



**Hewlett Packard**  
**Enterprise**

# **Unlock business value with HPC & Artificial Intelligence**

**FORUM TERATEC – June 19, 2018**

**José RODRIGUES – HPC Sales Manager**

# Investment in HPC delivers compelling financial returns

## Financial Services



Accelerating trading and portfolio management

## Oil and Gas



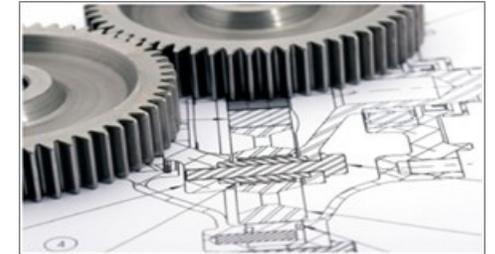
Transforming energy exploration and production

## Life Sciences



Enhancing patient care and predictive medicine

## Manufacturing



Automating product lifecycle management

Average Revenue

\$834M

\$419M

\$160M

\$83M

Avg. Profit or Cost Savings

\$81M

\$46M

\$41M

\$20M



**\$551** in revenue and **\$52** in profits or cost savings for every dollar invested<sup>1</sup>



# Maximize your High Performance Computing ROI with HPE



**Purpose-built, secure platforms**  
for extreme performance, scale and  
efficiency delivering actionable insights

**Flexible consumption models**  
enables and accelerates outcomes and  
provides more choice and agility

**Software-defined architectures**  
optimized to simplify and scale  
massively parallel processing

**Expertise** to advise, deploy, integrate  
and support at every step

**Integrated solutions**  
Tailored for vertical industries and Edge  
to Core Data Center applications

**Technology partner ecosystem**  
for best of breed innovation for  
today and the future

# HPE purpose-built portfolio for HPC and AI

## HPC and AI 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, & efficiency

#### HPE Apollo 6000



Extreme Compute Performance in High Density

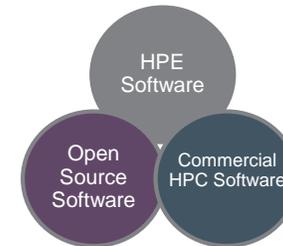
#### HPE Apollo 2000



The bridge to enterprise scale-out architecture

### HPE Performance Software Suite

#### HPE Performance Software Suite



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

### Emerging HPC

#### HPE Apollo 6500



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 and tiered data management platform

### Choice of Fabrics



Arista Networking

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

# Confidently partner with HPE, the proven HPC industry leader

## #1

provider of Top500 energy efficient supercomputers<sup>1</sup>



## #1

HPE is the global HPC market leader with 35.9 percent market share<sup>3</sup>



## #1

New HPE Apollo 6000 Gen10, the world's most secure industry standard server<sup>5</sup> for HPC



## #1

HPE has the most systems in the Top500, with 143 systems in the Top500 List, or 28.6%<sup>2</sup>



## #1

HPE announced world's first Memory-Driven Computing architecture: Major milestone for The Machine project



## #1

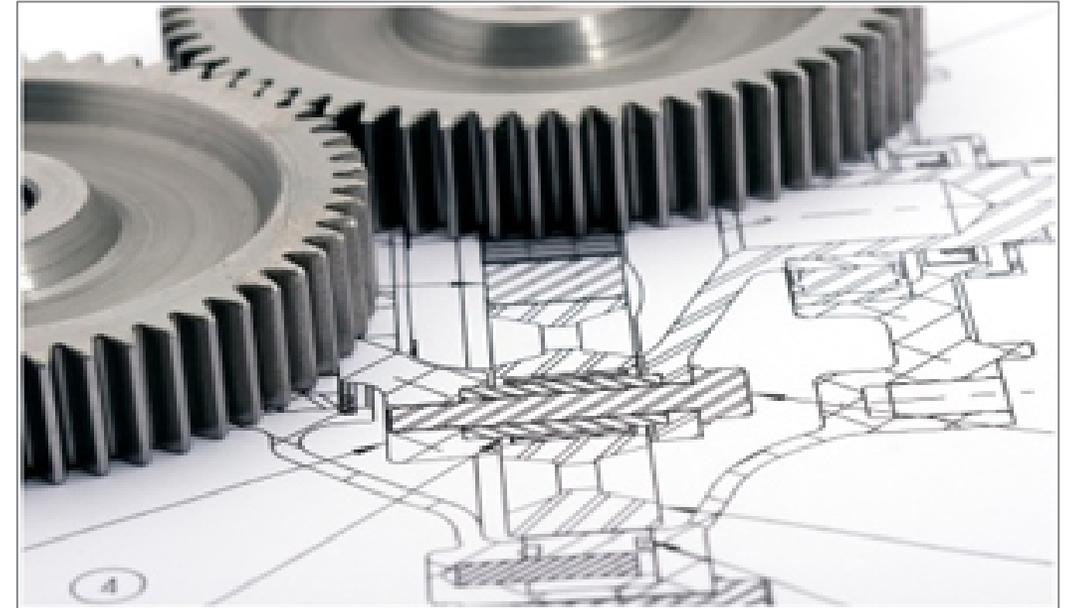
New HPE SGI 8600, delivers industry's fastest parallel processing performance<sup>6</sup>



# HPE is powering BASF's global digital transformation with HPC

## World's largest chemical company creates chemistry with HPC

- **Enables** the development of new plastics, colors or cosmetics
- **Simultaneously** simulates, calculates and evaluates thousands of test series with **HPE Apollo 6000**
- Calculations, which today take about **one year, will run through in only 10 days**<sup>1</sup>
- Classical laboratory experiments are much more targeted and effective than before **saving a lot of money**



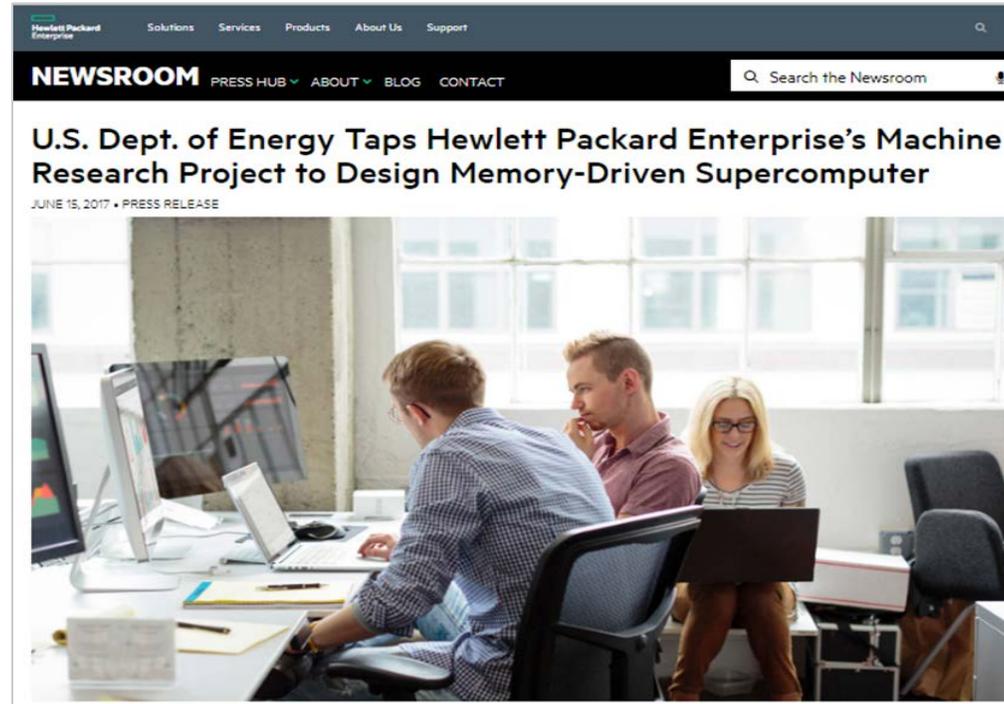
- **Reduces processing times** in chemical research from **months to days**<sup>2</sup>
- **Exceeds the performance** of the computers used to date by **up to 100 times**<sup>3</sup>

# HPE leads the charge into Exascale computing

## PathForward DoE win extends HPE's leadership in HPC

### Win highlights

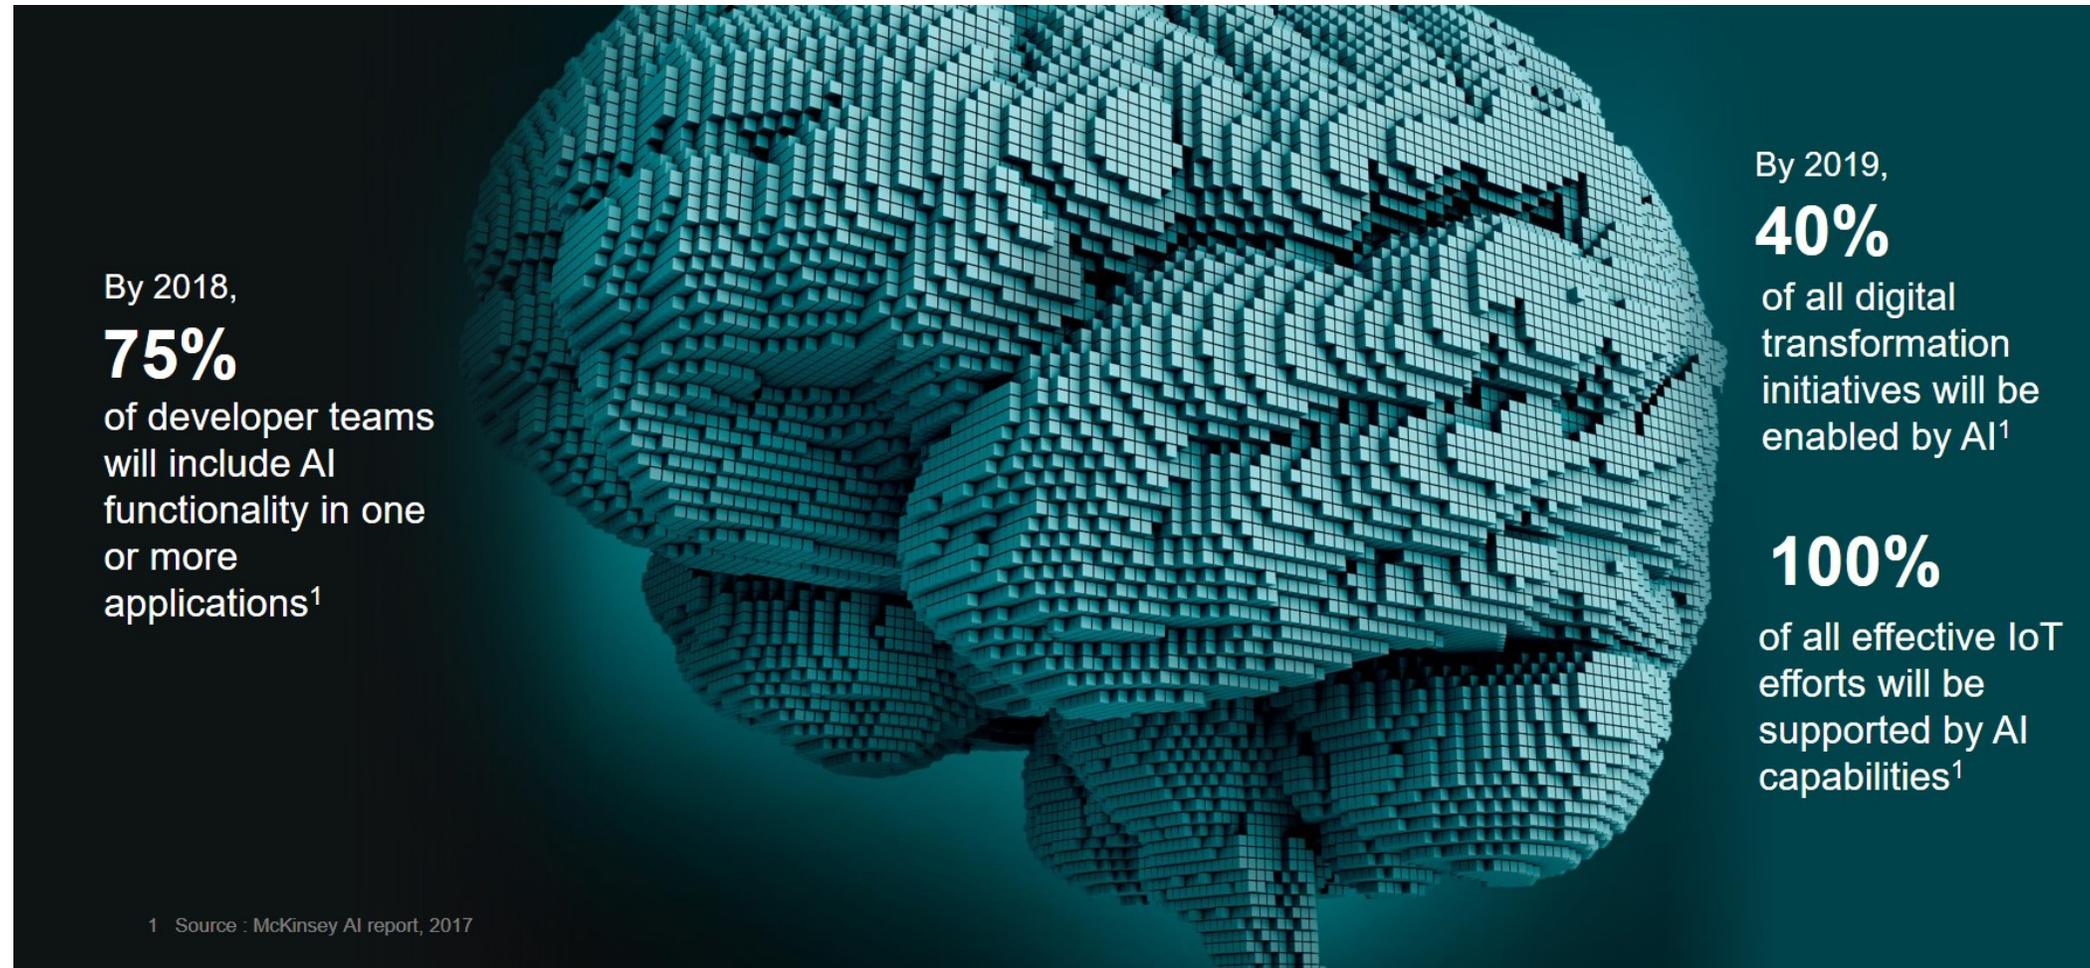
- HPE wins U.S. Department of Energy (DoE) research grant to **develop exascale supercomputer prototype**
- **Inspired by Hewlett Packard Labs** technology advancements
- **Innovative technologies** include industry's first Memory Driven Computing, Silicon photonics and Advanced Non Volatile Memory
- **Ecosystem innovation** with partners and Gen-Z consortium
- **World-class open architectures** based on open industry standards



“As the global HPC market leader, HPE is highly motivated and well positioned to accelerate the move toward this next milestone in computing capability.”

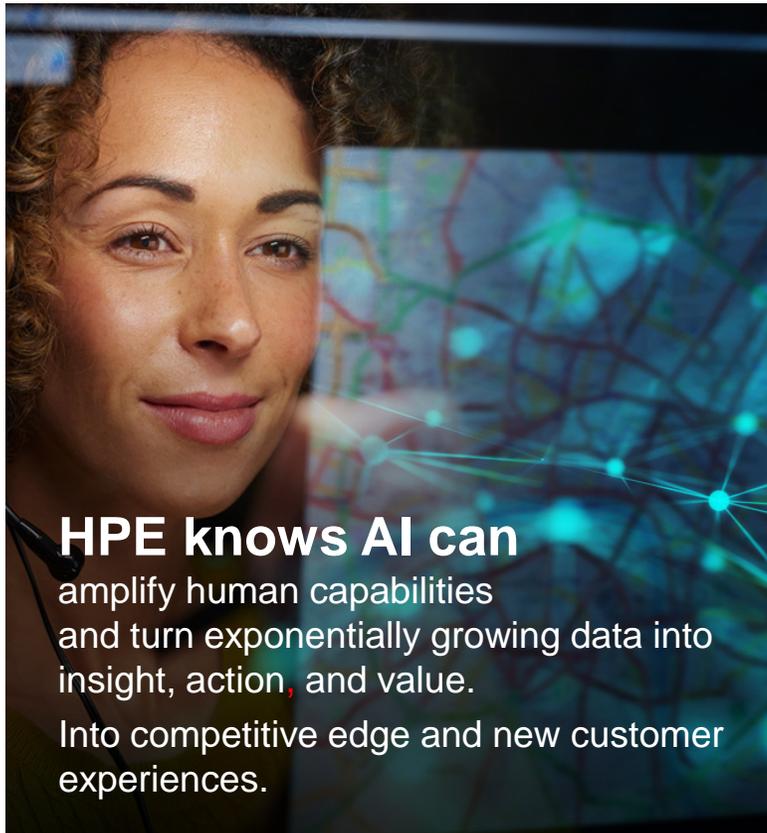
-Steve Conway, Hyperion Research VP of HPC

# What's all the “buzz” around AI?



**Gain competitive advantage using the vibrant new market of AI**

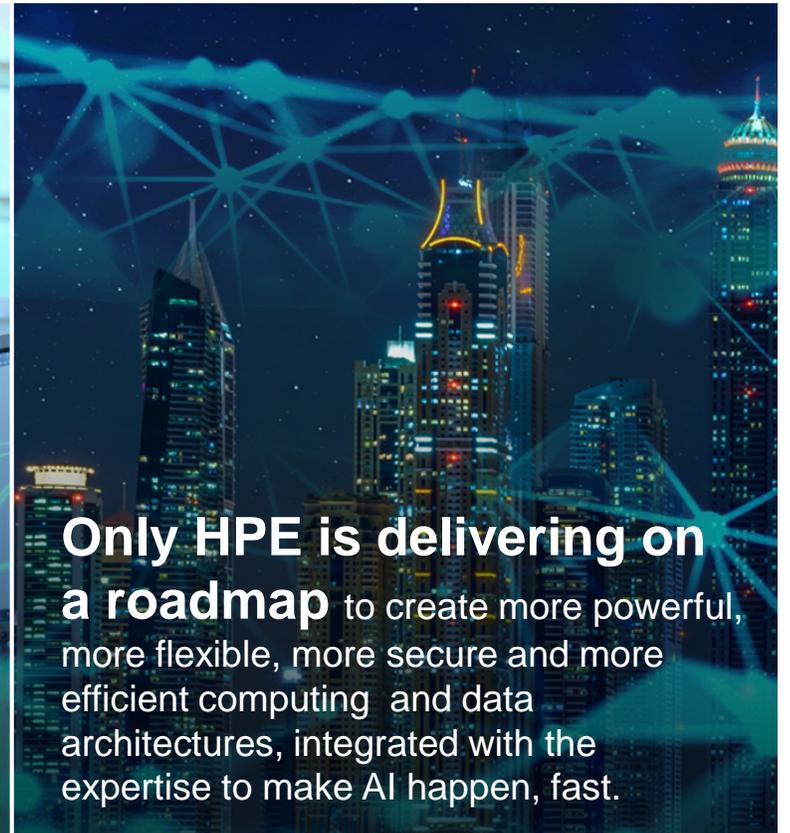
# Artificial intelligence – HPE’s point of view



**HPE knows AI can** amplify human capabilities and turn exponentially growing data into insight, action, and value. Into competitive edge and new customer experiences.



**HPE is the trusted, global partner** enterprises need to navigate this fast-changing field and realize the promise of AI and data analytics everywhere, from IoT to cloud to data center.



**Only HPE is delivering on a roadmap** to create more powerful, more flexible, more secure and more efficient computing and data architectures, integrated with the expertise to make AI happen, fast.

**HPE is powering AI from edge to cloud to core.  
Now and into the future.**

# HPE delivers competitive advantage to customers

## Solving complex AI challenges with hybrid cluster



東京工業大学  
Tokyo Institute of Technology



“Through our partnership with SGI, and now HPE, the Tokyo Institute of Technology has worked successfully to deliver a **converged world-leading HPC and Deep Learning platform....**”

**Satoshi Matsuoka**  
Professor and Tsubame Leader  
Tokyo Institute of Technology..



## Upping the ante on artificial intelligence



“The discoveries and insights our researchers are now uncovering will have direct effects on human lives by way of **advancing precision medicine, increasing energy efficiency, and improving policymaking for the economy.**”

**Dr. Nick Nystrom**  
Senior Director of Research  
Pittsburgh Supercomputing Center

## Libratus AI program defeats world's best poker players



“The best AI’s ability to do **strategic reasoning with imperfect information** has now surpassed that of the best humans.”

**Dr. Tuomas Sandholm**  
Professor of Computer Science Department  
Carnegie Mellon University

# Key challenges for businesses constraining AI adoption

## ★ AI for business

Who can help me understand and plan?



“I could use expert advice and guidance to uncover business opportunities with AI, and to plan and deploy my AI initiatives.”

**Lack of resources or expertise**

What is the optimal solution?



“I need an integrated and scalable solution and infrastructure to rapidly deploy my deep learning models to gain faster insights from my data.”

**Inability to integrate IT environment**

Who are the key AI partners?

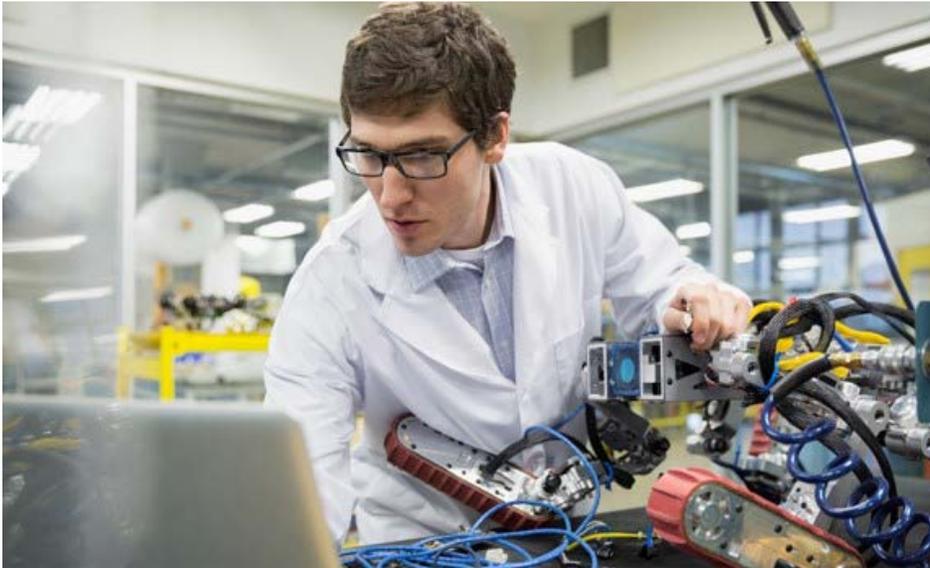


“I need strategic planning and technology help in developing my AI ecosystem.”

**Difficulty in navigating partner ecosystem**

# Tools to guide choice of the best hardware and software for DL

## HPE Deep Learning Cookbook



### Components

- **Deep Learning Benchmarking Suite:** automated benchmarking tool to collect performance measurements on various HW/SW configurations
- **NEW Deep Learning Performance Guide:** web-based tool with access to a knowledge base of benchmarking results to guide choice of optimal hardware / software
- **Reference Designs:** hardware / software recipes for selected workloads

Navigate deep learning ecosystem with confidence

### Customer benefits

Comprehensive set of tools to guide the choice of the best hardware and software environment for different deep learning workloads.

- **Eliminate the “guesswork”** with Deep Learning Performance Guide: tap into a massive pool of performance data to deliver the optimal hardware configuration for workloads
- **Validate the configuration** of customer hardware and software environment with Deep Learning Benchmarking Suite
- **Get started fast** with one of our reference designs as a default technology recipe

# One-day interactive workshop to help customers get started

## HPE Artificial Intelligence Transformation Workshop



HPE  
POINTNEXT

### Scope

- Select and analyze AI, advanced analytics, data use cases
- Identify desired outcomes – automated or manual
- Assess data characteristics – data readiness, requirements

Turn data into action and revenue

### Customer benefits

One-day visual and interactive session allows HPE experts to focus on customers' needs for next generation AI and help embark on a transformational journey to achieve business goals.

- Get started on an AI project fast
- Align business, data and IT operations teams
- Explore opportunities, priorities and select relevant use cases
- Identify dependencies, data sources, level of readiness
- Define a high-level roadmap for intelligent data strategy

# HPE has a comprehensive, purpose-built portfolio for deep learning

Government, academia and industries



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 ideal for training models in data center

## HPE SGI 8600

Petaflop scale for deep learning and HPC



## HPE Apollo 6500

The enterprise bridge to accelerated computing



## HPE Apollo sx40

Maximize GPU capacity and performance with lower TCO



Compute for both training models and inference at edge

## HPE Apollo 2000

The bridge to enterprise scale-out architecture



Edge analytics and inference engine

## HPE Edgeline EL4000

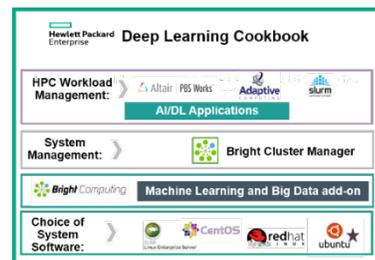
Unprecedented deep edge compute and high capacity storage; open standards



AI Software Framework

## Easy Setup and Flexible OS

Using Bright Computing's distribution of deep learning software development components and workload management tool integration



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

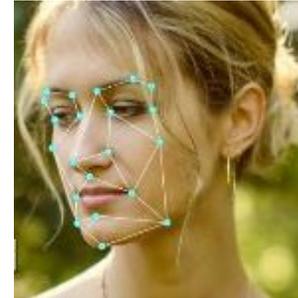
# Enterprise platform for accelerated computing

## Introducing the HPE Apollo 6500 Gen10 System



Reliable deep learning and HPC platform that delivers faster intelligence

### Use cases



- Identify vehicles, pedestrians, and landmarks; e.g., autonomous vehicles
- Monitor oil field drilling rigs to prevent disasters
- Pattern matching for fraud detection
- Speech recognition and translation
- Design drugs

Up to **3X faster**<sup>1</sup> deep learning training with new Gen10 system

### Customer benefits

HPE Apollo 6500 Gen10 System provides superior performance-per-dollar for GPU intensive workloads, with eight NVIDIA Tesla V100 GPUs per server and NVLink interconnect, delivering up to 125 TFlops single precision compute<sup>2</sup> for faster intelligence.

- **Unprecedented performance** delivering economical AI and deep learning
- **Rock-solid, enterprise-level reliability, availability, serviceability** – RAS features
- **Supports a wide range of workloads**, including deep learning and HPC workloads of complex simulation and modeling

<sup>1</sup> Benchmark tests by HPE, Houston 2/18. Additional details in speaker notes.

<sup>2</sup> NVIDIA test results. NVIDIA Tesla V100 provides 15.7 TFlops single precision performance with NVLink 2.0 per GPU x 8 GPUs = 125 TFlops, July 2017. <http://www.nvidia.com/content/PDF/Volta-Datasheet.pdf>

# Collaboration unveils new technology innovation

HPE and NVIDIA



## Shared engineering



NVIDIA NVLink on  
HPE Apollo 6500 Gen10



**New 8 NVIDIA NVLink GPUs on Apollo 6500 Gen10**

## Customer benefits

HPE and NVIDIA team to help customers find new ways to harness massive amounts of their data and leverage AI applications using the NVIDIA NVLink and HPE Apollo systems for deep learning and machine learning techniques.

- **Shared engineering delivering faster deep learning training** using NEW 8 NVIDIA® NVLink™ GPUs on HPE Apollo 6500 Gen10
- **Ultra-fast communication between GPUs** with NVIDIA® NVLink™
- **Up to 3X faster<sup>1</sup> deep learning training** with new Gen10 system



**Hewlett Packard**  
Enterprise

**Merci!**