



Scott Smith

IBM System x Private Cloud Offering, Advanced Configuration: Architecture and Implementation Guide

Introduction

The IBM® System x® Private Cloud Offering provides solutions that are part of the Microsoft Fast Track for Hyper-V program. This program has requirements for systems management, virtualized management, and fault tolerance through the use of the System Center management platform, clustering, and redundancy in the network and storage.

This document provides setup and configuration details for the IBM 8-Node Hyper-V Fast Track configuration based on the IBM System x3755 M3 servers. Guidance and references are given throughout this document to aid understanding.

This document is intended for experienced IT engineers who have familiarity with Windows, Hyper-V, Microsoft Clustering, Networking, and iSCSI storage with IBM System Storage® N Series products. You should already have hands-on experience configuring these systems.

This paper describes planning and configuration of the solutions. It has the following topics:

- ▶ “Racking and power distribution” on page 3
- ▶ “Networking and VLANs” on page 4
- ▶ “Active Directory” on page 10
- ▶ “Storage” on page 10
- ▶ “Setup of IBM System x3550 M3 for Active Directory” on page 14
- ▶ “Setup of x3650 M3 for Management Cluster” on page 15
- ▶ “Setup of x3755 M3 for Production Hyper-V Cluster” on page 21
- ▶ “Conclusion” on page 24
- ▶ “Appendix A: Reference documents” on page 24
- ▶ “Appendix B: Brocade 2-port 10GbE Converged Network Adapter” on page 26
- ▶ “Appendix C: IBM Ethernet Switch B24X” on page 31
- ▶ “Appendix D: Worksheets” on page 35
- ▶ “Appendix E: SnapVault sample code” on page 39

Overview of architecture solution

This solution is based on the following clustered infrastructure:

- ▶ 2-node cluster running Hyper-V with virtual machines (VMs) running management tools
- ▶ 8-node cluster running Hyper-V for the production VM environment

This configuration is based on the hardware components in Table 1.

Table 1 Configuration components

Component	Quantity
IBM System x3550 M3 1U server (optional Active Directory server)	1
IBM System x3650 M3 2U servers (management servers)	2
IBM System x3755 M3 2U servers (Hyper-V servers)	8
IBM System Storage N6210 storage controller	1
EXN3000 disk expansion trays	5
IBM Ethernet Switch B24X 10GbE switches	2
Brocade 10GbE Dual Port CNA for IBM System x Adapters	20

A diagram of the overall architecture is shown in Figure 1 on page 3.

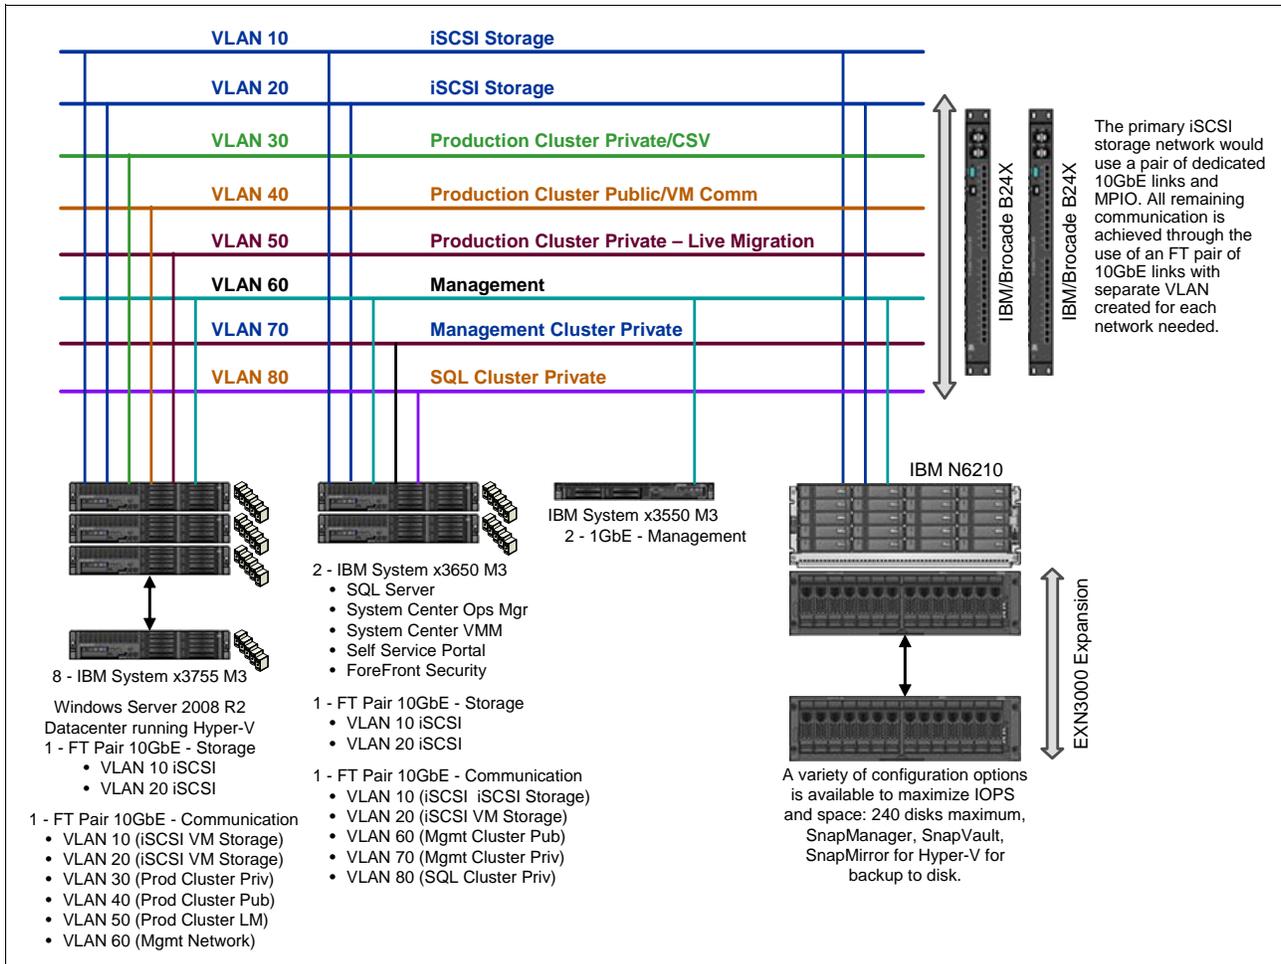


Figure 1 Solution architecture

Racking and power distribution

Installation of the power distribution units (PDUs) and their cabling must be performed before any of the systems or storage is placed in the rack. There should be sufficient PDUs to allow redundant power supplies to be fed from separate electrical circuits. Consider the following information when you cable the PDUs:

- ▶ Ensure sufficient separate electrical circuits and receptacles to support the PDUs needed.
- ▶ To minimize the chance of a single electrical circuit failure taking down a device, ensure that the PDUs, which feed redundant power supplies, are supplied from separate electrical circuits.
- ▶ For devices that have redundant power supplies, plan for one electrical cord from separate PDUs.
- ▶ Each 10 Gb Ethernet switch should be powered from a separate PDU.
- ▶ Locate Ethernet switches to allow the easiest cable runs possible.
- ▶ Locate storage in lower portions of the rack for added stability.

Networking and VLANs

A total of four Converged Network Adapter (CNA) ports will be available on each server running Hyper-V (two adapters in each server, each with two ports). The networks necessary to support this configuration are achieved with a combination of physical and virtually isolated networks.

Eight networks (VLAN IDs 10, 20, 30, 40, 50, 60, 70, and 80) are needed to build this configuration. All physical servers will use two CNA ports (one from each card) to connect to the iSCSI storage through a dedicated pair of 10GbE links. Use the IBM N Series device-specific multipath I/O (MPIO) driver for fault tolerance and load balancing. These storage connections are made on VLAN 10 and 20. The remaining networks are created by using the remaining two CNA ports and the Brocade adapter teaming software to create a fault tolerant (FT) NIC pair from the communications network. VLANs 30, 40, 50, 60, 70, and 80 are created by using the Brocade teaming utility to use this FT team as described in “Brocade network adapter teaming” on page 6.

For more technical information about the IBM B24X switch, see “Appendix C: IBM Ethernet Switch B24X” on page 31.

To assist in planning of the networks needed for this configuration see “Appendix D: Worksheets” on page 35, which lists each VLAN and the servers on each respective network.

VLAN description

The eight VLANs are described in Table 2.

Table 2 VLAN definitions

Network	Name	Description
VLAN 10	iSCSI Storage Network	Used for all iSCSI storage traffic. Physical servers maintain dedicated physical FT connection to this network. Virtual machines use a VLAN created from the communications network.
VLAN 20	iSCSI Storage Network	Used for all iSCSI storage traffic. Physical servers maintain dedicated physical FT connection to this network. Virtual machines use a VLAN created off the communications network
VLAN 30	Production Cluster Private & CSV Network	Used for cluster heartbeat and Cluster Shared Volume (CSV) traffic
VLAN 40	Production VM Communication Network	Used as primary communication network for production VMs (also known as Cluster Public Network)
VLAN 50	Production Live Migration	Second private network for production cluster used for live migration.
VLAN 60	Management Network	Network used for Active Directory (AD) and all communication for management of servers.
VLAN 70	Management Cluster Private Network	Used for cluster heartbeat and CSV traffic.
VLAN 80	SQL Cluster Private Network	Used for cluster heartbeat traffic.

Switch port layout

This configuration uses two IBM B24X 10 Gb Ethernet switches, each with 24 10GbE ports and four 1GbE ports. The first twelve ports are for iSCSI communication to the storage through VLAN 10 and VLAN 20. The best practice for configuring MPIO for iSCSI is to have separate VLANs for each iSCSI path. In addition, one port from each B24X switch can be used to connect the two switches. Port recommendations for B24X switches are as follows:

- ▶ Ports 1 - 2: N6210 Storage (for iSCSI storage network)
 - Untagged at the switch on VLAN 10 (Switch1)
 - Untagged at the switch on VLAN 20 (Switch2)
 - Storage must also be configured to send traffic out tagged on VLAN 10 and 20
- ▶ Ports 3 - 10: Production nodes (for iSCSI storage network)
 - Untagged at the switch on VLAN 10 (Switch1)
 - Untagged at the switch on VLAN 20 (Switch2)
- ▶ Ports 11 - 12: Management nodes (for iSCSI storage network)
 - Untagged at the switch on VLAN 10 (Switch1)
 - Untagged at the switch on VLAN 20 (Switch2)

- ▶ Port 13 (for cross-switch traffic)
 - Tagged at the switch with all VLANs
- ▶ Port 14: CorpNet uplink
 - Connectivity depends on CorpNet switch: Layer 2 and Layer 3 supported
- ▶ Ports 15 - 16: Management nodes
 - Tagged at the switch for VLANs 10, 20, 60, 70, and 80
 - Fault tolerant NIC team with interfaces for each VLAN (10, 20, 60, 70, 80) must be created on the server using Brocade adapter software. Enable jumbo frames on the devices before creating the team. See “Appendix B: Brocade 2-port 10GbE Converged Network Adapter” on page 26.
- ▶ Ports 17 - 24: Production nodes
 - Tagged at the switch on VLANs 10, 20, 30, 40, 50, and 60
 - Fault tolerant NIC team with interfaces for each VLAN (10, 20, 30, 40, 50, 60) must be created on the server by using Brocade adapter software. Enable jumbo frames on the devices before creating the team. See “Appendix B: Brocade 2-port 10GbE Converged Network Adapter” on page 26.
- ▶ Port 25: 1 GbE RJ-45 port – Optional AD server
 - Untagged at the switch on VLAN 60
- ▶ Port 26: 1 GbE RJ-45 port – N6210 management port
 - Untagged at the switch on VLAN 60
- ▶ Switch ports must allow the appropriate VLAN traffic as needed.

Figure 2 shows the switch port layout for B24X switch.

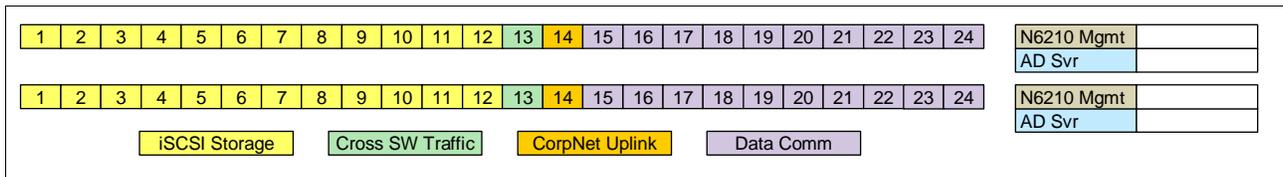


Figure 2 Switch port layout for B24X

Brocade network adapter teaming

Using the Brocade teaming software, the two CNA ports are combined as a fault tolerant team. This way, in combination with each port going to a separate switch, can provide fault tolerance for this configuration. In addition, the teaming software can be used to create individual VLAN devices off this team.

Before the team is created, both devices must be set to support 9000 byte jumbo frames under device manager. The VLANs created from these devices can then support jumbo frames if needed. The VLANs supporting iSCSI to a VM would be using jumbo frames from the teamed devices.

Figure 3 shows the creation of fault tolerant NIC team, and the associated MAC address for the team. Multiple VLANs can then be created from this team as shown.

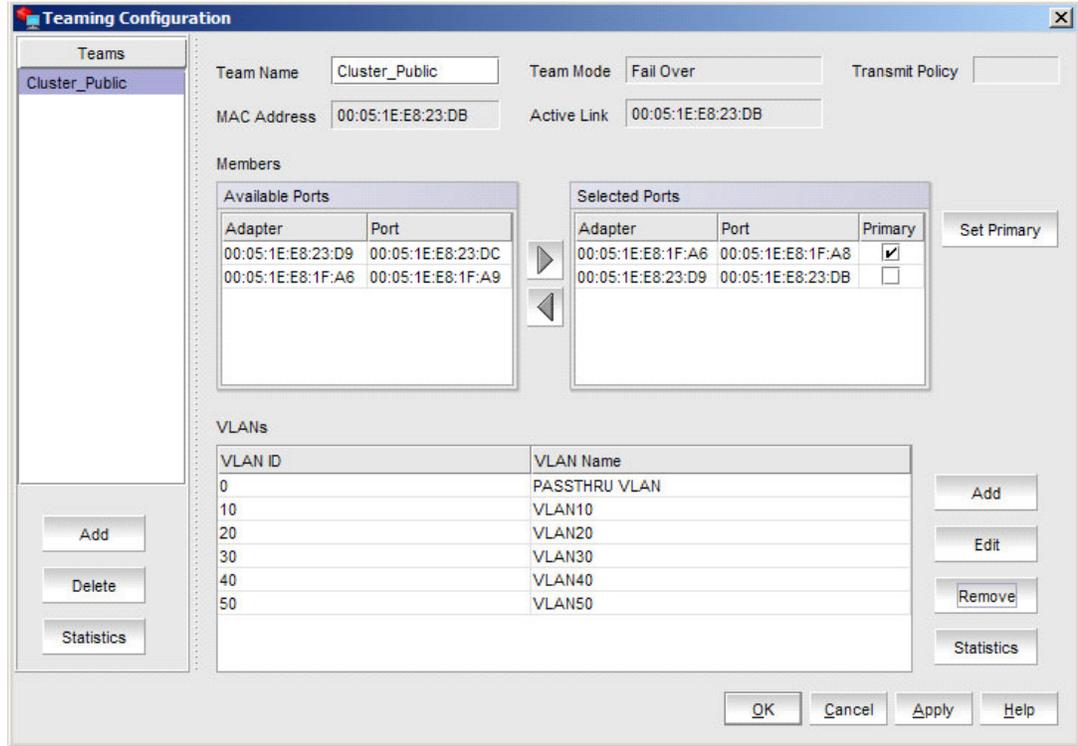


Figure 3 Brocade teaming software and VLAN creation

These devices are listed under Windows “Network Connections” and can be assigned an appropriate IP address, or visualized under Hyper-V for use by virtual machines. Consider setting a Windows quality of service (QoS) limit to bandwidth to an appropriate level for each network. See the *Hyper-V: Live Migration Configuration Guide*:

[http://technet.microsoft.com/en-us/library/ff428137\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/ff428137(WS.10).aspx)

The networks are each described in the following sections.

iSCSI storage networks (VLANs 10 and 20)

These networks are reserved for server access to the iSCSI storage network. All iSCSI traffic will be isolated on VLAN 10 and VLAN 20.

Storage controller access

Note the following information:

- ▶ There should be two Ethernet connections from each storage controller to the IBM B24X Ethernet switches. Each storage controller should maintain a 10 Gb Ethernet connection to VLAN 10 and VLAN 20. This way provides redundancy if one link fails.
- ▶ Set the switch and storage controller to support 9000 byte jumbo frames.

Physical host storage access

Note the following information:

- ▶ Each physical server must maintain two 10 Gb Ethernet connections to these storage networks, one connection each to VLAN 10 and VLAN 20. One port on each CNA is used to provide fault tolerance and terminate in the assigned port in each physical switch for redundancy.
- ▶ Ten ports on each of the two physical switches must be reserved for server connections to the dedicated iSCSI networks.
- ▶ Set the switch ports and CNA ports to use 9000 byte jumbo frames.
- ▶ Connect the following devices to the respective VLANs at each switch:
 - Each of the two management servers (two ports)
 - Each of the eight production Hyper-V servers (eight ports)
- ▶ The N6210 MPIO device-specific module must be installed and used to manage the two connections for fault tolerance and load balancing.

Load balancing can be achieved by setting the properties of the device to **Round Robin** from within the Server Manager as shown in Figure 4.

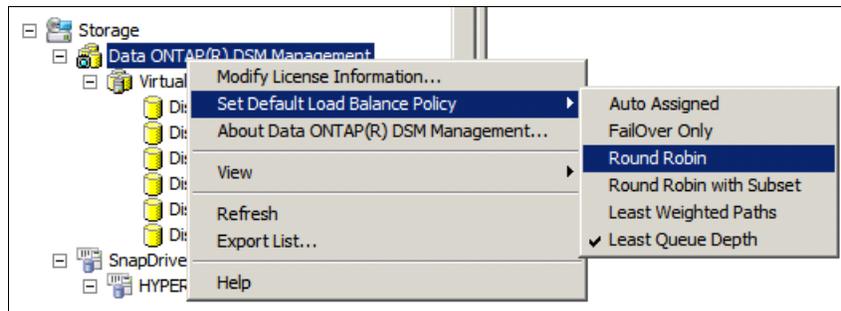


Figure 4 Setting MPIO load balancing policy

Virtual machine storage access

Note the following information:

- ▶ Virtual machines might require access to storage for either guest clustering or connection to pass-thru disks. These connections can be achieved in one of several ways:
 - A pass-thru disk is exposed to the parent partition and seen under disk manager of the physical machine. If left offline in the parent machine, it can then be made available to a virtual machine under the disk settings of the VM.
 - For direct connections to storage, such as for configuring guest clustering in virtual machines, the Microsoft iSCSI Initiator is available under Windows 2008 R2 and can be used to map available LUNs from the storage controller to specific virtual machines directly.

To support this approach, the VLAN 10 and VLAN 20 network connections must be made available to these virtual machines. These VLANs must be created under the NIC team that is created on the data communication network, with the remaining to the CNA ports as seen in “Brocade network adapter teaming” on page 6.

The IBM N Series DSM MPIO driver must also be loaded in the virtual machine.

- ▶ The Brocade adapter software should tag this network device, and the switch be configured to allow VLAN 10 and VLAN 20 traffic from this port to communicate on the iSCSI storage networks.

- ▶ At a minimum, the two physical management servers require virtual machine storage access to support the guest clustering of two virtual machine SQL servers.
- ▶ Consider making QoS settings on VM storage access to cap I/O traffic and not overwhelm network communication for the remaining virtual machines.

Cluster Heartbeat and Cluster Shared Volume Networks (VLANs 30, 70, and 80)

These networks are reserved for cluster private (heartbeat) communication between clustered servers. The management servers, Hyper-V production servers, and SQL servers have their own virtually isolated networks, created by using the Brocade adapter software to create separate VLANs. Switch ports must be configured to limit the scope of each of these VLANs as appropriate. There should be no routing on any VLAN that is being used for cluster private networks.

The following devices should access this network:

- ▶ Each of the two management servers: VLAN 70 should be unique for this cluster
- ▶ Each of the eight production Hyper-V servers: VLAN 30 should be unique for this cluster
- ▶ Each of the two SQL Server virtual machines: VLAN 80 should be unique for this cluster

Network interfaces for each of these VLANs are configured on the server through the use of the Brocade adapter teaming software.

Production Virtual Machine Communication Network (VLAN 40)

This network can support virtual machine communication for the Hyper-V production cluster. VLAN 40 is created from the FT NIC team as described in “Brocade network adapter teaming” on page 6. The production Hyper-V servers use VLAN 40 to create a virtual NIC through Hyper-V to be presented to production virtual machines. The production servers themselves do not have to maintain an IP address on this network.

Production Live Migration Network (VLAN 50)

Create VLAN 50 to support Live Migration for the production cluster. Create VLAN 50 as part of the FT team supporting the communication network using the Brocade adapter teaming software. Consider setting a QoS limit (see “Appendix B: Brocade 2-port 10GbE Converged Network Adapter” on page 26) to ensure plenty of bandwidth remains for VM communication during live migration.

Management Network (VLAN 60)

This network is reserved for management of the Hyper-V cloud. The System Center management servers communicate with the Hyper-V servers and supporting devices across this network. The two clusters of physical machines use their respective connections to support this communication, as needed.

The following devices and servers are on this network (VLAN 60):

- ▶ The N6210 storage controllers
- ▶ Each of the two physical management servers
- ▶ Management Virtual Machines (SCOM, SCVMM, SQL, SSP)

- ▶ Supporting these virtual machines requires the creation of virtual NICs from this device under Hyper-V for presentation to the management VMs
- ▶ Each of the eight production Hyper-V servers
 - This network is reserved for management and must not be exposed to the VMs

Routing summary

Configure routing on the switches to allow traffic between VLAN 50 and VLAN 60, and other corporate network resources as needed. Do not enable routing for VLAN 10, 20, 30, 40, 70, and 80 because these are private networks that support the operations of this configuration.

Active Directory

The Hyper-V Fast Track configuration must be part of an Active Directory domain. This way is required to form the clusters, and by some of the System Center management components. A domain controller is included in this configuration as shown in Figure 1 on page 3, but may be replaced by an existing domain structure if available.

A total of three distinct environments are created:

- ▶ An eight node cluster of Hyper-V production servers to host virtual machines to various users
- ▶ A two node Hyper-V cluster to manage the virtual environment
- ▶ A pair of clustered SQL virtual machines to provide DB services to the management cluster

Develop a consistent naming convention for the physical and virtual servers in this environment. For example the production Hyper-V servers might be named Hyper1, Hyper2, through Hyper8).

Storage

This configuration uses the IBM System Storage N6210 storage controller with optional 10GbE ports as the primary storage for the management and Hyper-V production clusters. The configuration that is used for this reference architecture consists of N6210 controller and five EXN 3000 expansion trays for a total of 120 disks. These disks will be configured and distributed across both storage controllers to balance the I/O workload

The storage subsystem used by this configuration will need to support the 2-node management cluster, and the 8-node Hyper-V production cluster. Storage I/O can vary widely depending on the applications being run. Therefore, be sure to profile your storage needs to ensure adequate space and I/O capacity for your needs. The storage layout shown in Figure 5 on page 11 provides a list of disk volumes needed on the N Series storage to support this configuration.

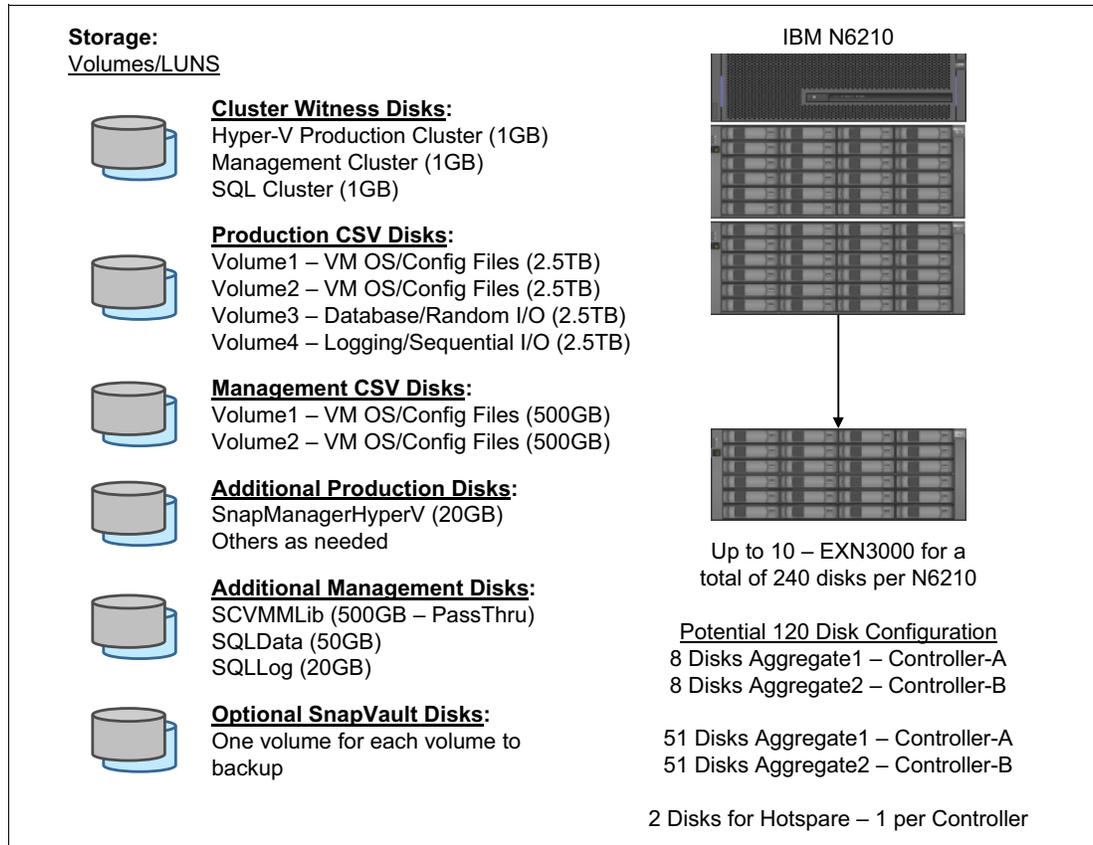


Figure 5 Storage layout

Additional storage I/O cards

The IBM N6210 has the following I/O cards installed to support operations:

- ▶ 10GbE PCIe cards installed (one card per controller for a total of two fault-tolerant 10GbE ports)
- ▶ Two 256GB Flash Cache Modules (one per controller)

Cabling

The IBM N6210 must be powered from two separate PDUs operating from unique electrical circuits to provide maximum fault tolerance. Each EXN3000 disk expansion drawer must be connected to PDUs in the same fashion as the N6210.

Each EXN3000 expansion drawer must be connected to the N6210 over a SAS connection. To maximize fault tolerance and performance, see the N Series documentation for cabling directions.

Each 10 Gb Ethernet connection must be plugged into its assigned port on the IBM Brocade B24X switch:

- ▶ Configure these ports as untagged ports on the switch in VLAN 10 and VLAN 20.
- ▶ Enable the 9000 byte jumbo frames on the controller, switch, and server adapters.

One connection from each 1GbE management port on the N6210 should be connected into the assigned 1GbE ports on each B24X switch. These ports are configured as untagged ports and assigned to VLAN 60 on the switch.

Management

Management of this storage is performed using a web-based management tool for the following actions:

1. Controller setup and license enablement
2. RAID Group configuration
3. Aggregate creation
4. Volume creation
5. LUN creation
6. iSCSI target assignments

To begin management of the IBM N6210, complete the following actions:

1. Verify one connection from each controller is connected to a 1 GbE RJ45 port on each of the B24X switches.
Set the 1 GbE management connections for each controller to the assigned IP address on VLAN 60.
2. If needed establish a serial connection to the controller (or controllers) to set the IP addresses for each management connection.
3. Ensure each of the 10GbE connections from each controller are connected to the appropriate ports on the switches.
4. Using a browser, connect to the N6210 storage unit and set the appropriate IP addresses for each controller on each of the two 10GbE ports:
 - The address is as follows:
`http://N6210-IP-address/na_admin`
 - The default user name is root, and no password (change the root password to something secure).
5. Enable the required licenses for the following features:
 - SnapDrive (storage management tool)
 - Snap Manager for Hyper-V (snapshot tool for Hyper-V)
 - SnapRestore (snap restore tool)
 - SnapVault (optional, if backing up to non-production storage)

Configuration

Before you configure the disks, note the following information:

- ▶ Verify all disks are available and online before you begin.
- ▶ Leave one disk available as a hot spare for each controller.

To configure the disks, use the following steps:

Note: With the numbers listed, the assumption is that 120 - 600 GB 15K SAS disks are available to the storage controllers.

1. Create RAID Groups that use RAID double parity (DP).

This enhanced RAID 6 (double parity) is optimized for the IBM N Series storage:

- RG1 = 8 Disks RAID DP
- RG2 = 8 Disks RAID DP
- RG3 = 17 Disks RAID DP
- RG4 = 17 Disks RAID DP
- RG5 = 17 Disks RAID DP
- RG6 = 17 Disks RAID DP
- RG7 = 17 Disks RAID DP
- RG8 = 17 Disks RAID DP

2. Create aggregates from the RAID Groups in the previous step.

Aggregates consist of multiple RAID groups and allow storage pools that consist of many more disks than is optimal for a single RAID group. I/O can then be spread across a greater number of spindles, increasing the performance of the storage system.

Create the following aggregates:

- AG1 = RG1 (approximately 1.7 TB)
 - This aggregate is reserved for volumes that support management cluster.
 - Assign Controller-A as the preferred controller.
- AG2 = RG2 (approximately 1.7 TB)
 - This aggregate is reserved for volumes supporting management cluster.
 - Assign Controller-B as the preferred controller.
- AG3 = RG3 + RG4 + RG5 (approximately 17.4 TB)
 - This aggregate is reserved for volumes supporting Hyper-V production cluster.
 - Assign Controller-A as the preferred controller.
- AG4 = RG6 + RG7 + RG8 (Approximately 17.4 TB)
 - This aggregate is reserved for volumes supporting Hyper-V production cluster.
 - Assign Controller-B as the preferred controller.

3. Create volumes.

Create the volume to support operations while distributing the I/O load across both controllers. Volumes are created from aggregates, and LUNS are created from volumes. Create one LUN from each volume. When volumes are created there is typically a 20% reserve set aside for snapshots. Sizes given here are after snapshot reserves are accounted for.

The following volume configuration can be used to support this reference configuration:

- From AG1 and AG2, the following volumes can be created and masked to the management cluster physical machines. The SQL cluster quorum will be masked to the IQNs of the SQL VMs, when known:
 - 500 GB management cluster CSV LUN
 - 500 GB SCVMM library LUN
 - 1 GB Management Cluster Quorum LUN
 - 1 GB SQL cluster quorum LUN
 - SQL DB disk (50 GB size varies by environment)
 - SQL Log Disk (20 GB size varies by environment)
- From AG3 and AG4, the following volumes can be created and masked to the Hyper-V production servers participating in the eight-node cluster:
 - Four 2.5 TB production CSVs for VM configuration and OS files
 - 20 GB production SnapInfo LUN
This LUN can be created using the SnapDrive utility as a clustered resource after the cluster has been created.
 - 1 GB production cluster quorum LUN
 - Optional staging LUN
- If SnapVault will be used to copy disk images from production disks, additional space is required:
 - SnapVault requires enough space to make the initial copy of the data, plus the incremental updates going forward
 - Volumes must be created as part of a qTree. A qTree-based volume provides additional flexibility and supports an extended setup of commands.
 - One volume is necessary for each volume to be backed up.

Multipath I/O (MPIO) fault-tolerance driver

The IBM N6210 provides fault tolerance with redundant controllers. It uses a device-specific MPIO driver to provide balanced I/O across the two redundant controllers. The driver to support this approach will be loaded in the physical servers making connections to the N6210. You may download this driver from the IBM N Series support site after properly registering:

<https://www.ibm.com/support/entdocview.wss?uid=ssg1S7003467>

An example of configuring the servers can be seen in Figure 4 on page 8.

Setup of IBM System x3550 M3 for Active Directory

An Active Directory (AD) server is required so that the servers can participate in a Microsoft cluster and is needed by several parts of the management solution. If an AD server is already available in your environment, the x3550 M3 offered as the AD server in this configuration is not needed.

The setup of the IBM x3550 M3 is straightforward because no optional devices are installed in this system.

Take the following actions before software configuration:

1. Confirm one onboard NIC port is connected to one of the 1 GbE RJ45 ports on each of the B24X switches.
2. Confirm that the B24X switch port is configured for VLAN 60.

Be sure the two local disks are configured as a RAID 1 set. For more information, see *ServeRAID M1015 SAS/SATA Controller User's Guide*:

ftp://ftp.software.ibm.com/systems/support/system_x_pdf/ibm_doc_sraidmr_m1015_2ndedition_userguide.pdf

See “Appendix A: Reference documents” on page 24 for links to other resources.

Next, complete the following actions:

1. Install Windows Server 2008 R2 SP1 Standard Edition.
2. Assign fixed TCP/IP address (VLAN 60).
3. Verify network connectivity.
4. Ensure the system UEFI code and integrated management module are updated to the latest supported version and latest device driver. The x3550 M3 driver matrix are located at the following address:

<http://www-933.ibm.com/support/fixcentral/systemx/selectFixes?parent=ibm~Systemx3550M3&product=ibm/systemx/7944&platform=Windows+2008+x64&function=all>

5. Promote to domain controller.
6. Run Windows Update.

System service accounts are added to the domain to support the System Center components that will be added later. At a minimum, the following accounts should be created:

- ▶ SVC_SCOM
- ▶ SVC_SCVMM
- ▶ SVC_VMMSSP
- ▶ SVC_SQL

Setup of x3650 M3 for Management Cluster

The management cluster consists of two IBM x3650 M3 systems with 24 GB of RAM each.

Complete the following actions before software configuration:

1. Confirm that two dual port CNA cards are installed.
2. Confirm one port from each CNA card is connected to the assigned storage network port (VLAN 10 and VLAN 20) on each respective B24X switch.
3. Confirm one port from each CNA card is connected to the communications network ports (VLAN 60, 70, and 80) on each B24X switch.

Set up involves installing Windows Server 2008 R2 Datacenter Edition on each server, validating networking, and storage connectivity, and enabling Hyper-V and Microsoft Clustering. The Datacenter version of the Windows product is preferable because of the unlimited number of Windows Server 2008 R2 virtual machines that can be run under the Datacenter without further licensing. The cost of the Datacenter version is based on the number of populated sockets on the server.

Configure the two local disks as a RAID 1 array. For more information, see the *ServeRAID M1015 SAS/SATA Controller User's Guide*:

ftp://ftp.software.ibm.com/systems/support/system_x_pdf/ibm_doc_sraidmr_m1015_2ndedition_userguide.pdf

“Appendix A: Reference documents” on page 24 has links to additional resources.

Next, complete the following actions:

1. Install and update Windows Server 2008 R2 SP1 Datacenter Edition.
2. Join Domain.
3. Enable Hyper-V.

When you create VMs, consider changing the VM memory allocation from static to dynamic to allow greater flexibility of system resources

4. Enable Windows failover clustering.
5. Run Windows Update
6. Ensure that the system UEFI code and integrated management module are updated to the latest supported version and latest device drivers. The x3650 M3 driver matrix is at the following web address:

<http://www-933.ibm.com/support/fixcentral/systemx/selectFixes?parent=ibm~Systemx3650M3&product=ibm/systemx/5454&platform=Windows+2008+x64&function=all>

7. Install N Series device-specific MPIO driver.

The N Series drivers and updates are available from your registered support link with IBM.

8. Install N Series Appliance Watch Pro driver.
9. Install N Series SnapDrive utility.
10. Install N Series SnapManager for Hyper-V.
11. IBM Director Platform agent for System Center upward integration.

Network configuration

In this section, we cover network configuration.

Physical server storage network

Complete the following steps to configure the network:

1. Assign fixed IP addresses to the two CNA connections that are attached to the iSCSI storage network. Use one port from each CNA; each port should be connected to a separate B24X switch.
2. Assign IP address for each device on VLAN10 & VLAN20 ID, respectively.
3. Validate connectivity for each device.

Communication networks

Complete the following steps to configure the network:

1. Using the Brocade teaming software in the parent partition, create a fault tolerant (FT) team from the remaining two CNA ports that are connected to the communications network.
2. Using the Brocade teaming software, create VLAN 10 and VLAN 20 interfaces over the FT team to support guest clustering:
 - Be sure no IP address for this device are on the physical servers because this address will be passed for VM direct connections to the storage only. To ensure this action does not happen, clear the management access check box when the virtual switch is created.
 - Create Hyper-V virtual switches to be passed to the SQL VMs to support clustering.
3. Using the Brocade teaming software, create VLAN 60 interface over the FT team to support management communications:
 - Assign fixed IP address to each physical server.
 - Create a Hyper-V virtual switch to be passed to VMs to support management operations.
4. Using the Brocade teaming software, create VLAN 70 interface over the FT team as a cluster private network for use by the two management servers:
 - Assign a fixed IP address to the physical servers
5. Using the Brocade teaming software. create VLAN 80 interface over the FT team for use as cluster private network for the two clustered SQL Servers:
 - There should be no IP address for this device on the physical servers.
 - Create a Hyper-V virtual switch to be passed to the SQL VMs to support guest clustering.
6. Using the PING command, validate network connectivity between the respective VLANs for the physical management servers on VLAN 60 and VLAN 70 for use in the management cluster creation.

Storage configuration

To ensure that the storage volumes are isolated to the specific servers assigned to them in the network, unique iSCSI qualified names (IQNs) are used to mask the management server's storage. To view this setting, select **Control Panel** → **iSCSI Initiator** → **Configuration** tab, as shown in Figure 6 on page 18.

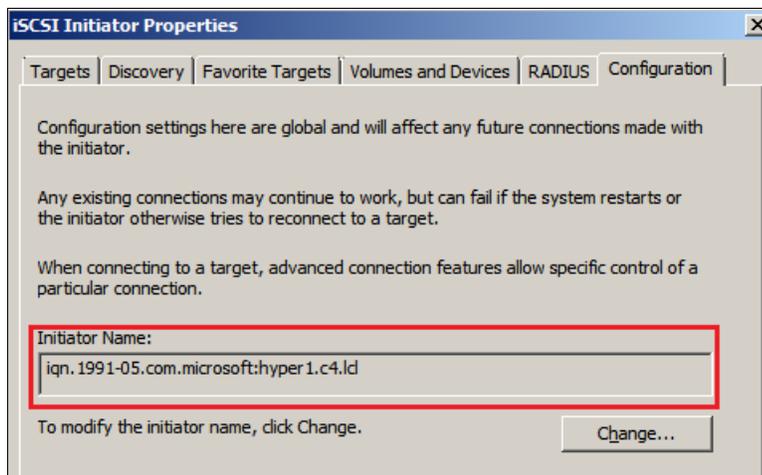


Figure 6 iSCSI Server IQN

Complete the following steps to configure the storage:

1. Record IQNs of each management server, being careful to ensure that the server has already joined the domain because not already joining the domain might change the IQN.
2. Use the IBM N6210 management console to add the recorded IQNs to storage volumes reserved for the management cluster:
 - Management Cluster CSV volume
 - Management Cluster Quorum volume
 - SCVMM library volume
3. Use the Windows iSCSI Initiator to discover the N6210 storage. A connection should exist from each CNA to each controller, for a total of four connections.
4. Use Disk Manager to validate that storage on N6210 is visible. Create partitions, and format the management cluster quorum and any management cluster CSV disks.

Cluster creation

To provide a highly available management environment, a cluster is created between the two management servers. The management servers will then run highly available virtual machines (HA VMs) on this cluster. The SQL Server virtual machines will not be HA VMs, but will form their own cluster to support SQL Server in a highly available environment. If you are not familiar with setting up a Microsoft cluster under Windows 2008, see “Appendix A: Reference documents” on page 24, which has links to more information.

First, use the Failover Clustering management console to complete the following steps:

1. Select the validate configuration option to validate the two physical management servers as a potential cluster, and address any errors found.
2. Create a cluster with the two physical management servers.

Next, complete the following actions:

1. Validate the cluster name and IP address.
2. Validate cluster networking.
3. Validate storage in cluster.
4. Enable Cluster Shared Volumes and add the designated volume as a CSV.
5. Set default paths for VM creation to use the CSV.

SQL Server VM

Several databases that are necessary to support the System Center servers (SCOM, SCVMM, SSP) will be hosted on a pair of clustered SQL Server virtual machines. To provide a fault tolerant environment a two-node SQL Server cluster is created by using guest clustering in virtual machines, one on each of the two management servers with Hyper-V installed. Although these virtual machines already reside on a pair of clustered servers, they must not be enabled as HA virtual machines. Running guest clustering in a HA virtual machine is not supported.

The following items are required for this configuration:

- ▶ 2 non-HA virtual machines (one on each management server)
- ▶ 4 virtual CPUs
- ▶ 6 GB of RAM
- ▶ 4 Virtual NICs: 1 client connection (VLAN 60), 1 cluster private (VLAN 80), 2 iSCSI (VLAN 10 and VLAN20). The N Series MPIO driver must be installed in the VM for storage access
- ▶ Windows Server 2008 R2 Enterprise Edition or Datacenter Edition installed and updated in the VMs. This version must be a clusterable version of Windows Server 2008 R2.

Storage being used for the SQL Server cluster must be allocated from the available storage in the management server RAID Group on the N6210, but the volumes must be isolated to the two SQL Server VMs by using IQN restrictions on the volume. Allocate the following storage:

- ▶ Operating system virtual hard disk (OS VHD) on CSV volume (30 GB)
- ▶ LUN1 – SQL databases
- ▶ LUN2 – SQL logging
- ▶ LUN3 – SQL cluster quorum (1 GB)

The SQL LUNs must be masked to the two virtual machines.

Use the SQL Service account created on the AD server when you install SQL Server. Guidance for installing and creating clustered SQL Server 2008 SP1 systems is at the following address:

[http://technet.microsoft.com/en-us/library/ms179530\(SQL.100\).aspx](http://technet.microsoft.com/en-us/library/ms179530(SQL.100).aspx)

The following databases are to be used by System Center:

- ▶ VMM <VMM_DB>
- ▶ VMM SSP (Self Service Portal) <SCVMMSSP>
- ▶ SCOM <Ops_Mgr_DB>
- ▶ SCOM <Ops_Mgr_DW_DB>
 - SCOM has a DB creation utility on the installation media to run on the SQL Servers to create the required databases.
 - SCVMM and SSP create the DB on the SQL Server during setup.

System Center Operations Manager

A System Center Operations Manager server must be available to support the management environment used in this configuration. To provide fault tolerance this server will be a highly available virtual machine.

This server is implemented as a virtual machine with the following specifications:

- ▶ HA virtual machine
- ▶ 2 virtual CPUs
- ▶ 4 GB of RAM
- ▶ 1 virtual NIC – Management Network Connection (VLAN 60)
- ▶ 1 OS VHD

To install System Center Operations Manager complete the following steps:

1. Install and update Windows Server 2008 R2 on VM.
2. Assign fixed IP address on VLAN 60.
3. Install SCOM 2007 R2.
4. Update SCOM.

The following roles are required by Operations Manager:

- ▶ Root Management Server
- ▶ Reporting Server (DB will reside on SQL Server)
- ▶ Data Warehouse (DB will reside on SQL Server)
- ▶ Operator Console
- ▶ Command Shell

The following Operations Management packs must be installed:

- ▶ Virtual Machine Manager 2008 R2 pack
- ▶ Windows Server Base Operating System pack
- ▶ Windows Failover Clustering pack
- ▶ Windows Server 2008 Hyper-V pack
- ▶ Microsoft SQL Server pack
- ▶ Microsoft Windows Server Internet Information Server pack
- ▶ System Center packs
- ▶ IBM Director Upward Integration pack (see “Appendix A: Reference documents” on page 24 for links to more information)
- ▶ IBM N Series integration pack

System Center Virtual Machine Manager

A System Center Virtual Machine Manager server must be available to support the management environment used in this configuration. To provide a fault tolerance, this server will be a highly available virtual machine.

The server is implemented as a virtual machine with the following specifications:

- ▶ HA virtual machine
- ▶ 2 virtual CPUs
- ▶ 4 GB of RAM
- ▶ 1 virtual NIC: Management Network Connection (VLAN 60)
- ▶ 1 OS VHD
- ▶ 1 data VHD or pass-thru disk for VMM Library

To install System Center Operations Manager complete the following steps:

1. Install Windows Server 2008 R2 on VM.
2. Assign fixed IP address on VLAN 60.
3. Update Windows.
4. Install SCVMM 2008 R2.
5. Update SCVMM.

The following roles are required by Virtual Machine Manager:

- ▶ VMM Server
- ▶ Administrative Console
- ▶ VMM Library
- ▶ SQL DB will reside on remote SQL cluster
- ▶ Command Shell

In addition, VMM will be integrated with Operations Manager. VMM uses System Center Operations Manager to monitor the health and availability of the virtual machines and virtual machine hosts that VMM is managing (see “Appendix A: Reference documents” on page 24 for links to integration guidance).

System Center Self Service Portal 2.0

System Center Self Service Portal 2.0 must be available to support the management environment and self-service provisioning of virtual machines. To provide a fault tolerance this server will be a highly available virtual machine. Microsoft installation and configuration instructions are available from the following web address:

<http://www.microsoft.com/downloads/en/details.aspx?FamilyID=fef38539-ae5a-462b-b1c9-9a02238bb8a7&displaylang=en#SystemRequirements>

The server is implemented as a virtual machine with the following specifications:

- ▶ HA virtual machine
- ▶ 2 virtual CPUs
- ▶ 4 GB of RAM
- ▶ 1 virtual NIC – Management Network Connection (VLAN 60)
- ▶ 1 OS VHD

To install System Center Self Service Portal 2.0, complete the following steps:

1. Install and update Windows Server 2008 R2 on VM.
2. Assign fixed IP address on VLAN 60.
3. Install and update IIS with required components.
4. Install and update SSP 2.0.

Setup of x3755 M3 for Production Hyper-V Cluster

The production cluster will consist of eight IBM x3755 M3 systems with 256GB of RAM each. Before configuring the software, complete the following steps:

1. Confirm that two dual port CNA cards are installed.
2. Confirm one port from each CNA card is connected to the assigned storage network port (VLAN 10 and VLAN20) on each B24X switch.
3. Confirm one port from each CNA card is connected to the assigned communications network port (VLAN 30 through VLAN 80) on each B24X switch.

Software installation

Setup involves installing Windows Server 2008 R2 SP1 Datacenter Edition on all servers, validating networking, setting up storage, and enabling Hyper-V and Microsoft Clustering.

First, the two local disks must be configured as a RAID 1 array. See the *ServeRAID M1015 SAS/SATA Controller User's Guide*:

ftp://ftp.software.ibm.com/systems/support/system_x_pdf/ibm_doc_sraidmr_m1015_2ndedition_userguide.pdf

“Appendix A: Reference documents” on page 24 has links to additional resources.

Next, complete the following steps:

1. Install and update Windows Server 2008 R2 Datacenter Edition.
2. Join Domain.
3. Enable Hyper-V.
4. Enable Windows Failover Clustering.
5. Ensure the system UEFI code, integrated management module, and latest device drivers are updated to the latest supported version. Go to the following web address:
<http://www-933.ibm.com/support/fixcentral/systemx/selectFixes?parent=ibm~Systemx3755M3&product=ibm/systemx/7164&platform=Windows+2008+x64&function=all>
6. Install the latest Windows updates.
7. Install IBM N Series device-specific MPIO driver.
8. Install IBM N Series Appliance Watch Pro driver.
9. Install IBM N Series SnapDrive utility.
10. Install IBM N Series SnapManager for Hyper-V.
11. Install Data ONTAP PowerShell Toolkit.
12. Install IBM Director level 1 agent for System Center upward integration.

Network configuration

Assign fixed VLAN 10 and 20 IP addresses to the two CNA connections that are attached to the iSCSI storage network. Use one port from each CNA. Connect each port to a separate B24X switch. Note the following information:

- ▶ No NIC teaming or VLAN assignment is required on the server for these two ports.
- ▶ Validate connectivity for each device.

Using the Brocade teaming software, create FT team from two remaining CNA ports that are connected to communications network, as follows:

1. Create a VLAN 10 interface from the FT team if iSCSI access is needed for any virtual machines on the production network.
A Hyper-V switch should be created for VM use only. Do not allow the physical server access.
2. Create a VLAN 20 interface from the FT team for the cluster private network.
A Hyper-V switch should be created for VM use only. The physical server should not be allowed access.

3. Create a VLAN 30 interface from the FT team for the cluster private network.
Assign the VLAN a static IP address, however no Hyper-V switch should be created on top of this because it is reserved for clustering of the physical machines.
4. Create a VLAN 40 interface on the FT team as a second cluster private network for use by Live Migration
Assign the VLAN a static IP address, however no Hyper-V switch should be created on top of this because it is reserved for Live Migration of virtual machines.
Consider setting a Windows QoS limit to bandwidth to an appropriate level for each network. See the *Hyper-V: Live Migration Configuration Guide*:
[http://technet.microsoft.com/en-us/library/ff428137\(ws.10\).aspx](http://technet.microsoft.com/en-us/library/ff428137(ws.10).aspx)
5. Create a VLAN 50 interface on the FT team to support virtual machine communications.
A Hyper-V switch should be created for VM use only. The physical server should not be allowed access.
6. Create a VLAN 60 interface on the FT team to support management communications.
Assign the VLAN a static IP address, however no Hyper-V switch should be created on top of this because it is reserved for management purposes.
7. Validate the network connectivity between physical servers on VLAN 30, 40, and 60 respectively.

Storage configuration

To properly enable the storage to be isolated from other servers in the network the unique IQNs are used to mask the management server's storage. To view this information, select **Control Panel** → **iSCSI Initiator** → **Configuration** tab.

Complete the following steps:

1. Record IQNs of each host server, carefully ensuring that the server has already joined the domain, because the process of joining the domain could change the IQN.
2. Use the IBM N6210 management console to add the recorded IQNs to storage volumes that are reserved for the production cluster:
 - Production Cluster CSV volumes
 - Production Cluster Quorum volume
3. Delay the creation of the SnapInfo volume until you have the SnapDrive utility installed and operational. This step makes the creation and masking of the this volume much easier.
4. Use the Windows iSCSI Initiator to discover the N6210 storage
5. Use Disk Manager to validate that storage on N6210 is visible.
6. Create partitions and format the production cluster quorum, and future production cluster CSV disks.

If configuring SnapVault to back up snap images to non-production disks, see “Appendix E: SnapVault sample code” on page 39, which shows a short PowerShell script that can be set to run as part of the SnapManager for Hyper-V task. Set it to run after the snaps are taken.

Cluster creation

To provide a highly available virtual machine environment, a cluster is made up of the eight production servers. The servers created from these clustered resources are then run as highly available virtual machines on this cluster. If you are not familiar with setting up a Microsoft cluster under Windows Server 2008, see the “Appendix A: Reference documents” on page 24 for links to more information.

First, use the Failover Clustering management console to complete the following steps:

1. Validate the eight physical production servers as a potential cluster, and address any errors found.
2. Create a cluster with the eight physical productions servers.

Next, complete the following actions:

1. Validate cluster name and IP address.
2. Validate cluster networking.
3. Validate storage in cluster.
4. Enable Cluster Shared Volumes and add the designated volumes as a CSV.
5. Set default paths for VM creation to use the Cluster Shared Volume.

Conclusion

After you complete these steps, you have two operational Hyper-V clusters, one for management, and one for production virtual machines. Fault tolerance is achieved at several points through the use of clustering, redundant network switches, and redundant storage controllers for the iSCSI storage.

System Center Operations Manager, with the help of the IBM upward integration module, monitors the health of the physical system; System Center Virtual Machine manager monitors the health of the virtual environment and provides the VM deployment mechanism to the Self Service portal.

System Center Self Service Portal can be used to allocate a subset of the resources available to this environment to one or more organizations. These organizations can then start or stop virtual machines as needed. Billing and metering information for the resources used can be collected and presented to assist in breaking down, and presenting costs to the organizations that are using these resources.

Appendix A: Reference documents

The following resources provide additional information regarding setup and configuration.

- ▶ IBM System x3550 M3 Installation and User's Guide
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5083597>
- ▶ IBM System x3650 M3 Installation and User's Guide
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5083595>
- ▶ IBM System x3755 M3 Installation and User's Guide
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5085682>

- ▶ IBM firmware update and best practices:
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/firmware_update_best_practices.pdf
- ▶ *IBM ServeRAID M1015 SAS/SATA Controller* guide
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/ibm_doc_sraidmr_m1015_2ndedition_userguide.pdf
- ▶ *IBM ServeRAID M5014/M5015 SAS/SATA Controllers* guide
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/serveraid-m5014_m5015_u_g_1st_edition_20090924.pdf
- ▶ IBM IMM and uEFI white paper
ftp://ftp.software.ibm.com/systems/support/system_x/transitioning_to_uefi_and_imm.doc
- ▶ Hyper-V and Microsoft clustering
[http://technet.microsoft.com/en-us/library/cc732181\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc732181(WS.10).aspx)
- ▶ Hyper-V Live Migration and QoS configuration
[http://technet.microsoft.com/en-us/library/ff428137\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/ff428137(WS.10).aspx)
- ▶ IBM Director Agent download (Platform Agent)
<http://ibm.com/systems/software/director/downloads/agents.html>
- ▶ IBM Director System Center integration (download integration modules)
<http://ibm.com/systems/software/director/downloads/integration.html>
- ▶ Microsoft Virtual Machine Manager Self Service Portal 2.0
<http://www.microsoft.com/downloads/en/details.aspx?FamilyID=fef38539-ae5a-462b-b1c9-9a02238bb8a7&displaylang=en#SystemRequirements>
- ▶ *NetApp SnapManager 1.0 for Hyper-V Best Practices*
<http://media.netapp.com/documents/tr-3805.pdf>
- ▶ Microsoft System Center Virtual Machine Manager Deployment Guide
<http://technet.microsoft.com/en-us/library/ee441285.aspx>
- ▶ Configuring Operations Manager Integration with VMM
<http://technet.microsoft.com/en-us/library/cc956099.aspx>
- ▶ Microsoft SQL Server 2008 documentation
[http://technet.microsoft.com/en-us/library/bb418470\(SQL.10\).aspx](http://technet.microsoft.com/en-us/library/bb418470(SQL.10).aspx)
- ▶ Microsoft ForeFront documentation
<http://technet.microsoft.com/en-us/library/ff630950.aspx>
- ▶ IBM white paper about SCOM integration
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/ibm_hw_mp_mgt_guide.pdf
- ▶ IBM website for Microsoft Systems Management integration
<http://ibm.com/support/entry/portal/docdisplay?lnocid=SYST-MANAGE>

Appendix B: Brocade 2-port 10GbE Converged Network Adapter

The Brocade 2-port 10GbE Converged Network Adapter for IBM System x is a two-port Converged Network Adapter (CNA) PCIe 2.0 x8 card that mounts inside supported IBM System x servers. This card combines the functions of a 10GbE network interface card (NIC) and Fibre Channel host bus adapter (HBA).

The CNA supports full Fibre Channel over Ethernet (FCoE) protocol offload and allows Ethernet and storage traffic to run simultaneously over a converged link. Advanced capabilities such as support for jumbo frames, VLAN tagging, TCP segmentation offload (TSO) and large send offload (LSO) also helps in iSCSI storage environments.

Note: Be sure you are running driver v2.3 or higher to support the configurations required for Fast Track.

Features

The Converged Network Adapter supports the following features:

- ▶ PCI Express x8 Gen 2 host interface
- ▶ IPv4 and IPv6
- ▶ Brocade Host Connectivity Manager (HCM) device management and Brocade Command-Line Utility (BCU) tools
- ▶ Unified management with IBM System Storage Data Center Fabric Manager (DCFM) and IBM Systems Director

The adapter has the following features that are related to Ethernet:

- ▶ 10 Gbps throughput per port full duplex
- ▶ 1.8 million packets per second per port in each direction (700 byte packets, latency < 2 μ s)
- ▶ Checksum/CRC offloads for FCoE packets, IPv4/IPv6 TCP and UDP packets, and IPv4 header
- ▶ VLAN support: up to 64
- ▶ Windows NIC teaming: up to eight teams per system, eight ports per team
- ▶ Jumbo frame support: up to 9600 bytes
- ▶ Header data split (HDS) feature for advanced link layer
- ▶ Receive side scaling (RSS) feature for advanced link layer
- ▶ TCP segmentation offload (TSO) and large send offload (LSO)
- ▶ Link aggregation (NIC teaming)
- ▶ Priority-based Flow Control (IEEE 802.1Qbb)
- ▶ Enhanced Transmission Selection (IEEE 802.1Qaz)
- ▶ Data Center Bridging eXchange Protocol (DCBX)

Managing CNAs

The Brocade CNAs for IBM System x can be managed by using either the Brocade Command-Line Utility (BCU) that is installed with the drivers on the system, or the GUI-based Host Connectivity Manager (HCM). HCM utilizes a client-agent based model. The HCM agent is also installed with the adapter drivers. The HCM GUI client can be installed either on the same system that the CNAs are installed on or on a separate management system.

Figure 7 shows HCM managing a Brocade dual-port CNA.

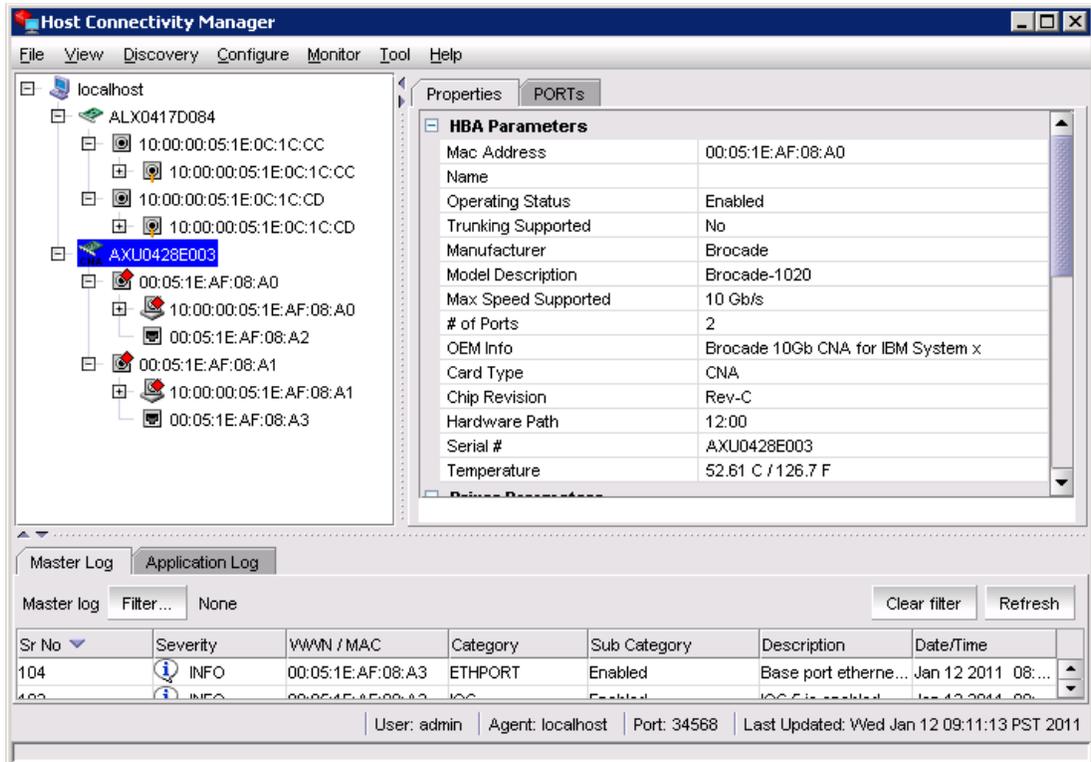


Figure 7 Brocade Host Connectivity Manager

Context-sensitive menus are used to configure devices within HCM.

Brocade CNAs support NIC teaming, which virtualizes two or more physical network interfaces into a single logical interface. In essence, the driver advertises the same MAC and IP address of all teamed interfaces. Although the following two types of teaming methods are available, this configuration uses active-passive:

- ▶ Active-Active

This method transmits across all devices in the team in a round-robin method; it provides redundancy and load balancing. This method can be used only if the teamed interfaces have uplinks to the same physical switch (or multiple switches) that support Brocade multi-chassis trunking (MCT)

- ▶ Active-Passive

One physical interface in a teamed group, called the *primary port*, is active while the other interfaces are on standby for redundancy. If the primary port stops, a secondary port is chosen to be the next primary.

The active-passive configuration can be in either of the following modes:

- Failback mode: If the original port designated to be the primary goes down and comes back up, it becomes the primary port again
- Failover mode: The secondary port that became the primary stays as the primary port until it fails.

The current Fast Track configuration supports only active-passive mode. Typically, all settings such as MTU size, link speed, port VLAN ID, and others on the interfaces to be teamed must match. Up to eight separate teams can be created on a single system with up to eight ports per team. All ports must be from Brocade CNAs.

Note: FCoE traffic is not supported when active-active NIC teaming or bonding is used.

To configure NIC teaming and VLANs over a NIC team, right-click on the server level, **localhost** in the example, and select **Teaming**, as shown in Figure 8.

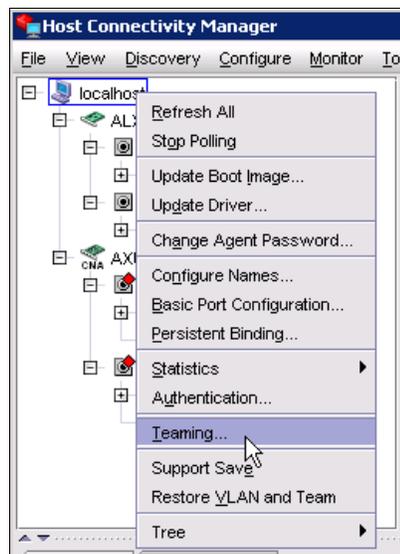


Figure 8 Creating a team

Use the Teaming Configuration window to create, destroy, or configure NIC teams.

VLANs for these teams can also be added and deleted. In the next example, a team named Test is created with two ports in failover mode. A pass-thru VLAN 0 always exists; this is the device interface that forwards untagged packets. In addition, a device interface for VLAN ID 10 has also been created for this NIC team. Packets sent out of this device are IEEE 802.1q, tagged for VLAN 10.

See Figure 9.

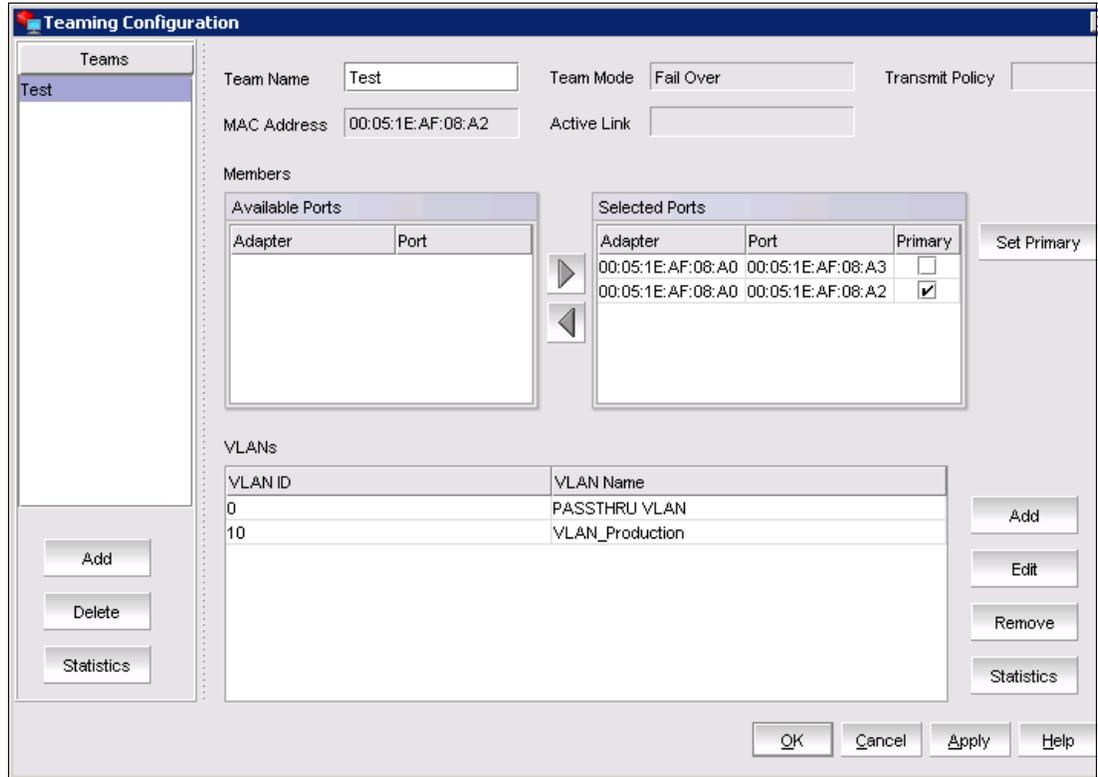


Figure 9 Configuring the teaming

In Windows, open Network Connections. Figure 10 shows that two new device interfaces are listed: Local Area Connection 7, which is TEAM#Test and represents the pass-thru VLAN 0; and Local Area Connection 8, which represents VLAN ID 10 for the NIC team we just created.

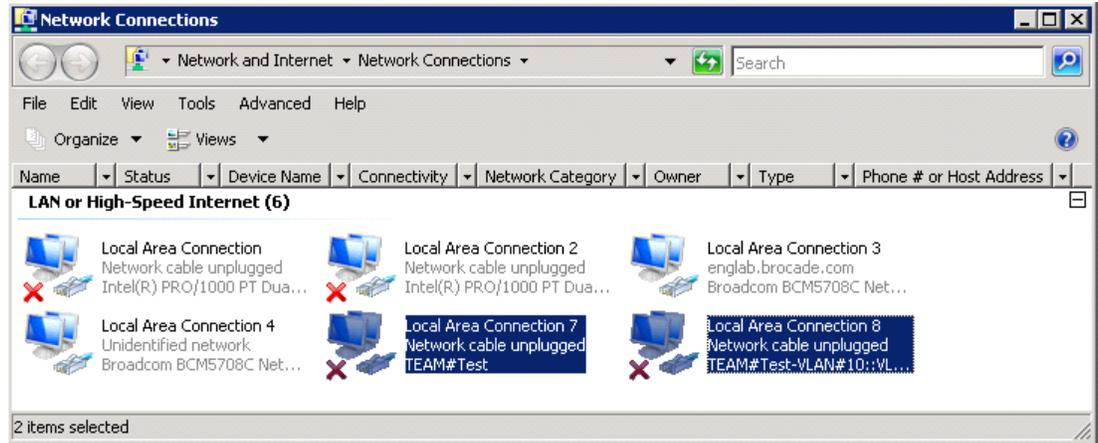


Figure 10 Network connections

Windows Device Manager

Jumbo frames and other options such as IPv4 Checksum Offload, TSO, LSO, RSS, HDS, VLAN support, and VLAN ID can be configured as follows, before building teams of the devices:

1. Open the Windows Device Manager, right-clicking the Brocade adapter you want to configure.
2. Select **Properties**.
3. Click the **Advanced** tab.

Although VLAN ID can be configured in the Advanced property tab; it is not recommended in the context of configuring your Microsoft Fast Track environment. Multi-VLAN configuration should be done using either BCU or HCM and not Windows Device Manager. See Figure 11.

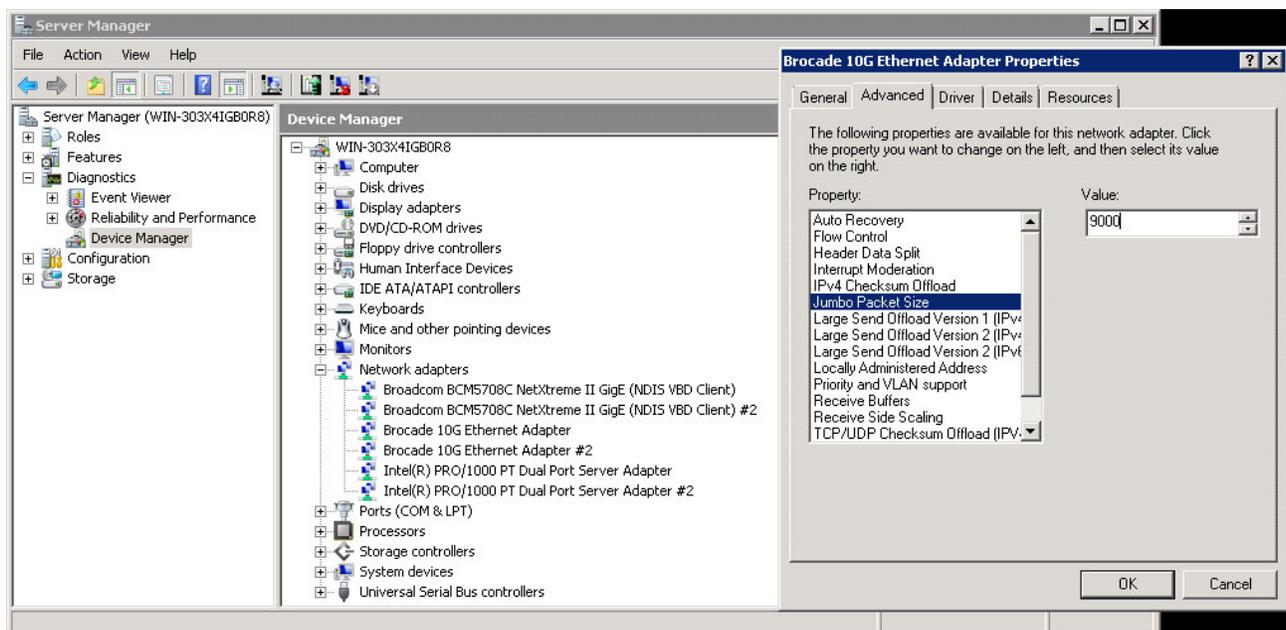


Figure 11 Device Manager

Related links

These documents provide additional information regarding setup and configuration.

- ▶ *Brocade 10Gb CNA for IBM System x at-a-glance guide*
<http://www.redbooks.ibm.com/abstracts/tips0718.html?Open>
- ▶ Brocade 2-port 10GbE CNA drivers and documentation for IBM
<http://www.brocade.com/services-support/drivers-downloads/CNA/IBM.page>
- ▶ IBM ServerProven® compatibility
<http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/xseries/lan/matrix.html>
- ▶ Brocade 2-port 10GbE CNA for IBM System x product page
<http://www.ibm.com/systems/storage/product/systemx/cna/index.html>

Appendix C: IBM Ethernet Switch B24X

The IBM Ethernet Switch B24X is a compact high-performance, high-availability, and high-density 1 RU switch specifically designed for mission-critical data center requirements. This switch provides 24 10/1 GbE (SFP+) ports plus four 10/100/1000 MbE (RJ45) ports of connectivity in an ultra-low-latency, cut-through, non-blocking architecture. The IBM Fast Track configuration uses the B24X for storage connectivity for a high-speed iSCSI storage network, eliminating the need for multiple aggregated 1 GbE networks that would normally be required.

Note: Make sure you are running firmware R04.2.00a or higher to support the configurations required for Fast Track.

Features are as follows:

- ▶ Capable of full Layer2 and IPv4 Layer3
- ▶ 24x ports of dual-speed 10/1 GbE SFP+ ports plus 4x 10/100/1000 MbE RJ45 ports
- ▶ Capable of line speed, full-duplex throughput on all ports—488 Gbps total switch capacity and ultra-low latency (1.5 micro-seconds)
- ▶ 1+1 redundant, hot-swappable and load-sharing 300 W ac power supplies for high availability
- ▶ Hot-swappable, resilient triple fan assembly with automatic speed control for efficient front-to-back cooling: 2+1 redundant fans in assembly
- ▶ Industry standard command-line interface (CLI)

For a full list of features, see the following resources:

- ▶ *IBM Ethernet Switch B24X IBM System x at-a-glance guide*
<http://www.redbooks.ibm.com/abstracts/tips0785.html?Open>
- ▶ IBM Ethernet Switch B24X product page
<http://www.ibm.com/systems/x/options/networking/ethernetb24x/index.html>

Managing B24X

The IBM Ethernet Switch B24X uses an industry-standard CLI for management. This CLI can be accessed over Telnet or SSHv2. Fault management and monitoring can also be accomplished by using IBM Systems Director software. In addition, the switch supports industry-standard SNMP.

For additional details about configuring the switch, see “Related links” on page 34. Also, see the following resources:

- ▶ *IBM Ethernet Switch B24X Installation and User's Guide*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5084455>
- ▶ Release Notes that are included with each firmware update for the switch

Fast Track configuration

Example 1 is a sample configuration file for an IBM Ethernet Switch B24X that is configured as outlined in this document. This configuration file is for one of two B24Xs that are deployed in a redundant fashion; the second B24X can be configured similarly.

Example 1 Fast Track configuration

```
Current configuration:
!
ver 04.2.00
!
!
global-stp
!
!
vlan 1 name DEFAULT-VLAN by port
 spanning-tree
!
vlan 10 name iSCSI by port
 tagged ethe 13 to 24
 untagged ethe 1 to 12
 spanning-tree
!
vlan 30 name Prod_Cluster_Private by port
 tagged ethe 13 to 14 ethe 17 to 24
 spanning-tree
!
vlan 40 name Prod_Cluster_Live_Migration by port
 tagged ethe 13 to 14 ethe 17 to 24
 spanning-tree
!
Vlan 50 name Prod_VM_Network by port
 tagged ethe 13 to 14 ethe 17 to 24
 router-interface ve 50
 spanning-tree 802-1w
 spanning-tree 802-1w priority 0
 spanning-tree 802-1w ethe 14 admin-pt2pt-mac
!
vlan 60 name Management_Network by port
 tagged ethe 13 to 24
 untagged ethe 25 to 26
 router-interface ve 60
 spanning-tree 802-1w
 spanning-tree 802-1w priority 0
 spanning-tree 802-1w ethe 14 admin-pt2pt-mac
!
vlan 70 name Management_Cluster_Private by port
 tagged ethe 13 to 16
 spanning-tree
!
vlan 80 name SQL_Cluster_Private by port
 tagged ethe 13 to 16
 spanning-tree
!
!
```

```

!
!
jumbo
hostname IBM_B24X_Bot
router vrrp-extended
snmp-server community ..... rw
interface ethernet 13
    spanning-tree 802-1w admin-pt2pt-mac
    link-aggregate configure key 10001
    link-aggregate active
!
interface ethernet 14
    link-aggregate configure key 10001
    link-aggregate active
!
interface ethernet 15
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 16
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 17
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 18
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 19
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 20
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 21
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 22
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 23
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 24
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 25
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 26
    spanning-tree 802-1w admin-pt2pt-mac
!
interface ethernet 27
    spanning-tree 802-1w admin-edge-port
!
interface ethernet 28

```

```
spanning-tree 802-1w admin-edge-port
!
interface ve 50
 ip address 192.168.50.3 255.255.255.0
 ip vrrp-extended vrid 50
  backup
  advertise backup
  ip-address 192.168.50.1
  enable
!
interface ve 60
 ip address 192.168.60.3 255.255.255.0
 ip vrrp-extended vrid 60
  backup
  advertise backup
  ip-address 192.168.60.1
  enable
!
!!
!
end
```

Related links

These documents provide additional information regarding setup and configuration.

- ▶ *IBM Ethernet Switch B24X at-a-glance guide*
<http://www.redbooks.ibm.com/abstracts/tips0785.html?Open>
- ▶ *IBM b-type Data Center Networking: Product Introduction and Initial Setup*
<http://www.redbooks.ibm.com/abstracts/sg247785.html?Open>
- ▶ *IBM b-type Data Center Networking: Design and Best Practices*
<http://www.redbooks.ibm.com/abstracts/sg247786.html?Open>
- ▶ IBM System Networking support page for IBM Ethernet Switch B24X (firmware and documentation)
[http://www.ibm.com/support/entry/portal/Overview/Hardware/Systems/Systems_networking/Ethernet_Switches/IBM_Ethernet_Switch_B24X_\(4002-X2A\)](http://www.ibm.com/support/entry/portal/Overview/Hardware/Systems/Systems_networking/Ethernet_Switches/IBM_Ethernet_Switch_B24X_(4002-X2A))
- ▶ IBM Ethernet Switch B24X product page
<http://www.ibm.com/systems/x/options/networking/ethernetb24x/index.html>

Appendix D: Worksheets

Use the worksheets in this section to help you plan the network configuration (Table 3 through Table 10 on page 38).

Table 3 VLAN 10

VLAN 10 addresses (iSCSI storage)	IP addresses
N6210 Controller-A	
N6210 Controller-B	
Management Server 1 - NIC1	
Management Server 2 - NIC1	
SQL Server VM1 (Coming off FT team VLAN10)	
SQL Server VM2 (Coming off FT team VLAN10)	
Hyper-V Production Server 1 - NIC1	
Hyper-V Production Server 2 - NIC1	
Hyper-V Production Server 3 - NIC1	
Hyper-V Production Server 4 - NIC1	
Hyper-V Production Server 5 - NIC1	
Hyper-V Production Server 6 - NIC1	
Hyper-V Production Server 7 - NIC1	
Hyper-V Production Server 8 - NIC1	
Production VMs (if needed from FT team VLAN10)	

Table 4 VLAN 20

VLAN 20 addresses (iSCSI storage)	IP addresses
N6210 Controller-A	
N6210 Controller-B	
Management Server 1 – NIC2	
Management Server 2 – NIC2	
SQL Server VM1 (Coming off FT team VLAN20)	
SQL Server VM2 (Coming off FT team VLAN20)	
Hyper-V Production Server 1 – NIC2	
Hyper-V Production Server 2 – NIC2	
Hyper-V Production Server 3 – NIC2	
Hyper-V Production Server 4 – NIC2	
Hyper-V Production Server 5 – NIC2	
Hyper-V Production Server 6 – NIC2	
Hyper-V Production Server 7 – NIC2	
Hyper-V Production Server 8 – NIC2	
Production VMs (if needed from FT team VLAN20)	

Table 5 VLAN 30

VLAN 30 addresses (Production Cluster Private)	IP addresses
Hyper-V Production Server 1	
Hyper-V Production Server 2	
Hyper-V Production Server 3	
Hyper-V Production Server 4	
Hyper-V Production Server 5	
Hyper-V Production Server 6	
Hyper-V Production Server 7	
Hyper-V Production Server 8	

Table 6 VLAN 40

VLAN 40 addresses (Production Live Migration)	IP addresses
Hyper-V Production Server 1	
Hyper-V Production Server 2	
Hyper-V Production Server 3	
Hyper-V Production Server 4	
Hyper-V Production Server 5	
Hyper-V Production Server 6	
Hyper-V Production Server 7	
Hyper-V Production Server 8	

Table 7 VLAN 50

VLAN 50 addresses (VM Production)	IP address range
Routing IP Address (Virtual Routing Address)	
Routing IP Address (Top Switch)	
Routing IP Address (Bottom Switch)	

Table 8 VLAN 60

VLAN 60 addresses (Management)	IP addresses
Routing IP Address (Virtual Routing Address)	
Routing IP Address (Top Switch)	
Routing IP Address (Bottom Switch)	
N6210 Controller-A	
N6210 Controller-B	
Management Server 1	
Management Server 2	
Management Cluster	
SQL Server VM1	
SQL Server VM2	
SQL Server Cluster	
SQL Server Cluster Service	
SQL MSDTC Service (Optional)	
System Center Operations Manager	
System Center Virtual Machine Manager	
System Center Self Service Portal 2.0	
Hyper-V Production Server 1	
Hyper-V Production Server 2	
Hyper-V Production Server 3	
Hyper-V Production Server 4	
Hyper-V Production Server 5	
Hyper-V Production Server 6	
Hyper-V Production Server 7	
Hyper-V Production Server 8	
Hyper-V Production Cluster	

Table 9 VLAN 70

VLAN 70 addresses (Management Cluster Private)	IP addresses
Management Server 1	
Management Server 1	

Table 10 VLAN 80

VLAN 80 addresses (SQL Cluster Private)	IP addresses
SQL Server VM1	
SQL Server VM2	

Appendix E: SnapVault sample code

Example 2 shows the code to initiate SnapVault transfer between volumes.

Example 2 Initiate SnapVault transfer between volumes

```
#Sample Code to initiate SnapVault transfer between volumes
#This can be called from within SnapManager for Hyper-V as a method to backup up
Snapshots off of the #production volumes.
#
Import-Module DataONTAP
$contr_N6210a=Connect-NaController -Name N6210a -RPC
$contr_N6210b=Connect-NaController -Name N6210b -RPC
$snaplist1=Get-NaSnapshot -Name vProd20S -Controller $contr_N6210a
Start-NaSnapvaultSecTransfer -SecondaryPath /vol/vProd20S_SV/qProd20S_SV
-PrimarySnapshot $snaplist1[0].Name -Controller $contr_N6210b
```

The team who wrote this paper



Scott Smith is an IBM System x Systems Engineer working at the IBM Center for Microsoft Technology. Over the past 15 years, Scott has worked to optimize the performance of IBM x86-based servers running the Microsoft Windows Server operating system and Microsoft application software. He has extensive experience in helping IBM customers understand the issues that they are facing, and developing solutions that address them.

Thanks to the following people for their contributions to this project:

- ▶ Steven Tong, Brocade
- ▶ Bryan Hoke, Brocade
- ▶ David Ye, IBM
- ▶ Daniel Ghidali, IBM
- ▶ Marco Rengan, IBM
- ▶ David Watts, IBM Redbooks®
- ▶ Linda Robinson, IBM Redbooks
- ▶ Diane Sherman, IBM Redbooks

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document REDP-4771-00 was created or updated on July 22, 2011.



Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.



Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM®

Redbooks®

Redbooks (logo) ®
ServerProven®

System Storage®

System x®

The following terms are trademarks of other companies:

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

SnapVault, SnapRestore, SnapManager, SnapDrive, Data ONTAP, NetApp, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.