Hewlett Packard Enterprise

探索HPE GPU/AI 工作負載儲存方案

蔡宜家

技術規劃處技術經理
tsai@hpe.com

HEP Taiwan
慧與科技股份有限公司

如何決定儲存方案

必須取決的因素

AI workflow stages

Data aggregation, data prep, training, tuning, inference

AI workloads

Gen AI, computer vision, natural language processing (NLP), reinforcement learning, speech recognition

Cost

符合運用需求
功能, 效能, 擴充性

Hewlett Packard Enterprise

敦新科技

AI WORKFLOW STAGES

	Stored data	Capacity	Write perf	Read perf	Price	Features
Inference	• Raw data • Prepped data • Model				**$$**	• Availability • GPU direct storage • Data reduction
Training and tuning	• Prepped data • Model • Checkpoints				**$$$$**	• GPU direct storage • Availability • Resiliency
Data prep	• Multiple copies of prepped data • Backups of model version				**$$**	• Unified access • Snapshots, clones • Data reduction
Data aggregation	• Raw data in different formats and sources				**$**	• Single namespace • Unified access • Resiliency

Hewlett Packard Enterprise

敦新科技

AI WORKLOADS

	Computer vision	Generative AI multimodal	Generative AI text	Reinforcement learning	NLP	Speech recognition	Traditional AI
	Detect patterns on images/videos.	Generate text, image, video, audio.	Generate text from user prompt.	Find optimal decision trees.	"Understand" human language.	Synthesize or transcribe speech.	Non machine learning (ML)–type AI
Model examples	Waymo	GPT-4, DALL-E, Stable Diffusion	GPT-3, Llama-2	AlphaGO	BERT	Kaldi	Expert systems like IBM DeepBlue
Use case examples	• Theft detection • Auto checkout • Medical imaging	• Marketing content • Game design	• Chatbot • Code generation • Knowledge mgmt.	• Fraud detection • Inventory optimization	• Predictive text • Sentiment analysis • Basic chatbot	• Meeting transcription	• Trend analysis on time series
Data type	**Image | Video** Relatively few large data points	**Text | Audio** **Image | Video** Relatively few large data points	**Text** Many small data points	**Text** Several or many small data points	**Text** Several small data points	**Text | Audio** Several medium data points	**Text** Relatively few small data points
Data prep size	PBs	PBs	TBs to low PBs	GBs to PBs	GBs to TBs	GBs	Low GBs
Training size	TBs	Low PBs	TBs	Low GBs to PBs	GBs to TBs	GBs	Low GBs
Parameters	<500M	Up to a trillion +		<500M			<10K

儲存應該與AI WORKLOADS進行最適當配置

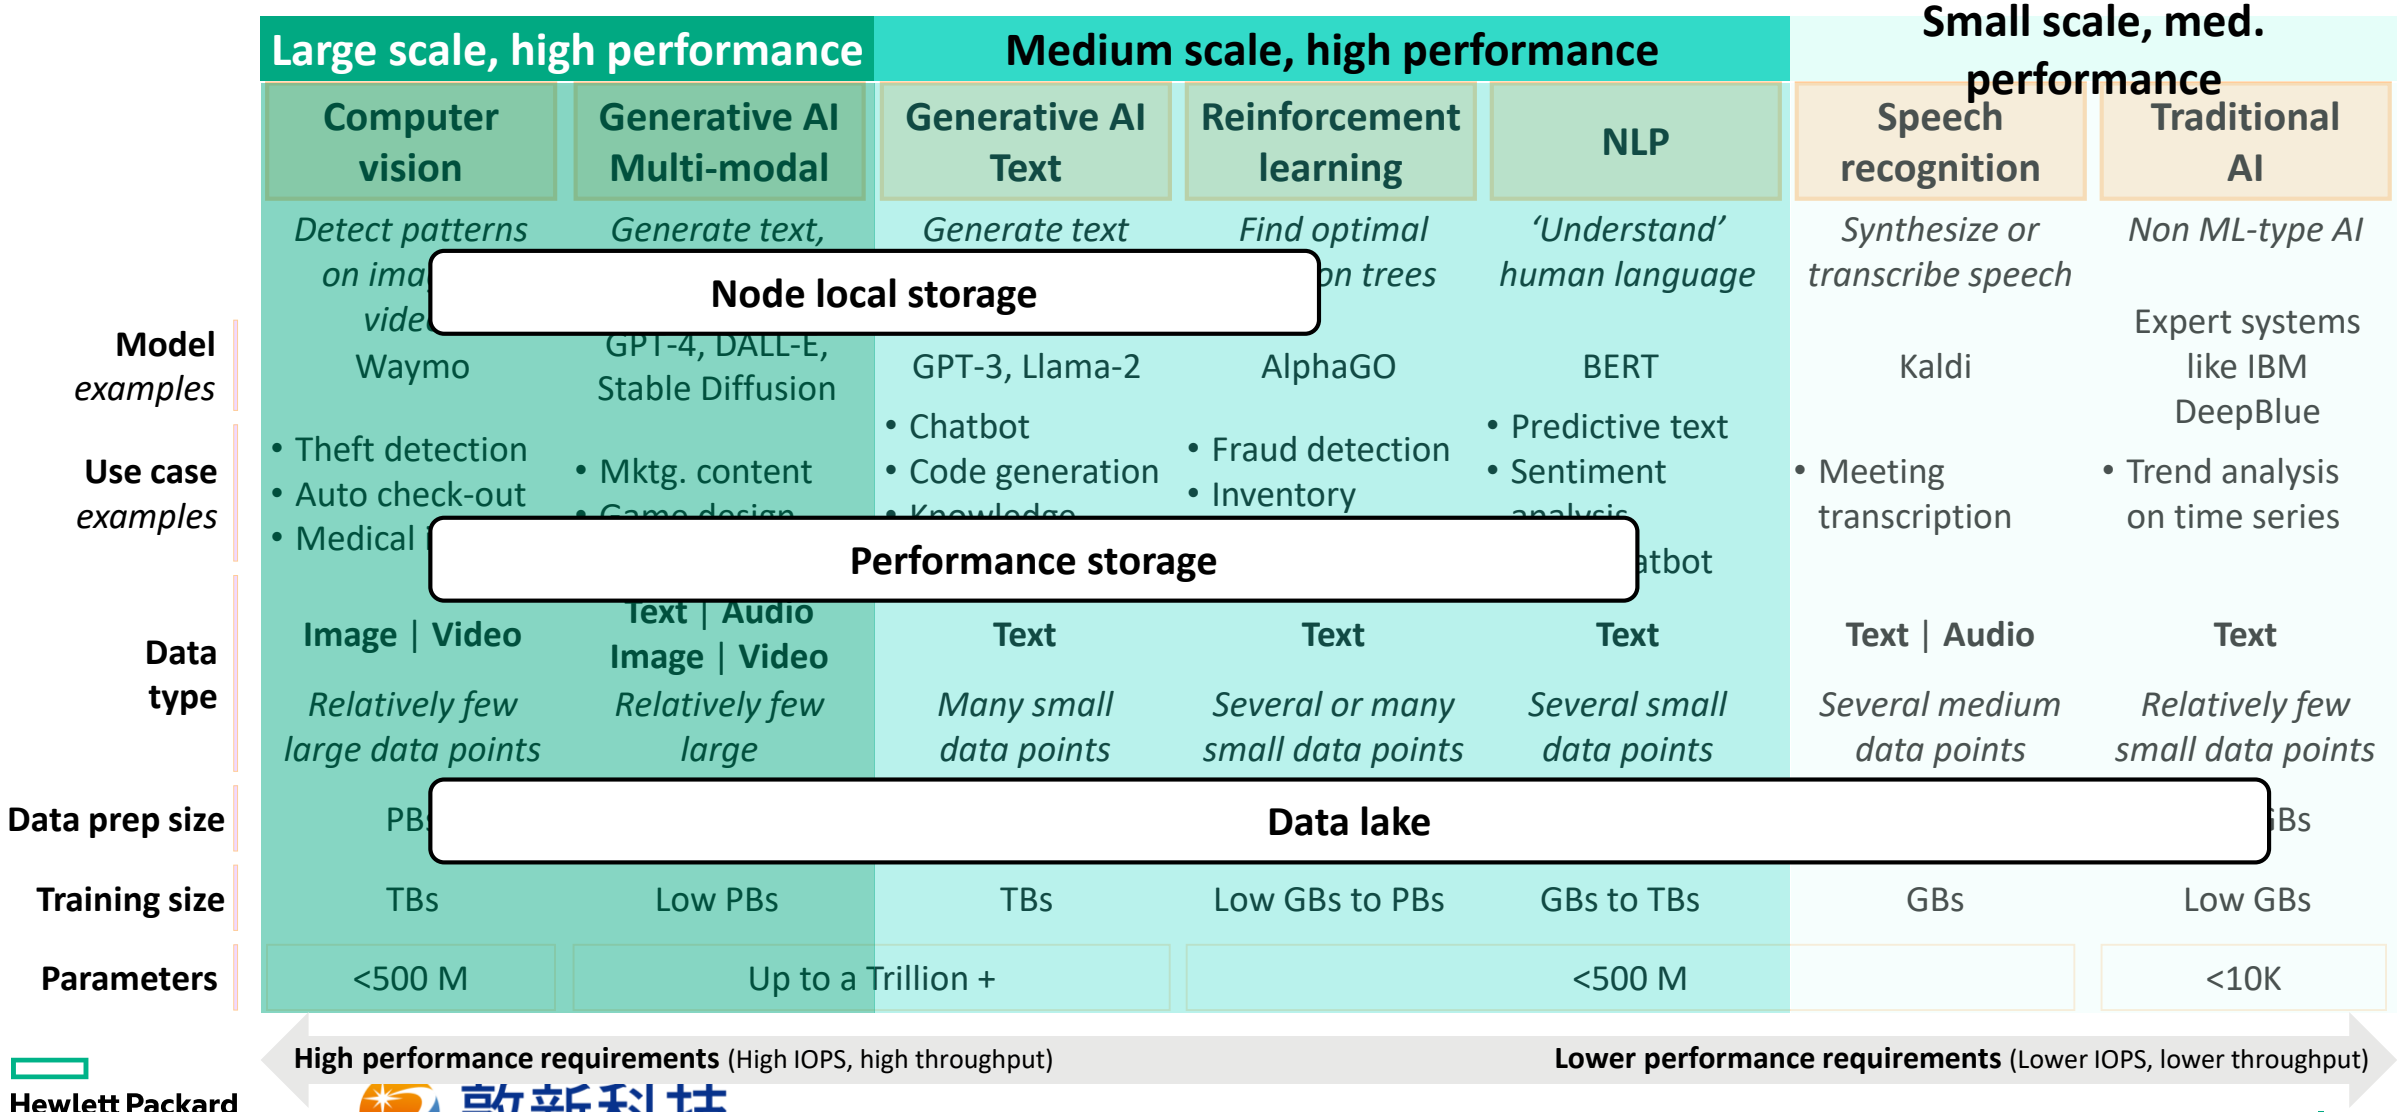

Hewlett Packard Enterprise

敦新科技

HPE 提供AI演算進行最佳儲存搭配

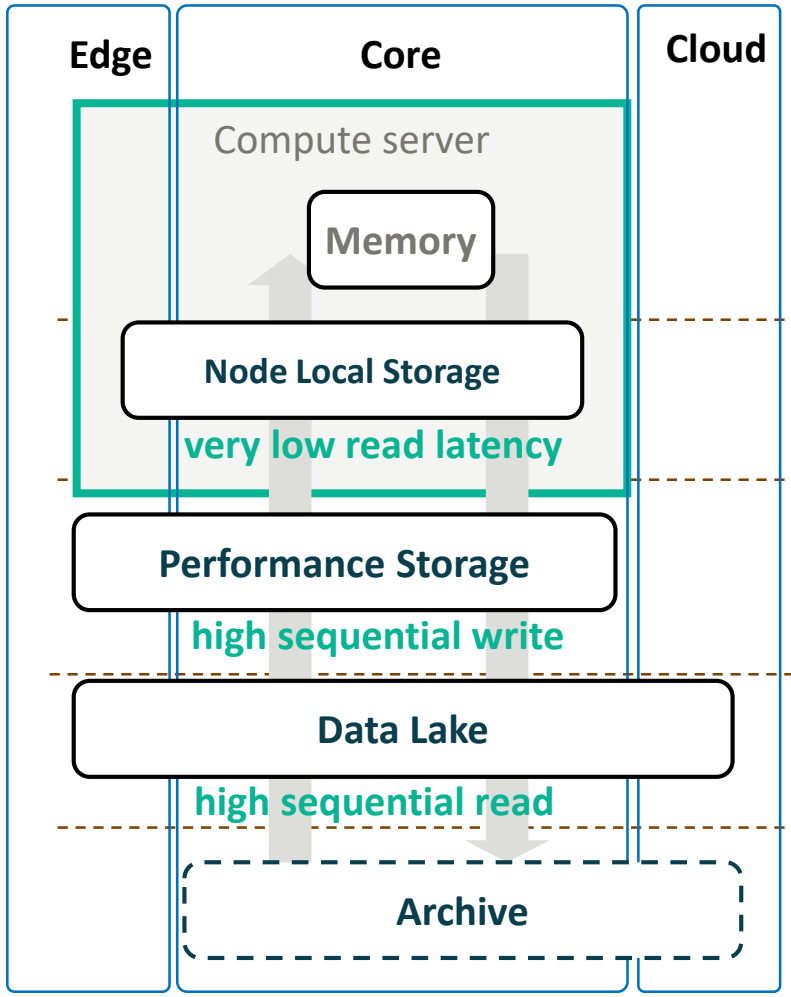

ClusterStor	WEKA	GL 4 File VAST	Qumulo	SCALITY
	Provides a converged option			
Provides parallel FS write and read performance	*Provides parallel FS write and read performance*	*Falls short in write performance, required during training*	*Falls short in write performance, required during training*	
	Expensive as a DL in large configs	*Expensive as a DL in large configs*	*Provides Flash and HDD solutions*	
			Provides Archive Node on HDD	*Scality Ring*

Hewlett Packard Enterprise

敦新科技

WEKA Overview

HPE Solutions For Weka Parallel File system

WEKA 一次解決AI儲存所有挑戰

- Improves productivity and faster time to market and value
  - Accelerate large scale data pipelines
- Reduced epoch times
- Fastest inferencing
- Highest images / secs benchmarks
  - Run entire pipeline on the same storage backend
  - Multi Protocol Support
  - Faster than local storage

Graphic courtesy: WEKA

Hewlett Packard
Enterprise

Weka 效能優勢讓AI演算從循序模式轉換為平行模式成為可能

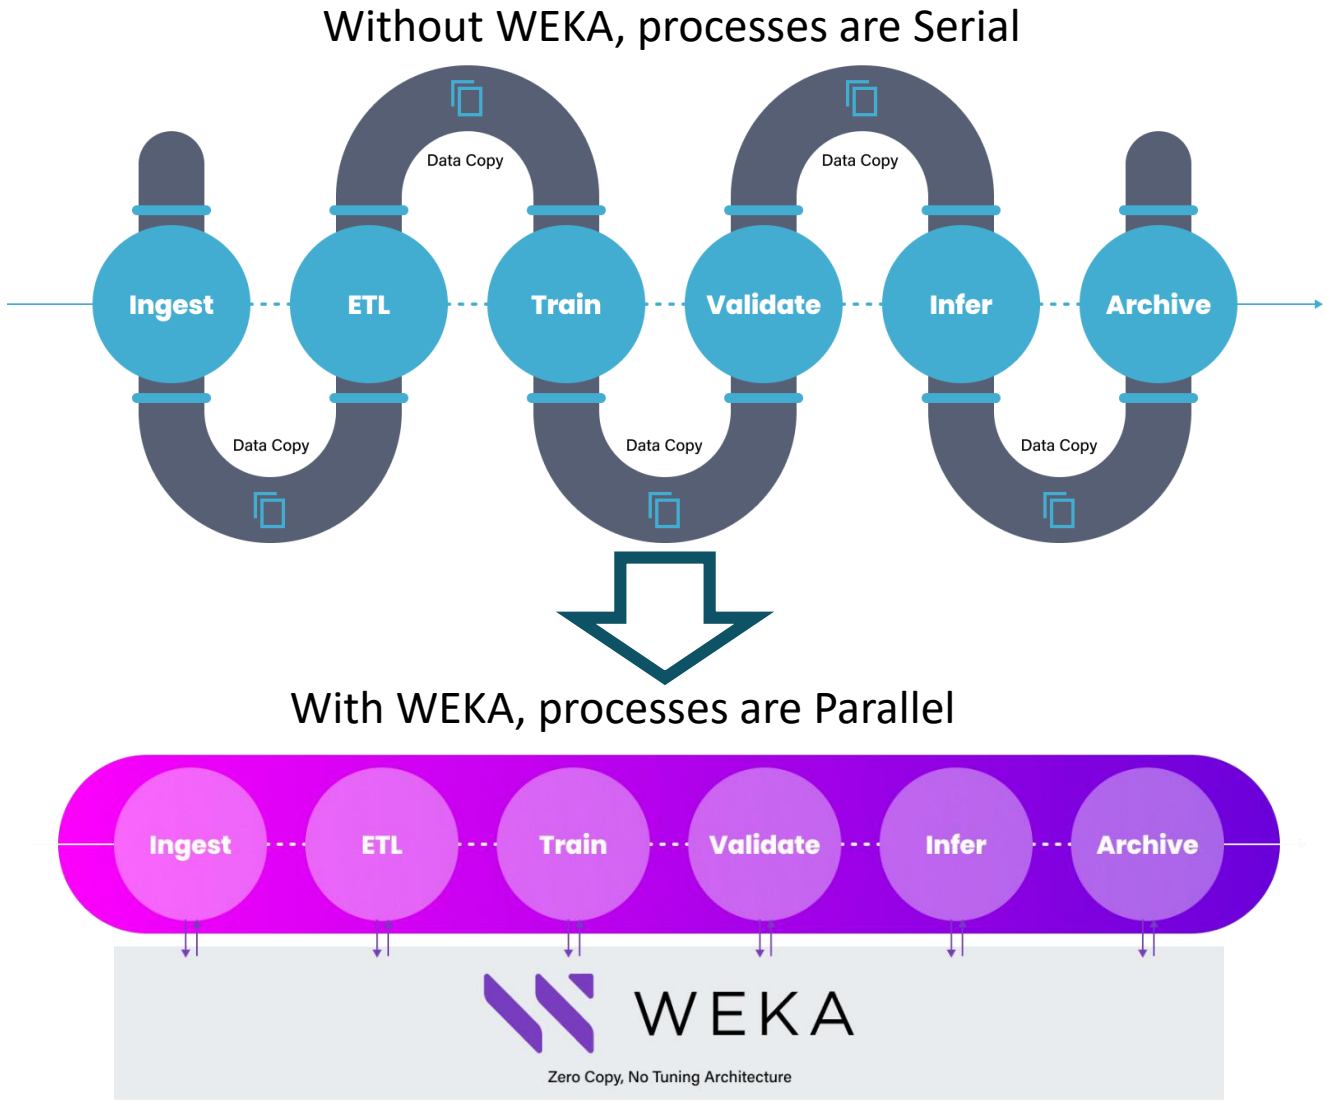

HPE 提供彈性佈署優勢

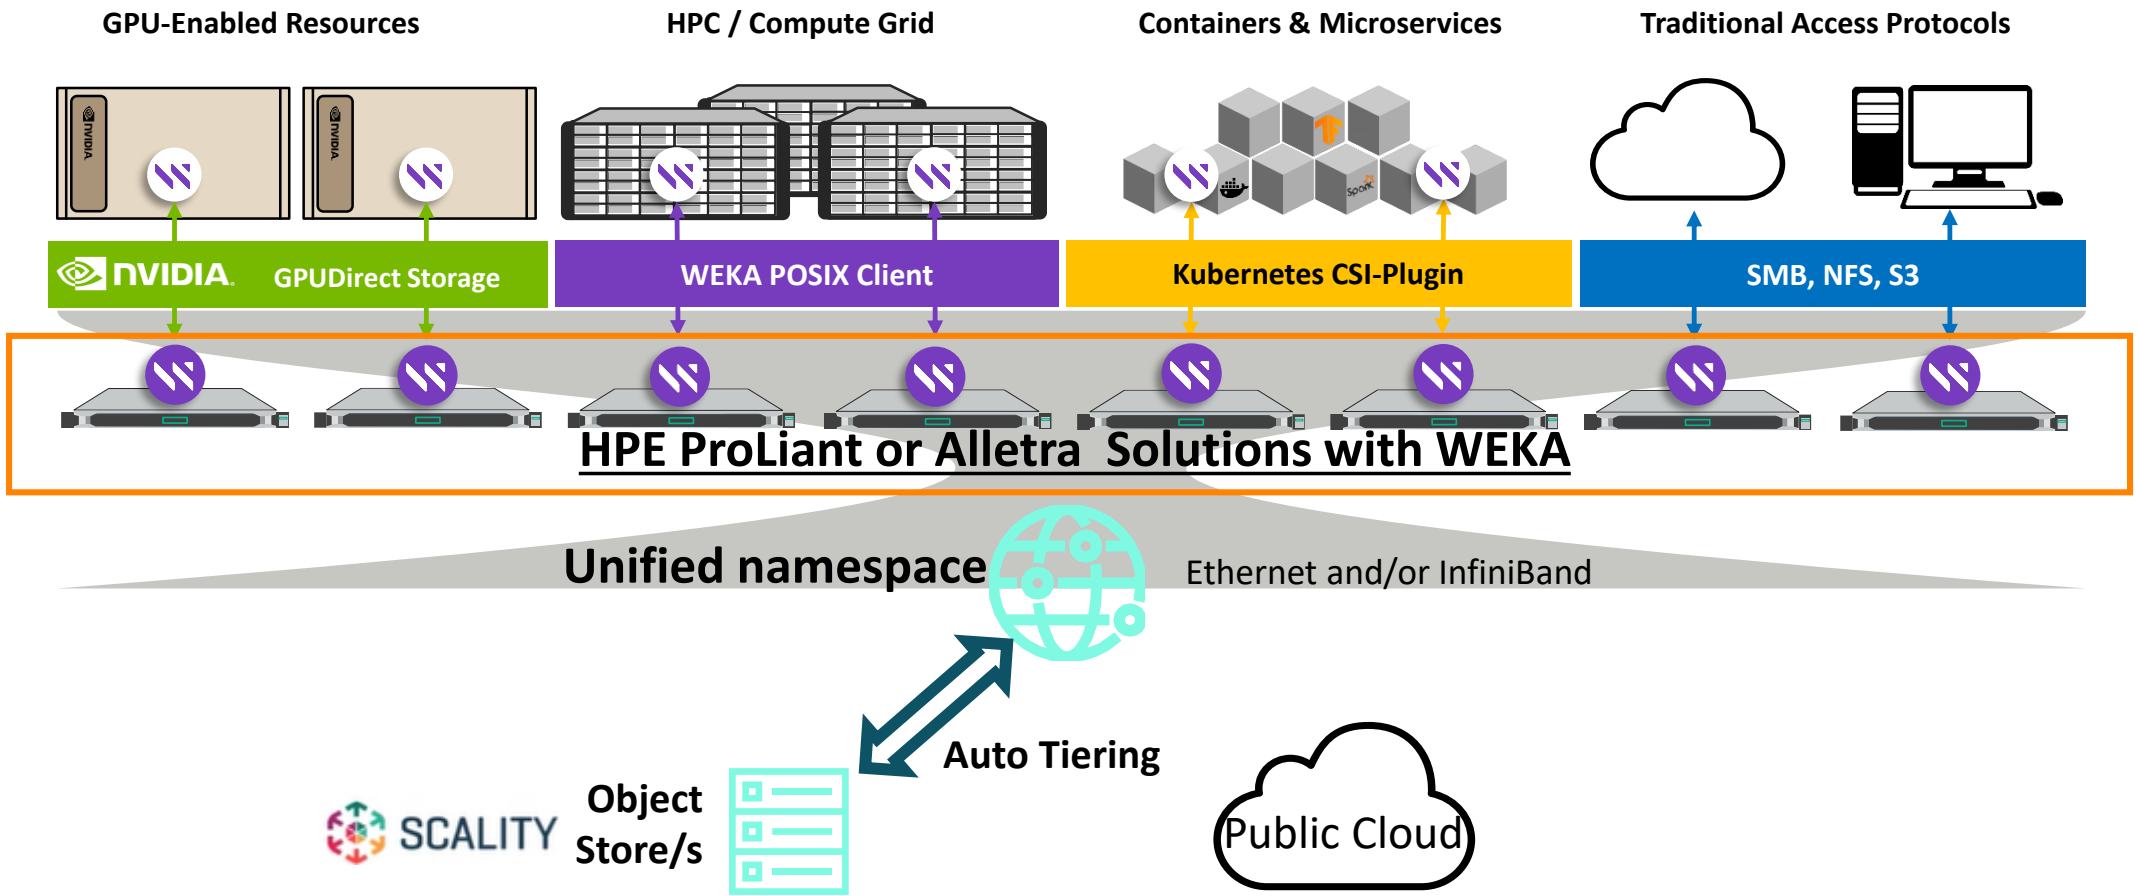

Scale Out to Cost-Effective Object Data Lake and Burst to Public Cloud

Hewlett Packard Enterprise

HPE GreenLake for File Storage (GL4F) Overview

HPE Parallel File system

Introducing HPE GreenLake for

File Storage

Enterprise grade, scale-out File storage to supercharge data-intensive workloads

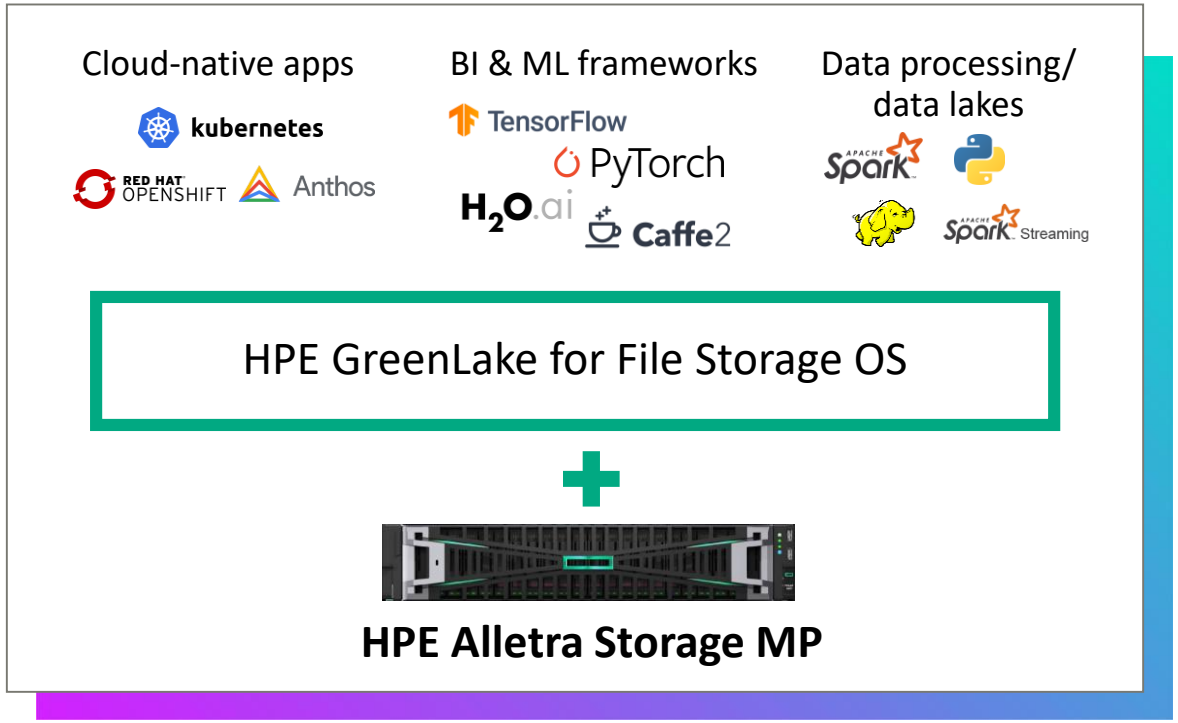

Accelerate
with enterprise performance at scale

Simplify
with an intuitive cloud experience

Enhance
productivity with faster time to insights

HPE GreenLake for File Storage (GL4F)

Performance spanning the scale of your data

Accelerate your most data-intensive applications, including HPC, Media, AI/ML, Life Sciences

Exabyte scale architecture

Future-proof your storage for data growth

Architecture designed for HA

Gain higher resiliency with no rebuild times for controller failures

Unique modular storage infrastructure

Scale performance and capacity independently

DASE™ architecture

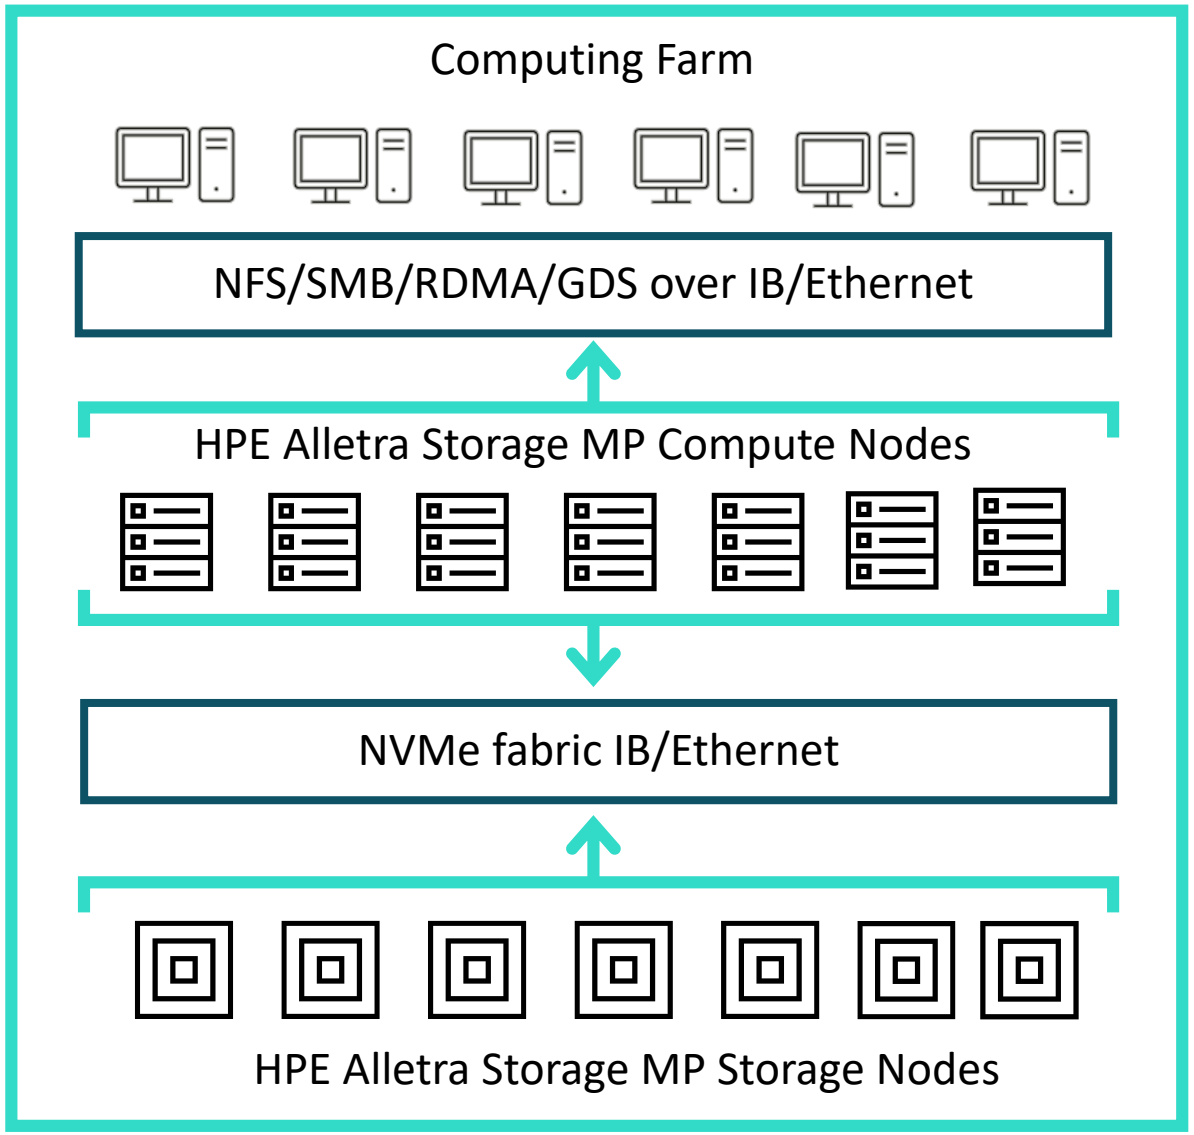

Hewlett Packard Enterprise

資料減量新紀元 - GL4F Similarity 演算法

Compression

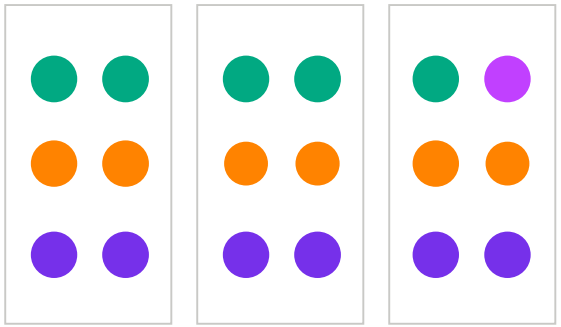

Fine-grained, but local

Deduplication

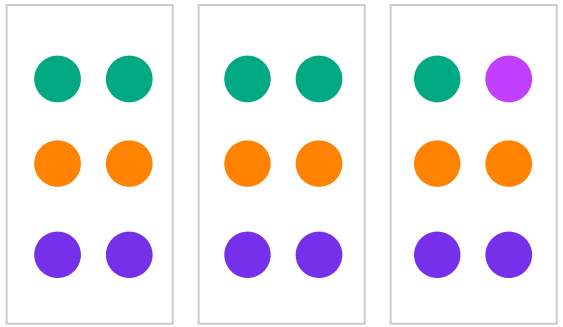

Global, but coarse

Similarity reduction

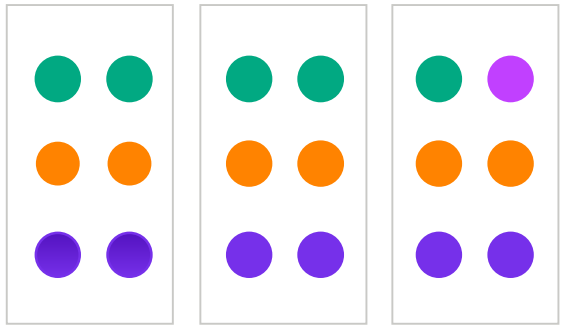

Global and fine-grained

Example savings from similarity

3:1 Pre-reduced backups

3:1 Pre-compressed Log files

2:1 Life science data

3:1 HPC data

3:1 Animation data

8:1 Uncompressed time-series data

HPE Storage Substantiation provides substantiation for all the data reduction ratios shown.

Overview

HPE Solutions For Qumulo Distributed File system

Qumulo Hybrid architecture

Distributed File System

- Real world 93%+ cache hit rate
  - Performance close to AFF
- Flexible node options in 2U
  - 36TB, 90TB, 480TB, 1.2PB, 1.6PB
- 4KB granular block design
- EC ratio adjustable by adding node
- 100% Space utilization guarantee
- Against Cluster hardware failures
  - Driver protect from 2 to 4
  - Node protect from 1 to 4
- Lowest Per TB/PB cost

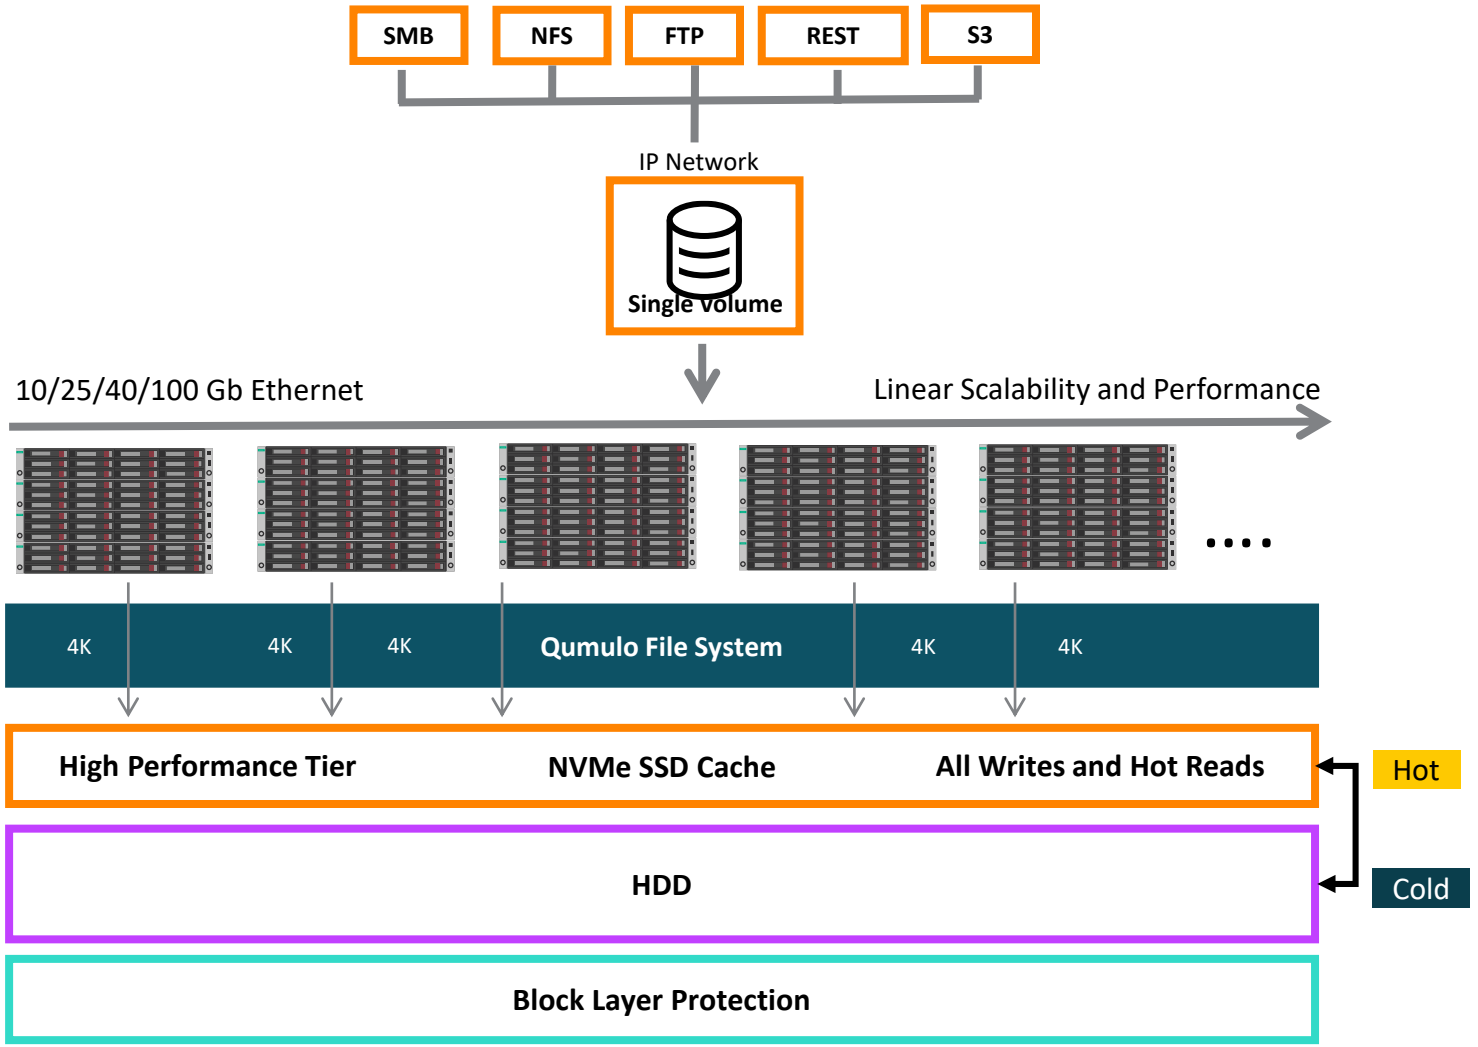

Qumulo All Flash architecture

Distributed File System

- Flexible node options in 1U
  - 38TB, 153TB, 307TB, 430TB, 737TB, 860TB
- 4KB granular block design
- EC ratio adjustable by adding node
- 100% Space utilization guarantee
- Against Cluster hardware failures
  - Driver protect from 2 to 4
  - Node protect from 1 to 4

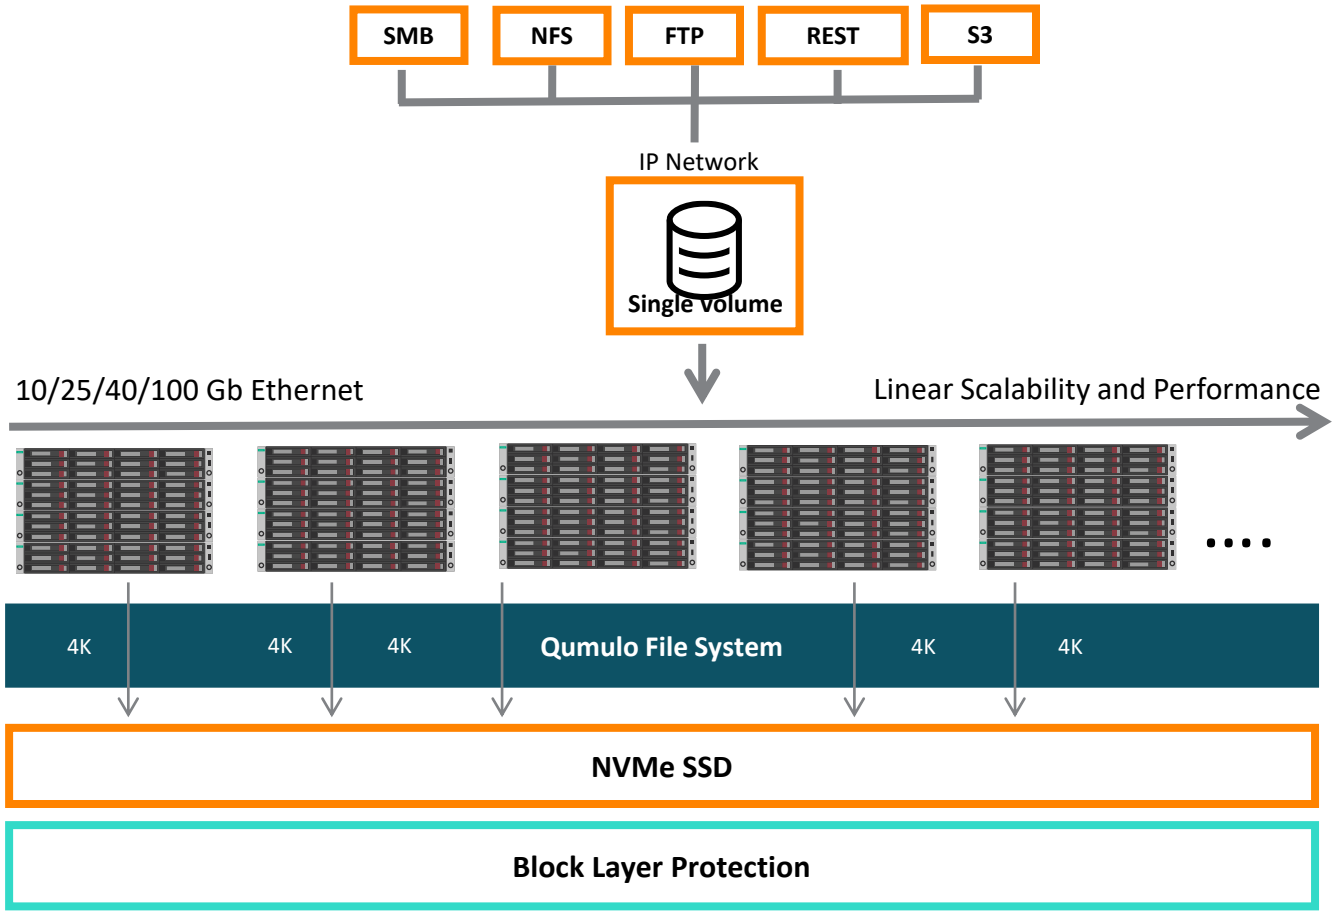

Qumulo無虛壓縮去重最佳小型檔案空間效益

Size on disk:

9.49 TB

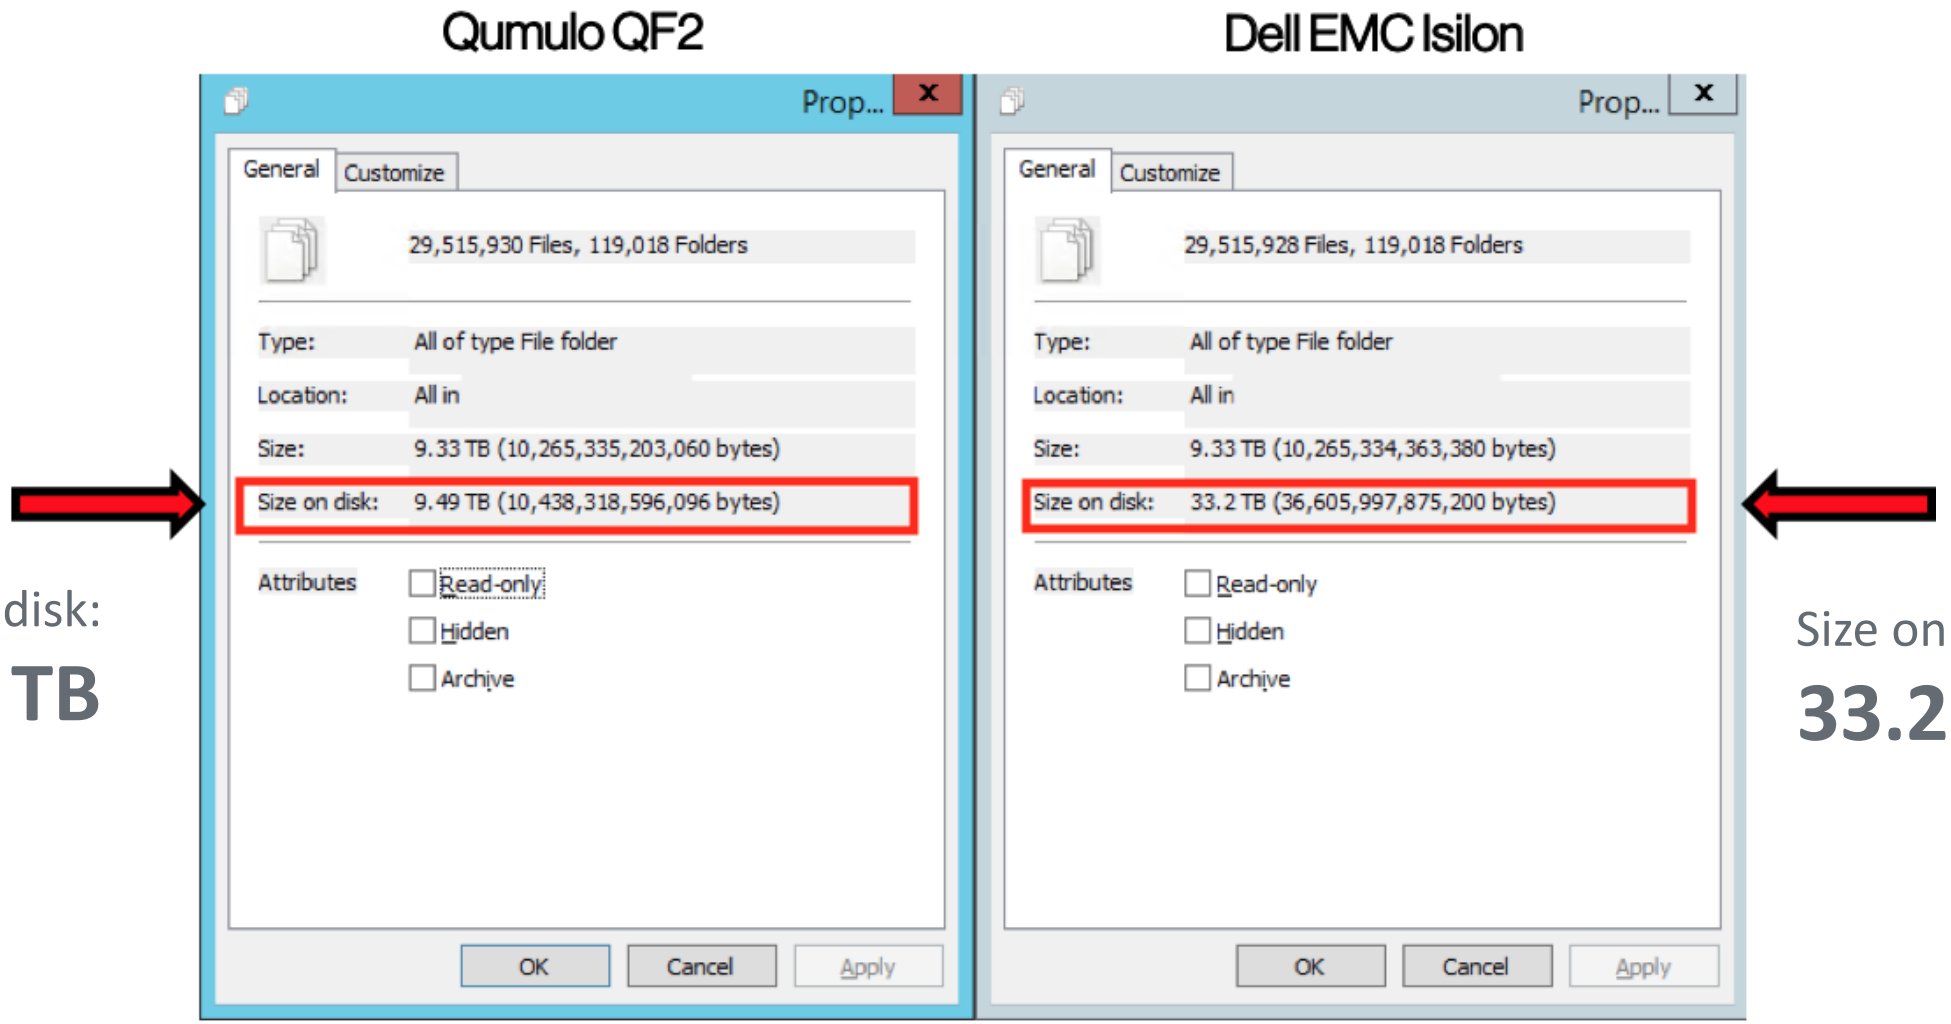

Size on disk:

33.2 TB

HPE 儲存滿足AI演算各階段效能及成本需求

Qumulo

- **Performance scale**
  - **AFF > 1,100GB/s**
  - **Hybrid > 790GB/s**
- 地端及多雲佈建全球存取
- NFS nconnect支援
- **Hybrid及All NVMe可選**
- **Hybrid單節點36TB-1.6PB**
- **100%空間使用保証**
- 即時效能分析
- 幾乎無限的檔案數量及大小
  - 18 Quintillion files
  - 9 **Exabyte file size**

HPE GL4F

- **Performance scale > 1,200GB/s**
- DASE架構
- NFS over RDMA driver
- **高密度設計 單一NVMe櫃1U可達1.3PB**
- GPU Direct支援
- **NVIDIA DGX SuperPOD™認證**
- NFS nconnect支援
- **支援Infiniband**
- **Similarity極致資料減量技術**
- **資料保護空間耗損最低僅3%**

WEKA

- **Performance scale > 17,000GB/s**
- **真正平行存取架構**
- **POSIX Client agent**
- **超低Latency設計**
- **地端及多雲佈建全球存**
- **支援S3 tiering架構模式**
- GPU Direct支援
- **支援NVIDIA DGX SuperPOD™**
- NFS nconnect支援
- **支援Infiniband**
- **支援去重壓縮**

HPE 提供極致效能演算儲存方案

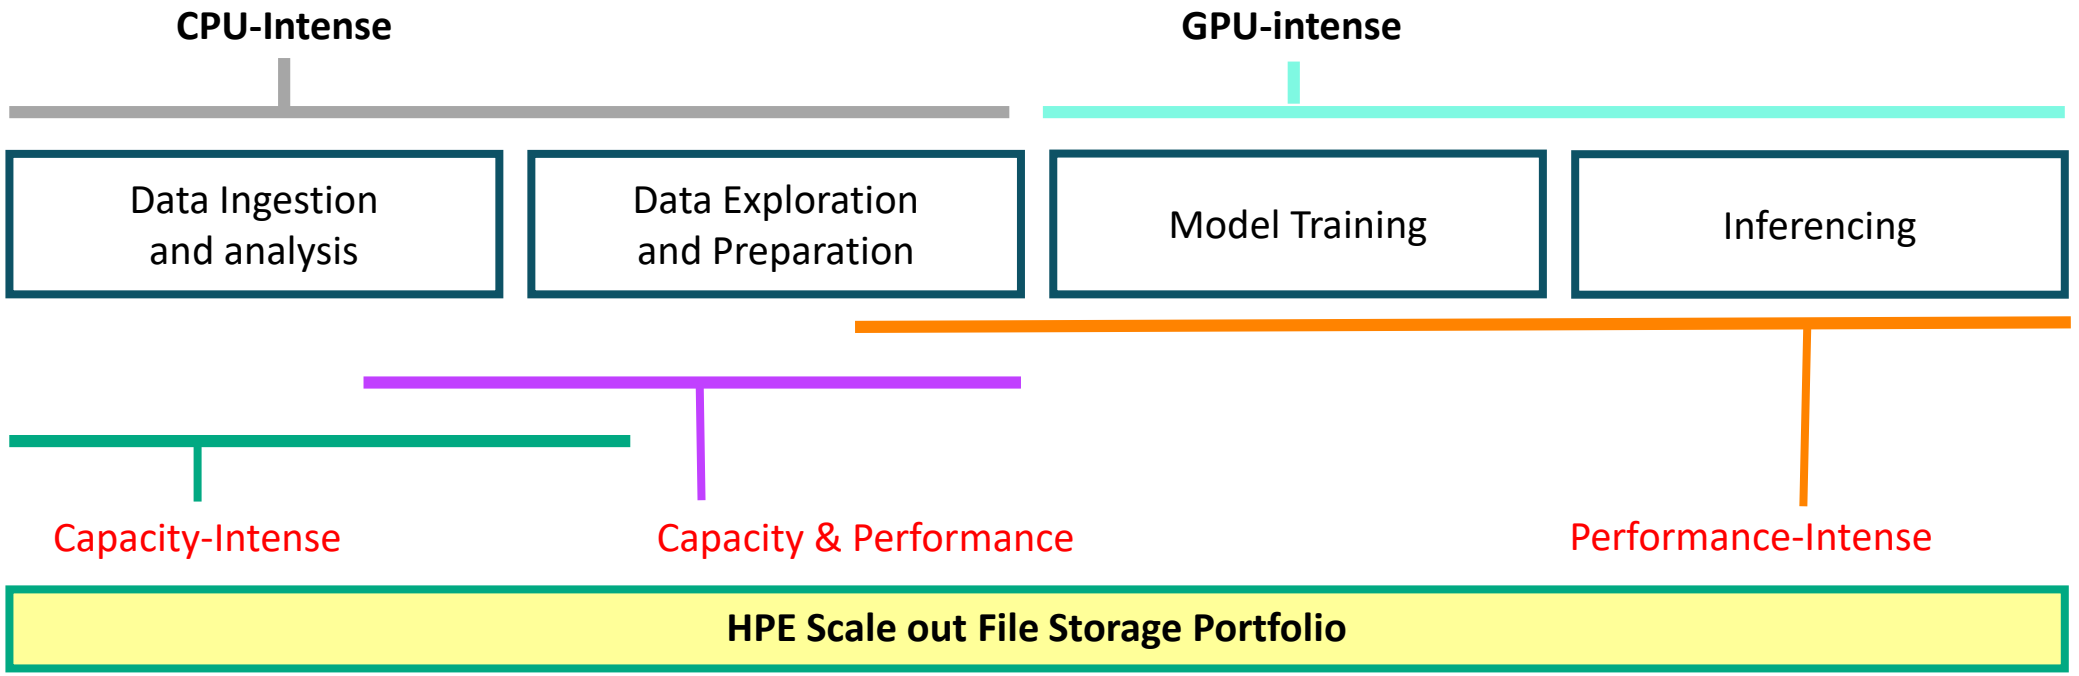

HPE GL4F

GDS ready

WEKA

GDS ready

Hewlett Packard Enterprise

敦新科技

HPE 提供效能成本最佳平衡演算儲存方案

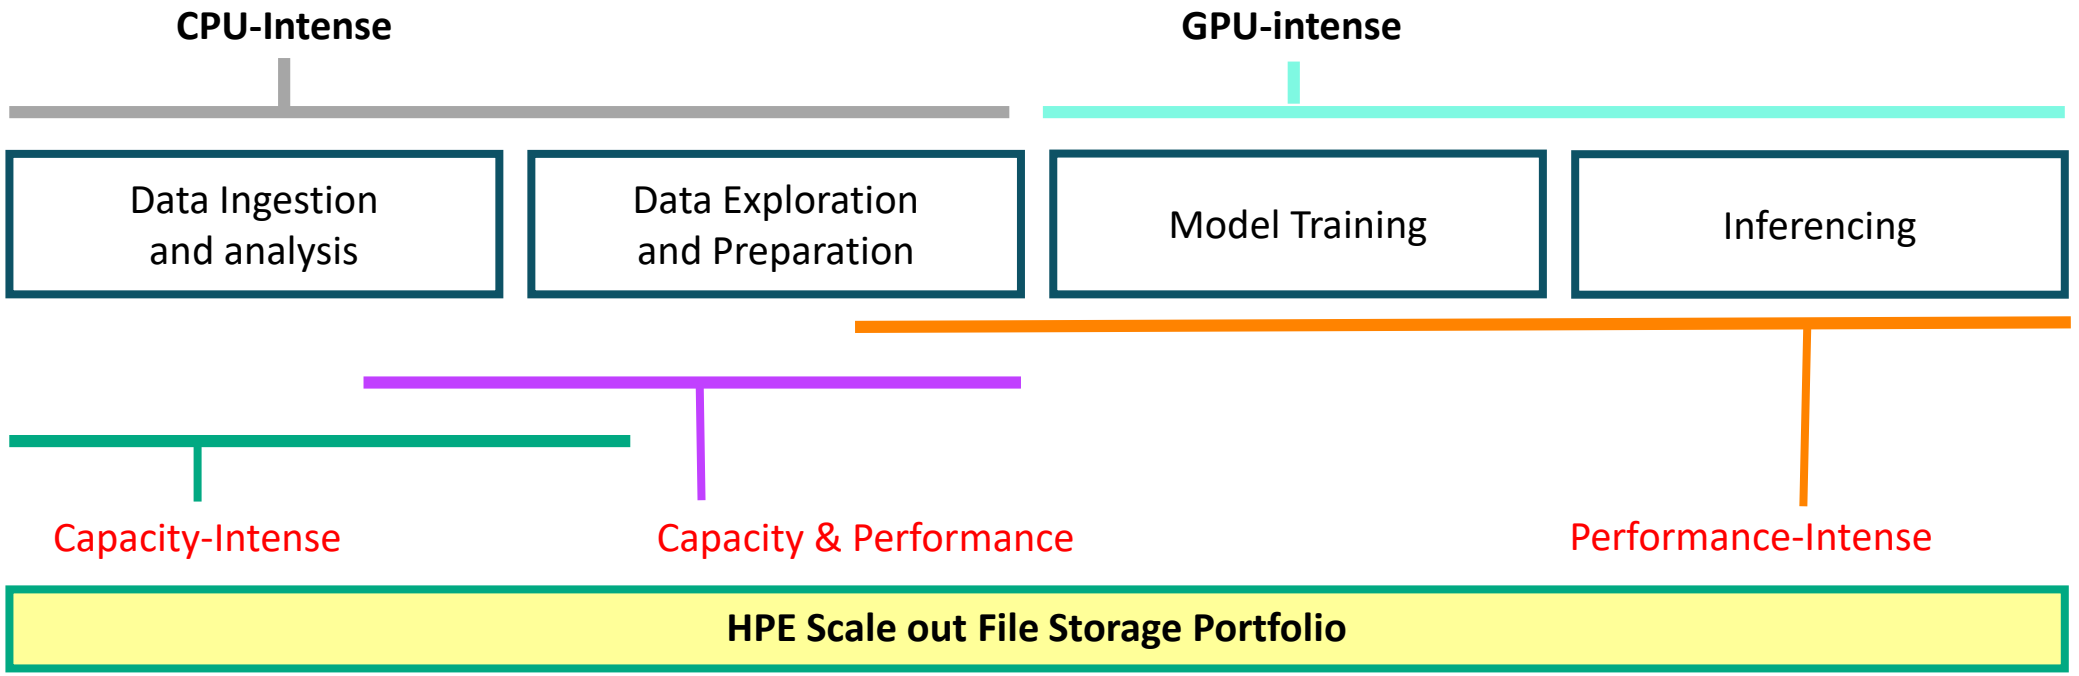

Qumulo

HPE GL4F

GDS ready

WEKA

GDS ready

Hewlett Packard Enterprise

敦新科技

HPE 提供最佳成本效益演算儲存方案

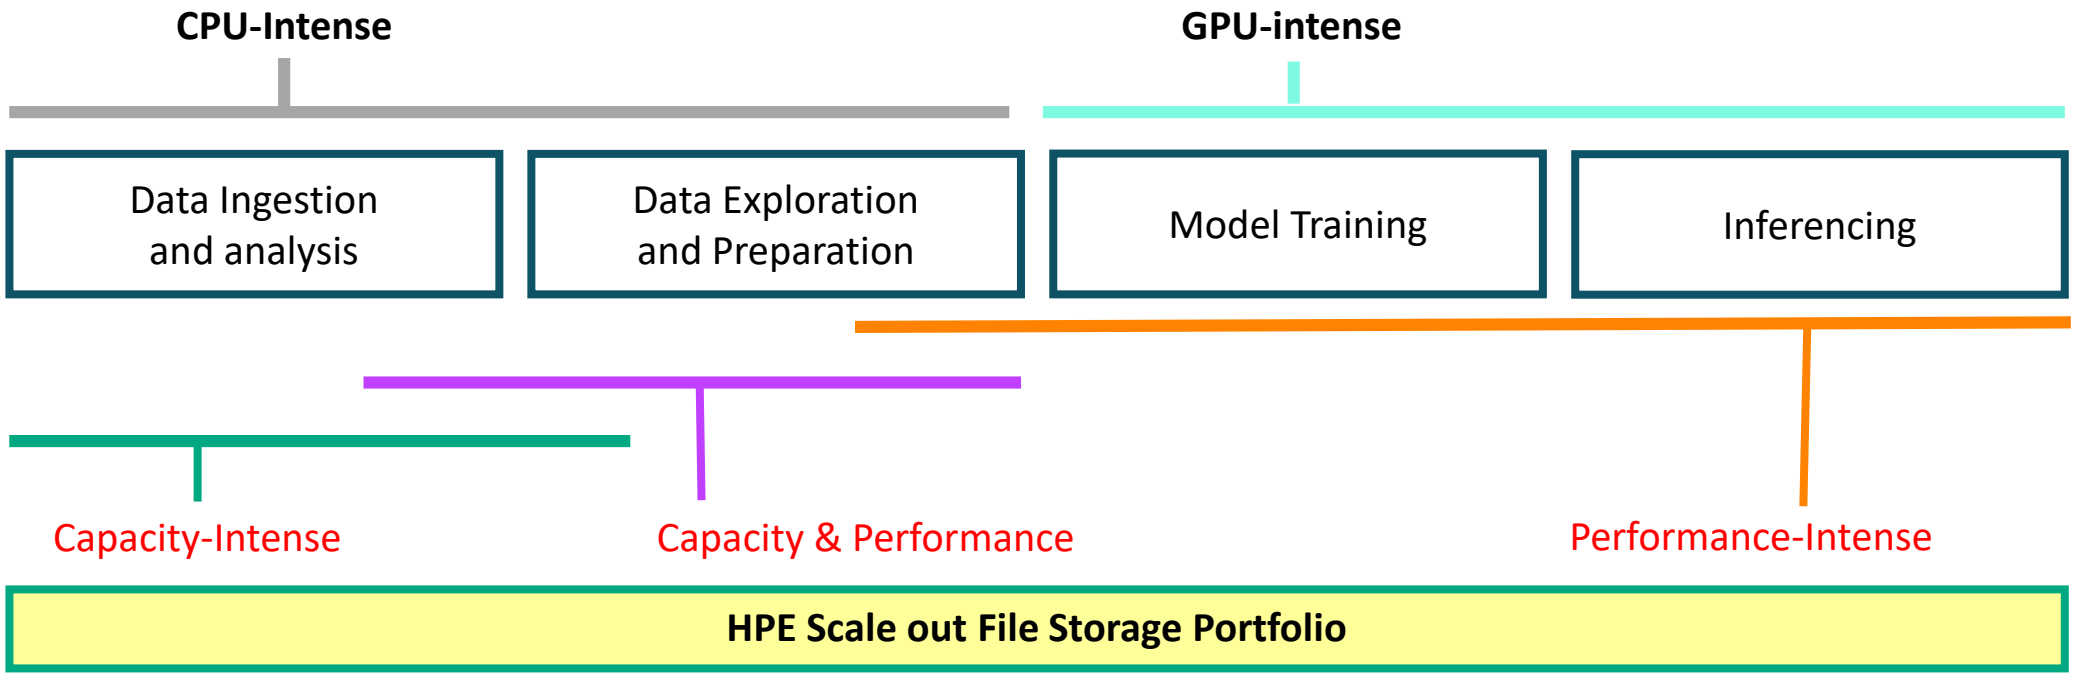

Qumulo

HPE GL4F

GDS ready

Hewlett Packard Enterprise

敦新科技

HPE 提供最低成本效能平衡演算儲存方案

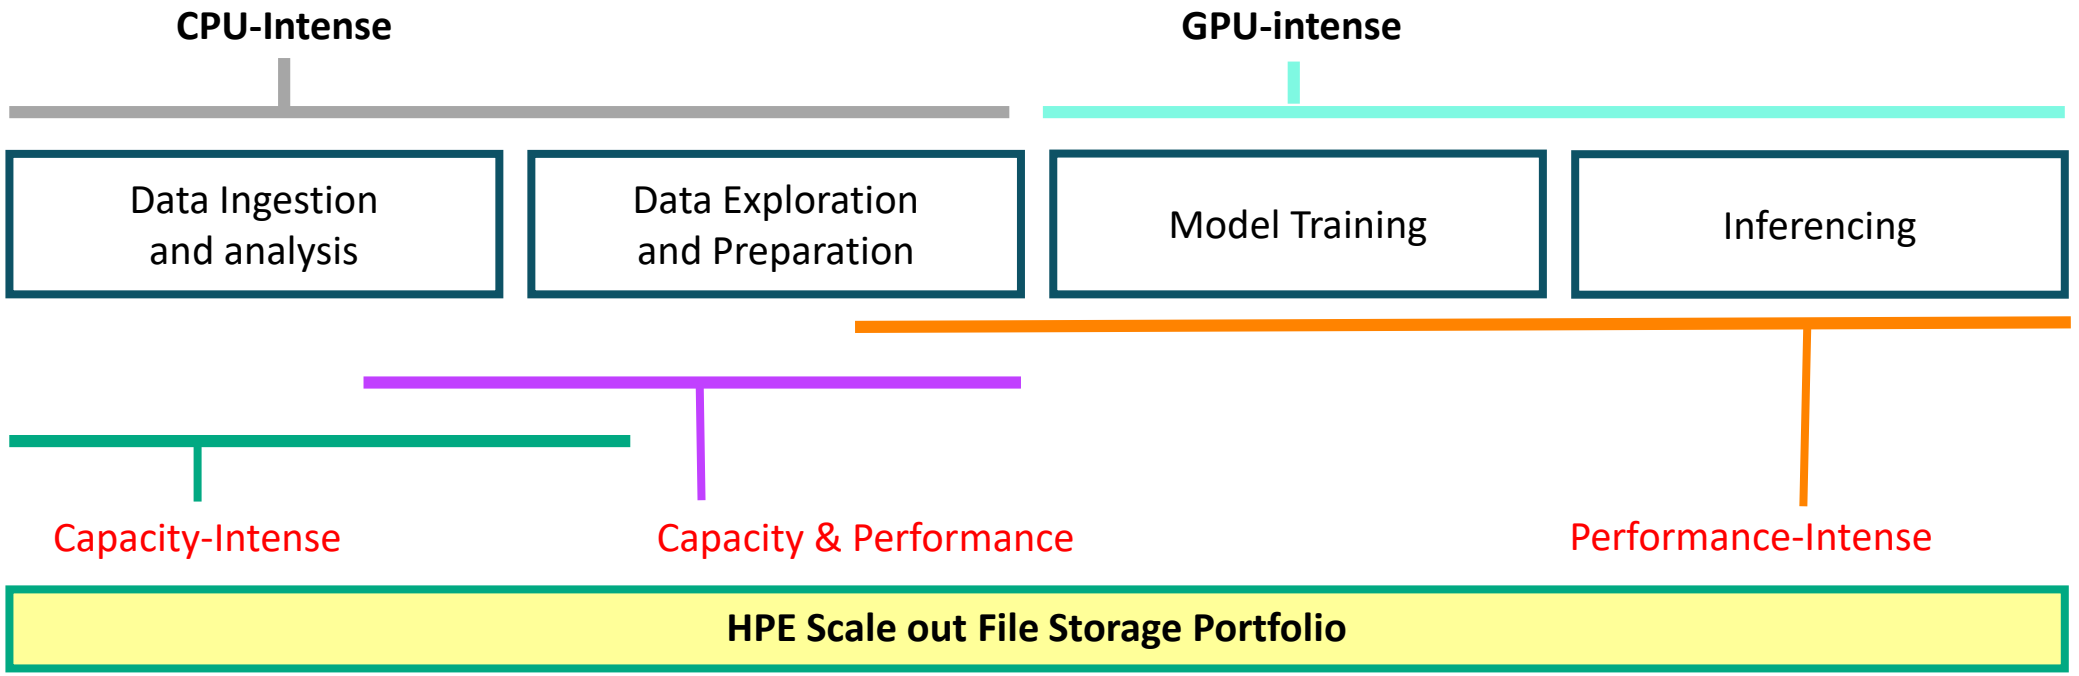

Qumulo

Hybrid

Qumulo

All NVMe

Hewlett Packard Enterprise

敦新科技


THANK YOU!