Hewlett Packard
Enterprise

BEYOND SUPER

MAKING APPLICATIONS READY FOR THE EXASCALE ERA

Gabriele Paciucci Sr. Solution Architect EMEA

June, 2020


BEYOND SUPER
POWERING THE EXASCALE ERA


MODELING & SIMULATION
+
ARTIFICIAL INTELLIGENCE
+
BIG DATA ANALYTICS
RUNNING ON ONE MACHINE IN MISSION-CRITICAL WORKFLOWS
EXASCALE ERA


PERFORMS LIKE A SUPERCOMPUTER
RUNS LIKE A CLOUD

ADAPTIVE SUPERCOMPUTING

HPE's Cray Shasta supercomputer is focused on delivering innovative next-generation systems that integrate diverse processing technologies at the node level into a unified architecture, allowing customers to meet their users' continued demand for higher sustained performance.

- Flexibility in node design.
- Full software and user programming environment.
- Scalable HPC and storage network.
- Predictable HPC performance at scale.
- Cloud service delivery models.
- Support for Multi-Tenancy.

2023

AMD + AMD

Genoa Next-Gen

2022

INTEL + INTEL

SPR PVC

2021

AMD + AMD

Trento MI200

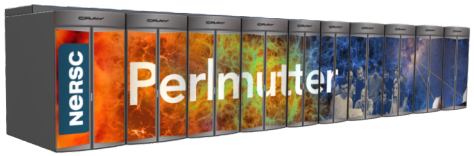

2020

AMD + NVIDIA

Milan Volta Next

WHAT MAKES CRAY SHASTA UNIQUE

SHASTA OLYMPUS 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 300KW capability per rack
- Scales to 100's of cabinets

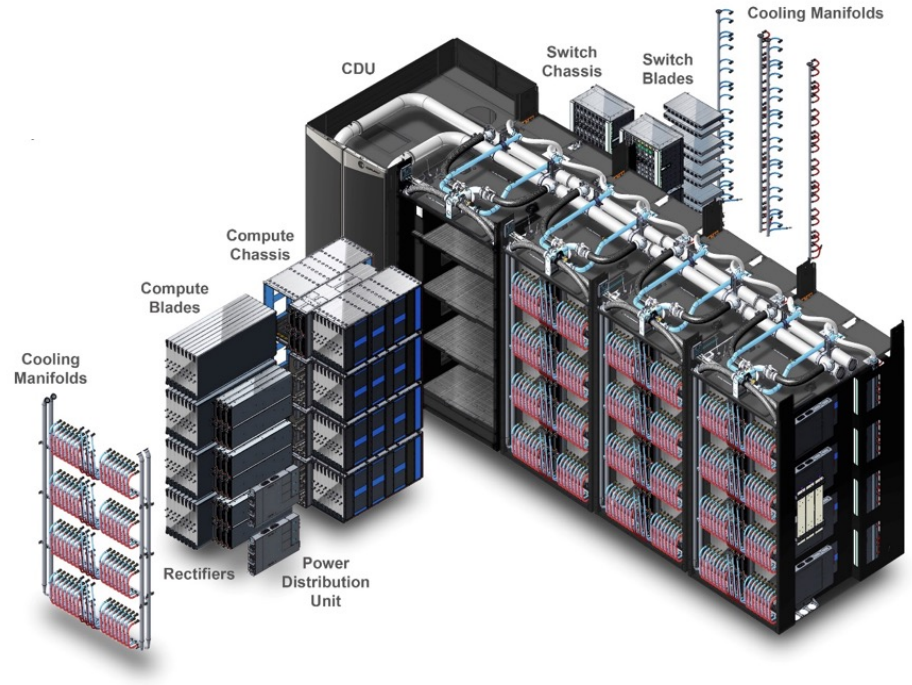

CRAY CLUSTERSTOR E1000 : VERY POWERFUL, HYBRID AND SCALABLE

- Powerful design
- Intelligent tiering software
- Hybrid, HDD and SSD
- Scalable
- Directly connected to Slingshot fabric

SLINGSHOT: INTERCONNECT FOR A DATA-CENTRIC WORLD

64 ports x 200 Gbps

Over 250K endpoints with a diameter of just three hops!

Ethernet Compatible

Easy connectivity to datacenters and third-party storage.

"HPC inside"

World class Adaptive Routing and QoS

High utilization at scale. Flawless support for hybrid workloads.

Highly Effective Congestion Control

Performance isolation between workloads.

Low, Uniform Latency

Focus on tail latency, because real apps synchronize.

CONGESTION MANAGEMENT PROVIDES PERFORMANCE ISOLATION

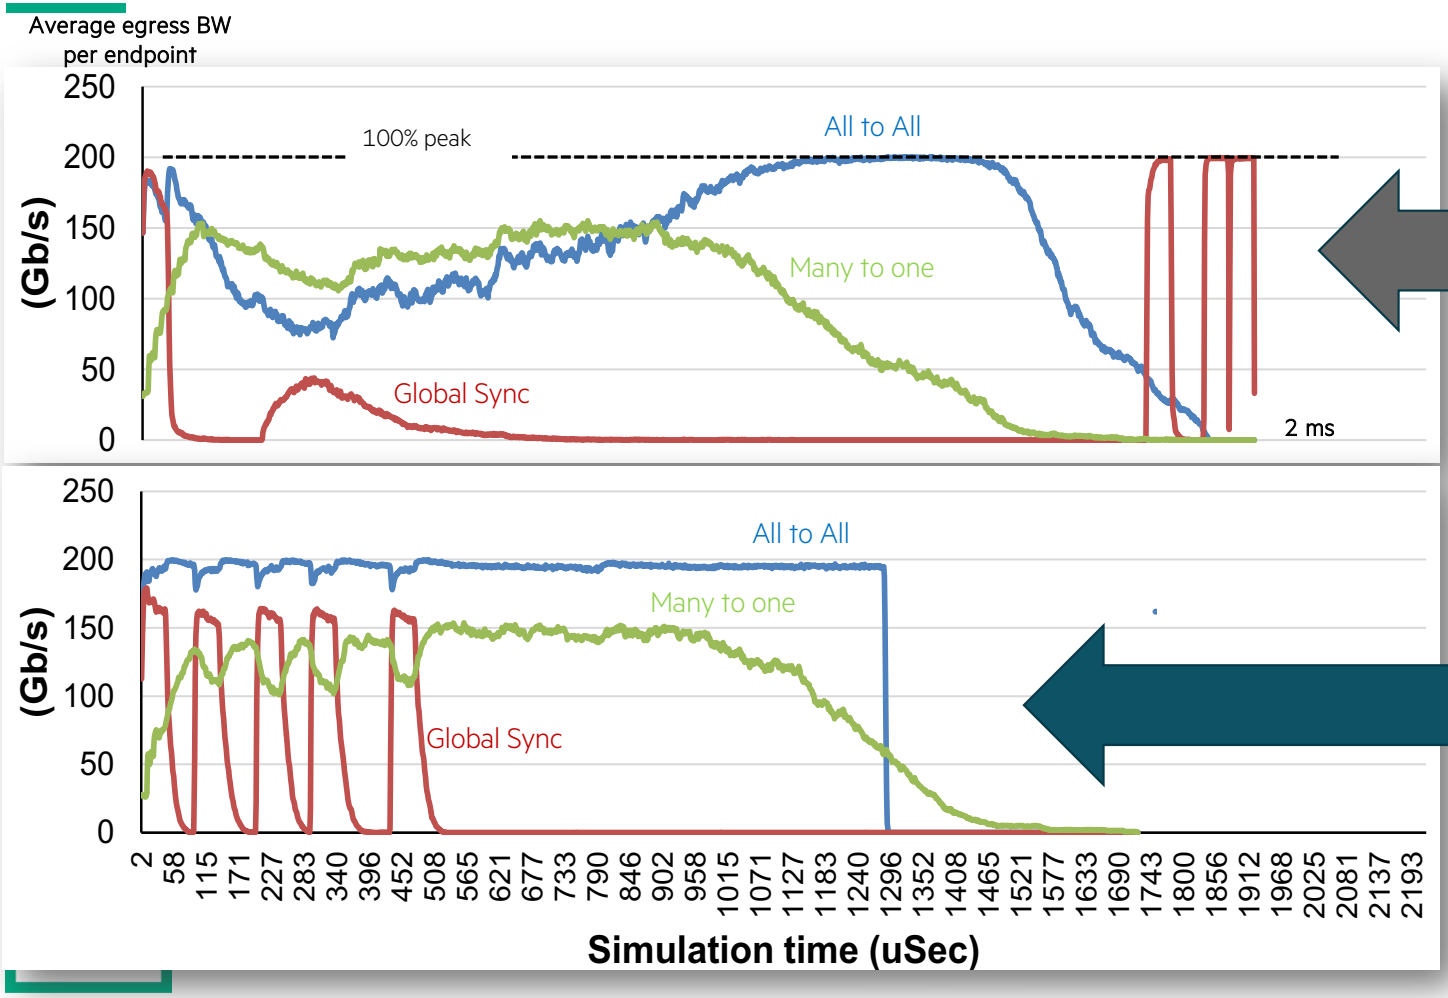

Job Interference in today's networks

Congesting (green) traffic hurts well-behaved (blue) traffic, and *really* hurts latency-sensitive, synchronized (red) traffic.

With Slingshot Advanced Congestion Management

GROMACS VARIABILITY STUDY

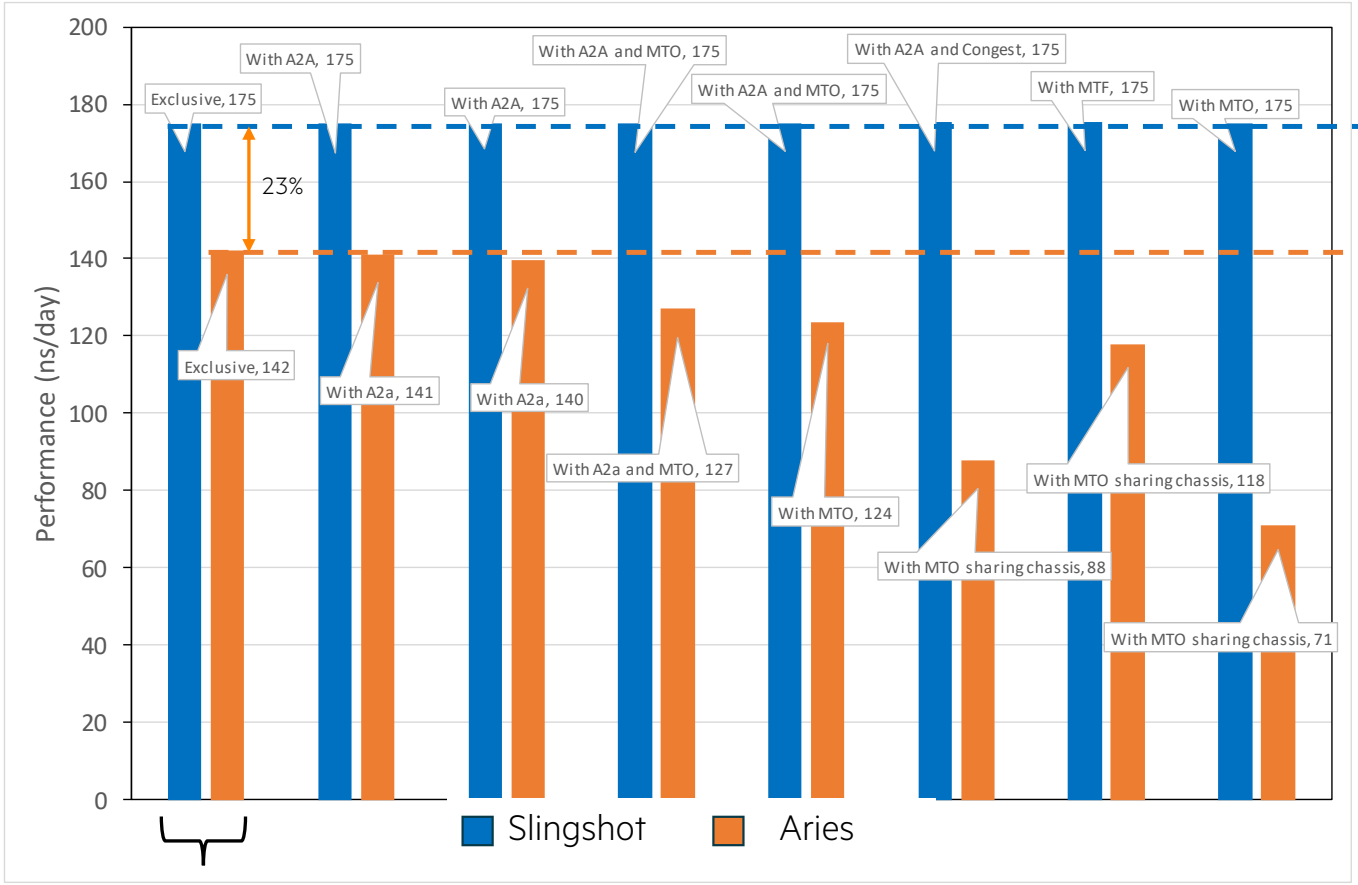

Note: Both systems running an identically configured 128x Skylake processor nodes

Highly Effective Congestion Control

Performance isolation between workloads.

Slingshot in Shasta delivers identical performance results across trials with every congestion type

CRAY SOFTWARE PLATFORM ARCHITECTURE

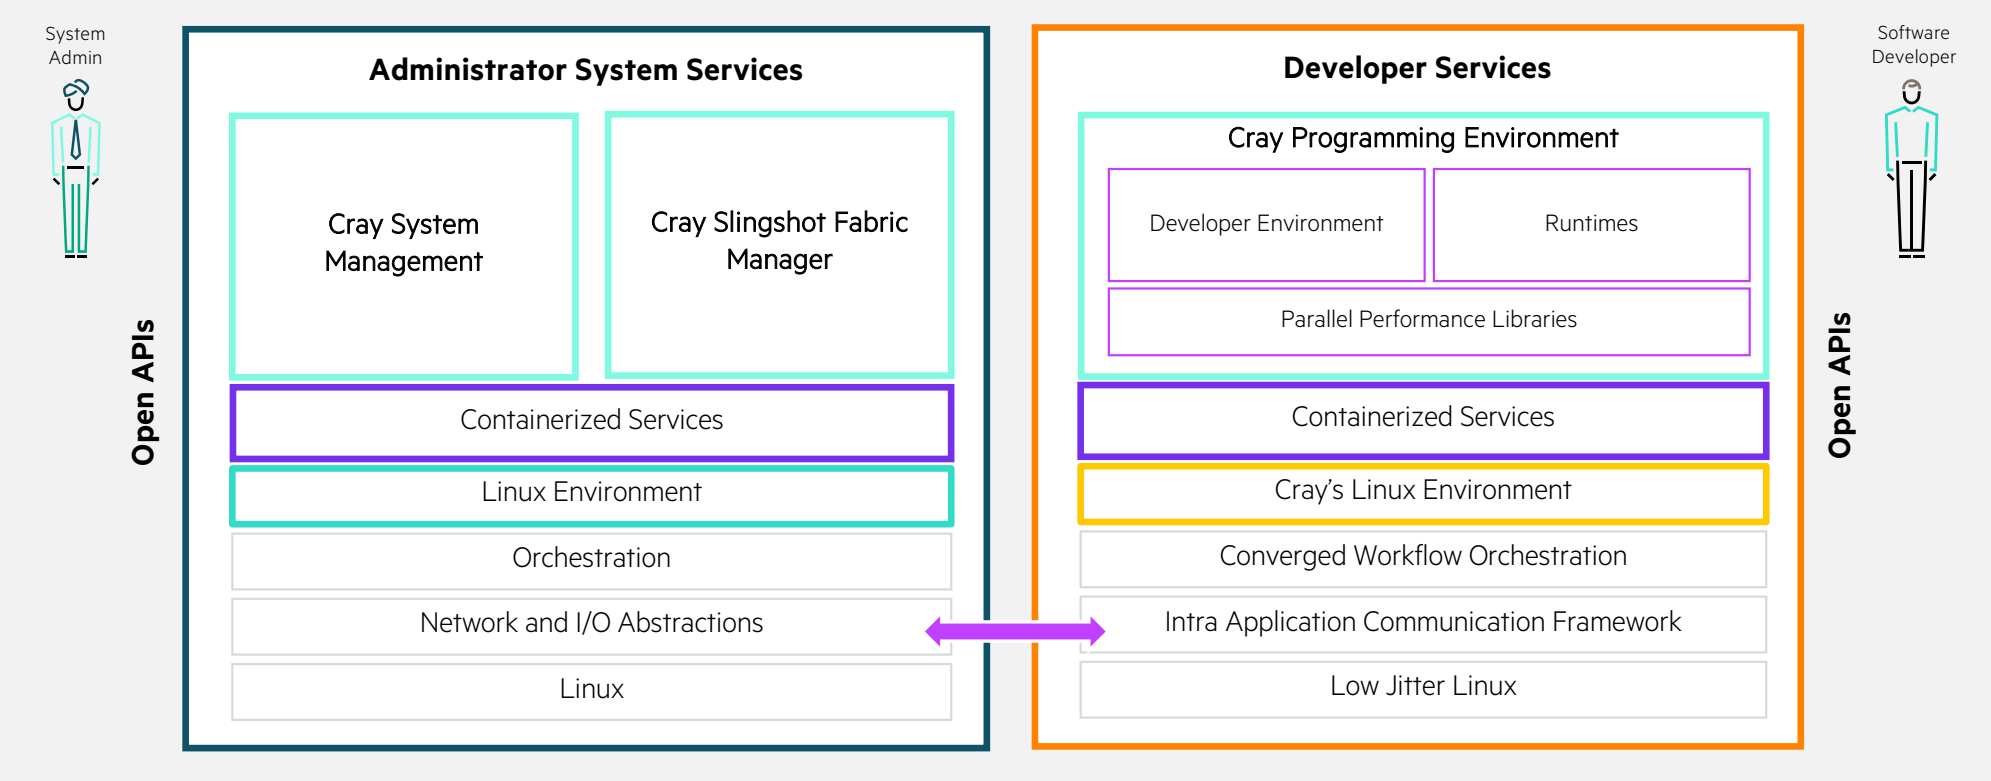

Expanding the power of supercomputing with the flexibility of cloud and full datacenter interoperability

PERFORMS LIKE A SUPERCOMPUTER RUNS LIKE A CLOUD

Flexible Workload Management

- **Batch environment**
  - Support for Slurm
- **Orchestrated workloads**
  - Leverage Kubernetes on compute nodes for orchestrated application (e.g. Urika for ML/DL)
  - Enables additional cloud-like provisioning of new workflows and enables cog-sim integration
- **Leverages built-in partitioning feature for Cray System Management**

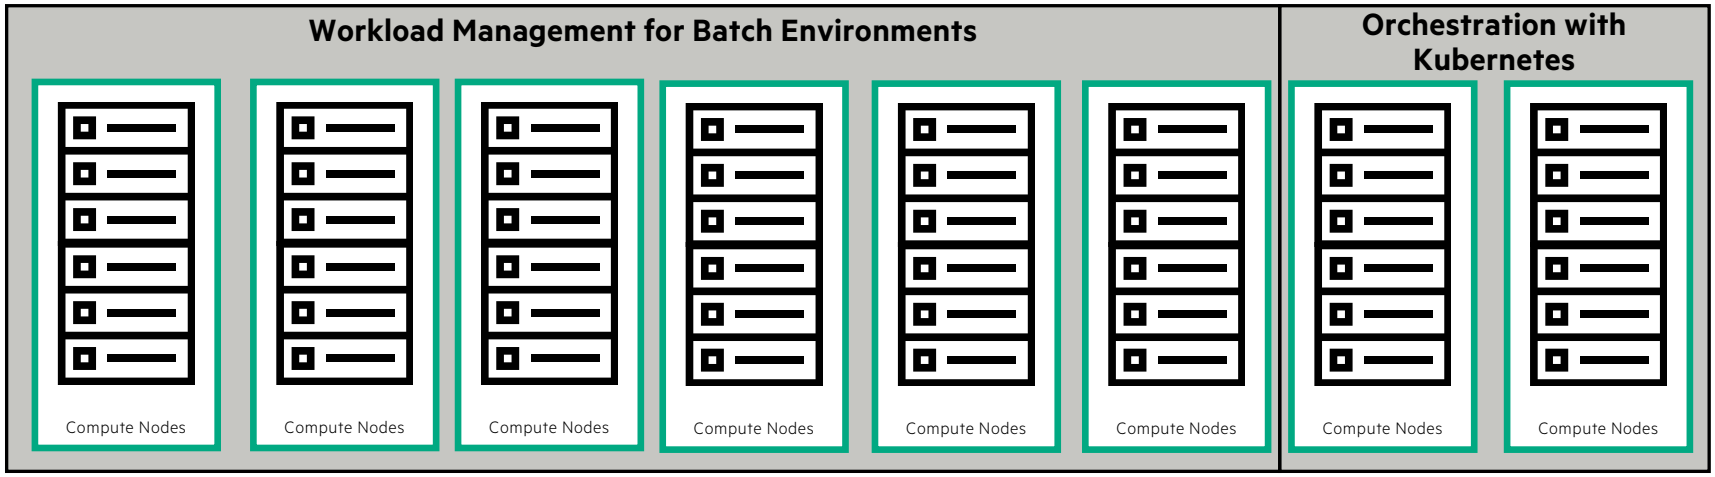

ONE EXPERIENCE ACROSS ARCHITECTURES

Capability	AMD Epyc (2019)	Intel Xeon (2020)	ARM CPU (2021)	NVIDIA GPU (2020)	AMD GPU (2021)	Intel GPU (2022)
Single management interface	●	●	●	●	●	●
Single monitoring framework and datastore	●	●	●	●	●	●
Cray OS	●	●	●	N/A	N/A	N/A
Portable development environment (containerized tools)	●	●	●	●	●	●
Slingshot Optimized Middleware: MPI, OpenMP, PGAS, SHMEM	●	●	●	●	●	●
Optimized Math Libraries	●	● ●	●	●	● ●	●
Code Optimization, Porting, and Debug	● ●	● ●	● ●	●	● ●	●
Fully supported compilation and execution environment (Fortran, C, C++, Python, and R)	● ●	● ●	● ●	●	● ●	●

Key

● Optimized or enhanced by Cray

● Industry or 3rd party – integrated by Cray

CRAY DEVELOPMENT ENVIRONMENT

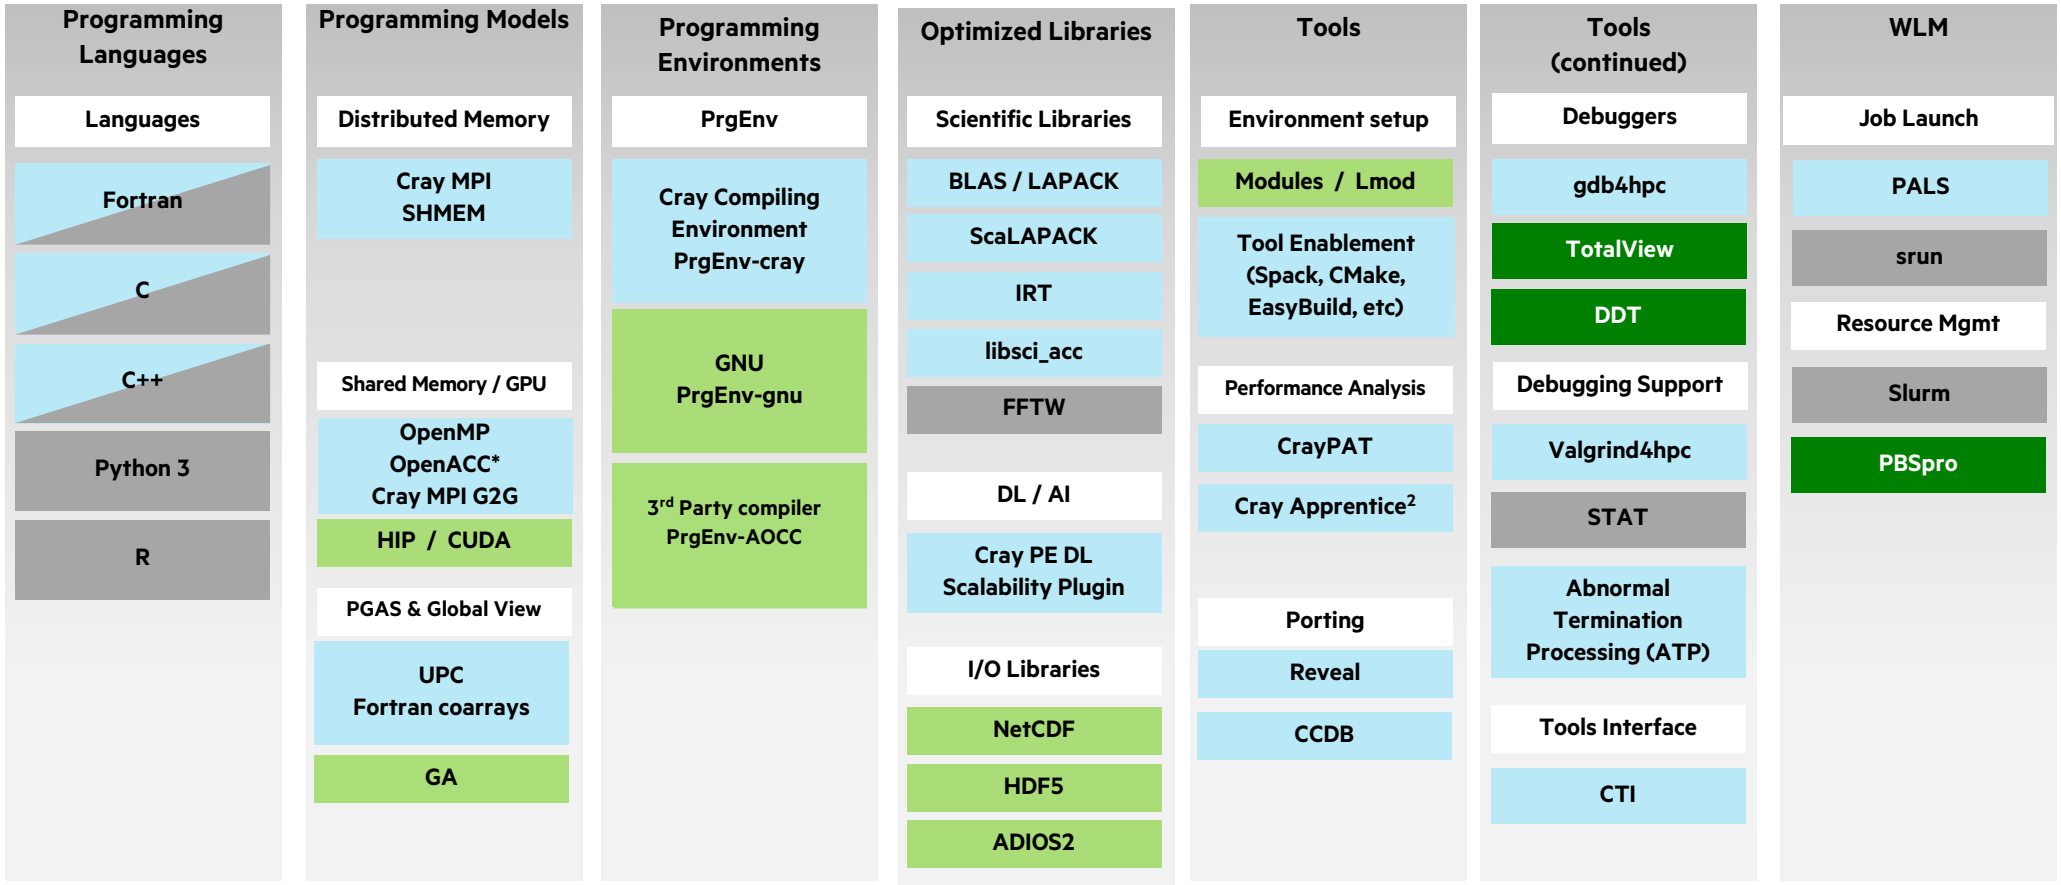

HIP: AMD OPEN PORTABLE GPU PROGRAMMING MODEL

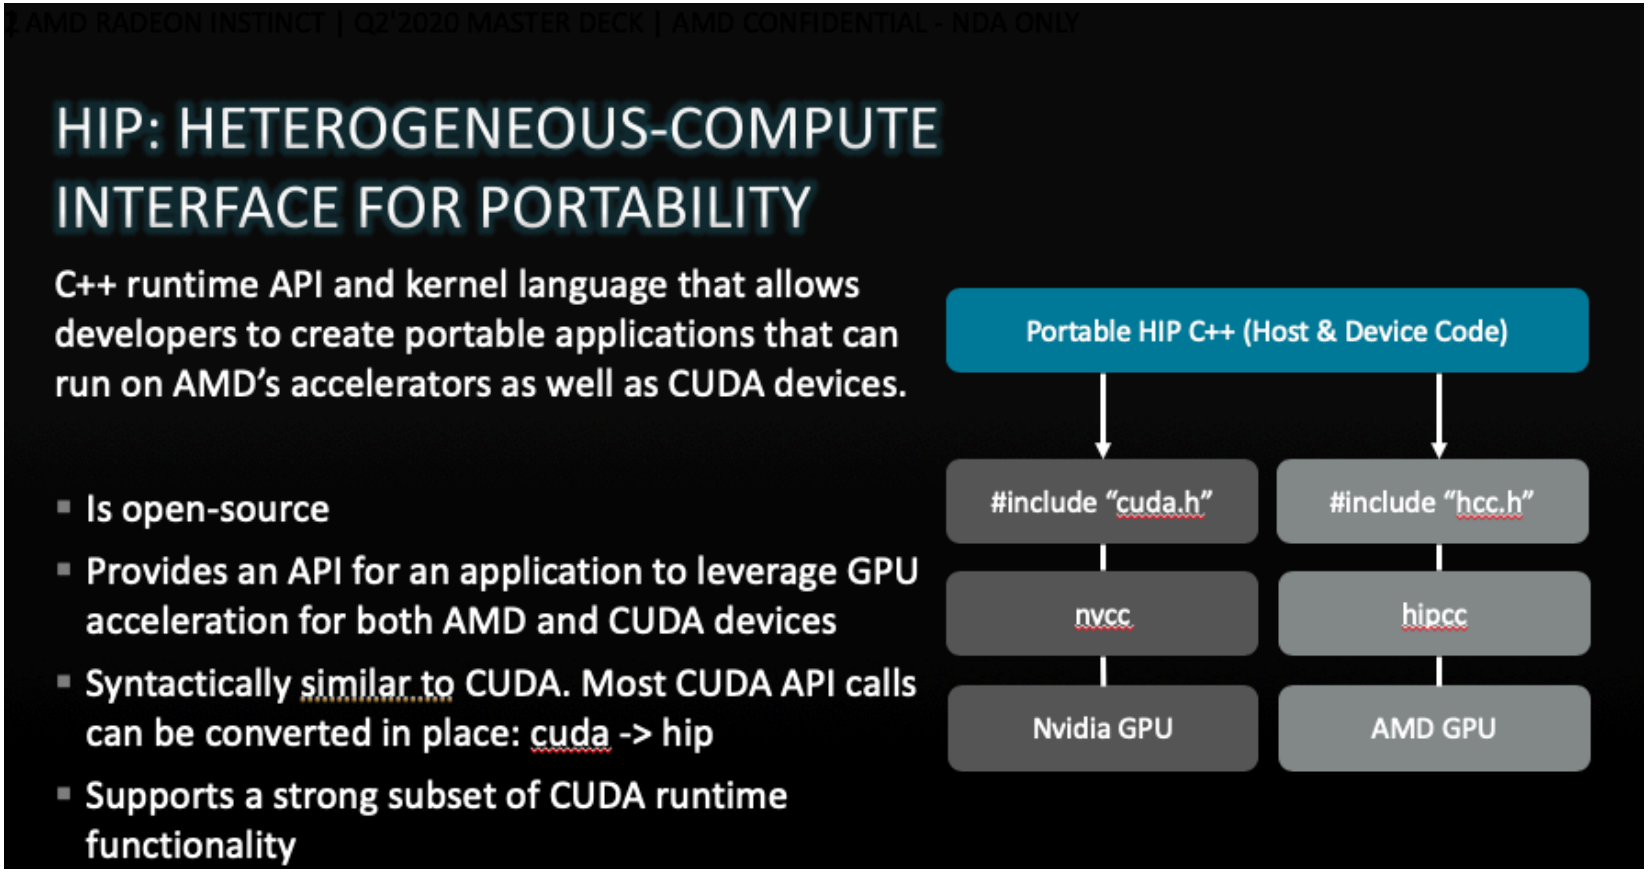

https://www.olcf.ornl.gov/wp-content/uploads/2019/09/AMD_GPU_HIP_training_20190906.pdf

KERNEL CODE IS IDENTICAL BETWEEN CUDA AND HIP

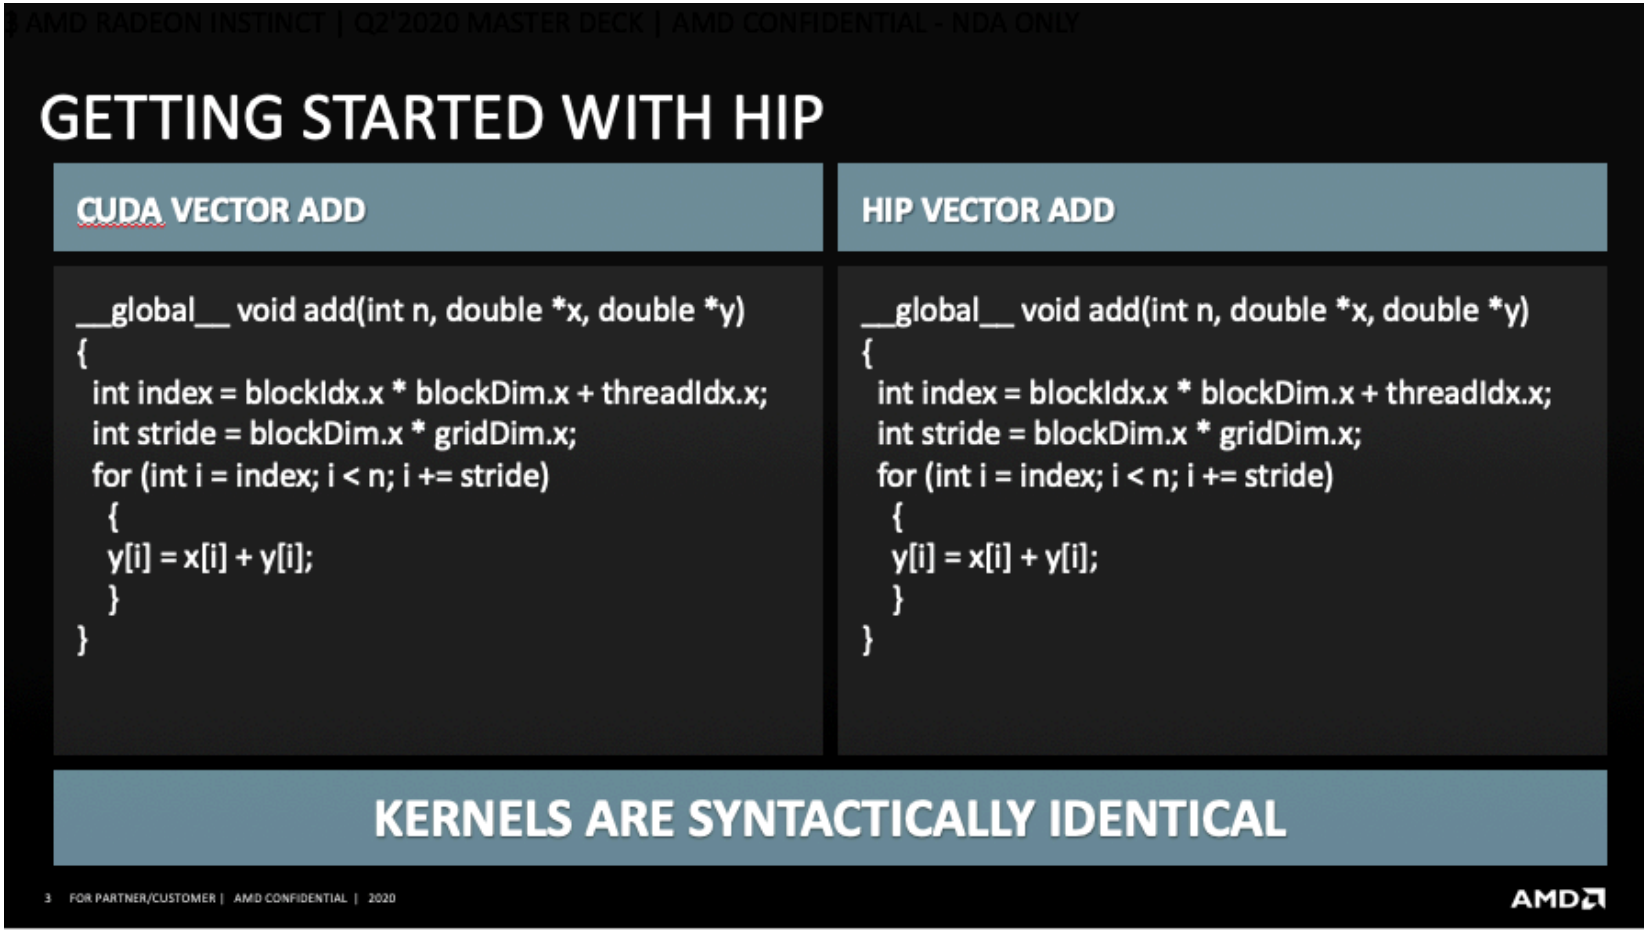

https://www.olcf.ornl.gov/wp-content/uploads/2019/09/AMD_GPU_HIP_training_20190906.pdf

HIP PORTABLE SCIENTIFIC LIBRARIES

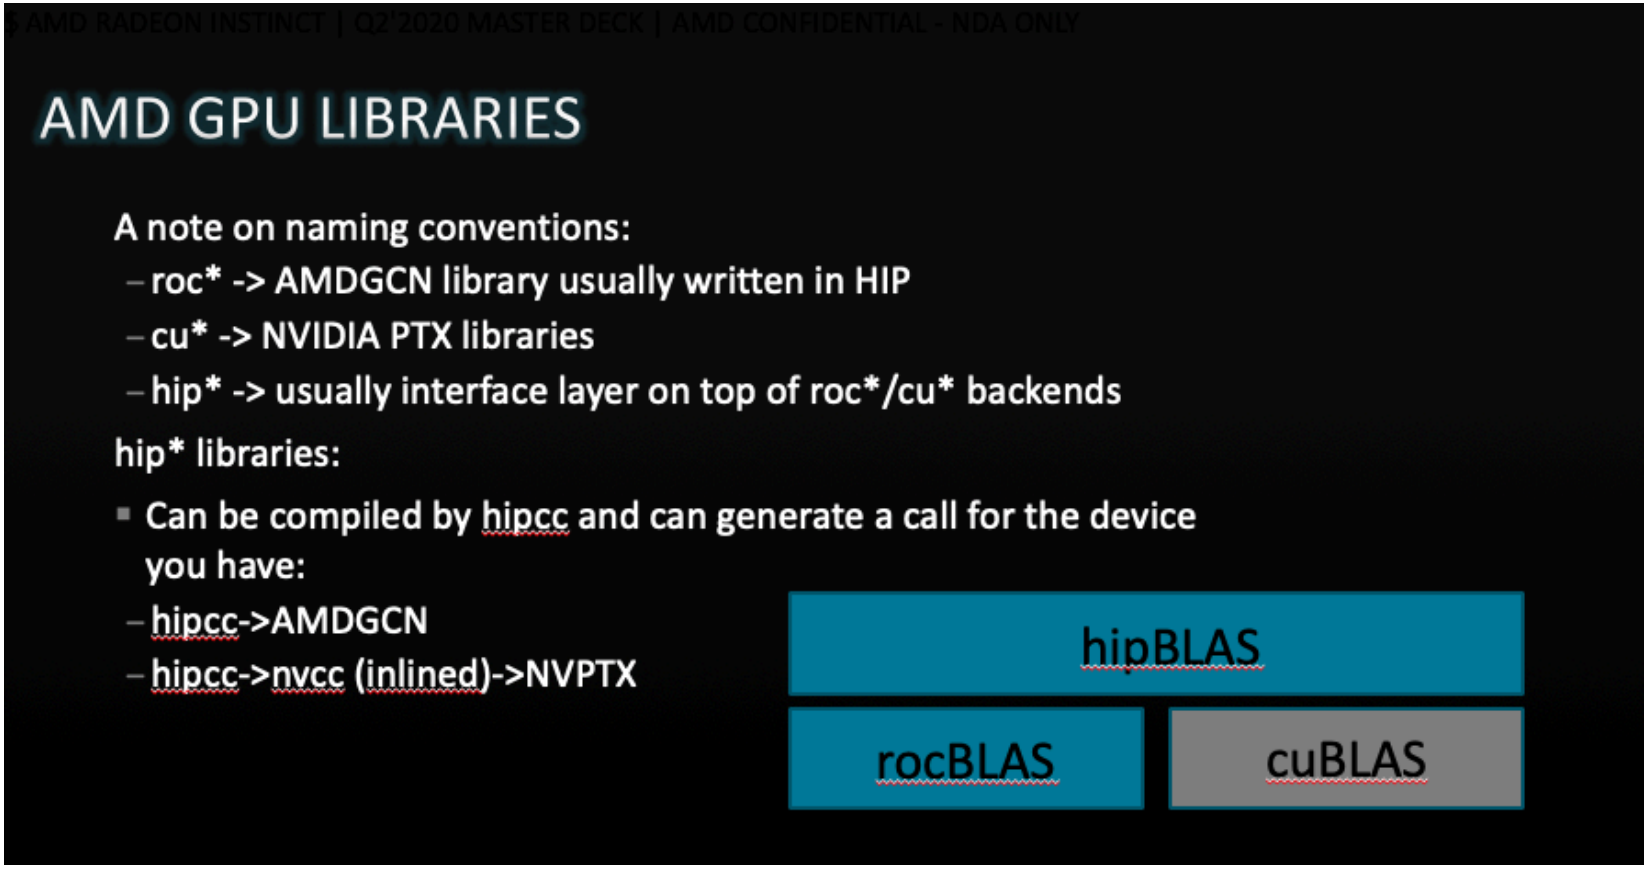

https://www.olcf.ornl.gov/wp-content/uploads/2019/09/AMD_GPU_HIP_training_20190906.pdf

PORTING THE QUDA LIBRARY TO HIP

- MILC collaboration code for lattice QCD calculations
- The MILC Code is a body of high performance research software written in C for doing SU(3) lattice gauge theory on high performance computers as well as single-processor workstations. A wide variety of applications are included.
- MILC implementation for GPU is largely based on the QUDA library
- QUDA depends on many additional CUDA libraries: Eigen, CuFFT, CuBLAS, CuRAND, Thrust, CUB
- AMD asked a local consulting company to port QUDA to HIP:
  - 10K lines of hand-tuned CUDA kernels -> hipify converted without problems
  - 35K lines of header code -> hipify mostly converted but needed manual switch to new library dependencies
  - 74K lines of library code -> mostly successful hipify conversion, reuired some manual changes
  - 34K lines of test suite code -> hipify converted without problems
  - Template use for tests -> no solution, manual porting
- 15 developer days
- Much much bigger effort without HIP tools

HIP MEASURED PERFORMANCE ON NVIDIA GPU

Performance (II)

- Average of normalized HIP performance was 99.8% with data transfer costs, 99.9% w/out

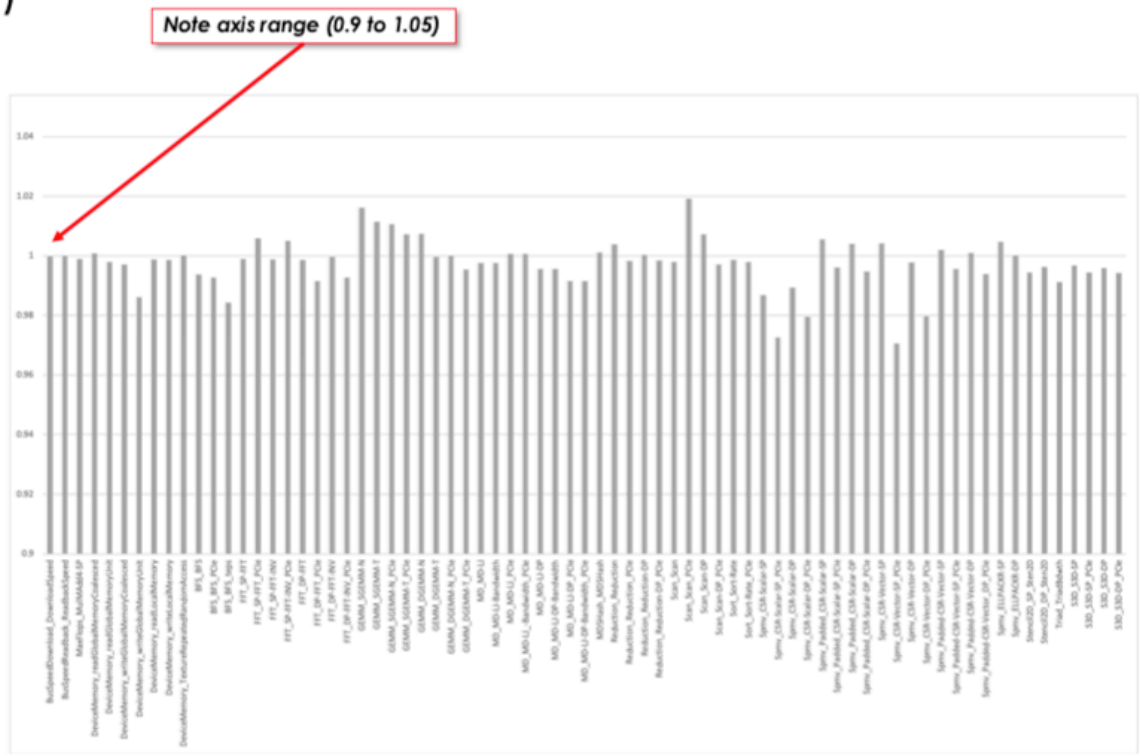

OAK RIDGE National Laboratory

https://www.olcf.ornl.gov/wp-content/uploads/2019/10/Roth-HIP-on-Summit-20191009.pdf

THANK YOU

gabriele.paciucci@hpe.com