Hewlett Packard Enterprise

GPU 서버를 통한 가장 효율적인 AI & HPC 구현 방안

HPE Korea AI/HPC Sales
박정현 부장


The world is replacing
programming with training
“The Artificial Intelligence market will
increase ten-fold by 2020.”
– BoA
Hewlett Packard
Enterprise

Hewlett Packard Enterprise

SK hynix

SEAGATE

SK Hynix partners with HPE Pointnext to transform to a containerized SSD firmware test environment

More solutions delivered faster with imp

Seagate optimizes manufacturing using edge systems and AI

Seagate built an edge-to-cloud architecture that makes real-time decisions, leading to greater efficiency and productivity in its manufacturing process.

→ Watch the video

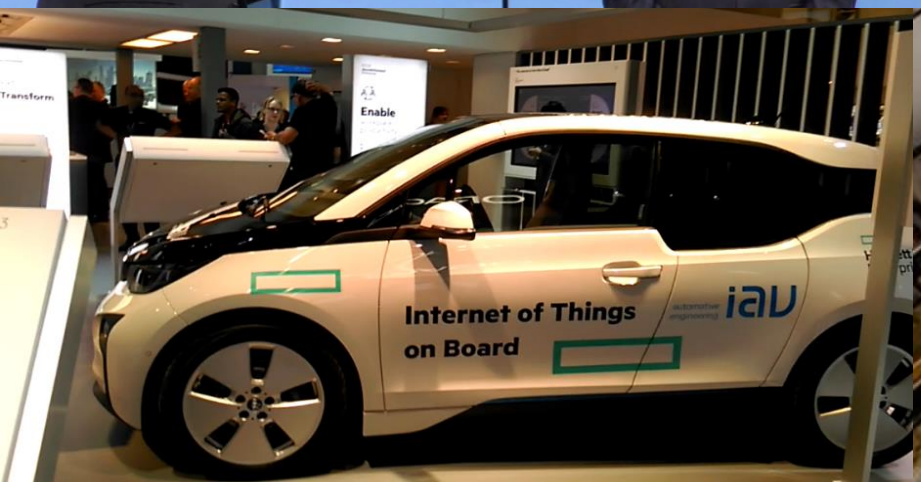

AI Development Trend

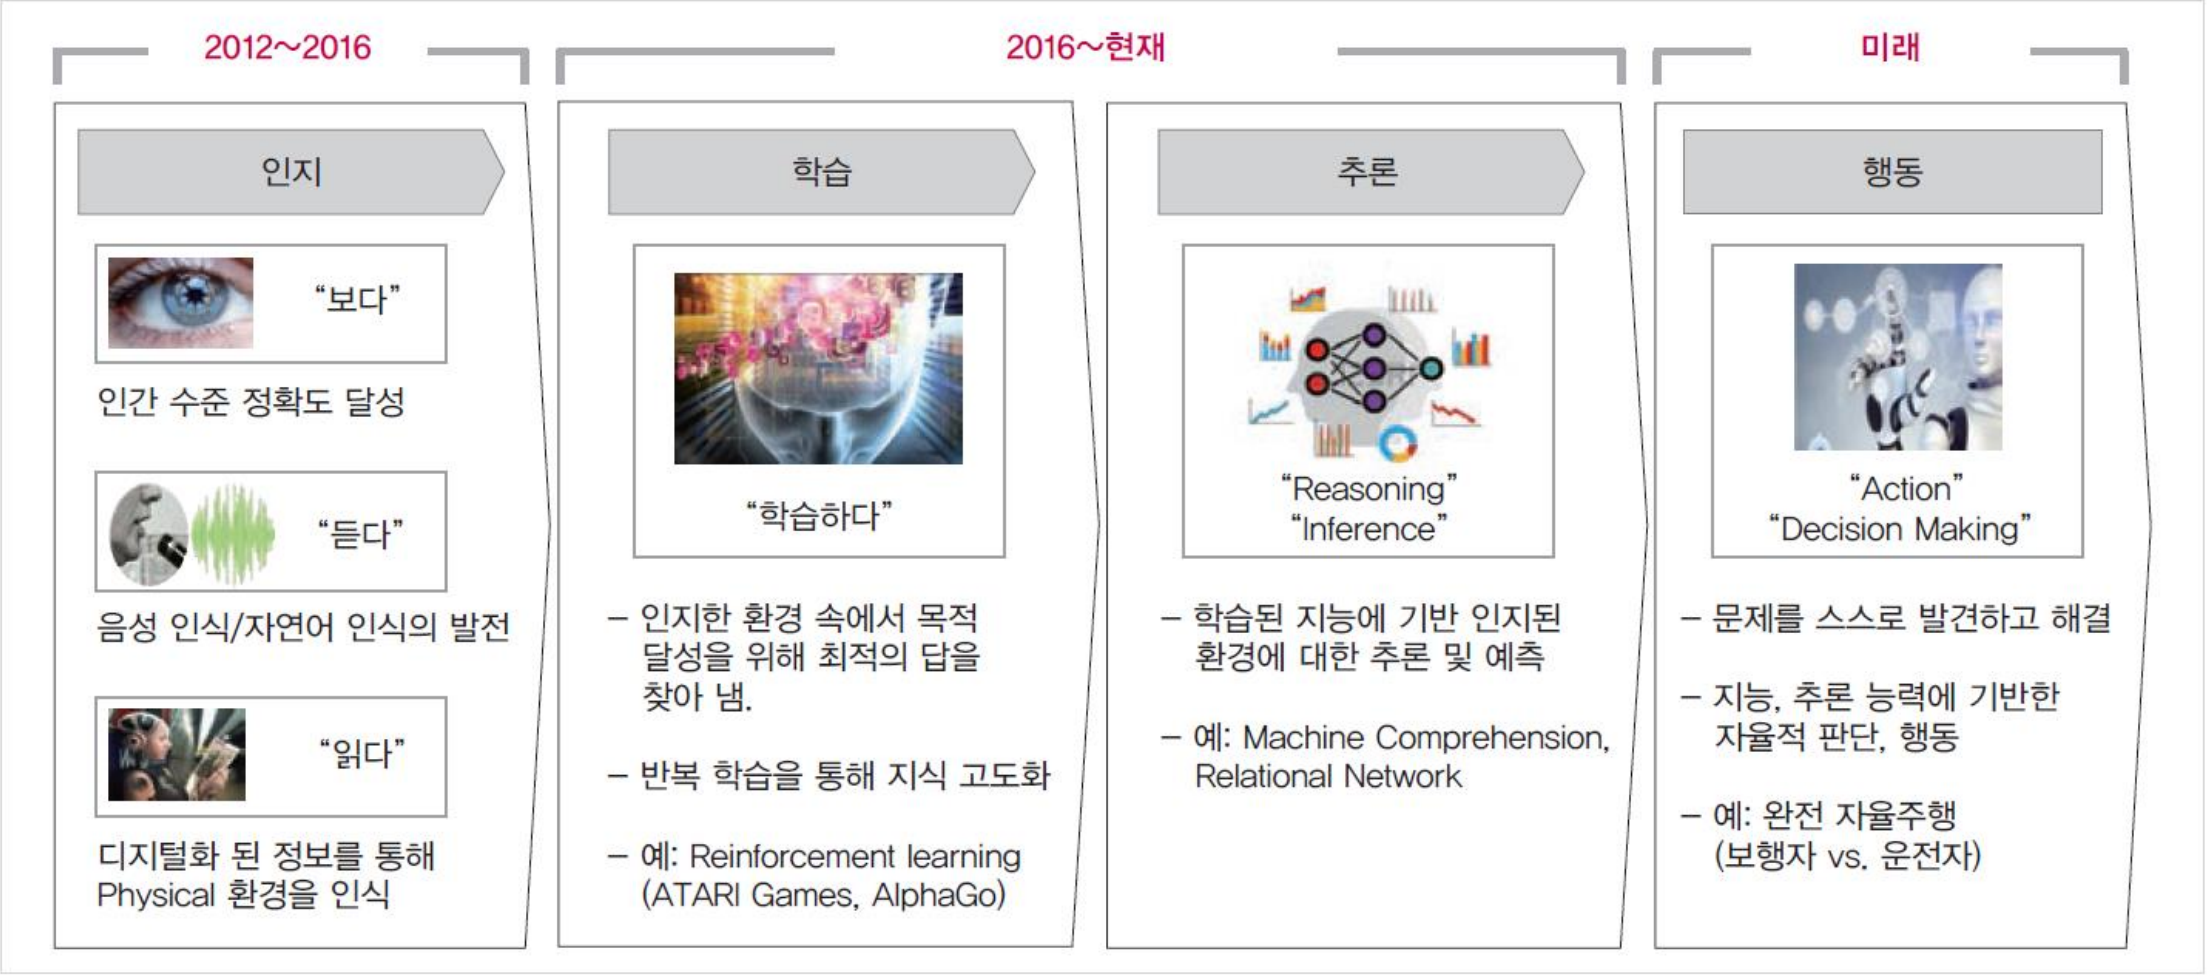

미래의 High-level machine intelligence

Driving a truck - 2027

Retail - 2031

Surgeon - 2043

Writing a bestseller – 2049

Math Research - 2060

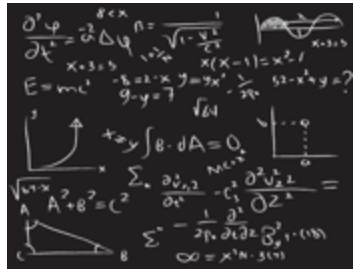

Full Automation of labor – 2140

Hewlett Packard Enterprise

Grace et al , When Will AI Exceed Human Performance? Evidence from AI Experts

HPE의 시각: What is the Market Size for AI?

IDC forecasts spending on AI-focused hardware, software, and services to reach $58bn by 2021

Figure 2: Global AI-focused spending* ($, bn)

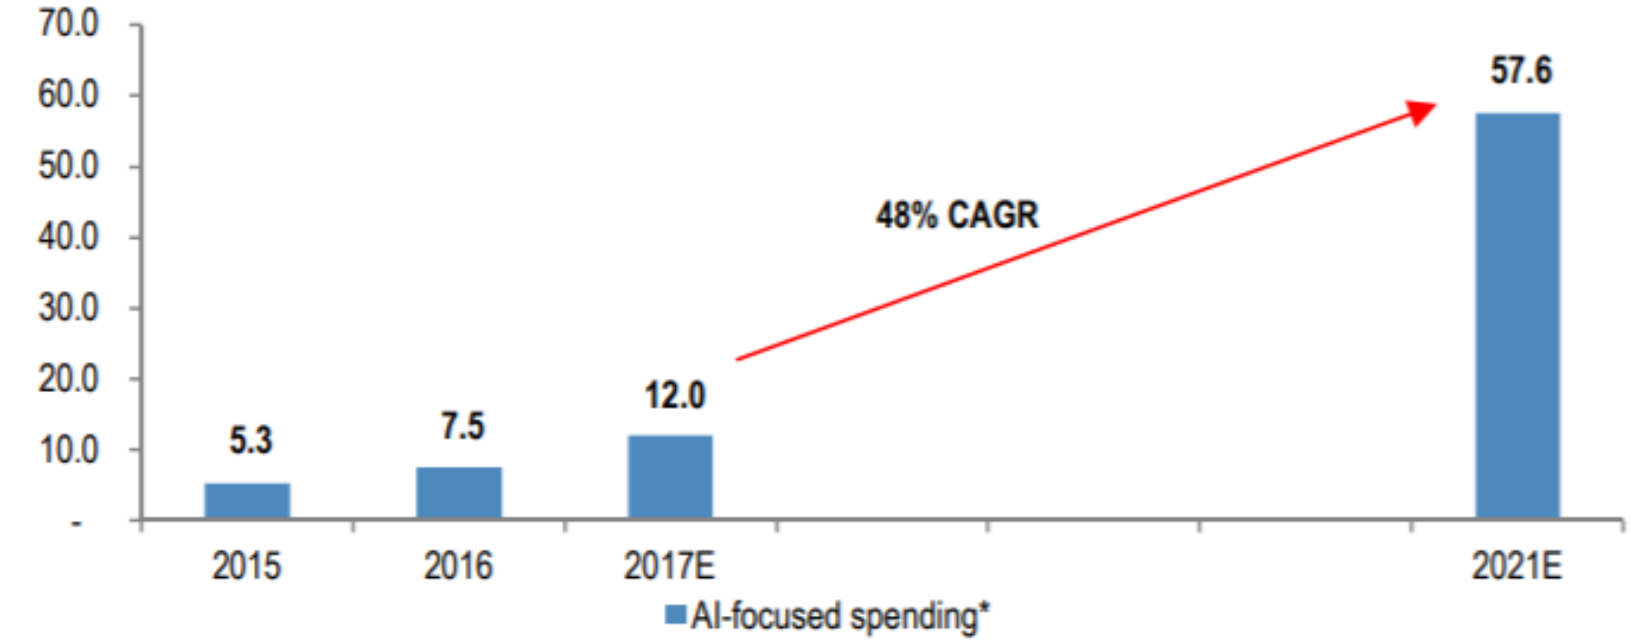

Source: AI-spending estimates from IDC. *Includes AI-focused spending on hardware, software (applications + software platforms), and services (IT consulting & system implementation).

HPE의 시각: AI가 주목하는 영역은?

Sample AI Use Cases Across Different Industry Verticals

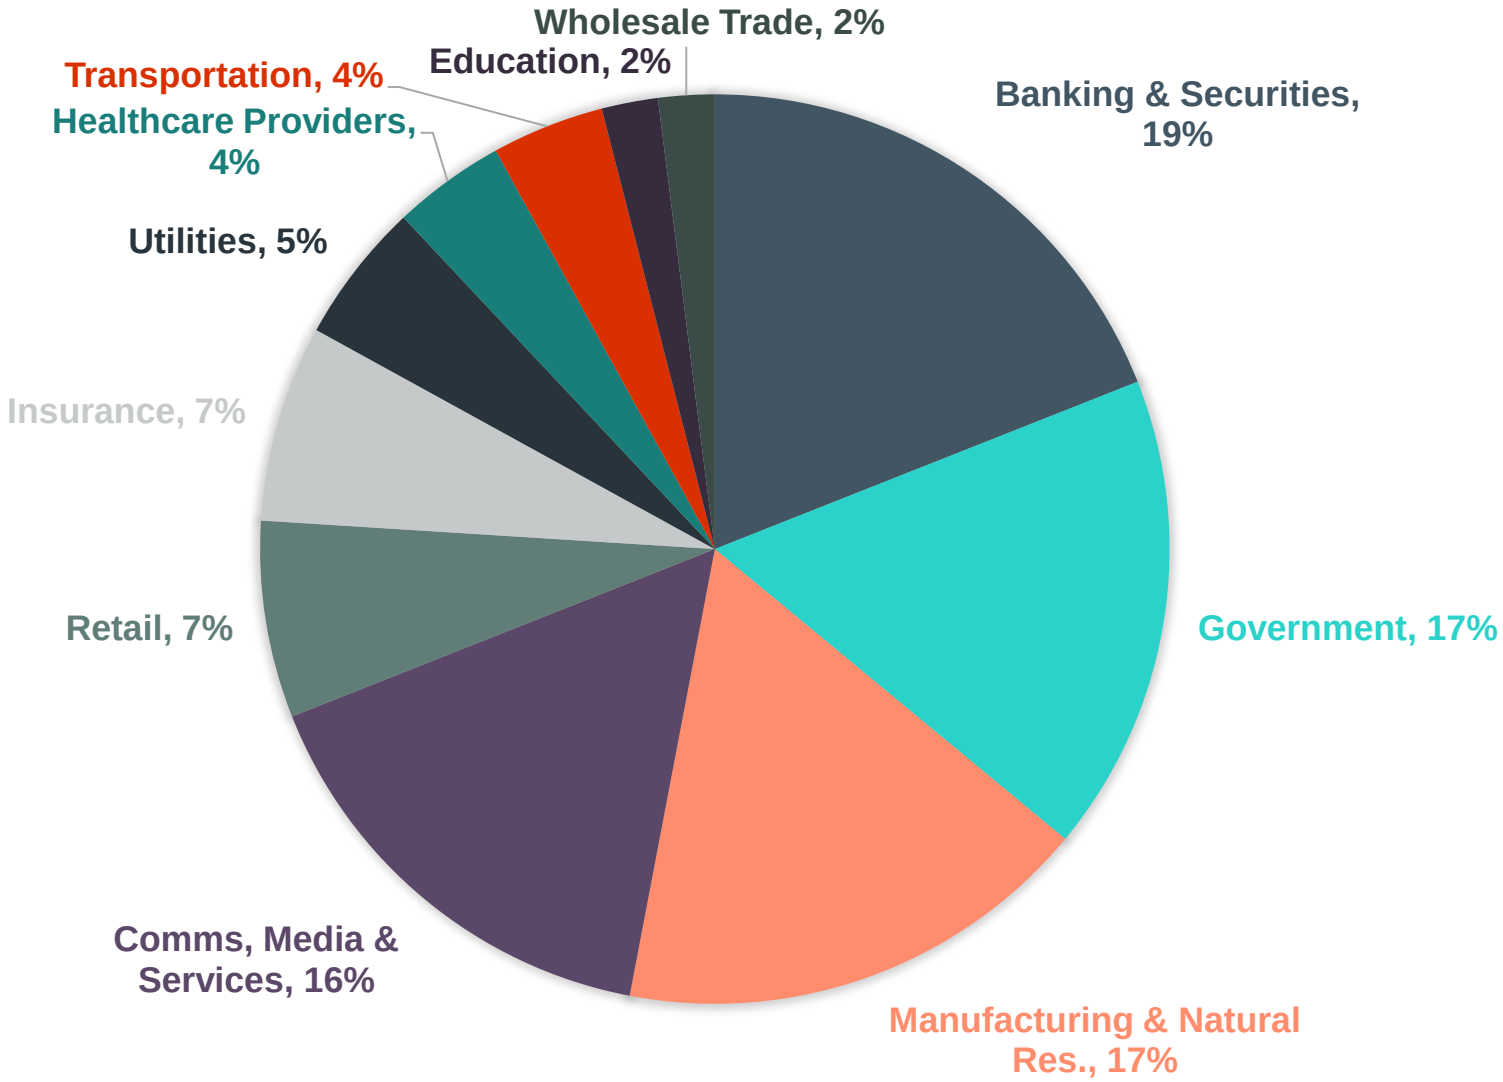

Hewlett Packard Enterprise

Infrastructure

An End to End Data Pipeline

Functional & Applications View

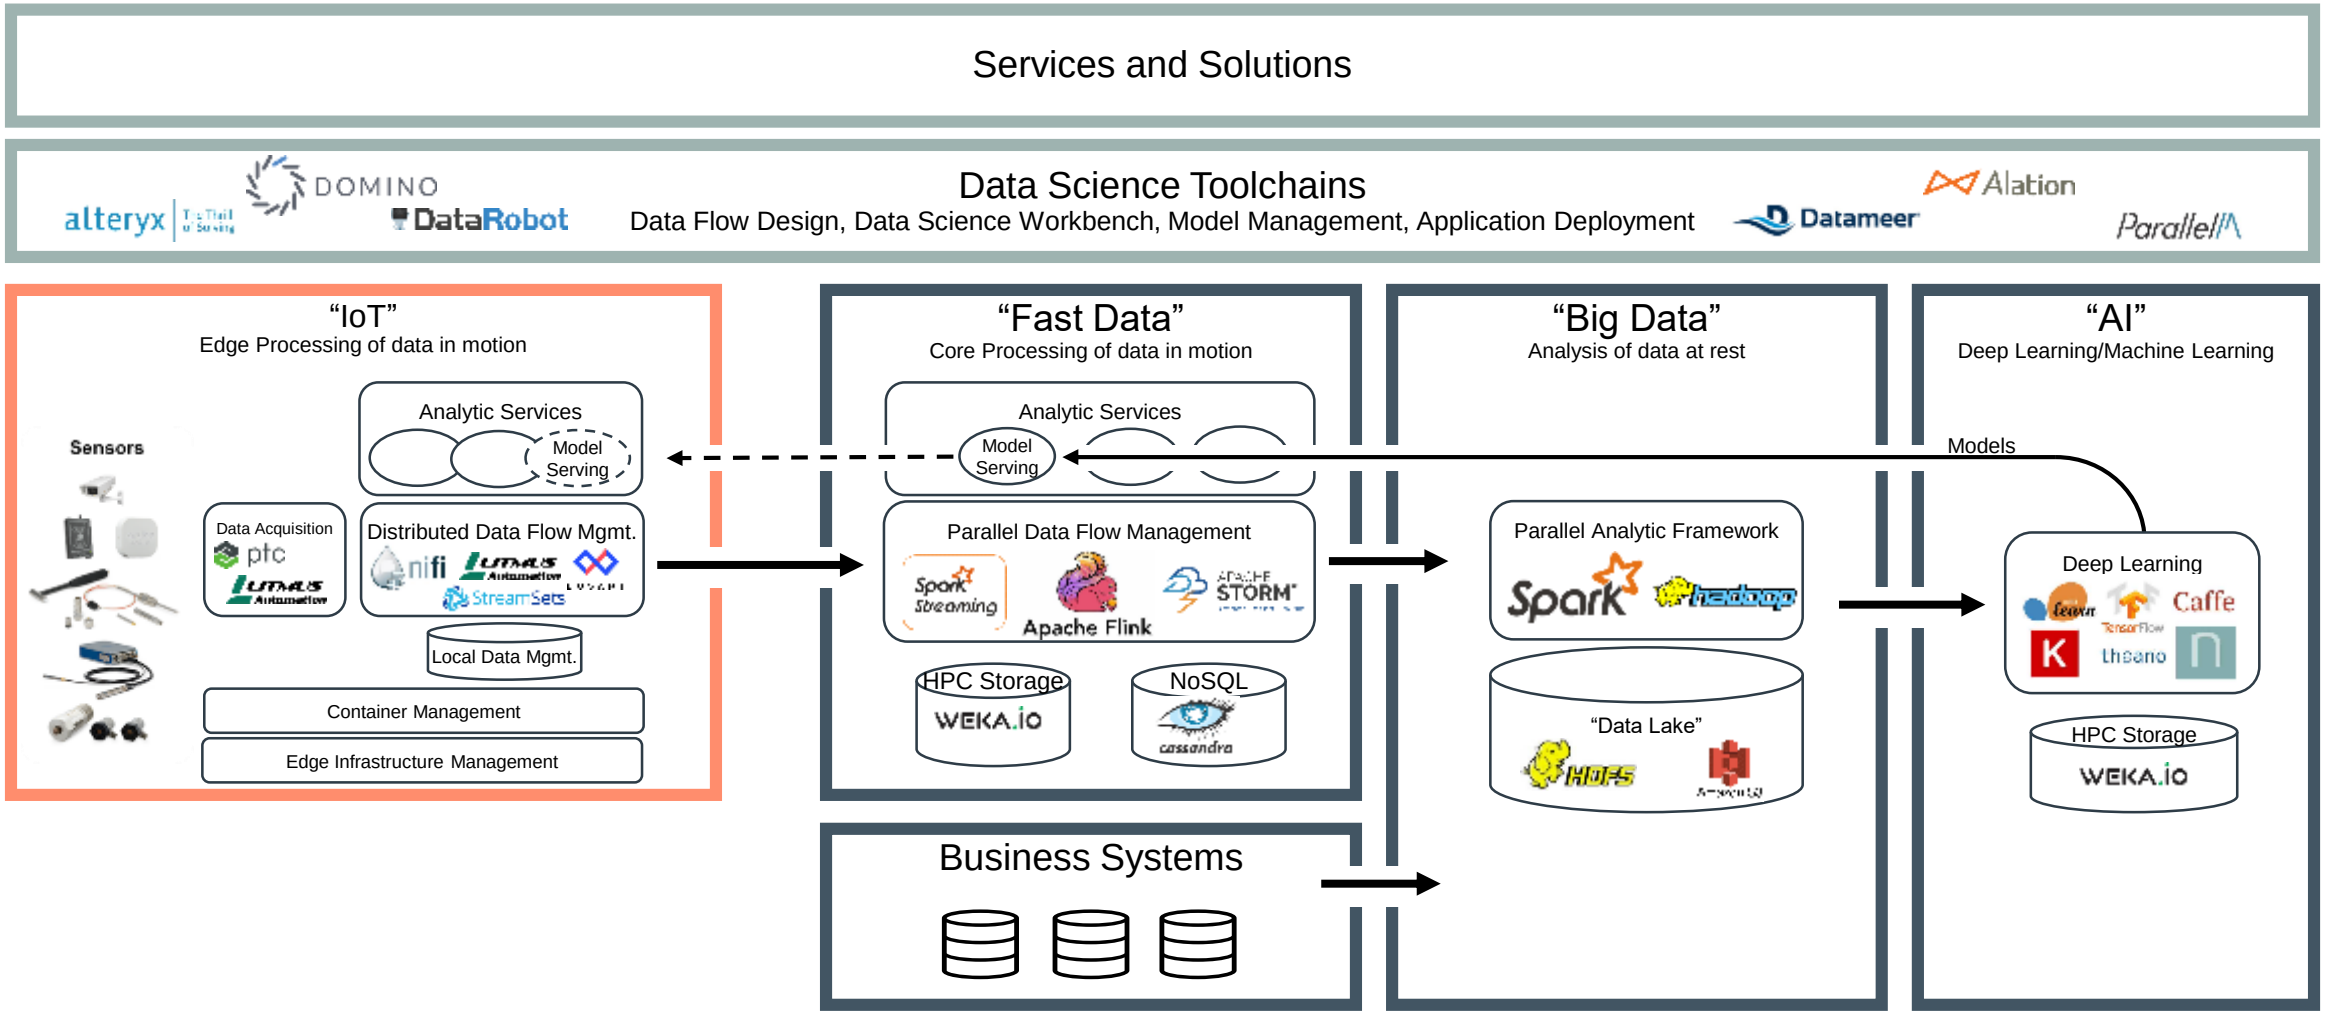

An End to End Data Pipeline

HPE Organizational View

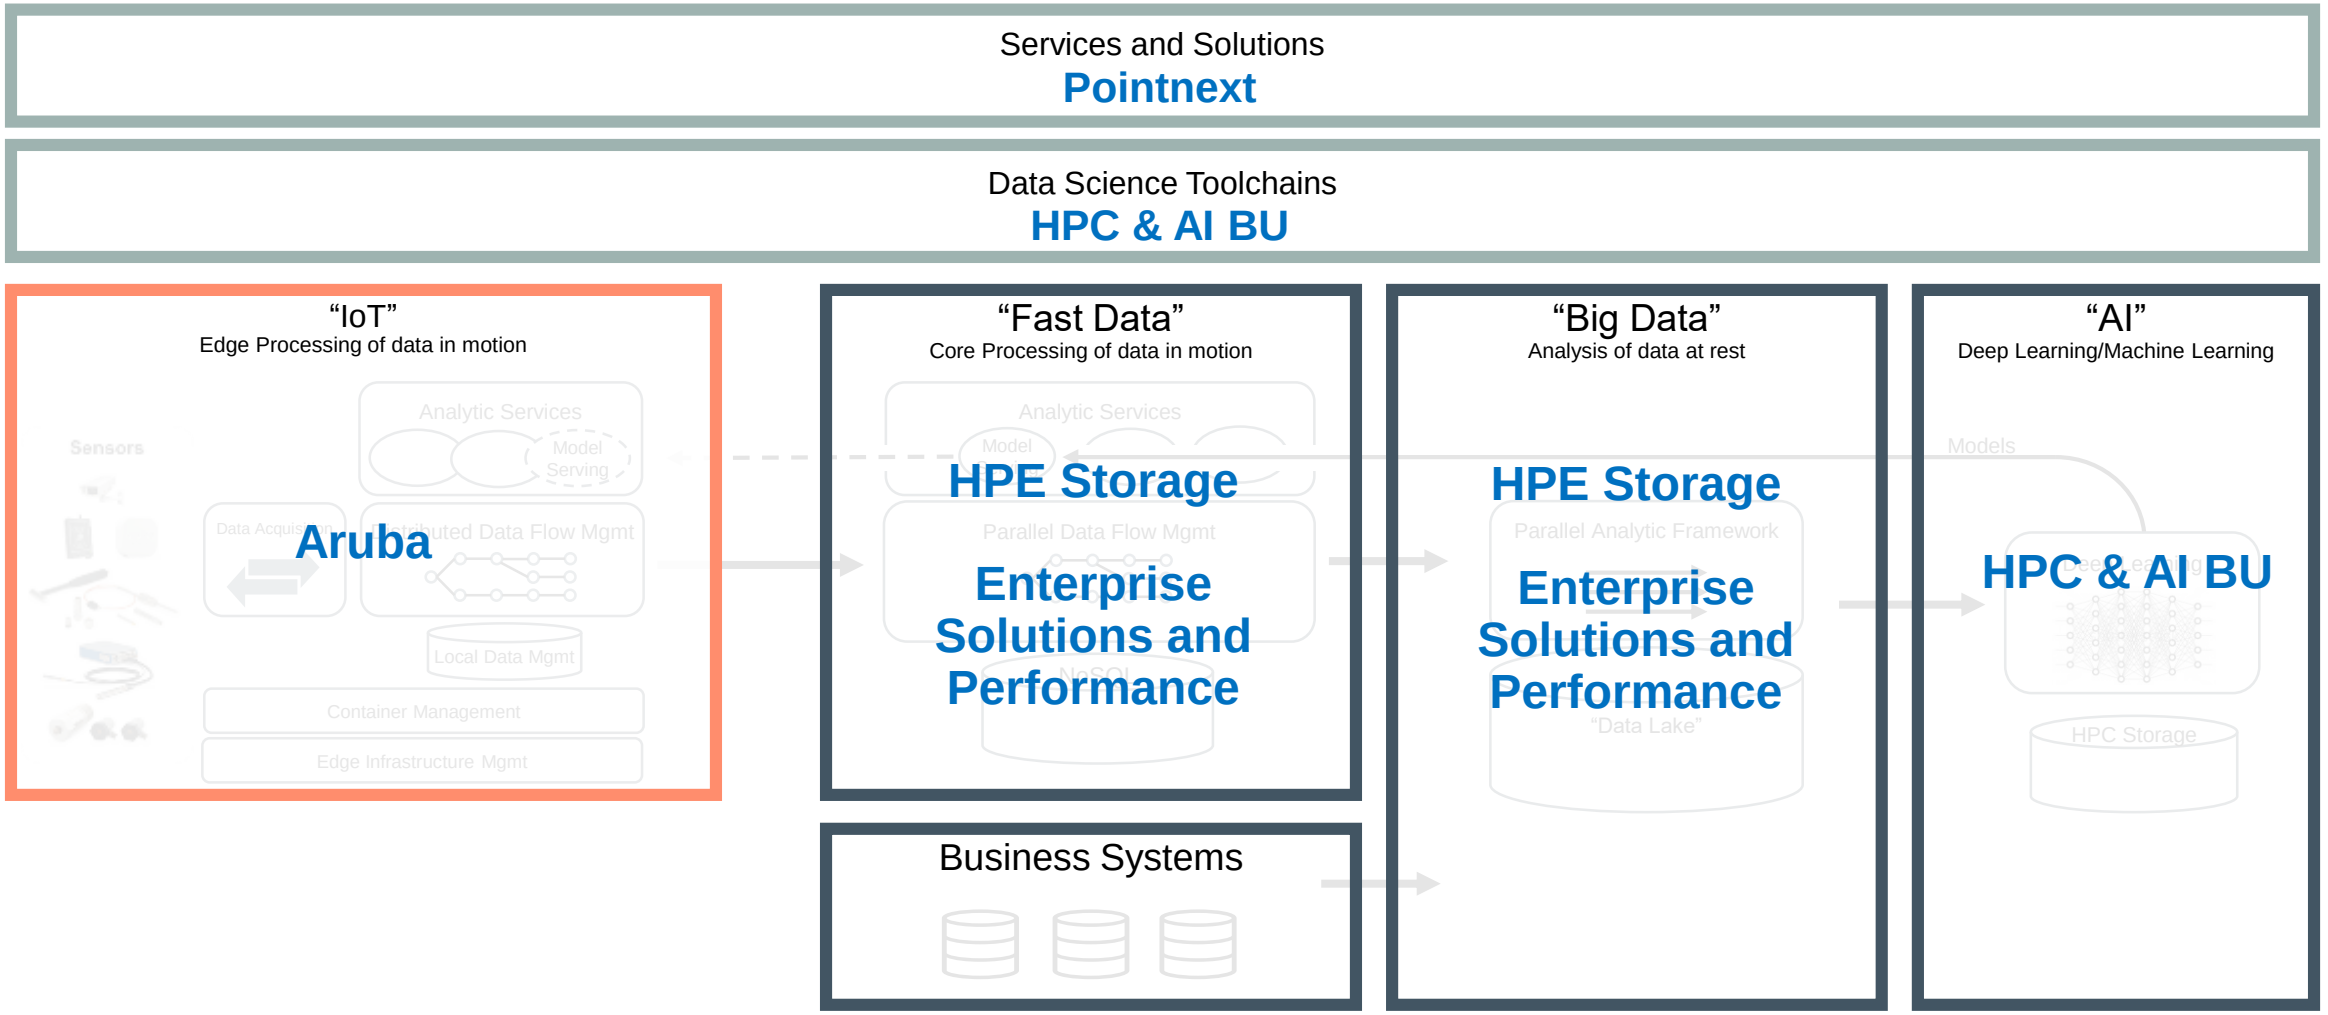

HPE enables AI from “intelligent edge to core data center”

HPE Pointnext

Intelligent Edge
(inference)

HPE OneView

HPE DMF

HPE System Management Software

HPE Aruba

HPE Networking

EDGE DATA

TRAINING DATA

Core data center
(deep learning)

HPE Edgeline

IOT Devices

Cost optimized Storage

HPE Apollo

Scality
Ceph

Performance optimized storage

InfiniBand

HPE Apollo

HPE ProLiant

WekaIO
Qumulo

HPE Apollo 6500 Gen10

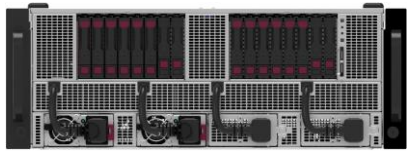

HPE Apollo

HPE ProLiant

HPE Synergy

NVIDIA®
Tesla® GPU
accelerators

Hewlett Packard Enterprise

HPE Deep Learning Cookbook

Deep learning 워크로드에 대한 적용 가이드

TensorFlow Caffe2 mxnet Caffe

– 8개의 Deep Learning 프레임워크 기반의 11개 워크로드에 대해서 8종의 HPE 하드웨어 구성에 대한 정보 제공

벤치마크 테스트 데이터 제공

– 워크로드에 대한 성능 예측치를 제공하여 최적의 시스템 사이징 근거 자료 제공

벤치마크 및 아키텍처 툴에 대한 오픈 소스화

GitHub

– Deep Learning 벤치마크 도구를 GitGub에 공개 예정

– Deep Learning 성능 분석 도구

– Hpe.com에 표준 아키텍처 정보 공개

Hewlett Packard Enterprise

Deep Learning Cookbook

Automatic Meeting Notes | Video Surveillance | Hospital Smart Care Unit | Custom

- [ ] Images
- [x] Videos
- [ ] Text
- [ ] Speech
- [ ] Sensor Data

- [ ] Classification
- [x] Detection
- [ ] Generation
- [ ] Anomaly Detection

- ( ) Training
  - ( ) Large
  - (•) Medium
  - ( ) Small
- ( ) Inference

Recommend

Data and Model

Data size	Epochs	Model
10000000	50	VGG19

Hardware

Server	Processor unit
Apollo 6500	NVIDIA P100

Count	Cluster size	Interconnect
8	2	InfiniBand FDR

Software

Framework	Batch size	Scaling
Caffe2	1024	strong

Add

Training performance

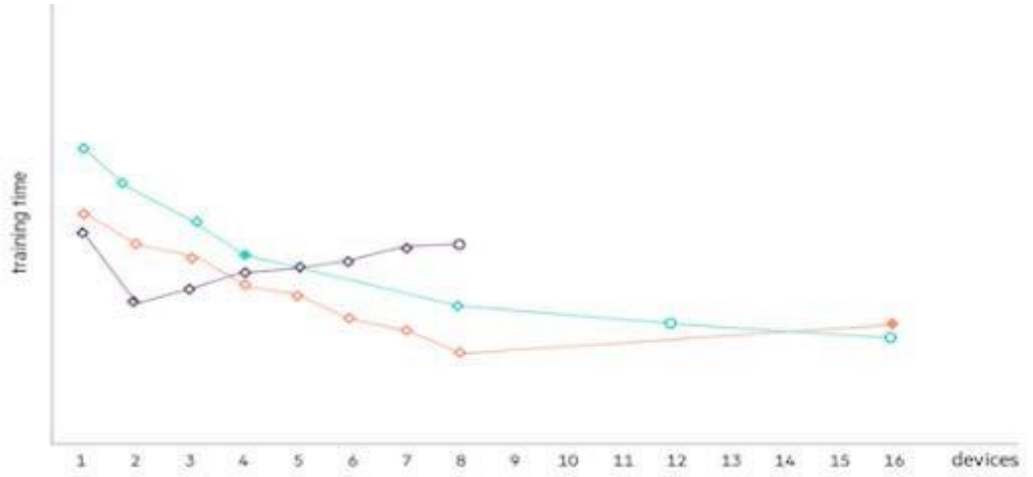

	Data			Hardware			Software	Time (hours)	
■	Size	Epochs	Model	Server		PU	Framework		
	10000000	50	AlexNet	Apollo 6500		NVIDIA P100	Caffe2		
				Count	Cluster size	Interconnect	Batch	22.5	✕
				8	2	IB	1024(strong)		
■	Size	Epochs	Model	Server		PU	Framework		
	10000000	50	GoogleNet	Apollo 6500		NVIDIA P100	Caffe2		
				Count	Cluster size	Interconnect	Batch	26.3	✕
				8	2	IB	1024(strong)		
■	Size	Epochs	Model	Server		PU	Framework		
	10000000	50	VGG19	Apollo 6500		NVIDIA P100	Caffe2		
				Count	Cluster size	Interconnect	Batch	147.4	✕
				8	2	IB	1024(strong)		

Remove all

Hewlett Packard Enterprise

Hewlett Packard Labs

각 Workload별 서버의 특징

HPE purpose-built portfolio for HPC and AI

HPE POINTNEXT

Advisory, Professional, Operational Services | HPE GreenLake Flex Capacity | HPE Hybrid HPC

Supercomputing / Enterprise / Commercial HPC

HPE SGI 8600

Liquid cooled, Petaflop scale for HPC and AI

HPE Apollo 6000 Gen10

Air-cooled, HPC at rack scale

HPE Apollo 2000 Gen10

The bridge to enterprise scale-out architecture

HPE Apollo sx40

Max GPU capacity and performance with lower TCO

Edge Compute

HPE Edgeline EL4000

Unprecedented edge compute and high capacity storage

Adaptive Cooling

HPE Adaptive Rack Cooling System

Higher power density with less datacenter heat

Software

HPC Software

- **HPE Performance Cluster Manager**
- **HPE Message Passing Interface (MPI)**
- 3rd party solutions sold by HPE

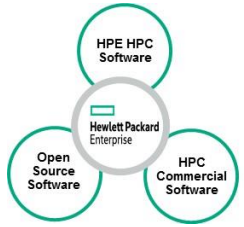

AI Software

- **Easy Setup and Flexible OS** using Bright Computing's distribution of deep learning software development components
- **NVIDIA GPU Cloud** Comprehensive catalog of GPU-accelerated containers for deep learning software

Emerging

HPE Apollo 6500 Gen10

Enterprise platform for accelerated computing

HPE Apollo 70

HPC cluster ready Arm Based server

HPE Apollo 35

Best AMD Performance in Dense HPC Platform

In-memory

HPE Superdome Flex Server

Scale-up, shared memory HPC, combines best of HPE and SGI technologies

Storage

HPE Data Management Framework

HPE Apollo 4000 Series

Networking

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

GPU의 필요성

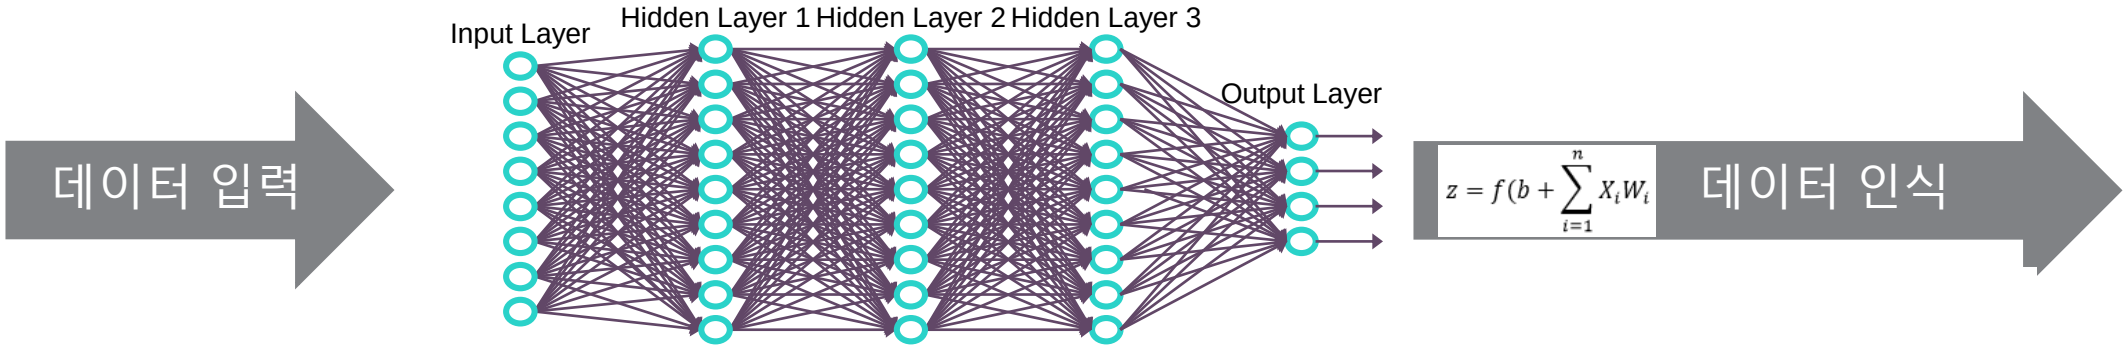

NVIDIA GPU 가 Weight Factor 연산을 위한 계산에 적합

몇 일 소요

8개의 GPU가 서로 통신하여 연산하면 연산 시간 단축

몇 시간 소요

GPU 간 통신 성능 향상을 통해서 추가적인 훈련 성능 가능

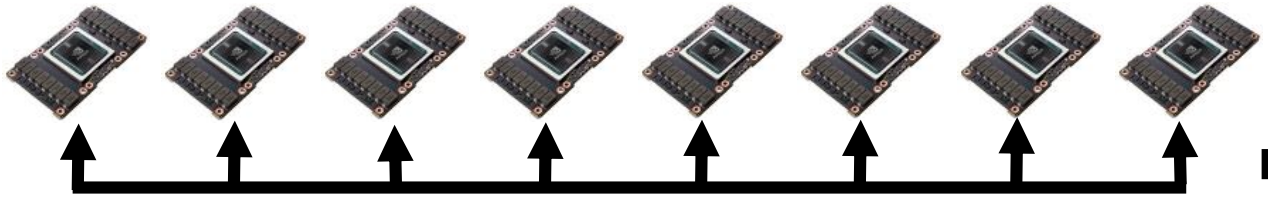

몇 분 소요

Hewlett Packard Enterprise

HPE Deep Learning용 GPU Server

Inference 또는 소규모 개발용 서버

HPE Apollo 40 SX40 Gen10

- 1U 크기
- 4 x NVLink GPU 지원
- Training용

HPE ProLiant DL380 Gen10

- 2U 크기
- 3 x PCIe GPU 지원
- Development, Inference

HPE Edgeline 4000

- 1U 크기
- 4 x PCIe GPU 지원 (NVIDIA P4,T4)
- Inference, Edge Computing용

Hewlett Packard Enterprise

HPE Apollo6500 Gen10 전면, 후면부

NVME 드라이브 장착

2개의 8-디스크 장착 베이
- 최대 16 SAS/SATA
- 최대 4개의 NVMe

2+2 파워이중화
- 4개의 2200W 전원 모듈

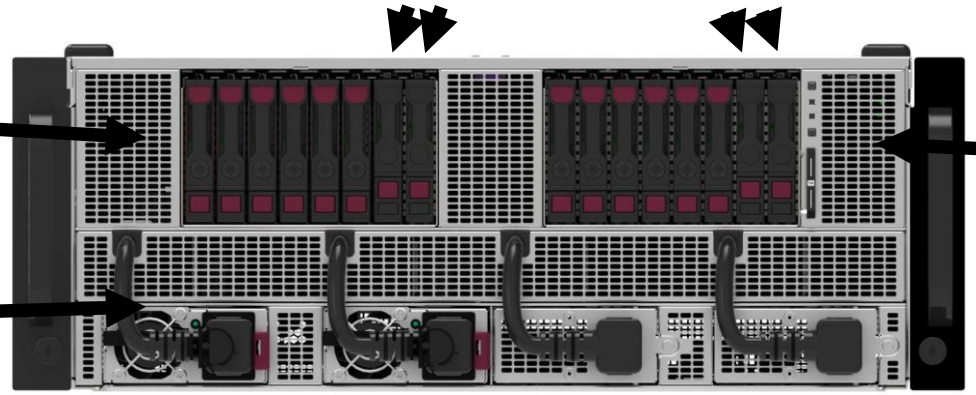

Front View

시스템 정보 표지
Serial Label Pull tab
전원버튼 / LED
System Health ID
NIC ID
Unit Identification (UID)
온도 센서

Rear View (SMX-2)

GPU Module SXM-2
- 최대 8 GPU
- 4개 별도 PCIe 슬롯

메인보드

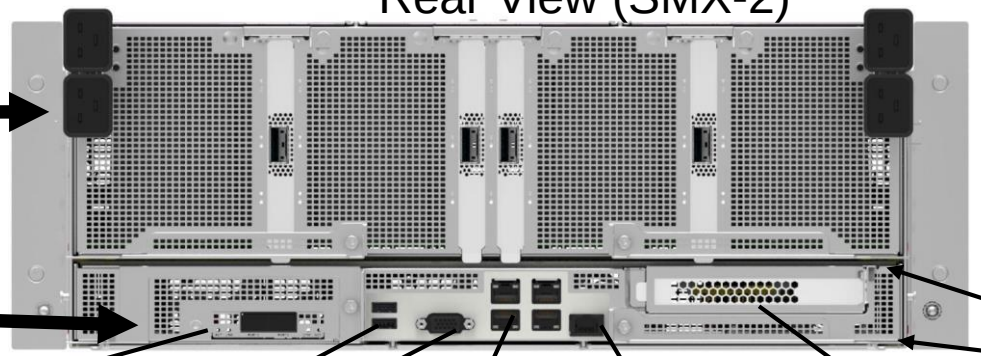

4U
1000mm 표준랙 장착

메인보드/GPU 장착용
케이스 레버

Hewlett Packard Enterprise

HPE의 AI 전용 서버 Apollo 6500 Gen10, 4th Generation

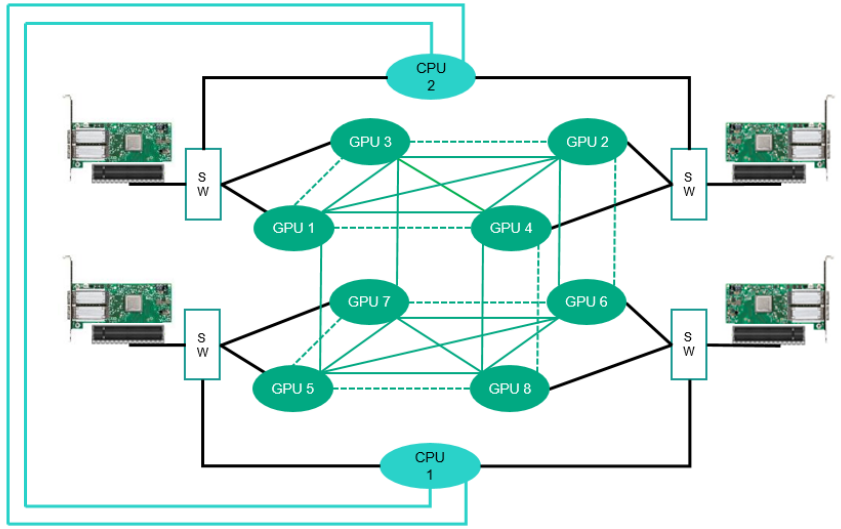

2018 최다판매 AI 시스템 in Korea
8GPUs NVLink / PCI NVIDIA-AMD-Intel FPGA
완벽한 온도 발열 테스트
HPE certified VBIOS / Driver

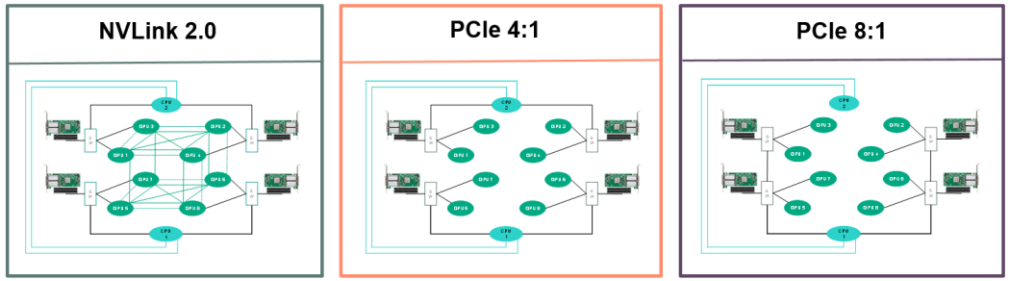

Hewlett Packard Enterprise

[1,2] Please see speaker notes for substantiation

Check Point: Environment – 4U 8GPU is answer for current

HPE ProLiant

Model: XL270d 8 SXM2.0 Module | Gen.: Gen10 | Thermal Report

Date: May 10, 2018

Configuration Description	Condition at 25°C: Power, idle Watts	Power, idle BTU/hr	Airflow, idle CFM	Airflow, idle m^3/hour	Average System ΔT, idle (see note 5) °Celsius	Average System ΔT, idle (see note 5) °Fahrenheit	Power, stressed Watts	Power, stressed BTU/hr	Airflow, stressed CFM	Airflow, stressed m^3/hour	Average System ΔT, stressed (see note 5) °Celsius	Average System ΔT, stressed (see note 5) °Fahrenheit	Weight[2] lbs	Weight[2] kg	Overall System Dimensions[2] inches width	inches depth	inches height	mm width	mm depth	mm height	rack U height
Minimum, redundant fans	560	1910.7	194.9	331.3	5.2	9.3	2260	7711.1	257.8	438.3	15.8	28.4									
Typical, redundant fans	550	1876.6	214.6	364.8	4.6	8.3	3510	11976.1	290.2	493.3	21.8	39.2									

Configuration Description	Condition at 35°C: Power, idle Watts	Power, idle BTU/hr	Airflow, idle CFM	Airflow, idle m^3/hour	Average System ΔT, idle (see note 5) °Celsius	Average System ΔT, idle (see note 5) °Fahrenheit	Power, stressed Watts	Power, stressed BTU/hr	Airflow, stressed CFM	Airflow, stressed m^3/hour	Average System ΔT, stressed (see note 5) °Celsius	Average System ΔT, stressed (see note 5) °Fahrenheit
Minimum, redundant fans	494	1685.5	230.6	392.0	3.9	6.9	2255	7694.1	384.6	653.8	10.6	19.0
Typical, redundant fans	555	1893.7	314.4	534.5	3.2	5.7	3693	12599.5	420.8	715.4	15.8	28.4

System Firmware	
ROM version:	U45 v1.40 (2/27/2018)
iLO version:	1.20 Pass 27
Thermal settings:	Optimal Cooling

ASHRAE Class[3]
A2

Cooling Airflow Diagram

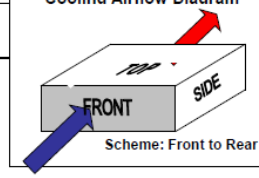

Scheme: Front to Rear [4]

Configuration Description		Feature Quantity and Type per Server: Processors	DIMMs	HDDs	Media	IO adapter 1	IO adapter 2	IO adapter 3	IO adapter 4	IO adapter 5	IO adapter 6	Feature Quantity and Type per System: Servers	Fans	PSUs	Other[9]
Minimum	Quantity	2	12	8	0	1	1	4	4			1	5	2	
	Description	G6142M 150W	128GB LR-DIMM	8x3.2TB SAS SSD	N/A	P816i HD Controller[9]	1GB 4-Port FLOM[9]	2P QSFP Card[9]	V100 SXM2.0 GPU				80&60mm Nidec	Artysan 2200W	
	Location(s)	P1 & P2	1 DPC	HDDs 1-16	N/A	AROC Slot	FLOM	Slots 9-12	Slots 1-8			N/A	Slots 1-5	Slots 1&2	
Typical	Quantity	2	12	16	0	1	1	1	4	8		1	5	4	
	Description	G6142M 150W	128GB LR-DIMM	12x3.2TB SAS SSD, 4x1.2TB	N/A	P816i HD Controller[9]	1GB 4-Port FLOM[9]	NVMe Controller	2P QSFP Card[9]	V100 SXM2.0 GPU			80&60mm Nidec	Artysan 2200W	
	Location(s)	P1 & P2	1 DPC	HDDs 1-16	N/A	AROC Slot	FLOM	Slot 21	Slots 9-12	Slots 1-8		N/A	Slots 1-5	Slots 1-4	

Notes:

1 Data represents system testing at idle in RHEL v7.3 and while stressed with NVQual,PTUgen,Vesper,Fabric Stress, and mPrime.

2 In case of discrepancies to information within QuickSpecs, QuickSpecs should be considered to have the most updated information.

3 This represents the closest ASHRAE class to this system's operating environmental requirements for the listed configurations, which are:

System inlet Temperature: 10° to 35°C (50° to 95°F) at sea level with an altitude derating of 1.0°C per every 305m (1.8°F per every 1000ft) above sea level to a maximum of 3050m (10,000ft), no direct sustained sunlight. Maximum rate of change is 10°C/Hr (18°F/Hr). The upper limit may be limited by the type and number of options installed.

Relative Humidity: 10 to 90% relative humidity (Rh), 28°C (82.4°F) maximum wet bulb temperature, non-condensing.

Maximum Altitude: 3050 m (10,000 ft). This value may be limited by the type and number of options installed. Maximum allowable altitude change rate is 457 m/min (1500ft/min).

For information regarding hardware configurations that support extended ambient environments, view the following URL: **http://www.hpe.com/servers/ashrae**

4 System airflow is as indicated within the diagram. Rack-level airflow is front-to-rear, unless stated otherwise.

5 The average temperature rise from the front to the exhaust of the system is calculated via the caloric equation using the measured power and airflow for a given test.

6 Due to the multitude of uses for this system, there is no one configuration nor application program that can truly be deemed as "typical".

7 All values shown are for reference only. Customers are encouraged to measure the airflow on their system configuration, running their software applications, under e conditions that represent their point of use, as all of these factors can influence fan speed and power consumption.

8 The listed measured power is pressumed to equate to the system's dissipated power.

9 Indicated adapter is 3-D Sea of Sensor compliant.

Table 18. Configuration K

CPU Power Dissipation / NVLink SXM2	2x 70 W	2x 85 W	2x 105 W	2x 125 W	2x 130 W	2x 140 W	2x 150 W	2x 165 W
300W	25	24	22	21	20	19	18	18

Table 19. Configuration M

CPU Power Dissipation / NVLink SXM2	2x 70 W	2x 85 W	2x 105 W	2x 125 W	2x 130 W	2x 140 W	2x 150 W	2x 165 W
300W	24	24	23	23	22	22	20	20

Hewlett Packard Enterprise

타사의 1U/4GPU서버 환경 조건

Check Point2: RAS

AI용 GPU 시스템의 시스템-GPU 지원특징

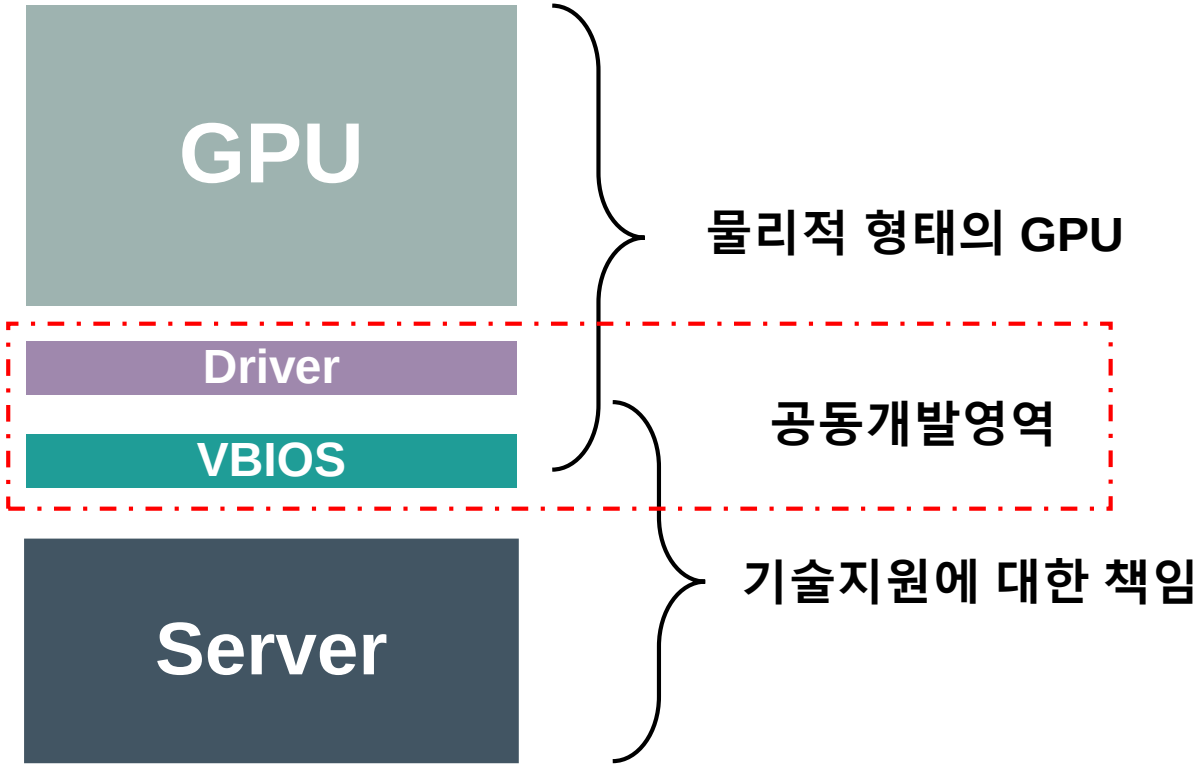

Disk Hotswap은 되는가?

Fan이 Hotswap이 되는가?

L1/L2/L3 기술지원이 본사에 준비되어있는가?

Eco System과의 기술 지원 체계?

Framework/Library에 대한 지식?

Hewlett Packard
Enterprise

Check Point3: 보안

1	BMC(Baseboard Management Controller) 칩을 서버 업체가 자체 제작해야 한다.
2	서버 부팅 후에도 주기적으로 주요 Firmware의 오염 여부를 확인할 수 있어야 한다.
3	BMC 칩의 Firmware가 오염되었을 경우 가장 안전한 버젼으로 자동 복구되어야 한다.

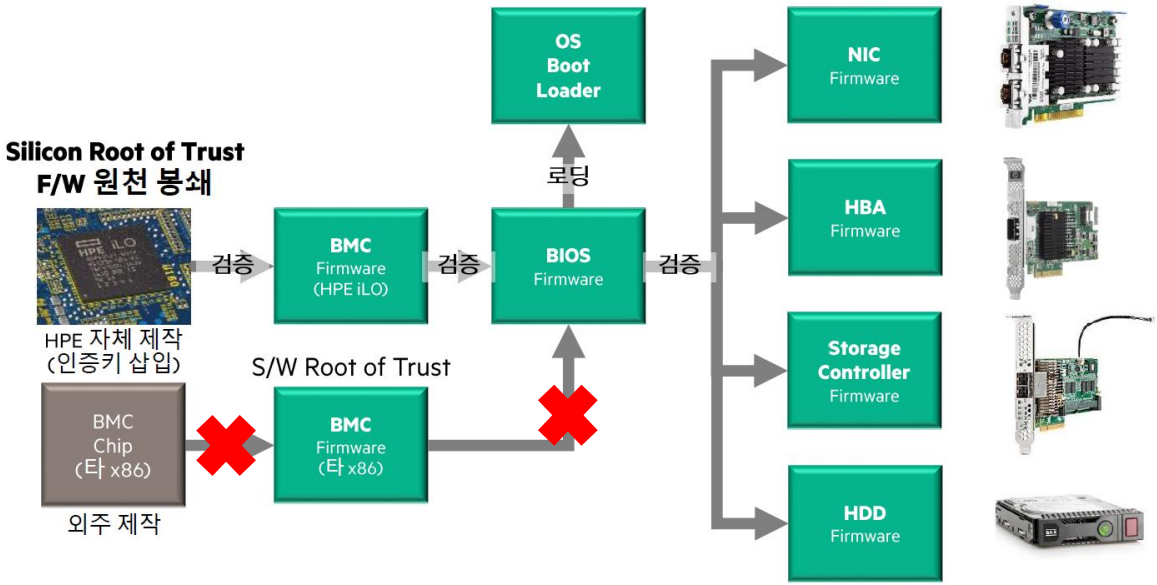

Gen10 Security 보안 기능으로 HPE를 새로 선택해주신 고객

Hewlett Packard Enterprise

HPE HPC and AI storage point-of-view

Optimal storage solutions for HPC and AI

Storage & Data Management

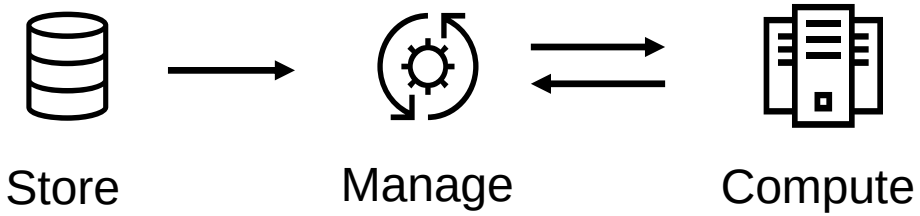

- HPE delivers **storage options** that map to various HPC and AI strategies and implementations, e.g. existing infrastructure, training goals, pre-processing applications, etc.
- CPU, GPU and application requirements drive storage **performance and capacity** management requirements
- HPE storage portfolio provides **flexibility** and **choice**

HPC & AI storage components

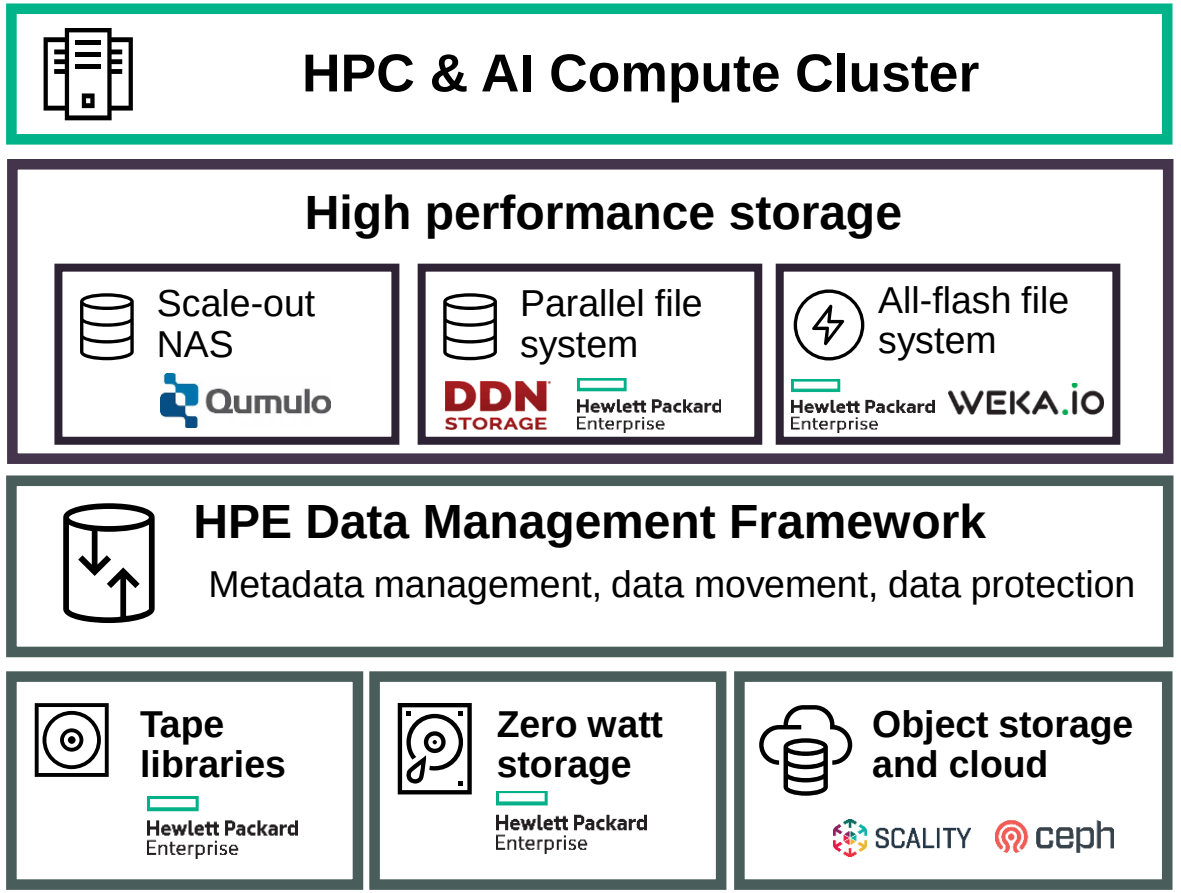

Reference Architecture & Technical White Papers

Reference Architecture

HPE Reference Configuration for HPE Apollo 4200 Gen10 with Hadoop 3

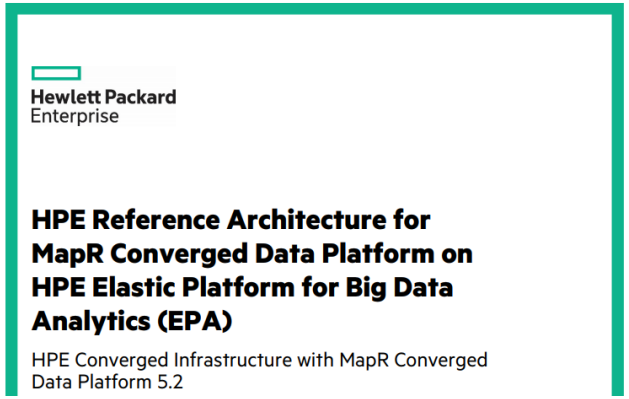

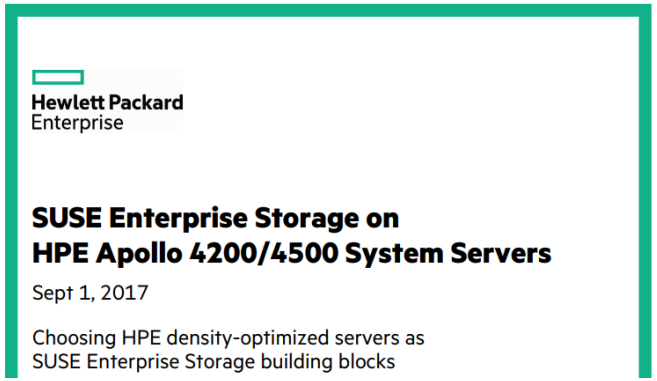

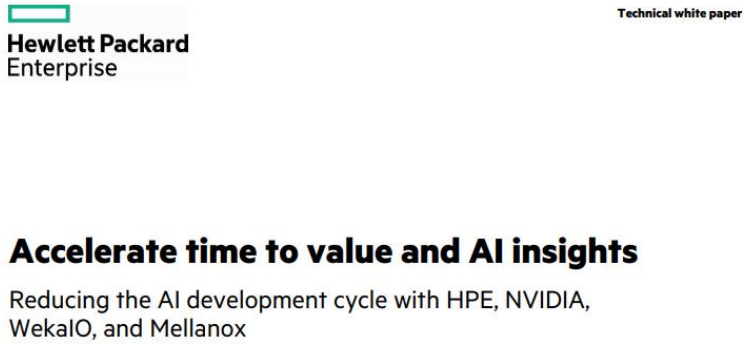

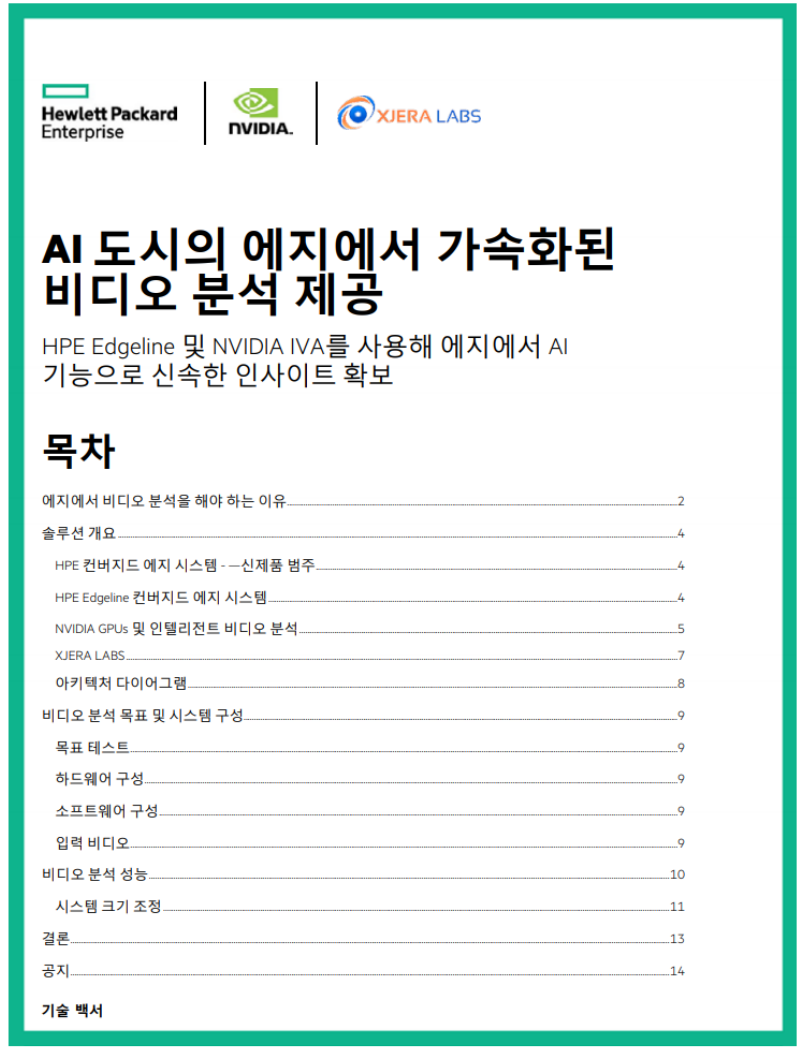

ONNX

OPEN NEURAL NETWORK EXCHANGE FORMAT

The new open ecosystem for interchangeable AI models

aws

Facebook Open Source

Microsoft

What is ONNX?

ONNX is a open format to represent deep learning models. With ONNX, AI developers can more easily move models between state-of-the-art tools and choose the combination that is best for them. ONNX is developed and supported by a community of partners.

AMD arm NVIDIA IBM

MEDIATEK Hewlett Packard Enterprise intel AI

HUAWEI Preferred Networks Qualcomm

SOPHON skymizer Baidu 百度

In HPC, we are making great progress….

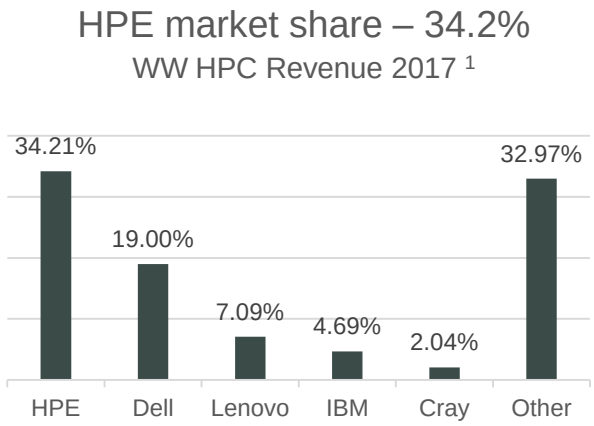

HPE's #1 share is 15 pts > #2 Dell

Hyperion Research

"HPE's next-generation and new Apollo systems will facilitate that adoption by providing easier integration and management while delivering extreme density to reduce data center footprint and extend the range of HPC and AI use cases."

Steve Conway,
SVP of Research

Creating value for customers

Rolls-Royce · Zentralanstalt für Meteorologie und Geodynamik ZAMG · PSC Pittsburgh Supercomputing Center · NREL National Renewable Energy Laboratory · Tokyo Institute of Technology · Notre Dame · NCAR National Center for Atmospheric Research · Sandia National Laboratories

HPE helps EPFL Blue Brain Project unlock the secrets of the brain

Blue Brain Project

World's largest Arm Supercomputer for U.S. Department of Energy[2]

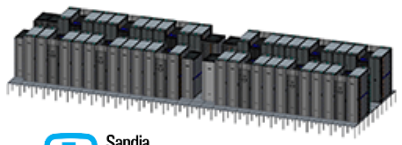

Sandia National Laboratories

Best HPC Server – 2017[3] HPE Apollo 6000 Gen10

Best HPC Server Product or Technology
HPC wire
Readers' Choice Awards 2017

Fastest parallel processing performance[4]

HPE SGI 8600

HPE GreenLake – best-in-class[5] consumption model

Hewlett Packard Enterprise

Thank you