

SHARE in Anaheim
March 2011



A Mainframe Guy Discovers Blades.....as in zEnterprise Blade Extension

Session ID: zZS18

Glenn Anderson, IBM Training



© 2011 IBM Corporation

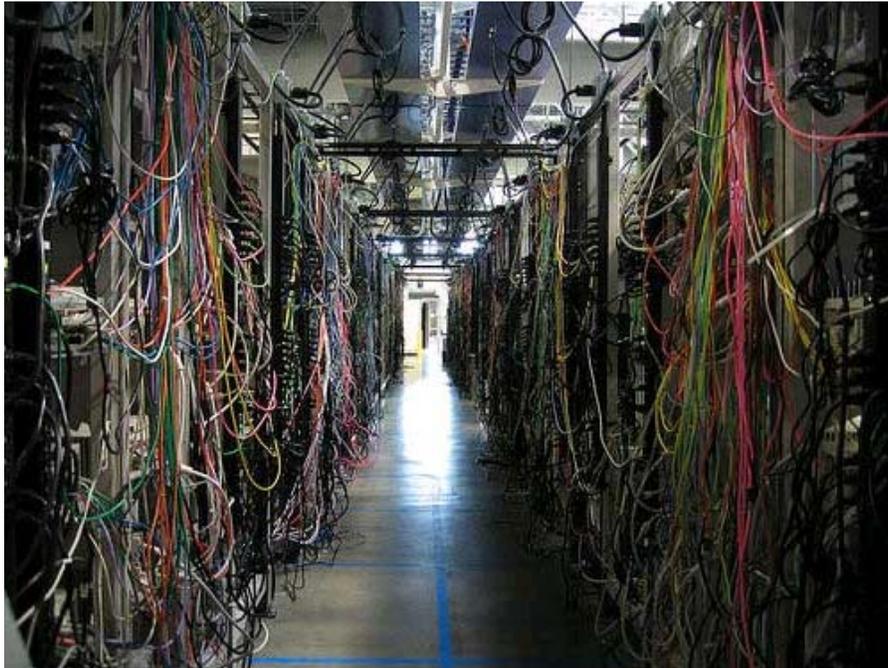






History of Blades.....

- 1999
 - Data center inefficiency
 - Easier way to deploy large numbers of rack mount web servers in data centers
 - A new server form factor...lower power consumption without sacrificing performance



History of Blades.....

- 1999
 - Data center inefficiency
 - Easier way to deploy large numbers of rack mount web servers in data centers
 - A new server form factor...lower power consumption without sacrificing performance
- 2001
 - RLX Technologies
 - New term... "server blade"
 - First generation target market – large internet data centers



History of Blades.....

- 2002 – 2003
 - As service provider market collapsed, blade manufacturers attempted to take products to broader enterprise data center market
 - HP, Compaq, Dell, IBM, Sun
- 2006 - 2007
 - New blade designs further address needs of data centers....better I/O management and thermal management



“Interesting Quotes”

- “Blade servers arrived before the advent of virtualization.....”
 - Greg Shields, IT Consultant



What is a Blade Server?

- A stripped down server computer with a modular design optimized to minimize the use of physical space and energy
 - A standard [rack-mount server](#) can function with (at least) a power cord and network cable
 - [Blade servers](#) have many components removed to save space, minimize power consumption and other considerations



What is a Blade Server?

- Blade servers offer a standardized method of deploying multiple processors, memory and I/O resources by placing those resources on plug-in boards that slide into a standard chassis



What is a blade server?

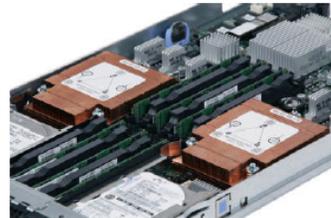
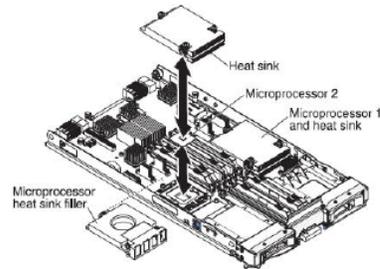
- A *blade server* contains the core components of a server
- Each blade server has:
 - Processor(s)
 - Memory
 - Internal storage (optional)
 - Network interface cards (NIC)
 - Optional plug-in components
- The blade server plugs into the midplane of a *chassis* that provides common functions:
 - Management console access (KVM)
 - Power
 - Cooling
 - Connectivity (LAN, SAN, NAS, HPC)
 - Shared media devices (DVDROM, USB)
 - Optional I/O modules to support additional function



© Copyright IBM Corporation 2010

Processor

- Processor architectures:
 - x86-64
 - IBM POWER
 - IBM PowerXCell 8i processor
- Single to four socket implementations
- Single to eight core processors



Storage

- Hard drives
 - SAS
- Hot-swap on selected blade servers
- RAID (0,1)
- Solid state disk (SSD)



Solid state disk (SSD)



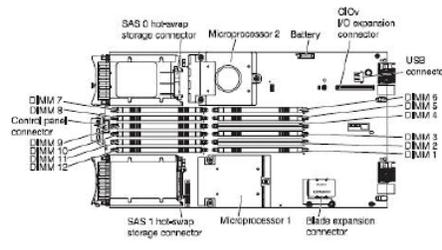
Solid state disk (SSD) for HX5



Hard disk

Memory

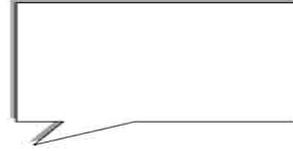
- System board-mounted DIMMs
- Speed and capacity
 - Up to PC2-6400 memory speed
 - System board capacity to 128 GB
- Error correction and redundancy
 - ECC detection and correction
 - Chipkill technology



Top of the Rack Switch (TOR)

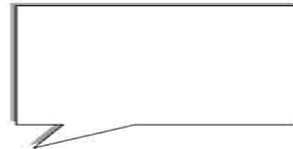
- A Top of the Rack or (TOR) switch is a small port count switch that sits on the very top or near the top of a Telco rack you see in [Data Centers](#) or Co-location facilities.
- Cisco and Brocade as well as other OEM (Original Equipment Manufacturers) describe their new FCoE (Fiber Channel over Ethernet) switches to be "Top of the Rack" thus differentiating them as being small [low-cost](#) low port count switches easy to deploy and manager etc.

Blades vs Racks: Industry Opinions



- Rack-mount servers have lower cost entry point than a blade-based installation
- Investment in blades carries a greater economic and technological risk
- Blades have an inherent lack of network interface cards (NICs)

Blades vs Racks: Industry Opinions



- Fill the chassis with blade servers, price per server comes way down
- Blades ease server management for administrator
- Blade infrastructure more efficient than racks with consolidated power and cooling
- Blades reduce cabling requirements

IBM Blade Servers



IBM BladeCenter Family

BladeCenter Chassis by model and type

IBM Blade Center E

High density, power efficient



- 7U design
- Up to 14 blade bays
- Up to four switch fabrics
- Low cost
- Support 10 GB uplinks
- Support 8 Gb FC
- SAS Connectivity Module

IBM BladeCenter T

Highly rugged, Telco, AC/DC, NEBS, air filtration



- 8U design
- Up to 8 blade bays
- Up to 4 switch fabrics
- AC or DC models
- NEBS compliant
- Rugged
- Support 10 GB uplinks
- Support 8 Gb FC
- Telco, military, dirty floor
- SAS Connectivity Module

IBM Blade Center H

Ultra high performance, and I/O flexibility



- 9U design
- Up to 14 blade bays
- Up to 10 GB midplane
- I/O flexibility up to 8 switch bays
- Support 30 mm blades with up to 8 ports
- Support 10 GB Ethernet
- Support 8 Gb FC
- Support 4x(QDR) InfiniBand
- SAS Connectivity Module

IBM BladeCenter HT

Highly rugged, Telco, AC/DC, NEBS, air filtration



- 12U design
- Up to 12 blade bays
- AC or DC models
- I/O flexibility up to 8 switch bays
- NEBS compliant
- Rugged
- Up to 10 GB midplane
- Support 10 GB Ethernet
- Support 8 Gb FC
- SAS Connectivity Module
- Telco, military, dirty floor

IBM BladeCenter S

Integrated disk, 110-240v power, BladeCenter outside the datacenter



- 7U design
- Up to 6 blade bays
- Integrated storage
- Up to 3 switch fabrics
- Power (100 - 240v)
- 950 with 1450 AC auto-sensing
- Support 10 GB uplinks
- Support 4 GB FC
- SAS RAID and Connectivity Modules

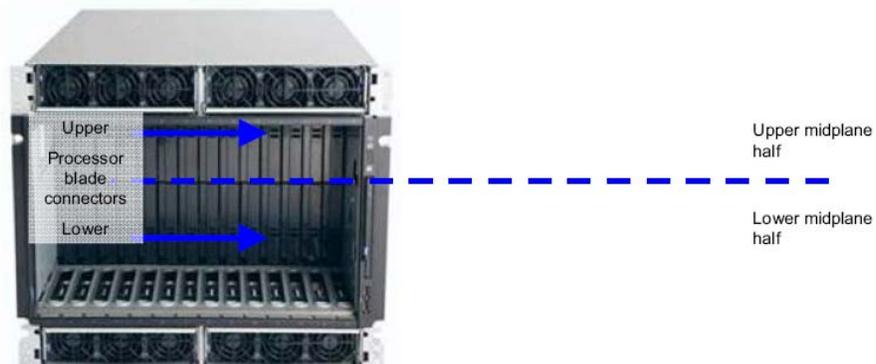
← Common blades, common switches, common management →

Chassis: IBM BladeCenter H

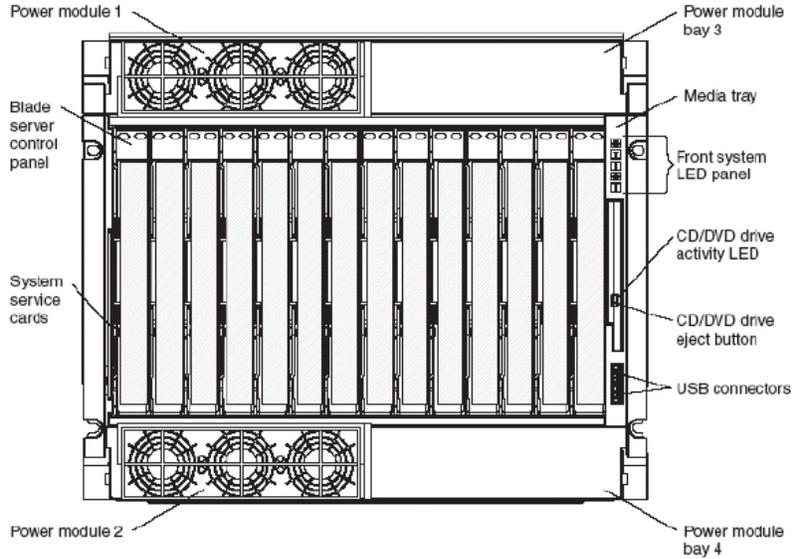
- 9U and 28 inches deep
- Supports up to 14 30mm blades
- Customer serviceable, hot swappable media tray
- 9.5mm combo drive (CD/DVD)
 - Two USB front inputs
 - Full light path diagnostics panel
- Rack mounted on rails
- Four front load 2900W power supplies
 - Each power supply includes a replaceable fan pack with three fans



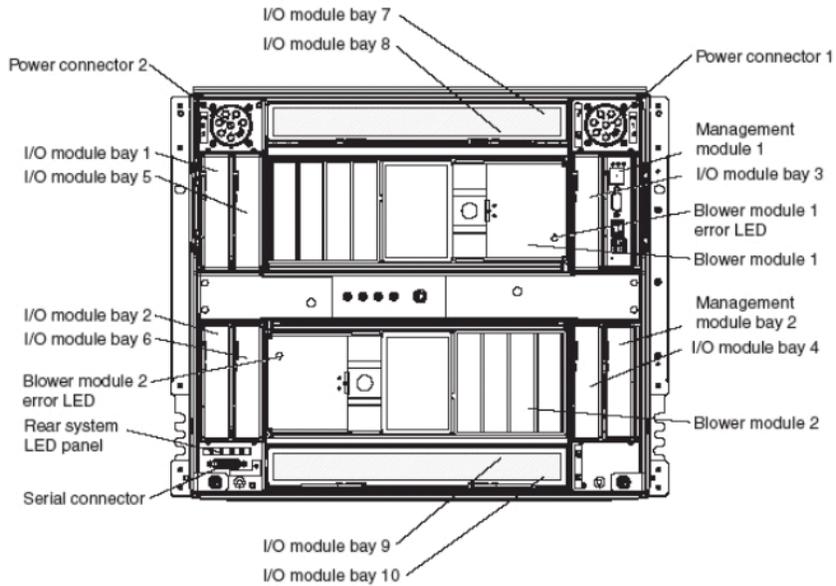
BladeCenter Interior View



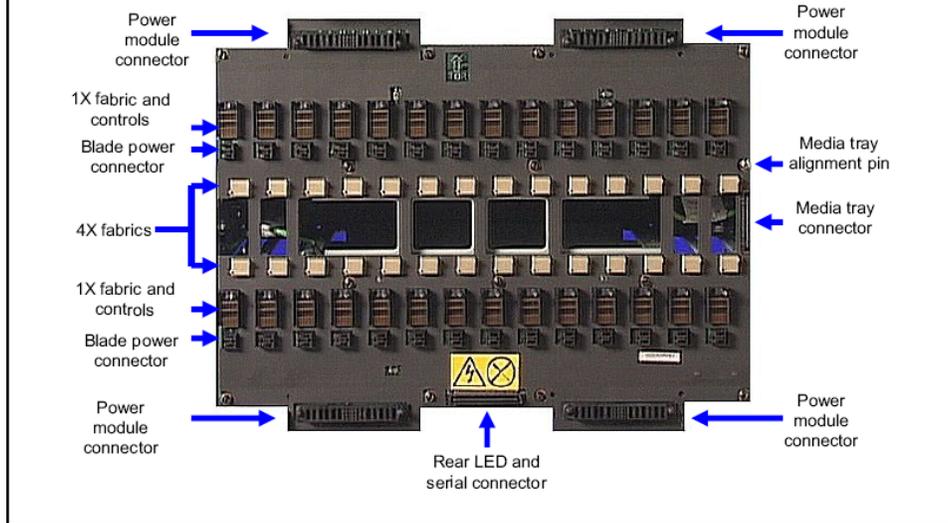
IBM BladeCenter H Front Panel



IBM BladeCenter H Rear View

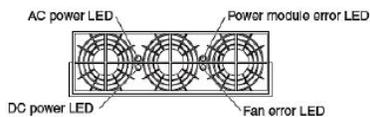
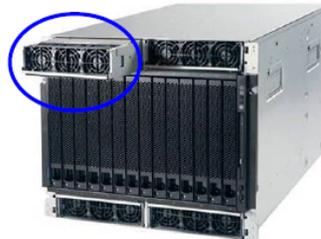


BladeCenter Midplane

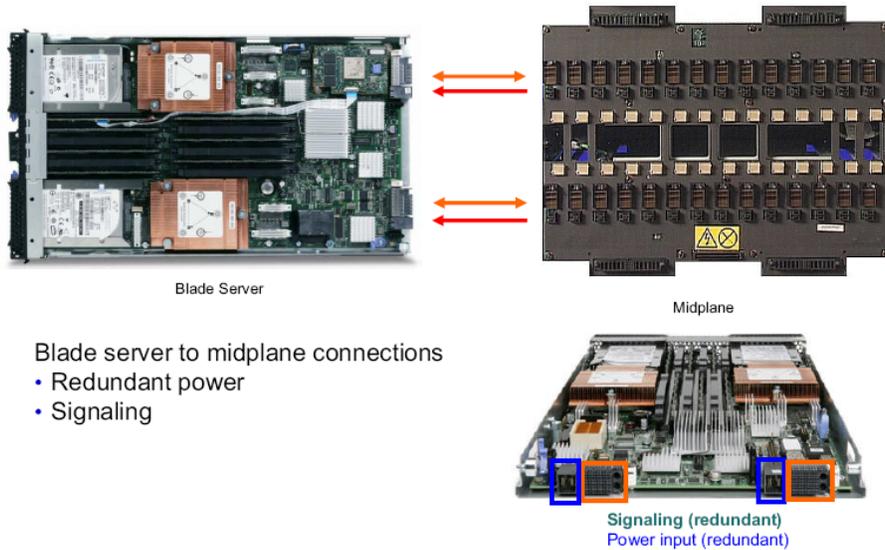


BladeCenter H Power Modules

- Power module bays
 - Maximum of four per chassis
 - Two ship standard with the chassis
 - Power modules 1 and 2 standard pair
 - Power modules 3 and 4 optional pair

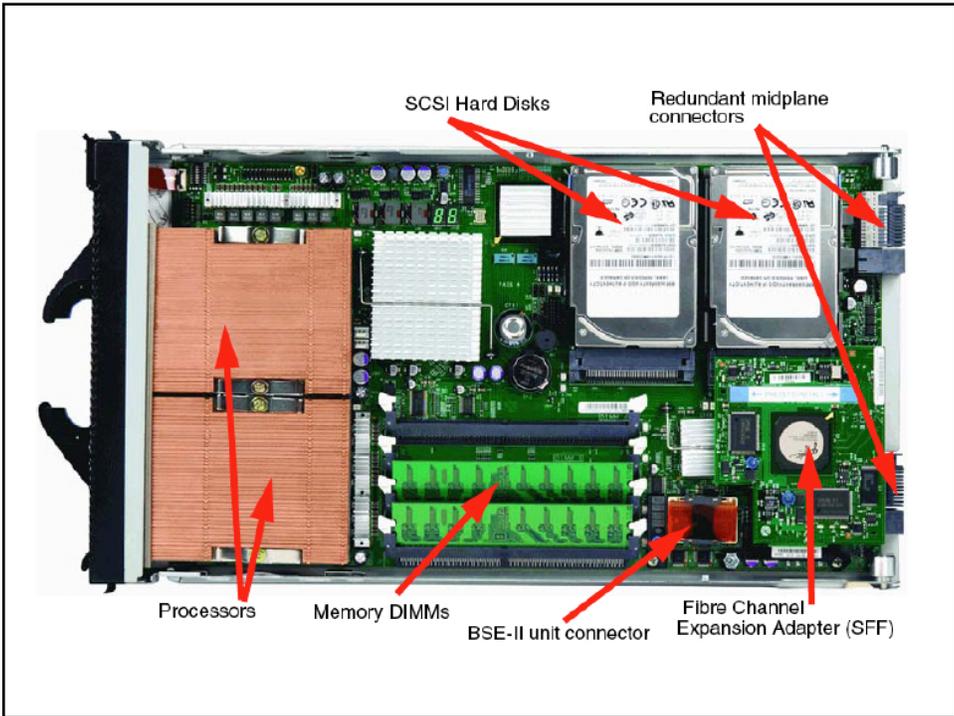
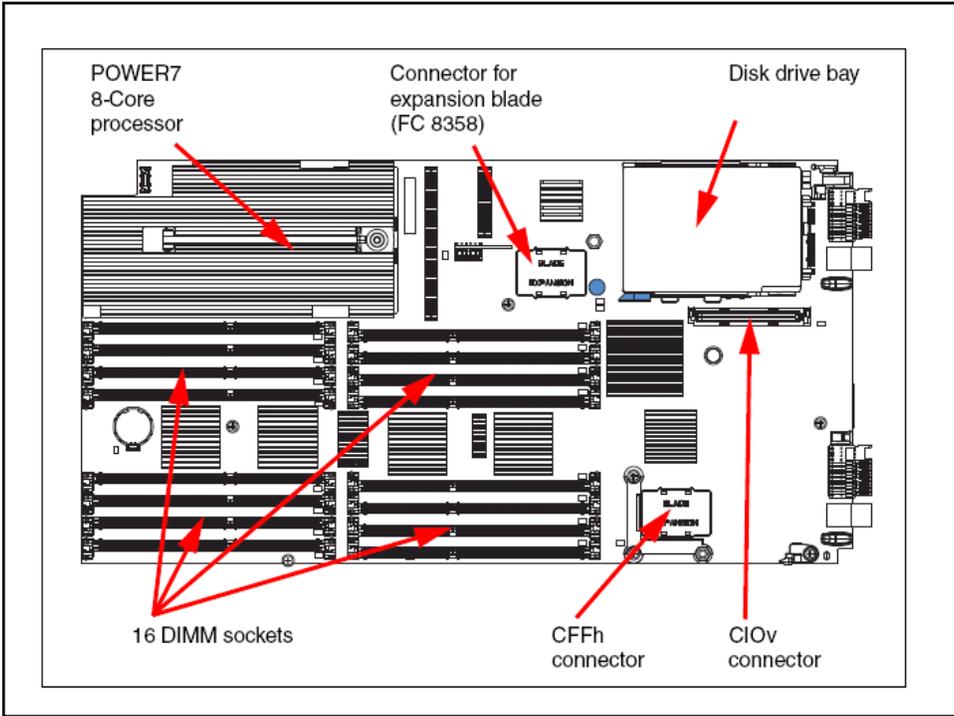


Power and Signal Paths



The POWER7 Product Family

Express	Scalable Midrange	Scalable High End
 <p>POWER 750 Express 6,8,12,16,18, 24,32 cores 8-512 GB Memory Up to 8 I/O Drawers</p>	 <p>POWER 780 8 or 16 Cores / Node TurboCore Mode 4-64 Cores at 3.8 GHz 4-32 Cores at 4.1 GHz 24x7 Warranty PowerCare Service Up to 2 TB memory Capacity on Demand</p>	 <p>POWER 795 Up to 128 @ 4.25 GHz (4) Up to 192 @ 3.72 GHz (6) Up to 256 @ 4.00 GHz (8) Up to 8 TB memory PowerCare Service Capacity on Demand</p>
 <p>POWER 740 Express 4, 6, 8, 12, 16 cores 8-256 GB Memory Up to 8 I/O Drawers</p>		
 <p>POWER 720 Express 4, 6 or 8 cores 8-128 GB Memory Up to 4 I/O Drawers</p>	 <p>POWER 770 12 or 16 Cores / Node 4-64 Cores at 3.1 GHz 4-48 Cores at 3.5 GHz Up to 2 TB memory Capacity on Demand</p>	<p>HPC</p>  <p>POWER6 575 32 4.7 GHz / Node Up to 256 GB / Node</p>
 <p>POWER 730 Express 8, 12, or 16 cores 8-128 GB Memory</p>		
 <p>POWER 710 Express 4, 6 or 8 cores 8-64 GB Memory</p>		 <p>POWER 755 32 Cores at 3.55 GHz Up to 256 GB Memory</p>
<p>POWER Blades PS700 – 4 core/64 GB max PS701 – 8 cores/128 GB max PS702 – 16 cores / 256 GB max</p>	<p>Selected for zEnterprise zBX</p>	



The PS701 blade server

The PS701 blade server (8406-71Y) is a single socket, single-wide 8-core 3.0 GHz POWER7 processor-based server. The POWER7 processor is a 64-bit, 8-core with 256 KB L2 cache per core and 4 MB L3 cache per core.

The PS701 blade server has 16 DDR3 memory DIMM slots. The industry standard VLP DDR3 memory DIMMs are either 4 GB or 8 GB running at 1066 MHz. The memory is supported in pairs, thus the minimum memory required for PS701 blade server is 8 GB (two 4 GB DIMMs). The maximum memory that can be supported is 128 GB (16x 8 GB DIMMs).

The PS701 blade server has two Host Ethernet Adapters (HEA) 1 GB integrated Ethernet ports that are connected to the BladeCenter chassis fabric (midplane). The PS701 also has an integrated SAS controller that supports local (on-board) storage, integrated USB controller and Serial over LAN console access through the service processor, and the BladeCenter Advance Management Module.

The PS701 has one on-board disk drive bay. The on-board storage can be one 2.5-inch SAS HDD or SSD drive. The PS701 also supports one PCIe CIOv expansion card slot and one PCIe CFFh expansion card slot. See 1.5.8, "I/O features" on page 24 for supported I/O expansion cards.

Advanced Management Module

Center point for IBM BladeCenter infrastructure intelligence

- Hot-swappable module
- Powerful and robust systems management
- Proxy for expansion modules
- Controls all aspects of power, connectivity and communication
- Reliability, availability, and serviceability (RAS)



AIX Version 7.1



- Latest generation of IBM's well-proven, scalable, open standards-based UNIX operating system
- Built on IBM POWER technology and virtualization
- Celebrating 25th anniversary
 - Originally ran on RISC processor, engineering workstation

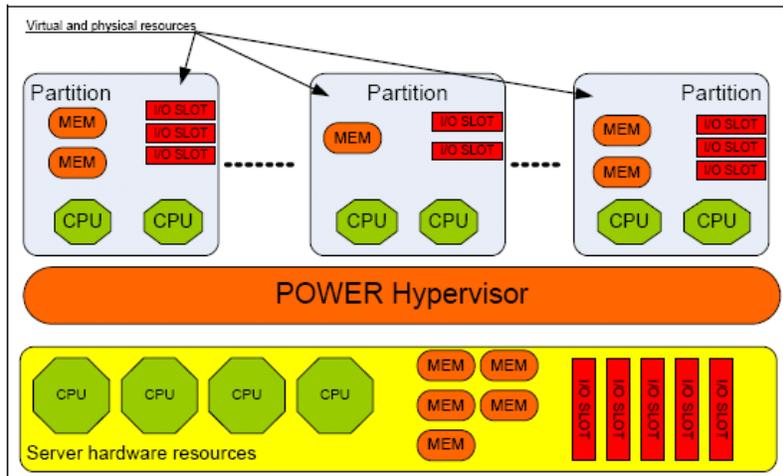
“Interesting Quotes”

- “AIX is the only operating system that leverages decades of IBM technology innovation designed to provide the highest level of performance and reliability of any UNIX operating system.”

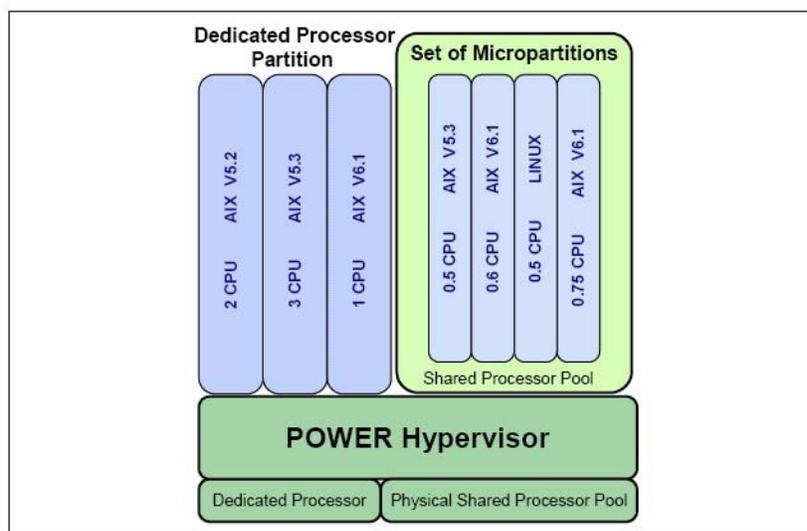
- IBM AIX Web Page



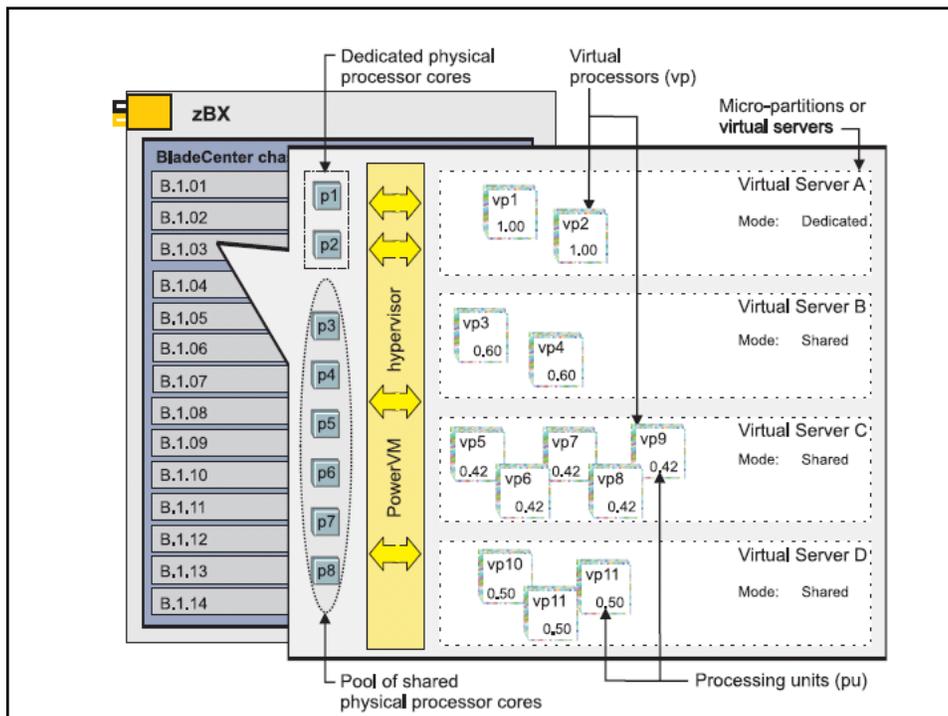
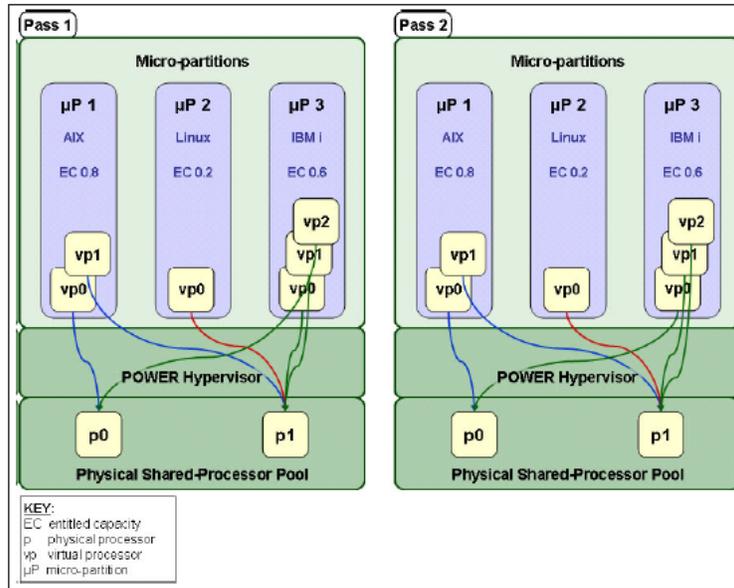
Power Hypervisor



Dedicated and Shared Processor Modes

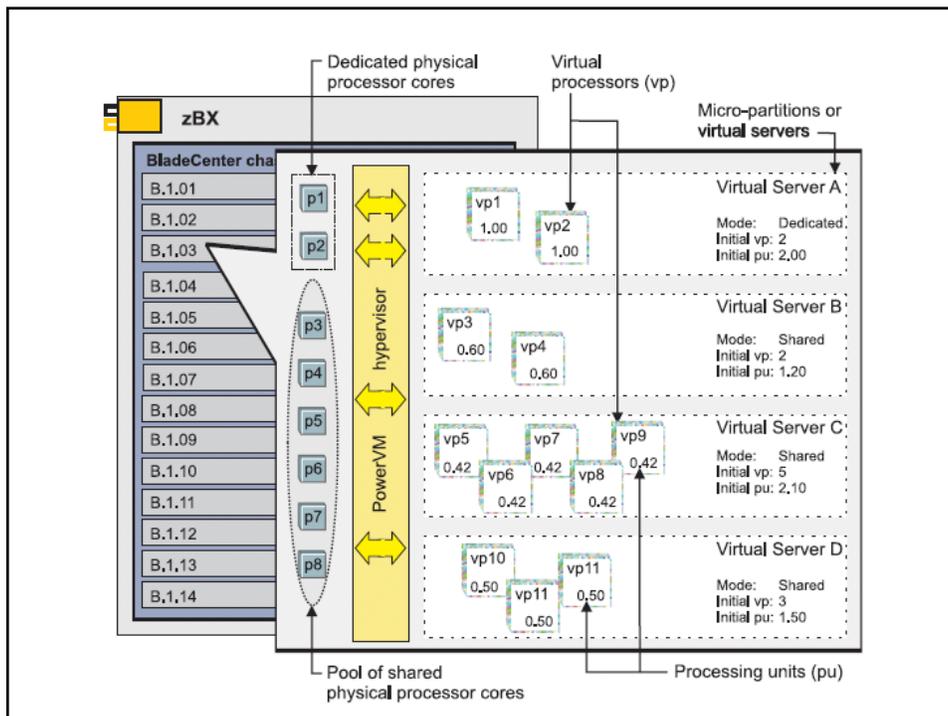


Virtual Processor Mapping

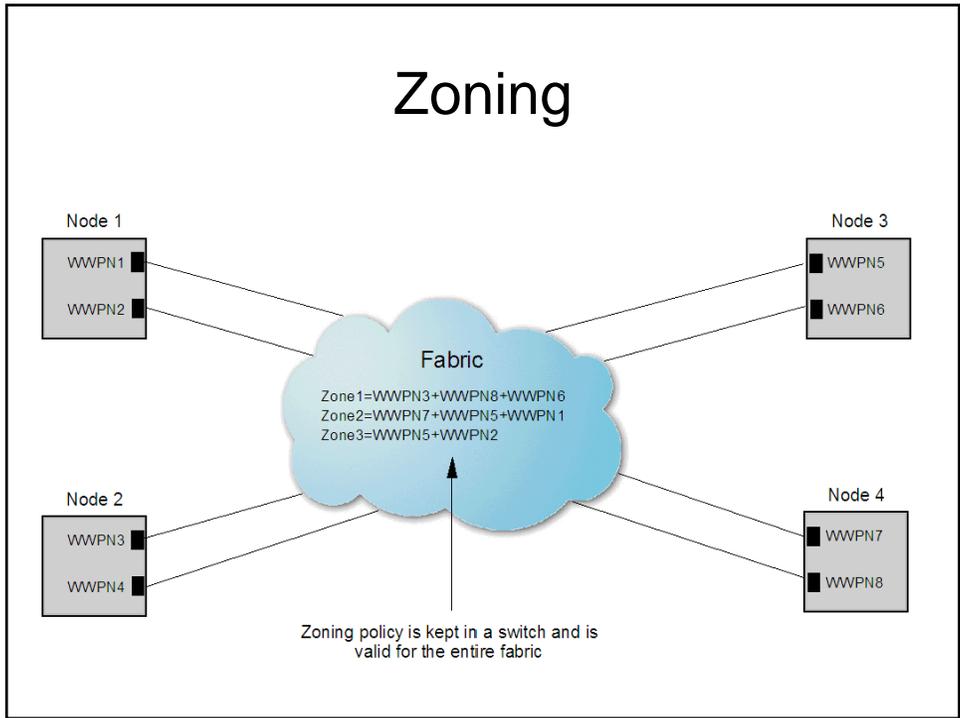


Sample Calculation for Configuring Virtual Servers

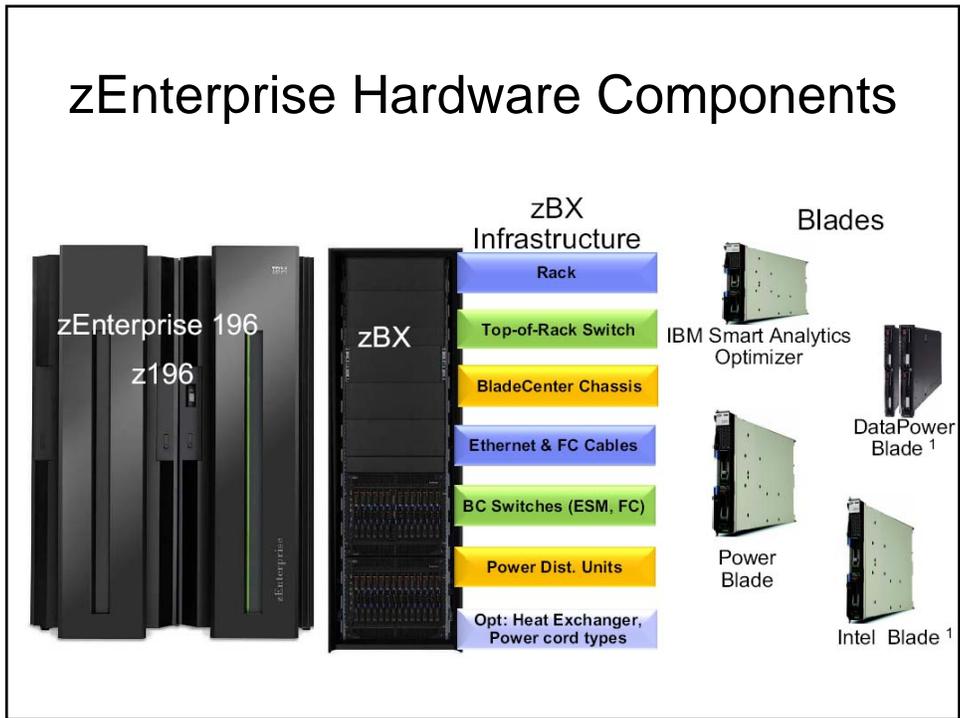
Configuration step	Available physical processor cores	Available processing units (pu)
Start with one PS701 blade that has no virtual servers defined	8	8.00
Subtract amount reserved for the blade hypervisor (0.80 pu)	7	7.20
Add virtual server A, which requires 2 dedicated processors assigned to two virtual processors	5	5.20
Add virtual server B, which requires 1.20 processing units assigned to two virtual processors	4	4.00
Add virtual server C, which requires 2.10 processing units assigned to five virtual processors	1	1.90
Add virtual server D, which requires 1.50 processing units assigned to three virtual processors	0	0.40



Zoning



zEnterprise Hardware Components



POWER7 Blade in the zBX

- General-purpose computing platform
 - Housed in standard BladeCenter H chassis inside IBM® zEnterprise™ BladeCenter® Extension enclosure
 - Up to 112 blades
 - 14 blades per BladeCenter
 - 2 BladeCenters per rack
 - 4 racks per zBX Model 2
- Managed by the IBM zEnterprise Unified Resource Manager
- Virtualized with firmware-supplied hypervisor
- Entitled through System z firmware



© 2011 IBM Corporation

zEnterprise Firmware

- PowerVM Enterprise Edition is the hypervisor
 - pHyp 2.1
 - VIOS 2.1.3
 - Extensions for configuration systems management
 - Hardware setup
 - FFDC
 - Heartbeat
 - PPM daemon
- Packaged and managed as firmware on Support Element
 - Loaded automatically at power on
 - Serviced via System z MCL process
 - Full set of operational controls
- Managed as a closed system from the Hardware Management Console (HMC)
 - No AMM access
 - No VIOS access, so no VIOS add-ons supported. (e.g., ITM hypervisor agent, or PCM drivers)
 - Enables creation and management of virtual servers via the Unified Resource Manager



© 2011 IBM Corporation

Hypervisor

- CPU (8-core PS701)
 - Each CPU has capacity of 1.0 processing units (PUs)
 - VIOS default capacity is 0.1 PUs per CPU (8 shared virtual CPUs on PS701)
 - Each dedicated CPU requires 1.0 PUs
 - Each shared virtual CPU requires at least 0.1 PUs
 - Up to 7.2 PUs across 8 shared CPUs or up to 7 dedicated CPUs (one for VIOS)
 - No over-commitment of PUs
- Memory
 - No over-commitment.
 - VIOS requires at least 4GB



© 2011 IBM Corporation

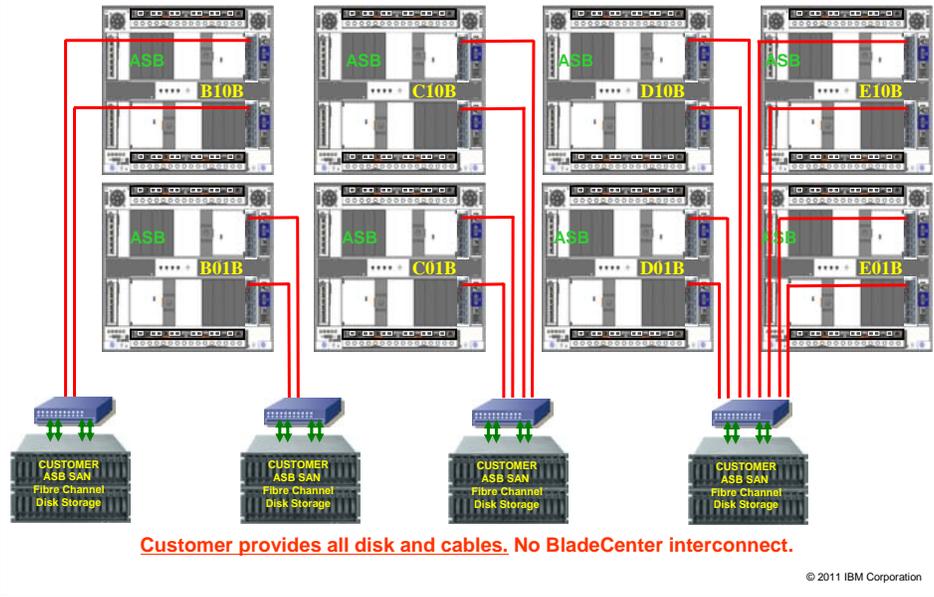
Provision Virtual Server

- Required Planning
 - Processors (virtual CPUs and capacity in PUs)
 - Memory
 - Storage (i.e., zones)
 - Network (i.e., VLANs)
 - Possible static migration targets (i.e., SAN zones, VLANs)
- To provision a virtual server
 - Create it
 - Install an operating system

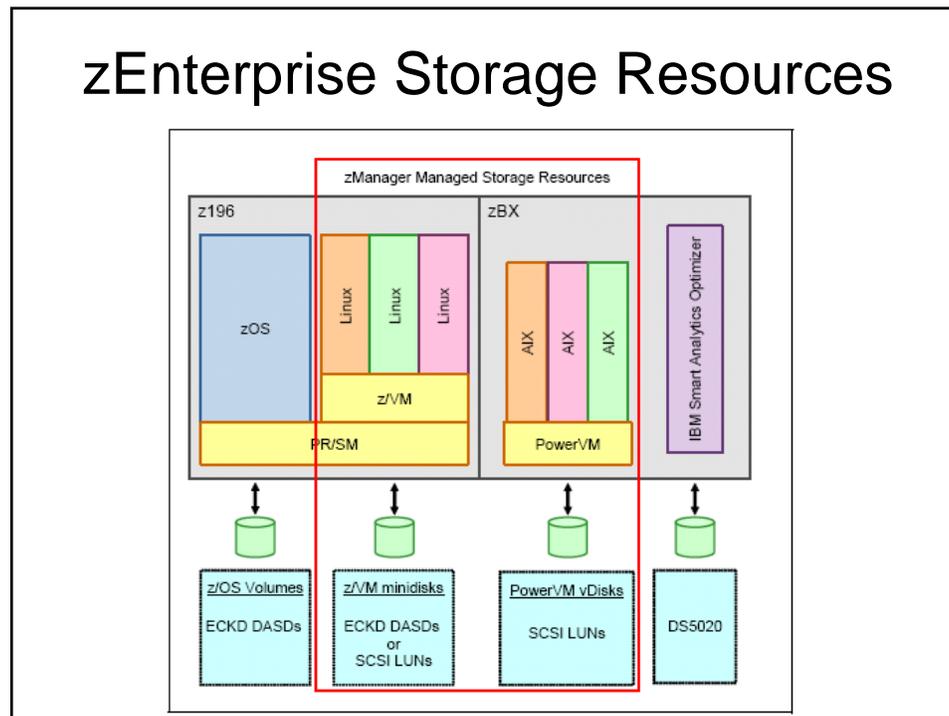


© 2011 IBM Corporation

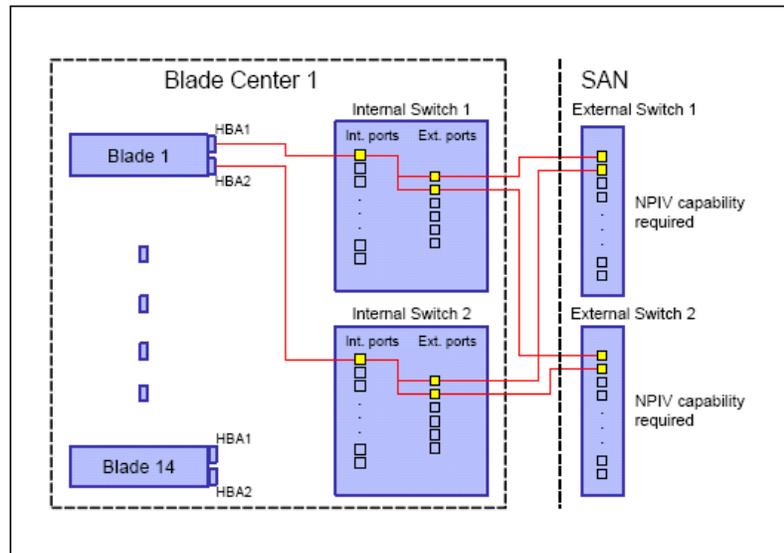
Blade FibreChannel Disk Storage Connections



zEnterprise Storage Resources



zBX Storage Topology



Similarities to Current POWER7 Blades

- zEnterprise POWER7 Blade offerings, PS701
 - 8 core, 3.0 GHz processors, 32G/64G/128G memory options
 - 2 1Gb Ethernet, 2 10 Gb Ethernet, 2 8Gbps fibre channel ports
- AIX 5.3 or higher
 - Installed via virtual media or NIM
- Virtualization via PowerVM Enterprise Edition
 - Built on POWER7 Hypervisor
 - Shared or dedicated processors
 - Capped / Uncapped LPARs
- Acquired through normal sales channels



© 2011 IBM Corporation

Differences from Current POWER7 Blades

- Only the 8 core, PS701 POWER7 blade is supported,
 - Memory, adapters, and disk options are pre-determined
- The POWER EE Hypervisor and VIOS are packaged, distributed, and installed from the System z196 Support Element using the zManager UI (Ensemble HMC)
 - The VIOS is required for all PS701 blades in zBX.
 - There is no need for AMM access.
 - There is no need for VIOS command line access.
- All LPAR lifecycle management is provided via zManager
- End to End Performance Management from the Hypervisor down is provided via zManager
 - Adjusts entitled CPU capacity to meet installation-specified objectives
 - Requires the Automate Feature Code (FC0020)
- SAN access via a pair of Qlogic 8Gbit ports which connect to the SAN via 8Gbit QLogic switches operating in Passthru Mode (NPIV). NAS support is also available.
- No Active Memory Sharing / Active Memory Expansion
- No Live Partition Mobility (at GA1)
- No Linux (PowerVM Lx86) or IBM i (formerly i5/OS)



© 2011 IBM Corporation

NEXTGEN: Hardware Management Console Workplace (Version 2.11.0) - Mozilla Firefox

9.60.92.193 https://9.60.92.193/hmc/connects/mainuiFrameset.jsp

Hardware Management Console

Manage Storage Resources | New Virtual Server | pedebug | Help | Logoff

Ensemble Management > My Ensemble

Ensemble Resources | Hypervisors | Virtual Servers

Table | Topology

Select	Name	Status	Description
<input type="checkbox"/>	Members	Exceptions	
<input type="checkbox"/>	PZBONZAI	Not Operating	Central Processing Complex (CPC)
<input type="checkbox"/>	BladeCenters	OK	All zBx BladeCenters of the Server
<input type="checkbox"/>	B.1	Operating	Represents one BladeCenter
<input type="checkbox"/>	B.2	Operating	Represents one BladeCenter
<input type="checkbox"/>	B.2.01	No Power	Represents one Power Application Server Blade
<input type="checkbox"/>	B.2.02	No Power	Represents one Power Application Server Blade
<input type="checkbox"/>	B.2.03	No Power	Represents one Power Application Server Blade
<input type="checkbox"/>	B.2.04	No Power	Represents one Power Application Server Blade
<input type="checkbox"/>	B.2.05	No Power	Represents one Power Application Server Blade

Max Page Size: 500 | Total: 22 | Filtered: 22 | Selected: 0

Tasks: My Ensemble

- Ensemble Details
 - Toggle Lock
- Configuration
 - Add Member to Ensemble
 - Delete Ensemble
 - Manage Ensemble MAC Prefixes
 - Manage Storage Resources
 - Manage Virtual Networks
 - New Workload
- Monitor

Status: Exceptions and Messages

R90HMC1: Workloads Report - Mozilla Firefox: IBM Edition

9.12.16.241 https://9.12.16.241/hmc/content?taskId=18&refresh=37

Hypervisor Report - r90f1b206v1

Report Interval: Starting 2/10/11 10:15:53 AM for 15 minutes (2/10/11 10:30:53 AM) [Modify](#)

Start: 15 minutes

Starting from: Date: 2/10/11 Time: 10:15:53 AM Duration: 15 minutes

OK Apply Cancel

Hypervisor Details:

Hypervisor: R 2.05 Processor count: 8 Total memory allocated for virtual servers: 40,960 MB
 Hypervisor type: PowerVM Total CPU consumption: 95.6% Total memory: 65,536 MB
 Total allocated processing units: 7.50

Virtual Servers:

Virtual Server	Processor Management Status	Processor Management Reason	Virtual Processors	Min Virtual Processors	Max Virtual Processors	Consumed Processors	Hypervisor Processing Unit Duty (%)	Allocated Memory (MB)	Dedicated	Capped	Processing Units	Initial Processing Units	Min Processing Units	Max Processing Units	Min Memory (MB)	Max Memory (MB)
r90f1b206v1	Active	None	2	1	7	1.02	37.4	4,096	--	--	1.00	0.50	0.10	7.00	4,096	4,096
r90f1b206v2	Active	None	2	1	7	1.87	11.0	4,096	--	--	1.00	0.90	0.10	7.00	4,096	4,096
r90f1b206v3	Active	None	2	1	7	0.09	37.9	4,096	--	--	0.71	1.20	0.10	7.00	4,096	4,096
r90f1b206v4	Active	None	2	1	7	1.00	37.7	4,096	--	--	0.90	0.90	0.10	7.00	4,096	4,096
r90f1b206v5	Active	None	2	1	7	0.99	50.7	4,096	--	--	0.35	0.90	0.10	7.00	4,096	4,096
r90f1b206v6	Active	None	2	1	7	0.02	29.0	4,096	--	--	1.00	0.90	0.10	7.00	4,096	4,096
r90f1b206v7	Active	None	2	1	7	0.99	49.2	4,096	--	--	0.34	0.60	0.10	7.00	4,096	4,096
r90f1b206v8	Active	None	2	1	7	0.97	51.4	4,096	--	--	0.34	0.80	0.10	7.00	4,096	4,096

Page 1 of 1 Total: 8 Filtered: 0 Displayed: 8

Resource Adjustments

Close Help

R90HMC1: Virtual Server Details - Mozilla Firefo...

9.12.16.241 https://9.12.16.241/hmc/wcl/Tb46

Virtual Server Details - r90f1b206v1

Name Status **Processors** Memory Network Storage Options Workloads Performance

Processing mode: Shared

Maximum assignable processing units: 7.2

Maximum assignable virtual processors: 64

Minimum processing units: * 0.1

Initial processing units: * 0.5

Maximum processing units: * 7.0

Minimum virtual processors: * 1

Initial virtual processors: * 2

Maximum Virtual processors: * 7

OK Apply Cancel Help

Done

Blade.org

About Blade.org

Blade.org is a collaborative organization and developer community focused on accelerating the development and adoption of open blade server platforms. The organization was established in February 2006 to increase the number of blade platform solutions available for customers and to accelerate the process of bringing them to market. From eight founding companies, Blade.org has grown to nearly 100 members including leading blade hardware and software providers, developers, distribution partners and end users from around the globe.

For more information, please visit: <http://www.blade.org>.