



# Virtualized High Availability and Disaster Recovery Solutions

Microsoft Hyper-V on the IBM System Storage™ DS5000, IBM System x™ Servers, & IBM BladeCenter

## Highlights

- *Windows Server 2008 and Hyper-V with MSCS failover cluster support provides efficient high availability*
- *Quick migration of Hyper-V guest machines between IBM HS21XM Blades with automatic fail-over*
- *Optimized disaster recovery with the flexibility of virtualization and efficiency of the IBM BladeCenter, System x 3850 M2 and the DS5000 modular storage system*
- *Reliable replication of data between sites with the IBM DS5000 and Remote Volume Mirroring*
- *A robust, yet easy to manage & configure HA/DR scenario*



IBM DS5000

Today's virtualization solutions help customers reduce datacenter complexity and management costs, providing dynamic and flexible configurations across multiple resources. Companies of all sizes are interested in energy and space efficiency, centralized management, and flexible disaster recovery options to reduce their total cost of ownership. High availability and disaster recovery are two key areas virtualization can offer significant flexibility and cost savings. By leveraging IBM servers and storage with Microsoft Hyper-V and Windows Server 2008, organizations are equipped to streamline their IT infrastructure for maximum efficiency.

**The Solution – Implementing Hyper-V for high availability and disaster-recovery scenarios**, using the IBM System Storage DS5000, System x3850 M2, and BladeCenter servers with Hyper-V. This reference architecture outlines one possible solution for intra-site high availability and inter-site disaster recovery scenarios. The configuration includes simulation of 2 geographically dispersed sites, with local site HA provided by 2 MSCS failover clustered HS21XM Blade servers hosting multiple Hyper-V virtual machines. The DR simulation consists of Remote Volume Mirror LUN replication between the two sites, with the capability of bringing the Hyper-V virtual machines online at the second site if the main datacenter goes down or is otherwise unavailable.

Remote Volume Mirroring is a premium feature of the IBM System Storage DS5000. Data replication and control is handled by the array controllers over replication links with little or no impact on host applications. Advanced features such as ordered-writes and consistency group support help ensure that complex application LUN layouts and database integrity are preserved. Figure 1 below shows the FC SAN logical design:

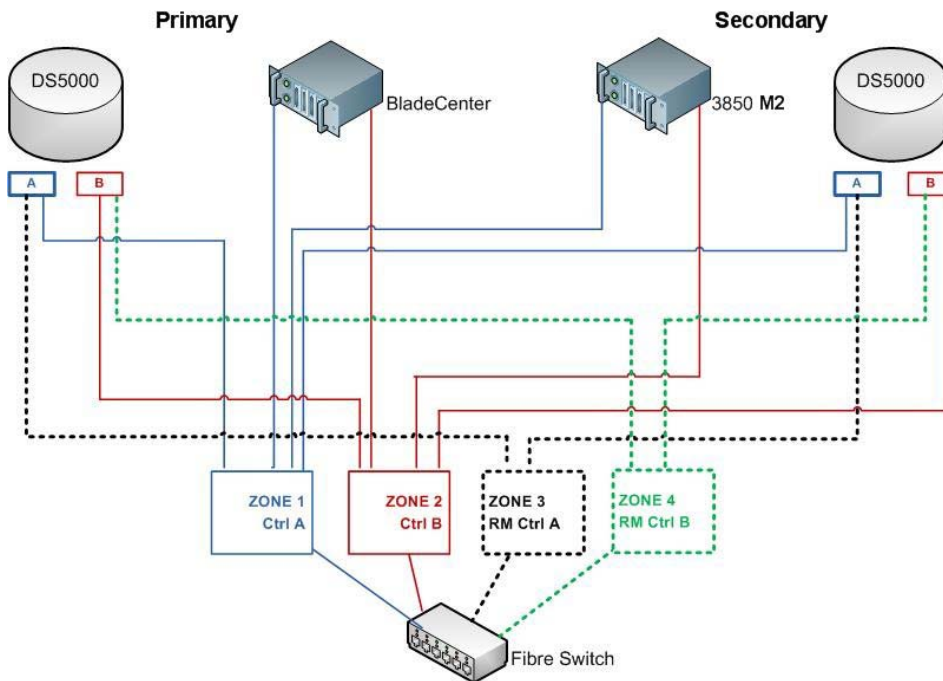


Figure 1) FC solution design for RVM.

Hyper-V virtual servers run as cluster resources, and automatically fail over if one host node is unavailable. These virtual machines can also be quickly migrated between the cluster nodes, pausing running applications and resuming automatically once online. In a DR event, replicated LUNS and guest are manually activated and brought online. Figure 2 below highlights the logical view of the solution:

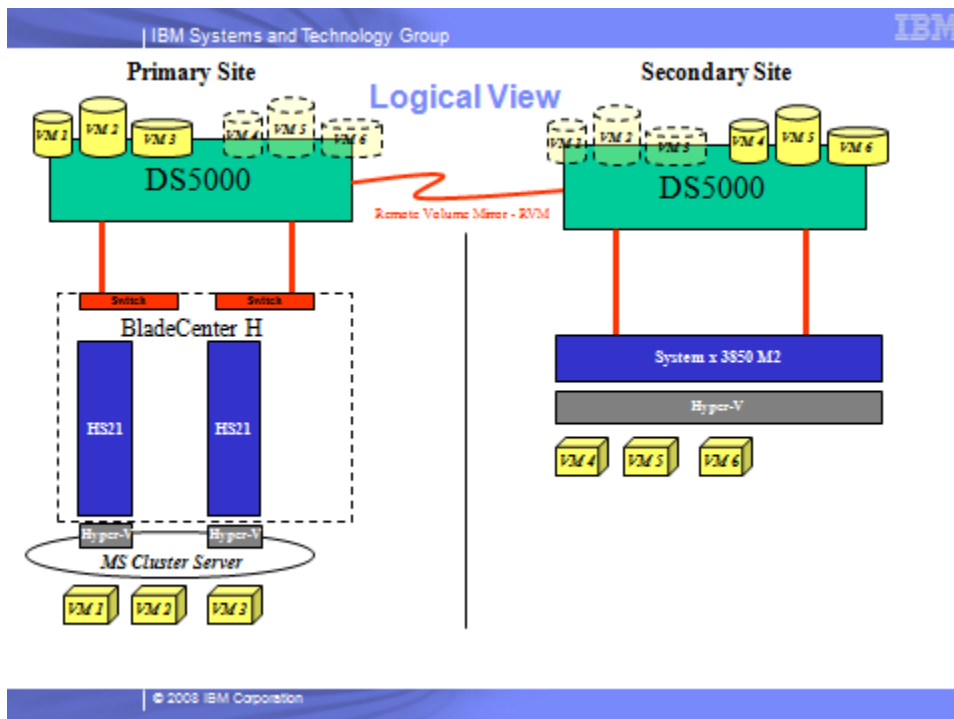


Figure 2) Hyper-V logical HA/DR solution view



Each of the two blades at the primary site host multiple Hyper-v guest machines (VM1, VM2, and VM3). The secondary site also runs Hyper-V guests (VM4, VM5, and VM6). In the event of a failure of one of the computer nodes at the primary site, the guests are automatically migrated to the surviving node. The primary site remains online. Guest LUNs on the primary site are replicated via RVM to the secondary site. Figure 3 below shows a localized server failure:

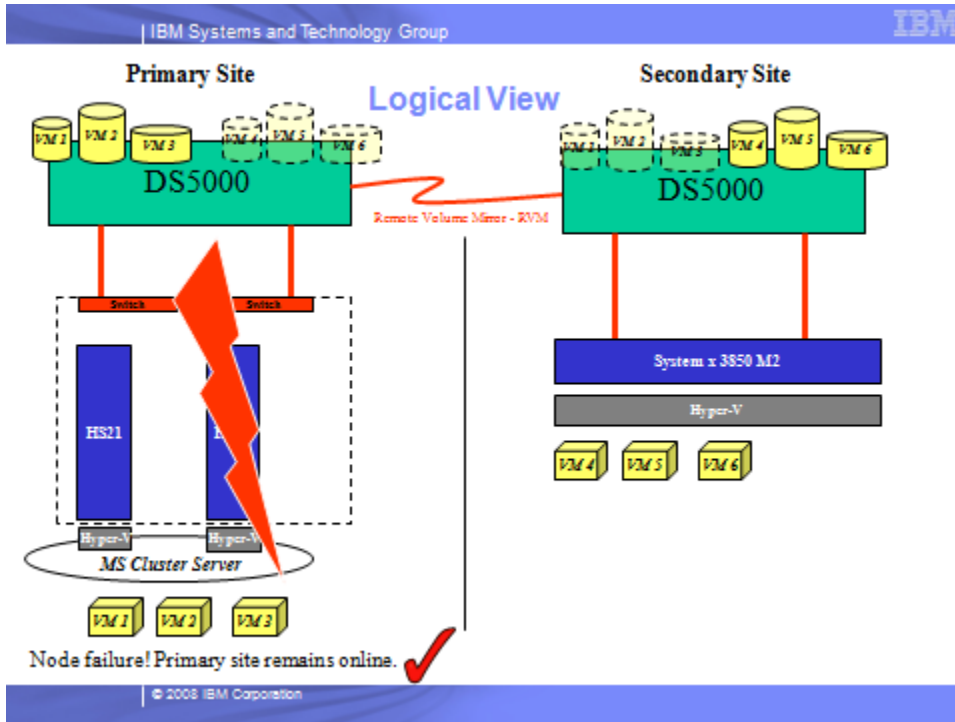


Figure 3) Localized server failure.



In the event of a local site failure, if the entire BladeCenter or DS5000 goes offline within the primary site, Hyper-V guest machines are restarted on the System x3850 M2 at the secondary site. Replicated data will be in a crash-consistent state, and application recovery methods (i.e. log replay) are activated. In a manual fail over (i.e. for maintenance), this would not need to occur. Once the LUN's are enabled read/write & mapped to the relocated Hyper-V guests, the applications are restarted and service resumes. Figure 4 below shows a complete site failure example.

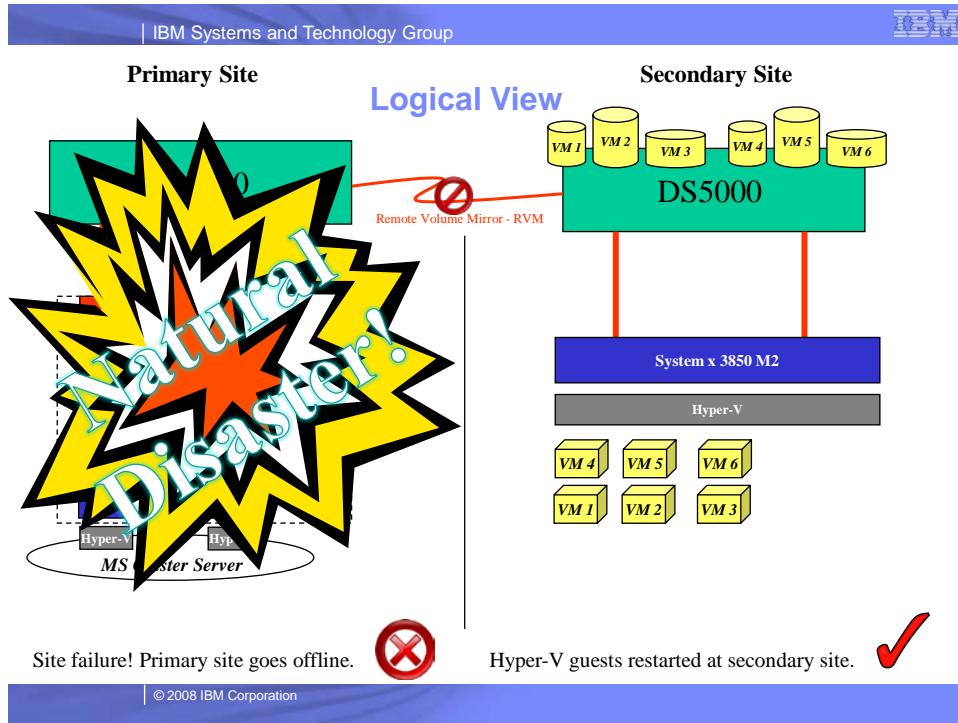


Figure 4) Primary site failure.



In the event of an outage at the secondary site, as shown in Figure 5 below, the reciprocal process to the primary site failure would occur, and Hyper-V guests would be manually restarted at the primary site.

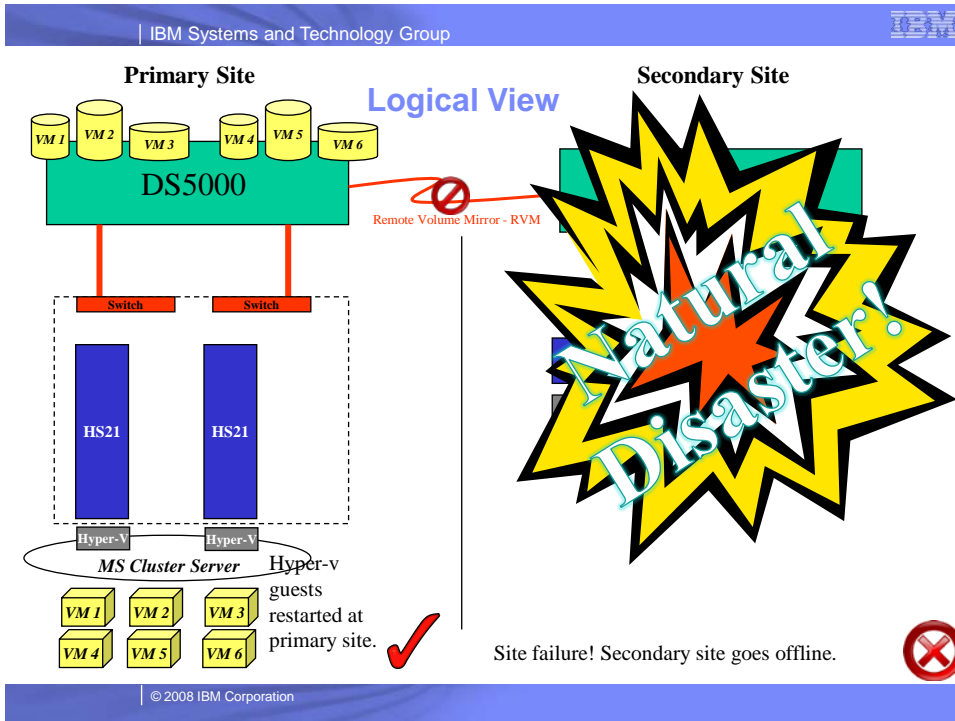


Figure 5) Secondary site failure.



Since Hyper-V uses the familiar MMC 3.0 framework, Windows administrators can manage the virtualized, clustered environment in a familiar environment. Furthermore, the DS5000 Storage Manager GUI is easy to configure & manage even in complex HA/DR scenarios such as RVM. Figure 6 below shows a clustered Hyper-V MMC:

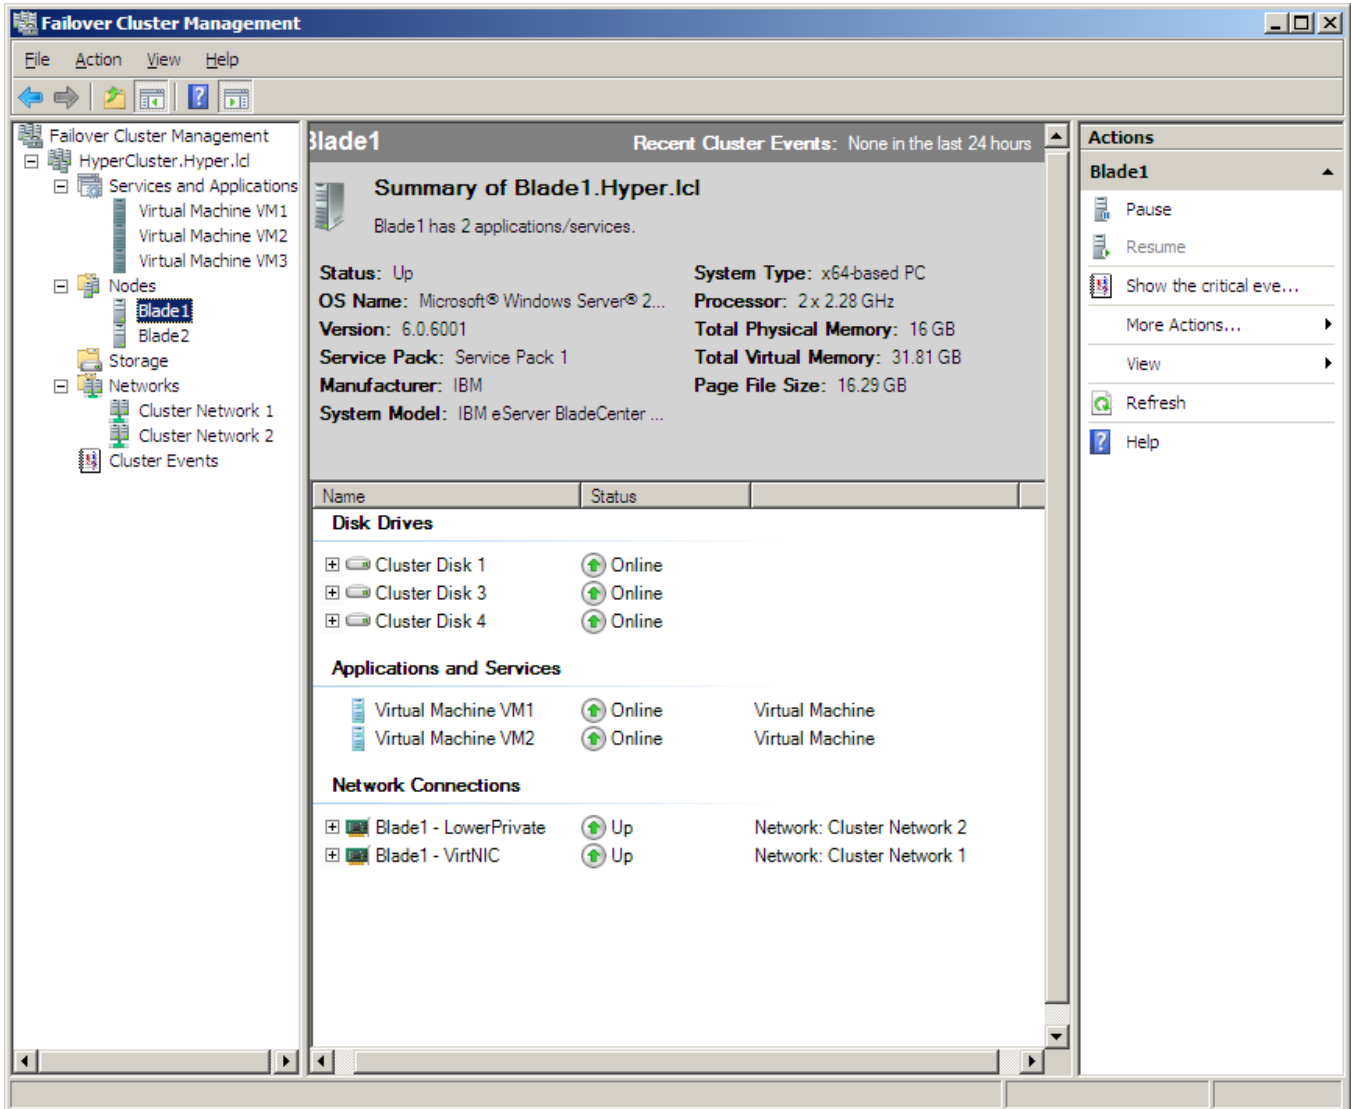


Figure 6) FCM MMC panel.

Microsoft's Hyper-V, included with Windows Server 2008, provides software infrastructure and management tools that you can use to create and manage a virtualized server computing environment for consolidated and scalable data centers.



### Key features of Hyper-V include:

- 64-bit native hypervisor-based virtualization
- Ability to run 32-bit and 64-bit virtual machines concurrently
- Large virtual machine memory support and Virtual LAN support
- Virtual machine snapshots, to capture the state of a running virtual machine so you can revert the virtual machine to a previous state quickly & easily.
- Runs on all roles including Server Core, of Windows Server 2008
- Hyper-V leverages Microsoft Cluster Services for failover cluster support
- Microsoft Management Console (MMC) 3.0 interface

**The NEW IBM System Storage DS5000** sets new standards for performance, scalability, reliability, availability, and flexibility for midrange storage systems. As IBM's most powerful midrange storage system, the DS5000 is the ideal platform for a virtualized environment that can keep pace with your business growth. Organizations can buy only the capacity needed initially, and can then dynamically upgrade and reconfigure additional capacity & features later to meet changing business requirements, all without any system downtime.

The DS5000 delivers class-leading performance and is equally adept at supporting transactional-applications, such as databases and OLTP, throughput-intensive applications, such as HPC and rich media, and concurrent workloads, well-suited for consolidation and virtualization. With its relentless performance and architected to provide the highest reliability and availability, the DS5000 storage system is comfortable supporting the most demanding service level agreements (SLAs). And when requirements change, the DS5000 can easily be reconfigured "on-the-fly" to add or replace host interfaces, increase performance, grow capacity, or add cache – ensuring it keeps pace with your growing company.

### DS5000 key features:

- Flexible host interface options are 8 Gb/s Fibre Channel and 10 Gb/s iSCSI ready
- Sixteen 4 Gb/s Fibre Channel drive interfaces for support up to 256 drives in initial release, with future support for 448 FC/SATA drives, using EXP5000/EXP810 drive expansion units.
- Up to 16 GB of dedicated data cache (8 GB per controller) in initial release, with future support for 32 GB. Dedicated cache mirroring channels and persistent cache backup in the event of a power outage
- Support for RAID 6, 5, 3, 10, 1, 0
- Two performance levels (base and high) with ability to field-upgrade
  - Base model is DS5100
  - High model is DS5300
- Remote Volume Mirroring and FlashCopy premium features for Volume Shadow Copy (VSS) supported backups and flexible DR scenarios
- Break-through performance levels over 5X greater than the DS4800





### System x3850 M2 and x3950 M2 key features:

- True 2-to-16-socket scalability up to 64 cores
- Revolutionary Intel Xeon dual-core and quad-core MP 7300 Series processors
- Up to 1TB of registered DIMM memory for better workload density and up to 20–30% less power consumption than competitors' fully buffered DIMM technology\*
- IBM Memory ProteXion™ with redundant bit-steering offers twice the memory resilience of the competition
- 4th generation snoop filter 4 times larger than the competition's best
- IBM Predictive Failure Analysis®, not just on hard drives and memory but, unlike competitors, also on processors, power supplies, fans, and voltage regulator modules
- 40% lower memory latency than the nearest competition
- More flexible memory configurations than competitors, at significantly lower costs



The IBM System x3850 M2 and x3950 M2 servers provide an uncomplicated, cost-effective and highly flexible solution. With the ability to scale while maintaining balanced performance between processors, memory and I/O, these servers can easily accommodate business expansion and the resulting need for additional application space.

### BladeCenter key features:

The IBM BladeCenter H delivers high performance, extreme reliability, and ultimate flexibility to even the most demanding IT environments. In 9 U of rack space, the BladeCenter H chassis can contain up to 14 blade servers, 10 switch modules, and four power supplies to provide the necessary I/O network switching, power, cooling, and control panel information to support the individual servers.

The chassis supports up to four traditional fabrics using networking switches, storage switches, or pass-through devices. The chassis also supports up to four high-speed fabrics for support of protocols like 4X InfiniBand or 10 Gigabit Ethernet. The built-in media tray includes light path diagnostics, two front USB inputs, and an optical drive.

- Dense, space-saving 9U chassis
- Up to 14 blades in a chassis
- IBM Cool Blue® technology
- Energy-efficient design
- IBM Open Fabric
- Powerful solutions management
- High-availability midplane
- Hot-swappable, redundant management, switch, power supply and blower modules
- New high-speed switch module support
- New high-speed bridge module bays
- Advanced server management
- IBM Remote Deployment Manager
- Light path diagnostics self-diagnosis panel
- Predictive Failure Analysis
- 9.5 mm UltraSlim DVD
- 3-year customer replaceable unit and onsite limited warranty







**Emulex LightPulse® IBM Qualified 43W6859 (LP1105-BCX) and 42D0494 (LPe12002) Fibre Channel host bus adapters (HBAs)** provide streamlined installation and management, outstanding scalability and industry-leading virtualization support well-suited for small-to-large enterprises and Microsoft Windows Server 2008 and Hyper-V storage area network (SAN) environments. With powerful management tools and broad System x and BladeCenter support, the LightPulse family of IBM-branded 4Gb/s and 8Gb/s HBAs (IBM Server Proven Validation) delivers high performance for a broad range of applications and environments.

**Emulex HBA key features:**

- Exceptional performance and full-duplex data throughput
  - Comprehensive virtualization capabilities with support for N-Port ID Virtualization (NPIV)
  - Simplified installation and configuration using AutoPilot Installer®
- Administration via HBAnyware® integrated with IBM Director

**Conclusion – Putting it all together** - Microsoft's Hyper-V ushers in a new era of application virtualization affordably to the masses, offering new opportunities for increased resource utilization, ease of management, and improved ROI. IBM has worked closely with Microsoft to ensure our products are optimized for Hyper-V deployments. Together with the new System Storage DS5000, the ability for companies of all sizes to implement a highly-available and disaster-tolerant computing environment has never been easier.

A full solution whitepaper will be available on the IBM ISV Solution website by Q408:

<http://www-03.ibm.com/systems/storage/solutions/isv/index.html#microsoft>



Copyright © 2008 by International Business Machines Corporation.

This document could include technical inaccuracies or typographical errors. IBM may make changes, improvements or alterations to the products, programs and services described in this document, including termination of such products, programs and services, at any time and without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document is current as of the initial date of publication only, and IBM shall have no responsibility to update such information.

Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing and such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS-IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR INFRINGEMENT.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM program or product in this document is not intended to state or imply that only that program or product may be used. Any functionally equivalent program or product, that does not infringe IBM's intellectually property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

The provision of the information contained herein is not intended to, and does not grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

IBM, the IBM logo, System x, and System Storage are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries or both.