

# HPE Services for HPC and Liquid Cooled solutions

---

EurOpen

Martin Papik, May 2023

# QUICK FACTS ABOUT

# HPE

**\$27.8B**

↑  
Revenue  
3%  
(2021 vs 2020)

**~60.400**

Employees  
(2021)

**#19**

Fortune 100  
Best Companies  
to Work For®  
(2022)




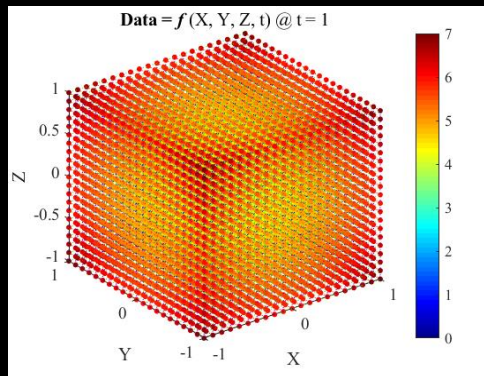





# ATLANTIS

1st November 2022 - Official opening ceremony







# HIGH PERFORMANCE COMPUTING MARKET

- Human genome sequencing
- Nuclear Stockpile Simulation
- Airplane/car manufacturers
- Military systems, Security/Intelligence
- Rendering farms, Entertainment industry
- Oil & Gas, reservoir simulation, seismic...
- Chemistry
- Banks, Financial markets
- Formula 1
- Weather forecast
- Universities – Theoretical Science



# WHAT DO THESE COMPUTE PROBLEMS HAVE IN COMMON?



- High complexity
  - High number of coupled / related tasks
  - Single node memory exhaustion
- => **Divide et Impera** is possible solution



Philip II of Macedon

Challenging to distribute across multiple nodes in an HPC cluster. Want to **solve „holistically“ as one virtual node.**

Compete for “fat nodes” (with more processors/memory), **compounding time and effort to discovery**

Or too big for the one node

If memory is exhausted, simulation fails (hours wasted)

Use smaller models, reducing accuracy and increasing prototype costs/time to production

# PARALLELISM

Parallelism means doing multiple things at the same time: you can get more work done in the same time.

Less fish ...



More fish!



# MPI

	Cray MPI	HPE MPI
Fabric	Slingshot <i>Support for InfiniBand in June 2020</i>	InfiniBand
Packaging	Cray Programming Environment	HPE Message Passing Interface (MPI)
Platforms	Cray Supercomputers <i>HPE Apollo &amp; HPE ProLiant DL in June 2020</i>	HPE Apollo, HPE ProLiant DL
More facts	Tuned for AMD, Intel, Arm and GPUs CrayPAT optimizes rank ordering Works with Cray's Linux Environment XPMEM and Huge Pages OS: Cray Linux Environment, RHEL	Tuned for AMD, Intel, Arm and GPUs MPI profiling and guided placement tools Boost runtime performance for any MPI app w/o recompile OS: RHEL, SLES, CentOS, TOSS



# Why HPE go into Liquid?



# Why Liquid Cooling?

## Performance

Reliable top-bin CPU/GPU operation  
Sustained turbo modes

## Density

More servers per rack  
Fewer racks required

## Efficiency

More effective heat capture  
Lower cooling power required



# HPE Server Solutions for Liquid Cooling systems

# HPC LIQUID COOLING - OVERVIEW

## Hybrid Cooling solutions

How it works:

Liquid to air cooling AND direct liquid cooling (slow convergence into Direct Liquid technologies)



ARC Rack +



Apollo 2000 DLC  
and XD (near  
future)

## Direct Liquid Cooling solutions

How it works:

Coolant flows through cold-plates to extract heat directly from all server components



Cray EX Lite



Cray EX

# HPE Server Cooling Options

## Air Cooling



100% heat to air  
0% heat to liquid

OR

## Direct Liquid Cooling



30% heat to air  
70% heat to liquid

**OPTIONAL**  
ADDITIONAL COOLING

*Can be used* →  
*with air cooled or liquid cooled servers*

## RDHX Rear Door Heat Exchanger



Uses facility water to neutralize exhaust hot air back to room neutral temperatures

OR

## ARCS Adaptive Rack Cooling Solution



Encapsulates the IT via sealed racks and uses facility water to cool air-cooling system

*Requires Facility Water*

# HPE HYBRID LIQUID COOLING – AIR/DLC

**DLC - Direct Liquid Cooling provide** to the hottest components in the server such as the CPU/GPU, memory DIMMs while using air cooling to cool the remaining low heat components


- Rack comes with pre-installed rack mounted CDU and manifolds
- Can be 100% liquid cooled when accompanied with RDHX or ARCS
- Racks can be populated partly and scale over time (Scale UP and OUT). They come pre-filled with coolant.
- Supports top or bottom feed facility water (universal)
- Easy to adapt DC



# HPE Cray XD 2000 + ARCS

The HPE ARCS uses a closed-loop, room-agnostic design capable of cooling fully populated racks, even with top bin CPUs. The horizontal airflow of the HPE Adaptive Rack Cooling system fully supports Apollo, XD, DL, and ClusterStor E1000, and industry-standard front-to-back hardware.

- Allows for in-rack IT cooling without adding to the heat load in the data center.
- Implementation of variable speed fans within ARCS enables improved energy efficiencies by providing the right volume of airflow to all IT regardless of mounting position or workload.
- Acoustic benefits and reduced number of water taps bring advantages over other rack-level cooling products.
- ARCS enables customer datacenter power density to increase without need for facilities upgrades.



HPE ARCS Cooling Unit Rack




# HPE CRAY EX LITE

**Cray EX Lite provides DLC (Direct Liquid Cooling) to all components in the server, including CPU/GPU, memory, voltage regulators, etc.**

- Mainly AMD based CPU and GPU design now
- Racks can be populated one chassis at a time and scale over time. Pre-filled with coolant and ready to connect to existing system.
- Dedicated in-rack CDU, allows additional Cray EX Lite cabinets to be added to the existing cluster over time
- Up to 96 nodes per cabinet
- **Integrated liquid cooled switches**
- **100% liquid cooling solution**
- Supports top or bottom feed facility water



# HPE CRAY EX

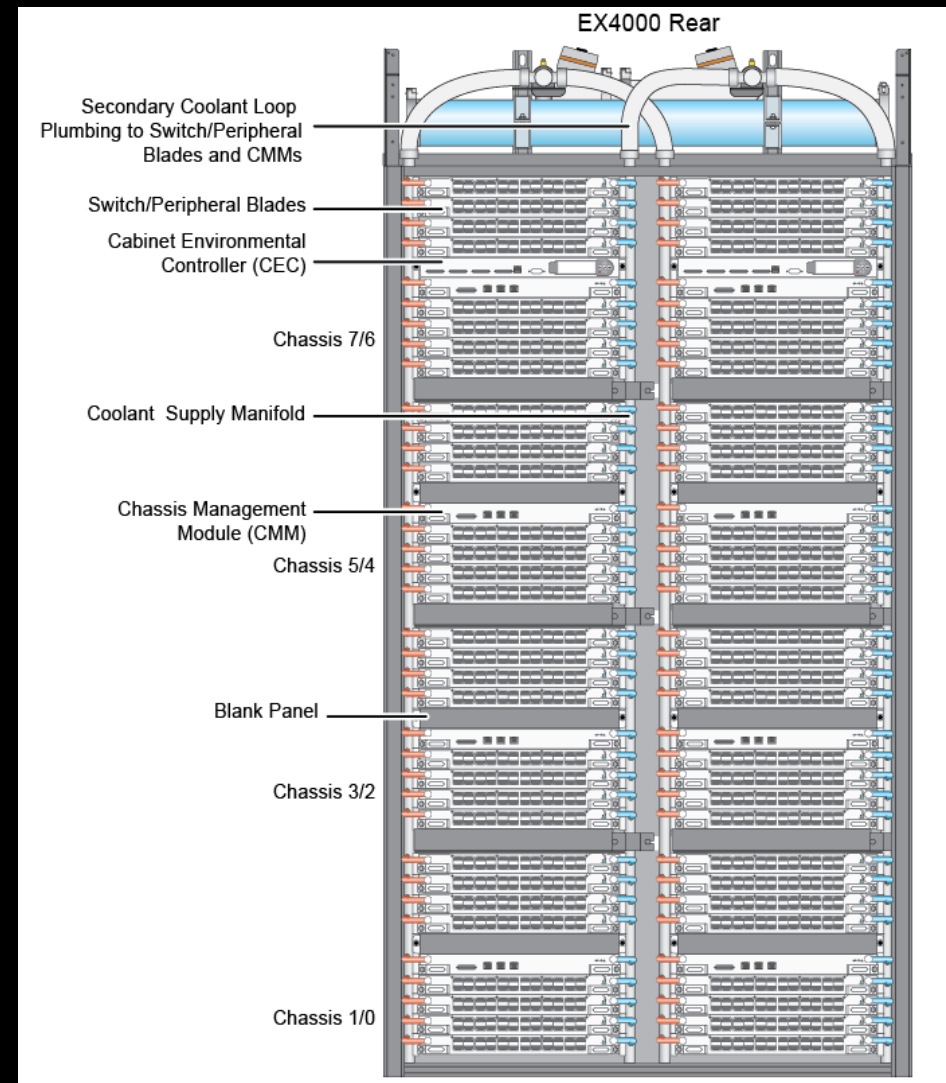



Cray EX provides DLC (Direct Liquid Cooling) to all components in the server, including CPU/GPU, memory, voltage regulators, etc.

- Supports the latest top bin CPU & GPU technology
- Best option for high performance computing deployments exceeding 1000+ nodes
- Large row-based CDUs manage coolant for multiple cabinets
- **Up to 256 nodes per cabinet**
- **1 CDU per 4 Compute racks**
- Racks are heavy and nonstandard dimensions – special DC requirements
- **Management rack with DL3x5 need to be supplied for HPCM and other sw.**




# HPE CRAY EX - Detail




## Blade Detail

---

- EX425
- AMD based
- 4 servers per blade



# EXASCALE HPC GOAL COMPLETED 2022!!!



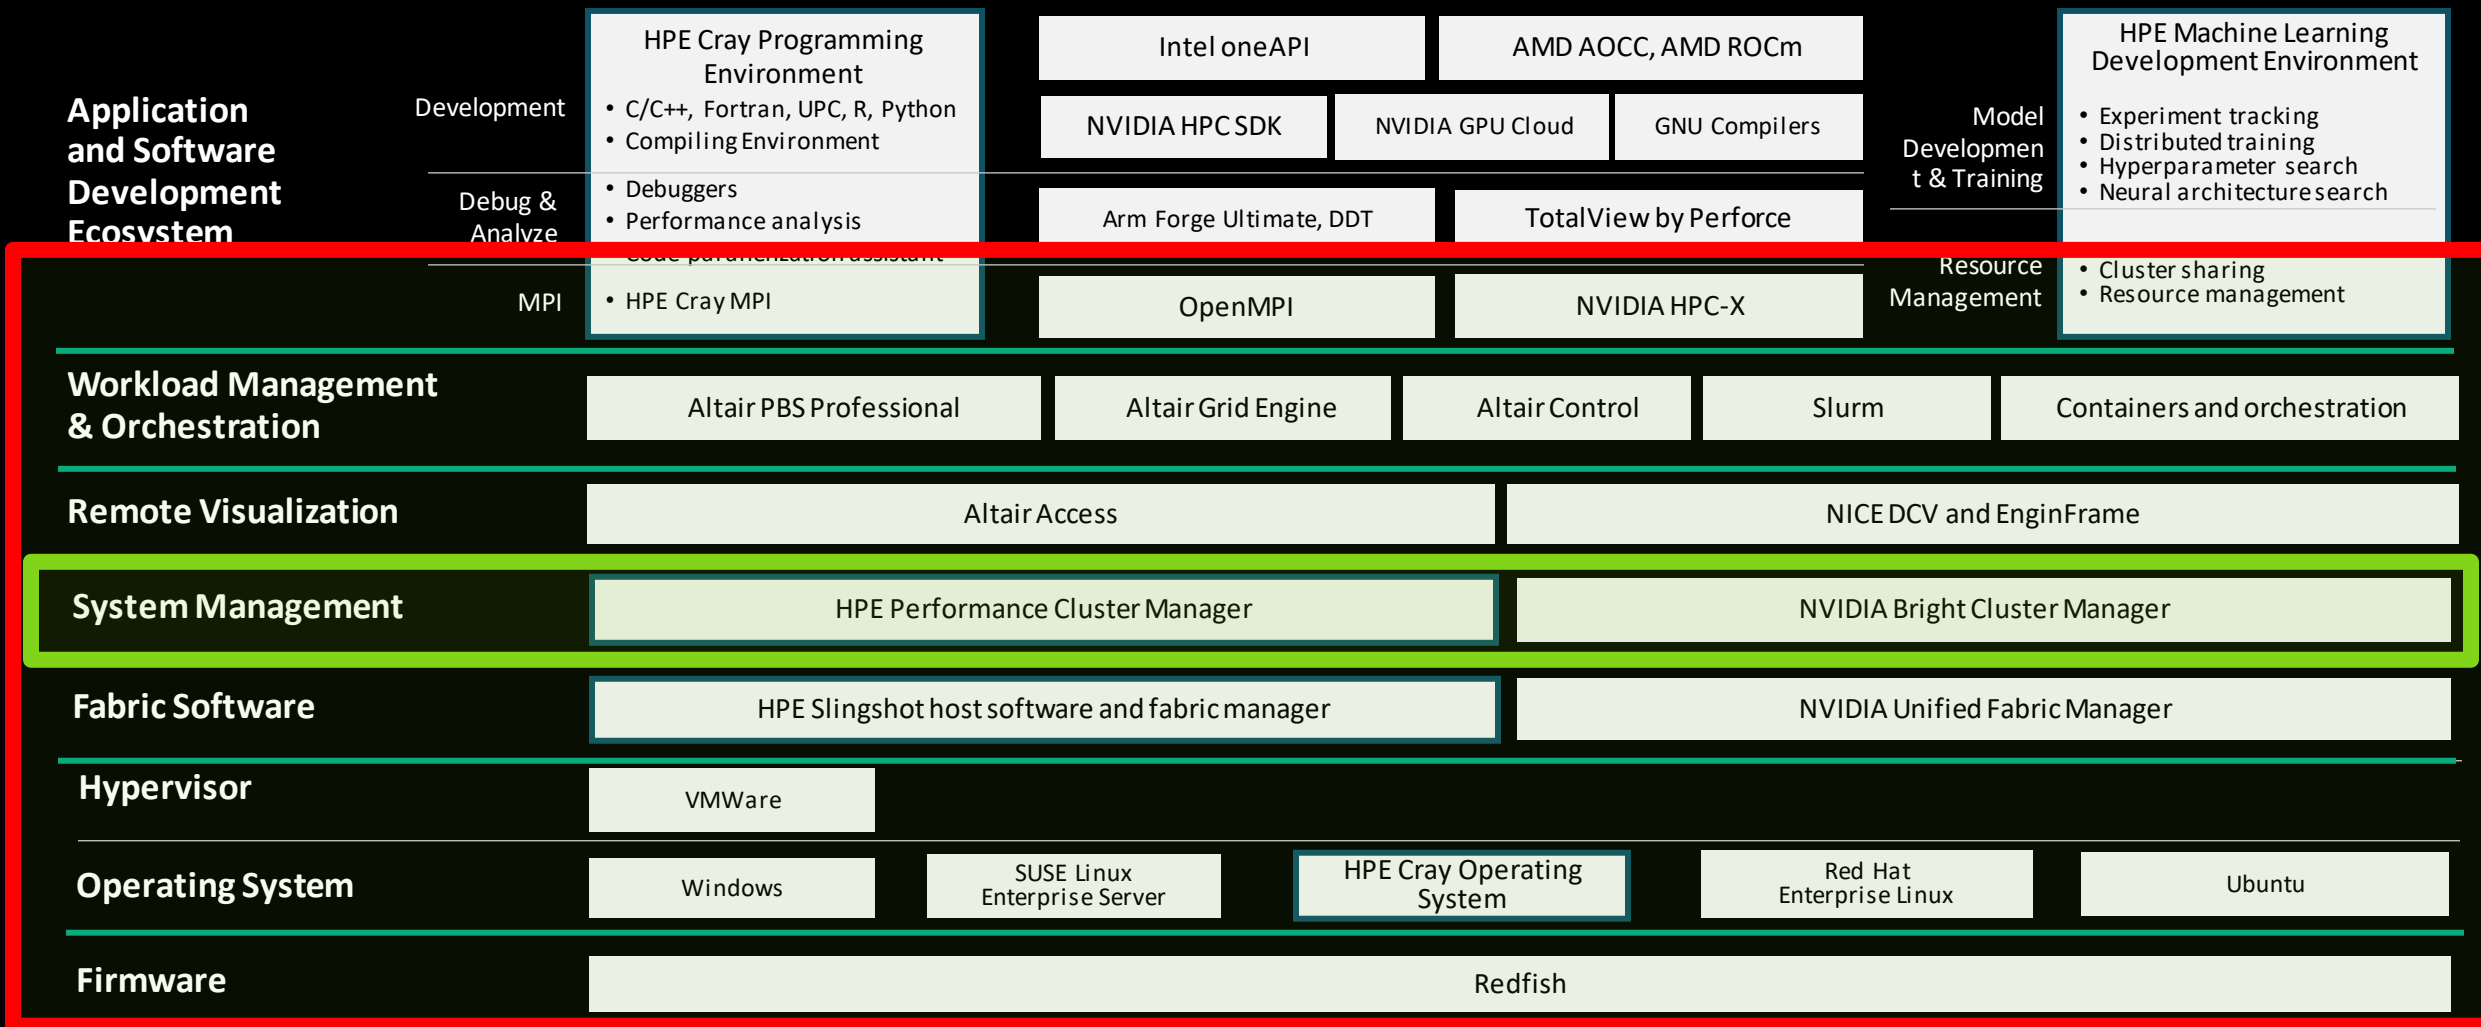
## HPE CRAY EX

- Up to 64 compute blades, and 512 processors per rack
- Flexible bladed architecture supports multiple generations of CPUs, GPUs, and interconnect
- 100% direct liquid cooling
- Scales to 100s of cabinets

El Capitan, an HPE-built, AMD-powered exascale system, is on track to deliver at least **2 exaflops** of peak performance.



# Comprehensive software portfolio for HPE Cray XD systems / GIC services for HPC



# NEXT MISSION: SPACEBORNE COMPUTER-2

LAUNCHED: 20-FEB 2021 ABOARD NORTHROP GRUMMAN'S NG-15



# Thank you

---

[martin.papik@hpe.com](mailto:martin.papik@hpe.com)