

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



VANGUARD

PRESENTED BY Andrew J. Younge

PIs: James H. Laros III, Kevin Pedretti, Si Hammond

Sandia National Laboratories



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.

Outline

- Overview of Vanguard
 - Prototype HPC Architectures
- Astra – Petascale ARM platform
- ATSE – Advanced Tri-lab Software Environment
- R&D Activities
- Conclusion




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 a system viable for future ATS/CTS platforms?
 - Increase architecture AND integrator choices
- Buy down risk and increase technology and vendor choices
 - Ability to accept higher risk allows for more/faster advancement
 - Lowers/eliminates mission risk and initialreduces investment
- Jointly address hardware and software challenges

Phase I: Expanding the ARM Ecosystem



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's NNSA/ASC ARM Platforms



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

7 | 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

8 GB DDR4-2666 DR
8 GB DDR4-2666 DR

Cavium Thunder-X2
ARM v8.1
28 cores @ 2.0 GHz

Cavium Thunder-X2
ARM v8.1
28 cores @ 2.0 GHz

8 GB DDR4-2666 DR
8 GB DDR4-2666 DR

x8

x8

PCIe Gen3

PCIe Gen3

Mellanox ConnectX-5
OCP Network Interface

Management
Ethernet
1 Gbps

1 Gbps

1 EDR link, 100 Gbps

Vanguard-Astra System Packaging



HPE Apollo 70 Chassis: 4 nodes



HPE Apollo 70 Rack



18 chassis/rack


72 nodes/rack

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

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

2592 compute nodes
(5184 TX2 processors)

3 IB spine switches
(each 540-port)




Hewlett Packard
Enterprise

VANGUARD
Astra

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





$108 \text{ L1 switches} * 24 \text{ nodes/switch} = 2592 \text{ compute nodes}$

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)				racks	total (kW)				
	wall	peak	nominal (linpack)	idle		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



STREAM TRIAD




TX2 DDR4-2400
SkyLake 8160

Trinity Haswell
Trinity KNL DDR

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






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 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
 - Experimental system software






ARM Tri-lab Software Environment (ATSE)





Open Source Limited Distribution Closed Source Integrator Provided ATSE Activity

Integrate Components from Many Sources



Acceptance Plan – Maturing the Stack



ATSE Deployed Beta Stack



- Setup local Open Build Service (OBS) build farm at Sandia
- Built set of software packages needed for Astra milestone 1
 - When OpenHPC recipe was available, we tried to use it, modifying as necessary
 - Otherwise, we built a new build recipe in same style as OpenHPC
- Installed on Arm testbed at Sandia, now using as the default environment
- Tested with STREAM, HPL, HPCG, ASC mini-apps
- Compiler toolchain support
 - GNU compilers, 7.2.0
 - ARM HPC compilers

ATSE Modules Interface, Mirrors OpenHPC


```
[ktpedre@mayer2 ~]$ module avail
----- /opt/atse/pub/moduledeps/gnu7-openmpi3 -----
phdf5/1.10.2    pnetcdf/1.9.0
----- /opt/atse/pub/moduledeps/gnu7 -----
hdf5/1.10.2    openblas/0.2.20    openmpi3/3.1.1 (L)
----- /opt/atse/pub/modulefiles -----
arm/18.3        binutils/2.30 (L)  gnu7/7.2.0 (L)    pmix/2.1.1        spack/0.11.2
atse (L)       cmake/3.11.1 (L)  hwloc/1.11.10    prun/1.2          zlib/1.2.11
autotools (L)  git/2.18.0 (L)   numactl/2.0.12   singularity/2.5.2

Where:
L: Module is loaded

Use "module spider" to find all possible modules.
Use "module keyword key1 key2 ..." to search for all possible modules matching any of the "keys".
```

R&D Areas

- Enable machine learning & novel computing environments
 - Provide greater flexibility when default stack doesn't fit
 - Enable evaluation of emerging new computing environments
 - Provide DevOps capabilities when applicable
- Task: Investigate containers and virtual machines
 - ARMv8.1 includes new virtualization extensions, SR-IOV
 - Singularity for full container solution for HPC, initial Docker on ARM
- 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 at scale
 - Exploring kernel spectrum: stock Linux distro to HPC-tuned kernels to non-Linux lightweight kernels and multi-kernels
 - Arm-specific optimizations



Conclusion



- Vanguard program allows us to take necessary risks to ensure a healthy HPC ecosystem
 - Increase technology choices
 - Prove ability to run multi-physics production applications at scale
 - Arm is our first priority
- Tri-lab software environment to mature ARM for ASC computing
 - Harden compilers, math libs, and tools
 - Optimize performance, verify expected results
 - Increase modularity and openness of software stack
 - Support traditional HPC with emerging AI + ML workloads

Scale will drive future R&D efforts and demand in ARM HPC Ecosystem



VANGUARD



Exceptional Service in the National Interest



Questions?

`ajyoung@sandia.gov`