



Scott Smith

IBM System x Private Cloud Offering: Architecture and Implementation Guide

Overview

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 System x Private Cloud Offerings based on the IBM System x3650 M3 servers. The intended audience for this document is IT professionals who are familiar with the setup and configuration of hardware and software products.

This document is intended for experienced IT engineers with familiarity with Windows®, Hyper-V, Microsoft clustering, networking, and iSCSI storage with IBM N Series storage products.

The topics in this paper are:

- ▶ “Configuration overview” on page 2
- ▶ “Planning actions” on page 2
- ▶ “Racking and power distribution” on page 2
- ▶ “Networking” on page 3
- ▶ “Storage” on page 8
- ▶ “Setup of the x3550 M3 for AD and DNS” on page 12
- ▶ “Setup of the x3650 M3 for management cluster” on page 13
- ▶ “Setup of the x3650 M3 for the production Hyper-V cluster” on page 18
- ▶ “Conclusion” on page 20
- ▶ “Appendix A: Related documents” on page 21
- ▶ “Appendix B: Brocade 2-port 10 GbE CNA” on page 22
- ▶ “Appendix C: IBM Ethernet Switch B24X” on page 28
- ▶ “Appendix D: IP address worksheet” on page 32
- ▶ “The team who wrote this paper” on page 34

Configuration overview

This configuration is based on the following:

- ▶ One IBM x3550 M3 1U server (optional Active Directory server)
- ▶ Ten IBM x3650 M3 2U servers
- ▶ One IBM N3600 storage controller and disk expansion trays
- ▶ Two IBM B24X switches
- ▶ Twenty Brocade 10 GbE dual-port Converged Network Adapter (CNA) for IBM System x adapters

This hardware will be used to create the following:

- ▶ Two-node cluster running Hyper-V with Microsoft System Center running management tools
- ▶ Eight-node cluster running Hyper-V for the production virtual machine environment

Planning actions

Before setting up the environment, some planning should take place to ensure consistent and professional results.

Items to consider are:

- ▶ Racking location and power distribution
- ▶ Network cabling plan and port assignment
- ▶ Naming convention
- ▶ Active Directory domain
- ▶ Management servers and cluster
- ▶ Hyper-V servers and cluster
- ▶ Virtual machine naming convention
- ▶ TCP/IP addresses, subnets, and virtual LANS (VLANs)
- ▶ Network naming conventions

Racking and power distribution

Installation of the power distribution units (PDUs) and PDU cabling should be performed *before* any of the systems or storage are placed in the rack. There should be sufficient PDUs to allow redundant power supplies to be fed from separate electrical circuits.

When cabling the PDUs, keep the following in mind:

- ▶ Ensure that there are sufficient separate electrical circuits and receptacles to support the PDUs.
- ▶ To minimize the possibility of a single electrical circuit failure taking down a device, ensure that the PDUs that 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 GbE switch should be powered from a separate PDU.

- ▶ Locate GbE switches to allow the easiest cable runs possible.
- ▶ Locate storage in lower portions of rack for added stability.

Networking

There will be a total of four CNA ports available on each server running Hyper-V. The networks needed to support this configuration are achieved with a combination of physical and virtually isolated networks.

There are seven networks needed to build this configuration. All physical servers will utilize two CNA ports (one from each card) to connect to the iSCSI storage through a dedicated pair of 10 GbE links using the IBM N Series device-specific multipath I/O (MPIO) driver for fault tolerance and load balancing. The remaining networks will be created using the last two CNA ports and the Brocade adapter teaming software to create a fault tolerant (FT) Network Interface Card (NIC) pair off of the communications network. VLANs will be created to utilize this FT team as described in Figure 1.

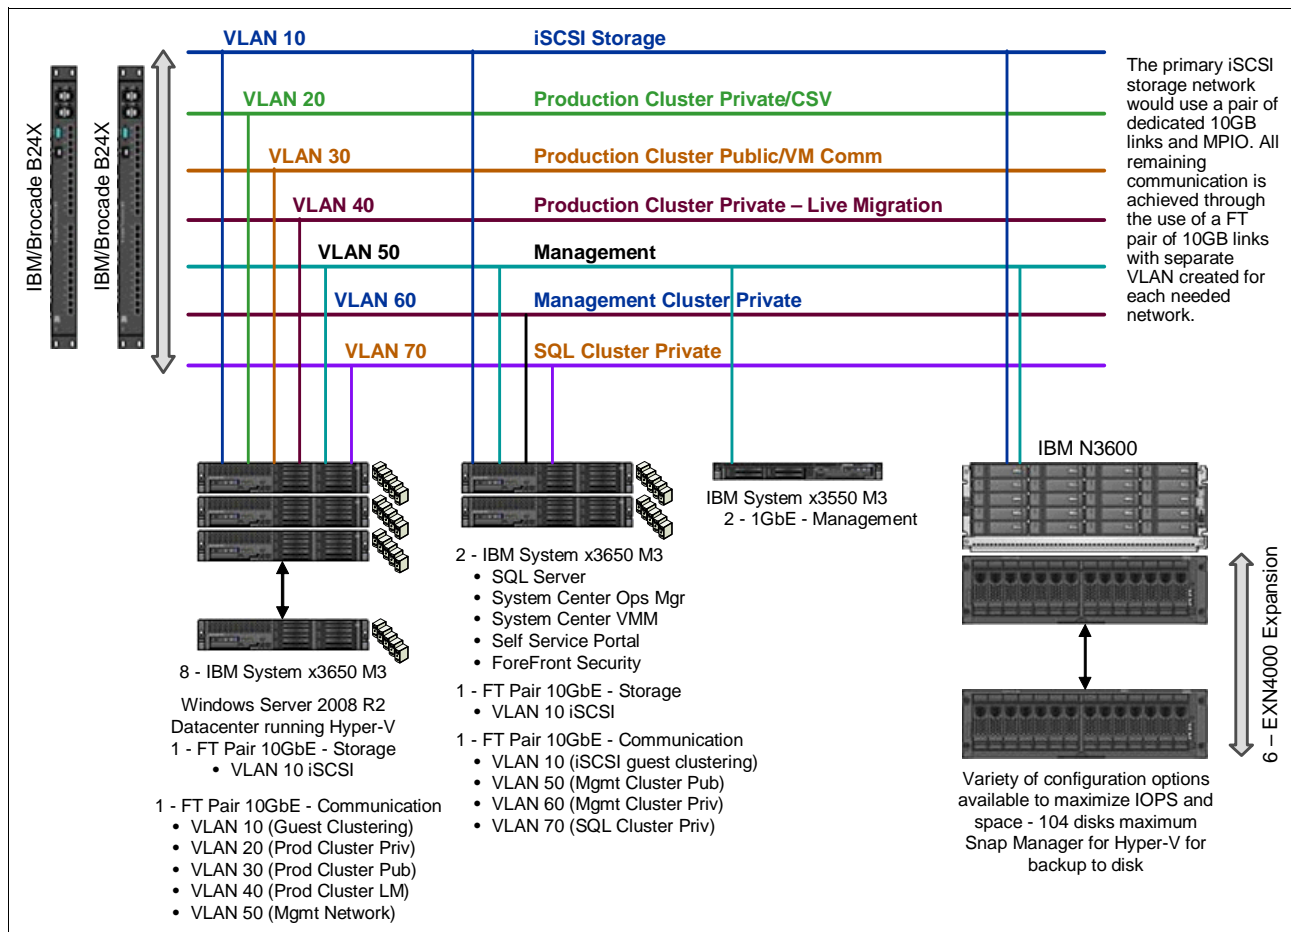


Figure 1 Diagram of VLAN and iSCSI storage

Table 1 further describes VLAN 10 and iSCSI storage. When these have been assigned IP addresses in your environment, you can use the worksheet that we have provided in “Appendix D: IP address worksheet” on page 32 to list your addresses.

Table 1 VLAN 10 and iSCSI specifications

Network	Name	Description
VLAN 10	iSCSI storage network	Used for all iSCSI storage traffic, these physical servers maintain a dedicated physical, FT connection to this network. Virtual machines utilize a VLAN created off of the communications network.
VLAN 20	Production cluster private and cluster-shared volume (CSV) network	Used for cluster heartbeat and CSV traffic.
VLAN 30	Production VM communication network	Used as the primary communication network for production VMs (also known as Cluster Public Network).
VLAN 40	Production Live Migration	The second private network for the production cluster, this network is used for Live Migration.
VLAN 50	Management network	Network used for AD and all communication for management of servers.
VLAN 60	Management cluster private network	Used for cluster heartbeat and CSV traffic.
VLAN 70	SQL cluster private network	Used for cluster heartbeat traffic.

This configuration uses two IBM B24X 10 GbE switches. One port from each B24X switch is used to connect the two switches. The following are port recommendations for the B24X switches:

- ▶ Ports 1 through 8: Production nodes (for iSCSI storage network)
 - Untagged at the switch on VLAN 10
- ▶ Ports 9 through 10: Management/Structured Query Language (SQL) nodes (for iSCSI storage network)
 - Untagged at the switch on VLAN 10
- ▶ Port 11: iSCSI storage array
 - Tagged at the switch on VLAN 10
 - Storage array must also be configured to send traffic out, tagged on VLAN 10
- ▶ Port 12: For cross-switch traffic
 - Tagged at the switch with all VLANs
- ▶ Ports 13 and 14: CorpNet uplink
 - Connectivity depends on the CorpNet switch. Layers 2 and 2 are supported.
- ▶ Ports 15 and 16: Management/SQL nodes
 - Tagged at the switch on VLANs 10, 50, 60, and 70
 - Network interfaces for each VLAN must be created on the server using Brocade adapter software. See “Appendix B: Brocade 2-port 10 GbE CNA” on page 22.
- ▶ Ports 17 through 24: Production nodes
 - Tagged at the switch on VLANs 10, 20, 30, 40, and 50
 - Network interfaces for each VLAN must be created on the server using Brocade adapter software. See “Appendix B: Brocade 2-port 10 GbE CNA” on page 22.

- ▶ Port 25: 1 GbE RJ-45 port; optional AD server
Untagged at the switch on VLAN 40
- ▶ Port 26: 1 GbE RJ-45 port; N3600 management port
Untagged at the switch on VLAN 50

Switch ports must allow the appropriate VLAN traffic as needed (Figure 2).




Figure 2 Switch port layout for the two IBM B24X 10 GbE switches

Brocade network adapter teaming

Using the Brocade teaming software, the two communication CNA ports are combined as a fault-tolerant team. This, in combination with each port going to a separate switch, provides fault tolerance for this configuration. In addition, the teaming software can be used to create individual VLAN devices off of this team (Figure 3).

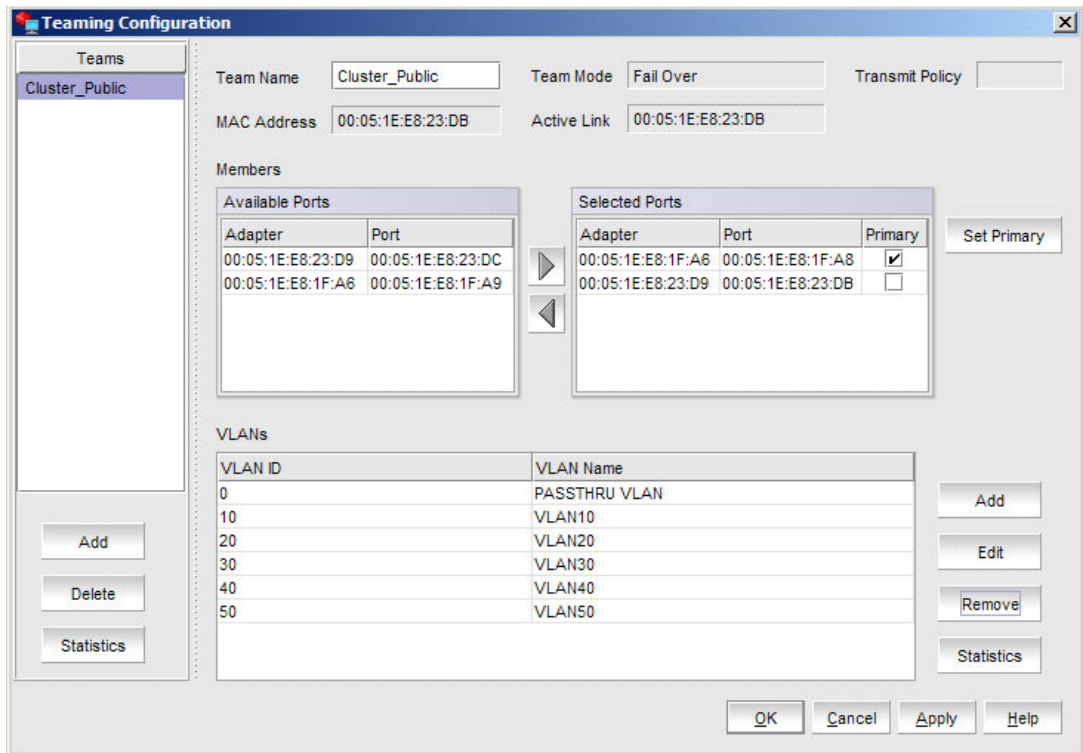


Figure 3 Teaming configuration

These devices display under Network Connections in Windows networking and can be assigned an IP address or virtualized under Hyper-V for use by virtual machines. Consider setting a Windows Quality of Service (QoS) limit to an appropriate bandwidth level for each network. See “Appendix A: Related documents” on page 21.

The networks are described in the following sections.

iSCSI storage network

This network is reserved for server access to the iSCSI storage network. All iSCSI traffic should be isolated on VLAN 10.

Storage access

Each of the two storage controllers should maintain two 10 GbE connections to the VLAN. For redundancy, one port from each controller terminates in the assigned port in each physical switch.

The switch and storage controller should be set to support 9000 byte jumbo frames.

Physical host access

Specifications for accessing the physical hosts are:

- ▶ Each physical server should maintain two 10 GbE connections to this VLAN. For redundancy, one port on each CNA is used to provide fault tolerance and terminate in the assigned port in each physical switch.
- ▶ Ten ports on each of the two physical switches should be reserved for the dedicated iSCSI network.
- ▶ The switch ports and CNA ports should be set to use 9000-byte jumbo frames.
- ▶ The following devices should be connected to this VLAN at the switch:
 - Each of the two management servers (4 ports)
 - Each of the eight production Hyper-V servers (16 ports)
- ▶ Use the N3600 MPIO device to manage the two connections for fault tolerance and load balancing.

Virtual machine access

Specifications for accessing virtual machines follow:

- ▶ Virtual machines that might require access to storage for either guest clustering or connection to pass thru disks. These connections are achieved using the Brocade teaming software to create the appropriate VLAN 10 device off of the FT team on the communication network.
- ▶ The Brocade adapter software should tag this network device and the switch that is configured to allow VLAN 10 traffic from this port to communicate on the iSCSI storage network.
- ▶ At a minimum, the two physical management servers require this to support the guest clustering of two virtual machine SQL servers.
- ▶ Consider defining QoS settings (see “Appendix A: Related documents” on page 21) to cap I/O traffic and not overwhelm network communication for the remaining virtual machines.

Cluster private networks and cluster shared volumes

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

- ▶ Each of the two management servers
VLAN 60 should be unique for this cluster.
- ▶ Each of the eight production Hyper-V servers
VLAN 20 should be unique for this cluster.
- ▶ Each of the two SQL Server virtual machines
VLAN 70 should be unique for this cluster.

The network interface for each VLAN is configured on the server using the Brocade adapter teaming software (see “Appendix B: Brocade 2-port 10 GbE CNA” on page 22).

Live Migration

Create separate VLAN to support Live Migration for the production cluster. This is done using the Brocade adapter teaming software to create VLAN 40 off of the FT team that supports the communication network. Consider setting a QoS (see “Appendix A: Related documents” on page 21) limit to ensure that plenty of bandwidth is left for VM communication during Live Migration.

Virtual machine communications network

This network supports virtual machine communications for the production (Hyper-V) cluster. VLAN 30 is created from the FT team, as described in “Live Migration” on page 7.

The production Hyper-V servers utilizes this to create a virtual NIC through Hyper-V to be presented to production virtual machines. The production servers themselves should not maintain an IP address on this network.

Management network

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 the network. The two clusters of physical machines utilize their respective connections to support this communication as needed.

The following devices are on this network (VLAN 50):

- ▶ The N3600 storage controller
 - ▶ Each of the two physical management servers
 - Management virtual machines (System Center Operations Manager, System Center Virtual Machine Manager, and so on).
 - Supporting the virtual machines requires the creation of virtual NICs from these devices under Hyper-V for presentation to the management VMs.
 - ▶ Each of the eight production Hyper-V servers
- This network is reserved for management and should not be exposed to the virtual machines.

Routing summary

Routing should be configured on the switches to allow traffic between VLAN 30 and VLAN 50 and other corporate network resources as needed. No routing should be enabled for VLANs 10, 20, 50, 60, or 70, as these are private networks used to support the operations of this configuration.

Storage

This configuration uses the IBM N3600 storage controller with optional 10 GbE ports as the primary storage for the management and Hyper-V production clusters.

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. It is good practice to profile your storage needs to ensure adequate space and I/O capacity for your needs.

The IBM System Storage® N3600 configuration used for this reference architecture consists of an N3600 controller and six EXN 4000 expansion trays. These disks are configured and distributed across both storage controllers to balance the I/O workload (Figure 4).

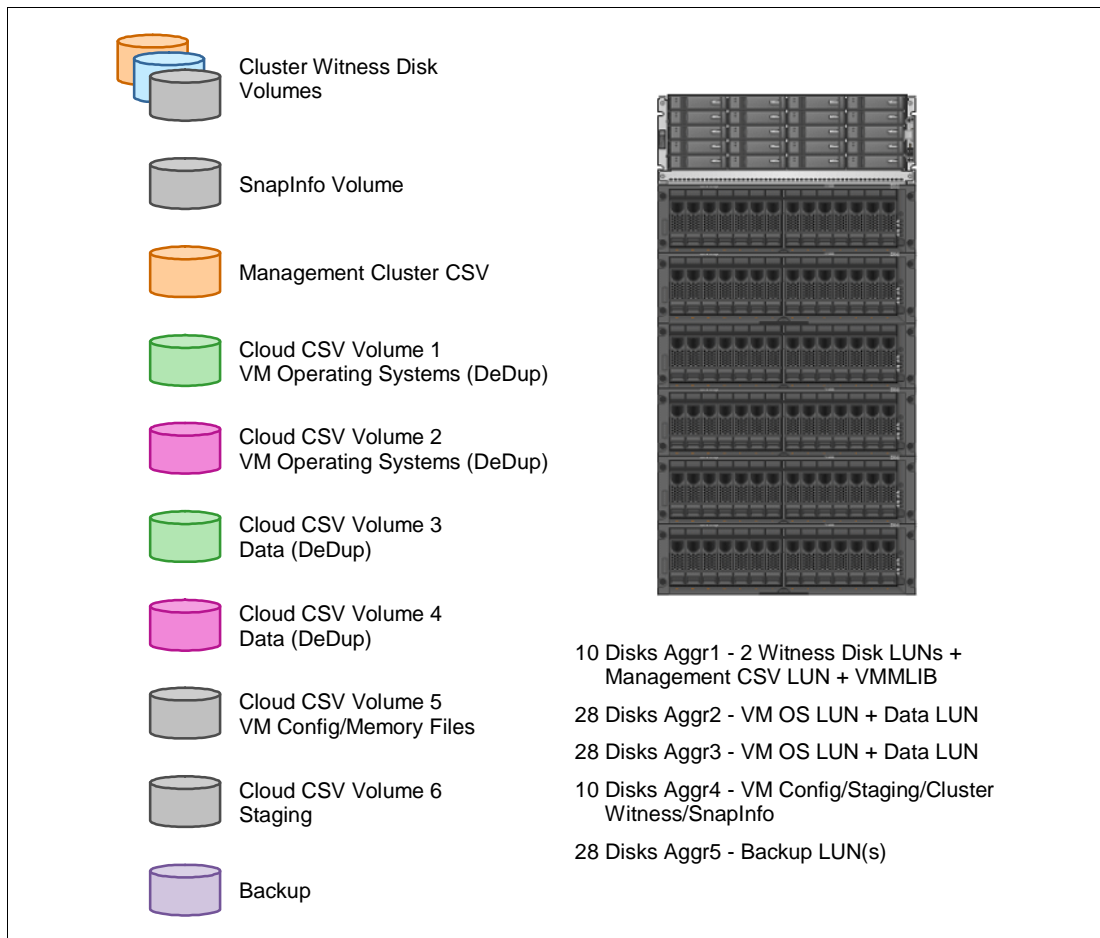


Figure 4 IBM System Storage N3600 configuration

The IBM N3600 needs to have two additional 10 GbE PCIe cards installed. One card is needed per controller for a total of two FT 10 GbE ports.

Cabling

The IBM N3600 should be powered from two separate PDUs operating off of unique electrical circuits to provide maximum fault tolerance. Each of the EXN4000 disk expansion drawers needs to be connected to PDUs in the same fashion as the N3600.

Each of the EXN4000 expansion drawers needs to be connected to the N3600 over a fiber connection. To maximize fault tolerance and performance, refer to N Series documentation for cabling instructions.

Each of the two 10 GbE connections should be plugged into its assigned ports on the IBM Brocade B24X switch. These ports are configured as tagged ports on the switch in VLAN 10. Enable 9000-byte jumbo frames on the controller, switch, and server adapter (see "Appendix C: IBM Ethernet Switch B24X" on page 28).

One connection from each controller 1 GbE management port on the N3600 should be connected to one of the four 1 GbE ports on each of the B24X switches. These ports are configured as untagged ports and assigned to VLAN 50 on the switch.

Management

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

- ▶ Controller setup and license enablement
- ▶ RAID group configuration
- ▶ Aggregate creation
- ▶ Volume creation
- ▶ Logical unit number (LUN) creation
- ▶ iSCSI target assignments

To begin management of the IBM N3600, perform the following:

1. Verify that 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 50.
 - If needed, establish a serial connection to the controllers to set the IP address for each management connection.
2. Ensure that each of the 10 GbE connections from each controller is connected to the appropriate port on the switches.
3. Using a browser, connect to the N3600 storage unit and set the appropriate IP address for each controller on each of the two 10 GbE ports. The URL to use is of the form:
`http://netapp_ip/na_admin`
Use user ID is *root*. The default password is blank (no password).
4. Enable the required licenses and features:
 - SnapDrive®
 - Snap Manager for Hyper-V
 - SnapRestore®
 - SnapVault® (optional, if backing up to non-production storage)
 - Enable jumbo frames for each 10 GbE network interface
 - A-SIS (deduplication) (for appropriate volumes)

Configuration

To configure the disks, do the following:

1. Verify that all disks are available and online.
2. Leave one disk for each controller available as a hot spare.
3. Create RAID groups :
 - RG1 = 9 Disks RAID DP
 - RG2 = 9 Disks RAID DP
 - RG3 = 14 Disks RAID DP
 - RG4 = 14 Disks RAID DP
 - RG5 = 14 Disks RAID DP
 - RG6 = 14 Disks RAID DP

- RG7 = 14 Disks RAID DP
 - RG8 = 14 Disks RAID DP
4. Create aggregates from the RAID groups from the previous step:
- AG1 = RG1 (approximately 3 TB)
 - Reserved for volumes supporting the management cluster.
 - Assign Controller-A as the preferred controller.
 - AG2 = RG2 (approximately 3 TB)
 - Reserved for volumes supporting the Hyper-V production cluster.
 - Assign Controller-B as the preferred controller.
 - AG3 = RG3 + RG4 (approximately 12 TB)
 - Reserved for volumes supporting the Hyper-V production cluster.
 - Assign Controller-A as the preferred controller.
 - AG4 = RG5 + RG6 (approximately 12 TB)
 - Reserved for volumes supporting the Hyper-V production cluster.
 - Assign Controller-B as the preferred controller.
 - AG5 = RG7 + RG8 (approximately 12 TB)
 - Reserved for volumes supporting backup operations.
 - Assign Controller-B as the preferred controller.

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

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

1. From AG1, create the following volumes and mask them to the management cluster physical machines. The SQL cluster quorum will be masked to the iSCSI qualified names (IQNs) of the SQL VMs when known:
 - 1 TB management cluster CSV LUN
 - 500 GB SCVMM library LUN
 - 1 GB management cluster quorum LUN
 - 1 GB SQL cluster quorum LUN
 - SQL database disk (size varies)
 - SQL log disk (size varies)
2. From AG2, create the following volumes and mask them to the Hyper-V production servers participating in the eight-node cluster:
 - 1 TB production CSV for VM configuration files
 - 10 GB production SnapInfo LUN

This should be created using the SnapDrive utility as a clustered resource after the cluster has been created.
 - 1 GB production cluster quorum LUN
 - 500 GB staging LUN

3. From AG3, create the following volumes and mask them to the Hyper-V production servers participating in the eight-node cluster:
 - 4 TB LUN (production VM operating system virtual hard disk file format (VHD)).
 - 4 TB LUN (production VM data VHD files).
 - 2.5 TB will be used as Snap reserve space.
4. From AG4, create the following volumes and mask them to the Hyper-V production servers participating in the eight-node cluster:
 - 4 TB LUN (production VM operating system VHD).
 - 4 TB LUN (production VM data VHD files).
 - 2.5 TB will be used as Snap reserve space.
5. From AG5, create volumes to support the offloading of VM snap images from the production disk volumes. This uses SnapVault to make the copies.

Multipath I/O (MPIO)

The IBM N3600 provides fault tolerance with redundant controllers and uses a device-specific MPIO driver to provide balanced I/O across the two redundant controllers. The driver to support this will be loaded in the physical servers making connections to the N3600.

Setup of the x3550 M3 for AD and DNS

This optional server will serve the purpose of providing an active directory and Domain Name System (DNS) environment for your servers. This is required for the servers to participate in a Microsoft cluster and several pieces of the management solution. If this is already available in your environment, then this server is not needed. The setup of the x3550 M3 is straightforward, as there are no optional devices installed in this system.

Before configuring the software, carry out the following steps:

1. Verify the processor and memory configuration for systems.
2. Confirm that one on board port is connected to one of the 1 GbE RJ45 ports on each of the B24X switches and that the switch port is configured for VLAN 50.

The two local disks should be configured (see “Appendix A: Related documents” on page 21) as a RAID-1 set. Complete the following steps:

1. Install Windows Server 2008 Standard Edition.
2. Assign fixed TCP/IP address (VLAN 50).
3. Verify network connectivity.
4. Ensure that the system UEFI code and the integrated management module are updated to the latest supported version.
5. Verify that latest drivers are installed for devices at <http://www.ibm.com/support>. A firmware update guide is referenced in “Appendix A: Related documents” on page 21.
6. Promote to the domain controller.
7. Run Windows Update.

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

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

Setup of the x3650 M3 for management cluster

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

Before configuring the software, carry out the following steps:

1. Verify the processor and memory configuration for the systems.
2. Confirm that two dual-port CNA cards are installed.
3. Confirm that one port from each CNA card is connected to the assigned storage network port (VLAN 10) on each B24X switch.
4. Confirm that one port from each CNA card is connected to the communications network port (VLAN 50, ...) on each B24X switch.

Setup involves installing Windows Server 2008 R2 Datacenter Edition on each server, validating networking and storage connectivity, and enabling Hyper-V and Microsoft clustering. Windows Server Datacenter Edition is suggested because there are an unlimited number of Windows Server 2008 R2 virtual machines that can be run under Datacenter Edition without further licensing. The cost of Datacenter Edition is based on the number of processors installed in the server (sockets, not cores).

The two local disks should be configured (see “Appendix A: Related documents” on page 21) as a RAID1 set. Then do the following:

1. Install Windows Server 2008 R2 Datacenter Edition.
2. Join the domain.
3. Enable Hyper-V.
4. Enable Windows failover clustering.
5. Ensure that the system UEFI code and the integrated management module are updated to the latest supported version.
6. Install the latest Windows updates.
7. Verify that the latest drivers are installed.
 - a. Install the N Series device-specific MPIO driver.
 - b. Install the N Series Appliance Watch Pro driver.
 - c. Install the N Series SnapDrive utility.
 - d. Install the N Series SnapManager® for Hyper-V.
 - e. Install the IBM Systems Director level 1 agent for System Center upward integration.

Physical server storage network

Assign fixed IP addresses to the two CNA connections attached to the iSCSI storage network. Use one port from each CNA, and connect each port to a different B24X switch. No teaming or VLAN assignment is required on the server for these two ports.

These steps are:

1. Validate connectivity for each device.
2. Install the IBM N3600 Device Specific Module (DSM) MPIO driver.

Communication networks

Use the Brocade adapter software in the parent partition to create a fault-tolerant team from the remaining two CNA ports connected to communications network.

The steps are:

1. Using the Brocade teaming software, create a VLAN 10 interface over the FT team to support guest clustering. There should be no IP address for this device on the physical servers, as this will be passed for VM direct connections to the storage only.

Create a Hyper-V virtual switch to be passed to the SQL VMs to support clustering.

2. Using the Brocade teaming software, create a VLAN 50 interface over the FT team to support management communications:
 - a. Assign fixed IP addresses to the physical servers.
 - b. Create a Hyper-V virtual switch to be passed to the VMs to support management operations.

3. Using the Brocade teaming software, create a VLAN 60 interface over the FT team as a cluster private network for use by the two management servers.

Assign fixed IP addresses to the physical servers.

4. Using the Brocade teaming software, create a VLAN 70 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.

5. Validate network connectivity between the physical management servers on VLAN 50 and the VLAN 60 for use in the management cluster creation.

Storage configuration

To properly enable the storage to be isolated from other servers on the network, the unique IQN names are used to mask storage to the management servers. This can be found in the Control Panel, under iSCSI initiator, on the Configuration tab:

1. Record IQN names of each management server, ensuring that the server has already joined the domain, as this might change the IQN name.
2. Use the IBM N3600 management console to add the recorded IQN names to storage volumes reserved for the management cluster:
 - Management cluster CSV volume
 - Management cluster quorum volume
 - SCVMM library volume
3. Using the Windows iSCSI initiator, discover the N3600 storage. There should be a connection from each CNA to each controller, for a total of four connections.
4. Using Disk Manager, validate that storage on N3600 is visible.
5. Create partitions and format on the management cluster quorum and on future management cluster CSV disks.

Cluster creation

To provide a highly available management environment, we create a cluster between the two management servers. The management servers then run as highly available virtual machines on this cluster. The SQL server virtual machines are not highly available virtual machines, but form their own cluster to support SQL in a highly available environment. For more information about setting up a Microsoft cluster under Windows 2008, see “Appendix A: Related documents” on page 21.

To create the cluster, do the following using the Failover Clustering management console:

1. 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.
3. Validate the cluster name and IP address.
4. Validate cluster networking.
5. Validate cluster storage.
6. Enable cluster shared volumes and add the designated volume as a CSV.
7. Set default paths for VM creation to use the cluster shared volume.

SQL server virtual machine

A SQL Server system must be available to support the System Center environments to be used in this configuration. To provide a fault-tolerant environment, a two-node SQL cluster will be created using guest clustering in virtual machines and with the two management servers with Hyper-V installed.

While these two virtual machines already reside on a pair of clustered servers, they should not be enabled as high availability (HA) virtual machines.

The following are required for this configuration:

- ▶ Two non-HA virtual machines, one on each management server
 - Four virtual CPUs
 - 6 GB of RAM
 - Three virtual NICs
 - One client connection (VLAN 50)
 - One cluster private (VLAN 70)
 - One iSCSI (VLAN 10)
- ▶ Windows Server 2008 R2 Enterprise/Datacenter installed in VMs
 - The version of Windows Server 2008 R2 used here might vary according to the version installed on the physical machine.
 - Must be a clusterable version of Windows Server 2008 R2.
- ▶ Storage
 - LUN1: OS VHD on CSV Volume (30 GbE)
 - LUN2: SQL Databases
 - LUN2: SQL Logging
 - LUN4: SQL cluster quorum (1 GB)

Note: The SQL volumes should be masked to the two virtual machines.

The storage being used for the SQL cluster can be allocated from the available storage in the management server RAID group on the N3600, but the volumes should be isolated by using IQN restrictions on the volume.

We suggest using the SQL service account created on the active directory server when installing SQL. Guidance on installing and creating clustered SQL 2008 SP1 systems is located at:

[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>

System Center Operations Manager

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

The HA virtual machines have the following specifications:

- ▶ Two virtual CPUs
- ▶ 4 GB of RAM
- ▶ One vNIC - management network connection (VLAN 50)
- ▶ One OS VHD

To install System Center Operations Manager, perform the following steps:

1. Install Windows Server 2008 R2 on VM.
2. Assign a fixed IP address on VLAN 50.
3. Update Windows.
4. Install SCOM 2007 R2.

The following roles are required by the operations manager:

- ▶ Root management server
- ▶ Reporting server (The database will reside on SQL Server.)
- ▶ Data warehouse (The database will reside on SQL Server.)
- ▶ Operator console
- ▶ Command shell

The following operations management packs should be installed:

- ▶ Virtual Machine Manager (VMM) 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
- ▶ IBM N Series integration pack

System Center Virtual Machine Manager

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

The HA virtual machines have the following specifications:

- ▶ Two virtual CPUs
- ▶ 4 GB of RAM
- ▶ One vNIC - management network connection (VLAN 50)
- ▶ One OS VHD
- ▶ One data VHD or pass-through disk for the VMM library

To install System Center Virtual Machine Manager, perform the following steps:

1. Install Windows Server 2008 R2 on VM.
2. Assign a fixed IP address on VLAN 50.
3. Update Windows.
4. Install System Center Virtual Machine Manager 2008 R2.

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, Virtual Machine Manager will be integrated with Operations Manager. Virtual Machine Manager uses System Center Operations Manager to monitor the health and availability of the virtual machines and virtual machine hosts that Virtual Machine Manager is managing (see “Appendix A: Related documents” on page 21 for 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 will be a highly available virtual machine. Microsoft installation and configuration instructions are in the download link found in “Appendix A: Related documents” on page 21.

The HA virtual machines have the following specifications:

- ▶ Two virtual CPUs
- ▶ 4 GB of RAM
- ▶ One virtual NIC (Management Network Connection (VLAN 50))
- ▶ One OS VHD

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

1. Install Windows Server 2008 R2 on VM.
2. Assign a fixed IP address on VLAN 50.
3. Update Windows.
4. Install IIS.
5. Install System Center Self Service Portal.

System Center Forefront Security

Antivirus and intrusion detection can be provided by the System Center Forefront Security product. The physical host systems and the virtual machines should run Microsoft System Center Forefront Endpoint Protection 2010. This can be managed using the Microsoft Forefront Protection Server Management console. To provide fault tolerance, this will be a highly available virtual machine.

The HA virtual machines have the following specifications:

- ▶ Two virtual CPUs
- ▶ 4 GB of RAM
- ▶ One virtual NIC - Management Network Connection (VLAN 50)
- ▶ One OS VHD

To install System Center Forefront Security, perform the following steps:

1. Install Windows Server 2008 R2 on VM.
2. Assign a fixed IP address on VLAN 50.
3. Update Windows.
4. Install Forefront Protection Server Management Console.

Setup of the x3650 M3 for the production Hyper-V cluster

The production cluster consists of eight IBM x3650 M3 systems with 96 GB of RAM each.

Take the following actions before configuring the software:

1. Verify the processor and memory configuration for the systems.
2. Confirm that two dual-port CNA cards are installed.
3. Confirm that one port from each CNA card is connected to the assigned storage network port (VLAN 10) on each B24X switch.
4. Confirm that one port from each CNA card is connected to the assigned communications network port (VLAN 50, ...) on each B24X switch.

Setup involves installing Windows Server 2008 R2 Datacenter Edition on all servers, validating networking and storage, and enabling Hyper-V and Microsoft Clustering. The two local disks should be configured (see “Appendix A: Related documents” on page 21) as a RAID-1 pair.

Then complete the following steps:

1. Install Windows Server 2008 R2 Datacenter Edition.
2. Join the domain.
3. Enable Hyper-V.
4. Enable Windows failover clustering.
5. Ensure that the system UEFI code and integrated management module are updated to the latest supported version.
6. Install the latest Windows Updates.
7. Verify that the latest drivers are installed. A firmware update guide is referenced in “Appendix A: Related documents” on page 21.
8. Install the IBM N Series device-specific MPIO driver.

9. Install the IBM N Series Appliance Watch Pro driver.
10. Install the IBM N Series SnapDrive utility.
11. Install the IBM N Series SnapManager for Hyper-V.
12. Install the IBM Director level 1 agent for System Center upward integration.

Network configuration

Assign a fixed VLAN 10 IP addresses to the two CNA connections attached to the iSCSI storage network. Use one port from each CNS. Each port should be connected to a different B24X switch. Note that 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 a FT Team from the two remaining CNA ports connected to the communications network:

1. Use the Brocade teaming software to create a VLAN 10 interface from the FT team if iSCSI access is needed for any virtual machines on the production network.
2. Using the Brocade teaming software, create a VLAN 20 interface from the FT team for the cluster private network. Assign a fixed IP address to the physical servers.
3. Using the Brocade teaming software, create a VLAN 30 interface from the FT team for public VM communication. There should be no IP address for this device on the physical servers. Create a Hyper-V virtual switch to be passed to the VMs to support communications.
4. Using the Brocade teaming software, create a VLAN 40 interface on the FT team as a second cluster private network for use by Live Migration. Assign a fixed IP address to the physical servers.
5. Consider setting a Windows QoS (see “Appendix A: Related documents” on page 21) to limit bandwidth to an appropriate level for each network.
6. Using the Brocade teaming software, create a VLAN 50 interface on the FT team to support management communications. Assign a fixed IP address to the physical servers.
7. Validate network connectivity between physical servers on VLAN 20, VLAN 40, and VLAN 50.

Storage configuration

To properly enable the storage to be isolated from other servers on the network, the unique IQN names are used to mask the storage on the management server. You can do this in Control Panel, under iSCSI initiator, on the Configuration tab.

1. Record the IQN names of each host server, ensuring that the server has already joined the domain, as this might change the IQN name.
2. Use the IBM N3600 management console to add the recorded IQN names to the storage volumes reserved for the production cluster:
 - Production cluster CSV volumes
 - Production cluster quorum volume

Reserve the creation of the SnapInfo volume until you have the SnapDrive utility installed and operational. This makes the creation and masking of the this volume much easier.

3. Using the Windows iSCSI Initiator, discover the N3600 storage.
4. Using Disk Manager, validate that storage on N3600 is visible.

5. Create partitions and format the production cluster quorum and future production cluster CSV disks.

Cluster creation

To provide a highly available virtual machine environment, we will create a cluster between the eight production servers. The servers created off of these clustered resources will then run highly available virtual machines on this cluster. For more information about setting up a Microsoft cluster using Windows 2008, see “Appendix A: Related documents” on page 21.

Using the Failover Clustering management console, validate the eight physical production servers as a potential cluster, and address any errors found.

Using the Failover Clustering management console, create a cluster with the eight physical production servers. When this has been completed, perform the following actions:

1. Validate the cluster name and IP address.
2. Validate cluster networking.
3. Validate storage in the 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

With these steps completed, you will have two operational Hyper-V clusters, one for management and one for production virtual machines. Fault tolerance is achieved at several points by using clustering, redundant network switches, and redundant storage controllers for the iSCSI storage.

Appendix A: Related documents

These documents and web pages are useful as references:

- ▶ *IBM x3550 M3 Setup and Configuration Guide*
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/81y6102.pdf
- ▶ *IBM x3650 M3 Setup and Configuration Guide*
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/81y6194.pdf
- ▶ *IBM Firmware Update and Best Practices Guide*
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/firmware_update_best_practices.pdf
- ▶ *IBM ServeRAID M1015 User Guide*
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/ibm_doc_sraidmr_m1015_2ndedition_userguide.pdf
- ▶ *IBM ServeRAID M5014 ad M5015 User Guide*
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/serveraid-m5014_m5015_ug_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 Guide*
[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 (common agent)
<http://www.ibm.com/systems/software/director/downloads/agents.html>
- ▶ IBM Director System Center Integration (download integration modules)
<http://www.ibm.com/systems/software/director/downloads/integration.html>
- ▶ Configuring Operations Manager Integration with Virtual Machine Manager
<http://technet.microsoft.com/en-us/library/cc956099.aspx>
- ▶ 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>
- ▶ IBM Nseries Support and Documentation (registration required)
<http://www.ibm.com/support>
- ▶ NetApp® SnapManager 1.0 for Hyper-V
<http://media.netapp.com/documents/tr-3805.pdf>
- ▶ *Microsoft System Center Operations Manager Deployment Guide*
<http://technet.microsoft.com/en-us/library/bb310604.aspx>
- ▶ *Microsoft System Center Virtual Machine Manager Deployment Guide*
<http://technet.microsoft.com/en-us/library/ee441285.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 on SCOM Integration
ftp://ftp.software.ibm.com/systems/support/system_x_pdf/ibm_hw_mp_mgt_guide.pdf
- ▶ IBM home page for Microsoft Systems Management Integration
<http://ibm.com/support/entry/portal/docdisplay?lnocid=SYST-MANAGE>

Appendix B: Brocade 2-port 10 GbE CNA

The Brocade 2-port 10 GbE CNA for IBM System x is a two-port CNA PCIe 2.0 x8 card that mounts inside supported IBM System x servers. This card combines the functions of a 10 GbE 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: Ensure that you are running driver v2.3 or later to support the configurations required for Fast Track.

Features

Features of the CNA, Ethernet, and FCoE are described in this section.

The CNA supports the following features:

- ▶ PCIe x8 Gen 2 host interface
- ▶ Communication module: Two 400MHz processors ASIC
- ▶ Support for 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

Ethernet features are:

- ▶ 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 (802.1Qbb)
- ▶ Enhanced Transmission Selection (802.1Qaz)
- ▶ Data Center Bridging eXchange Protocol (DCBX)

FCoE features are:

- ▶ 10 Gbps throughput per port full duplex
- ▶ Full hardware-based FCoE protocol offload
- ▶ 500,000 IOPS per port for maximum I/O transfer rates
- ▶ Target rate limiting
- ▶ N_Port ID Virtualization (NPIV)
- ▶ SAN Boot
- ▶ Fabric-based boot LUN discovery
- ▶ Persistent binding
- ▶ Fibre Channel-Security Protocol (FC-SP)
- ▶ FCoE Initialization Protocol (FIP)

Managing converged network adapters (CNAs)

The Brocade CNAs for IBM System x can be managed using either the BCU that is installed with the drivers on the system or the GUI-based Host Connectivity Manager (HCM). BCU is accessible from the Windows command shell and useful to administrators who wish to script activities. However, most configuration can easily be done using HCM, which is what we suggest for setting up the configuration initially. We discuss the HCM in this section.

The 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 on which the CNAs are installed or on a separate management system.

The HCM client uses a discovery mechanism to manage CNAs through the HCM agent. The host name can either be *localhost*, if managing CNAs installed on the same system on which HCM is installed, or the IP address of a remote system. The default user name and password are admin and password. Figure 5 shows the HCM discovery dialog box.

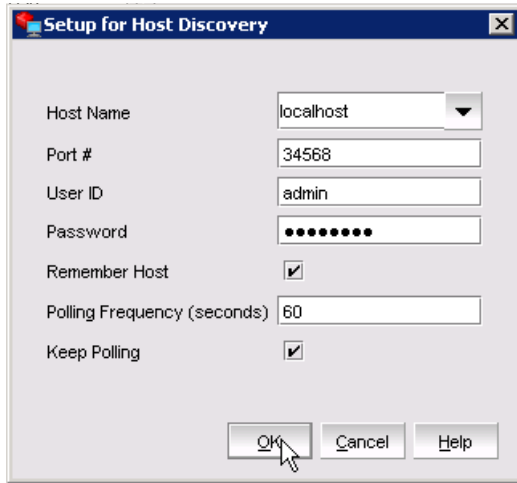


Figure 5 Setup for Host Discovery

Figure 6 shows HCM managing both a Brocade dual-port CNA and a Brocade dual-port HBA.

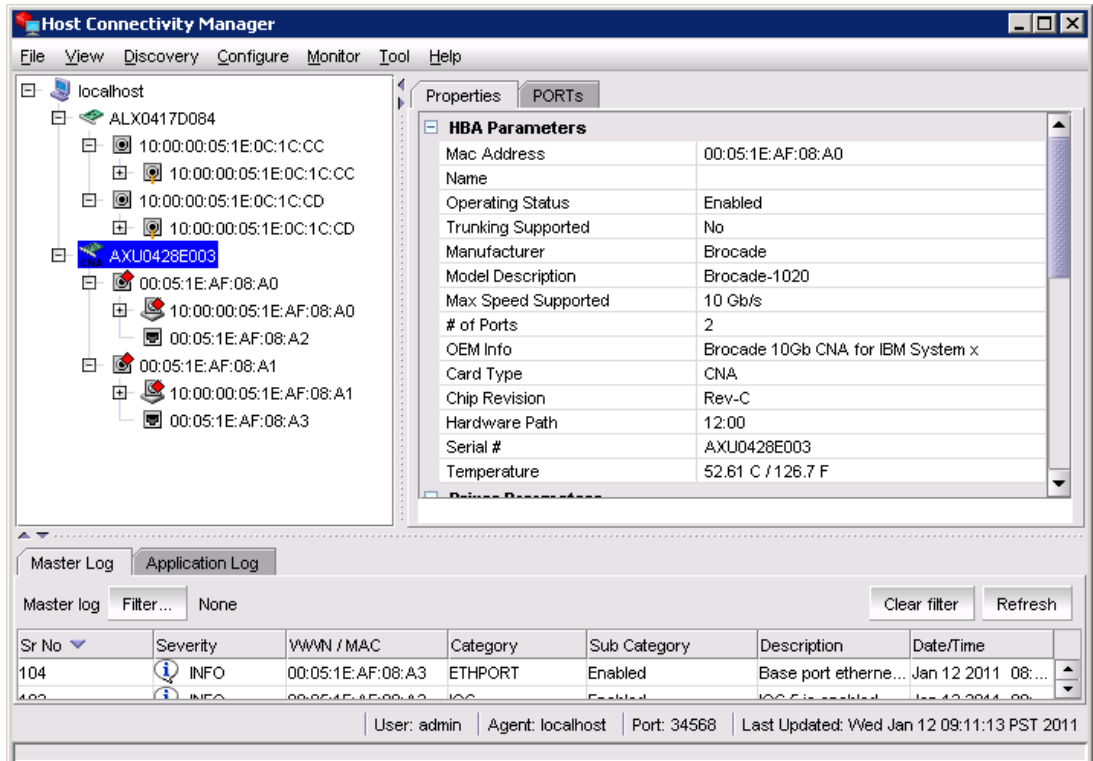


Figure 6 Host Connectivity Manager interface

Configuration within HCM is done by right-clicking elements in the tree and is context-sensitive.

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 out of all teamed interfaces. Two types of teaming methods are available:

- ▶ Active-active

All interfaces in a teamed group are active, and all data is sent out using a round-robin or similar type of algorithm. An active-active connection provides redundancy while providing load balancing and typically uses IEEE 802.3ad Link Aggregation Control Protocol (LACP).

An active-active connection can only be used if all teamed or bonded interfaces have uplinks to the same physical switch or a switch that supports LACP across multiple systems, for example, a switch in a stacked configuration or a switch in running Brocade Multi-Chassis Trunking (MCT), Cisco Virtual Switching System (VSS), or Cisco virtual Port Channel (vPC).

- ▶ Active-passive

Only 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 goes down, a secondary port is chosen to be the next primary.

The active-passive configuration can either be in failback mode or failover mode. In failback mode, if the original port designated as primary goes down and comes back up, it becomes the primary port again. In failover mode, the secondary port that becomes primary stays primary until it fails.

The current Fast Track configuration only supports Active-passive mode. In general, all settings such as MTU size, link speed, port VLAN ID, and so on, 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 using active-active NIC teaming/bonding.

To configure NIC teaming and VLANs over a NIC team, right-click the server level (localhost in Figure 6 on page 24) and select **Teaming**. The Host Connectivity Manager window displays (Figure 7).

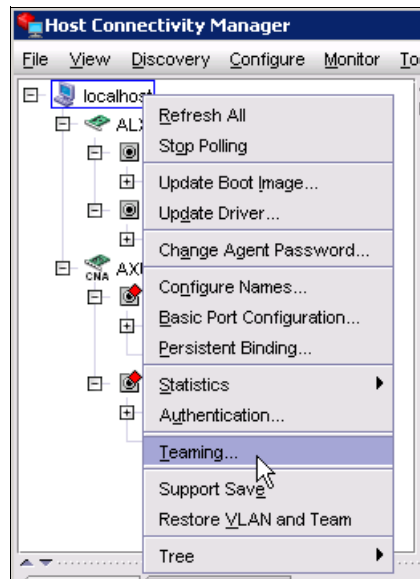


Figure 7 Configuring NIC teaming

In the Teaming Configuration dialogue box that appears (Figure 8), NIC teams can be created, destroyed, or configured.

VLANs for these teams can also be added and deleted. In Figure 8, a team named Test has been created with two ports in failover mode. A pass-through VLAN 0 always exists, as this is the device interface that forwards untagged packets. In addition, a device interface for VLAN ID 10 also has been created for this NIC team. Packets sent out of this device are IEEE 802.1q tagged for VLAN 10.

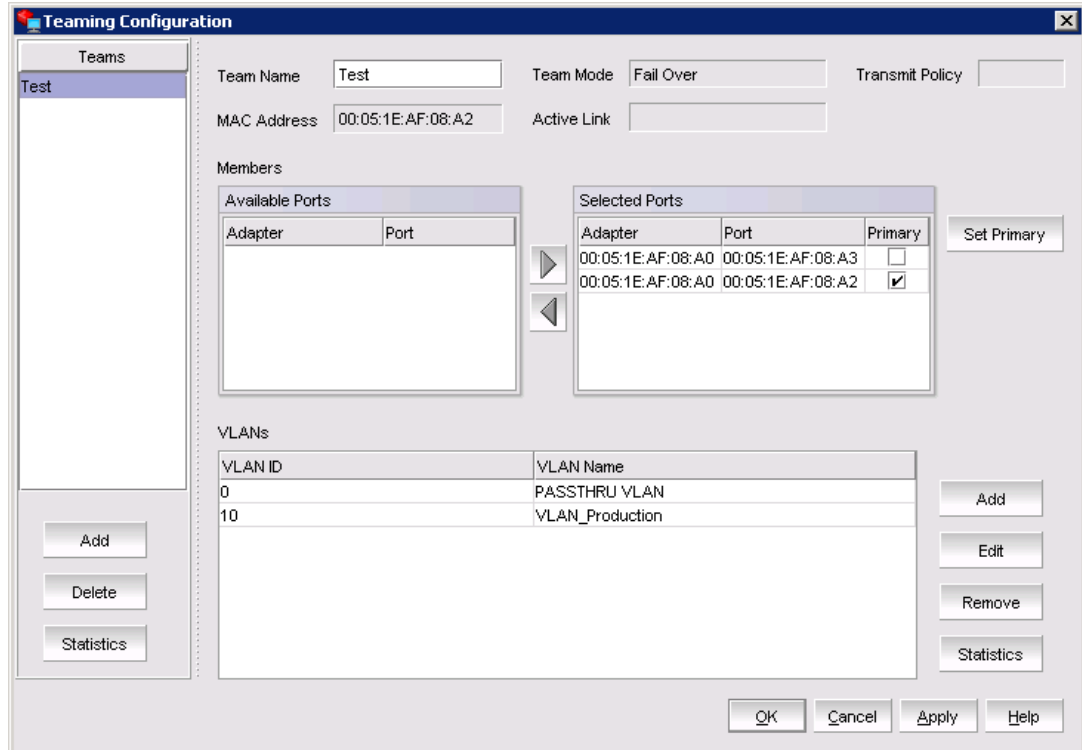


Figure 8 Teaming configuration

Opening the Network Connections window (Figure 9), we see two new device interfaces:

- ▶ Local Area Connection 7, which is TEAM#Test and represents the pass-through VLAN 0
- ▶ Local Area Connection 8, which represents VLAN ID 10 for the NIC team just created

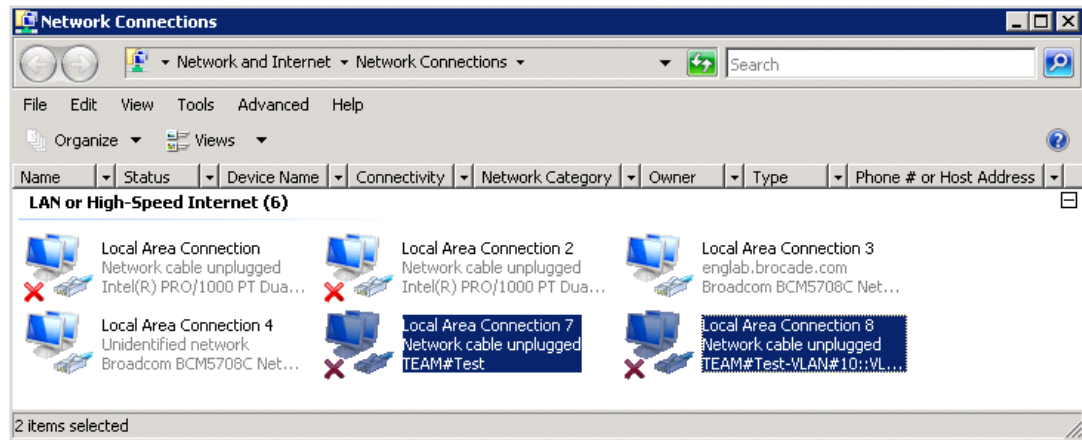


Figure 9 Network Connections in Windows

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 in the Windows Device Manager by right-clicking the Brocade adapter to configure, selecting **Properties**, and displaying the **Advanced** tab (Figure 10).

Note: While VLAN ID can be configured in this menu, we do not suggest doing so in the context of configuring a Microsoft Fast Track environment. Multi-VLAN configuration should be done using either BCU or HCM, rather than Windows Device Manager.

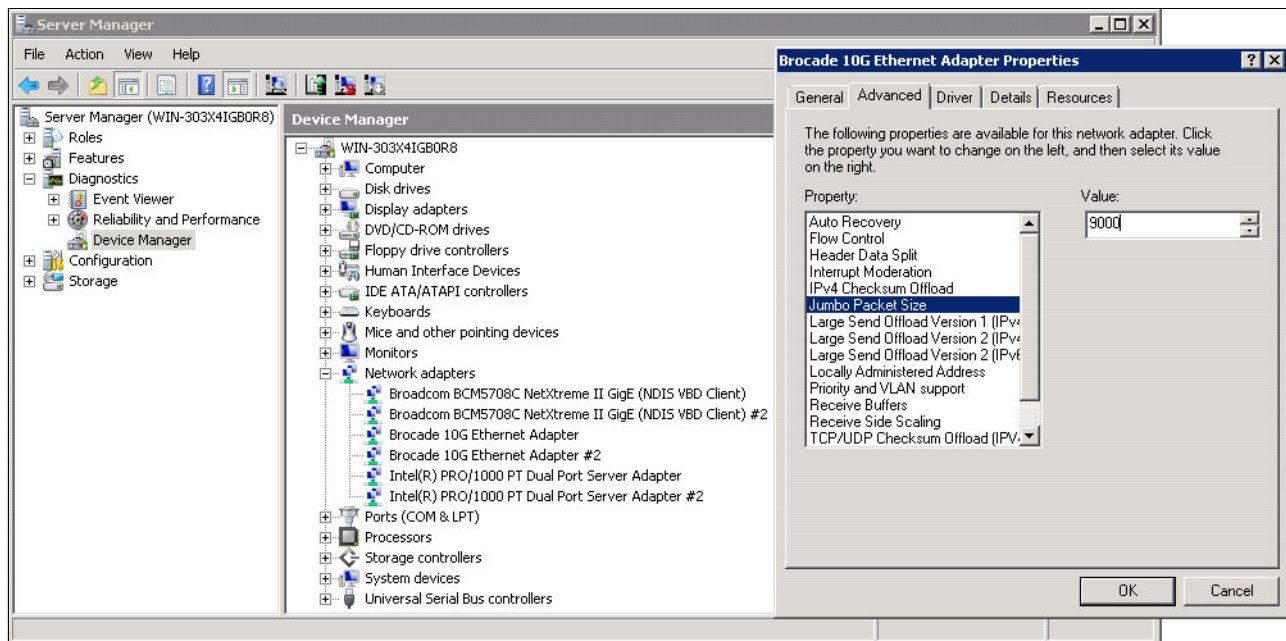


Figure 10 Brocade adapter properties in Device Manager

Related links

See the following resources for more information:

- ▶ *Brocade 2-port 10 GbE CNA for IBM System x At-A-Glance Guide*
<http://www.redbooks.ibm.com/abstracts/tips0718.html?open>
- ▶ Brocade 2-port 10 GbE 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 10 GbE 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, highly available 10 GbE solution that meets mission-critical data center and High Performance Computer (HPC) requirements. It can support 1GbE connections and 10GbE connections, thereby simplifying the migration to an all-10GbE server farm. In the IBM-Microsoft Fast Track configuration, the B24X provides connectivity for a high-speed iSCSI storage network and collapses multiple networks that would normally require dedicated 1 GbE links into a single 10 GbE link, securely separated using VLANs.

Note: Ensure that you are running firmware R04.2.00a or later to support the configurations required for Fast Track.

Features

The IBM Ethernet switch B24X is a fully featured switch with full Layer 2 and Internet Protocol version 4 (IPv4) Layer 3 capabilities. This provides optimal flexibility when connecting out of the Fast Track pod to the greater corporate network. The switch runs a feature-rich and mature IronWare Operating System with an industry-standard command-line interface (CLI), minimizing re-training and operational costs. Supporting a full range of RFC and IEEE standards for interoperability, along with compatibility for vendor-specific protocols such as Cisco Discovery Protocol (CDP) and Per-VLAN Rapid Spanning Tree Plus (PVRST+), the B24X integrates seamlessly into existing network infrastructures.

For a full list of features, visit the IBM Ethernet Switch B24X at-a-glance guide or product page listed in “Related links” on page 31.

Managing the B24X

The IBM Ethernet Switch B24X utilizes an industry-standard CLI for management. This CLI can be accessed over Telnet or SSHv2. For network-wide management and monitoring, use the GUI-based Brocade Network Advisor and Brocade Ironview Network Manager (INM) applications. Fault management and monitoring can also be accomplished using IBM Systems Director software. In addition, the switch supports industry-standard Simple Network Management Protocol (SNMP).

For additional details on configuring the switch, see the related IBM Redbooks® publications included in “Related links” on page 31. The following publications can also be referenced, and are all available for download through the IBM support pages for the switch:

- ▶ *IBM Ethernet Switch B24X Installation and User Guide*
- ▶ *FastIron and Turbolron Configuration Guide*
- ▶ Release Notes for the applicable firmware release on the switch

Fast Track network configuration example

Example 1 is a sample configuration file for an IBM Ethernet switch B24X that is configured as outlined in this IBM-Microsoft Fast Track document. This configuration file is for one of two B24X switches, which are deployed in a redundant fashion. The second B24X would be configured similarly.

Example 1 Sample configuration file

```
Current configuration:
!
ver 04.2.00a
!
!
!
!
vlan 1 name DEFAULT-VLAN by port
!
vlan 10 name iSCSI by port
    tagged ethe 11 to 12
    untagged ethe 1 to 10
!
vlan 20 name Private_Cluster_CSV by port
    tagged ethe 12 ethe 17 to 24
!
vlan 30 name Public_VM_communication by port
    tagged ethe 12 ethe 17 to 24
    router-interface ve 1
!
vlan 40 name VM_Live_Migration by port
    tagged ethe 12 ethe 17 to 24
    router-interface ve 2
!
vlan 50 name Management by port
    tagged ethe 12 ethe 15 to 24 ethe 26
!
vlan 60 name Private_Cluster_Management by port
    tagged ethe 12 ethe 15 to 16
!
vlan 70 name Private_Cluster_SQL by port
    tagged ethe 12 ethe 15 to 16
!
!
!
!
jumbo
hostname Brocade_B24X_Switch_A
ip route 10.0.0.0 255.0.0.0 10.18.228.1
!
router vrrp-extended
interface management 1
    ip address 10.18.228.160 255.255.255.0
!
interface ethernet 1
    port-name iSCSI-Production_Node_1
!
interface ethernet 2
    port-name iSCSI-Production_Node_2
!
interface ethernet 3
```

```

    port-name iSCSI-Production_Node_3
    !
interface ethernet 4
    port-name iSCSI-Production_Node_4
    !
interface ethernet 5
    port-name iSCSI-Production_Node_5
    !
interface ethernet 6
    port-name iSCSI-Production_Node_6
    !
interface ethernet 7
    port-name iSCSI-Production_Node_7
    !
interface ethernet 8
    port-name iSCSI-Production_Node_8
    !
interface ethernet 9
    port-name iSCSI-Mgmt_SQL_Node_1
    !
interface ethernet 10
    port-name iSCSI-Mgmt_SQL_Node_2
    !
interface ethernet 11
    port-name iSCSI_Storage_Array
    !
interface ethernet 12
    port-name To_Other_B24X
    !
interface ethernet 13
    port-name To_CorpNet-1
    !
interface ethernet 14
    port-name To_CorpNet-2
    !
interface ethernet 15
    port-name Comm-Mgmt/SQL_Node_1
    !
interface ethernet 16
    port-name Comm-Mgmt/SQL_Node_2
    !
interface ethernet 17
    port-name Comm-Production_Node_1
    !
interface ethernet 18
    port-name Comm-Production_Node_2
    !
interface ethernet 19
    port-name Comm-Production_Node_3
    !
interface ethernet 20
    port-name Comm-Production_Node_4
    !
interface ethernet 21
    port-name Comm-Production_Node_5
    !
interface ethernet 22
    port-name Comm-Production_Node_6
    !
interface ethernet 23

```

```
    port-name Comm-Production_Node_7
    !
interface ethernet 24
    port-name Comm-Production_Node_8
    !
interface ethernet 25
    port-name Mgmt-AD
    !
interface ethernet 26
    port-name Mgmt-N3600
    !
interface ve 1
    ip address 192.168.30.2 255.255.255.0
    ip vrrp-extended vrid 1
        backup priority 110
        advertise backup
        ip-address 192.168.30.1
        enable
    !
interface ve 2
    ip address 192.168.40.2 255.255.255.0
    ip vrrp-extended vrid 2
        backup priority 110
        advertise backup
        ip-address 192.168.40.1
        enable
    !
    !
    !
    !
    !
    !
    !
    !
end
```

Related links

See the following resources for more information:

- ▶ *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: IP address worksheet

List the IP address assignments in your environment using Table 2.

Table 2 VLAN IP address assignments by network: VLAN 10

VLAN 10 addresses (iSCSI storage)	IP addresses
N3600 Controller-A	
N3600 Controller-B	
Management Server 1 - 1	
Management Server 1 - 2	
Management Server 2 - 1	
Management Server 2 - 2	
SQL Server VM1 (coming off FT team VLAN10)	
SQL Server VM2 (coming off FT team VLAN10)	
Hyper-V Production Server 1 - 1	
Hyper-V Production Server 1 - 2	
Hyper-V Production Server 2 - 1	
Hyper-V Production Server 2 - 2	
Hyper-V Production Server 3 - 1	
Hyper-V Production Server 3 - 2	
Hyper-V Production Server 4 - 1	
Hyper-V Production Server 4 - 2	
Hyper-V Production Server 5 - 1	
Hyper-V Production Server 5 - 2	
Hyper-V Production Server 6 - 1	
Hyper-V Production Server 6 - 2	
Hyper-V Production Server 7 - 1	
Hyper-V Production Server 7 - 2	
Hyper-V Production Server 8 - 1	
Hyper-V Production Server 8 - 2	
Production VMs (if needed off FT team VLAN10))	

Table 3 VLAN IP address assignments by network: VLAN 20

VLAN 20 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 4 VLAN IP address assignments by network: VLAN 30

VLAN 30 addresses (Production)	IP address range
Production VMs as needed	

Table 5 VLAN IP address assignments by network: VLAN 40

VLAN 40 Addresses (Management)	IP addresses
N3600 Controller-A	
N3600 Controller-B	
Management Server 1	
Management Server 2	
Management Cluster	
SQL Server VM1	
SQL Server VM2	
System Center Operations Manager	
System Center Virtual Machine Manager	
System Center Self Service Portal 2.0	
Forefront Security Manager	
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	

VLAN 40 Addresses (Management)	IP addresses
Hyper-V Production Cluster	

Table 6 VLAN IP address assignments by network: VLAN 50

VLAN 50 addresses (Management Cluster Private)	IP addresses
Management Server 1	
Management Server 2	

Table 7 VLAN IP address assignments by network: VLAN 60

VLAN 60 addresses (SQL Cluster Private)	IP addresses
SQL Server VM1	
SQL Server VM2	

Table 8 VLAN IP address assignments by network: VLAN 70

VLAN 70 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	

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
- ▶ Karen Lawrence, IBM Redbooks

- ▶ Linda Robinson, 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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-4732-00 was created or updated on May 13, 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 (logo)  ®	System x®
Redbooks®	ServerProven®	
Redpaper™	System Storage®	

The following terms are trademarks of other companies:

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

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

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