HP/NVIDIA Solutions for HPC Compute and Visualization Performance

Ed Turkel
Group Manager, HPC Business Development
Apollo, HPC and Big Data Business, HP Servers

The most exciting shifts of our time are underway

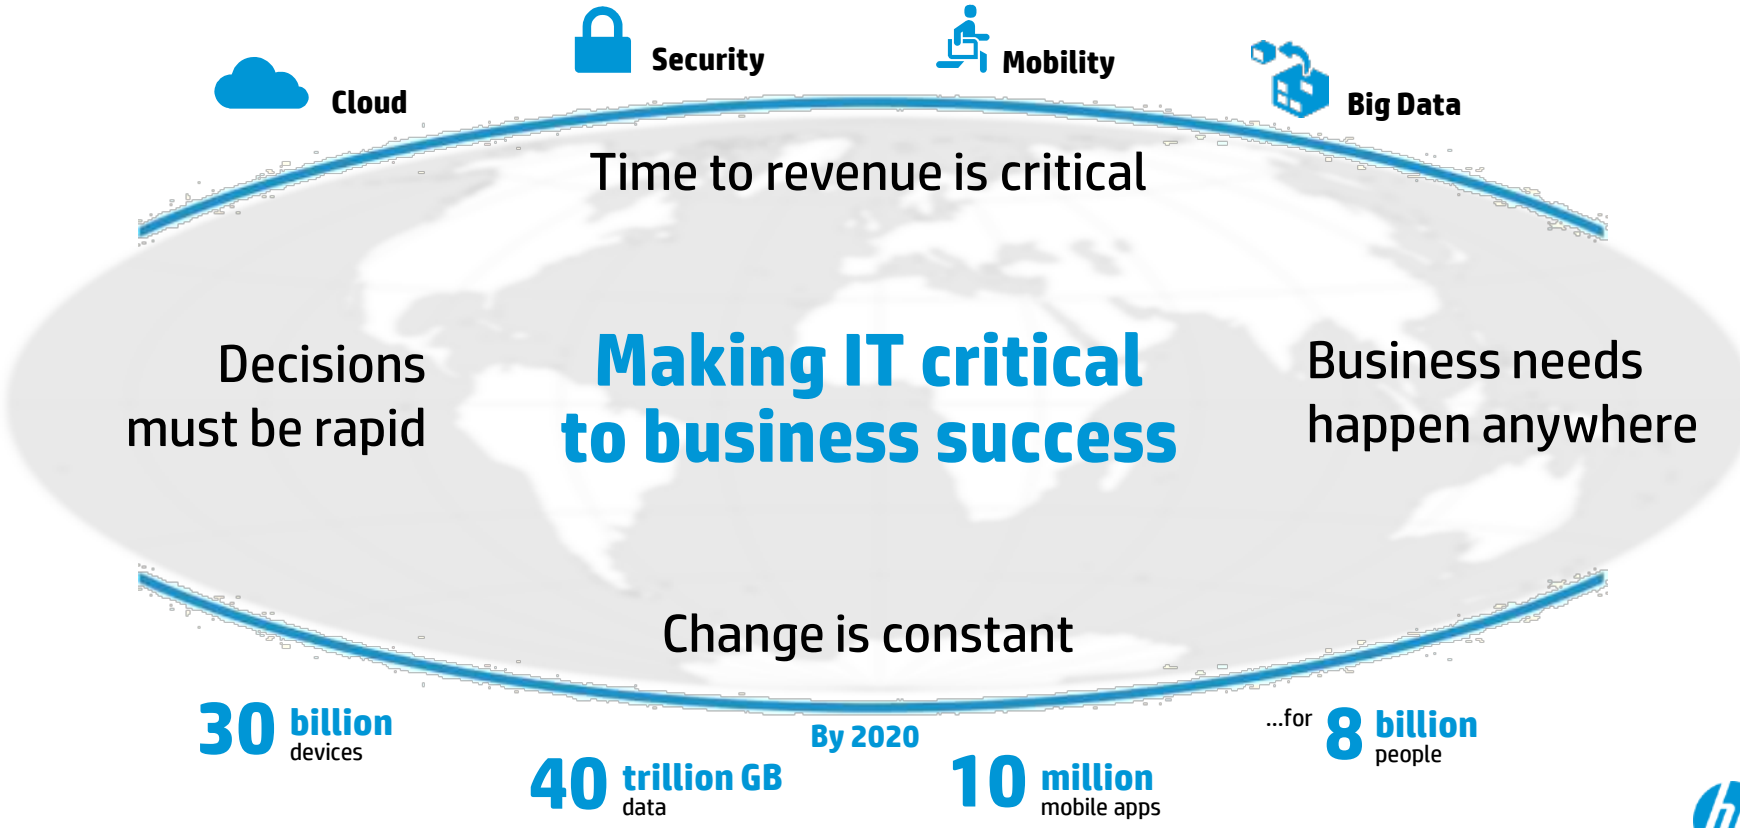

HP's compute portfolio for better IT service delivery

Software-defined and cloud-ready

HP OneView

HP Helion

OpenStack

RESTful APIs

Workload-optimized

Mission-critical environments

HP ProLiant scale-up

HP Integrity blades & Superdome

HP Integrity NonStop

Availability
for continuous business

Virtualized &cloud workloads

HP BladeSystem

Convergence
to accelerate IT service delivery

Core business applications

HP MicroServer

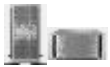

HP ProLiant ML

HP ProLiant DL

Intelligence
to increase productivity

Big Data, HPC, & web scalability

HP ProLiant SL

HP Moonshot

HP Apollo Family

Density and efficiency
to scale rapidly

SP Workloads

HP Cloudline

Lowest Cost
built to Scale

Converged

Converged network
HP Networking

Converged management
HP OneView

Converged storage
HP StoreVirtual VSA

Common modular architecture

Global support and services | Best-in-class partnerships | ConvergedSystem

Hyperscale compute grows 10X faster than the total market

HPC is a strategic growth market for HP

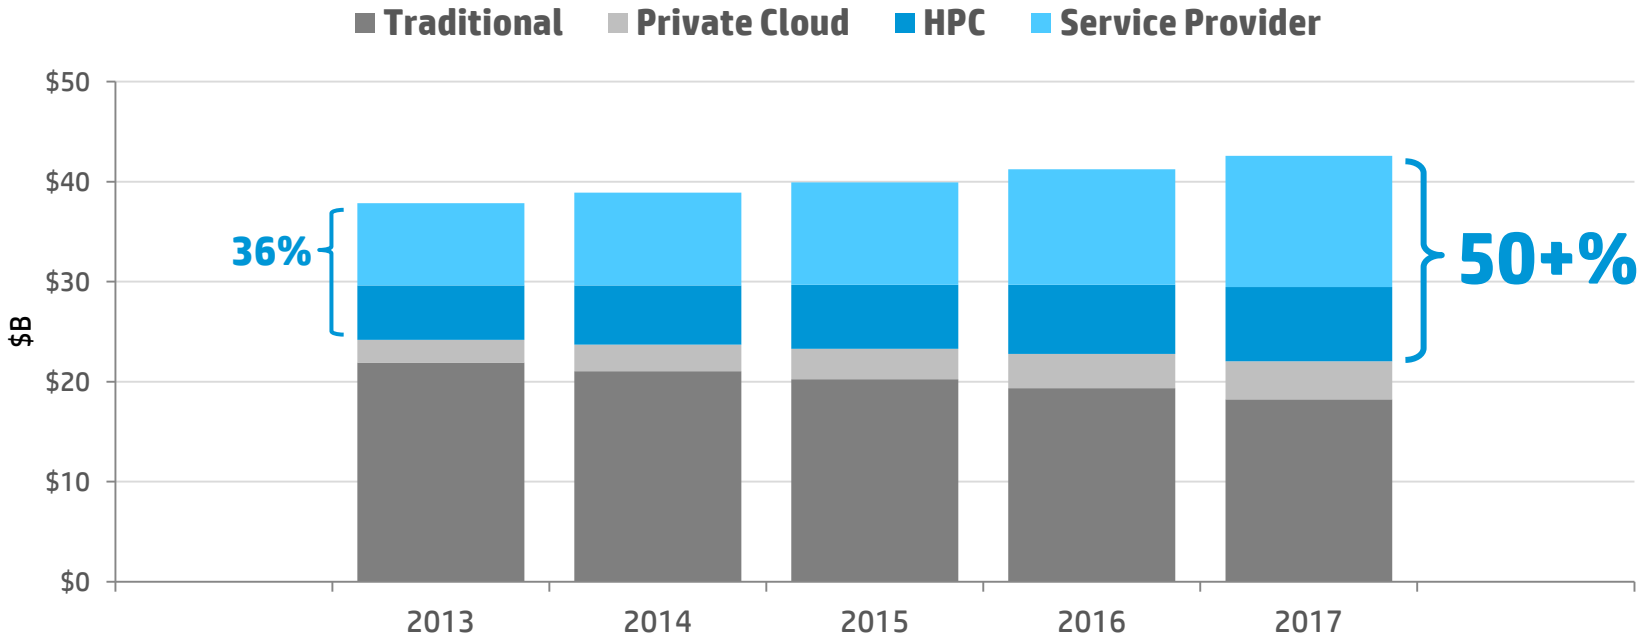

Source: IDC and HP analysis

HP is the HPC market leader

HP in the TOP500 list

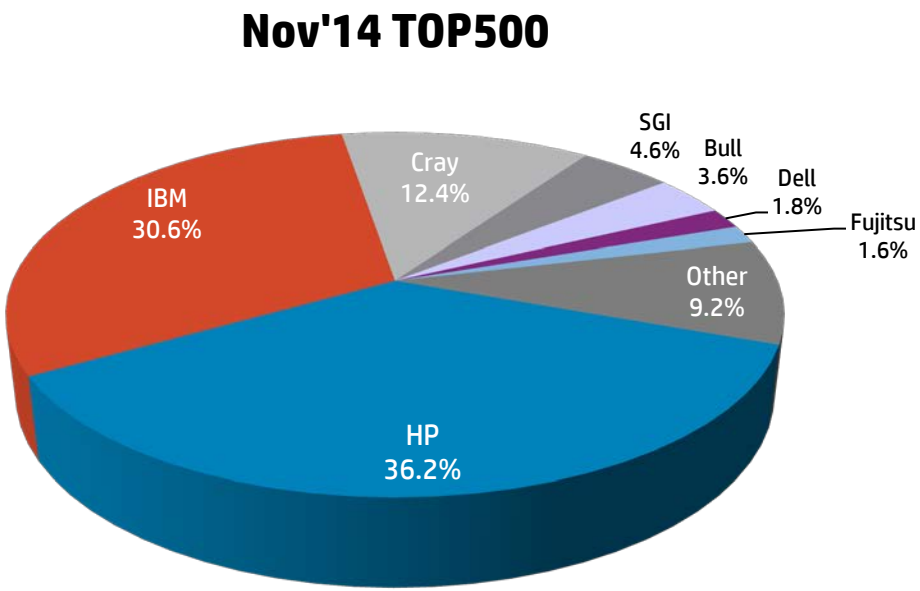

Top100 Systems

15. Tokyo Tech Tsubame 2.5
89. Clemson Palmetto2
94. USC HPCC

Source: www.top500.org

Solving global problems requires greater…

Genetics

Digital Content

Medical research

Weather modeling

Geophysics

Manufacturing/ Engineering

Finance

Government, Academia, Enterprises

Performance

Efficiency

Accessibility

Delivering a complete HPC solution

Servers

HP Apollo Systems,
HP BladeSystem,
HP ProLiant Gen9 and
HP Integrity Superdome X
Servers

Storage

HP SL4500 System
with Object-Storage
and HPC Storage
Solutions,
3rd party Lustre Solutions

Accelerators

NVIDIA Grid and Tesla GPUs,
HP Accelerator-enabled Servers

Remote Desktops

NVIDIA Grid GPUs,
HP GPU-enabled Servers,
HP Workstation Blade Servers,
HP Remote Graphics and
3rd party S/W

Network

Intel and Mellanox
InfiniBand,
Low-Latency Ethernet

Power & Cooling

HP Modular Cooling System,
HP Performance Optimized DataCenter,
HP Apollo 8000 System

Management

HP Insight CMU,
Insight CMU Connector Partners,
HP OpenView

Cloud

HP Helion Self-Service
HPC Solution,
OpenStack

Services

HP Services for HPC

HP ProLiant Gen9 Servers with Integrated NVIDIA GPUS

Accelerating Performance

New NVIDIA accelerators designed for HPC

Customer Benefits

- Fully integrated accelerators, installed, tested and supported in HP Apollo solutions
- Accelerate application performance, choosing the right accelerator for each application
- Enable maximum performance/ft² and performance/watt in HP Apollo solutions

Key features

- Support for NVIDIA Grid K1/K2 GPUs
- Support for NVIDIA Tesla K40 and New K80 GPUs
- **Unique** to HP: Support for NEW NVIDIA Tesla K40d GPUs

4x teraflops per square foot

Industry's most complete portfolio for HPC

Workload optimized, engineered for any demand

ProLiant BL family

Cloud-ready Converged Infrastructure

ProLiant DL family

Versatile, rack-optimized servers

ProLiant SL family

Purpose-built density-optimized servers

HP Moonshot

The world's first software defined server

HP Apollo

Optimized rack-scale computing for HPC

HP ProLiant WS460c Gen9 Server Blade

Broadest range of high performance, high density professional graphics

Proven and qualified Gen8 graphics cards support with new additions post-launch

Up to 70% performance increase with the new Intel® Xeon® E5-2600 v3 processors

New flexible embedded storage controller options for the ideal balance of performance and price

HP ProLiant WS460c Gen9 Graphics Server Blade

Two form factors with MXM & fullsize PCIe graphics

NVIDIA Quadro K3100M Graphics

WS460c Gen9 Graphics Server Blade

16 per 10U enclosure

HP MultiGPU Carrier with 6X K3100M

NVIDIA Quadro K6000/K5000/K4000, GRID K2/K1 GPU

WS460c Gen9 Graphics Server Blade with Expansion

8 per 10U enclosure

HP ProLiant Gen9 Operating & Virtualization Support

Microsoft Windows

Windows Server 2012 R2
Windows 7 Pro (64-bit), Enterprise (64-bit)
Windows 8.1 Pro (64-bit)

Linux

Red Hat Enterprise Linux Desktop 6.5 or later)

Virtualization

VMware Horizon View 6, vSphere 5.5 or later
Citrix XenServer 6.5 or later, XenDesktop 7

Content Subject to change – please check the OS Support Matrix for the latest www.hp.com/go/ossupport

HP ProLiant WS460c Graphics Server Blade

Business efficient and high-performance end-user computing delivered from the datacenter for superior user experience

HP MultiGPU Carrier

GPU accelerated end-user computing	Right sizing GPU performance per user type	Operational efficiency
▪ NVIDIA professional graphics ▪ MXM and fullsize PCIe graphics ▪ Client OS on bare-metal	▪ NVIDIA GRID vGPU HW virtualization ▪ Entry to Ultra-high GPU class ▪ Software configurable performance	▪ High host & GPU density per rack ▪ Simplified via Convergence ▪ Power, Cooling and Cost efficiency

HP innovation for a new style of IT

Extending modularity and workload optimized product developments

New!

Apollo 2000

Scalable Multi-node

SL4500

Storage Density

SL6500

Compute Intensive

Apollo 6000

Rack-scale Efficiency

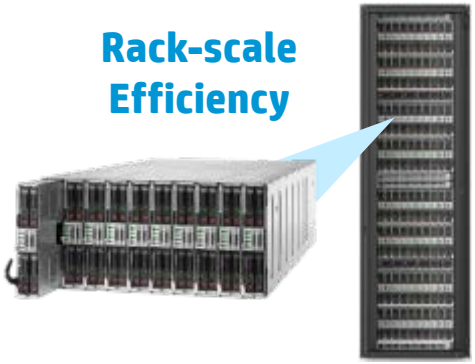

Apollo 8000

Performance Density

Delivering breakthrough efficiency at scale

New!
HP Apollo 2000 Servers

HP Apollo 2000 System

The Enterprise Bridge to Scale-Out Architecture

NEW

Density optimized for traditional datacenters

- Up to **4 powerful servers in 2U chassis** – 2X the density of 1U servers
- **Traditional racks and cabling** for existing datacenters
- **Cost effective** in any configuration

Configuration flexibility for a variety of workloads

- **Mix and Match servers** for workload optimization
- **HPC performance** with accelerators, top bin CPUs, fast HPC clustering
- **Storage flexibility** and a **broad range of I/O options** for workload optimization

Simple at scale – it's ProLiant Gen9

- **ProLiant** enterprise-class management and operational tools
- **HP iLO Management** saves administration time and cost
- **HP Advanced Power Manager** (APM) enables more efficient capacity per rack
- **HP Insight CMU** to monitor, manage and optimize compute clusters of any size

2X Density **7.9 TeraFLOPS** in 2U **Flexible Familiar Fast**

Apollo 2000 System offerings

Up to 4 servers in 2U, traditional rear cabling and front hot plug storage

Apollo r2200

Apollo r2600

Apollo r2800

ProLiant XL170r

ProLiant XL190r

Four powerful hot-pluggable servers in 2U

HP ProLiant XL170r - Gen 9 1U Node

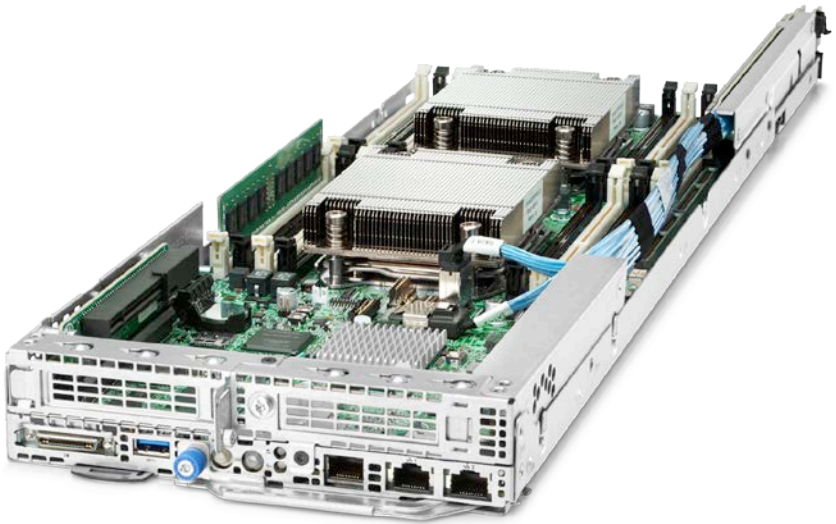

30-40% more performance
from dual Intel® Xeon® E5-2600 v3 processors , Top Bin Support

High performance memory
with HP SmartMemory at speeds of up to 2133 MHz and 512GB maximum

Individual node serviceability
with 4 independent serviceable servers

Flexible I/O
Two I/O expansion slots with flexibleLOM support

Expandable 2U Compute Trays

HP ProLiant Apollo XL190r - Gen 9 2U Node

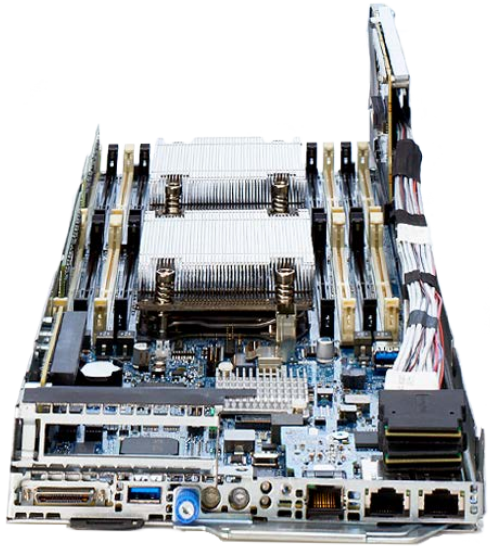

Accelerated application performance

- Support for up to 2 **NVIDIA Tesla K40 GPUs** per 2U node

Increased I/O capacity

- up to four PCIe 3.0 slots with FlexibleLOM support
- Multiple I/O options for workload customization

Endless Workload Possibilities

HPC

- Top bin processors
- **NVIDIA Tesla GPUs**
- I/O Expansion
- Storage Flexibility

Service providers

- Flexible configurations
- Low cost options
- Shared infrastructure efficiency
- Density
- Solutions Approach

SMB / Enterprise / Remote Site

- Flexible configurations
- Singe node service domain
- Redundant power and fans
- Traditional rear cabling and front hot plug drives

HP Apollo 6000 System

HP Apollo 6000 System

Rackscale performance and efficiency for High Performance Computing

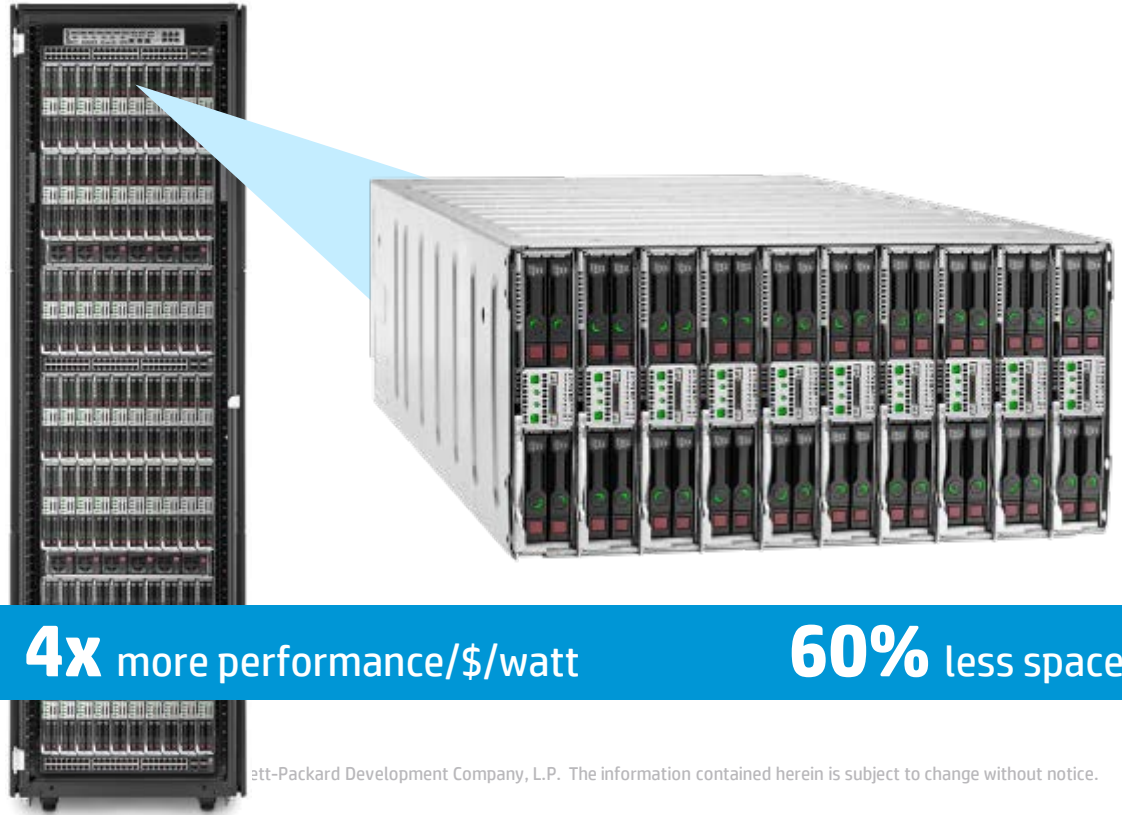

Leading performance

Get up to 4x more performance per $ per watt using 60% less rack space in 5U

Rack scale efficiency

External power shelf and HP Advance Power Manager helps maximize energy efficiency

Lower TCO for right workload

Various server trays and networking to fit workload needs while increasing cost savings

4x more performance/$/watt

60% less space

$3M TCO savings

The Apollo 6000 System

Density

Efficiency

HP Apollo 6000 Power Shelf

Serviceability

HP Apollo a6000 Chassis

Manageability

HP APM

HP ProLiant XL220a
2x1P servers

HP ProLiant XL230a
2P server

Flexibility

HP ProLiant XL250a
2P + 2 NVIDIA Tesla GPUs

HP Apollo a6000 Chassis

Standard 1m rack, rear cabled

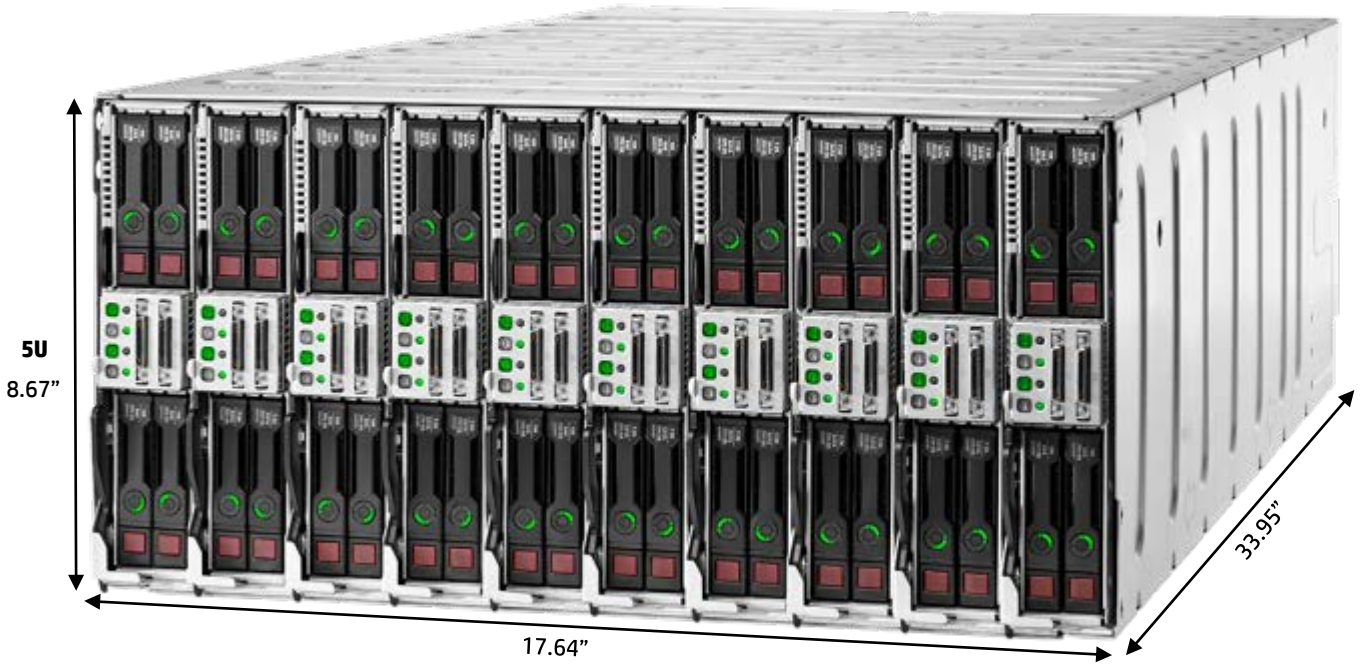

Features

- 1 slot and 2 slot tray support
  - -10 single slot trays
  - -5 double slot trays
- Mix-n-match trays
- Shared cooling
- 12V DC power distribution
- 5U tall infrastructure
- Up to 5700W per chassis

Serviceability

- Front serviceable trays
- Standard rear cabling
- Front serviceable hot plug drives
- Redundant, hot plug fans

5U (H) x 44.81cm (W) x 86.23cm (D)
5U (H) x 17.64 in (W) x 33.95 in (D)
Max Airflow < 330 CFM per chassis

HP Apollo a6000 Chassis

Easy Serviceability Support

HP Innovation Zone

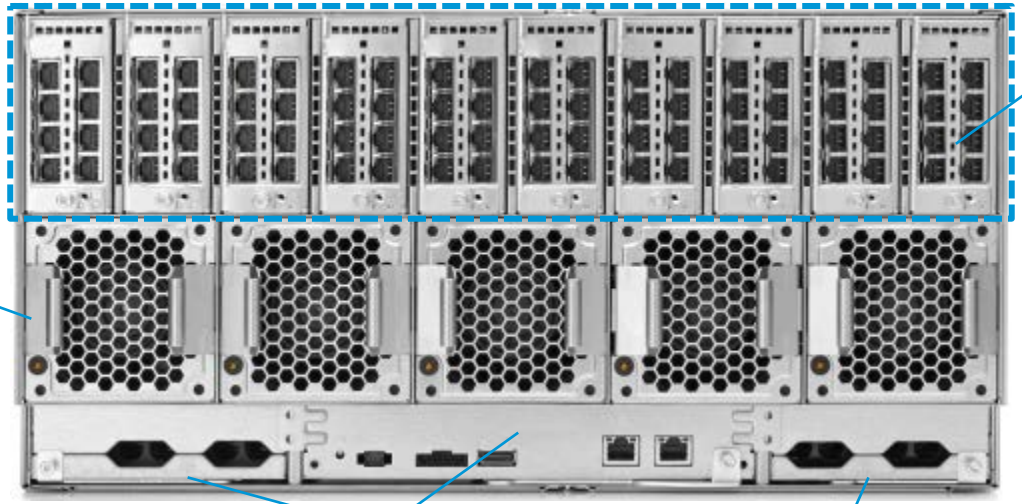

(5) Fans

-Hot pluggable

-80mm Redundant

-Dual rotor fans

Network I/O Modules

Two 1GbE 1-port module

OR

Two FlexibleLOM riser:
supports 1GbE, 10GbE or IB

OR

June ’15:

One standard LP PCIe (x16)
& one FlexibleLOM card (x8)

Management Module

Power Shelf HP APM Aggregated iLO

(2) DC Power Cable Cages per chassis

Supports up to 4 x 12V DC cables per chassis

HP Apollo 6000 Power Shelf

Pooled Power Efficiency

Front View

Back View

Efficiency

- External pooled power shelf
- Fits up to 6 power supplies
- 2400W or 2650W power supplies
- Up to 15.9kW non-redundant
- Single or 3-phased AC input
- Up to twelve 12V DC cables

1.5U (H) x 44.81cm (W) x 78.44cm (D)
1.5U (H) x 17.64 in (W) x 30.88 in (D)

HP Apollo 6000 Infrastructure

Rack Level Power Efficiency

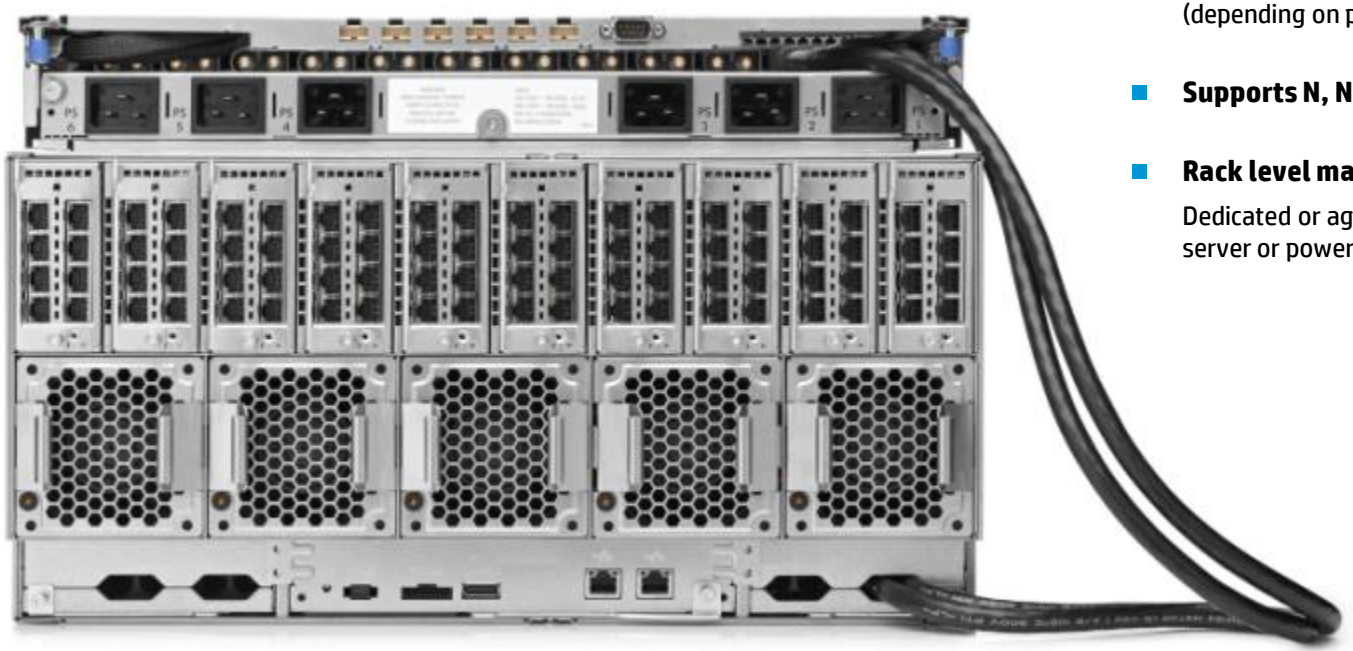

Serviceability meets Efficiency

- **Each power shelf can support 3 to 6 chassis**
  (depending on power and redundancy configuration)
- **Supports N, N+1, and N+N Redundancy**
- **Rack level management with HP APM module**
  Dedicated or aggregated iLO to manage at chassis, server or power level

HP ProLiant XL250a Gen9 Server

2P Compute with Accelerators

HP ProLiant XL250a Server (with Accelerators)	
Rack	Ideal fit for standard 1.0m rack depth
Chassis	HP Apollo a6000 Chassis (max 5 trays/servers per chassis)
Power	HP Apollo 6000 Power Shelf
Processor	Intel Xeon E5-2600 v3 family
Memory	16 DDR4 DIMMs (Registered or Load Reduced) 2133 MHz, 512GB max (16 x 32GB)
Network	-Two 1GbE single-port module[1] -Two IO Slot module: InfiniBand, Ethernet (1GbE or 10GbE), FC • 1 x16 PCIe Gen3 (HH LP PCIe Card or FlexibleLOM) • 1 x8 PCIe Gen3 FlexibleLOM
IO Slots	1 x8 PCIe Low Profile Gen3 for Smart Array or HBA 2 x16 PCIe Gen3 for FH/FL Accelerators
Storage	Up to 6 SFF hot plug SAS/SATA/SSD - front accessible storage
Accelerator	**NVIDIA Tesla: K40, K80**
Management	HP iLO 4 *(Each server will have its own dedicated iLO)* HP Advance Power Manager *(rack level mgmt)*

HP Apollo 8000 System

The New HP Apollo 8000 System

Advancing the science of supercomputing

Leading teraflops per rack for accelerated results

- **4X** teraflops/sq. ft. than air-cooled systems
- > **250 teraflops/rack**

Efficient liquid cooling without the risk

- **40% more FLOPS/watt** and **28% less** energy than air-cooled systems
- **Dry-disconnect** servers, intelligent Cooling Distribution Unit (iCDU) monitoring and isolation

Redefining data center energy recycling

- Save up to **3,800 tons** of CO2/year (790 cars)
- **Recycle water** to heat facility

4X teraflops/sq. ft. **40%** more FLOPS/watt **3,800** tons of CO2

HP Apollo 8000 System

Leading performance density

HP Apollo f8000 Rack

Dry disconnect server trays

HP ProLiant XL730f
2x2P Servers

HP ProLiant XL750f
2P+2 NVIDIA Tesla GPUs

HP InfiniBand Switch for
Apollo 8000 System

HP Apollo 8000 Cooling Circuit

Efficient liquid cooling without the risk

HP Apollo 8000 iCDU Rack

Apollo 8000 System Technologies

Advancing the science of supercomputing

Intelligent Cooling Distribution Unit
- 320 KW power capacity
- Integrated controls with active-active failover

Dry disconnect servers
- 100% water cooled components
- Designed for serviceability

Management infrastructure
- HP iLO4, IPMI 2.0 and DCMI 1.0
- HP Apollo 8000 System Manager

Warm water
- Closed secondary loop in CDU
- Isolated and open facility loop

Power infrastructure
- Up to 80kW per rack
- Four 30A 3-phase 380-480VAC

4 f8000 + 2 iCDU + under-floor Plumbing Kit

Server Tray alternatives

Two Dual Socket Server Tray

Features

- 2 Intel Xeon E5-2600v3 processors
- 1 SFF SSD per server
- 1 IB FDR/10GbE adaptor per server
- ~750 Watts power consumption

Dual Socket Dual GPU Server Tray

Features

- 2 Intel Xeon E5-2600v3 processors
- 2 NVIDIA Tesla K40 XL (K40d) GPUS
- 1 SFF SSD
- 1 IB FDR/10GbE adaptor
- ~1000 Watts power consumption

HP f8000 Power Block Diagram

High-Voltage AC to the Rack: Limiting conversion steps to improve efficiency

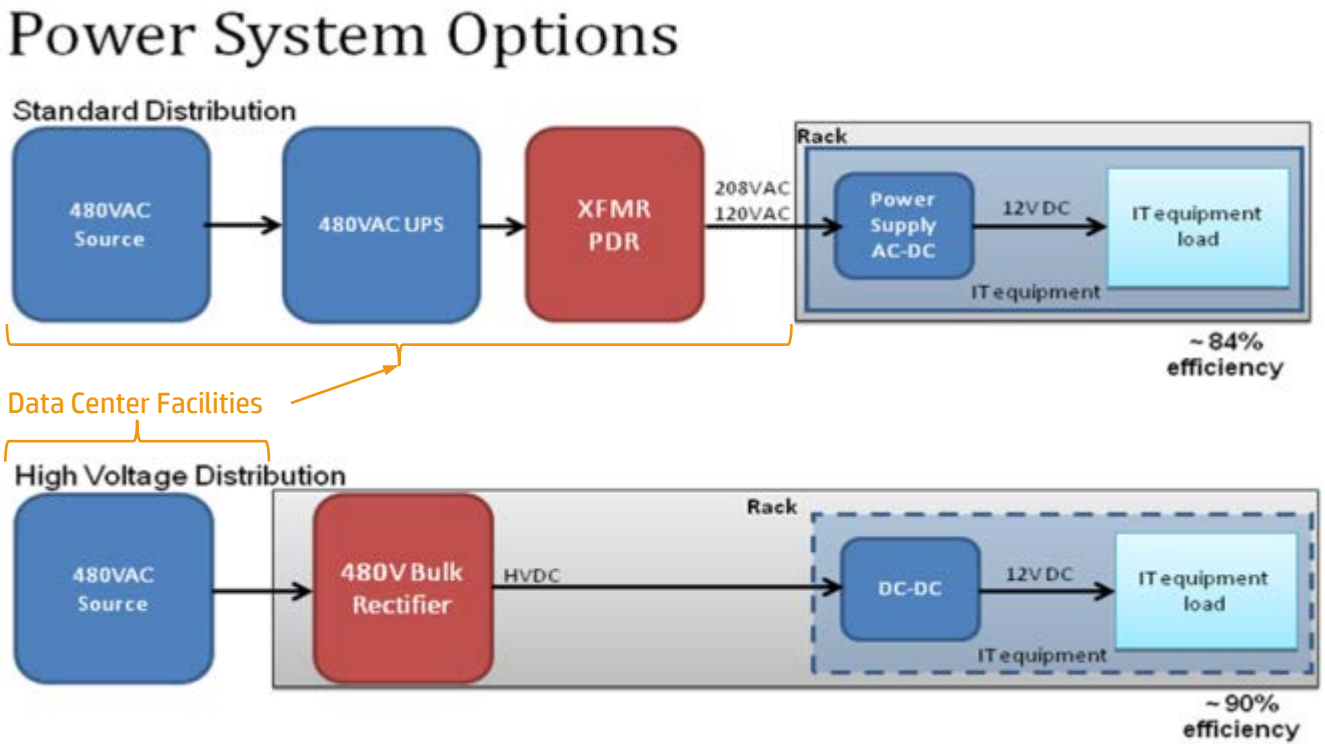

Energy Reuse

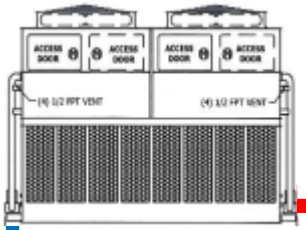

Evaporative Cooling Tower

Heat Buildings

Energy Reuse Loop

Warm Water Cooled Servers

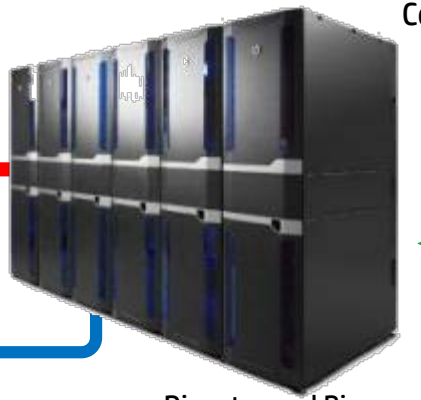

Digester and Biogas Handling

Bio Gas Generator Output

Melt Snow on sidewalks

HP Apollo Systems Management

Built-in intelligence that maximizes every hour, watt and dollar

Rack Level Environmental

HP Advanced Power Manager

- See and manage shared infrastructure.
- Flex to meet workload demands

Simplify, and save >80%

Solution Level

HP Insight Cluster Management Utility (CMU)

- Hyperscale lifecycle management software
- Fast and scalable cloning

Easy, friction-less control of remote servers

On System

HP iLO Management Engine

Intelligent management on every HP server enabling health and alerting, firmware maintenance and support automation

5x faster problem analysis

HPC Services

HP Services for HPC solutions

Data center planning and design services to achieve business outcomes

- Data center facilities
- Workload migration
- Big data, mobility, virtualization design, planning and implementation

Services to speed startup and build capabilities with new technology

- Factory Express
- Onsite installation
- HP Education

Ongoing support for business continuity

- Datacenter Care supports your environment
- Proactive Care helps prevent problems
- Foundation Care helps solve problems faster

Flexible payment plan and terms

- Available globally where HP Financial Services conducts business[1]
- Technology refresh approach to allow for future scalability and upgrades

[1]Financing and service offerings available through Hewlett-Packard Financial Services Company and its subsidiaries and affiliates (collectively HPFSC) in certain countries and is subject to credit approval and execution of standard HPFSC documentation. Rates and terms are based on customer's credit rating, offering types, services and/or equipment type and options. Not all customers may qualify. Not all services or offers are available in all countries. Other restrictions may apply. HPFSC reserves the right to change or cancel this program at any time without notice.

Reinventing HPC today to accelerate the world of tomorrow

Accelerating performance
to speed up answers

4x teraflops per square foot

Maximizing efficiency
for sustainability and savings

4x density per rack per dollar

Unleashing HPC
to enterprises of any size

Years to days
for new innovations

HP Apollo family
Optimizing rack-scale computing for HPC

Thank you

S5825 - HP/NVIDIA Solutions for HPC Compute and Visualization Performance (Presented by HP)

Ed Turkel *Group Manager, HPC Business Development, HP Servers, HP*

Ed manages the worldwide product marketing team for the High Performance Computing (HPC) business at Hewlett Packard. The HPC business delivers integrated solutions for HPC with maximum performance and efficiency, enabling innovative research, engineering and analytics. Ed's team is responsible for developing HP's solutions and go-to-market strategy for HPC, working closely with HP's customers to develop the solutions that enable them to best achieve their business and research outcomes. Ed has almost 35 years experience in HPC, including 30 years with HP, in various technical, marketing and business roles.

High Performance Computing is characterized by user demand for increasing levels of computational performance, combined with exploding volumes of data, to accomplish their science, engineering, or analytics workloads. Demands for performance growth are becoming increasingly limited by the power, space and cost of deployment of new systems, while exploding data volumes challenge traditional client/server computing models. For years, HP has partnered with NVIDIA to develop HPC solutions that are purpose-built for compute and visualization performance and scalability, while delivering innovative energy and space efficiency, with a focus on customer ROI. This session will showcase HP and NVIDIA's latest technologies and solutions in use today by leaders in the HPC community, plus trends for the future.