SAND2019-3446C

VANGUARD

Vanguard Astra - Petascale ARM Platform for U.S. DOE/ASC Supercomputing

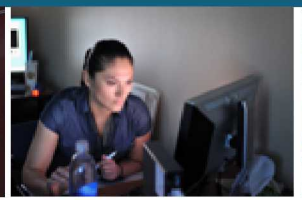

PRESENTED BY

Andrew J. Younge

Sandia National Laboratories

SAND2019-XXXX C
Unclassified Unlimited Release

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Vanguard Overview

Vanguard Program: Advanced Architecture Prototype Systems

- Prove viability of advanced technologies for NNSA integrated codes, at scale
- Expand the HPC-ecosystem by developing emerging unproven technologies
  - Is it viable for future ATS/CTS platforms?
  - Increase technology AND integrator choices
- Buy down risk and increase technology and vendor choices for future platforms
  - Ability to accept higher risk allows for more/faster technology advancement
  - Lowers/eliminates mission risk and significantly reduces investment
- Jointly address hardware and software challenges

- First prototype platform targeting ARM

Where Vanguard Fits

Test Beds

- Small testbeds (~10-100 nodes)
- Breadth of architectures
- **Brave users**

Vanguard

- Larger-scale experimental systems
- Focused efforts to mature new technologies
- Broader user-base
- **Demonstrate viability for production use**
- NNSA Tri-lab resource

ATS/CTS Platforms

- Leadership-class systems (Petascale, Exascale, ...)
- Advanced technologies, sometimes first-of-kind
- Broad user-base
- **Production use**

Sandia has a history with Arm - NNSA/ASC testbeds

2014

Hewlett Packard Enterprise

ARM

Hammer

Applied Micro
X-Gene-1
47 nodes

ARM

Sullivan

Cavium ThunderX1
32 nodes

2017

Hewlett Packard Enterprise

ARM

Mayer

Pre-GA Cavium
ThunderX2
47 nodes

2018

Hewlett Packard Enterprise

ARM

Vanguard/Astra

HPE Apollo 70
Cavium ThunderX2
2592 nodes


ASTRA
"Per aspera ad astra"
RM SUPERCOMPUTER
VANGUARD
ENERGY
NNSA

per aspera ad astra

through difficulties to the stars

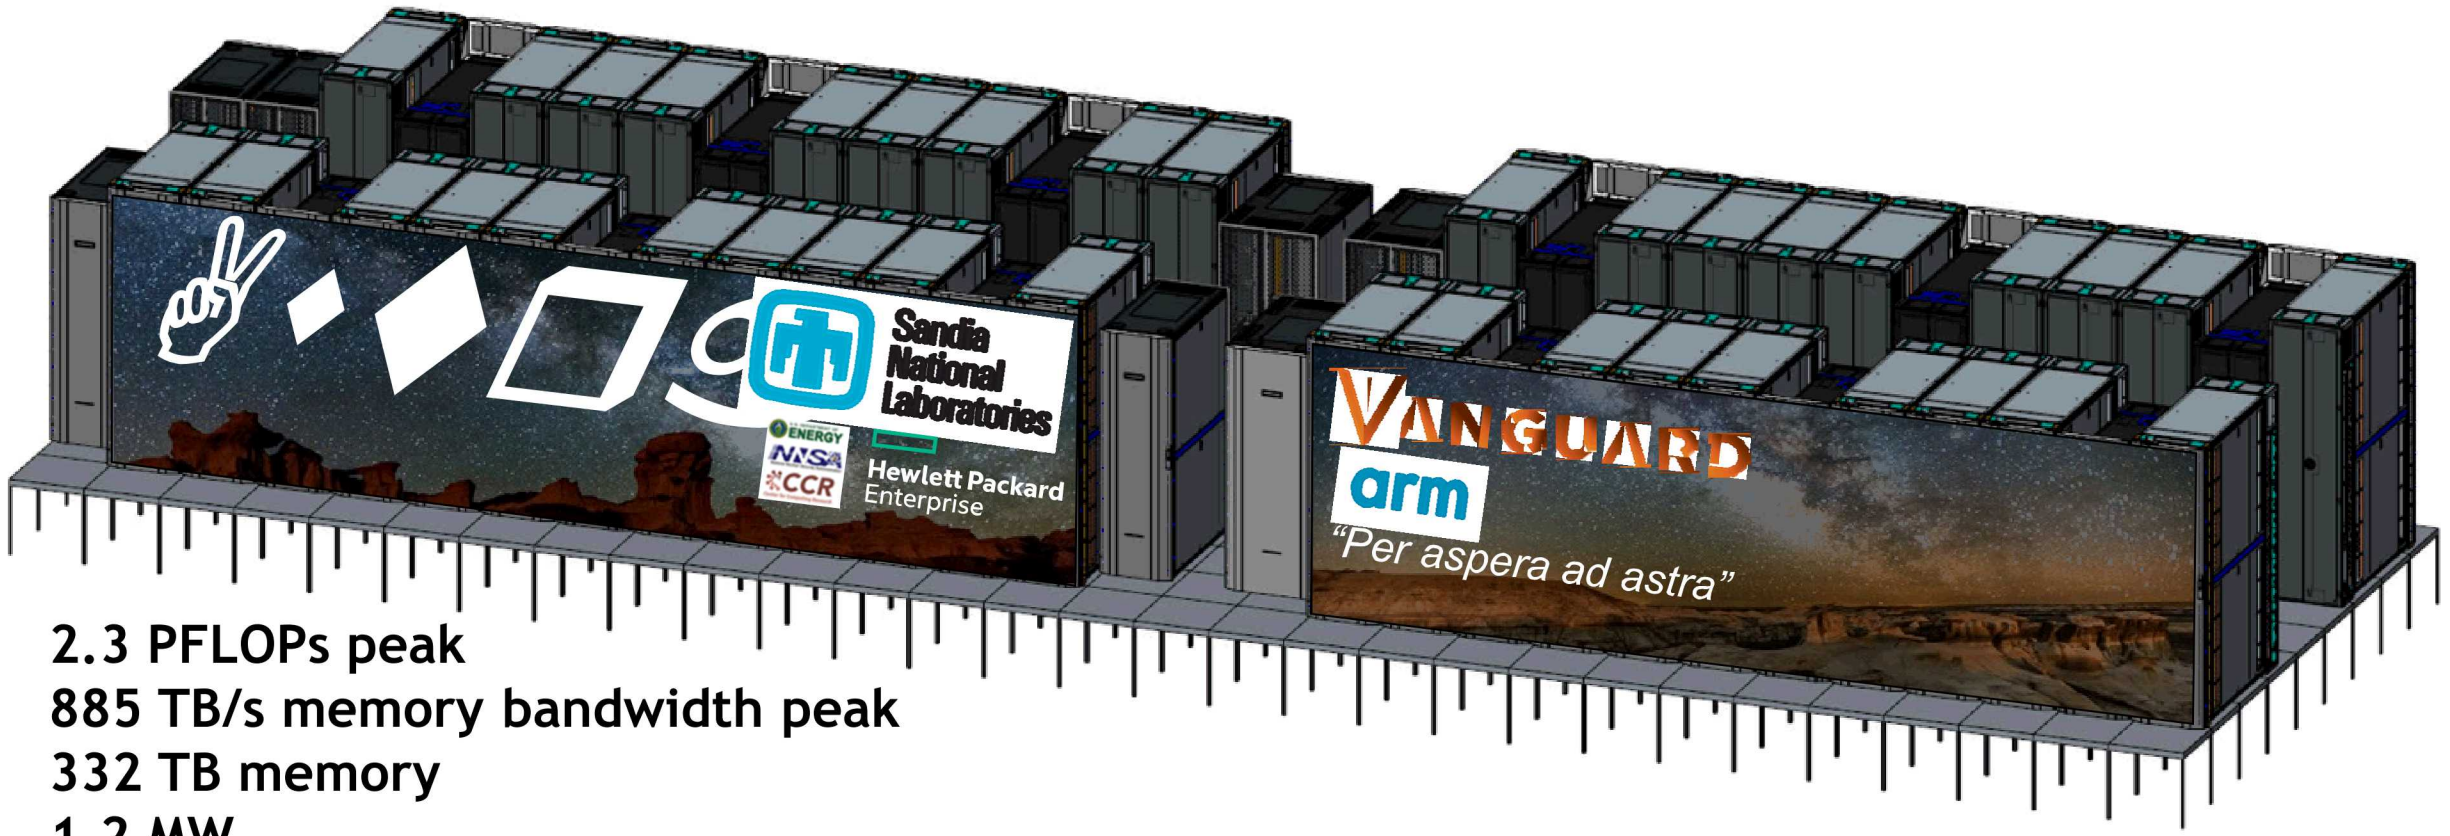

2.3 PFLOPs peak
885 TB/s memory bandwidth peak
332 TB memory
1.2 MW

Demonstrate viability of ARM for U.S. DOE Supercomputing

Vanguard-Astra System Packaging

HPE Apollo 70 Chassis: 4 nodes

18 chassis/rack

72 nodes/rack

3 IB switches/rack
(one 36-port switch
per 6 chassis)

HPE Apollo 70 Rack

Hewlett Packard
Enterprise

36 compute racks
(9 scalable units, each 4 racks)

2592 compute nodes
(5184 TX2 processors)

3 IB spine switches
(each 540-port)

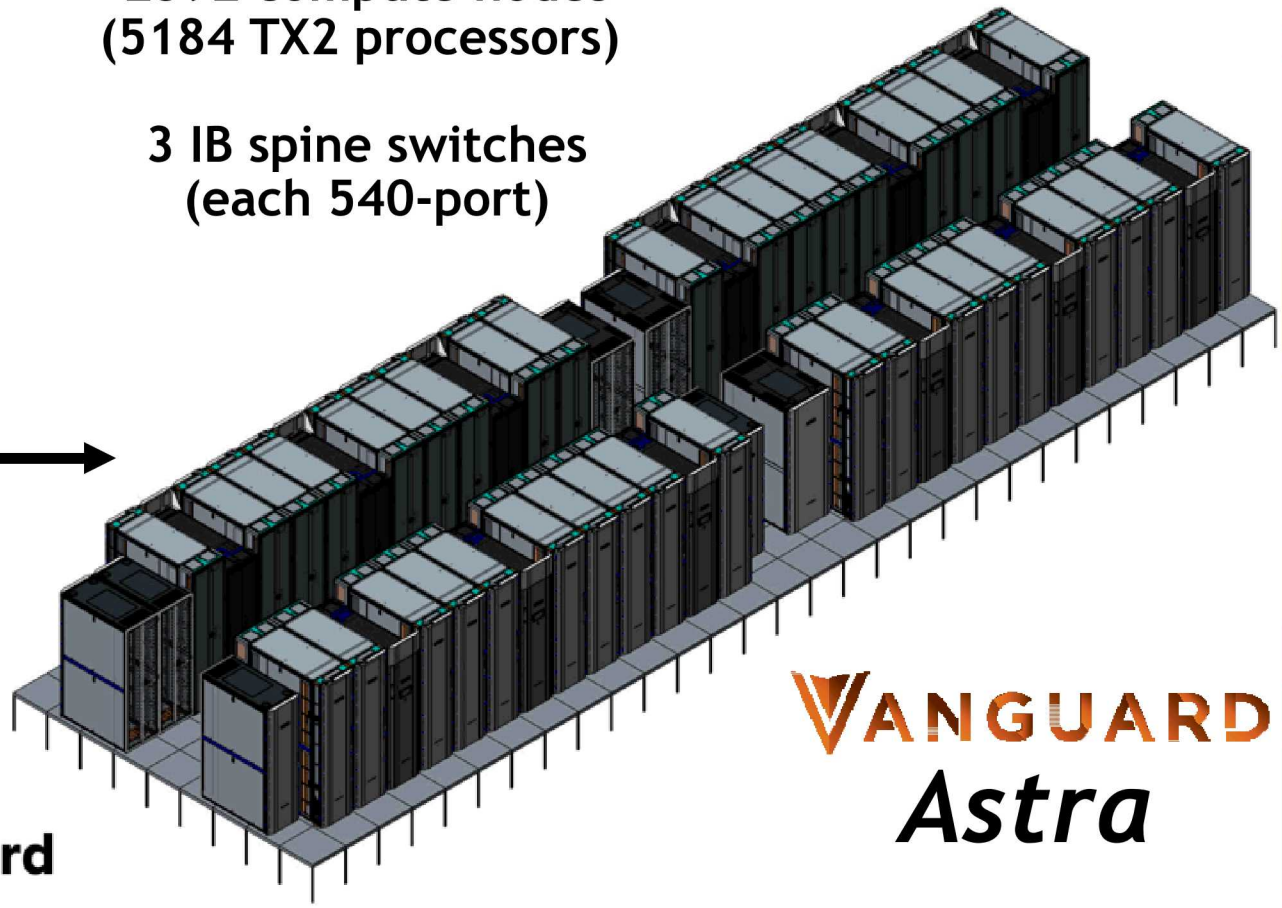

VANGUARD
Astra

Vanguard-Astra Compute Node Building Block

Hewlett Packard Enterprise · arm · CAVIUM · Mellanox TECHNOLOGIES · redhat

- Dual socket Cavium Thunder-X2
  - CN99xx
  - 28 cores @ 2.0 GHz
- 8 DDR4 controllers per socket
- One 8 GB DDR4-2666 dual-rank DIMM per controller
- Mellanox EDR InfiniBand ConnectX-5 VPI OCP
- Tri-Lab Operating System Stack based on RedHat 7.5+

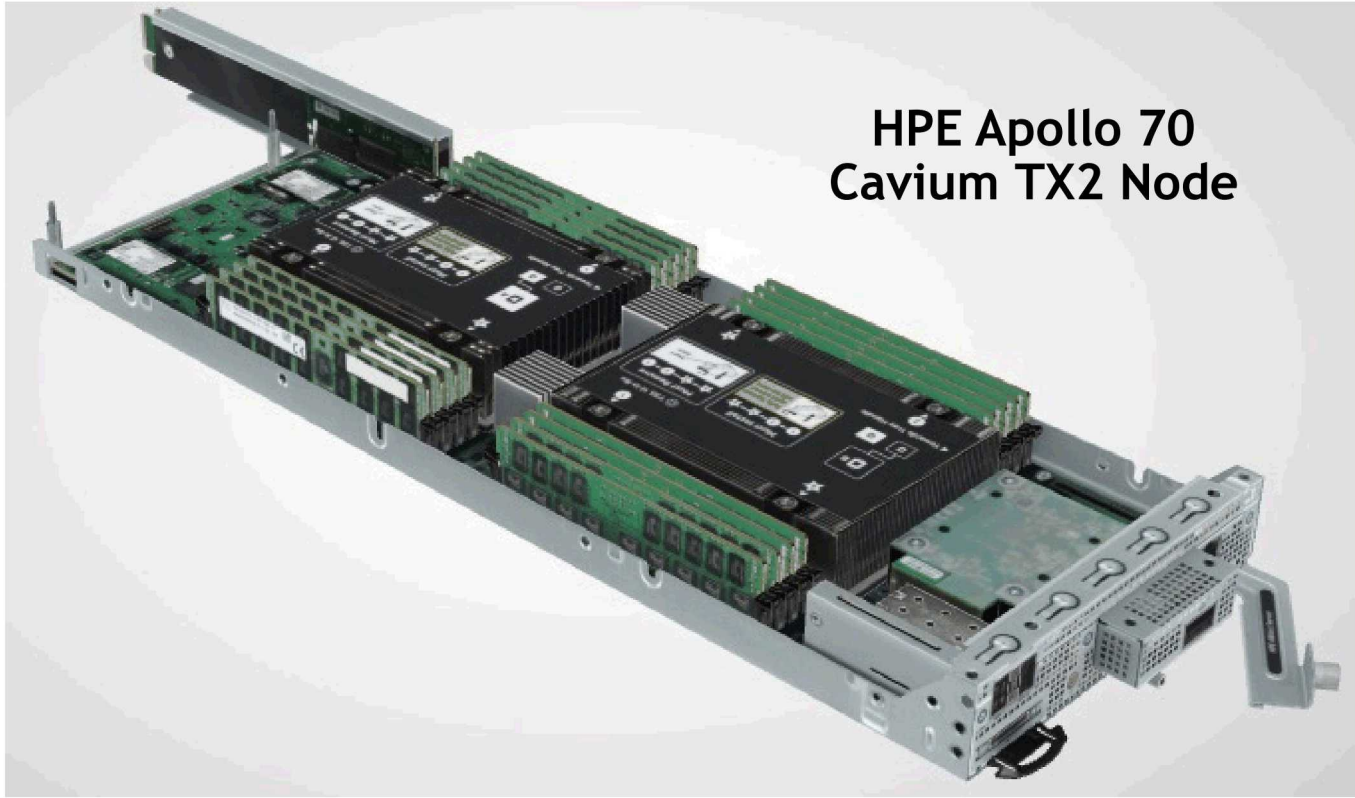

HPE Apollo 70
Cavium TX2 Node

Vanguard-Astra Compute Node

8 DDR4 channels/socket, 1 DIMM/channel
Each socket has its own PCIe x8 link to NIC

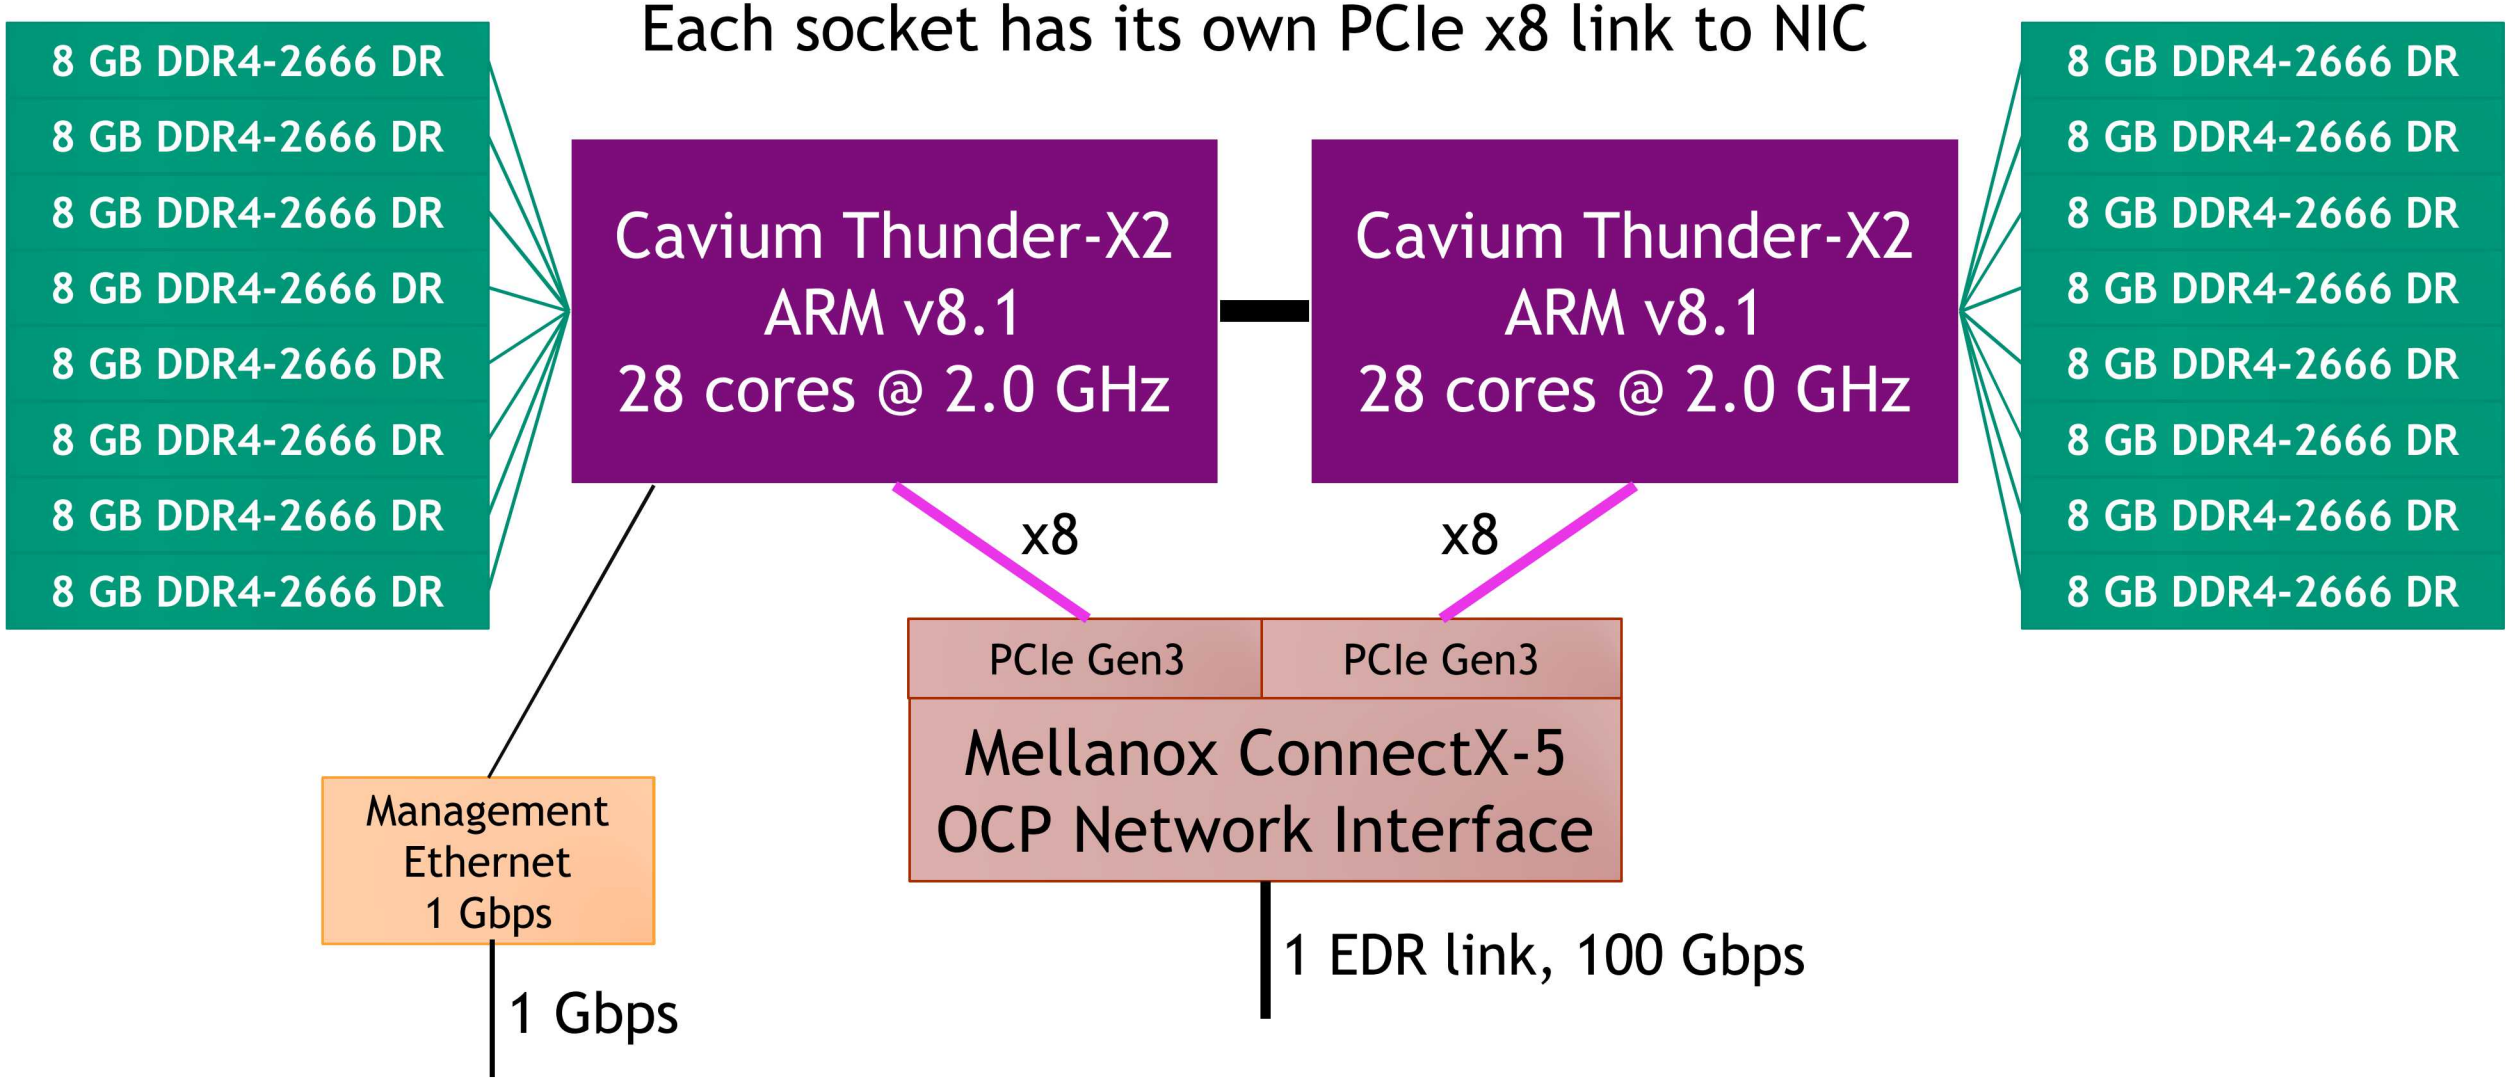

Vanguard-Astra Infrastructure

Login & Service Nodes	4 login/compilation nodes 3 Lustre routers to connect to external Sandia filesystem(s) 2 general service nodes
Interconnect	EDR InfiniBand in fat tree topology 2:1 oversubscribed for compute nodes 1:1 full bandwidth for in-platform Lustre storage
System Management	Dual HA management nodes running HPE Performance Software – Cluster Manager (HPCM) Ethernet management network, connects to all nodes One boot server per scalable unit (288 nodes)
In-platform Storage	All-flash Lustre storage system 403 TB usable capacity 244 GB/s throughput

Network Topology

Mellanox Switch-IB2 EDR, Radix 36 switches, 3 level fat tree, 2:1 taper at L1

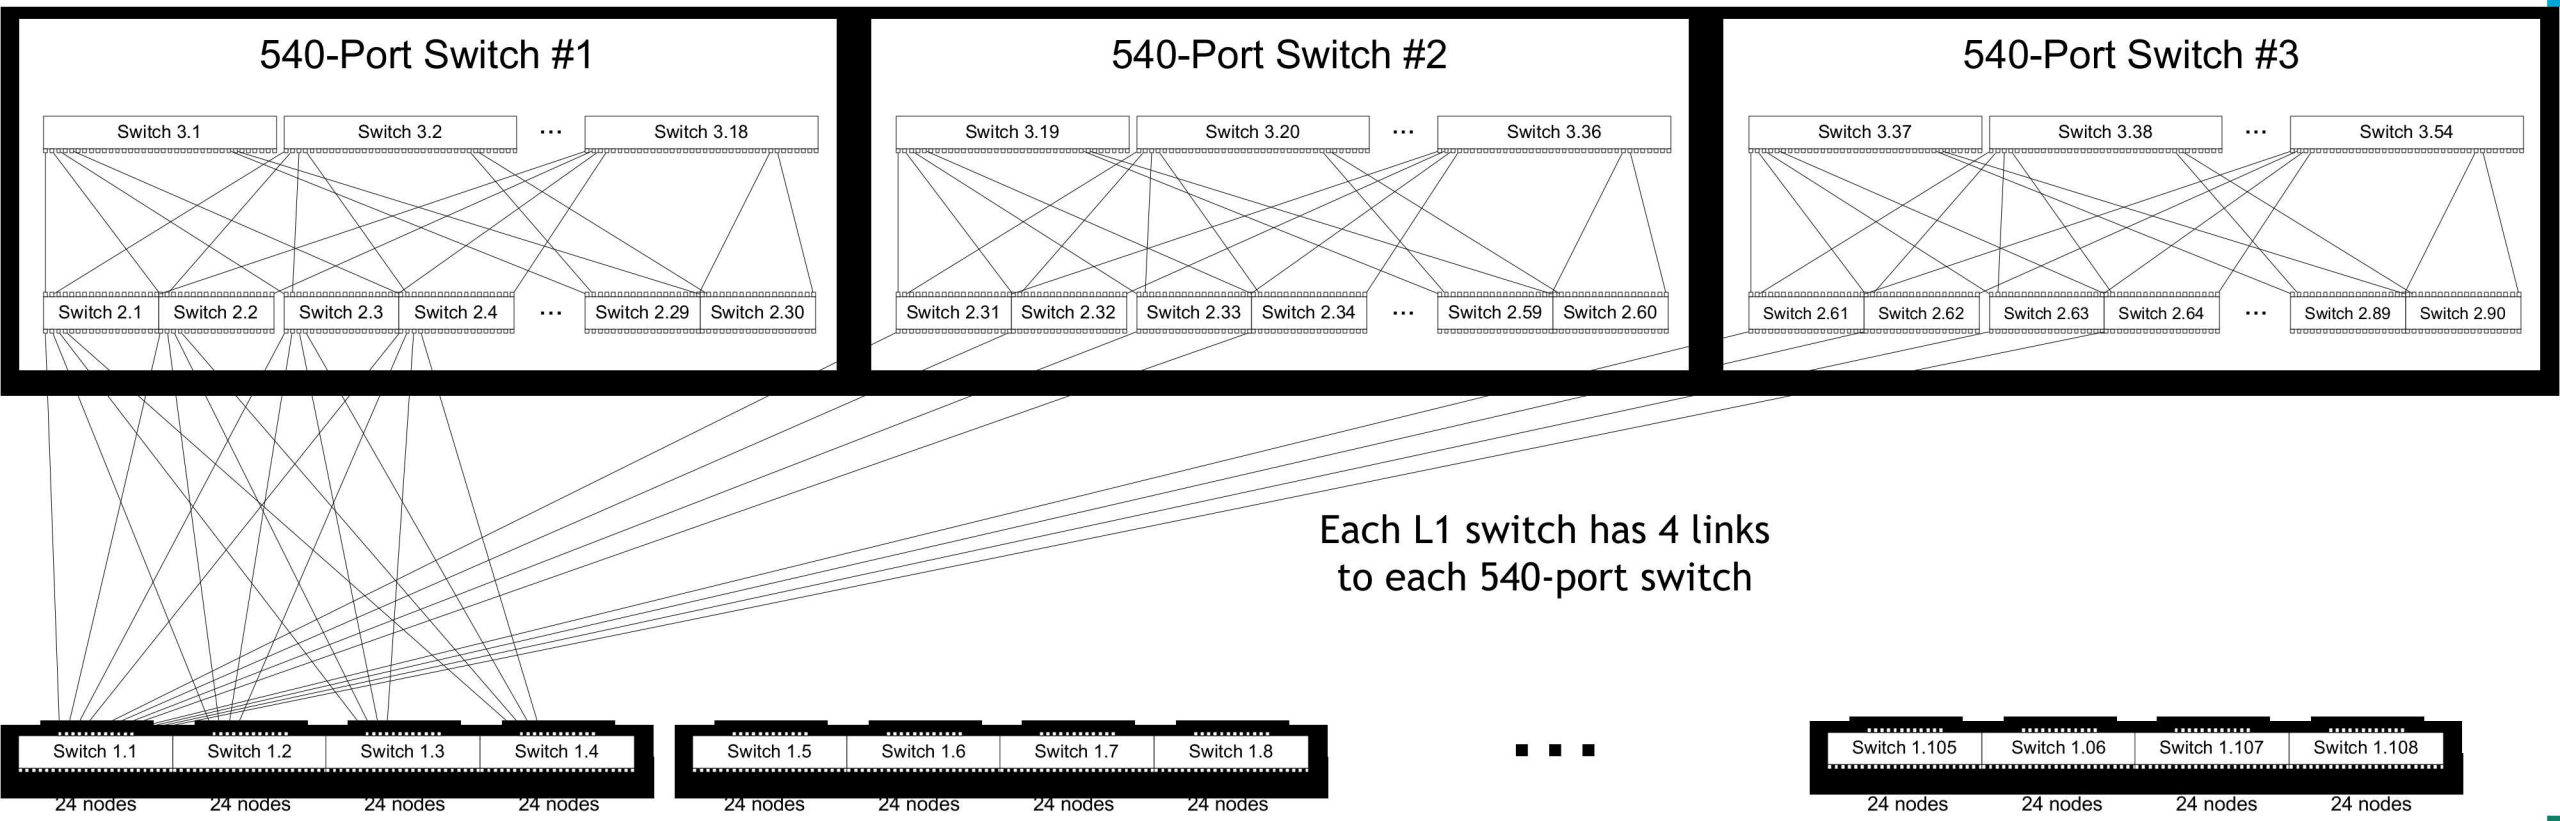

108 L1 switches * 24 nodes/switch = 2592 compute nodes

Vanguard-Astra Advanced Power & Cooling

Extreme Efficiency:

- Total 1.2 MW in the 36 compute racks are cooled by only 12 fan coils
- These coils are cooled without compressors year round. No evaporative water at all almost 6000 hours a year
- 99% of the compute racks heat never leaves the cabinet, yet the system doesn't require the internal plumbing of liquid disconnects and cold plates running across all CPUs and DIMMs

- Builds on work by NREL and Sandia: https://www.nrel.gov/esif/partnerships-jc.html

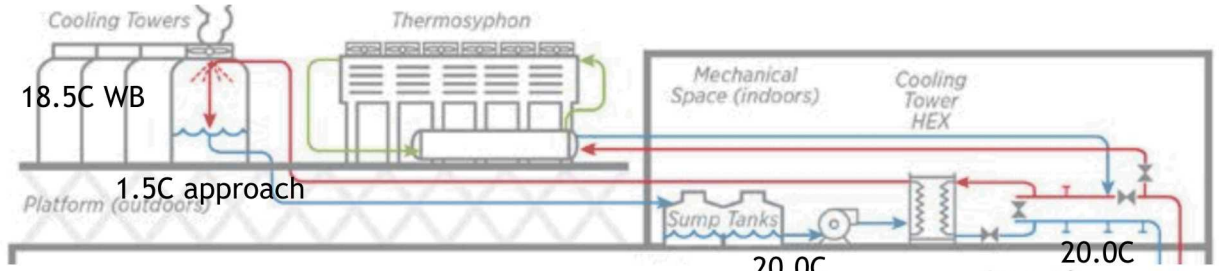

Projected power of the system by component									
	per constituent rack type (W)					total (kW)			
	wall	peak	nominal (linpack)	idle	racks	wall	peak	nominal (linpack)	idle
Node racks	39888	35993	33805	6761	36	1436.0	1295.8	1217.0	243.4
MCS300	10500	7400	7400	170	12	126.0	88.8	88.8	2.0
Network	12624	10023	9021	9021	3	37.9	30.1	27.1	27.1
Storage	11520	10000	10000	1000	2	23.0	20.0	20.0	2.0
utility	8640	5625	4500	450	1	8.6	5.6	4.5	0.5
						1631.5	1440.3	**1357.3**	274.9

ATSE – Advanced Tri-lab Software Environment

Tri-Lab Software Effort for ARM

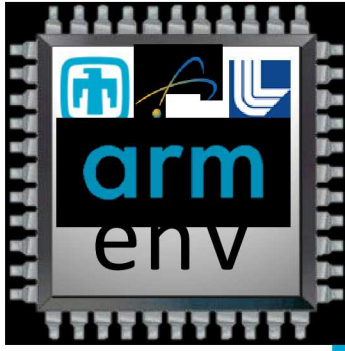

- Accelerate ARM ecosystem for ASC computing
  - Prove viability for ASC integrated codes running at scale
  - Harden compilers, math libraries, tools, communication libraries
    - Heavily templated C++, Fortran 2003/2008, Gigabyte+ binaries, long compiles
  - Optimize performance, verify expected results
- Build integrated software stack
  - Programming environment (compilers, math libs, tools, MPI, OMP, SHMEM, I/O, ...)
  - Low-level OS (optimized Linux, network, filesystems, containers/VMs, ...)
  - Job scheduling and management (WLM, app launcher, user tools, ...)
  - System management (boot, system monitoring, image management, ...)

Improve 0 to 60 time... ARM system arrival to useful work done

Advanced Trilab Software Environment (ATSE)

- Advanced Trilab Software Environment
  - Sandia leading development with Trilab Arm team
  - Partnership across the NNSA/ASC Labs and with HPE
  - Provide a user programming environment for Astra
    - Initial focus on ARM, have x86_64 port
  - Partnership across the NNSA/ASC Labs and with HPE
- Lasting value beyond just ARM
  - Documented specification of:
    - Software components needed for HPC production applications
    - How they are configured (i.e., what features and capabilities are enabled) and interact
    - User interfaces and conventions
  - Reference implementation:
    - Deployable on multiple ASC systems and architectures with common look and feel
    - Tested against real ASC workloads
    - Community inspired, focused and supported
    - Leveraging OpenHPC effort
    - Inform & improve vendor supplied software stack

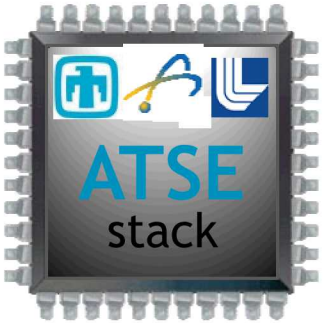

Hewlett Packard Enterprise

Sandia National Laboratories

Lawrence Livermore National Laboratory

openHPC

Los Alamos NATIONAL LABORATORY

ATSE is an integrated software environment for ASC workloads

Advanced Tri-lab Software Environment Goals

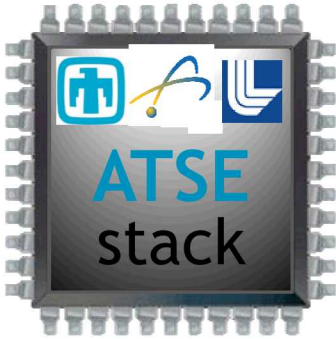

- Build an open, modular, extensible, community-inspired, and vendor-adaptable ecosystem
- Prototype new technologies that may improve the DOE ASC computing environment (e.g., ML frameworks, containers, VMs, etc)
- Leverage existing efforts
  - Tri-lab OS (TOSS)
  - OpenHPC & other programming environments
  - Exascale Computing Project (ECP) software technologies

ARM Tri-lab Software Environment (ATSE)

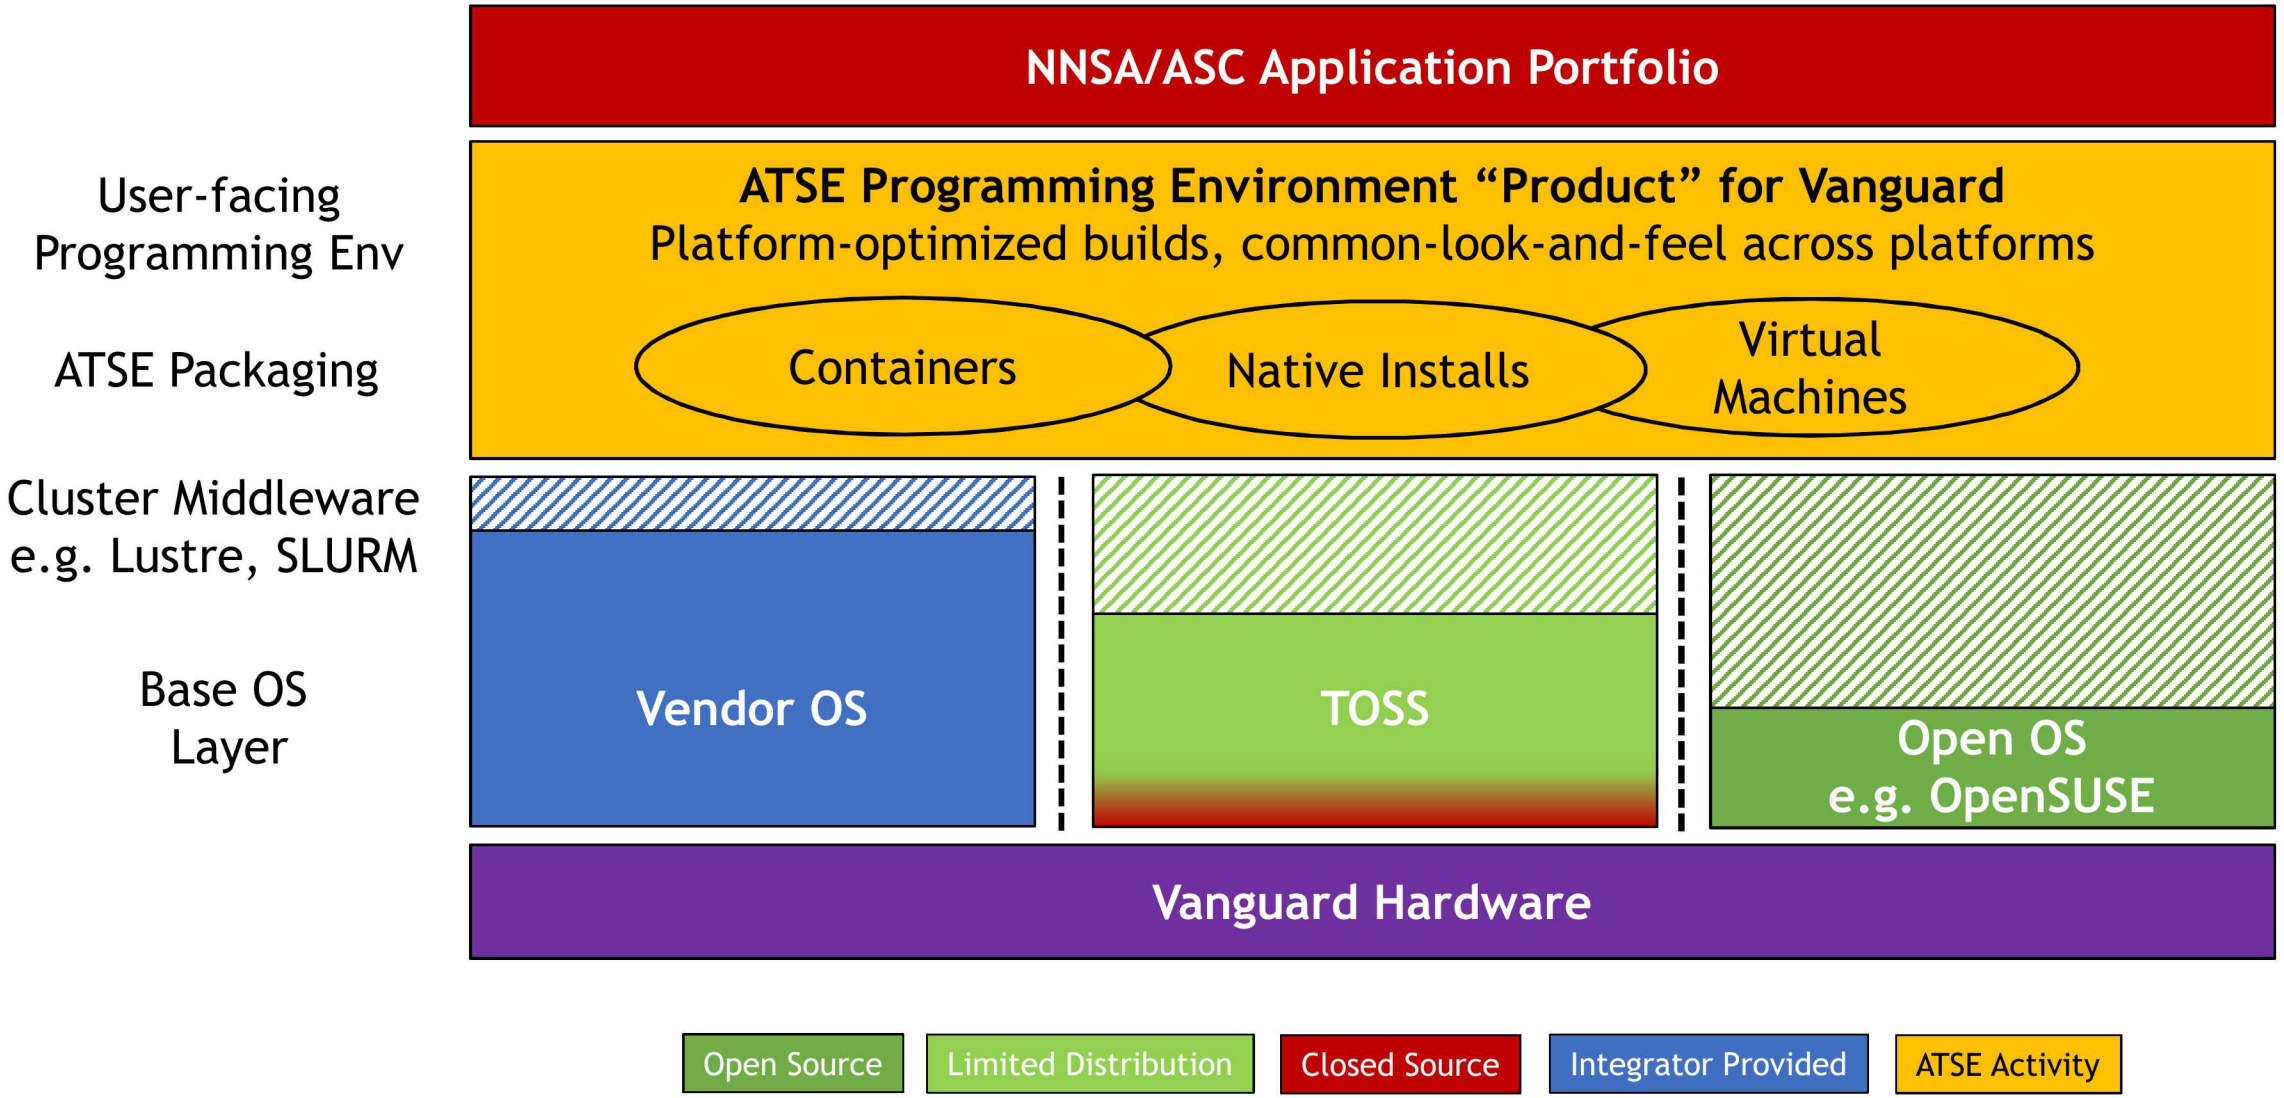

Research Directions

Early Results from Astra

- System has been online for around two weeks , incredible team
- Already running full application ports and many of our key frameworks
- First Petascale ARM platform, designed for production workloads
  - HPL: 1.5 Pflops Rmax, 2 Pflops Rpeak on Top500
  - HPCG: 67 Tflops, 36th on Top500

Baseline: Trinity ASC Platform (Current Production), dual-socket Haswell

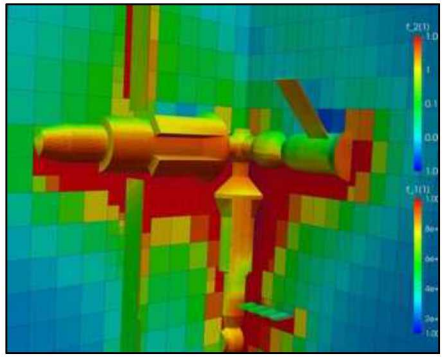

Monte Carlo

1.60X

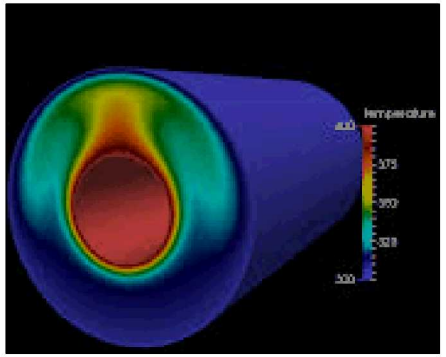

CFD Models

1.45X

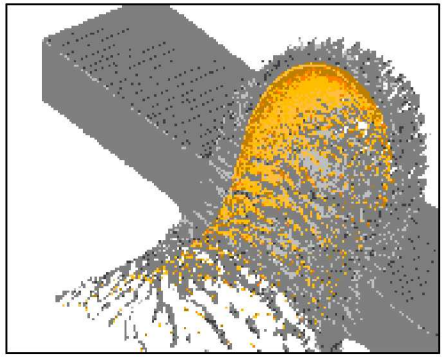

Hydrodynamics

1.30X

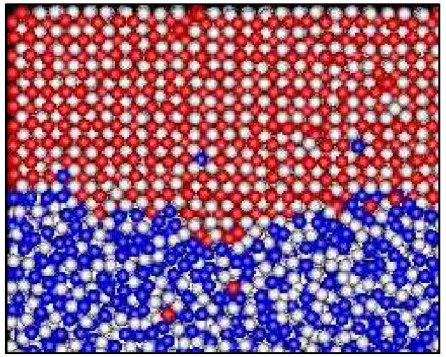

Molecular Dynamics

1.42X

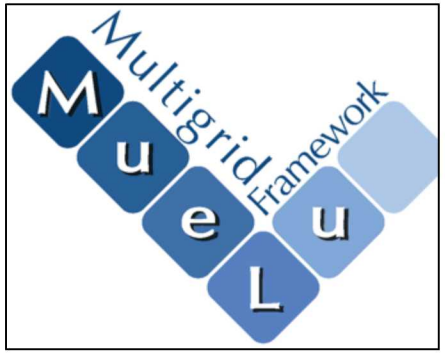

Linear Solvers

1.87X

Network Performance at scale

- **IMB Benchmark**
- **Mellanox SHARP**

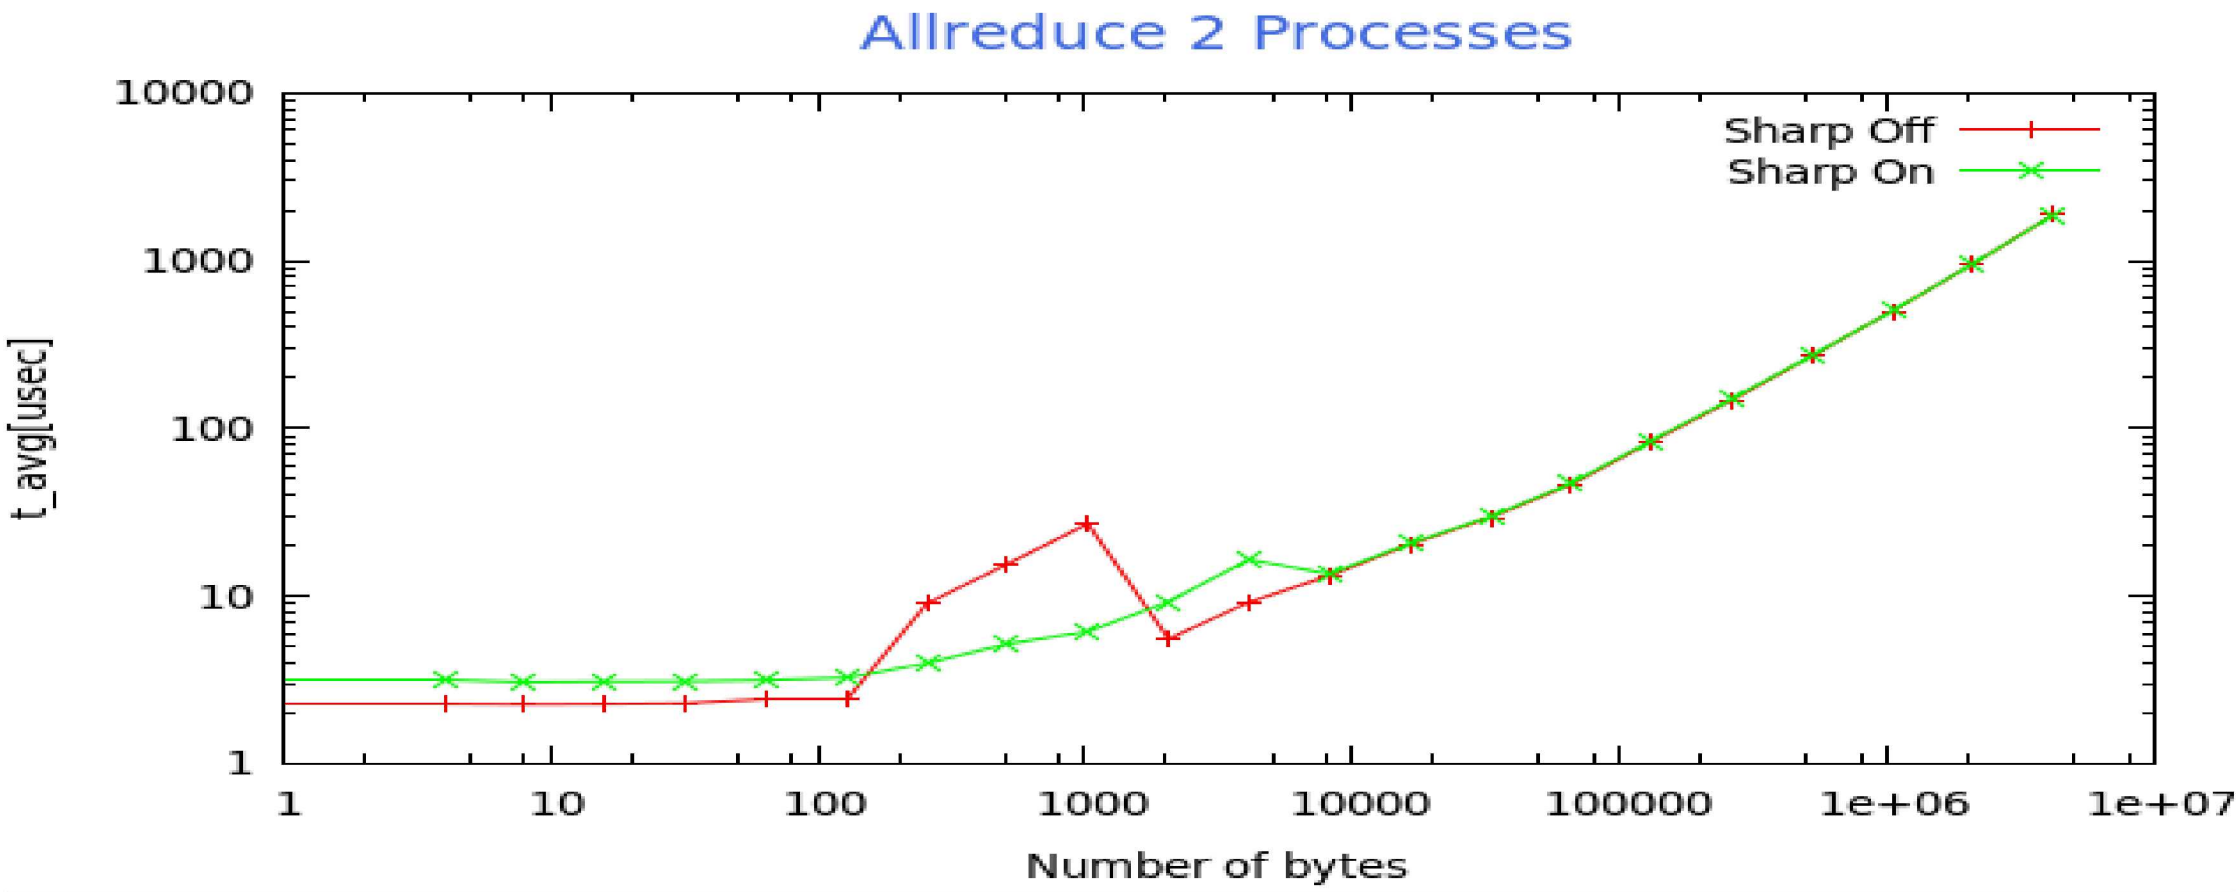

From: M. Aguilar et al, OpenFabrics Alliance Workshop 2019

Containers on ARM

- Leverage containers and virtual machines on ARM
- Singularity Containers
  - Working with Sylabs on full container solution
  - Support emerging ML/AI frameworks
  - Leverage remote builder, library, and secure signing services
  - Evaluate container scalability

- KVM Virtual Machine support
  - ARMv8.1 includes virtualization extensions, SR-IOV
  - Optimize and tune with libvirt for TX2

Concurrent R&D Areas

- Evaluating parallel filesystems + I/O systems @ scale
  - GlusterFS, Ceph, BeeGFS, Sandia Data Warehouse, …
- Resilience studies over Astra lifetime
- Improved MPI thread support, matching acceleration
- OS optimizations for HPC @ scale
  - Exploring HPC-tuned Linux kernels
  - non-Linux lightweight kernels and multi-kernels
  - Arm-specific optimizations?

McKernel

BUGS

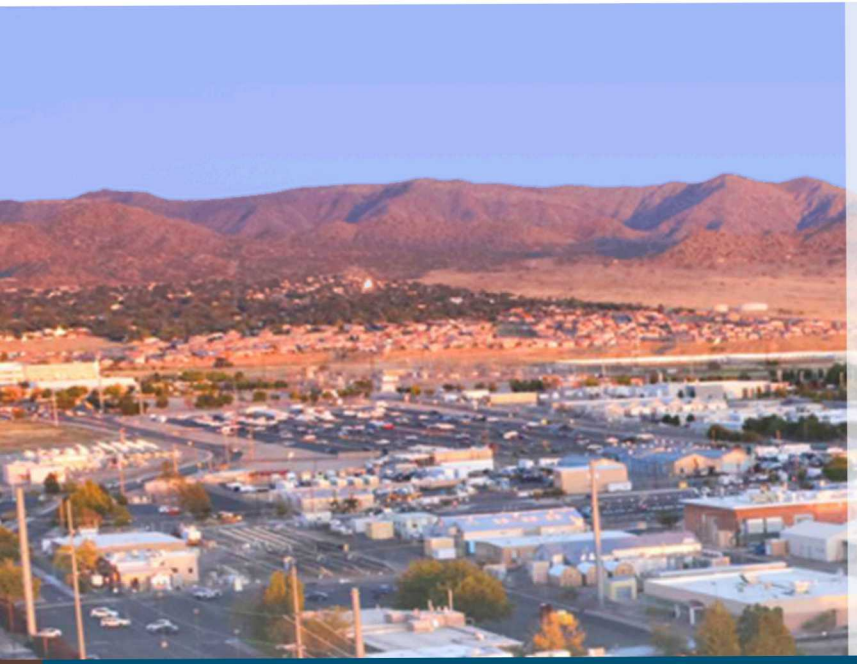

Questions?

ajyoung@sandia.gov