SAND2018-8504C



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

Ron Brightwell

Scalable System Software

Center for Computing Research

Cetraro Conference, July 2, 2018




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

- Expand the HPC ecosystem by developing emerging, yet-to-be-proven, technologies
  - Is technology viable for future production platforms supporting ASC integrated codes?
  - Increase technology choices
- Address hardware and software technologies together
  - If hardware technology is new, gaps in software stack are certain
- Buy down risk before commitment on capability/capacity class investment

# Vanguard: Large-scale Prototype Systems

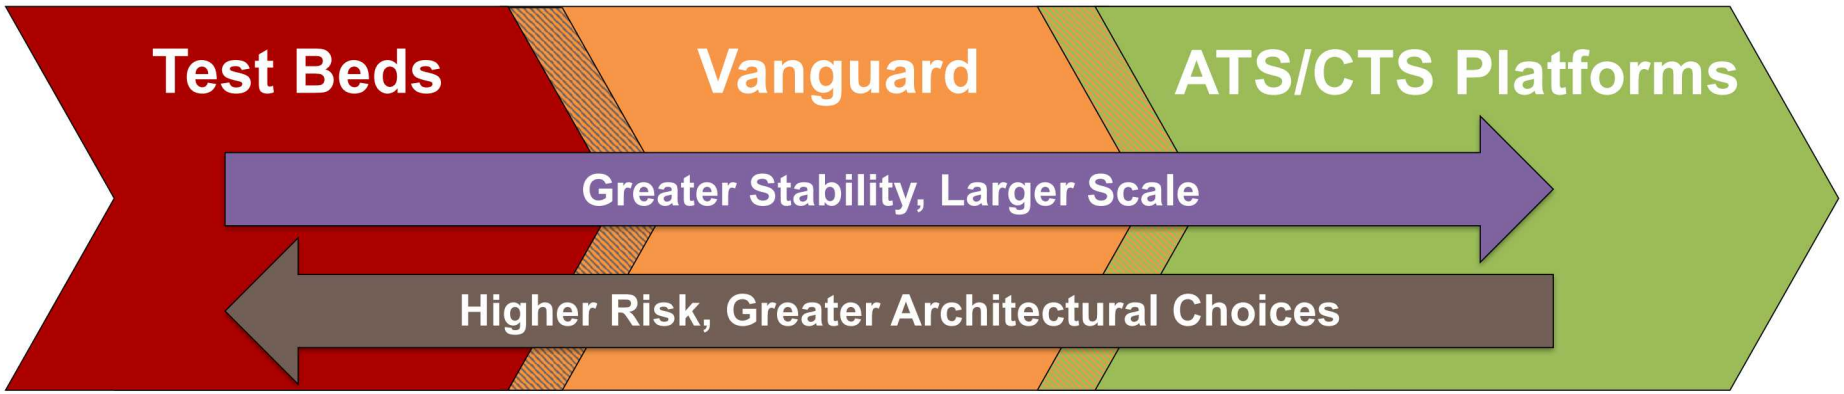


**Test Beds**

- 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**

# Advanced Prototypes Are Integrated into Future NNSA/ASC Platform Strategy

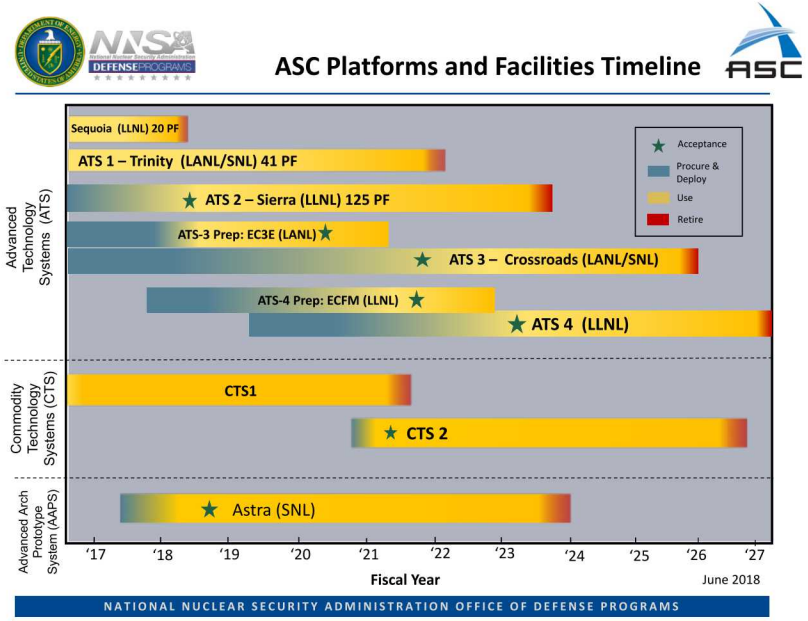


**ASC Program Targets from FY19 Implementation Plan:**

- **ATDM-2. Deploy advanced architecture prototype systems and develop accompanying tri-lab software environment.**
  - Deploy the Vanguard-Phase 1 system, named Astra, with a functioning software stack, leveraging FOUS Common Computing Environment (CCE) effort. The target date is FY19.
  - Deploy the Vanguard-Phase 2 system with a functioning software stack, leveraging FOUS CCE effort. The target date is FY21.

# Sandia's NNSA/ASC ARM Platforms

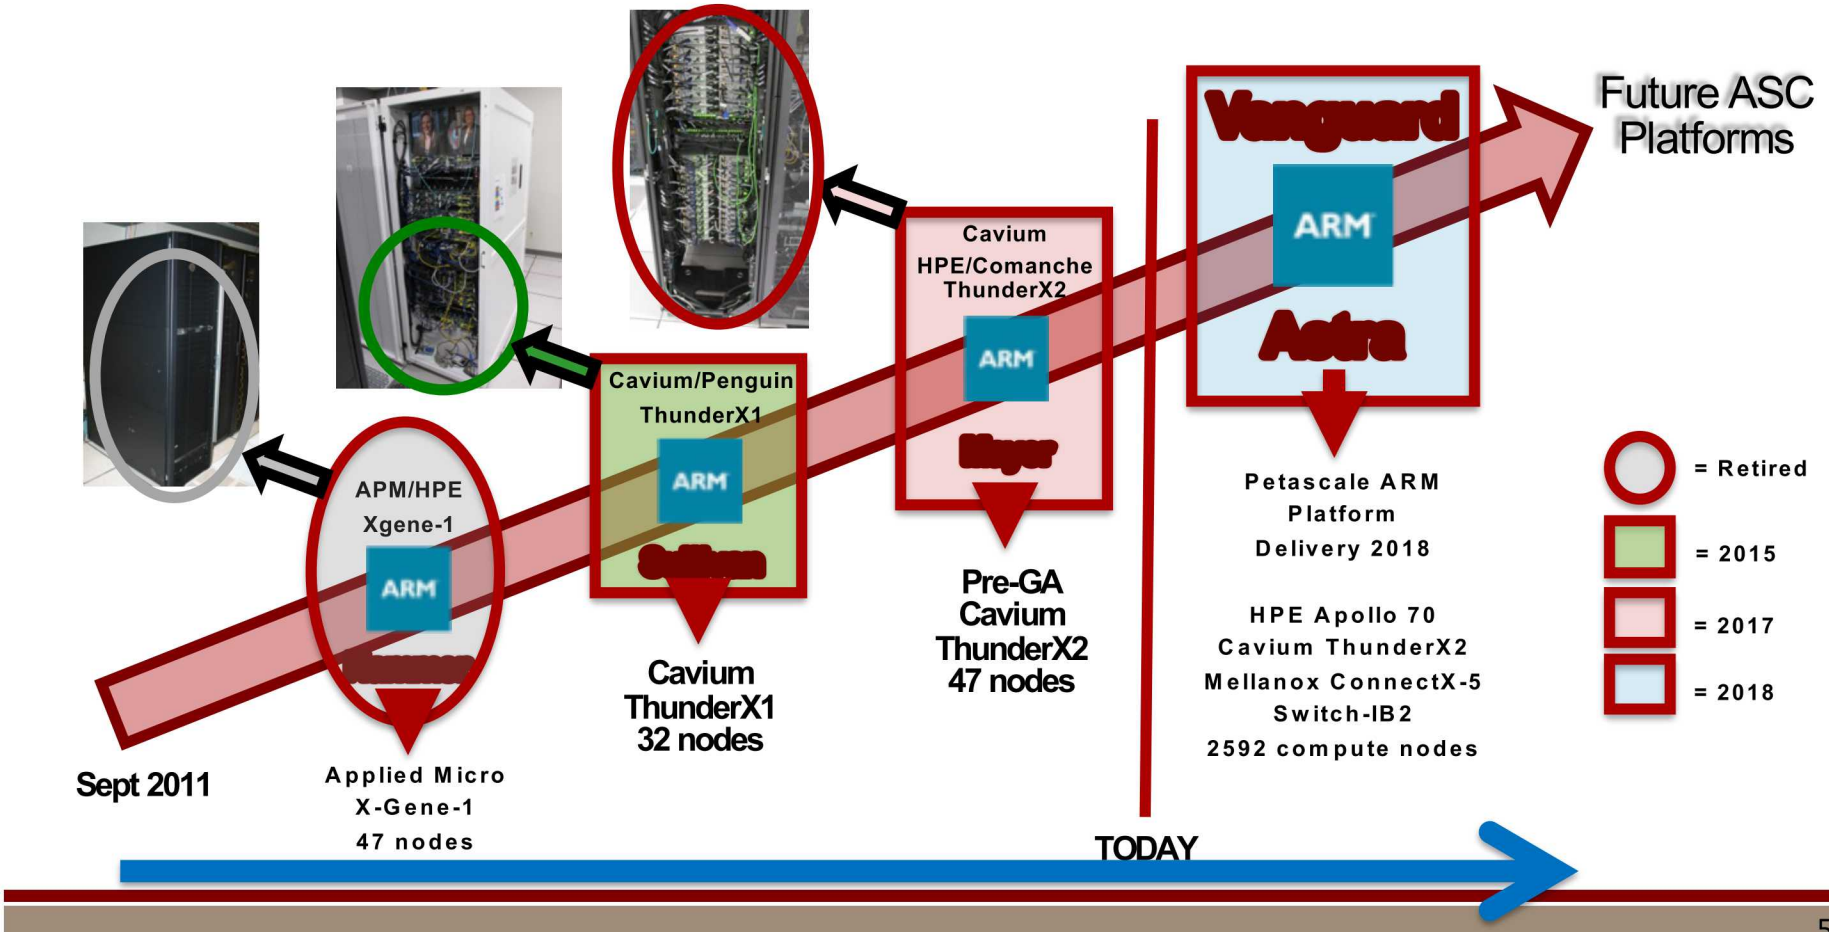

# *per aspera ad astra*

Sandia National Laboratories

through difficulties to the stars

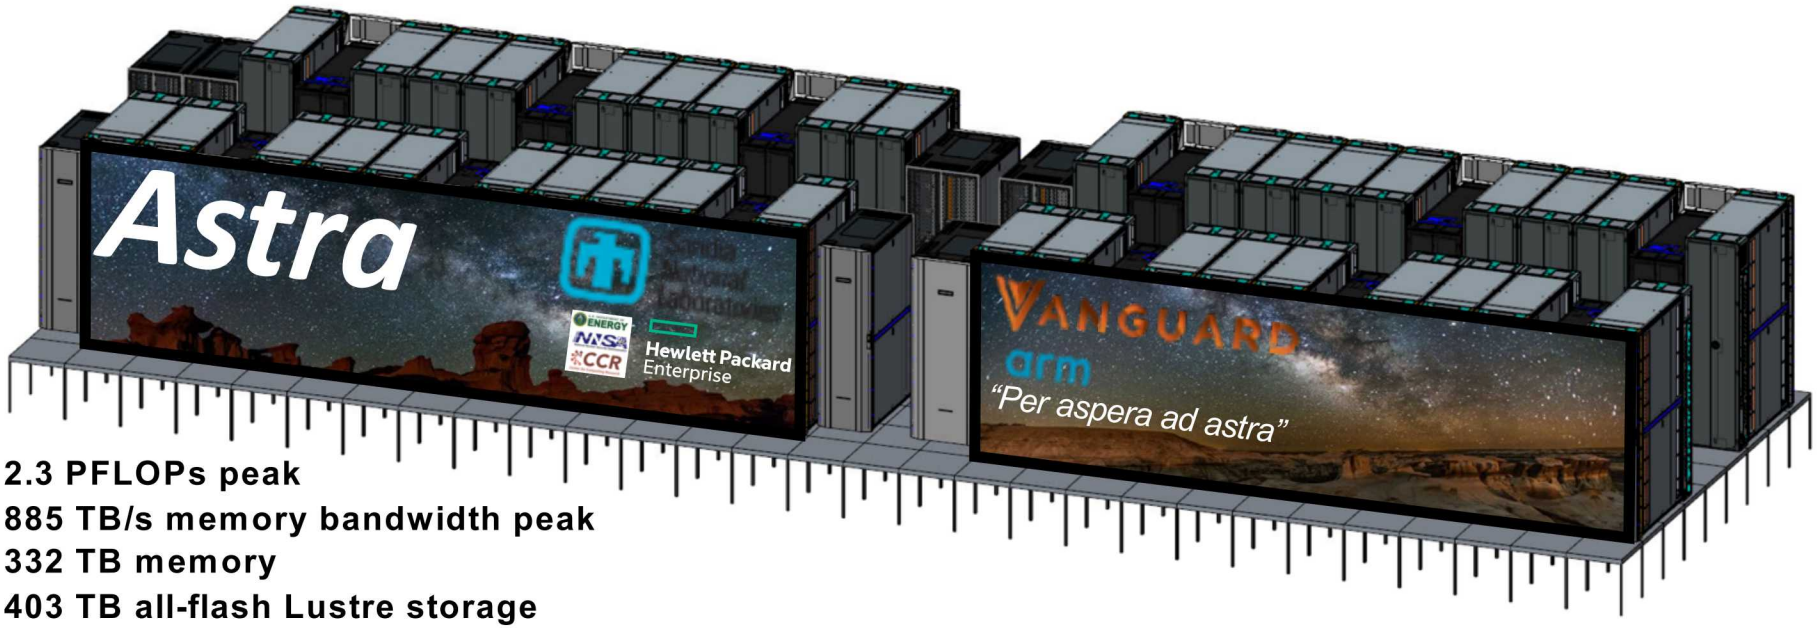


**2.3 PFLOPs peak**
**885 TB/s memory bandwidth peak**
**332 TB memory**
**403 TB all-flash Lustre storage**
**1.2 MW**

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+

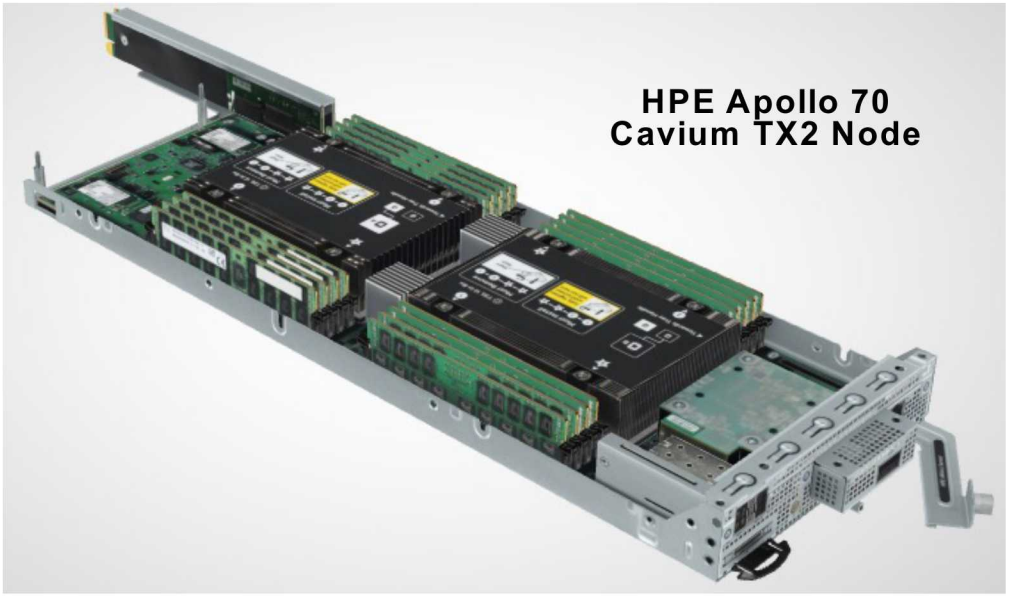

# Vanguard-Astra Compute Node

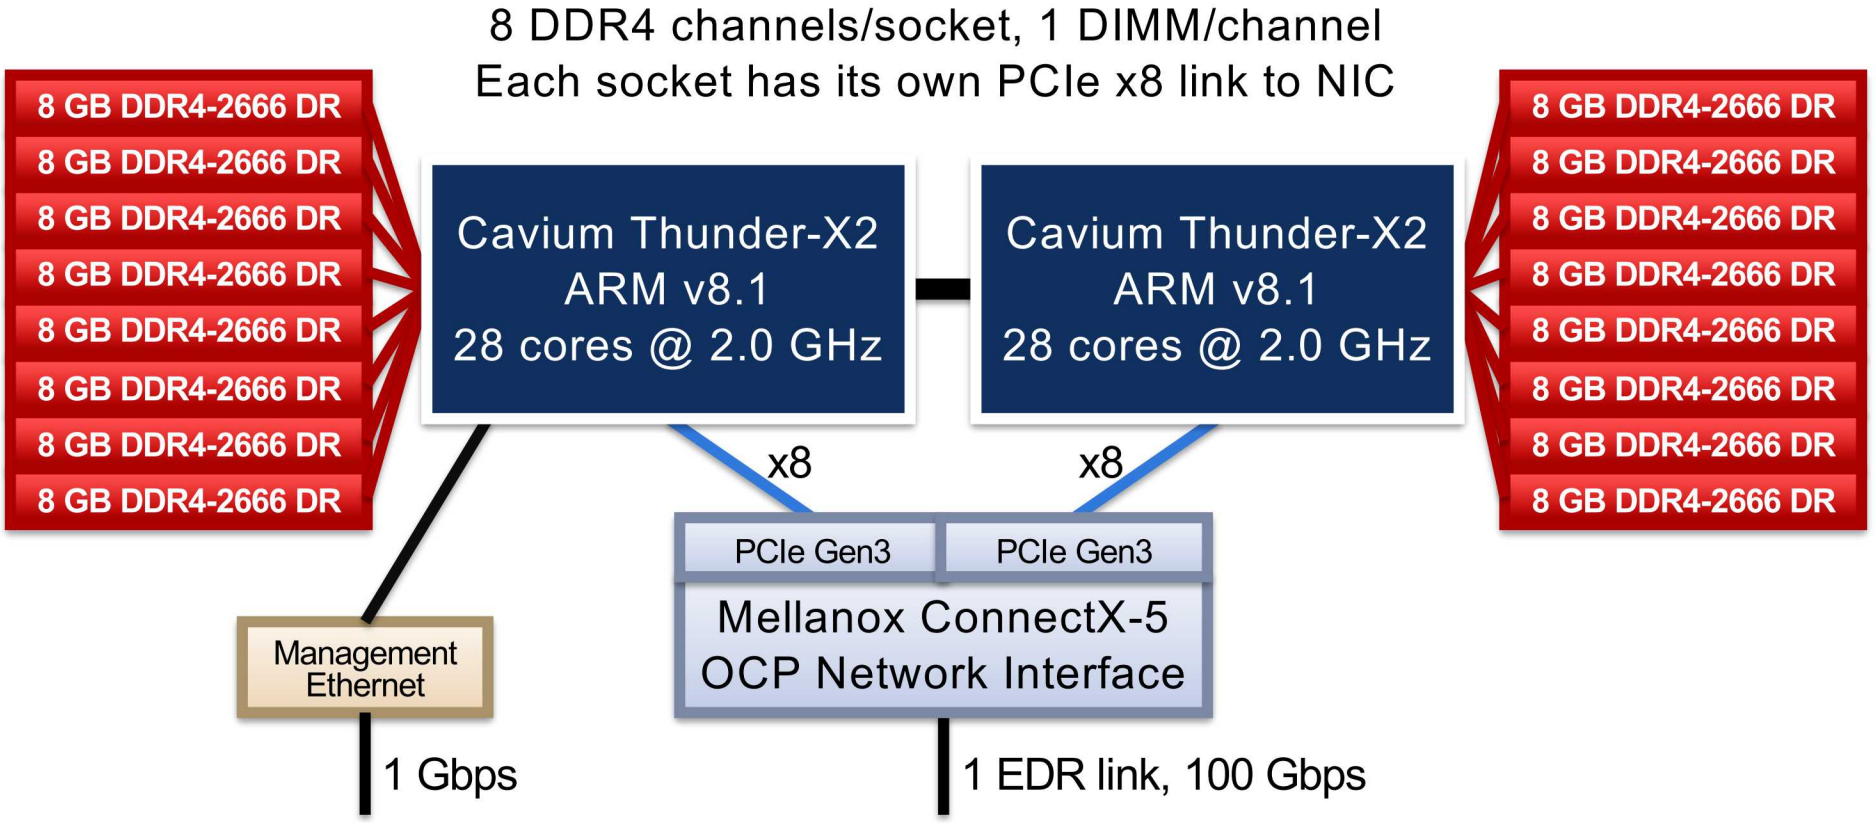

# Vanguard-Astra System Packaging

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

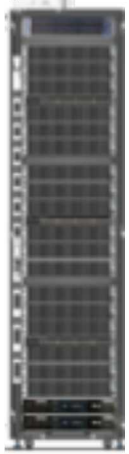

36 compute racks
(9 scalable units, each 4 racks)

2592 compute nodes
(5184 TX2 processors)

3 IB spine switches
(each 540-port)

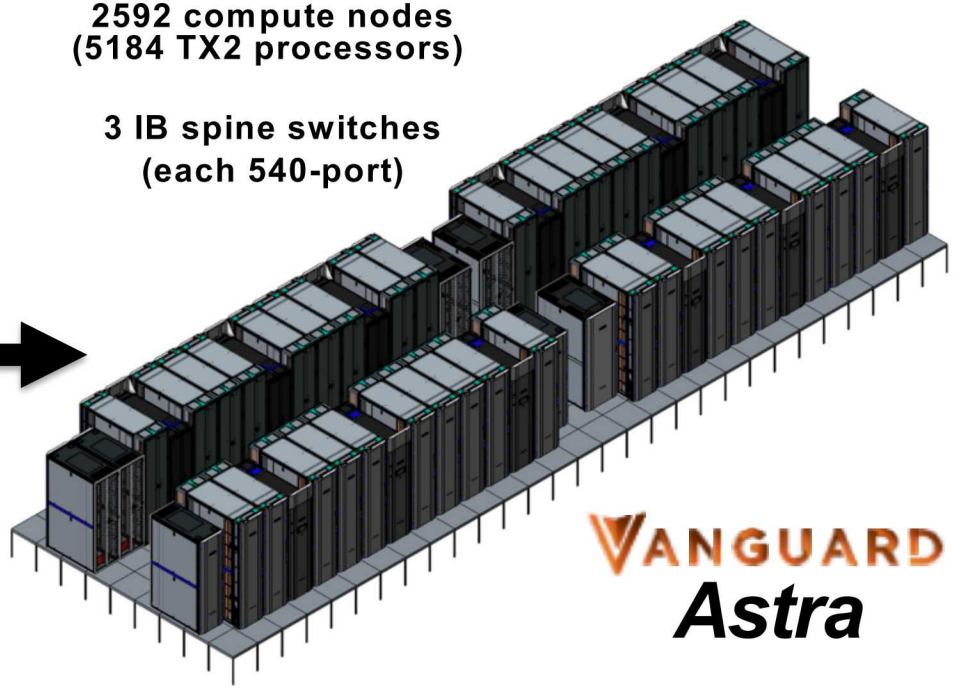

VANGUARD
Astra

# Network Topology Visualization

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

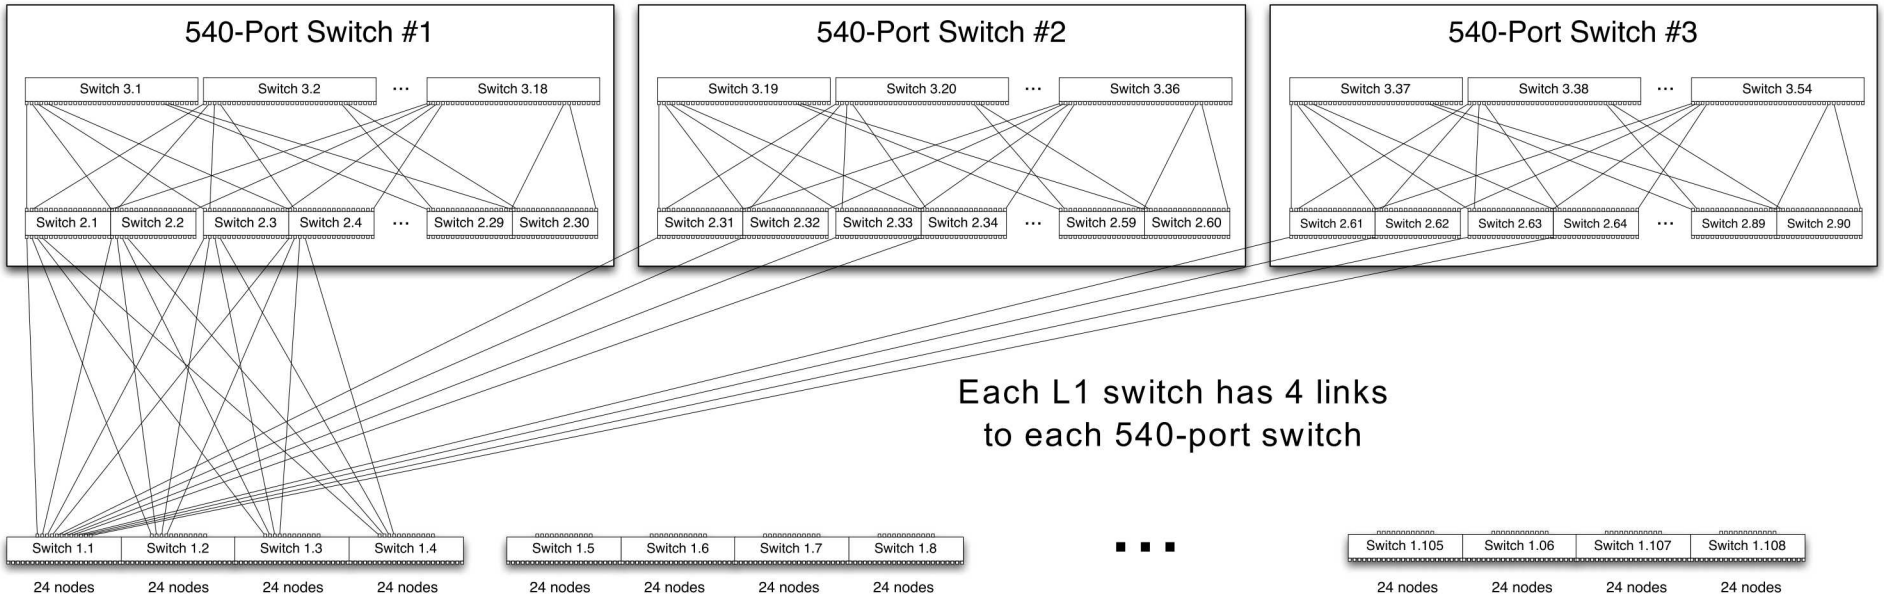


108 L1 switches * 24 nodes/switch = 2592 compute nodes

# 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-ic.html

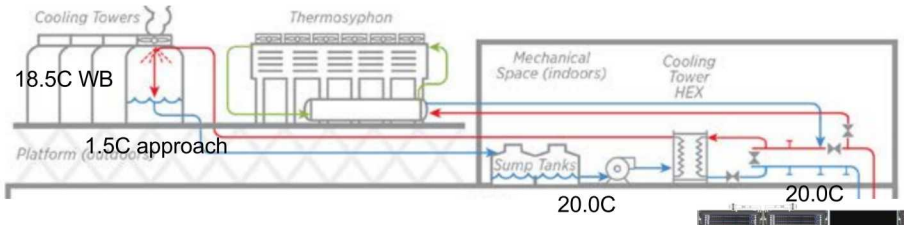


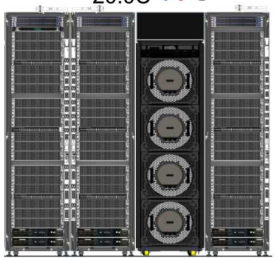

| Projected power of the system by component | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | per constituent rack type (W) | | | | | total (kW) | | | |
| | wall | peak | nominal (linpack) | idle | racks | wall | peak | nominal (linpack) | idle |
| Node racks | 39888 | 35993 | 33805 | 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 |

# Expanded Computing Facility

- Institutionally Funded
- Design/Build Contract Awarded
- All Permits Received – Site Preparation underway
- Groundbreaking Event 9/28/17
- 40% and 90% Design Reviews Completed
- 100% Design Review by 12/8/17
- Completion Data 7/15/18
- Will feature 90% liquid cooling 10% air cooling
- Thermosyphons & Air-Side Economization for Water/Energy Savings
- Solar Farm for LEED Certification
- Non-load-bearing, movable west wall for expansion (14,000 – 20,000sf)
- 7 MW power expandable to 15 MW

# Cavium Arm64 Providing Best-of-Class Memory Bandwidth

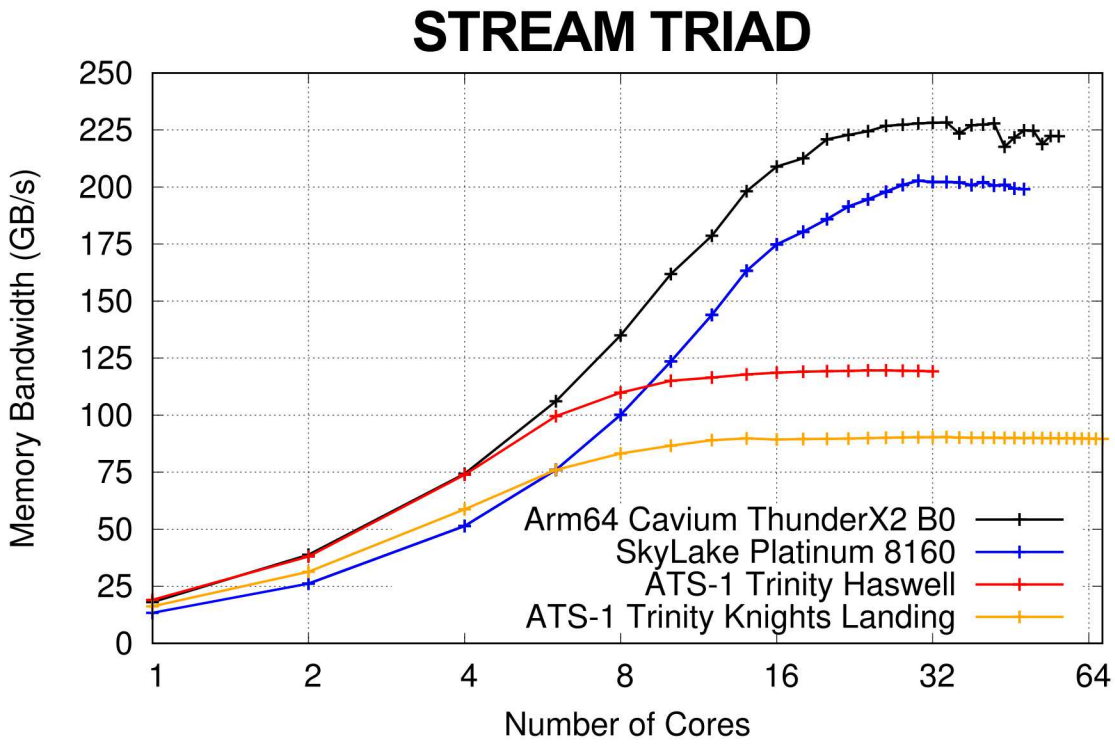


TX2 DDR4-2400

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

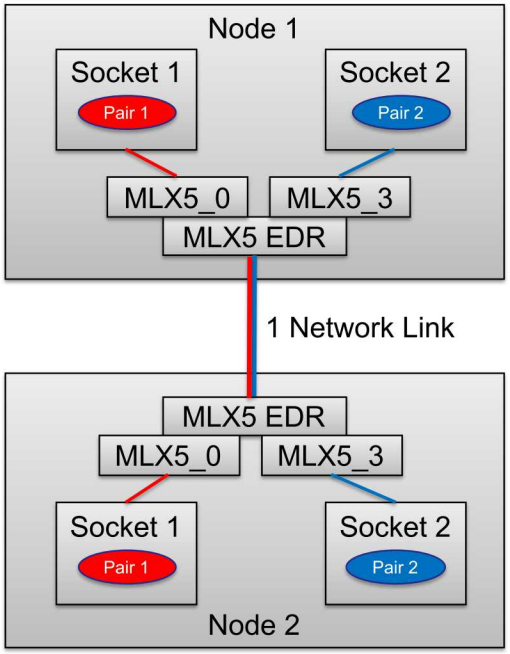


## OSU MPI Multi-Network Bandwidth

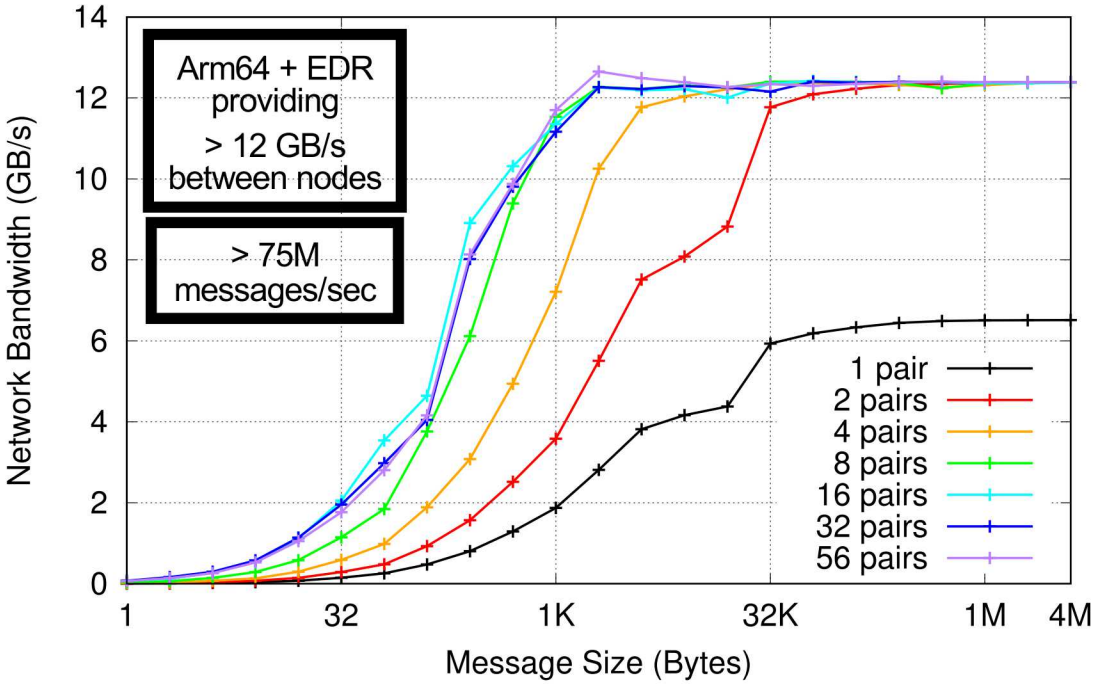

# Mini-App Performance on Cavium ThunderX2

- 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

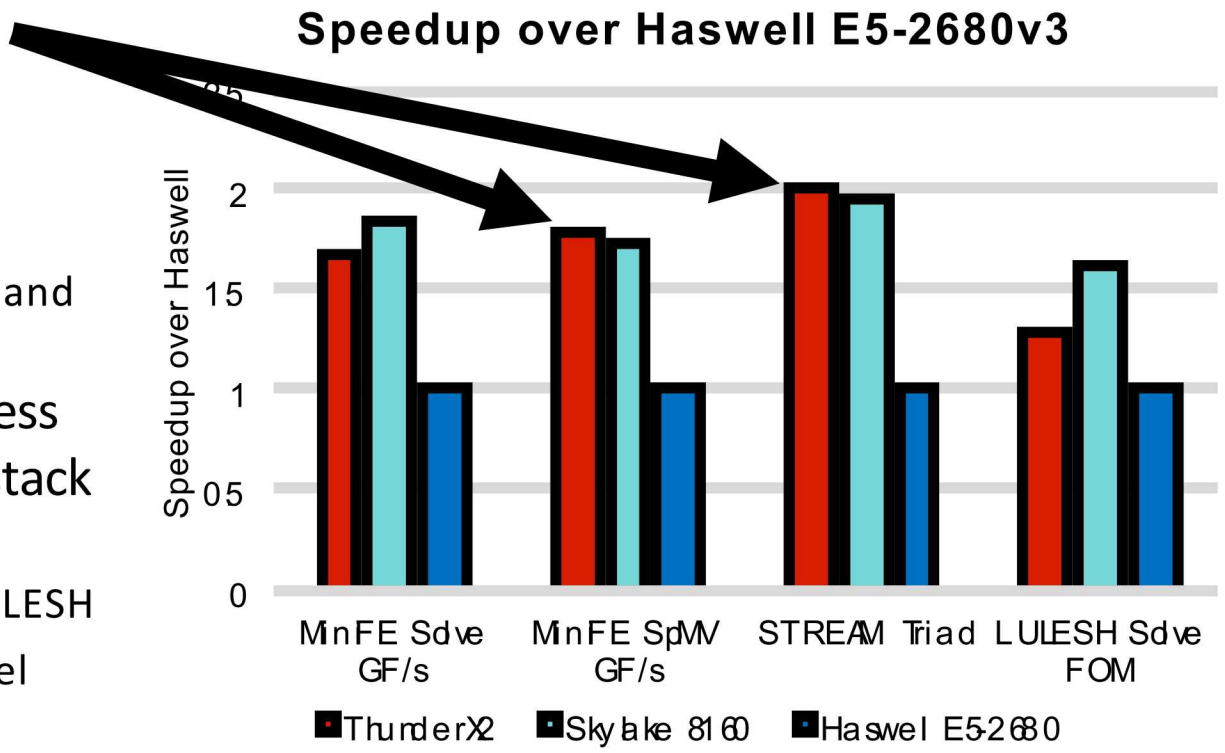

# Vanguard-Astra Schedule – Past and Projected

- ✓ September 22nd – 2nd Draft RFI released
  - Note: We are releasing the RFI via the procurement vehicle more frequently, possibly weekly, to maintain equality in the communication process.
- ✓ Week of September 25th – Prime F2F presentations
- ✓ RFP release no later than January 12th
  - Possibly before Christmas holidays ACTUAL:RFP released December 20th
- ✓ RFP responses due no later than February 8th
  - ✓ Hopefully worse case scenario ACTUAL: Vendors now have ample time, extensions won't be entertained.
- ✓ January 11th - Vendor pre-proposal brief at Sandia NM (9-12 CSRI/90)
- ✓ RFP responses distributed to technical team members February 9th (Actual: February 8th)
- ✓ February 20th – Opportunity for groups to meet Face to Face (up to group lead and teams)
- ✓ February 21st – Technical review (SNL Albuquerque)
- ✓ February 21st – Source Selection (with Tri-lab members to minimize travel)
  - ✓ February 22nd – Follow on Source Selection only if necessary
- ✓ February/March 2018 – Negotiations and SOW development
- ✓ April/May 2018 – SOW development and contract placement
- **July 2018 – Facility (725-E) completed**
- **August 2018 – Astra hardware delivery begins**
- **September 2018 – Astra hardware delivery completed**

# Astra Acceptance and Milestones

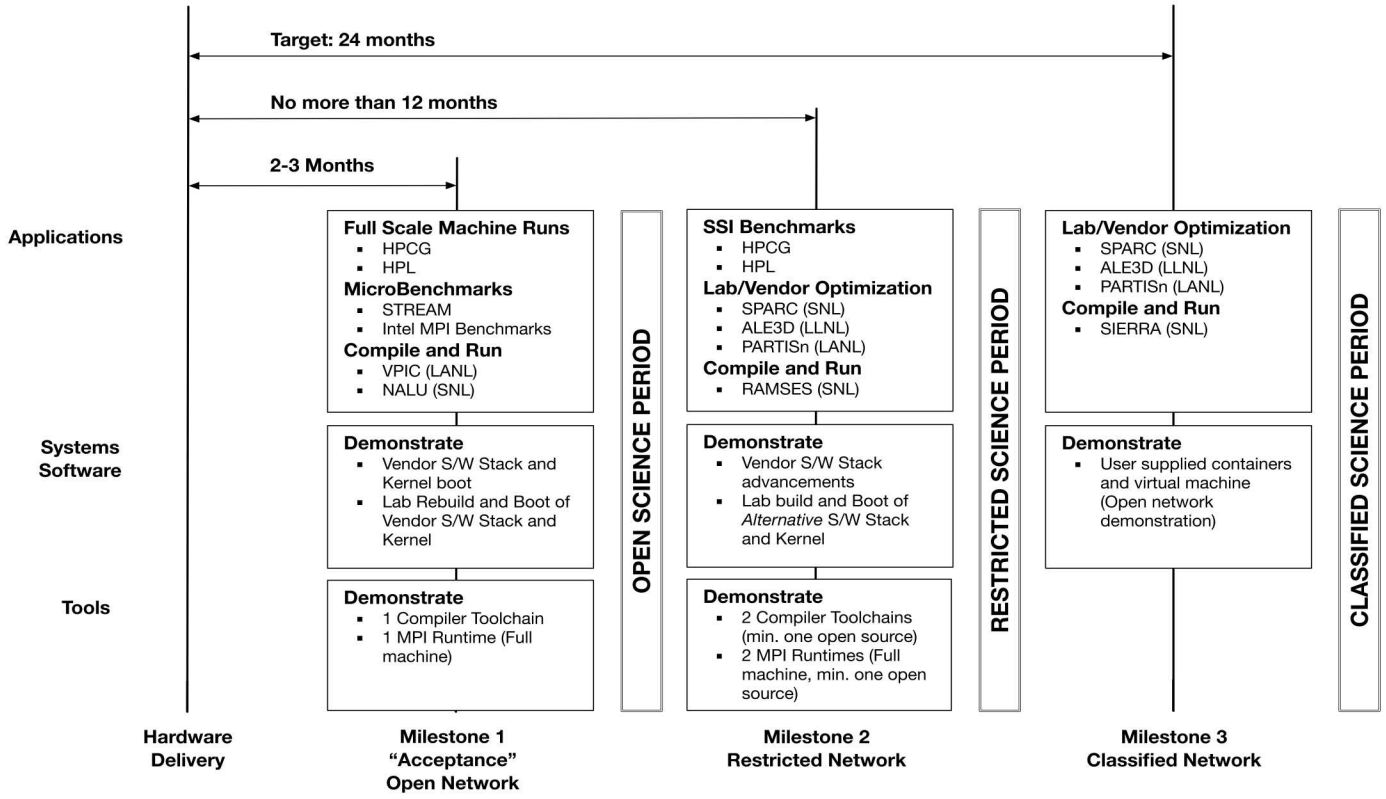

# Vanguard Tri-Lab Software Effort

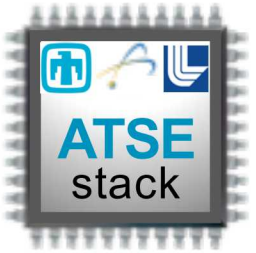


- Accelerate maturity of ARM ecosystem for ASC computing
  - Prove viability for NNSA integrated codes running at scale
  - Harden compilers, math libraries, tools, communication libraries
    - Heavily templated C++, Fortran 2003/2008, Gigabyte+ binaries, long compiles
  - Optimize performance, verify expected results
- Build integrated software stack
  - Programming env (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... Astra arrival to useful work done

# Advanced Tri-lab Software Environment High-level Goals

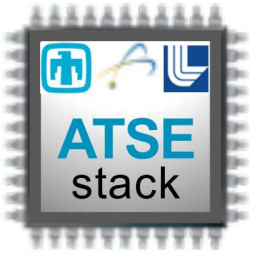


- Build an open, modular, extensible, community-engaged, and vendor-adaptable ecosystem
- Prototype new technologies that may improve the DOE ASC computing environment (e.g., ML frameworks, containers, VMs, OS optimizations)
- Leverage existing efforts such as Tri-lab OS (TOSS), programing environments,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

# Vanguard-Astra Software Stack

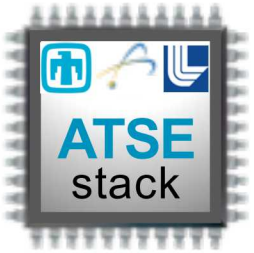


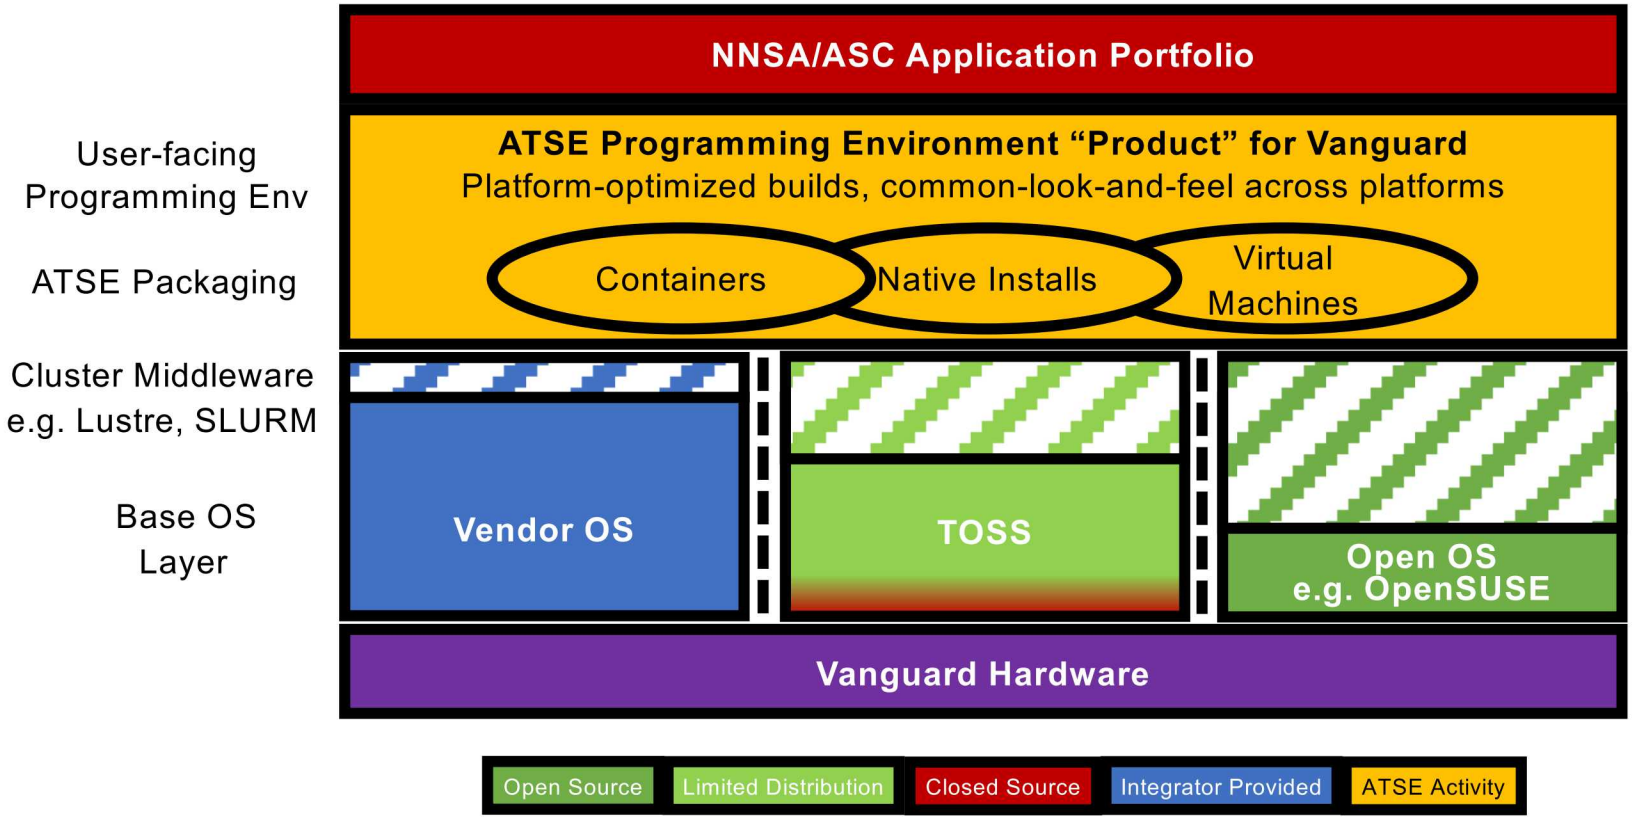

# Integrate Components from Many Sources

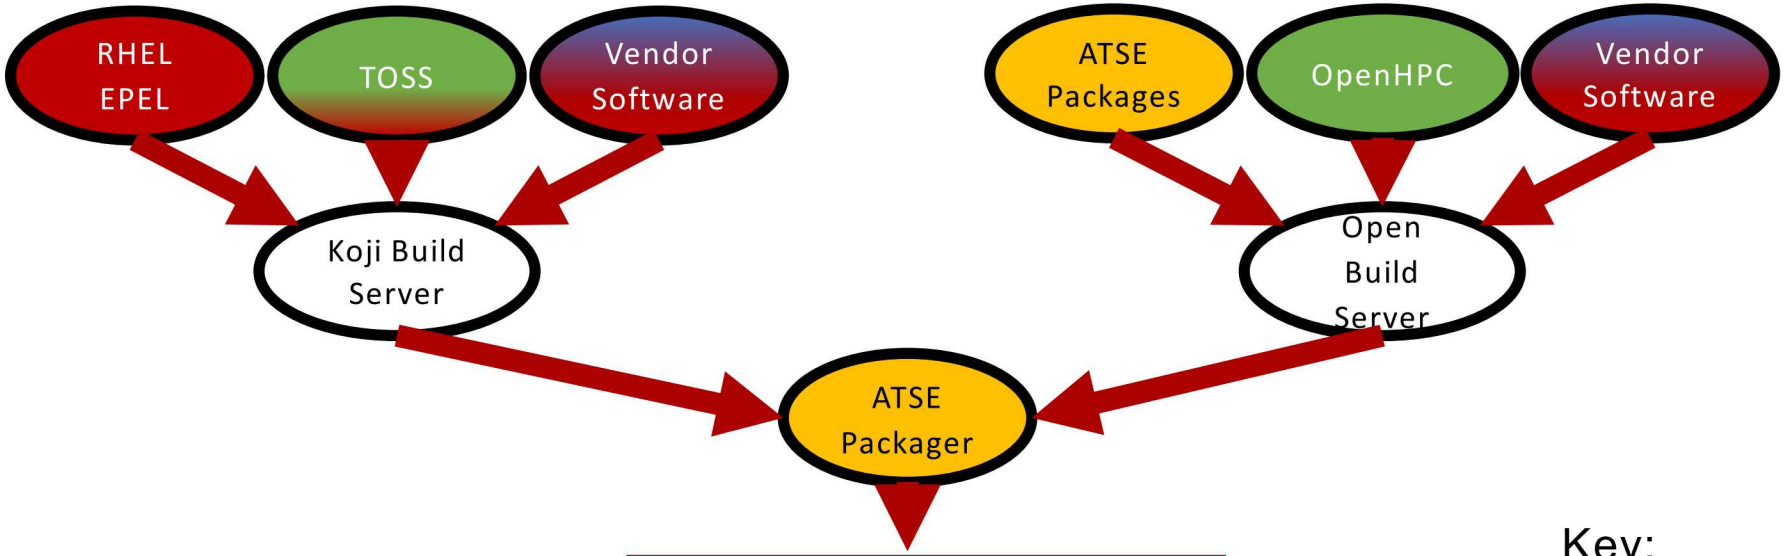


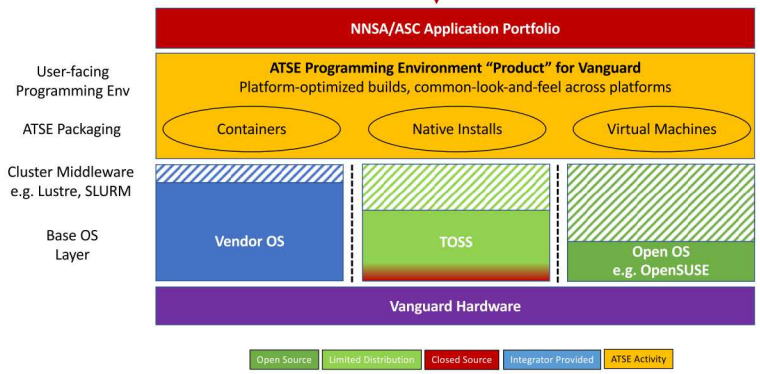


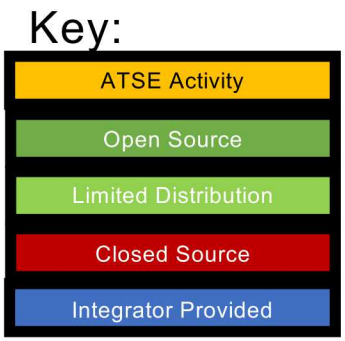

# Close Collaboration with HPE Open Leadership Software Stack (OLSS) Effort

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

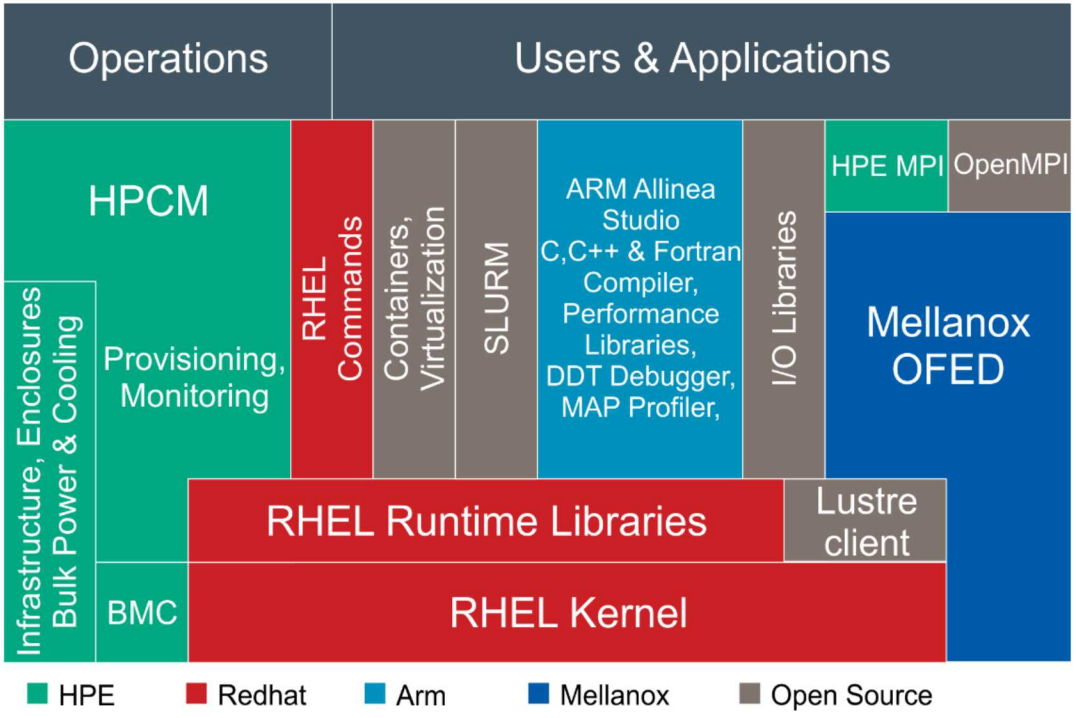

# Early Application Porting

| Workload | GCC 7.2.0 | Arm HPC Compilers |
|---|---|---|
| LAMMPS | | |
| SPARTA | | |
| SPARC | | |
| NALU | | |
| CTH | | FORTRAN issue |
| Drekar | | |
| Xyce-UUR | | |
| VPIC | | |
| SNAP | | |

*Most codes build without trouble, optimization work remains*

**Placing collaborative vendor contracts to harden Arm64 compilers, math libraries, and tools – both for Astra and Arm ecosystem in general**

# 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

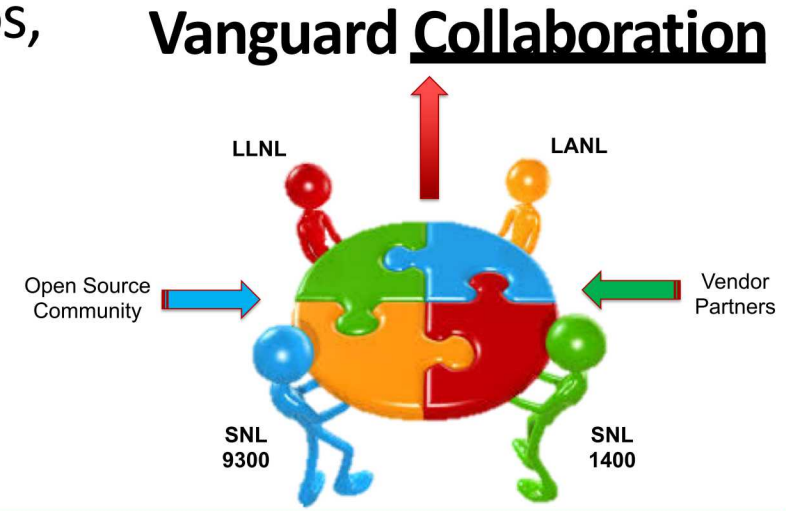

# Acknowledgments


- Kevin Pedretti
- Si Hammond
- Jim Laros

Sandia National Laboratories

# Backup Slides

# TOSS – Tri-lab Operating System Stack

- Targets commodity technology systems at NNSA Tri-labs
  (Lead: Livermore, Los Alamos, Sandia)
  - RHEL7 based, supports x86_64, ppc64le, and aarch64 from single source
  - ~4K packages on all archs, 200+ built for TOSS by LLNL (compilers, MPI, ...)
  - Baseline not optimized for particular system, labs optimize
  - Partnership with RHEL to add support for new hardware pre-GA
- Concerns
  - Distribution of TOSS restricted due to licensing, vendors-only for lab use
  - Not focused on user-facing programming environment
  - Missing on ARM: security scanners, backup tools, firmware tools, third party software

# OpenHPC


- Targets HPC Linux clusters
  - Community effort
  - Common ingredients needed to deploy and manage an HPC cluster
  - Goal to enhance modularity and interchangeability of key components
  - Current release 1.3.3, builds on Centos 7.4 or SLES12, arm64 + x86_64
- Concerns
  - Lack of integration
  - Small cluster focused, lack of hierarchy needed for scalability
  - Ability to do optimized builds

# Exascale Computing Project Software, Support on ARM

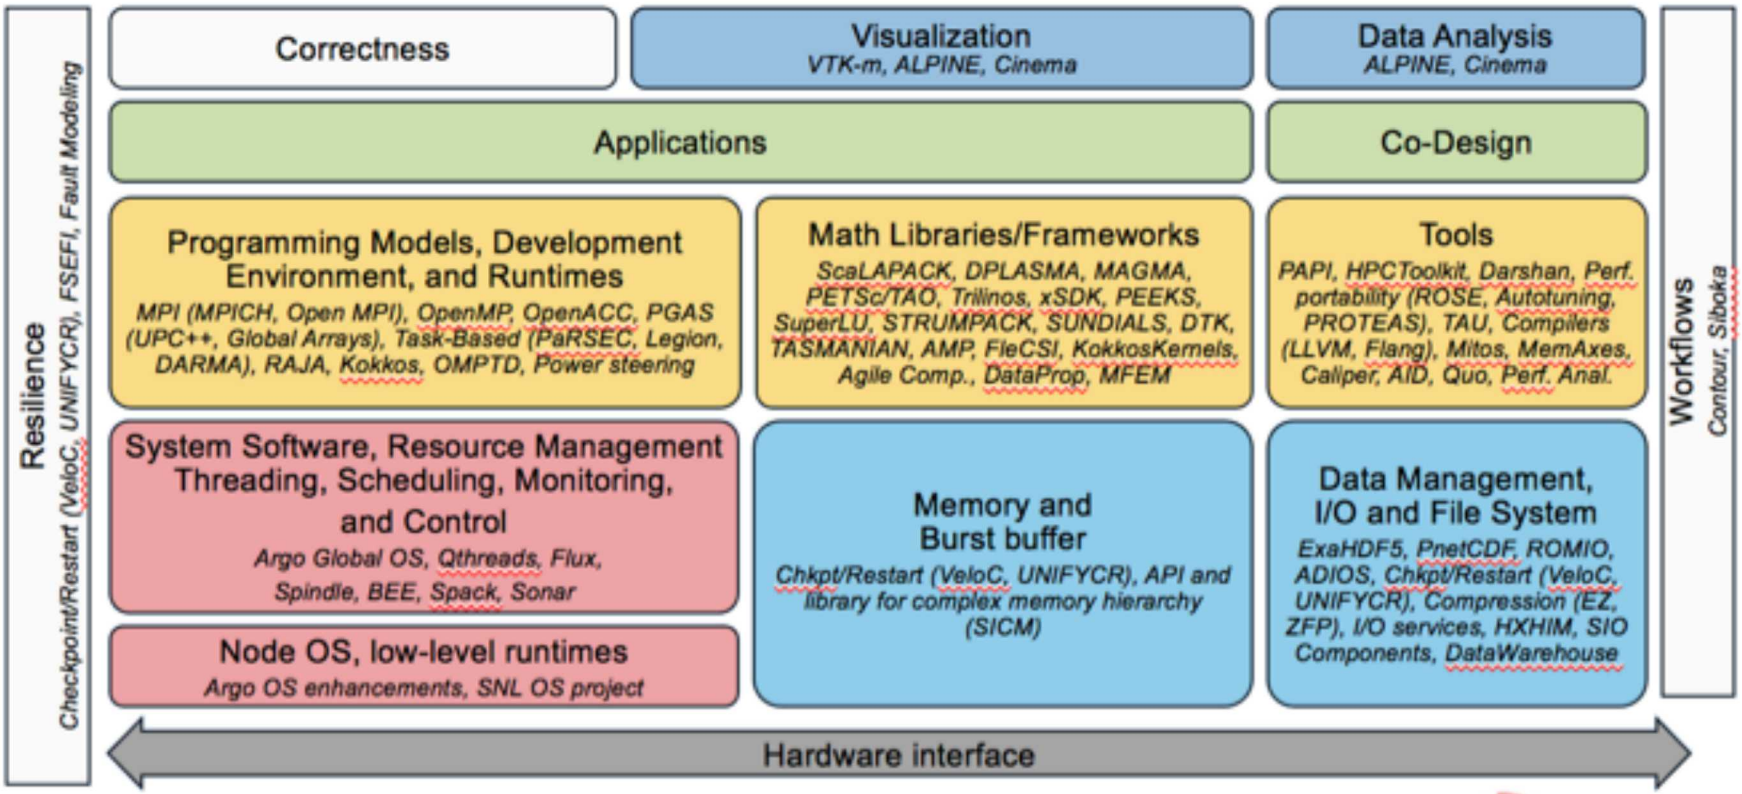

# Vanguard-Astra Infrastructure

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

# 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