



Accelerating HPC and AI innovation for today and tomorrow

Mike Vildibill

VP, Advanced Technologies & Exascale Development

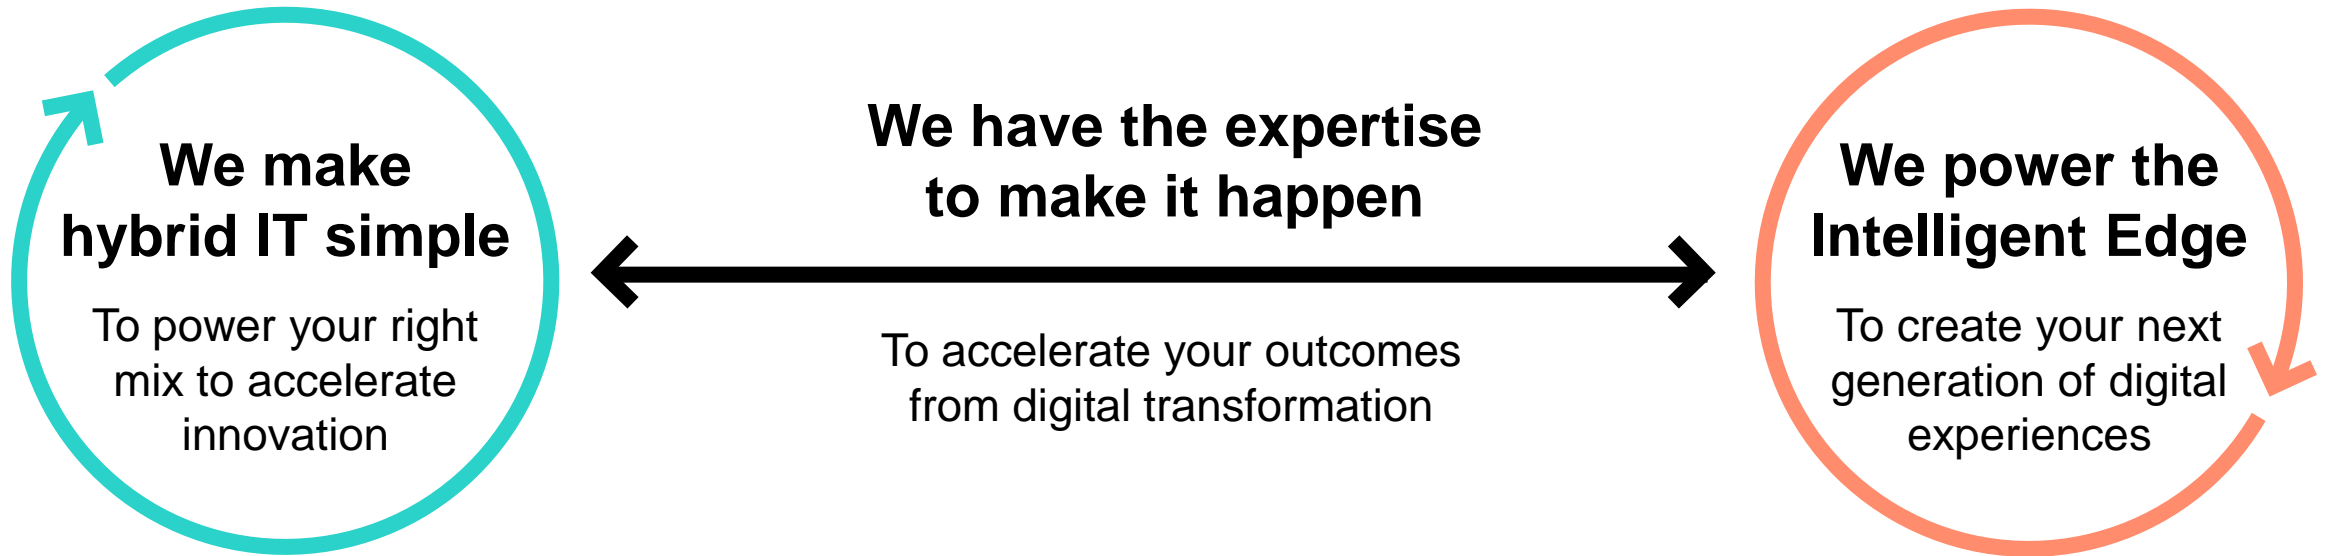
Steve Tolnai

HPC Asia Pacific & Japan

19 October 2017



Hewlett Packard Enterprise - Who we are



SGI and HPE integration update



Overall Status

- **Careful approach** to integrating the companies **without business disruption**
- All engineering, product management, marketing, manufacturing and back office functions have been integrated
- Legal merge in Japan on **1 August 2017**
- Full integration will **complete on 31 October 2017**




Technology and portfolio

- **Product line portfolio consolidated** with fully integrated roadmap
- **“Best of the best”** technologies of platforms and software for HPC and AI
- HPE SGI 8600, UV and Enhanced Hypercube Network Technology Integration



Services and Support

- **Comprehensive global services & support** for HPC and AI under Pointnext Services
- **SGI services teams will integrate on 1 August** and remain intact under Pointnext
- Commitment to continued exceptional customer experience and satisfaction



Sales and Account Coverage

- SGI sales teams will fully integrate on **31 October 2017**
- HPC focus and account coverage will **deepen and grow**
- SGI and HPE sales teams are **already working together**



HPC and AI context, HPE strategy and portfolio

HPC, big data analytics and AI empower the data-driven enterprise

Simulation and modeling for science, research and business



Deep Learning and Cognitive computing

Pattern recognition
and Artificial
Intelligence



Analysis
of high volume
machine, human
and business data

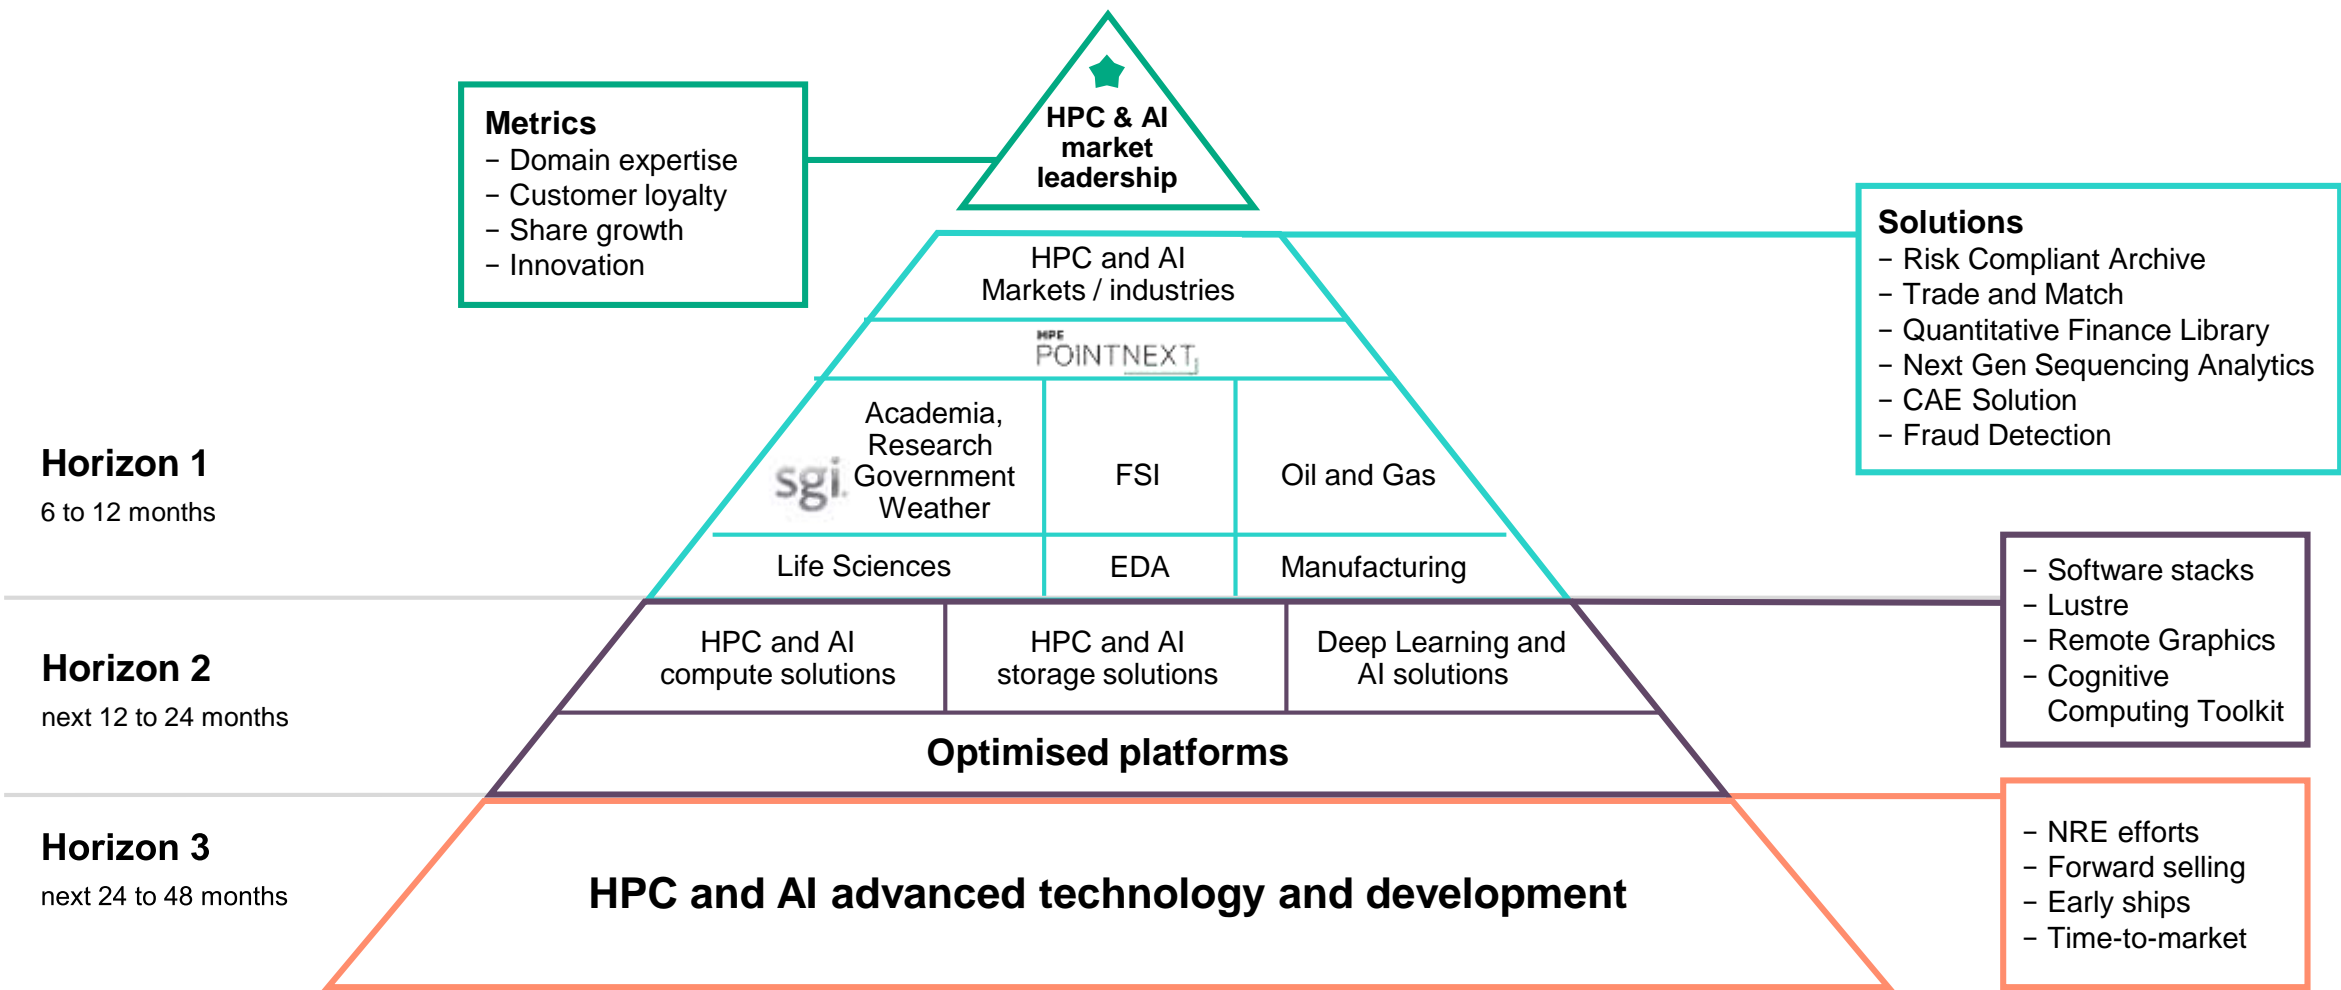


Managing data

unstructured data;
images, audio and
video, backup and
recovery



HPE Strategy: Accelerate HPC and AI leadership today and into the future



HPE is a strong leader in the HPC Market

TOP500: 49th edition – vendor trends from 2012 to 2017



HPE delivering Green HPC for Academic Research

Tokyo Institute of Technology

Tsubame 3.0, a modified HPE ICE XA System at the
GSIC Center, Tokyo Institute of Technology



東京工業大学
Tokyo Institute of Technology



#1 on Green500,
June 2017

HPE purpose-built portfolio for HPC

HPC Industry Solutions



Financial Services



Academia,
Research, Gov't



Life Sciences,
Health



EDA / CAE
Manufacturing



Oil and Gas,
Energy



Weather and
Climate Research



Advisory, Professional and Operational Services – HPE Flexible Capacity for HPC, HPE Datacenter Care for Hyperscale

Supercomputing / Enterprise / Commercial HPC

HPE SGI 8600



Liquid cooled, delivering industry leading performance, density and efficiency

HPE Apollo 6000 Gen10



Extreme Compute Performance in High Density

HPE Apollo 6000 Gen9




Rack-scale HPC

HPE Apollo 2000 Gen9



The bridge to enterprise scale-out architecture

HPE Performance Software Suite



- HPE Performance Software - Core Stack
- HPE Insight Cluster Management Utility
- HPE SGI Management Suite
- HPE Performance Software – Message Passing Interface

Emerging HPC

HPE Apollo 6500 Gen9



Rack-scale GPU Computing

In-memory HPC

HPE Integrity MC990 X



Scale-up, shared memory HPC, UV Technologies

HPE Integrity Superdome X



HPC Storage

HPE Apollo 4520



Additional Storage Options available

HPC Data Management Framework Software

Large-scale, storage virtualization & tiered data management platform

Choice of Fabrics



Arista Networking

- Intel® Omni-Path Architecture
- Mellanox InfiniBand
- HPE FlexFabric Network

HPE has a comprehensive, purpose-built portfolio for Deep Learning

Deep Learning solutions



Financial services



Government and academia



Life Sciences, Health



Autonomous vehicles/mfg.

HPE
POINTNEXT

Advisory, professional and operational services | HPE Flexible Capacity, HPE Datacenter Care for Hyperscale

Compute for Core data center Training model

HPE SGI 8600

Petaflop scale for Deep Learning and HPC



Liquid cooled, delivering industry leading performance, density and efficiency

HPE Apollo 6500

Scalable, automated real-time intelligence



Rack-scale GPU computing with up to 8 GPUs per compute node

Compute for Core data center Inference engine

HPE Apollo 2000

The bridge to enterprise scale-out architecture



High compute density, ease of use and simplicity

Edge analytics and Inference engine


HPE Edgeline EL4000 Converged Edge System

Unprecedented deep edge compute and high capacity storage, based on open standards



Right-sized and more portable servers on the "Intelligent Edge"

HPE Performance Software Suite



- HPE Performance Software - Core Stack
- HPE Insight Cluster Management Utility
- HPE SGI Management Suite
- HPE Performance Software – Message Passing Interface

HPC Storage

HPE Apollo 4520



HPC Data Management Framework Software

Large-scale, storage virtualization & tiered data management platform

Choice of Fabrics



Arista Networking

- Intel® Omni-Path Architecture
- Mellanox InfiniBand
- HPE FlexFabric Network

Simplifying HPC and AI adoption with packaged solutions

Industry Solutions across FSI, Life Sciences and Manufacturing

Risk Compliant Archive Solution



Simplify enterprise storage while adhering to regulatory compliance standards

Trade and Match Server Solution



Minimize system latency for high-frequency trading operations

HPC for Trader Workstation



Reliable, density optimized, high performance solution with greater throughput, low latency

HPE ANSYS CAE Solution



Turn-key, fully managed HPC clusters for CAE

HPE Fraud Detection Solution leveraging Compute with GPUs



Automated, real-time fraud detection Deep Learning solution

HPE Next Generation Sequencing Solution



Highly scalable, reliable solutions for genome analytics

HPE Quantitative Finance Library Solution



Enhance FSI application software performance by enabling accelerators

HPC and AI technology partner ecosystem

Technology and go-to-market innovation with key industry leaders



HPC Solutions Alliance



- HPE and Intel HPC Alliance since Jul 2015
- Centers of Excellence
- Code Modernization
- Intel Xeon Phi, OPA, software
- Product development, benchmarking, engineering



HPC Fabric collaboration



- HPE and Mellanox collaboration announced Jun
- InfiniBand and Gen-Z open Standards
- In-network computing
- HPC Fabric for HPE SGI 8600 and New HPE Apollo 6000 Gen10
- Joint Development with HPE Advanced Dev Team - Exascale



Deep Learning collaboration



- HPE and NVIDIA collaboration forged May 2017
- Centers of Excellence
- Deep Learning Institute
- Product development, benchmarking, engineering

Accelerating HPC and AI innovation for today and tomorrow

Introducing optimized solutions, a new Compute experience and advanced technologies for the future

**Workload optimized
for extreme performance**



**More performance, efficiency
and scale to accelerate
business innovation**

**Secure, agile, flexible
Compute experience**



**A better way to protect
data, manage infrastructure
and control economics**

**Exascale and advanced
technology programs**



**More expertise and
flexibility to plan future
technology roadmap**

Powered by the world's most secure industry-standard servers¹ for HPC and AI

Accelerating HPC and AI innovation for today and tomorrow

Workload optimised for extreme performance

New HPE SGI 8600

Next gen petaflop scale, liquid cooled supercomputer

- **Greater performance, scale and efficiency**



Secure, agile, flexible Compute experience

A new experience in IT security and protection

*World's Most Secure Servers¹
for HPC and AI – HPE Apollo
6000 Gen10*



Exascale and advanced technology programs

DoE PathForward Exascale Program


- New Exascale program to create reference designs
- Inspired by Memory-Driven Computing and Hewlett Packard Labs technologies



New HPE Apollo 6000 Gen10

Next gen air cooled, purpose built enterprise HPC solution

- **Best in class performance, rack scale efficiency**



New HPE Performance Software Suite: Out-of-the-box HPC stack, enhanced cluster system management and acceleration tools

New HPE Apollo 10 Series

- **Cost effective** platforms for AI and emerging applications



- **NEW** collaboration for AI application in precision medicine

New Services and Consumption Model

- New Advisory, Professional and Operational Services
- HPE Flexible Capacity for HPC



New disruptive technology based system architecture

- ARM processor based system
- Proof of concepts with select customers



A close-up photograph of a server rack. A hand is shown plugging a yellow Ethernet cable into one of the ports on a server unit. The server units are stacked vertically, and the rack has a perforated metal door. The background is slightly blurred, showing more of the server environment.

**New optimized solutions to enhance
performance, efficiency and scale**

Solution for the world's most complex supercomputing problems

HPE SGI 8600 System optimized for performance, scale and efficiency



Leading performance, density, scale, and efficiency coupled with robust system tools for quickest times to solutions

Leading performance

- ***Fastest Distributed Memory Systems on the planet¹*** for message passing with performance validated on SPECmpiM_2007 and SPECmpiL_2007 both peak and base results.
- ***Legacy of leading benchmark*** and real world application performance²

Ease of use

- **HPE SGI Management Suite:**
 - Provisioning of **thousands of nodes in minutes³**
 - From detailed system health monitoring to fine-grained power management
- ***Quick time to solution*** with off the shelf OS and applications

Density / Scale / Efficiency

- ***Scaling to >10,000 nodes without additional switches⁴*** using integrated switches and hypercube technology
- ***Substantial savings in cooling costs⁵*** with “Closed-loop Airflow” which ensures no air within the cell is mixed with data center air

 SPEC and the benchmark name SPEC MPI are registered trademarks of the Standard Performance Evaluation Corporation (SPEC). All rights reserved. The stated results are published as of May 12, 2017; see spec.org.

Solving the Unsolvable

TSUBAME 3.0 delivering 47.2 Petaflops at half-precision set for AI Computing



Massively Parallel

540 Compute Nodes, 2160 GPUs
>7.7 Million NVIDIA CUDA Cores
4 Ports of Intel® Omni-Path per node

Tokyo Institute of Technology deploys the most advanced
& largest Open Platform for Big Data analysis

- Deep Learning platform to monitor and observe data processing for **pattern recognition and anomaly detection**
- **Largest Tesla P100 SXM2 deployment to date** with 2,160 NVLink-enabled GPUs
- Powerful enough to support significant AI and scientific HPC workloads providing **unprecedented ability to analyze large data sets.**

Solution for complex commercial HPC challenges

HPE Apollo 6000 Gen10 System purpose-built for performance and rack scale efficiency



Fast, secure and resilient compute, storage and fabric technologies built with rack level efficiencies to deliver **exceptional price performance**

Leading performance

- Uses the new *Intel® Xeon® Processor Scalable family*
- Increase performance and future proofing with persistent storage on memory bus; support for Intel® Xeon® Phi, Integrated OPA Switch and Intel® NVMe SSD drives

Rack-scale efficiency

- *Improve rack level RASM features* through system integration
- *Quickly deploy, service, and manage* with cold aisle front accessible nodes

Purpose-built for HPC

- *Optimize full network switch utilization* with node to fabric alignment
- *Automate task scheduling and management* with Insight Cluster Management Utility

World's largest chemical company creates chemistry with HPC

HPE supercomputer enables global digital transformation at BASF



BASF Supercomputer



- Debuted at #65 on June TOP500
- Designed to be one of the **world's largest supercomputers**
- **1st Supercomputer to use Intel® Xeon® Processor Scalable family**
- Drive digitalization of **BASF's worldwide research**
- Shorten modeling / simulation times (**months to days**)
- Solve complex problems while **decreasing discovery time**
- Run virtual experiments to **reduce time-to-market, lower costs**

Key features

- HPE Apollo 6000 Gen10 System
- > 1 Petaflop using Next Gen platform
- Intel® Omni-Path architecture
- Multitude nodes
- Work simultaneously on highly complex tasks
- Dramatically reduce processing time



“The new supercomputer will promote the application and development of **complex modeling and simulation approaches**, opening up completely new avenues for our research at BASF.”

– **Dr. Martin Brudermueller**, Vice Chairman of the Board of Executive Directors and CTO, BASF

HPC FSI Trade & Match 2.0: Unique HPE/Intel Gold 61xx processor

2 Tiered approach to address both **FSI & EDA**

HPE Gen 10 Servers Trade & Match Offering

2 Tiered Approach



Electronic Design Automation applications

- **Apollo 6000 Gen10 1P & 2P configuration**
- Max density for high frequency configs



Financial Services Industry

- ProLaint DL380: 1P & 2P configurations
- Standard form factor for co-lo deployments



Up to 2.5X
performance increase
over Gen 9 Trade &
Match Solution

Optimised for
applications that run
best at high frequency
and low core count

Gen9 vs. Gen10 Benchmark Comparison


| | SPECint_base 2006 | SPECint_2006 | SPECint_rate_base2006 | SPECint_rate_2006 | LINPAK |
|------------------|-------------------|--------------|-----------------------|-------------------|--------|
| G9 E5-1680v3 | | | 336 | 346 | 408 |
| G10 61xx Gold 1P | 78.6 | 82.4 | 547 | 572 | 647.5 |
| G10 61xx Gold 2P | 80 | 84.1 | 1120 | 1180 | 1105 |

*Spec Int rate result compares Gen 9 vs. Gen 10

Global Energy Company uses supercomputer to pinpoint 200 million-barrel cache

Supercomputer for energy exploration

- **\$2 billion in recoverable oil** in 7,000 feet of water 150 miles from the coast
- Long obscured by a salt dome, the **oil reserves were revealed by using a supercomputer and mathematical algorithm** to interpret the seismic data in a new way.
- **1st large scale deployment** of HPE Apollo 6000 with Intel® Xeon Phi™



2 weeks

and a more
accurate model
using HPC

VS.

1 year

using
traditional
analytics methods

A person wearing a yellow shirt is seen from the side, working on a laptop. The laptop screen displays a dashboard with several circular charts and graphs, likely representing system performance or resource usage. The background shows a server rack with multiple server units, indicating a data center or server room environment. The text "A new Compute experience for HPC and Artificial Intelligence" is overlaid on the image in a large, white, sans-serif font.

A new Compute experience for HPC and Artificial Intelligence

AI versus Competitors



Imagine accelerated apps and faster business insights

Introducing “A New Compute Experience” for HPC and AI

Security



**Protect against Cyber attacks
with HPE Secure Compute
Lifecycle**

Software-defined



**Enhance agility with
comprehensive software
environment**

Economic control



**“Pay as you go” with on-
demand consumption
model**

Powered by the world’s most secure industry-standard servers¹

Security embedded firmware to protect against cyber attacks



World's most secure
industry-standard
servers¹ for HPC and AI

Silicon Root of Trust

- **Only HPE** offers industry standard servers with **firmware anchored into the silicon**
- **Impenetrable protection** through entire supply chain: manufacturing, distribution, shipping, configuration and installation.

Commercial National Security Algorithms

- **Only HPE** offers the **highest level of security**
- Typically used for handling the **most confidential and secret information**
- Uses the **highest level of cryptography** in the industry
- **No increase in server latency**
























Secure Recovery

- **Recovering firmware to known good state** after detection of compromised code
- **Options to recover** to factory settings or last known good or not recovering at all taking server off-line
- Ability to recover other server settings

Firmware Runtime Validation

- **Daily checking** of firmware every 24 hours verifying validity and credibility of UEFI, CPLD, iLO, IE, and ME.
- Valid and secure Firmware copy **stored in lock-box**
- **Alert of compromised code** through iLO audit logs

Comprehensive, integrated software portfolio to enhance agility

| | | | |
|----------------------|--|--|--|
| Software Stack |  Hewlett Packard Enterprise | HPE Performance Software - Core Stack for HPE Apollo, HPE ProLiant | |
| System Management |  Hewlett Packard Enterprise | HPE SGI Management Suite for HPE SGI 8600 HPE Insight Cluster Management Utility for HPE Apollo, HPE ProLiant |  |
| Workload Management |    | | |
| Software Development |  Hewlett Packard Enterprise | HPE Performance Software - Message Passing Interface* |  Intel® Parallel Studio XE |
| | |  NVIDIA CUDA Toolkit |  OpenACC |
| | | |     Mellanox HPC-X: OpenSHMEM, Berkeley UPC, MXM, FCA  |
| Visualization |  Hewlett Packard Enterprise | Remote Visualization Software | |
| Data Management |  Hewlett Packard Enterprise | HPE Clustered Extents File System (CXFS) HPE Data Management Framework (DMF) |  Intel® Enterprise Edition for Lustre* |
| Fabric Software |  Intel® Omni-Path Intel Omni-Path Fabric Software |  Mellanox OFED | System Software |
| | | |    |

HPE Pointnext Services innovation overview

Advisory Services



Create a blueprint of your future state environment allowing for **innovation and growth** with Advisory, Design, Implementation, and Integration Services

Professional and Operational Services



Flawless and on-time implementation, on-budget execution, and creative configurations to **get the most out of software and hardware**

Flexible infrastructure consumption model

By the end of 2018:

50%

Pay as you go models will account for 50% of on-premises and off-premises physical IT and datacenter asset spending by 2018¹

New consumption-based IT providing **scalable capacity on demand, pay only for what you use**; servers, storage, networks, software & services

Faster service delivery to your key business stakeholders with HPE Pointnext Services from the Core data center to Intelligent Edge

HPE
POINTNEXT

A person with blonde hair, wearing a grey jacket, is in the foreground on the left, looking towards the right. The background is a server room with multiple racks of servers. The servers have glowing green lights, and a small screen is visible on one of the racks. The overall lighting is dim, with the green light from the servers providing the main illumination.


**Advanced technology
programs for innovation
roadmap guidance, choice and
flexibility**

HPE wins Department of Energy grant to develop Exascale prototype

HPE leads the charge into Exascale computing




Research grant over three years to develop a reference design for Exascale




Funding for Hewlett Packard Enterprise's best and brightest technology talent



Extends HPE's market leadership in HPC



Validates HPE's strategy to double down on infrastructure and builds on the SGI acquisition



Inspired by Hewlett Packard Labs technology advancements
(i.e. Memory-Driven Computing, silicon photonics and Advanced Non Volatile Memory)



Ecosystem innovation with partners and Gen-Z consortium

Let's Talk. Join us at the HPE Exhibit.

- Join HPE on the Showfloor to learn more about HPE High Performance Computing Solutions
- Complete the HPE Survey Form for your chance to win* a Titan Drone
- Prize will be drawn on Friday 20 October 2017 during the morning tea break at the HPE Exhibit #04



HPE HPC innovation with NASA...





Thank You

