



**Hewlett Packard
Enterprise**

HPE AI Day



October 10, 2024

Forward-looking statements

This presentation contains forward-looking statements within the meaning of the safe harbor provisions of the Private Securities Litigation Reform Act of 1995. Such statements involve risks, uncertainties, and assumptions. If the risks or uncertainties ever materialize or the assumptions prove incorrect, the results of Hewlett Packard Enterprise and its consolidated subsidiaries (“Hewlett Packard Enterprise”) may differ materially from those expressed or implied by such forward-looking statements and assumptions. The words “believe”, “expect”, “anticipate”, “guide”, “optimistic”, “intend”, “aim”, “will”, “estimates”, “may”, “could”, “should” and similar expressions are intended to identify such forward-looking statements. All statements other than statements of historical fact are statements that could be deemed forward-looking statements, including but not limited to any anticipated financial or operational performance; any projections, estimations, or expectations of addressable markets and their sizes, revenue (including annualized revenue run rate), and margins; any projections or estimations of future orders; any statements of the plans, strategies, and objectives of management for future operations, as well as the execution and consummation of corporate transactions or contemplated acquisitions (including our proposed acquisition of Juniper Networks, Inc.) and dispositions (including disposition of our H3C shares and the receipt of proceeds therefrom), research and development expenditures, and any resulting benefit, cost savings, charges, or revenue or profitability improvements; any statements concerning the expected development, performance, market share or competitive performance relating to products or services; any statements concerning technological and market trends, the pace of technological innovation, and adoption of new technologies, including artificial intelligence-related and other products and services offered by Hewlett Packard Enterprise; any statements regarding current or future macroeconomic trends or events and the impact of those trends and events on Hewlett Packard Enterprise and our financial performance, including but not limited to supply chain, demand for our products and services, and access to liquidity, and our actions to mitigate such impacts to our business; and any statements of assumptions underlying any of the foregoing.

Risks, uncertainties, and assumptions include the need to address the many challenges facing Hewlett Packard Enterprise’s businesses; the competitive pressures faced by Hewlett Packard Enterprise’s businesses; risks associated with executing Hewlett Packard Enterprise’s strategy; the impact of macroeconomic and geopolitical trends and events, including but not limited to supply chain constraints, the use and development of artificial intelligence, the inflationary environment (though easing), the ongoing conflicts between Russia and Ukraine and in the Middle East, and the relationship between China and the U.S.; the need to effectively manage third-party suppliers and distribute Hewlett Packard Enterprise’s products and services; the protection of Hewlett Packard Enterprise’s intellectual property assets, including intellectual property licensed from third parties and intellectual property shared with its former parent; risks associated with Hewlett Packard Enterprise’s international operations (including from public health crises, such as pandemics or epidemics, and geopolitical events, such as those mentioned above); the development and transition of new products and services and the enhancement of existing products and services to meet customer needs and respond to emerging technological trends; the execution of Hewlett Packard Enterprise’s transformation and mix shift of its portfolio of offerings, the execution and performance of contracts by Hewlett Packard Enterprise and its suppliers, customers, clients, and partners, including any impact thereon resulting from macroeconomic or geopolitical events such as those mentioned above; the prospect of a shutdown of the U.S. federal government; the hiring and retention of key employees; the execution, consummation, integration, and other risks associated with business combination, disposition, and investment transactions, including but not limited to the risks associated with the disposition of H3C shares and the receipt of proceeds therefrom and completion of our proposed acquisition of Juniper Networks, Inc. and our ability to integrate and implement our plans, forecasts, and other expectations with respect to the consolidated business; the impact of changes to privacy, cybersecurity, environmental, global trade, and other governmental regulations; changes in our product, lease, intellectual property, or real estate portfolio; the payment or non-payment of a dividend for any period; the efficacy of using non-GAAP, rather than GAAP, financial measures in business projections and planning; the judgments required in connection with determining revenue recognition; impact of company policies and related compliance; utility of segment realignments; allowances for recovery of receivables and warranty obligations; provisions for, and resolution of pending investigations, claims, and disputes; the impacts of tax law changes and related guidance or regulations; and other risks that are described in Hewlett Packard Enterprise’s Annual Report on Form 10-K for the fiscal year ended October 31, 2023, subsequent Quarterly Reports on Form 10-Q, Current Reports on Form 8-K, and in other filings made by Hewlett Packard Enterprise from time to time with the Securities and Exchange Commission.

Speakers



Antonio Neri
President & CEO



Fidelma Russo
EVP & GM, Hybrid Cloud
& HPE Chief Technology Officer



Neil MacDonald
EVP & GM, Server



Today's takeaways

1. Our portfolio comprises the critical building blocks to deliver on the promise of AI and hybrid cloud.
2. HPE's innovation and expertise in designing, manufacturing, installing, and servicing air-cooled and direct-liquid-cooled AI systems at scale, and for enterprises, set us apart.
3. HPE's unique and compelling AI value proposition across networking, storage, servers, services and financing, positions us to create sustainable value for shareholders.



AI-at-scale manufacturing leadership

200K+ Units

Direct-liquid-cooled server nodes
deployed since 2018

22K Units

Direct-liquid-cooled network
switches deployed since 2020

86MW

Power capacity

18K Tons

Chilled water liquid cooling*

*Slide was revised on October 11, 2024 to correct for: 18K Tons chilled water liquid cooling

Delivering on our strategy

HPE GreenLake cloud

HPE Aruba Networking Central (Networking Ops) | Morpheus (DevOps) | OpsRamp (IT Ops) | AI Ops | Security | Marketplace
Data Protection Ops | Sustainability Insight Center | Consumption Analytics (FinOps)

HPE Hybrid Cloud

HPE ProLiant Servers [CPU | GPU]
HPE Alletra Storage [Block | File | Object]
HPE Private Cloud [Virtualization | Containers | BM]
Backup & Recovery
Zerto Disaster Recovery
HPE GreenLake Flex Solutions

HPE AI

100% Fanless Direct Liquid Cooling System Architecture
HPE Cray XD Servers [Service Providers | Sovereign | DLC]
HPE Cray EX Systems [Supercomputing | Sovereign | DLC]
HPE AI Software
HPE Private Cloud AI [Enterprise | NVIDIA]

HPE Networking

Wireless LAN | Switching (Campus, Branch & Data Center) | SD-WAN | SASE | Private 5G | IoT

EDGE | DATA CENTER | COLO | PUBLIC CLOUD

Addressing the \$171B AI market opportunity

	Hyperscaler & model builders ¹	T2/T3 Service providers ²	Sovereigns ³	Enterprises	Total
AI WW TAM '27	\$66B	\$44B	\$19B	\$42B	\$171B
CAGR '23-'27	29%	30%	18%	19%	25%
Liquid-cooled servers	2023 → 2027 \$2B → \$18B	2023 → 2027 \$1B → \$9B	2023 → 2027 \$1B → \$4B	2023 → 2027 \$1B → \$4B	2027 \$35B
Ethernet network	\$1B → \$8B	\$0.3B → \$4B	\$0.2B → \$1B	\$0.1B → \$2B	\$15B
Advisory & support	\$2B → \$8B	\$3B → \$8B	\$3B → \$5B	\$5B → \$11B	\$31B
Air-cooled servers & storage	\$19B → \$33B	\$12B → \$23B	\$6B → \$9B	\$15B → \$25B	\$90B
AI use cases	<ul style="list-style-type: none"> • Foundational model training • Large scale inference by API 	<ul style="list-style-type: none"> • Foundational model training • Large scale inference by API 	<ul style="list-style-type: none"> • AI model training • Generative model fine-tuning • Large scale inference by API 	<ul style="list-style-type: none"> • Fine-tuning models • API inference • Agent (customer, code-gen, creative) run-time 	
AI workload characteristics	<ul style="list-style-type: none"> • 10k-500k cluster size • Optical / Ethernet back-end • Dense, scalable servers • Fragmented silicon options • Air-to-Liquid cooled systems 	<ul style="list-style-type: none"> • Up to 10k cluster size • IB / Ethernet back-end • Dense, scalable servers • Air-to-Liquid cooled systems • Services attach 	<ul style="list-style-type: none"> • 1k to 10k cluster size • Ethernet back-end • Dense clusters • Mostly liquid-cooled systems • Storage, services attach 	<ul style="list-style-type: none"> • 100 - 1k cluster size • IB / Ethernet back-end • Server clusters / appliances • Mostly air-cooled systems • Storage, SW, services attach 	

* TAM numbers may not sum due to rounding

(1) Includes Top 5 U.S. and Top 3 China hyper scalers as well as some model builders (2) Includes T2/3 cloud providers and some model builders, co-los, ISVs, MSPs and telcos (3) Government departments, education and R&D incl. HPC Exascale

Source: HPE / IDC / Dell'Oro. 2024



Delivering on our strategy

HPE GreenLake cloud

HPE Aruba Networking Central (Networking Ops) | Morpheus (DevOps) | OpsRamp (IT Ops) | AI Ops | Security | Marketplace
Data Protection Ops | Sustainability Insight Center | Consumption Analytics (FinOps)

HPE Hybrid Cloud

HPE ProLiant Servers [CPU | GPU]
HPE Alletra Storage [Block | File | Object]
HPE Private Cloud [Virtualization | Containers | BM]
Backup & Recovery
Zerto Disaster Recovery
HPE GreenLake Flex Solutions

HPE AI

100% Fanless Direct Liquid Cooling System Architecture
HPE Cray XD Servers [Service Providers | Sovereign | DLC]
HPE Cray EX Systems [Supercomputing | Sovereign | DLC]
HPE AI Software
HPE Private Cloud AI [Enterprise | NVIDIA]

HPE Networking

Wireless LAN | Switching (Campus, Branch & Data Center) | SD-WAN | SASE | Private 5G | IoT

EDGE | DATA CENTER | COLO | PUBLIC CLOUD


AI requires a modern, secure network

Complexity of network,
applications, data

Convergence of
networking
and security

Sustainability goals
and legislation

Generative AI
development
and rollout



AI FOR NETWORKS: leveraging AI to optimize
experience for operators and users



NETWORKS FOR AI: building datacenters
optimized for data learning and inference



Creating the secure, AI-native network

HPE GreenLake cloud

Security and management | Multivendor AI ops | Single point of control | Client to cloud

← SIMPLIFIED NETWORK AI OPERATIONS →

Wireless



Wired



SD-WAN



Routing



Data Center



SECURE, AI-NATIVE, CLOUD-NATIVE



Users, Devices



Edge



Data Center



Co-Lo



Public Cloud,
Telcos

Delivering on our strategy

HPE GreenLake cloud

HPE Aruba Networking Central (Networking Ops) | Morpheus (DevOps) | OpsRamp (IT Ops) | AI Ops | Security | Marketplace
Data Protection Ops | Sustainability Insight Center | Consumption Analytics (FinOps)

HPE Hybrid Cloud

HPE ProLiant Servers [CPU | GPU]
HPE Alletra Storage [Block | File | Object]
HPE Private Cloud [Virtualization | Containers | BM]
Backup & Recovery
Zerto Disaster Recovery
HPE GreenLake Flex Solutions

HPE AI

100% Fanless Direct Liquid Cooling System Architecture
HPE Cray XD Servers [Service Providers | Sovereign | DLC]
HPE Cray EX Systems [Supercomputing | Sovereign | DLC]
HPE AI Software
HPE Private Cloud AI [Enterprise | NVIDIA]

HPE Networking

Wireless LAN | Switching (Campus, Branch & Data Center) | SD-WAN | SASE | Private 5G | IoT

EDGE | DATA CENTER | COLO | PUBLIC CLOUD

Making AI simple for the enterprise



Enterprise AI use cases

Financial services

Top 5 US bank enables AI platform-as-a-service across the enterprise with HPE and NVIDIA

Pharmaceutical

Canadian university driving drug discovery using molecular design for more powerful antibiotics

Manufacturing

Aerospace company using AI for break-through jet engine with more power and fewer emissions

Customer support

Global airline enhancing airport kiosk response times through edge computing

Legal compliance

International law firm redefined recordkeeping with instant insights from decades of legal cases

Health care

UK company drive improved patient care through intelligent risk assessment

Invoice and expense management

International bank fine-tunes open-source LLM for faster invoice response


Retail productivity

Queue-free shopping in autonomous stores powered by computer vision

Content creation

Helping marketing generate personalized materials using AI-driven tools

A new category of AI systems



HPE PRIVATE CLOUD AI

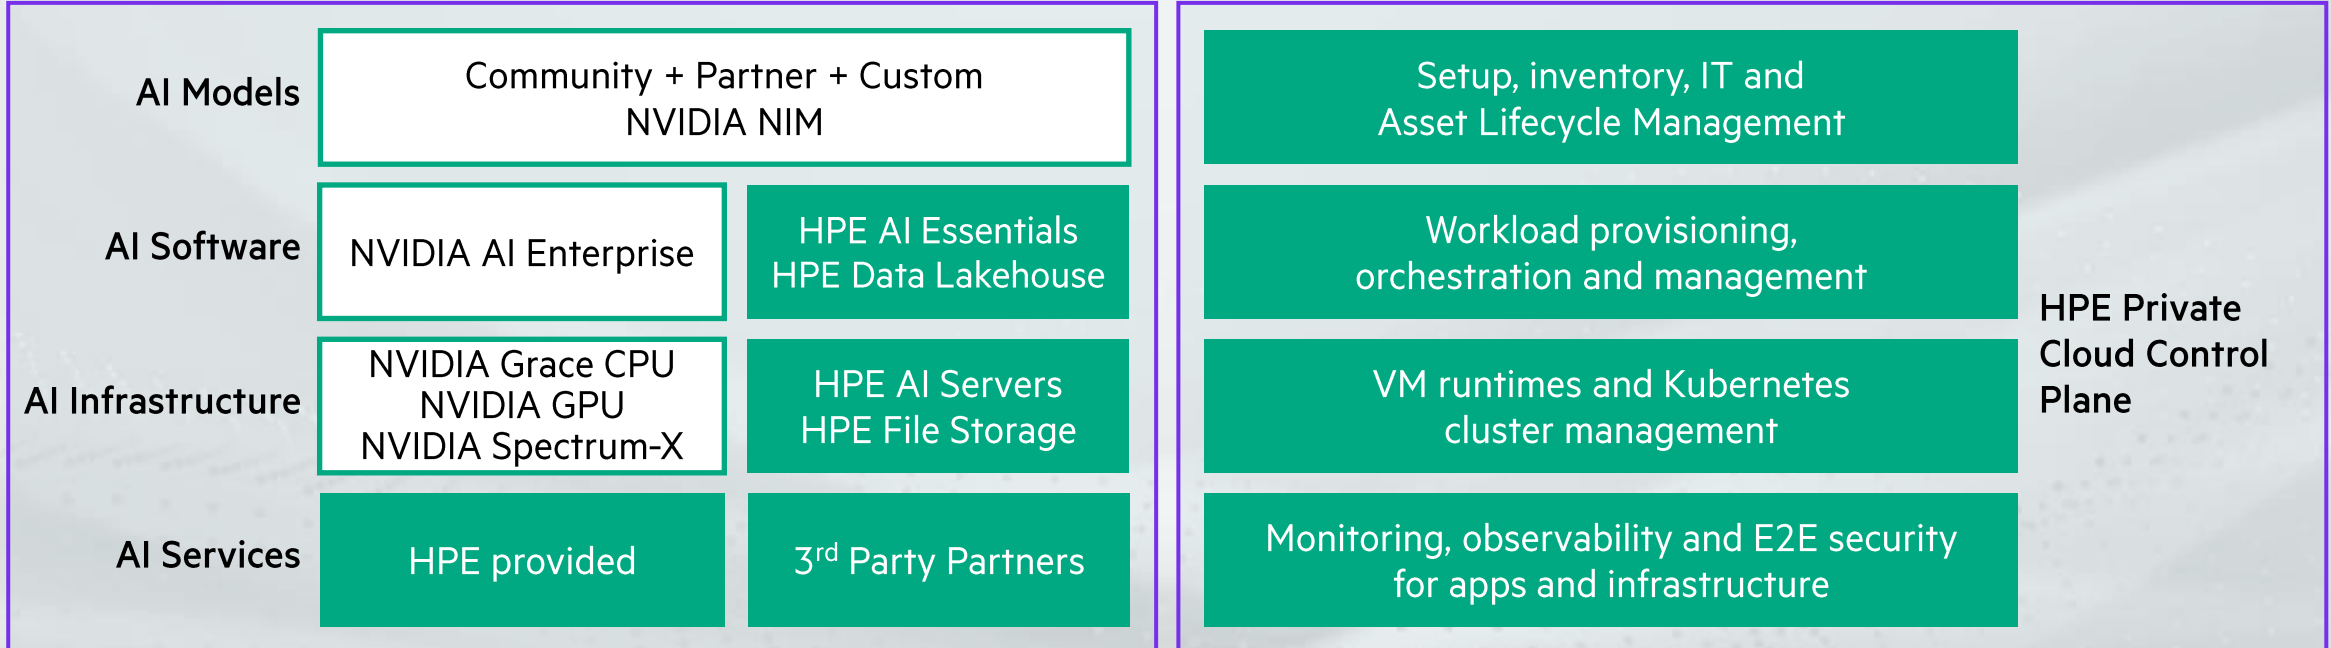
75% less expensive versus public cloud



NVIDIA AI Computing by HPE

HPE Private Cloud AI

HPE GreenLake cloud



“Unleash AI” ecosystem of partners



HPE Alletra Storage MP for Gen AI: Unique disaggregated architecture

Modular, multi-protocol storage
architecture for block or file


Block

File

HPE Alletra Storage MP Controller Nodes




NVME fabric



HPE Alletra Storage MP Capacity Nodes

HPE Private Cloud AI

Scalable performance
AI-ready storage
Comprehensive protection



AI made simple for the enterprise

1. HPE GreenLake is the leading hybrid cloud
2. HPE Private Cloud AI establishes a new category of AI systems
3. Up to 75%¹ better TCO vs. public cloud

¹HPE Analysis: HPE Private Cloud AI vs Public Cloud, Oct. 2, 2024*

*Footnote on slide was revised on October 11, 2024 to correct for: October 2, 2024



Delivering on our strategy

HPE GreenLake cloud

HPE Aruba Networking Central (Networking Ops) | Morpheus (DevOps) | OpsRamp (IT Ops) | AI Ops | Security | Marketplace
Data Protection Ops | Sustainability Insight Center | Consumption Analytics (FinOps)

HPE Hybrid Cloud

HPE ProLiant Servers [CPU | GPU]
HPE Alletra Storage [Block | File | Object]
HPE Private Cloud [Virtualization | Containers | BM]
Backup & Recovery
Zerto Disaster Recovery
HPE GreenLake Flex Solutions

HPE AI

100% Fanless Direct Liquid Cooling System Architecture
HPE Cray XD Servers [Service Providers | Sovereign | DLC]
HPE Cray EX Systems [Supercomputing | Sovereign | DLC]
HPE AI Software
HPE Private Cloud AI [Enterprise | NVIDIA]

HPE Networking

Wireless LAN | Switching (Campus, Branch & Data Center) | SD-WAN | SASE | Private 5G | IoT

EDGE | DATA CENTER | COLO | PUBLIC CLOUD

History informing future innovation

1970s-2010s


2010s

2020s



Cray 1

Refrigerant cooled



Cray 2

Pumped Single Phase
Immersion Cooled
(Fluorinert)



Cray XT

Vertical
Refrigerant
Cooling



**HPE Apollo
8000**

Liquid Cooling



**HPE ProLiant
DL365 Gen11**
DLC



**HPE ProLiant
DL385 Gen11**
DLC



**HPE ProLiant
DL360 Gen11**
DLC



**HPE ProLiant
DL380 Gen11**
DLC



**HPE Cray
XD670**
DLC



**HPE Cray
XD665**
DLC



HPE Cray 2000
DLC




**HPE Slingshot
interconnect**
100 % fanless DLC



**HPE Cray
EX2500**

100% fanless DLC



First to exascale
100% fanless DLC



Isambard AI
100% fanless DLC



Why cooling innovation is needed

Transistors
per CPU/GPU

500x

2007

2024

200M

100B

Single server
power consumption

33x

2007

2024

.3kW

10kW



Direct liquid cooling delivered reliably at scale

Air Cooling

Fans, air conditioning, and vents circulate air and remove heat from computing equipment



Liquid to Air Cooling

Chilled water supply from the facility cools down the air-cooling system positioned close to the servers



70% Direct Liquid Cooling

Combined direct liquid cooling and air cooling



100% Fanless Direct Liquid Cooling


Coolant flows through a network of tubes and cold plates to extract heat directly from all components on the server



Cooling efficiency and capacity (kW/rack) increases from left to right

100% Fanless direct liquid cooling system architecture

Cooling Design



System Design

High-density, high-performance system design, rigorous testing, monitoring software, and on-site services to support successful deployment.

Fabric Design

Integrated network fabric design based on dragonfly topology.

Open Design

Open system design to offer accelerator flexibility.

**100% fanless
direct liquid cooling**

Impact

90%

reduction in cooling power
consumption vs air cooled

50%

reduction in floor
space required

37%

reduction in cooling power
consumption vs hybrid DLC

at least **50%**

reduction in network
connectivity power
consumption



Advanced capabilities across customer segments

Model Builders

- Reliable large-scale training and inferencing infrastructure
- Expertise in at-scale deployments
- Data storage infrastructure & networking

AI Service Providers

- Speed to deploy
- GPU time to market
- Service response time
- Training / tuning / inferencing infrastructure

Sovereigns

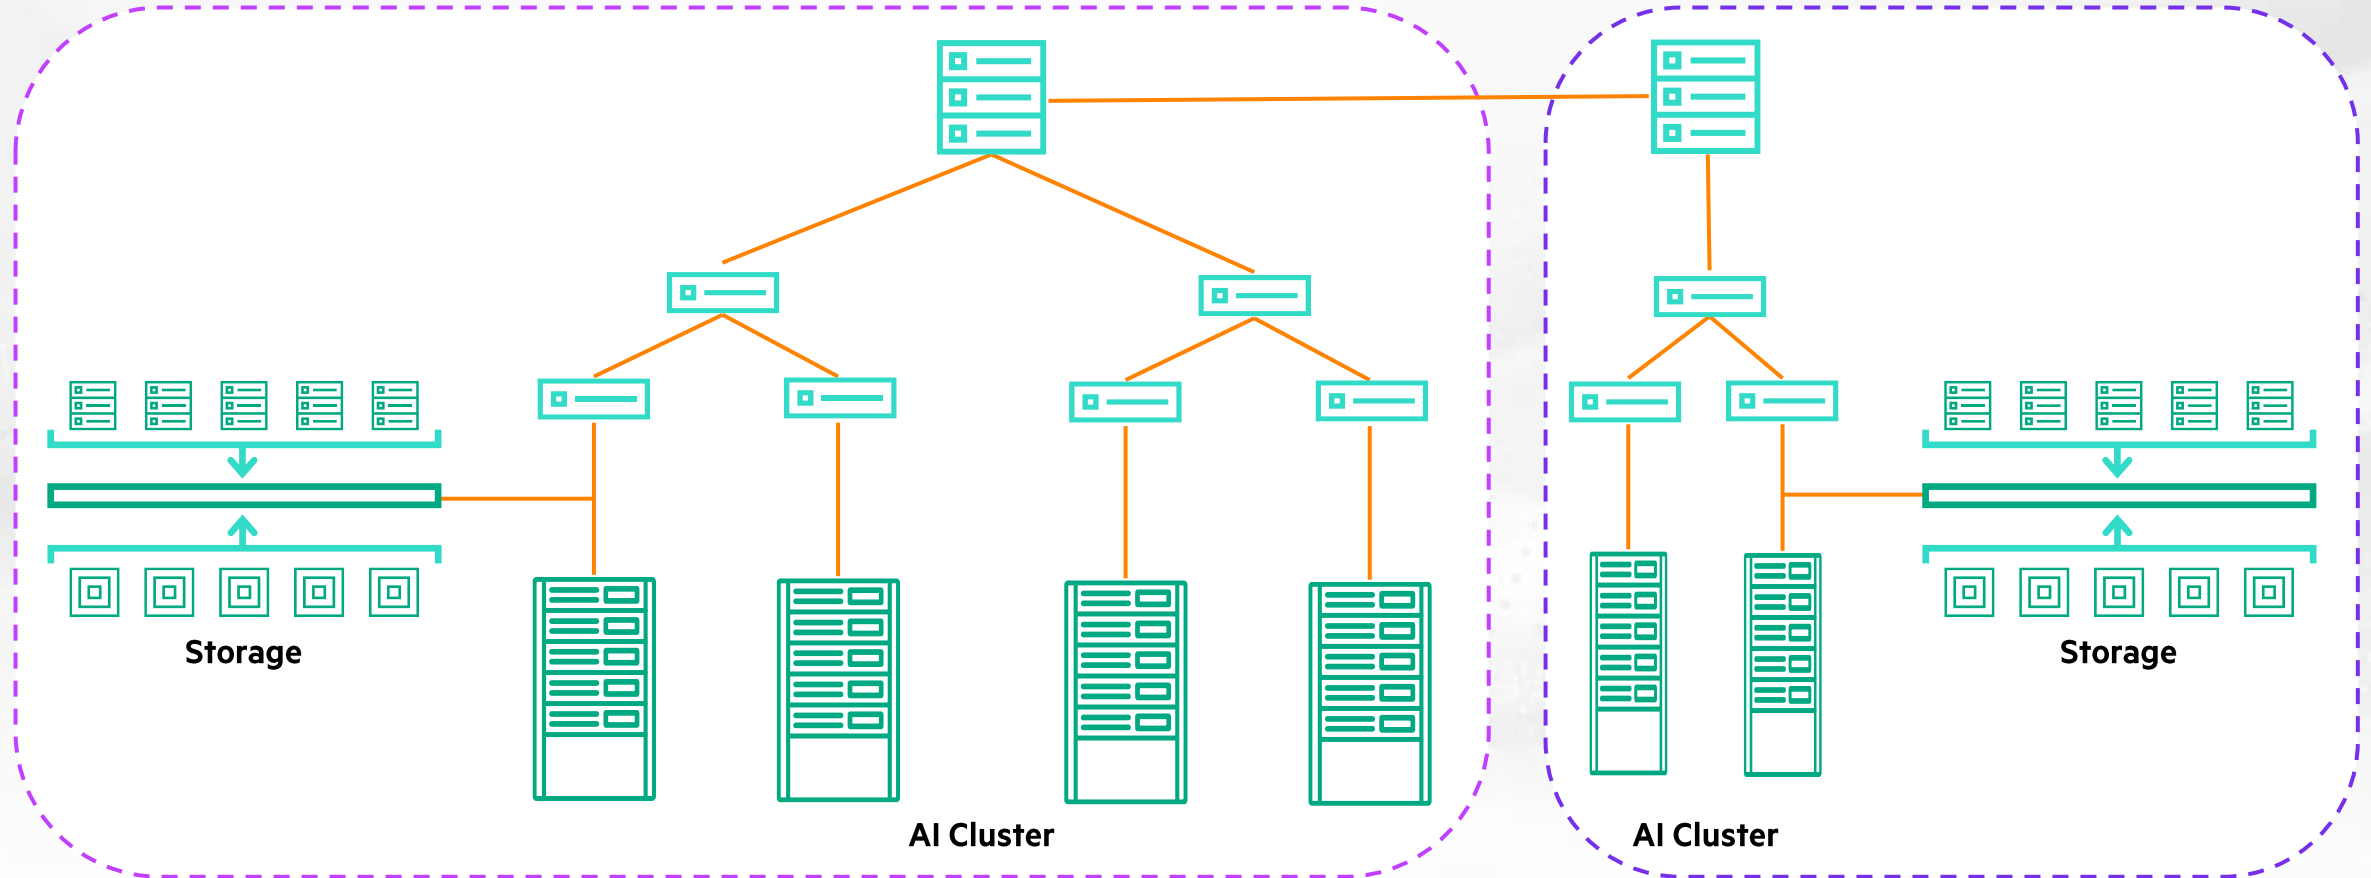
- Systems design expertise
- Reliability
- Fully integrated technology
- Full suite of services from installation through management

Systems design | Manufacturing excellence | 100% Fanless DLC | Global deployments | Modular DC
File & Object Storage | Hybrid Cloud | Networking | Software | Consulting & Services

AI drives the need for compute, storage, and networking

Site 1

Site 2



Leading AI at scale

Power consumption and heat increase with performance, HPE leads the industry with direct liquid cooling solutions.

Deep experience in end-to-end AI systems design and testing is critical to customer success, especially for Sovereigns and Enterprise.

HPE's diversity of accelerator options and time-to-market enables Tier 2/3 service providers.

Networking is key to optimizing performance at many levels: between GPUs, pods, racks, and physical locations.





Key takeaways

1. Our portfolio comprises the critical building blocks to deliver on the promise of AI and hybrid cloud.
2. HPE's innovation and expertise in designing, manufacturing, installing, and servicing air-cooled and direct-liquid-cooled AI systems at scale, and for enterprises, sets us apart.
3. HPE's unique compelling AI value proposition across networking, storage, servers, services and financing, positions us to create sustainable value for shareholders.

Thank you

