SAND2013-2990P

Advanced Systems Technology Test Beds

Sue Kelly, smkelly@sandia.gov

Presented to SST Hack-a-thon Attendees
April 12, 2013

Sandia National Laboratories

Exceptional service in the national interest

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Purpose

- R&D and Pathfinding explorations of:
  - Alternative Programming Models
  - Architecture-aware algorithms
  - Low-energy Runtime & System SW
  - Advanced Memory sub-system development
- Computer hardware architecture simulation studies and validation
- Performance analysis
- Vendor interactions
- Outreach for collaboration
- Recruiting

Philosophy

Sandia National Laboratories

- These systems are not for production capability/capacity computing cycles - but can be provided to Test Pilot users
- Both hardware and software is intended to be highly dynamic. They are closer to prototypes and development drivers.
- It is more important to explore a diverse set of architectural alternatives, than push large scale
  - Node-level and small scale O(100) multi-node for network testing

Pipeline of events

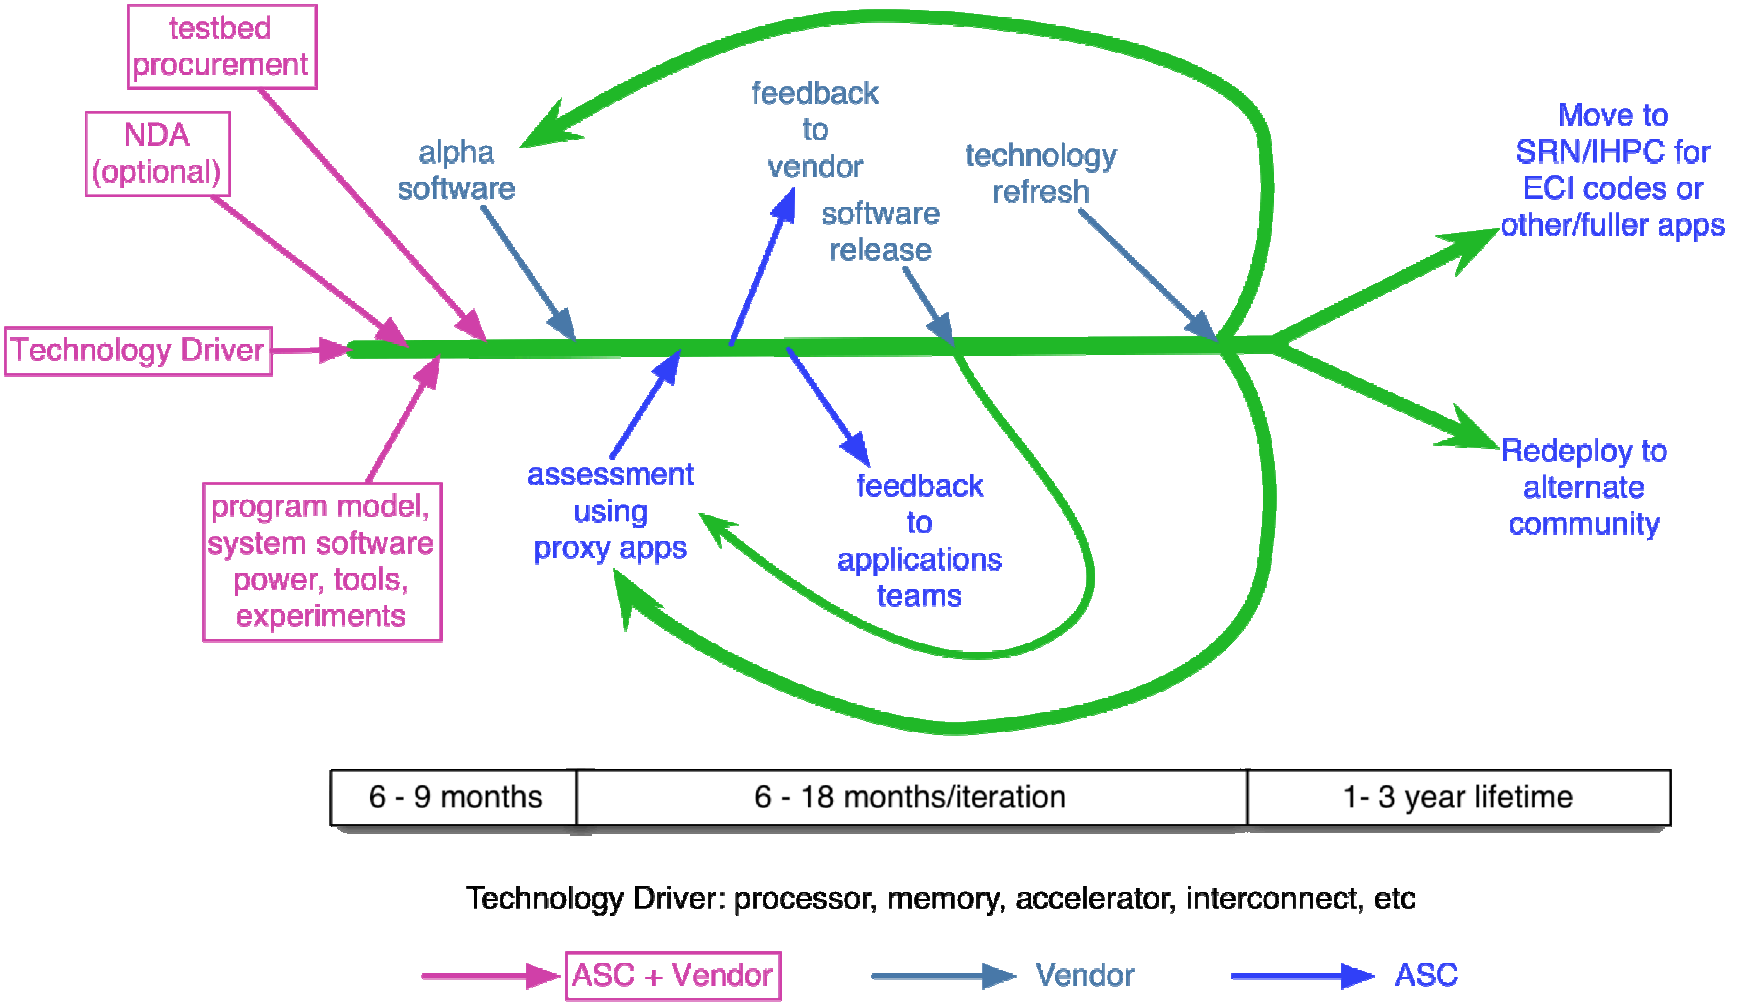

Community

- ASC co-design project members
- ASC tri-lab staff
- ASC university partners
- ASC vendor partners
- ASCR co-design center researchers
- On a case by case basis, determined by the value added to an Exascale future:
  - Non-ASC Sandia staff
  - Non-ASC collaborators sponsored by a Sandia staff member
  - CRADA partners
- As long as the hardware is not export controlled, foreign nationals have access to the newest test beds until they are moved to the SRN (approximately two years)

Hardware strategy

- Once a year refresh for individual technologies
- New technology – one or two per year
- Support approximately 5-7 distinguishing architectures
  - In some cases, purchase small production systems so that riskier software (e.g. OS or power) can be vetted here before being run on large systems elsewhere.
- Test beds are distributed between two network environments.
  - External collaboration network (maybe OHPC) – which allows foreign nationals
    - New machines will typically be placed here first
  - Sandia's Restricted Network or IHPC – which allows export controlled information
    - Older test beds will be moved here for use by application developers who need more mature systems to build and run complex applications.

High Performance Computing at SNL/NM

- **ESNET – Our link to the outside world recently upgraded to 10Gb service**
- **IHPC and OHPC – new intersite collaboration environments**
- **HPC enclave of Sandia Restricted Network (SRN) – standard protected environment for unclassified**
- **ECN – Houses Testbeds and provides open collaborations with external entities – Intel, AMD, Foreign Nationals**

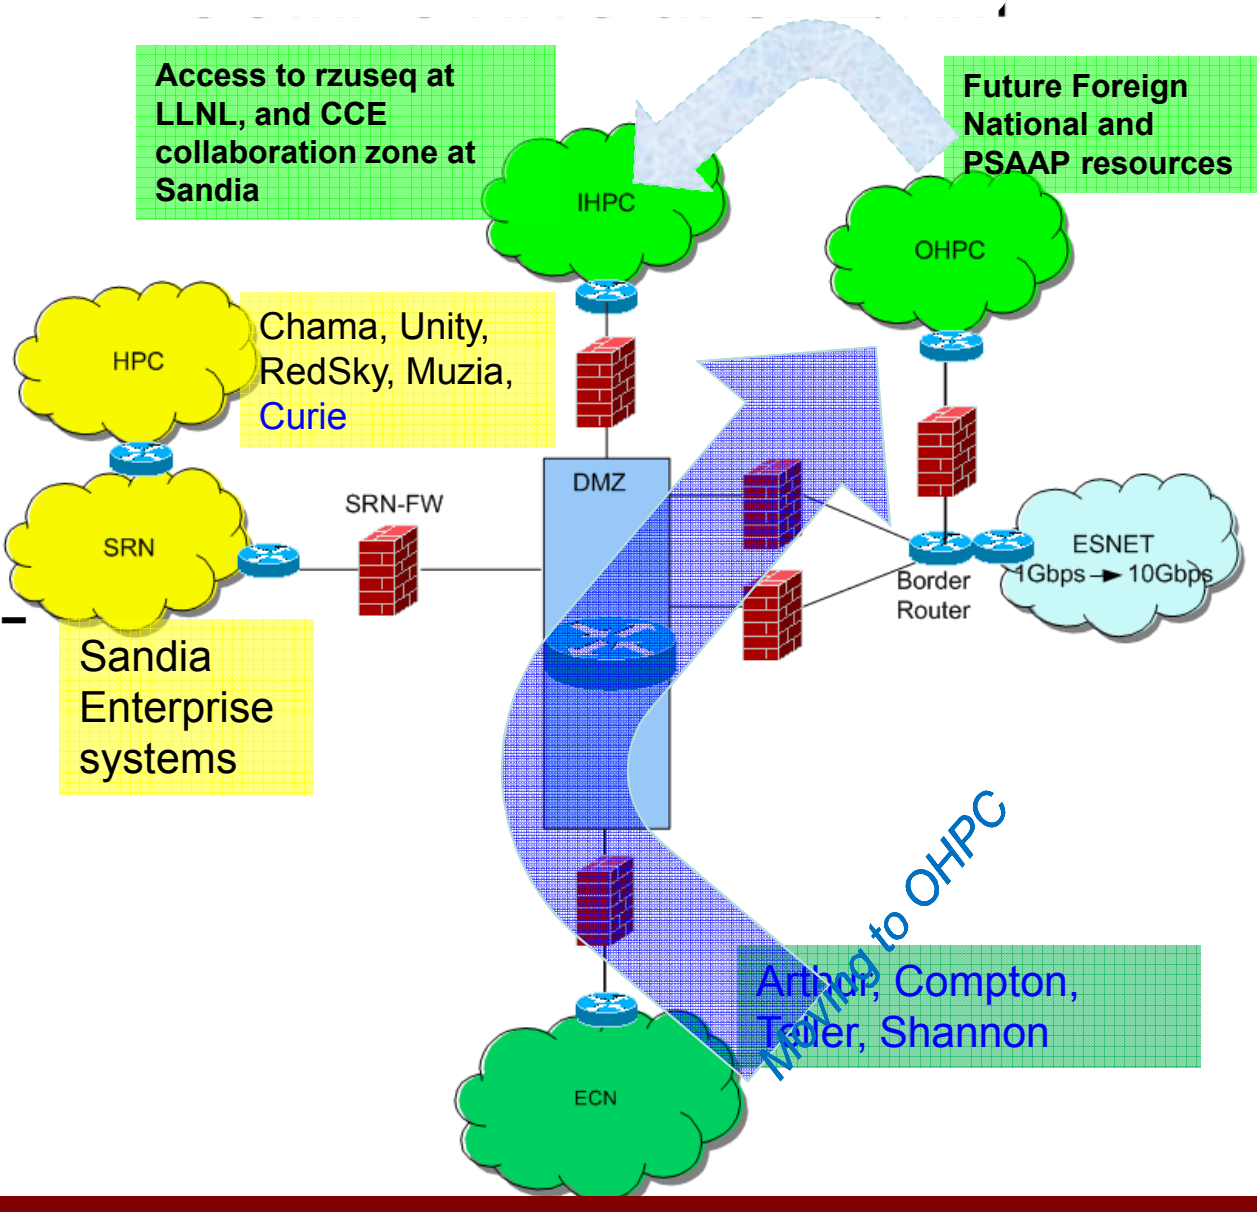

Advanced Architecture ECN Test Beds

Sandia National Laboratories

Feature	Intel MIC Phase 1	Intel MIC Phase 2	AMD Fusion Phase 2	**NVIDIA Phase 2**
Hostname	Arthur	Compton	Teller	Shannon
CPU	Dual Socket Intel X5690 (Westmere) 3.47 GHz 6-core	Dual Socket Intel E5-2670 (Sandy Bridge) 2.6 GHz 8-core	AMD A10-5800K (Piledriver) 3.8GHz Quad-core	Sandy Bridge
Accelerator	Pre-Production Intel Xeon Phi Knights Ferry (KNF) 2 per node	Pre-Production Intel Xeon Phi Knights Corner (KNC) 2 per node	Radeon Northern Islands (on die integration)	Nvidia Kepler K20X 2 per node
GPU cores	30-32 1.1 GHz cores	57 1.1GHz cores	384 800MHz cores	2688 732 MHz cores
Nodes	42	42	104	32
Interconnect	Mellanox QDR IB	Mellanox QDR IB	QLogic QDR IB	Mellanox QDR IB
Other	80GB SSD per node	80GB SSD per node	Integrated CPU/GPU+ 256GB SSD/node custom-designed power monitoring capability	Full PCI Gen 3 NVIDIA GPU Direct

Advanced Architecture Test Beds: Overview

Sandia National Laboratories

Feature	Intel MIC Phase 1	Intel MIC Phase 2	AMD Fusion Phase 1	AMD Fusion Phase 2	Cray XK6	Cray XK7 Phase 2	**NVIDIA Phase 2**
Hostname	Arthur	Compton	Teller	Teller	Curie	Curie	Shannon
CPU	Dual Socket Intel X5690 (Westmere) 3.47 GHz 6-core	Dual Socket Intel E5-2670 (Sandy Bridge) 2.6 GHz 8-core	AMD A8-3850 (Husky) 2.9GHz Quad-core	AMD A10-5800K (Piledriver) 3.8GHz Quad-core	Interlagos	Interlagos	Sandy Bridge
Accelerator	Pre-Production Intel Xeon Phi Knights Ferry (KNF) 2 per node	Pre-Production Intel Xeon Phi Knights Corner (KNC) 2 per node	Radeon HD Evergreen (on die integration)	Radeon Northern Islands (on die integration)	Nvidia Fermi 2090x	Nvidia Kepler K20X	Nvidia Kepler K20X 2 per node
GPU cores	30-32 1.1 GHz cores	57 1.1GHz cores	400 600MHz cores	384 800MHz cores	512 650 MHz cores	2688 732 MHz cores	2688 732 MHz cores
Nodes	42	42	104	104	52	52	32
Interconnect	Mellanox QDR IB	Mellanox QDR IB	QLogic QDR IB	QLogic QDR IB	Gemini	Gemini	Mellanox QDR IB
Other	80GB SSD per node	80GB SSD per node	Integrated CPU/GPU+ 256GB SSD/node	Integrated CPU/GPU+ 256GB SSD/node custom-designed power monitoring capability	Full featured RAS sys	Full featured RAS sys	Full PCI Gen 3 NVIDIA GPU Direct

Software

- Aggressively update software that is in line with the purpose of the platform, in particular compilers, runtime libraries, development tools , and performance analysis tools
- Support multiple versions and/or images at the same time to allow for different experiments (e.g. experimental kernels)

Documentation

Sandia National Laboratories

- Vendors typically provide almost no user documentation
- Test Bed administrators provide a Wiki.
- Administrators document the hardware configuration on the wiki.
- The earliest test pilot users are asked to create “how to run hello_world” documentation.
- Further wiki documentation is provided by administrators and users as they are inspired.

Support

- Account requests are via SARAPE for tri-lab users or via Sandia's "WebCars" account request system for users not from LANL or LLNL. This production process takes approximately one week.
- An email help service (<testbed>-help@sandia.gov) and phone number is provided.
  - Help email is first processed (or created from a phone call) by Sandia's "One-Stop" service, sometimes called level 1 tech support.
  - The test bed system administrators and the test bed managers are level 2 support. They receive the emails at the same time as level 1 and will respond immediately if the request is best handled by them.
  - For application issues, users are to send email to the test bed community at <testbed>-users@sandia.gov. No computational scientists or consultants are provided.

Gaps

Sandia National Laboratories

- Limited SRN/IHPC presence
- Enable transition from proxy apps to real codes with the more mature platforms