
Hewlett Packard Enterprise
NVIDIA
THE NEXT GENERATION IN SUPERCOMPUTING
슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는
HPE CRAY EX SUPERCOMPUTER
2021. 09. 29(수), 10:30 ~ 11:40
ALLSHOW TV

Hewlett Packard Enterprise | NVIDIA

HPE Cray EX시스템 레퍼런스를 통한 HPC 및 AI 시장 동향

Beom-Soo Kim | HPE

WHO IS USING HPC?

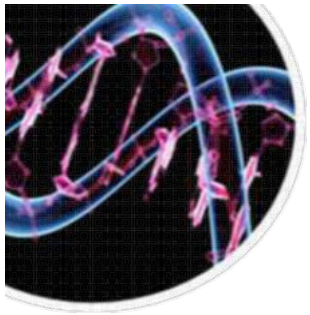

- **Human Genome Sequencing**
- **Nuclear Stockpile Simulation**
- **Airplane/ Car Manufacturers**
- **Military Systems**
- Rendering Farms
- **Oil & Gas, Reservoir Simulation, Seis mic Processin**
- **Chemistry**
- Banks
- Formula 1
- **Weather forecast**
- Universities
- .. and many more

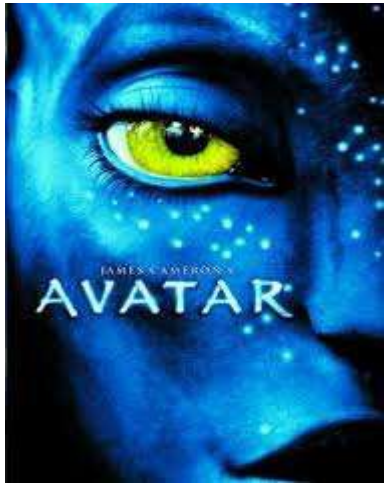

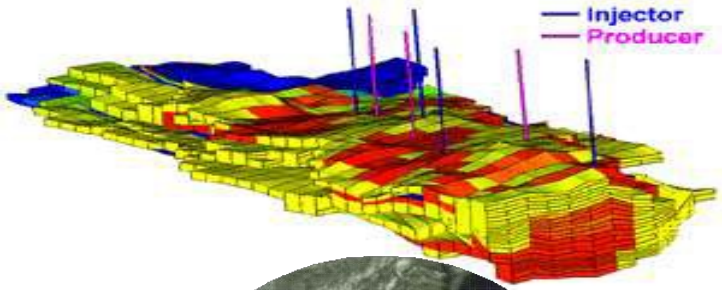

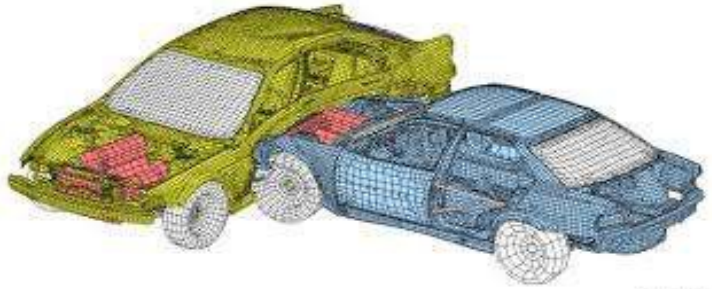

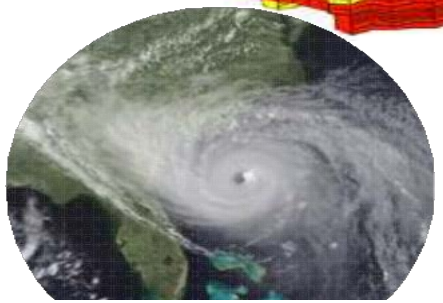

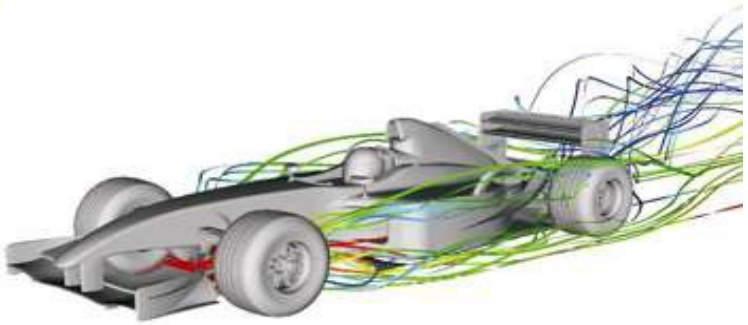

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

WHO IS USING AI?

Pointnext Services

Computer Vision / Video & Image Analytics	Security | Safety | Object Detection, Tracking, Classification
Natural Language Processing	Text Analytics | Speech to Text | Call Categorization

Healthcare & Life Sciences
- Cryo-EM
- Genomics

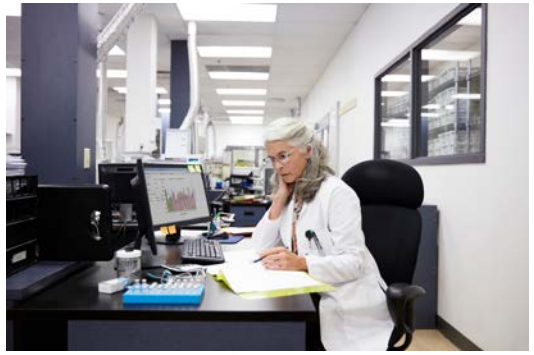

Manufacturing
- Condition Monitoring
- Quality Assurance
- Digital Prescriptive Maintenance
- Autonomous Driving

Financial Services
- Fraud Management
- Call Analytics
- Compliance

Other
- Signal Identification
- Intelligent Building
- Back to Work
- Weather

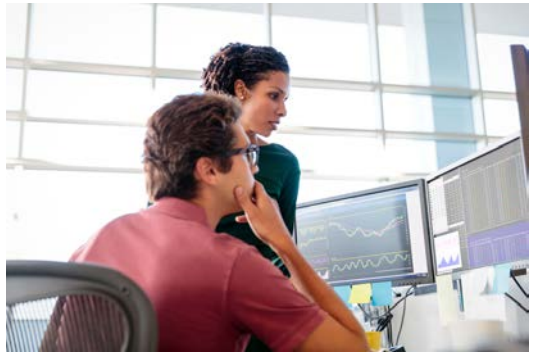

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPC ARCHITECTURE AND HPE PRODCUT PORTFOLIO

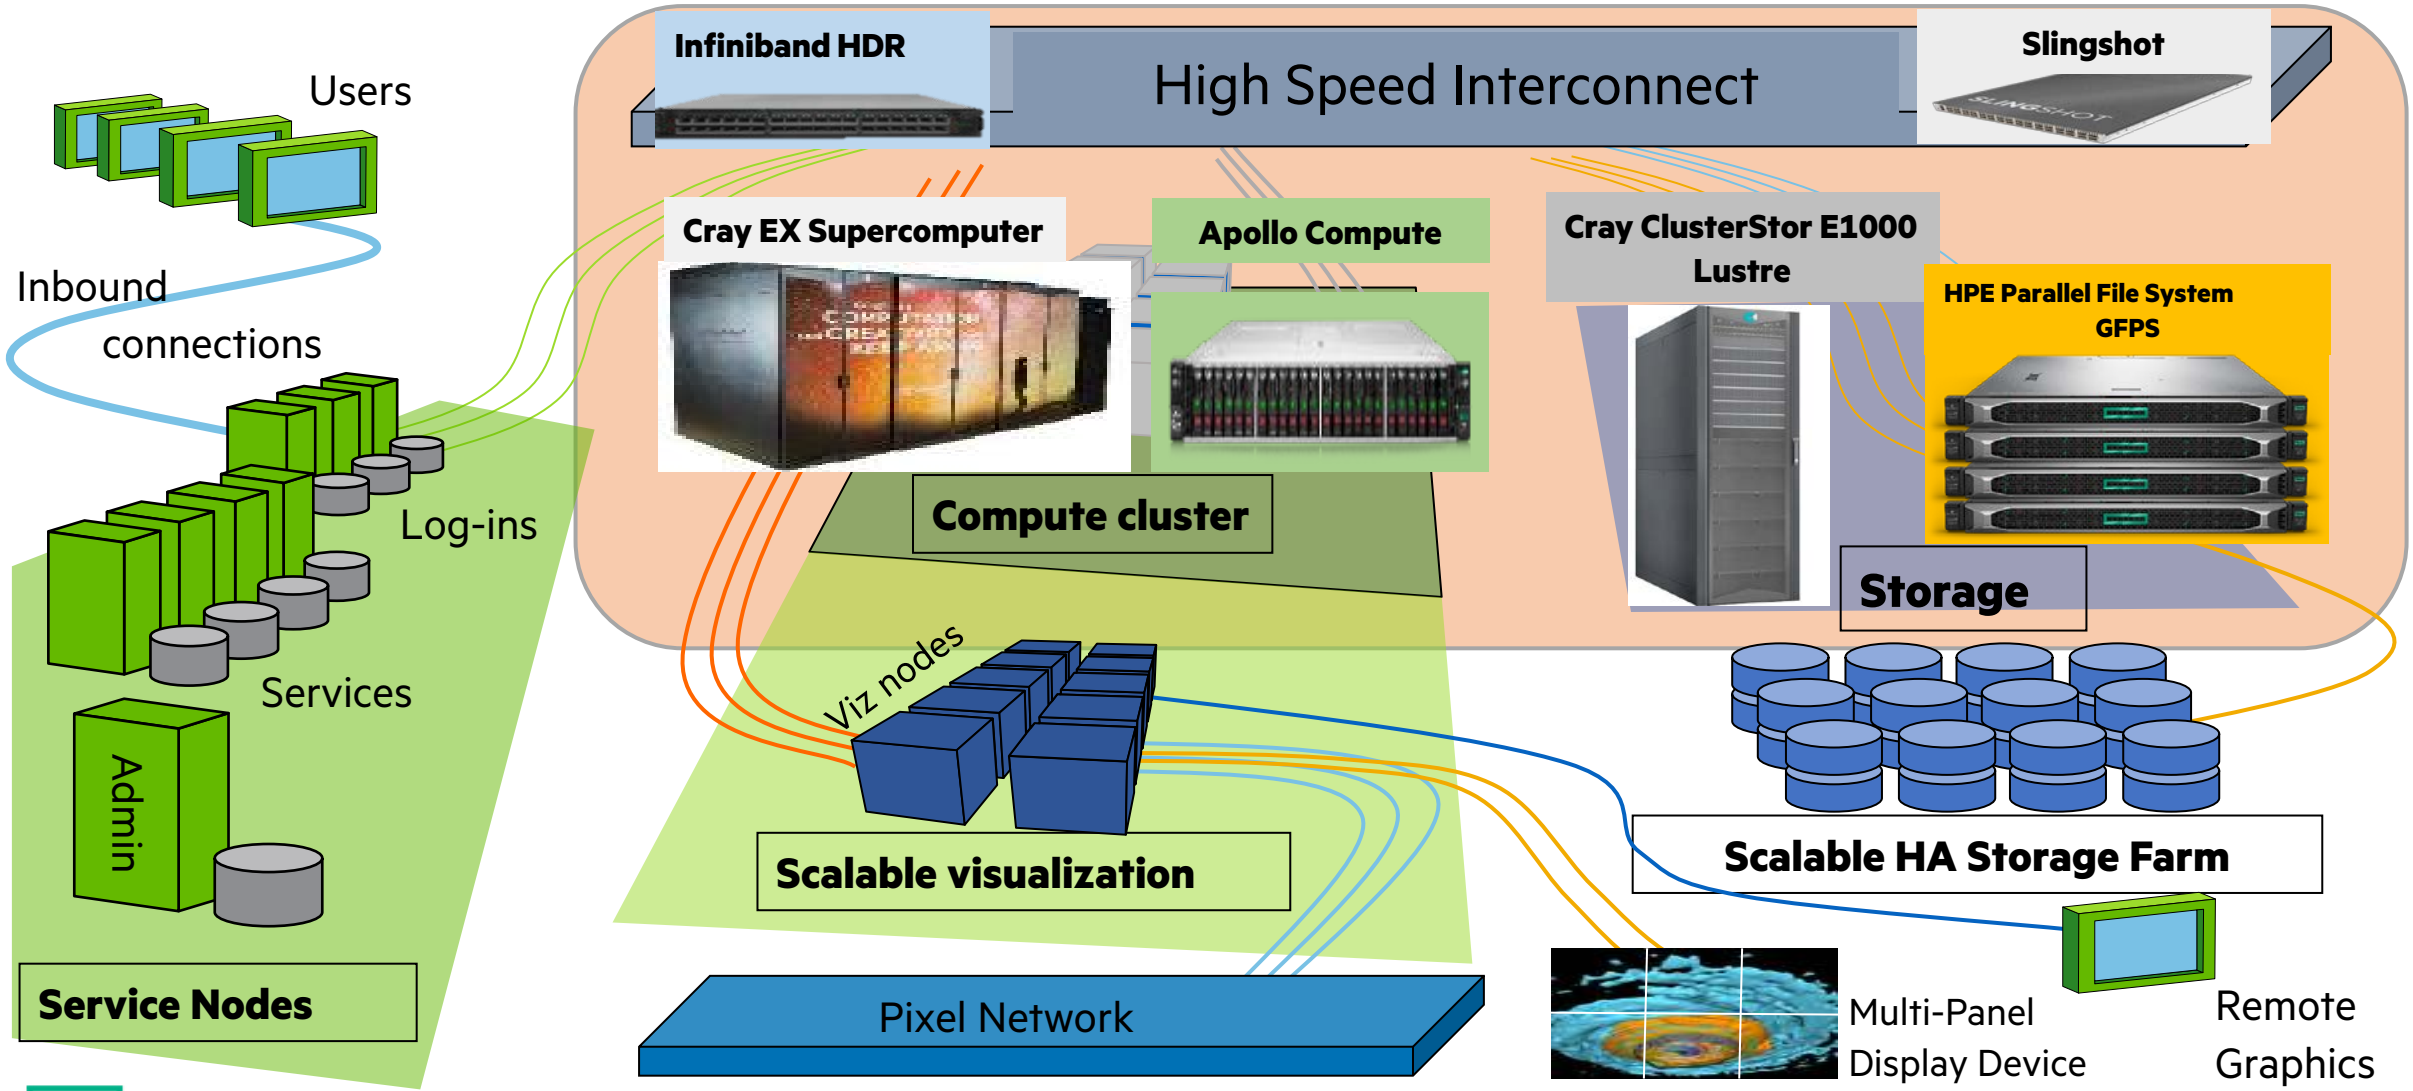

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

— 슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE PURPOSE BUILT HPC/AI SYSTEMS

HPE Apollo 6500 Gen10 Plus System

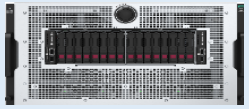

Accelerated compute platform for HPC and AI workloads

HPE Apollo 2000 Gen10 Plus System

Density-optimized, scale out compute for HPC and AI workloads

HPE Superdome Flex

Near limitless scale-up compute/memory for HPC and AI

HPE Cray Supercomputers Air-Cooled

Air-cooled, Converged HPC-AI

HPE Cray EX Liquid-Cooled Supercomputers

Liquid-Cooled Converged HPC-AI

Complexity of workloads / workflows

HPC

SCALE

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY SUPERCOMPUTERS WIN THE "EXASCALE TRIPLE-CROWN"

Argonne National Laboratory "Aurora"

- More than **1.5 EF** sustained performance
- Future Intel® Xeon® CPU and Intel Xᵉ architecture and Slingshot switch
- Mixed AI and HPC workload

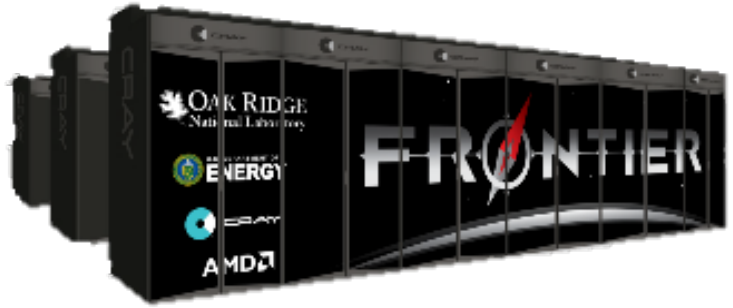

Oak Ridge National Laboratory "Frontier"

- More than **1.5 EF** sustained performance
- Future AMD® EPYC™ CPU and Radeon GPU and Slingshot switch
- Mixed AI and HPC workload

Lawrence Livermore National Laboratory "El Capitan"

- More than **2.0 EF** sustained performance
- Future AMD® GPU and Slingshot switch
- Mixed AI and HPC workload

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY SUPERCOMPUTERS WIN

National Energy Research Scientific Computing Center (NERSC)

- AMD® EPYC™ CPU and Nvidia A100 GPUs with NVLink
- Mixed AI and HPC workload

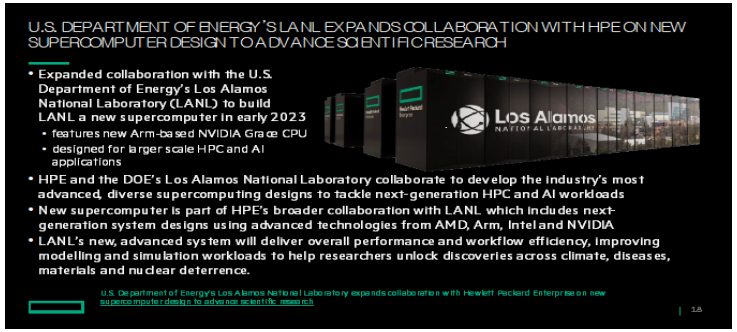

Los Alamos National Laboratory (LANL)

- More than **13.3 PF** sustained performance
- Features new Arm-based NVIDIA Grace CPU
- AMD, Intel and NVIDIA
- Mixed AI and HPC workload

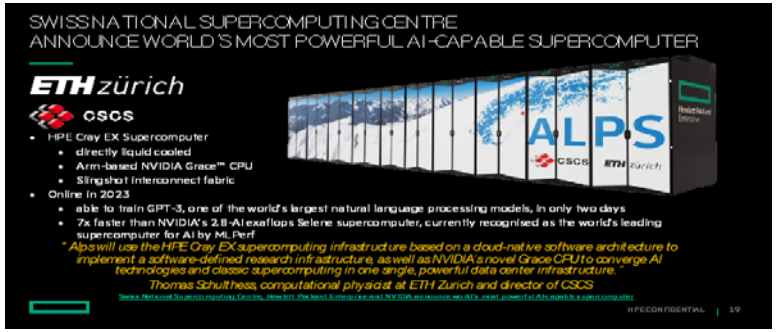

Swiss National Supercomputing Centre (CSCS)

- Arm-based NVIDIA Grace CPU and NVIDIA GPU and Slingshot switch
- Mixed AI and HPC workload

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

— 슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY EX

HPE CRAY EX SUPERCOMPUTER

Performance

- **Highest power CPUs supported via direct liquid cooling**
- **Up to 16 Slingshot injection ports per compute blade**
- **Hardware & Software scalable to Exascale class systems**

TCO

- **Warm water cooling (W3 and W4 temps supported)**
- **Efficient power conversion from mains to point-of-load**
- **Upgradeable for multiple technology generations**

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY EX INFRASTRUCTURE

Architected for maximum performance, density, efficiency, and scale

- Up to 64 compute blades, and **512 processors per rack**
- Flexible bladed architecture supports **multiple generations** of CPUs, GPUs, and interconnect
- **Cableless interconnect** between switches and nodes inside chassis
- **100% direct liquid cooling** enables 300kVA capability per rack
- Scales to 100's of cabinets

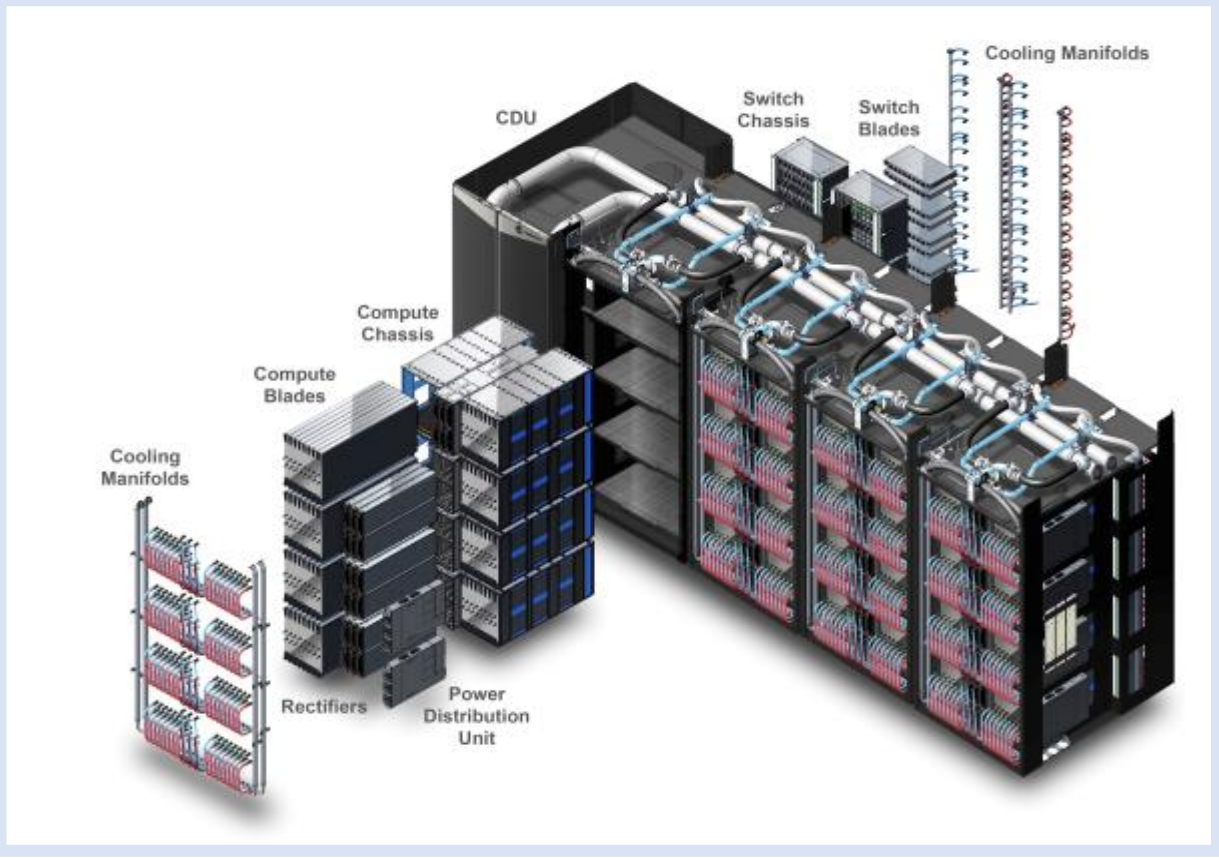

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

— 슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY EX

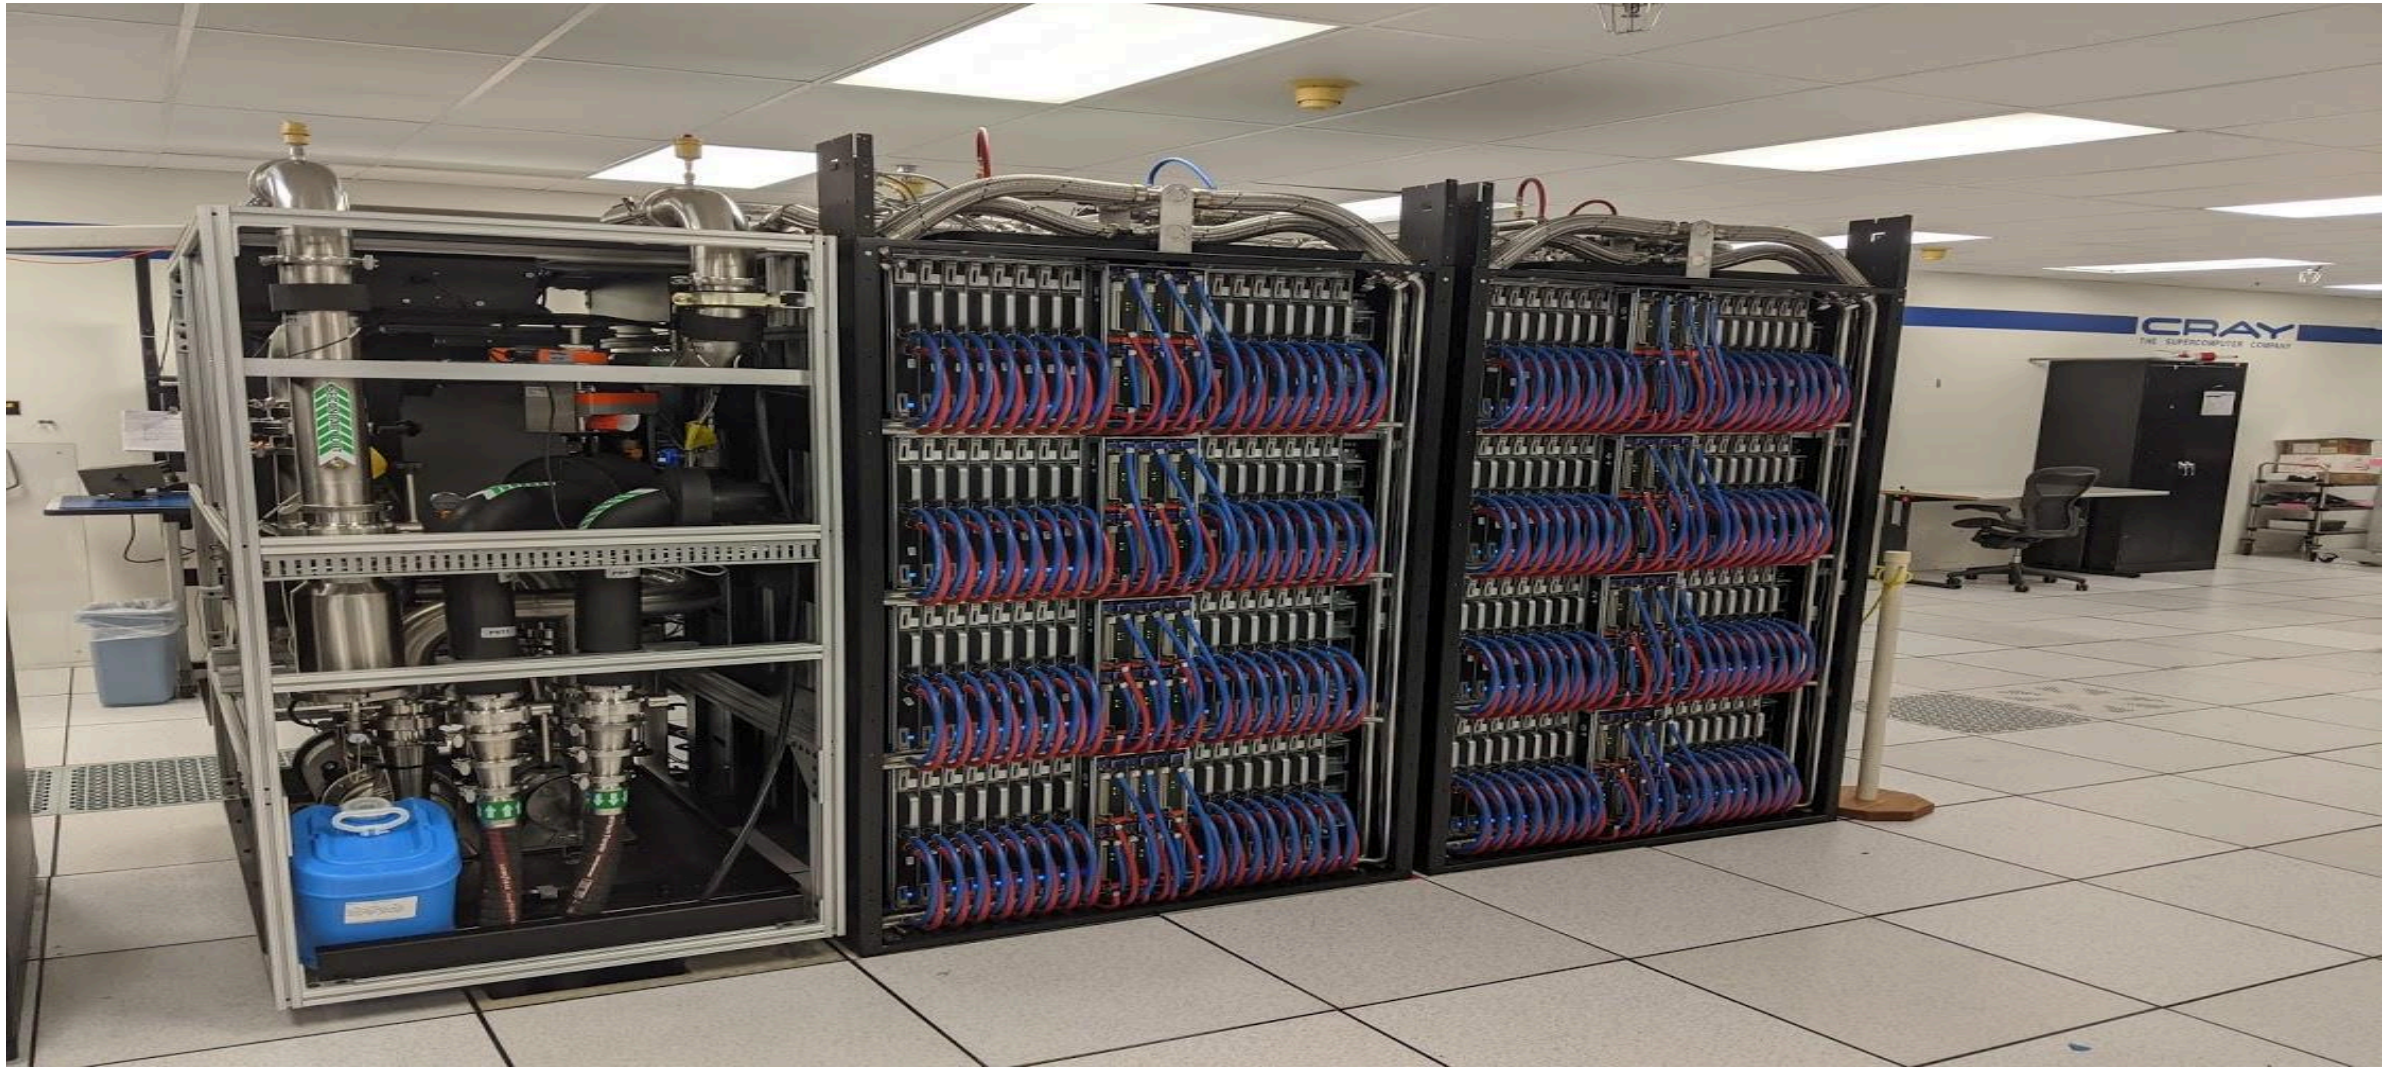

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY EX CABINETS

Front

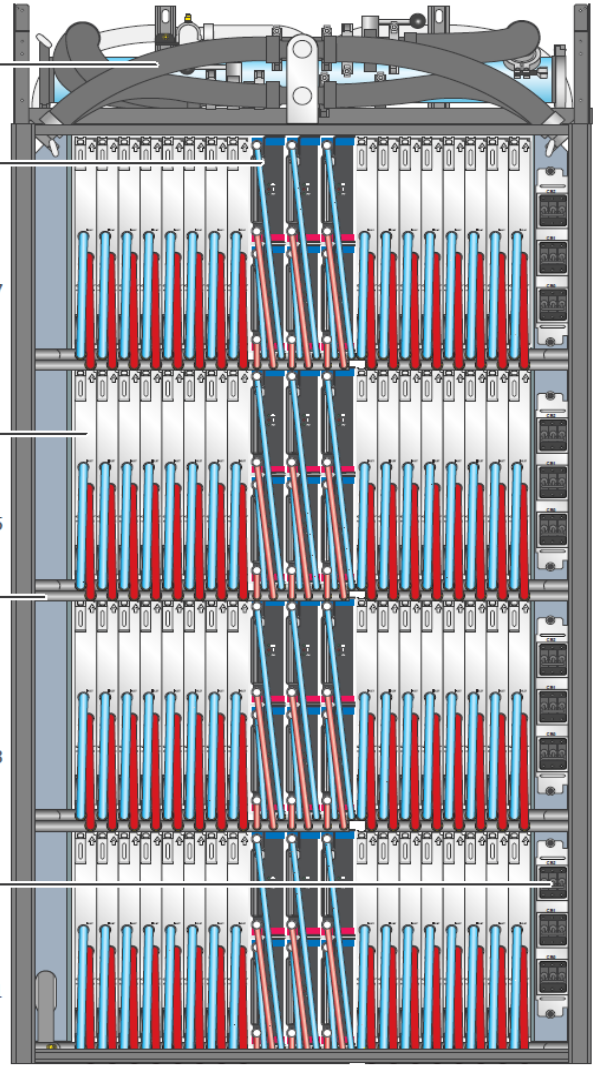

Rear

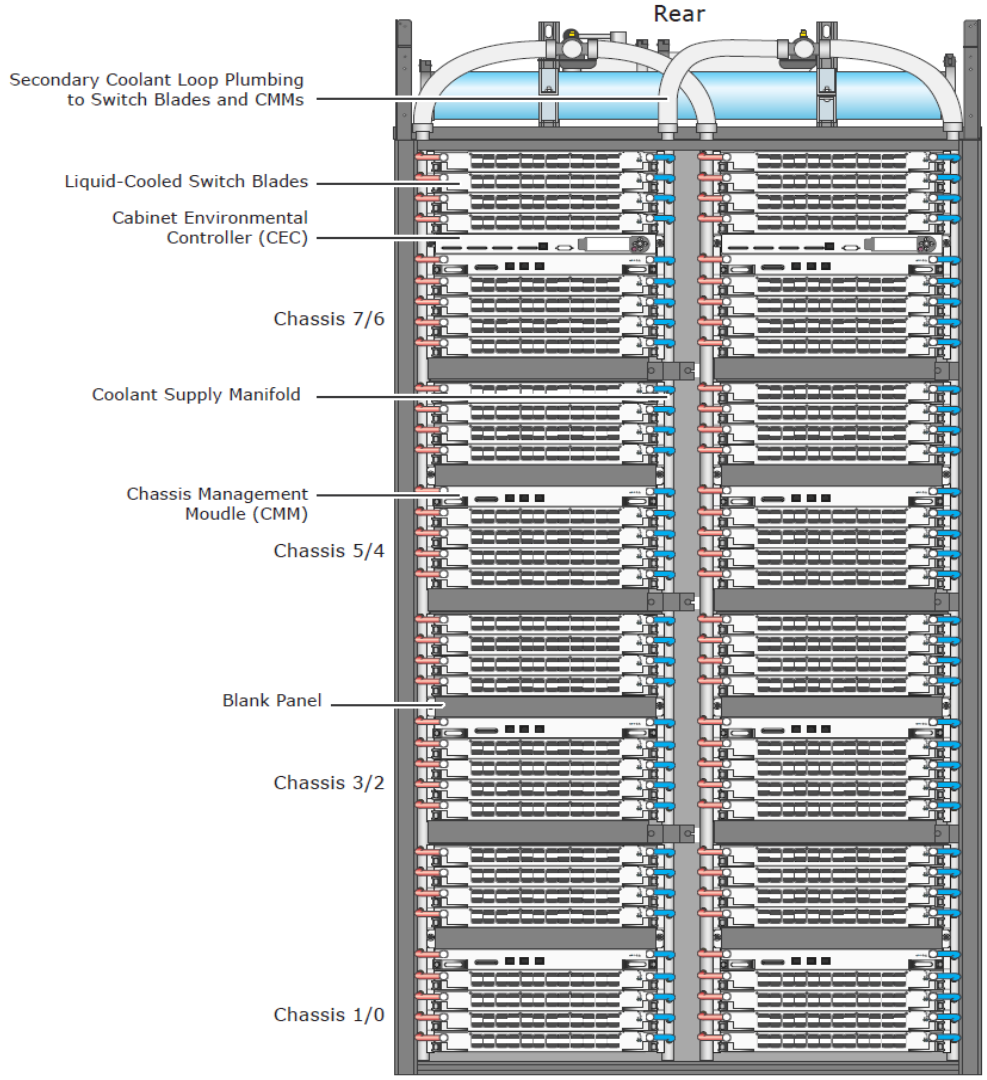

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HIGH-LEVEL ARCHITECUTRE

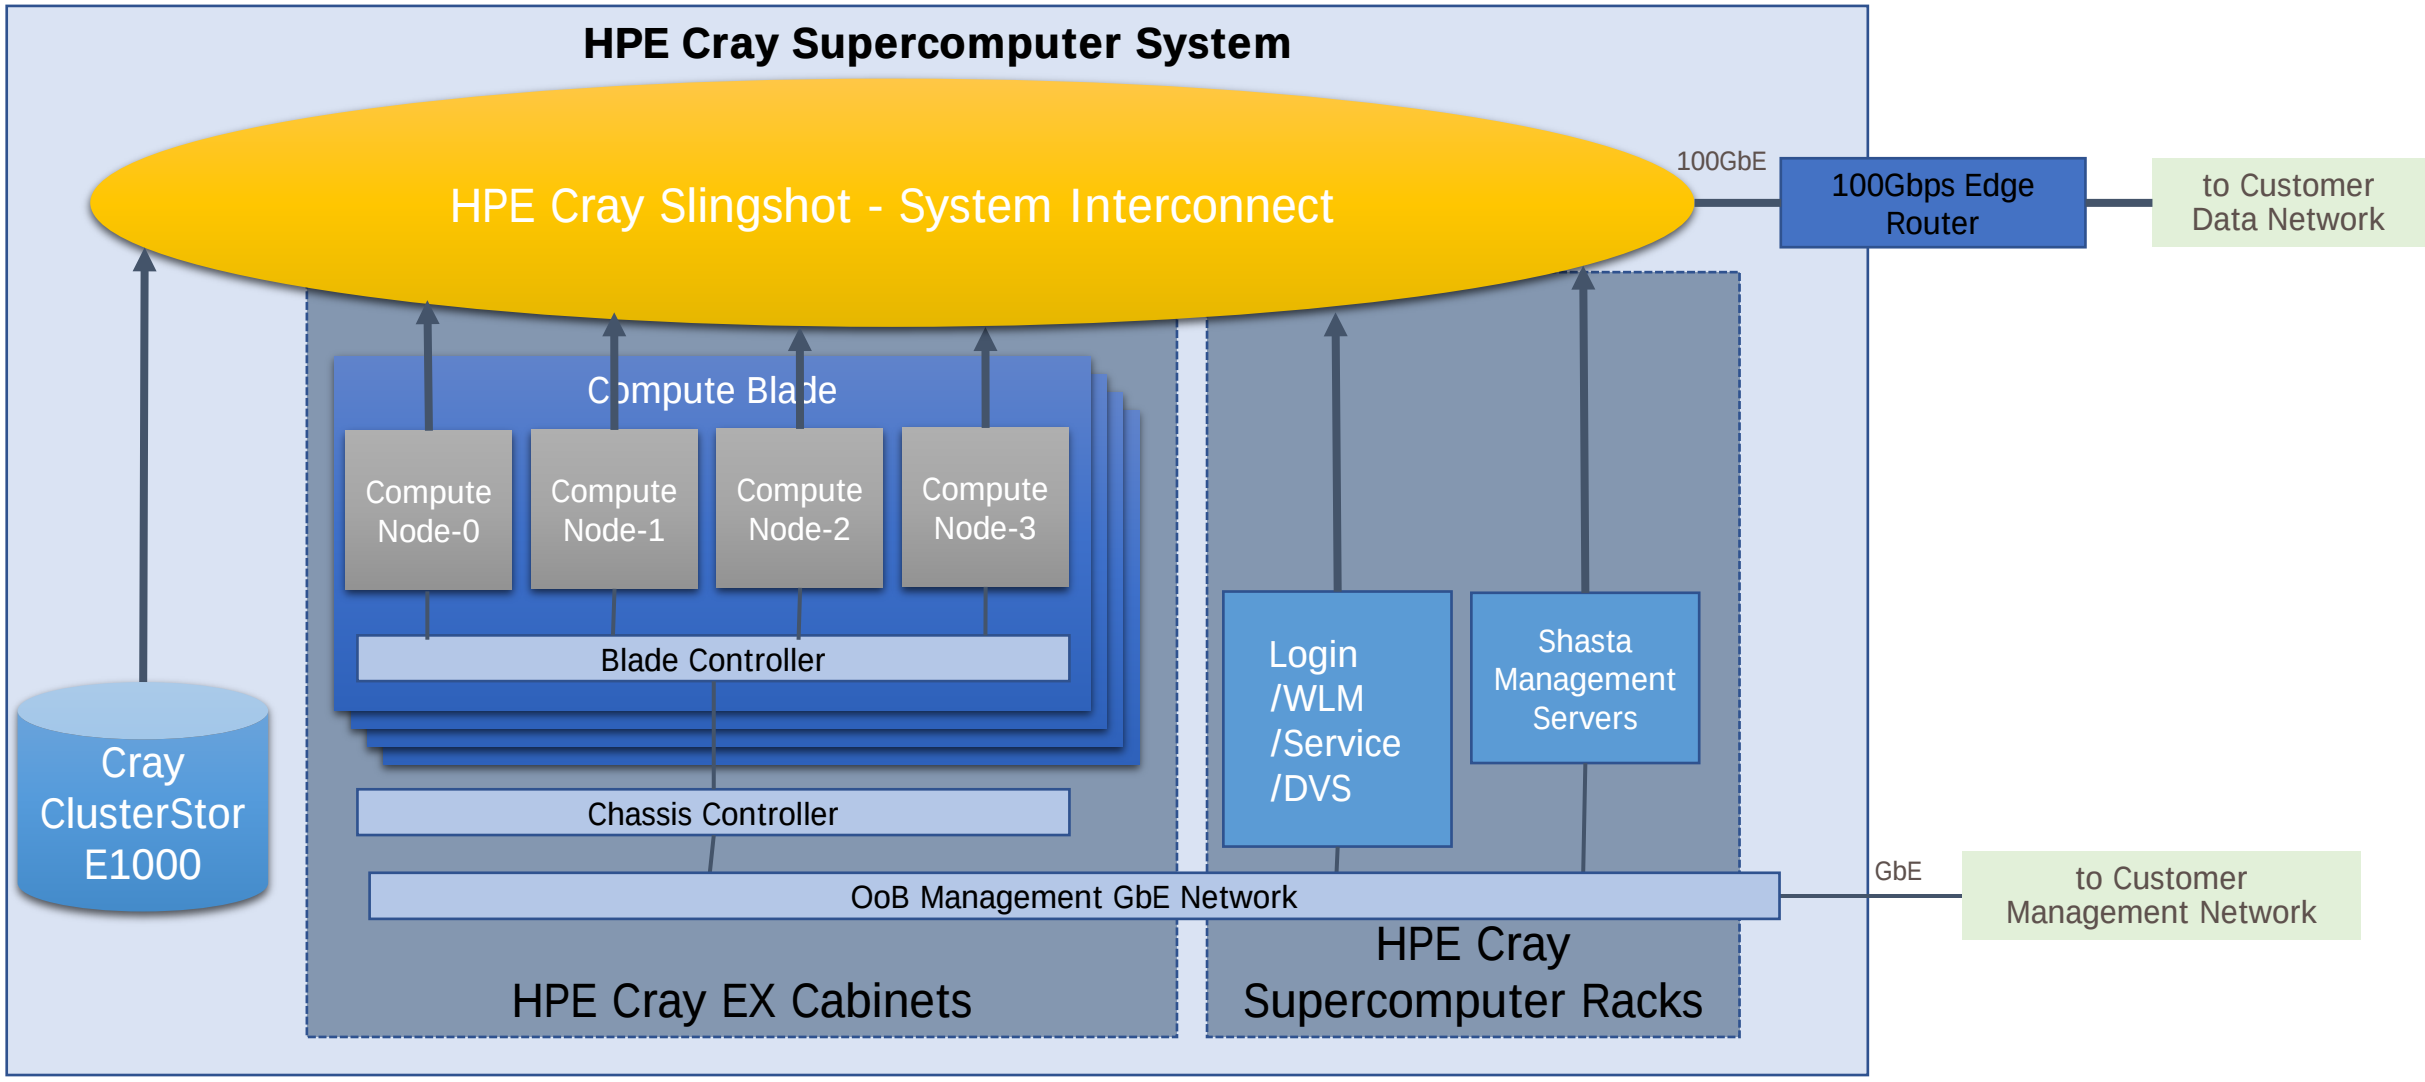

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

ROME/MILAN COMPUTE NODES

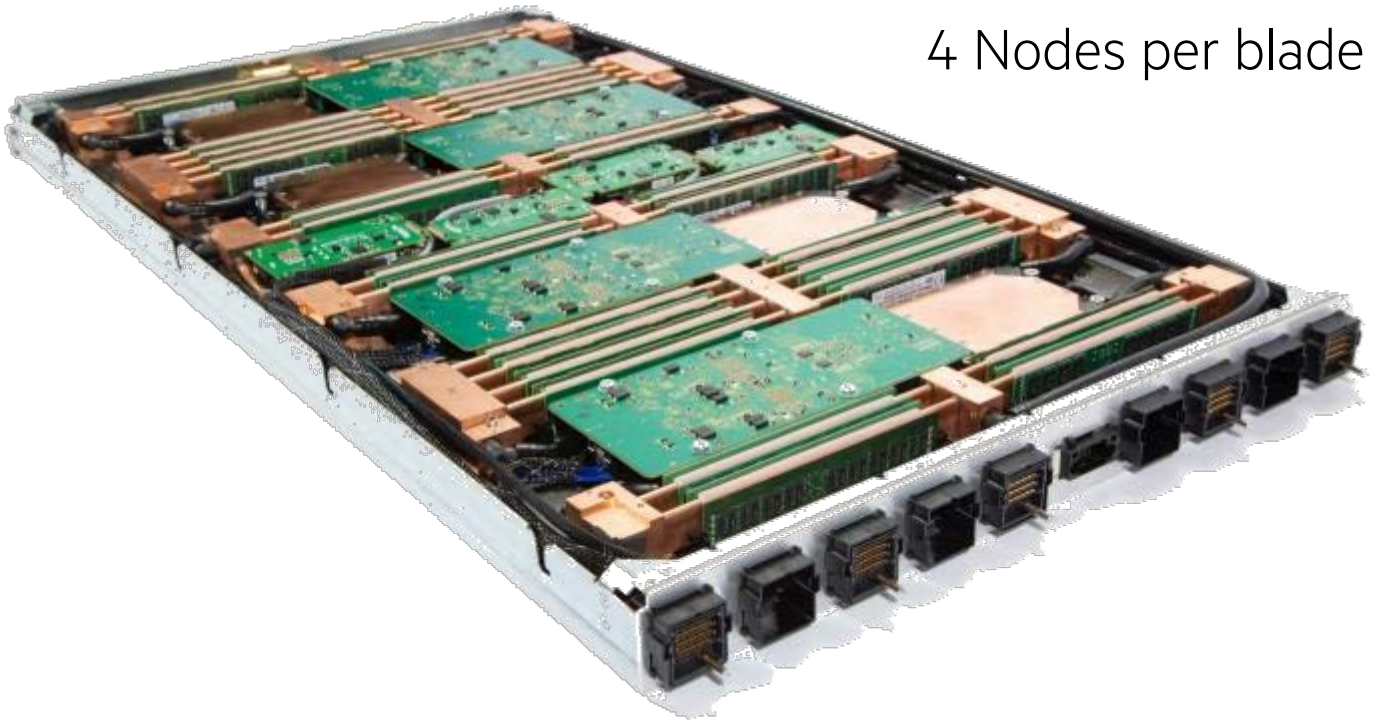

4 Nodes per blade

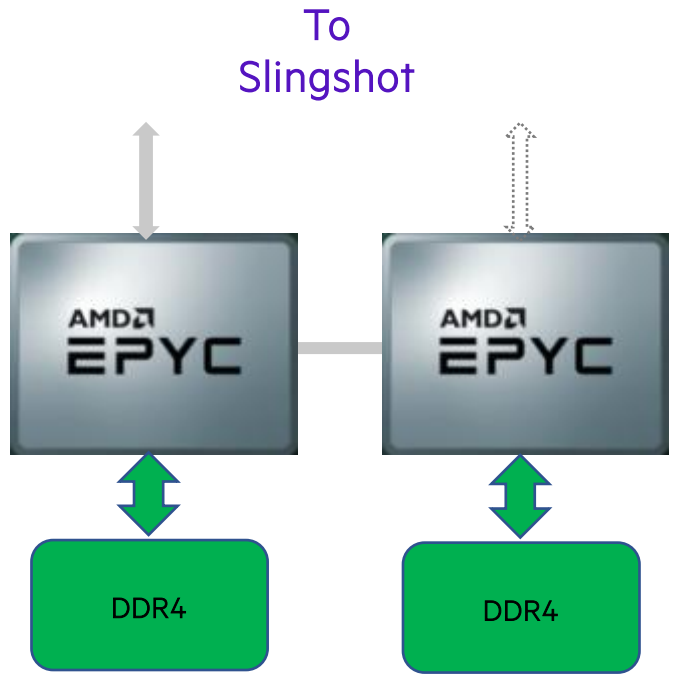

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

GRIZZLY PEAK BLADE

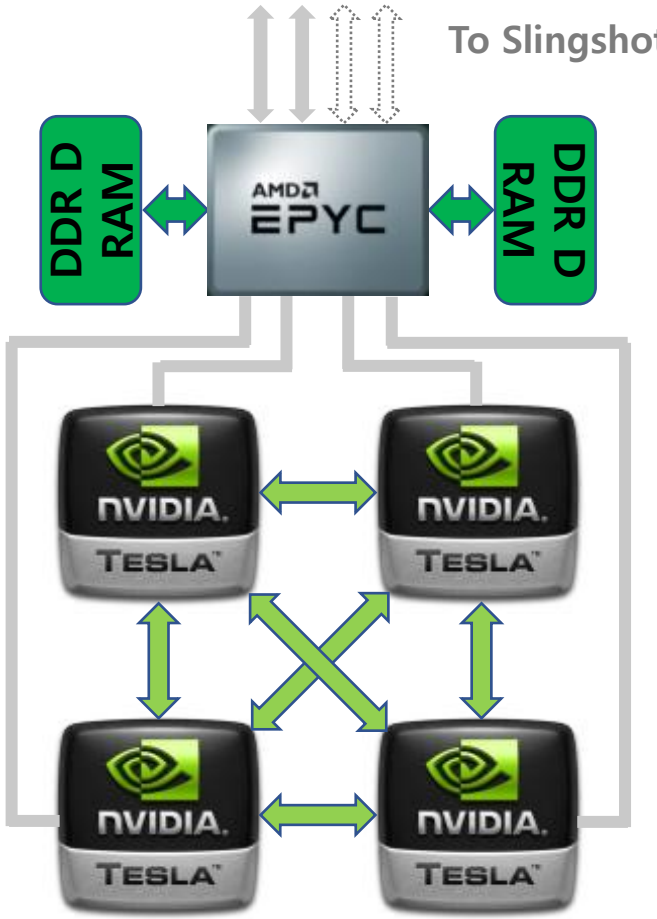

X 2 per blade

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

— 슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY EX COOLING - CDU

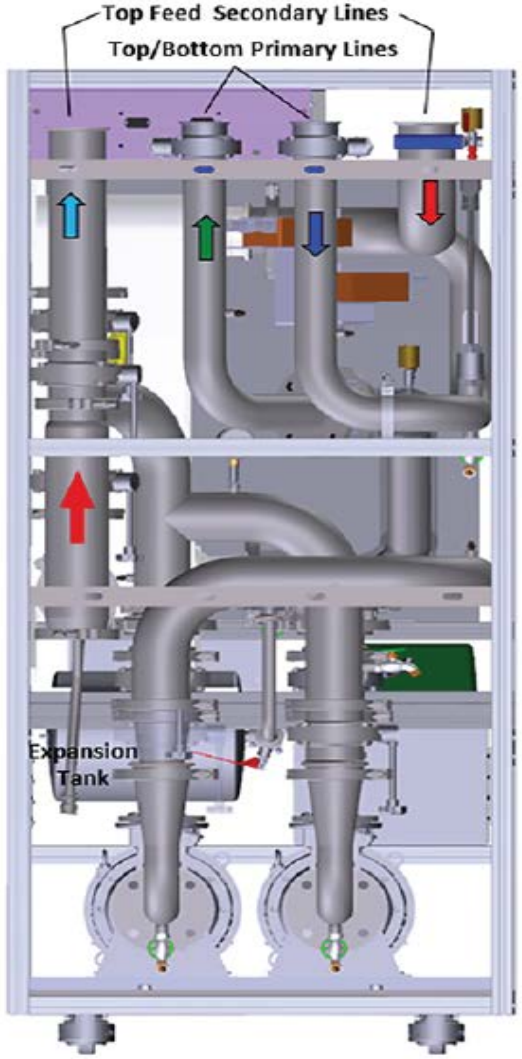

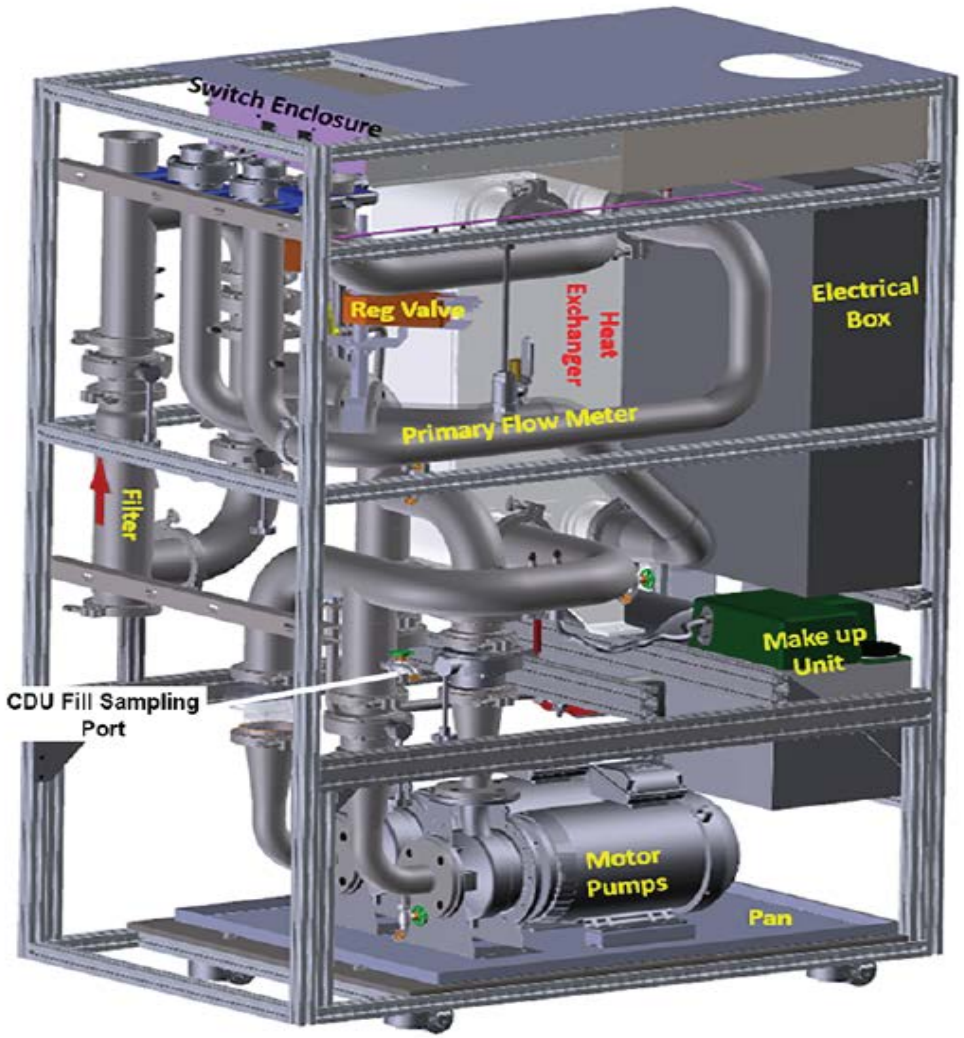

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY SLINGSHOT

High Performance Ethernet Network

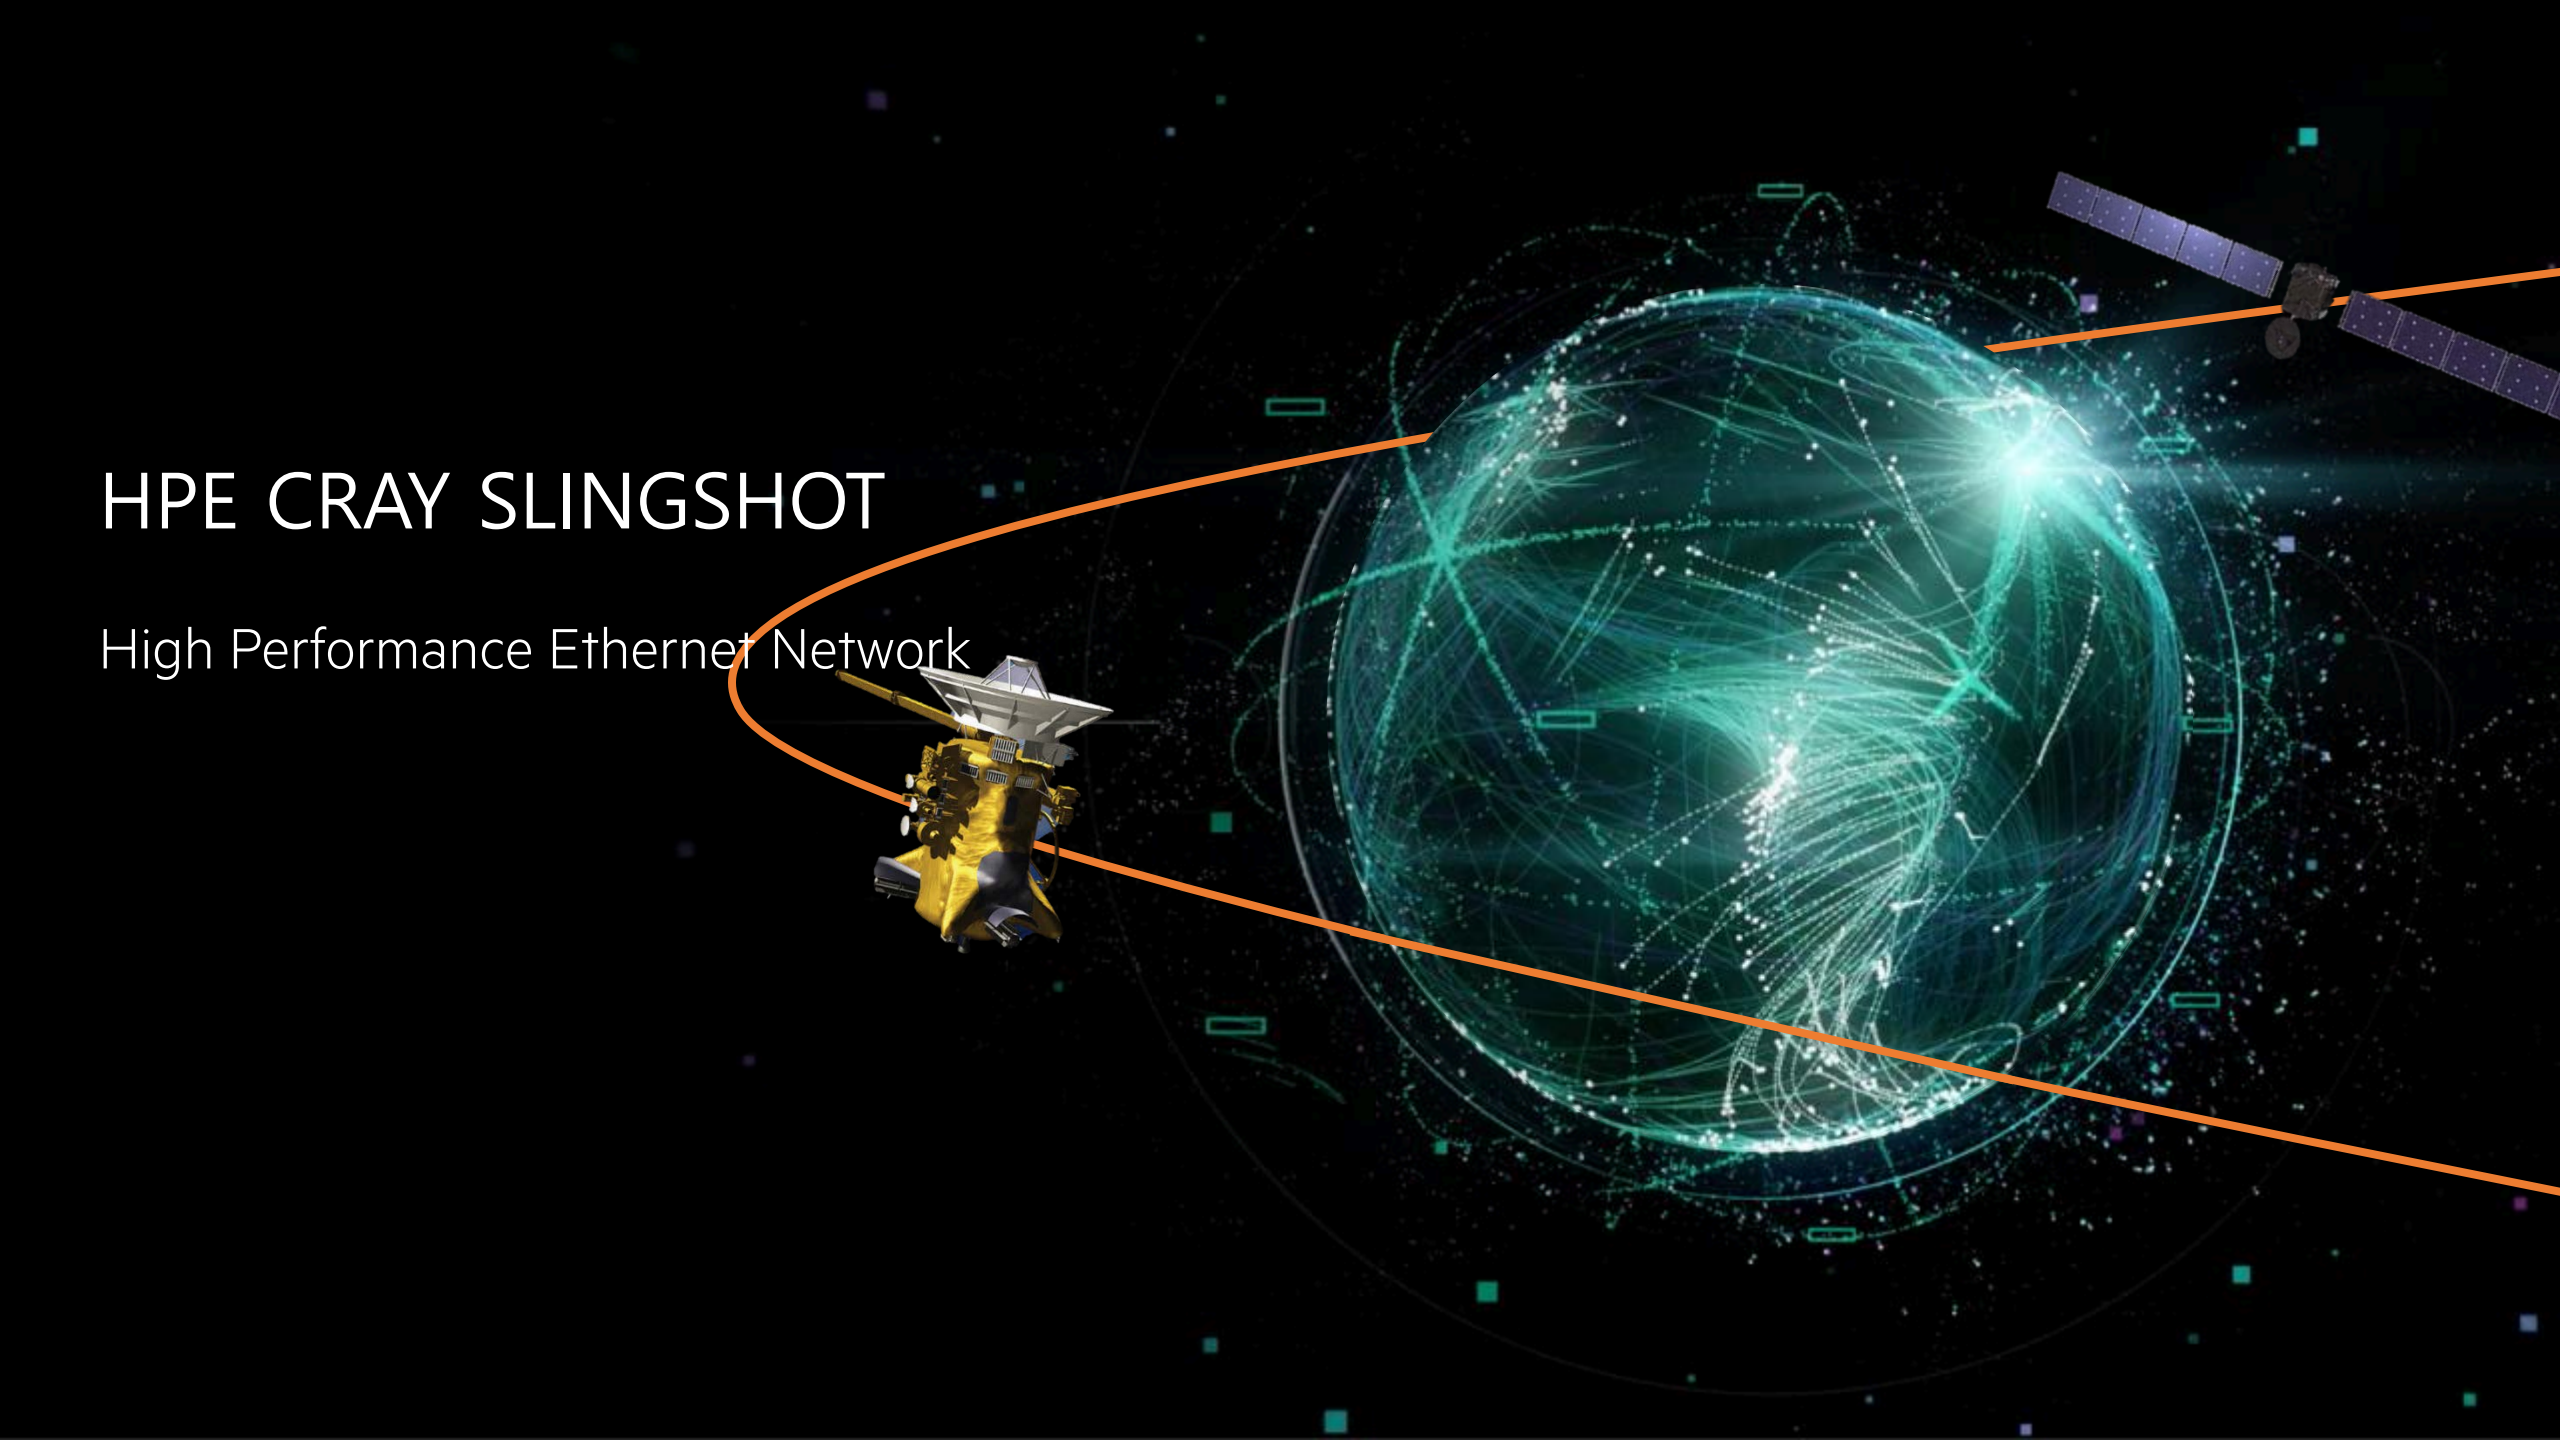

A NEW APPROACH FROM OUR PRIOR GENERATIONS

Cray Aries Chip

Switch

NICs

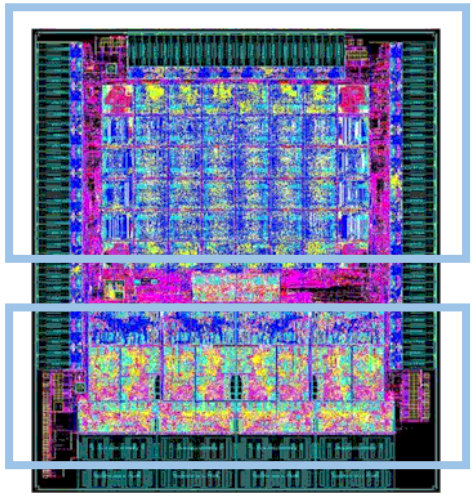

- **Low latency**
- **Great performance**
- **Cost effective at scale**
- **Great solution…**

Slingshot Chips

Switch

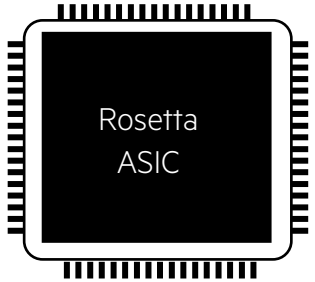

NICs

New Flexibility Enables:

- ✓ **More Value**
- ✓ **More Opportunity**

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

SLINGSHOT – ENABLING THE NEXT ERA OF COMPUTING

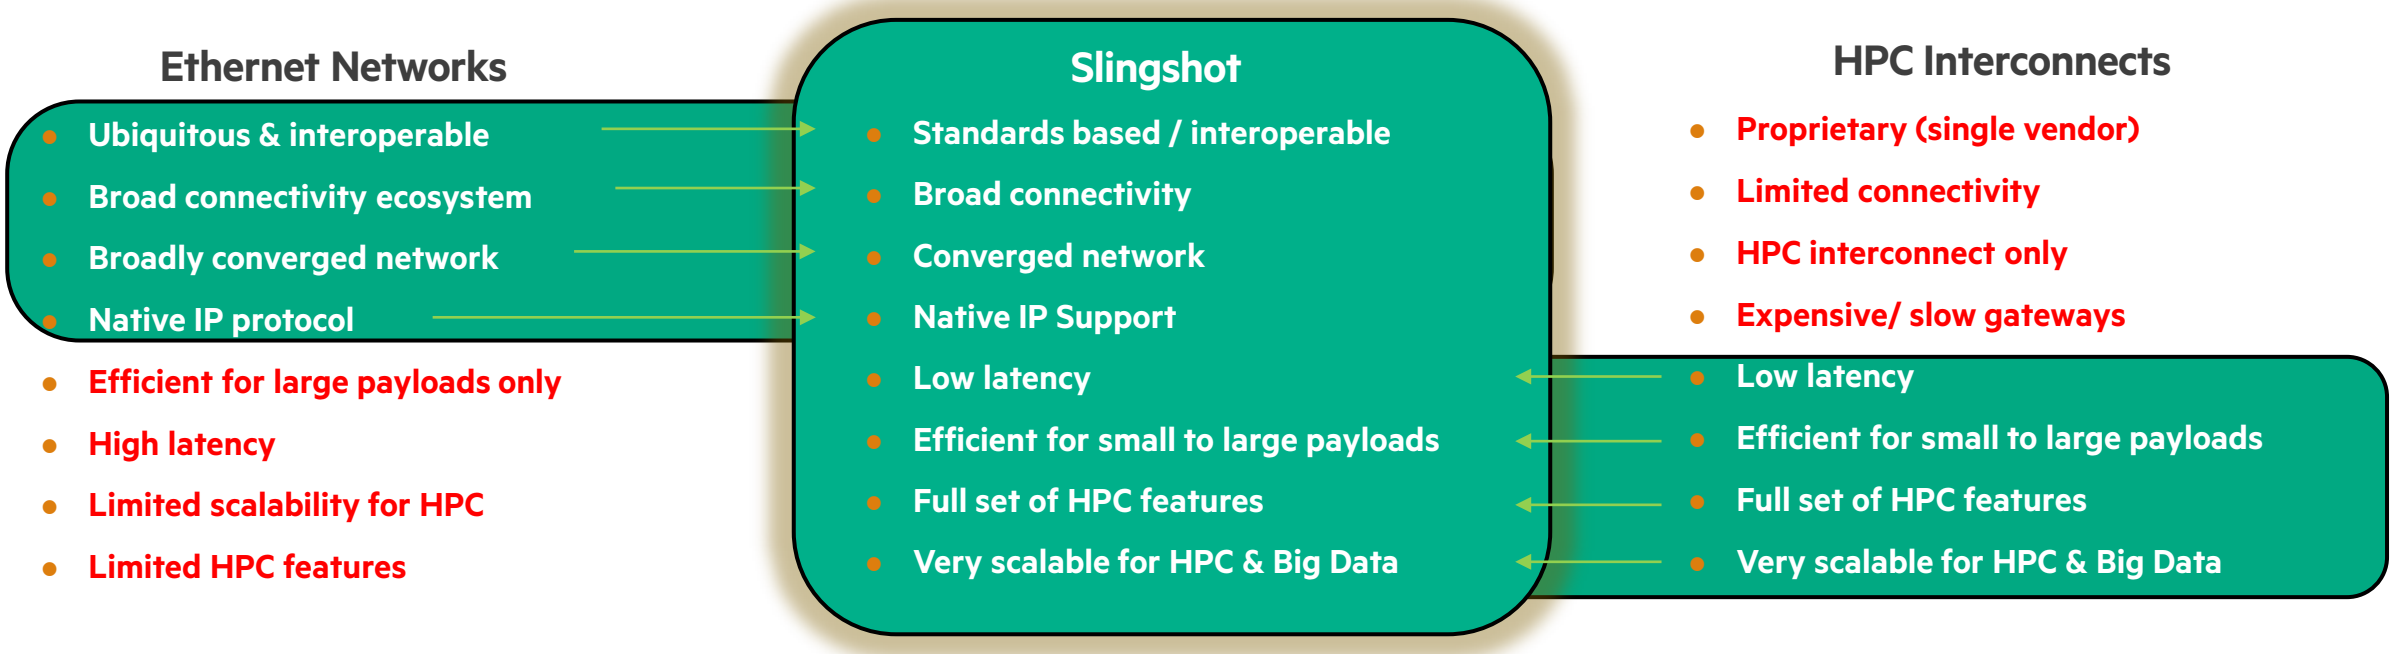

- ✓ **Consistent, predictable reliable** high performance - from one rack to exascale
- ✓ **Native ethernet connectivity** to data center resources
- ✓ **Excellent for emerging infrastructures** that mix tightly coupled HPC, AI, and cloud workloads

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

— 슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

SLINGSHOT ARCHITECTURE

- ✓ **Consistent, predictable reliable** high performance - from one rack to exascale
- ✓ **Excellent for emerging infrastructures** that mix tightly coupled HPC, AI, and cloud workloads
- ✓ **Native connectivity** to data center resources

64 ports x 200 Gbps	Ethernet Compliant	World Class Adaptive Routing and QoS	Efficient Congestion Control	Low, Uniform Latency
Over 250K endpoints with a diameter of just three hops	Easy connectivity to datacenters and third-party storage	High utilization at scale; flawless support for hybrid workloads	Performance isolation between workloads	Focus on tail latency, because real apps synchronize

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

— 슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

INHERITED FROM ARIES: FINE-GRAIN ADAPTIVE ROUTING

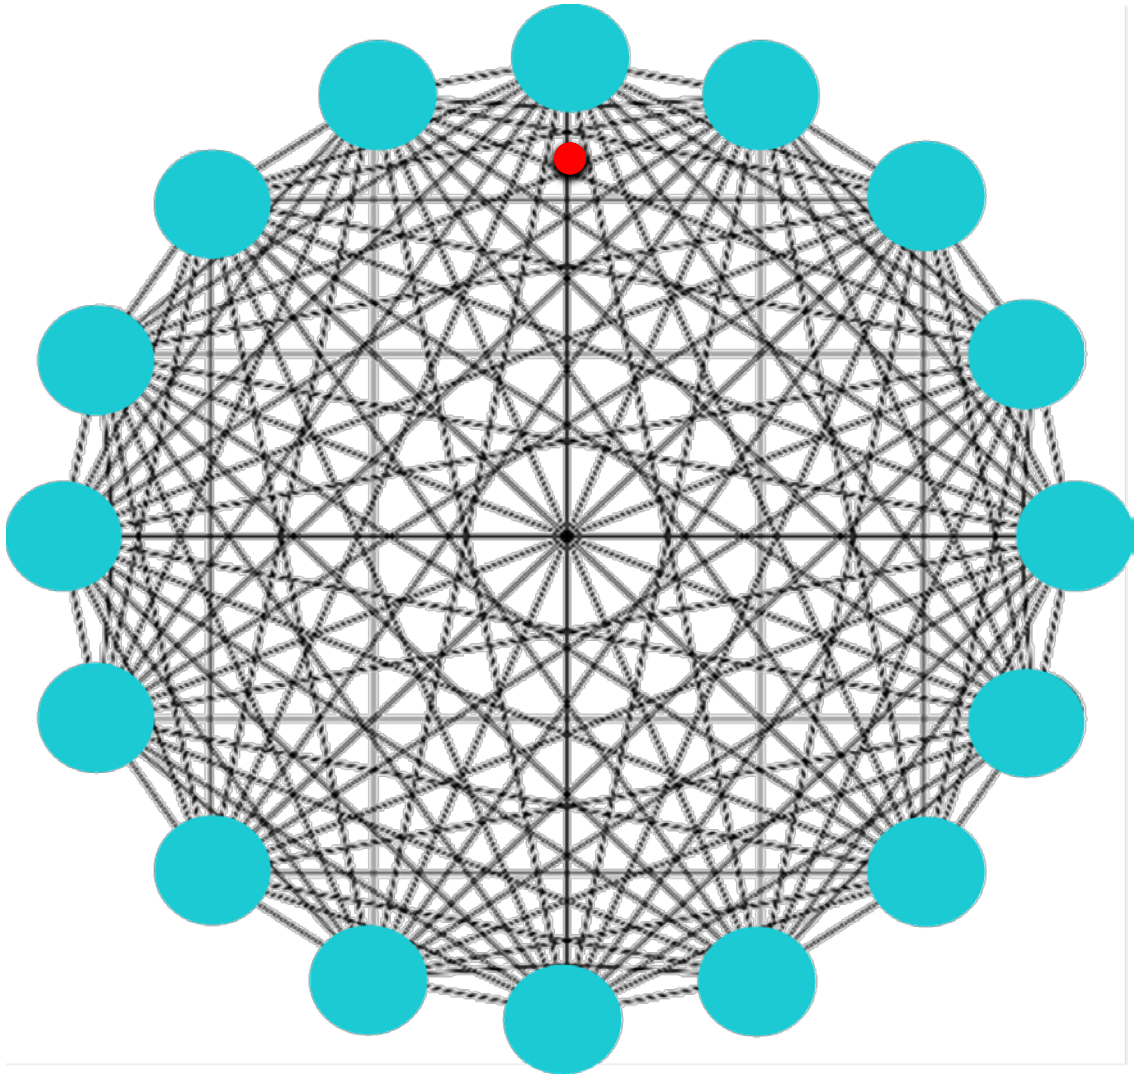

- 16 Switch, 256 Node Group

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

EXTREME SCALE AND PERFORMANCE

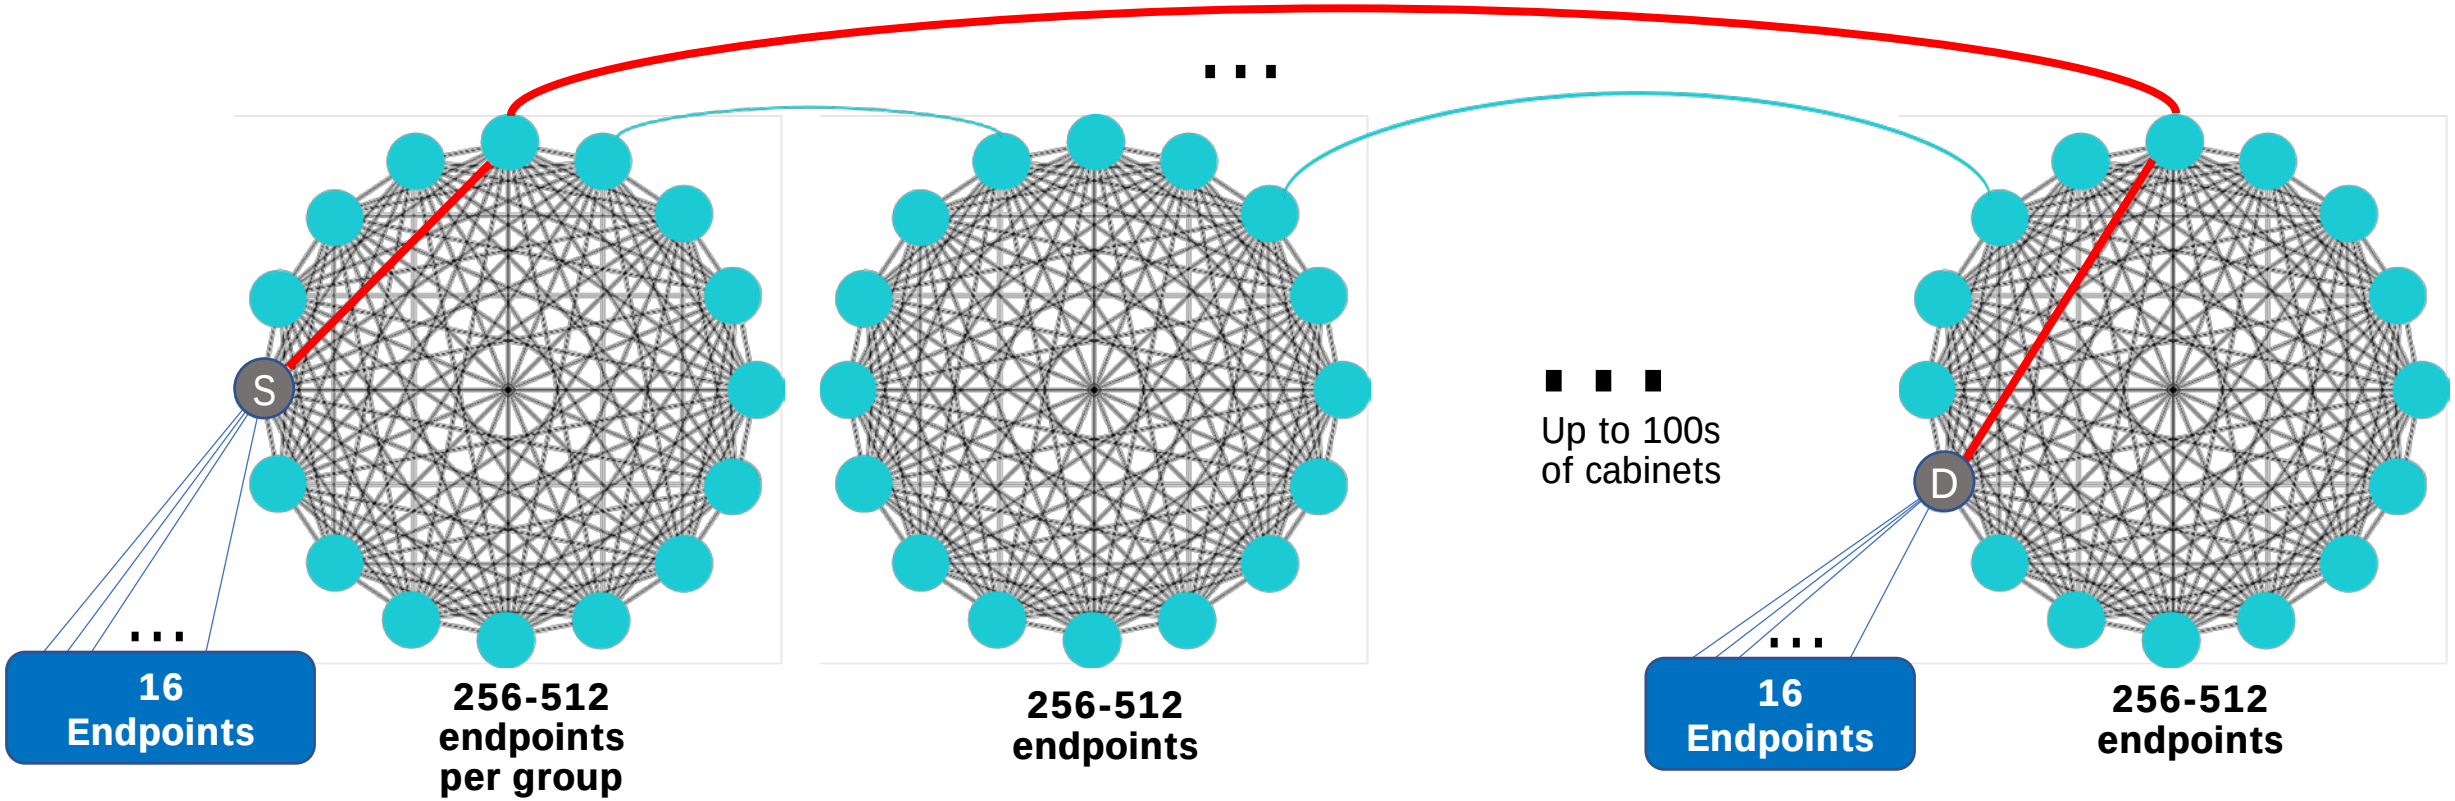

Each Slingshot “Rosetta” Switch is 64 ports x 200 Gbps using 32 cables

Scales to over 250K endpoints with just ***three*** switch-switch hops!

The Cray network behaves as a single, large, logical switch

Adaptive routing achieves > 90% efficiency at scale

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY EX SOFTWARE

Hewlett Packard Enterprise | NVIDIA

HPE HIGH PERFORMANCE COMPUTING SOFTWARE PORTFLOLIO

Application and Software Development Ecosystem	**Development Environments**	**HPE Cray Programming Environment** • C/C++, Fortran, UPC, R, Python • Compiling Environment	**Intel® oneAPI Base and HPC Toolkit (w/Intel MPI)**	**Arm® Allinea Studio**	**AMD AOCC**
			NVIDIA HPC SDK	**NVIDIA GPU Cloud**	**GNU Compilers**
	Debug & Performance	• Debuggers • Performance analysis & optimization tools • Code parallelization assistant	**Arm® Forge Ultimate & Arm DDT**	**TotalView™ by Perforce**	**Vampir**
	MPI	• HPE Cray MPI	**HPE Message Passing Interface (MPI)**	**Open MPI**	**Mellanox® HPC-X™**
Workload Management & Orchestration			**Altair® PBS Professional®**	**Slurm®**	**Kubernetes®**	**Containers: Docker®, Singularity**
Remote Visualization			**NICE DCV and EnginFrame**		
Storage File Systems			**Cray ClusterStor E1000 Storage Solution (Lustre-based)**		
Data Management			**HPE Data Management Framework (DMF)**		
System Management		**HPE Cray supercomputer software** • HPE Cray System Management	**HPE Performance Cluster Manager**	**Bright Cluster Manager®**	
Fabric Software		• HPE Slingshot fabric manager	**Mellanox® Unified Fabric Manager™**	**Intel® Omni-Path Fabric Software**	
Operating System		• HPE Cray Operating System	**SUSE® Linux Enterprise Server**	**Red Hat® Enterprise Linux™**	**TOSS**

HPE Apollo, HPE ProLiant DL, HPE SGI **HPE Apollo 80** **HPE Cray supercomputer** **Cray ClusterStor**

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

— 슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY EX SUPERCOMPUTER SOFTWARE STACKS

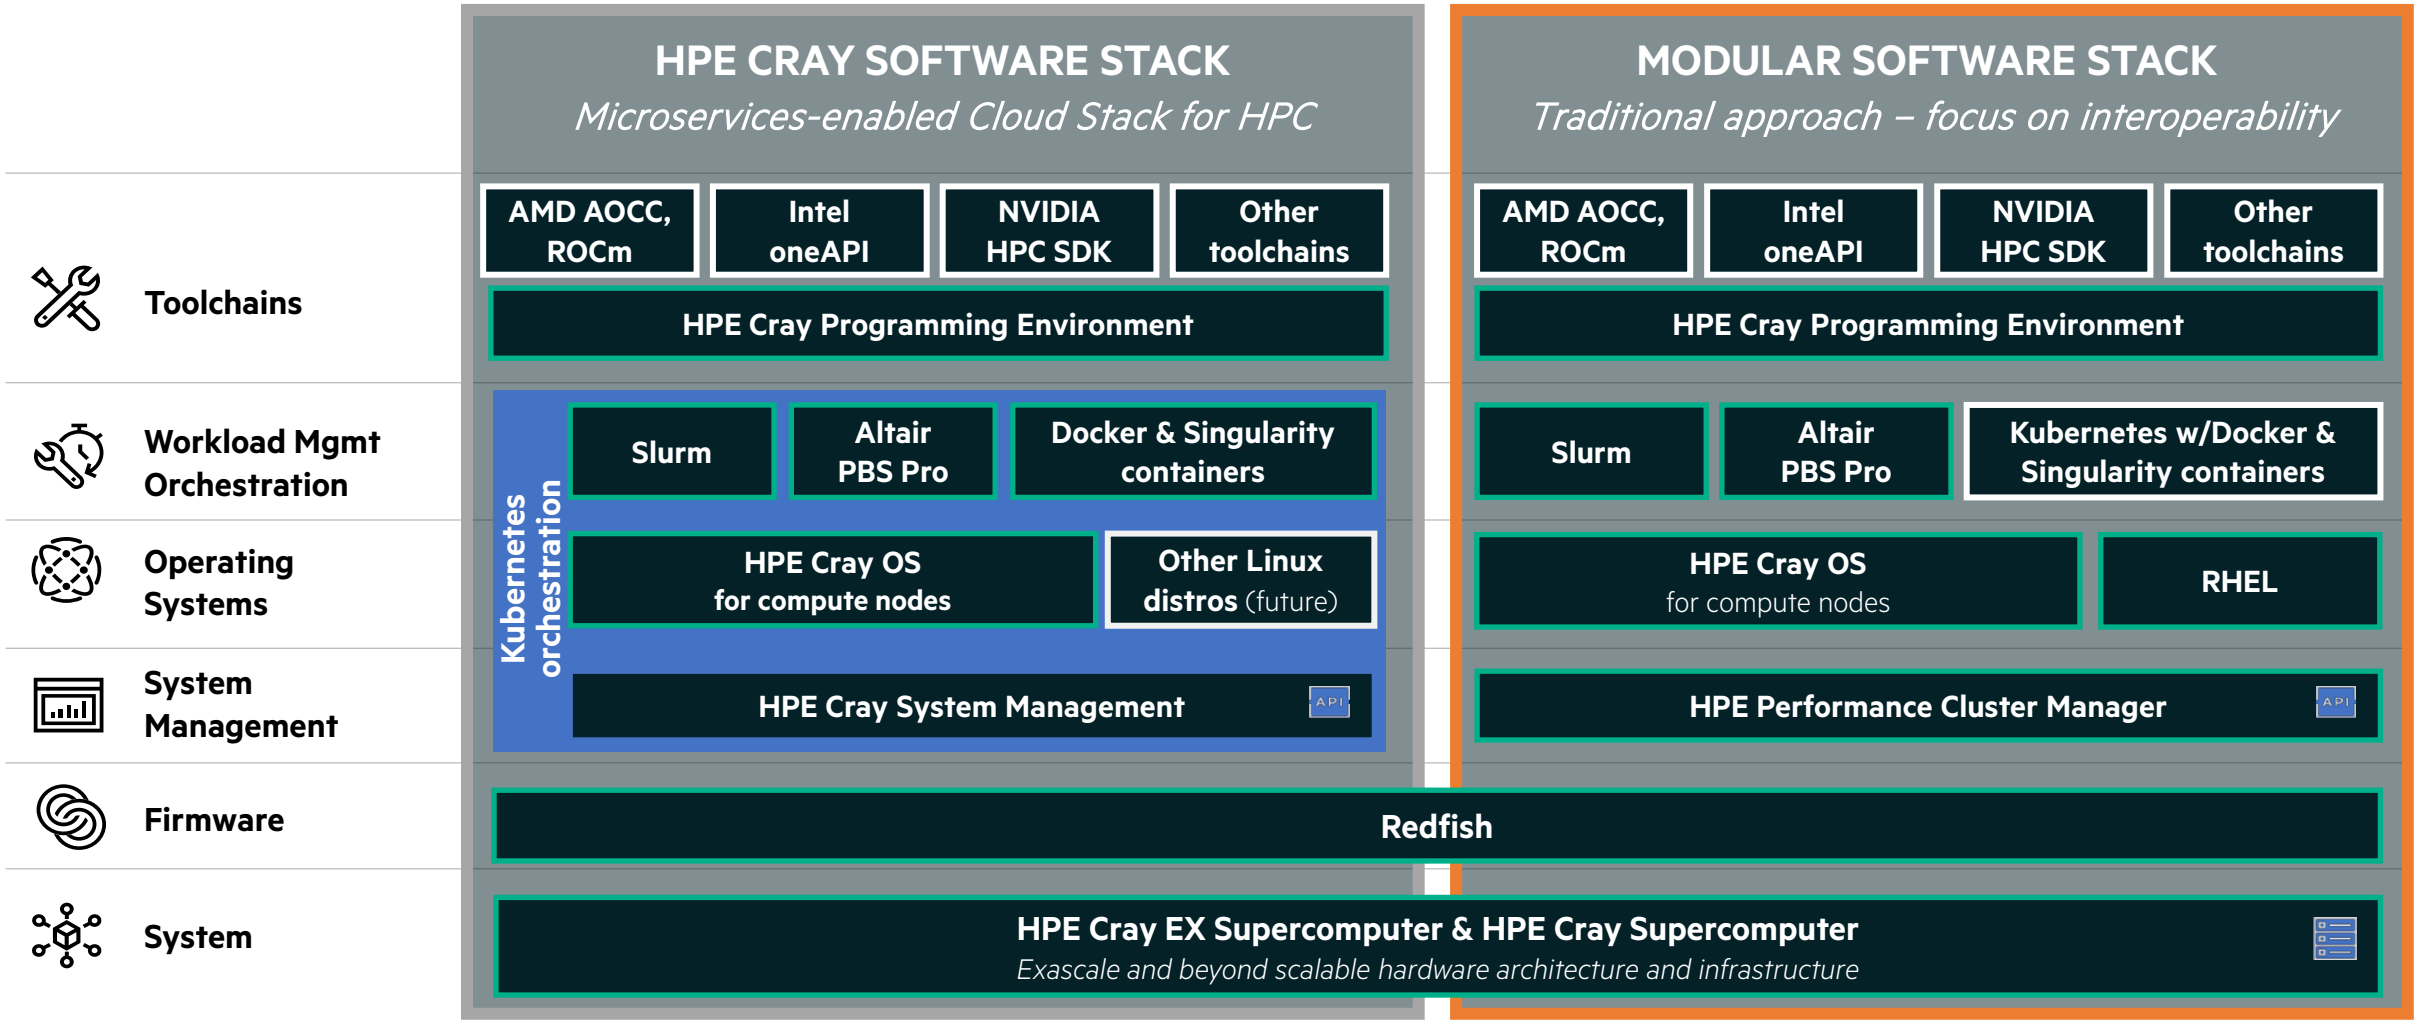

Hewlett Packard Enterprise

NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

HPE CRAY PROGRAMMING ENVIRONMENT

Comprehensive set of tools for developing, porting, debugging, and tuning of HPC applications on HPE & HPE Cray systems

Development

Programming Models

- HPE Cray MPI
- SHMEM
- OpenMP | OpenACC 2.0
- AMD ROCm HIP | NVIDIA CUDA
- UPC | Fortran co-arrays
- Global Arrays

Programming Environments

- Compiling Environment
- GNU
- Intel Programming Environment
- AMD Programming Environment
- NVIDIA HPC SDK

Programming Languages

- C
- C++
- R
- Python
- Fortran

Optimized Libraries

- LibSci (BLAS)
- LAPACK & ScaLAPACK
- LibSci_ACC
- IRT
- FFTW

I/O Libraries

- NetCDF
- Adios 2
- HDF5

DL / AI Tools

- Deep Learning Plug-in

Debugging

Comparative Debugger
Compare two versions of an application

GDB for HPC
Parallelized gdb for HPC

Valgrind for HPC
Memory debugging at scale

Stack Trace Analysis Tool
Stack tracing at scale

Tool for Abnormal Termination Processing
Manage core files at scale

TotalView

DDT

Performance Analysis & Optimization

Performance Analysis Tool (PAT)
Whole program performance analysis, exposing wide set of indicators, identifying bottlenecks and automatically generating suggestions to improve performance.

Visualization Tool
Quick assessment of severity of issues

Code Parallelization Assistant
Reveal hidden potential of an application via code restructuring

Setup & Runtime

Environment Setup

- Tool Enablement (for Spack, CMake, EasyBuild, etc)
- Modules / Lmod

Supported systems:

- HPE Cray supercomputers
- HPE Apollo 2000 & Apollo 80 systems
- HPE ProLiant DL
- Legacy Cray systems

HPE –authored | HPE Added-value to 3rd party | 3rd party

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

Hewlett Packard Enterprise | NVIDIA

THE NEXT GENERATION IN SUPERCOMPUTING

슈퍼컴퓨팅, 클라우드, 데이터센터를 잇는 HPE CRAY EX SUPERCOMPUTER

Thank you

Contact: beom-soo.kim@hpe.com