

# Improving the mission impact of HPC systems through CO-DESIGN

Rob Hoekstra (+ CSSE/FOUS Team)

The logo for Sandia National Laboratories, featuring a stylized green and blue circular emblem with a grid pattern, followed by the text "Sandia National Laboratories" and the date "October 29, 2019".

Sandia National Laboratories  
October 29, 2019



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



Strategic Approach/Goals



The Value Proposition of Co-Design



Topical Areas



Progress



High Level Challenges/Questions

# Sandia's Mission Computing Strategy

## 1. We must out-innovate our adversaries in computational weapons-engineering and attract, train, and retain future talent

- This includes novel problem formulations, advanced algorithms, code-acceleration, smart-resource allocation, data-fusion, advanced engineering workflows, automation & continuous computing, etc. all to enable more efficient and effective use of computing resources.

## 2. We require larger, more agile capacity computing systems to meet increased mission demands

- Increased investments to support *ramp-up* of modernization programs and help counter slowing of computer performance improvements,
- Data analytics clusters tied to HPC and NW databases for *data-centric* deterrence,
- Data pipelines, workflow, and automation to support *full-system life-cycle models*.

## 3. We have the expertise and capabilities to lead Tri-lab activities in advanced architecture prototypes for co-design, vendor influence, workforce engagement, and to promote tri-lab partnerships in HPC technologies and system design.

## 4. We must engage in the design and deployment of Tri-lab capability computing resources to support Sandia's unique weapons-engineering missions (non-nuclear components in all environments: re-entry, hostile, safety-security, etc.).

## 5. We commit to full usability of Tri-lab capability computing resources, regardless of siting location, providing a level-of access, usability, and code/user support equivalent to a locally sited resource.



# The growing role of co-design in HPC



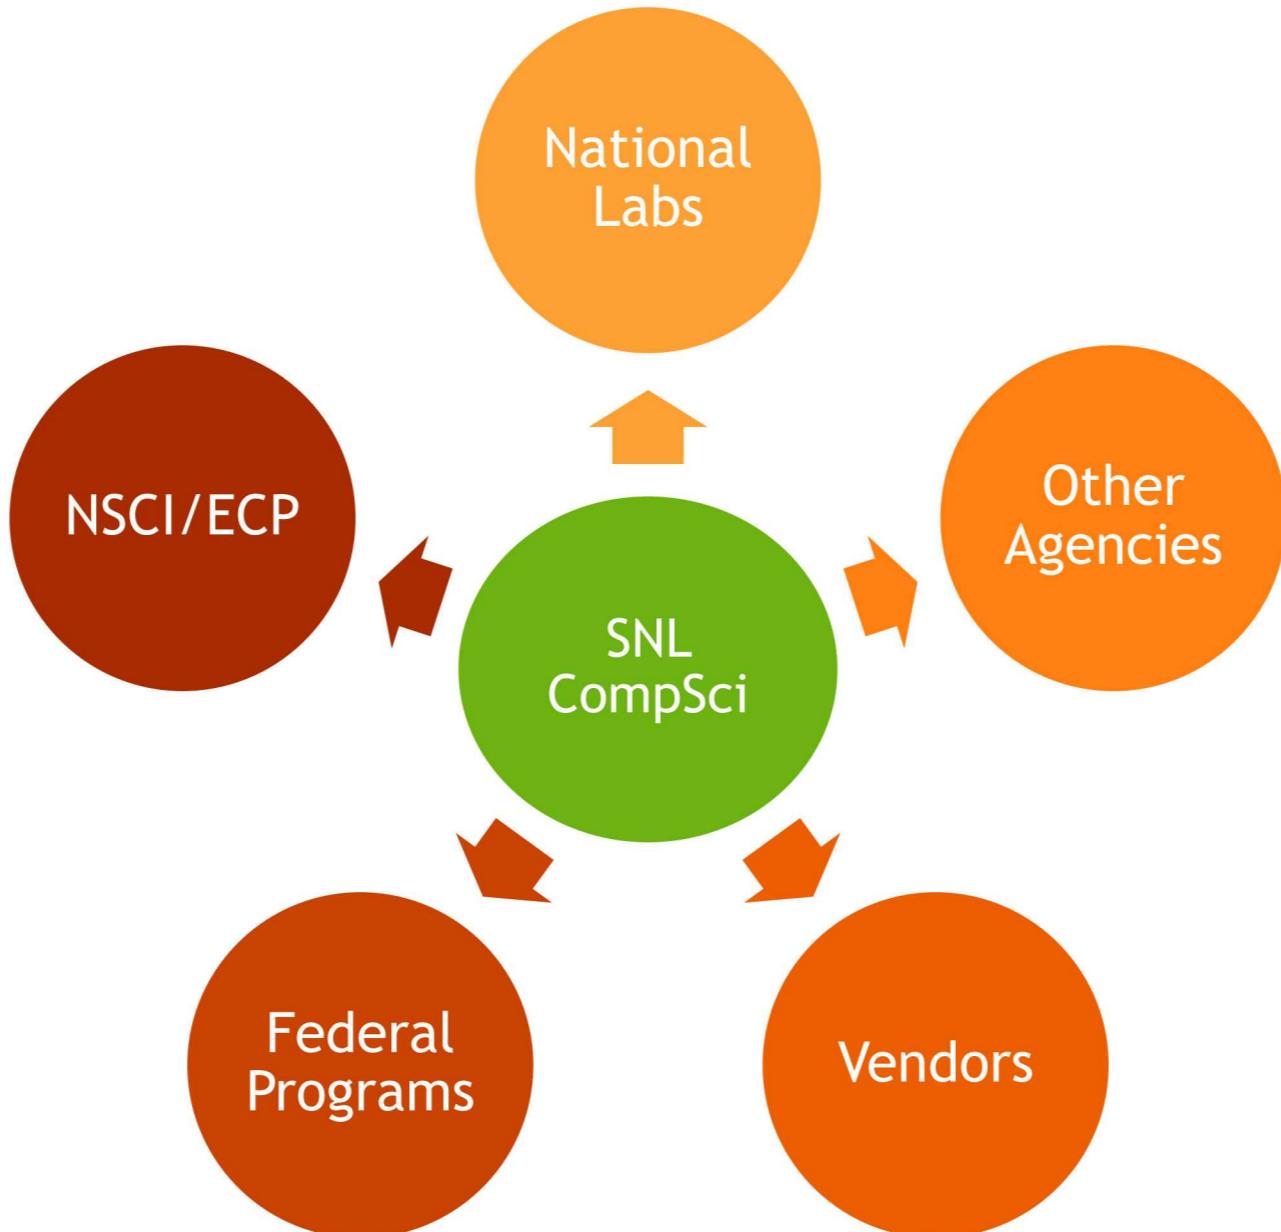
Historically the connections between HPC communities have been tenuous at best/non-existent at worst

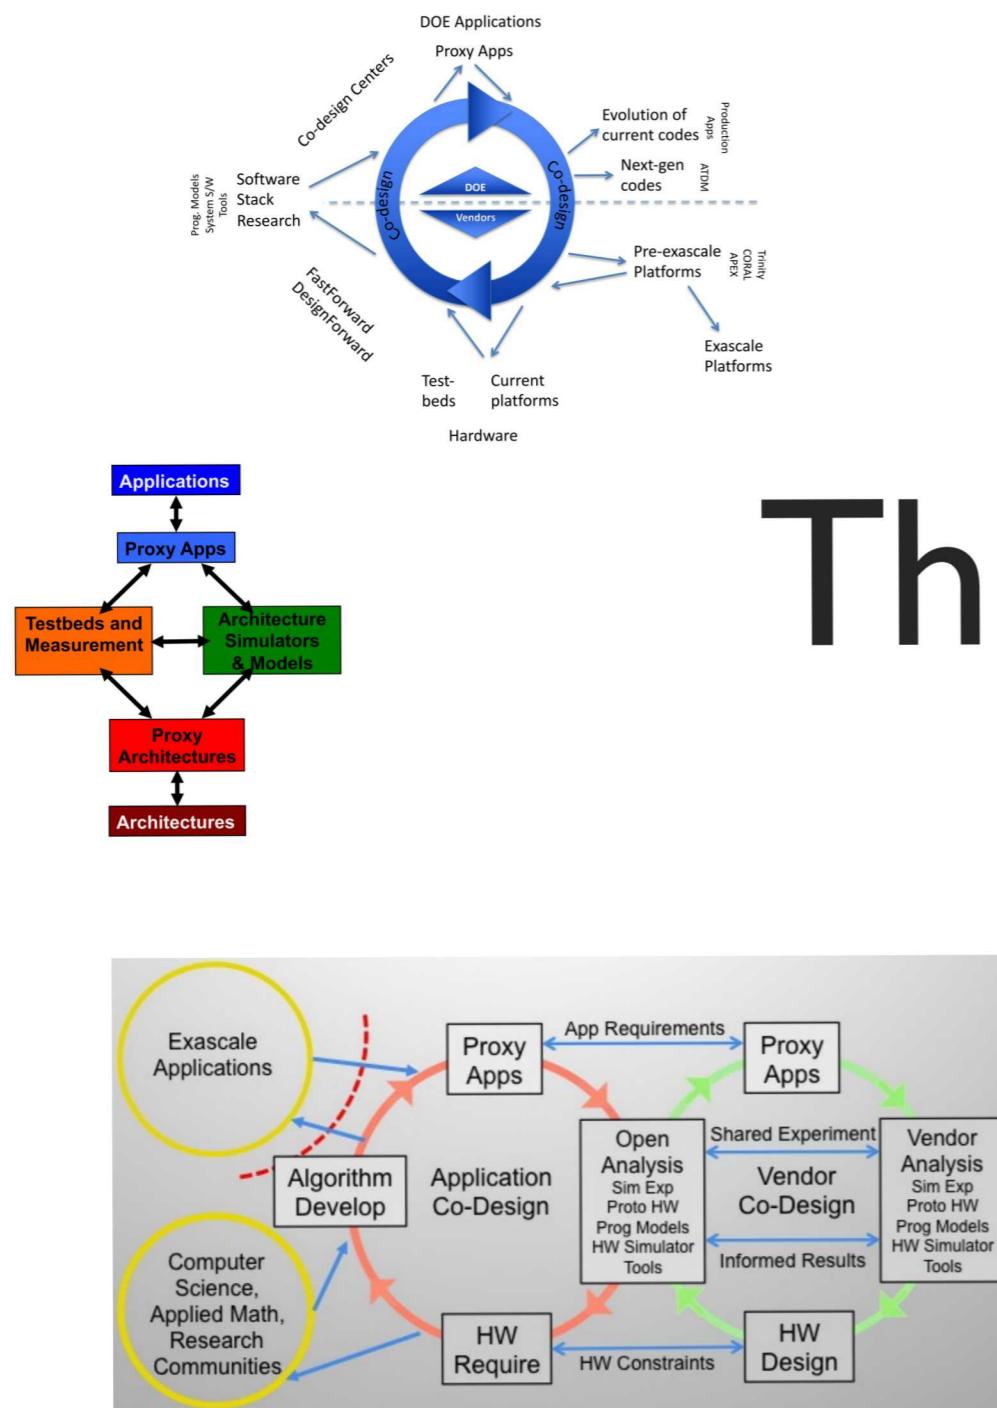
We have been leaders in trying to change that and benefit from collaboration across traditional boundaries

The broader HPC community has embraced this approach

What role do we see ourselves playing going forward?

# Crucial External Interactions



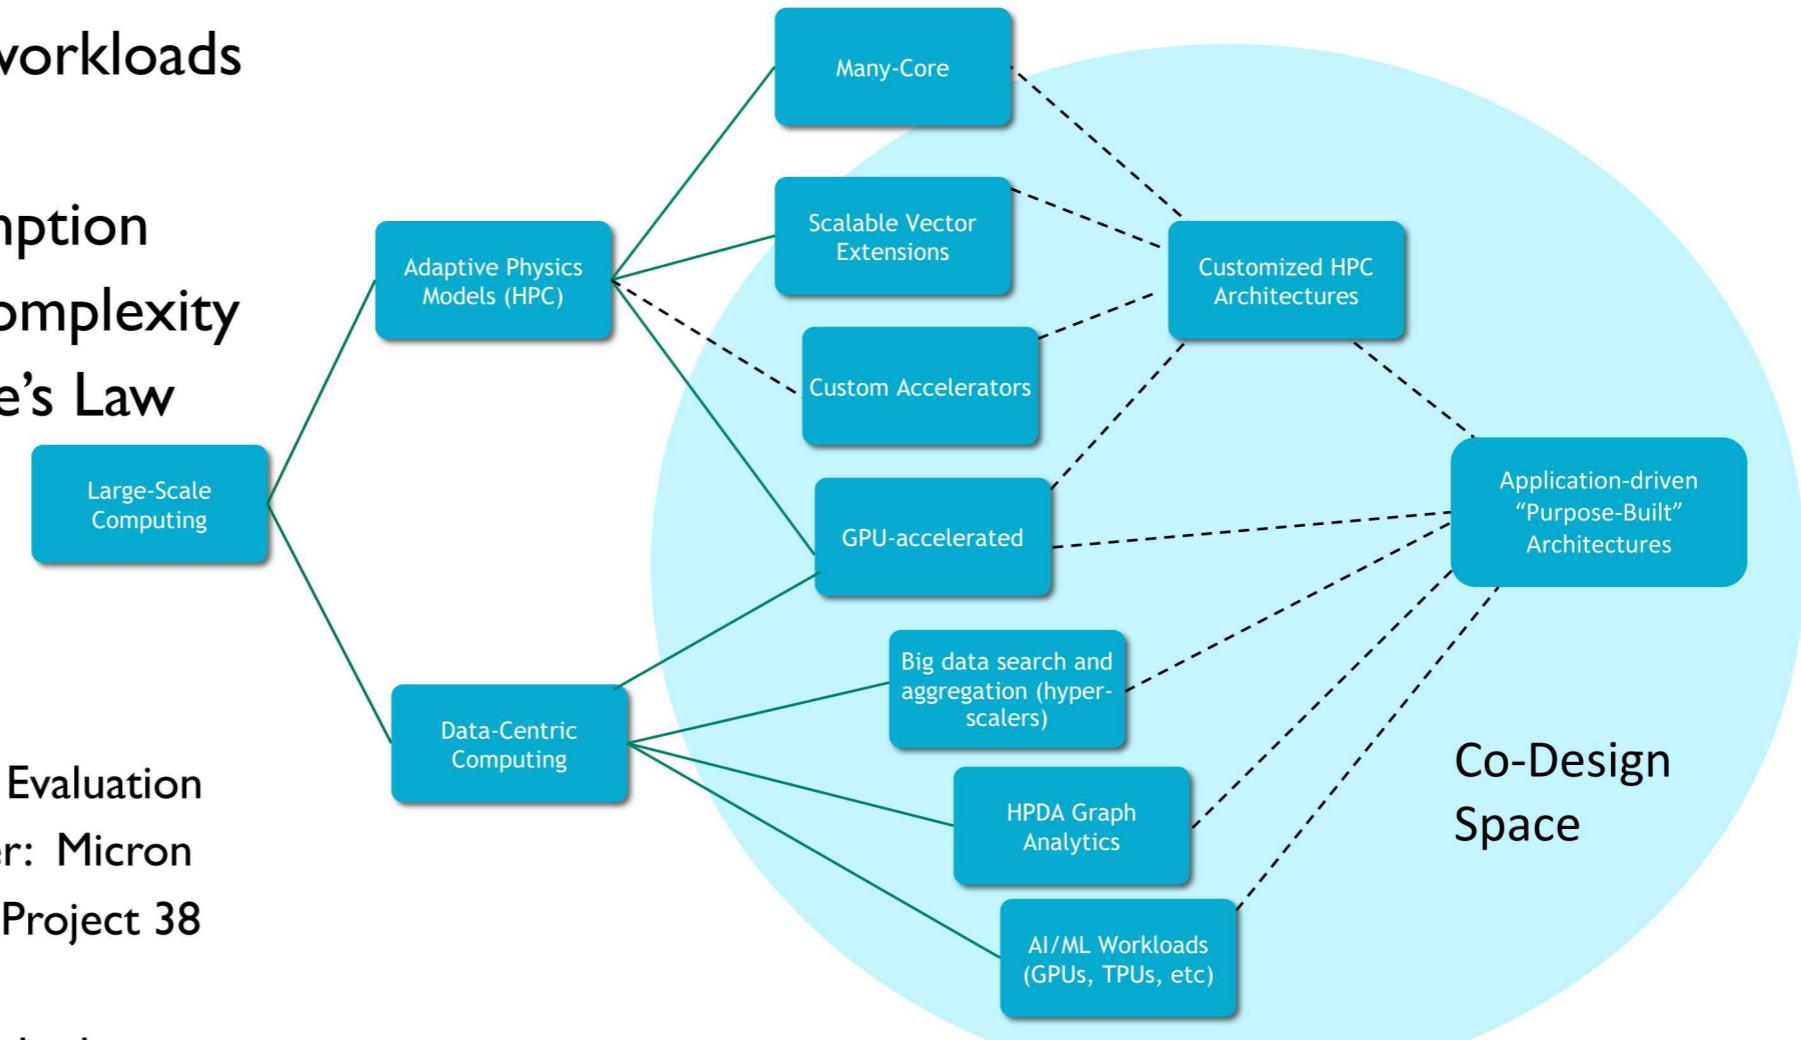


# The many faces of Co-design

# 6 Node-Level Heterogeneity – Commodity and Custom Accelerators

## Goals:

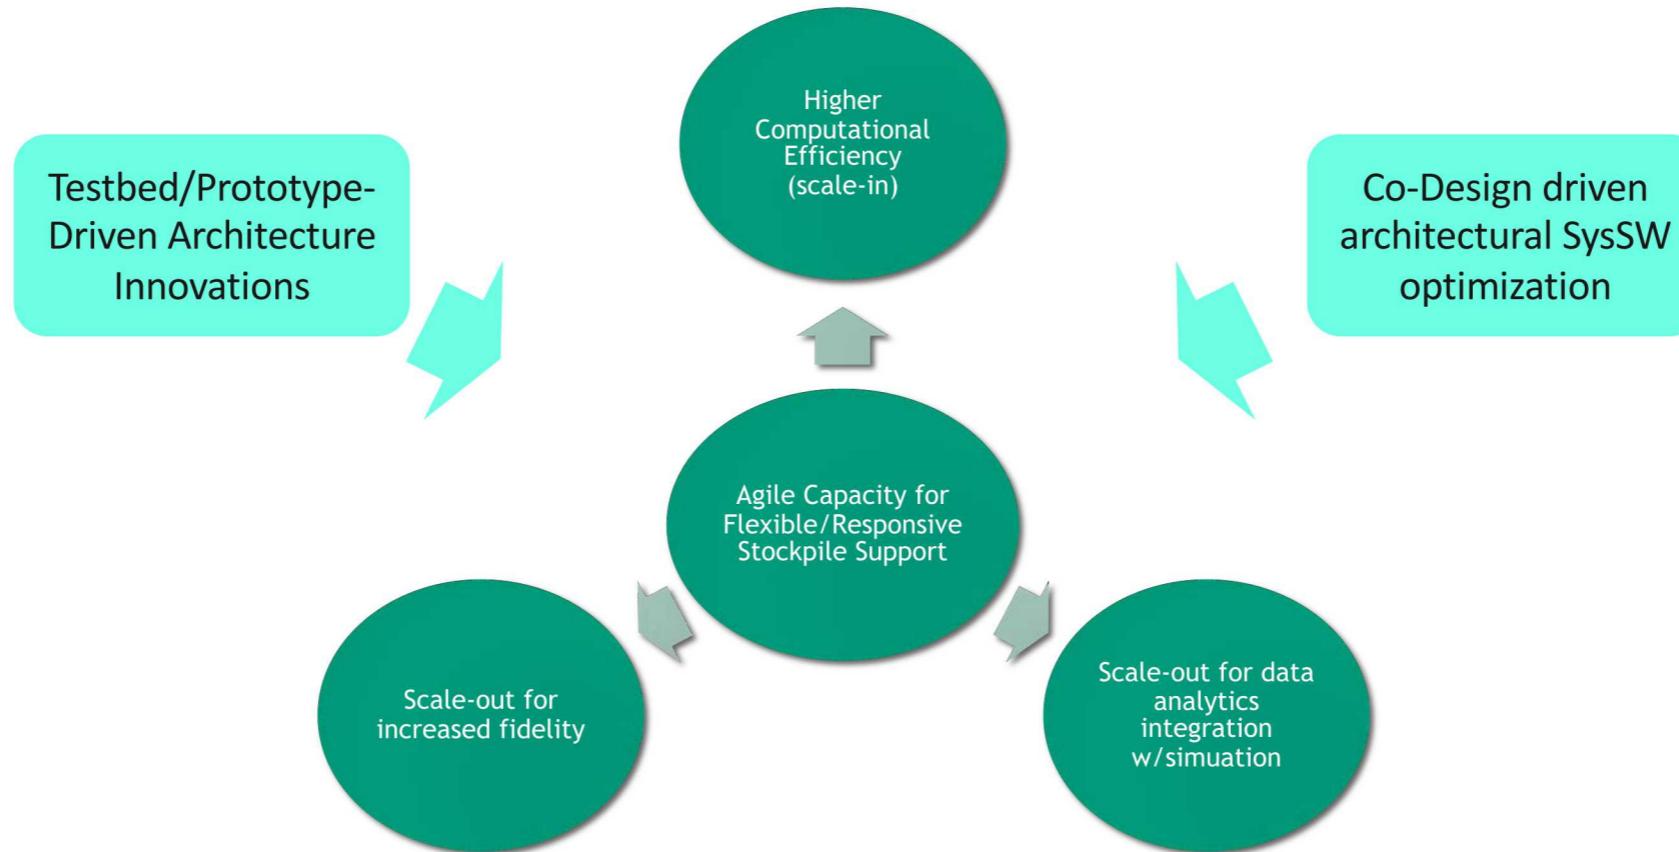
- Match architectures to workloads
- Increase node efficiency
- Decrease power consumption
- Reduce code rewrites/complexity
- Sustain or exceed Moore's Law



## Mechanisms:

- ECP: PathForward & HW Evaluation
- Memory Innovation Center: Micron
- Multi-agency exploration: Project 38
- HW Simulation: SST
- Advanced Architecture Testbeds
- Architectural Prototypes: Vanguard (Astra)

# System-scale Extreme Heterogeneity: “Agile Capacity” for Data-Centric Engineering



***Key strategy: Leverage commodity and semi-custom solutions to deliver next generation heterogeneous systems for MIXED WORKLOADS for Nuclear Deterrence Engineering Mission***

# What do we want to accomplish?



LDRD  
Adv Arch  
Testbeds

Accelerate HPC innovation which has the highest impact on our workloads

DOE Procurements  
Interagency Collaborations

Vanguard Prototypes  
ECP HW Eval

Optimize the HPC platforms acquired by NNSA, DOE and the nation

Vendor Partners  
ACES/APEX

SST HW Sim  
ECP Pathforward

Assure mission computing has the best possible HPC resources

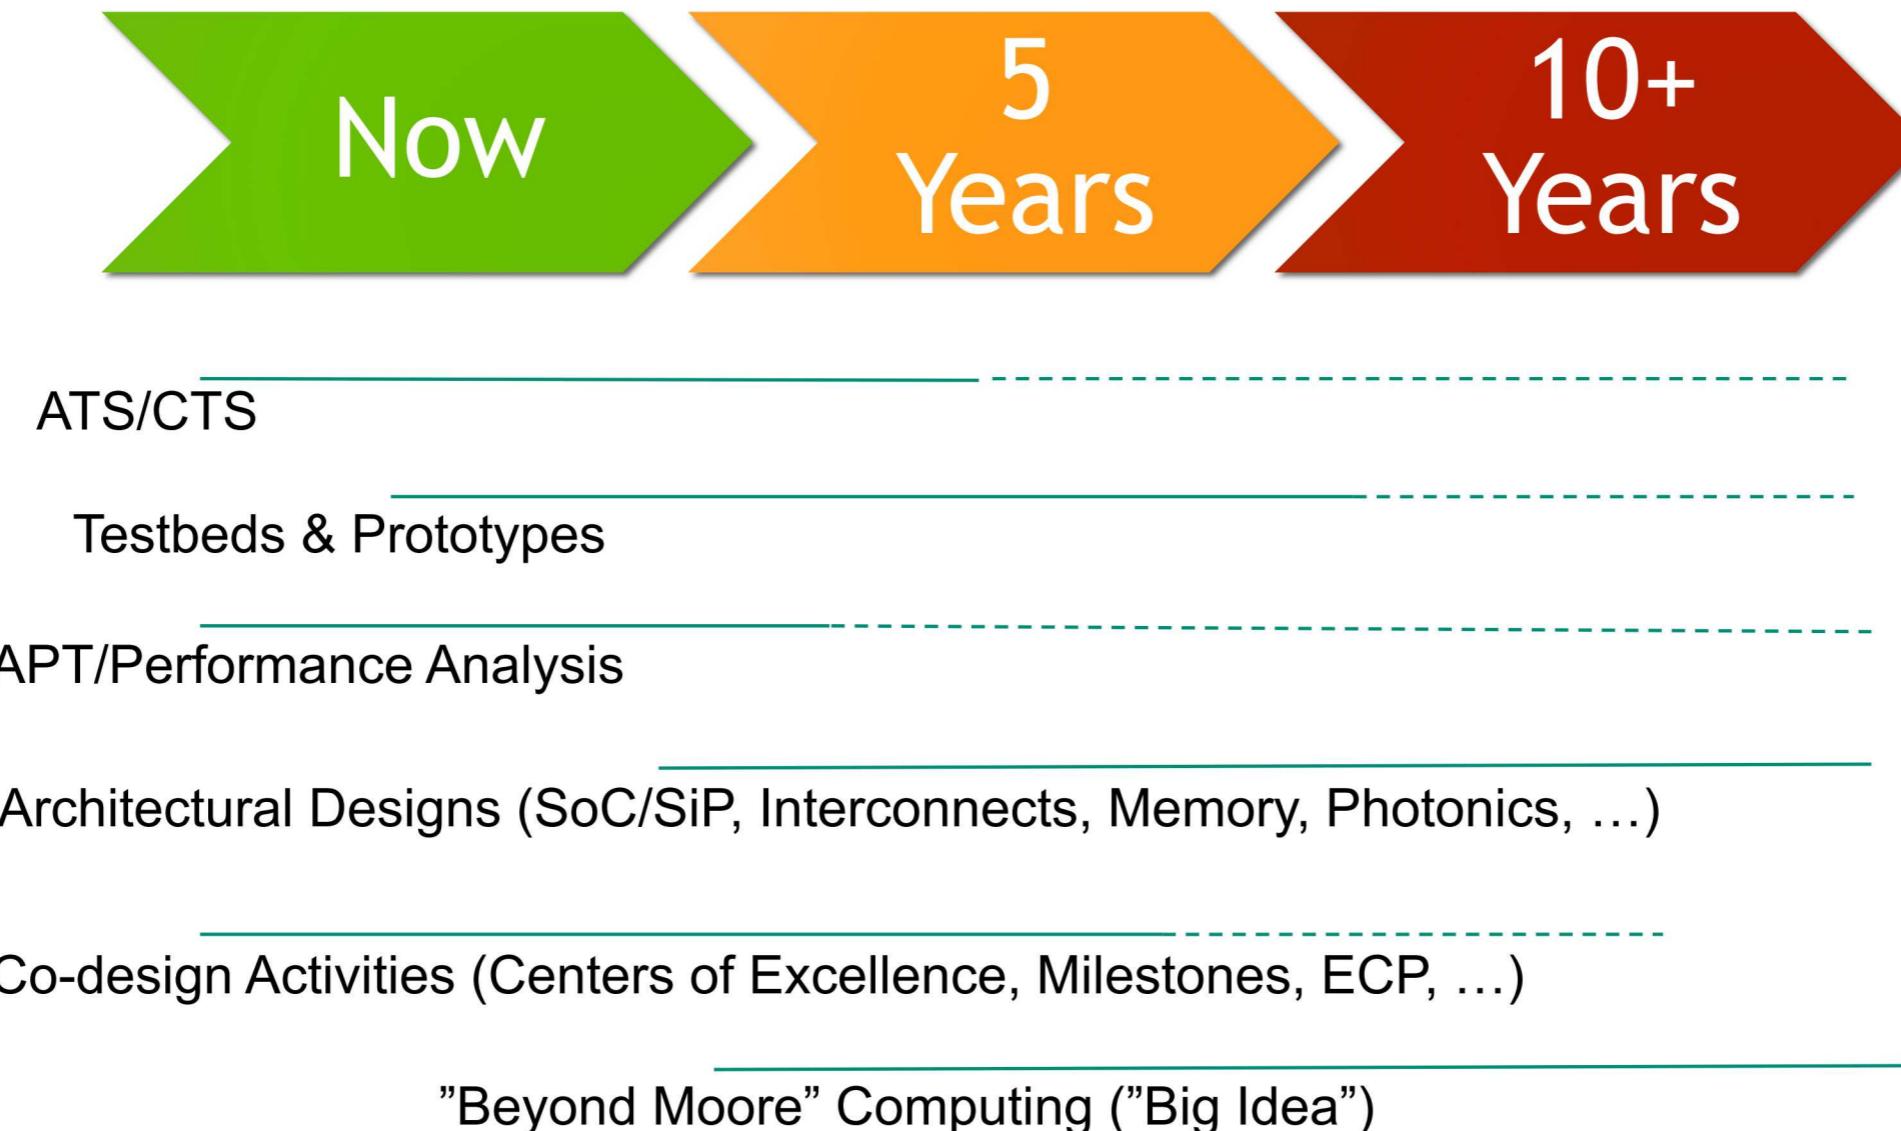
Arch. Office  
Comp. Infra.

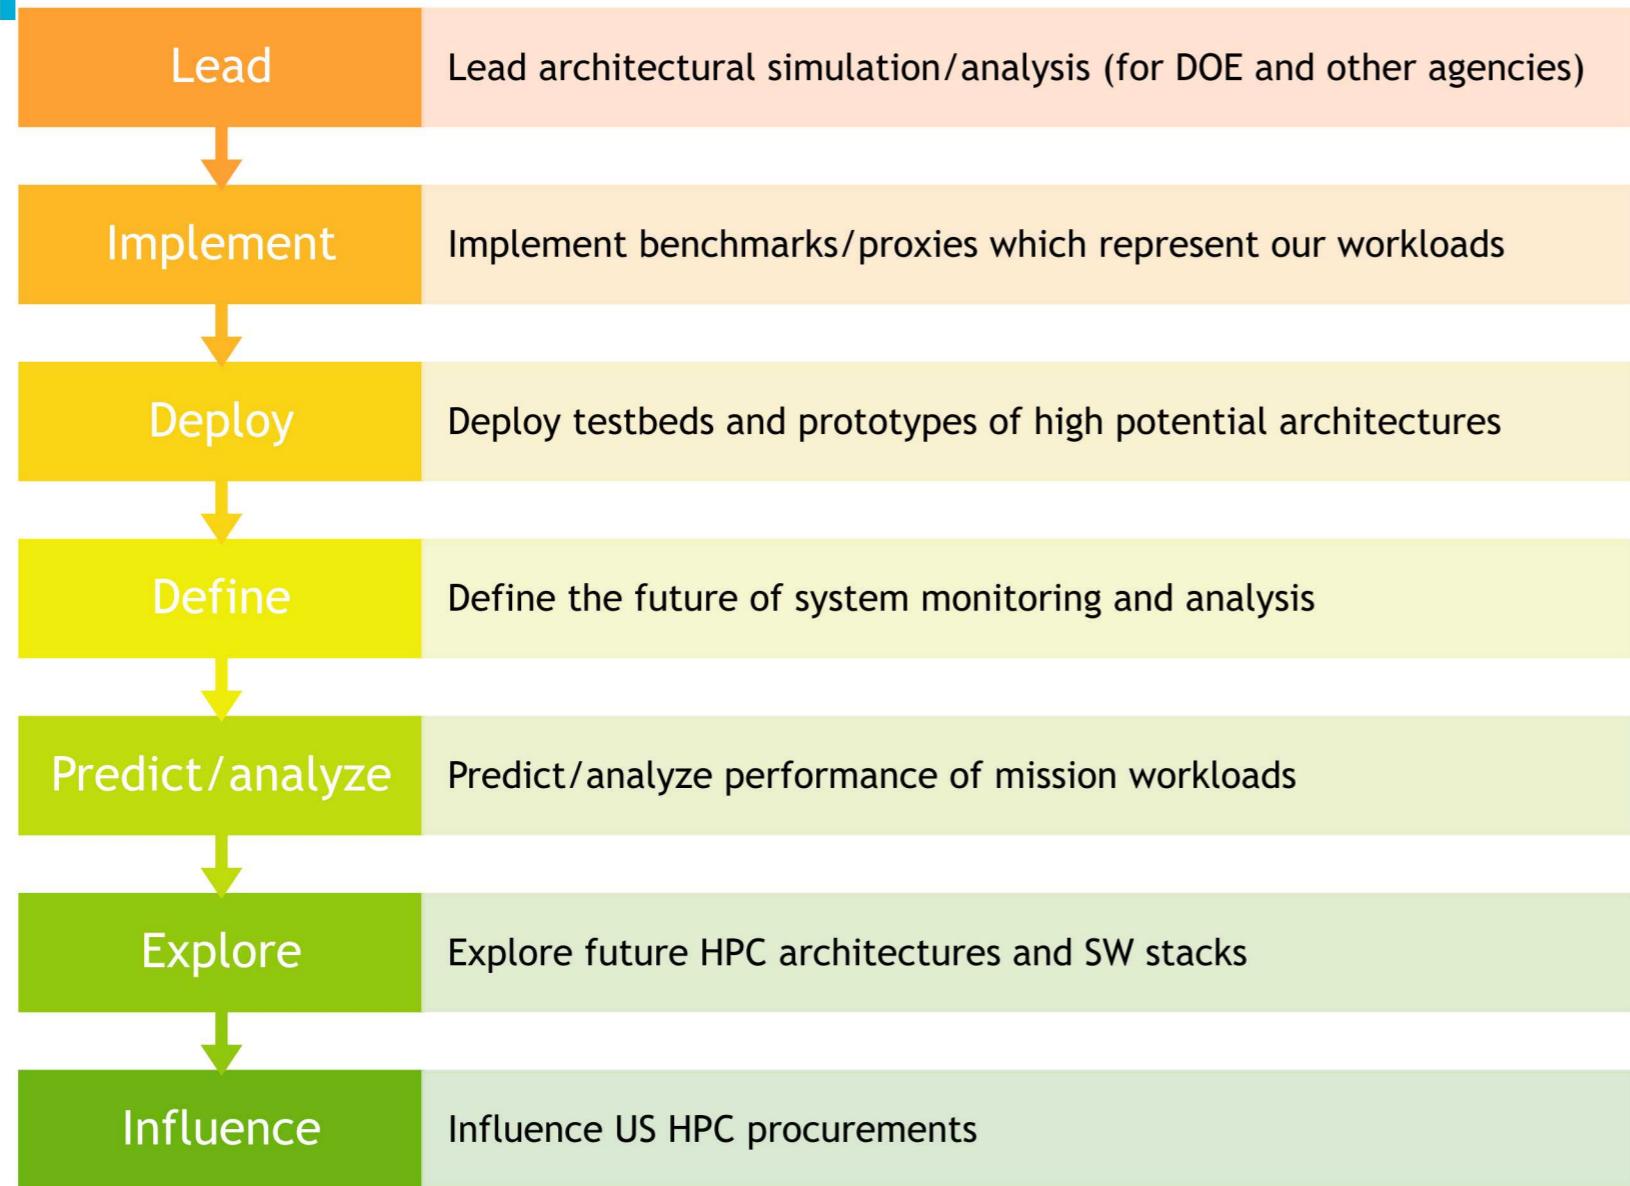
Tri-lab Co-design  
Proxy Apps

Make sure our codes are ready for this future

Performance Tools  
App Perf Team

# Impact the Future of HPC





Drive the future of HPC architectures and system SW to maximize their national impact (especially for national security)

# Areas of Co-design Focus (not exhaustive 😊)



TESTBEDS  
PROTOTYPES



SYSTEM  
MONITORING



HW SIMULATION



BENCHMARKING  
PERFORMANCE EST.



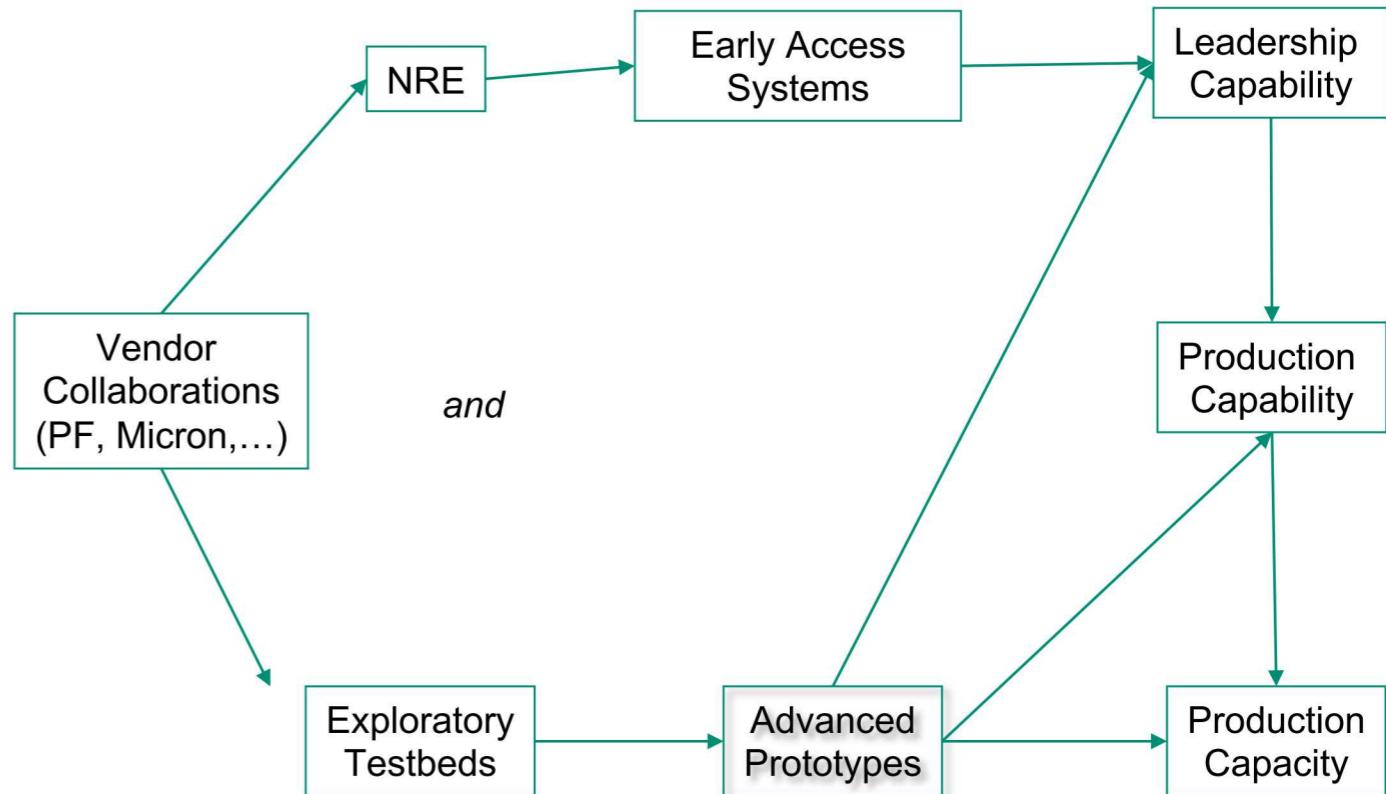
SYSTEM SW



PROGRAMMING  
MODELS

# Innovative pathways to production: the value of testbeds and prototypes

- Complementary approach to NRE for fostering innovation
- Cost-effective approach to expand the HPC ecosystem and manage innovation risks
- More than “take a sip”, but less than “bet the farm”
- “Admiral’s Test” approach



**Enables innovation and intellectual leadership  
for a site that doesn't host the “big iron”**

# Advanced Architecture Testbeds and Vanguard Prototypes



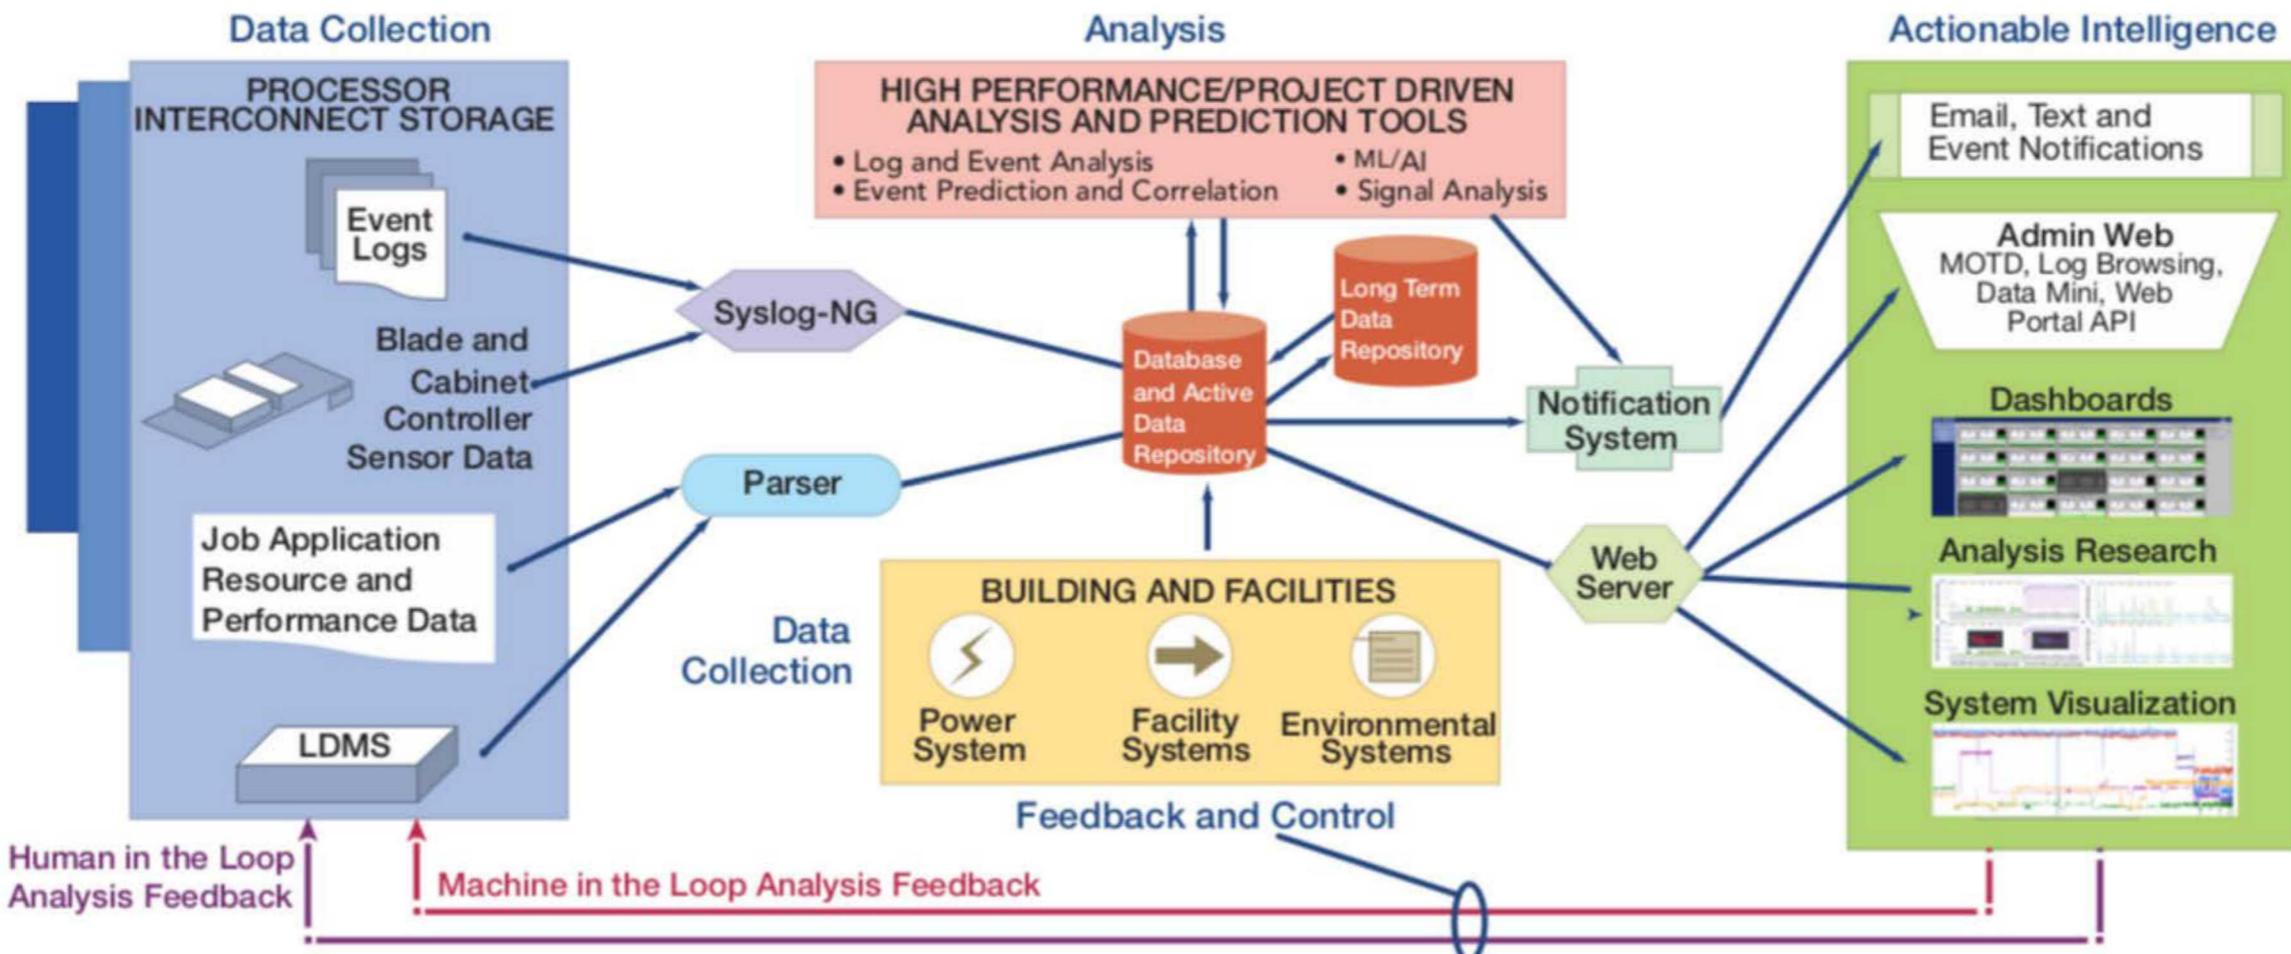
ASC Advanced Architecture Testbeds Program



ASC Astra: First System in Vanguard Advanced Prototype Program



# Putting Performance Into HPC Through Monitoring, Analysis, and Feedback: LDMS



Identify new instrumentation available on new architectures

- Create appropriate data samplers, validate utility of data, measure collection overhead
- Provide advanced insights on how efficiently applications are utilizing new architectural features

Discrepancies between actual and simulated resource utilization:

- Valid simulation capability and help identify errors in assumptions
- Drive more accurate simulations of application resource utilization and performance for both current and future architectures

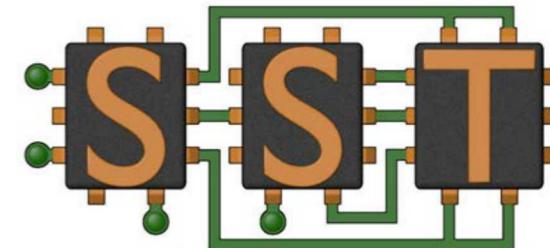
# The Structural Simulation Toolkit



*Using supercomputers to design supercomputers*

## Goals

- Create a standard architectural simulation framework for HPC
- Ability to research & evaluate future systems on DOE/DOD workloads
- Facilitate hardware-software codesign for future architectures
- Communication and collaboration tool for community



## Status

- Parallel Simulation Framework “Core”
- Integrated component libraries “Elements”
- Current Release 9.0

## Consortium

- Bring together labs, academia, & industry



# Benchmarking and Proxy Apps

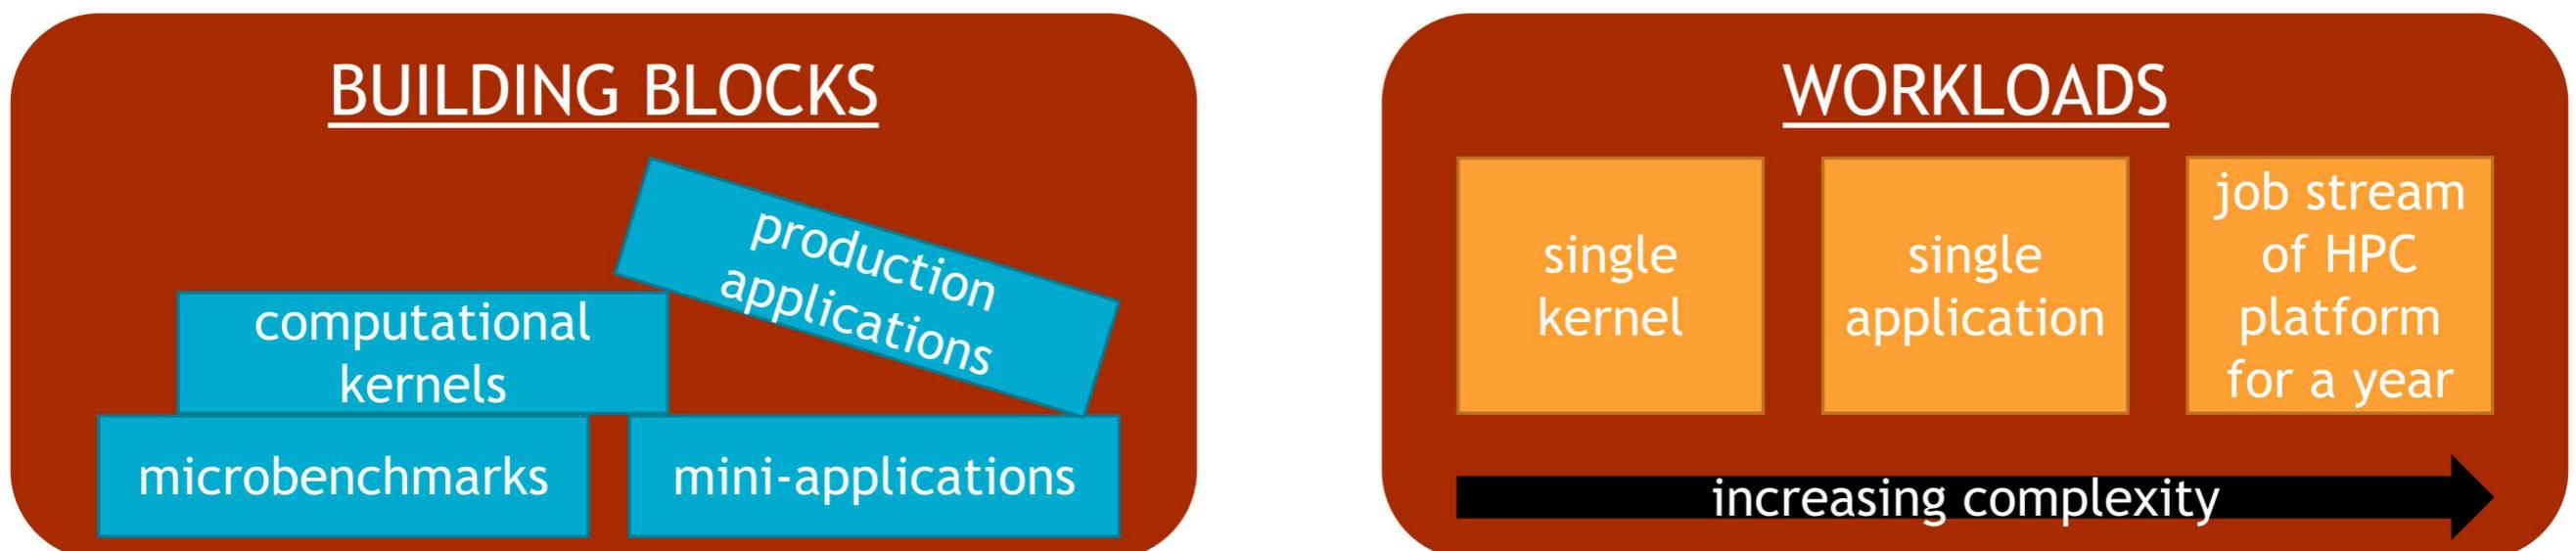


Early leadership in Proxy Apps and benchmarks led by Mike Heroux

- Mantevo Proxy App Suite
- HPCG Top500 Benchmark

Novel approach being explored in Symphony

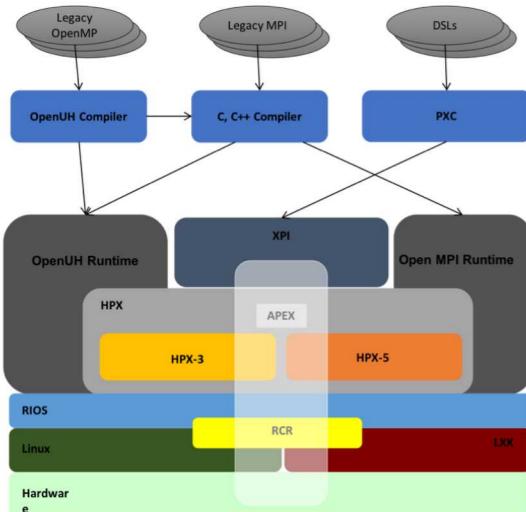
- Symphony is an internal tool and methodology developed by the HPC Application Readiness Team (HPCART).
- Symphony approximates a complex **workload** using a collection of simpler **building blocks**.
- This is potentially useful for many tasks, e.g., when testing performance, generating system loads, when trying to understand the most important components of a workload.



# System SW for next generation systems



*Containers for HPC:  
Partnering with Sylabs  
on Singularity*



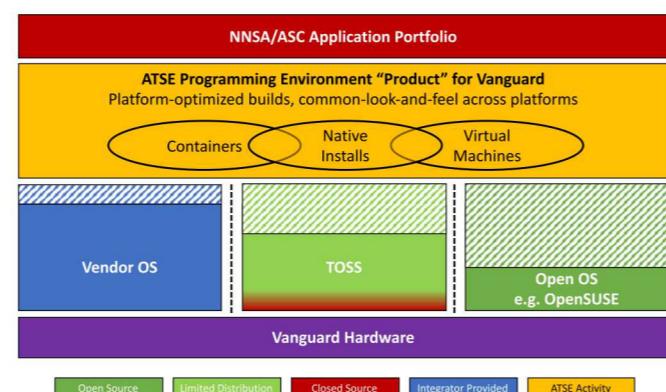
*XPRESS: eXtreme-scale  
Programming Environment and  
System Software*



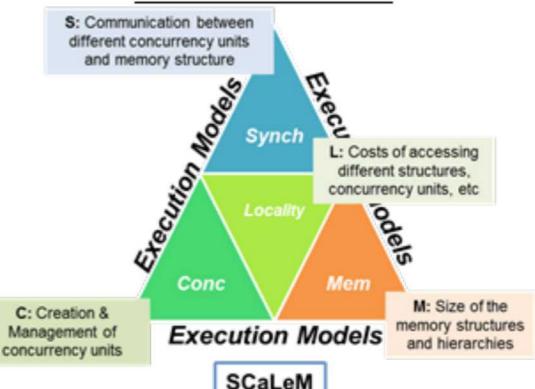
*Portals Communication Library advances  
Interprocessor Communication research  
(Bull interconnect HW support for  
Portals)*



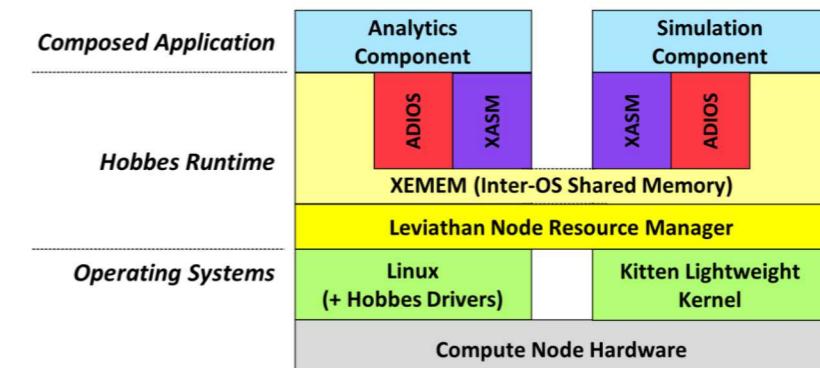
*Engagement with MPI  
Forum and OpenMPI*



**Foundations**



*Kitten Lightweight Kernel (LWK)*

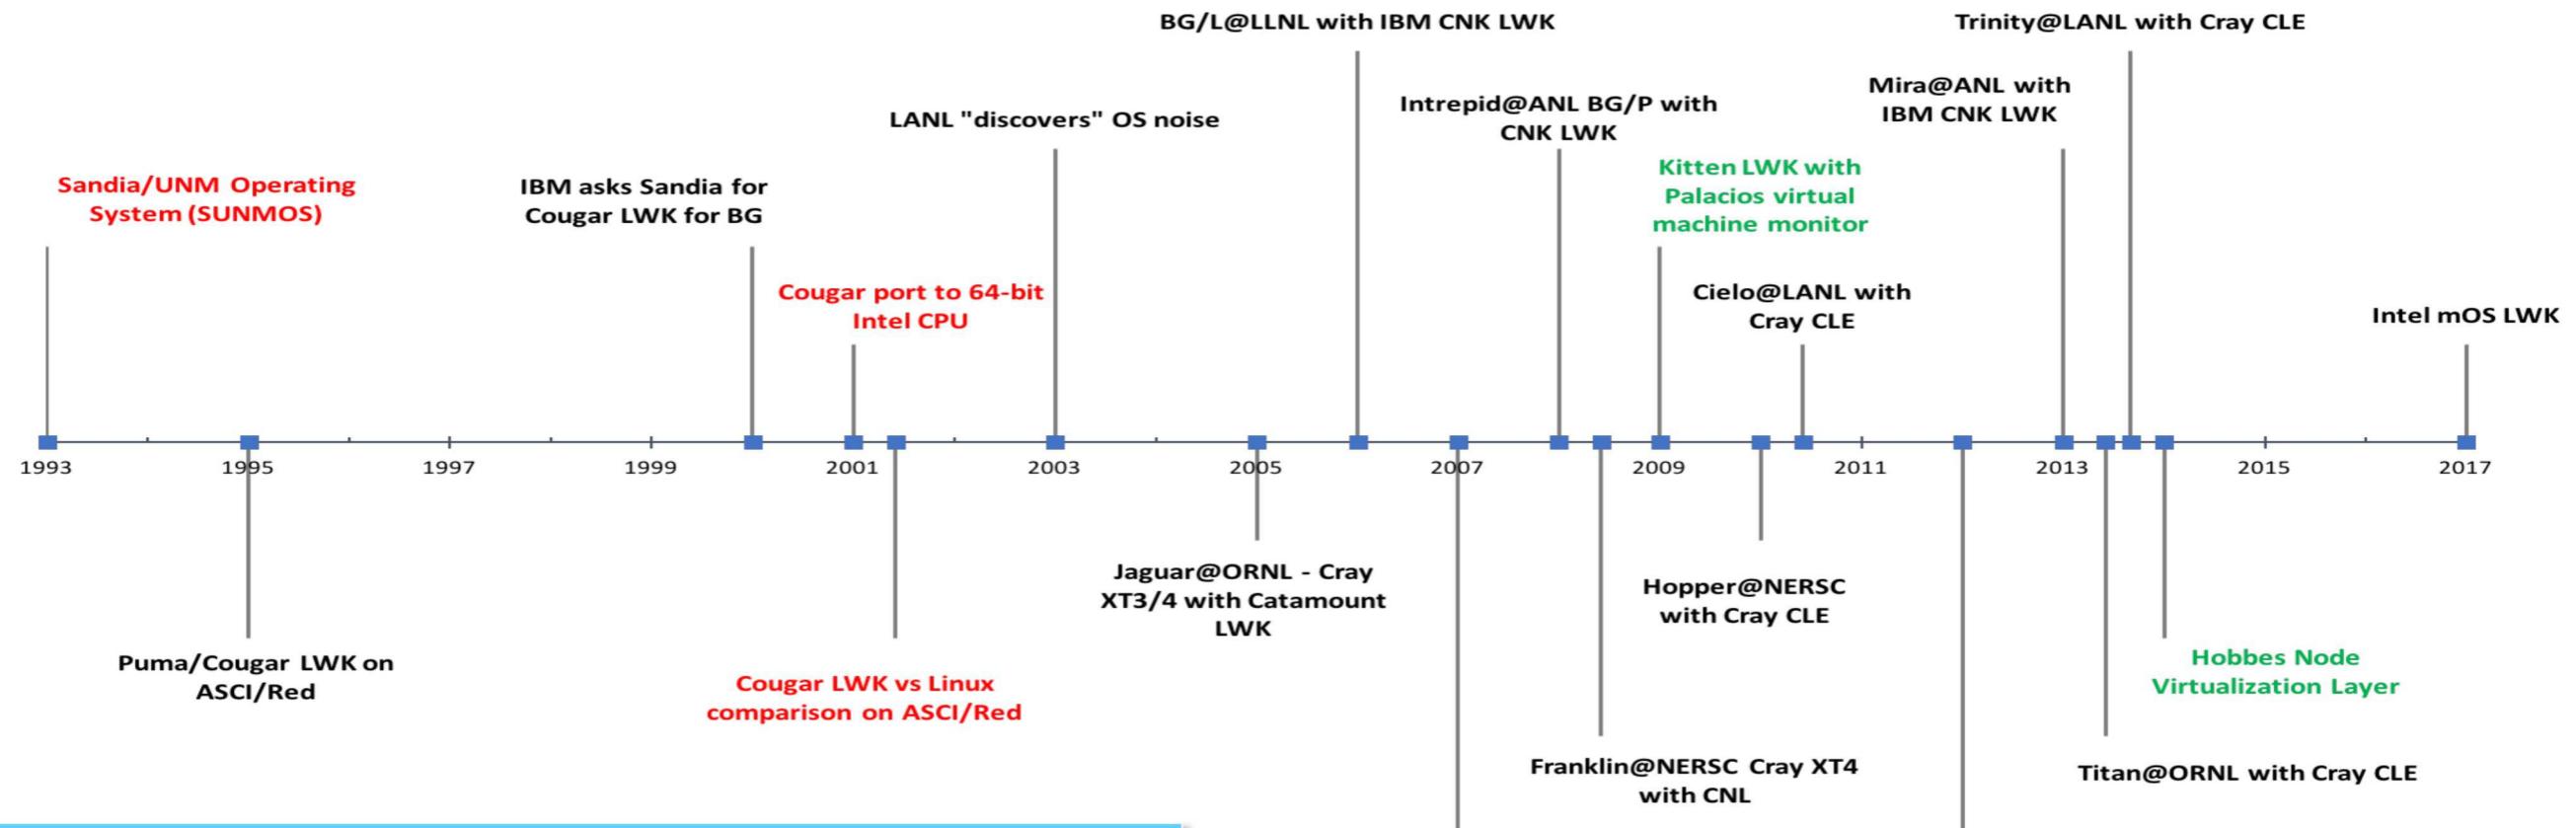


*Hobbes enabling multiple SW stacks to unify  
simulation and analytics for workflows on  
compute nodes*

# Sandia's LWK Approach Has Had Broad Impact

Sandia has partnered with vendors to deploy a custom OS for multiple production systems

- SUNMOS LWK on Intel Paragon; Cougar LWK on ASCI/Red; Catamount on Cray Red Storm
- Other vendors have followed the LWK model: IBM CNK for BG/{L,P,Q}; Cray's Linux Environment



Every large-scale DOE distributed memory machine in the past 25 years has deployed a lightweight OS

Cray develops lightweight Compute Node Linux (CNL)

Sequoia@LLNL with IBM CNK

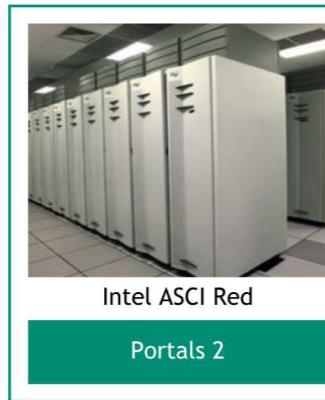
# Vendor Impact of Sandia's Portals Networking Technology

All of these production vendor-supported systems used Portals as the network hardware programming interface. Portals enabled the first TeraFLOPS platform (ASCI Red) and the first non-accelerated PetaFLOPS platform (Jaguar).



Intel Paragon

Portals 0



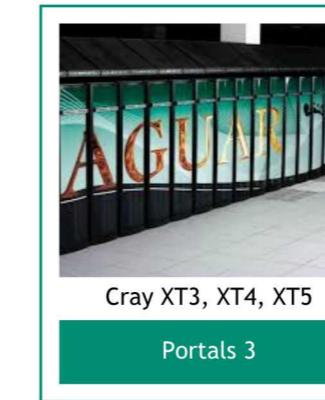
Intel ASCI Red

Portals 2



Cray Red Storm

Portals 3



Cray XT3, XT4, XT5

Portals 3



Atos Tera1000

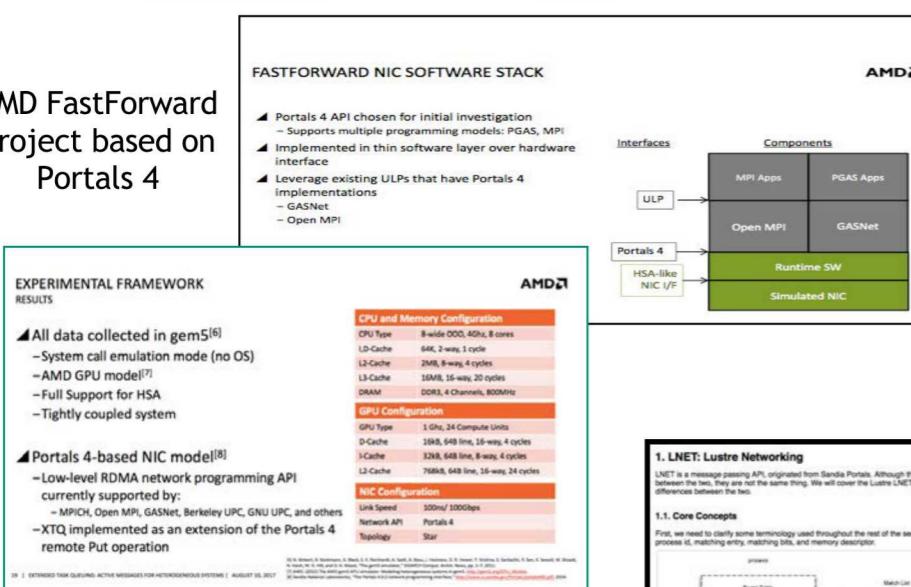
Portals 4



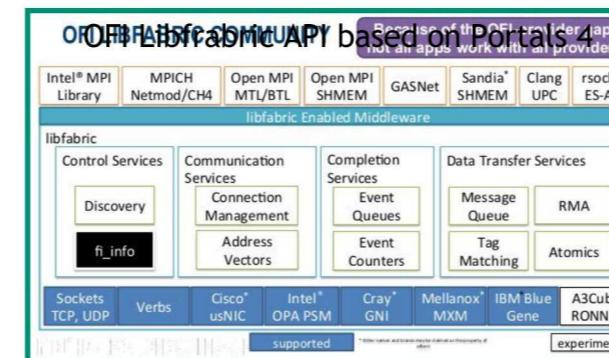
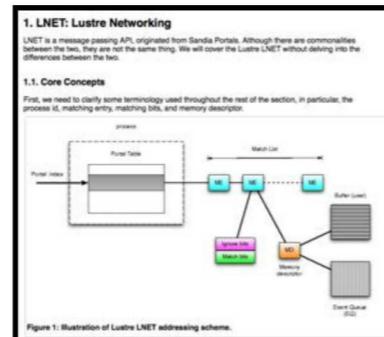
Unlike other low-level network programming interfaces, Portals is intended to enable co-design rather than serve as a portability layer.

The influence and impact of Portals can be seen in vendor co-design activities, other low-level network programming interfaces, and emerging network hardware.

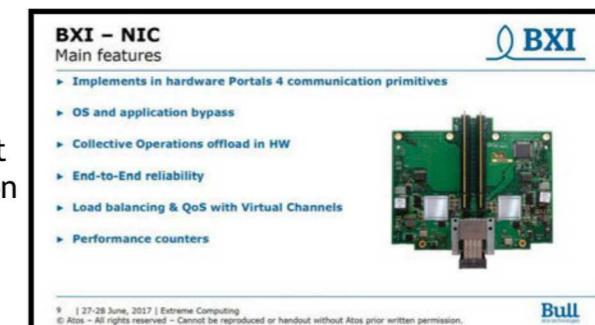
## AMD FastForward Project based on Portals 4



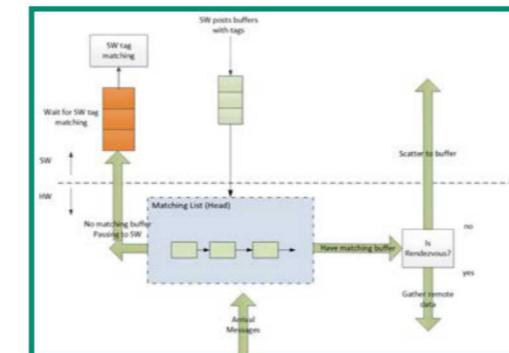
Lustre File System network based on Portals 4



## Atos Bull eXascale Interconnect (BXI) based on Portals 4

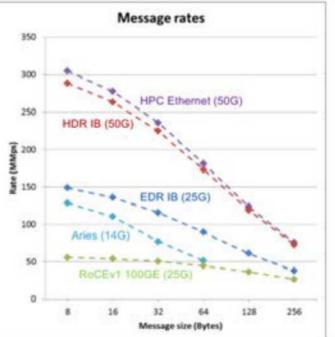


## Cray Slingshot Supports Portals 4 header



Mellanox ConnectX-5 MPI tag matching in hardware

- Slingshot speaks standard Ethernet at the edge, and optimized HPC Ethernet on internal links
- Reduced minimum frame size
  - Remove Ethernet's 64B minimum frame size
  - Target a 40B frame rate but allow 32B frames + sideband
- Removed inter-packet gap
- Optimized header
  - Reduced preamble
  - IPv4 and IPv6 packets can be sent without an L2 header
  - Portals uses modified IPv4 header without an L2 header
- Credit-based flow control
- Protocol also provides resiliency benefits
  - Low-latency FEC (see 25Gb Ethernet Consortium)
  - Link level retry to tolerate transient errors
  - Lane degrade to tolerate hard failures



# Extreme Heterogeneity Presents Significant Resource Management Challenges



**OS/RTS Design:** HW resources will become more complex/diverse.

- OS/RTS must be efficient and sustainable for an increasingly diverse set of hardware components
- Must provide capability for dynamic discovery of resources as power/energy constraints impose restrictions on availability

**Decentralized resource management:** New scalable methods of coordinating resources must be developed that allow policy decisions and mechanisms to co-exist throughout the system.

- HW resources are becoming inherently adaptive, making it increasingly complex to understand and evaluate optimal execution and utilization
- System software must be enhanced to coordinate resources across multiple levels and disparate devices in the system
- Must leverage cohesive integration of performance introspection and programming system abstractions to provide more adaptive execution

**Autonomous resource optimization:** Responsibility for efficient use of resources must shift from the user to the system software.

- Need more automated methods using machine learning to optimize the performance, energy efficiency, and availability of resources for integrated application workflows
- More sophisticated usage models beyond batch-scheduled, spaced-shared nodes adds significant complexity to the management of system resources

**Map the machine to the application rather than vice-versa**

1

How do we maximize impact on the future of HPC?

2

What gaps do we need to fill and when should we leverage other organizations?

3

How do we balance near-term vs. long-term impact?

4

How do we help craft and support the program's strategies?

Questions about strategy...

# Challenges & Opportunities



1. Deploy and support community tools such as SST and LDMS
2. Exploratory modeling of future architectures which incorporate extreme heterogeneity and beyond Moore components
3. Maximize impact of our testbeds and prototype projects on the program
4. Prediction of application performance on these future systems
5. Resource management and monitoring in the face of extreme heterogeneity
6. Influencing vendor-owned SW stacks and driving towards more open community stacks
7. Leverage multi-agency collaborations targeting innovative government solutions
8. Deepen our partnerships with other labs, academia and industry



*Exceptional service in the national interest*



**Sandia  
National  
Laboratories**

## Next Steps: Tri-lab Mission Applications and Vanguard Phase 2



- **Based on the importance of Astra and Arm technology for future architectures, and request from NNSA, we accelerated our schedule for moving to the classified network**
  - Astra is now a platform element of the ASC ATDM Level 1 milestone in FY20, which involves Tri-lab mission codes running on next-generation platforms
  - Tri-lab applications are already being ported, tested and optimized on Astra prior to the migration to classified
- **Released RFI for Vanguard Phase 2 in coordination with the Tri-labs**
- **We are engaging Los Alamos and Lawrence Livermore early in this process**
  - Review of Request for Information (RFI) responses
  - Participate in follow-on meetings with vendors
  - Participate in down-select and definition of technology targets for Vanguard 2
  - Review development of RFP (if necessary)
- **The RFI will seek both near-term (2020-2022) technologies to prototype, as well as 3-5 years-out target opportunities**
  - Leverage exploration of “Pathforward2” technologies as well

# SST: Differentiating HW Sim

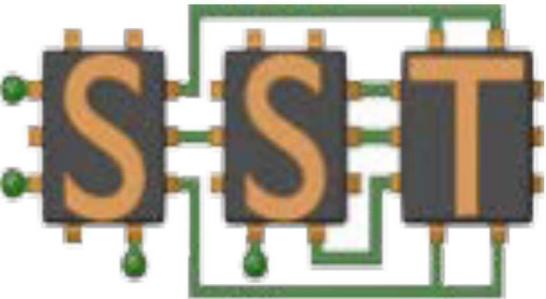
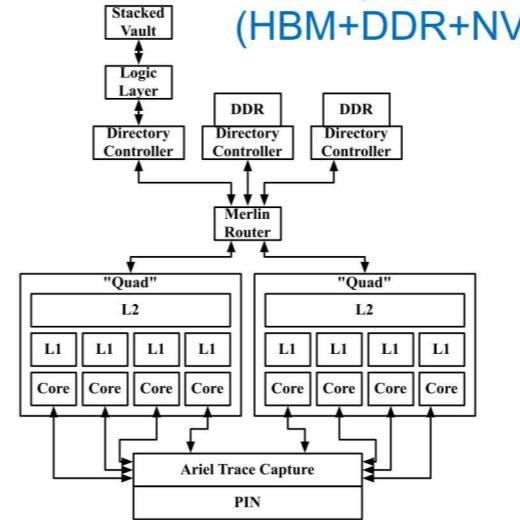
## New opportunities

- IARPA project with LBL and PNNL
- Multi-Agency 'Project 38'
- Expanding Vendor Use
- University Collaborators

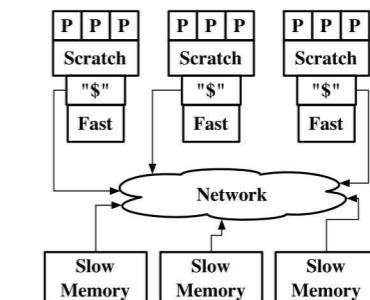
## Pointing to the future

- Next Gen. Interconnects
- Accelerators (GPGPU-Sim)
- Beyond Moore
- Design Optimization

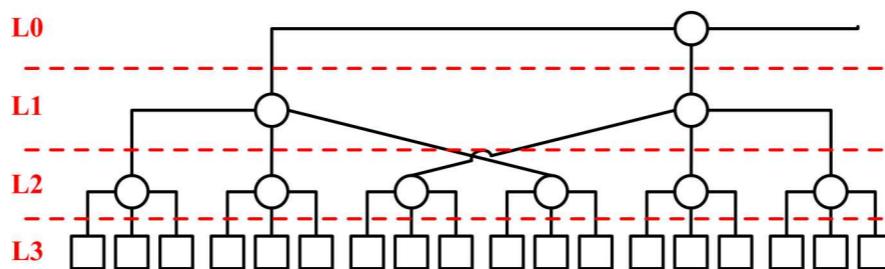
## Multi-Level Memory (HBM+DDR+NV)



## Disaggregated Memory



## Photonic Network Topology & Routing



# ASC Co-design Foundational Principles



## Mission

Ensure the nuclear weapons codes provide preeminent support toward NNSA mission-critical activities ensuring a stockpile that is safe, secure, and effective, through aggressive advancements in our ability to use advanced computational resources

## Vendor Engagement

Proactively engage the U.S. computer industry to influence commodity hardware and software capabilities and gain a deeper understanding of architectural trends and their implications for the nuclear weapons code base

## Research

Develop a focused research agenda among designers of hardware, applications, and programming environments to tackle the interdependent challenges that next-generation extreme-scale platforms present to ASC applications

## Partnerships

Leverage the strength of vendors, academia, and the national laboratories in pursuit of a sustainable High Performance Computing eco-system

# Symphony Feedback



Symphony produces a set of weightings (i.e., a recipe) to use the building blocks as an approximation of the target workload. This recipe can be used to:

1. Understand the relative importance of applications in cluster cycle usage, for prioritizing applications in test suites, and for outreach to software development teams for collaboration
2. Understand which of various mini-apps, computation kernels, and microbenchmarks is a better proxy for gauging workload performance
3. Understand the accuracy of using combinations of building blocks as proxies for the performance of individual applications or workloads
4. Reproduce the workload of a cluster over a period of time using a simplified combination of building blocks for operating system patch testing, performance tool testing, future production cluster procurements, etc.

Symphony data can be used for many exploratory research topics.