

Standards-based, Energy-efficient HPC Systems: Trends, Implementations and Solutions.


BAI Yun Fei

HP ISS APJ /SP&HPC/Business Development Manager



CHALLENGES TO INNOVATION

CUSTOMER OBJECTIVES



Deeper insight

Smarter decisions

Create profit

OBSTACLES

- Increasing scale and complexity of deployments and applications
- Exploding volumes of data
- Pressure to provide instantaneous results
- Affordability
- Power and space constraints
- Technology and market risks and uncertainty
- Adapting to new business models and technologies

HP = HPC LEADERSHIP FOR A CHANGING WORLD

Accelerating innovation through a converged infrastructure for the HPC data center

HPC solutions,
purpose-built for scale

Holistic datacenter portfolio

Adaptive and scalable software
for HPC

Worldwide HPC expertise and
engagement

Performance

Efficiency

Agility

Confidence

Purpose-built HPC Servers



The Most Successful Architecture Ever to Enter the TOP500 – the **BL-Series (c-Class)**



HP BladeServer c-Class 2p servers (subset)

	BL280c G6	BL460c G6	BL490c G6	BL465c G7 (June)
Processor	Up to 2P, up to 6c Intel Xeon 5500/5600 series			Up to 2, up to 12-Core AMD Opteron 6100 Series
Max Memory	12 DDR3 slots Max memory: 192GB		18 DDR3 slots Max memory: 192GB	16 DDR3 Sockets Max memory: 256GB
Storage	2 non-hot plug SFF SATA/SAS/SSD drives	Up to 2 Hot Plug SFF SAS/SATA	Up to 2 non-hot plug SSD	Up to 2 Hot Plug SFF SATA/SAS/SSD
Networking	2 integrated Multifunction GbE ports	2 integrated Multifunction 10GbE ports with Flex-10 support		
Form factor	16 per 10u enclosure 8 per 6u enclosure			
Usage	General Compute	General Compute, with hot-plug drives & 10GbE	Large Memory, with 10GbE	General Compute, with hot-plug drives & 10GbE



New Performance/Density for HPC: HP ProLiant BL2x220c G6



BL2x220c G6	
Processor	Two 80W or 60W dual- or quad-core Intel Xeon 5500 Series processors per server node*
Memory	Registered or Unbuffered DDR3 6 DIMM Sockets per server 96GB max per server
Internal Storage	1 Non-Hot Plug SFF SATA HDD per server
Networking	2 integrated 1GbE Ethernet ports per server
Mezzanine Slots	1 PCIe Gen2 x8 mezzanine expansion slot per server
Additional Features	Internal USB 2.0 connector Optional internal SD Card slot (consumes the USB slot)
Management	ProLiant Onboard Administrator powered by iLO2
Density	32 server nodes in 10U enclosure

* 95W processors available through the SCI Private-plus process




PURPOSE DRIVEN SCALE-OUT PRODUCT LINES



Density optimized for the data center



Shared infrastructure for accelerated service delivery



Extreme scale out datacenters with lean management

	DL	BL	SL
Design center	Rack	Blade enclosure in rack	Rack
Design focus	Versatility & value	Integrated & optimized, maximum redundancy	Cost & features optimized for extreme scale out
Application	General purpose	General purpose / private cloud / scale out	Web 2.0 / cloud / scale out
Management	Essential and advanced management HP Insight Dynamics	Advanced management-accelerated service delivery & change in minutes	Home grown management Basic management via IPMI/DCMI

Purpose-built HPC Nodes



Highly Flexible SL Chassis



Benefits: Low cost, high efficiency chassis

- 4U Chassis for deployment flexibility
- Standard 19" racks, with front I/O cabling
- Unrestricted airflow (no mid-plane or I/O connectors)
- Reduced weight
- Individually Serviceable Nodes
- Variety of optimized Node Modules
- Ability to mix and match nodes

Multi-node, Shared Power & Cooling Architecture

- Shared Power & Fans
- Optional Hot-Plug Redundant PSU
- Energy efficient Hot Plug fans
- 3 Phase Load Balancing
- 94% Platinum Common Slot Power Supplies
- N +1 Capable Power Supplies (up to 4)



SL Advanced Power Manager Support


- Power Monitoring
- Node Level Power Off/On

HP ProLiant SL Scalable System

Next generation breakthrough server family optimized for

SL390
CPU+GPU

- Affordable scale
 - Lower acquisition cost than traditional rack servers
 - Right-sized dense server solutions
 - Based on Industry Standards
- Leading performance and efficiency
 - Concentrated compute power
 - Shared high efficiency power and cooling components
 - Lower your operating costs
- Flexible and serviceable solutions
 - Modular design allows tailoring
 - Serviceability and storage capable designs
 - Works in existing data center infrastructure



Purpose-built HPC Storage



Complementary Scalable Storage Solutions for High Performance Computing

X9000 Network Storage System

- Scalable performance and capacity
 - Scalable aggregate bandwidth
 - Scalable metadata, ideal for small files
- Shared datacenter multipurpose storage
 - Linux and Windows clients
 - NFS & CIFS support
- Ideal for applications in media, FSI, bioinformatics, web/cloud

DDN Storage with Lustre

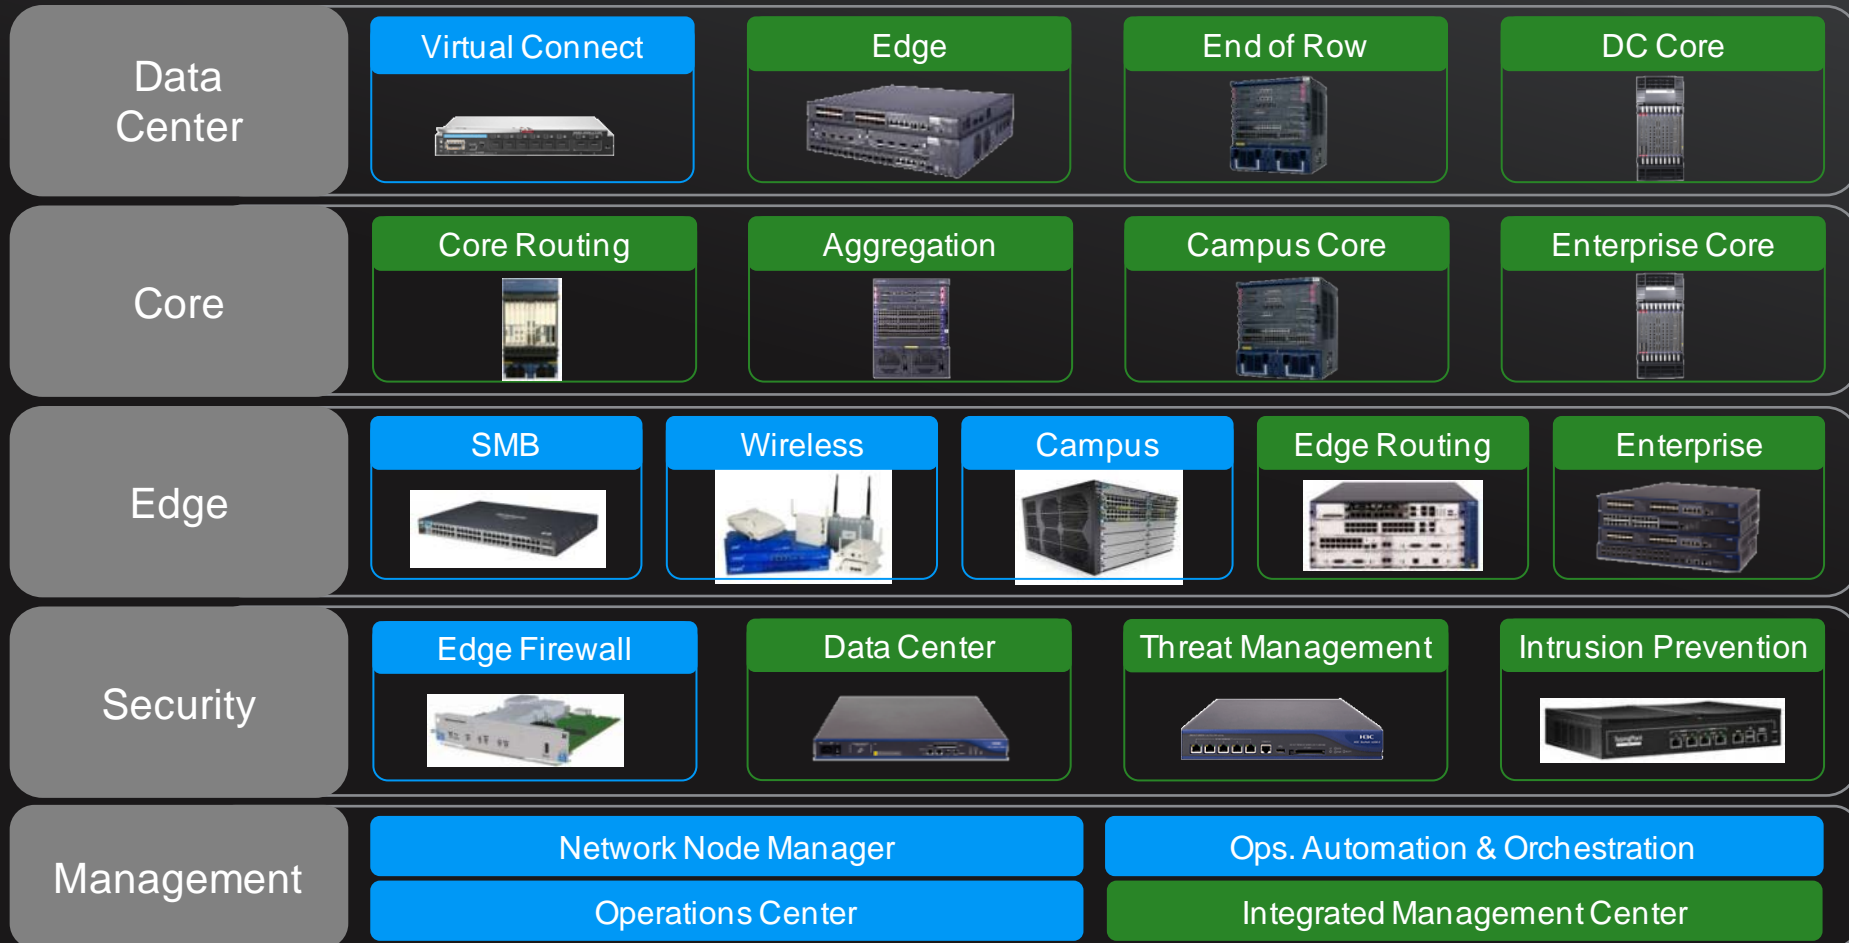
- Scalable performance and capacity
 - Scalable single-file bandwidth, with multiple writers
 - demanding bandwidth requirements
- Tightly coupled to HPC Linux clusters
- Ideal for parallel applications in traditional HPC



Purpose-built HPC Fabrics



HP + 3Com – Leadership from Edge to Data Center Core



HPC Software Infrastructure



Unified Cluster Portfolio

HPC Technical and Enterprise

HPC application, development and cloud software portfolio

Advanced and specialty options

(Accelerators, Visualization, other)

Scalable data management

(HP x9000 NSS, Lustre Cluster FS)

Cluster management layer

HP CMU

Partner and Open
Source choice

Microsoft Windows
HPC Server 2008

Operating environment and OS extensions

Linux

Windows

HP cluster platforms

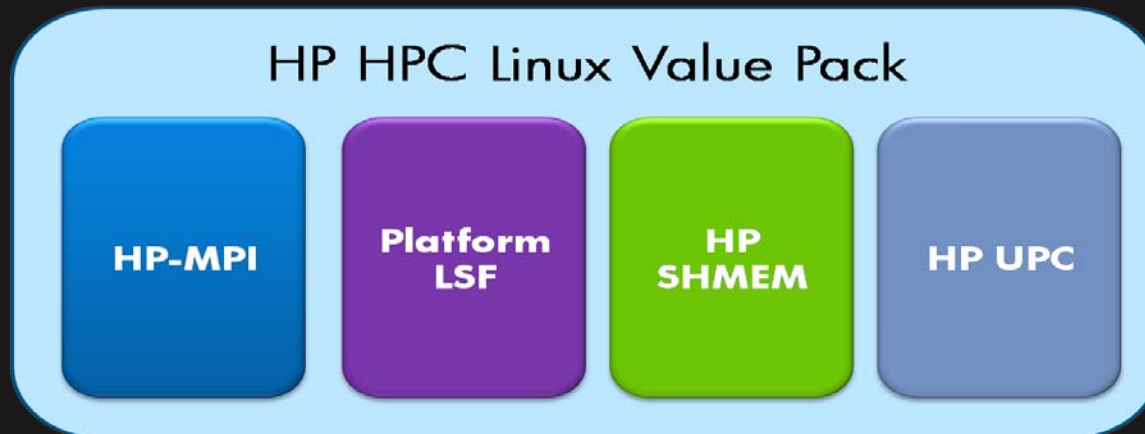
HP ProLiant servers, HP BladeSystem, multiple interconnects

HP Datacenter Products & Services



A la Carte cluster options for HP Clusters

- Operating systems: RHEL, SLES, or customer-supported community distributions; Microsoft Windows HPC Server 2008
- Cluster Management: HP CMU, or third party, via SLMS or customer installed (e.g., ROCKS, Platform Cluster Manager)
- MPI: HP-MPI, or third party/open source; Windows MPI
- Workload manager: Platform LSF (via SLMS now), Altair PBS Pro (HP SKU), Adaptive Computing Moab (via SLMS)



Datacenters – PODs



Trend? Are Next Generation Data Centers Ugly?



HP POD products and concepts

- 22 50U racks 40ft
- 600kW power capacity
- Designed for high density deployments – max 34kW per rack
- Flexible for redundant or non-redundant deployments



- 10 50U racks 20ft
- Modular design for better supply chain efficiency
- Flexibility to customize



- Rugged exterior
- EMI shielding
- Designed for portability



Peta-scale Implementation Example: TITECH Tsubame 2.0



TSUBAME 2.0 Overview

- Compute nodes: 2.4PFlops (CPU+GPU)
 - **New SL-node** >>1408)thin nodes, each with 2 Westmere-EP and 3 NVIDIA M2050
 - 1347 with 54GB and SSD 60GB, 41 with 96GB and SSD 120GB
 - **Suse Linux Enterprise Server or Windows HPC Server**
 - **DL580 G7** Medium (24) and Fat (10) nodes, with 2 NVIDIA S1070
 - Medium: 128GB plus SSD 120GB x4
 - Fat: 256GB plus SSD 120GB x4
- QDR InfiniBand, full bisection, non-blocking
 - Spine: **Voltaire Grid Director 4700** 12 x 324port
 - Edge: **Voltaire Grid Director 4036** 179 x 36 port and **4036E** 6 x 34port/10GbE 2 port
- Storage: 5.93PB
 - Lustre file system 5.93PB: **DDN SFA 10000** (10/rack, 5 racks) and DL360 G6 (30)
 - Home file system: 1.2PB: **DDN SFA 10000** (10/rack, 1 racks), BlueArc Mercury 100 (2) and DL360 G6 (30)
- Press release (Japanese):
 - <http://www.gsic.titech.ac.jp/sites/default/files/pdf/TSUBAME/press.pdf>



FOR MORE INFORMATION

– Our web sites

- HPC: www.hp.com/go/hpc
- HP BladeSystem: www.hp.com/go/bladeSystem
- x9000 NSS: www.hp.com/go/x9000
- POD: www.hp.com/go/pod
- MOAB ACS: www.hp.com/go/max

– Our booth at SC09

- Purpose-built solutions for scale: new SL series, BL2x220c, x9000 , 10GbE/IB networking
- Holistic datacenter portfolio: POD
- Adaptive and scalable software: new Moab ACS, CMU, HP-MPI

Thank You

