies on and builds upon that foundation. To gain access to the XIV storage, perform these steps:

1. Create and configure zones on the fiber channel fabric such that each host has connectivity to the IBM XIV storage system
2. Install the IBM XIV Host Attachment Kit for Windows on all participating hosts. As part of the installation, the IBM XIV Host Attachment Kit will walk through connecting and configuring both the server and storage for optimal connectivity to the array.
3. As part of step 1, the IBM XIV Host Attachment Kit will create a host object in the XIV Storage Management UI on the Hosts and Clusters view. Now create a cluster in the XIV Storage Management UI from the Hosts and Clusters view by right clicking and selecting **Add cluster**. Then add all of the participating hosts to that cluster.
4. Run the XIV Host Attachment Wizard – Follow the Prompts; See **Figure 5** below:

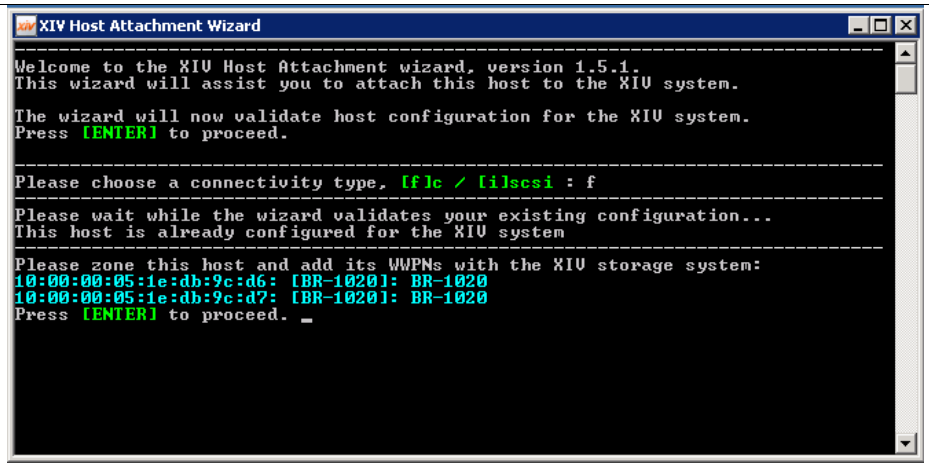

Figure 5: XIV Host Attachment Wizard

5. Add the ports shown in the XIV Host Attachment Wizard to the Host Created in Step 2. See **Figures 6** and **7** below:

Figure 6: XIV Storage Management Graphical User Interface

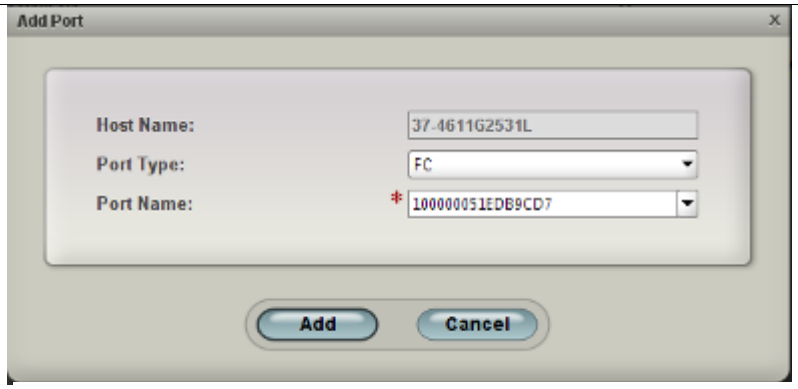

Figure 7: Adding a Port in the XIVGUI

6. Install the IBM MPIO Management Console to verify storage path connectivity. **Figures 8** and **9** below reveal screenshots of the successful MPIO Management Console installation:

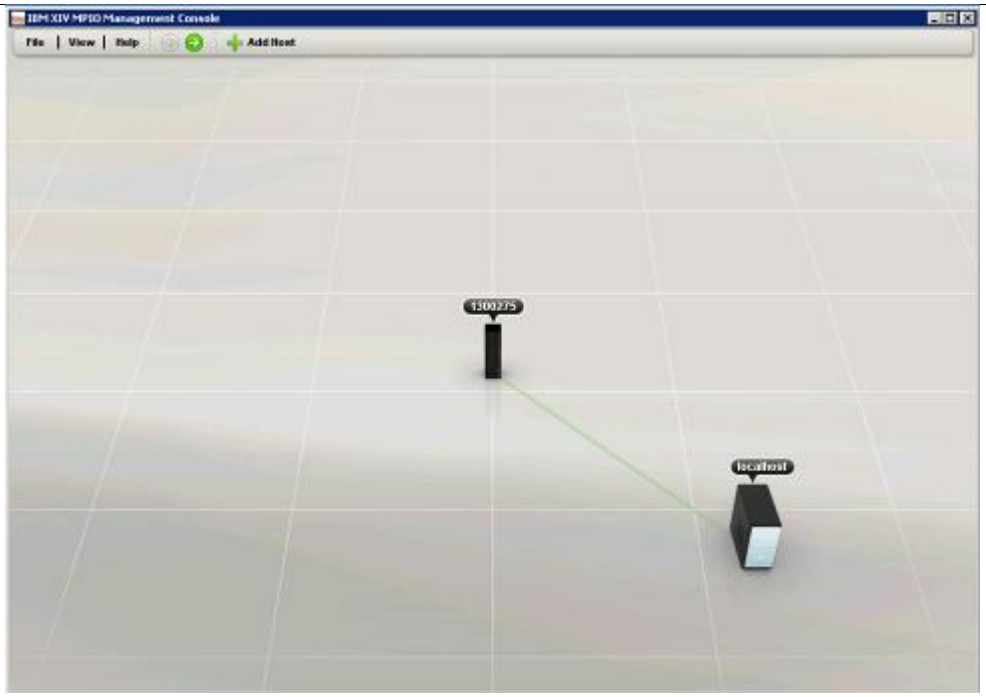

Figure 8: IBM MPIO Management Console

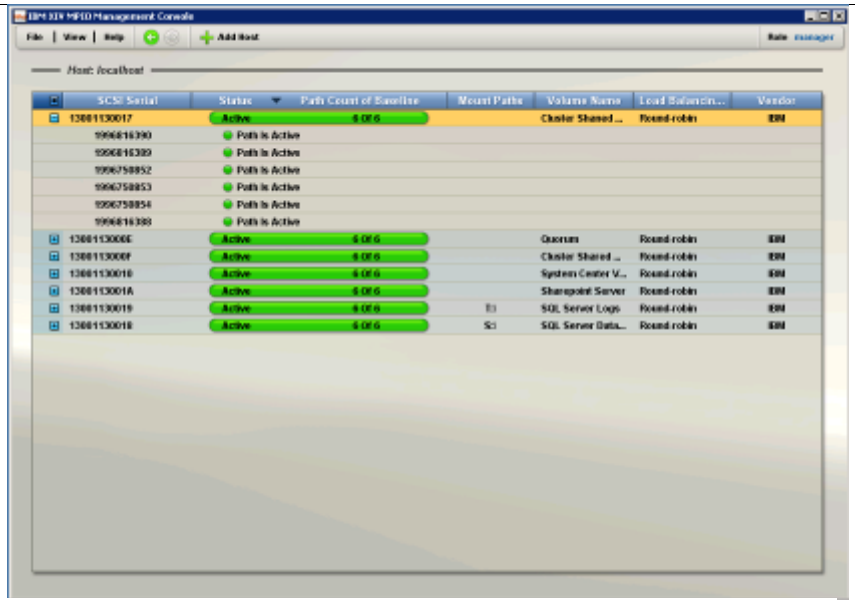

Figure 9: IBM MPIO Management Console

7. If required, create and map a quorum volume to the cluster using the XIV Storage Management Graphical User Interface (XIVGUI).
8. Create a Windows Failover Cluster verifying the Cluster Validation Wizard passes successfully.
9. Enable Cluster Shared Volumes in the Microsoft Windows Failover Cluster Management Interface or with the Failover Clustering PowerShell Interfaces. For more information on requirements and recommendations regarding CSVs, reference the links below:

   Requirements Using Cluster Shared Volumes in a Failover Cluster in Windows 2008 R2
   Recommendations for Cluster Shared Volumes in a Failover Cluster in Windows 2008 R2

10. Using the IBM XIV Management Interface, create and present a volume to the cluster to be used as a Cluster Shared Volume.
11. From one of the online cluster nodes, initialize and format the volume (without assigning a drive letter or path).
12. Add the volume to the available storage group in the Microsoft Windows Failover Cluster Management or Failover Clustering Power Shell Interfaces.
13. Add the volume as a cluster shared volume in the Microsoft Windows Failover Cluster Management or Failover Clustering Power Shell Interfaces.

    Note: It is recommended to re-run the storage portion of the cluster validation when adding new storage.

VM Guest Cluster Configuration

Windows Server Clustering offers robust application mobility. Utilizing iSCSI connected storage, this functionality extends into Microsoft Hyper-V's virtual machines. Keep in mind if the hosts are clustered, all virtual machines running on a host will failover to the other host(s) if the physical server has an unplanned or unexpected failure. Using failover cluster manager with clustered guests spread across physical hosts, applications can simply be moved between virtualized servers with minimal or no down time. This allows customers more flexibility when it comes to application server management. To help determine which clustering model to pursue, download the very informative Microsoft podcast which details the differences between guest vs. host clustering.

One thing to keep in mind is there are a few unique prerequisites when it comes to guest vs. host clustering. The following outlines the unique guest prerequisites:

Prerequisites for guest clustering

- Install Windows Server 2008 R2 using the FC VHD's
- Install Hyper-V™ integration components
- Ensure all software is updated with current service packs and/or software updates
- Join the domain
- Create a minimum of Four Guest Network Connections (one for cluster public, one for cluster private, two or more for iSCSI storage)
- Add the Failover Clustering Feature
- Install and configure IBM XIV Host Attachment Kit for Windows on all participating hosts.
- Install the IBM XIV Management Tools
- Install IBM XIV MPIO Management Console (optional)

Note: IBM XIV Storage Software can be located at the following link:

http://www.IBM.com/systems/support/storage/XIV

Guest Configuration

1. Install the IBM XIV Host Attachment Kit for Windows on all participating virtual machines. As part of the installation, the IBM XIV Host Attachment Kit (HAK) will walk through connecting and configuring both the server and storage for optimal connectivity to the array. Additionally, the IBM XIV HAK will create a host object in the XIV Storage Management UI in the Hosts and Clusters view.
2. Now create a cluster in the XIV Storage Management UI. From the Hosts and Clusters view, right click and select **Add cluster.** Then add all of the participating hosts to that cluster.
3. Install the IBM MPIO Management Console to verify storage path connectivity.
4. If required, create and map a quorum volume to the cluster using the XIV Management Interface.
5. Create a Windows Failover Cluster verifying the Cluster Validation Wizard passes successfully.

Test Objectives

The goal of this testing project was to determine how well do IBM servers, XIV storage and Brocade hardware perform together with Hyper-V? Testing focused on the strongest set of Hyper-V features and benefits. All of the following functions were tested:

- Configuring Cluster Shared Volumes (CSV)
- Planned and Unplanned VM cluster failovers
- VM Quick Migrations between cluster nodes
- VM Live Migrations between cluster nodes
- Hot add storage (VHDs) on running VM's
- Quick Storage Migration (requires Microsoft SCVMM 2008 R2)

Enabling Cluster Shared Volumes

The CSV option becomes available in the Failover Cluster Manager after the Hyper-V role is installed. To fully appreciate CSVs, revisit how typical clusters operate. In a traditional cluster, only one node at a time can access a SAN based physical disk resource. Multiple cluster nodes cannot read from or write to the disk simultaneously to prevent data corruption.

However, with CSVs, multiple nodes of a cluster can access the same physical disk resource, which can serve to facilitate the transparent move of online Hyper-V VM's between host nodes without any perceived downtime via Live Migration. CSV's make it easy to set up Live Migration, because you can simply give each VM access to the shared storage. Performing live migrations with CSV is fast because both VMs can access the same stored files and only the memory needs to be moved between the Hyper-V hosts. Performing live migrations without CSV takes a bit longer because the storage must be moved between the nodes. It should be noted however that CSV's are *not* a requirement for Live Migration.

Here are the steps to enable CSV's in a Failover Cluster:

1. In the Failover Cluster Manager, select the cluster to manage.
2. In the center pane, under **Configure**, select **Enable Cluster Shared Volumes**.
3. Select the checkbox in the corresponding dialog box.

   Note: By enabling CSVs, all nodes in the cluster will be able to use them.

Once any disks are added as CSV's, they appear on each node's system drive. They are located under the \ClusterStorage\ folder (i.e. C:\ClusterStorage\Volume1, etc…). When one creates new Virtual Machines in the Hyper-V Manager, this path must be specified for the location of the VHD files to enable Live Migration.

Virtual Machine Cluster Failover

Failover Cluster Manager can move virtual machines from one node to another using several different methods. Both planned and unplanned failovers can be periodically tested to confirm the health and viability of the cluster. Planned failovers accommodate a graceful move of the virtual machine to another node. This type of failover is commonly used for maintenance windows when not wanting to update an active cluster host or physical server. If desiring to simulate a server failure, opt to use the unplanned failover method. These are basic cluster functions which the intended audience should know; for additional information, reference the TechNet article on testing failovers.

Virtual Machine Quick Migration

Quick Migration a feature released with Windows Server 2008, allows VM's to be moved or migrated to other physical nodes in the cluster. This is accomplished by stopping the VM, saving the state of the VM followed by the memory to disk. Ownership is then transferred to the other node, which mounts the VM, and restarts it. Quick Migration typically occurs faster than Live Migration; however a loss of availability is incurred. The failover time is directly related to the amount of memory assigned to the virtual machine, along with the connection speed between the server and shared storage. Quick Migration can also be used to automatically restart a virtual machine on another cluster node in the event of a failure.

You can perform a Quick Migration with the following steps:

1. Using the Failover Cluster Manager, select the virtual machine to migrate.
2. Click on **Quick Migrate virtual machine(s) to another node**.
3. The VM will migrate and come online on the other node.

Virtual Machine Live Migration

With the release of Windows Server 2008 R2, Microsoft introduced the new Live Migration high availability feature. Its great appeal is derived from the ability to perform VM migration with zero downtime. As previously mentioned, live migration offers the ability to migrate data seamlessly between server data pools in a virtualized environment. Live Migration is accomplished by transferring memory pages from the source VM to the target via the Live Migration network, then changing ownership. This process can take longer than Quick Migration; however the VM is available during the entire transition. There is a 'brownout' period, but this occurs within the TCP/IP timeout window and is all but imperceptible

You can perform a Live Migration via the following steps:

1. Within Failover Cluster Manager, highlight the virtual machine to migrate.
2. Click on **Live Migrate virtual machine to another node**.
3. The virtual machine will remain online, and migrate to the other node.

Hyper-V Hot Add Storage

Another new high availability feature in Windows Server 2008 R2 is the ability to dynamically add or remove disks with an online virtual machine. This feature includes adding storage as passthrough disks, VHD's on dedicated LUN's, and CSV's, as long as one specifies the CSV location for the VHD. The LUNs must be zoned and available to all nodes of the cluster.

The Hot Add Storage procedure can be performed via the following steps:

1. Within the Failover Cluster Manager, open the Settings panel for the virtual machine.
2. Click on the **SCSI** controller (not the IDE controller which does not support this feature).
3. Select **Hard Drive**, and click **Add**.

Responding To Growth

The ability to hot add storage is just one facet when it comes to the subject of growth. Ideally, proper planning measures are taken prior to any deployment. However, there are moments of growth which require adaptability and change no matter how much planning.

The flexibility and ease of management that Hyper-V and XIV affords makes administration, particularly in these cases, incredibly simple. As more servers are virtualized on Hyper-V with IBM XIV storage, growing both the cluster as well as the storage may be required. With Microsoft System Center Virtual Machine Manger and IBM XIV storage, this is an elegant process. One is no longer exposed to many of the common problems such as application downtime incurred in order to back up and restore data to a new, larger LUN caused by unplanned growth.

In **Figure 10** below, an example of unforeseen growth surfaces in the form of a dialog box warning within. SCVMM indicating that the cluster is “Over-committed” requiring the addition of Hyper-V servers. This suggests it’s time to respond to growth. The appropriate response is simply to add another Hyper-V cluster node.

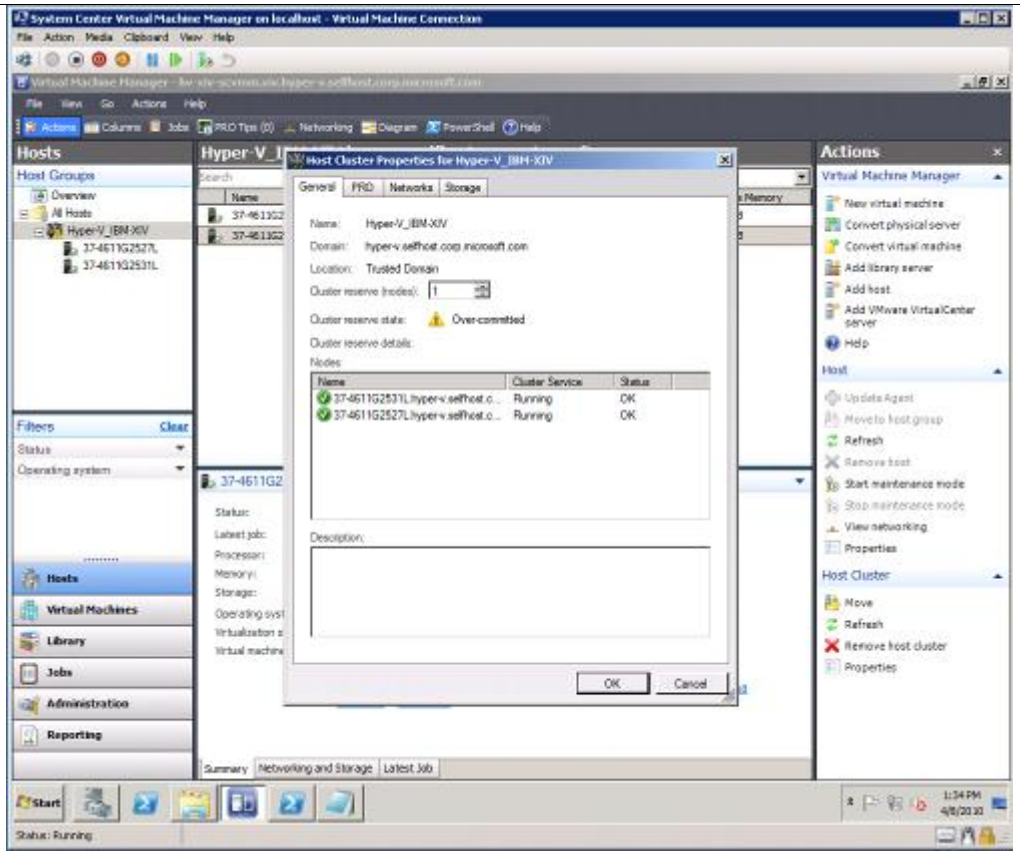

Figure 10: Microsoft SCVMM 2008 Host Cluster Properties

Figure 11 below reveals how easy it is to add another host to the cluster from an IBM XIVGUI perspective.

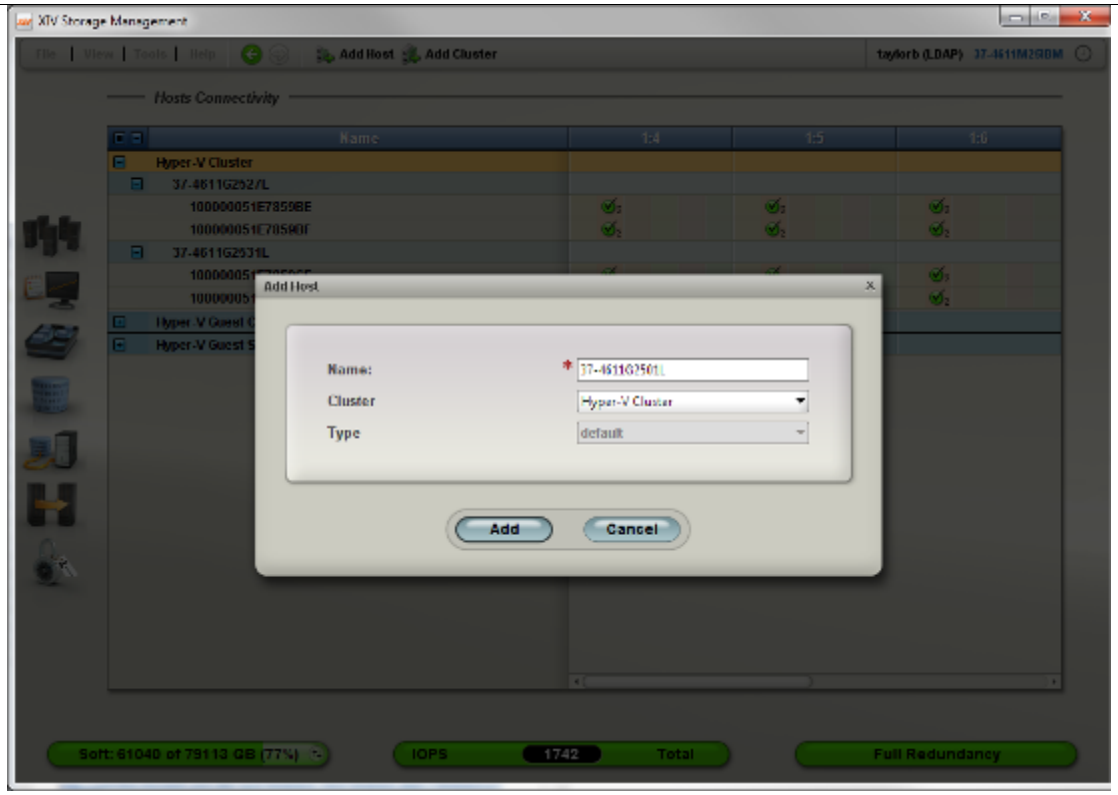

Figure 11: XIV Storage Management Add Host Dialog Box

Using the IBM XIV Storage Management Interface to ensure the newly added node is properly mapped is as simple as adding it to the failover cluster.

Below are the steps for accomplishing both:

1. Install and configure Windows.
2. Zone the storage.
3. Run the XIV host attachment wizard.
4. Use Add Host to the cluster in the XIV Storage Management UI.
5. Prior to adding a server to the cluster in Windows, don't forget to run the Validate a Configuration Wizard.

   Note: Step 5 is optional but highly recommended.

6. Add the server using the Windows Failover Cluster Manager or the Windows PowerShell.

Adding a new host to the failover cluster as referenced in step 6 above is easily accomplished via SCVMM. The virtual machine can now be live migrated or created on the new node.

Note: **Figure 12** below reveals a newly added node as it appears within the Microsoft SCVMM.

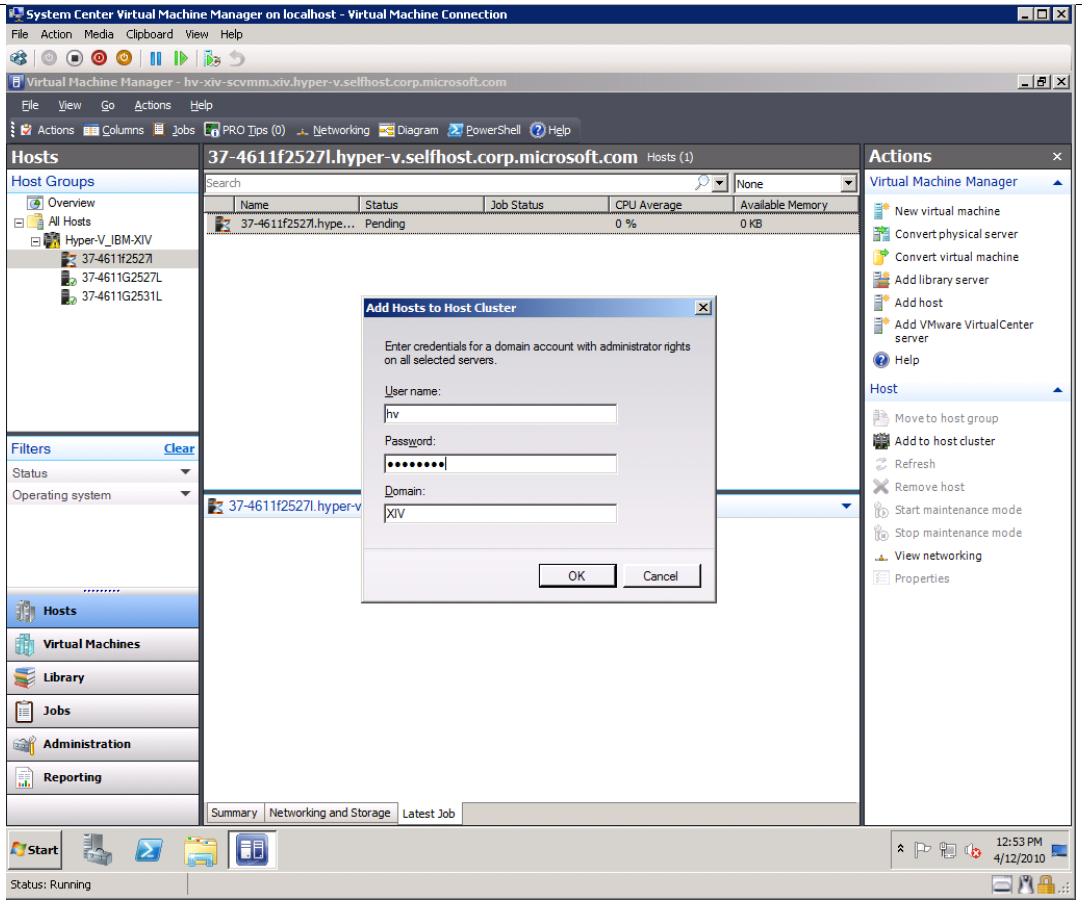

Figure 12: Adding a host to an existing cluster via System Center Virtual Machine Manager

Monitoring VMs

VM monitoring is multi-dimensional. Most of the same rules that apply to storage, networking, and physical servers apply to Hyper-V hosts and guests. Most administrators think of performance monitoring when asked about this topic, but server health monitoring is equally important.

As a best practice, one should begin by capturing baseline statistics during normal activity. That way if encountering performance issues in the future, the baseline statistics can be referenced to help troubleshoot, isolate and resolve any problems. Take note of the fact that Windows performance monitor counters for hosts vs. guests are unique. The following Microsoft article provides information on disk I/O, memory, networking, and processor performance considerations:

Measuring Performance on Hyper-V

With the IBM XIV Storage System, one can use the XIVGUI or XIVCLI to capture I/O performance statistics. The following XIVCLI commands reflect performance statistics and volume or storage pool history:

- Performance Statistics:

```
statistics_get [host=H |
host_iscsi_name=initiatorName | host_fc_port=WWPN |
target=RemoteTarget | remote_fc_port=WWPN |
remote_ipaddress=IPAdress | vol=VolName |
ipinterface=IPInterfaceName | local_fc_port=ComponentId ]
< start=TimeStamp | end=TimeStamp >
[ module=ComponentId ]
count=N
interval=IntervalSize
resolution_unit=<minute | hour | day | week | month>
```

- Volume or Storage Pool History:

```
usage_get < vol=VolName | pool=PoolName > [ start=StartTime | start_in_seconds=StartTime
]
[ end=EndTime ] [ max=MaxEntries ]
```

Additionally, one can measure Brocade switch performance via telnet for such things as individual por statistics or via a variety of Brocade GUI monitoring solutions.

Lastly, the management features of SCVMM are designed for centralized VM monitoring & management as well. SCVMM also has the flexibility of expanding its management and monitoring functions by installing additional management packs such as the Brocade HBA Management Pack.

The Brocade HBA Management Pack for SCVMM works with the Brocade HBA to monitor the health and performance of data center SAN infrastructure. It works with the Performance and Resource Optimization (PRO) framework within SCVMM to provide intelligent event-based recommendations that allow IT administrators to dynamically optimize the performance of their virtualized workloads.

These are but a few of the available options to customers seeking monitoring solutions for their virtualization environment. After reading this paper and the various provided links, one should have a solid monitoring foundation as well as a solid understanding of how the combination of Hyper-V, System x servers and the XIV storage system can provide a flexible, highly-available and easy-to-manage virtualization platform.

Summary

Implementing virtualization solutions with Microsoft Hyper-V, implemented on IBM System x servers and the System Storage XIV storage array results in a technology platform designed to address the ever increasing demands of businesses. The value to the IT organization ultimately boils down to reducing risk, cost, and IT platform complexity thereby providing high quality of service to their internal and external customers. This joint IBM, Microsoft and Brocade solution implementation guide provides a highly available, flexible, centrally-managed virtualization solution. IT customers know that in order to remain competitive, they must expedite their adoption of virtualization. The benefits of high scalability, resiliency, advanced technologies such as thin provisioning, lower TCO, and ease of manageability are the primary drivers. These benefits can also help support the green computing through footprint reduction and reduced power consumption.

---

[i] IBM Systems & Technology Group, Kirkland, WA ebjohn@us.ibm.com
[ii] Microsoft Corporation, Redmond, WA taylor.brown@microsoft.com
[iii] IBM Systems & Technology Group,, Raleigh, NC billwats@us.ibm.com
[iv] IBM Systems & Technology Group, Raleigh, NC dhartma@us.ibm.com