




Hewlett Packard
Enterprise

實現數位轉型的AI基礎建設藍圖





資料可以帶動企業轉型

透過 AI 和進階分析技術釋放資料 價值

為什麼應該關注 AI/進階分析？

我們的客戶想要 AI 和進階分析...

AI 與進階分析是 CIO 的
3 大優先事項

在 2021 年，AI 與進階
分析基礎架構
將佔市場規模的
15-20%¹

...但面臨諸多挑戰

使用案例
新角色、技能差距
文化與變革
資料準備
舊有基礎架構

世界正以電腦訓練來取代電腦程式設計

4 成的公司表示
缺乏分析技能
是一項重大挑戰

到 2020 年，**50%** 的組織將
缺乏創造商業價值所需的
足夠 AI 和資料素養技能

2300 萬
個開發者遍佈全球

少數資料科學家

讓我們從頭說起...什麼是 AI？

什麼讓機器變得如此聰明？

人工智慧 (AI)

模仿人類行為。讓機器用類似人類行為的方式解決任務的任何技巧。



範例：
自動駕駛車

深度學習

深度學習 (DL)

ML 的子集，使用深度神經網路做為模型，啟發自人類大腦的結構和功能。



範例：Siri

機器學習

機器學習 (ML)

讓電腦無需經過明確的程式設計，即可從典範中學習的演算法。



範例：
Google 地圖

人工智慧

AI 與進階分析正落實於企業中


尤其是這 8 種垂直行業



金融服務
詐騙偵測、ID 驗證



政府
網路安全性、智慧城市與公用事業



能源
震波與儲油模型



零售
影像監控、購物型態



醫療
個人化醫藥、影像分析



消費者科技
聊天機器人



服務供應商
媒體傳輸




製造
防禦性和目標性維護


我們的 AI 與分析願景，涵蓋從邊緣到雲端的資料



我們的承諾：無論您在資料旅程上的哪一個階段，HPE 都能助您一臂之力



整個資料生命週期的分析 功能圖




如今 HPE 開發了企業級的 AI 技術...



感應器資料會在試車和賽車過程中進行分析，
然後發出應對汽車或
駕駛策略做出改變的通知




包括安全城市、智慧水網、
智慧護欄、智慧廢棄物管理、智慧停車系
統和員工追蹤



融合 AI、IT 與 OT 技術
大幅增進產能與品質

...同時協助企業執行其進階分析發展方案



利用 SAP HANA 和 Hadoop
將資料帶進每項決策

使用in-memory分析技術
Alzheimer 的研究工作加快 100 倍

我們打造了使用案例專屬解決方案的產品組合

使用案例和垂直行業 AI 解決方案

使用案例：

使用
影像分析進行監控




垂直應用：

- 機場監控
- 臉部辨識
 - 佇列監控
 - 無人看管項目

對企業的影響：
因資源利用率提高，縮減高達 **28%** 的人力

品質控管
與目標性維護



- 製造品質控管
- 識別瑕疵

瑕疵偵測率提高至 **90%**

- 目標性工業
維護根據設備狀況

減少高達 **30%** 的維護費用
系統故障減少高達 **75%**

HPE Apollo purpose-built portfolio for HPC, Big Data, and AI

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

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

Storage

**HPE Apollo 4200 Gen9
(2U / 28 LFF or 54 SFF)**



Purpose-built for data analytics
and as enterprise bridge to
scale-out storage

**HPE Apollo 4510 Gen10
(4U / 1 node x 60+8 LFF)**







Purpose-built for
object storage

使用 NVIDIA GPU 的解決方案 最適合用於訓練與推論

適用於 AI/ML 的
水平 AI 基礎架構和服務

包括原廠安裝、實作和 3 年全年無休的 HPE Pointnext 積極主動型服務
包括適用於資料科學的 NVIDIA GPU 和 Bright Computing 軟體

入門級	多 GPU	多節點	生產
		 	 
  ProLiant DL380 Gen10	  HPE Apollo 6500 Gen10	 Mellanox 交換器   2 個 HPE Apollo 2000 Gen10，含 Weka.IO    2 個 HPE Apollo 6500 Gen10	 Mellanox 100GbE 交換器   2 個 HPE Apollo 2000 Gen10，含 Weka.IO      4 個 HPE Apollo 6500 Gen10

Accelerated performance for GPU intensive workloads

Reliable deep learning platform with industry leading GPUs and interconnects

HPE's highest GPUs per server

- Up to 125 Tflops¹ single precision performance
- Eight GPUs per server

Powerful host

- 2x the high speed fabric adapters²
- NVMe drives
- High speed DDR4 SmartMemory



Enterprise AI

Leading accelerator technology

- NVLink 2.0 enables dedicated GPU-to-GPU communication

Dependable system performance

- Power and cooling headroom for top bin accelerators
- Robust signal integrity provides higher bandwidth, lower latency

Flexibility for deep learning and HPC environments

Range of configuration and topology options to suit your workloads

Accelerator choice

- NVLINK 2.0 for maximum bandwidth
- PCIe for traditional GPU support

Flexible support and options

- Many high-speed fabric options
- Ubuntu and Enterprise OS choice
- HPE Pointnext support flexibility



Deep learning and HPC

Accelerator topologies

- Efficient hybrid cube-mesh for NVLink
- 4:1 or 8:1 GPU:CPU flexibility in PCIe

Storage options

- Up to 16 devices
- SAS/SATA SSDs, up to 4 NVMe drives¹

Resilient, secure and simple for lower TCO


Designed for ease of use by enterprises

Easy to service, upgrade

- Easy access modular design
- Rear cabled fabrics

Efficient management, security

- Save time and cost, gain improved user productivity with HPE iLO5
- World's most secure industry standard server using HPE iLO5



Designed for you

Resilient power, cooling


- 2+2 power for resiliency
- Redundant, dual rotor GPU fan modules for zoned cooling

Standard rack mount

- Designed for 1m deep rack
- No special flooring, plumbing needs

HPE Software environment for deep learning

Ease of deployment and flexibility for deep learning solutions



Solution Components

- Bright provides comprehensive software to spin up an effective deep learning environment
- Bright Cluster Manager provides easy installation and setup of the following on an HPE cluster:
 - Linux OS of choice
 - Bright Machine Learning and Big Data software
 - Workload management tool of choice
- Bright Machine Learning and Big Data Software supports Ubuntu, SLES, RHEL and CentOS
- Bright validates the necessary software components for the deep learning environment so developers can focus on the application

Specification Comparison (1/2)


Specifications	NVIDIA DGX-1		HPE Apollo 6500 Gen10	HPE Benefits/Other Comments
GPUs	8X Tesla V100	8X Tesla P100	Up to 8 PCIe or SXM-2 NVLink 2.0 GPU	HPE Apollo 6500 Gen10 offers a broader accelerator support with NVIDIA V100, P100, P40, PCIe. NVIDIA DGX-1 supports just V100 and P100.
CPU	Dual 20-Core Intel Xeon E5-2698 v4 2.2 GHz		Intel® Xeon® Scalable Processors (8100 and 6100 series)	HPE Apollo 6500 is fabricated to support a higher wattage and the latest CPU technology with Xeon Scalable Processors-up to 28 cores or 3.0GHz while the DGX-1 can go up to just 20 cores or 2.2GHz. HPE offers greater flexibility by providing various processor options supporting the Intel Xeon Scalable Processor 8100, and 6100 series. NVIDIA supports only the Intel Xeon E5-2698 v4 processor.
Maximum Power Requirements	Up to 4 load balancing 3,200 W power supplies (3+1)		Up to 4 HPE 2200W Platinum 80 Plus (2+2)	HPE Apollo 6500 Gen10 offers more power efficiency with just 2200W and Platinum 80 Plus class power supplies
System Memory	512 GB 2,133 MHz DDR4 LRDIMM		Supports up to 2666MT/s DDR4 SmartMemory 3TB Max, 24 DIMM slots	HPE offers a superior performance as compared to NVIDIA as HPE Apollo 6500 offers 3TB of memory which DGX-1 offering just about 512GB
Storage	4X 1.92 TB SSD RAID 0		Up to 16 SAS ¹ /SATA SSD, or up to 4 NVME SSD (optional), M.2	HPE offers more drive options with SAS ¹ /SATA/SSD/NVME SSD drives and NVIDIA supports only SSD. Also HPE offers more drive capacity leveraging the popular DL380 Gen10 drive cage.

Specification Comparison (2/2)

Specification	NVIDIA DGX-1	HPE Apollo 6500 Gen10	HPE Benefits/Other Comments
Network	Dual 10 GbE, 4 IB EDR	Four 1 GbE single-port RJ-45 Four x 16 PCI Express slots supporting InfiniBand EDR or Intel OPA or high- speed Ethernet up to 100 Gbps	HPE offers a broad networking option than NVIDIA
Bandwidth	300 GB/s NVLINK	300 GB/s NVLINK 2.0	
Software	Ubuntu Linux Host OS	Ubuntu, Red Hat, SLES or CentOS	HPE supports a broad set of Operating Systems while NVIDIA support is restricted to just one OS which is often not considered “Enterprise class”
System Weight	134 lbs.	130 lbs. / 60 kg	
System Dimensions	866 D x 444 W x 131 H (mm)	440 x 175.8 x 850.3 (mm)	
Operating Temperature Range	10 to 35 °C	10° to 35° C	
Warranty	1yr.	3/3/3 yrs.	HPE offers better protection to customers with a 3/3/3 yrs. warranty coverage for parts, labor and onsite support respectively. On the contrary NVIDIA offers only one year warranty period for its customers.

iLO5 System Management for ProLiant Gen10 Systems provides enterprise level RAS features

- **Security:** Raising our industry-leading bar even higher
- **Firmware update technology:** Update everything through iLO network
- **Agentless management:** Retiring OS-based agents
- **RESTful everything:** Redfish
- **At Server Management:** New innovation
- **Performance:** Everything faster



NVIDIA DGX

NVIDIA lacks an equivalent system management suite

使用 Intel CPU 和工具的解決方案 最適合用於推論

適用於 AI/ML 的
水平 AI 基礎架構和服務

可選擇的作業系統



CentOS



ubuntu®



系統管理 (內含)



Hewlett Packard
Enterprise

HPE Performance
Cluster Manager



Altair | PBS Works™

最佳化工具和資源 (內含)



Intel 的深度學習 Software
Optimization MKL-DNN
OpenSource 架構



Hewlett Packard
Enterprise

HPE 深度學習實踐手冊

AI 推論組合

- 特定搭售組合 (S、M、L)，搭配自選運算¹
- Intel® Xeon® 可擴充處理器
- Intel® Omni-Path 架構



使用記憶體內運算的進階分析解決方案

HPE Superdome Flex

- 配備 Intel® Xeon® 可擴充處理器，每個處理器最高可達 28 個核心
- 提供 768GB-48TB 的共用記憶體
- 包括最高可達 16 個 NVIDIA GPU

記憶體內資料庫



機器學習

- 強大的資料互依性
- 串流分析

技術 HPC 應用程式

- 基因體學
- 網路安全性
- 電腦輔助工程
- 詐騙偵測
- 風險管理

透過單一系統簡潔性，全面性解決複雜問題，並具有無可比擬的擴充性

適用於 AI/ML 的
水平 AI 基礎架構和服務

以單一系統形式，可
將插槽從 4 個無縫擴
充至 32 個

5U，4 插槽機箱




使用 Apollo 的巨量資料分析解決方案

適用於 AI/ML 的水平 AI 基礎架構和服務



可選擇代理商、ISV



系統管理 (內含)




最佳化工具和資源 (內含)



Elastic Platform for Analytics
建置區塊


- 以模組化、彈性的架構/願景，打造邊緣到雲端資料管道
- 驗證並最佳化工作負載
- 彈性、獨立擴充



運算區塊
Apollo 2600 機箱搭載
Proliant XL170r Gen10 伺服器




儲存區塊
Apollo 4200 Gen10



AI 區塊
Apollo 6500 Gen10

HPE's Elastic Platform for Analytics (EPA)


HPE's Architecture for Big Data Analytics






Hewlett Packard
Enterprise

HPE INFOSIGHT







Apps

Data


Loading...




HPE INFOSIGHT: INDUSTRY'S MOST ADVANCED AI FOR INFRASTRUCTURE



SEE ONCE, PREVENT FOR ALL



FULL-STACK ANALYTICS FOR VMWARE ENVIRONMENTS



Noisy Neighbor

Determine if VMs are hogging resources from another VM

Host & Memory Analytics

Visibility into host CPU and memory metrics

Latency Attribution

Identify root cause across host, storage, or SAN

Inactive VMs

Visibility into inactive Vms to repurpose/reclaim resources

Top Performing VMs

Visibility into top 10 VMs by IOPs and Latency

HPE INFOSIGHT COVERAGE ACROSS THE HPE PORTFOLIO



AUTONOMOUS AND CONTEXT-AWARE

Storage

Servers

Networking

Integrated



HPE Nimble Storage



HPE 3PAR



HPE Primera



HPE Servers



HPE Synergy



HPE SimpliVity



HPE Nimble Storage dHCI





PROACTIVELY INFORM CUSTOMERS ABOUT CUSTOMER ADVISORIES

Customer Experience: Supportability

I want to see an alert in OneView when HPE issues a Customer Advisory or Security Bulletin that is relevant to my specific infrastructure

Reduces downtime and testing time, as well as simplifying troubleshooting by freeing customers from manual review of every HPE advisory to determine relevance in their environment.



謝謝