HPE 在高效能運算與人工智慧領域的佈局
HPE's Presence in HPC & AI

D. I. Tsai
Solution Architect
HPC & AI, APAC

Sept 04, 2025

Why HPE?

An industry leader

Global reach

Unique expertise

Optimized, end-to-end solutions

Sustainability track record

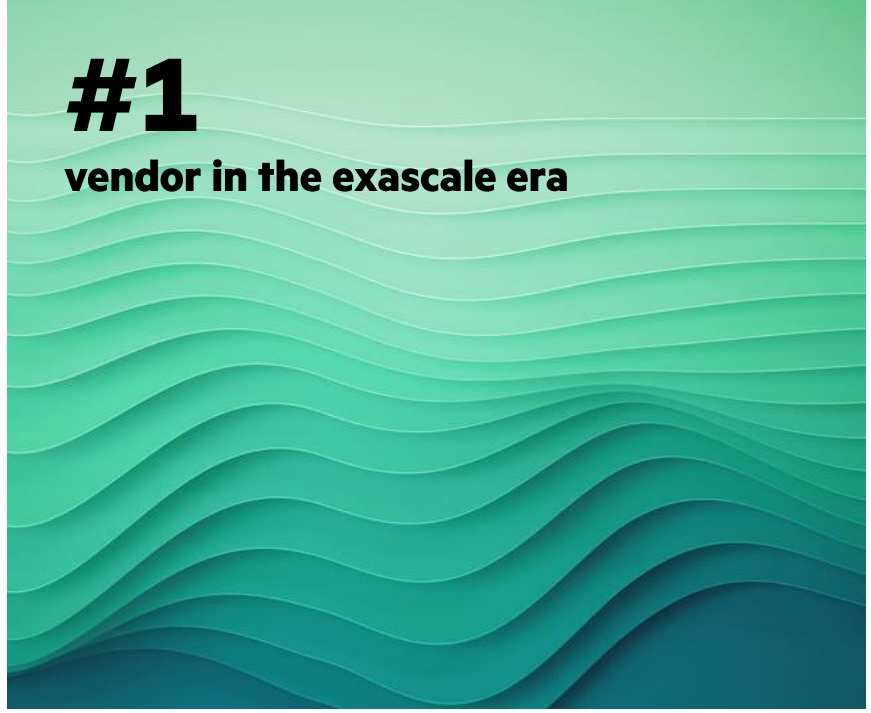

1,2) Source: Nov 2024 Top500

HPE has delivered the three world's fastest, verified supercomputers

AMD

ranked
SUPERCOMPUTER
in the world.
at 1.742 exaflops.

TOP 500 The List.

AMD

ranked
SUPERCOMPUTER
in the world.
at 1.353 exaflops.

TOP 500 The List.

ranked
SUPERCOMPUTER
in the world.
at 1.012 exaflops.

TOP 500 The List.

Sources: Nov 2024 Top500

HPE delivers leadership-class

SUPERCOMPUTING

innovation

#1 #2 #3

HPE has built the
THREE FASTEST VERIFIED SUPERCOMPUTERS
in the world

7

HPE has built
7 of the TOP10
performing
supercomputers
in the world

16

HPE has built
16 of the TOP25
most energy
efficient
supercomputers
in the world

Sources: Nov 2024 Top500

LEGACY OF LIQUID COOLING INNOVATION

Cray Research Inc formed

Silicon Graphics acquires **Cray Research Inc.**

Terra Computing Buys **SGI** vector Processing **Cray Inc**, Is formed

SGI acquired by **HPE**

Cray acquired by **HPE**

1970s

Cray 1
Refrigerant Cooled

1980s

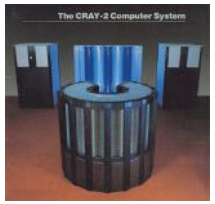

Cray 2
Pumped Single Phase Immersion Cooled (Fluorinert)

1990s

Cray YMP
Cold Plate Pumped Fluorinert Cooled

Cray C-90
Cold Plate (Fluorinert)

2000s

SGI ICE
Room Neutral Cooling

Cray XT
Vertical Refrigerant Cooling

2010s

HPE SGI 8600
ICE Hybrid Cell Cooling

Cray XC
Horizontal Chilled Water Cooling

HPE Apollo 8000
Liquid Cooling

Cray EX
Fanless DLC

HPE Apollo Gen10+
DLC

HPE Cray XD and ProLiant Compute
DLC

HPE ARCS
Room neutral cooling

HPE LIQUID COOLING TECHNOLOGY

MULTIPLE OPTIONS TO MEET WORKLOADS AND FACILITY CONSTRAINTS

HPE ProLiant Compute and HPE Cray XD

HPE Cray EX

Liquid to Air Cooling

70% Direct Liquid Cooling

100% Direct Liquid Cooling

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

HPE COOLING SOLUTIONS THERMAL RANGE

COOLING OPTIONS COVERING 10KW - 400KW

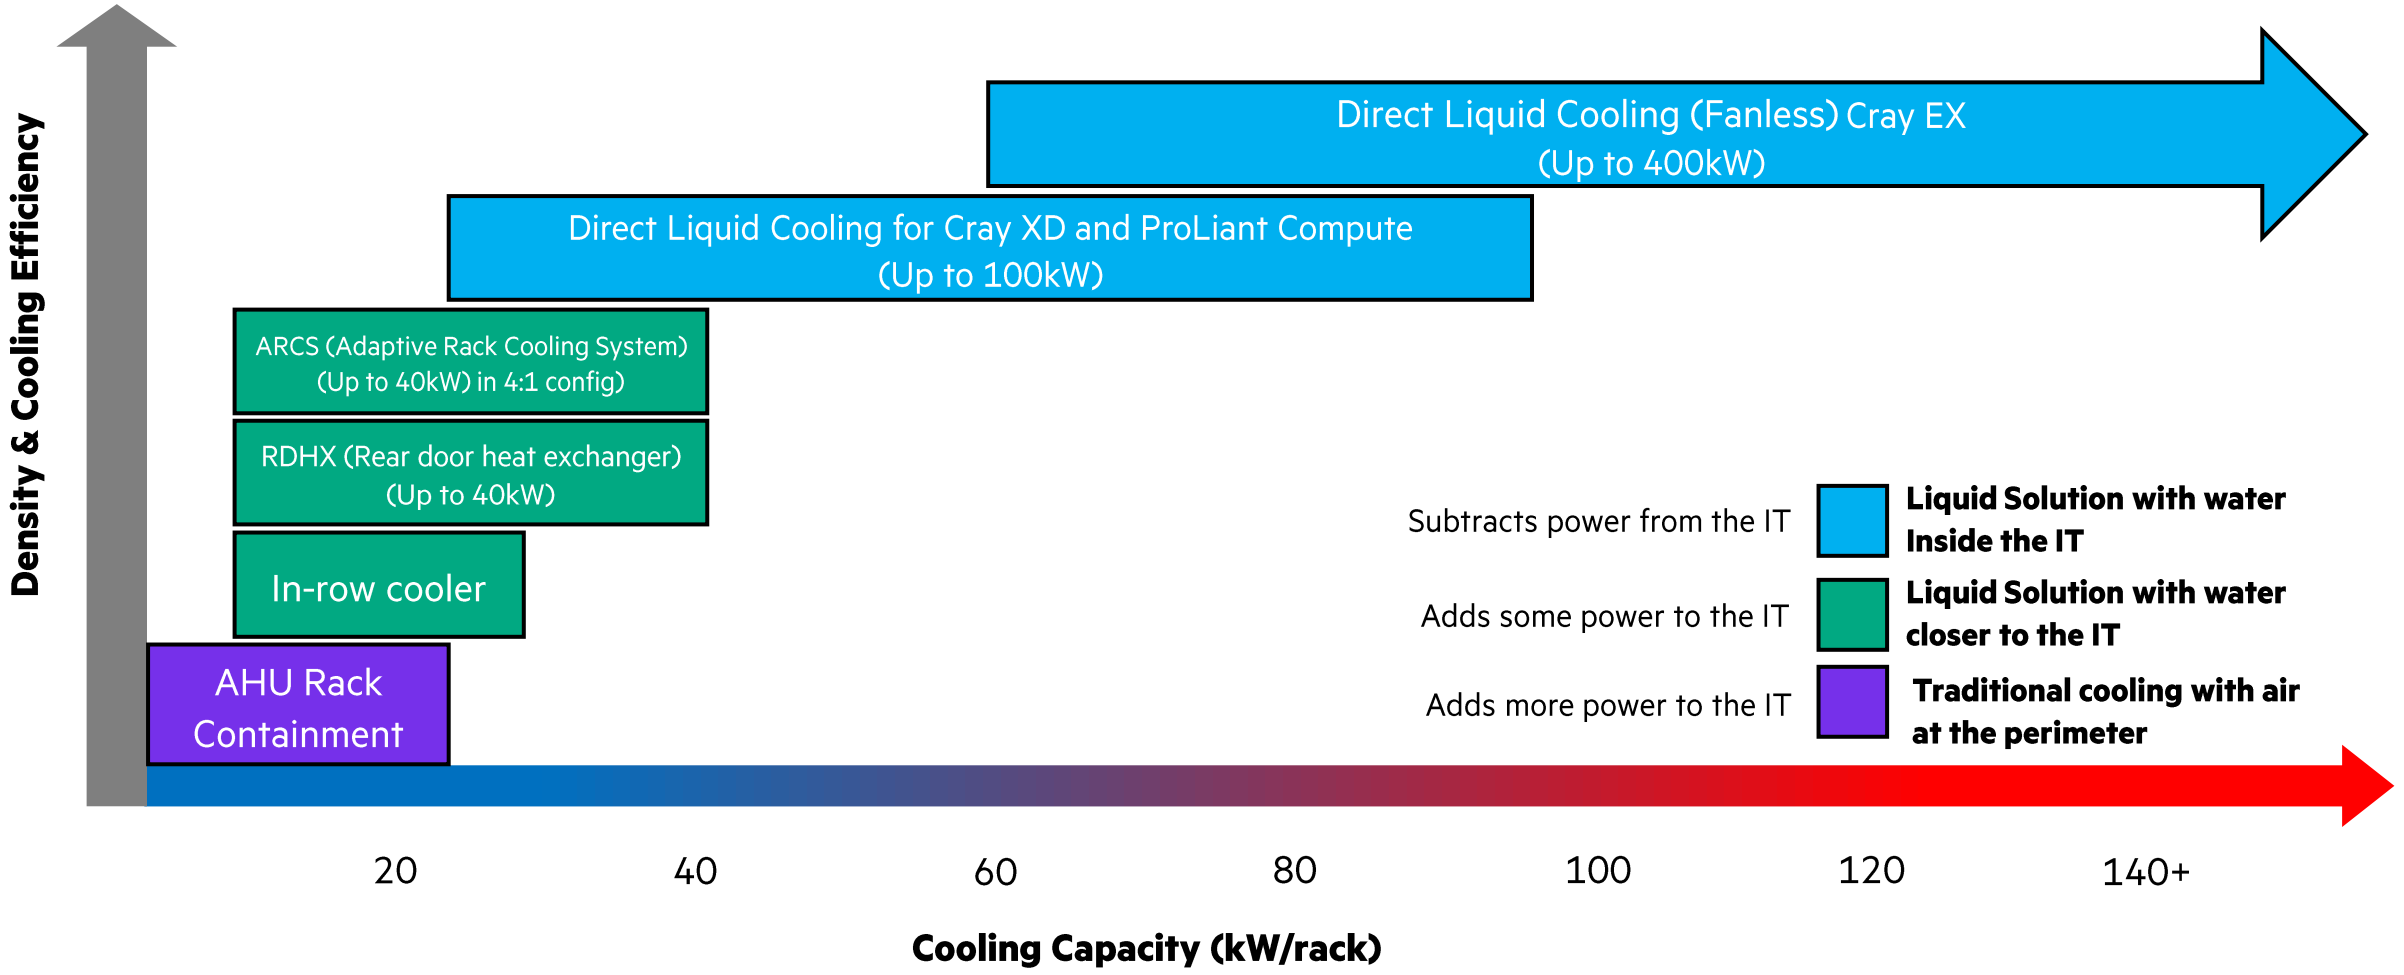

REUSING ENERGY FROM LIQUID-COOLED SYSTEMS

MULTIPLE PATHS TO REDUCING CARBON FOOTPRINT

NREL Transforming ENERGY

QScale

Green Mountain

hima

Warming buildings and facilities

Greenhouses to grow tomatoes

World's largest land-based trout farm

Powering 7 of the world's top 10* most powerful supercomputers

Cray ClusterStor E1000 is the de-facto standard for fast and cost-effective storage for leadership computing

Lawrence Livermore National Laboratory

#1

STORAGE
400 PB
Cray ClusterStor E1000

STORAGE STATUS:
IN PRODUCTION

OAK RIDGE National Laboratory

#2

STORAGE
690 PB
Cray ClusterStor E1000

STORAGE STATUS:
IN PRODUCTION

Argonne NATIONAL LABORATORY

#3

STORAGE
200 PB
Cray ClusterStor E1000

STORAGE STATUS:
IN PRODUCTION

eni

#5

STORAGE
125 PB
Cray ClusterStor E1000

STORAGE STATUS:
IN PRODUCTION

CSCS

#7

STORAGE
116 PB
Cray ClusterStor E1000

STORAGE STATUS:
IN PRODUCTION

EuroHPC Joint Undertaking

#8

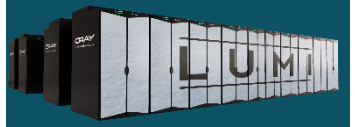

STORAGE
87 PB
Cray ClusterStor E1000

STORAGE STATUS:
IN PRODUCTION

Lawrence Livermore National Laboratory
(Tuolomne)

#10

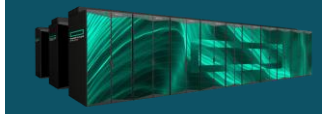

STORAGE
400 PB
shared with El Capitan

STORAGE STATUS:
IN PRODCUTION

Top 500 list, November 2024

Selected customers who have invested in our storage systems in FY24

Storage solutions for any performance requirements and compute cluster sizes

Powered By AMD

I/O performance

HPE Cray Supercomputing Storage Systems
E2000

SC24 Atlanta, GA | hpc creates.

HPE Cray Supercomputing Storage Systems
E1000

HPE Cray Storage Systems
C500

A Primary GPU Vendor

TotalEnergies eni

VOLVO

slb NTT

HLRS

# of supported compute nodes

10s | 100s | 1,000s | 10,000s

*Scalable Storage Unit

E2000 can feed GPUs and write checkpoints more than twice as fast

Powered By AMD

Cray ClusterStor E1000 Storage Systems

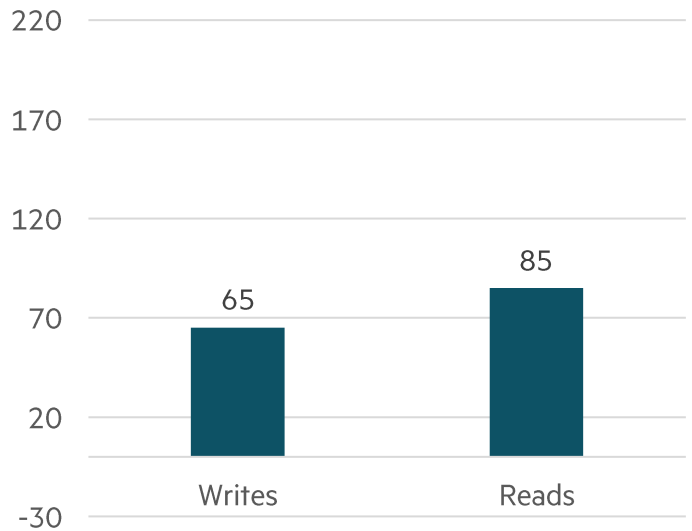

HPE Cray Supercomputing Storage Systems E2000

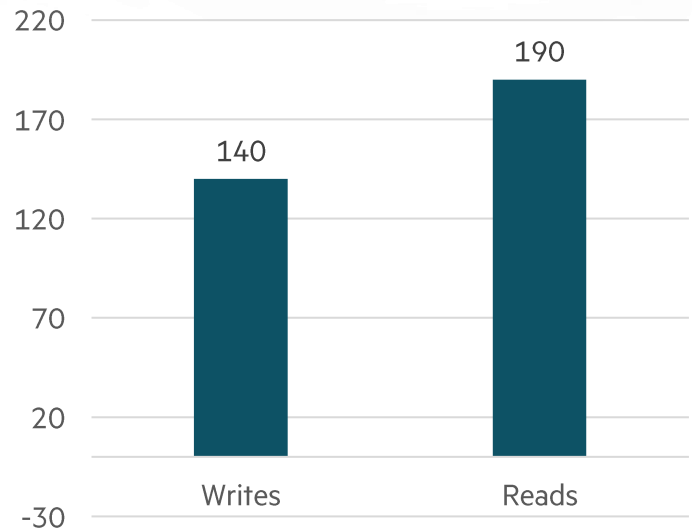

Earlier this year, E1000 supported 30 #1 MLPerf Inference v5.0 benchmarks

HPE Cray XD670 with Cray ClusterStor E1000:
Leadership with 8 SXM-based GPUs on AI inference workloads

#1 in 30 scenarios[1,2]

#1 Computer vision

Image classification, Object detection, 3D Medical Imaging:
Resnet50[1], Retinanet[1], 3D-Unet[2]

Commerce

#1 Recommender:
DLRMv2 99.0, 99.9 - Server & Offline[2]

LLMs

#1 Large Language Model:
GPT-J 6B[1] and Llama2-70B-Interactive[1]

Breaking news: MLPerf Storage benchmark v2.0 results from August 4th
Checkpointing massive AI models (Meta Llama 3.1 405B model) to 160 E1000 HDD OSTs in 24 seconds

Source: https://www.alcf.anl.gov/news/auroras-daos-and-lustre-excel-mlperf-storage-benchmark-large-scale-ai

1 Source: "MLPerf Inference: Datacenter v5.0 results" MLCommons, April 2025 – Top performance for Resnet50-Offline, Retinanet-Offline, GPTJ 99.0-Offline, GPTJ 99.9-Offline, Llama2-70B-Interactive 99.0-Server, Llama2-70B-Interactive 99.9-Server compared to other H100-SXM systems (submission IDs: 5.0-0039, 5.0-0040), and 7% better in Llama2-70B-Interactive 99.0/99.9 Server scenarios compared to other H100-SXM results
2 Source: "MLPerf Inference: Datacenter v5.0 results" MLCommons, April 2025 – HPE Cray XD670 is #1 top-performing server for 3D medical imaging (3D-Unet 99 and 99.9), Recommenders (DLRMv2 99.0 & 99.9 Offline/Server) for all H200-SXM and H100-SXM systems (submission IDs: 5.0-0039, 5.0-0040, 5.0-0041), and 8% better in DLRMv2 99.0-Server scenario compared to other H100-SXM results

HPE Slingshot is The Breakthrough Interconnect For HPC and AI

The Best of Traditional HPC Interconnects
- Low latency
- Efficient for small to large payloads
- MPI acceleration

+

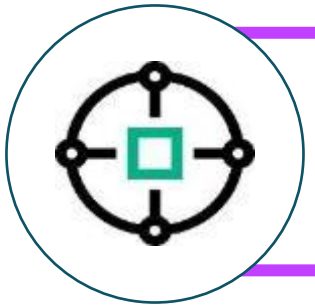

The Best of Ethernet Networks
- Ubiquity
- Interoperability
- Native software support

+

Unique and Powerful Innovations
- Congestion Control
- Fine Grain Adaptive Routing
- Extreme Scale

- ✓ **Efficient:** fewer optical cables
  - ✓ Use about ½ the optical cables versus fat tree for a given global bandwidth
  - ✓ Copper is much less expensive & more reliable, saves ~6-8W per AOC end
- ✓ **Lower latency:** no more than 3 switch-to-switch hops, up to 100,000s of endpoints
- ✓ **Reduced congestion:** Free up expensive computes from congestion-induced delays
- ✓ **Interconnect:** Native Ethernet connectivity
- ✓ **Optimization:** Performance under load, with better scaling, consistency, and utilization

HPE Slingshot Interconnect

Dragonfly Topology

- No more than three switch-to-switch hops on the minimal path even on the worlds largest Exascale systems

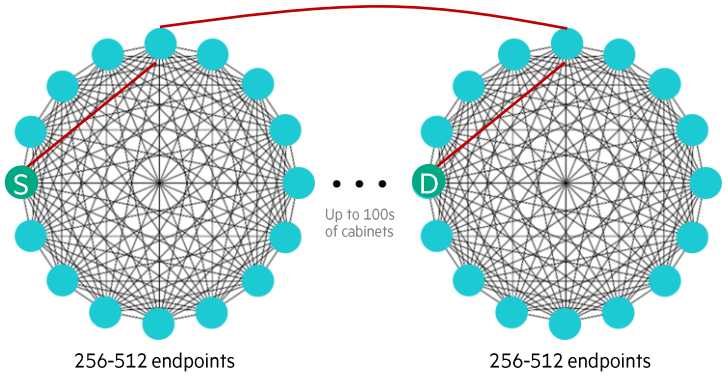

Fine-Grained Adaptive Routing

- Dynamically load balancing both ordered and unordered traffic across many paths to avoid mid-fabric congestion

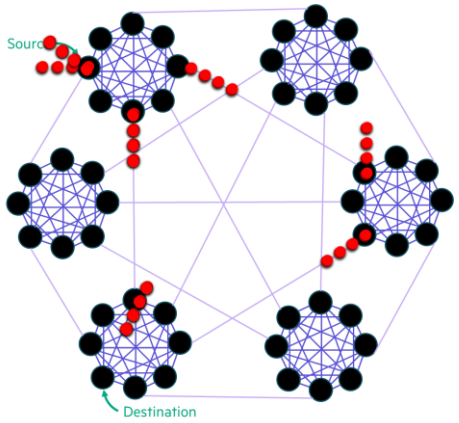

Flow-based Congestion Control

- Achieve consistent and much lower tail latency under load by rapidly detecting and mitigating edge congestion, even on short-lived, HPC-sized messages

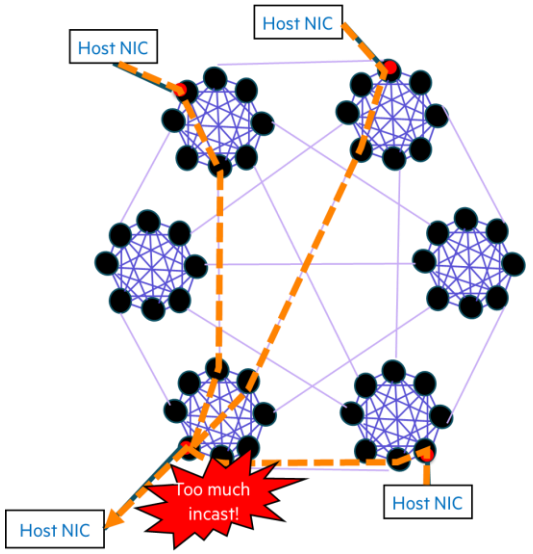

GUI Based Monitoring & Advanced Analytics

- Real-time synthesis of the HPE Slingshot fabric health with friendly graphical interface and deployment options
- Advanced analytics to provide insights into performance of distributed applications

Industry Leading Consortiums Partner with HPE

- Ultra-Ethernet Consortium (UEC)
  - The industry is moving to high-performance Ethernet, HPE Slingshot is leading the trend

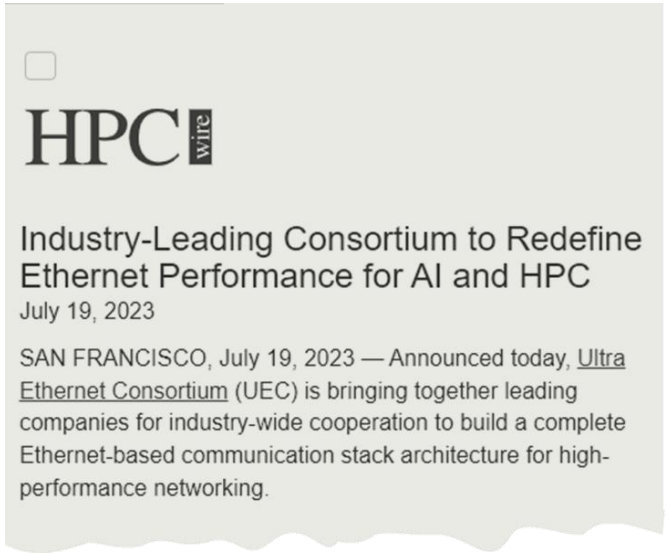

HPC wire

Industry-Leading Consortium to Redefine Ethernet Performance for AI and HPC

July 19, 2023

SAN FRANCISCO, July 19, 2023 — Announced today, Ultra Ethernet Consortium (UEC) is bringing together leading companies for industry-wide cooperation to build a complete Ethernet-based communication stack architecture for high-performance networking.

- Ultra Accelerator Link (UAL)
  - An open specification for a die-to-die interconnect and serial bus between AI accelerators

HPC wire

Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

Home
Topics
Sectors
Exascale
Specials
Resource Library
Podcast
Events
Job Bank
About
Subscribe

Everyone Except Nvidia Forms Ultra Accelerator Link (UALink) Consortium

By Doug Eadline

May 30, 2024

Consider the GPU. An island of SIMD greatness that makes light work of matrix math. Originally designed to rapidly paint dots on a computer monitor, it was then found to be quite useful in large numbers by HPC practitioners. Enter GenAI, and now these little matrix mavens are in huge demand, so much so that we call it the GPU Squeeze.

The well-known and dominant market leader, Nvidia, has charted much of the pathway for GPU technology. For HPC, GenAI, and a raft of other applications, connecting GPUs provides a way to solve bigger problems and improve your application's performance.

There are three basic ways to "connect" GPUs.

HPE HPC software solutions

Unlock the full potential of your HPC architecture to address the world's most complex problems.

System management

Provides monitoring and management capabilities to keep systems running at peak performance

User services software

Improves user experience for running diverse complex workflows, including boosting system efficiency

Development environment

Includes comprehensive toolchain for developing software optimized for HPC systems

Ensure smooth day-to-day operations

Flexible, easy-to-use system management solution offering system administrators tools to manage even the most complex HPC configurations capable of accommodating a growing variety of workloads

System setup

Covers entire HPC system, including login nodes, service nodes, compute nodes, and infrastructure

Hardware management

Supports HPE ProLiant DL, HPE ProLiant Compute XD, HPE Cray XD, and HPE Cray Supercomputing EX systems

Image management and software updates

Supports variety of Linux operating systems for both diskful and diskless systems

Cluster health management

Allows for monitoring, management, and diagnostics of node and system health

Power management

Enables management and analysis of power consumption for the system and workloads

ISV and open-source integration

Facilitates node health and power-aware job scheduling with HPC workload managers

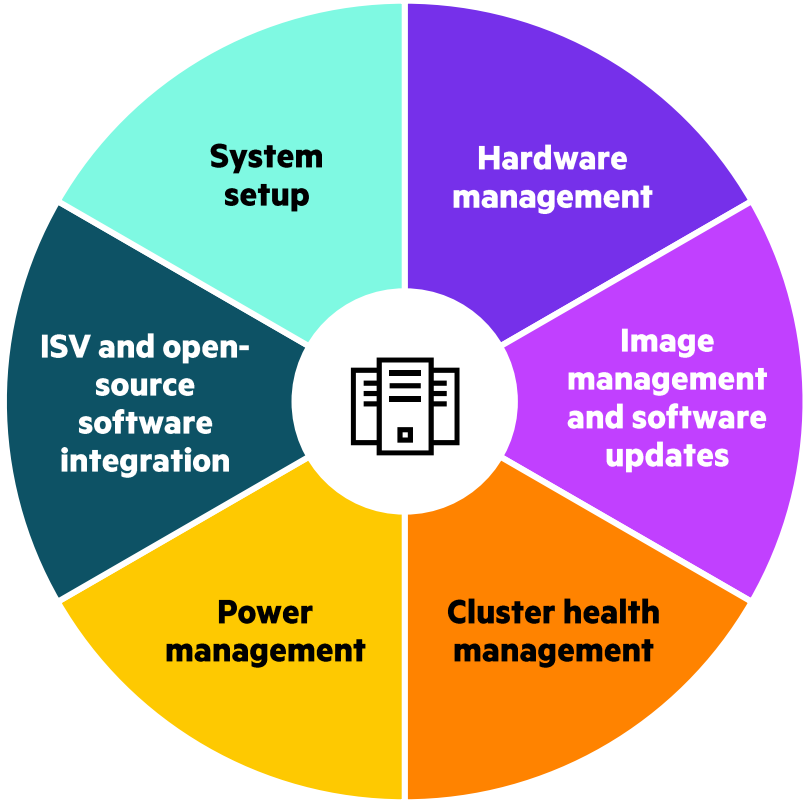

Flexibility

Diverse workflows
can be run with organizations' existing infrastructure.

Multiple Linux distributions
including SUSE Linux Enterprise Server and Red Hat Enterprise Linux are supported.

Power savings

Turbo boost limiting
throttles energy consumption and costs.

Power measurement
keeps users aware of power utilization and distribution.

HPE Cray Supercomputing User Services Software **revolutionizes user experience.**

Efficiency

Low noise mode
focuses resources on running compute workloads.

File system I/O operations
reduce bottlenecks associated with data I/O operations.

XPMEM
creates shared memory capabilities that enhance scalability.

HPE Cray Supercomputing Programming Environment Software

Fully integrated software suite with compilers, tools, and libraries designed to increase programmer productivity, application portability, and performance at scale for HPC organizations developing HPC code in-house

Compilers

Generate highly optimized code with an environment that supports multi-vendor CPUs, GPUs, and accelerators, and is fully integrated with other programming tools.

Libraries

Leverage our collection of numerical routines designed for best performance on HPE Cray systems, including highly tuned kernels and algorithms for a range of architectures.

MPI

Seamlessly run parallel applications on HPC systems with HPE Cray MPI, which establishes a practical, efficient, and flexible standard for message passing.

Debugging tools

Unlock 360° troubleshooting at scale with tools for deconstructing large-scale defects, sanitizing GPU applications, and comparisons.

Performance tools

Optimize programs for faster execution and more efficient computing resource usage by collecting and analyzing data with these tools.

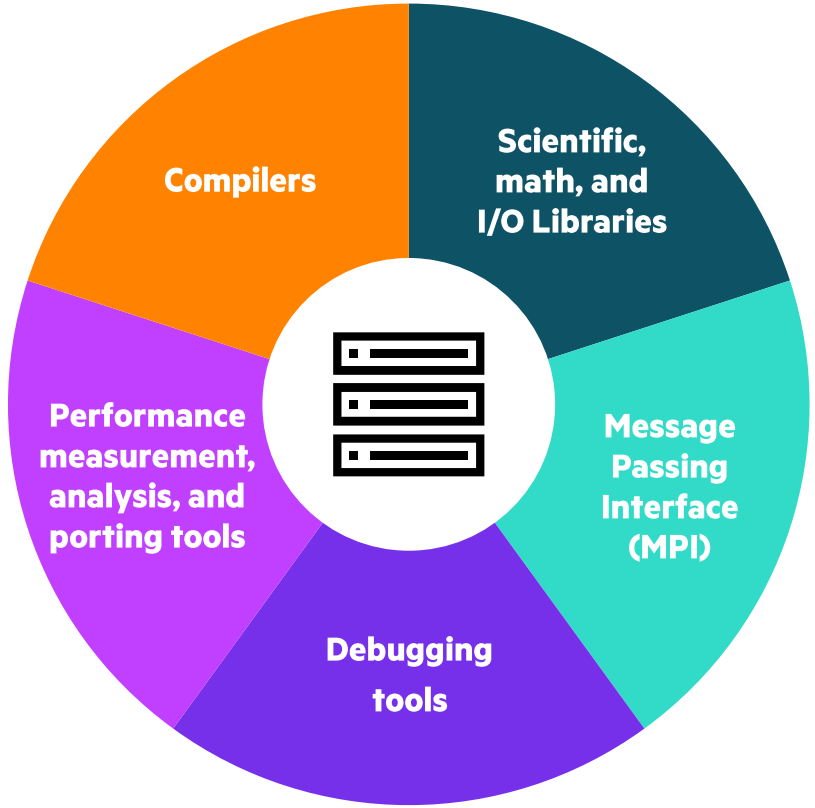

Next steps

If you want to learn more about the world's leading supercomputing solution for your next gen system..

..I am more than happy to schedule a technical deep dive with your experts and our experts.

Thank you