POWERING THE NEXT ERA OF SUPERCOMPUTING

Gabriele Paciucci

CRAY®
a Hewlett Packard Enterprise company

SAFE HARBOR STATEMENT

This presentation may contain forward-looking statements that are based on our current expectations. Forward looking statements may include statements about our financial guidance and expected operating results, our opportunities and future potential, our product development and new product introduction plans, our ability to expand and penetrate our addressable markets and other statements that are not historical facts.

These statements are only predictions and actual results may materially vary from those projected. Please refer to Cray's documents filed with the SEC from time to time concerning factors that could affect the Company and these forward-looking statements.

Together, Shaping and Leading the Next Generation of High Performance Computing (HPC)

CRAY
a Hewlett Packard Enterprise company

Hewlett Packard Enterprise

Global leader focused on developing intelligent solutions to capture, analyze and act upon data seamlessly from edge to cloud

CRAY®
a Hewlett Packard Enterprise company

Premier provider of high-end supercomputing solutions, addressing customers' most challenging data-intensive workloads for critical decisions

Executing Exascale

ANL “Aurora”

- ~$100M system subcontract
- More than 1 EF Sustained performance
- Future Intel Xeon CPU and Intel X^e architecture and Slingshot interconnect
- Mixed AI and HPC workload

ORNL “Frontier”

- ~$600M system contract
- More than 1.5 EF Sustained performance
- Future AMD EPYC CPU and Radeon GPU and Slingshot interconnect
- ClusterStor Lustre FS >10TB/s with 1 EB+
- Mixed AI and HPC workload

LLNL “El Capitan”

- ~$600M system contract
- More than 1.5 EF Sustained performance
- Processor and Accelerator Vendor TBD; Slingshot interconnect
- Mixed AI and HPC workload

Shasta: Value at Any Scale

NERSC-9 “Perlmutter”

- $146M system contract
- 100 PF pre-exascale system
- AMD Rome CPU and NVIDIA Tesla V100 GPU
- All-flash ClusterStor Lustre FS >4TB/s with 30 PB+
- Mixed AI and HPC workload

Indiana U “Big Red 200”

- ~$10M system contract
- 5.9 PF system
- AMD Rome CPU and NVIDIA Tesla V100 GPU
- AI system

U.S. Air Force “HPC11”

- $25M system contract
- Specs not disclosed
- AMD Rome CPU
- Weather modeling

Major Trends Impacting HPC going Exascale

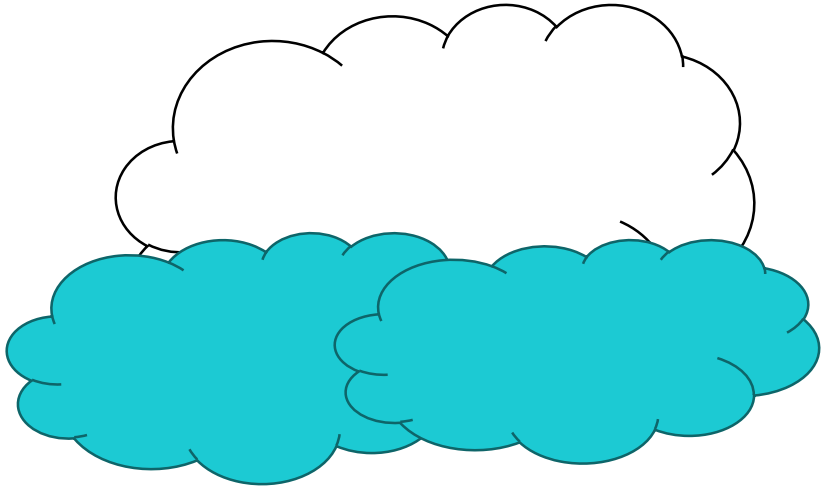

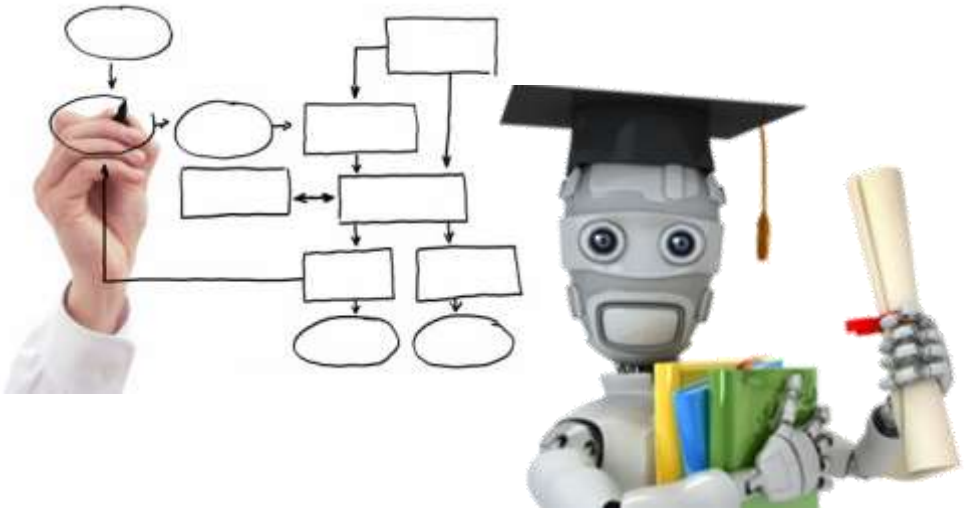

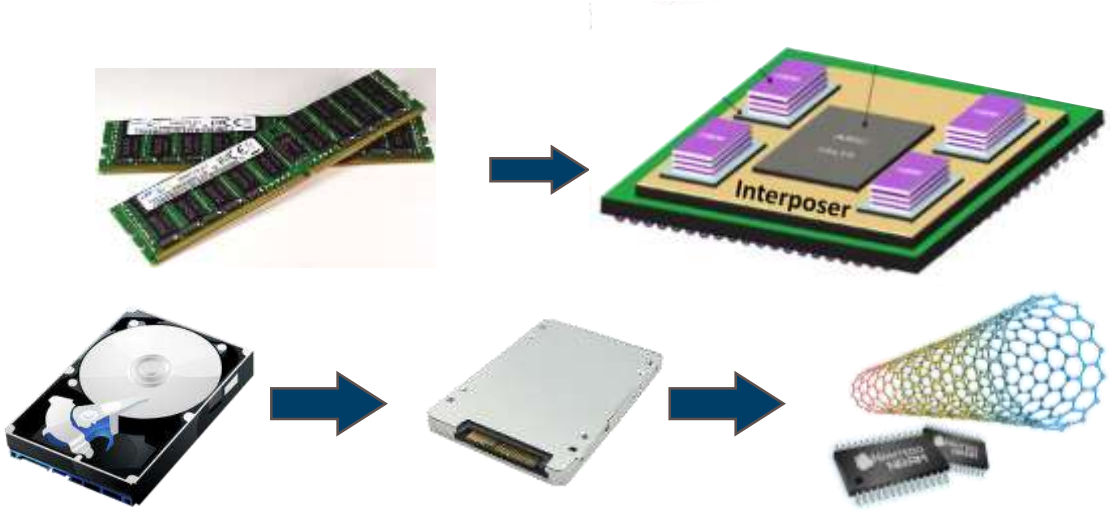

Recipe for Exascale Computing

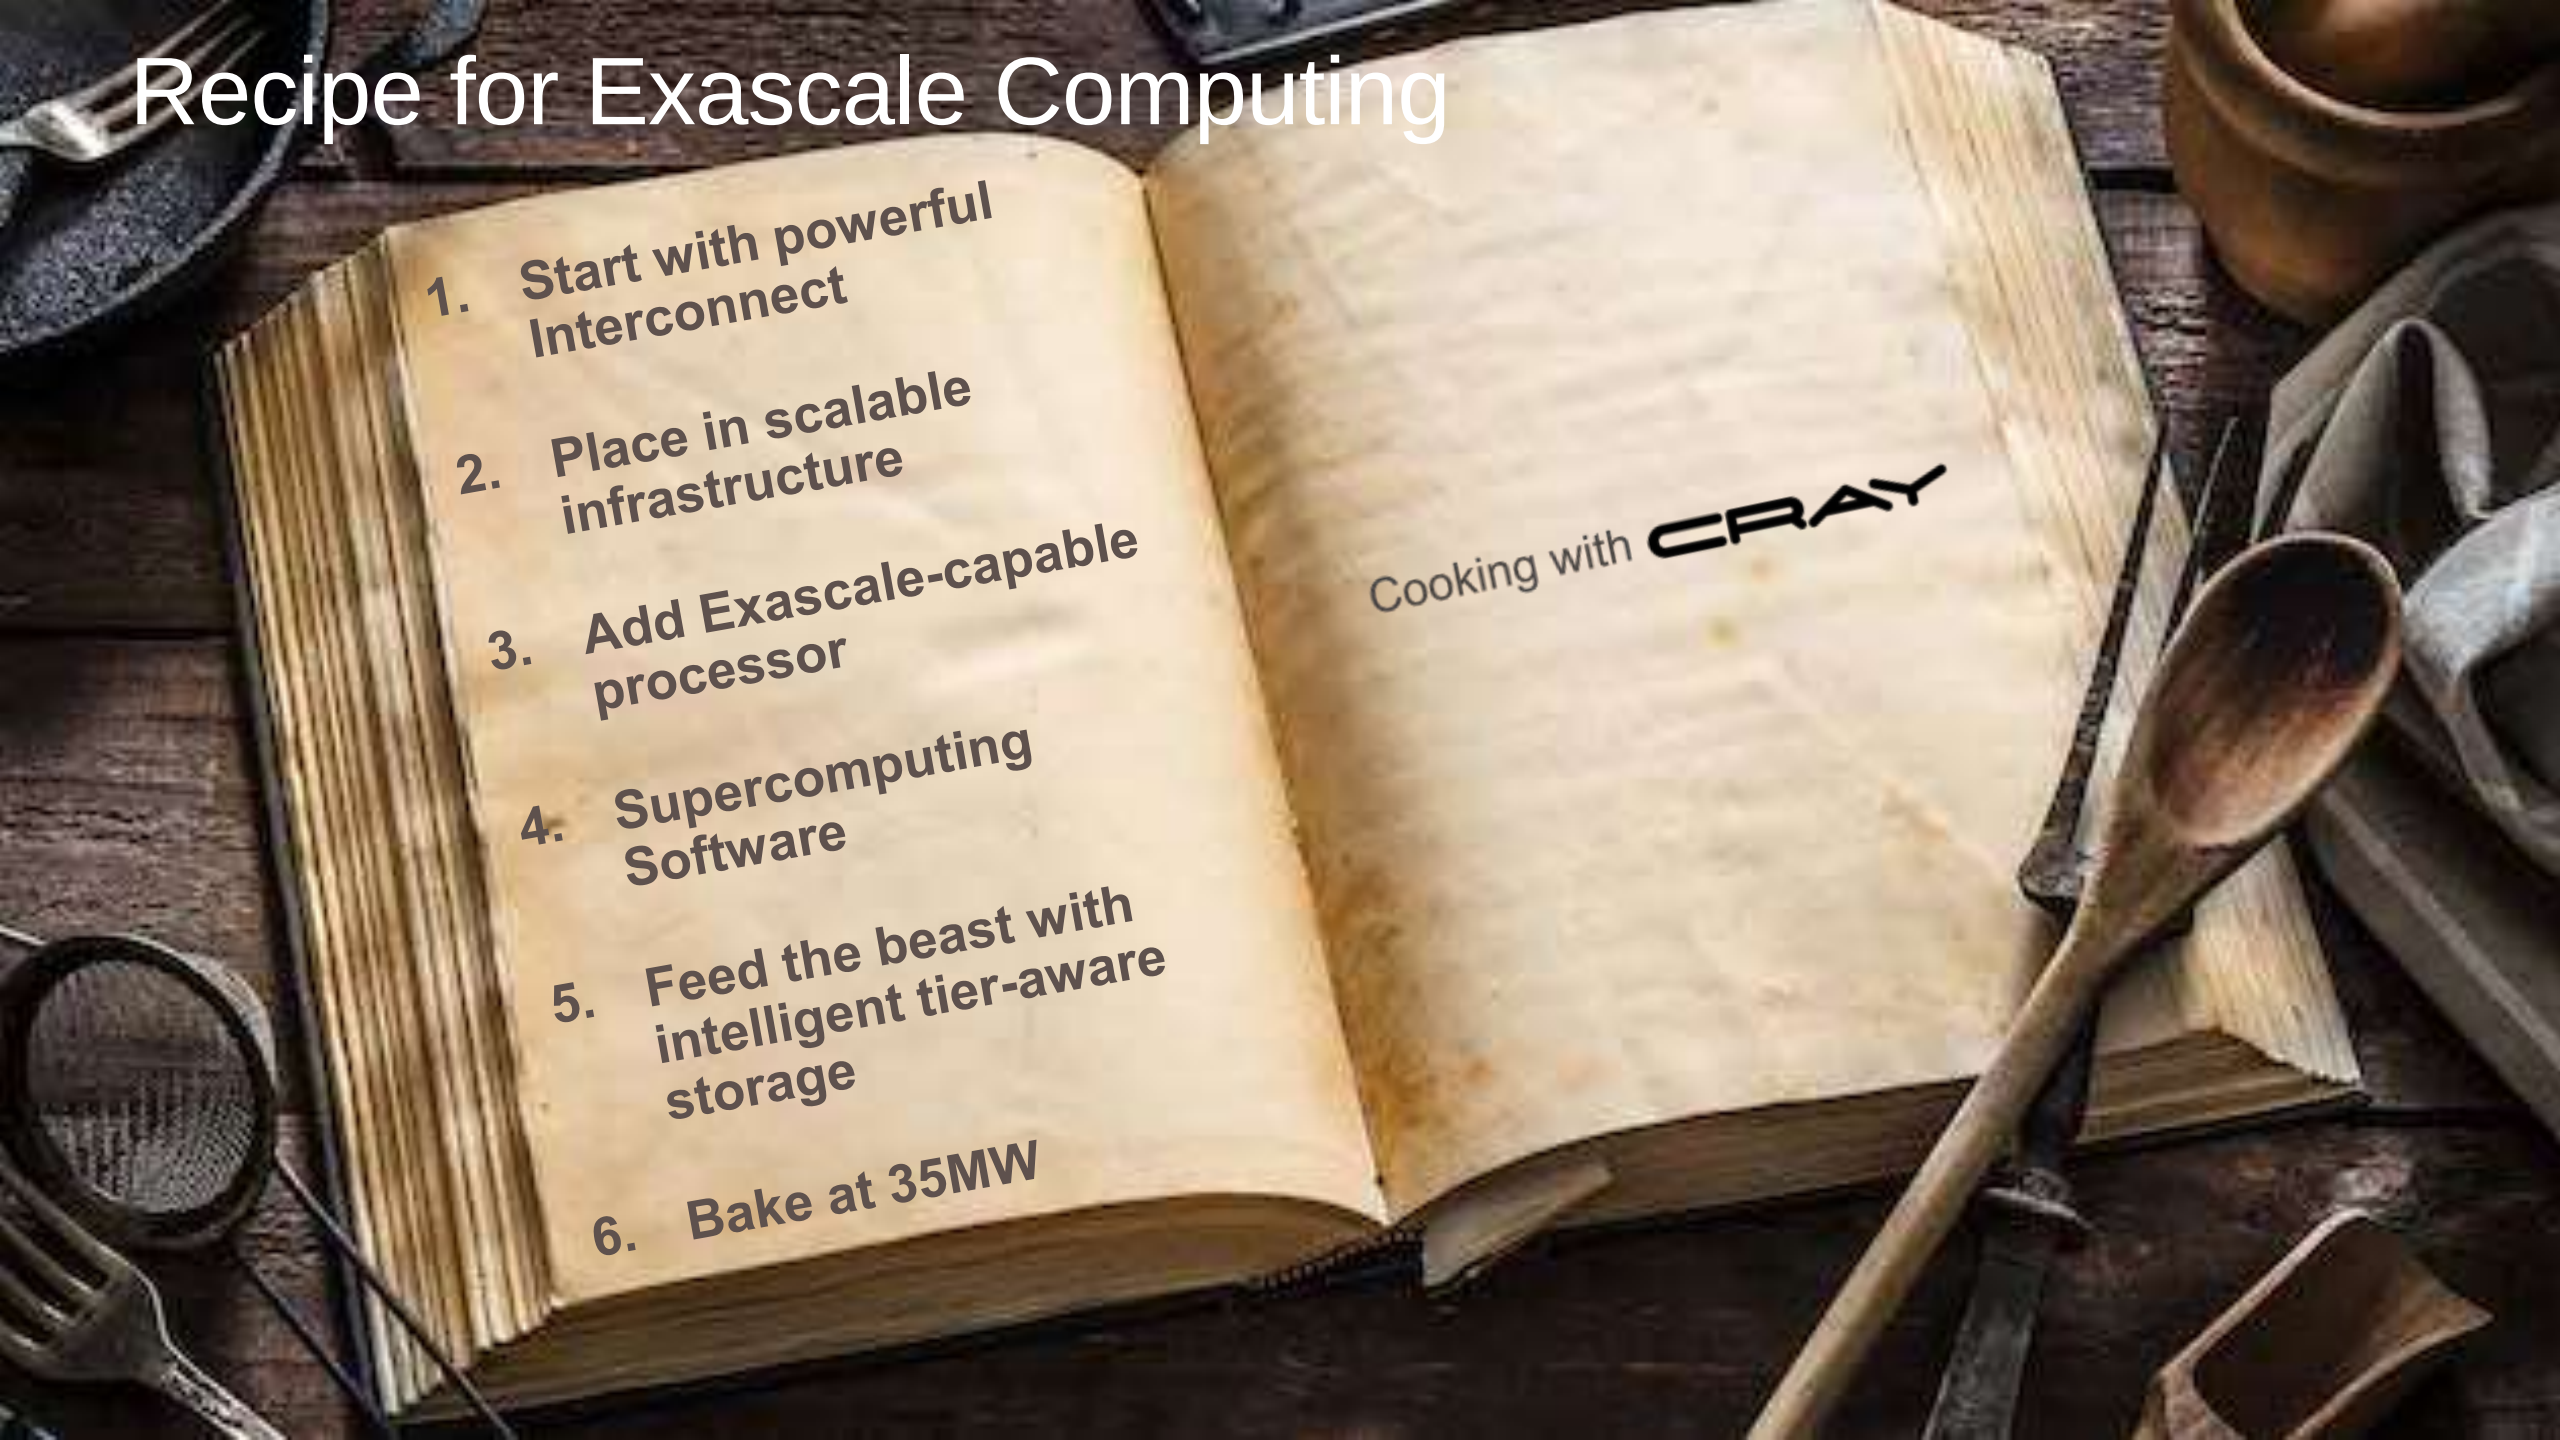

Slingshot: Interconnect for a Data-Centric World

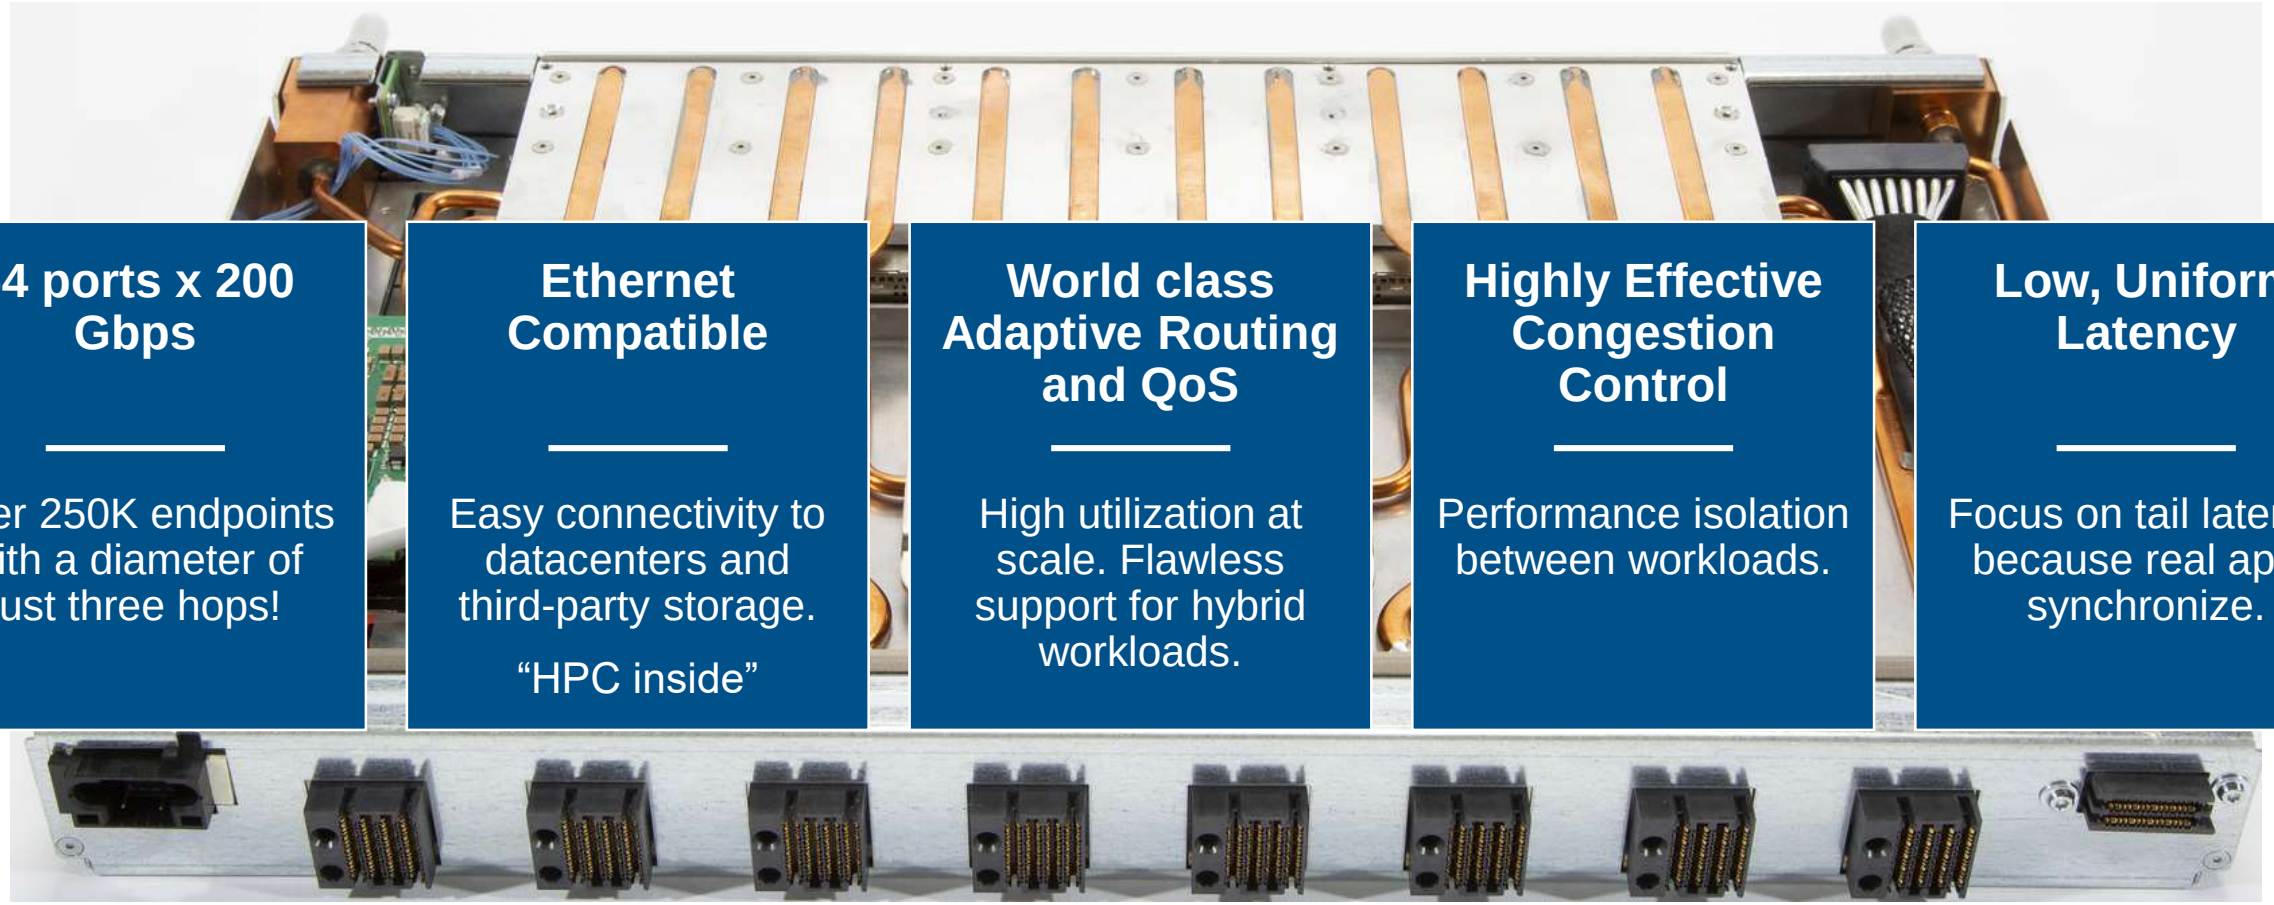

64 ports x 200 Gbps

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.

Slingshot: sustained performance

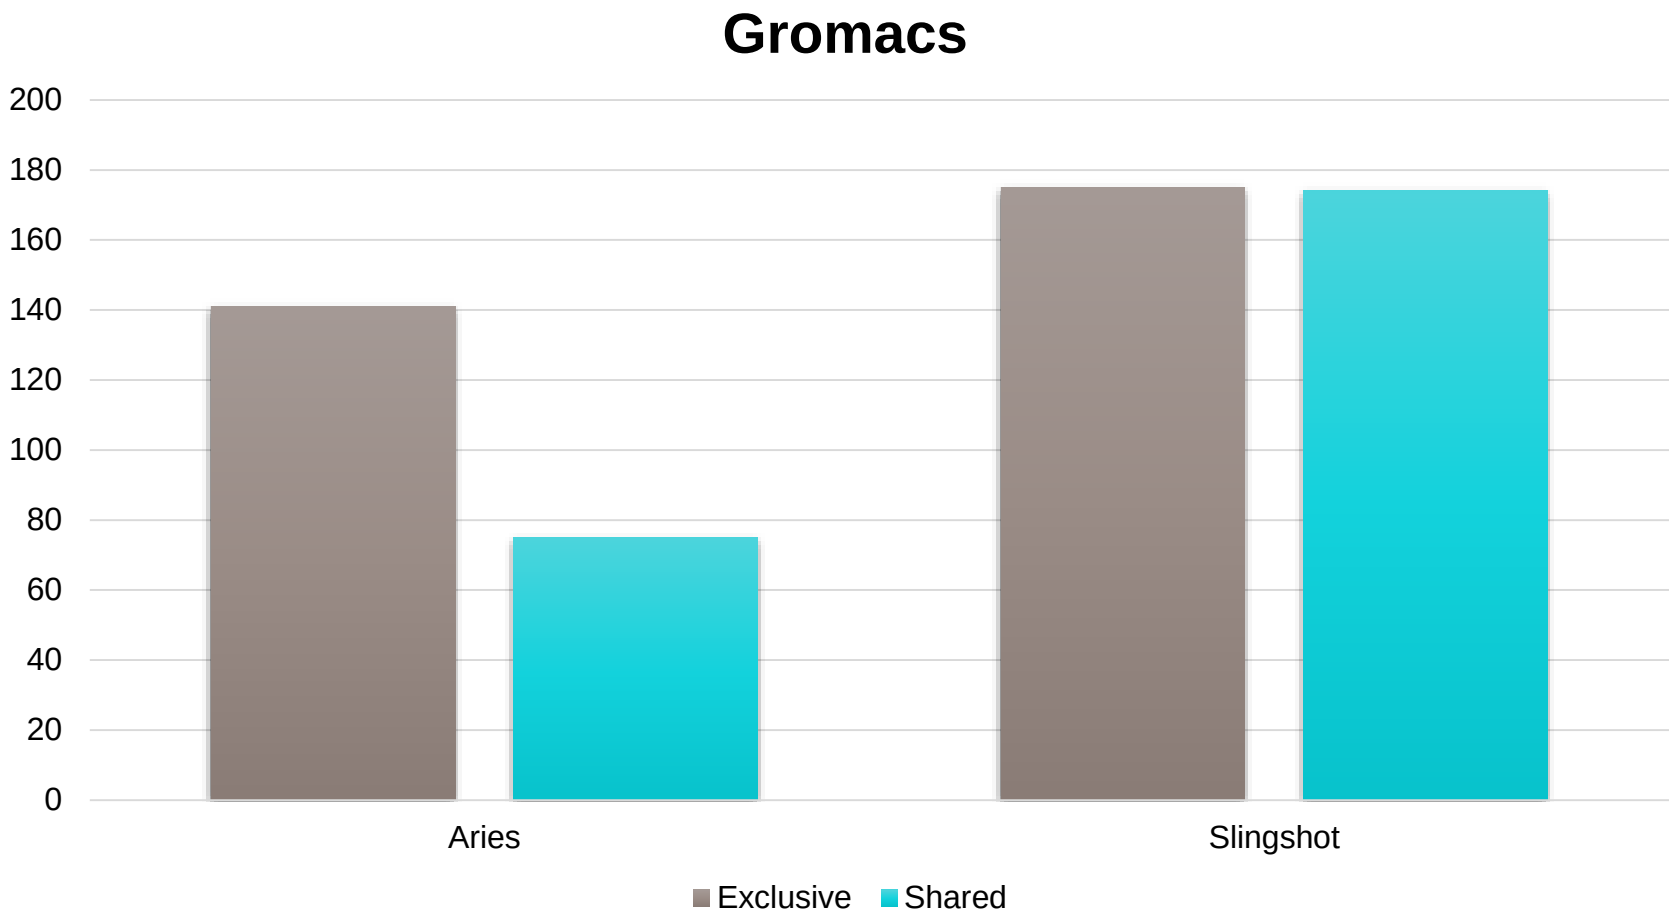

Gromacs: Cellulose test case on 128 Intel Skylake nodes

Shasta: Flexible and Scalable Infrastructure

Ultra-dense, scale-optimized

- 250-300KW with warm water cooling
- 512+ high-performance processors
- Flexible, high-density interconnect

Standard 19” Rack

- Air cooled with liquid cooling options
- Wide range of available compute and storage
- Deploy in wider range of datacenters

Same Interconnect - Same Software Environment

AMD Rome

(NERSC)

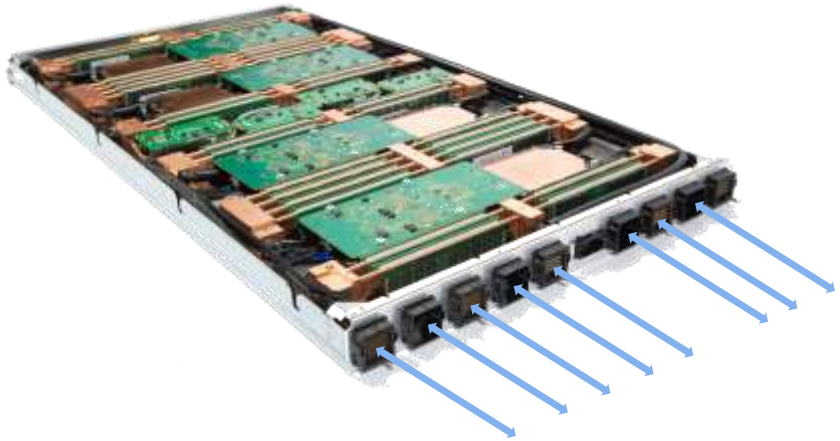

To Slingshot

Cray Programming Environment

Intel Xe

(ANL)

To Slingshot

Cray Programming Environment

Intel One API

Nvidia GPU

(NERSC)

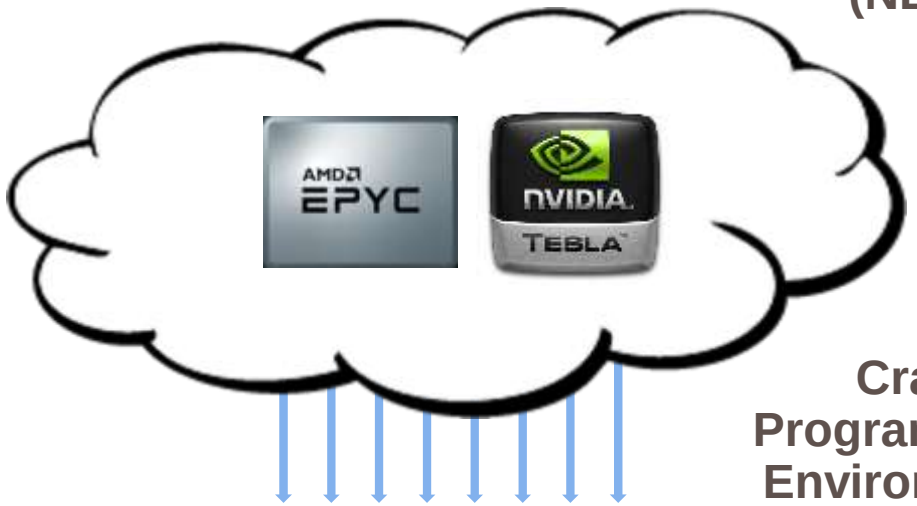

To Slingshot

Cray Programming Environment

PGI/CUDA

AMD GPU

(ORNL)

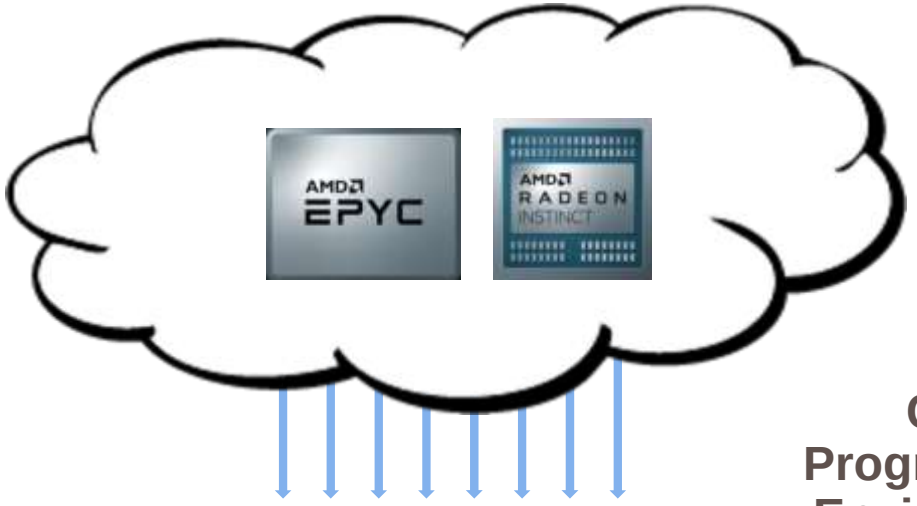

To Slingshot

Cray Programming Environment

Shasta: Stack for Converged Workloads

Traditional software for HPC (performance , scale, reliability)		Cray Shasta Software for Converged Workloads
Batch workflow scheduling for modeling and simulation	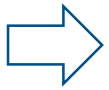	Scheduling for hybrid HPC/analytics/AI workflows with Kubernetes container orchestration
System partitioning for HPC jobs	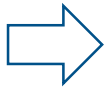	Multi-tenancy between HPC and AI partitions, and sub-partitioning within AI jobs
Tightly integrated supercomputer, requiring coordination and downtime for upgrades	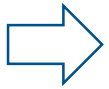	Containerized, self-healing management services; Separate compute and management systems for seamless upgrades
Tightly controlled system interfaces and limited access to system data	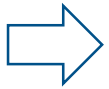	Open, RESTful APIs for integration, data access, and datacenter interoperability
Varied and discrete logs to track system and user activities and performance	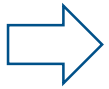	Fully integrated system and user level telemetry to quickly correlate and remediate issues

ClusterStor value – End to End Integration

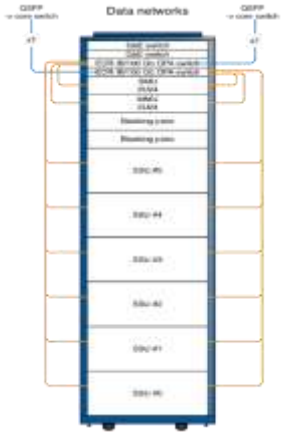

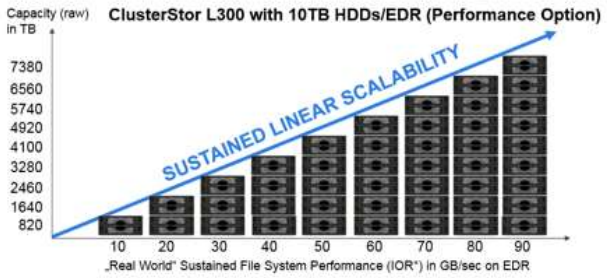

Performance Efficiency

1

Non-disruptive I/O Acceleration

2

Engineered Solution

3

Management and Support

4

Scalability

5

Reliability

6

CLUSTERSTOR™

THANK YOU

QUESTIONS?

gpaciucci@cray.com

CRAY®
a Hewlett Packard Enterprise company

cray.com
@cray_inc
linkedin.com/company/cray-inc-/