Systems and Technology Group

BladeCenter & System x Demo Days

- 2009 v.6 -

Alessandro Malosio

System x Technical Sales Support Team

Common x86 landscape

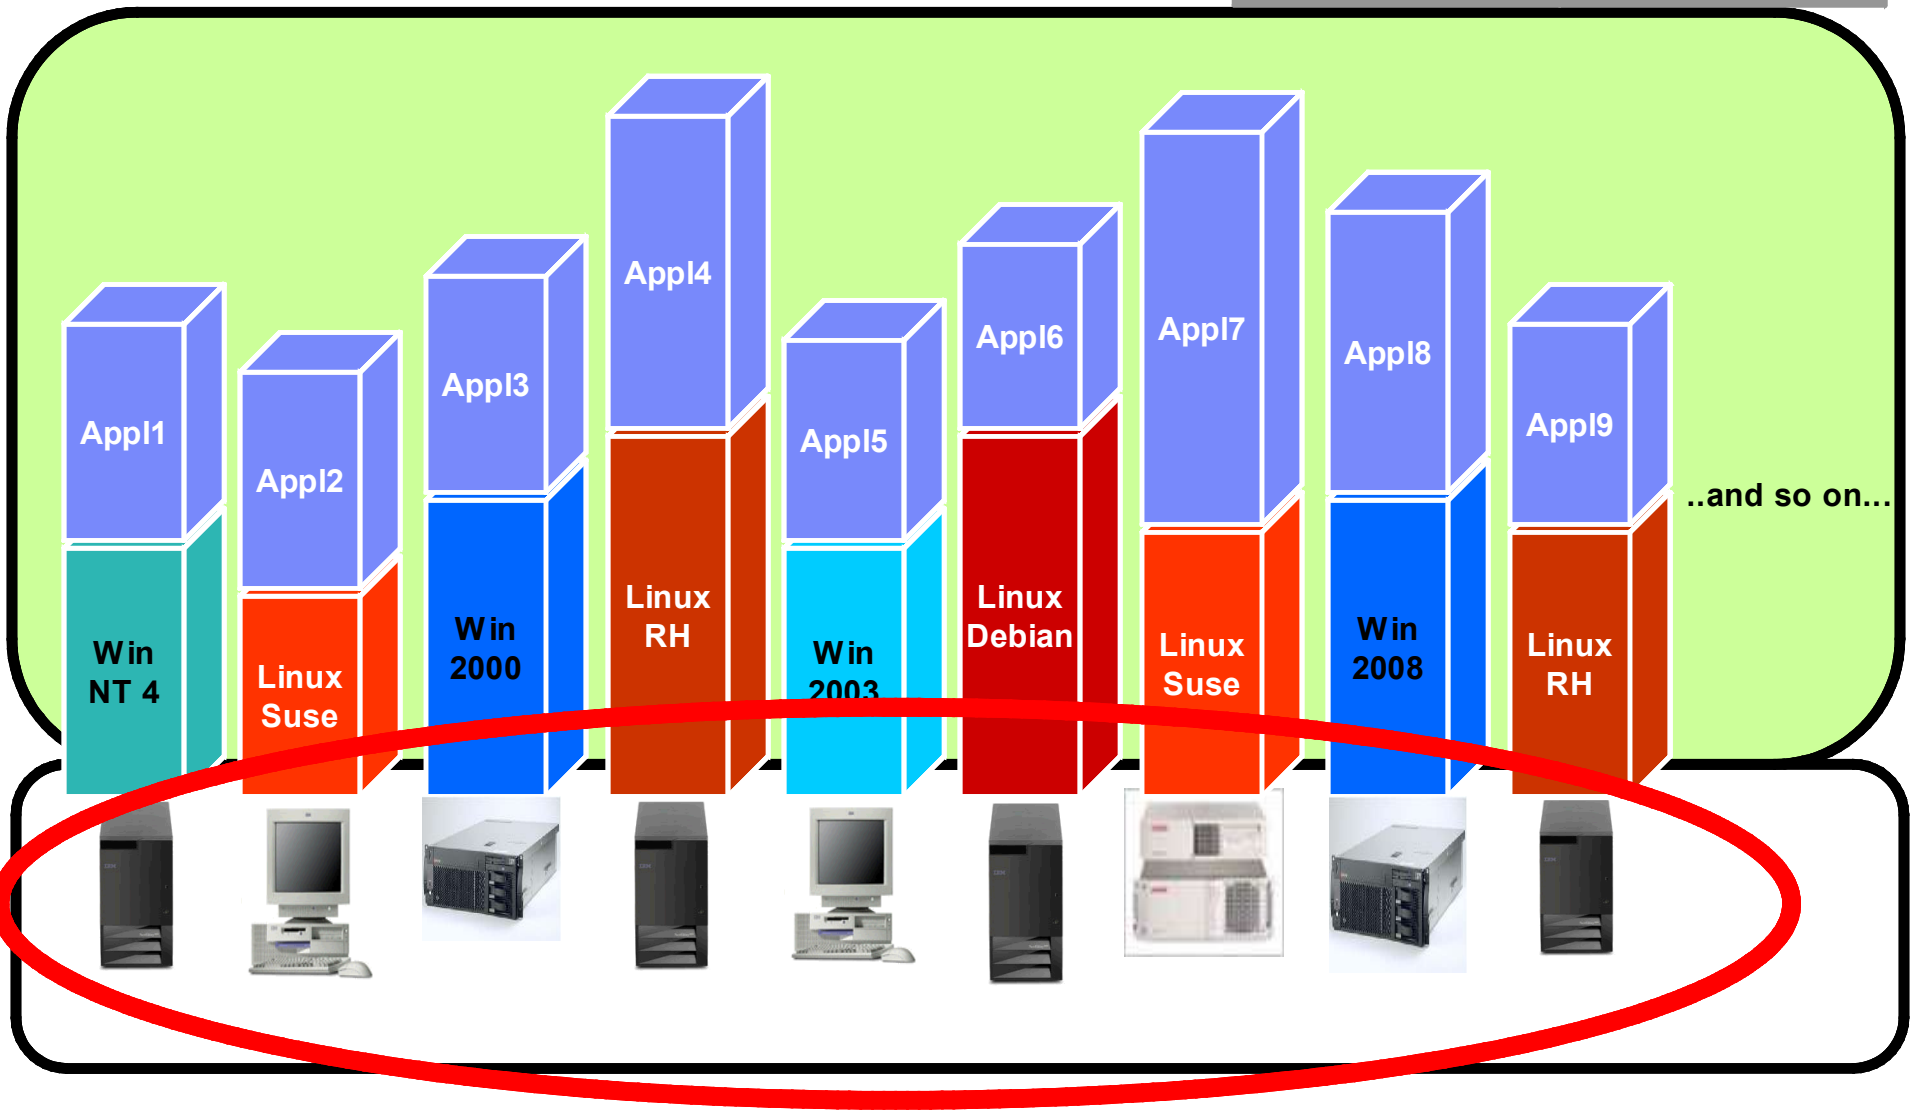

TRADITIONAL

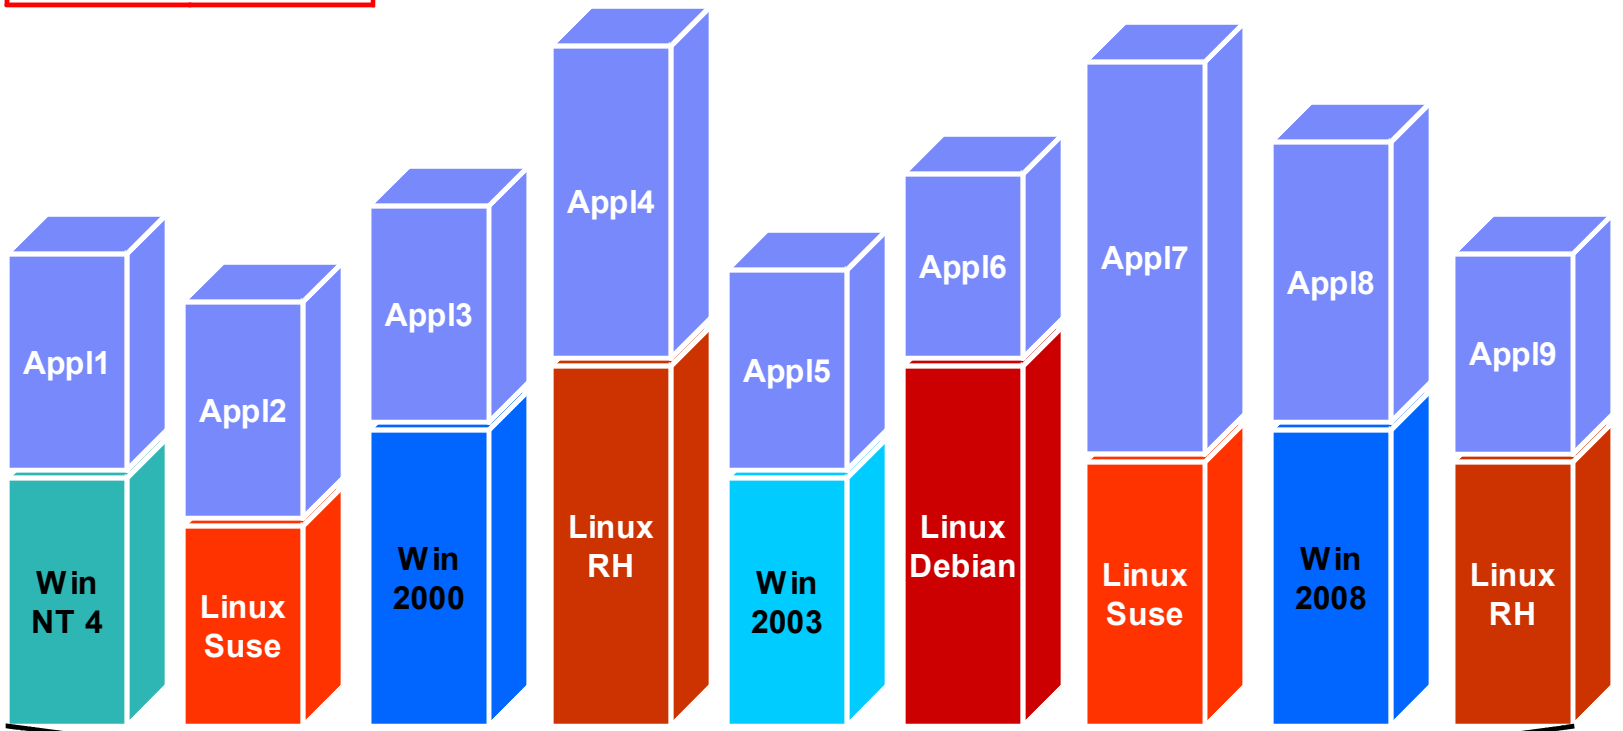

Rack Optimized servers

EVOLUTION

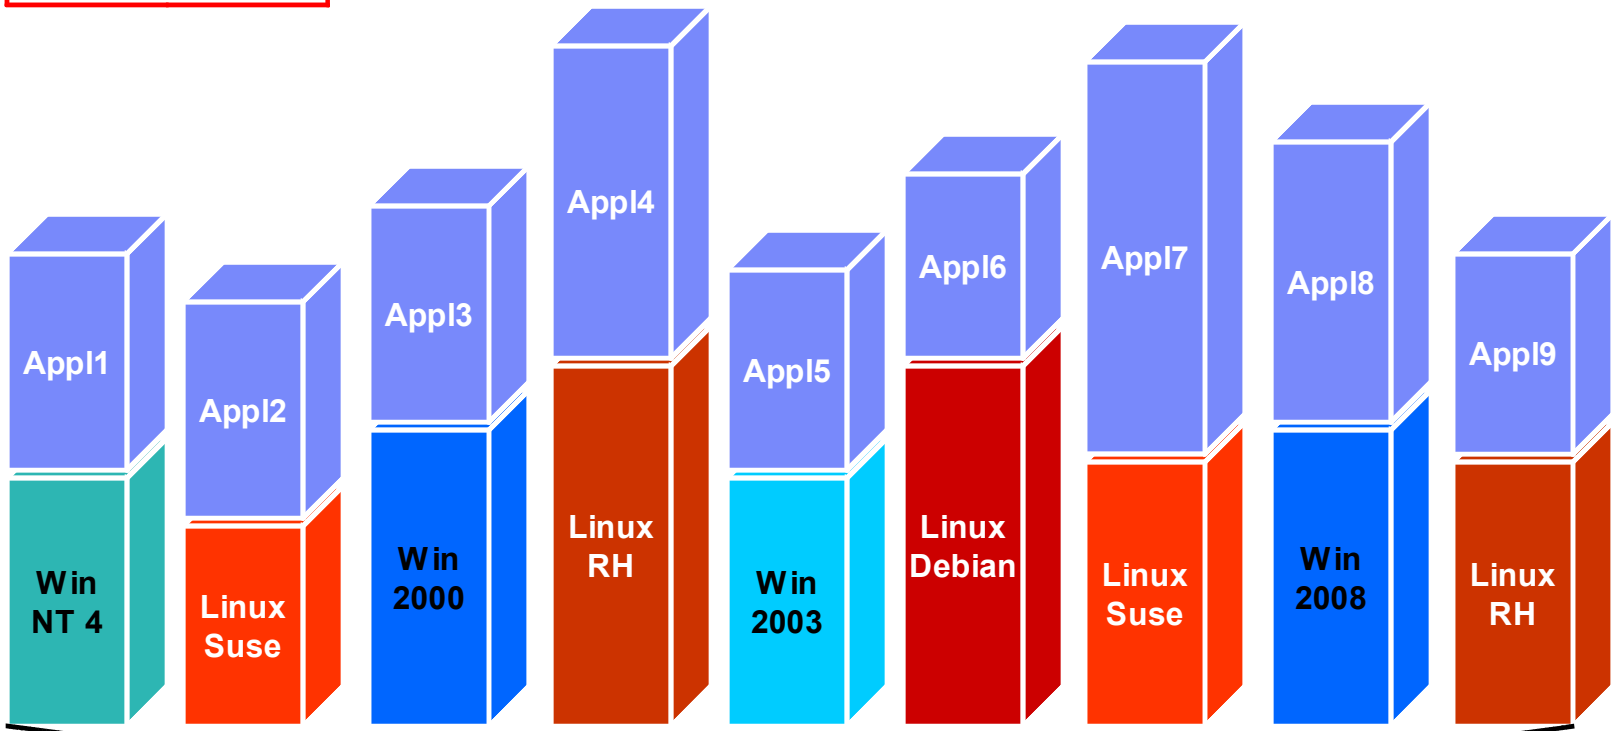

BladeCenter

SAN

LAN Eth

REVOLUTION

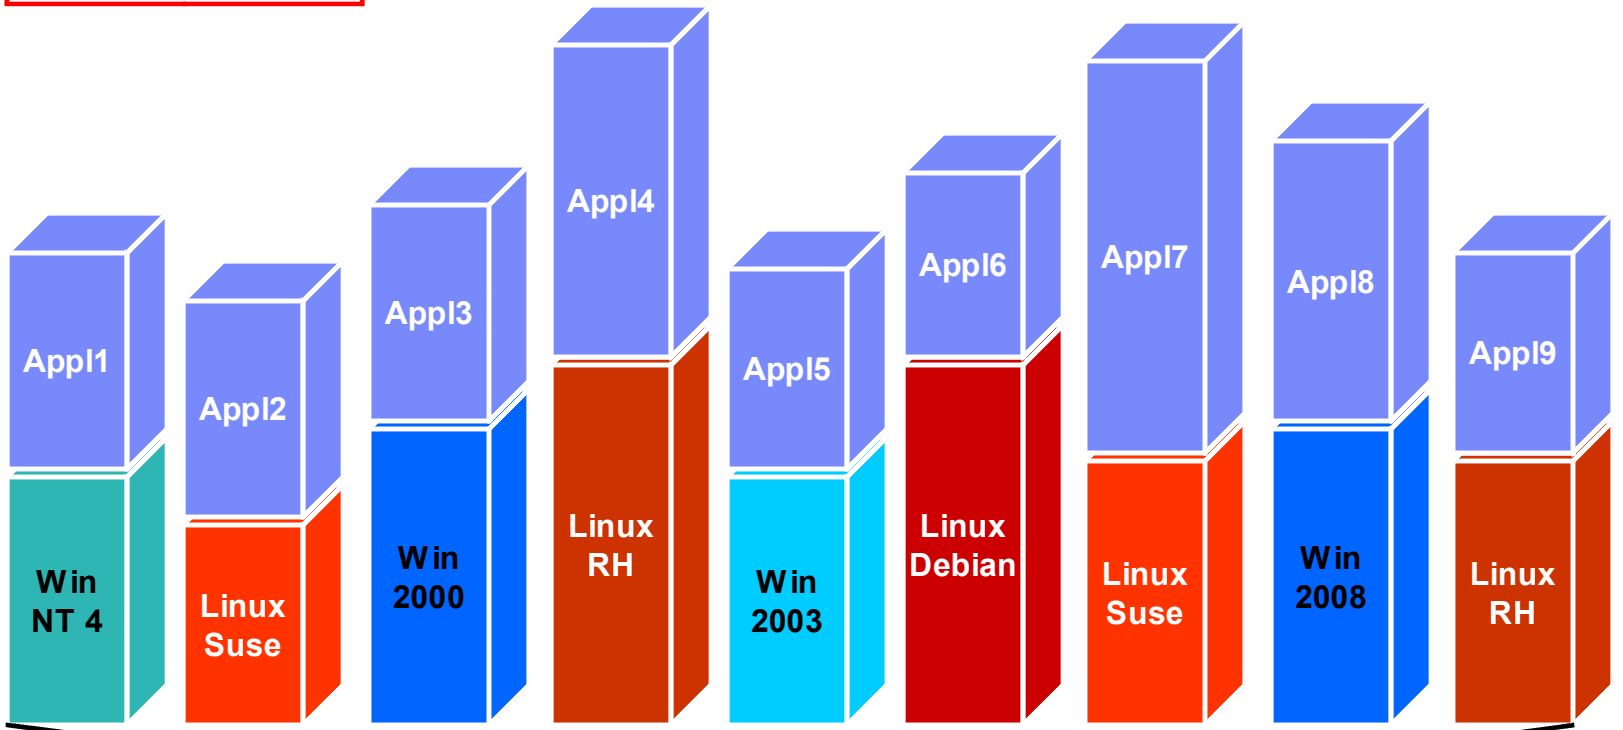

..and so on...

SW partitioning (i.e.VMware - Xen)

x4 architecture servers

System x

3850M2

3950M2

SAN

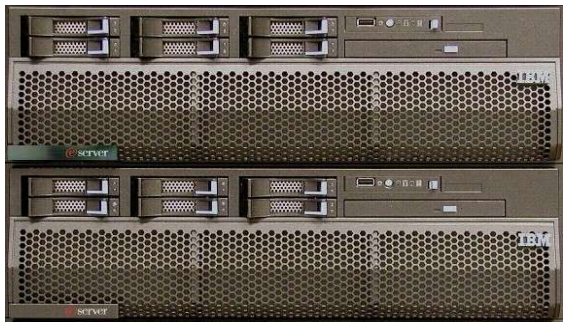

LAN Eth

Virtualization concept..

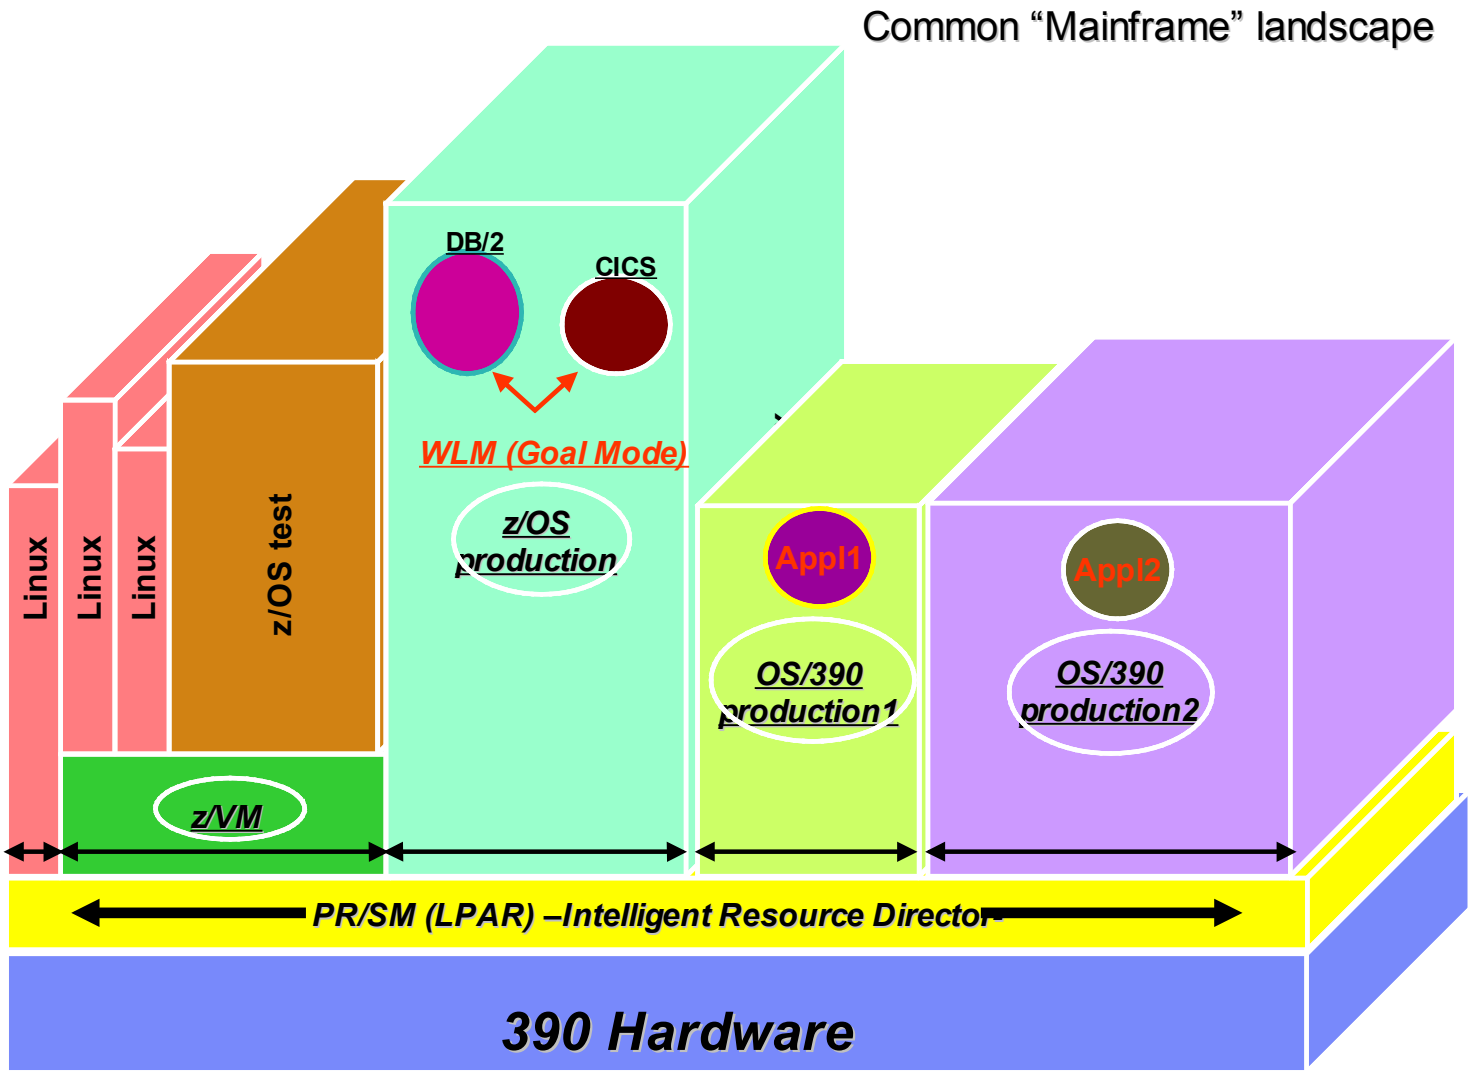

.x86 virtualization: VMware ESX Architecture

"standard" computer vs ESX architecture

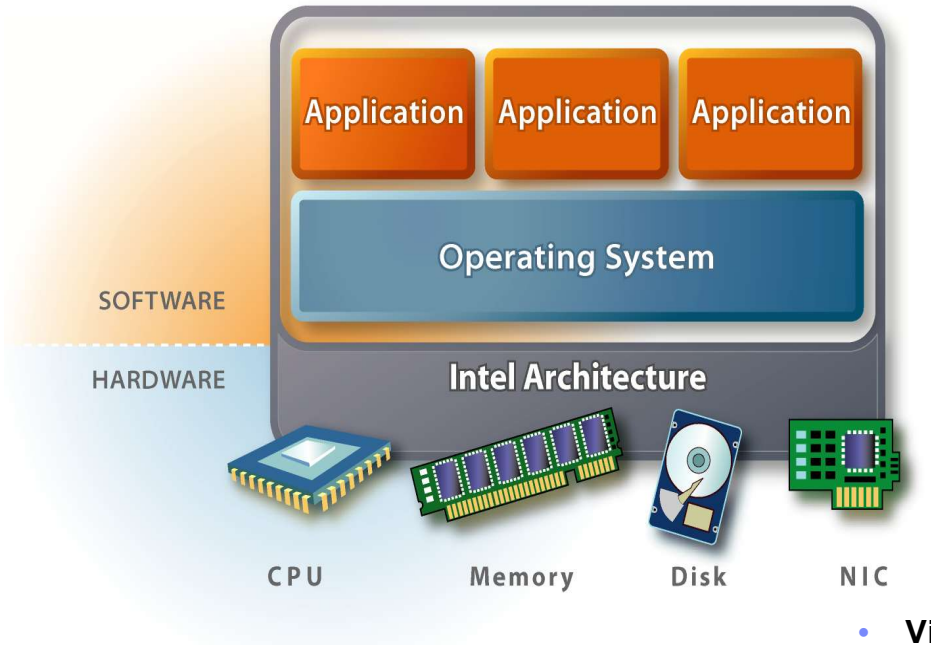

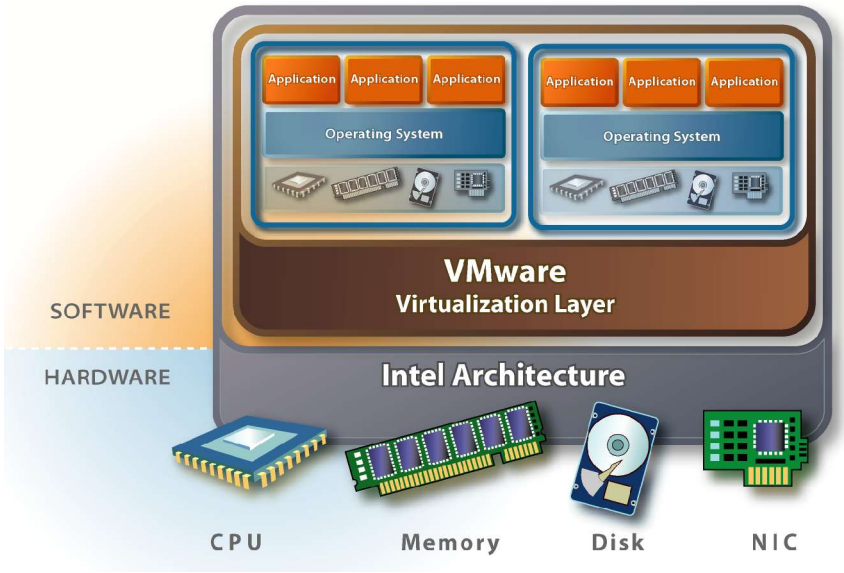

- **Virtualization layer maps virtual hardware to real hardware.**
- **Can multiplex several virtual hardware to single real HW.**
- **High Performance – map directly on hardware.**
- **Run multiple operating systems concurrently**
- **Fault, performance, security isolation**
- **Encapsulation**
- **Hardware-independent**

VMware vSphere 4

VMware vSphere 4.0 delivers new levels of performance and scalability, for running workloads in a virtualized environment:

- Expanded support for up to 64 cores per server
- Expanded support for up to 1 TB of physical memory per server
- Expanded support for up to 255 GB of RAM per virtual machine
- Expanded support for up to 320 virtual machines per server

vSphere IBM p/n now available:

Announce: 30 June 2009

Availability: 3 July 2009

"…the IBM System x3950M2 is the reference platform for vSphere scalability."

vSphere 4 Editions: Core Platforms

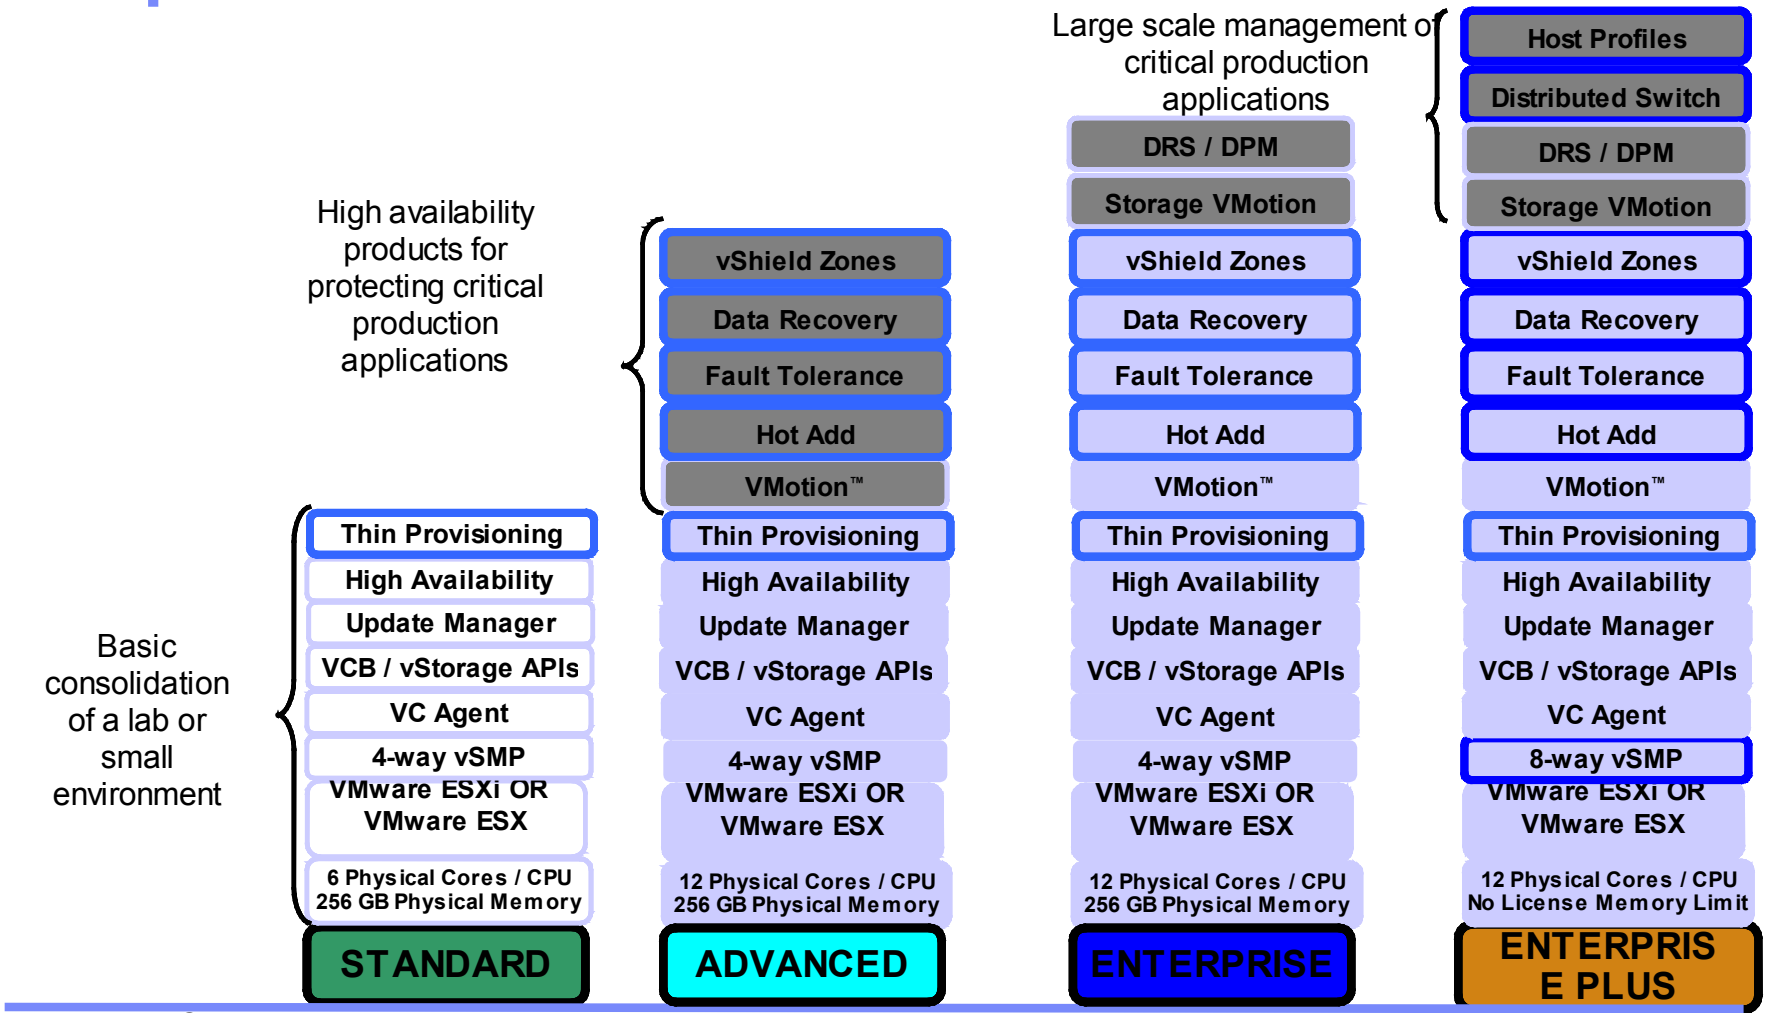

License (1 CPU)	$795	$2,245	$2,875	$3,495
SnS (Gold/Plat)	$275 / 325	$471 / 561	$604 / 719	$734 / $874

Note: US list pricing in USD – International pricing may vary

Note: Most features require one (1) instance of vCenter Server
Note: At least 1 year of Subscription and Support is mandatory

vCenter Server

vCenter Server is required for most vSphere features

VI Management	License
vCenter Server Standard -No host limit -Includes Linked Mode -Includes Orchestrator	$4,995
vCenter Server Foundation -3 ESX host limit	$1,495
vCenter Heartbeat	$9,995

- vCenter Server Standard and vCenter Server Foundation are management servers
- vCenter Server Foundation is limited to 3 nodes of any size / suite
- One (1) instance is required for most vSphere Standard, Advanced, or Enterprise Plus features

New! VMware vCenter Heartbeat

Ensures Availability for the VMware vCenter Server Management Platform

- Optimize availability and resiliency for VMware vCenter Server in any situation.
- VMware vCenter Server Heartbeat delivers the maximum uptime for your virtual datacenter infrastructure, ensuring consistent operation, even when VMware vCenter Server is threatened by unplanned or planned downtime.
- Protect VMware vCenter Server and its database against hardware, OS, application and network downtime.
- Ensure seamless failover and fail-back of VMware vCenter Server.
- Provide reliable disaster recovery for VMware vCenter Server in local and remote locations
- VMware vCenter Server Heartbeat has no hardware configuration dependencies and automatically detects standard VMware vCenter components upon installation, providing instant monitoring and protection.
- VMware vCenter Server Heartbeat delivers high availability and disaster recovery for VMware vCenter Server and all of its components across the LAN or WAN, including the database and licensing server and plug-ins like Vmware vCenter Update Manger eliminating costly, complex outages. VMware vCenter Server Heartbeat protects and recovers the VMware vCenter Server database instance, even if it's installed on a separate server.

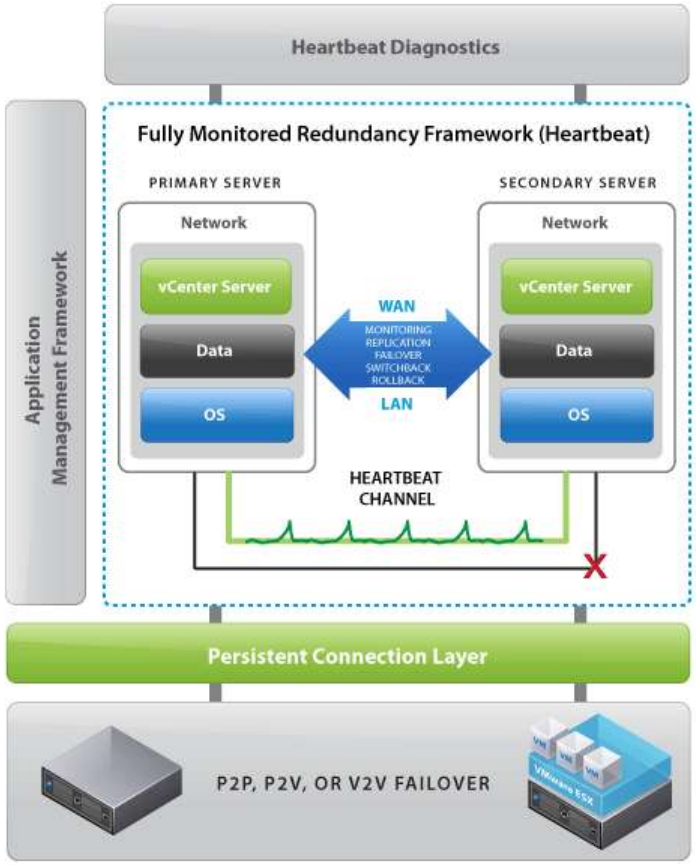

VMware vSphere Essentials and Essentials Plus Packages

"Always on IT" for mid-market businesses

Packages

	Data Recovery	x 6 CPUS
	High Availability	x 6 CPUS
Update Manager	Update Manager	x 6 CPUS
VC Agent	VC Agent	x 6 CPUS
4-way vSMP	4-way vSMP	x 6 CPUS
VMware ESXi OR VMware ESX	VMware ESXi OR VMware ESX	x 6 CPUS
vCenter Server for Essentials (up to 3 servers)	vCenter Server for Essentials (up to 3 servers)	x 1
ESSENTIALS ($879)	ESSENTIALS PLUS ($2,995)	

Architecture

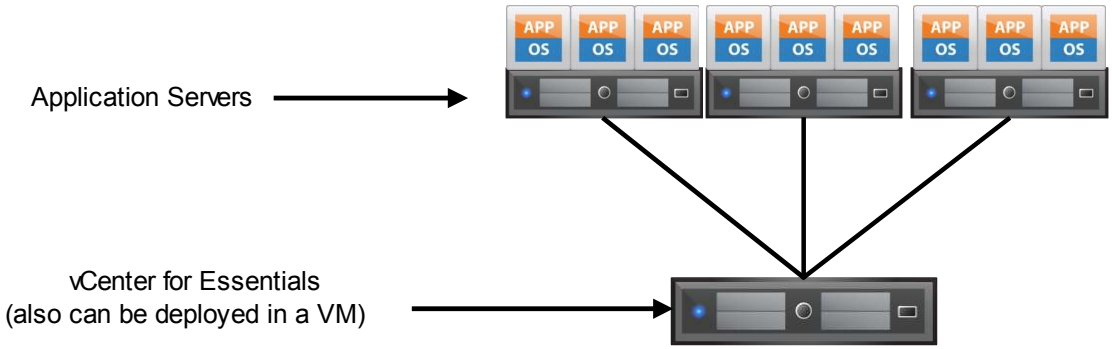

Features

- **High Availability:** Provides rapid and automated restart and failover of virtual machines without the cost or complexity of solutions used with physical infrastructure. For server failures VMware HA automatically and intelligently restarts affected virtual machines on other production servers.
- **Data Recovery** is a new disk-based backup and recovery solution that delivers quick, simple and complete data protection for your virtual machines. It's fully integrated with VMware vCenter Server to enable centralized and efficient management of backup jobs and also includes data de-duplication to save on disk storage for your backups.
- **All inclusive packages:**
  - **Licenses for 3 physical servers (up to 2 processors each)**
  - **License for central management server**
- Packages cannot be decoupled or combined; components cannot be used with other vSphere editions

EVOLUTION + **REVOLUTION** = **INNOVATION**

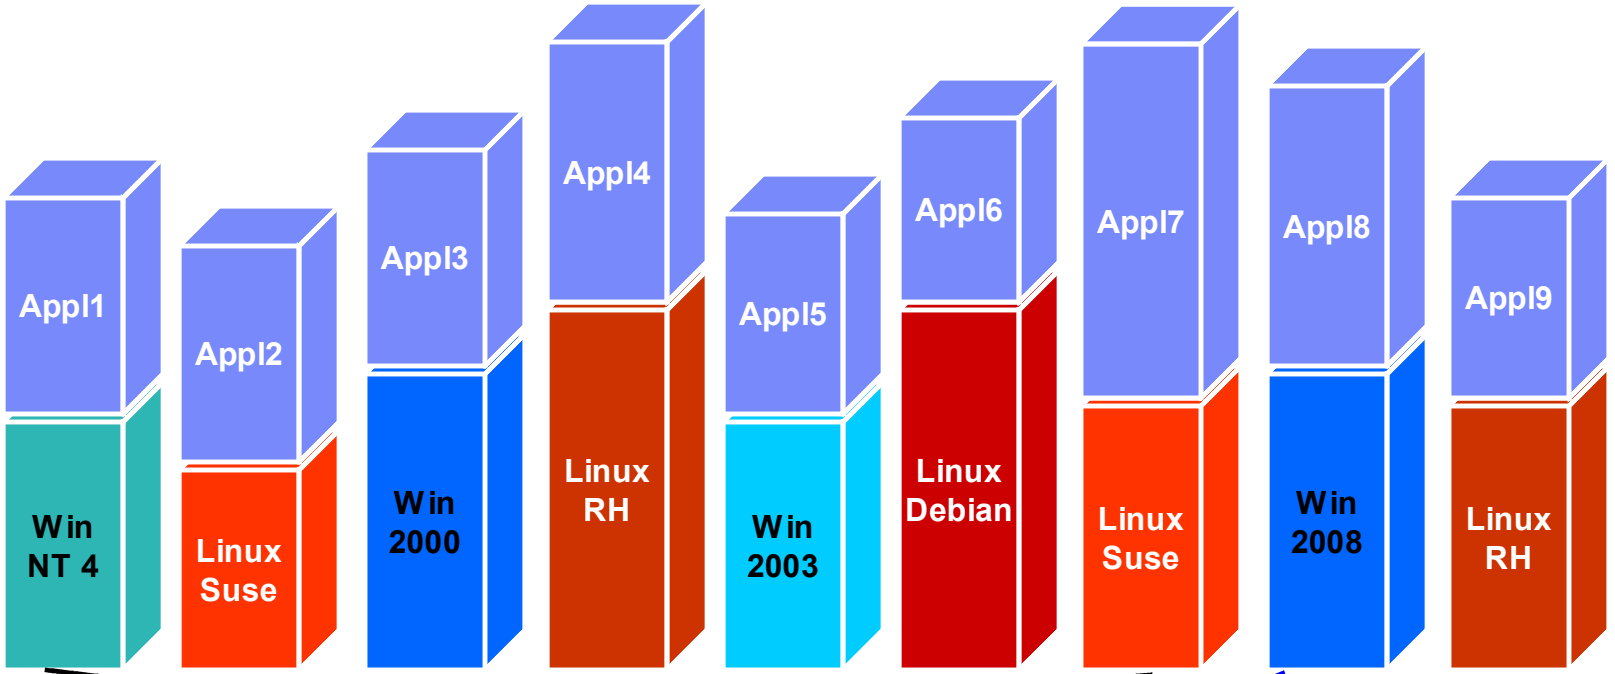

What's a "BladeCenter" ? and what's a "Blade"?

A "server on a card" - each "Blade" has its own:

- ➔processor
- ➔ethernet
- ➔memory
- ➔optional storage
- ➔etc.

IBM Blade ready for insertion into the BladeCenter

The chassis provides shared:

- ➔management console (KVM)
- ➔power supply
- ➔cooling
- ➔network switches
- ➔CD-ROM drive
- ➔diskette drive
- ➔etc.

IBM BladeCenter chassis - 7U rackable

YOU CAN ALWAYS COUNT ON FAMILY. ESPECIALLY THIS ONE.

THE IBM BladeCenter FAMILY





IBM BladeCenter S
- 6 blades, 7U
- Extra internal storage room
- Ideal for SMB server consolidation

IBM BladeCenter E
- 14 blades, 7U
- Mainstream applications
- Ideal for data centers and remote sites

IBM BladeCenter H
- 14 blades, 9U
- Ideal for extreme I/O, data-intense environments
- Excellent platform for virtualization or high perf requirements

IBM BladeCenter T
- 8 blades, 8U
- NEBS characteristics
- Ruggedized chassis
- Ideal for telco, military, medical-imaging applications

IBM BladeCenter HT
- 12 blades, 12U
- NEBS characteristics
- Ideal for services, control and transport planes
- Ruggedized chassis

Common Blades, Common Switches


Blade Center Chassis - front view

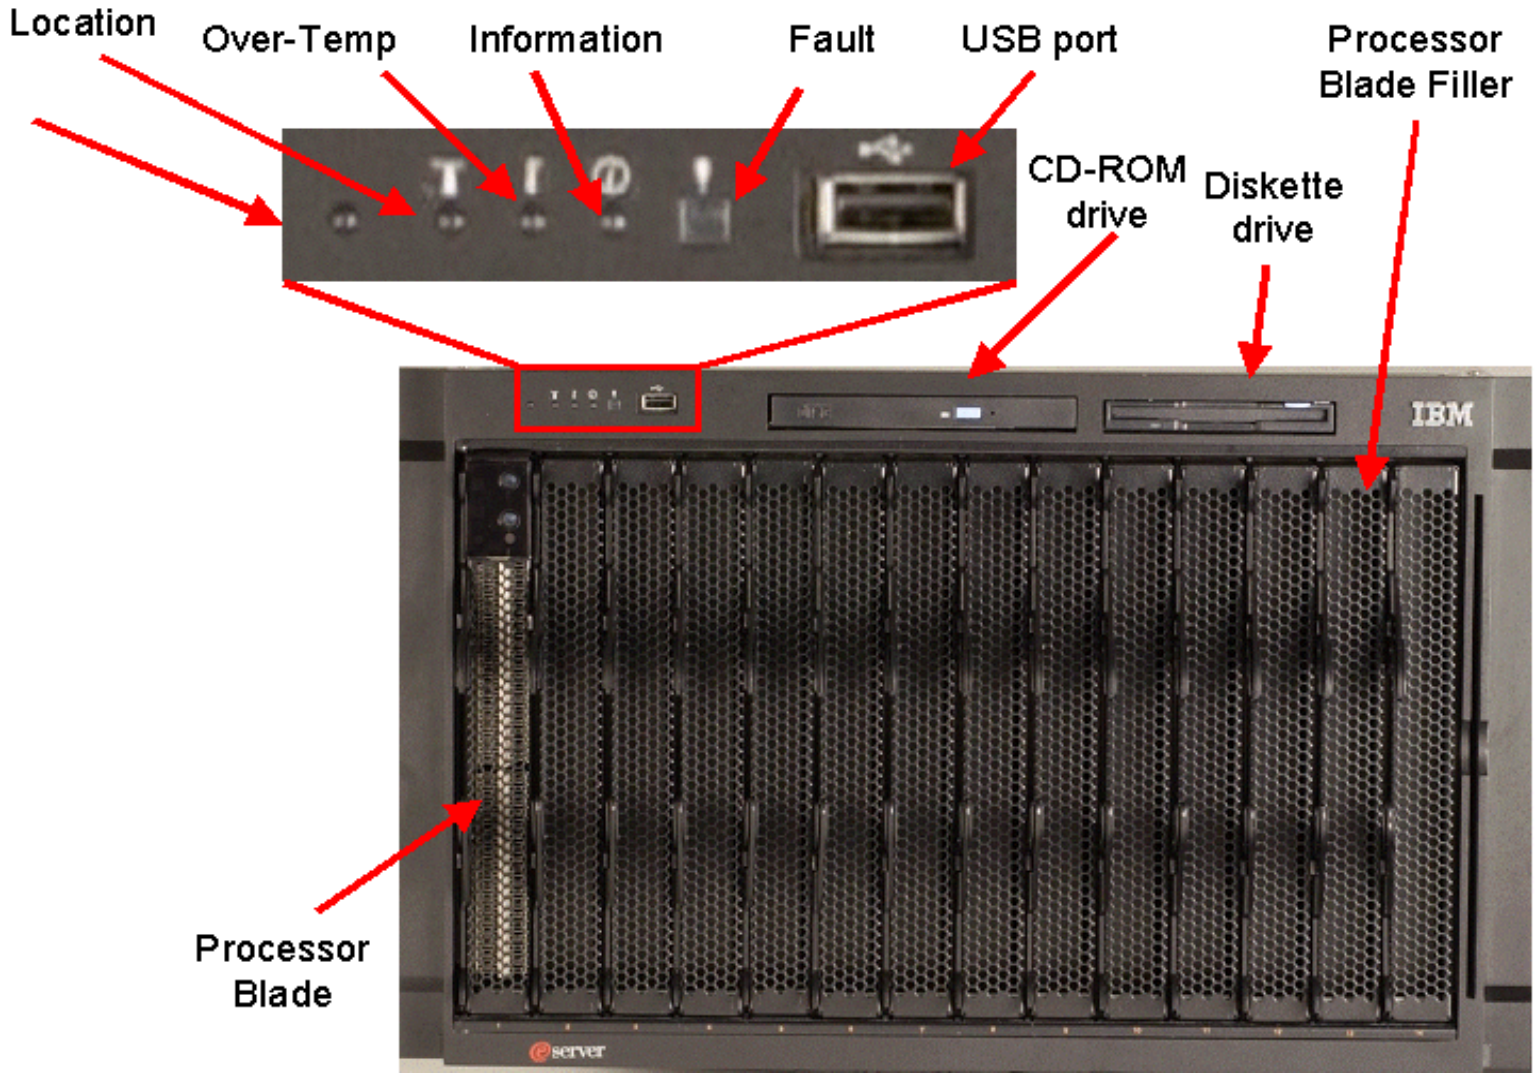

BladeCenter technology advantage: *dual midplane*

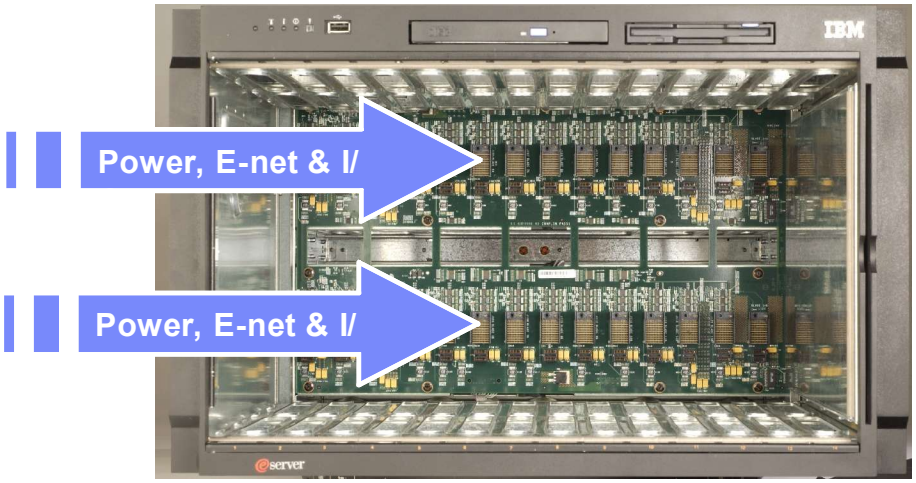

- At least two connections from each blade server to the midplane
- Dual communication paths to the *passive* midplane for Ethernet, Fibre Channel, KVM, Power, and Management signals
- Chassis can be configured with dual Ethernet, Fibre Channel Modules providing two active paths to your external network
- Multiple paths from blade to chassis components helps to protect you from potential failed connectors or traces

BladeCenter E chassis – rear view

- **Gigabit Ethernet Switches**
  - Portfolio of switches (Cisco, BNT)
  - Lower cost via Integration
  - Functions range from Layer 2 thru Layer 7
- **Fibre Channel Switches (FC Fabric)**
  - Portfolio of Switches (Qlogic, Brocade, MCdata)
  - Potentially lower cost via integration
  - Full support of FC-SW-2 standards
- **Power Subsystem**
  - Upgradeable as required
    - Now 2300 Watts/each;
    - Redundant and load balancing for high availability
- **Calibrated, Vectored Cooling™**
  - Highly fault tolerant
  - Allow maximum processor speeds
- **BladeCenter Management Modules**
  - Full remote video redirection
  - Out-of-band / lights out systems management
  - Concurrent Serial connectivity

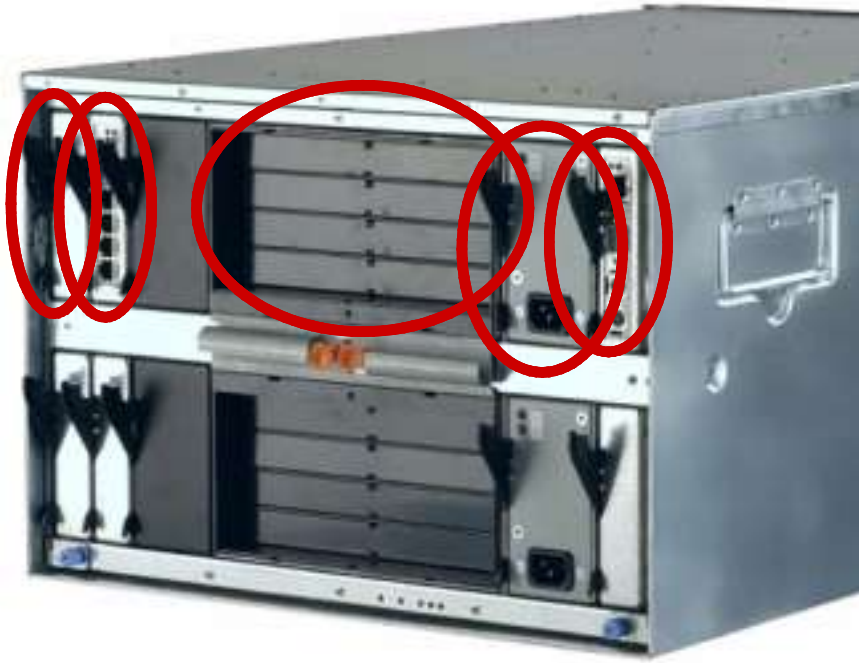

Internal Switch Connectivity - Ethernet

External Network Connections:
4 10/100/1000 Gigabit ports

Internal Ethernet Connections:
14 Gigabit ports (1 per blade)

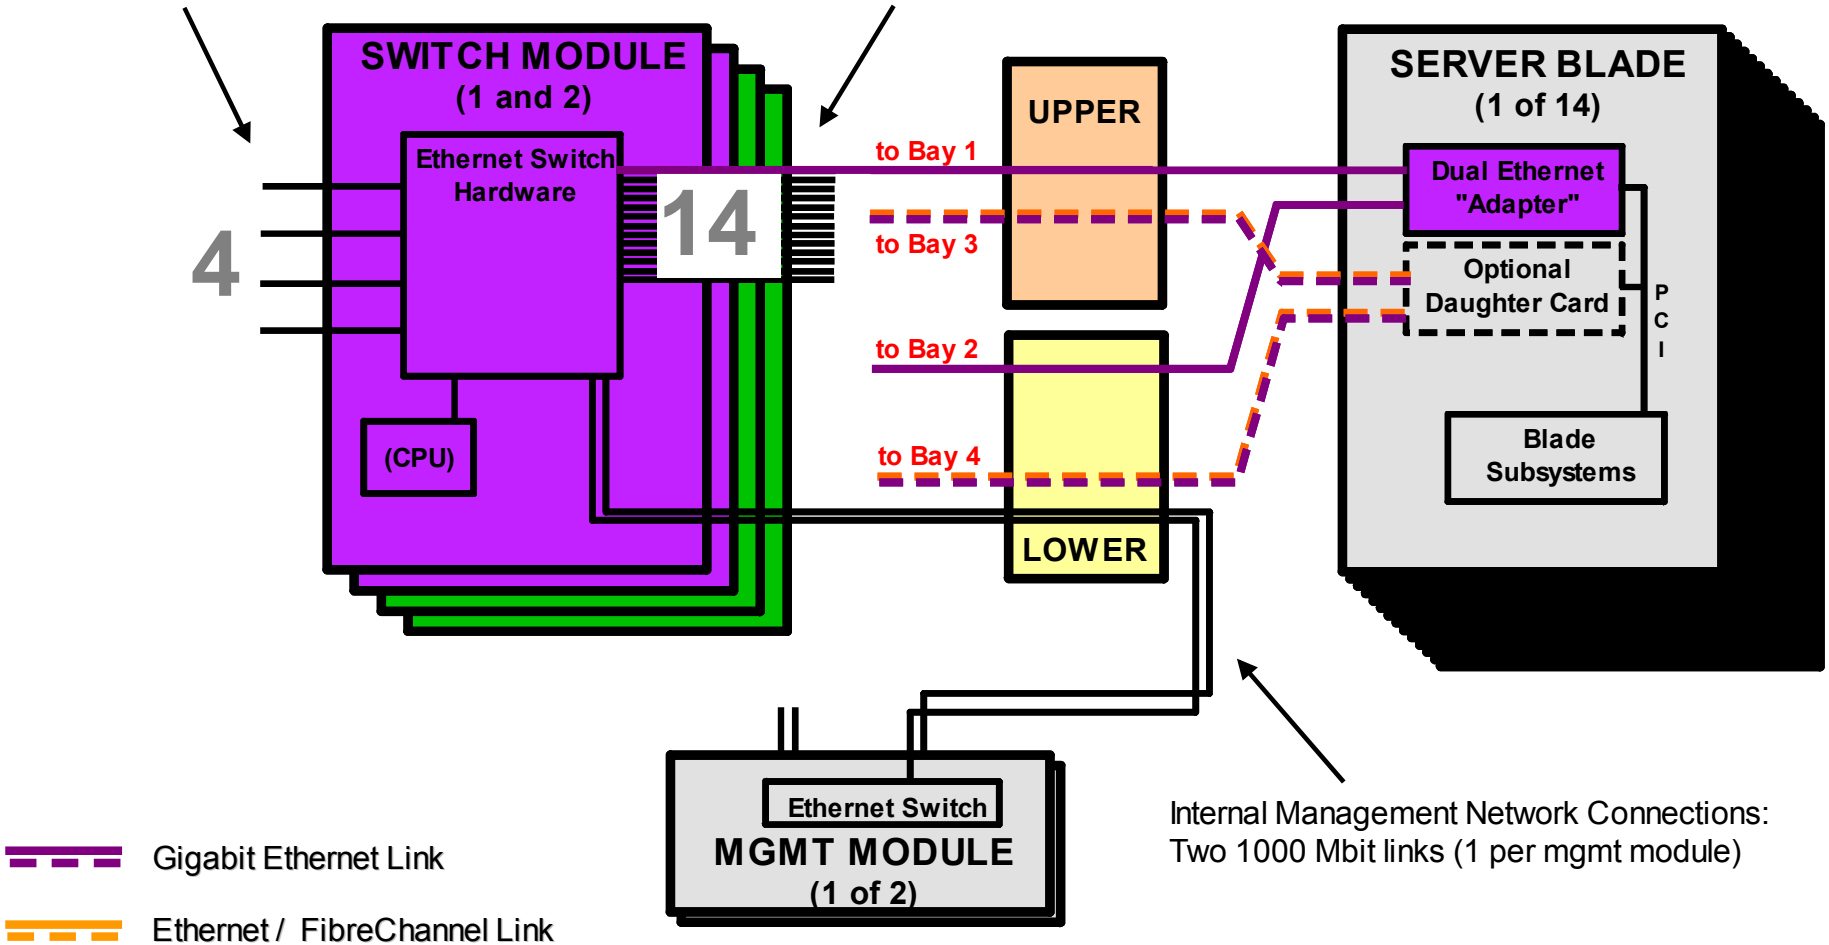

BladeCenter Ethernet Components

Cisco Systems® Intelligent Gigabit Ethernet Switch Module

BNT® Layer 2/3 (Fiber) Gigabit Ethernet Switch Module

Server Connectivity Module for IBM BladeCenter

BladeCenter Telco

BNT® Layer 2-3 Gigabit Ethernet Switch Module

BNT® Layer 2/3 10GbE Uplink Switch Module

- (1) 10 Gb MM Fiber Ports
- (2) 10 Gb Copper Ports

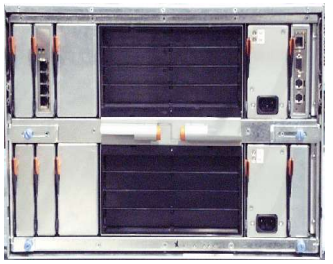

BladeCenter

Cisco Systems® (Fiber) Intelligent Gigabit Ethernet Switch Module

BNT® Layer 2-7 Gigabit Ethernet Switch Module

Intelligent Copper Pass-thru

BladeCenter H

Cisco Catalyst 3012/3110

- Integrates Cisco Catalyst networking technology into BladeCenter – under 'Catalyst' umbrella
- Three (3) new Cisco Switch Modules
  - 1Gb Switch without Stacking (***3012***)
  - 1 Gb Switch with Stacking (***3110g***)
  - 10Gb Uplink Switch with Stacking (***3110x***)
- Software Keys
  - Base Services - layer 2, static routing, RIP – included with all 3 switches
  - IP Services - Layer 2/3, OSPF, BGP (optional for 3110)
  - Advanced IP Services - IPv6 (optional for 3110)
- First to market with next generation switch technology
- First blade solution in industry with embedded Cisco switch stacking: ***Virtual Blade Switch*** technology (VBS)

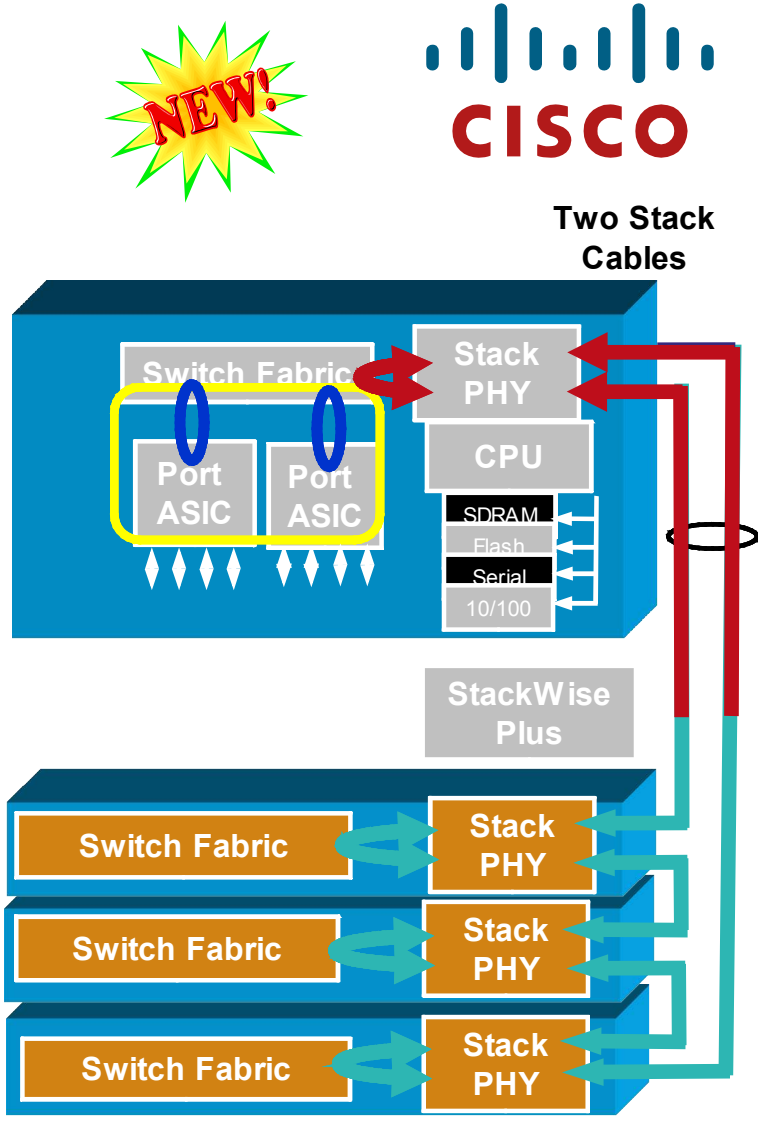

BNT® 1/10 GB ETHERNET SWITCH

- **Designed for IBM BladeCenter form factor**
- **Full Layer 2 Switching and Layer 3 Routing - Stackable**
- **Ports**
  - **22x 1G copper RJ-45**
    - **6x 1G RJ45 External uplinks**
    - **14x 1G Internal server links**
    - 2x Internal MM ports
  - **3x 10G ports**
    - **3 SFP+ uplink ports (SR/LR)**
- **Line rate performance with no packet drop**
- **Power**
  - **Meets IBM BladeCenter requirements: less than 45W**

Intelligent Copper Pass Through Design

Requirements

- Simpler cabling
- Support 100/1000 switching
- Support SOL and cKVM with pass through
- Maintain low cost
- Support across all chassis

Technical offering

- **Mode 1: 14-port copper pass-thru mode**
  - Direct 1-to-1 mapping
  - 100/1000 Mbps operation on uplinks
  - Setup & monitoring via aMM
- **Mode 2: Server Connectivity Mode**
  - Similar function suite as existing IBM SCM
  - 14 internal ports to 7 external uplinks
  - Simplified Layer 2 interconnect within chassis
  - Simple GUI and CLI for setup & configuration
  - No-loop design for external uplinks
  - Setup & monitoring via direct access (CLI or Browser)

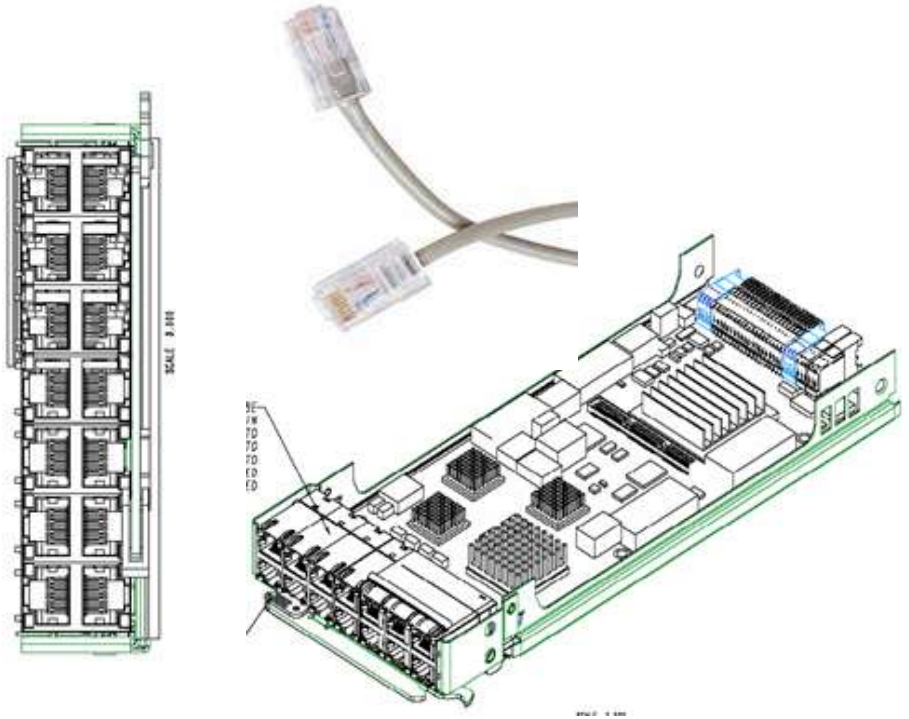

BladeCenter SAN Components

QLogic® 10 and 20 ports 4Gb Fibre Channel Switch Module

- NB: be aware, 10-port upgrade now only available from Qlogic support

McDATA® 10 and 20 ports 4Gb Fibre Channel Switch Module

Brocade® 10 and 20 ports 4Gb SAN Switch Modules

Cisco® 10 and 20 ports 4Gb Fibre Channel Switch Module

- SANOS
- Utilizes Cisco SFPs (also IBM p/n)

QLogic® Intelligent Pass-thru Module

- 6 external shared connections - 4Gb I/O
- NPIV technology (N_Port_ID Virtualization)
- Presents multiple N-Ports to the fabric
- Needs external switch NPIV capable

Optical Pass-thru Module

- Supplier: IBM
- Provides unswitched / unblocked optical connection
- Up to 14-optical connections to external SAN (requires breakout cable option)
- Max 2Gb bandwidth

BladeCenter SAN Components - new

QLogic®
20 ports 8Gb Fibre Channel Switch Module
(P/N 44X1905)

- 20-port 8 Gb Fibre Channel Switch with 14 internal/ midplane ports for server connectivity and six external/ SFP enabled ports for target or fabric connectivity
- 14 internal copper SERDES connections to the blade servers with F-port 2 Gb and 4 Gb autodetect speed types
- Six external client-accessible F/FL/E, 2, 4, and 8 Gb autodetect SFP ports

QLogic®
8Gb Intelligent Pass-thru Module (P/N 44X1907)

- 6 external shared connections - 4Gb I/O
- NPIV technology (N_Port_ID Virtualization)
- Presents multiple N-Ports to the fabric
- Needs external switch NPIV capable
- Eliminates E_Port Switch to switch incompatibility

BladeCenter H Tour
A Look At The Front

Server Blades
- Same Blades
- Same I/O Feature Cards
- Same Server Blade Features
- Same Dual Slot PCI adapter 'sidecar'

- BC-H and new Blades add:
  - Additional Blade Power/Thermal Capacity
  - Additional High Speed I/O Options
    - – PCI-Express x8 chipset link
    - – 4 4x Switch Module links

Power Modules (2 or 4)
- Hot Swap, Redundant
- 200-240V 50/60Hz AC (worldwide)

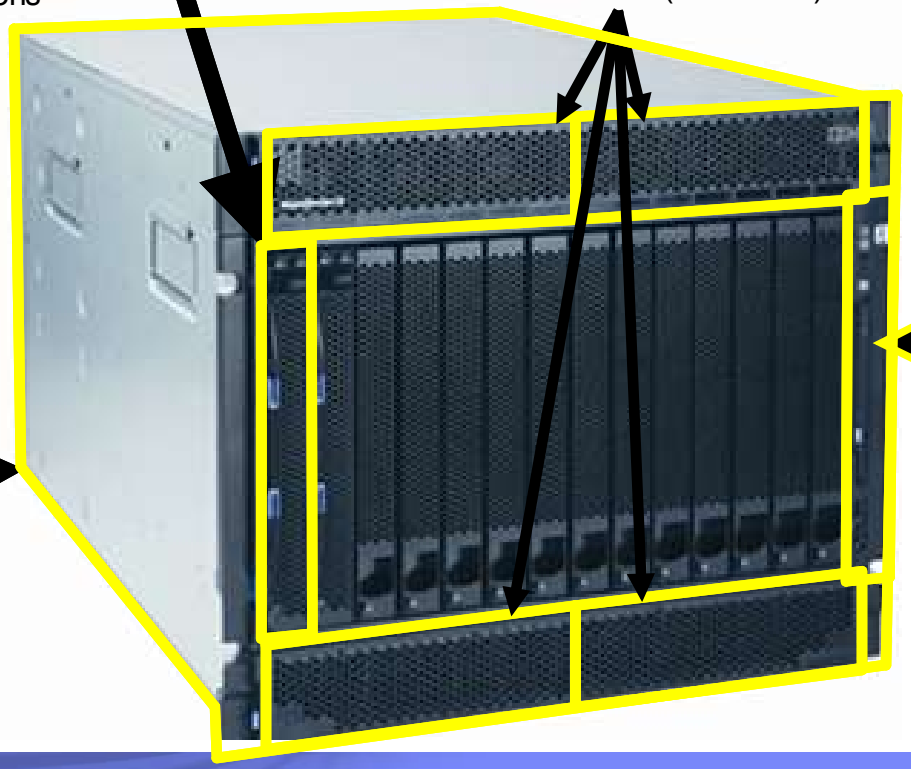

Op Panel & Media
- Chassis level LEDs-
  - Power, Alert, Info
  - Chassis 'Locate'
- 2 USB Ports
- Removable storage media
  - DVD

BC-H Chassis
- 18 inch rack mount
- Front to rear airflow
- Front/rear service
- Rear cabling
  - 14 Server Bays
  - 9U high, 28" deep

BladeCenter H Tour
What is Where?

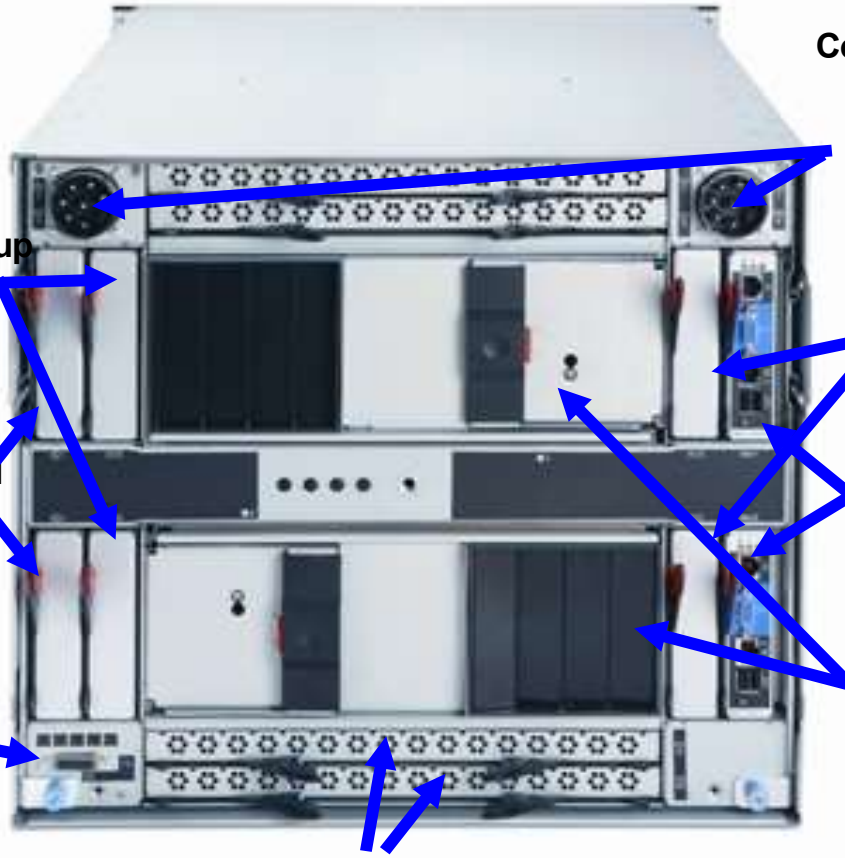

IBM BladeCenter H Midplane Internal Detail

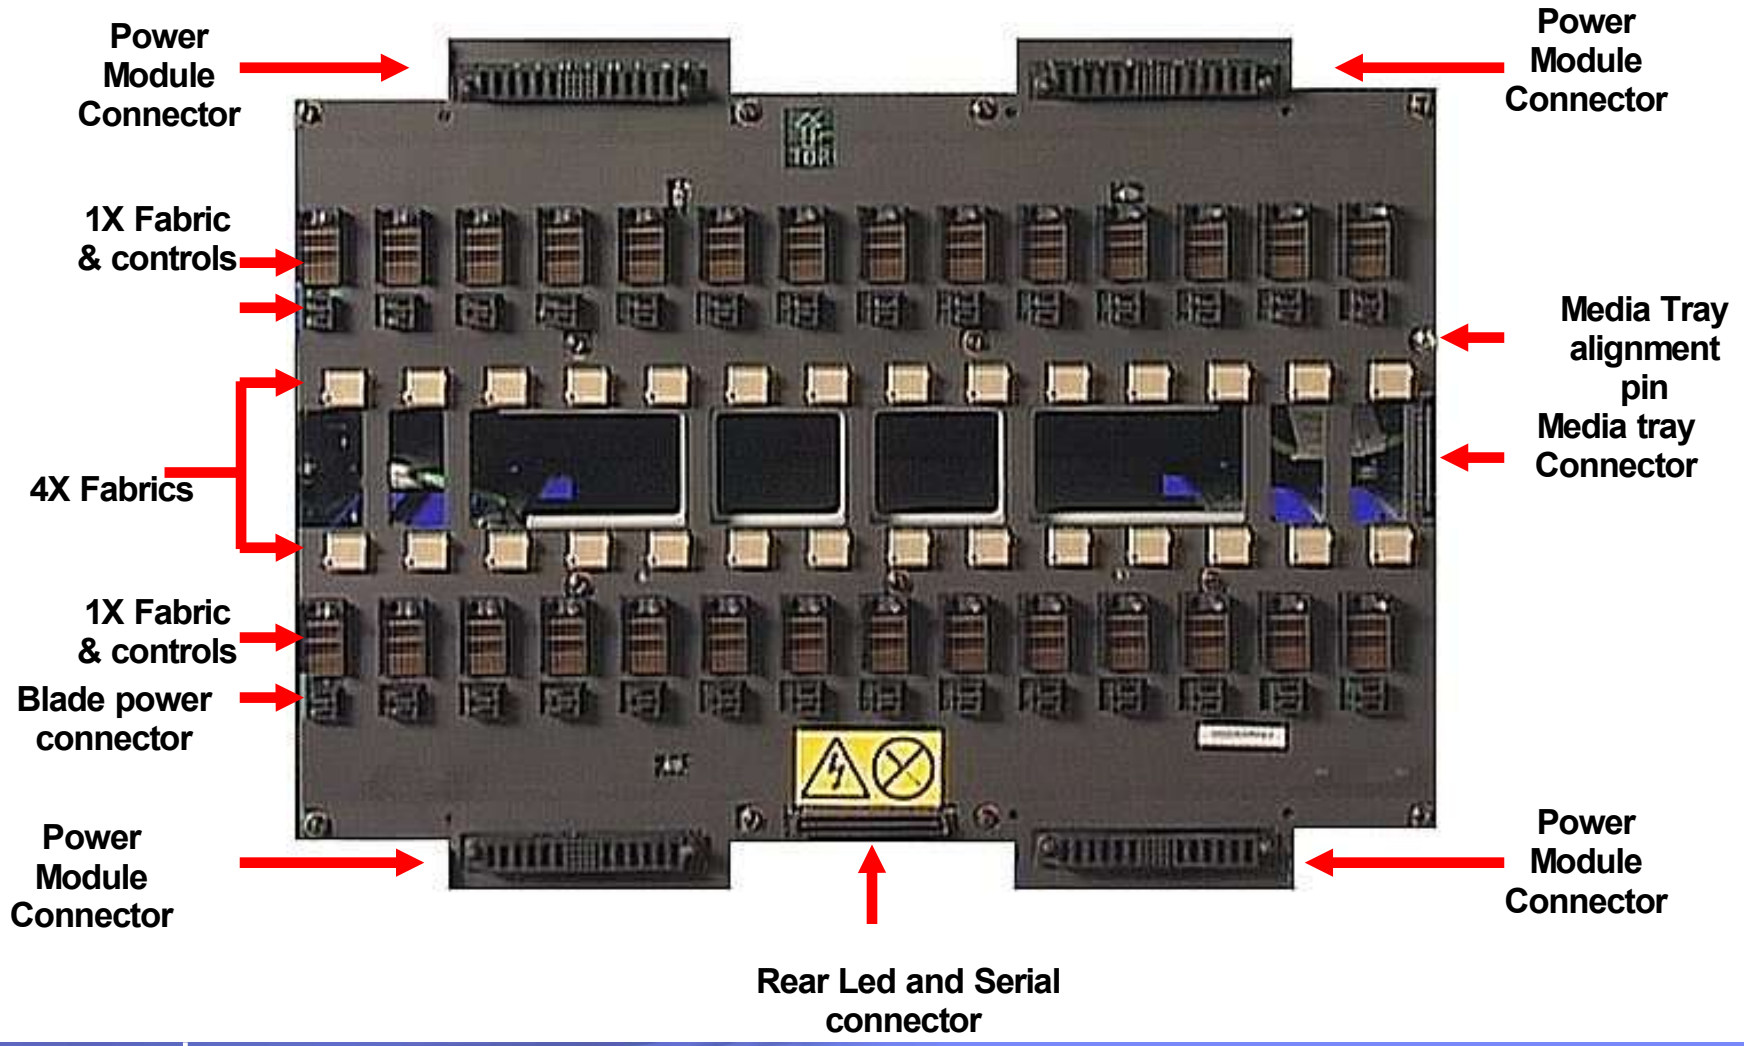

IBM BladeCenter H – connection paths

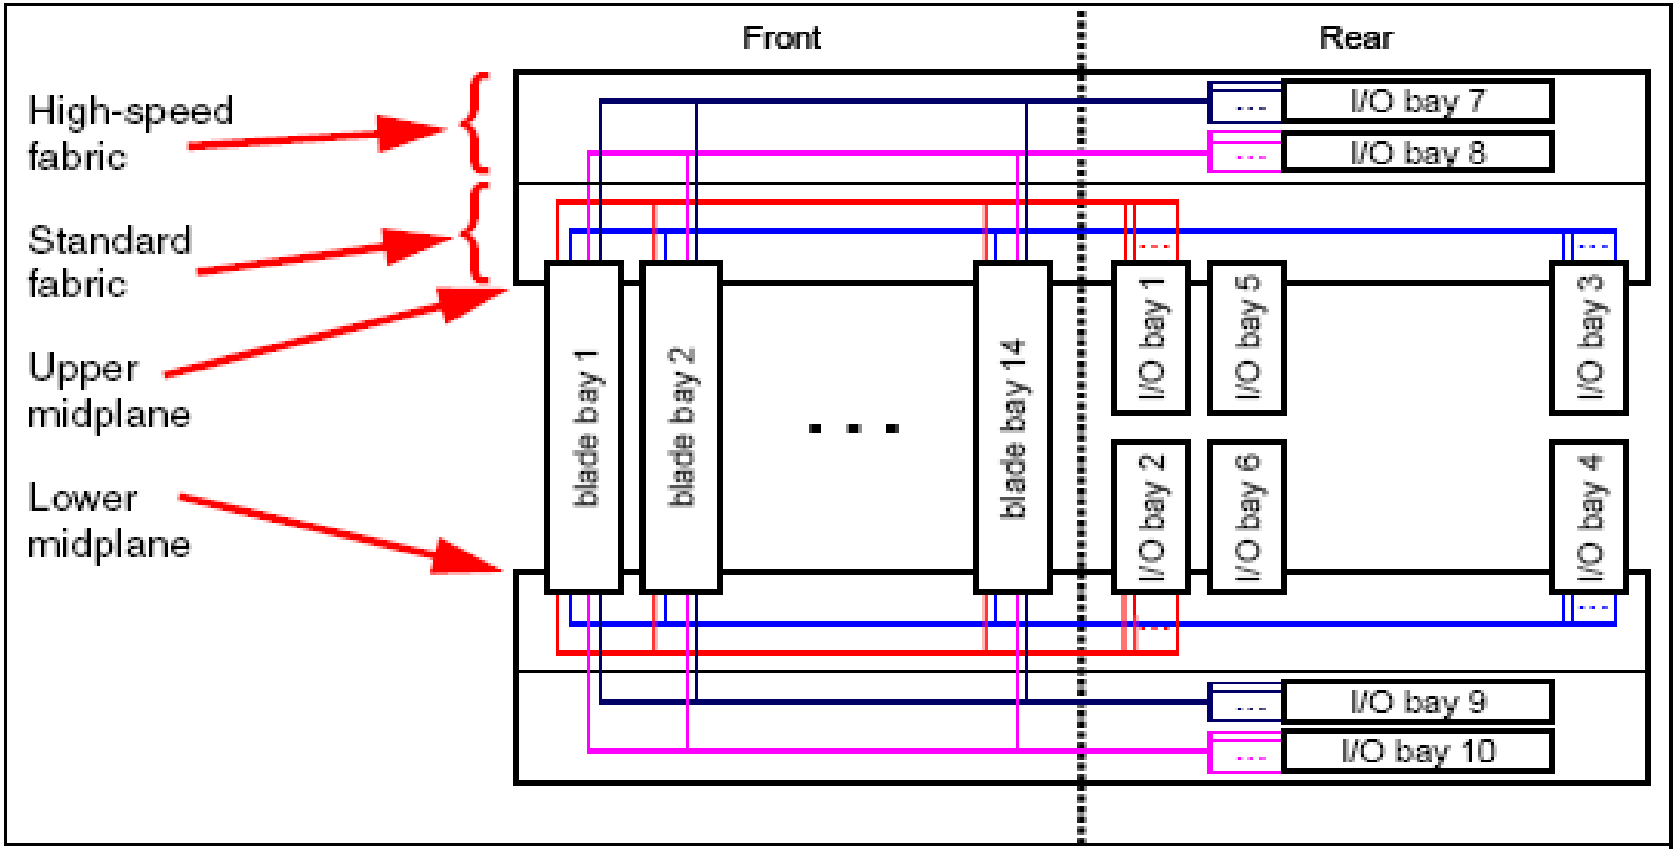

BladeCenter H – High Speed Network

- **Support for high speed switches**
  - **4X InfiniBand,**
  - **10G Ethernet**
- **Support for high speed DC's on blades**
  - **4 high-speed fabrics**
  - **Still have access to legacy fabrics**

InfiniBand on BladeCenter H

..first announced component of

Virtual Fabric Architecture

- **IBM and Cisco jointly developed a 4X (10Gb) InfiniBand solution for BladeCenter H**
  - Daughter Cards: Provide dual 4X connectivity to high speed switch modules - use PCI-Express (PCIe) connection on next generation Blades (e.g., JS21, HS21)
  - InfiniBand Switch Module: (14) 4X ports interfacing to blades (with daughter card) and (2) 4X and (2) 12X (30Gb) ports to network
- **Virtualized I/O via VFrame (Cisco) software**
- **Ship Support: July 25**

What can BladeCenter H be used for?

Installation of a 4X IB card and 4X switch allows blades to be added to high speed IB fabric. 80G per switch of bandwidth.

Customers can build a complete 4X IB fabric using only internal BladeCenter switches. Allows for redundancy, and high bandwidth.

For the ultimate virtualized solution the bridges inside the chassis can be used to deliver traditional ethernet and fibre connections at the rack level.

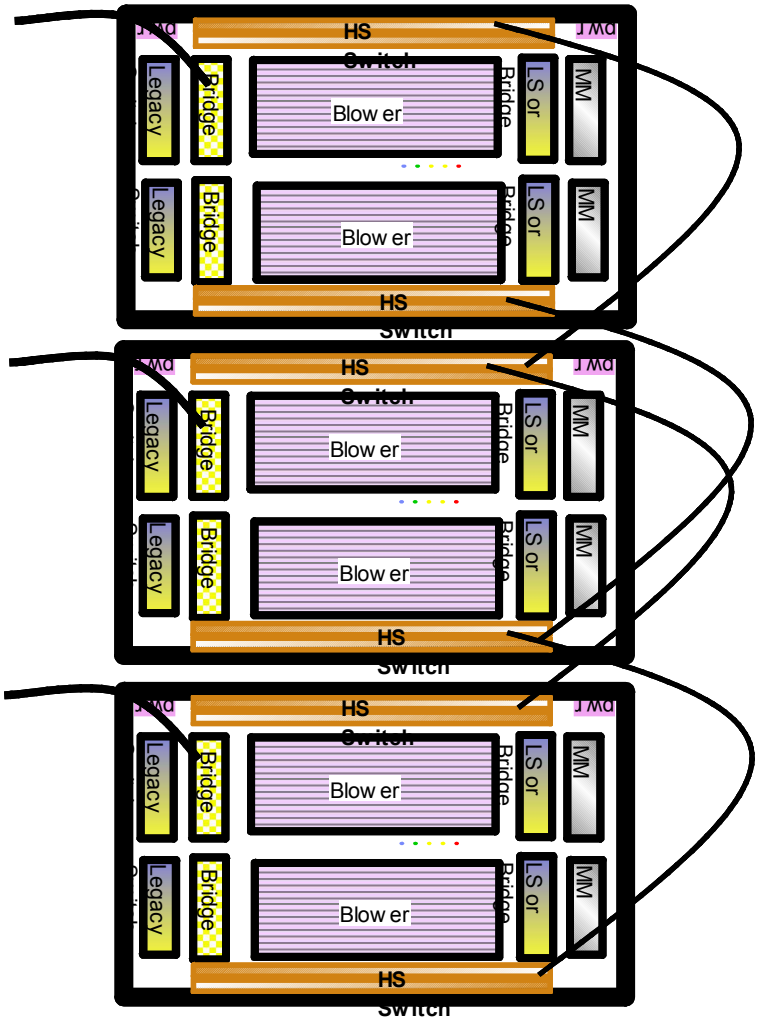

Introducing the BladeCenter BNT 10G Ethernet

Unsurpassed 10G Throughput

Ann: 26/05/09

GA: 17/06/09

Figure 1. BNT 10-port 10Gb Ethernet Switch Module for IBM BladeCenter

NEW!

Figure 1. The 10Gb Ethernet Pass-Thru Module for IBM BladeCenter

QLogic 2-port 10Gb Converged Network Adapter (CFFh)

Broadcom 2-port 10Gb Ethernet Expansion Card (CFFh)

Broadcom 4-port 10Gb Ethernet Expansion Card (CFFh)

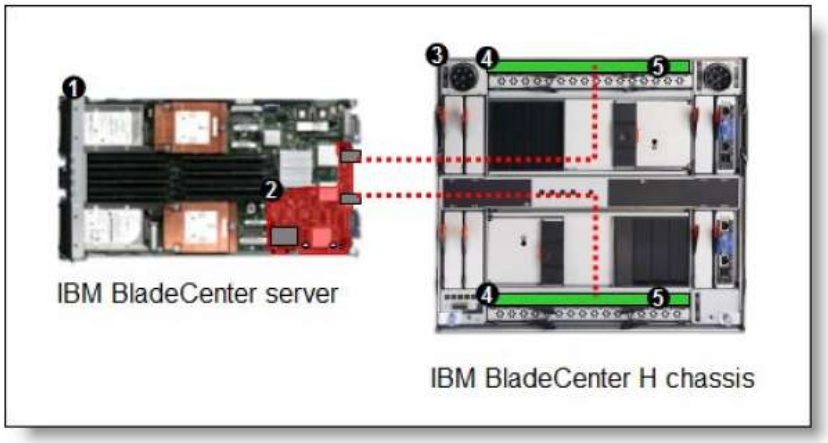

Figure 2. A 20Gb solution using two BNT 10-port 10Gb Ethernet Switch Modules

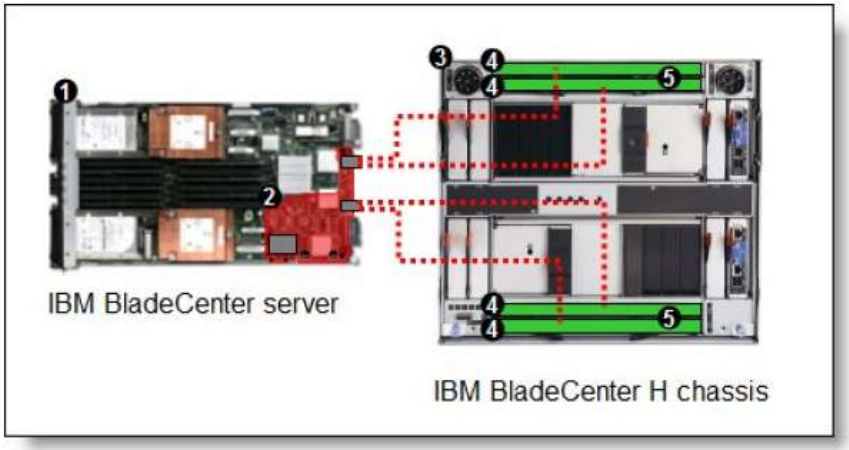

Figure 3. A 40Gb solution using four BNT 10-port 10Gb Ethernet Switch Modules

BNT 10-port 10G ESM

BLADE NETWORK TECHNOLOGIES

	Date
Announce	**May 19th**
GA	**Early June**
Part #	**46C7191**
IBM ESP	**$10,150**

- **HSSM Form Factor supported by IBM BladeCenter H and HT**
- **Ports**
  - **24x 10G ports**
    - **10x 10G SFP+ External uplinks**
    - **Ports can be configured as 10G SFP+ or 1G SFP**
    - **Ports support low cost DAC (direct attached cables) or twinax**
      - **Up to 1/7th the cost of transceiver and up to 1/10th the power**
  - **14x 10G Internal server links**
    - **Server links can be configured as 10G or 1G depending on the NIC card used**
    - **Up to 8x10G can be connected to bridge module for CEE***
  - **2x Internal MM ports**
- **Line rate performance with no packet drop**
- **FCoE** Ready™ **->** Support IEEE **CEE** (Conv Enh Eth)
- **Power -Meets IBM BladeCenter requirements: ~75-80W**
- **Future software enhancements planned**
  - **3Q09 – SmartConnect with VMready &vNIC™**

NIC	PN
Broadcom 2-port 10Gb Ethernet Exp.Card	44W4466
Broadcom 4-port 10Gb Ethernet Exp. Card	44W4465
Chelsio 2-port 10GbE Exp.Card * Cluster BOM	46M1817
2/4 Port 1G Ethernet Expansion Card	44W4479

Introducing VMready

NEW!

BLADE NETWORK TECHNOLOGIES

- VMready is switch-resident software – no stealing cycles from servers
- VMready automates network settings for VM movement
  - Ensures network settings migrate with the VM
- VMready is open
  - Works with all leading VM providers
    - VMware VMotion, Xen XenMotion, Microsoft Hyper-v Live Migration, Sun, Oracle, etc.
  - Standards-based – avoiding expensive vendor lock-ins
- Interfaces to management platforms
  - Export data on VM usage for accounting, change management, etc.

VMready keeps VMs connected and protected as they move

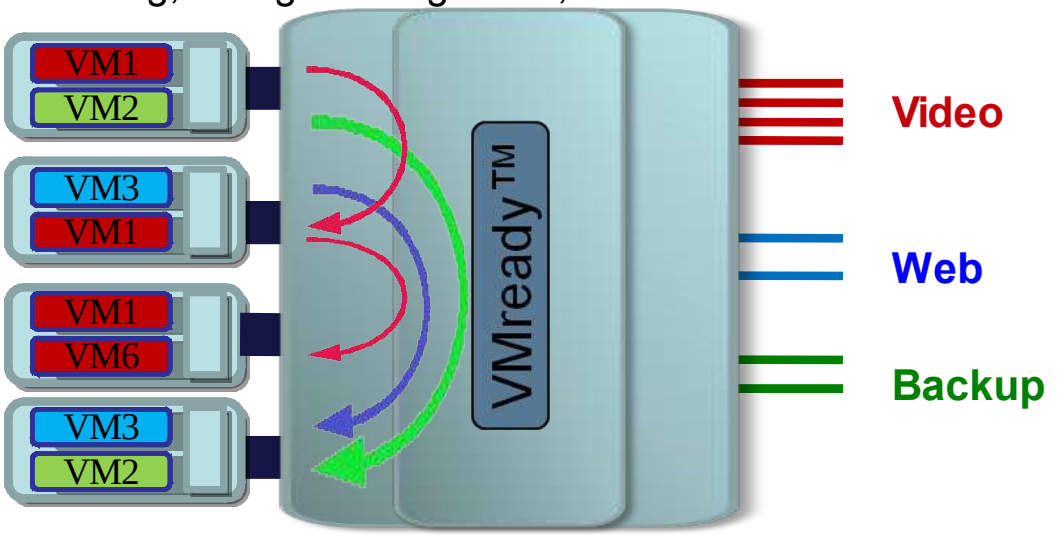

VMready in Action

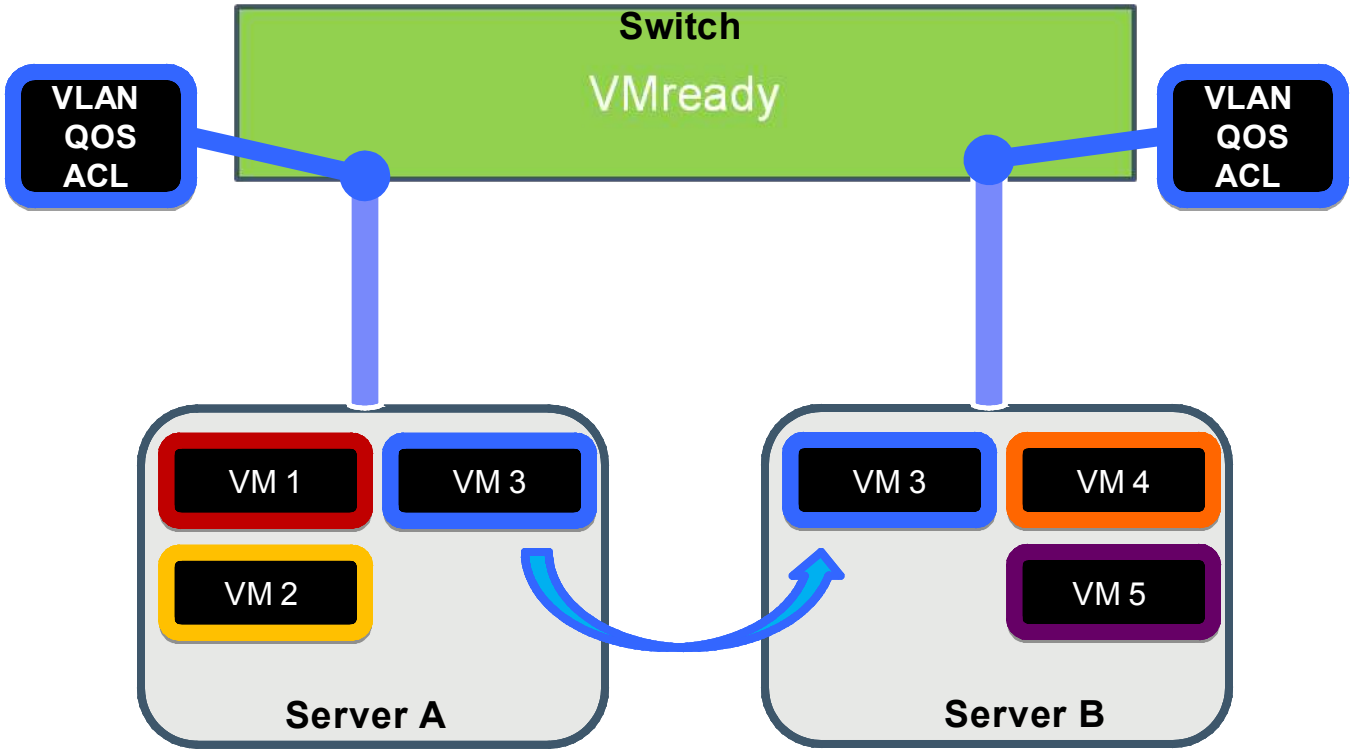

- Specific Virtual Machine (VM) network characteristics are associated with the switch “virtual” port.
- VMready detects a Vmotion process and “moves” the network characteristics of the VM to the new physical switch port.

SFP+ → Gaining Industry Momentum

- Flexibility in every port – choose as you go
  - **Low cost using copper cable – perfect for in-rack use**
    - **Known as "Direct Attach Cable" or "Twinax"**
  - **Longer distance using proven optical transceivers**
  - **Mix and match – 1G, 10G, copper, fiber – as you need**

SFP+ Copper (Twinax) cable

SFP+ Transceivers
- IBM 10GBase-SR SFP+ (PN 44W4408)
- BLADE 10GBase-LR SFP+ (BNT PN BN-CKM-SP-LR)

SFP+ Copper Direct Attach Cables (up to 10m) – no transceiver
- 1M cable (BLADE P/N: BN-SP-CBL-1M)
- 3M cable (BLADE P/N: BN-SP-CBL-3M)
- 7M cable (BLADE P/N: BN-SP-CBL-7M)
- 10M cable (BLADE P/N: BN-SP-CBL-10M)

IBM BladeCenter **Virtual Fabric Solution** – Components

Flexibility to support multiple workloads..1Gb, 10Gb, virtual NIC, IO Convergence

1. IBM BladeCenter H & HS22 blades
   - BladeCenter Open Fabric Manager address management
   - Future BOFM plug-in of BNT virtual NIC management

2. Emulex Virtual Fabric Adapter (CFFh) – PN 49Y4235 (Ann. 20 Oct 09 – G.A. 29 Nov 09)
   - 2 port 10Gb adapter capable of multiple I/O functions
   - Up to 8 virtual Ethernet NICs/adapter, appear as individual adapters to the OS
   - Future upgrade – combination of Ethernet and Fibre Channel virtual ports - benefits of both vNIC and Universal CNA on single adapter.

3. BNT 10-port 10G Ethernet Switch Module – PN 46C7191
   - Switch recognizes the vNIC and assigns bandwidth in increments of 100MB
   - The switch is already CEE/FCoE Ready – ability to support with NIC upgrade
   - Ability to group virtual NIC's helping simplify and reduce management time

Virtual Pipe and Virtual Grouping!

- Flexibility to fine-tune the speed required
  - **Assign the bandwidth required**
  - **Virtual NIC – carve up the 10G pipes**
    - Presented as individual NIC to OS/Hypervisor
    - Minimum of 100Mb
    - Maximum of 10Gb
  - **Ability allocate uplink bandwidth**
    - 100G of uplink bandwidth
    - Assign virtual groups bandwidths in increments of 10G

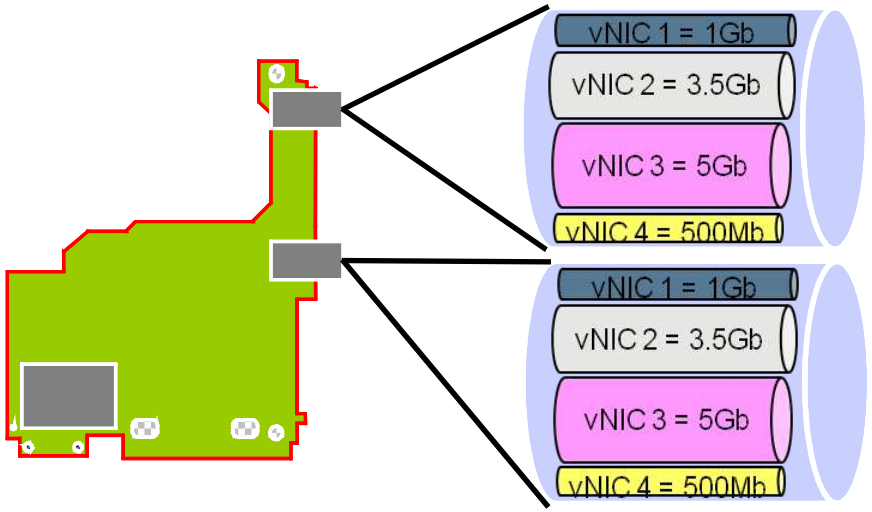

- High Availability and Security
  - **Ability to provide isolation between various virtual adapters**
  - **A failure in one virtual adapter or virtual Group are isolated by others**
  - **VMready can migrate networking characteristics with VM Migration**

- Simpler Management through Grouping
  - **Manage 4 virtual Groups per chassis vs. 112 virtual NICs**

EMULEX

Ann. 20 Oct 09
G.A. 29 Nov 09

IBM BladeCenter Virtual Fabric

Emulex Virtual Fabric Adapter

- Run More VM/CPU
- 20Gb/s Bandwidth
- Provision Performance
- Up to 8 vNICs

Phase 2
Up to 8 Virtual Ports
Ethernet/CNA per adapter
(planned 2H 2010)

Map vNICs to Virtual Port Groups – Corporate LAN, Management, VMotion, VMaware, NMotion

Physical 10GbE Ports

Physical 10GbE Ports

Virtual NICs

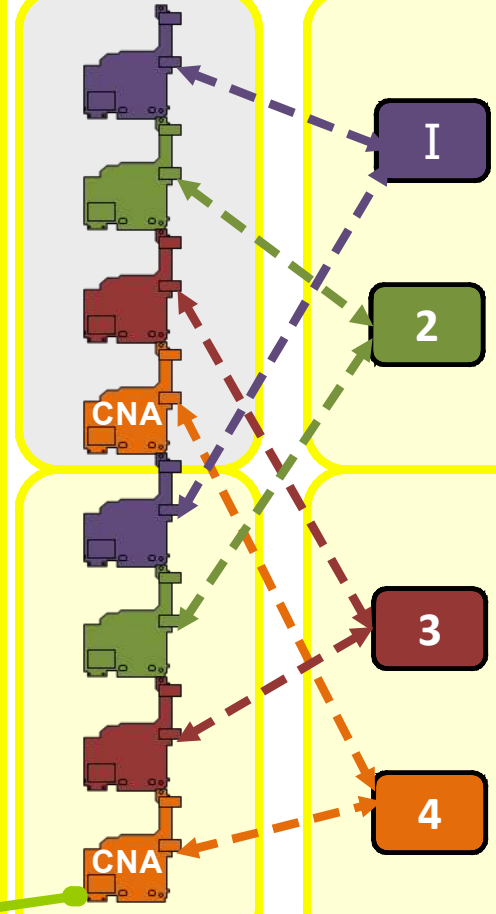

Virtual Groups

Physical BNT Switch

Physical BNT Switch

BladeCenter Open Fabric Manager
CLI, BNT Switch Interface with VMready

BNT 10-port 10Gb Ethernet Switch

- BNT Switch is FCoE Ready
- Simplifies Deployment
- Reduces Management
- Assigns Bandwidth in 100Mb Increments

MSIM - The Evolution of I/O Connectivity (1st)

- Virtualization, multi-core, and clustering are all driving additional port requirements per blade
- BladeCenter H has inherent hardware capability for eight I/O paths per blade
- **Multi-Switch Interconnect Module** provides a means for exploiting existing 8 paths to each blade (HS21, LS21, JS21)
- MSIM fits into the high speed switch slots in BladeCenter H
- Supports standard BladeCenter switch portfolio

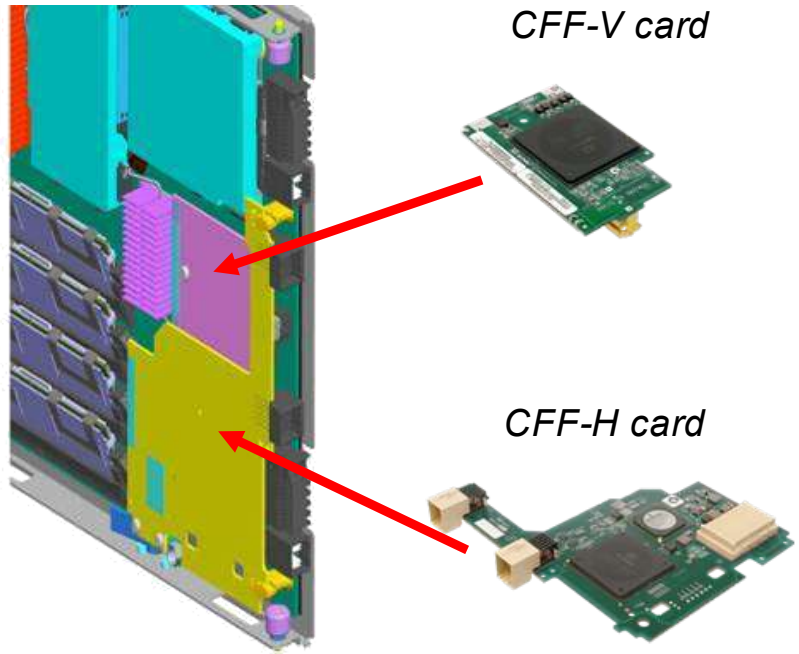

MSIM - The Evolution of I/O Connectivity (2nd)

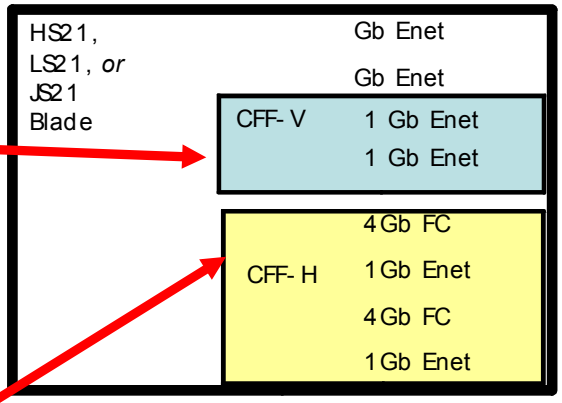

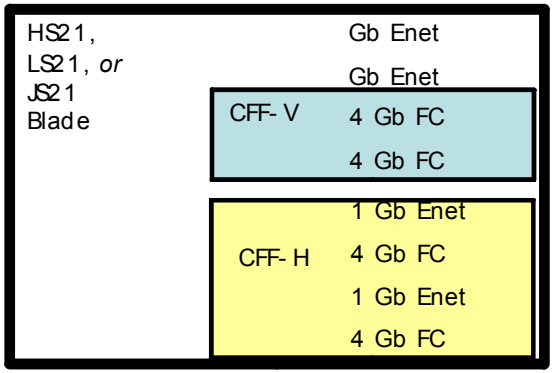

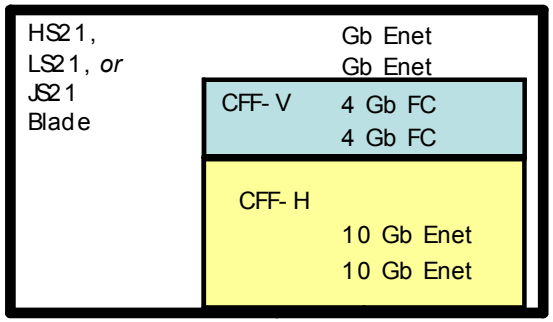

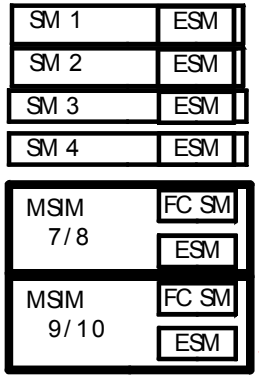

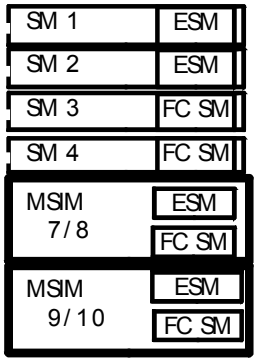

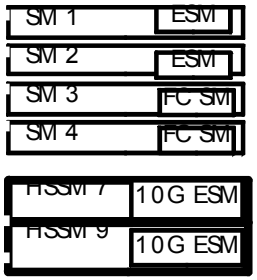

6 1G Ethernet
2 4G Fibre Channel

Base NICs
(1) 39Y9310 Gb Enet CFFV
(1) 39Y9306 Combo CFFH
(1) 39Y9314 MSIM

4 1G Ethernet
4 4G Fibre Channel

Base NICs
(1) 41Y8527 4Gb FC CFFV
(1) 39Y9306 Combo CFFH
(1) 39Y9314 MSIM

2 1G Ethernet
2 4G Fibre Channel
2 10G Ethernet Channel

Base NICs
(1) 41Y8527 4Gb FC CFFV
(1) 39Y9271 10Gb Enet CFFH

2/4 Port Ethernet Expansion Card (CFFh)

Ann:12/08/08
GA: 29/08/08

p/n: 44W4479

Meets customer needs of additional IO Bays on BC-H

Enables IO Bays 7,8,9&10 in BladeCenter H/HT for:

- Additional switch module slots available.
- Increased Bandwidth for Virtualization
- Up to 8 Ethernet port combinations possible

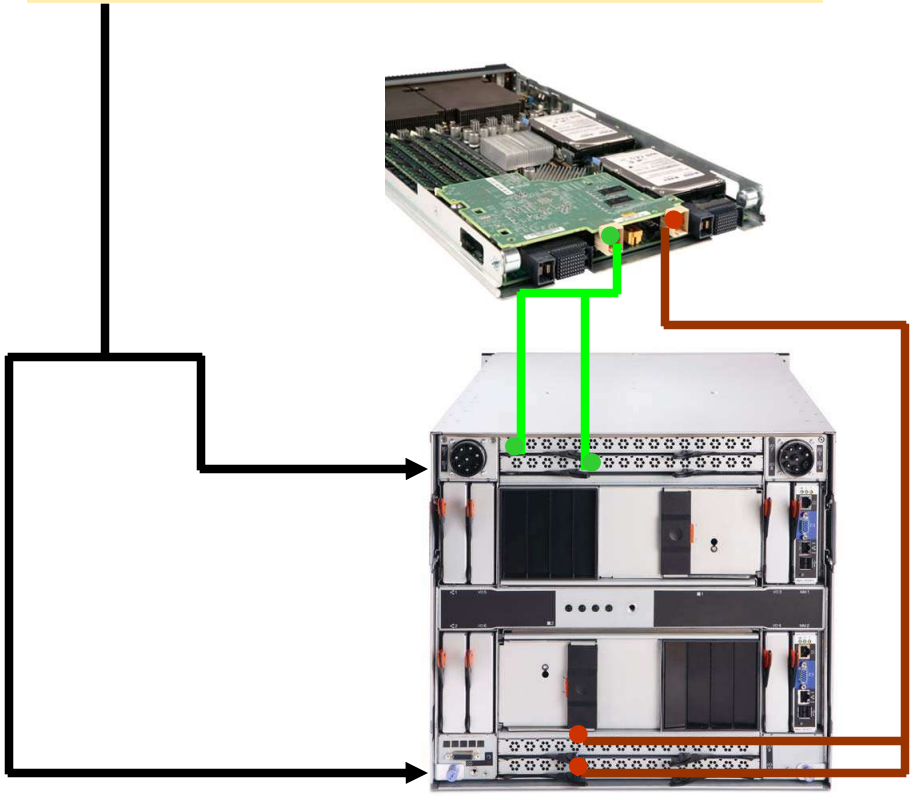

Dual Port 8Gb Fibre Channel & Dual Port 1Gb Ethernet PCIe Expansion Card

Enable 8Gb FC & extra Eth on BC-H + MSIM

- Features a highly desirable combination adapter that enables two 8Gb Ethernet ports, along with two 1Gb Ethernet ports
- Supported blade platforms include HS21, HS21xm, LS21, LS41, LS42, HS12
- Uses QLogic 2532 ASIC and Broadcom 5709S ASIC

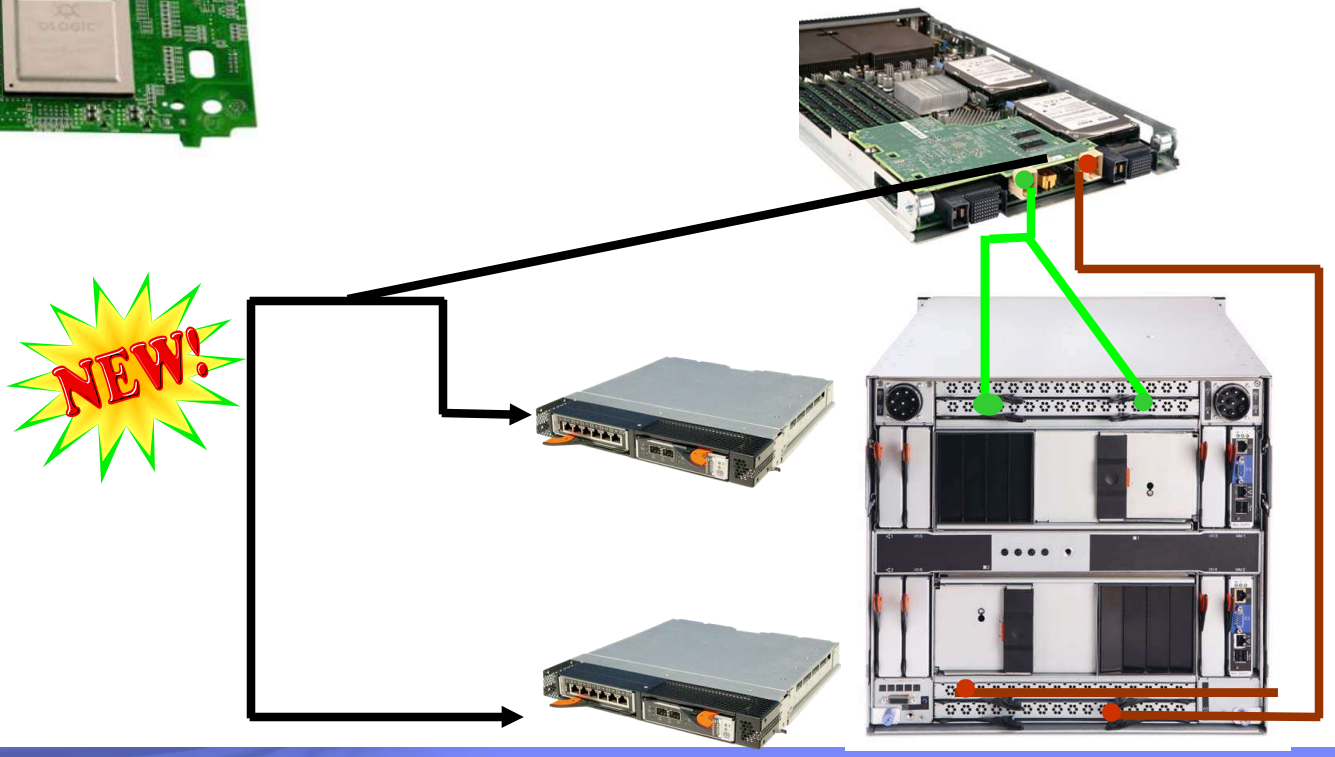

Ann:6/1/2009

GA: 30/1/2009

p/n: 44X1940

Introducing BladeCenter S

1. First BladeCenter chassis focused on distributed LE & SME markets
2. First BladeCenter chassis to incorporate integrated shared storage
3. Significant focus on improving usability and customer experience
4. Attractive alternative to racks and high availability towers
5. Greatly improved BladeCenter platform for small office environment

By tailoring BladeCenter S for distributed LE & SME, IBM is calling the same highly successful play that stormed BladeCenter into dominant market share in the Data Center

BladeCenter S Product Summary

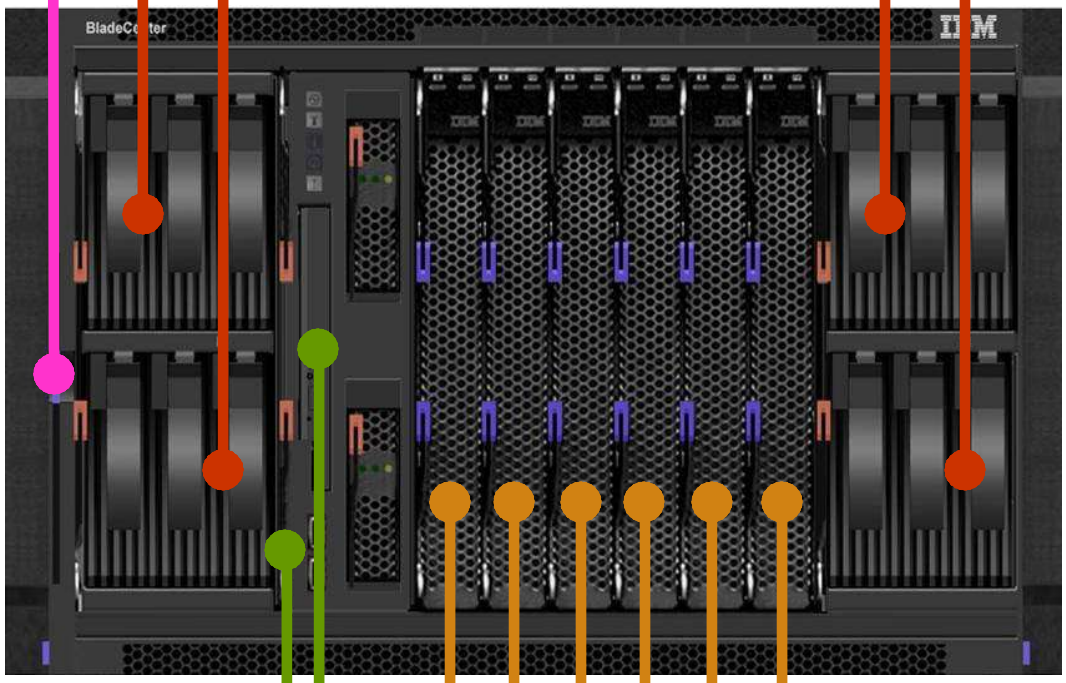

BladeCenter S Product Summary

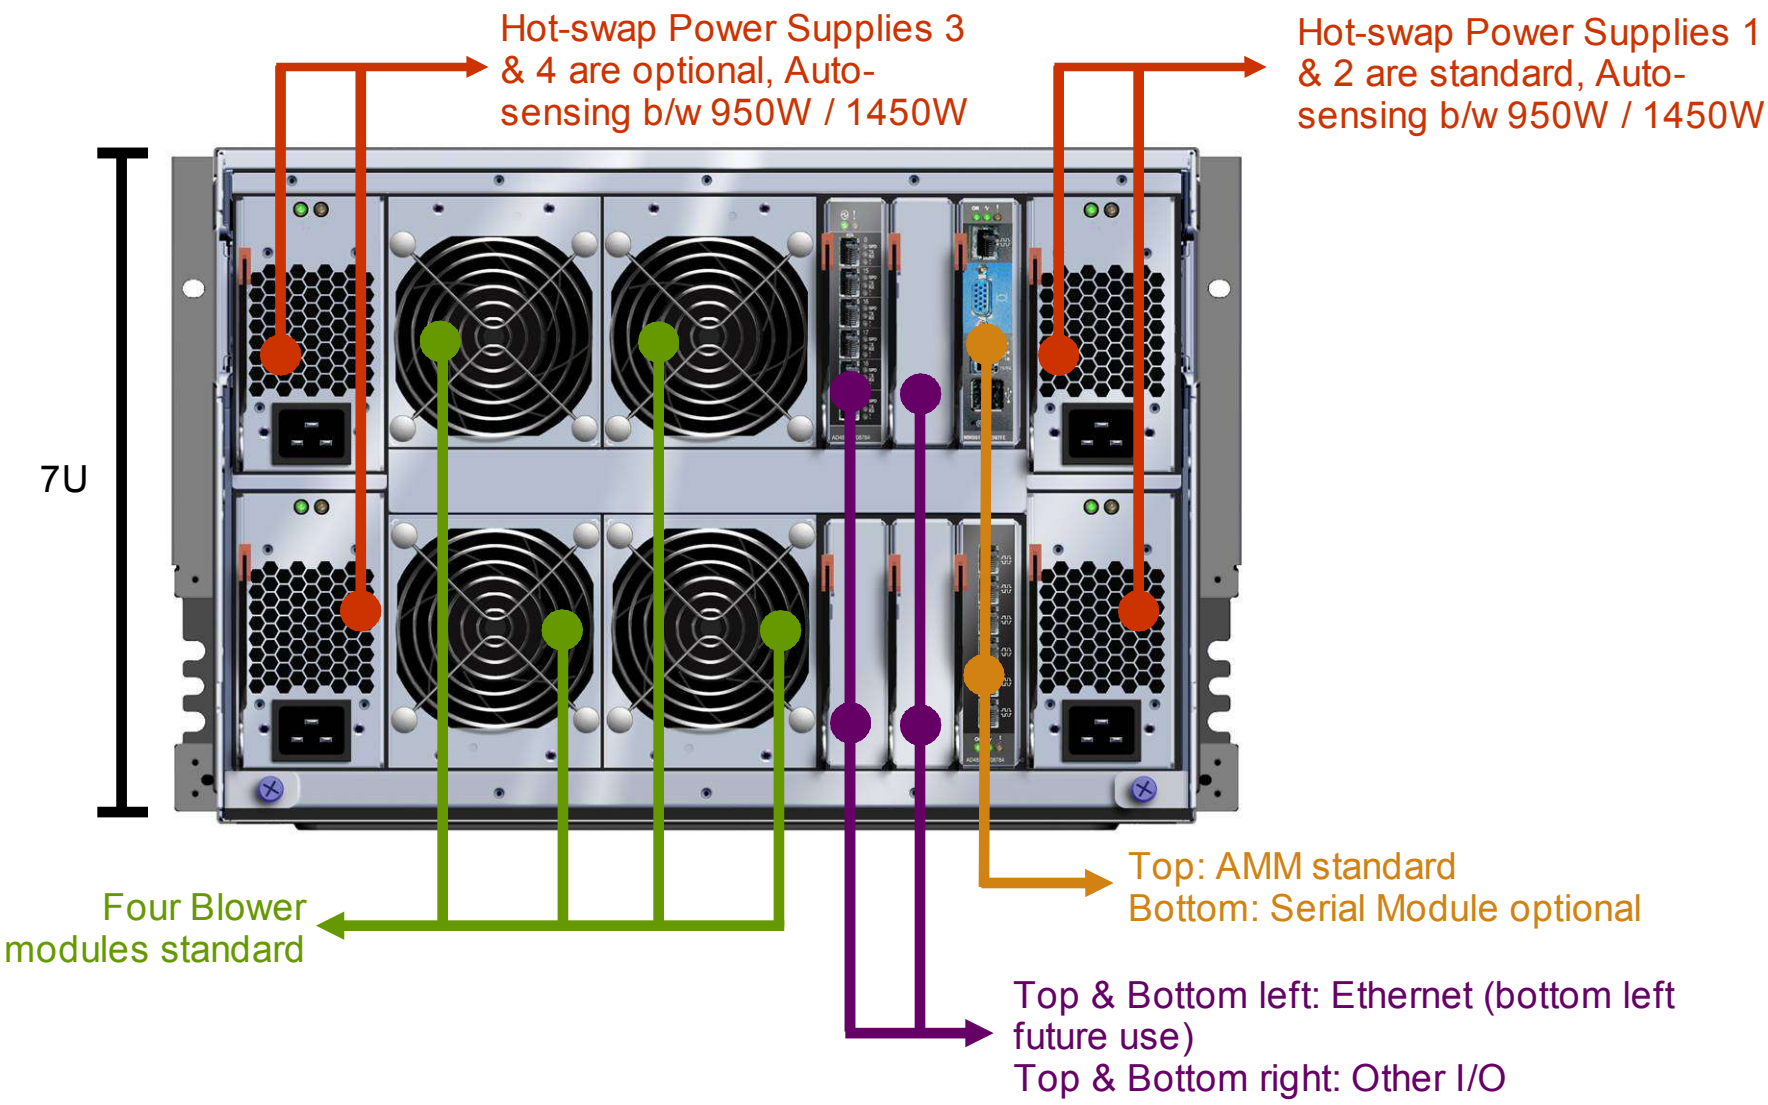

Interoperability

I/O Bays 1 and 2

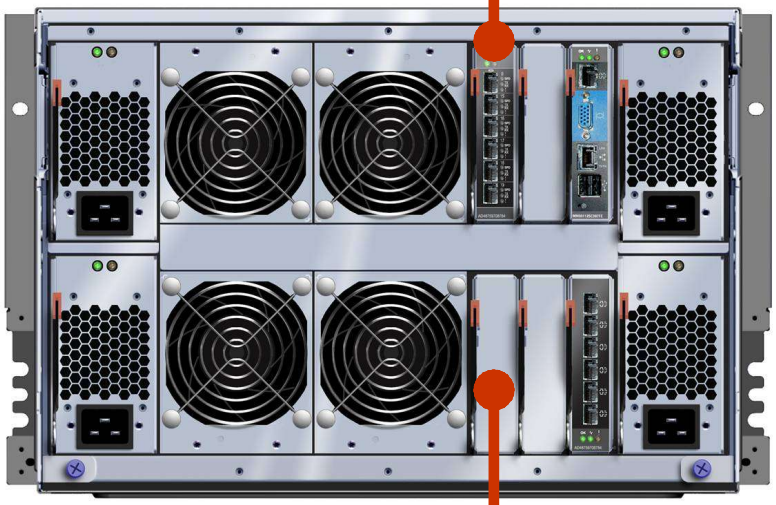

I/O Bay 1: For Ethernet switch

This switch bay is wired to both Ethernet ports of the 6 blade slots

This bay accepts the following switches:

Part #	Description
32R1783	BNT 10Gb Uplink Ethernet Switch Module
32R1860	BNT Layer 2/3 Copper GbE Switch Module
32R1861	BNT Layer 2/3 Fibre GbE Switch Module
39Y9324	Server Connectivity Module
39Y9320	IBM BladeCenter Copper Pass-thru Module
39Y9316	IBM BladeCenter Optical Pass-thru Module (Ethernet only)
32R1859	BNT Networks Layer 2-7 Gigabit Ethernet Switch Module

NEW!

I/O Bay 2:

Redundant Ethernet switch

with 4 port 1Gb CFFh exp card (44W4479)

2/4 Port Ethernet Expansion Card (CFFh)

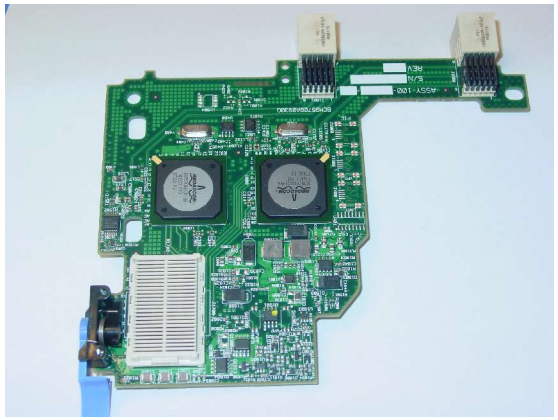

Ann:12/08/08
GA: 29/08/08

p/n:44W4479

Meets Customer needs for IO redundancy on BC-S

Enables 2nd Ethernet switch in BladeCenter S for:

- Switch Module level failure tolerance and redundancy
- Increased Ethernet bandwidth for Virtualization workloads
- Up to 6 Ethernet port combinations possible

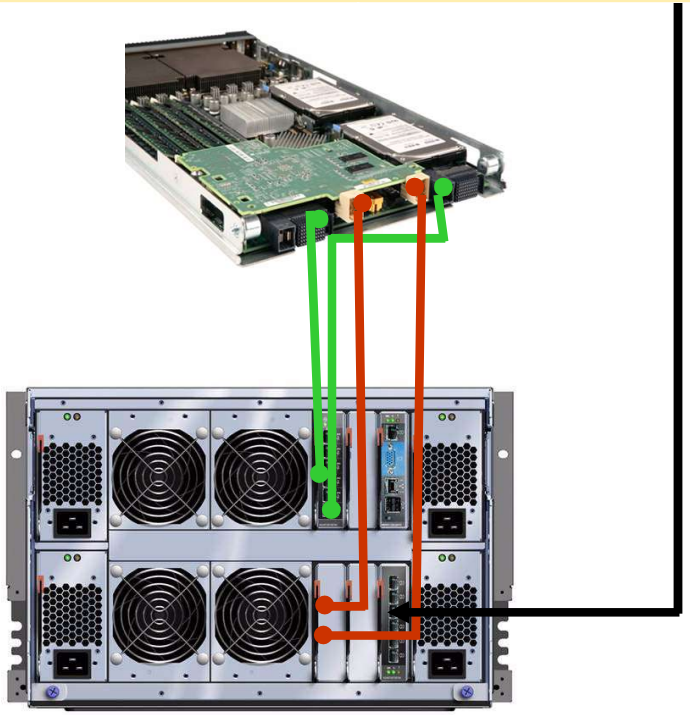

Interoperability

I/O Bays 3 and 4

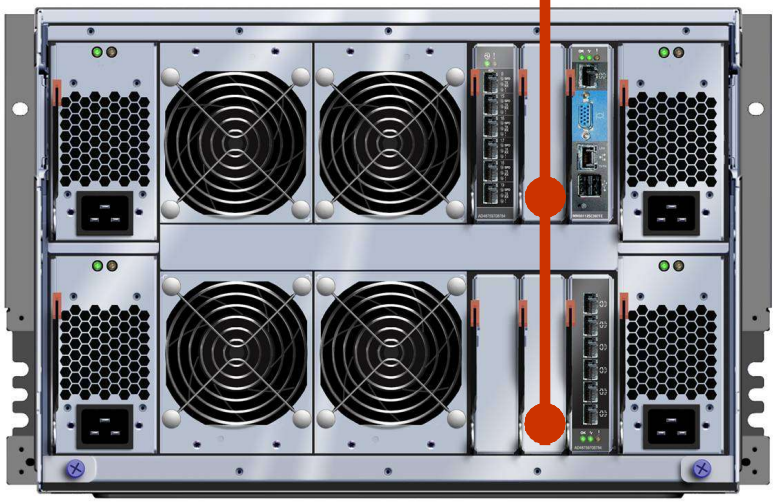

SAS RAID Controller Module

Ann: 09/09/08

GA: 24/10/08

NEW!

- Redundant configuration (dual RAID Controller, dual-switch and dual-BBU)
- RAID 0, 1, 0+1 & 5
- Supports two disk storage modules (DSM) with up to 12 x 3.5in SAS drives
- Need of Eth switch for internal communication

I/O Bay 3 & I/O Bay 4:
For Ethernet, Fibre Channel, or SAS
Must be of same type, mixing is not supported

If DSS is installed then SAS Connectivity Module or SAS RAID Controller is/are required.

The following are accepted:

Part #	SAS RAID controller
43W3584	IBM BladeCenter S SAS RAID Controller Module
Part #	SAS switch
39Y9195	IBM BladeCenter SAS Connectivity Module
Part #	Ethernet switches
32R1783	BNT 10Gb Uplink Ethernet Switch Module
32R1860	BNT Layer 2/3 Copper GbE Switch Module
32R1861	BNT Layer 2/3 Fibre GbE Switch Module
39Y9324	Server Connectivity Module
39Y9320	IBM BladeCenter Copper Pass-thru Module
39Y9316	IBM BladeCenter Optical Pass-thru Module (Ethernet only)
32R1859	BNT Networks Layer 2-7 Gigabit Ethernet Switch Module
Part #	Fibre Channel switches
32R1813	Brocade 10-port 4Gb SAN Switch Module
39Y9284	Cisco Systems 4Gb 10-port Fibre Channel Switch Module
43W6724	QLogic 10-port 4Gb SAN Switch Module
43W6723	QLogic Intelligent Pass-Thru Module
39Y9316	IBM BladeCenter Optical Pass-thru Module

Interoperability

DSM

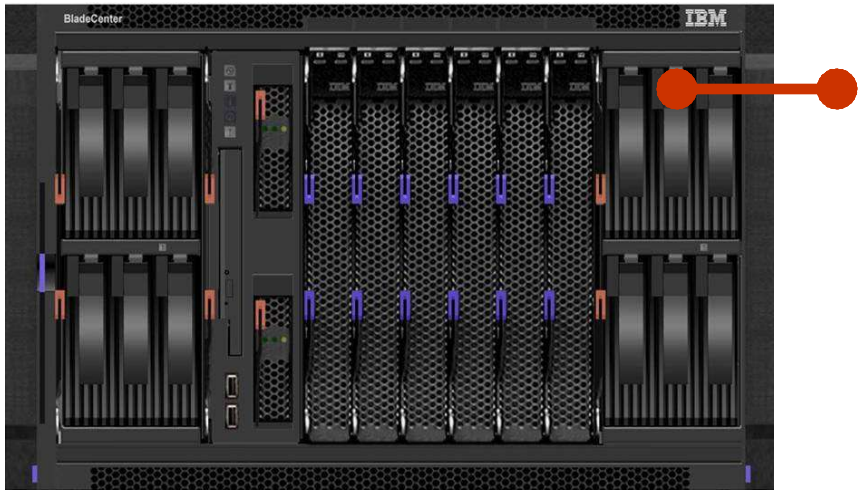

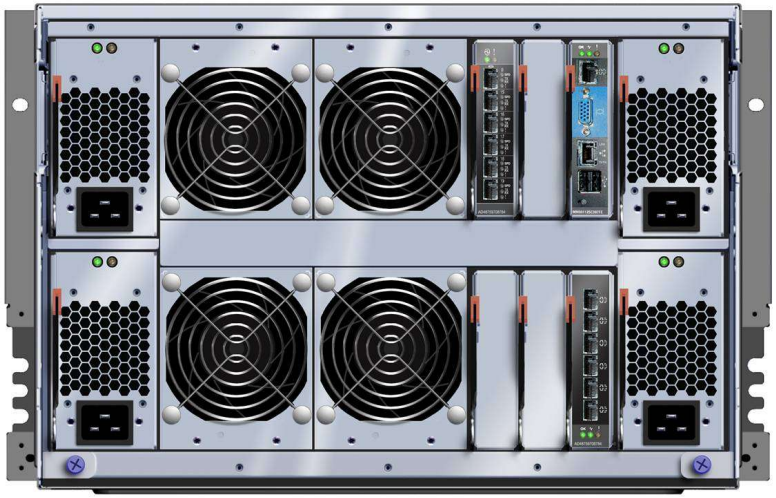

NEW!

DSM 1 & DSM 2:
Any SAS or SATA disk or combination is valid (note: no intermix with RAIDed switch, only SAS disks allowed).

SAS Switch is required if DSM are populated with disks.

Power Supply 3 & 4 are required if DSM 2 is installed.

The following disks are supported:

Part #	HDDs
40K1043	73GB 15K SAS
40K1044	146GB 15K SAS
43X0802	300GB 15K SAS
42D0519	450GB 15K SAS
42D0546	750GB 7.2K SAS
42D0547	1 TB 7.2K NL SAS
39M4530	500GB SATA
43W7576	750GB SATA

Ann: 2 dic 2008
Avail: 27 feb 2009

When to use the Optional Power Supplies 3 & 4

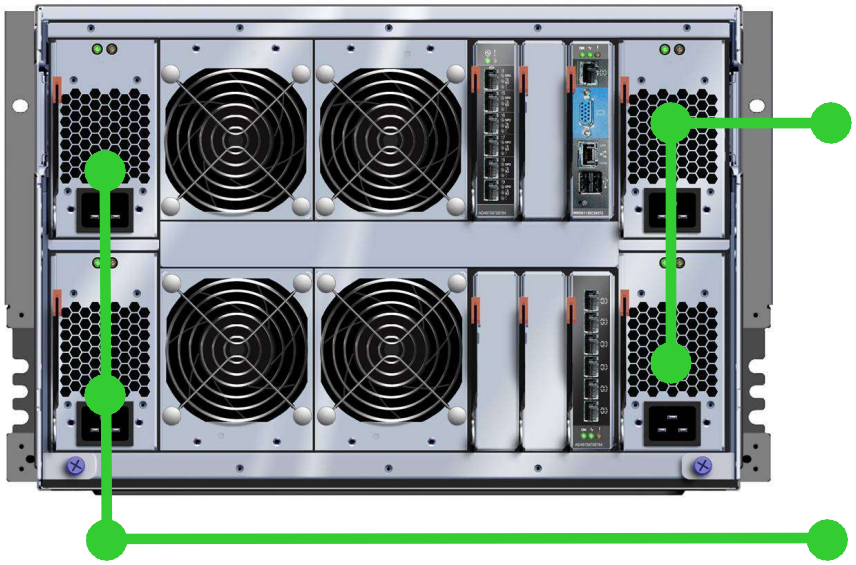

Standard Power Supplies supply provide power for:

- All blade servers,depending from electric load
- nr.1 Disk Storage Module with disks

Optional Power Supplies 3 & 4 needed to power all slots:

- When blades are power demanding
- nr.2 Disk Storage Modules installed (with disks)

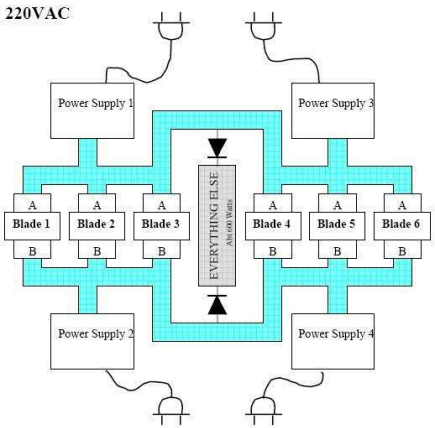

- Need 4 separate AC circuits, wired to 4 separate "breakers"
- Each circuit receptacle must be within 6 ft of the chassis
- Each power cord is a standard 6 ft length
- each supply 1450W capable
- total pwr limited to 2900W per Power Domain

BladeCenter S and the NEW O.E.K.

- Ideal way to deploy BladeCenter S into office environments
- Includes Acoustical Module built into the back
- Optional Air Filter on the front
- Locking door for security
- Mobile with rollers
- 33% (4U) extra room to grow

Incredibly Quiet
Dust Filter
Up to 9TB
Shippable WW

Perfect workgroup solution foundation

BladeCenter Server Portfolio

HS12
New Function, new price point

HS22
"No Compromise"
Enterprise Blade

LS22
High-performance

LS42
Scalable, enterprise performance

JS12 / 22
Power6

JS23 / 43
Power 6+

QS22
High-performance

T2BC
Themis
UltraSPARC

Announcing… HS22 !

Versatile, easy to use blade optimized for performance, power and cooling

What is announcing…

- New 2-socket, 30mm blade based on Intel's Xeon 5500 processor
- Next workhorse server for IBM BladeCenter
- Follow-on to both HS21 and HS21 XM

Announcement importance…

- Delivers a robust set of features that does many things exceptionally well:
  - 2 processor sockets supporting all of Intel's latest server processor SKUs
  - 2 hot-swap internal storage bays supporting both SAS and solid state
  - 12 memory DIMMs
  - PLUS MORE… (embedded hypervisor, battery-backed cache, etc.)
- Time-to-market with Intel

NEW!

HS22

Announce & Availability:

31 March 2009

Intel "Tick-Tock" Development Model

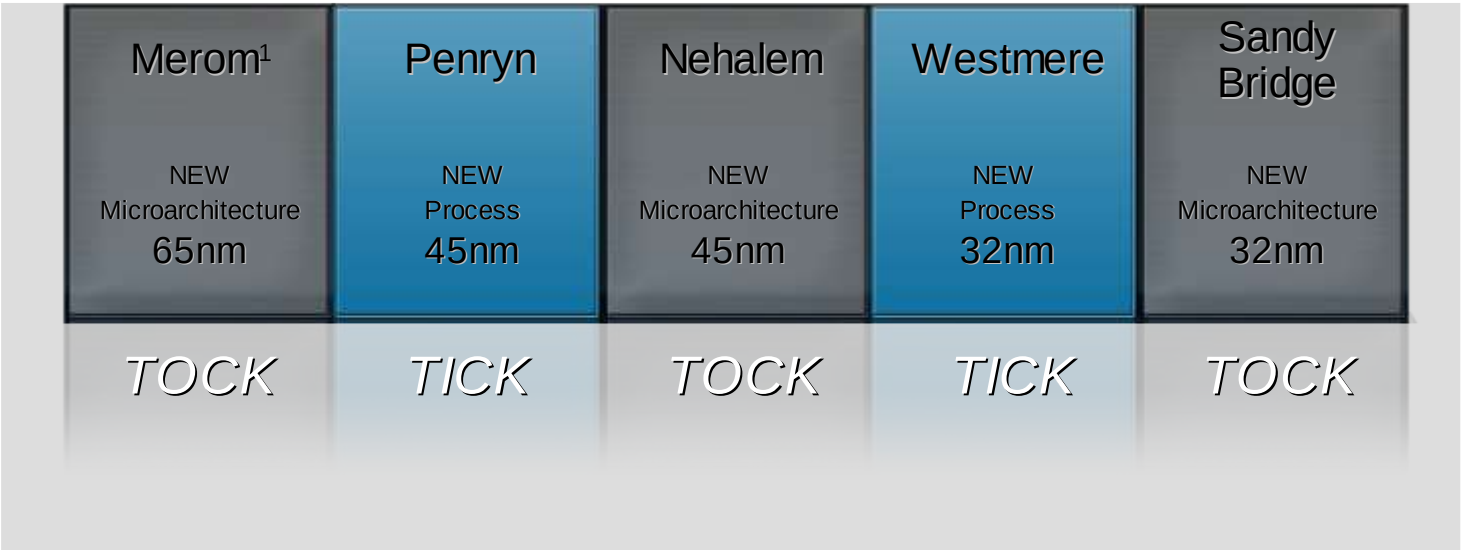

[1]*Intel® Core™ microarchitecture (formerly Merom)*
45nm next generation Intel® Core™ microarchitecture (Penryn)
Intel® Core™ Microarchitecture (Nehalem)
Intel® Microarchitecture (Westmere)
Intel® Microarchitecture (Sandy Bridge)

56

Nehalem EP Overview

- Micro-architecture enhancements (Core i7)
- 45nm Manufacturing Process
- Integrated three channel DDR3 memory controller
- 2 Quick path interconnect links
- Single die quad core
- Shared 8MB L3 cache
- Return of SMT (hyperthreading)
- High core frequencies (up to 2.93GHz for server)
- Significant (>2x) improvement in SpecFP and STREAM (>3x) benchmarks

- **Announcing March 30th**

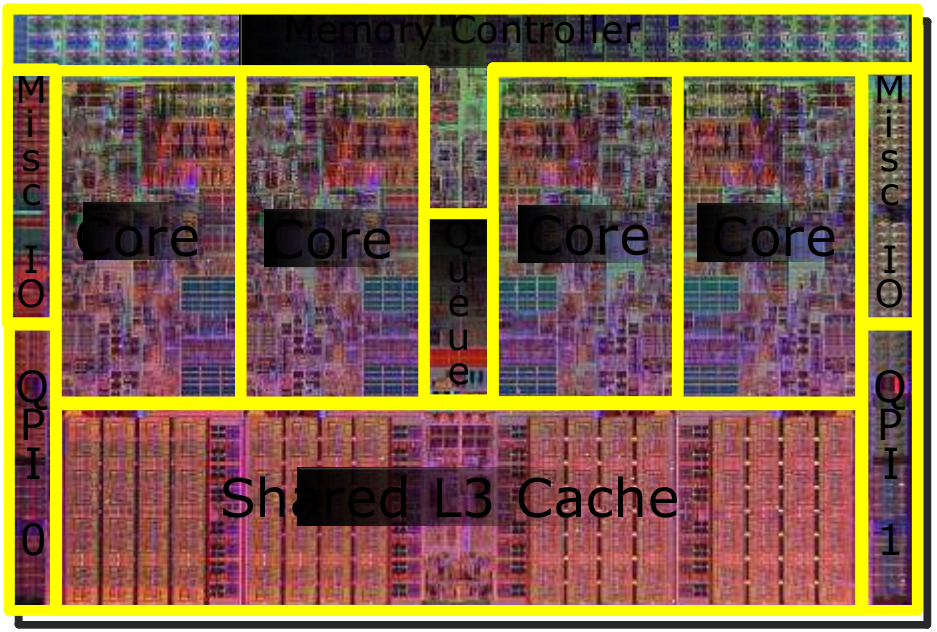

Intel® Xeon™ 5500 Platform Architecture

- **Integrated Memory Controller**
  - **3 DDR3 channels per socket**
  - **Massive memory** *bandwidth*
  - **Memory Bandwidth scales with # of processors**
  - **Very** *low memory latency*
- **Intel® QuickPath Interconnect (Intel® QPI)**
  - **New point-to-point interconnect**
  - **Socket to socket connections**
  - **Socket to chipset connections**
  - **Build** *scalable* **solutions**

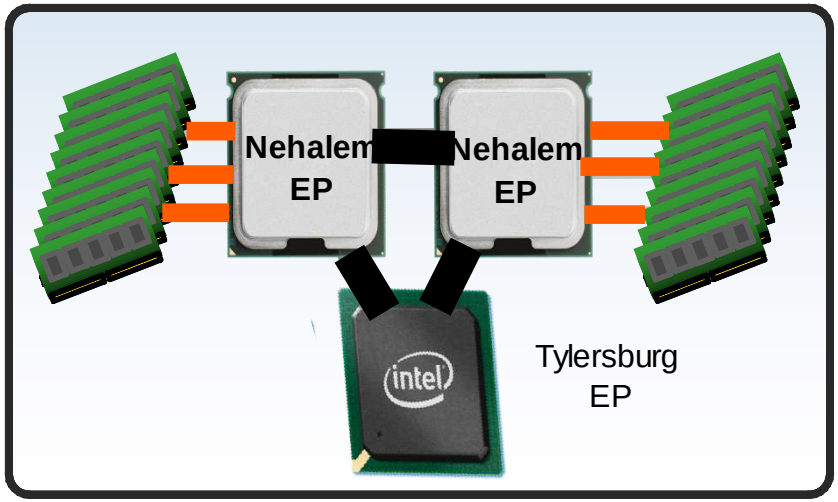

Significant performance leap from new platform

Intel® Core™ microarchitecture (Nehalem-EP)
Intel® Next Generation Server Processor Technology (Tylersburg-EP)

Intel® Xeon™ 5500: Integrated Memory Controller (IMC)

- **Memory controller optimized per market segment**
- **Initial Intel® Xeon™ 5500 Microarchitecture products**
  - **Native DDR3 IMC**
  - **Up to 3 channels per socket**
  - **Massive *memory bandwidth***
  - **Designed for *low latency***
  - **Support RDIMM and UDIMM**
  - **RAS Features**
- **Future products**
  - ***Scalability***
    - **Vary # of memory channels**
    - **Increase memory speeds**
    - **Buffered and Non-Buffered solutions**
  - **Market specific needs**
    - **Higher memory capacity**
    - **Integrated graphics**

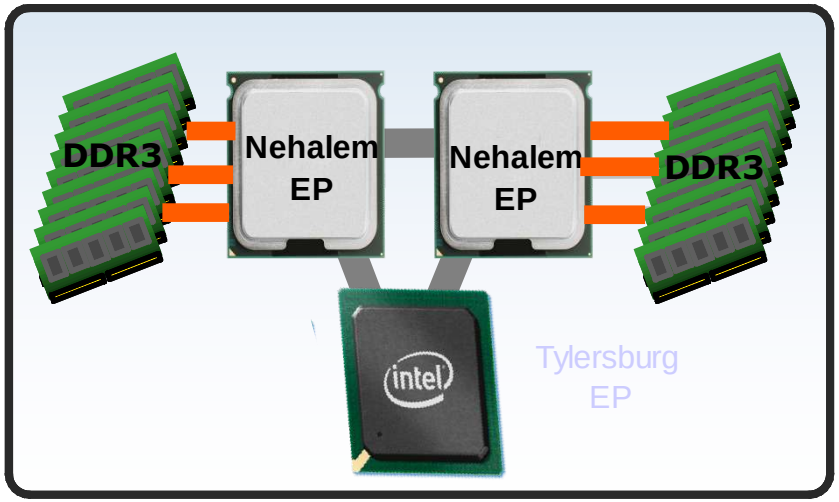

Significant performance through new IMC

roarchitecture (Nehalem-EP)

Intel® Xeon™ 5500: Intel® QuickPath Interconnect

- **Intel® Xeon™ 5500 Microarchitecture introduces new Intel® QuickPath Interconnect (Intel® QPI)**
- ***High bandwidth*, *low latency* point to point interconnect**
- **Up to 6.4 GT/sec initially**
  - **6.4 GT/sec -> 12.8 GB/sec**
  - **Bi-directional link -> 25.6 GB/sec per link**
  - **Future implementations at even higher speeds**
- **Highly *scalable* for systems with varying # of sockets**

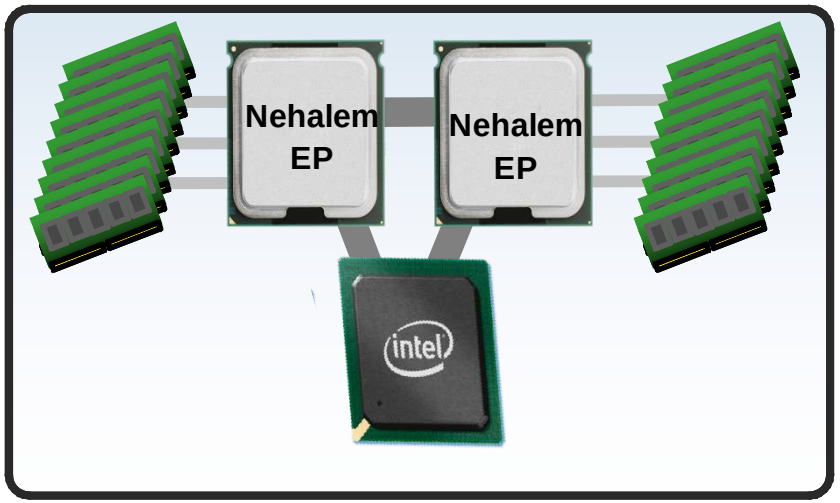

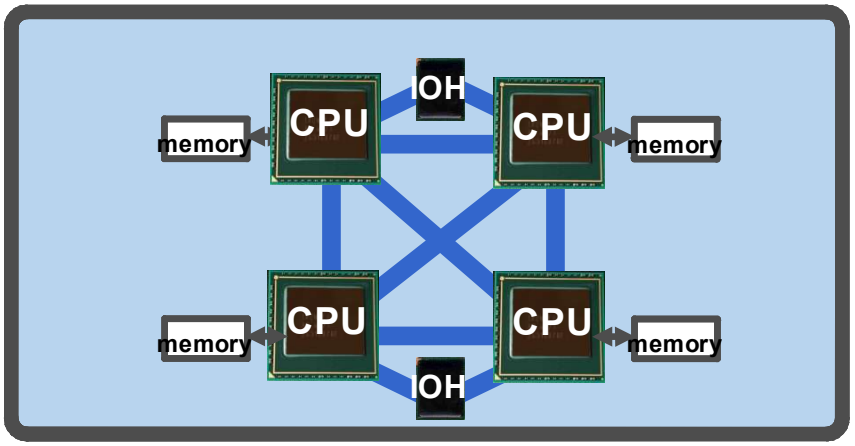

Intel® Core™ microarchitecture (Nehalem-EP)

Nehalem EP Features

- Turbo Boost
  - Allows you to turn off cores inside the processor and then increase the clock speed on the remaining cores
  - Also allows for short term increases in overall clock speed providing power and cooling is available (without turning off cores)

Intel Nehalem EP Turbo Mode Overview

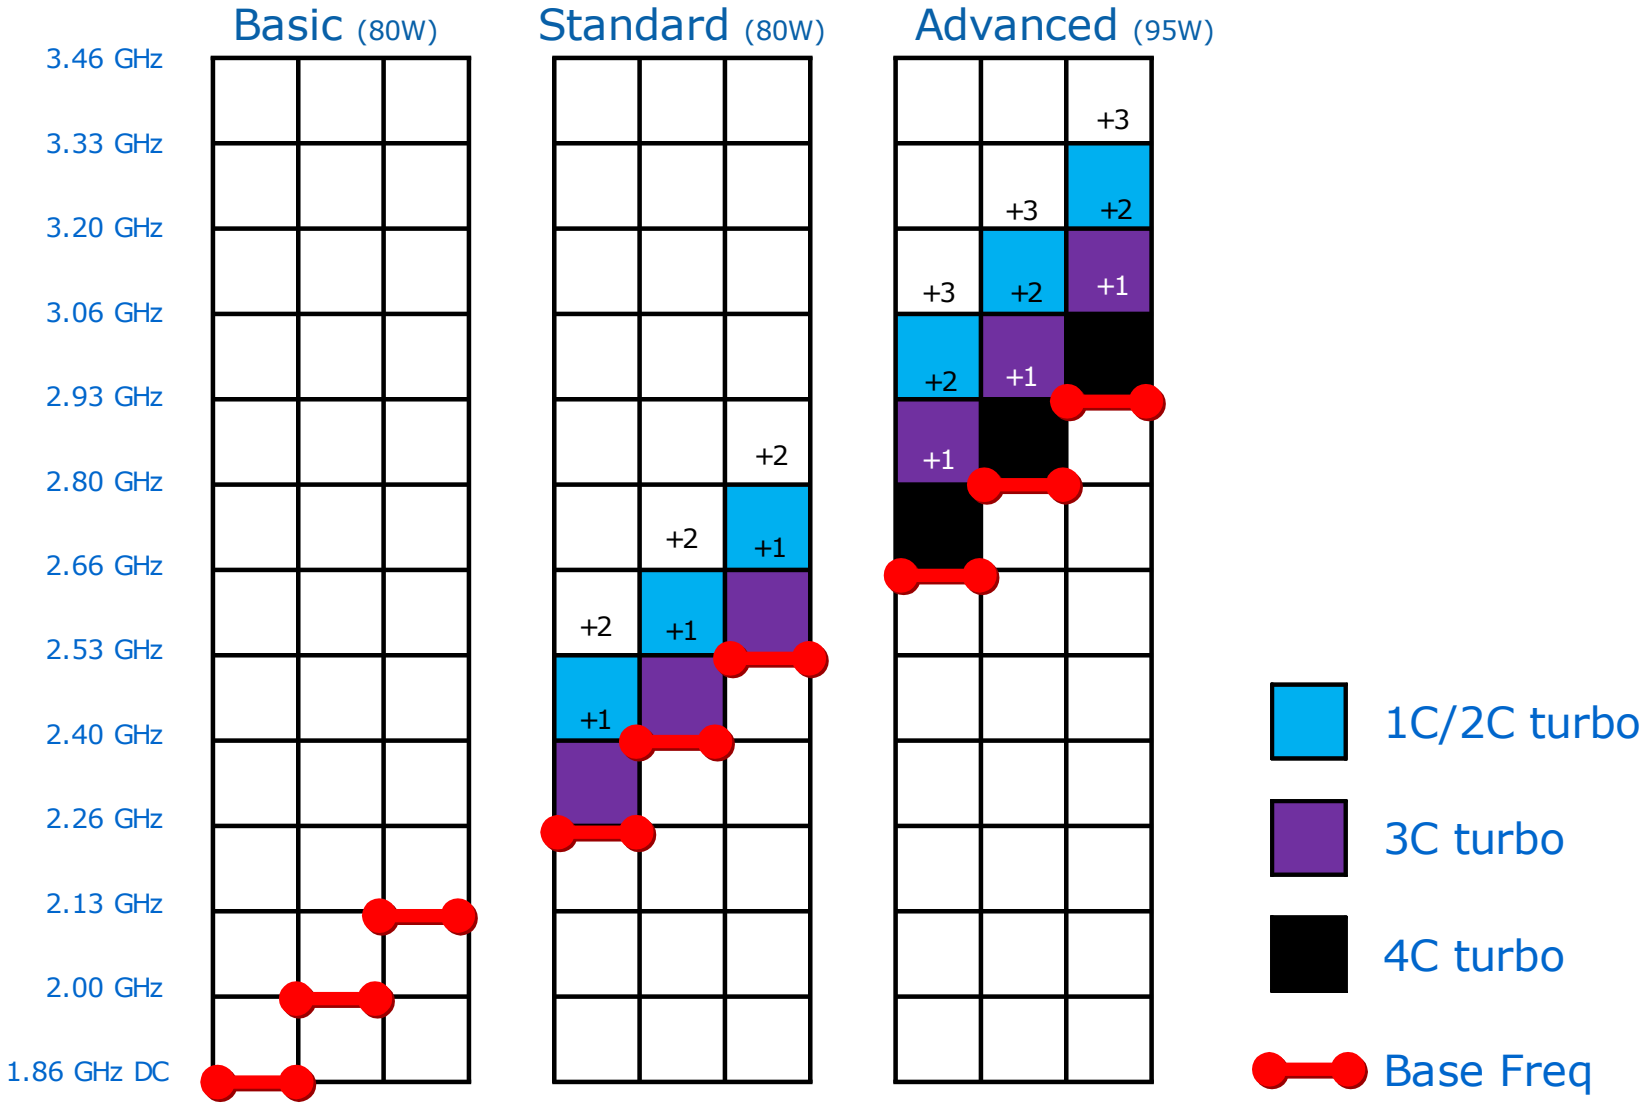

Intel® Turbo Boost Technology

OR

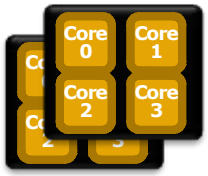

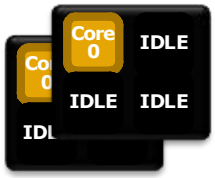

Cores / Threads	TURBO Freq BASE Freq	Benefit
16 (2 socket/ HT on)	3.20 GHz → 2.93 GHz	up to 6%† for 16 concurrent software threads
2 (2 socket /HT off)	3.33 GHz → 2.93 GHz	up to 10%†† for 2 software threads

Improves application responsiveness
Delivers higher processor frequency on demand

† Source: Intel Internal measurements Nov 2008., based on SPECjvm*2005. See backup for additional details
†† Source: Intel Internal measurements Nov 2008., based on SPECjbb*2005. See backup for additional details.

HS22 Hardware Overview

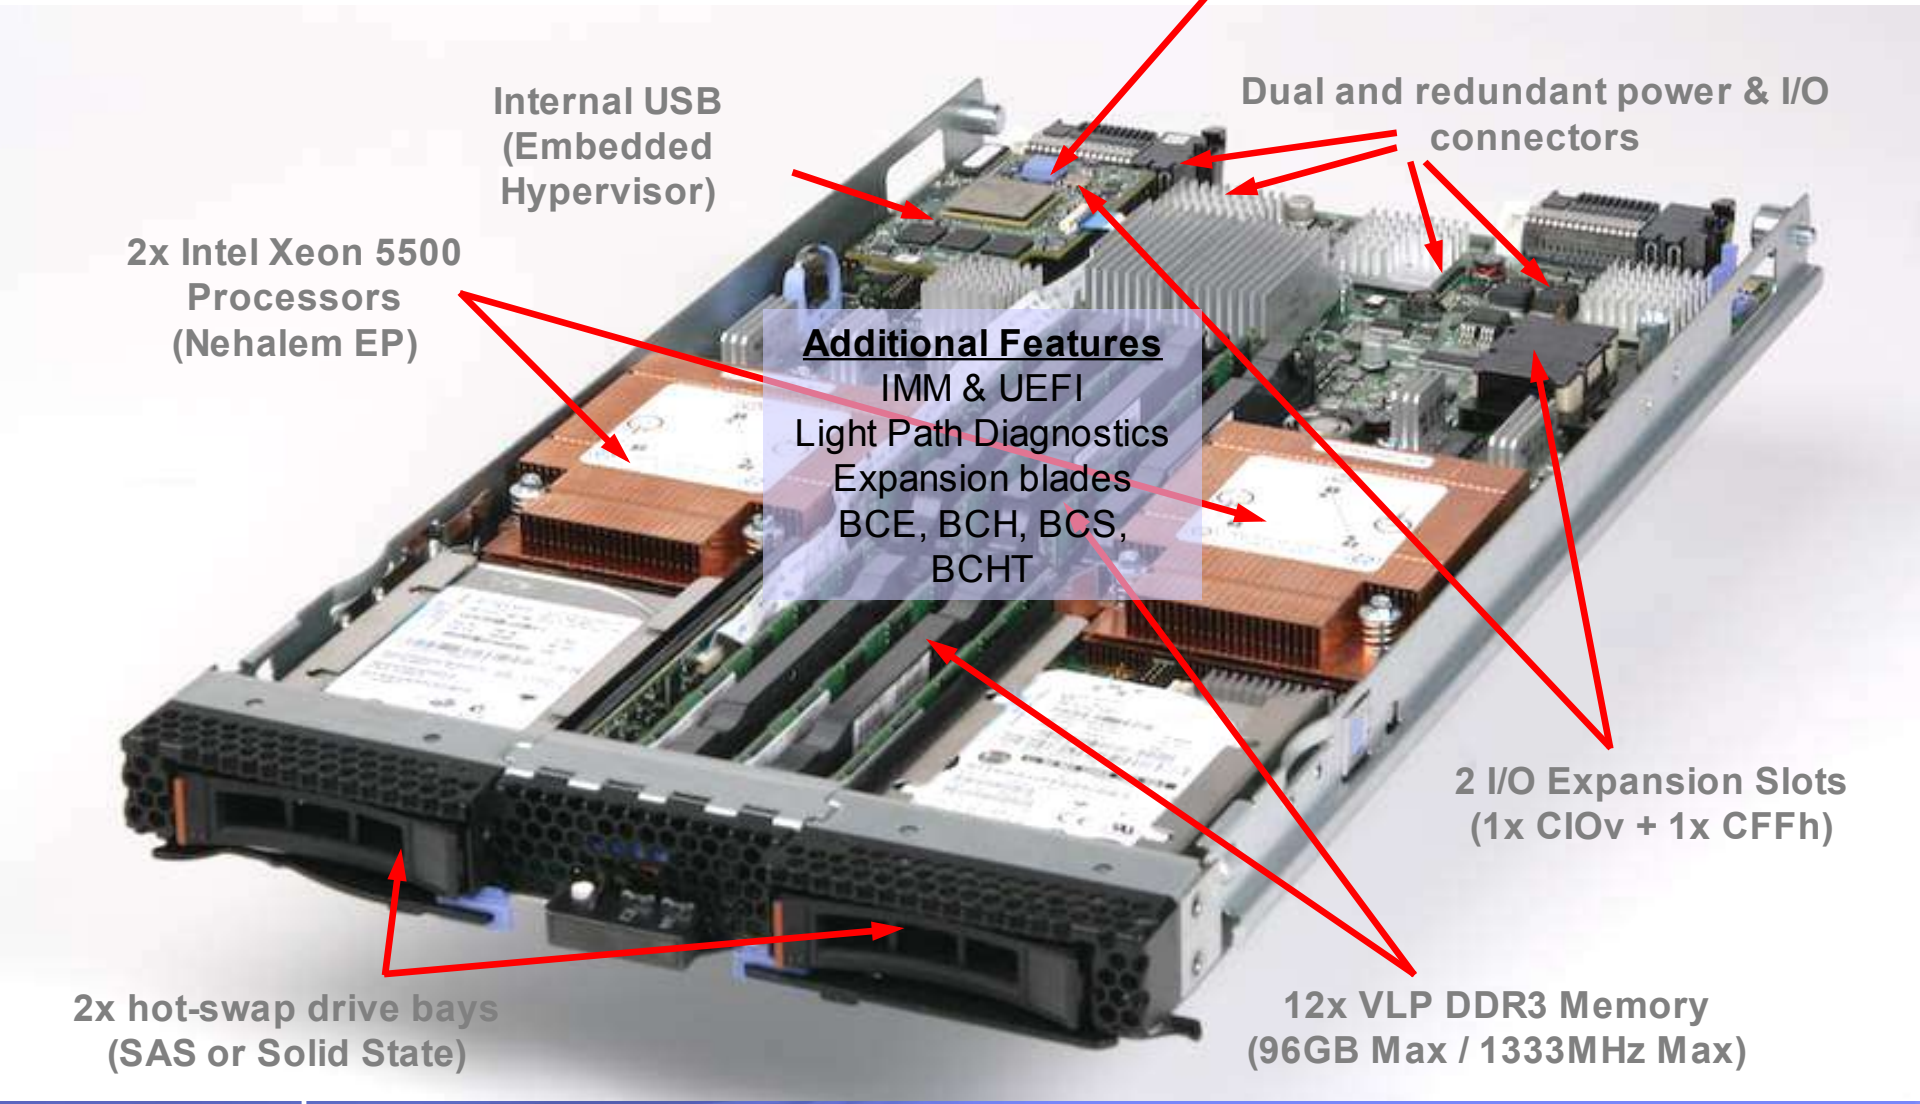

Unified Extensible Firmware Interface (UEFI)

More functionality, better user interface, easier management for users

- More functionality
  - Adapter vendors can add more features in their options (e.g., IPv6)
  - Modular designs allows faster updates as new features are introduced
  - More adaptors can be installed and used simultaneously
  - Fully backwards compatible with legacy BIOS
- Better user interface
  - Replaces ctrl key sequences with a more intuitive human interface
  - Moves adaptor and iSCSI configuration into F1 setup
  - Creates event logs that are more easily decipherable
- Easier management
  - Eliminates “beep” codes; all errors can now be covered by Lightpath
  - Reduces the number of error messages and eliminates out-dated errors
  - UEFI settings can be managed both in-band and out of band

UEFI OS Support	
Windows 2008	Today
SLES 11	2Q09
RHEL 6	1Q10

Prior operating system versions will operate in legacy BIOS mode

Common across all new IBM x86 servers!

HS22 I/O Expansion Daughter Cards

Low Speed Daughter Cards

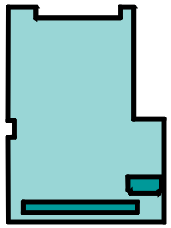

Std DC

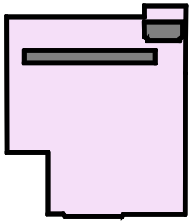

SFF DC

CFFv (PCIx)

CIOv (PCIe)

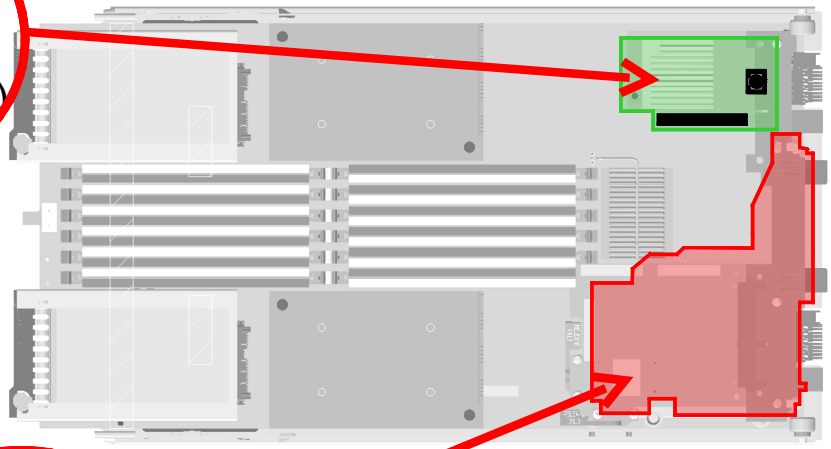

High Speed Daughter Cards

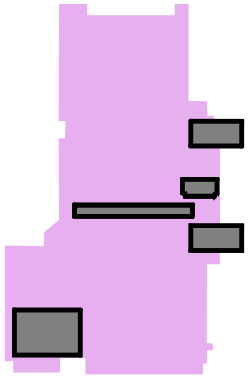

LFF HSDC

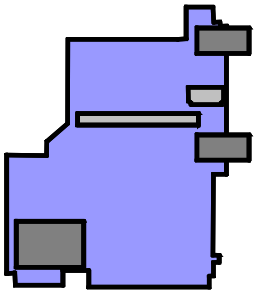

SFF HSDC

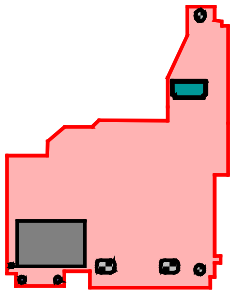

SvFF

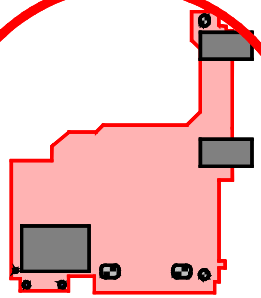

CFFh

HS22 I/O Expansion Daughter Cards

- Two Daughter card formats
  - CFFh
  - CIOv

- CFFh
  - All current cards supported
  - PCIExpansion connector enhanced to allow x16 cards
  - PCIExpress x1 added to connector for future cards

- CIOv
  - Dual Gigabit Ethernet Expansion Card – (44W4475)
  - 4Gb Fibre Channel– (46M6065)
  - SAS Connectivity Card – (43W4068)

Note: new form factor replaces CFFv

ServeRAID-MR10ie Controller

High-Performance BladeCenter RAID adapter with BBC for HS22

Target Market

- SMB market
- Citrix workloads using local disk
- Digital Video Surveillance solutions

NEW!

Now compatible w. HS22

Announce: 2 June 2009

Availability: 19 June 2009

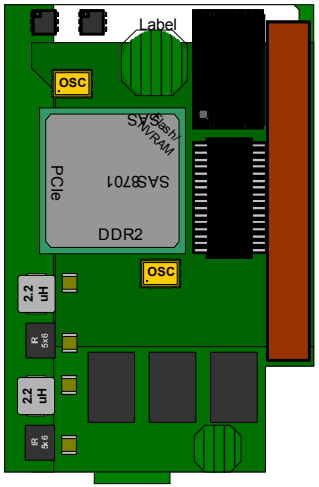

Battery Backed Cache features

- HS22 “CIOv” daughter card option (46C7167)
- One LSI 1078 RAID card w/256MB Cache
- 24 hour back-up of cache
- On-board charging unit
- Battery Unit installs in standard DIMM slot
- Supports hot swap SAS drives (no SATA or SSD)
- Comparable to our rack based solutions

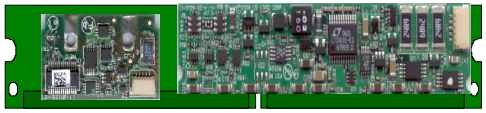

ServeRAID-MR10ie applications

- Solutions Using On-Blade HDDs [Citrix]

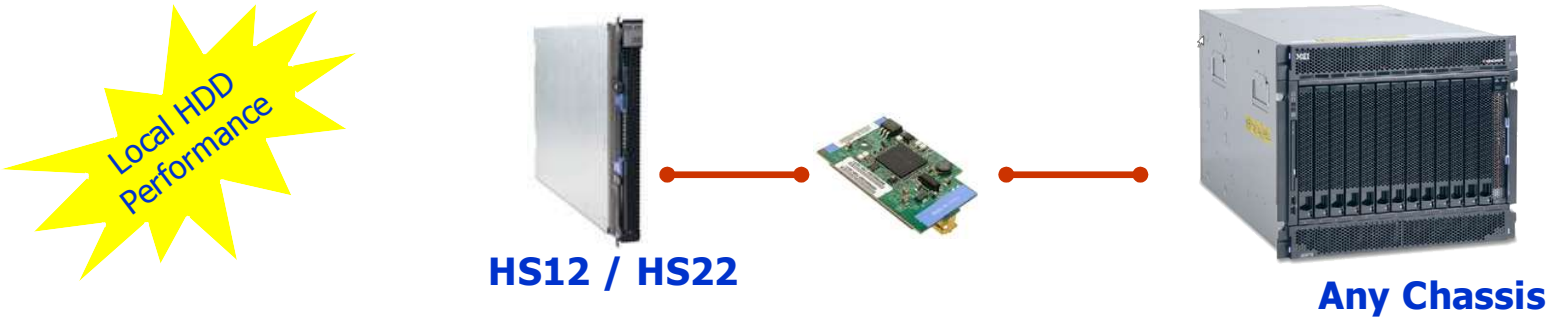

- Solutions Using Off-Blade HDDs [DVS]

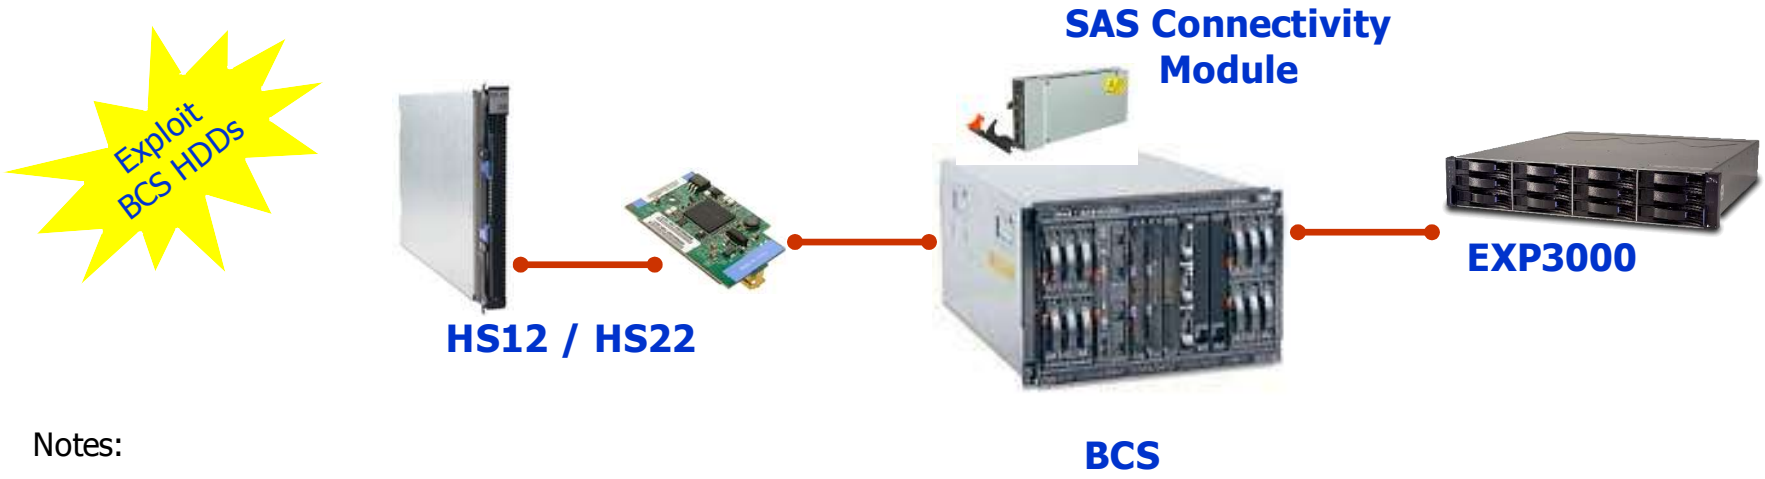

Notes:

- Maximum 26 HDD attachment per ServeRAID MR10ie Controller
- Maximum one EXP3000 attachment per ServeRAID MR10ie Controller

ServeRAID-MR10ie Controller on HS22 & battery

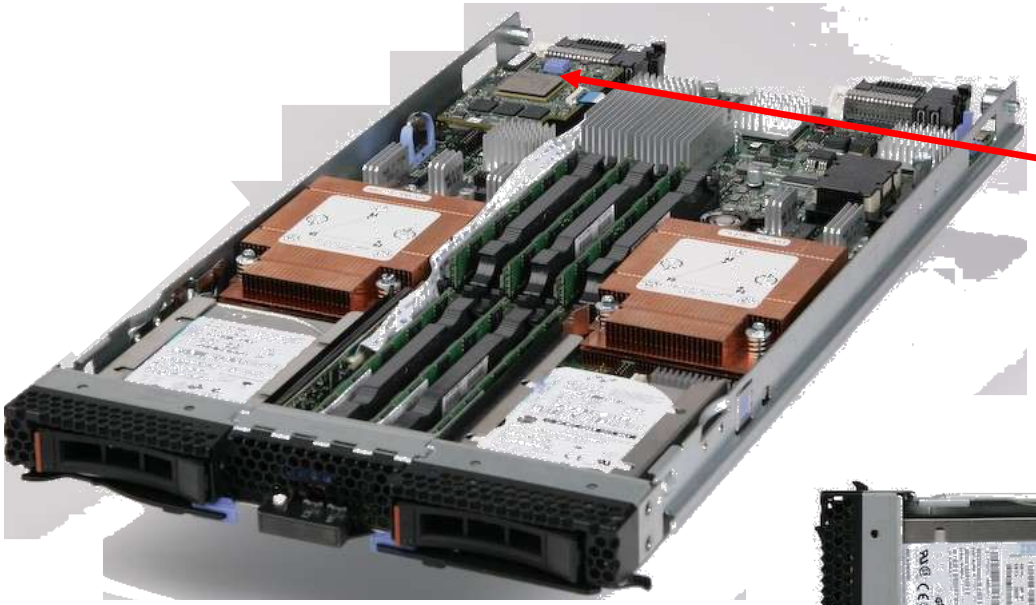

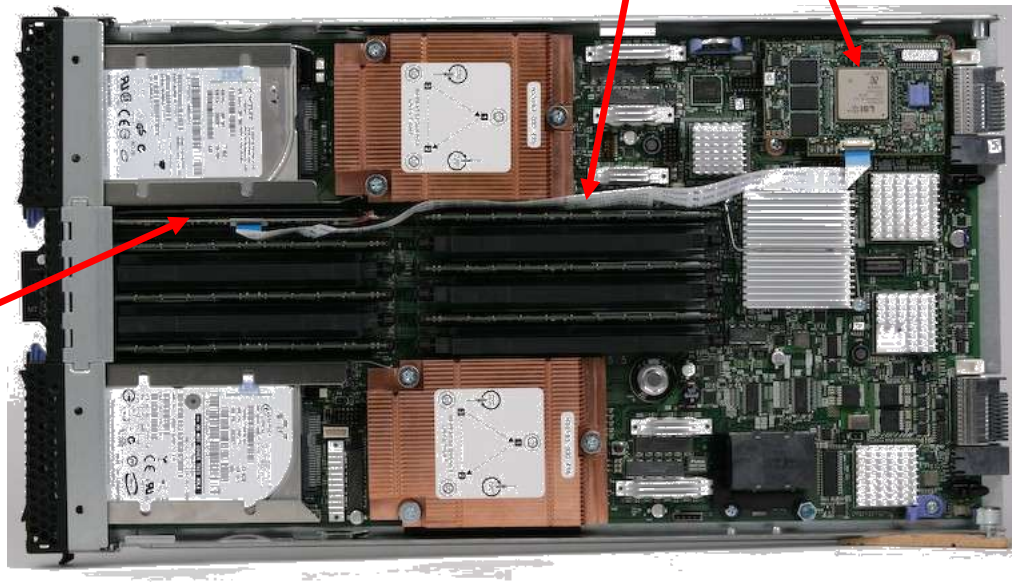

- SAS disks only….**No SATA or SSD**
- Raid Battery goes in DIMM slot 7

New SAS 2.5” SFF HDD

- New 73/146GB 15K 6Gbps SAS 2.5” SFF HDDs
  - 6Gbps ready – requires new adapter support to provide 6Gbps capability
  - Announce: 10 mar 09 Availability: 31 mar 09
    - Supported on HS12, HS22 (only ”Slim H/S”) and new 3550M2/3650M2
      - (42D0672) 73GB Slim-Hot Swap HDD
    - Supported on HS12, HS22 (only ”Slim H/S”) new 3550M2/3650M2 and ”Legacy” 2.5” platforms
      - (42D0677) 146GB Slim Hot Swap HDD
      - (42D0667) 146GB Non Hot Swap HDD
      - (42D0652) 146GB Hot Swap HDD

- New 146/300GB 10K 6Gbps SAS 2.5” SFF HDDs
  - 6Gbps ready – requires new adapter support to provide 6Gbps capability
  - Announce: 10 mar 09 Availability: 31 mar 09
    - Supported on HS12, HS22 (only ”Slim H/S”) and new 3550M2/3650M2
      - (42D0632) 146GB Slim Hot Swap HDD
    - Supported on HS12, HS22 (only ”Slim H/S”) new 3550M2/3650M2 and ”Legacy” 2.5” platforms
      - (42D0637) 300GB Slim Hot Swap HDD
      - (42D0627) 300GB Non Hot Swap HDD
      - (42D0612) 300GB Hot Swap HDD

IBM USB Key for VMware ESXi

- Now with **VMware ESXi 3.5 Update 4** (P/N 41Y8269)
  - installable on HS22
  - announced on 9 june 2009; available on 29 june 2009

NEW!

- **Positioning**
  - Simple and intuitive start-up experience for the new virtualization user providing easy on-ramp to virtualization
  - All the sophistication of VMware ESX and VI3 for existing virtual infrastructure customers

- **Initiatives driving the need for VMware ESXi**
  - Strengthen Security & Improve Reliability
  - Simplify Management of Hardware Resources & Virtual Machine Hosts
  - Enable Plug-&-Play Capacity Management in the Data Center

- **Merrill Lynch CIO Survey – April 2007**
  - "Adoption of x86 virtualization is growing rapidly"
    - Customers deploying x86 Virtualization grows from 8% to 18%
    - Intent to use increases from 68% to 85%
    - Expected deployment within 2 year moves from 29% to 58%
  - "Virtualization also goes hand in hand with server consolidation"
    - "74% of CIOs say they are moving toward bigger "scale up" machines and thus away from scale out architectures"

Part #: 41Y8268

Announce: Feb 17

GA: March 9

New High IOPS SSD

IBM Solid State Drive

- **50 GB SATA 3.5" HS High IOPS SSD (43W7698)**
- **50 GB SATA 2.5" SFF NHS High IOPS SSD (43W7706)**
- **50 GB SATA 2.5" SFF Slim-HS High IOPS SSD (43W7714)**
- **50 GB SATA 2.5" SFF HS High IOPS SSD (43W7722)**

Ann:10/03/09
GA: 31/03/09

NEW!

Blade servers that will be supported in

- HS21 (8853)
- LS22 (7901)
- LS42 (7902)
- HS21 XM (7995)
- HS12 (8014)

SATA compliance enables the IBM 50 GB SATA 2.5-inch HS SSD and 3.5-inch HS SSD to attach to **System x supported systems.**

- x3850M2/x3950M2 (7233,7241)
- x3755 (8877,7163)
- x3655 (7943)
- x3650 (7979)
- x3650 (7980)
- x3550 (7978)
- x3500 (7977)
- x3455 (7940,7941)
- x3400 (7973,7974,4192,4193)
- x3250M2 (4190,4191,4194)
- x3200 (4367,4368)

PERFORMANCE

50 GB performance characteristics include:

- Formatted capacity: Up to 50,000 MB
- Drive to host interface that supports 1.5 Gb/s burst rate
- Internal transfer read rate: Up to 80 MB/s
- Internal transfer write rate : Up to 50 MB/s
- Combination 67% Read/33% Write transaction processing rate: up to 2600 IOPS
- Energy saving with as little as 2.1 W of power consumption per drive

This drive is ideal for delivering high performance IOPS and thereby relieving storage bottlenecks.

HS12 – New Function. New Price Point.

Same Trusted BladeCenter Solution

- 6 DIMMs DDR2 memory
  - Low cost 12GB, 24GB max
- Flexible disk options
  - 2 Not-Hot-Swap SATA (up to 80GB each)
  - 2 Hot Swap SAS, RAID 0 &1 support (up to 146GB each)
  - 2 Hot Swap SSD, RAID 0 & 1 support (up to 31.4GB each)

- Supports single, dual, and quad core Intel CPU's: from single-core 1.86Ghz, to 2.83Ghz quad-core (or NEW 3.0Ghz dual-core)
- All the same I/O as HS21
- Supports cKVM
- Supported in all chassis

- Pricing starts at $999

Great Feature Set
Attractive pricing
Perfect for the SMB &
Non multi threaded apps

- **Departmental Print Server**
- **File server**
- **Application / Web server**
- **Light Mail**

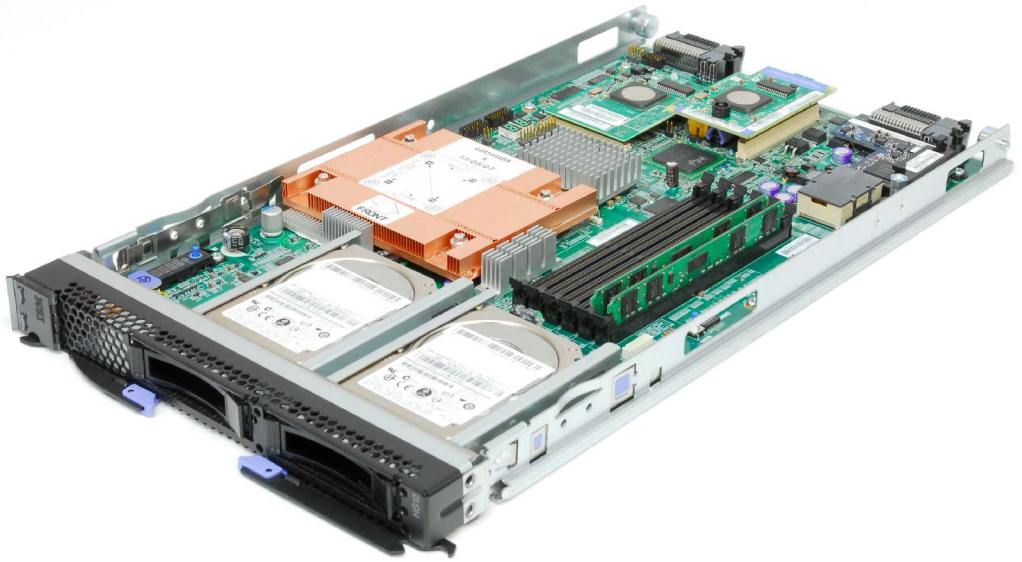

IBM Power System Blades

NEW!

NEW!

	IBM JS12	IBM JS22	IBM JS23	IBM JS43
				
Footprint, Packaging	Blade	Blade	Blade	Blade
Processor	POWER6	POWER6	POWER6+	POWER6+
# of processors (# of cores)	2	4	4	8
GHz clock	3.8	4.0	4.2	4.2
L3 Cache	0	0	32MB	32MB
DDR2 GB memory	4 to 64	4 to 32	4 to 64	8 to 128
Internal storage*	73GB – 600TB	73GB – 600TB	69GB – 600TB	69GB – 600TB
Maximum rPerf	14.71	30.26	36.28	68.2
PCIe PCI-X slots	1 1	1 1	2	4
Max I/O drawers	N/A	N/A	N/A	N/A
Max micro-partitions	40	40	40[1]	80[1]
IBM i Operating System	5.4 & 6.1	5.4 & 6.1	6.1	6.1
AIX® support	5.3, 6.1	5.3, 6.1	5.3, 6.1	5.3, 6.1
Linux® support	RHEL 4.6 / 5.1 SLES 10 / 11	RHEL 4.6/ 5.1 SLES 10 / 11	RHEL 4.6 / 5.1 SLES 10 / 11	RHEL 4.6 / 5.1 SLES 10 / 11

(1) Requires purchase of optional feature to support micro-partitions

*With maximum I/O drawers

Optional

LS22 & LS42 Value Proposition

Next generation AMD blades built for **quad & six core**

Introducing the industry's only snap-n-scale blade that can grow as your business grows…

…and won't break the bank!

LS22 / LS42

..now with "Istanbul" Six-Core processor (up to 2.6 Ghz)

Announce: 28 July 2009

Availability: 04 Sept 2009

AMD “Istanbul” six-core cpu

- **The LS22 blade server offers the following new processor:**
  - **AMD Opteron Model 2435 six-core processor with 3 MB L2 cache, a 6 MB L3 cache, and a speed of up to 2.6 GHz for energy efficiency**
- **The LS42 blade server offers the following new processors:**
  - **AMD Opteron Model 8431 six-core processor with 3 MB L2 cache, a 6 MB L3 cache, and a speed of up to 2.4 GHz to help with energy efficiency**
  - **AMD Opteron Model 8435 six-core processor with 3 MB L2 cache, a 6 MB L3 cache, and a speed of up to 2.6 GHz to help with energy efficiency**

Notes:

- AMD announced an increase of 34% in performance-per-watt over the previous generation of quad-core Opterons.
- The chip is rated at a thermal design power of 75 watts

LS22 Layout Diagram

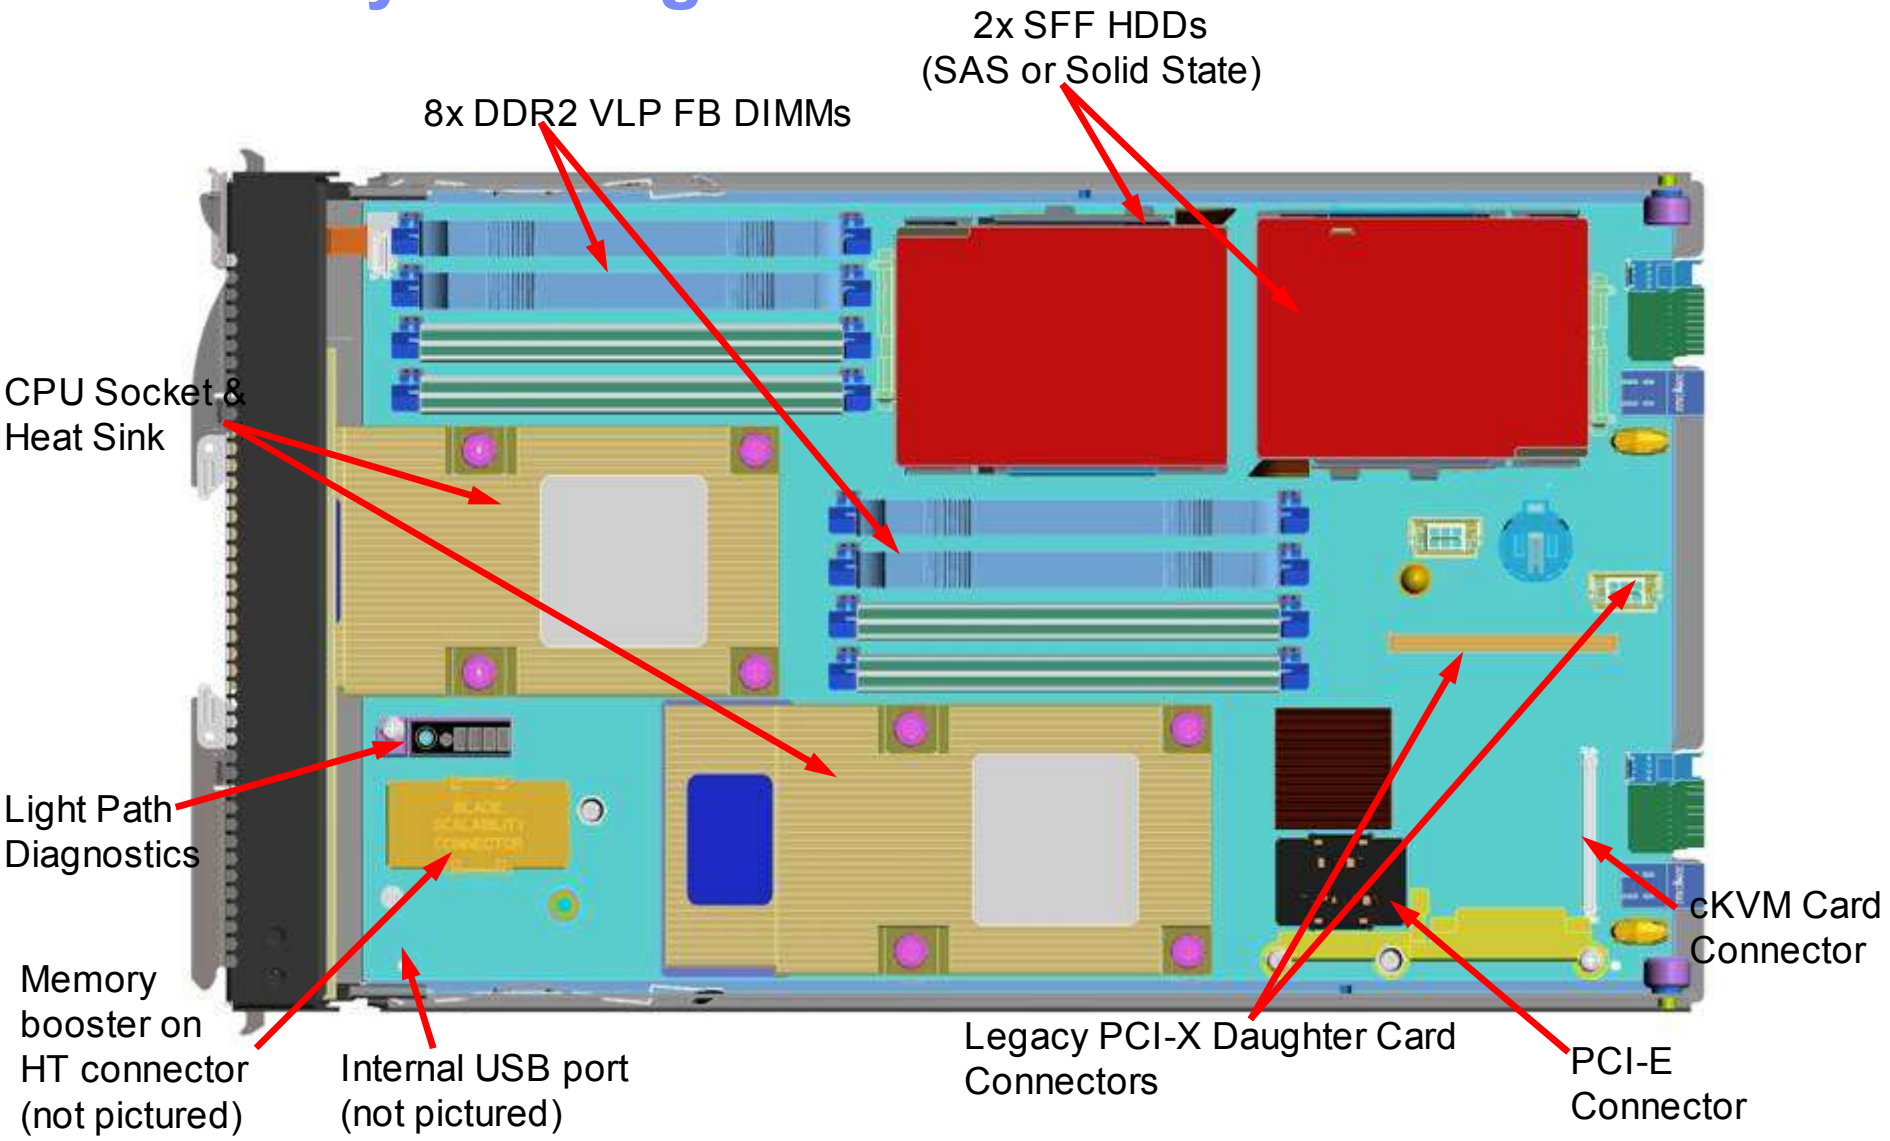

LS42 2+2 = investment protection you can afford Six-Core processor

Introducing the industry's only snap-n-scale blade that can grow as your business grows…

…and won't break the bank!

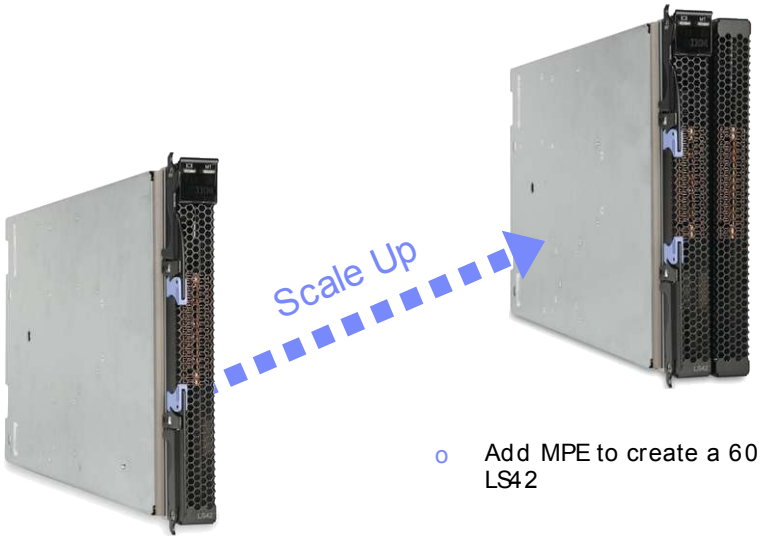

- Start with a 30mm, 2-socket scalable LS42
- **2x AMD 8431 or 8435 series six-core processors** (up to 2.6Ghz)
- Up to 64GB on 8 DIMMs
- 2x fixed SAS or SSD
- Ready to scale to 4 sockets with additional MPE

- Add MPE to create a 60mm, 4-socket LS42
- **4x AMD 8431 or 8435 series six-core processors** (up to 2.6Ghz)
- **Up to 128GB on 16 DIMMs trough 8GB dimm**
- 2x fixed SAS or SSD
- **embedded VMWare ESXi3.5 through an optional USB key**

LS22 / LS42

..now with "Istanbul" Six-Core processor (up to 2.6 Ghz)

Announce: 28 July 2009

Availability: 04 Sept 2009

Themis Joint Development with IBM and Sun: T2BC Blade Server

- **Themis has a long history building UltraSPARC workstations on a "Blade-like" form-factor**
- **IBM approached Themis about building an UltraSPARC based Blade Server for the BladeCenter**
- **Themis, IBM and Sun Engineering teams worked closely together to insure the T2BC integrated seamlessly into the BladeCenter Environment.**

UltraSPARC T2 Processor

Solaris Operating System

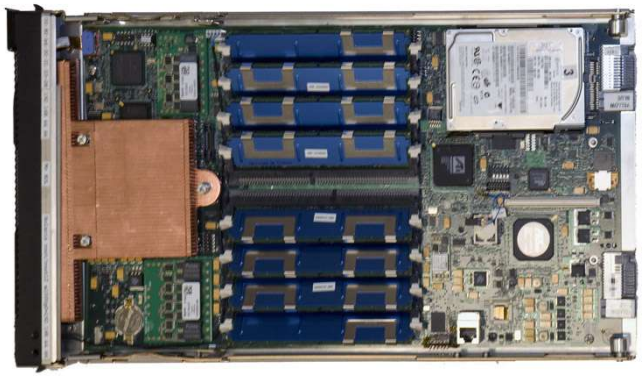

BladeCenter

Themis T2BC Blade Server

- One 4/8 cores 1.2 GHz UltraSPARC T2 Processor
- Up to 32 GB of FBDIMM memory in 8 DIMM Slots
- Network Fabric
  - Two Gbit Ethernet Ports Standard
  - Fibre Channel or InfiBand ports with Optional CFFV Daughter Cards
  - Two 10 Gbit Ethernet Ports with Second Optional Daughter Card
- 2.5” Internal Hard Drives: Two SATA or One SAS
- Front Panel Serial Port and Status LEDs
- System Manager/Environmental Monitor
  - Voltage, Temp
- Runs Solaris 10
- Based on Reference Design for Sun T5120 Server

IBM BladeCenter QS22

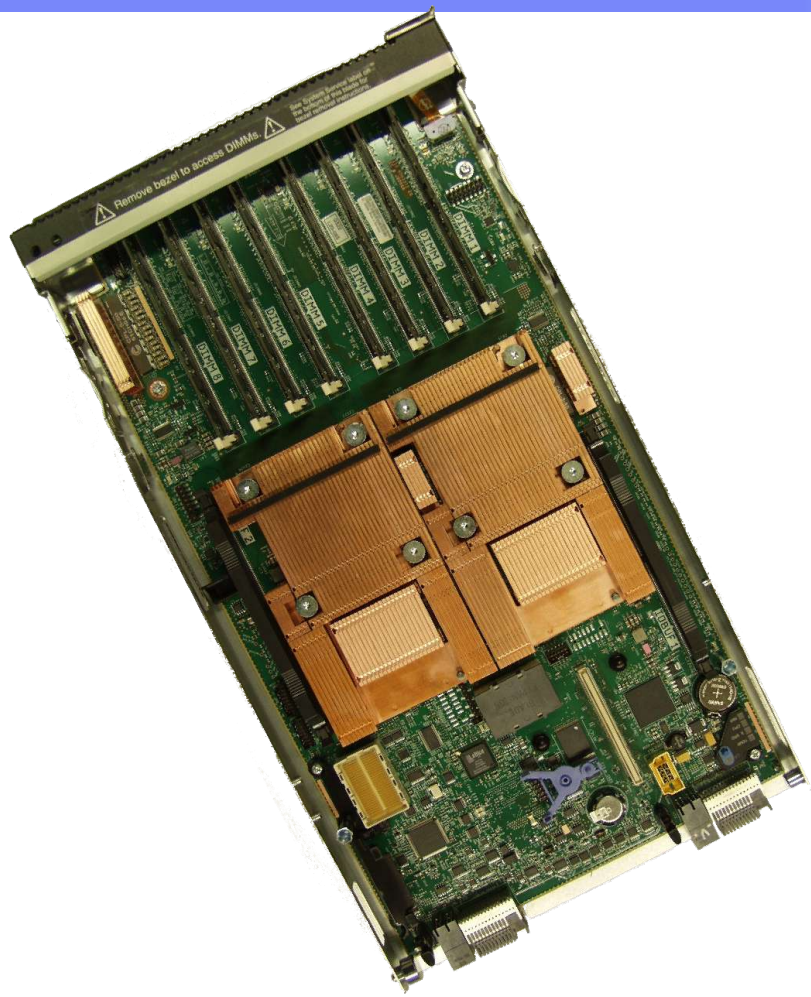

■ Third generation of IBM Cell blade is designed to drive production deployments for selected HPC applications and workloads that require double precision, large memory, parallel processing and/or streaming data

■ Target industries/workloads include: Medical Imaging, Electronic Design Automation (EDA), Aerospace and Defense, DVS, Seismic, Financial Markets, Healthcare and Life Sciences, Digital Media, etc.

■QS22 is complementary to other IBM systems based on Intel, AMD and POWER processors

PowerXCell 8i processor on QS22 offers:

- 10 times or more faster processing than traditional processors for targeted workloads
- 5 times faster double precision processing than previous generation blades
- 16 times more memory (maximum) than the previous generation blade

QS22

Announce: 13 may 2008

Availability: 6 june 2008

Cell/B.E. architecture reaches wide and deep – from consumer products to high performance computing

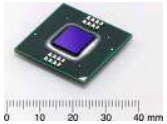

Toshiba SpursEngine (SPU's. + Host)

SCE PS3 (Cell/B.E. + GPU)

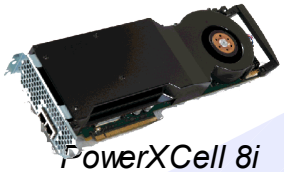

PowerXCell 8i PCI card (Cell/B.E. + Host)

Mercury 1u Dual Cell Sony Cell/B.E. Computing Unit (Cell/B.E. + GPU + AV I/O)

IBM BladeServer (2 Cell/B.E. or PowerXCell 8i)

Mini-Roadrunner Custom

Roadrunner (16,000 PowerXCell 8i. + AMD)

Consumer ***Business*** ***Enterprise*** ***High Performance Computing***

Common OS's, Infrastructure, Tools, Libraries, Code…
the SAME SPE code runs from *end to end*

Infrastructure Savings

- **Significant Infrastructure Savings with BladeCenter**
  - More than twice the density of 1U servers
  - up to 83% fewer cables than typical 1U
  - up to 64% fewer cables than our competitors

Per 42 Servers (IDE, dual SAN, dual enet, KVM, redun pwr)	1U Servers	Blades	Reduction / Addition
Rack Space	42U	21U	-50%
Ethernet Cabling	84	6 / 24	-71%
Fibre Channel Cabling	84	12	-86%
KVM Cabling	42	0 / 3	-93%
Systems Mgmt Cabling	42	6	-86%
Power Cords	84	12	-86%
PDU's	8	4	-50%
KVM Switches	6	0 / 1	-83%

Additional Notes: Networking, SAN and KVM Switching takes 10+ Power Cords and 2U to 8U in Rack
Power cord and PDU figures assume equivalent function of redundant power in 1U server

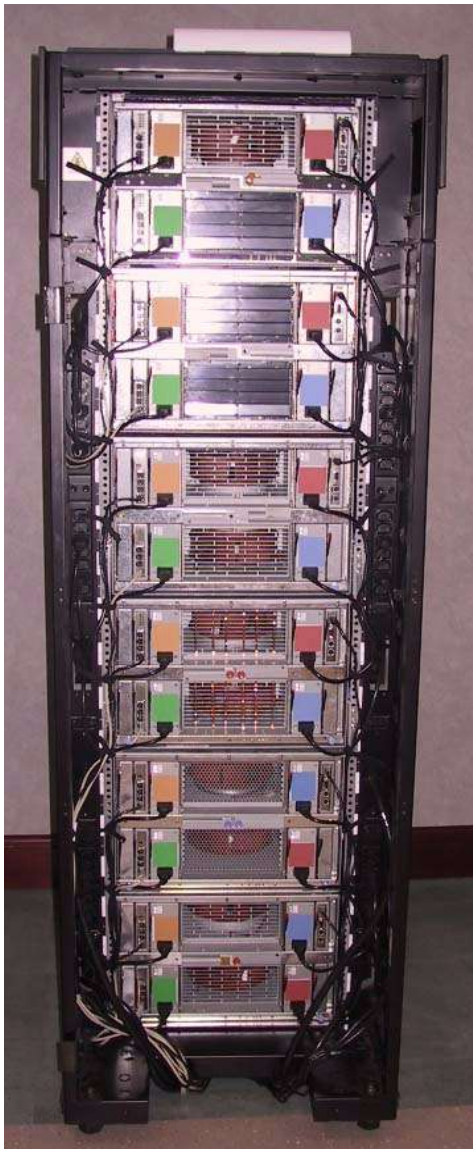

BladeCenter management

2. Director Console

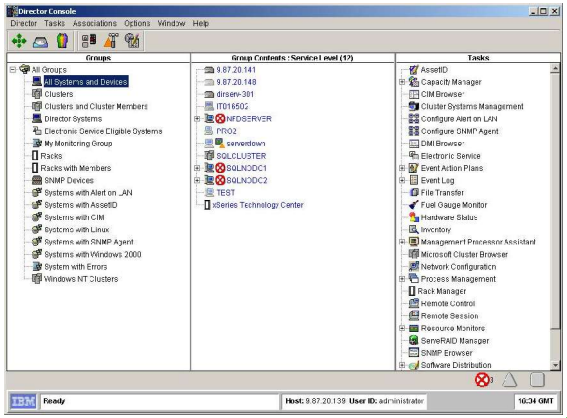

1. Web Interface

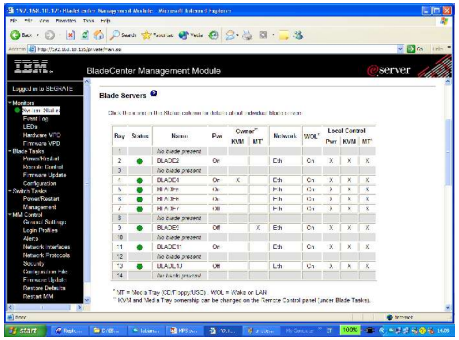

3. Telnet Interface

4. Command Line Interface (MPCLI)

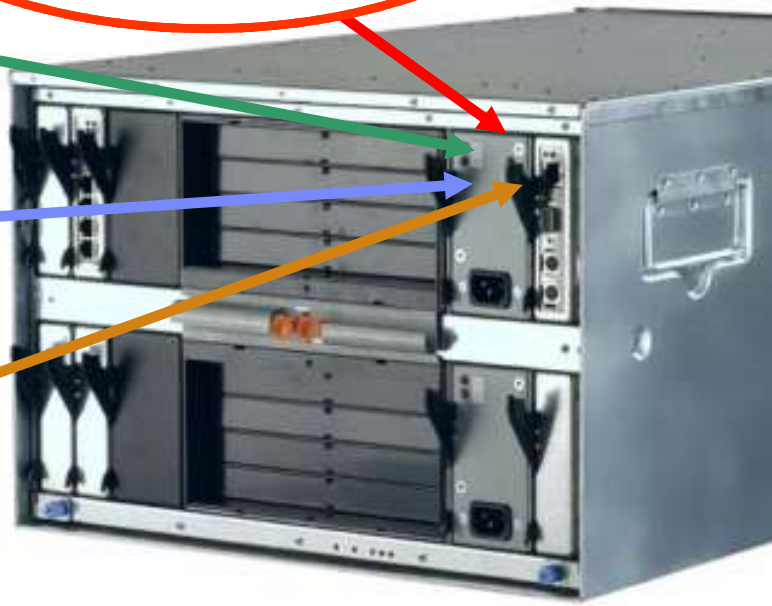

IBM Systems Director on x86, v6.1

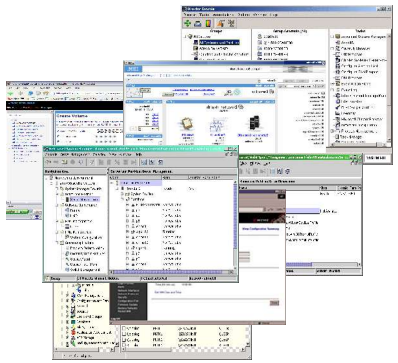

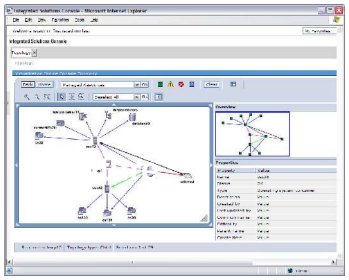

NEW!

Ann: 7 october 2008

Avail:

- web, 21 november 2008

-phys, 12 december 2008

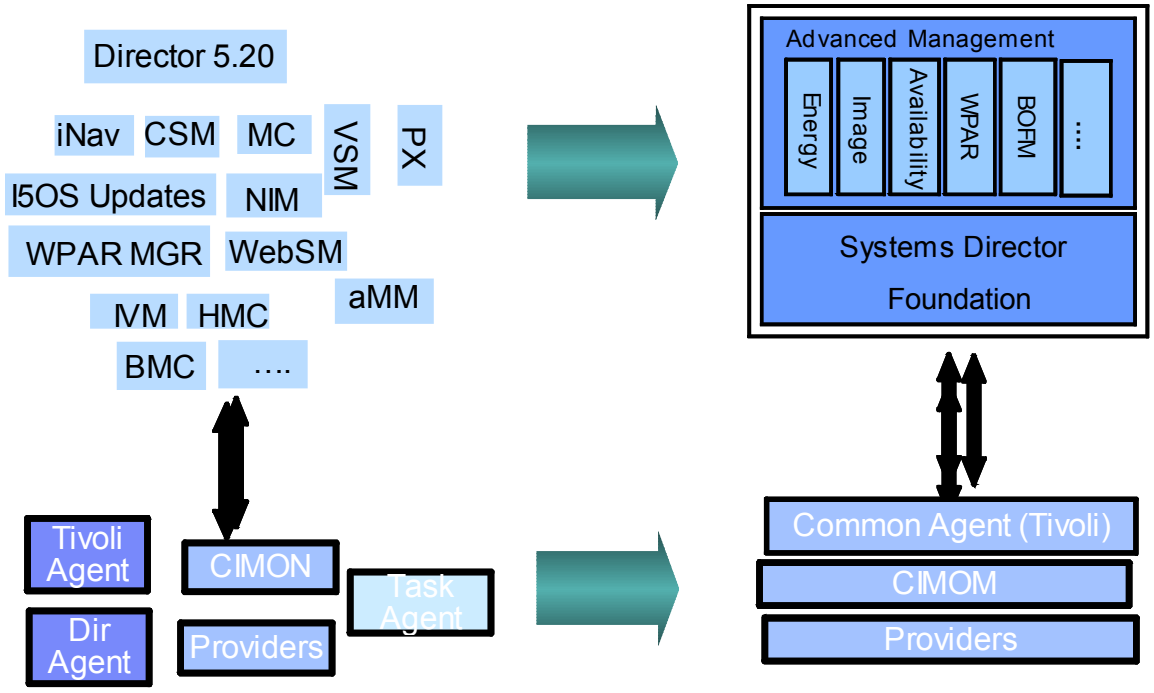

- Simplified Web User Interface with stream lined task integration and simplification
- Modular and extensible foundation enabling advanced platform management capabilities.
- Standards based data model and instrumentation delivering with our Systems and Storage
- Shared Common Agent with TPM 7.1

IBM Systems Director Foundation

Platform Lifecycle Management

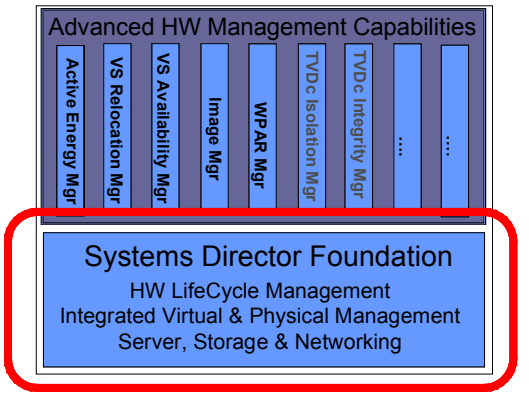

- **Consolidation of Platform Management Tools**
  - Single consistent cross-platform management tool
  - Simplified tasks via Web based interface
- **Integrated Physical and Virtual Management**
  - Discovery and Inventory of physical and virtual resources
  - Configuration and provisioning of platform resources
  - Status, Health, and Monitoring of platform resources
  - Visualization of server resource topologies
- **Platform Update Management**
  - Simplified consistent cross-platform tools to acquire, distribute and install firmware and OS updates
- **System Planning and Deployment Tools**
  - Systems and Virtualization planning for new systems
  - Initial system and OS deployment

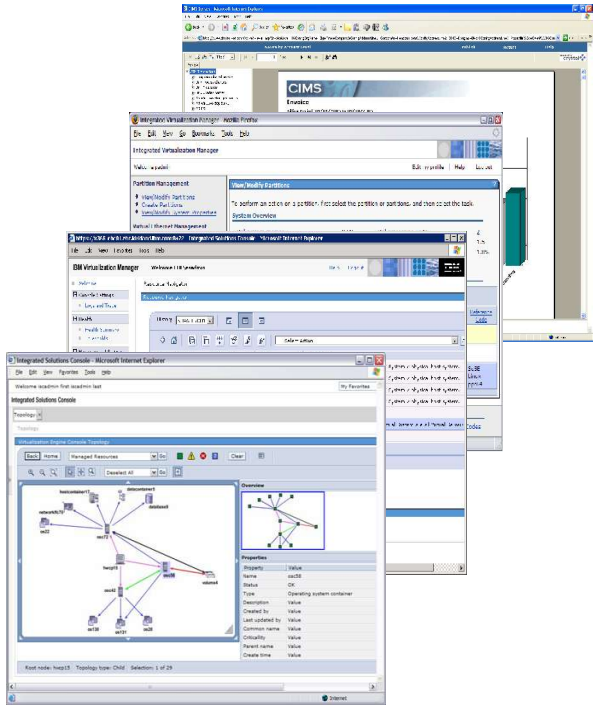

No Charge for Managers Included with Director Foundation

Platform Management Blueprint in Action:

E2E Management

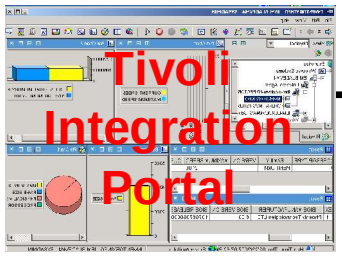

Tivoli Monitoring

Tivoli Provisioning Manager

Microsoft · ca · altiris · bmcsoftware SOLUTION PARTNER

Integration

Seed and Grow

3rd Party Tools

Basic Management

IBM® Systems Director

Physical and Virtual Management
Server, Storage, Networking

Foundation		Advanced Management	
Health	Updates	Image	Energy
Inventory	Virtualization	Availability	....

Required Management

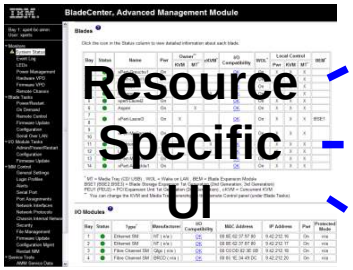

Operating Systems: z/OS · Windows · Linux · i5/OS · AIX L

Virtualization Software: vmware · Xen · Advanced POWER Virtualization · IBM TotalStorage SAN Volume Controller · Microsoft Virtual Server · z/VM

Hardware

Simplified Web Based User Interface

- **Setup and Configuration**
  - Start, manage, learn
- **Focus on Health, Automation**
  - Health Summary
  - Monitors
  - Thresholds
- **Easy to Navigate IBM Systems**
  - Power Systems summary
  - Intuitive drilldown
  - Topology map
- **Simply Manage Virtual Systems**
  - Create virtual server
  - Edit virtual resources
  - Relocate

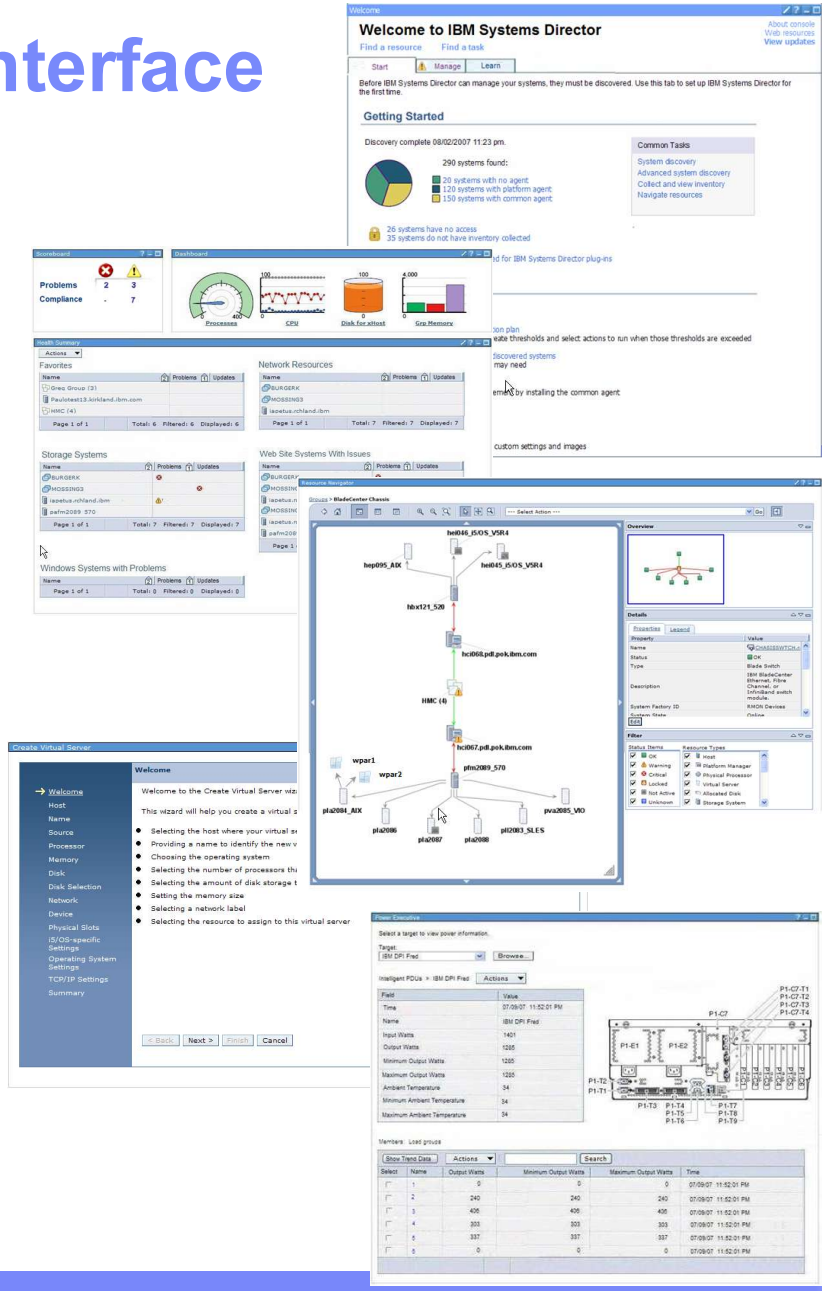

Advanced BOFM Upgrade Rel. 3.0

- Version for IBM Director 6.1
- Web Based User Interface
- Welcome page
- Template Wizards
  - Blade Address Manager Configuration Template Wizard
  - Standby Blade Pool Configuration Template Wizard
- Import/Export configuration as a CSV file
- SAS Support
- Evaluation and permanent License Keys Support
- Products key upgrade installer

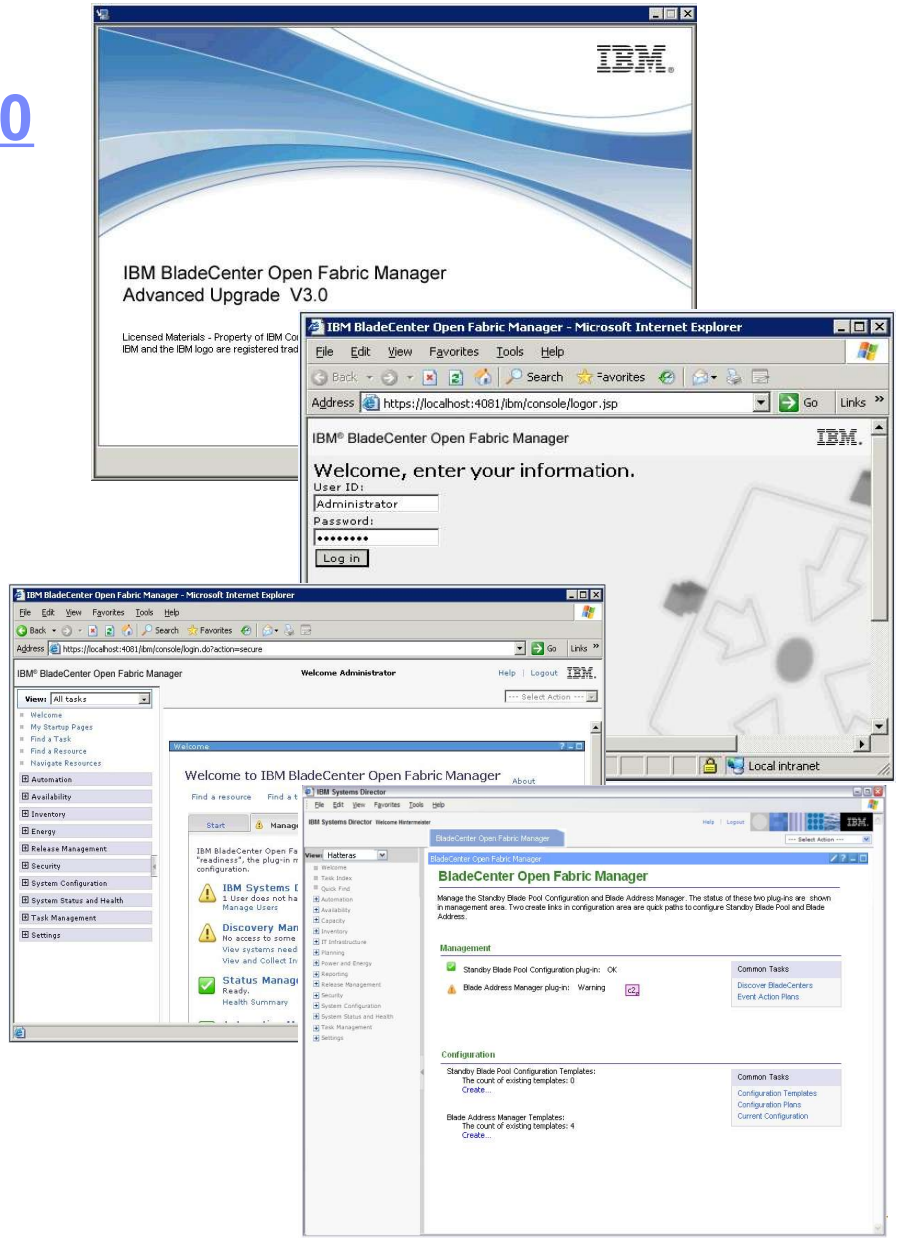

Announce: 16 june 2009

Availability: 26 june 2009

SAS failover support for NEW storage market!

Manual Rip-N-Replace (Basic BOFM):

1. Addresses are sent to the blade from the AMM as configured by BOFM
2. Blade fails….
3. Remove failed blade…
4. Replace with spare blade…
5. BOFM assigned I/O addresses are automatically sent to the new blade from the AMM thereby restoring SAN and LAN connections automatically.

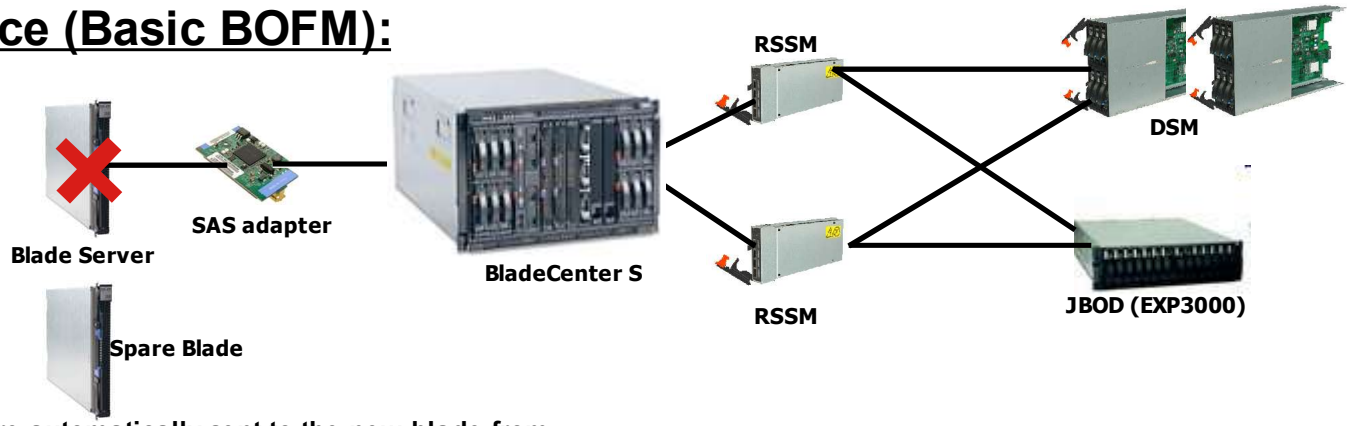

Automatic Failover (Advanced BOFM):

1. Addresses are sent to the blade from the AMM as configured by BOFM
2. Blade fails….
3. Advanced BOFM moves Ethernet MAC addresses and WWNs over to new blade, moves VLAN config from Ethernet switch to switch.

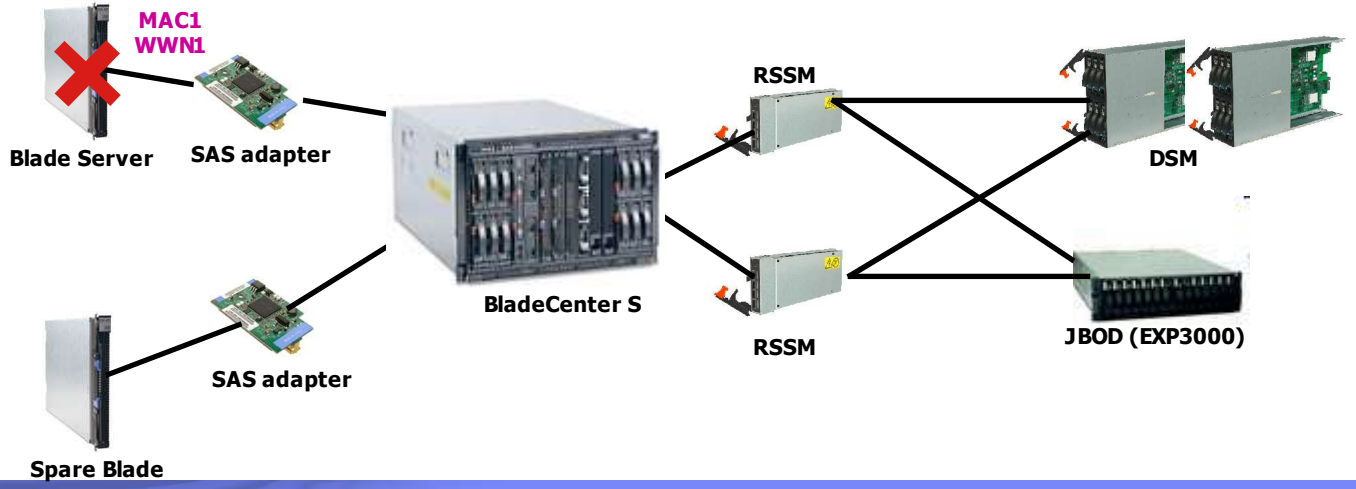

Supported Hardware

- Chassis:
  - BladeCenter S
  - BladeCenter E
  - BladeCenter T
  - BladeCenter H
  - BladeCenter HT
- Blades:
  - HS12 (1916, 8014, 8028)
  - HS21 (8853, 1885)
  - HS21XM (7995, 1916)
  - LS21 ((7971) / LS41 (7972)
  - LS22 (7971) / LS42 (7002) *
  - JS22
  - HS22
- On-board NIC:
  - Broadcom 1G Ethernet
- Expansion Cards:
  - Broadcom 1G Ethernet Cards (SFF & CFFv)
  - Qlogic 4Gb FC Cards (SFF, CFFh & CFFv)
  - Qlogic/Broadcom FC/Enet combo CFFh card
  - Emulex 4 Gb FC
  - Netxen 10 Gb dual port*
  - 2-Port SAS CFFv Expansion Card*
  - SAS Connectivity Card CFFv For IBM BladeCenter*
- External Storage Devices:
  - DS3400
  - DS4200, DS4700 & DS4800
  - DS8000
  - TS3500
  - TS3100/3200/3310
  - DS8000
  - SVC
- Ethernet Switches:
  - Server Connectivity Module
  - Nortel L2/3 ESM
  - Cisco Intelligent ESM
  - Nortel L2/3 10Gb Uplink ESM
  - Nortel Layer 2-7 ESM
  - Nortel 10Gb Switch Module
  - Cisco Catalyst 1 Gb*
  - Cisco Catalyst 10 Gb Uplink*
  - Copper Passthru Module (CPM)
- Fibre Channel Switches:
  - Qlogic 4 Gb
  - Brocade 4 Gb
  - Cisco 4 Gb
  - Optical Passthru Module (OPM)
- SAS Switches:
  - RSSM*
  - NSSM*

Energy innovation: CoolBlue strategy

- IBM systems architecture, chip design and software
- Systems Director Active Energy Manager
- Rear Door Heat Exchanger
- Calibrated Vector Cooling
- 90% efficient power supplies
- Power Configurator
- Thermal Analyzer

Worldwide Server Market (IDC) : il reale costo dei server

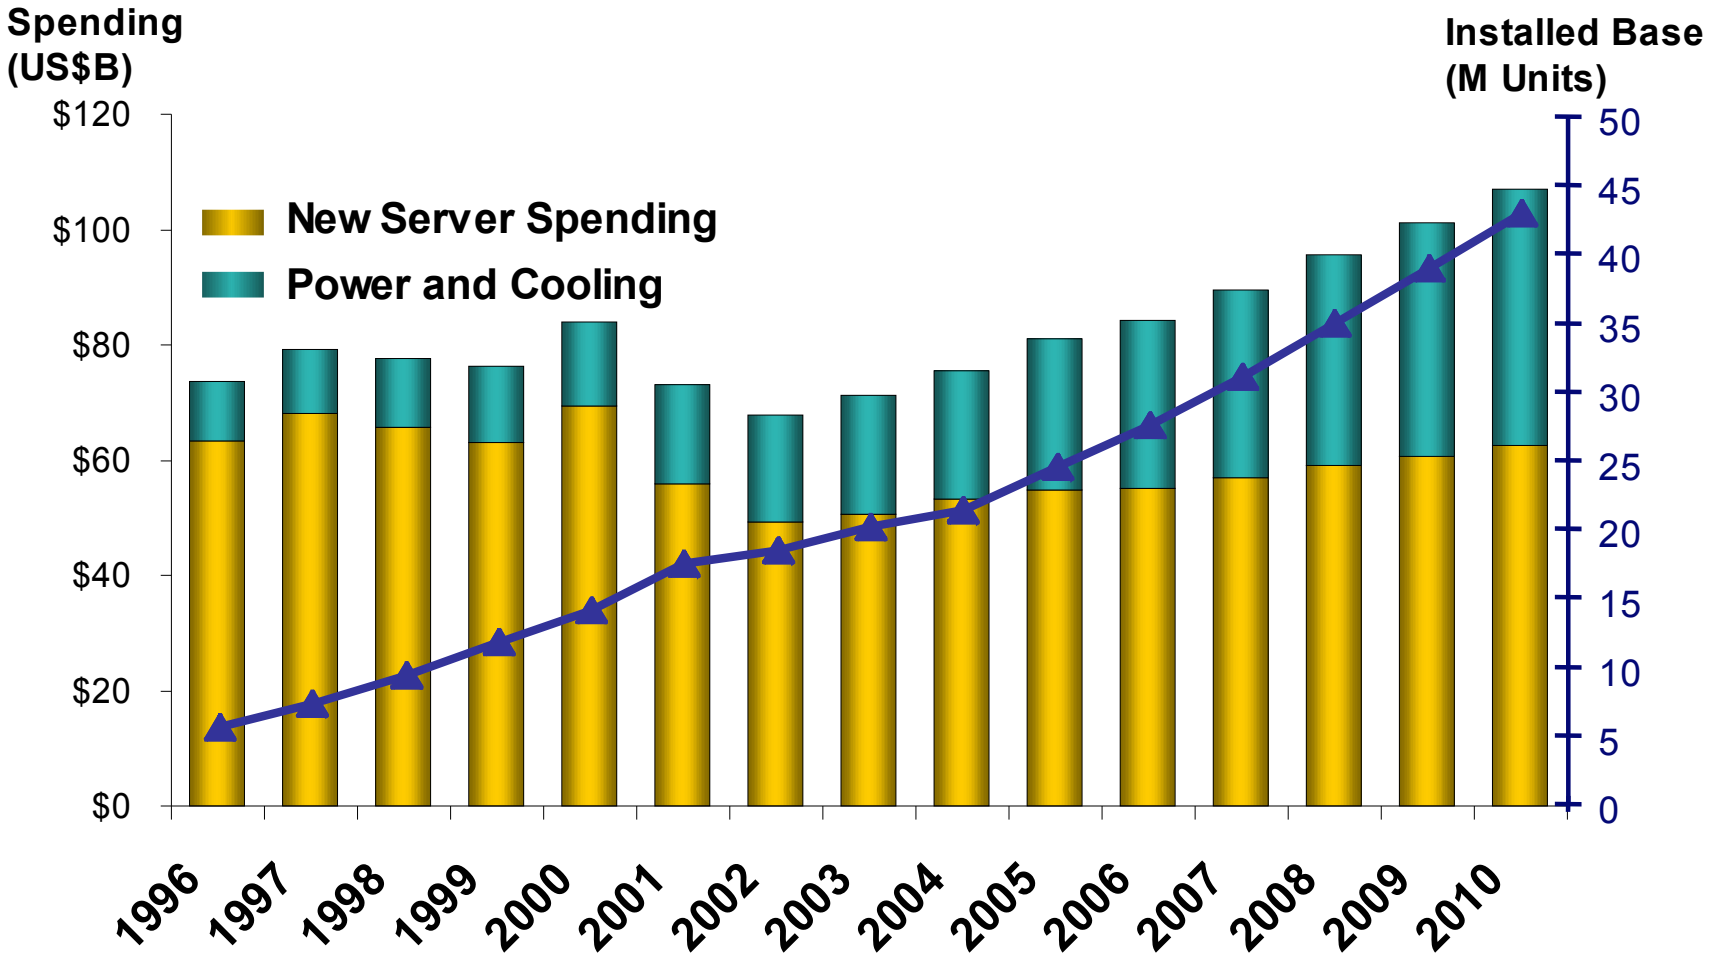

IDC Presentation, The Impact of Power and Cooling on Data Center Infrastructure, Doc #201722, May 2006

Question:
What is the greatest facility problem with your primary data center?

(Gartner 2006)

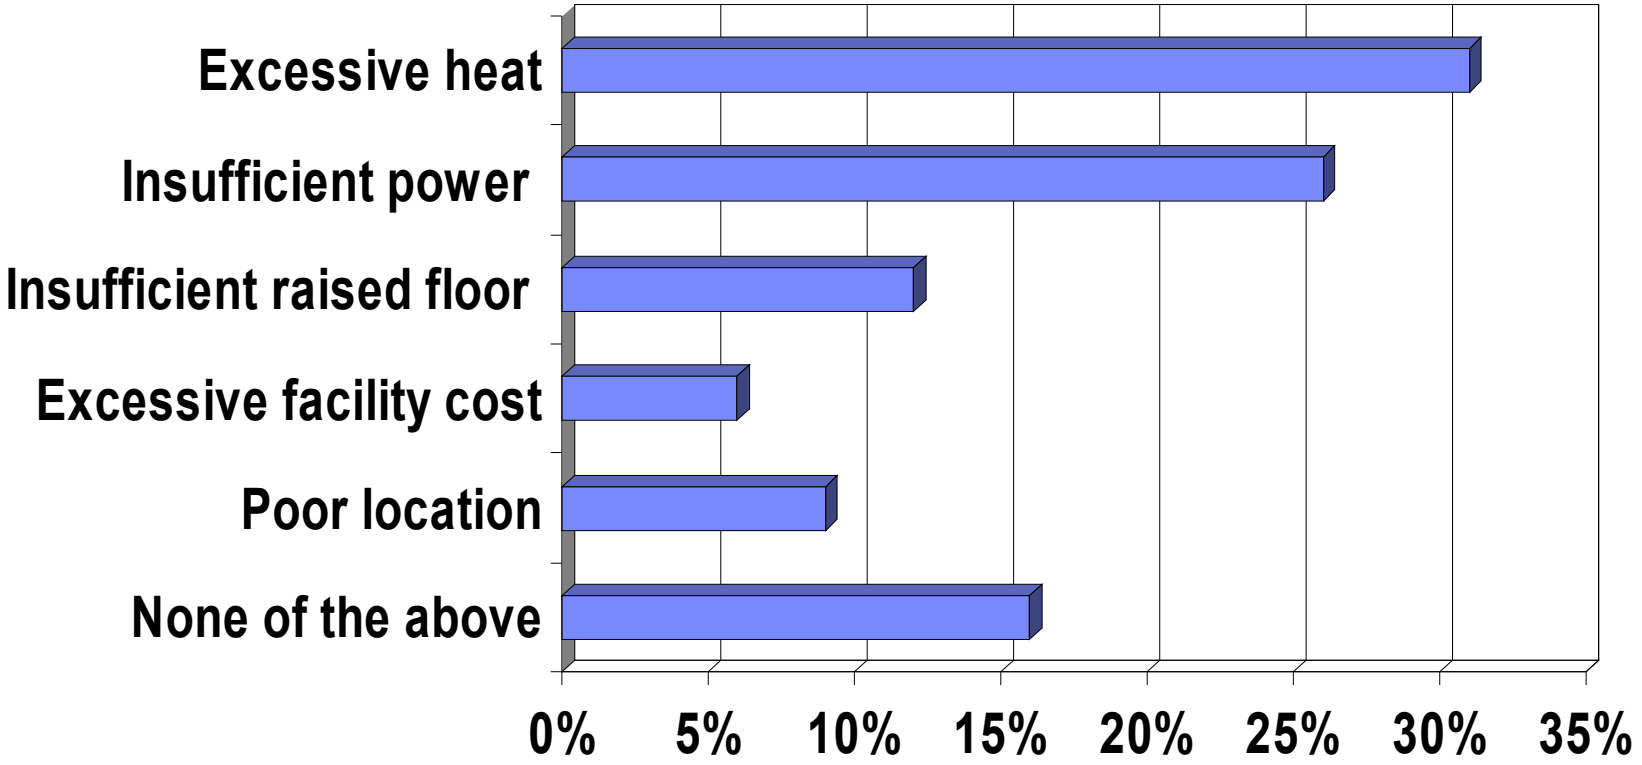

What's using the power?

The processor power growth is the largest single contributor but there are many other areas - the more you pack into a server the more power it needs!

BladeCenter helps in this area

Low Voltage processors help in this area

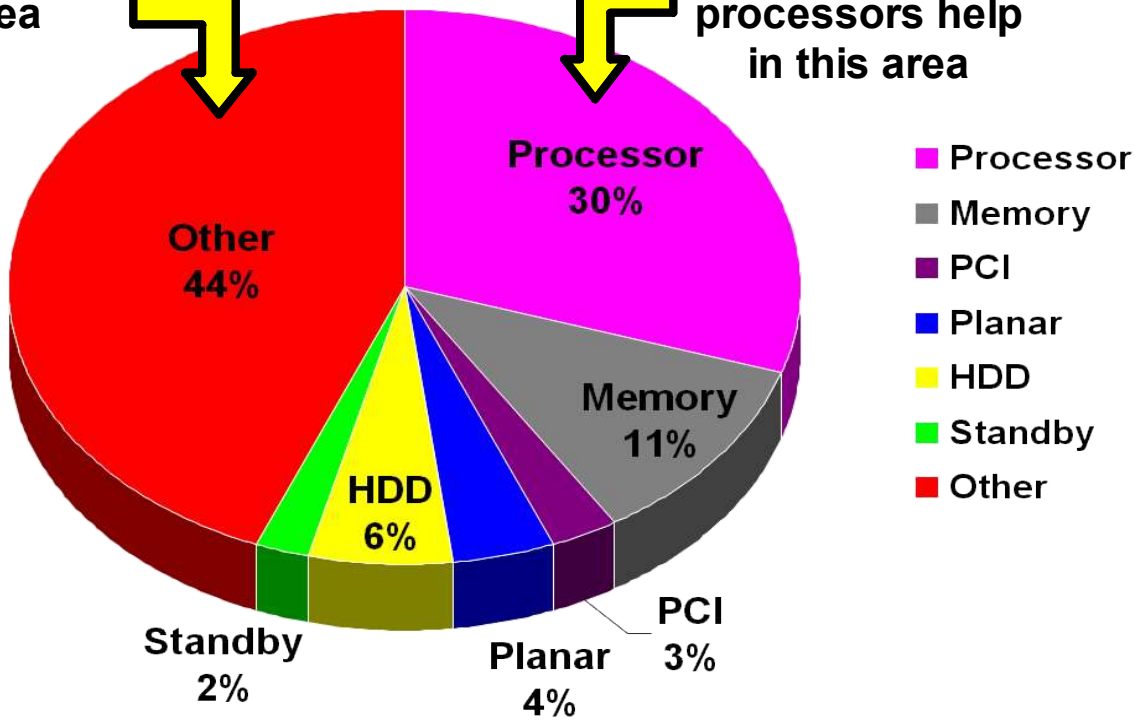

OTHER?

- AC to DC Transitions
- DC to DC Deliveries
- Fans and air movement

Increasing Power Efficiency

- **Power supplies typically work better at high loading/utilization**
  - **A problem for servers with N+N power**
  - **It's a cost issue, expensive to design supply components that work efficiently across range of load**
  - **BC has four supplies can afford and chose to invest in efficiency of our supplies**
    - **91% peak efficiency**
    - **nearly 85% efficiency at 20% load**

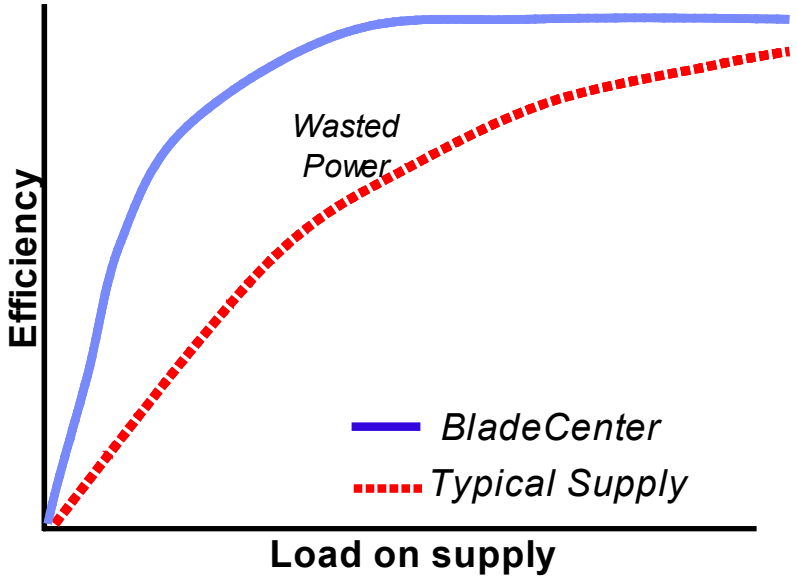

BladeCenter Cooling is Hands Down Better

(New HP cClass Changed Everything and Nothing)

IBM BladeCenter H

Two N+N Hot Swap Blowers + Power Supply/Fabric Fans

- **Few moving parts**
- **All Hot Swap**
- **Requires 100W of power in normal operation (<25C)**
- **Same design as BC - years of proven reliability**

Which Rack Do You Want in Your Data Center?

42 1U servers

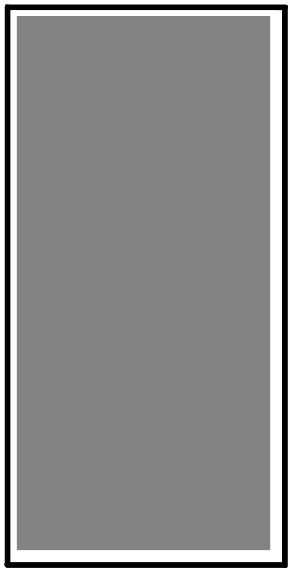

Same Number of servers

Same performance

Same function

OR

25-40% more servers at no additional power or cooling

42 HS21 Blades

21U white space

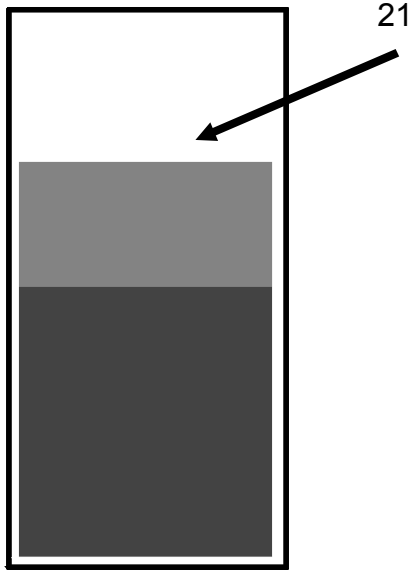

25-40% less power

25-40% less heat

Nearly 50% less air flow

40% less weight

IBM® *Systems Director Active Energy Manager* ™ in action!

Manage Power at the rack and server level

Monitor:
- **Power supply**
- **Blades**
- **I/O modules**
- **Mgmt module**
- **Blowers**
- **Media tray**
- **Midplane**

Compare actual vs. name plate power at system level

View inlet and exhaust temperatures

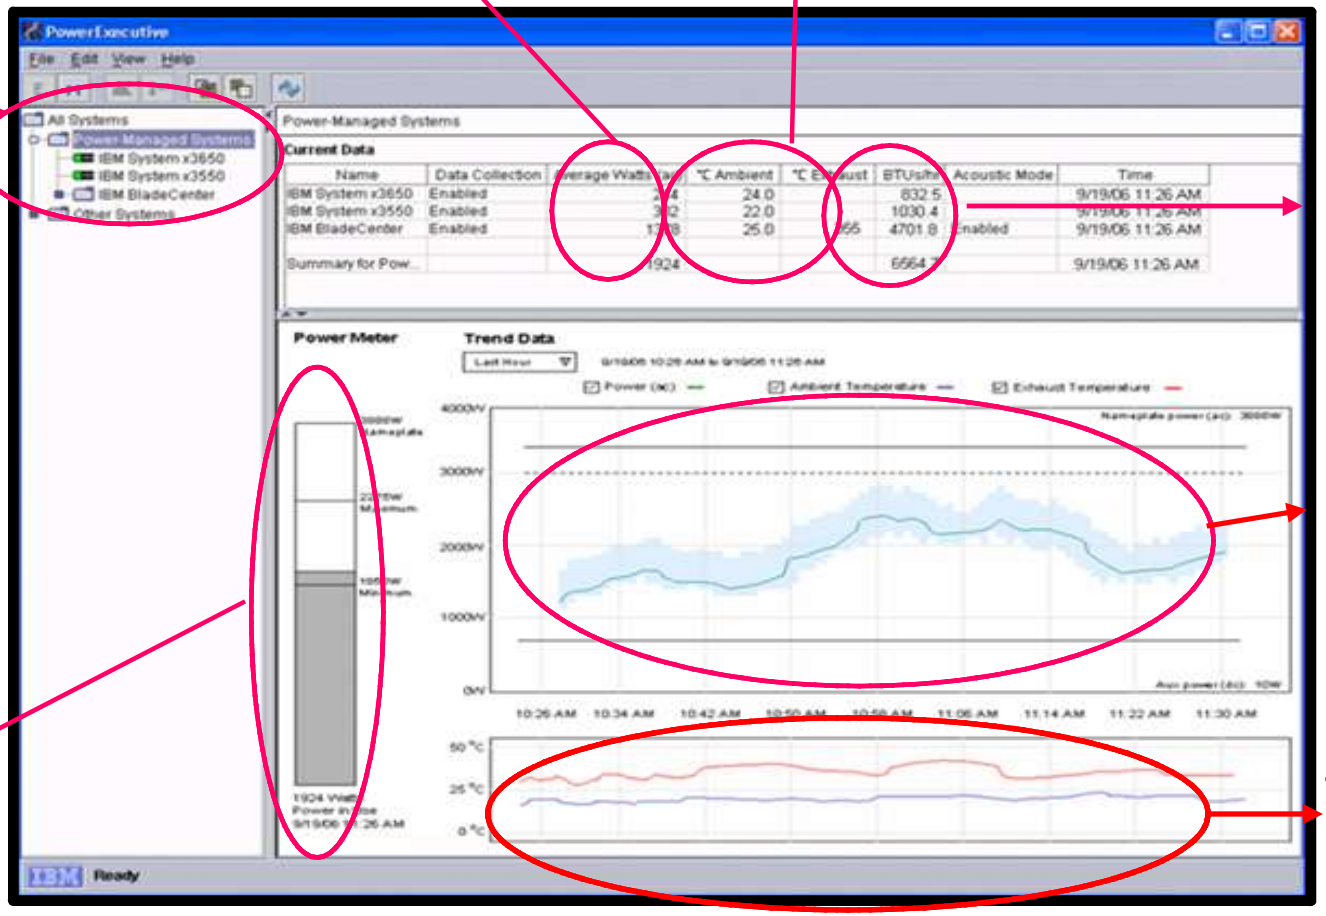

Track heat emitted

Trend power use over time

Trend temperature over time

Compare rack actual power vs. Label Power

High Performance Systems : x3850 M2 – x3950 M2

- **Embedded Hypervisor Models - 3i**
  - Announce: 25 march 08
  - Availability: 9 may 2008
    - Allows system to boot as virtualized platform
    - Eliminates hours of primary partition installation activities

7 Steps to Success - Benefits of X4 over X3-

1) Separate FSB for each processor
2) FSB frequency increased from 667Mhz to 1066Mhz
3) 2.6x aggregate FSB bandwidth improvement for higher performance
4) Memory bandwidth increased from 333Mhz to 533 Mhz
5) 4X memory capacity (2X from chipset, 2X from DRAM technology)
6) Internal chipset bandwidth and queues increased to support 4 core processors
7) Almost 2X increase in scalability port bandwidth for improved scaling

Evolving X3 into X4

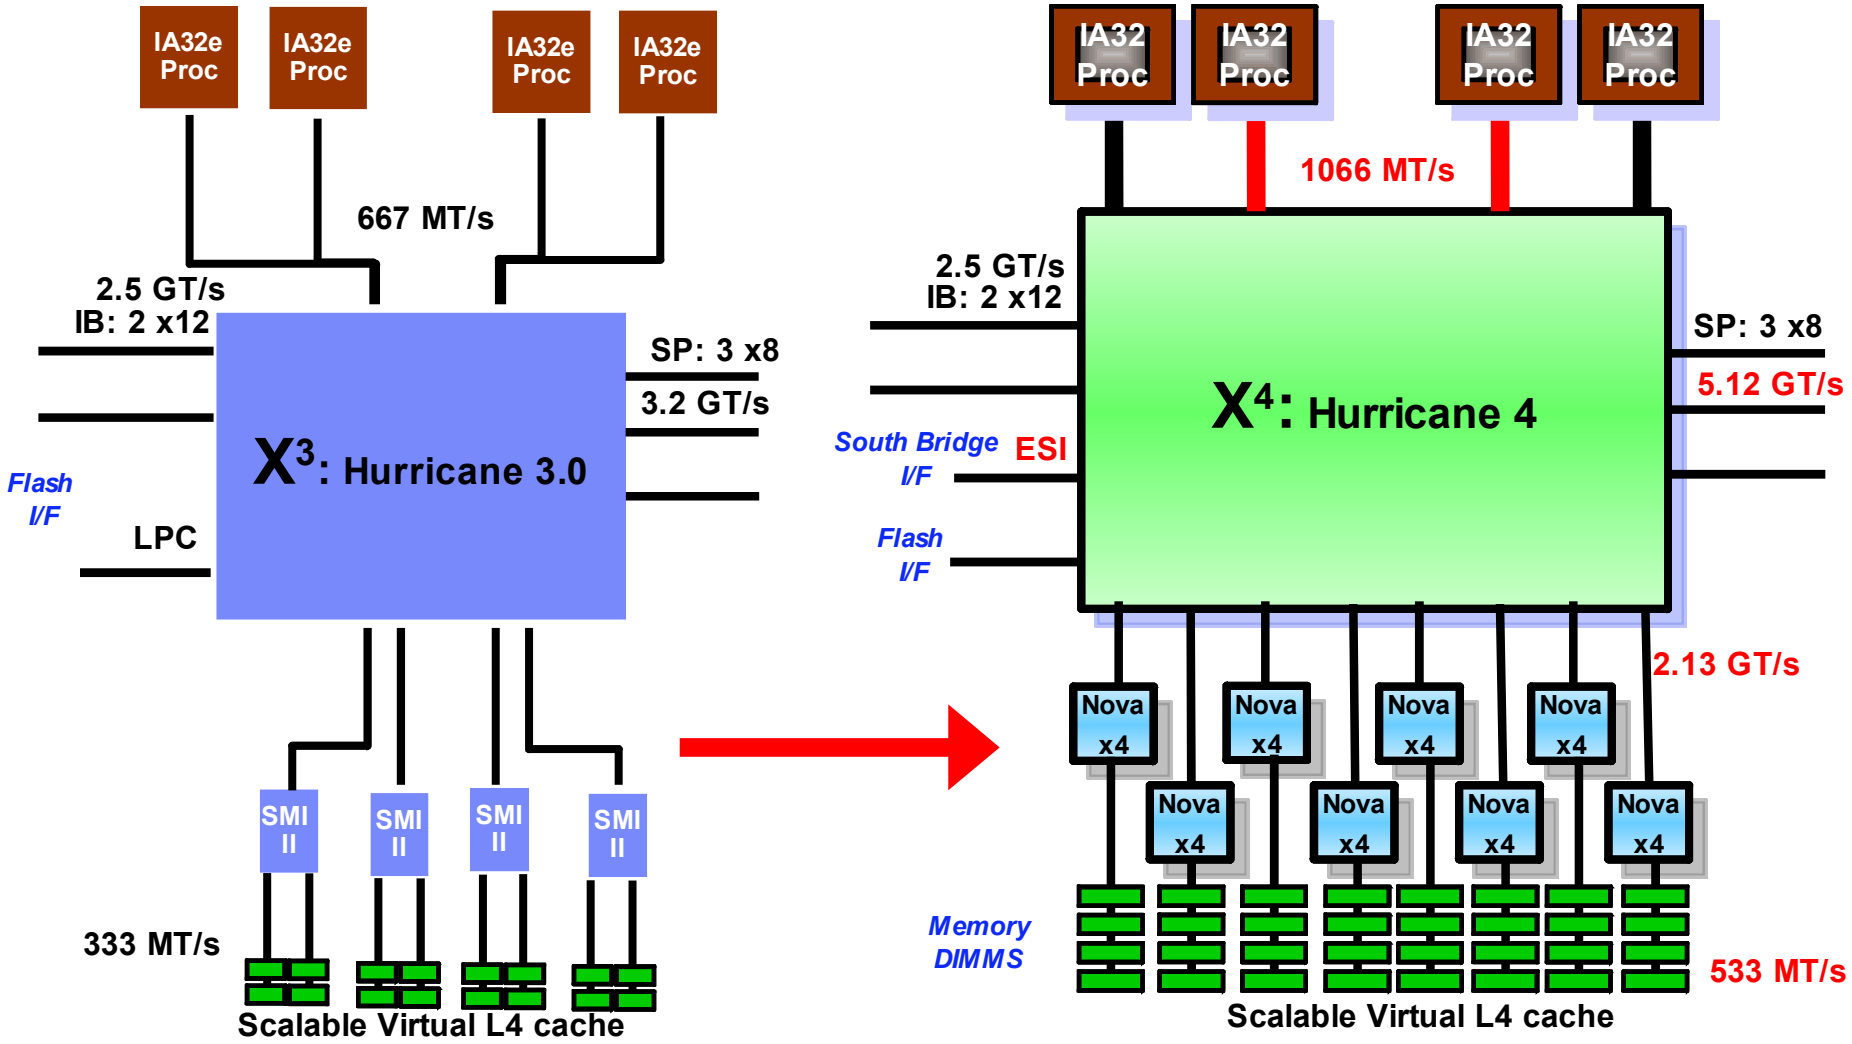

CPU's & Chipset Hurricane 4 Description

Processors

Intel Xeon Processor "Tigerton"

- x7350 (2.93GHz 8MB L2 Quad Core)
- E7330 (2.40GHz 6MB L2 Quad Core)
- E7320 (2.13GHz 4MB L2 Quad Core)
- E7310 (1.60GHz 8MB L2 Dual Core)

Intel Xeon Processor "Dunnington"

- E7420 (2.13 GHz 6 MB L2 Quad Core)
- E7450 (2.4 GHz 9 MB L2 Six Core)
- x7460 (2.66 GHz 9 MB L2 Six Core)
- L7445 (2.13 GHz 6 MB L2 Quad Core)

Ann:16/09/08
GA:10/10/08

Processor Interface

- Quad 1066 FSB - 16GB/s
  - 2.6x increase over Hurricane 3.0
- Snoop Filter for Quad FSB coherency tracking

Memory

- 8 Novax4 EI3 channels
- 34GB/s Read, 17GB/s Write
- DDR2/3-533, 1GB, 2GB, 4GB, 8GB DIMMs
- X4 Chipkill w/RBS, Memory mirroring, Hot Add/Replace

Direct connect I/O

- 2 x12 IB Ports – 12GB/s
- Connects to Calgary (PCIX) and Calioc2 (PCIE)
- ESI support for ICH-8 South Bridge

Scalability

- 3 Scalability Ports – 30GB/s
- Up to 16 sockets

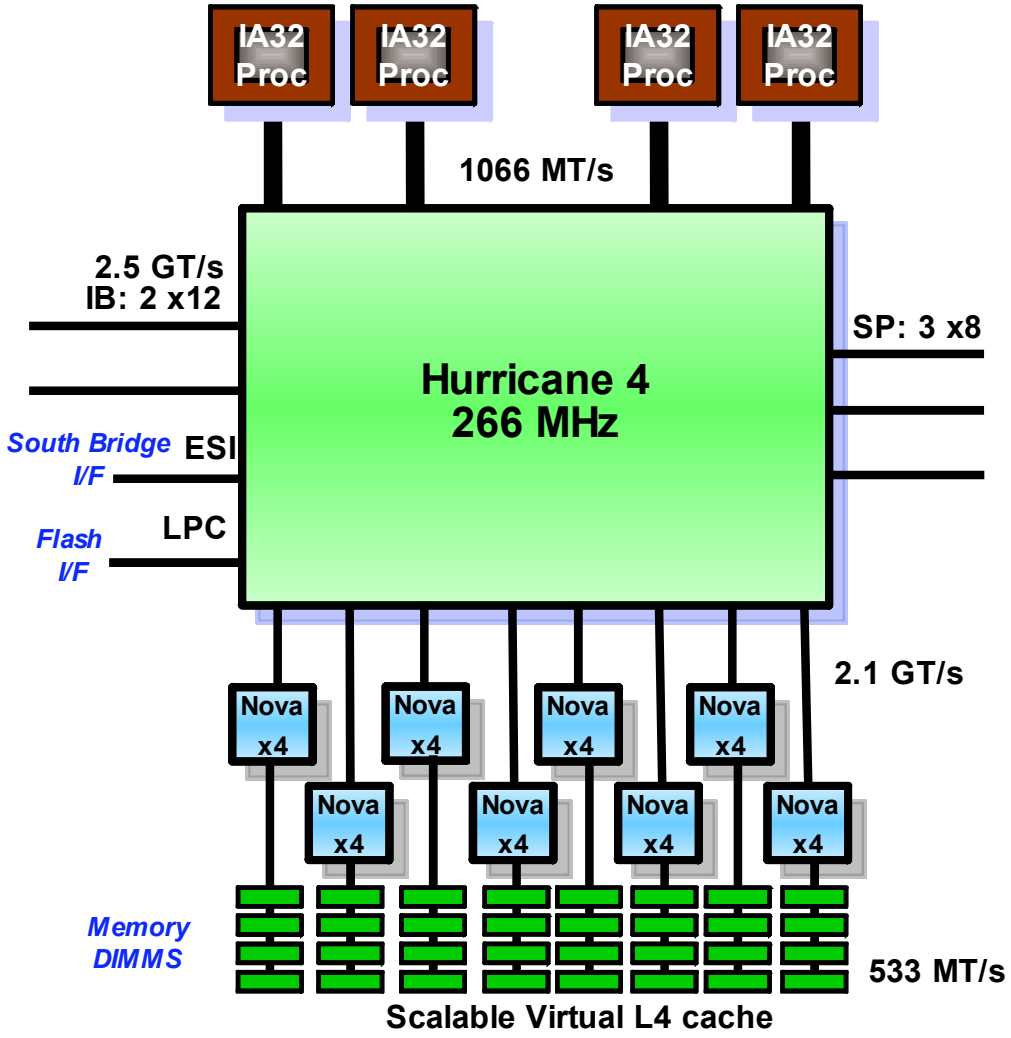

Multi Chassis configuration

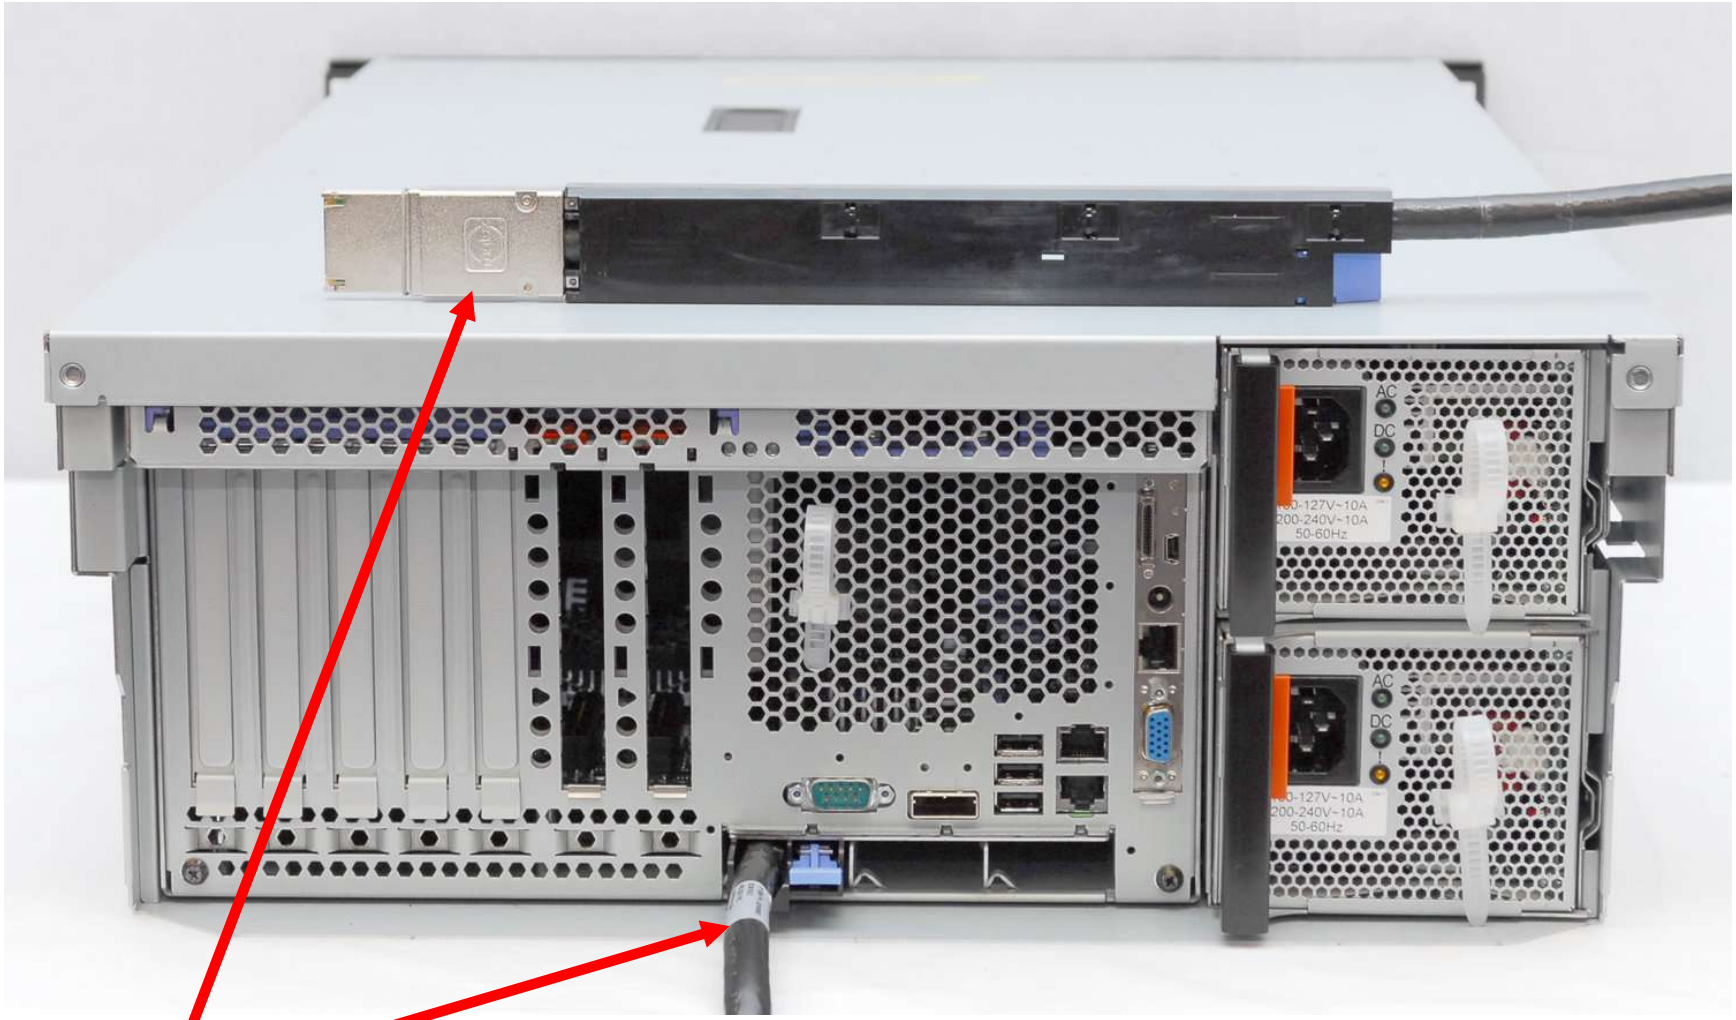

"Athena" Scalability Upgrade: ..from x3850 M2 to x3950 M2

x3850 M2

A 4U rack-optimized model, expandable to 16U (..and 16 CPU sockets!) using a "ScaleXpander Option Kit" and three additional chassis.
Installing the ScaleXpander Option Kit turns the x3850 M2 into an x3950 M2 (including a new bezel)

ScaleXpander Option Kit

- P/N 44E4249
- Available from 15 feb 2008
- Scalability Cable 3.08m
- Entry Cable Management Arm
- ScaleXpander chip

x3950 M2

IBM System x3950 M2-specific: ScaleXpander Kit *standard*

Scalability Icon lights up when active

IBM eX4 Unmatched Reliability, Scalability and Performance

Mainframe Inspired Solutions for the x86 Market

For more information visit eX4 Sales Support

Top 5 Reasons to purchase an IBM x3850 M2 / x3950 M2

1. Most Reliable x86 Platform for your Mission Critical Apps
   - Includes more reliability features then any industry x86 server
     - IBM Chipkill™ Memory
     - Memory ProteXion™
     - Hot Swap PCI & Memory
     - Predictive Failure Analysis
2. Low power and high utilization for maximum efficiency
   - Multi-processor utilization, DDR2 memory technology, and lower wattage energy efficient power supplies for lower total cost of ownership
3. Greatest top end performance x86 platform
   - Large 64 Dimm Memory Capacity, IBM eX4 Technology, and scalability to 16 sockets delivers leadership performance
     - #1 tpc-c x86 score
     - #1 VMmark 24-core score
4. Lowest cost for large memory and socket licensed software
   - 10% cost savings for high memory capacity, combined with 40% lower memory latency and 60% more memory bandwidth ability to have 64 Dimms with only two processor, creates ideal memory rich platform
     - 22% lower cost per SAP SC
     - 22% lower cost per VMWare VM
5. Investment protection with 'pay as you grow' scalability
   - Modular scaleable design allows your system to grow as your business grows up to 4 times an original fully loaded configuration, without any additional initial cost expenditure

High Performance Systems

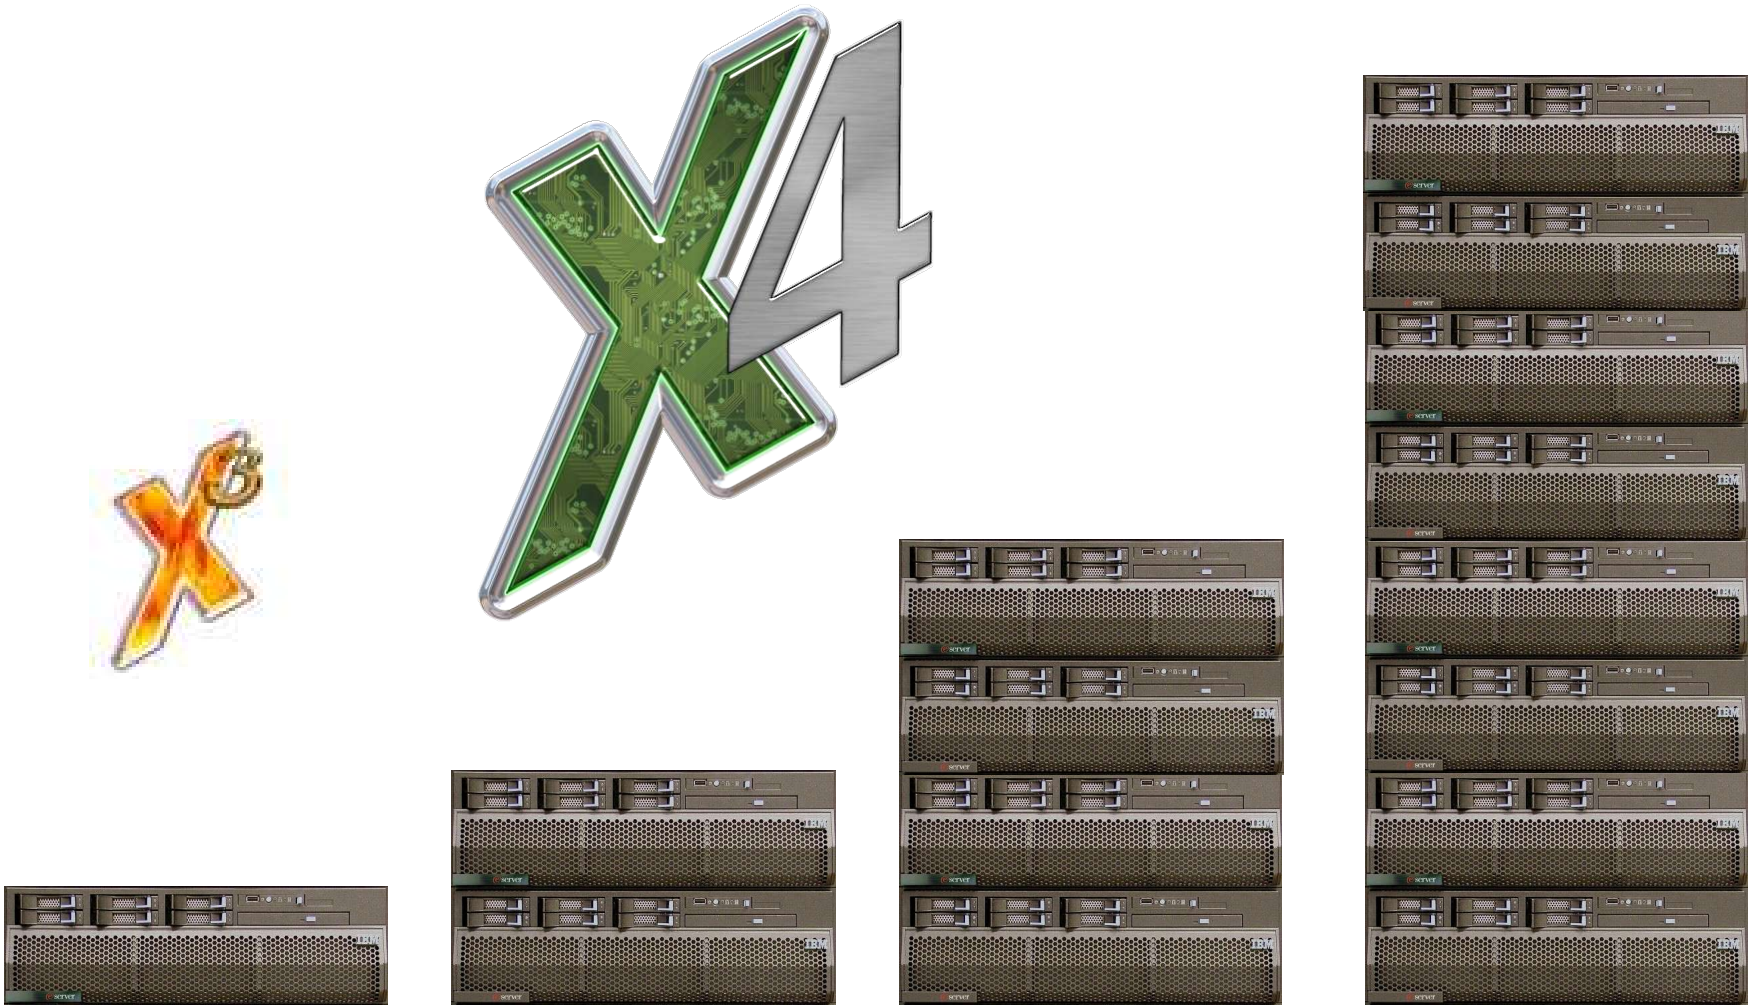

Decade Of Innovation

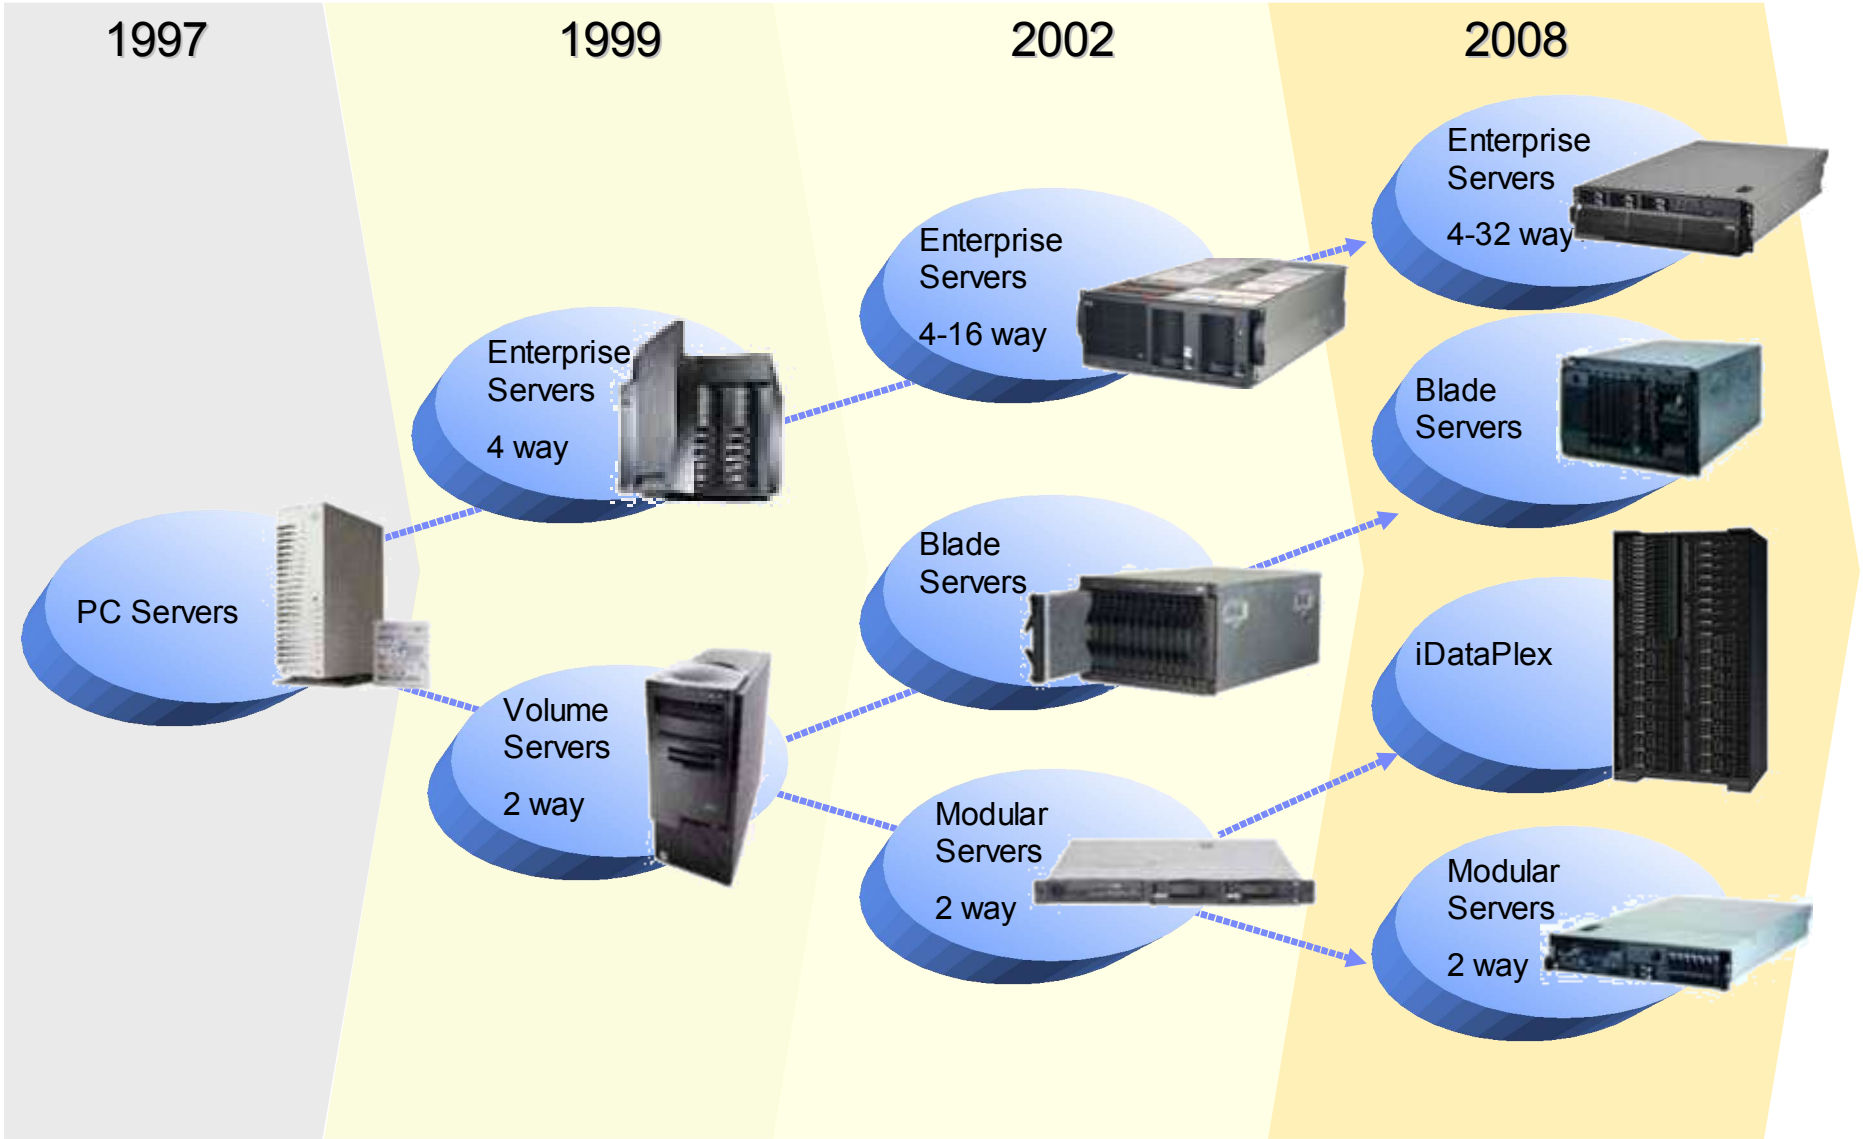

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Many of the features described in this document are operating system dependent and may not be available on Linux. For more information, please check: http://www.ibm.com/systems/p/software/whitepapers/linux_overview.html

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised January 19, 2006

Special notices (cont.)

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, AIX/L, AIX/L (logo), alphaWorks, AS/400, BladeCenter, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, ClusterProven, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DirectTalk, Domino, DYNIX, DYNIX/ptx, e business(logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, i5/OS, IBM, IBM(logo), ibm.com, IBM Business Partner (logo), Informix, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, Passport Advantage, POWERparallel, Power PC 603, Power PC 604, PowerPC, PowerPC(logo), PowerPC 601, Predictive Failure Analysis, pSeries, PTX, ptx/ADMIN, RETAIN, RISC System/6000, RS/6000, RT Personal Computer, S/390, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SpaceBall, System/390, The Engines of e-business, THINK, Tivoli, Tivoli(logo), Tivoli Management Environment, Tivoli Ready(logo), TME, TotalStorage, TURBOWAYS, VisualAge, WebSphere, xSeries, z/OS, zSeries.

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: Advanced Micro-Partitioning, AIX 5L, AIX PVMe, AS/400e, Chipkill, Chiphopper, Cloudscape, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, e-business(logo), e-business on demand, eServer, Express Middleware, Express Portfolio, Express Servers, Express Servers and Storage, GigaProcessor, HACMP, HACMP/6000, I5/OS (logo), IBMLink, IBM TotalStorage Proven, IMS, Intelligent Miner, iSeries, Micro-Partitioning, NUMACenter, ON DEMAND BUSINESS logo, OpenPower, POWER, Power Architecture, Power Everywhere, Power Family, Power PC, PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, System i, System i5, System p, System p5, System Storage, System z, System z9, S/390 Parallel Enterprise Server, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Virtualization Engine, Visualization Data Explorer, X-Architecture, z/Architecture, z/9.

A full list of U.S. trademarks owned by IBM may be found at: http://www.ibm.com/legal/copytrade.shtml.
UNIX is a registered trademark in the United States, other countries or both.
Linux is a trademark of Linus Torvalds in the United States, other countries or both.
Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States and/or other countries.
Intel, Intel Xeon, Itanium and Pentium are registered trademarks or trademarks of Intel Corporation in the United States and/or other countries.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States and/or other countries.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.
AltiVec is a trademark of Freescale Semiconductor, Inc.
Other company, product and service names may be trademarks or service marks of others.

Revised January 19, 2006

Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM System p5, ~ p5, pSeries, OpenPower and IBM RS/6000 Performance Report at http://www.ibm.com/servers/systems/p/hardware/system_perf.html.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, AIX Version 4.3 or AIX 5L were used. All other systems used previous versions of AIX. The SPEC CPU2000, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C Enterprise Edition V7.0 for AIX, XL C/C++ Enterprise Edition V7.0 for AIX, XL FORTRAN Enterprise Edition V9.1 for AIX, XL C/C++ Advanced Edition V7.0 for Linux, and XL FORTRAN Advanced Edition V9.1 for Linux. The SPEC CPU95 (retired in 2000) tests used preprocessors, KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors. Other software packages like IBM ESSL for AIX, MASS for AIX and Kazushige Goto's BLAS Library for Linux were also used in some benchmarks.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
NotesBench	http://www.notesbench.org
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle Applications	http://www.oracle.com/apps_benchmark/
PeopleSoft - To get information on PeopleSoft benchmarks, contact PeopleSoft directly	
Siebel	http://www.siebel.com/crm/performance_benchmark/index.shtm
Baan	http://www.ssaglobal.com
Microsoft Exchange	http://www.microsoft.com/exchange/evaluation/performance/default.asp
Veritest	http://www.veritest.com/clients/reports
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised January 19, 2006

Notes on Performance Estimates

rPerf

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

rPerf estimates are calculated based on systems with the latest levels of AIX 5L and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM @server pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration.

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

Revised January 19, 2006