SAND2021-14649C

Sandia National Laboratories

Exceptional service in the national interest

SANDIA'S EXPERIENCES WITH ARM

ARM BOF @ SC21

PRESENTED BY

ANDREW J. YOUNGE, SI HAMMOND,

KEVIN PEDRETTI, JAMES LAROS III

SAND2021-7786 C

UNCLASSFIED UNLIMITED RELEASE

WHY SPEND TIME WITH ARM?

- The previous decade has seen Intel and AMD dominate in the HPC and server market
- Important for DOE/NNSA to have several options when selecting HPC systems
  - Tailor to our specific needs
  - Help promote effective competition
  - Benefit from broader vendor engagement and ecosystem development

- Sandia HPC teams using Arm since 2012
  - Vanguard Astra prototype system
  - New testbeds with A64FX and NVIDIA GPUs

- Several timelines to consider in our work with **Arm** systems:
  - **Short Term** – more options for system deployments
  - **Medium Term** – help select HPC specific IP & SKUs by customizing cores, memory systems
  - **Longer Term** – ability to influence and optimize for HPC-appropriate cores and other IP components

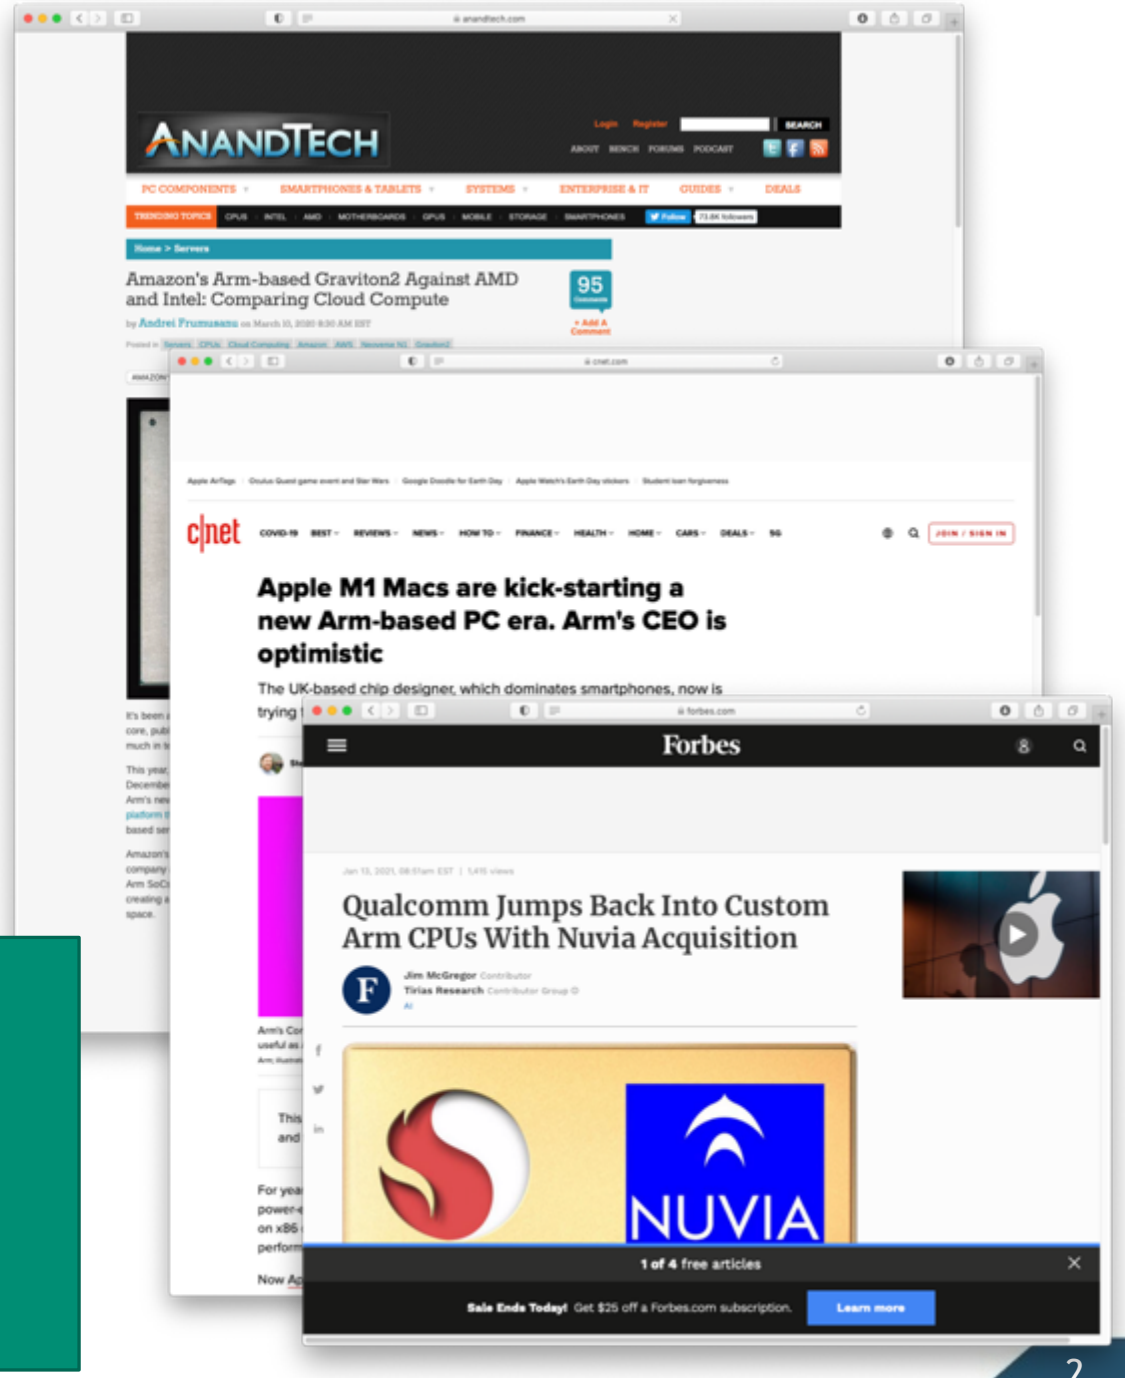

ASTRA AT A GLANCE – FIRST PETASCALE ARM

- Deployed in 2018
- 2,592 HPE Apollo 70 compute nodes
  - 5,184 CPUs, 145,152 cores, 2.3 PFLOPS (peak)
  - 36 compute racks arranged in two rows
  - Full system power ~ 1.2 MW
- Marvell ThunderX2 ARM SoC, 28 core, 2.0 GHz
  - 2 sockets per node
- Memory per node: 128 GB
  - 16 channels of 8 GB DDR4-2666 DIMMs
  - Aggregate capacity: 332 TB, 885 TB/s (peak)
    - 247 GB/s per node STREAM TRIAD
- Mellanox IB EDR, ConnectX-5
  - 112 36-port leaf, 3 648-port spine switches
  - SocketDirect capability
- ATSE software stack
  - Extensible user programming environment
  - TOSS Base Operating system
- HPE Apollo 4520 All–flash Lustre storage
  - Storage Capacity: 990 TB (usable)
  - Storage Bandwidth: 250 GB/s
    - 400 GB/s stunt mode, 432 GB/s peak

More details in SC20 paper – [Chronicles of Astra: Challenges and Lessons]

SOFTWARE ECOSYSTEM IS KEY

- Need robust software ecosystem to meet NNSA computing requirements
- ATSE is a ecosystem enablement and collaboration mechanism
  - HPE, Arm, Red Hat, OpenHPC, Linaro, and other DOE labs
  - Full programming environment for codes
  - Enabling vendor innovation when available
  - Leverage open when possible
- Invested directly in the Arm HPC compiler toolchain
  - Seen significant performance improvements
  - Ongoing collaboration with Arm
- Enabled containerized DevOps model
- Ensured apps can utilize full system scale
- Extending ATSE to NVIDIA and A64FX (and x86)

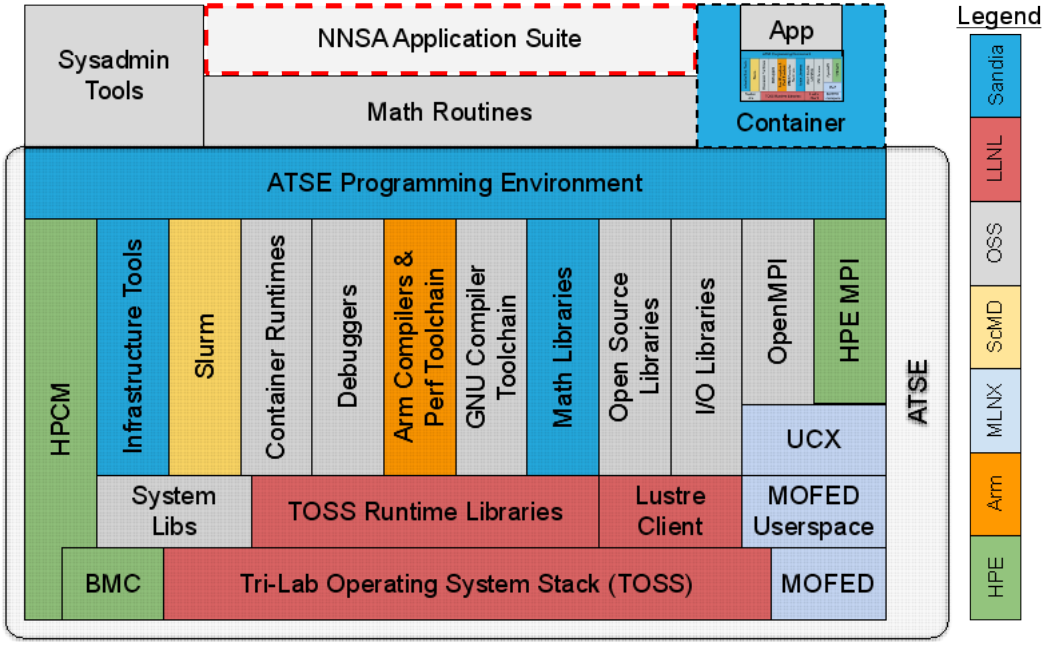

ASTRA TODAY

- Astra is in production and has been well utilized since release to users
  - Many production users are finding the scale and turnaround on the system appealing
  - Running mission workloads

- Plan to run Astra as a production resource for the rest of its hardware lifetime
- NNSA ATDM L1 code milestone results (2020) helped to drive fixes and improvements
  - And to show code teams the system really was a viable platform for their codes
  - Helped to get configurations for Kokkos, Trilinos, etc ready for production users

OVERVIEW

So is Arm ready for production use cases?

- ✓ Availability of production-class compilers (that can compile production applications and pass relevant tests for correctness, functionality, etc)
- ✓ Broad ecosystem to support HPC development (availability of debuggers, profilers, libraries, etc)
- ✓ Provide competitive performance for NNSA systems (or strong case on performance/$)
- ✓ Reliability that can now meet production-class system needs

Platform enablement of important libraries and software environment for NNSA is complete

Early results for GPUs also show software environment is ready for accelerator-enabled Arm systems

FUTURE DIRECTIONS

- Arm-ing the Top500
- Los Alamos deploying Arm resource in 2023 in a new partnership with NVIDIA
  - Means the next-generation of Arm processors will be in production use at the NNSA
- The ecosystem of Arm in HPC will continue to mature
- Will also see Alps (CSCS, Switzerland) deploy an Exascale-class NVIDIA system in 2023

- Future of Arm to engage in custom IP and heavily optimized HPC processors isn't as clear
  - This needs much more careful study and investigation
  - Still may require billions in investment
- Not clear where the diversity of Arm processor cores may come from in the future
  - Depends on NVIDIA-Arm acquisition
  - Options in Arm ecosystem don't automatically give us all the advantages we want

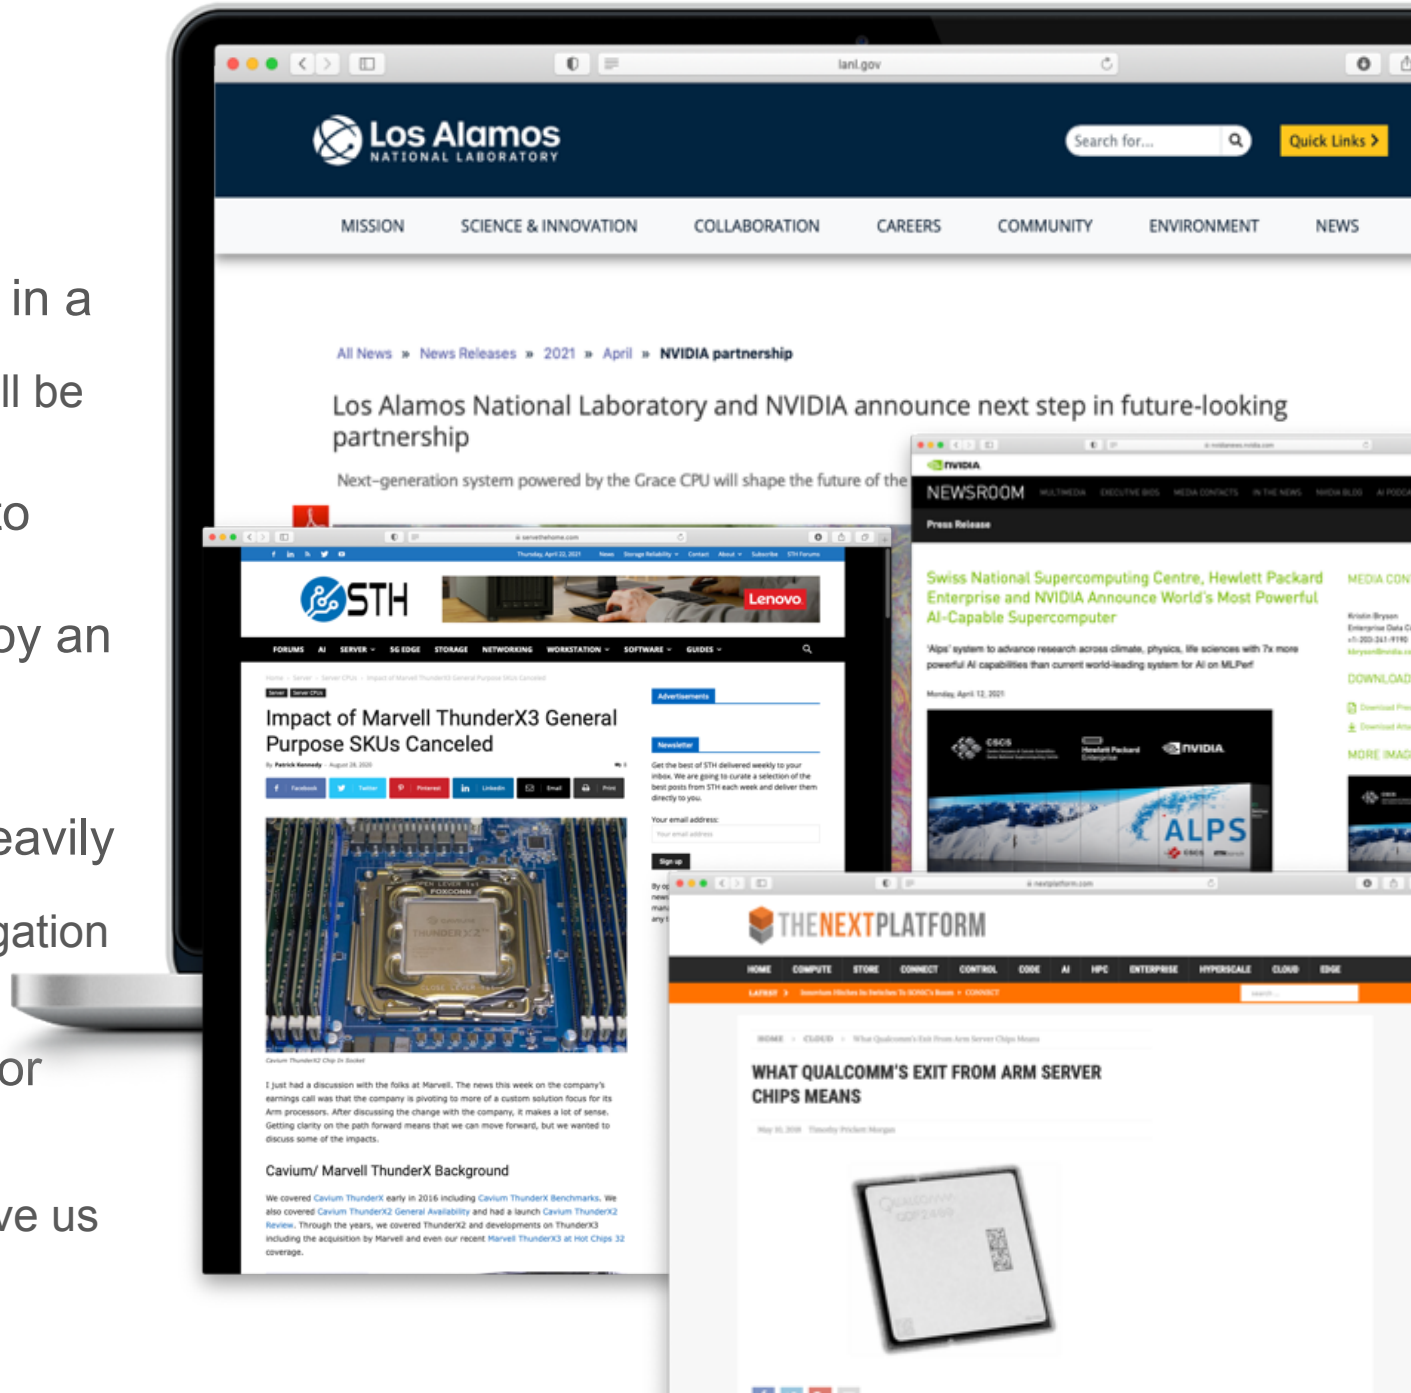


Sandia
National
Laboratories