



Marco Rengan
David Ye

IBM System x Private Cloud Offering: Solution and Component Guide

The Private Cloud Offering (PCO) on IBM® System x® with Hyper-V is a ready-to-use, easy-to-manage, high performance, ultra-reliable computing solution for use in data centers and for any business that wants to participate in the benefits of cloud computing. The offering consists of a rack of servers, storage, and networking that can be purchased from a distributor as a complete, deployable unit.

This IBM Redpaper™ document leads you through the process of setup, installation, and configuration to enable quick and easy deployments of a private cloud in a data center. New in this second edition is the Advanced Configuration, an eight-node Hyper-V Fast Track offering based on IBM System x3755 M3 servers.

This paper is intended to provide an overview of the PCO solution for management teams and provide a detailed parts list for the IT teams responsible for integrating and deploying a private cloud infrastructure based on this offering.

Topics covered in this paper are:

- ▶ “Introduction”
- ▶ “Architecture” on page 2
- ▶ “Rack and power infrastructure” on page 6
- ▶ “Host/Compute nodes: Elite and Mainstream” on page 7
- ▶ “Host/Computer node: Advanced” on page 11
- ▶ “Management node” on page 13
- ▶ “Active Directory server” on page 14
- ▶ “Storage” on page 16
- ▶ “Networking” on page 20
- ▶ “Topology and architecture” on page 22
- ▶ “Scale out methodology” on page 23
- ▶ “Operating system software” on page 24
- ▶ “Private cloud operational software” on page 26
- ▶ “Management software” on page 30

Introduction

IBM brings a powerful yet affordable platform that can be rapidly virtualized into multiple secure and completely segregated environments so that each environment acts as though it has complete ownership of the entire platform and all of its resources. The resources incorporated into PCO on System x with Microsoft Hyper-V solutions are leading edge, from Brocade 10 Gb Ethernet network products to the latest Microsoft Self-Service Portal and Microsoft Hyper-V. IBM System x, Microsoft, and Brocade continue the tradition of delivering simple, easy-to-use solutions, incorporating the IBM System Director agent into System Center to provide a *single pane of glass* (SPOG, a single console) for setup, provisioning, management, and administration.

System x PCO with Hyper-V combines high-quality servers and networking hardware with robust tools and software to create an innovative, comprehensive, and usable cloud-computing offering. System x PCO with Hyper-V maximizes client value and performance per square foot of data center space, lowering the barrier of entry for a cloud solution. Therefore, midmarket businesses can reap cloud-computing benefits without a large IT staff or a significant financial investment.

System x PCO with Hyper-V is a pre-tested, end-to-end cloud solution that is ready to deploy and run. It consists of:

- ▶ IBM System x rack servers, which form the reliable foundation
- ▶ Microsoft Hyper-V Cloud Fast Track software, which provides the latest in virtualization and cloud management software
- ▶ IBM Systems Director Upward Integration Module (UIM), which simplifies management by way of Systems Director
- ▶ Brocade top-of-rack switches and adapters, which provide the data center with converged connectivity to storage and networking

System x PCO with Hyper-V alleviates key IT private cloud deployment inhibitors by centralizing computer power, pooling resources, maximizing utilization, and seamlessly responding to and measuring fluctuating demand for server features. Integrated virtualization with the Windows Server 2008 R2 Hyper-V Server enables cost savings by reducing hardware, power, cooling, and space requirements, and increasing operational efficiencies. Capacity needs can be predefined within the integrated management tools, simplifying planning of the private cloud, and VM and infrastructure deployment. The Self-Service Portal software allows for rapid and automated VM deployment.

Hyper-V, the Microsoft hypervisor-based server virtualization technology, and flexible licensing policies make it easy to take advantage of the cost savings of virtualization through Windows Server 2008 R2. Windows Server 2008 R2 with Hyper-V offers a complete, integrated virtualization solution that includes Live Migration and Cluster Shared Volumes.

Architecture

The IBM System x PCO offers an architecture that can be easily sized to fit the needs of the SMB. Each selected configuration is a predefined, robust solution with immediate usability. System x PCO with Hyper-V is a tested and Microsoft Cloud Fast Track architecture certified, end-to-end solution that comprises all parts of a future-ready data center in a modular, extensible, flexible, and scalable rack server format. This format optimizes value in terms of providing dynamically changeable capability based on changing user needs and ease of

deployment, as well as administration, support, expansion, and, most importantly, affordability.

To support business agility, end users must have the ability to add or reduce resources, such as memory and disk storage capacity, when required. IBM System x supports this dynamic scaling of resources and with the help of the Self-Service Portal, customers can autonomously make these requests. System x PCO with Hyper-V does just that with the hardware that is part of this cloud architecture:

- ▶ IBM System x3755 M3, x3550 M3, and x3650 M3 servers are optimized to meet the virtualization management and advanced workload demands of private cloud data centers.
- ▶ Brocade 10 Gbps Ethernet components form the backbone for data and network connectivity. The combination of high-performance adapters and low latency, cut-through switches enables the high-speed infrastructure that is critical for resource utilization and load balancing within the cloud.
- ▶ The IBM N series N6210 Storage System was selected to bring high-performance SAN and NAS features in a unified system with efficiency features like deduplication, thin replication, and no-performance-impact snapshots.

The structural elements of the offering consist of a set of functional groups, which we discuss in the following sections:

- ▶ “Rack and power infrastructure” on page 6
- ▶ “Host/Compute nodes: Elite and Mainstream” on page 7
- ▶ “Host/Computer node: Advanced” on page 11
- ▶ “Management node” on page 13
- ▶ “Active Directory server” on page 14
- ▶ “Storage” on page 16
- ▶ “Networking” on page 20
- ▶ “Topology and architecture” on page 22
- ▶ “Scale out methodology” on page 23
- ▶ “Operating system software” on page 24
- ▶ “Private cloud operational software” on page 26
- ▶ “Management software” on page 30

The IBM System x PCO with Hyper-V structural components have been selected primarily to contribute a combination of performance, reliability, and value. Performance is a key consideration to support the virtualization and elasticity requirements. Reliability and high availability are of paramount importance when offering cloud services. Up time and fault tolerance are among the features that are essential ingredients of a successful deployment. Finally, for businesses of all sizes, but especially for the SMB market, value means offering leading edge technologies at a mainstream price point.

The three configurations in the Private Cloud Offering are:

- ▶ Elite configuration (formerly named SMB Pro)
- ▶ Mainstream configuration
- ▶ Advanced configuration

PCO configurations and components

The set of components consists primarily of software from Microsoft and hardware from IBM. The Elite and Mainstream configurations use IBM System x3650 M3 compute nodes. The Elite configuration is approximately half the compute capacity of the Mainstream configuration. Businesses starting with the Elite configuration can scale up to the Mainstream

configuration in response to an increase in demand. The Advanced configuration uses IBM System x3755 M3 compute nodes for higher performance and support for a larger number of virtual machines.

As shown in Figure 1 (which shows the Mainstream configuration), the hardware building blocks consist of:

- ▶ Networking
- ▶ Management node
- ▶ Compute node
- ▶ Storage subsystem

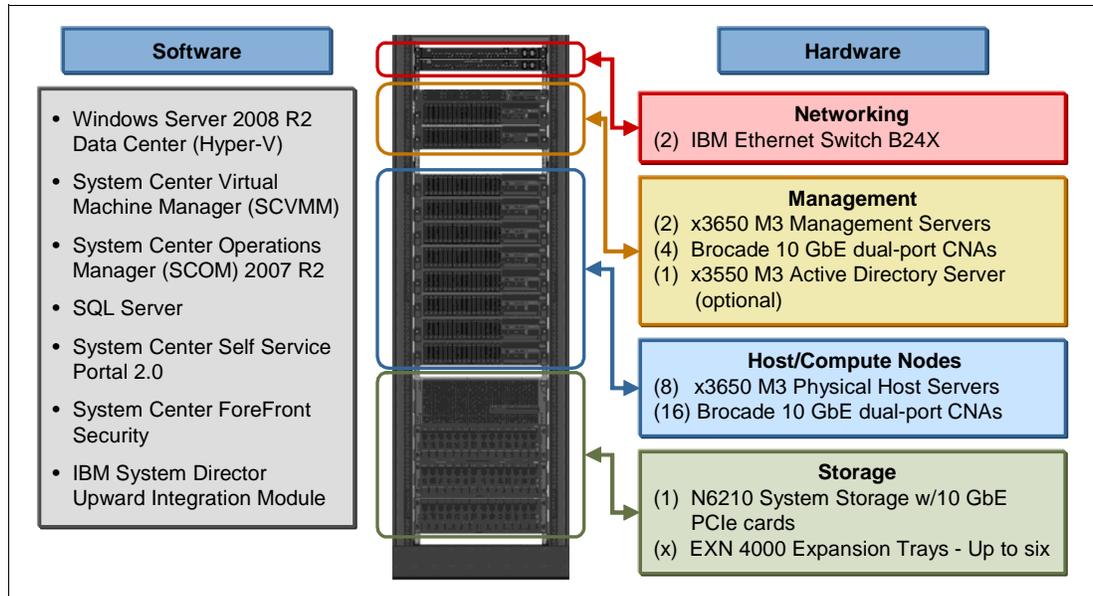


Figure 1 Major components of the System x Private Cloud Offering (Mainstream configuration)

The three main configurations in the IBM System x Private Cloud Offering are as follows:

- ▶ Elite configuration (referred to in other documents as the Base configuration or PCO/ei configuration)

The Elite configuration is a subset of the next level up, Mainstream configuration, which makes for a seamless scaleup for environments that require additional resources, such as VM support or storage.

The configuration consists of the same components as the Mainstream configuration, but with half the number of computer servers and less storage capacity. In addition, the Active Directory server is not required.

- ▶ Mainstream configuration (referred to in other documents as the PCO/mi configuration)

The Mainstream configuration is a super-set of the Elite configuration, with a seamless scaleup of the Elite configuration for environments that require additional resources. The Mainstream configuration supports about twice the number of VMs as does the Elite configuration (depending on the workload, number of users, and other parameters).

The component building blocks of the Mainstream and Elite configurations are the same. However, the Mainstream configuration entails four additional computer servers, one additional Active Directory server, and more SAN storage capacity.

- ▶ Advanced configuration (referred to in other documents as the PCO/za configuration)

The Advanced configuration shares the same Management node structure as the Elite and Mainstream configurations, but uses the x3755 M3 servers as the host node servers.

It supports about four times the number of VMs as the Mainstream configuration. The production node can be added to an existing configuration, but the incremental addition of x3755 M3 servers to an Elite or Mainstream host node is not supported.

As a stand-alone implementation, the Advanced configuration uses the latest N6210 storage subsystem. Figure 2 shows the Advanced configuration.

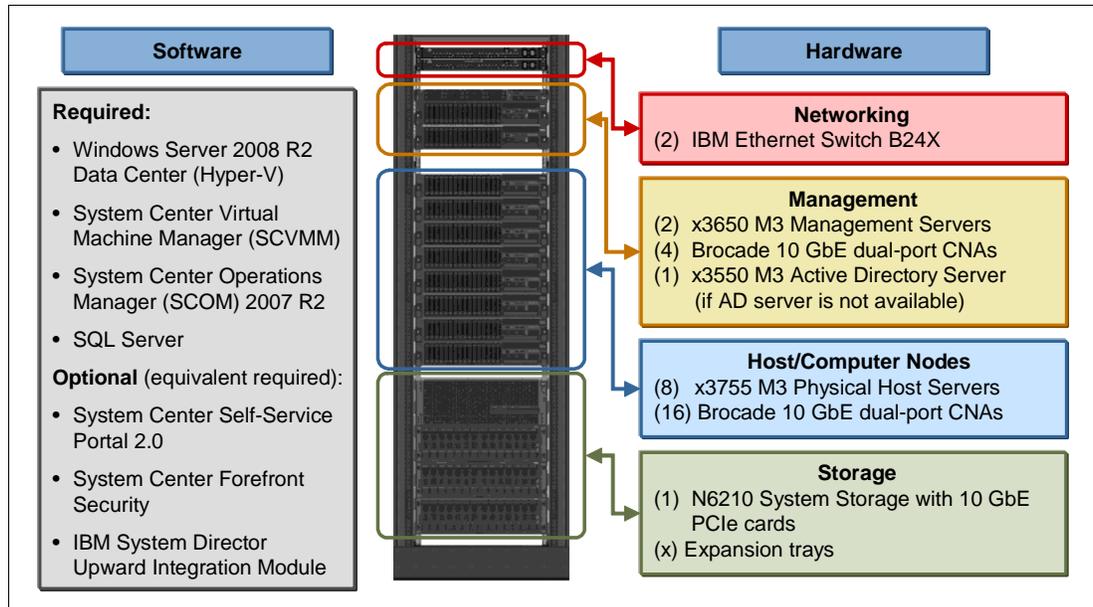


Figure 2 Major components of the System x Private Cloud Offering (Advanced configuration)

Table 1 shows the key differences between the configurations.

Table 1 Virtualization and Private Cloud Offering - feature comparison

| Feature | Elite | Mainstream | Advanced |
|------------------------------|----------|------------|----------|
| Host server family | x3650 M3 | x3650 M3 | x3755 M3 |
| Active Directory server | Optional | x3550 M3 | x3550 M3 |
| Host CPU sockets/server | 2 | 2 | 4 |
| Number of host servers | 4 | 8 | 8 |
| Number of host cluster cores | 24 | 48 | 384 |
| SAN family | N6210 | N6210 | N6210 |
| Range of VMs (estimated) | 90-120 | 100-240 | 200-768 |

Tip: We specify the IBM System Storage® N6210 in Table 1. However, both the N3600 and N6210 have been tested and certified, and are interchangeable in the Elite and Mainstream configurations. The N6210 is the optimal storage platform.

Rack and power infrastructure

Optimized infrastructure equipment is critical to drive improved IT efficiency and availability for the data centers of today and tomorrow. The IBM rack and power infrastructure offerings are custom designed for IBM System x servers and result in:

- ▶ Improved data center efficiency
 - Increased power efficiency
 - Increased space efficiency (avoid over-design!)
 - Lower cost through better data center utilization
- ▶ Improved IT availability
 - Improved uptime
 - Act before downtime impacts business
 - Match utilization, power resources, and capacity planning

In addition, IT availability and efficiency are primary drivers to data center spending:

- ▶ Servers per rack are up 50% since the year 2000
- ▶ Energy consumption is up 20% due to more memory, along with improved utilization due to virtualization
- ▶ Higher power densities at the server and rack levels

In today's online environment, even minutes of downtime can have a significant impact on an organization's operations, customer satisfaction, and financial results, making high-availability an essential feature. The technology fundamentals for today's data center require a solid foundation of rack and power infrastructure that delivers the ability to securely manage and control power resources, servers, and appliances in the data center and across the network. This is imperative to maintain the highest levels of IT availability and drive operational efficiencies.

IBM has announced over 40 new products, refreshing the offerings across the entire rack and power options portfolio, including the following:

- ▶ Three new racks that are 1200 mm deep. This new lineup includes a new 47U tall rack and new 42U versions including a "dynamic" rack that is ship loadable.
- ▶ An IBM lineup of optional universal power supply (UPS) units that includes new rack-mounted and tower units, supporting voltages and configurations not previously available, with new 1500, 2200, 3000, and 6000 volt-ampere (VA) units.
- ▶ A new line of 0U Strip Power Distribution Units, designed for tool-less installation in the new racks. These PDUs have 24 outlets for today's server-dense rack installations.
- ▶ IBM is also offering new Local and Global Console Managers that support unique cabling options ("conversion options") to enable chaining up to 1,024 managed devices that can be managed from a single console.

These offerings are shown in Figure 3.

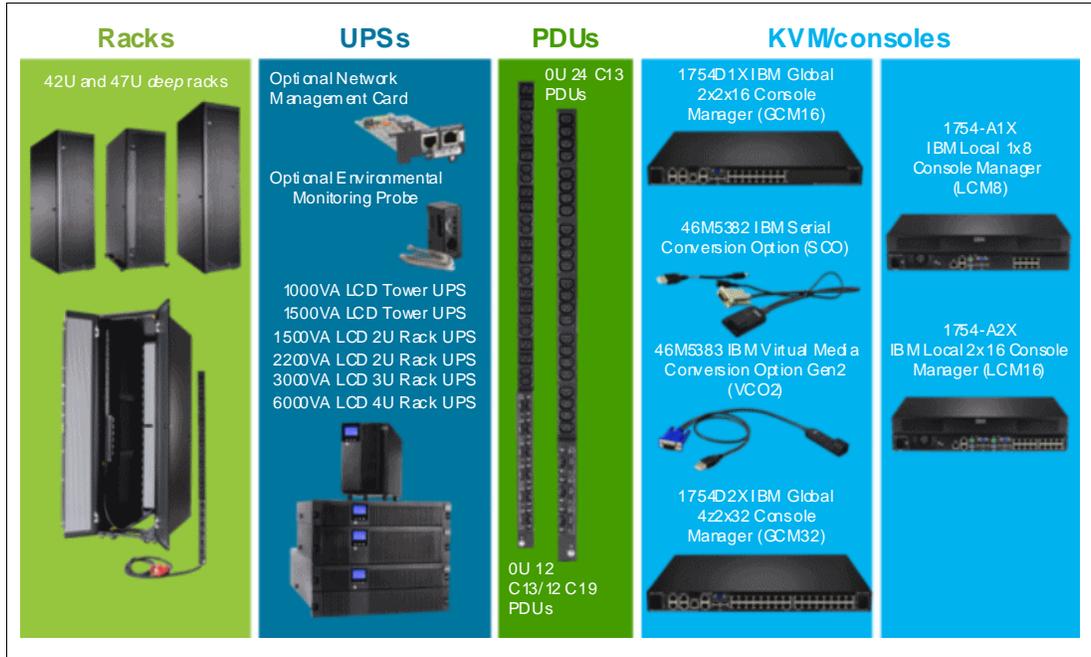


Figure 3 System x rack and power choices

Table 2 shows the rack and power items used in all of these offerings.

Table 2 System x PCO with Hyper-V rack and power parts list

| Part number | Description | Quantity |
|-------------|--|----------|
| 1754D1X | IBM Global 2x2x16 Console Manager | 1 |
| 46M5383 | IBM Virtual Media Conversion Option VCO2 | 11 |
| 172319X | 1U 19 inch Flat Panel Monitor Console Kit with DVD | 1 |
| 40K5372 | IBM Keyboard with Integrated Pointing Device | 1 |
| 53956KX | IBM 6000VA 4U Rack UPS (230V) | 2 |
| 46M4110 | IBM LCD UPS Network Management Card | 2 |
| 46M4004 | IBM 1U 12 C13 Switched & Monitored DPI® PDU | 4 |
| 40K9614 | IBM DPI 30A Cord (NEMA L6-30P) | 4 |
| 93604PX | IBM 42U 1200 mm Deep Dynamic Rack | 1 |

Host/Compute nodes: Elite and Mainstream

The compute nodes are the processing elements for the offering. Virtualization technology in the Hyper-V hypervisor allows each user to see the compute node as a dedicated resource even though it is shared among other users.

The Elite configuration uses four compute nodes whereas the Mainstream configuration has eight compute nodes:

- ▶ Base Elite (PCO/ei) has four x3650 M3 servers
- ▶ Mainstream (PCO/mi) has eight x3650 M3 servers

IBM is a leader in technology and innovation, and has a deep understanding of virtual environments. With substantial investments in green initiatives and energy-smart designs, IBM not only provides high-performing, easy-to-manage servers, but can also help minimize costs on power and cooling.

The IBM System x3650 M3 rack servers address current business requirements while delivering the scalability to manage future requirements. Key capabilities of these servers include performance, reliability, and simplified management. System x servers provide a robust, reliable platform that clients can trust to run their business.

IBM System x3650 M3

Built on the latest six-core Intel Xeon processor 5600 series, the IBM System x3650 M3 two-socket server offers a highly available, high-performance 2U platform for Windows Server 2008 R2 with Hyper-V. With integrated 6 Gbps RAID adapters and double the I/O performance, the x3650 M3 provides a resilient architecture that is ideal for virtualization environments.

The Intel Xeon processor 5600 series expands the benefits of virtualization with innovations that boost performance, increase consolidation ratios, and enable servers of different generations to be combined in the same virtualized server pool. The result can be a more flexible and robust infrastructure with improved load-balancing, virtual machine failover, and disaster recovery capabilities.

IBM System x3650 M3 servers provide a strong foundation for virtualization with large memory and I/O capacity to satisfy the high-speed processing and reliability requirements of running software in a virtualized environment.

The following list provides a quick summary of the x3650 M3 features:

- ▶ Innovative, energy-smart 2U design that helps to lower operational costs.
- ▶ Up to two Intel Xeon 5600 series processors.
- ▶ A maximum of 18 DIMMs of DDR3 memory for up to 192 GB of memory. Chipkill, ECC, and memory mirroring are also supported.
- ▶ Included 6 Gbps SAS RAID adapter.
- ▶ Four PCIe x8 2.0 adapter slots are supported. An optional embedded hypervisor for managing virtual workloads is also supported.
- ▶ Up to 16 hot-swap HDDs with an internal storage capacity of 8.0 TB (using 2.5-inch hot-swap SAS or SATA drives). Alternatively, up to 16 solid-state drives (SSDs) are also available, offering up to 800 GB of storage.
- ▶ Enhanced systems management with an integrated management module and UEFI, and software tools including ToolsCenter, IBM Systems Director, and Active Energy Manager.
- ▶ Best-in-Class RAS (highly redundant design, PFA, LightPath Diagnostics).
- ▶ Open architecture and common tools for quick and simple installation.
- ▶ Better serviceability with a highly scalable, flexible design that is easy to deploy, integrate, service, and manage.

Figure 4 shows the front view of the IBM system x3650 M3 server.

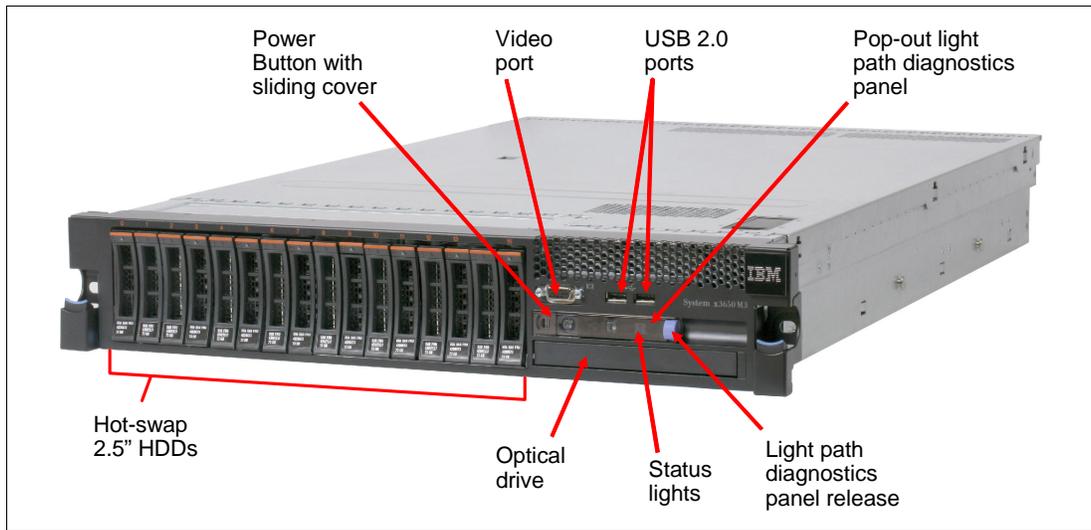


Figure 4 The IBM System x3650 M3

Figure 5 shows the rear view of the x3650 M3 server.

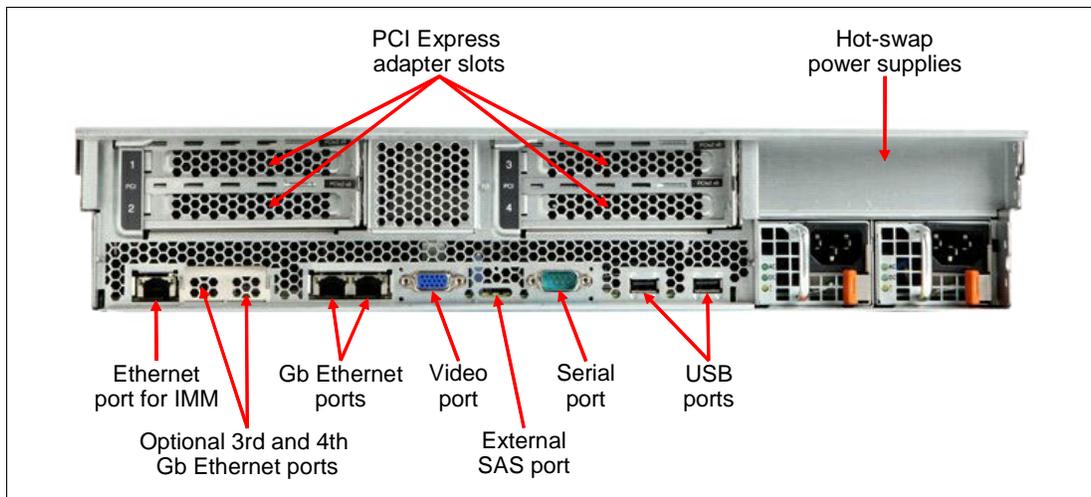


Figure 5 x3650 M3 Rear View

For additional details about the x3650 M3, refer to the *IBM System x3650 M3 Product Guide*, available at the following URL:

<ftp://public.dhe.ibm.com/common/ssi/ecm/en/xso03095usen/XS003095USEN.PDF>

Compute node cluster

The compute node cluster is formed using x3650 M3 servers. The Mainstream configuration uses 8 nodes, and the Elite configurations uses 4 nodes.

Each node in the cluster has the following components:

- ▶ IBM System x3650 M3
- ▶ Two Intel Xeon Processor X5670 6C 2.93 GHz 12 MB Cache 1333 MHz 95w
- ▶ 96 GB RAM (made up of twelve 8-GB DIMMs)

- ▶ One ServeRAID M1015 SAS/SATA Controller (battery not included)
- ▶ Two IBM 73 GB 15K 6 Gbps SAS hot-swap disk drives
- ▶ Two Brocade 10 Gb Dual Port Converged Network Adapters

The Mainstream offering can be ordered using the components list in Table 3.

Table 3 Mainstream configuration parts list

| Part number | Feature | Description | Quantity |
|-------------|---------|--|----------|
| 7945AC1 | None | 3650 M3 Host Config: IBM System x3650 M3 | 8 |
| 46M1072 | 5086 | PCI-Express (2 x8 slots; 1 FH/FL, 1 LP) Riser Card 1 | 8 |
| 49Y1397 | 8921 | 8 GB (1x8GB, 2Rx4, 1.35V) PC3L-8500 CL7 ECC DDR3 1066 MHz LP RDIMM | 96 |
| 46M0901 | 4161 | IBM UltraSlim Enhanced SATA DVD-ROM | 8 |
| 59Y2975 | 3734 | PCI-Express (2 x8 slots; 1 FH/FL, 1 FH/HL) Riser Card 2 | 8 |
| None | 5694 | IBM System x3650 M3 Base with 675 W AC power supply | 16 |
| 42C1820 | 1637 | Brocade 10 Gb Dual Port CNA for IBM System x | 16 |
| 59Y4025 | 7711 | Addl Intel Xeon Processor X5670 6C 2.93 GHz 12 MB Cache 1333 MHz 95w | 8 |
| 42D0672 | 5522 | IBM 73 GB 15K 6 Gbps SAS 2.5 inch SFF Slim-HS HDD | 16 |
| 59Y4039 | 4591 | Intel Xeon Processor X5670 6C 2.93 GHz 12 MB Cache 1333 MHz 95w | 8 |
| 46M2982 | 6311 | 2.8 m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable | 16 |
| 49Y1013 | 2306 | Rack Installation >1U Component | 8 |
| 59Y3846 | 5663 | System Common Planar for 1U/2U | 8 |

The Elite offering can be ordered using the components list in Table 4.

Table 4 Elite configuration - parts list

| Part number | Feature | Description | Quantity |
|-------------|---------|--|----------|
| 7945AC1 | None | 3650 M3 Host Config: IBM System x3650 M3 | 4 |
| 46M1072 | 5086 | PCI-Express (2 x8 slots; 1 FH/FL, 1 LP) Riser Card 1 | 4 |
| 49Y1397 | 8921 | 8 GB (1x8GB, 2Rx4, 1.35V) PC3L-8500 CL7 ECC DDR3 1066 MHz LP RDIMM | 48 |
| 46M0901 | 4161 | IBM UltraSlim Enhanced SATA DVD-ROM | 4 |
| 59Y2975 | 3734 | PCI-Express (2 x8 slots; 1 FH/FL, 1 FH/HL) Riser Card 2 | 4 |
| None | 5694 | IBM System x3650 M3 Base with 675W AC power supply | 8 |
| 42C1820 | 1637 | Brocade 10 Gb Dual Port CNA for IBM System x | 8 |
| 59Y4025 | 7711 | Addl Intel Xeon Processor X5670 6C 2.93 GHz 12 MB Cache 1333 MHz 95w | 4 |
| 42D0672 | 5522 | IBM 73 GB 15K 6 Gbps SAS 2.5 inch SFF Slim-HS HDD | 8 |
| 59Y4039 | 4591 | Intel Xeon Processor X5670 6C 2.93 GHz 12 MB Cache 1333 MHz 95w | 4 |
| 46M2982 | 6311 | 2.8 m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable | 8 |
| 49Y1013 | 2306 | Rack Installation >1U Component | 4 |

Host/Computer node: Advanced

The Advanced configuration has servers with four sockets. In essence, this doubles the number of CPUs on the host node from 16 to 32. Because of this increase in the number of sockets, the Advanced configuration offers an almost doubling of the number of VMs, making for a rich mix of scaleup possibilities in an actively growing data center.

IBM System x3755 M3 (4-socket computer server)

The IBM System x3755 M3 is a 4-socket server that provides outstanding performance and capacity with a data center-friendly, 2U footprint. Based on the 8- or 12-core AMD Opteron 6000 Series platform, the x3755 M3 helps organizations to scale as workload demands increase, accommodating up to 48 processor cores and 512 GB of memory for cost-effective scaling.

The x3755 M3 is an ideal server for business workloads, including database, virtualization, Java, and enterprise applications (such as enterprise resource planning (ERP)). The increased processor density helps to reduce networking complexity and cost for high-performance computing environments, and the available 16 TB of internal storage facilitates data-intensive applications, such as business intelligence.

The following is a quick summary of the x3755 M3 features:

- ▶ Innovative, energy-smart 2U design that helps to lower operational costs
- ▶ Up to 2.5 GHz 12-core or 2.6 GHz 8-core AMD Opteron 6000 Series processors
- ▶ Up to 512 GB DDR-3 RDIMM memory or 128 GB DDR-3 UDIMM memory using 32 DIMM slots
- ▶ 6 Gbps SAS RAID adapter included
- ▶ Four PCIe x8 2.0 adapter slots are supported
- ▶ Up to 8 hot-swap SAS/SATA, and optionally up to 6 simple-swap SATA (up to 16 TB of internal storage)
- ▶ IBM Systems Director, light path diagnostics panel, keyboard, video, and mouse (KVM) switch over IP for remote management

Figure 6 shows the front view of the IBM System x3755 M3 server.



Figure 6 The IBM System x3755 M3

For additional details about the x3755 M3, refer to the *IBM System x3755 M3 Product Guide*, available at the following URL:

<http://public.dhe.ibm.com/common/ssi/ecm/en/xso03119usen/XS003119USEN.PDF>

Details of the host/computer cluster

In this section, the components of the host node cluster are explored as family. A key goal in the architecture and design is ensuring reliability and failover support, which require redundant and high-speed communications between nodes in the cluster.

Each node in the cluster has the following components:

- ▶ IBM System x3755 M3
- ▶ AMD Opteron processor Model 6176 12C 2.3 GHz 12 MB Cache 115w
- ▶ 96 GB RAM (12x 8 GB (1x8 GB, 2Rx4, 1.5V) PC3-10600 CL9 ECC DDR3 1333 MHz LP RDIMM)
- ▶ One ServeRAID M1015 SAS/SATA controller
- ▶ Two IBM 300 GB 15K 6 Gbps SAS 3.5" Hot-Swap HDD
- ▶ Two Brocade 10 Gb dual-port Converged Network Adapters (CNAs)

The Advanced offering can be ordered using the components list in Table 5.

Table 5 Advanced configuration parts list

| Part number | Feature | Description | Quantity |
|-------------|---------|--|----------|
| 7164AC1 | | Server1: IBM System x3755 M3 | 1 |
| | 0095 | ServeRAID M1015 SAS/SATA controller | 1 |
| | 8937 | 8 GB (1x8 GB, 2Rx4, 1.5V) PC3-10600 CL9 ECC DDR3 1333 MHz LP RDIMM | 32 |
| | 5977 | Select Storage devices - no IBM-configured RAID required | 1 |
| | 2593 | 1100W Redundant Power Supply | 2 |
| | 6689 | x3755 M3 Base configuration | 1 |
| 42C1820 | 1637 | Brocade 10 Gb CNA for IBM System x | 2 |
| | A1F6 | Additional AMD Opteron Processor Model 6176 12C 2.3 GHz 12 MB Cache 115w | 2 |
| | 9970 | System Documentation and Software-US English | 1 |
| | 5311 | IBM 300 GB 15K 6 Gbps SAS 3.5" Hot-Swap HDD | 2 |
| | 5716 | Hot swap backplane | 1 |
| | A1F3 | AMD Opteron Processor Model 6176 12C 2.3 GHz 12 MB Cache 115w | 2 |
| | 5718 | System FAN redundant | 5 |
| | 6311 | 2.8 m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable | 3 |
| | 0745 | System Packaging - WW | 1 |
| | 2591 | 2U Bracket for ServeRAID M5015 SAS/SATA Controller | 1 |
| | 8027 | No Optical Selected | 1 |
| | 6443 | Code GBM - x3755 M3 | 1 |
| | 6442 | DVD filler for x3755 M3 | 1 |
| | 6445 | Rail kit | 1 |

| Part number | Feature | Description | Quantity |
|-------------|---------|---------------------------------|----------|
| | 6462 | 3.5" Hot Swap Hard Drive Filler | 6 |
| | 6444 | Cable management assembly | 1 |
| | 9206 | No Preload Specify | 1 |

Management node

The management node for all three configurations consists of a pair of IBM System x3650 M3 rack servers. The management node architecture is designed to provide redundancy and fault tolerance, and is an essential component in managing the cloud and virtualization. Because it serves such a vital role, the management node is configured the same across all three PCO configurations. It manages the complete data center of PCO server computer cluster, regardless of whether the host is made up of Elite, Mainstream, or Advanced clusters.

This management cluster runs the following Microsoft applications:

- ▶ Windows Server 2008 R2 Data Center
- ▶ System Center Virtual Machine Manager (SCVMM)
- ▶ System Center Operations Manager (SCOM) 2007 R2
- ▶ SQL Server
- ▶ System Center Self-Service Portal 2.0
- ▶ System Center Forefront Security

Of particular note is Microsoft System Center Virtual Machine Manager (SCVMM), which manages the overall virtualized infrastructure. It enables increased physical server utilization, centralized management, and rapid provisioning of new VMs. There is no need to rip-and-replace current VMware investments. SCVMM allows you to manage your Hyper-V, virtual server, and VMware environments, all from a centralized platform. This is the essence of SPOG management: A key ingredient to the IBM ease-of-use feature that simplifies management of the data center.

Microsoft System Center Virtual Machine Manager provides:

- ▶ Live migration
- ▶ Clustered shared volumes
- ▶ Hot addition and removal of storage
- ▶ Maintenance mode
- ▶ Expanded support for Internet Small Computer System Interface (iSCSI) SANs
- ▶ SAN migrations in and out of clustered hosts
- ▶ Support for third-party cluster file systems
- ▶ Support for quick storage migration

Management server parts list

The pair of servers constitutes a management cluster configuration that provides reliability, fault tolerance, and peace of mind.

Each of the two management servers has the following components:

- ▶ IBM System x3650 M3
- ▶ Two Intel Xeon X5670 processors: 6C 2.93 GHz 12 MB Cache 1333 MHz 95w
- ▶ 24 GB RAM (composed of six 4-GB DIMMs)
- ▶ One ServeRAID M1015 SAS/SATA Controller (optional battery not included)

- ▶ Two IBM 73 GB 15K 6 Gbps SAS 2.5 inch SFF Slim-HS disk drives
- ▶ Two Brocade 10 Gb Dual Port CNAs

Table 6 provides parts ordering information for the management nodes for the Mainstream and Elite configurations.

Table 6 Management node parts list for the Mainstream and Elite configurations

| Part number | Feature | Description | Quantity |
|-------------|---------|--|----------|
| 7945AC1 | None | 3650 M3 Mgmt Node: IBM System x3650 M3 | 2 |
| 46M1072 | 5086 | PCI-Express (2 x8 slots; 1 FH/FL, 1 LP) Riser Card 1 | 2 |
| 49Y1404 | 3894 | 4 GB (1x4GB, 2Rx4, 1.35V) PC3L-10600 CL9 ECC DDR3 1333MHz LP RDIMM | 12 |
| 59Y3824 | 1745 | IBM System x3650 M3 8 HDD Kit | 2 |
| 46M0901 | 4161 | IBM UltraSlim Enhanced SATA DVD-ROM | 2 |
| 59Y2975 | 3734 | PCI-Express (2 x8 slots; 1 FH/FL, 1 FH/HL) Riser Card 2 | 2 |
| None | 5694 | IBM System x3650 M3 Base with 675W AC power supply | 4 |
| 42C1820 | 1637 | Brocade 10 Gb CNA for IBM System x | 4 |
| 42D0672 | 5522 | IBM 73 GB 15K 6 Gbps SAS 2.5 inch SFF Slim-HS HDD | 4 |
| 59Y4037 | 4589 | Intel Xeon Processor X5650 6C 2.66 GHz 12 MB Cache 1333 MHz 95w | 2 |
| 46M2982 | 6311 | 2.8 m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable | 4 |
| 49Y1013 | 2306 | Rack Installation >1U Component | 1 |
| 59Y3846 | 5663 | System Common Planar for 1U/2U | 2 |

Active Directory server

This section describes the components, configuration specifications, and parts ordering information for Active Directory servers.

Note: The Active Directory server is standard only for the Mainstream and Advanced offerings.

This optional server provides an Active Directory and DNS environment for your servers. This is required for the servers to participate in a Microsoft cluster and to work with several pieces of the management solution. If an Active Directory server is already available in your environment, this server is not needed.

In the Elite configuration, having a physical Active Directory server is optional. The Active Directory server can be virtualized with minimal impact on performance. If additional capacity is required (depending on workload), an Active Directory server can be added, as is done with the Mainstream and Advanced configurations.

The IBM System x3550 M3 builds on the latest Intel Xeon processor technology with extreme processing power and superior energy management and cooling features. With twice the performance of previous generations and a flexible, energy-smart design that integrates low-power components, the x3550 M3 can help you meet demanding workloads at a lower cost per watt.

Figure 7 shows the x3550 M3 server.



Figure 7 IBM System x3550 M3 server

Each Active Directory server in the Mainstream offering is configured with the following specifications:

- ▶ IBM System x3550 M3
- ▶ Two Intel Xeon X5640 processors - 4C 2.66 GHz 12 MB Cache 1066 MHz CPUs
- ▶ 4 GB RAM (configured as one, 4 GB DIMM)
- ▶ ServeRAID M1015 SAS/SATA controller
- ▶ Two IBM 73 GB 15k 6 Gbps SAS 2.5 inch SFF Slim-HS disk drives

Table 7 provides parts ordering information for the Active Directory server.

Table 7 Active Directory server parts list

| Part number | Feature | Description | Quantity |
|-------------|---------|---|----------|
| 7944AC1 | None | IBM System x3550 M3 | 1 |
| 49Y1406 | 8936 | 4 GB (1x4GB, 2Rx4, 1.5V) PC3-10600 CL9 ECC DDR3 1333 MHz LP RDIMM | 1 |
| 59Y3975 | 1791 | IBM System x3550 M3 4 HDD Kit | 1 |
| 59Y2953 | 4375 | PCI-Express (1 x16; FH/HL) Riser Card 2 | 1 |
| None | 5693 | IBM System x3550 M3 Base with 675W AC power supply | 1 |
| 46M1075 | 2100 | IBM 675W Redundant AC Power Supply | 1 |
| 46M0831 | 0095 | ServeRAID M1015 SAS/SATA Controller | 1 |
| 42D0672 | 5522 | IBM 73 GB 15 K 6 Gbps SAS 2.5 inch SFF Slim-HS HDD | 2 |
| 59Y4036 | 4588 | Intel Xeon Processor E5640 4C 2.66 GHz 12 MB Cache 1066 MHz 80w | 1 |
| 46M2982 | 6311 | 2.8 m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable | 2 |
| 59Y2955 | 4039 | IBM Gen-II Slides Kit | 1 |
| 59Y2294 | 4069 | 2.5 inch HDD Filler Bezel | 2 |
| 59Y3894 | 1787 | IBM System x3550 M3 Riser Filler | 1 |
| 59Y3848 | 5777 | System code Group BoM | 1 |
| 59Y3895 | 2134 | IBM System x3550 M3 dual HDD filler | 2 |
| 59Y3846 | 5663 | System Common Planar for 1U/2U | 1 |
| 59Y2938 | 2577 | System Packaging-WW | 1 |
| 49Y1063 | 9206 | No Preload Specify | 1 |

Storage

The IBM System Storage N6210 storage subsystem used in the IBM System x PCO with Hyper-V has key features required to provide performance and reliability in a fault-tolerant and highly available environment in which the System x PCO with Hyper-V operates, including:

- ▶ High availability: Takes advantage of proven features, including a high performance, scalable operating system, data management software, and redundancy features.
- ▶ Simple replication, backup, and recovery: Designed to support disk-based backup, with file-level or application-level recovery using the SnapMgr, Snapshot, and SnapRestore software features.
- ▶ Management simplicity: Self-diagnosing systems designed to enable dynamic provisioning.
- ▶ System Management Integration: Integrates with Microsoft System Center Operations Manager and Self-Service Portal. The Powershell toolkit provides cmdlets to easily manage storage using Powershell scripting automation.
- ▶ Boost performance and cut costs: The optional Flash Cache speeds access to data by caching recently read user data and metadata in the storage controller. With Flash Cache, you can increase I/O throughput by up to 75% and use up to 75% fewer disk drives without compromising performance.

Figure 8 shows the IBM System Storage N6210

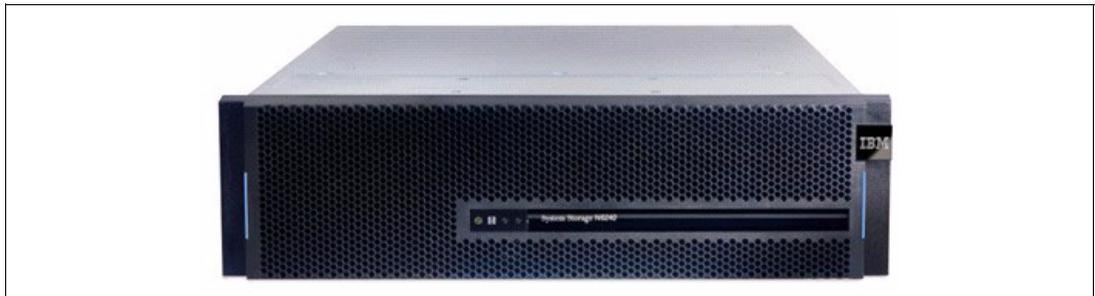


Figure 8 IBM System Storage N6210

The IBM System Storage N6210 is compatible with the entire family of N Series unified storage systems, which features a comprehensive lineup of hardware and software designed to address a variety of possible deployment environments. The N6210 series scales up to 480 TB by supporting up to 240 disk drives.

Storage architecture is a critical design consideration for Hyper-V cloud solutions. This topic is challenging because it is rapidly evolving in terms of new standards, protocols, and implementations. Storage and supporting storage networking are critical to the overall performance of the environment, and also to the overall cost because storage tends to be one of the more costly items.

Storage architectures today have several layers, including one or more storage arrays, the storage network, the storage protocol, and for virtualization, the clustered file system utilizing the physical storage. One of the primary objectives of the Private Cloud solution is to enable rapid provisioning and de-provisioning of VMs. Doing so on a large scale requires tight integration with the storage architecture and robust automation. Provisioning a new VM on an already existing logical unit number (LUN) is a simple operation. However, provisioning a new

LUN, adding it to a host cluster, and so on, are relatively complicated tasks that also greatly benefit from automation.

The N6210 has a modular expansion system that enables the addition of capacity with multiple System Storage EXN3000 expansion trays. Each tray can support 24 additional drives.

Tip: The N6210 is fully backward compatible with the N3600, and this configuration and parts list can be used in lieu of the N3600 and its components.

iSCSI Storage Area Network

The N6210 supports the iSCSI communications protocol using a 10 GbE iSCSI controller. In doing so, it acts as a SAN with resources that can be shared across computer nodes on the same network. This allows for graceful and seamless scaleup of computer resources, independent of the storage unit.

iSCSI is an increasingly popular storage networking architecture that enables the use of the SCSI protocol over TCP/IP network infrastructure. iSCSI enables the use of standard Ethernet networking components, such as network interface cards (NICs), switches, and routers to build a storage area network. Typically, iSCSI SANs are less expensive to implement than traditional Fibre Channel SANs. The storage array used in the iSCSI architecture is usually a low- to mid-range array that is shared by multiple host servers. Redundant, dedicated gigabit Ethernet NICs are preferable for iSCSI connectivity.

As with storage connectivity, high input/output operations per second (IOPS) and low latency are more critical than maximum sustained throughput when it comes to host server sizing and guest performance. When selecting drives, this translates into selecting those with the highest rotational speed and lowest latency possible. Utilizing 15k RPM drives over 10k RPM drives can result in up to 35% more IOPS per drive.

IBM System Storage N6210 parts list

This section lists the parts for the storage subsystem. The storage capacity of the N6210 and the expansion units as used in the System x PCO with Hyper-V configurations are:

- ▶ IBM System Storage N6210 with two controllers
- ▶ System Storage EXN3000 Expansion: Each unit contains 14.4 TB using 24x 600 GB 15K SAS drives

The System x PCO with Hyper-V configurations offer the following storage:

- ▶ Advanced configuration: One N6210 and five EXN3000 units, 72 TB total
- ▶ Mainstream configuration: One N6210 and five EXN3000 units, 72 TB total
- ▶ Elite: One N6210 and one EXN3000 14.4 TB total

N6210 Base configuration

Table 8 provides the parts list for the IBM System Storage N6210 base unit for all of the offerings.

Table 8 N6210 Base configuration

| Product | Description | Qty |
|----------|------------------------------------|-----|
| 2858-C20 | IBM System Storage N6210 Model C20 | 1 |

| Product | Description | Qty |
|---------|-------------------------------|-----|
| 1065 | 2-Port 10-Gbps Ethernet NIC | 2 |
| 1058 | 256-GB Flash Cache (Optional) | 2 |
| 5705 | iSCSI initial | 1 |
| 5710 | Data ONTAP | 1 |
| 5711 | SnapRestore | 1 |
| 8050 | SnapDrive for Windows | 1 |
| 8135 | SnapMgr Hyper-V - Host | 1 |
| 9000 | PDU Power Cord, 42U Rack | 1 |
| 9202 | Field install rack mount kit | 1 |
| 9560 | Dual-path FC cabling | 1 |

The N6210 licenses are shown in Table 9. They include features discussed earlier such as SnapDrive for Windows.

Table 9 N6210 Licenses

| Product | Description | Qty |
|----------|--|-----|
| 2870-58E | IBM System Storage N6210 58E Functions Authorization | 1 |
| 5700 | Data ONTAP Essentials | 1 |
| 5705 | iSCSI initial | 1 |
| 5710 | Data ONTAP | 1 |
| 5711 | SnapRestore | 1 |
| 8050 | SnapDrive for Windows | 1 |
| 8135 | SnapMgr Hyper-V - Host | 1 |

EXN3000 expansion unit

The EXN3000 expansion unit is included in the Advanced and Mainstream offerings. It is optional for Elite. Table 10 provides the product offerings for each EXN3000 expansion unit.

The Mainstream configuration has five expansion units. It can be can be configured with additional expansion units, allowing for future growth.

Table 10 EXN3000 expansion unit parts list

| Product | Description | Qty |
|----------|--|-----|
| 2857-003 | IBM System Storage EXN3000 SAS/SATA Expansion Unit | 1 |
| 1102 | SAS disk power supply | 1 |
| 2054 | 5.0 m SAS QSFP-QSFP Cable | 2 |
| 2065 | 5.0 m SAS RJ-45 ACP Cable | 2 |
| 4017 | 600 GB, 15K RPM SAS HDD | 24 |
| 9000 | PDU Power Cord, 42U Rack | 2 |

| Product | Description | Qty |
|----------|--|-----|
| 9202 | Field install rack mount kit | 1 |
| 2857-003 | IBM System Storage EXN3000 SAS/SATA Expansion Unit | 4 |
| 1102 | SAS disk power supply | 4 |
| 2053 | 2.0 m SAS QSFP-QSFP Cable | 4 |
| 2064 | 2.0 m SAS RJ-45 ACP Cable | 4 |
| 4017 | 600 GB, 15K RPM SAS HDD | 96 |
| 9000 | PDU Power Cord, 42U Rack | 8 |
| 9202 | Field install rack mount kit | 4 |

N6210 software licenses – All offerings

Table 11 shows the software licensing options for the N3600 Storage System.

Table 11 N3600 software licenses

| Product | Description | Qty |
|----------|---|-----|
| 5773-NAS | Support Line for Network Attached Storage (N Series) 3 yr | 1 |
| 1206 | Per Storage Device Support 3 yr 7X24 | 1 |
| 5809 | IBM BP CustLtr SL-Storage | 1 |

Services: All offerings

Table 12 provides the services offerings for the N6210 Storage System.

Table 12 Service offerings

| Product | Description | Qty |
|----------|----------------------------------|-----|
| 6942-25B | Warranty Service Upgrade | 1 |
| 0050 | Request Proposal for Prices | 1 |
| 0105 | 2857-003 24x7x4 WSU | 5 |
| 4524 | WSU Same Day 24x7x4 Response | 1 |
| 6778 | 2858-C20 24x7x4 WSU | 1 |
| | | |
| 6942-87M | On-Site Problem Determination | 1 |
| 0050 | Request Proposal for Prices | 1 |
| 8011 | On-Site Problem Determination | 1 |
| | | |
| 6942-CON | Contract Coverage (N/C codes) | 1 |
| 0096 | ServiceElite | 1 |
| | | |
| 6942-HMA | Hardware Maintenance (N/C codes) | 1 |

| Product | Description | Qty |
|----------|-------------------------------|-----|
| 0025 | Hardware Maintenance Selected | 1 |
| 0050 | Request Proposal for Prices | 1 |
| | | |
| 6942-TCS | Terms (N/C codes) | 1 |
| 0050 | Request Proposal for Prices | 1 |
| 0053 | Prepay Billing 3 Year Term | 1 |

Disk usage

The storage subsystem used by System x PCO with Hyper-V will need to support the 2-node management cluster and the Hyper-V production cluster (8-node cluster for the Mainstream and Advanced configuration). Storage I/O can vary widely depending on the applications being run. It is a good practice to profile your storage needs to ensure that adequate space and I/O capacity are available.

Clustered Hyper-V servers can use storage in several ways:

- ▶ The primary use of storage is to present an iSCSI LUN to the host servers. This LUN would be available to store VM files being run on the server.
- ▶ VHD files are the virtual hard disks used by VMs to contain the operating system and data files under typical circumstances. Use of virtual hard disks allows maximum flexibility in terms of VM migration and backup options. These VHD files would reside on an iSCSI LUN that has been made available to the host server.
- ▶ Clustered Shared Volumes (CSV) is the technology available to clustered Hyper-V servers to allow simultaneous access to LUNs from all servers in the cluster for the purpose of hosting VM files. This can greatly simplify your storage configuration by providing a common namespace for all Hyper-V servers in a cluster.
- ▶ Certain circumstances might require a VM to have direct access to a storage LUN. Examples of this are the use of guest clustering in the VM directly, and disk I/O-intensive applications. A pass-through disk is taken offline from the server, and the VM can take direct control of this.

Networking

Using a solution built around iSCSI SANs with high bandwidth 10 Gbps Ethernet makes advanced networking and storage architectures accessible to the value-conscious installation. In order to create a high bandwidth 10 Gbps Ethernet backbone, the offering uses a pair of IBM Ethernet Switch B24X top of rack switches from Brocade and 10 GbE Brocade CNAs working in tandem to provide a high performance Ethernet network.

Figure 9 shows the IBM Ethernet switch B24X.

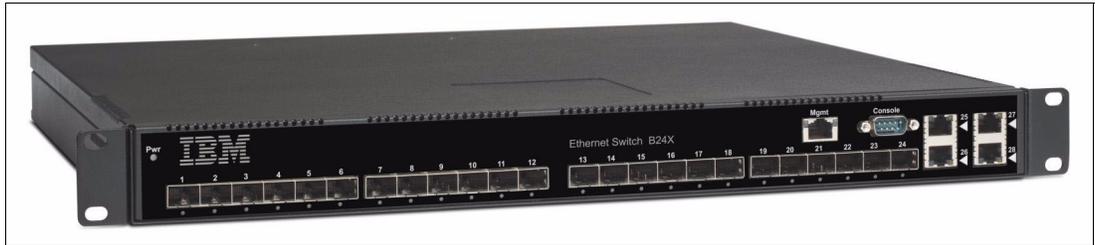


Figure 9 IBM Ethernet switch B24X

The IBM Ethernet switch B24X is a compact, high-performance, highly available 10 GbE solution in a power-efficient package. Features of the switch include:

- ▶ 24x 10/1 GbE dual-speed SFP+ ports plus four 10/100/1000 MbE RJ-45 ports
- ▶ Full Layer 2 and Layer 3 capabilities with jumbo frame support
- ▶ High performance, 488 Gbps line speed, full-duplex throughput with low latencies
- ▶ Dual redundant, hot-swappable power supplies and highly available fan unit
- ▶ Mature and consistent IronWare OS found across all IBM b-type system networking devices, providing a common, industry-standard command line interface (CLI)
- ▶ Upwards integration into Systems Director Network Control with inventory and monitoring capabilities

Two dual-port, 10 GbE CNAs, which act as the NICs, are used to provide a highly available, redundant communications pathway and to physically separate iSCSI storage traffic from other traffic.

Figure 10 shows the Brocade dual-port GbE CNA.

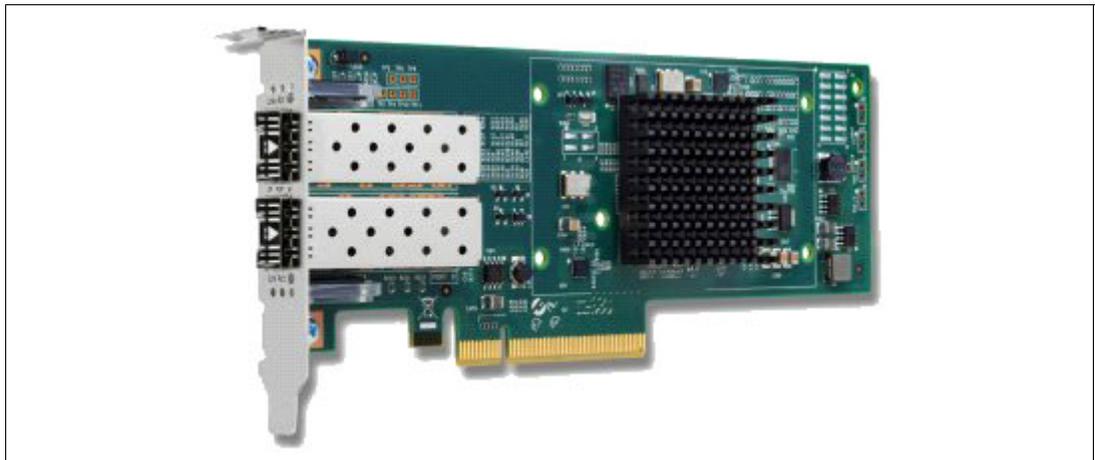


Figure 10 Brocade dual-port 10 GbE CNA for IBM System x

Features of the CNA include:

- ▶ PCIe x8 Gen 2 host interface
- ▶ Industry leading, 500k IOPS performance per port
- ▶ NIC teaming and multi-VLAN capabilities integrated in Windows NDIS stack
- ▶ Full Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) capabilities
- ▶ Jumbo frame support

- ▶ TCP Segmentation Offload (TSO), Large Send Offload v1/v2 (LSO), Header Data Split (HDS), and Receive Side Scaling (RSS) offload capabilities

Topology and architecture

The architecture of the offering consists of the following structures:

- ▶ Computer node
- ▶ Management node
- ▶ Storage subsystem
- ▶ Networking

System x PCO with Hyper-V structures these components so that configurations can easily build on common elements within the architecture. The building block architecture is facilitated by the use of the high-speed 10 Gb Ethernet backbone created by a pair of redundant IBM B24X top of rack switches on which intranet, network, and connectivity to the iSCSI SAN reside.

Two 10 GbE ports from each server are used to provide a high speed, isolated path for iSCSI traffic over a segregated VLAN. Multipath I/O is used to load balance the traffic and provide fault tolerance. The remaining two ports are teamed together using Brocade software. NIC teaming virtualizes multiple physical ports into a single logical interface and provides fault tolerance in this solution. Additional logical interfaces can be created using advanced multi-VLAN technology, thereby allowing for the separation of Hyper-V clustering, Live Migration, Management, and VM Access traffic, which are all unified over a single logical link.

Figure 11 shows a high-level illustration of the network topology.

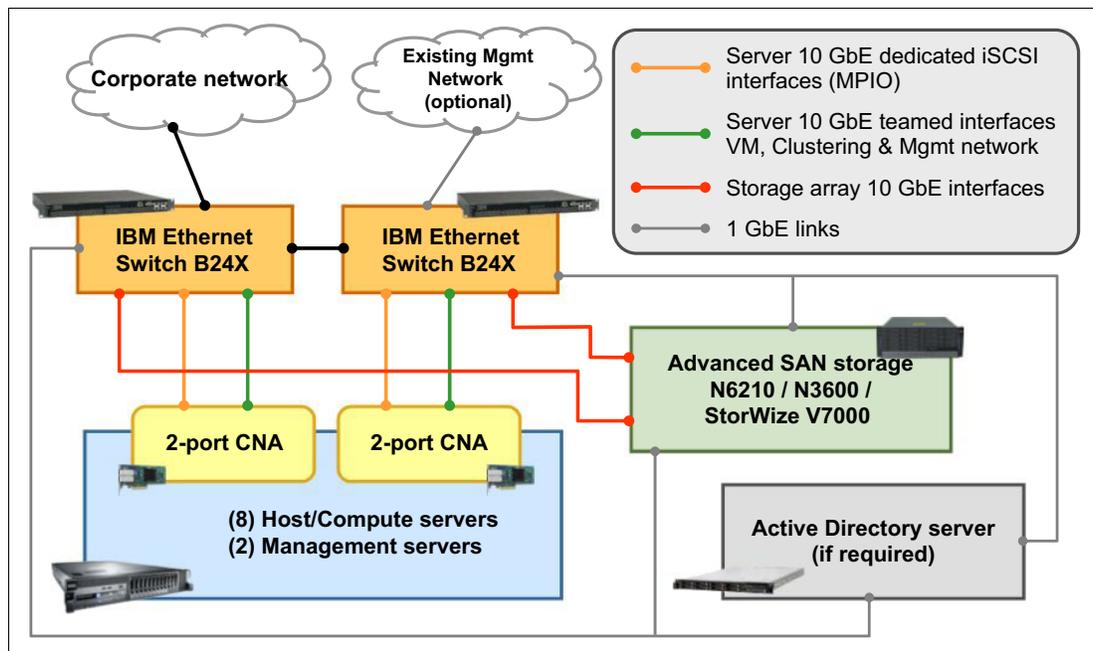


Figure 11 High-level network topology

Brocade networking components

Table 13 provides parts ordering information for Brocade networking components.

Tip: The CNAs are ordered as part of the servers.

Table 13 Brocade networking components

| Part number | Description | Quantity |
|-------------|---|----------|
| 0563012 | IBM Ethernet Switch B24X | 2 |
| 69Y0373 | 300 W AC Power Supply (redundant) | 2 |
| 45W2398 | 10 GbE SFP+, Active Direct Attached Cables, 1 m | 6 |
| 45W2408 | 10 GbE SFP+, Active Direct Attached Cables, 3 m | 32 |
| 69Y0389 | 10 GbE SFP+, 10GBASE-SR, 300 m, multimode LC | 6 |

Scale out methodology

The System x PCO architecture is designed to accommodate growth. This growth will most typically be for one or more of the following reasons:

- ▶ Need for additional VMs to accommodate more users or workloads
- ▶ Need to isolate specific hardware elements from a group of users

Depending on the starting point, the scale-out approach will be one of the following approaches:

- ▶ **Elite:** Additional x3550 M3 or x3560 M3 servers can be added to the production cluster. The management cluster can be complemented with an Active Directory server based on an additional x3550 M3 server, if one was not included at the outset, because this is an optional element in the Elite. In addition, depending on the extent of the scale out, expansion trays can be added to the storage sub-system to increase total storage.
- ▶ **Mainstream:** Scale out in this case is simply a matter of adding more x3550 M3 or x3560 M3 servers to the production cluster.
- ▶ **Advanced:** Here the scale-out will require the addition of x3755 M3 servers.

In either case, the architecture allows for the same management servers to administer the additional host clusters that are added while providing the required security or performance driven by customer or business needs. This is shown in Figure 12.

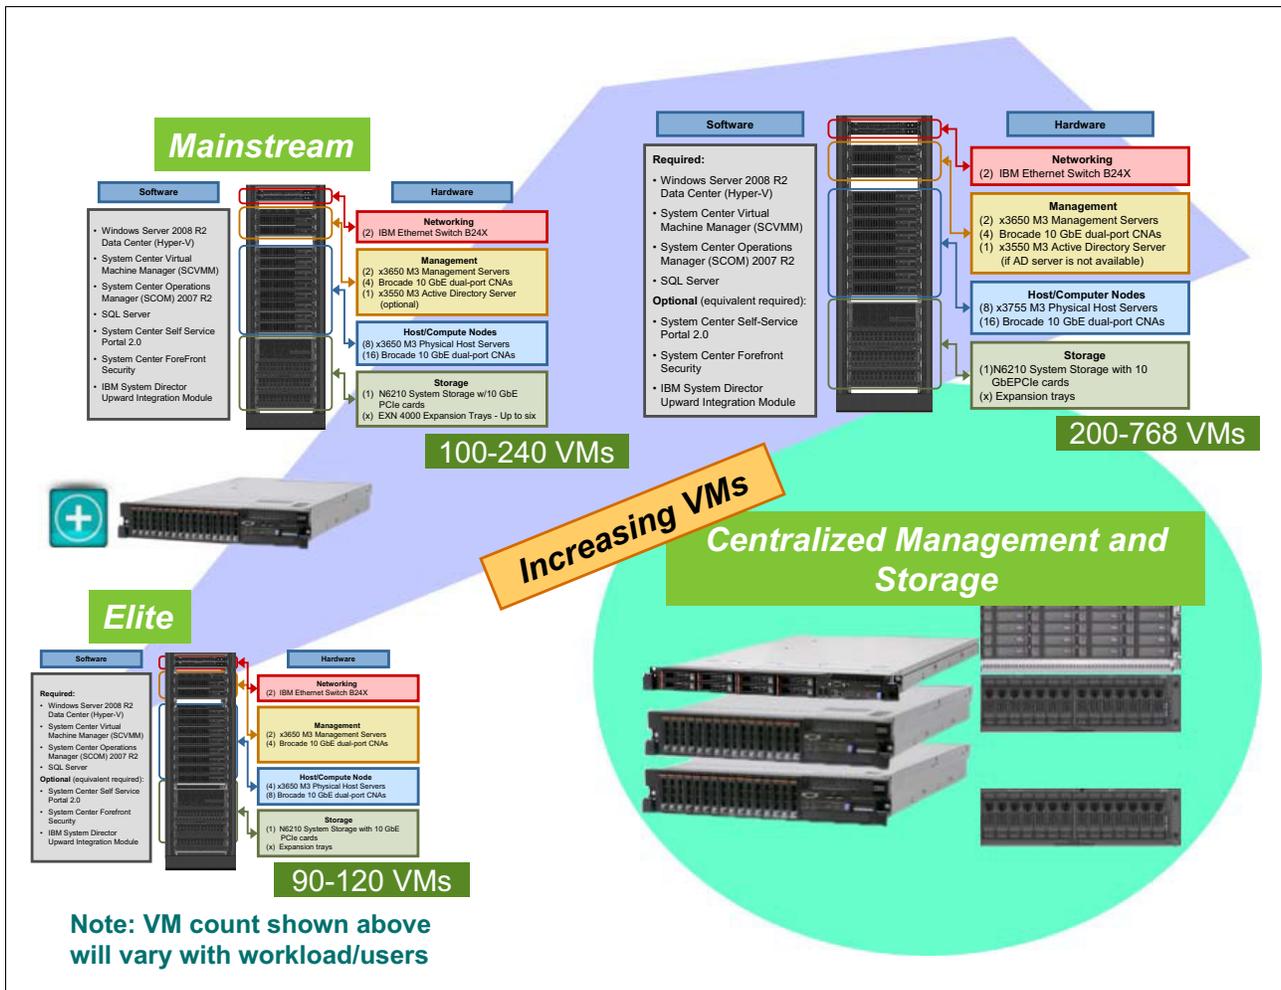


Figure 12 Scale-out of production clusters

There are additional considerations that should be taken into account when scaling-out. In general, when starting with Intel-based servers, continue with the Intel processor-based server family. Similar thinking applies to the x3755 M3 AMD-processor-based server family. Regarding the physical attributes of the rack and the number of ports in the top-of-rack switch, scaling out requires additional cables, switches, racks, and rack components when their initial capacity is exceeded.

Operating system software

The host servers are one of the critical components of a dynamic, virtual infrastructure. The host servers, running Windows Server 2008 R2 with Hyper-V technology, provide the foundation for running VM guests and also provide the management interface between the guests and Microsoft System Center Virtual Machine Manager.

In this section, we describe a detailed host server design and sizing methodology, and present a set of reference server architectures. The reference server architectures are

intended to be a starting point for the design process and provide a foundation for documenting the final design.

Microsoft Windows Server 2008 R2 with Hyper-V

In the IBM System x PCO, the server infrastructure is distributed among three groups:

- ▶ Computer (host) node
- ▶ Management node
- ▶ Directory node (Active Directory)

The underlying operating system for each of these nodes is Microsoft Windows Server 2008 R2. The preferred edition (Standard, Enterprise, or Data Center) depends on which type of node is under consideration. Because the computer node hosts the virtual guests that reflect the users, a multiple guest license is suggested. Similarly, because the Management node is the control point for the computer node, the best license will depend on how many computer node servers (and CPUs) are in the environment. Finally, because the Directory node works fairly independently, it does not need an elaborate license.

The software components used in System x PCO with Hyper-V are as follows:

- ▶ Windows Server 2008 R2 with Hyper-V (x64)
- ▶ System Center Virtual Machine Manager (SCVMM)
- ▶ System Center Operations Manager 2007 R2 (SCOM)
- ▶ SQL Server
- ▶ System Center Self-Service Portal 2.0
- ▶ System Center Forefront Security
- ▶ IBM Systems Director Upward Integration Module

The essential software components in the System x PCO with Hyper-V solution are:

- ▶ Windows Server 2008 R2 with Hyper-V
- ▶ System Center Virtual Machine Manager (SCVMM)
- ▶ System Center Operations Manager (SCOM) 2007 R2

Windows Server 2008

As we mentioned previously, the various editions (licenses) of Windows Server 2008 allow support for a varying number of users, and are designed to optimize scalability of the infrastructure. There are three groupings of the OS, which are:

- ▶ Standard Edition
- ▶ Enterprise Edition
- ▶ Datacenter Edition

Each of these groupings provides a tiered level of *virtualization use rights*. Each group allows for a specific number of Windows-based VMs:

- ▶ Windows Server 2008 R2 Standard edition includes one running VM
- ▶ Windows Server 2008 R2 Enterprise Edition includes up to four VMs
- ▶ Windows Server 2008 R2 Datacenter Edition includes unlimited virtualization rights

The Enterprise Edition license allows up to four Windows guests. For more than four, additional Windows Server licenses are required for the additional VMs.

The license option you should use for PCO is Windows Server 2008 R2 Datacenter Edition because it includes unlimited virtualization use rights and allows as many guests as are needed on the physical server.

Virtual machine guests with Hyper-V

Hyper-V greatly increases the scalability of guest VMs. As a consequence, there is a much larger set of viable workloads that can be consolidated, including multi-processor, multi-core servers, servers with large disk or IO requirements, and so on.

In general, it is prudent to configure each guest with no more than the resources needed. This is nominal guidance that applies to most resource allocation scenarios to ensure equitable distribution for other guests and future growth.

Private cloud operational software

Microsoft System Center is the central operational and control element that completes the cloud and virtualization management. Microsoft Systems Center has two key components:

- ▶ System Center Virtual Machine Manager (SCVMM)
- ▶ System Center Operations Manager (SCOM)

For detailed guidance on how to install and configure System Center Virtual Machine Manager 2008 R2, refer to:

<http://technet.microsoft.com/en-us/systemcenter/vmm/default.aspx>

Host server parts list - Windows Server 2008 R2 Datacenter Edition 2-CPU

Table 14 provides parts ordering information for the Windows Server 2008 R2 Datacenter Edition, licensed for two CPUs. Quantities to order are:

- ▶ Mainstream offering: 8
- ▶ Elite offering: 4
- ▶ Advanced offering: 8

Table 14 Windows Server 2008 R2 parts list

| IBM part number | Geography | Description |
|-----------------|--------------|--|
| 4849-MGN | World-wide | Windows Server 2008 R2 Data Center (2 CPU, 5 CAL) ROK - Multilang (BR,EN,FR,SP) US, Canada and LA |
| 4849-MGM | EMEA/APK/APJ | Windows Server 2008 R2 Data Center (2 CPU, 5 CAL) ROK - Multilang (CS,EN,FR,DE,IT,JA,KO,PL,RU,SP,TR) |
| 4849-MGX | APGCG/ASEAN | Windows Server 2008 R2 Data Center (2 CPU, 5 CAL) ROK - Chinese Simplified |
| 4849-MGT | APGCG/ASEAN | Windows Server 2008 R2 Data Center (2 CPU, 5 CAL) ROK - Chinese Traditional |

Management Server parts list – Windows Server 2008 R2 Enterprise 10 CAL

Table 15 on page 27 provides parts ordering information for the Windows Server 2008 R2 Enterprise 10 CAL. With two x3650 M3 servers, each with 2 CPUs, the quantities to order are:

- ▶ Advanced offering: 2
- ▶ Mainstream offering: 2
- ▶ Elite offering: 2

Table 15 Windows Management Server 2008 R2 parts list

| IBM part number | Geography | Windows Server 2008 R2 Enterprise 10 CAL |
|-----------------|--------------|--|
| 4849-MTN | World-wide | Windows Server 2008 R2 Enterprise (1-8 CPU, 10 CAL) ROK - Multilang (BR,EN,FR,SP) US, Canada, and LA |
| 4849-MTM | EMEA/APK/APJ | Windows Server 2008 R2 Enterprise (1-8 CPU, 10 CAL) ROK - Multilang (CS,EN,FR,DE,IT,JA,KO,PL,RU,SP,TR) |
| 4849-MTX | APGCG/ASEAN | Windows Server 2008 R2 Enterprise (1-8 CPU, 10 CAL) ROK - Chinese Simplified |
| 4849-MTT | APGCG/ASEAN | Windows Server 2008 R2 Enterprise (1-8 CPU, 10 CAL) ROK - Chinese Traditional |

Active Directory server parts list – Windows Server 2008 R2 Standard

Table 16 provides parts ordering information for the Windows Server 2008 R2 Standard. Quantities to order are:

- ▶ Advanced offering: 1
- ▶ Mainstream offering: 1
- ▶ Elite offering: Not required

Table 16 Windows Active Directory Server 2008 R2 parts list

| IBM part number | Geography | Windows Server 2008 R2 Standard |
|-----------------|--------------|---|
| 4849-MSN | World-wide | Windows Server 2008 R2 Standard (1-4 CPU, 5 CAL) ROK - Multilang (BR,EN,FR,SP) US, Canada, and LA |
| 4849-MSM | EMEA/APK/APJ | Windows Server 2008 R2 Standard (1-4 CPU, 5 CAL) ROK - Multilang (CS,EN,FR,DE,IT,JA,KO,PL,RU,SP,TR) |
| 4849-MSX | APGCG/ASEAN | Windows Server 2008 R2 Standard (1-4 CPU, 5 CAL) ROK - Chinese Simplified |
| 4849-MST | APGCG/ASEAN | Windows Server 2008 R2 Standard (1-4 CPU, 5 CAL) ROK - Chinese Traditional |

Management server parts list – Virtual Machine Manager

Table 17 provides parts ordering information for the Virtual Machine Manager. Quantities to order are:

- ▶ Advanced or Mainstream offering: No IBM part number at the time of this publication. See “Ordering Information for Additional Microsoft components” on page 28 for alternate ordering of the part.
- ▶ Elite offering: Two x3650 M3 servers, each with two CPUs, can manage up to five servers.

Table 17 Virtual Machine Manager parts list

| IBM part number | Geography | Virtual Machine Manager 2008 R2 Workgroup Edition |
|-----------------|------------|---|
| 4849-MVN | World-wide | MS Sys Ctr Virtual Machine Mgr 2008 R2 Workgroup Ed ROK - Multilang (Xcc) - US, Canada, and Latin America |
| 4849-MVN | World-wide | MS Sys Ctr Virtual Machine Mgr 2008 R2 Workgroup Ed ROK - Multilang (Standard Seller)-WW |

| IBM part number | Geography | Virtual Machine Manager 2008 R2 Workgroup Edition |
|-----------------|-----------|--|
| 4849-MVM | EMEA | MS Sys Ctr Virtual Machine Mgr 2008 R2 Workgroup Ed ROK - Multilang (Express Seller)-EMEA only |

Ordering Information for Additional Microsoft components

Microsoft licenses provided with IBM Hyper-V Cloud Offering (as specified in the parts list in “Host server parts list - Windows Server 2008 R2 Datacenter Edition 2-CPU” on page 26”) include:

- ▶ Elite Offering
 - Windows Server 2008 R2 Enterprise Edition for Management Server
 - Windows Server 2008 R2 Data Center for Host Servers
 - SCVMM 2008 R2 - Workgroup Edition for Management Servers
- ▶ Mainstream and Advanced Offerings
 - Windows Server 2008 R2 Enterprise Edition for Management Server
 - Windows Server 2008 R2 Data Center for Host Servers

Additional Microsoft licenses that need to be procured outside of the IBM provided Microsoft Licenses include:

- ▶ Elite Offering
 - System Center Operations Manager 2007 R2
 - SCVMM Self-Service Portal v2.0 Solution Accelerator
 - Microsoft SQL Server 2008 R2
- ▶ Mainstream and Advanced Offerings
 - Microsoft SQL Server 2008 R2
 - System Center Operations Manager 2007 R2
 - System Center Virtual Machine Manger 2008 R2
 - System Center Data Protection Manager 2010
 - Forefront Identity Manager 2010
 - Forefront Endpoint Protection 2010
 - Microsoft Deployment Toolkit 2010 Update 1

Summary

Table 18 (for US and Canada) and Table 19 on page 29 (for EMEA and the rest of the world) summarize the software components being used in the three defined offerings and the licenses needed.

Table 18 Part number list for the US and Canada

| Software | IBM part | Where used | Server type | Quantity of licenses in each offering | | |
|--|----------|---|-------------|---------------------------------------|------------|----------|
| | | | | Elite | Mainstream | Advanced |
| Windows Server 2008 R2 Data Center 4-CPU | 4849-MHN | All computer production servers – unlimited users | x3755 M3 | 0 | 0 | 8 |
| Windows Server 2008 R2 Data Center 2-CPU | 4849-MGN | All computer production servers - unlimited users | x3650 M3 | 4 | 8 | 0 |

| Software | IBM part | Where used | Server type | Quantity of licenses in each offering | | |
|--|----------|-------------------------|-----------------|---------------------------------------|----------------|----------------|
| | | | | Elite | Mainstream | Advanced |
| Windows Server 2008 R2 Enterprise 1-8 CPU 10CALs | 4849-MTN | All management servers | x3650 M3 | 2 | 2 | 2 |
| Windows Server 2008 R2 Standard 1-8 CPU 10CALs | 4849-MSN | Active Directory server | x3550 M3 | Optional | Optional | 1 |
| System Center Virtual Machine Manager 2008 R2 | None | Management servers | VM on x3650 M3 | 1 ^a | 1 ^a | 1 ^a |
| Windows System Center Operations Manager 2007 R2 | None | Management servers | VM on x3650 M3 | 1 ^a | 1 ^a | 1 ^a |
| Windows SQL Server Standard Edition 2008 R2 | None | Management servers | VMs on x3650 M3 | 1 ^b | 1 ^b | 1 ^a |

- a. A single instance (and license) of the software running in a VM on one of two management servers in the Hyper-V Management cluster
- b. A single license shared between two cluster-aware VMs in an active-passive cluster running on the two management servers

Table 19 Part number list for EMEA and the rest of the world

| Software | IBM part | Where used | Server type | Quantity of licenses in each offering | | |
|--|----------------------|---|-------------|---------------------------------------|----------------|----------------|
| | | | | Elite | Mainstream | Advanced |
| Windows Server 2008 R2 Data Center 2-CPU | 4849-MGM | All computer production servers – unlimited users | x3650 M3 | 4 | 8 | 0 |
| Windows Server 2008 R2 Data Center 4-CPU | 4849-MHM | All computer production servers – unlimited users | x3755 M3 | 0 | 0 | 8 |
| Windows Server 2008 R2 Enterprise 1-8 CPU, 10 CALs | 4849-MTM | All management servers | x3650 M3 | 2 | 2 | 2 |
| Windows Server 2008 R2 Standard, 1-8 CPU, 10 CALs | 4849-MSM | Active Directory Server | x3550 M3 | Optional | 1 | 1 |
| MS System Center Virtual Machine Manager 2008 R2 Workgroup Ed. | 4849-MVW (EMEA only) | One Management Server (**License shared between two MGMT servers) | x3650 M3 | 1 ^a | 1 ^a | 1 ^a |
| Windows System Center Operations Manager 2007 R2 | None | One Management Server | x3650M3 | 1 ^a | 1 ^a | 1 ^a |
| Windows SQL Server Standard Edition 2008 R2 | None | One Management Server | x3650M3 | 1 ^b | 1 ^b | 1 ^b |

- a. A single instance (and license) of the software running in a VM on one of two management servers in the Hyper-V Management cluster

- b. A single license shared between two cluster-aware VMs in an active-passive cluster running on the two management servers

Management software

This section describes the management software that forms part of the PCO offerings.

IBM Hardware Management Pack for Microsoft System Center Operations Manager 2007

The IBM Hardware Management Pack for Microsoft System Center Operations Manager (MSCOM) 2007 is an upward integration module that plugs into Microsoft SCOM Server 2007 R2. With this management pack, you can use the enhanced features of Microsoft System Center Operations Manager to manage the health of the IBM hardware in the offering.

Key features include:

- ▶ Extensive monitoring of the health of hardware components for IBM System x servers and BladeCenter® x86/x64 blade servers running Microsoft Windows
- ▶ Comprehensive monitoring of the health of software stacks for managing IBM hardware
- ▶ Easy determination of system health problems by aggregating monitors for hardware health

To download the latest version, visit the following URL (IBM ID required - follow the link on the page to register) and select the UIM as shown in Figure 13:

https://www14.software.ibm.com/webapp/iwm/web/reg/download.do?source=dmp&S_PKG=director_uim&S_TACT=sms&en_US&cp=UTF-8



Figure 13 Downloading the IBM Hardware Management Pack

Systems Director agent

The operating system agents serve as the control point for accessing operating system and host information. These agents, in conjunction with the IBM Hardware Management Pack for SCOM 2007 R2, can be discovered and managed by Microsoft System Center Operations Manager Server.

Go to the following website for information about downloading Upward Integration Module agents:

<http://ibm.com/systems/software/director/downloads/agents.html>

The team who wrote this paper

This paper was produced by a team of specialists from around the world.



Marco Rengan is a long-time IBM employee and has seen many transitions in the PC and server marketplace over the length of his career. He has had his own IT company and was the CTO for a small chain. He is an engineer with an MBA and has been focused on new product marketing for IBM. In his current role as Worldwide Marketing Manager for the System x High Volume Server team, Marco has been instrumental in defining and bringing the IBM System x Private Cloud Offering to market along with Microsoft and Brocade.



David Ye is a Senior Solutions Architect and has been working at IBM Center for Microsoft Technologies for 13 years. He started his career at IBM as a Worldwide Windows Level 3 Support Engineer. In this role, he helped IBM customers solve complex problems and was involved in many critical customer support cases. He is now a Senior Solutions Architect in the IBM System x Enterprise Solutions Technical Services group, where he works with customers on Proofs of Concept, solutions sizing, performance optimization, and solution reviews. His areas of expertise are Windows Server, SAN storage, virtualization, and Microsoft Exchange Server.

Thanks to the authors of the first edition of this paper:

- ▶ Marco Rengan
- ▶ Dan Ghidali
- ▶ Scott Smith
- ▶ Steven Tong
- ▶ David Ye

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

- ▶ David Watts, IBM Redbooks®
- ▶ Karen Lawrence, IBM Redbooks
- ▶ Linda Robinson, IBM Redbooks
- ▶ Bryan Hoke, Brocade

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-4731-01 was created or updated on August 31, 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:

BladeCenter®
DPI®
IBM®

Redbooks®
Redpaper™
Redbooks (logo) ®

System Storage®
System x®

The following terms are trademarks of other companies:

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

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

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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.