Vanguard Astra: Maturing the ARM Software Ecosystem for U.S. DOE/ASC Supercomputing

GoingArm @ ISC18
June 28, 2018

Kevin Pedretti, Jim H. Laros III, Si Hammond
ktpedre@sandia.gov

SAND2018-7067 C

Sandia National Laboratories

Exceptional service in the national interest

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: Large-scale Prototype Systems

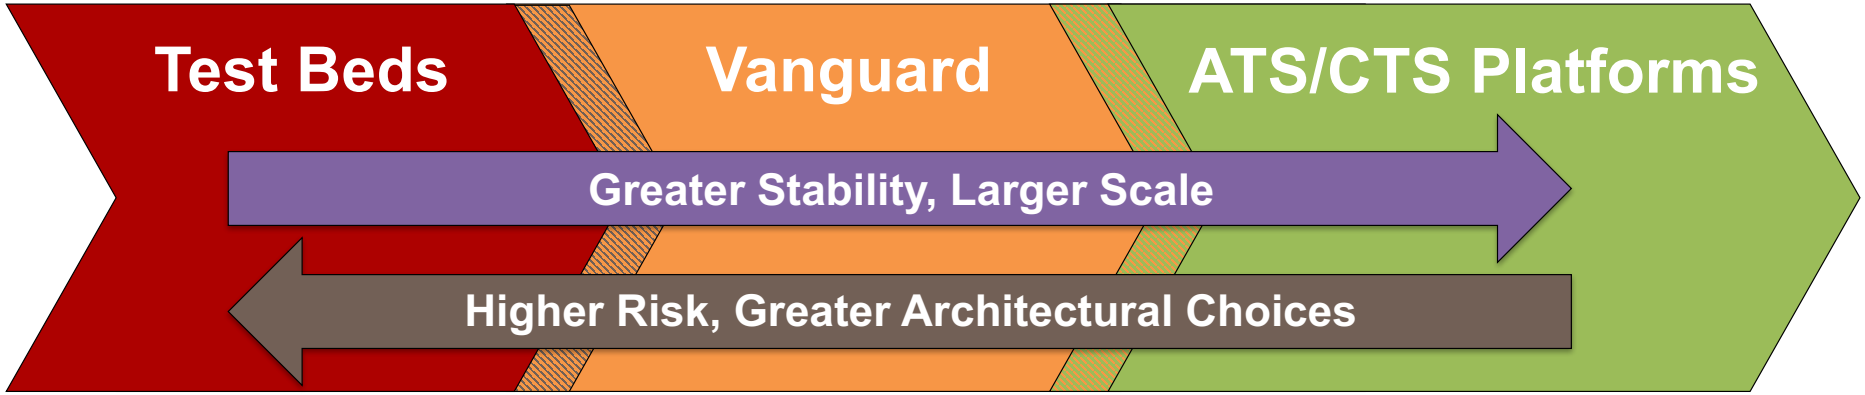

Test Beds

- Small testbeds (~10-100 nodes)
- Breadth of architectures Key
- **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's NNSA/ASC ARM Platforms

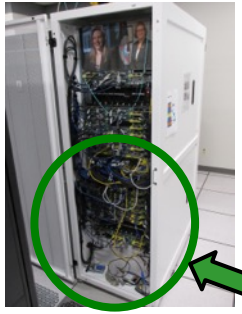

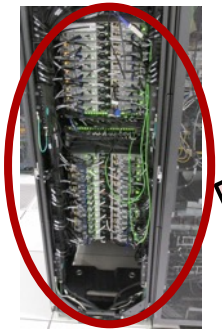

Sept 2011

APM/HPE
Xgene-1
ARM
Hammer

Applied Micro
X-Gene-1
47 nodes

Cavium/Penguin
ThunderX1
ARM
Sullivan

Cavium
ThunderX1
32 nodes

Cavium
HPE/Comanche
ThunderX2
ARM
Mayer

Pre-GA
Cavium
ThunderX2
47 nodes

TODAY

Vanguard
ARM
Astra

Petascale ARM
Platform
Delivery 2018

HPE Apollo 70
Cavium ThunderX2
Mellanox ConnectX-5
Switch-IB2
2592 compute nodes

Future ASC
Platforms

= Retired
= 2015
= 2017
= 2018

per aspera ad astra

Sandia National Laboratories

through difficulties to the stars

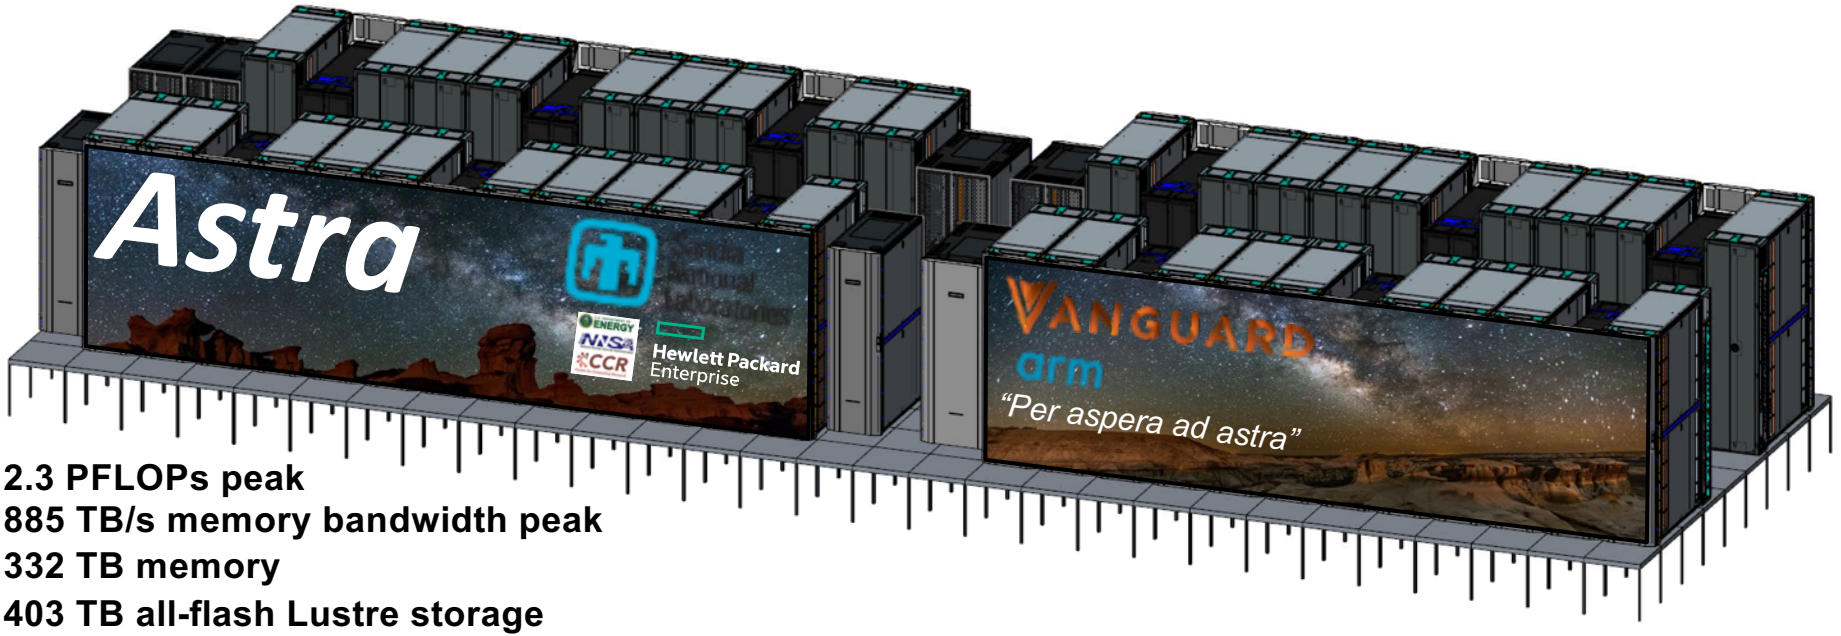

2.3 PFLOPs peak
885 TB/s memory bandwidth peak
332 TB memory
403 TB all-flash Lustre storage
1.2 MW

Demonstrate viability of ARM for U.S. DOE NNSA Supercomputing

Vanguard-Astra Compute Node Building Block

Hewlett Packard Enterprise · arm · CAVIUM · Mellanox TECHNOLOGIES · redhat · Sandia National Laboratories

- 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+

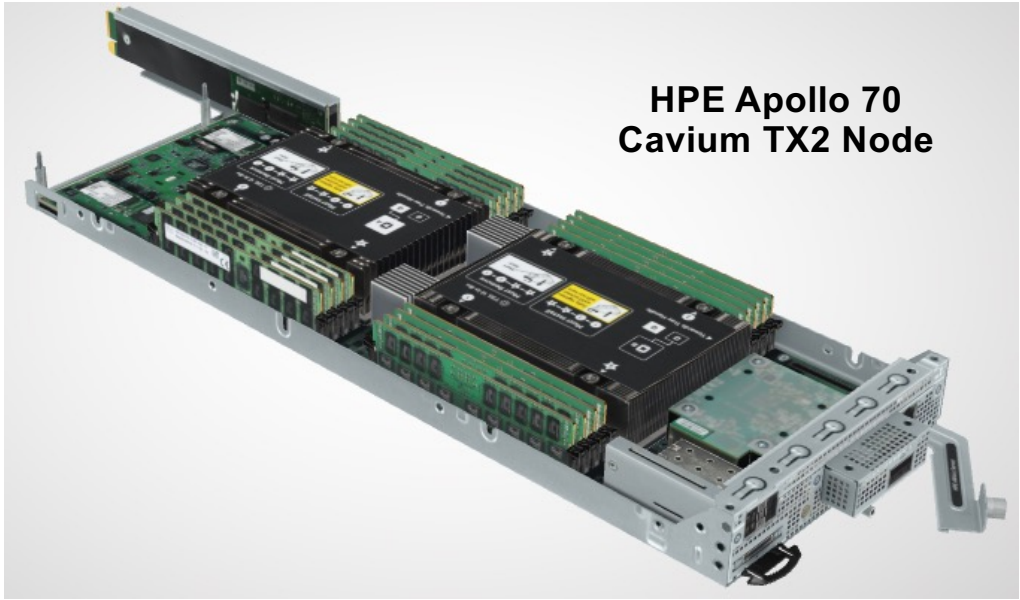

Vanguard-Astra System Packaging

Sandia National Laboratories

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

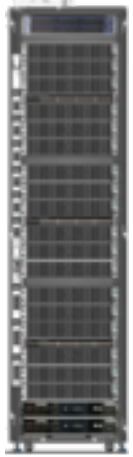

36 compute racks
(9 scalable units, each 4 racks)

2592 compute nodes
(5184 TX2 processors)

3 IB spine switches
(each 540-port)

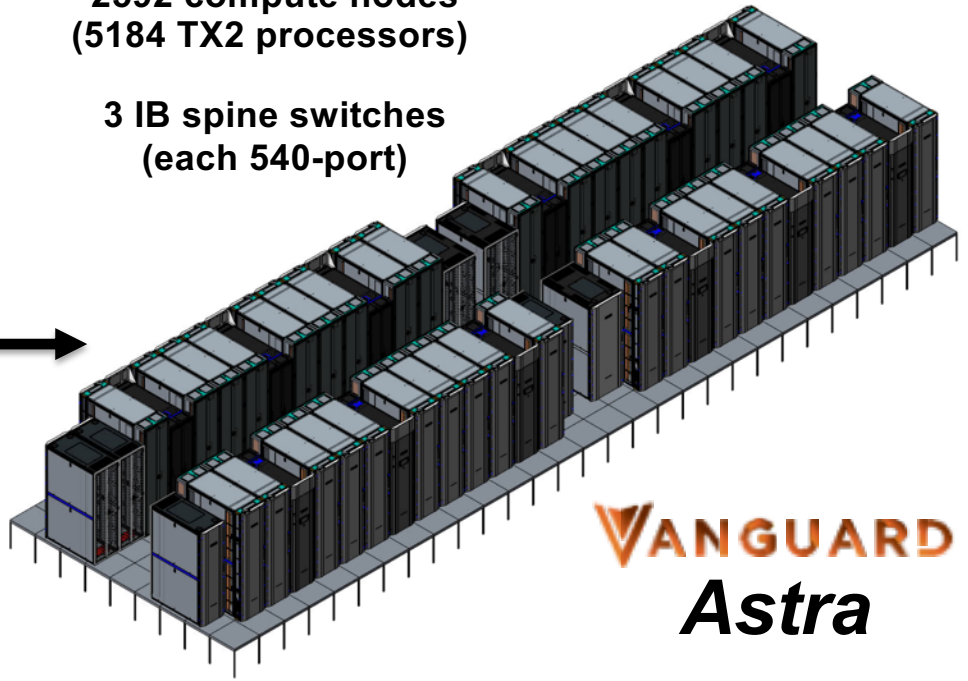

VANGUARD
Astra

Network Topology Visualization

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

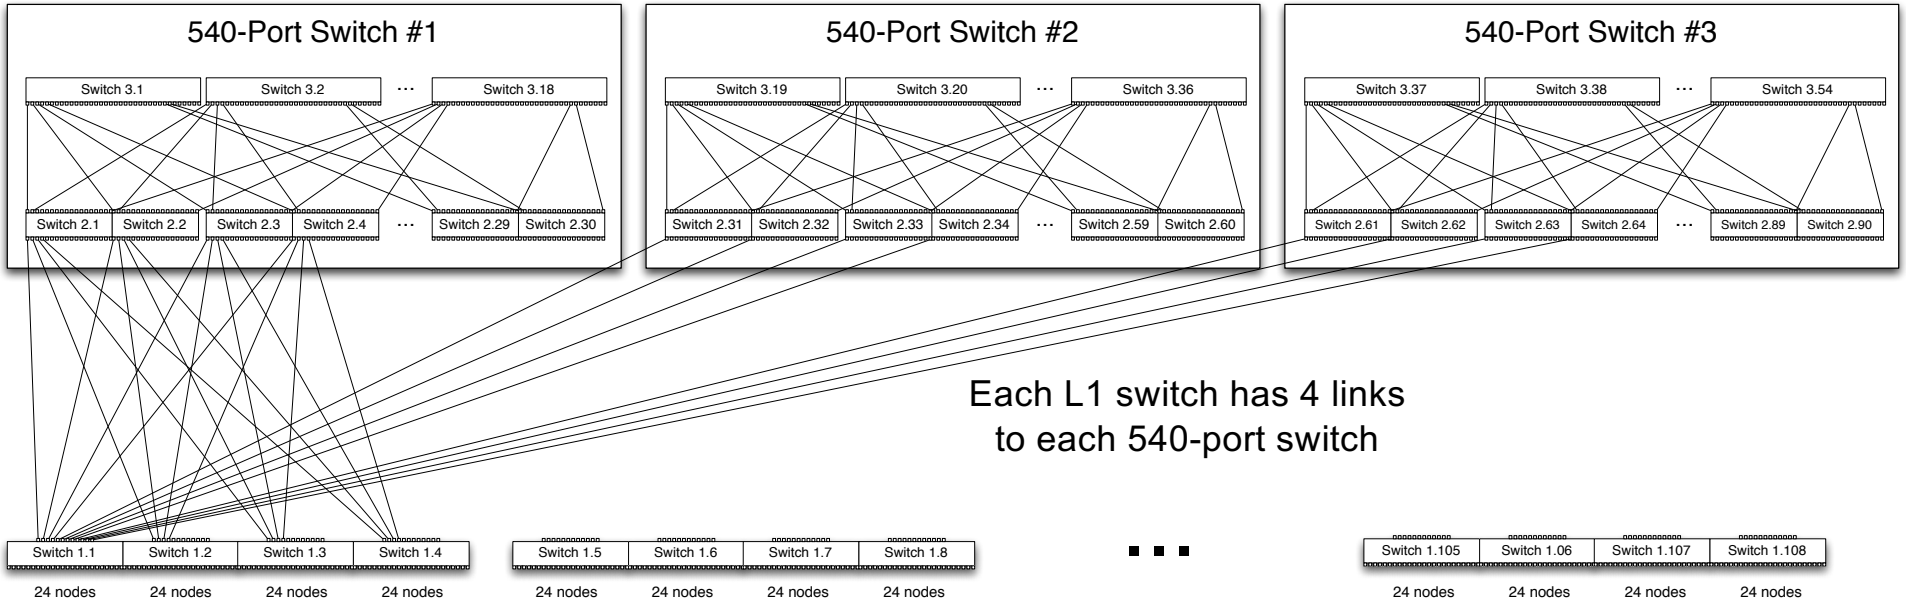

108 L1 switches * 24 nodes/switch = 2592 compute nodes

Vanguard-Astra Advanced Power & Cooling

Power and Water Efficient:

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

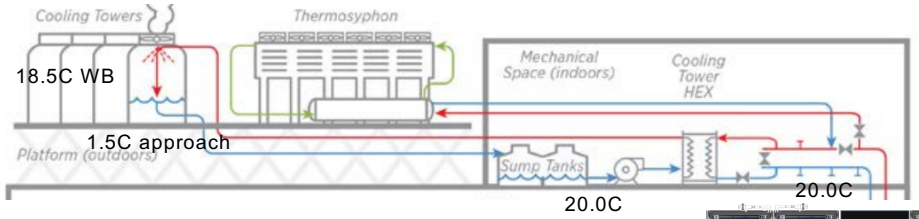

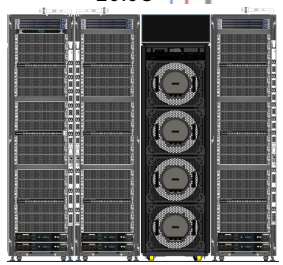

Projected power of the system by component										
	per constituent rack type (W)					total (kW)				
	wall	peak	nominal (linpack)	idle		racks	wall	peak	nominal (linpack)	idle
Node racks	39888	35993	33805	6761		36	1436.0	1295.8	1217.0	243.4
MCS300	10500	7400	7400	170		12	126.0	88.8	88.8	2.0
Network	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

Cavium Arm64 Providing Best-of-Class Memory Bandwidth

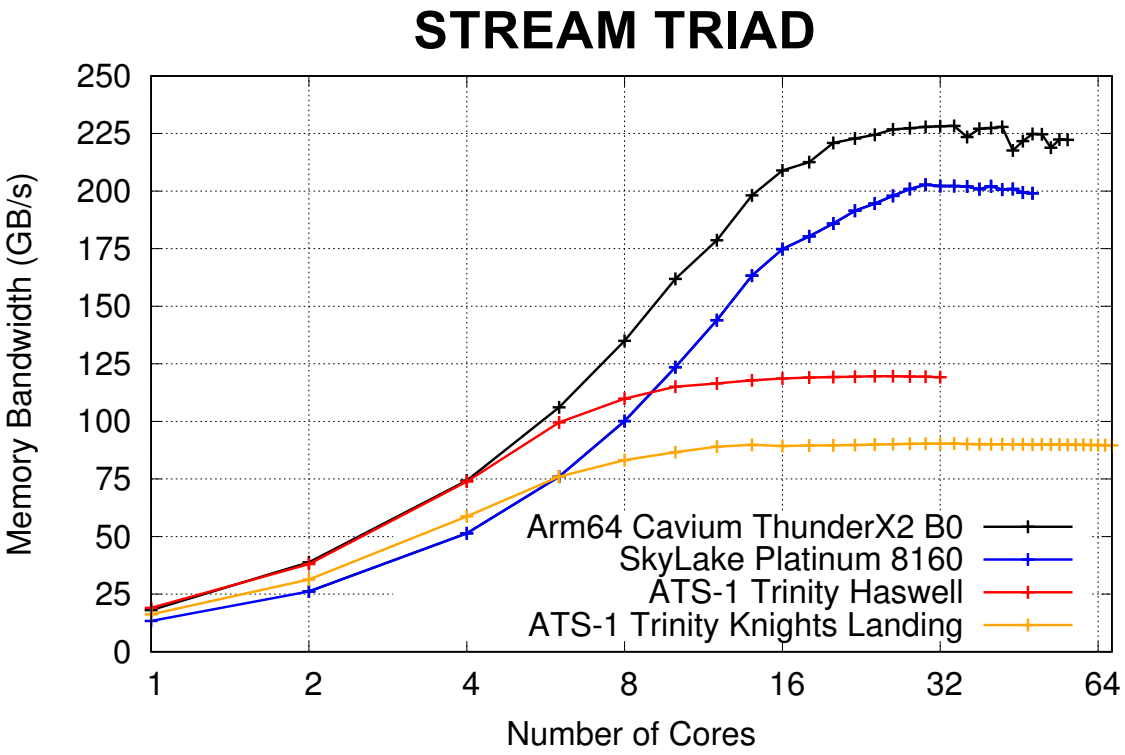

TX2 DDR4-2400

SkyLake 8160

Trinity Haswell

Trinity KNL DDR

Network Bandwidth on ThunderX2 + Mellanox MLX5 EDR with Socket Direct

Socket Direct – Each socket has dedicated path to the NIC

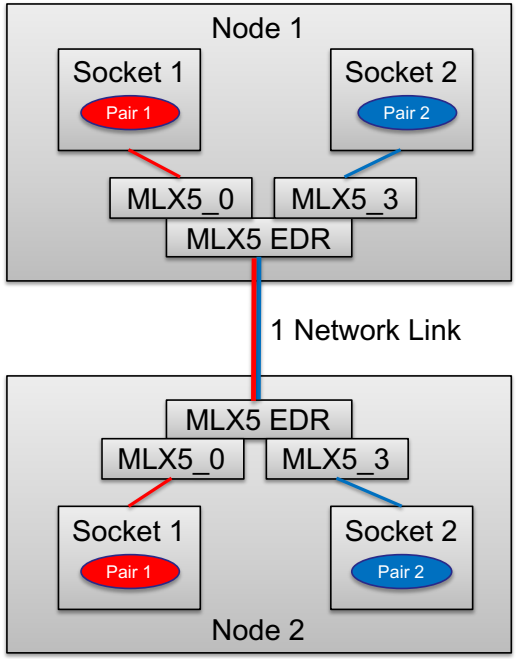

OSU MPI Multi-Network Bandwidth

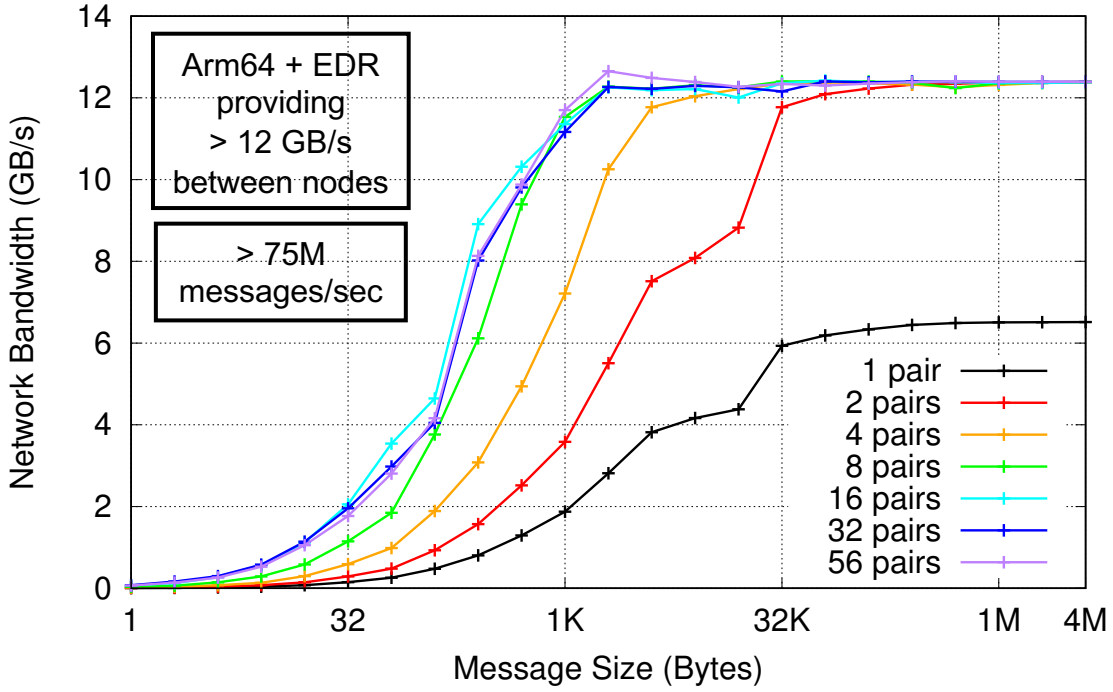

Mini-App Performance on Cavium ThunderX2

- ThunderX2 providing high memory bandwidth
  - 6 channels (Skylake) vs. 8 in ThunderX2
  - See this in MiniFE SpMV and STREAM Triad
- Slower compute reflects less optimization in software stack
  - Examples – Non-SpMV kernels in MiniFE and LULESH
  - GCC and ARM versus Intel compiler

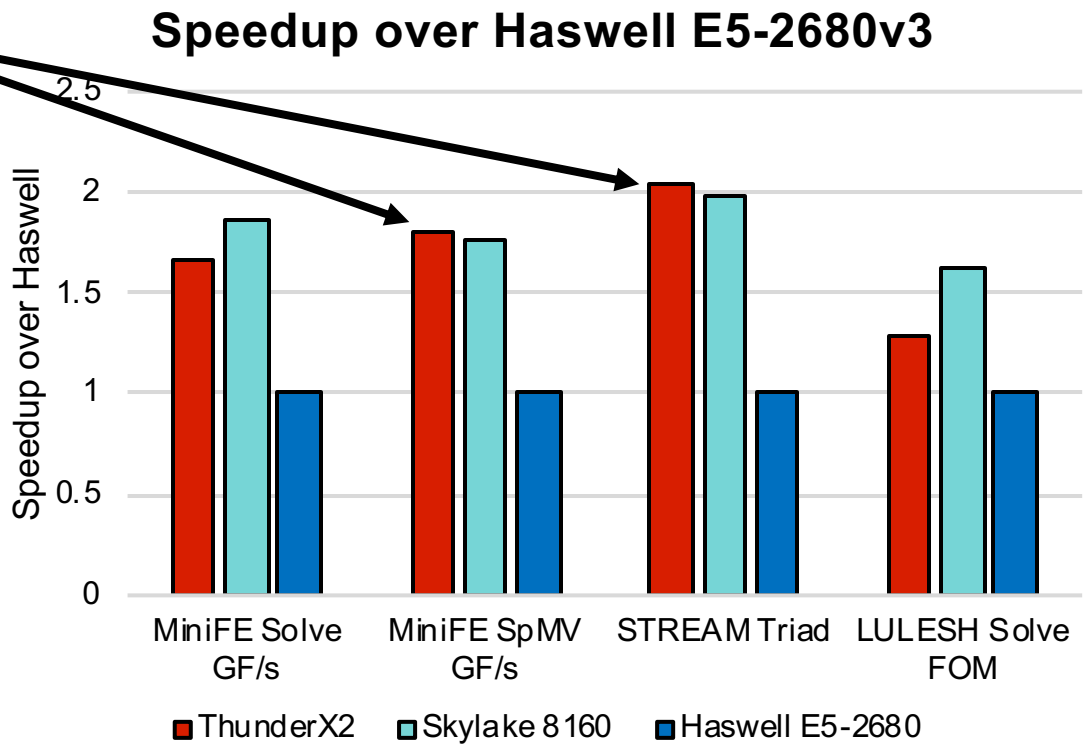

Advanced Tri-lab Software Environment High-level Goals

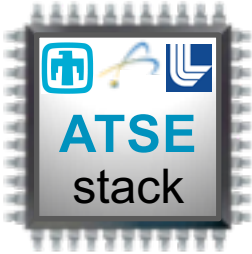

- Build an open, modular, extensible, community-engaged, and vendor-adaptable ecosystem
- Prototype new technologies that may improve the DOE ASC computing environment (e.g., ML frameworks, containers, VMs, OS optimizations)
- Leverage existing efforts such as Tri-lab OS (TOSS), programing environments, and Exascale Computing Project software technologies

Vanguard-Astra Software Stack

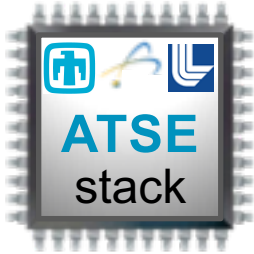

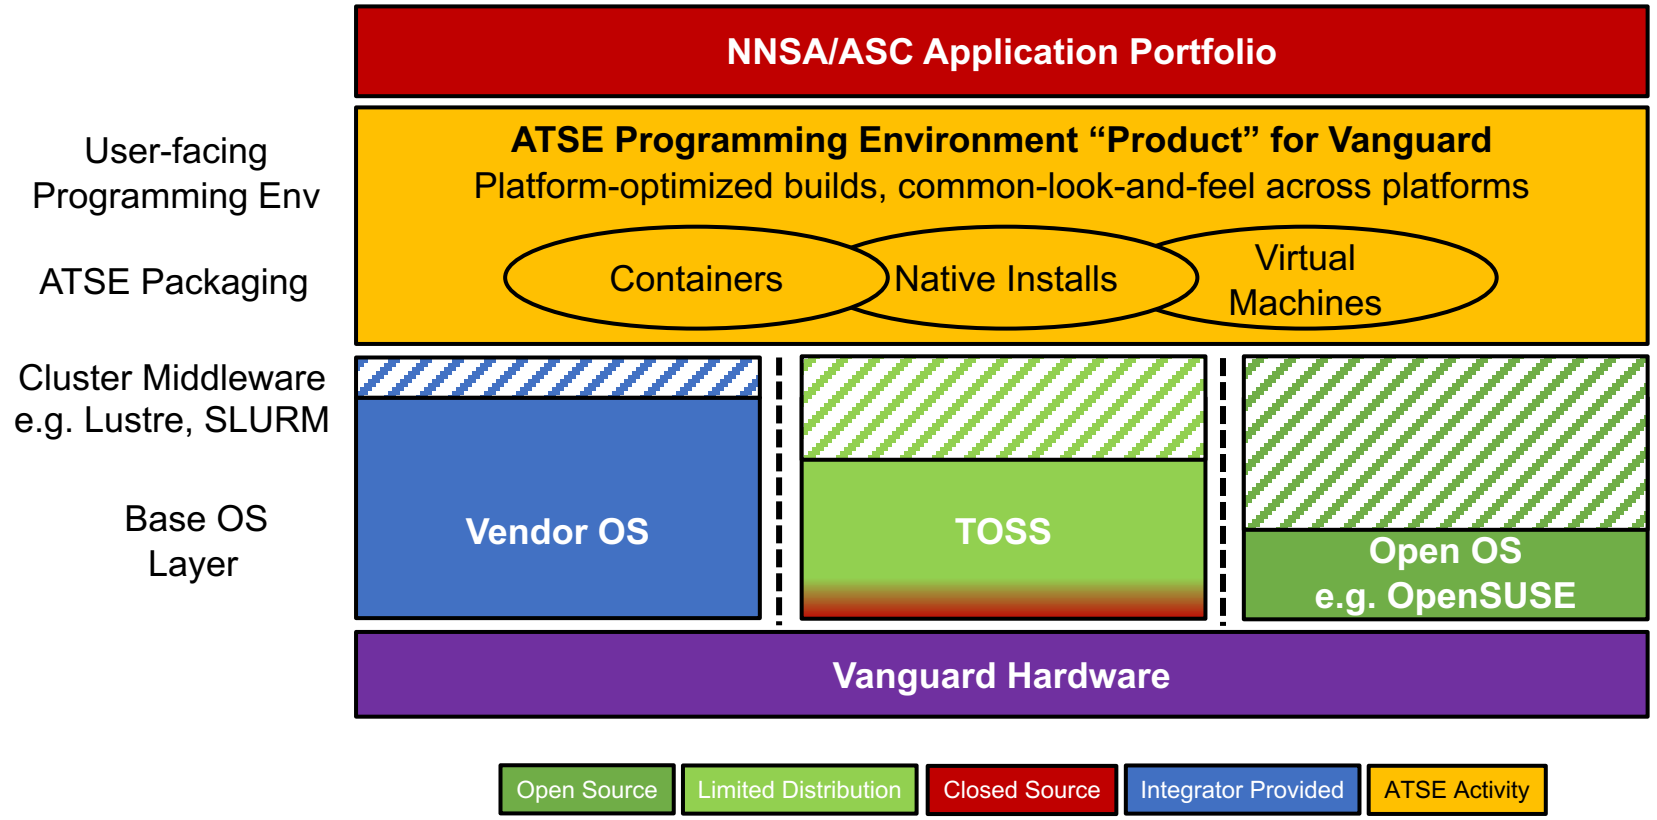

R&D Areas

- Workflows leveraging containers and virtual machines
  - Support for machine learning frameworks
  - ARMv8.1 includes new virtualization extensions, SR-IOV
- 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 spectrum from stock distro Linux kernel to HPC-tuned Linux kernels to non-Linux lightweight kernels and multi-kernels
  - Arm-specific optimizations

Conclusion

- Vanguard expanding HPC ecosystem by developing emerging, yet-to-be-proven, technologies, taking appropriate risk
  - Mature new technologies for NNSA ASC integrated codes
- Vanguard-Astra will be one of the first Arm-based supercomputers
- NNSA Tri-lab team (Sandia, Los Alamos, Lawrence Livermore) is working in partnership with HPE, Arm, Cavium, RedHat, Mellanox, and others to develop the ATSE software stack for Astra

Vanguard Collaboration

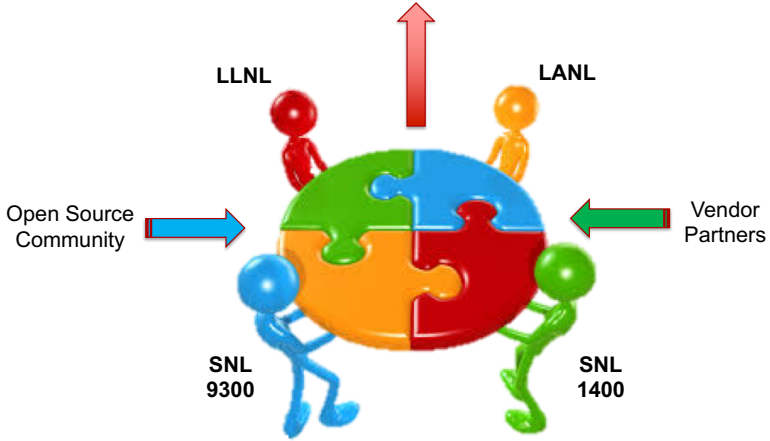

per aspera ad astra

Sandia National Laboratories

through difficulties to the stars

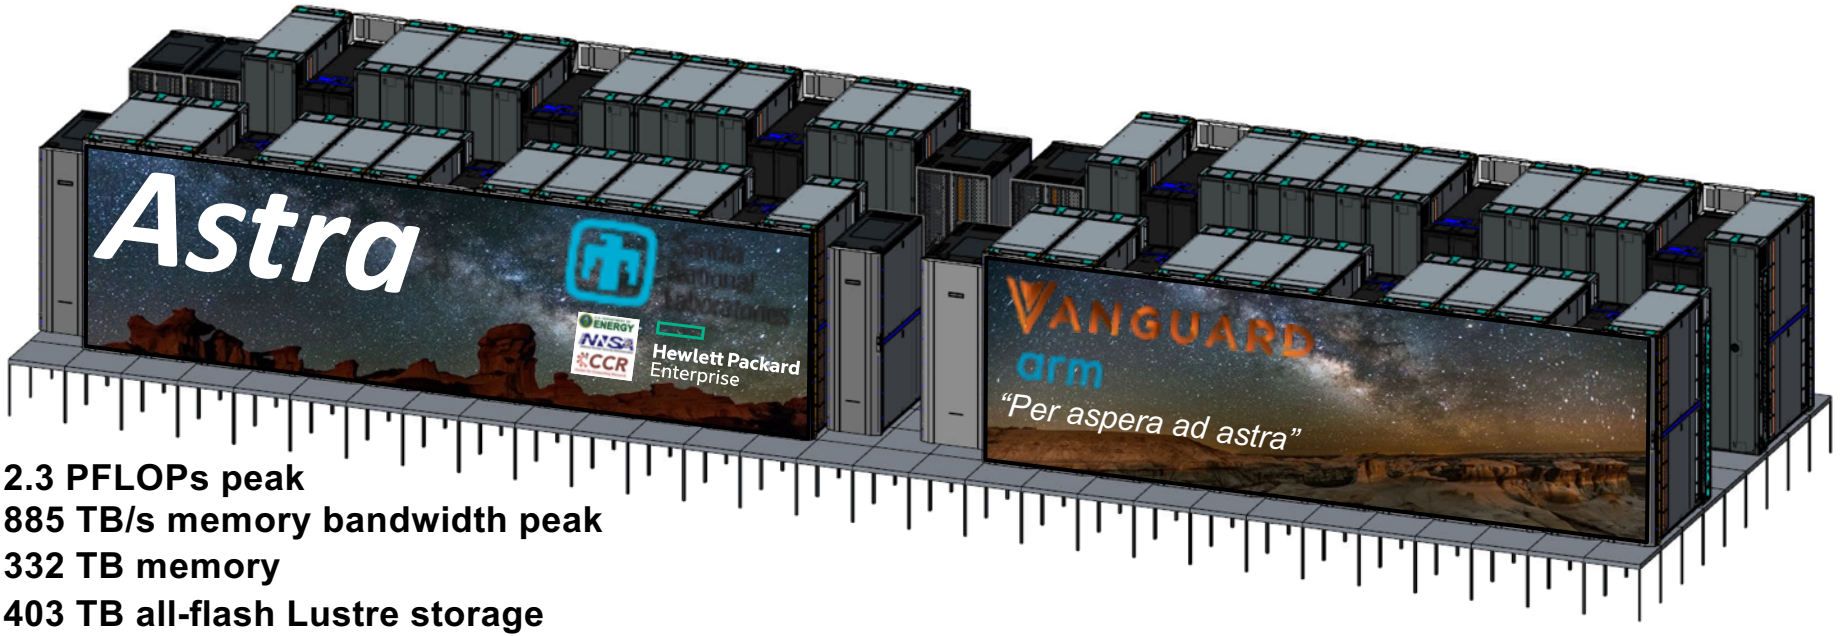

2.3 PFLOPs peak
885 TB/s memory bandwidth peak
332 TB memory
403 TB all-flash Lustre storage
1.2 MW

Demonstrate viability of ARM for U.S. DOE NNSA Supercomputing

Backup Slides

Vanguard-Astra Infrastructure

Login & Service Nodes	Four login/compilation nodes Three Lustre routers to connect to external Sandia filesystem(s) Two 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

Vanguard-Astra Compute Node

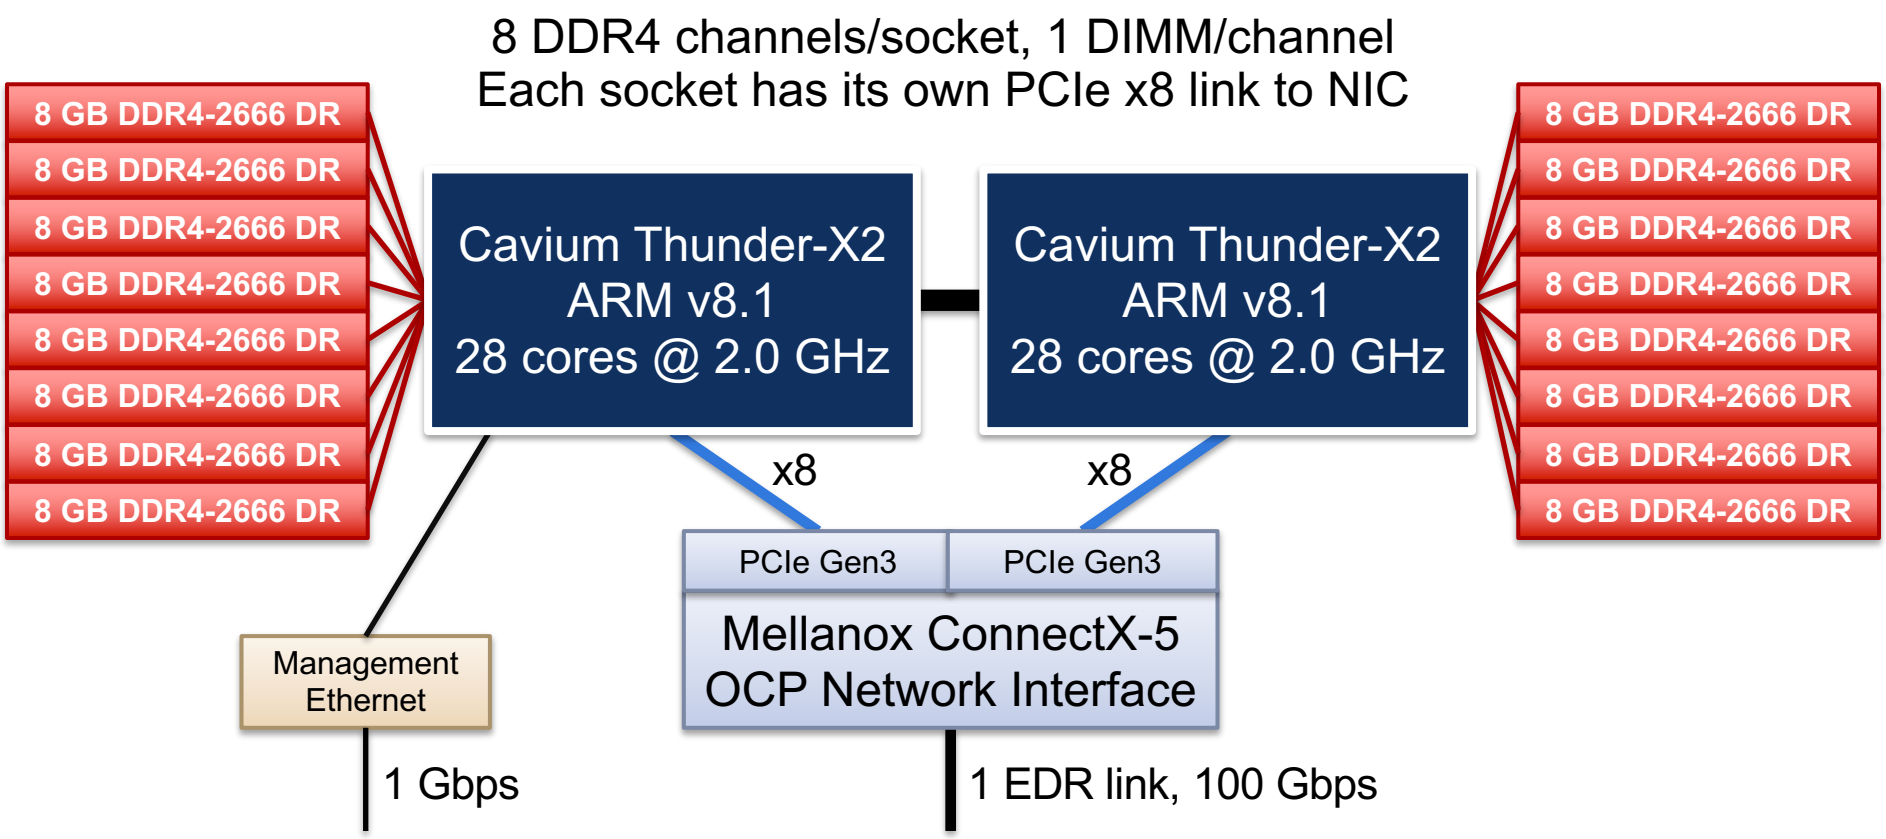

Vanguard-Astra Acceptance Plan

Sandia National Laboratories

Milestone 1
Open Science
2-3 months

Full Scale Machine Runs
- HPCG
- HPL

Micro-benchmarks
- STREAM
- Intel MPI Benchmarks

Compile and Run
- **NALU (SNL)**
- **VPIC (LANL)**
- **PF3D (LLNL)**

Milestone 2
Restricted Science
< 12 months

SSI Benchmarks
- HPCG
- HPL

Lab/Vendor Optimization
- **SPARC (SNL)**
- **PARTISn (LANL)**
- **ALE3D (LLNL)**

Compile and Run
- **RAMSES (SNL)**

Milestone 3
Classified Science
Remainder of Life

Lab/Vendor Optimization
- **SPARC (SNL)**
- **PARTISn (LANL)**
- **ALE3D (LLNL)**

Compile and Run
- **SIERRA (SNL)**

Demonstrate
- **User-specified containers and virtual machines**

Close Collaboration with HPE Open Leadership Software Stack (OLSS) Effort

Hewlett Packard Enterprise

- HPE:
  - HPE MPI (+ XPMEM)
  - HPE Cluster Manager
- Arm:
  - Arm HPC Compilers
  - Arm Math Libraries
  - Allinea Tools
- Mellanox-OFED & HPC-X
- RedHat 7.x for aarch64