

HPE AI Solutions

From the Tactical Edge to the Core



Hewlett Packard
Enterprise

Connection[™]
we solve IT[™]

Best-in-class AI technology






Hewlett Packard
Enterprise

AI Terminology

VERY Abridged Version


Deep Learning

Untrained
Neural Network Model



Deep Learning


Untrained
Neural Network Model




Training
Learning a new capability
from existing data

Deep Learning

Untrained
Neural Network Model




Training
Learning a new capability
from existing data




Deep Learning

Untrained
Neural Network Model




Training
Learning a new capability
from existing data




Deep Learning

Untrained
Neural Network Model




Training
Learning a new capability
from existing data




Deep Learning


Untrained
Neural Network Model



Training
Learning a new capability
from existing data




Trained Model
New capability
optimized for
performance




Deep Learning


Untrained
Neural Network Model



Training
Learning a new capability
from existing data




Trained Model
New capability
optimized for
performance




Inference
Applying this capability
to new data

Deep Learning


Untrained
Neural Network Model




Training
Learning a new capability
from existing data



Trained Model
New capability
optimized for
performance




Inference
Applying this capability
to new data

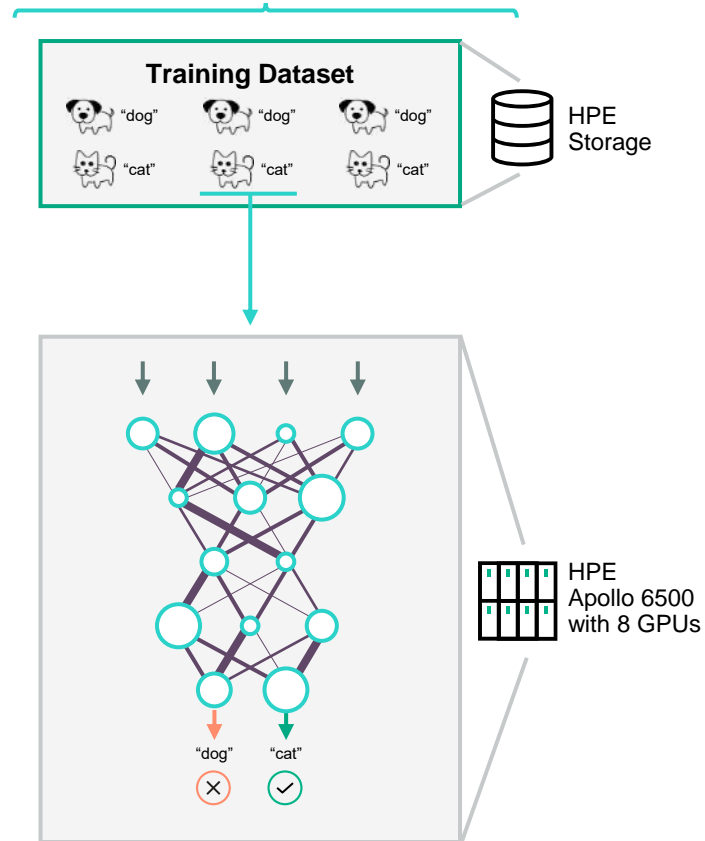


Deep Learning


Untrained
Neural Network Model



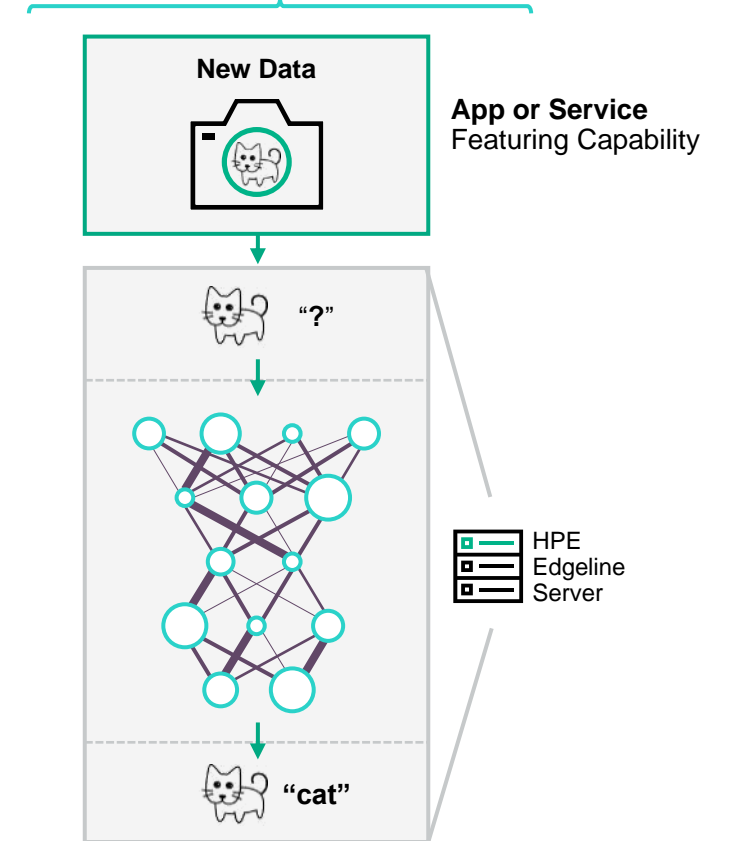
Training
Learning a new capability
from existing data



Trained Model
New capability
optimized for
performance




Inference
Applying this capability
to new data

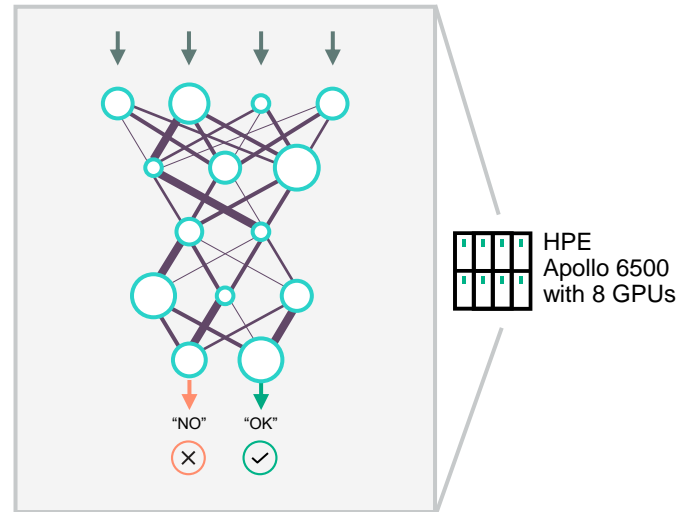



Deep Learning


Untrained
Neural Network Model



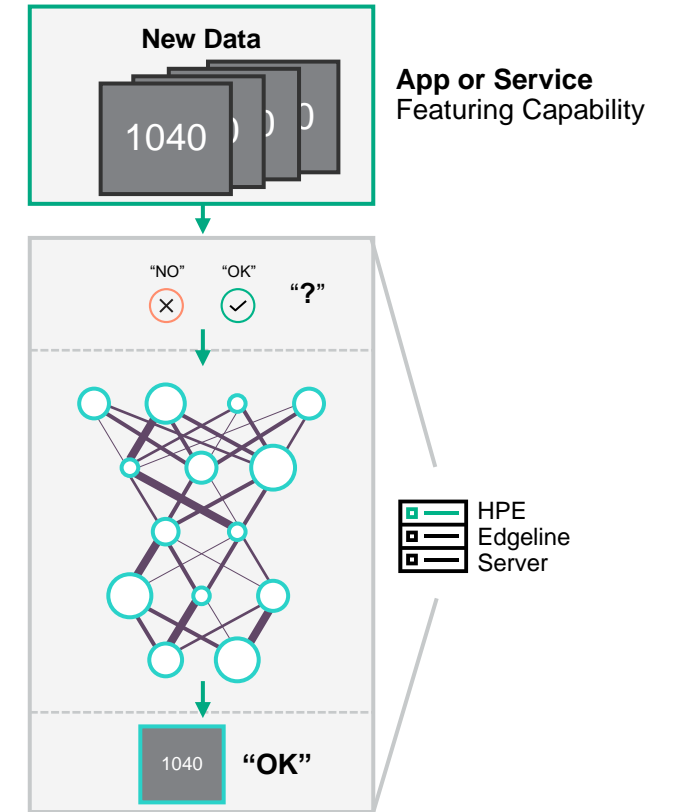
Training
Learning a new capability
from existing data



Trained Model
New capability
optimized for
performance

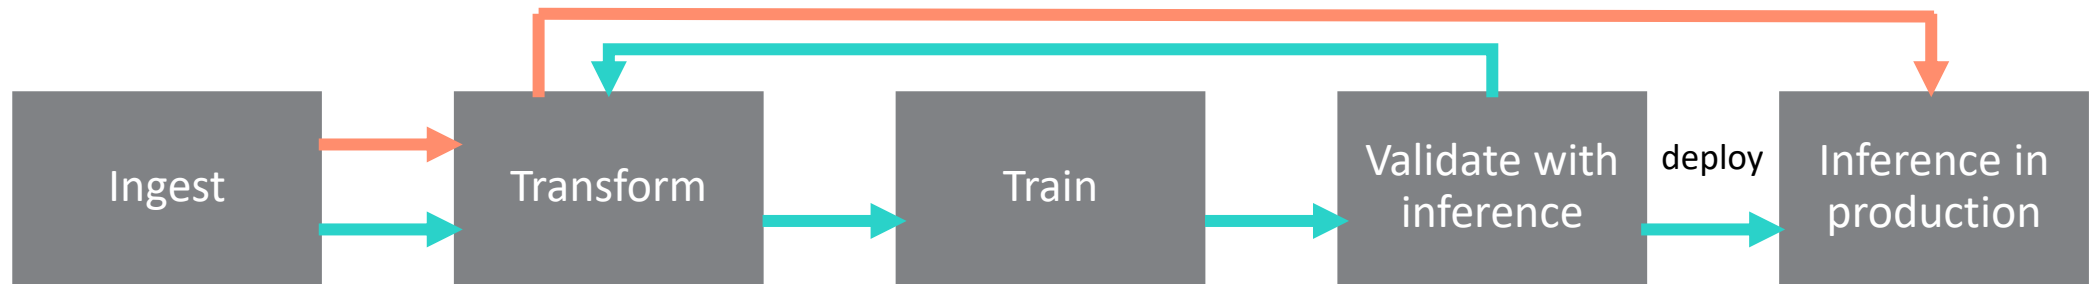


Inference
Applying this capability
to new data




Phases of Machine Learning

— training
— inference

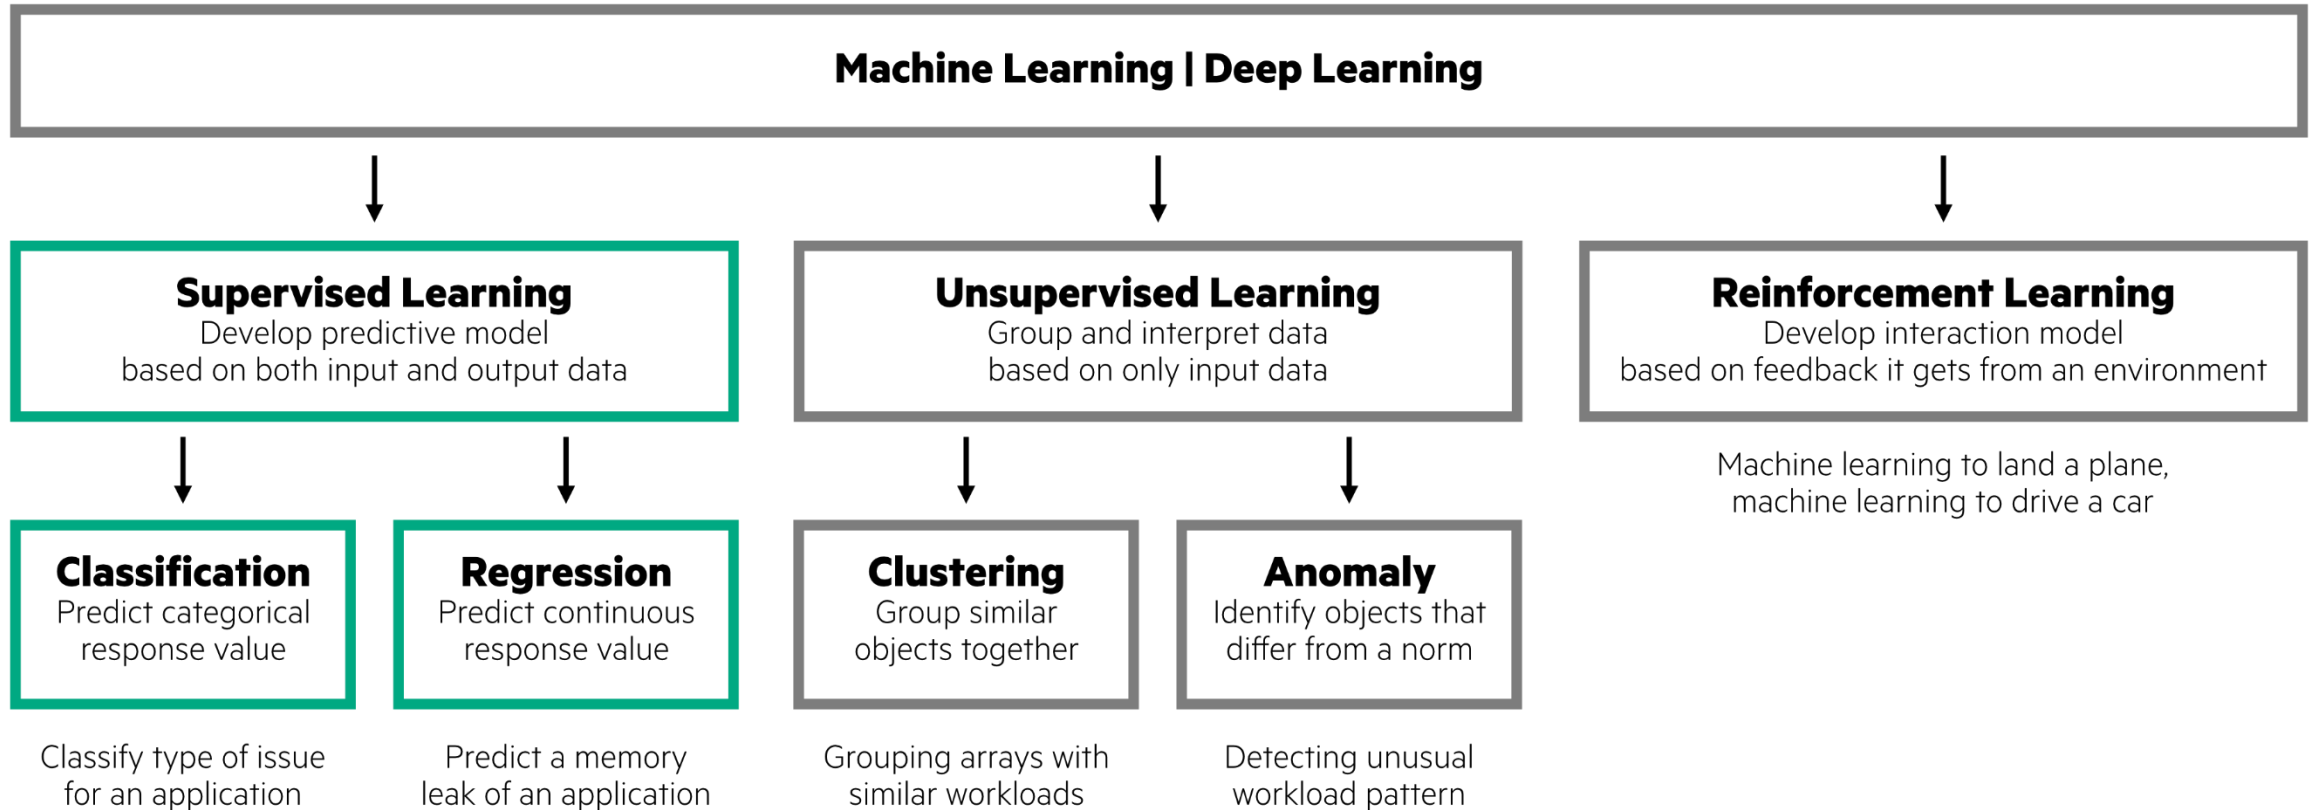


NVIDIA TensorRT



The Microsoft Cognitive Toolkit (CNTK)

Machine learning categories



Basic Things AI Can Do Today

Recognizing – What is this object?

Detecting – Did something change?

Classifying – Where does this fit?

Searching – Where is something?

Identifying – Who is this?

Grouping – What things are alike?

Ranking – What's best/worst?

Recommending – Try this based on what's known.

Filtering – Removing unwanted content.

Predicting/forecasting – Based on the past data.

Optimizing – These actions will give you better outcome



**Hewlett Packard
Enterprise**

HPE AI Focus Areas

Use Cases

IDENTIFYING OPPORTUNITIES

Look for these initiatives

Quality control using prescriptive maintenance and image, video analytics

Integrating IT and OT data to improve quality and processes

- Asset maintenance
- Defect identification
- Parts traceability
- Robotics analytics



Surveillance using video analytics

Operational insights based on captured video data, ranging from:

- Facial recognition
- Queue monitoring
- Unattended items



Speech to Text natural language processing

Communication surveillance


- Speech to Text
- Biometric search
- Live Call monitoring
- Chatbot



Highly-autonomous driving

Implementation of level 3 and level 4 autonomous driving

- Simulations, testing
- In-vehicle analytics
- Road, car condition prediction



PRESCRIPTIVE MAINTENANCE AND QUALITY CONTROL

Improve efficiency and quality in the production, manufacturing, retail processes through predictive, prescriptive and image analytics/ML.

Benefits

- Decrease cost of asset maintenance
- Faster defect identification
- Parts traceability

Industries

Manufacturing, energy and utilities, aerospace, defense, transportation, logistics, retail

HPE Solution

Through open framework for IT/ OT integration the solution prescribes best maintenance, decisions automates actions and enhances quality control processes by implementing AI-based computer vision techniques.

References

[Seagate video](#)

[Seagate case study](#)

[Seagate blog](#)

[Texmark](#)- predictive maintenance

Resources

[Technical Whitepaper](#)

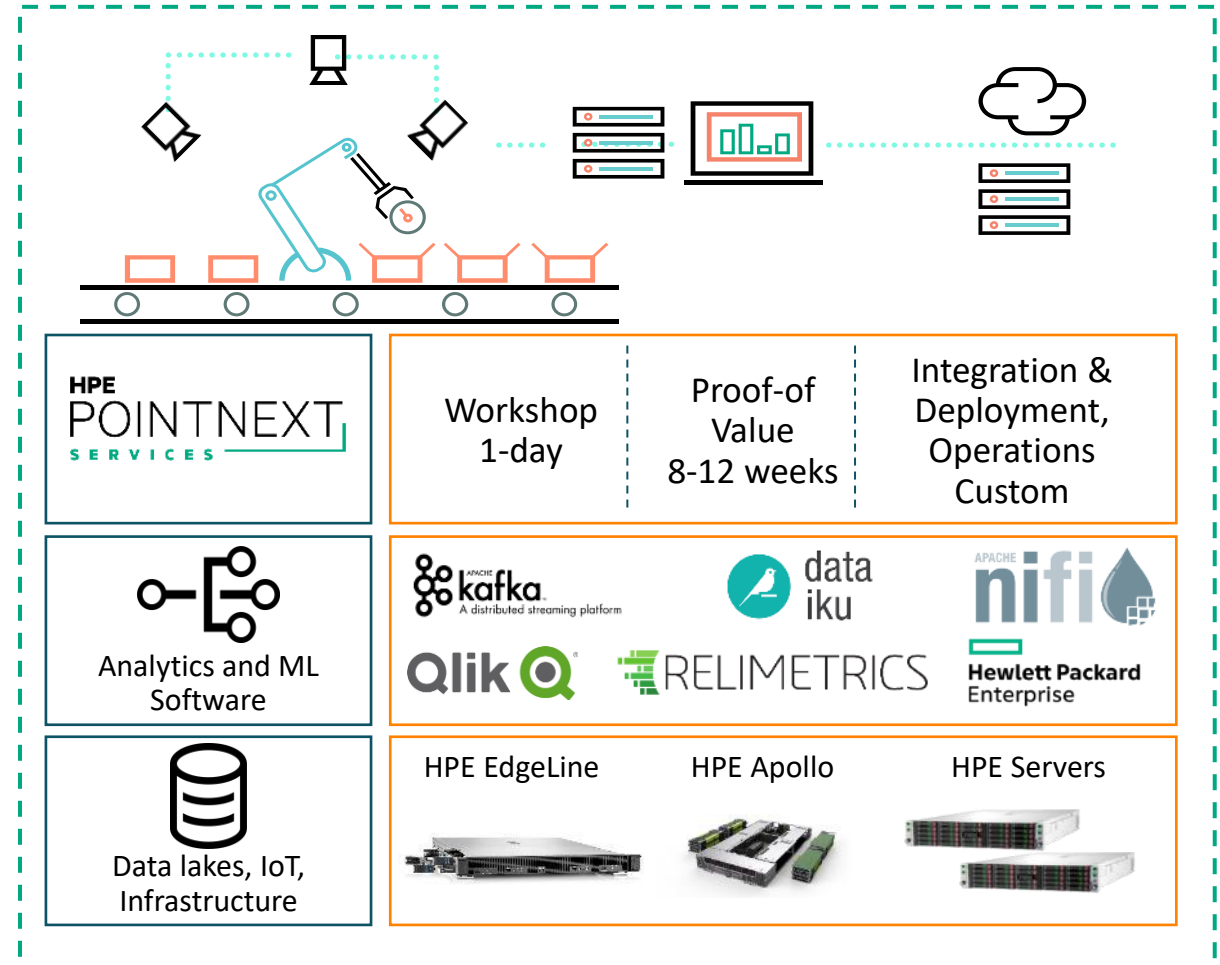
[Customer presentation](#)

[Sales play](#)

[Service brief](#)

[Demo \(youtube\)](#)

[Sales training replay](#) & [slides](#)



SPEECH AND LANGUAGE PROCESSING

Automate the understanding of spoken and written language by machines applying natural language processing, speech to text analytics, biometric search or live call monitoring.

Benefits

- Fraud detection, better compliance- voice recognition
- Improved customer service-call center monitoring
- Faster emergency, law enforcement services
- Legal discovery

HPE Solution

Secure, optimized speech and NLP solutions tuned to transform unstructured audio data into a rich set of semantic data for instant insight and intelligence.

Industries

Financial services, healthcare, government, sports/entertainment, publishing

Resources

[Technical whitepaper](#)

[Customer deck](#)

[Solution brief](#)

[Video/ demo](#)

[HPE Deep Learning with NVIDIA NGC tuning guidelines, white paper](#)



Workshop 1-day	Proof-of Value 8-12 weeks	Integration & Deployment, Operations Custom
<ul style="list-style-type: none"> • High-speed automatic speech recognition (ASR) • Domain-specific language model building • Biometric identification and search • Live call monitoring • PCI redaction • Advanced language detection for English 	<ul style="list-style-type: none"> • and non-English • SmartTranscript • Topic extraction • Sentiment analysis • Super search • API based integration 	

Training	Training/Inference		Inference
<p>HPE Apollo 6500 System: Enterprise data center training supporting NVIDIA NVLink technology</p>	<p>HPE ProLiant DL380 Server: Entry data center training and inference HPE Apollo 2000 G10 Server: training and inference</p>	<p>HPE Edgeline EL8000 Converged Edge System: Edge training and inference</p>	<p>HPE Edgeline EL1000/4000 Converged Edge Systems: Edge inference</p>
<p>NVIDIA® Tesla® V100 SXM2 GPUs: Ultimate deep learning for training performance and scale-up</p>	<p>NVIDIA Tesla V100 PCIe GPUs: Training with light inference</p>	<p>or</p> <p>NVIDIA® T4 GPUs: Inference with light training</p>	<p>NVIDIA T4 GPUs: Low power, low profile, optimized for scale-out and mainstream enterprise acceleration</p>



VIDEO ANALYTICS

Automatically analyzes video to detect events, uncover identity, environment, people and obtain operational insights based on captured video data for safety and quality assurance

Benefits:

- Decrease human error-real-time video analytics
- Shorten wait times -queue monitoring
- Increase building, facility security
- Detection of unattended items- increased safety
- Reduction in defects, rework, improved inspection

HPE Solution

Edge-to-core video analytics systems for a wide variety of workload and operating conditions.

Industries

Airports, buildings, schools, government/smart city, manufacturing, retail

References

- [Foxconn Upgrades Quality with Video Analytics](#)
- [Purdue University](#)

Resources

- [Technical whitepaper](#)
- [Customer deck](#)
- [EdgeLine brochure](#)



Workshop 1-day	Proof-of Value 8-12 weeks	Integration & Deployment, Operations Custom
-------------------	------------------------------	---



- **Rugged, edge-optimized compute platforms: HPE Edgeline EL300, EL1000, EL4000 and EL8000 Converged Edge Systems**
- **Scalable storage platforms for video retention: HPE MSA 2050, HPE Nimble Storage, HPE 3PAR, HPE Apollo with Qumulo, HPE Scalable Object Storage with Scality RING**

HIGHLY-AUTONOMOUS DRIVING (HAD)

Provide the data platform and infrastructure enabling developers to build the optimum HAD solutions for the robust and safe transportation grids of tomorrow based on machine learning and deep learning neural networks, cutting edge accelerated compute, high-bandwidth networks and interconnect fabrics.

Benefits

- Reduce amount of traffic fatalities
- Reduce congestion enabling automated driving
- Enable faster in-car data analytics

Industries: Transportation, Automotive

HPE Solution

A scale-out data ingestion platform built and tuned for; open source services, machine learning and deep learning neural networks. Hardware consists of cutting edge accelerated compute, high-bandwidth networks and interconnect fabrics.

Resources


[Technical whitepaper](#)



Workshop
1-day

Proof-of Value
8-12 weeks

Integration &
Deployment, Operations
Custom



Infrastructure

- Onboard data logging and terminal systems for data offloading
 - HPE Edgeline EL8000
- Compute platforms for training, simulation and re-simulation
 - HPE Apollo 6500 Gen10, 6000 Gen10
 - HPE ProLiant DL380 Gen10
- HIL/SIL simulation systems
 - HPE Apollo 2000 Gen10 with XL170 CPU and XL190 GPU
- Archive and backup solutions
 - HPE Data Management Framework (DMF)
 - HPE Tfinity Exascale Tape Library

AI SOLUTION TECHNICAL WHITEPAPERS

Hewlett Packard
Enterprise

Technical white paper

DIGITAL PRESCRIPTIVE MAINTENANCE AND QUALITY CONTROL

Building an AI solution from edge to cloud



[Prescriptive maintenance and quality control](#)

Hewlett Packard
Enterprise

Technical white paper

DEVELOPMENT OF A HIGHLY AUTONOMOUS DRIVING INFRASTRUCTURE

A technology blueprint overview



[Highly Autonomous Driving Infrastructure](#)

Hewlett Packard
Enterprise

Technical white paper

SPEECH AND NATURAL LANGUAGE PROCESSING SOLUTION

Building an AI solution from edge to core



[Speech to Text natural language processing](#)

Hewlett Packard
Enterprise

Technical white paper

VIDEO ANALYTICS AND SURVEILLANCE

Building a scalable end-to-end AI solution




[Video analytics and surveillance](#)

VIDEO ANALYTICS FOR SAFETY

New scenarios for our current Video Analytics value proposition

3 VIDEO ANALYTICS USE CASES – SINGLE TECHNOLOGY STACK

Social Distancing




- Leverage existing cameras or deploy new
- Triangulate individuals' location in 3D space
- Measure distance between individuals in the frame
- Count people within a defined zone
- Trigger events when distances or counts are breached
- Store data for later trend analysis for future planning

Fever Detection



- Leverage heat sensitive cameras to determine hot spots within a crowd of people.
- Using Machine Learning algorithms to determine heat temperature is from a person and not from other objects.
- Use advanced analytics to determine temperature difference vs. the rest of the people to remove environmental influence and isolate them.

PPE Detection



- Leverage existing cameras or deploy new.
- Using object identification to find individuals and faces, arms, torso in a frame.
- Use AI models to determine PPE compliance.
- Deploy in locations where PPE such as masks are mandatory including health centers, public transport hubs and large venues

COMPLIANCE – VIDEO ANALYTICS SOLUTION

Social Distancing : Crowd Counting

Objectives:

Generate alerts when a crowd breaches the certain threshold .
(Thresholds are based on Zone wise)

Approach:

Train a person detector AI model and then estimate the crowd counts.

Generate real-time alerts where people are estimated to form close clusters and exceed crowd density thresholds. Option to mark and exclude certain hotspots like major displays, to make them out of scope for alarms.

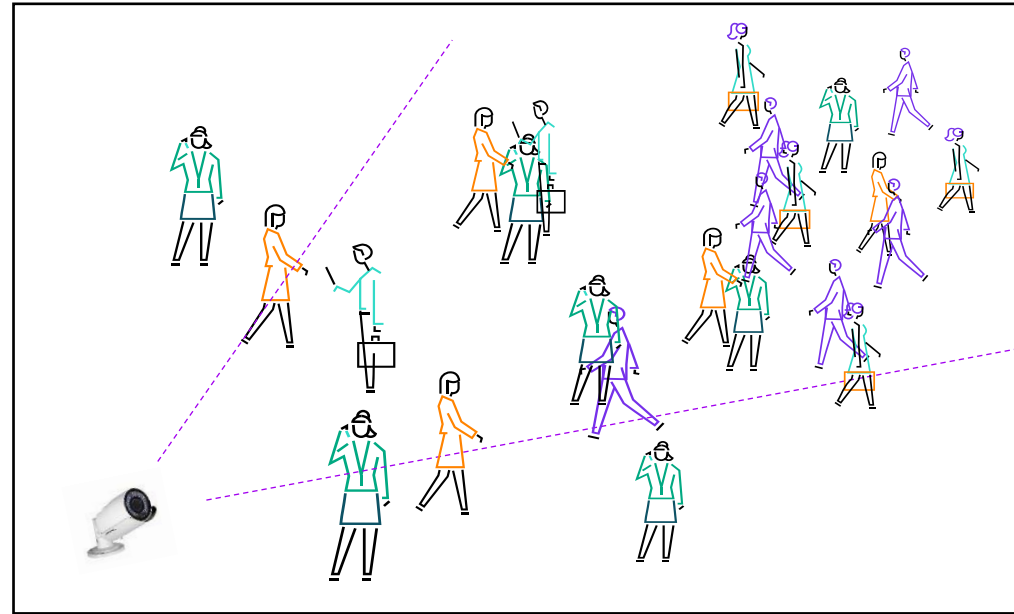
Requirements:

- For initial prototype, needs video from multiple camera feeds monitoring several zones of interest inside the Airport

Target accounts

- Airports, buildings, Educational Institutes, government / smart city, manufacturing, retail, hotels, shopping centers

Each camera will adaptively estimate crowd in it's visual field



COMPLIANCE – VIDEO ANALYTICS SOLUTION

Face Mask Detection

Objectives:

- Generate alerts where a person is not wearing a face mask in the field of view of a camera for a zone

Approach:

Creating Synthetic mask datasets with existing data

Train a classifier model to detect mask/without mask and generate alerts.


Requirements & Limitations

- Feasibility and visual field will depend on the height at which a camera is mounted
- Accuracy and Non-detection will depend on the angle of the face wrt to the camera and the distance
- Certain mask types may result in non-detection specially at poor angles

Target accounts

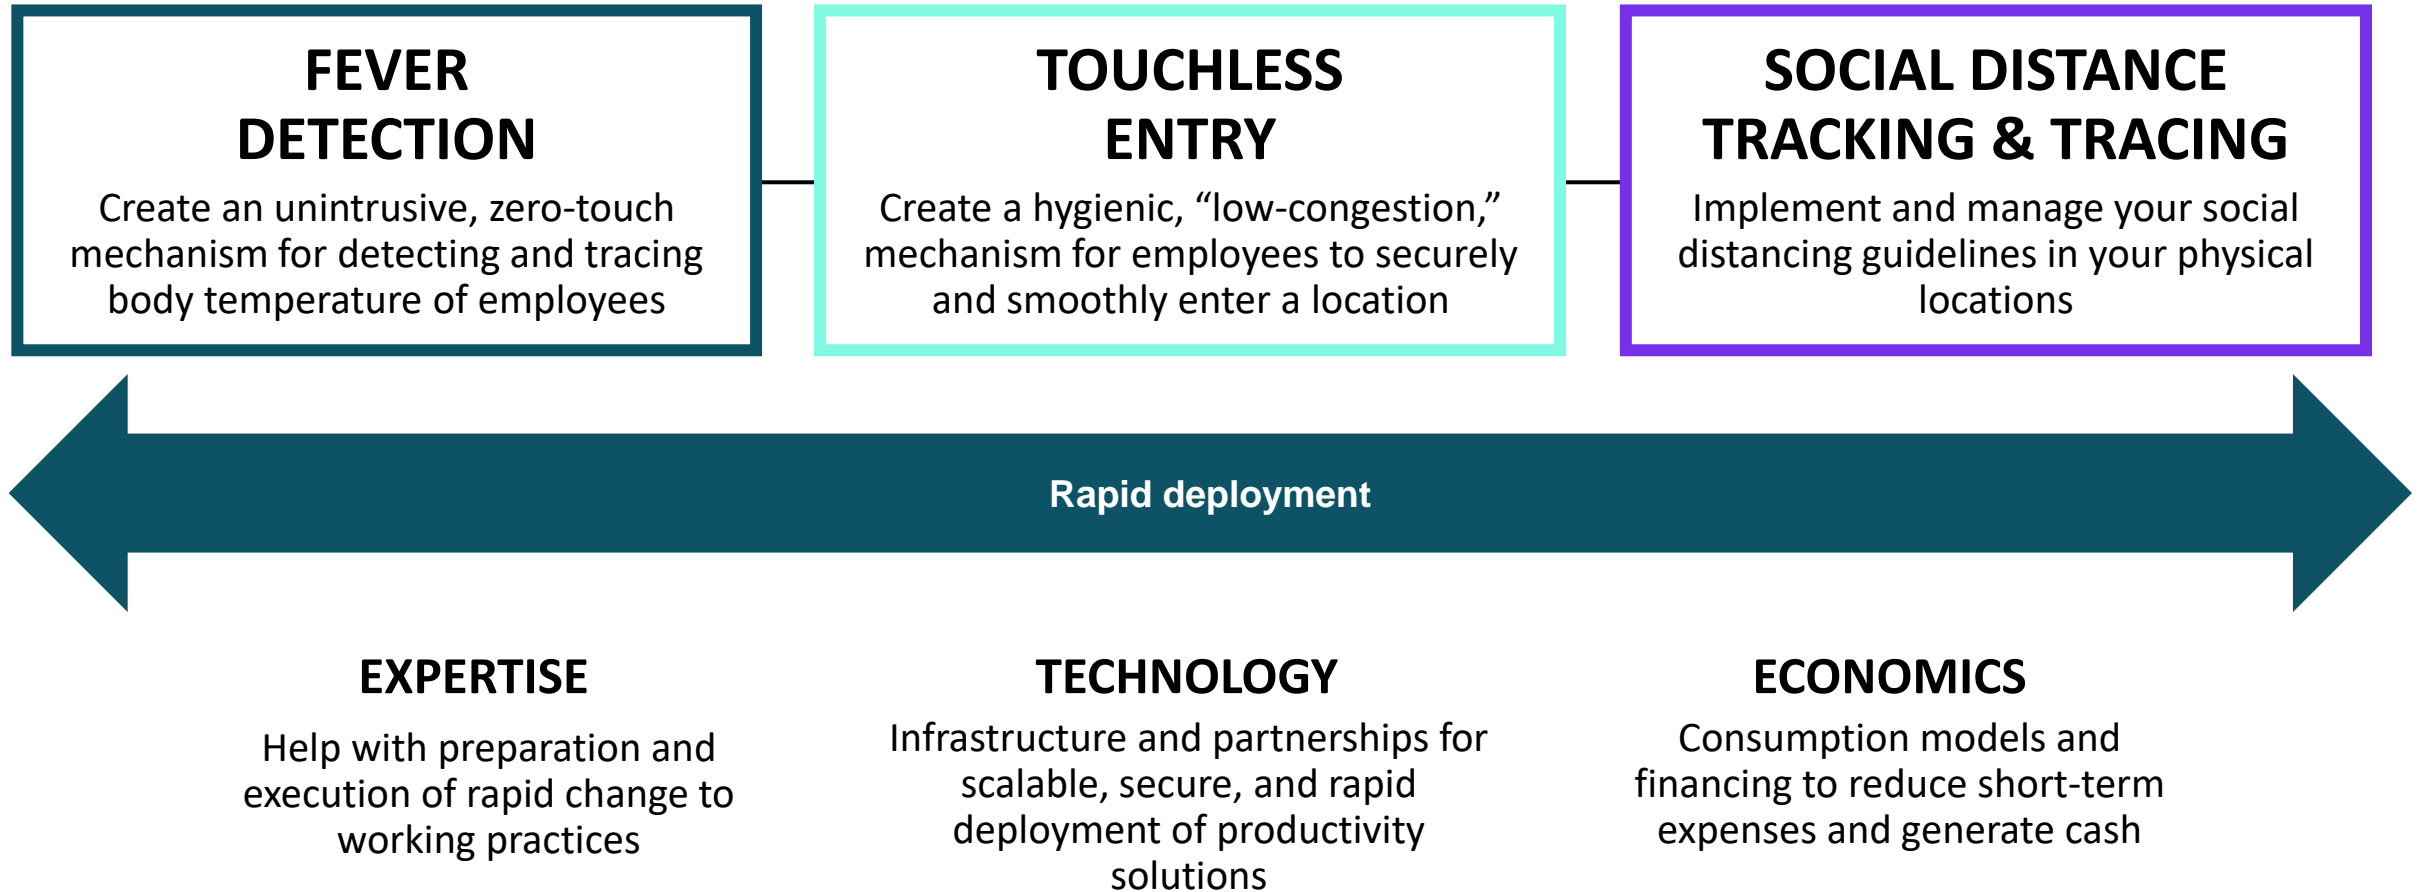
- Airports, buildings, Educational Institutes, government / smart city, manufacturing, retail, hotels, shopping centers

Height of the camera mount, resolution for the visual field to be covered should be optimal
Angle of the face and the distance will affect the detection / non-detection of facial mask



HPE PARTNERING TO CREATE INTELLIGENT WORKPLACE SOLUTIONS

Helping your employees get back to work



FEVER DETECTION SOLUTION FOR KRAFT HEINZ

Solution overview

Per Entrance



KraftHeinz

\$5M Deal

40 Manufacturing Sites

HPE
POINTNEXT
SERVICES

- Advisory and design
- Professional services
- Integration and deployment solution (CDIS)
- End-user training
- Solution support
- Managed services

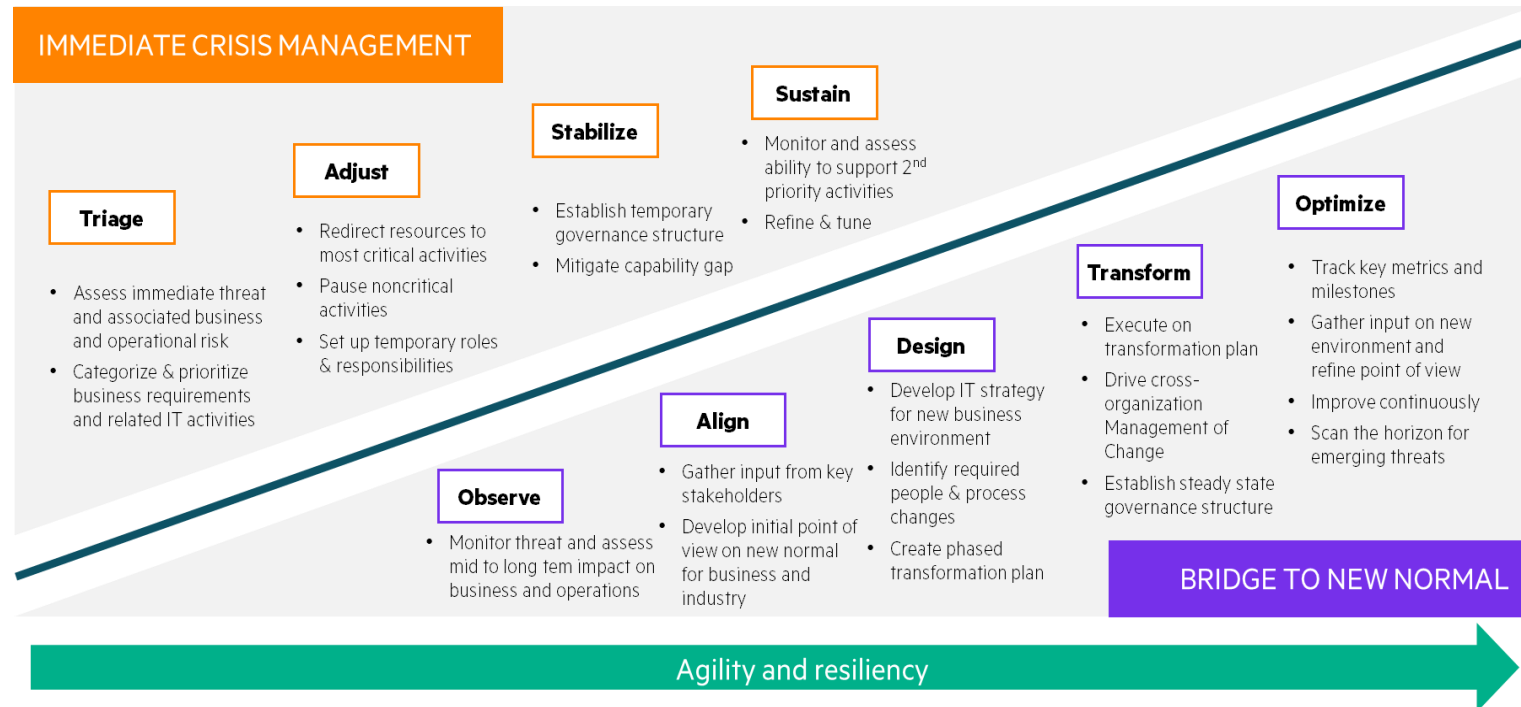
Per Site



HPE Compute/Storage portfolio
(DL380, Edgeline, Aruba
Switches)

THOUGHT LEADERSHIP: CRITICAL ELEMENT OF ADVISORY CAPABILITY

30 thought leadership articles giving insights and expertise



Published articles

- [Nine steps to the new normal](#) [Last mile challenges to remote work](#)
- [Security for the rapidly expanding remote workforce](#)
- [Flexible ERP platforms critical to the new normal](#)
- [How Covid-19 accelerated the move to hybrid cloud](#)
- [From crisis management to the new normal for data centers](#)
- [Expert advice on containerization in enterprise IT](#)
- [Data aggregation during a pandemic](#)
- [How Smart Cities handle pandemics](#)
- [Why IT practitioners are going to lead the way to the new normal](#)
- [How can building skills today prepare your business and employees for the new normal?](#)
- [5 business continuity priority shifts for remote working](#)
- [Operational Resiliency in Crisis Management](#)
- [IT Governance best practices: Practical examples to support the business during crises](#)

Coming soon:

- Management of change and 9 steps to the new normal
- Business continuity, remote working, and changed priorities
- Digital Journey Map to create bridge to new normal
- Optimizing your new minimum-viable operating model in a time of crisis
- Etc.

- Sales Collateral in Seismic Briefcase: [Selling in Uncertain Times](#)
- Thought leadership articles published on an ongoing basis and broadly shared through social channels
- Special edition of The Doppler with all articles to kick-off Discover Virtual Event

AI will produce ~~business~~ *organization* value everywhere

Creation of insights that personalize the customer experience

New automated processes that reduce friction and improve business efficiency

New business models, new industries

HEALTH

- Self-testing
- Diet/supplement recommendations
- Self-triage at home
- Home monitoring
- Fall monitoring
- Habit monitoring
- Personalized drugs
- Disease prediction
- Personal bots for the elderly
- Interactive register & triage kiosks
- Robot-assisted surgery

RETAIL

- User characterization
- Upsell / cross-sell
- Facial recognition
- Computer vision
- Deep learning for auto-checkout
- Bundle design
- Real-time store & online pricing
- Quality & dissatisfaction monitoring
- Supply chain optimization
- Store design
- Continuous inventory tracking
- Counterfeiting source detection

UTILITIES

- Field gear predictive maintenance
- Energy source forecasting
- Self-generation use
- Selling of power
- Fault detection via drones/robots
- Micro-adjustments for efficiency
- Automating call centers

AGRICULTURE

- Precision farming identification
- Precision farming application
- Physical security
- Optimum lighting
- Optimum nutrition
- Crop picking and processing
- Self-driving farm machinery
- Autonomous drones

Crossover Applications for Government Use



Hewlett Packard
Enterprise

HPE Data Analytics Strategy

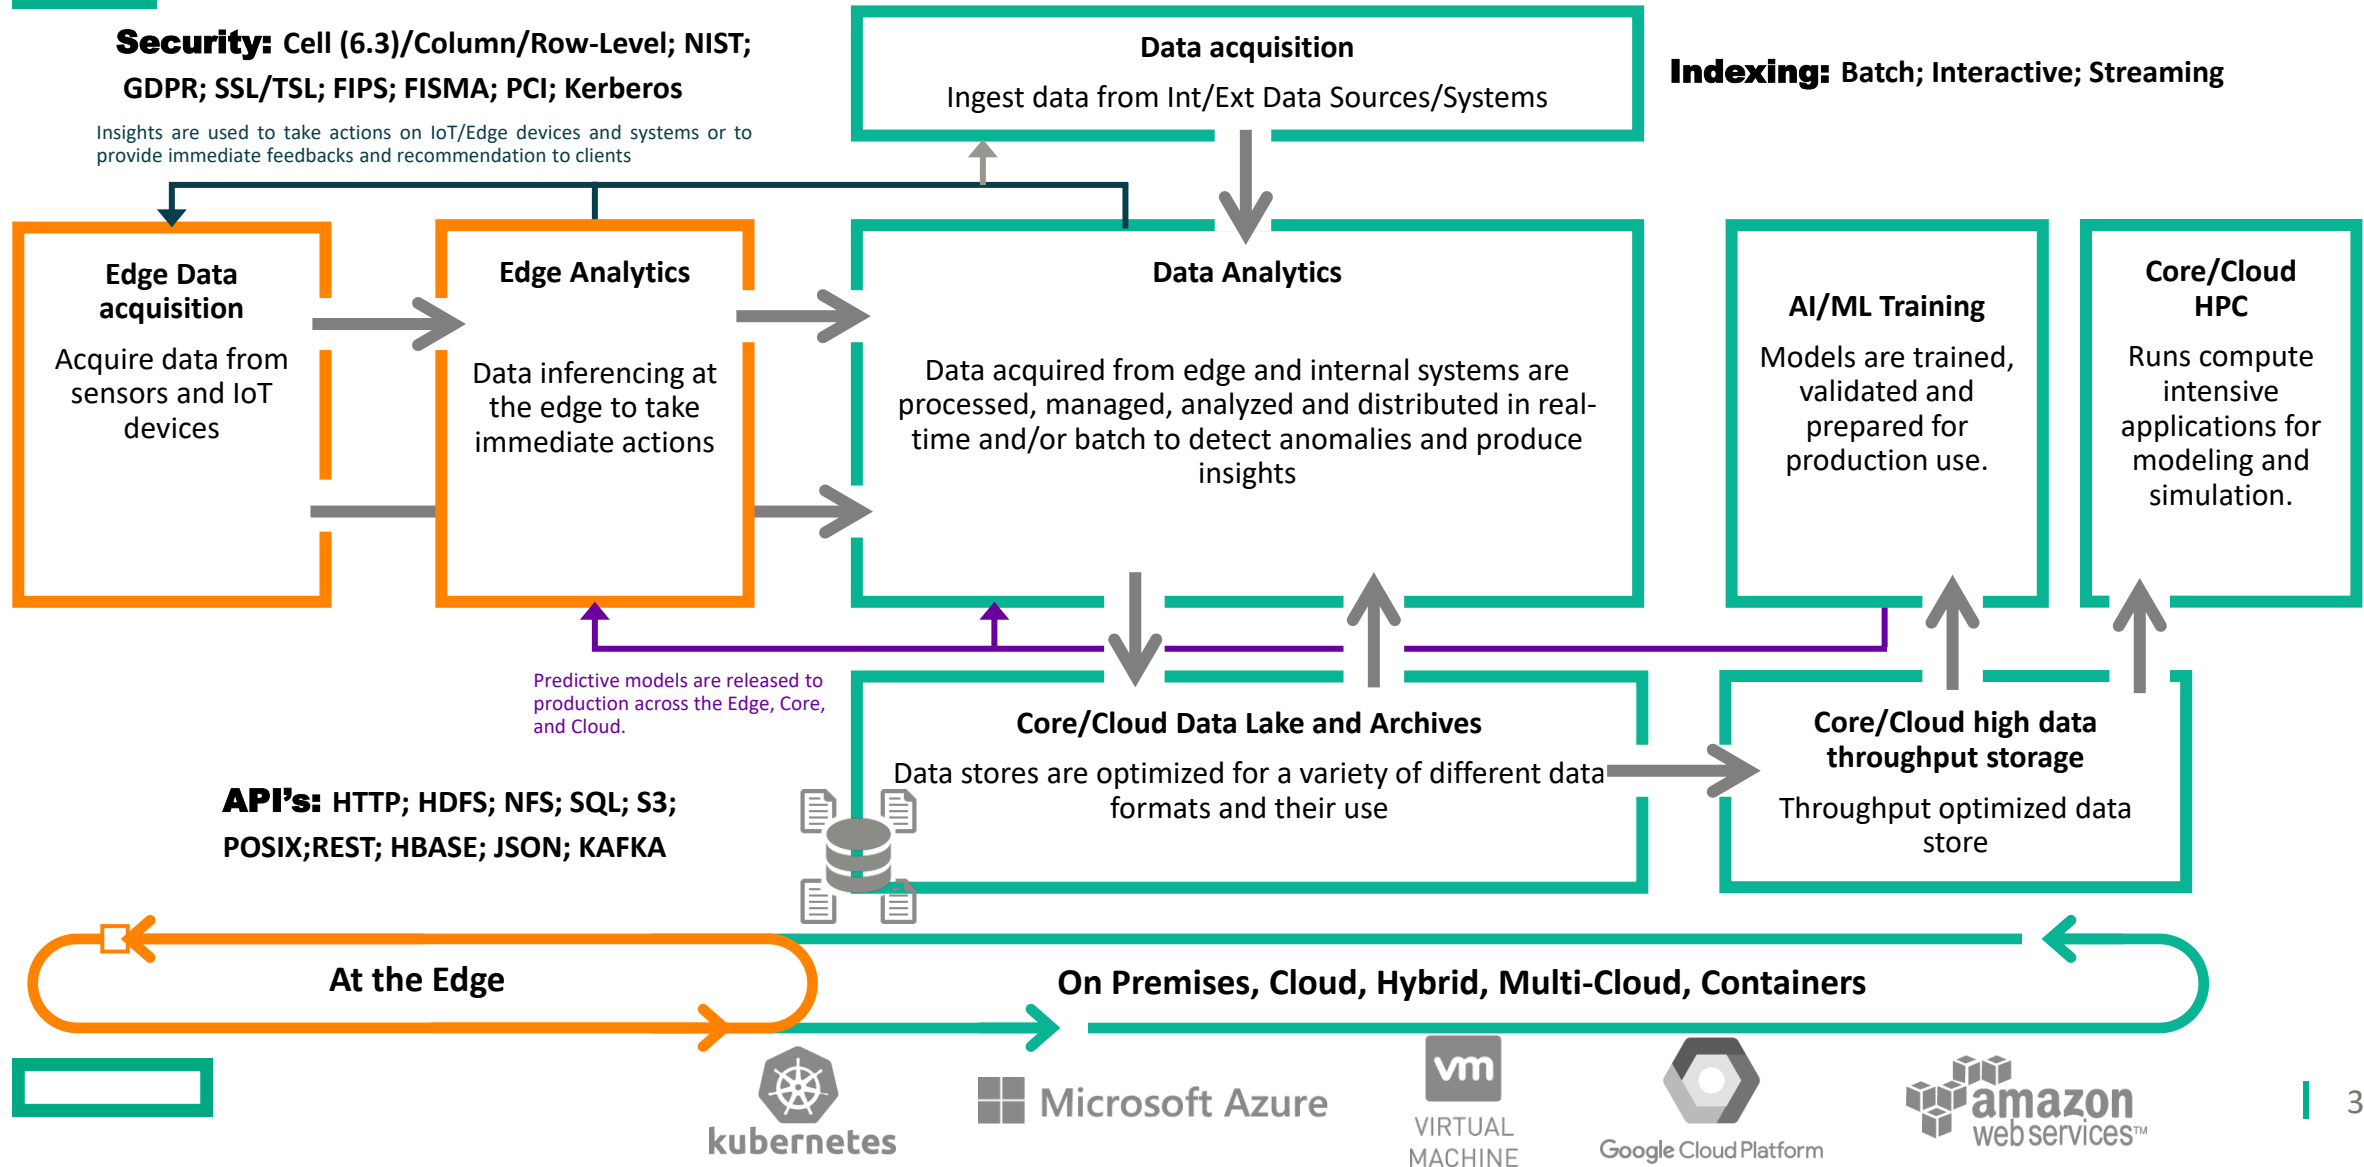
End to End, Top to Bottom

Best-in-class AI technology

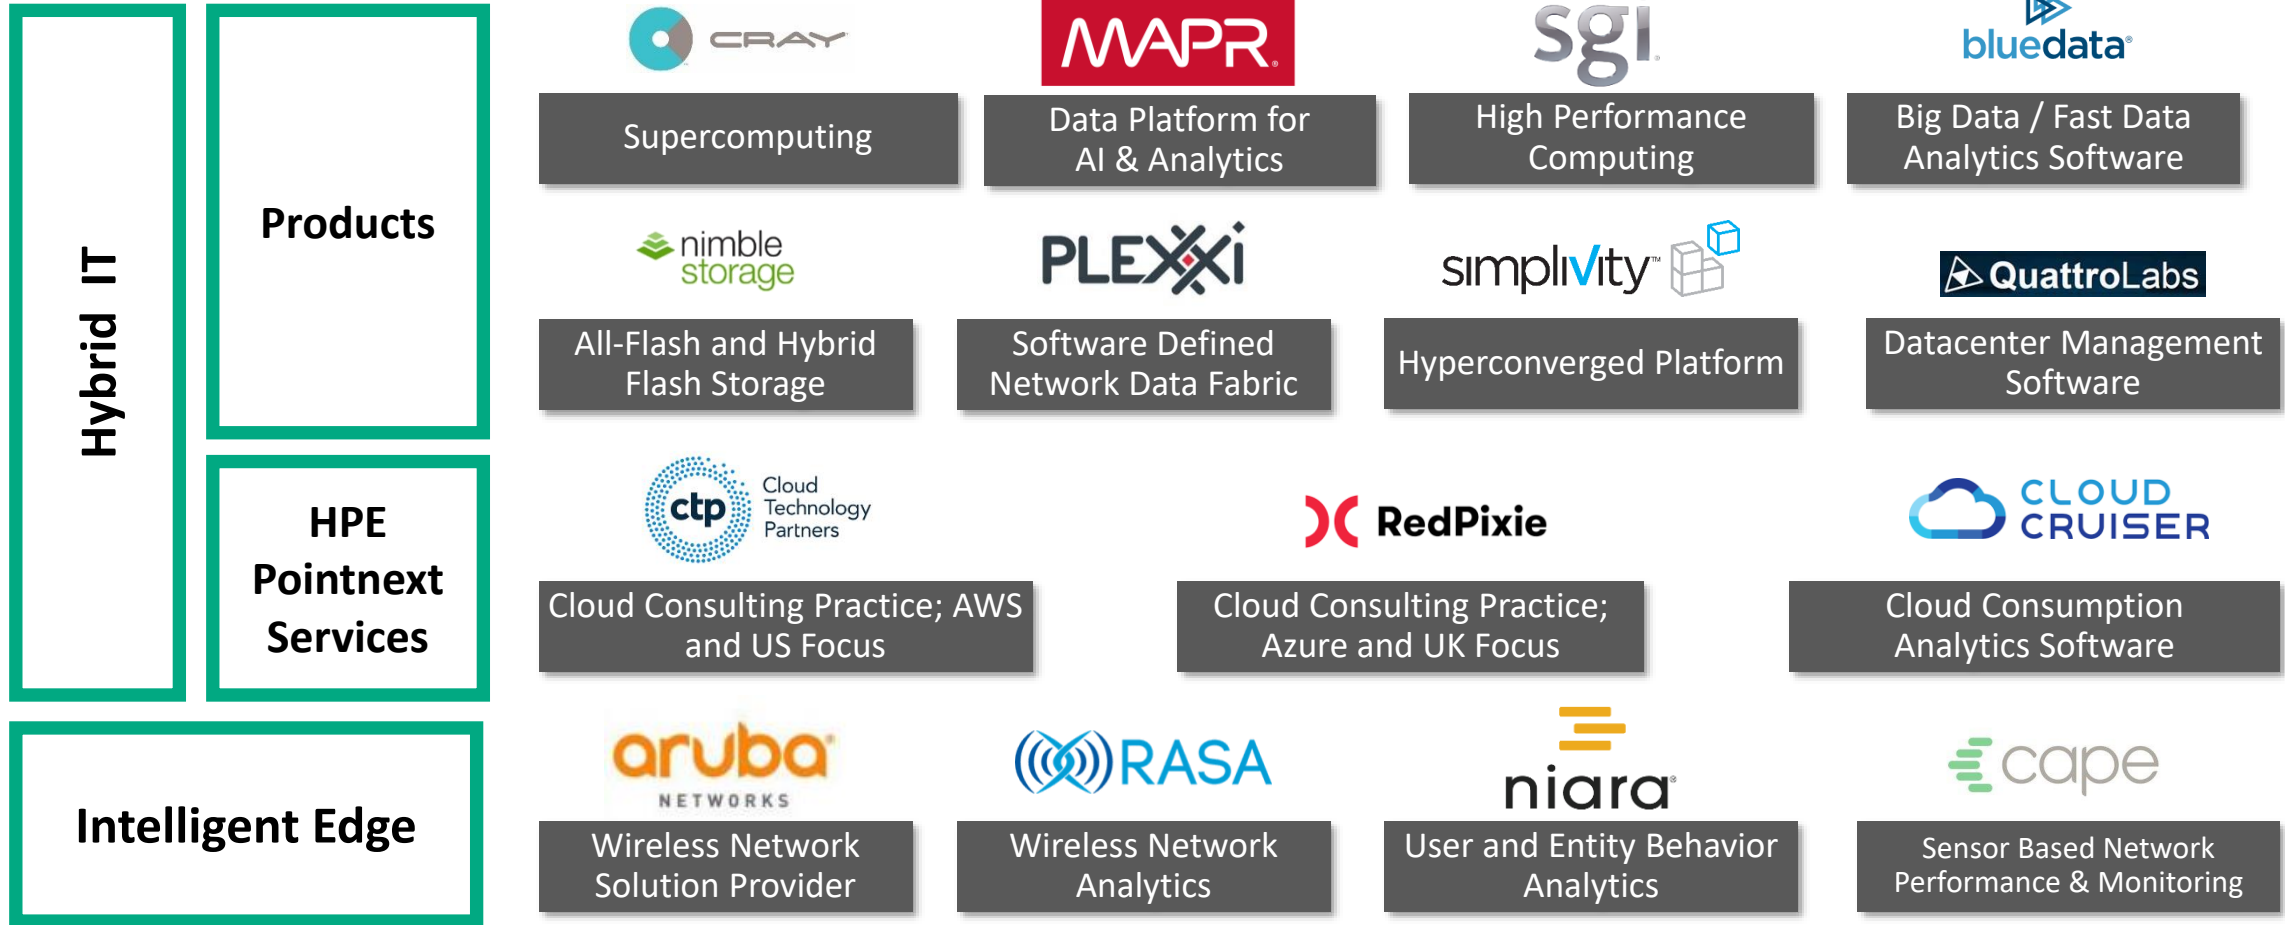


AI AND ANALYTICS DATA PIPELINE

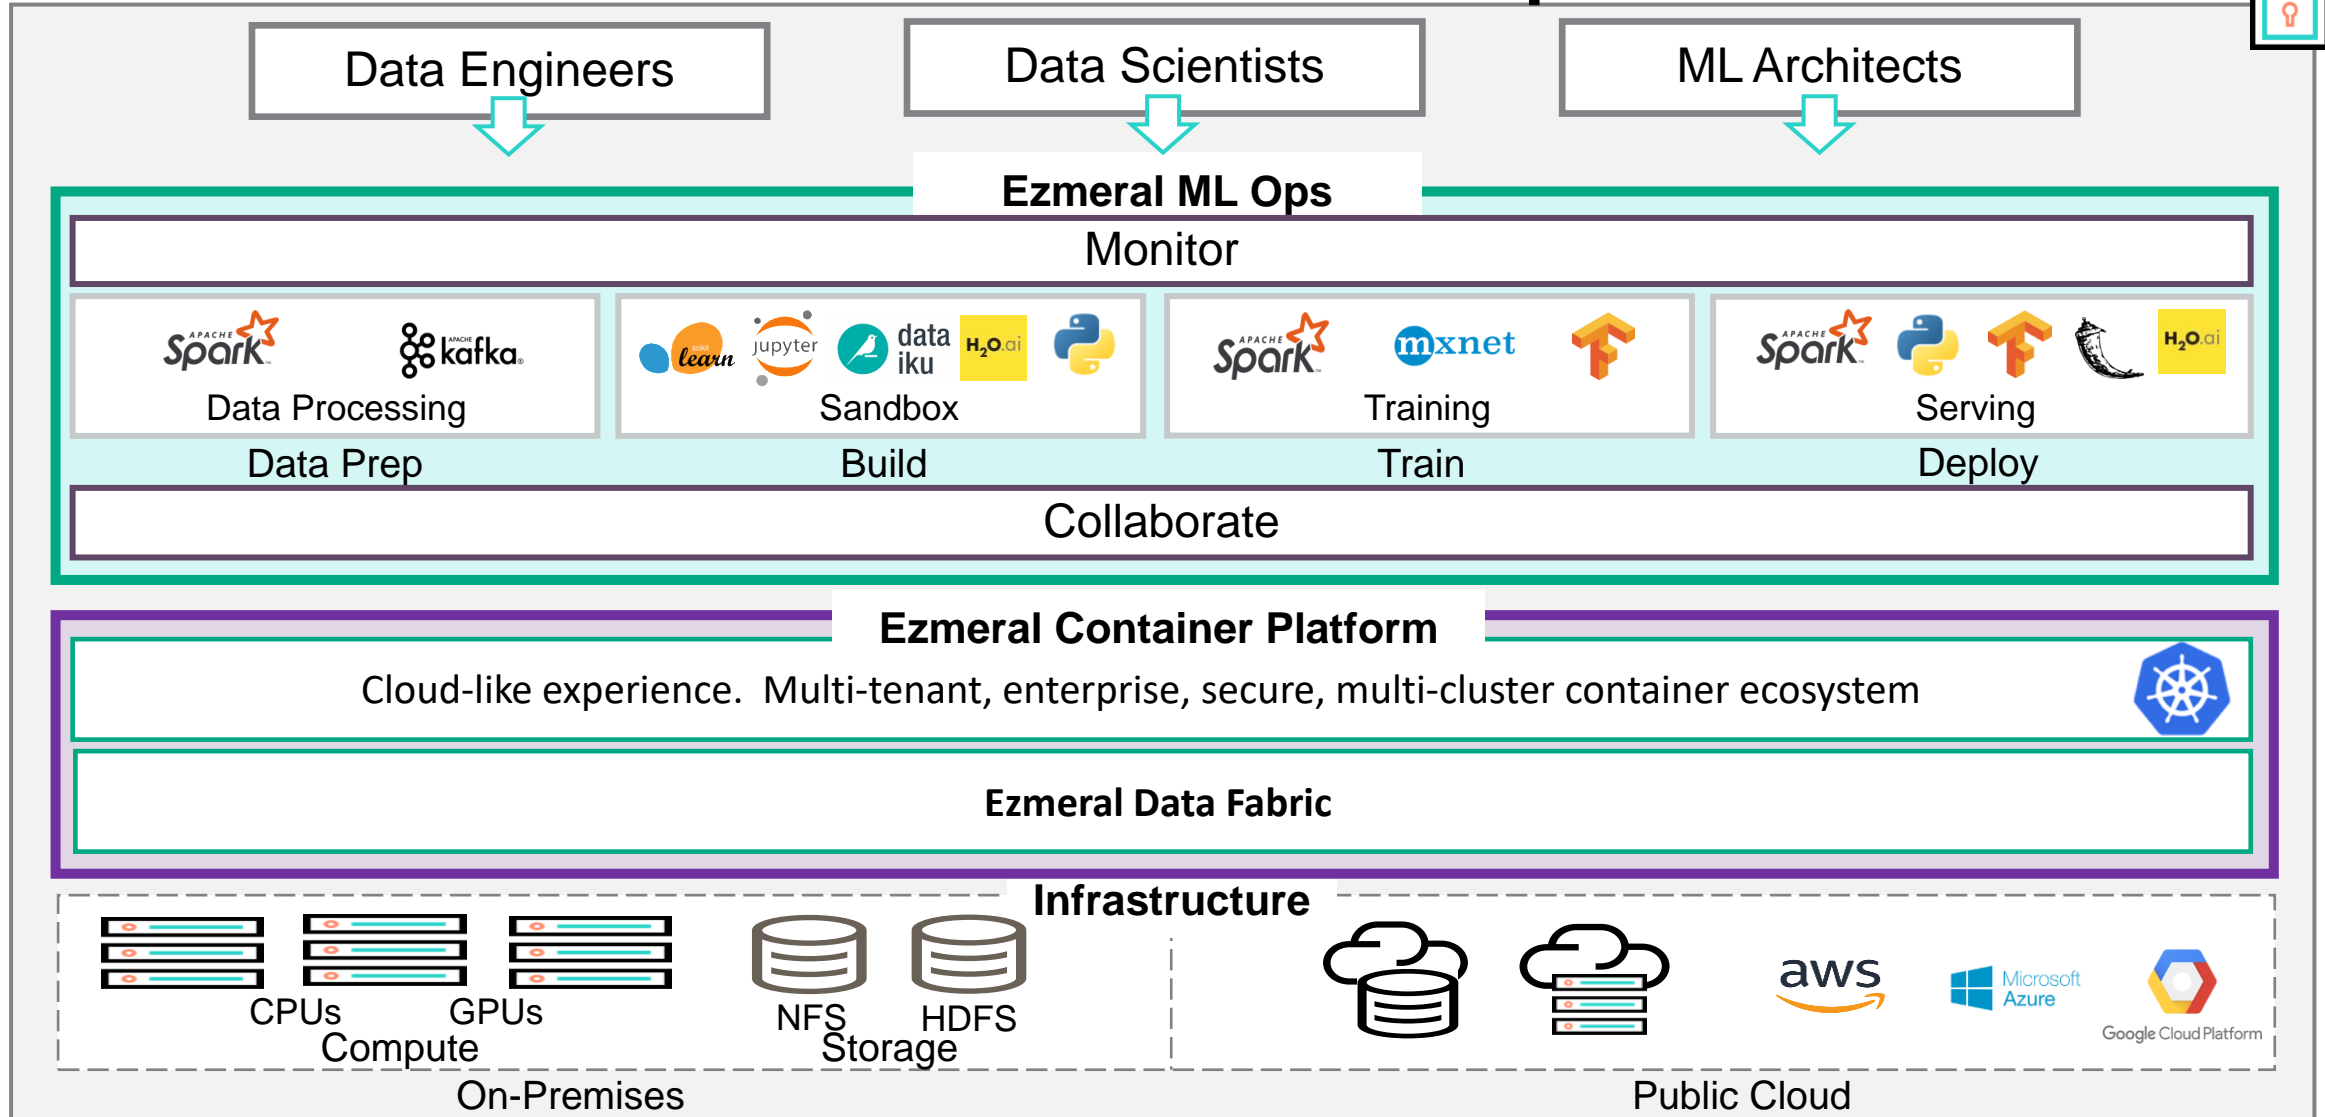
A simple view



Successfully accelerating strategy and enhancing competitive position through acquisitions



HPE Ezmeral Container Platform and ML Ops Architecture





Hewlett Packard
Enterprise

HPE IoT / Edge Strategy

Growing Family of Solutions

The Edge is a spectrum of use cases – Versatility is key

* OT = Operational Technology



Manufacturing

Legacy Interworking/PLC bridging, Asset Management, Condition Monitoring/Predictive, Augmented Reality, Edge Video Analytics



Energy & Utilities

Asset Management, Condition Monitoring/Predictive, Augmented Reality, AI/Deep-Learning




Transportation & Automotive

Connected Vehicles, Asset Management, Condition Monitoring/Predictive, Edge Video Analytics



Campus, Branch, Retail

Edge Workspace Appliance for Workstation/Desktop/App delivery, Edge Video Analytics (Campus Surveillance/CCTV)



Oil & Gas

Asset Management, Condition Monitoring/Predictive, Augmented Reality, Edge Video Analytics, AI/Deep-Learning



Defense & Intelligence

Connected Battlefield, Signals Intelligence, Edge Video Analytics, AI/Deep-Learning



Distributed Telco, Media, Comms

Multi-Access Edge Computing (MEC), 5G, vRAN, vCPE, Edge CDN/Caching, Video Transcoding

HPE created a new product space: “Converged Edge Systems”

HPE Edgeline Converged Edge Systems

3 Points of Convergence all in one box: Edge compute, Converged OT*, Remote systems management



HPE Edgeline EL300



HPE Edgeline EL1000



HPE Edgeline EL4000



HPE Edgeline EL8000

Unprecedented Edge Compute

Powerful open-standards performance



- Up to 112 Intel Xeon cores
- Up to 6TB Memory
- GPU, FPGA and VPU Accel
- Up to 96 TB SSD storage
- 100GbE & IB networks

Unique integration of Operational Technology (OT)

Data Acquisition, Control Systems



HPE OTLink HW & SW

Data center-class security, reliability & systems management

Remote monitoring and maintenance over wired or wireless networks



HPE Integrated Lights Out (iLO) and edge-optimized Integrated System Manager (iSM)

Engineered for the harsh edge environment:

Compact, low energy, rugged – shock, vibration, temperature, various mounting options



EdgeLine 8000 ~ the new standard for the Intelligent Edge



Data Center on the Edge

EL8000 – SWaP – Size Weight and Power

5U Chassis



Small Enough To Take With You



Up to four e910 Blades – That's up to 112 Cascade Lake cores, 6TB of DDR4 memory, 56TB of NVME flash storage and HPE Enterprise Server management in a carry-on case!



**Hewlett Packard
Enterprise**

HPE AI Servers

Growing Family of AI/ML/DL Solutions

Flexible choices for Big Data Analytics and AI

HPE Elastic Platform for Analytics

ISVs



cloudera



H₂O.ai



System Management



HPE Performance Cluster Manager



HPE OneView



HPE Ezmeral

EPA Building Blocks

Workload-optimized compute



ProLiant DL380



Apollo 2000



Apollo 6500



Superdome Flex

High speed networking

HPE FlexFabric

Plexxi

InfiniBand

Aruba

Tiered storage



Apollo 4200



Apollo 4510

Consumption models



HPE GreenLake Big Data



HPE GreenLake Flex Capacity



Capital Expense



HPE Apollo 6500 Gen10 System: An elegant, powerful solution

The maximum performance and efficiency for GPU computing from HPE



Maximum Power in an Enterprise Server


Traditional rack style server
Highly leveraged design based on DL380 Gen10
Integrated power supplies

NVLink 2.0 via SXM-2 GPU Module

Maximum GPU:GPU bandwidth for Deep Learning
Up to 8 HPE NVIDIA Tesla V100 for almost 1 PetaFlop HP

Traditional GPU support via PCIe Module

Supports your tested and qualified GPU
Configurable topologies



HPE SUPERDOME FLEX

Power critical apps, accelerate data analytics, tackle HPC & AI workloads holistically

Outpace evolving data demands with optimum flexibility


- Unique modular scale-up architecture utilizing 4-socket building block
- Designed for environments of all sizes – never outgrow, no over-provisioning
- HPE OneView management, Redfish APIs

Process and analyze growing data at extraordinary speed

- 4–32 sockets Intel® Xeon™ Scalable Processors as a single system
- 768GB–48TB shared memory for in-memory computing
- Unbounded I/O to support a wide variety of workloads

Safeguard mission-critical workloads

- Extreme Superdome RAS to achieve highest service levels
- HPE Serviceguard for Linux for high availability and disaster recovery
- HPE Pointnext Services and expertise to lower risk



Designed with Memory-Driven Computing Principles



**Hewlett Packard
Enterprise**

All the Best Accelerator Partners

GPUs, FPGAs, VPUs, ASICs

HPE and NVIDIA



Edge to core accelerated compute

Flexible choices using Industry-leading NVIDIA accelerators on HPE



NVIDIA® Tesla® V100 GPU



NVIDIA® Tesla® P100 GPU



NVIDIA® Tesla® P4 GPU



NVIDIA® T4 GPU



ProLiant DL380



ProLiant ML350



Apollo 6500



Edgeline 4000

Improved productivity

using integrated, and validated software / solutions



NGC-Ready Program¹
NGC Support Service²



RTX Server Solutions³



Better resource Utilization

Using efficient shared resources and tools from HPE and NVIDIA



[HPE Server-NVIDIA matrix](#)



[HPE and NVIDIA GPU QuickSpecs](#)

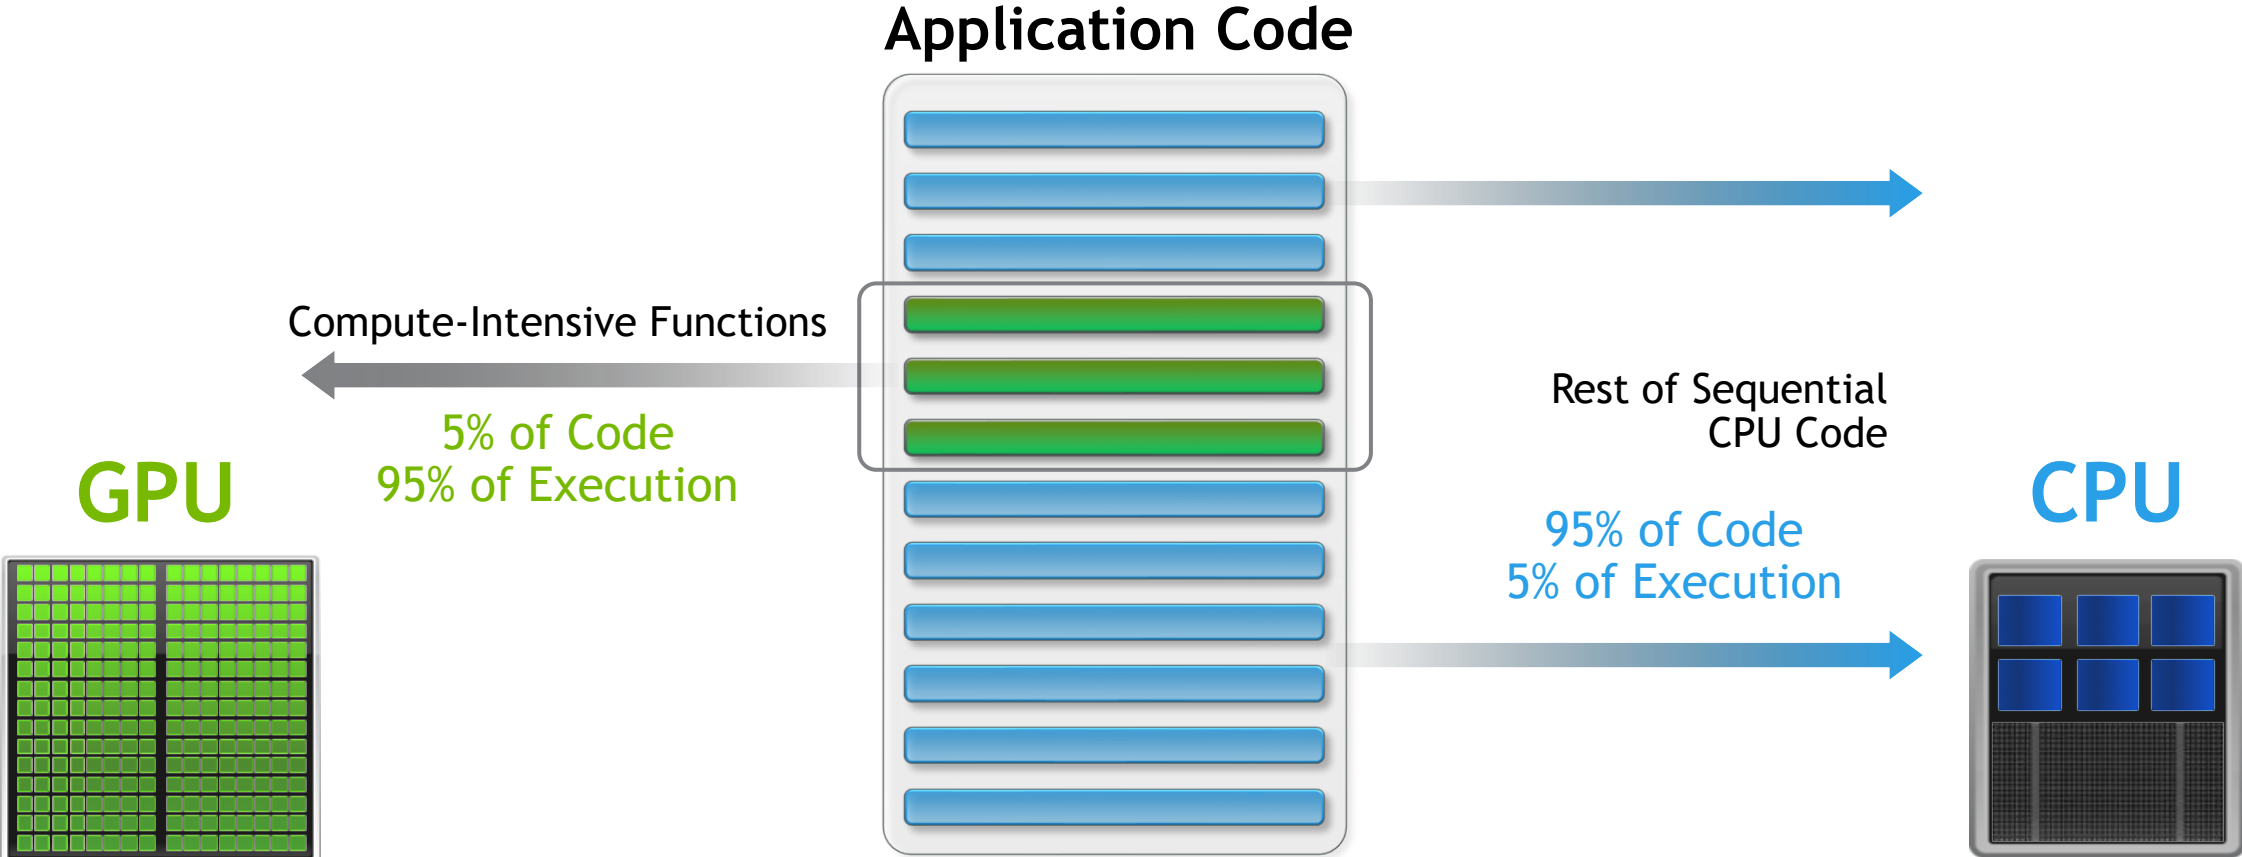
Deep Learning Cookbook:

- Deep Learning Benchmarking Suite
- Deep Learning Performance Guide



DEEP LEARNING INSTITUTE


How GPU Acceleration Works





**Hewlett Packard
Enterprise**

US FEDERAL AI READY CONFIGURATIONS



THE AI IMPLEMENTATION CHALLENGE

Data Scientist

Challenges

- Latest tools, models and techniques
- Accelerate model development time
- Rapid access to necessary resources
- Deliver the promise of AI




Want:

- Easy access to the latest data sciences tools (GPUs, Frameworks, open data science tools, etc.)
- Access to a model development/training infrastructure

IT Manager/Admin

Challenges

- Get control of AI projects across the organization
- Maximize utilization of resources
- Accelerate Speed of Deployment and time to value



Need:

- Assure environment is maintained to IT standards
- Leverage existing vendor relationships
- Best practices design, deployment & support model
- Ability to scale

What do business and project leaders want?

The sooner we start, the sooner we see results

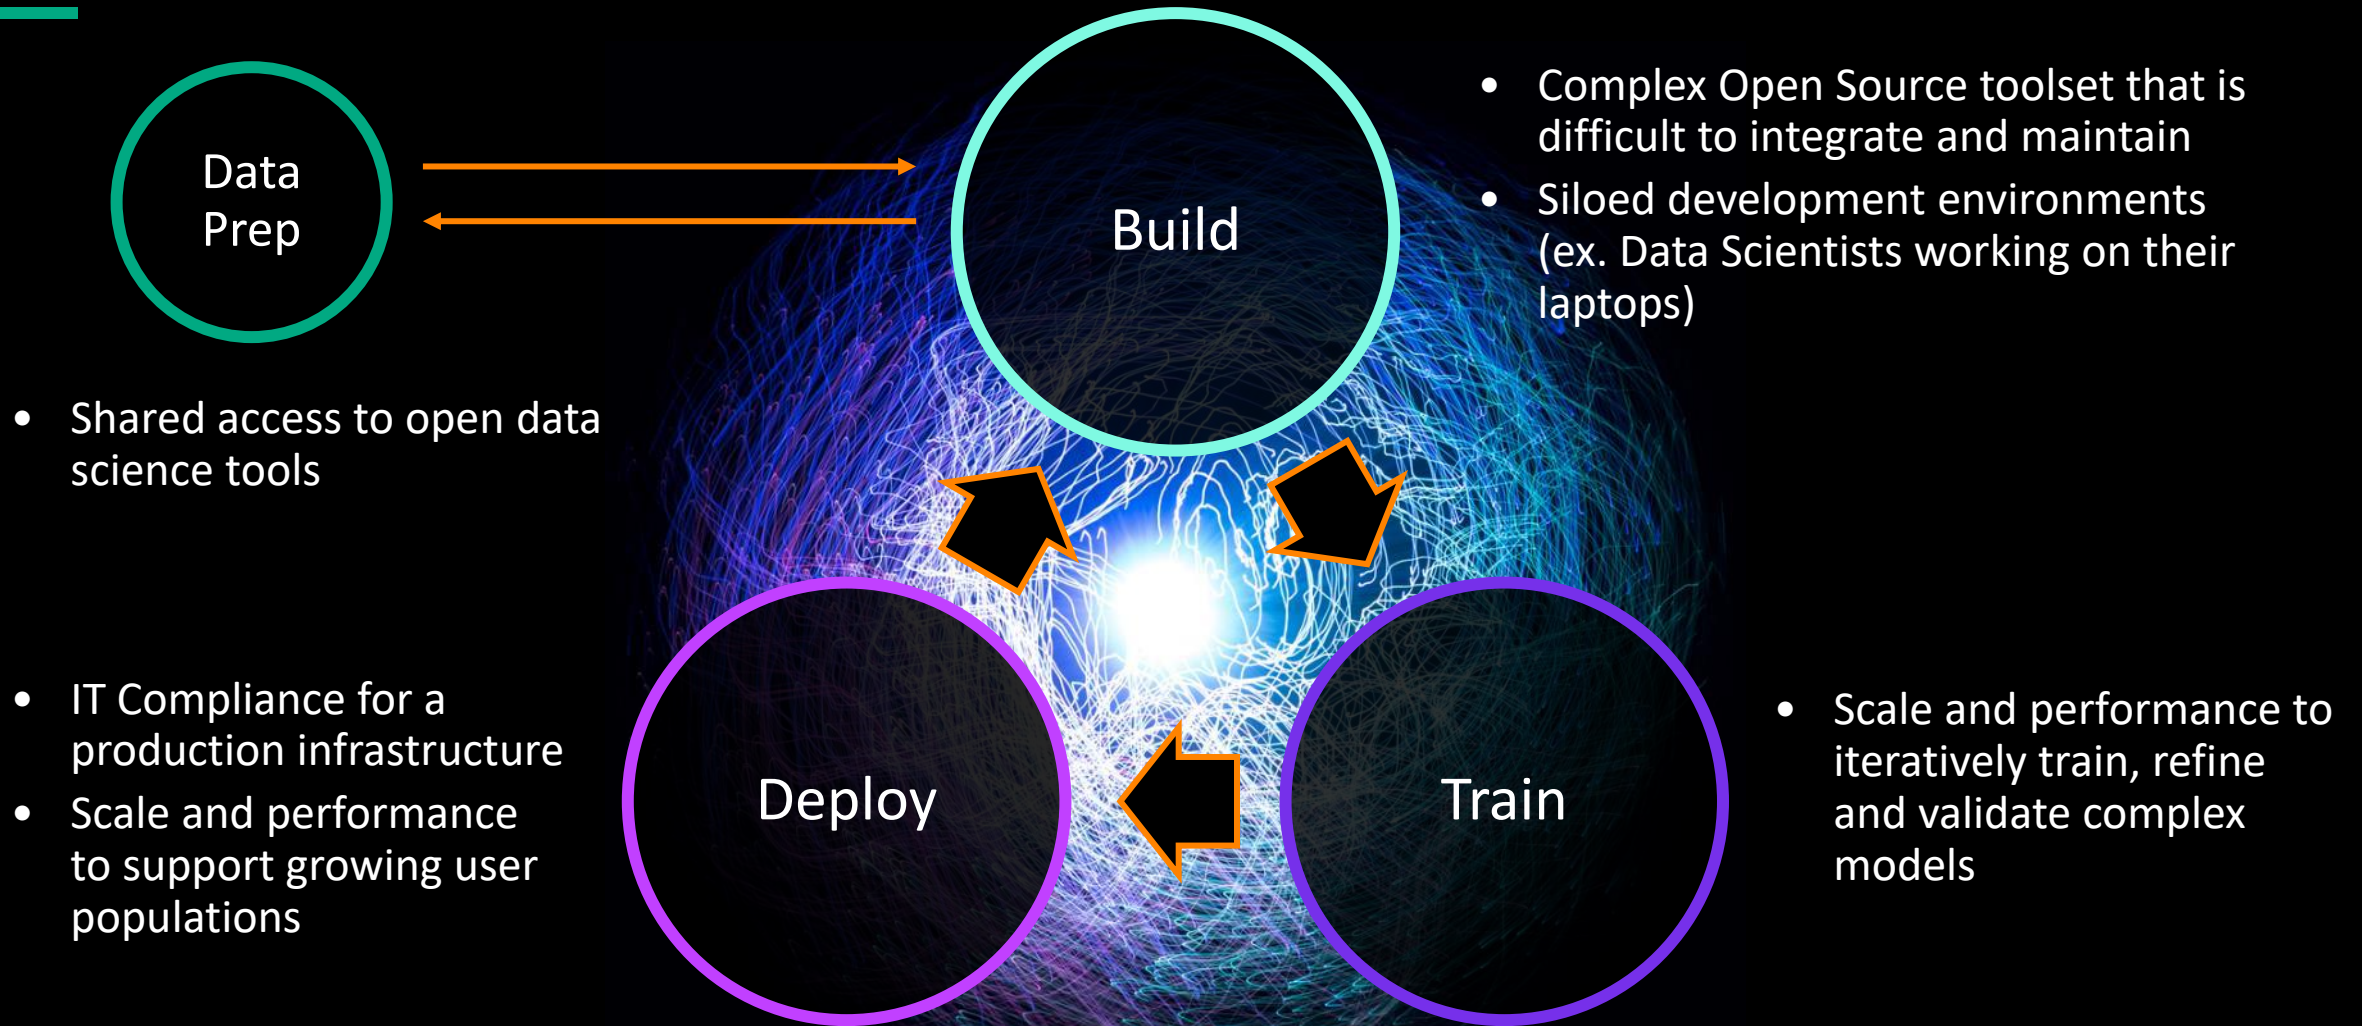
Federal mandate¹ to bring AI to the Agency. ...but this is uncharted territory. Needs:

- Satisfies both Data Scientist and IT Team
- A solution, not a kit
- Doesn't break the bank

Federal AI Bundles

COMPLETE: SW, HW, & Support
HPE / NVIDIA Certified Platform
Perfect for getting started
Aggressive Pricing

AI MODEL LIFECYCLE



APPROACHES TO AI INFRASTRUCTURE

Do It Yourself approach to AI Infrastructure

- Open Source Frameworks and Tool Sets
- Accelerated GPU Performance
- **Complexity introduced with never-ending testing and integration**
- **Community Support for the software environment**
- **Partial support for the whole system**

Fixed approach to AI Infrastructure

- Open Source Frameworks and Tool Sets
- Accelerated GPU Performance
- Integrated AI solution
- **System incompatible with IT standards**
- **Consistent support and services for entire AI infrastructure**

HPE approach to AI Infrastructure

- Open Source Frameworks and Tool Sets
- Accelerated GPU Performance
- Integrated AI solution
- IT Compliant Systems
- Services and Support
- Enterprise Scale

THE HPE APPROACH TO AI INFRASTRUCTURE

Integrated AI Solution

- Built for machine learning/deep learning workloads
- Featuring NVIDIA® Tensor Core GPUs
- HPE performance reliability, manageability, scalability and security
- Built to easily integrate into your IT Infrastructure

The Right Software for your AI Environment


- Red Hat Enterprise Linux to meet your corporate standards
- NVIDIA GPU-optimized software for DL, ML and HPC Workloads
- Delivered as containers to eliminate deployment complexity

Services for Long-term Success

- Deployment and Installation for a fast start
- Container installation and deployment service for assured operation
- HPE Pointnext Services AI specialists are there to help you choose your right AI path

AI READY CONFIGURATIONS

Crafted to meet the compute demands of the AI model lifecycle and meet the needs of Data Science and IT teams



Training

HPE Apollo 6500 Gen10 Server for GPU-intensive machine learning workloads with eight high-performance NVIDIA V100 Tensor Core GPUs.



Inference

HPE ProLiant DL380 Gen10 Server for dedicated AI and inferencing with two NVIDIA T4 GPUs (expandable to 7).

AI SOLUTIONS FOR TRAINING & INFERENCE

AI Training Solution

Built for **machine learning/deep learning** workloads delivering extreme compute for AI model training.

HPE Apollo 6500 Gen10 Server



NVIDIA V100 GPU



- NGC-Ready
- Two Intel® Xeon®-Gold 6248 (2.45GHz/20-core/150W)
- Eight NVIDIA® V100 32GB SXM2 GPUs
- NVIDIA® NVLink™ 2.0 GPU Cross-connect
- 768 GB RAM
- 7.6 TB Storage
- Two 10/40 GB Ethernet Ports
- Four InfiniBand EDR/Ethernet 100 GB Ports
- HPE iLO Advanced Pack & Silicon Root of Trust
- Three-year Proactive Care Next Business Day
- Red Hat® Enterprise Linux® License

AI Inference Solution

Built for **AI inferencing and light training** workloads and optimized for high performance results.

HPE ProLiant DL380 Gen10 Server



NVIDIA T4 GPU



- NGC-Ready
- Two Intel® Xeon®-Gold 6254 (3.1 GHz/18-core/200W)
- Two NVIDIA® T4 GPUs
- 384 GB RAM
- 2.4 TB Storage (SAS SSD)
- One 10/25 GB Ethernet Port
- HPE iLO Advanced Pack & Silicon Root of Trust
- Three-year Proactive Care Next Business Day
- Red Hat® Enterprise Linux® License

NVIDIA GPU CLOUD (NGC)

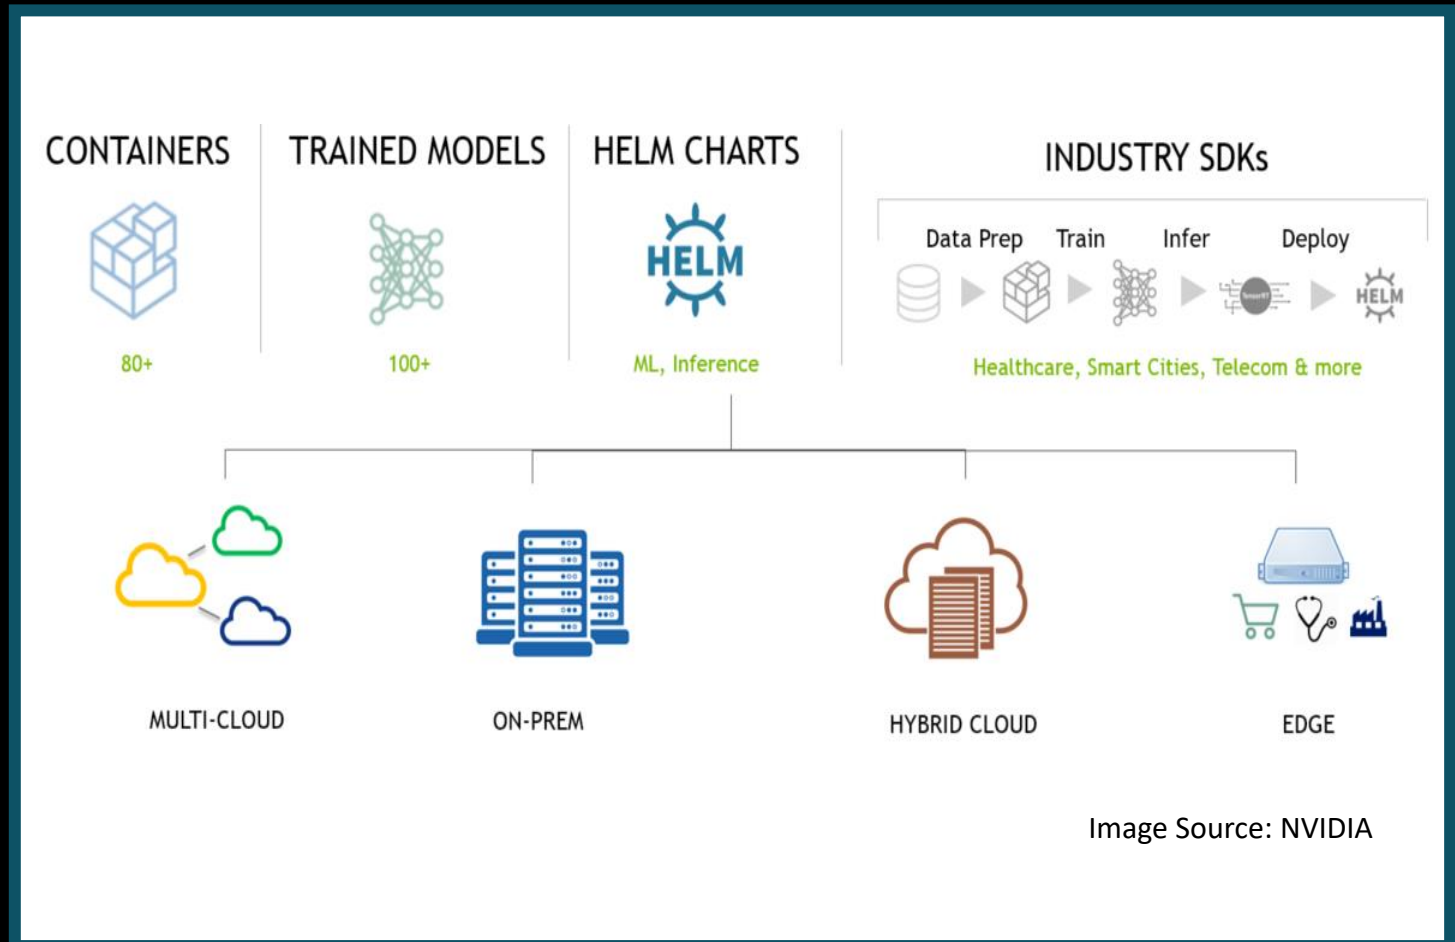
NGC is the GPU-Optimized Software Hub Simplifying DL, ML and HPC Workloads

NGC simplifies and accelerates

- Removes all the DIY complexity of managing dependencies, updates
- Tuned and tested software at no additional cost
- Same containers across on-premises servers, public cloud and the desktop


Containers, Models, SDKs

- RAPIDS End-to-end Data Science
- High Performance Computing (HPC)
- Deep Learning / Machine Learning
- Inference
- HPC Visualization
- GPU Infrastructure




AI READY CONFIGURATIONS - NATIONWIDE

Training



HPE Apollo 6500 Gen10 Server for GPU-intensive machine learning workloads with eight high-performance NVIDIA V100 Tensor Core GPUs

Inference



HPE ProLiant DL380 Gen10 Server for dedicated AI and inferencing with two NVIDIA T4 GPUs (expandable to 7)

	Training	Inference
NGC Ready	✓	✓
Speed and Performance Flexibility Peace of Mind Lower TCO	✓	✓
NGC Deployment and Integration Services	OPT	OPT

NVIDIA GPU Cloud (NGC) Ready

- Community GPU-optimized software containing validated, enterprise-grade containers for AI and HPC applications, pre-trained models, model scripts, helm charts and industry SDKs

NVIDIA GPU Cloud (NGC) Support Services (Optional)

- Direct NVIDIA support for NGC containers with expert advice available from Subject Matter Experts


NGC Deployment and Integration Services (Optional)

- Pointnext Services to install, configure and validate new GPU-enabled, NGC-ready system
- Valuable knowledge transfer to quickly take ownership of your new HPE NVIDIA solution




AI READY CONFIGURATIONS – FEDERAL BUNDLES

Training



HPE Apollo 6500 Gen10 Server for GPU-intensive machine learning workloads with eight high-performance NVIDIA V100 Tensor Core GPUs

Inference



HPE ProLiant DL380 Gen10 Server for dedicated AI and inferencing with two NVIDIA T4 GPUs (expandable to 7)

	Training	Inference
NGC Ready	✓	✓
Speed and Performance Flexibility Peace of Mind Lower TCO	✓	✓
NGC Deployment and Integration Services	✓	✓

NVIDIA GPU Cloud (NGC) Ready

- Community GPU-optimized software containing validated, enterprise-grade containers for AI and HPC applications, pre-trained models, model scripts, helm charts and industry SDKs

NVIDIA GPU Cloud (NGC) Support Services ~~(Optional)~~

INCLUDED

- Direct NVIDIA support for NGC containers with expert advice available from Subject Matter Experts

NGC Deployment and Integration Services ~~(Optional)~~

INCLUDED

- Pointnext Services to install, configure and validate new GPU-enabled, NGC-ready system
- Valuable knowledge transfer to quickly take ownership of your new HPE NVIDIA solution



NGC SUPPORT SERVICES

HPE NGC Certified Platforms



HPE Apollo 6500
Gen10



HPE Apollo 2000
Gen10



HPE ProLiant DL380
Gen10



HPE Edgeline EL4000



HPE Edgeline EL8000

Without NGC Support service contract:

- Customer service requests on NVIDIA containers are routed to NVIDIA community-based support option, DevTalk
- No SLA – best effort response
- No Standard service definition
- No access to SME



With NGC Support service contract:

- End-customers have a direct support contract directly with NVIDIA
- Customer requests service on NVIDIA containers by logging a case that flows directly to NVIDIA support
- Clearly defined support definition and escalation path
- Direct access to container SME

NGC DEPLOYMENT AND INTEGRATION SERVICES

Services include:

- **Service planning** – a detailed plan and schedule for AI Ready GPU-enabled infrastructure.
- **Infrastructure configuration and integration** – integrates the HPE server infrastructure, with NVIDIA's GPU accelerators, drivers, and container repository
- **NGC Configuration, integration and testing** – service includes configuration of the NGC service, deployment of a predefined set of containers, and demonstration and documentation of the container utilizing the GPU accelerator.
- **Knowledge transfer** – session to facilitate a transfer of operational control to the Customer



COMPLETE SOLUTION – FEDERAL BUNDLE ONLY

HPE GPU
Accelerated Platforms



NVIDIA NGC





NGC Support Services



Pointnext Services Deployment
and Implementation Services



Deep Learning
Institute Training
Credits




NVIDIA DEEP LEARNING INSTITUTE (DLI) TRAINING CREDITS

Training Bundle includes:

Two 8-hour★self-paced course credits
Four 2-hour self-paced course credits

Inference Bundle includes:

One 8-hour★self-paced course credit
Two 2-hour self-paced course credits



DEEP LEARNING FUNDAMENTALS

Fundamentals of Deep Learning for Computer Vision

Explore the fundamentals of deep learning by training neural networks and using results to improve performance ...

Optimization & Deployment of TensorFlow Models with TensorRT

Learn how to optimize TensorFlow models to generate fast inference engines in the deployment ...

AI Workflows for Intelligent Video Analytics with DeepStream

Learn how to build hardware-accelerated applications for intelligent video analytics ...

ACCELERATED COMPUTING FUNDAMENTALS

Fundamentals of Accelerated Computing with CUDA C/C++

Discover how C/C++ CPU-only applications to leverage the power of ...

Over 20 online courses ...

FEDERAL AI BUNDLES

SIMPLIFIED Federal AI Bundle End-User Price Estimator. Use only for conversational purposes!

Choose	?	QTY	Part Number	Fed AI Bundle Component TRAINING A6500	Unit Price	List \$	Discount**	Street Price*	BD Indirect	BD Direct	UCID
Mandatory	1	1	See BD#	Apollo 6500 Training Bundle	NA	NA	NA	\$ 100,000	97087209	97090717	5115528831-12
Optional	0	8	R2Y22AAEADH	NGC v100 1 GPU 1yr Sup E-LTU	\$ 1,999	\$ 15,992	0%	\$ 15,992	NVIDIA GPU Cloud software support. Priced per GPU.		
Pick 1	1	8	R2Y23AAEADH	NGC v100 1 GPU 3yr Sup E-LTU	\$ 4,799	\$ 38,392	0%	\$ 38,392			
Optional	0	1	HA113A1#5VW	HPE Install HPE Apollo 6500 Gen10 Service	\$ 600	\$ 600	0%	\$ 600	Physical install only		
Pick 1	1	1	HA114A1#5VW	HPE Install and Startup HPE Apollo 6500 Gen10	\$ 1,900	\$ 1,900	0%	\$ 1,900	Physical install and OS install		
Optional	1	90	HL3U9A1	Onsite NGC Deployment	\$ 100	\$ 9,000	0%	\$ 9,000	Turnkey planning, installation, and integration of NGC software		
Pick 1***	0	60	HL3U9A1	Remote NGC Deployment	\$ 100	\$ 6,000	0%	\$ 6,000			



TOTAL> \$ 149,292

Choose	?	QTY	Part Number	Fed AI Bundle Component INFERENCE DL380	Unit Price	List \$	Discount**	Street Price*	BD Indirect	BD Direct	UCID
Mandatory	1	1	See BD#	ProLiant DL380 Inference	NA	NA	NA	\$ 30,000	97492004	97491722	5117552688-04
Optional	0	2	R2Y25AAEADH	NGC T4 1 GPU 1yr Sup E-LTU	\$ 1,999	\$ 3,998	0%	\$ 3,998	NVIDIA GPU Cloud software support. Priced per GPU.		
Pick 1	1	2	R2Y26AAEADH	NGC T4 1 GPU 3yr Sup E-LTU	\$ 4,799	\$ 9,598	0%	\$ 9,598			
Optional	0	1	HA113A1#5A6	HPE Install ProLiant DL38x(p) Service	\$ 375	\$ 375	0%	\$ 375	Physical install only		
Pick 1	1	1	HA114A1#5A6	HPE Installation and Startup DL38x(p) Service	\$ 1,150	\$ 1,150	0%	\$ 1,150	Physical install and OS install		
Optional	1	90	HL3U9A1	Onsite NGC Deployment	\$ 100	\$ 9,000	0%	\$ 9,000	Turnkey planning, installation, and integration of NGC software		
Pick 1***	0	60	HL3U9A1	Remote NGC Deployment	\$ 100	\$ 6,000	0%	\$ 6,000			

TOTAL> \$ 49,748

* APPROXIMATE STREET PRICE

AI READY BUNDLE VALUE PROPOSITION



1 Shorter Time To Value for AI

Accelerated GPU-powered AI using powerful systems and software for machine and deep learning training and inference

2 Enterprise Ready

IT-compliant server meeting the performance, manageability, security and service standards of IT organizations for production-grade AI

3 Expertise

Faster time to proof-of-concept, production and ROI with services that provide rapid tangible outcomes based on best practices

BUNDLE SELLING RESOURCES - CUSTOMER

Solution Brief - HPE AI Training Solution – 2pg – Customer handout for the Training Bundle.

https://hpe.seismic.com/Link/Content/DCW2rdXVzNxUigRtCS_-hlpQ

Solution Brief - HPE AI-ML Inference Solution – 2pg – Customer handout for the Inference Bundle.

<https://hpe.seismic.com/Link/Content/DCIZGRAvOhMk266PO5aJnuWA>

Technical White Paper - Accelerating AI adoption for the modern data-driven enterprise – 6pg – HPE/NVIDIA Marketecture white paper for Customers.

https://hpe.seismic.com/Link/Content/DC4usDnBy_2k2zAsAhfgVnTg



BUNDLE SELLING RESOURCES – RESELLER & INTERNAL HPE

At A Glance - AI Bundles from HPE Made Easy – 1pg - Describes the partner advantages of the AI Bundle Program.
https://hpe.seismic.com/Link/Content/DCmW53enM4_kK2ZrBjfUqGbw

Competitive Battlecard - AI Solutions Bundles – 5pg – Elevator pitch, selling points, qualifying questions, competition. https://hpe.seismic.com/Link/Content/DCj_rAd8eVPE2-6ktKcHLTOg

FAQ - HPE Apollo and HPE ProLiant AI Solutions Bundles – 6pg – Frequently asked questions including how to order, the BOM, and Pointnext optional services.
<https://hpe.seismic.com/Link/Content/DC3coZg2avM0q1Vsyvo08AUQ>

Sales Summary Sheet - HPE AI Training Solution – 2pg – Internal selling document summarizing the Training Bundle (Apollo 6500). <https://hpe.seismic.com/Link/Content/DCati6lV6ksEGeZPjV6qNRmg>

Sales Summary Sheet - HPE AI-ML Inference Solution – 2pg – Internal selling document summarizing the Inference Bundle (ProLiant DL 380) <https://hpe.seismic.com/Link/Content/DCrqu7FqQMiEmOAGhEKHoS2A>



BUNDLE SELLING RESOURCES – RESELLER & INTERNAL HPE

ONLINE TRAINING

Training Materials - AI Ready Bundles for North American Sales and Partners – AI Bundle online training course. Requires partner login to HPE learning management system.

https://content.ext.hpe.com/sites/LMS/LMS.html?deeplink=/Saba/Web_spf/HPE/common/ledetail/0001132967

POWERPOINT PRESENTATIONS

Customer Presentation with Speaker Notes - AI Ready Configurations – 16 slides - PowerPoint Presentation. Read and follow Note to Presenters before sharing with customer.

<https://hpe.seismic.com/Link/Content/DC1sJNjJB-EKNaF-YwWMEIw>

Customer Presentation - Deployment and Integration Services for NVIDIA GPU CLOUD – 16 slides – PowerPoint slide describing Pointnext NGC installation and integration services.

<https://hpe.seismic.com/Link/Content/DC30JYFxYrQUm-zkwoSxSYFA>



GSA CENTERS OF EXCELLENCE

coe.gsa.gov



GSA COE SERVICES



Service Offerings

- AI Assessment Of Governance & Enablers
- Use Case Discovery & Selection
- Workflow Mapping & Process Automation
- Lean Innovation Process Design
- AI Solutions Identification & Implementation

Service Catalog

- AI Maturity Assessment & Strategy Development
- Path-to-ATO Analysis
- Use Case Prioritization Models
- Rapid AI Solution Prototyping
- Training Data Sourcing
- AI Explainability & Traceability Processes
- AI Technology Roadmap

Four steps to your first win

1

Learn

2

Hunt

3

Ask
for help

4

Get the
first PO

Business development
manager

HPE Pointnext
Service Advisor

Pre-sales

AI ambassadors



Seize the opportunity now, before the competition!

Where can I learn more?

Explore sales briefcases supporting AI and Analytics

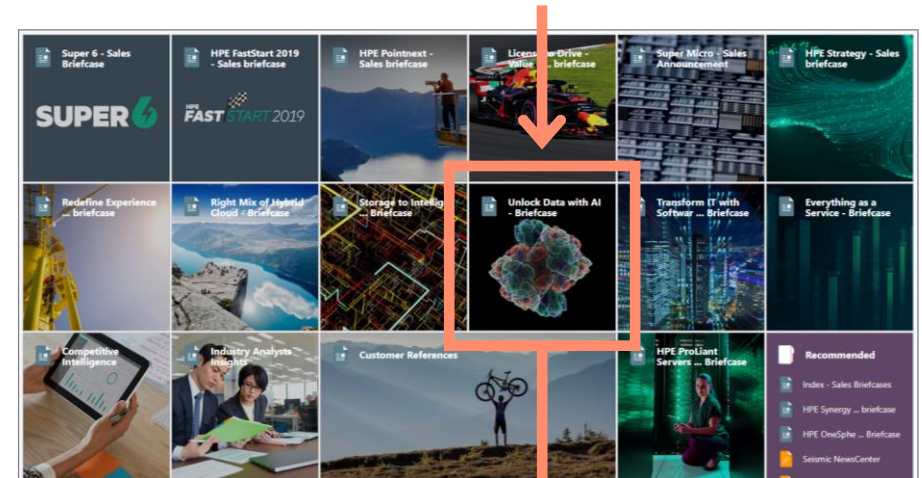
- [Unlock Data with AI Sales Briefcase](#)
- Customer decks, press releases, blogs
- Customer testimonials / case studies / FAQs
- Talking Points, Customer letters, Dummies Guides
- Links to tools, training and resources

Training resources for AI and advanced analytics


- HPE Pointnext Sales Trainings
- How to sell services for Artificial Intelligence, Data and Analytics [replay](#) & [slides](#).
- Selling HPE Digital Prescriptive Maintenance [replay](#) & [slides](#).
- Mission Critical Solutions Training Webinar for Generalists [replay](#) & [slides](#).

Big Data Analytics – Sales Materials

- Edge to Core data pipeline – [Customer Presentation](#)
- HPE EPA vs AWS S3 Storage – [Cost Analysis](#)
- Data Pipeline Use Cases
 - Retail – [Infographic](#)
 - Manufacturing – [Infographic](#)
 - Automotive – [Infographic](#)



UNLOCK DATA WITH AI - BRIEFCASE



Hewlett Packard
Enterprise

Unlock Data with AI

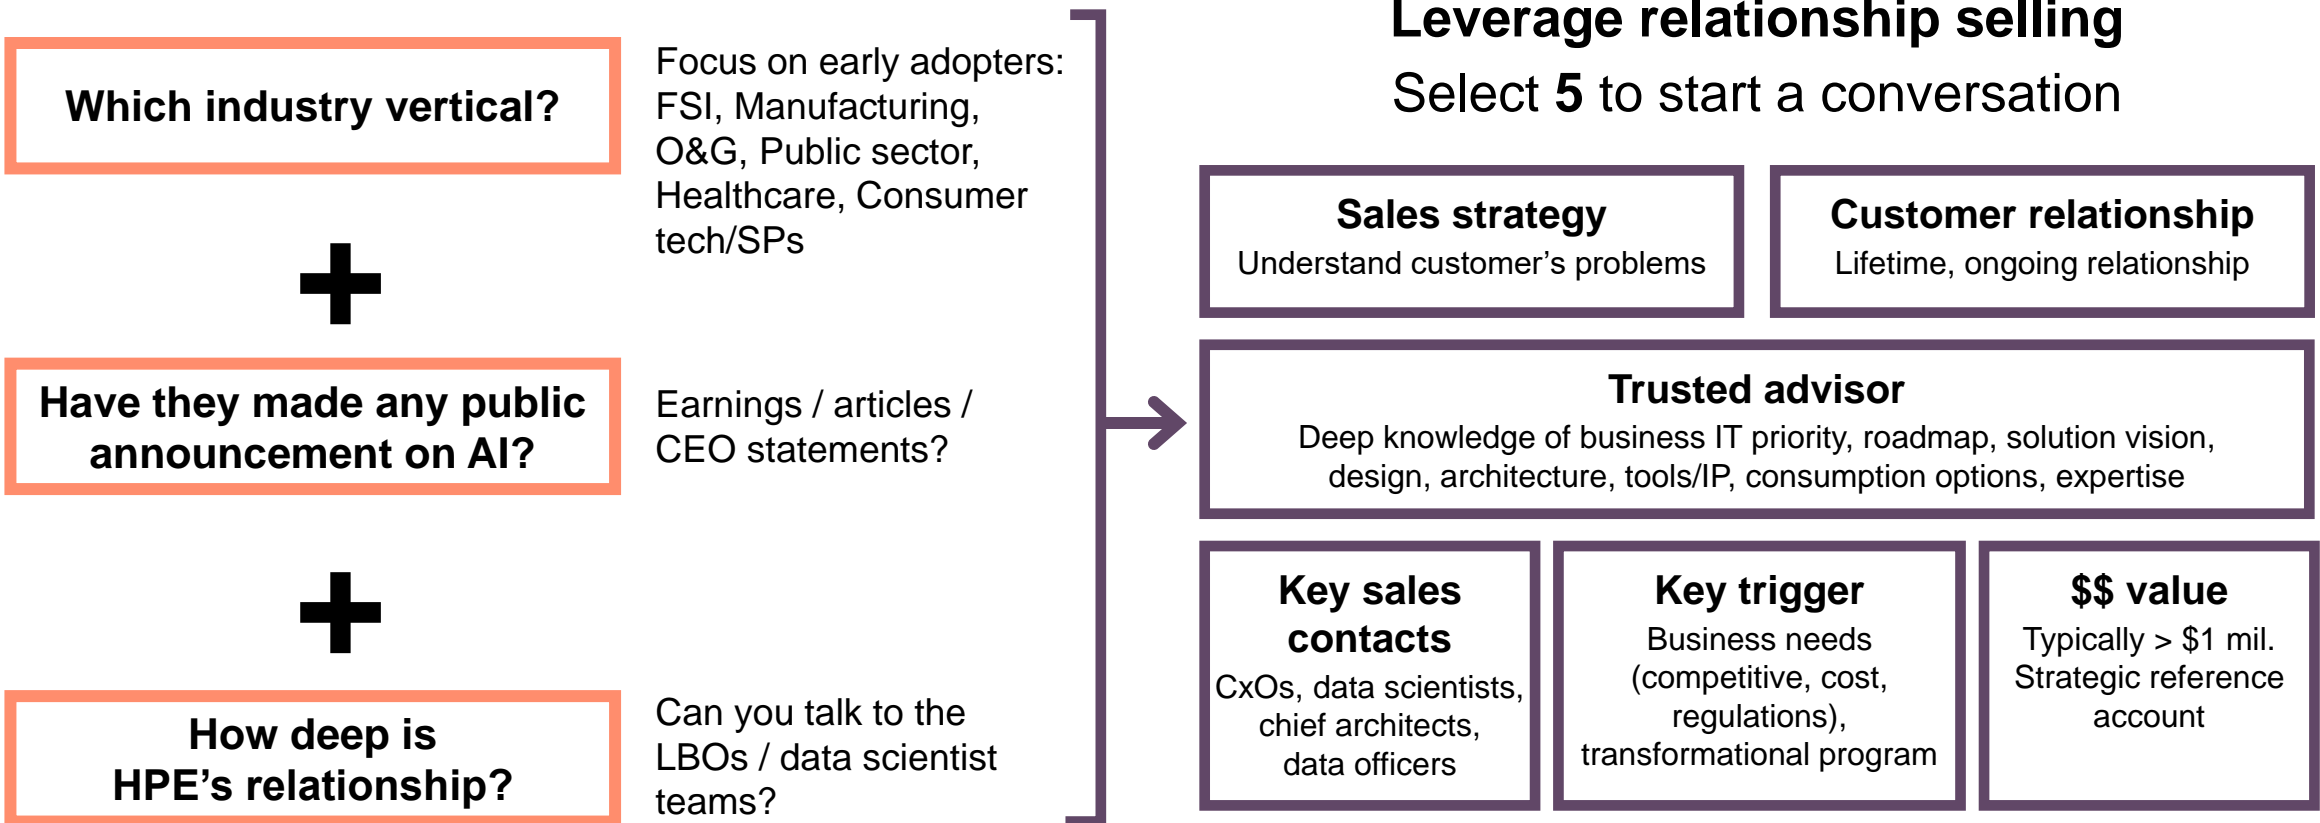
Sales and Partner Briefcase: **Unlock Data with Artificial Intelligence**

November 5, 2018

Data is the new currency, but to extract insights and drive the instantaneous actions required at the edge, we are investing in, partnering for and inventing new forms of computing to power an Artificial Intelligence (AI) driven world. Taking advantage of our industry leadership in both scale-out HPC architectures as well as large-memory architectures with our advanced research in Memory-Driven Computing and HPE Superdome Flex, HPE is at the forefront of what's next for data. Initiate customer conversations with AI and explore the potential to share HPE's portfolio and expertise unlocking the value of data, covering areas such as, HPE Pointnext Centers of Excellence (CoE), High Performance Computing, data lake architectures, edge computing, HPE Greenlake, hybrid cloud expertise, HPE Financial Services superchargers, and more.

Steps to identify your target customers

2



Target customers: Who to talk to

2



**CIO/Chief
Data Officer**

Questions / Challenges

- Edict for Business Transformation
- AI is the path to Digital transformation
- AI is a Board / CEO level discussion now – ‘Why AI’ moving to ‘Why not AI’

What are they looking for?

- ROI & Business Value on AI
- How/where to start?
- Defining the use case
- First to market with proven results
- Expand AI use in multiple domains



**Data
Scientist**

Questions / Challenges

- Latest models and techniques
- How to select the right model and use it
- Data management
- Deliver the promise of AI

What are they looking for?

- modeling and data management advice
- Quick access to training / inference architecture with proven benchmarks
- Infrastructure needed?



**IT Manager/
Admin**

Questions / Challenges

- ‘Get control’ of AI IT projects across the org
- Maintain infrastructure standards
- Accelerate Speed of Deployment
- Minimize downtime
- ‘Look smart’ in front of developers/data scientists

What are they looking for?

- Leverage existing vendor installations
- Automated management
- Expertise in design, deployment & support
- Ability to scale, boundless architecture
- Help on ‘Cloud vs. on-prem’ question



**ML
Architect/
Developer**

Questions / Challenges

- Accelerate Business Model Development
- Keep up with the Technology Frameworks

What are they looking for?

- Shortest path to learning & deployment
- Options to move “cloud deployed models” to on-premises
- Automated Cloud based offerings on a pay-per-use basis

Influencers: We need to find a way to talk to them

Help them look good, and execute fast

Effective hunting requires good listening...

2

Ask the question – Is your organization planning or working on any AI projects?

Deep learning or AI

Target audience:

AI Leads
CDO's
Data Scientists
Big Data Experts
ML Developers...

Listen for these key words and phrases

- Big Data v2.0 (Cloudera, Hadoop, MapR etc.)
- Data Pipeline
- Internet of Things (IoT)
- Deploy intelligence at the edge
- Automate translation, diagnosis
- Speech, character, facial recognition
- Virtual Digital Assistant
- Video Analytics, Computer Vision, Image Processing
- Predictive Maintenance
- Chatbots
- GPUs, FPGAs
- AI platforms (TensorFlow,...)

Traditional HPC

Target audience:

Data Scientists
R&D Teams
EDA Leads...

Listen for these key words and phrases

- Financial algorithms (Monte-Carlo etc.)
- Oil and Gas / Seismic analysis
- Manufacturing – complex product models and simulations
- EDA/CAE applications
- National Labs/Research
- Weather simulations
- Parallel File Systems
- FPGAs, GPUs


Questions to Ask



CIO/Chief Data Officer



Data Scientist



IT Manager/Admin



ML Architect/Developer

LEVEL 1

- Top data, analytics priorities for next 4 quarters?
- What does success look like for you?
- Business problem are you trying to solve?
- Currently using/planning to use AI / analytics tools to solve?
- Top Priorities for next 4 quarters? Where does AI / analytics fit into these priorities?
- What does success look like for you?
- Business problem are you trying to solve with AI / analytics?
- Challenges (if any) are you currently facing in your AI / analytics projects?

LEVEL 2

- Considering AI / analytics for any of your projects within the next 6 months? If YES, go to L3.
- If NO, why not? Can we share industry best practices?
- Current challenges (i.e. data quality/sources, training time / latency)?
- Managing / collecting any data on the edge?
- How and where are you deploying AI / analytics projects?
- Challenges are you facing today to deploy AI / analytics?
- Tools to deploy AI / analytics?
- How does infrastructure play a role in your AI / analytics deployment?
- Factors led you to deploy your AI / analytics projects in the Cloud (assuming it's running on Cloud)

LEVEL 3

- Defined business problem or identified a use case on AI / analytics?
- Biggest challenge(s) you're facing?
- Adequate tools/resources to solve your AI / analytics problems today?
- If Yes, What are they?
- Deploying any "accelerators" today? How are you using these?
- Concerned about Security in your AI / analytics deployment?
- Size of current data set?
- How long does it take to train your data set?
- Types of algorithms/neural network models are you currently using for AI?
- Adequate tools/resources to solve your AI / analytics problems today?

LEVEL 4

- What resources (people, tools, infra, services etc.) are you currently using/planning to use?
- Familiar with HPE's AI offerings? Yes/No – Can we set up some time to share?
- If yes, can HPE share some insights on how we could help you scale your AI / analytics ecosystem?
- Who do we need to engage to help you build a business case in addressing your challenge?
- Familiar with HPE's AI offerings? Yes/No – Can we set up some time to share our Solutions?
- If running in Cloud – Did you know that HPE "as-a-service" offerings could save you \$\$\$?
- Aware that HPE offers "as-a-service" model for AI workloads?

CTA: When can we set up the next conversation?

Reach out to an expert to go deeper

3

HPE Pointnext Contacts: AI & Analytics

APAC Weng Keng Chew chew@hpe.com	Europe North Nikolaj Mann nikolaj.mann@hpe.com	India Kulkarni, Abhish abhish.kulkarni@hpe.com	North America Eric Johnson (Rainier) e.johnson@hpe.com	France Laurence Morestin-Cadet laurence.morestin-cadet@hpe.com
DACH Christoph Huschke christoph.huschke@hpe.com	Europe South Carlos Sanchez-Largo carlos.sanchez-largo@hpe.com	Japan Rieko Nakano rieko.nakano@hpe.com	CEMA George Papaioannou gpa@hpe.com	UK&I Mark Tritton mark.tritton@hpe.com

WW HPE Pointnext AI, analytics practice (20+ experts)

coe-data-ai-core@hpe.com

Beena Ammanath
Luis Buezo

AI

APAC Steve Tolnai tolnai@hpe.com	Japan Ken Takahashi ken.takahashi@hpe.com	WW Pankaj Goyal pgoyal@hpe.com Seyed Mirsepassi seyed.mirsepassi@hpe.com Greg Schmidt Greg.schmidt@hpe.com
North America Steve Heibein steve.Heibein@hpe.com Manoj Suvarna manoj.suvarna@hpe.com	EMEA Edmondo Orlotti edmondo.orldotti@hpe.com	

HPC

APAC Laurent Herviou laurent.herviou@hpe.com	EMEA Lee Rand lee.rand@hpe.com	NA Tom Tyler thomas.tyler@hpe.com Manoj Suvarna manoj.suvarna@hpe.com Michael J Robinson Mj.robinson@hpe.com
WW Sammy Zimmerman sammyz@hpe.com	Japan Ken Takahashi ken.takahashi@hpe.com	

Reach out to an expert to go deeper (continued)

3

Big data

APAC

Nainika Taneja,
nainika.taneja@hpe.com

North America

Greg Cherry,
greg.cherry@hpe.com

CEMA

TBD

France

Fouad Bendris,
fouad.bendris@hpe.com

DACH

Andreas Bachmann,
andreas.bachmann@hpe.com

Aron Precht,
aron.precht@hpe.com

Dennis Geisse,
dennis.geisse@hpe.com

Europe North

Laurent van Veen,
laurent-van.veen@hpe.com

Europe South

Giorgio Raico,
giorgio.raico@hpe.com

UK&I

TBD

India

Gowri Shankar,
g.shankar@hpe.com

Japan

Yoji Inoue
yoji.inoue@hpe.com

WW

Don Wilson,
don.wilson@hpe.com

MCS

APAC

Wu Sheng Chen
wusheng.chen@hpe.com

Chang Hoon Lee
changhoon.lee@hpe.com

DACH

Uschi Becker-Haug
uschi.becker-haug@hpe.com

Heike Nill
heike.nill@hpe.com

North America

Nancy Davidson
nancy.davidson@hpe.com

CEEMA

Vassilis Garavelas
vassilis.garavelas@hpe.com

Europe North

Klaus Harder
klaus.harder@hpe.com

Chris van Maaren
chris.van.maaren@hpe.com

Europe South

Elena Sanchez (Spain)
elena.sanchez-prieto@hpe.com

Ehood Baratz (Israel)
ehood.baratz@hpe.com

France

Benoit Maillard
benoit.maillard@hpe.com

UK&I

Karl Graham
karl.graham@hpe.com

India

Sameer Nagar
sameer.nagar@hpe.com

Japan

Hajime Hino
hajime.hino@hpe.com

Getting to the first PO

Identify where the customer is in their journey

Stage 1 How to get started?

- Interested about AI
- Don't know how it can help their business
- No use cases
- Management discussions, but no action

Candidate for HPE AI Transformation Workshop (big data, advanced analytics)

Click [HERE](#) to download workshop brief
Contact: HPE Pointnext Services

Stage 2 Start a PoC or Proof of Value?

- Defined 1-2 use cases
- Has a team of developers/data scientists
- First POC being built or planned; IT may or may not be involved
- Commitment from management

Detailed architectural, use case discussions

Contact: HPE Pointnext for use case Point of View
Contact: HPE AI Ambassadors at artificial.intelligence@hpe.com
Use 'Opportunity Assessment Tool'

Stage 3 Scaling up?

- Early success
- Investment to scale up the solutions
- IT involved with data scientists
- Investment commitment from management

Ready for detailed RFPs

Contact: Account sales team and AI, HPC, BDA Pre-sales specialists and HPE Pointnext in your geo

Seize the opportunity and repeat!





Thank you

Follow the Data!

THANK YOU

Steve Heibein



AI-READY CONTAINER CONFIGURATIONS FROM HPE

	HPE Container Platform	NVIDIA GPU Cloud	Cray Urika	Do-It-Yourself
AI Content	<ul style="list-style-type: none"> HPE Container Platform App Store 	<ul style="list-style-type: none"> NVIDIA GPU Cloud (AI Frameworks, RAPIDS) 	<ul style="list-style-type: none"> Urika AI Suite 	<ul style="list-style-type: none"> AI Frameworks Bright Data Science (AI Frameworks)
Container Support	<ul style="list-style-type: none"> Docker Runtime 	<ul style="list-style-type: none"> Docker Runtime 	<ul style="list-style-type: none"> Docker Runtime Bare Metal 	<ul style="list-style-type: none"> Bare Metal Choice of Open Source Runtime
Job Scheduling & Orchestration	<ul style="list-style-type: none"> Kubernetes 	<ul style="list-style-type: none"> Kubernetes 	<ul style="list-style-type: none"> Kubernetes Slurm or Altair PBS Professional 	<ul style="list-style-type: none"> Kubernetes Slurm or Altair PBS Professional
Cluster Management	<ul style="list-style-type: none"> HPE Container Platform 	<ul style="list-style-type: none"> HPE Performance Cluster Manager Bright Cluster Manager 	<ul style="list-style-type: none"> Cray System Management 	<ul style="list-style-type: none"> HPE Performance Cluster Manager Bright Cluster Manager
Operating System	<ul style="list-style-type: none"> RHEL 7 CentOS 7 	<ul style="list-style-type: none"> RHEL 7 Ubuntu 18.04* 	<ul style="list-style-type: none"> Cray's Linux Environment 	<ul style="list-style-type: none"> RHEL 7 CentOS 7 Ubuntu 18.04*
Compute	<ul style="list-style-type: none"> HPE Servers 	<ul style="list-style-type: none"> HPE Apollo 6500 HPE ProLiant DL380 	<ul style="list-style-type: none"> Cray Supercomputers 	<ul style="list-style-type: none"> HPE Systems Cray CS Systems**

* Bright Cluster Manager only

** RHEL and CentOS only