

Update on the Exascale Computing Project (ECP)

Paul Messina, ECP Director

HPC User Forum

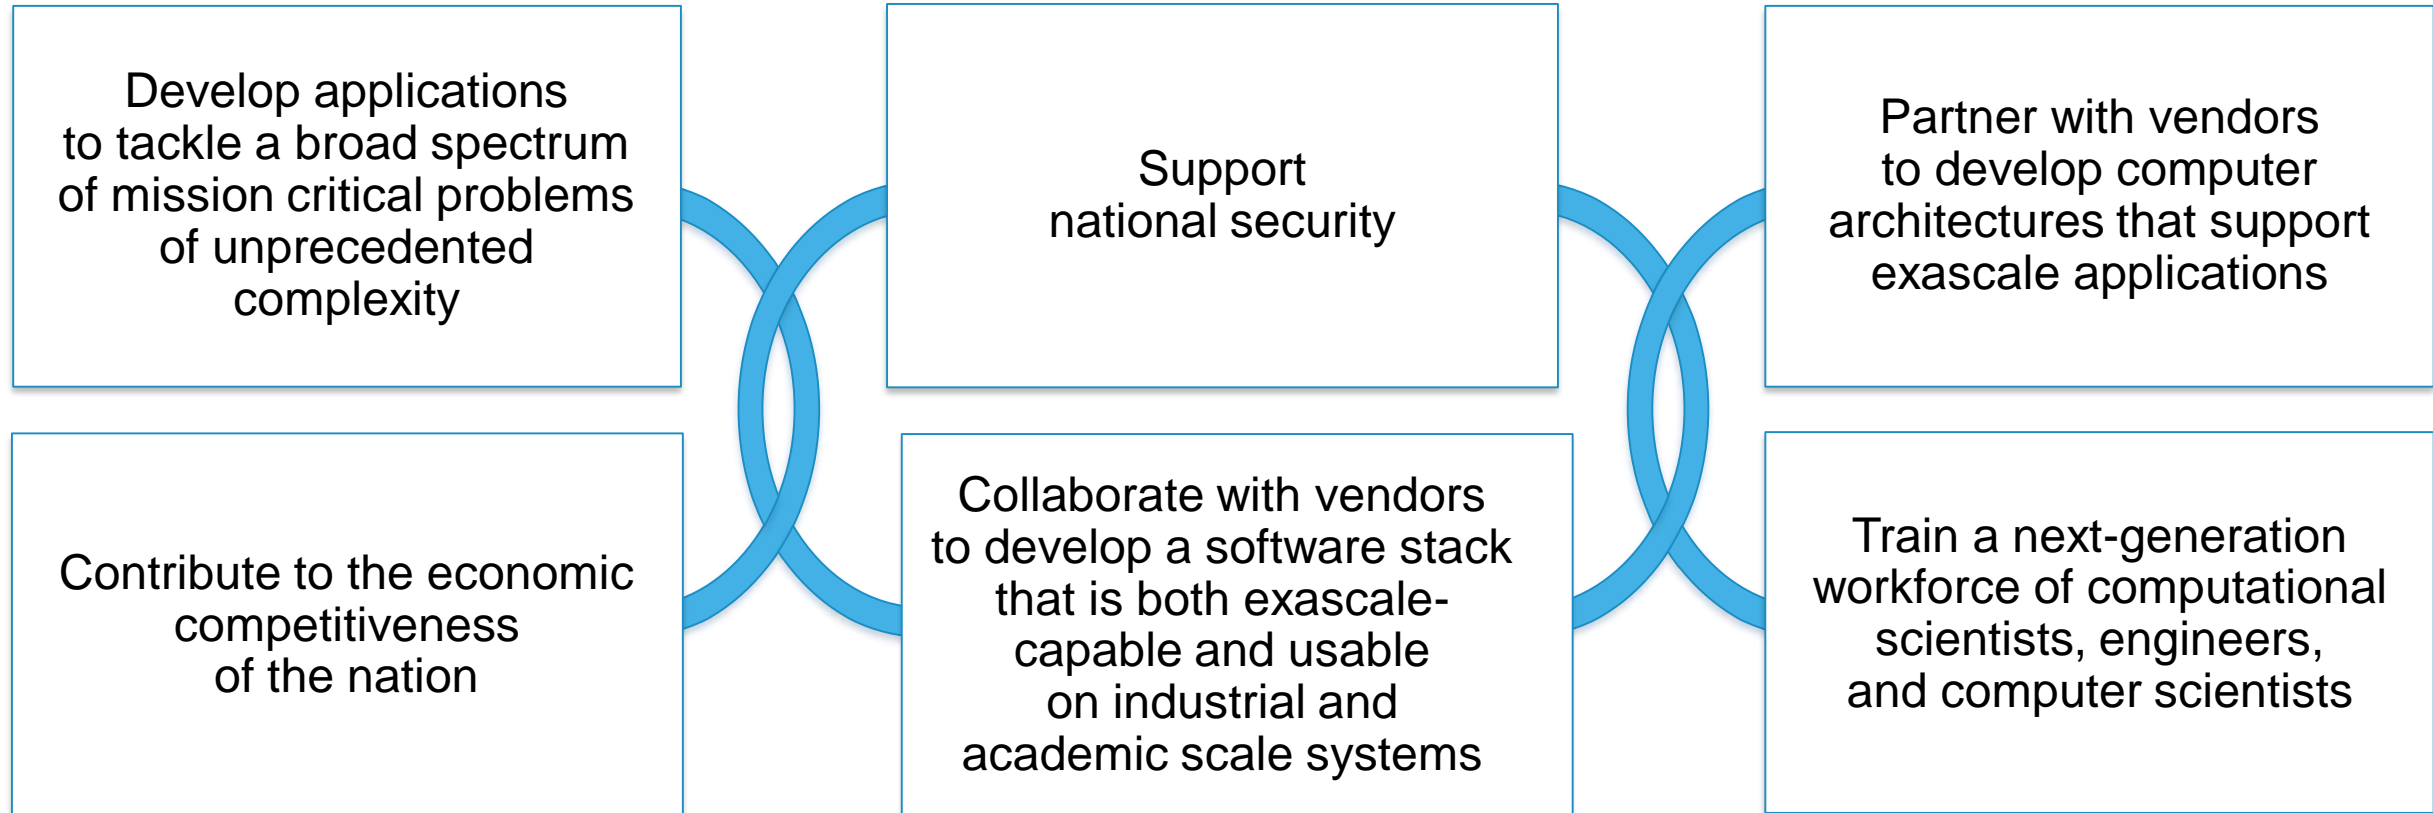
Santa Fe, New Mexico

April 18, 2017



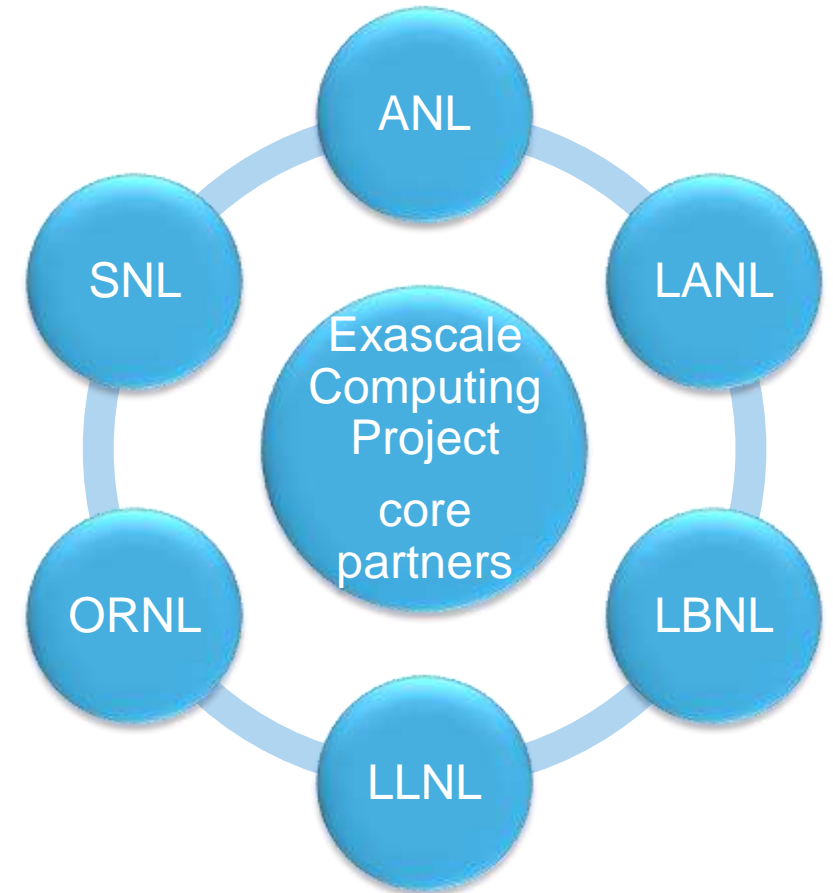
EXASCALE COMPUTING PROJECT

ECP Aims to Transform the HPC Ecosystem and Make Major Contributions to the Nation



ECP is a Collaboration Among Six Labs

- ECP project draws from the Nation's 6 premier computing national laboratories
- An MOA for ECP was signed by each Laboratory Director defining roles and responsibilities
- Project team has decades of experience advancing HPC and deploying first generation HPC systems
- Leadership team expertise spans all ECP activity areas



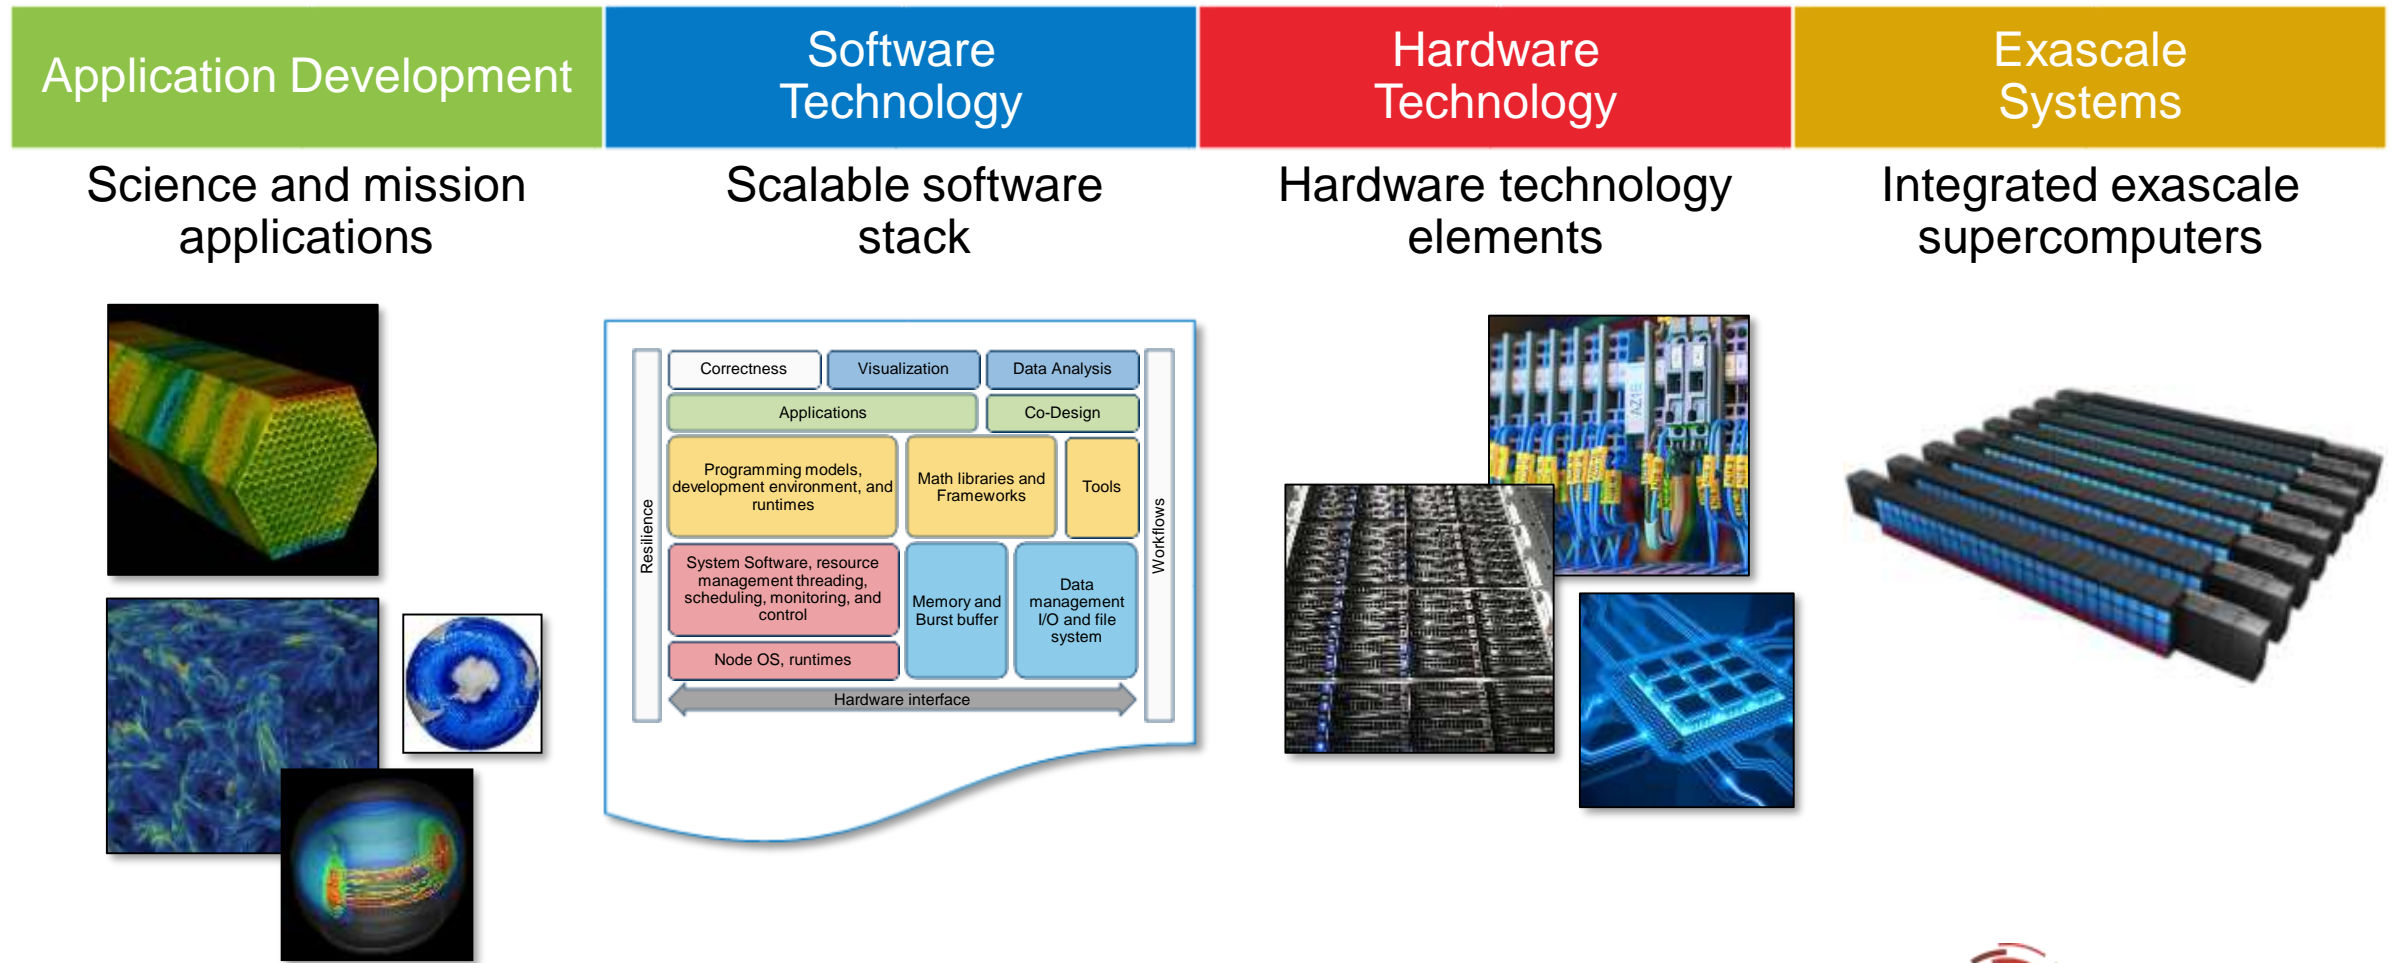
Four Key Technical Challenges Must be Addressed by the ECP to Deliver Capable Exascale Computing

- Parallelism a thousand-fold greater than today's systems
- Memory and storage efficiencies consistent with increased computational rates and data movement requirements
- Reliability that enables system adaptation and recovery from faults in much more complex system components and designs
- Energy consumption beyond current industry roadmaps, which would be prohibitively expensive at this scale

What Has Not Changed?

- Scope: ECP's work encompasses
 - applications,
 - system software,
 - hardware technologies and architectures, and
 - workforce development to meet scientific and national security mission needs.
- The project is executed with a holistic co-design and integration approach

ECP Has Formulated a Holistic Approach That Uses Co-Design and Integration to Achieve Capable Exascale



The New ECP Plan of Record

- A 7-year project that follows the **holistic/co-design** approach, that runs through 2023 (including 12 months of schedule contingency)
- Enable an initial exascale system based on **advanced architecture** delivered in 2021
- Enable **capable exascale** systems, based on ECP R&D, delivered in 2022 and deployed in 2023 as part of an NNSA and SC facility upgrades

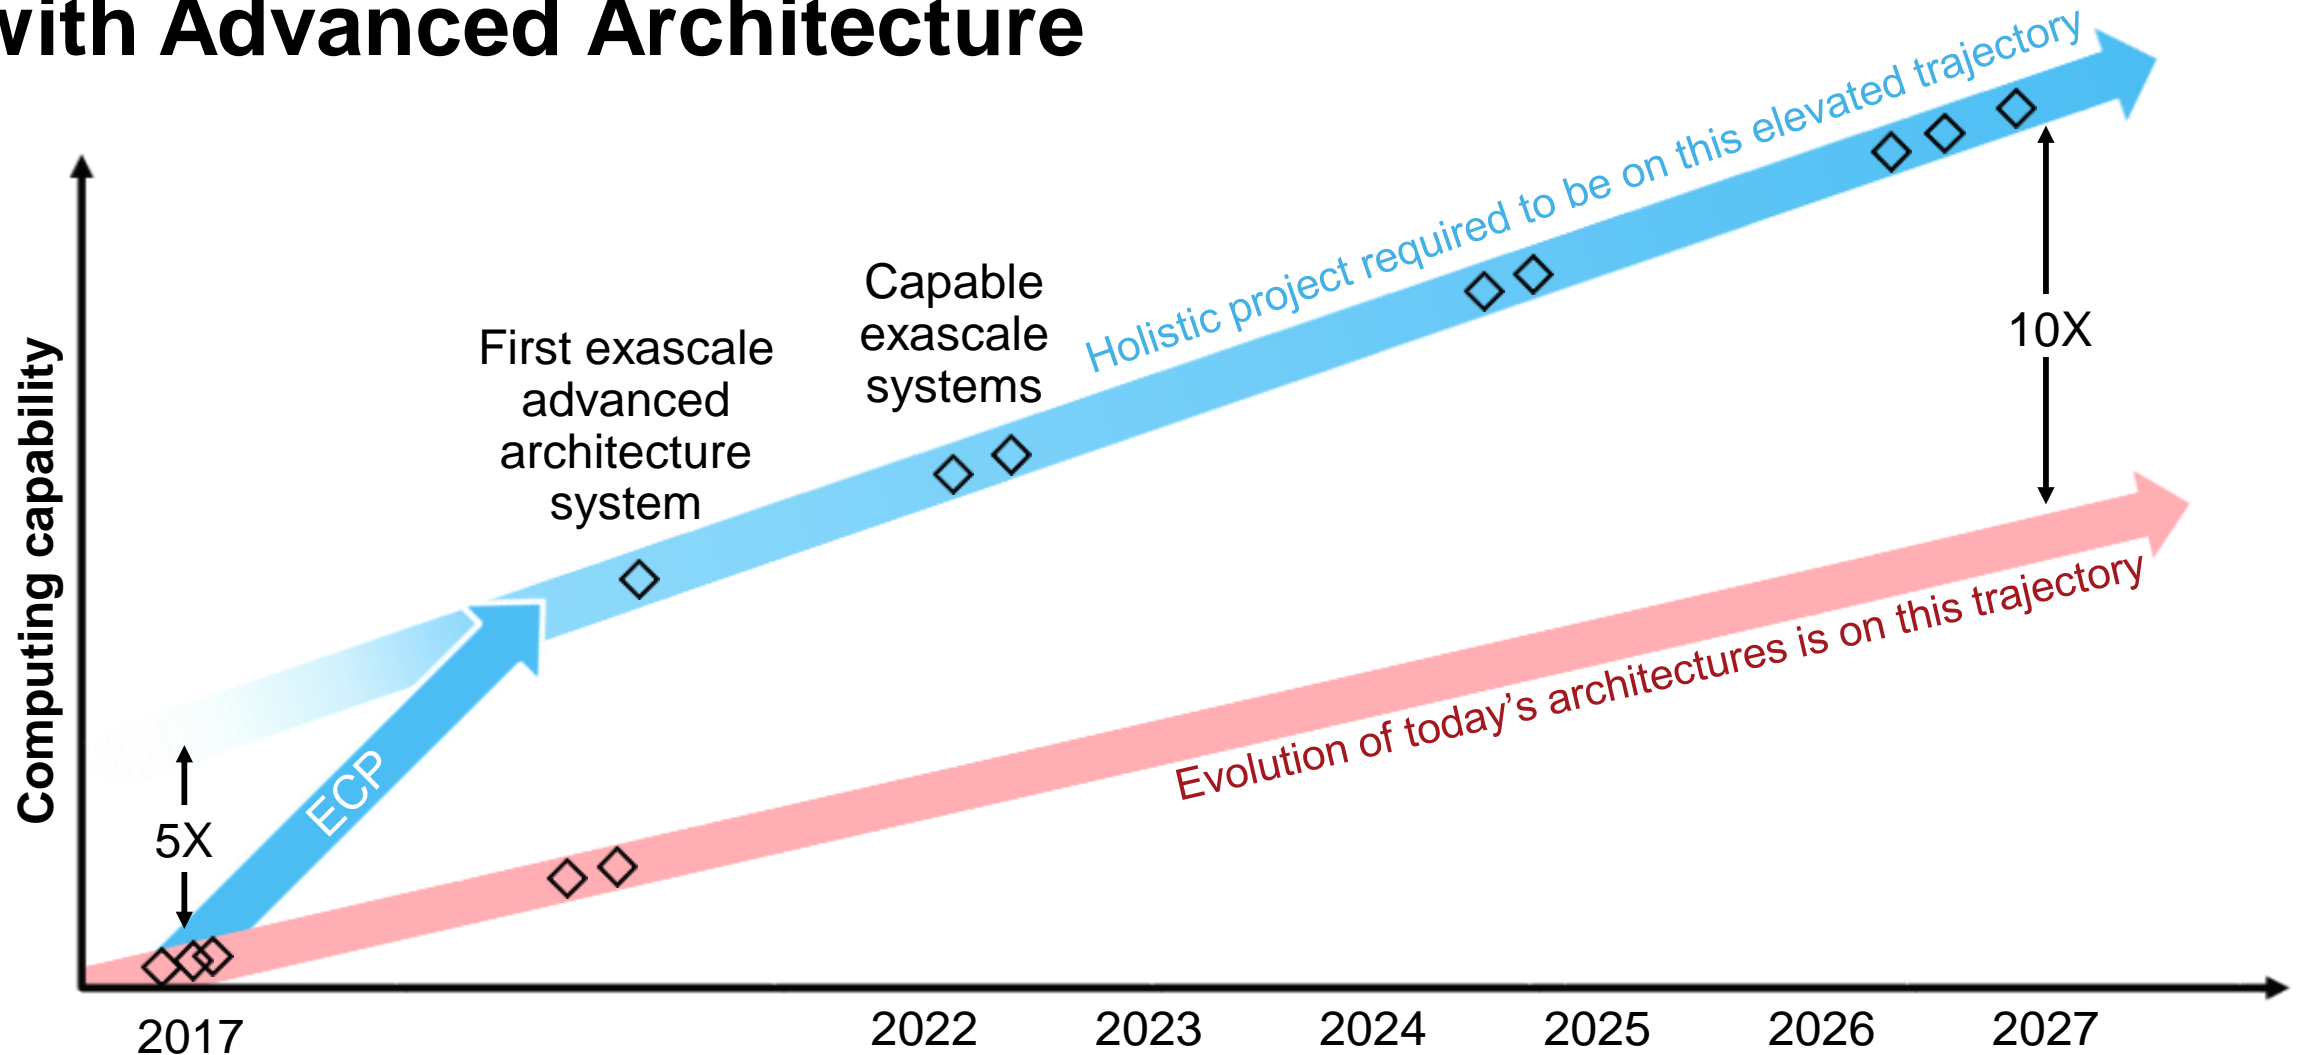
Acquisition of the exascale systems is outside of the ECP scope, will be carried out by DOE-SC and NNSA-ASC supercomputing facilities

What Is a **Capable** Exascale Computing System?

- Delivers 50× the performance of today's 20 PF systems, supporting applications that deliver high-fidelity solutions in less time and address problems of greater complexity
- Operates in a power envelope of 20–30 MW
- Is sufficiently resilient (perceived fault rate: $\leq 1/\text{week}$)
- Includes a software stack that supports a broad spectrum of applications and workloads

This ecosystem will be developed using a co-design approach to deliver new software, applications, platforms, and computational science capabilities at heretofore unseen scale


Transition to Higher Trajectory with Advanced Architecture



Capable Exascale System Applications Will Deliver Broad Coverage of 6 Strategic Pillars

National security

Stockpile stewardship



Energy security

Turbine wind plant efficiency

Design and commercialization of SMRs

Nuclear fission and fusion reactor materials design

Subsurface use for carbon capture, petro extraction, waste disposal

High-efficiency, low-emission combustion engine and gas turbine design

Carbon capture and sequestration scaleup

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Urban planning

Reliable and efficient planning of the power grid

Seismic hazard risk assessment



Scientific discovery

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Plasma wakefield accelerator design

Light source-enabled analysis of protein and molecular structure and design

Find, predict, and control materials and properties

Predict and control stable ITER operational performance

Demystify origin of chemical elements

Earth system

Accurate regional impact assessments in Earth system models

Stress-resistant crop analysis and catalytic conversion of biomass-derived alcohols

Metagenomics for analysis of biogeochemical cycles, climate change, environmental remediation


Health care

Accelerate and translate cancer research



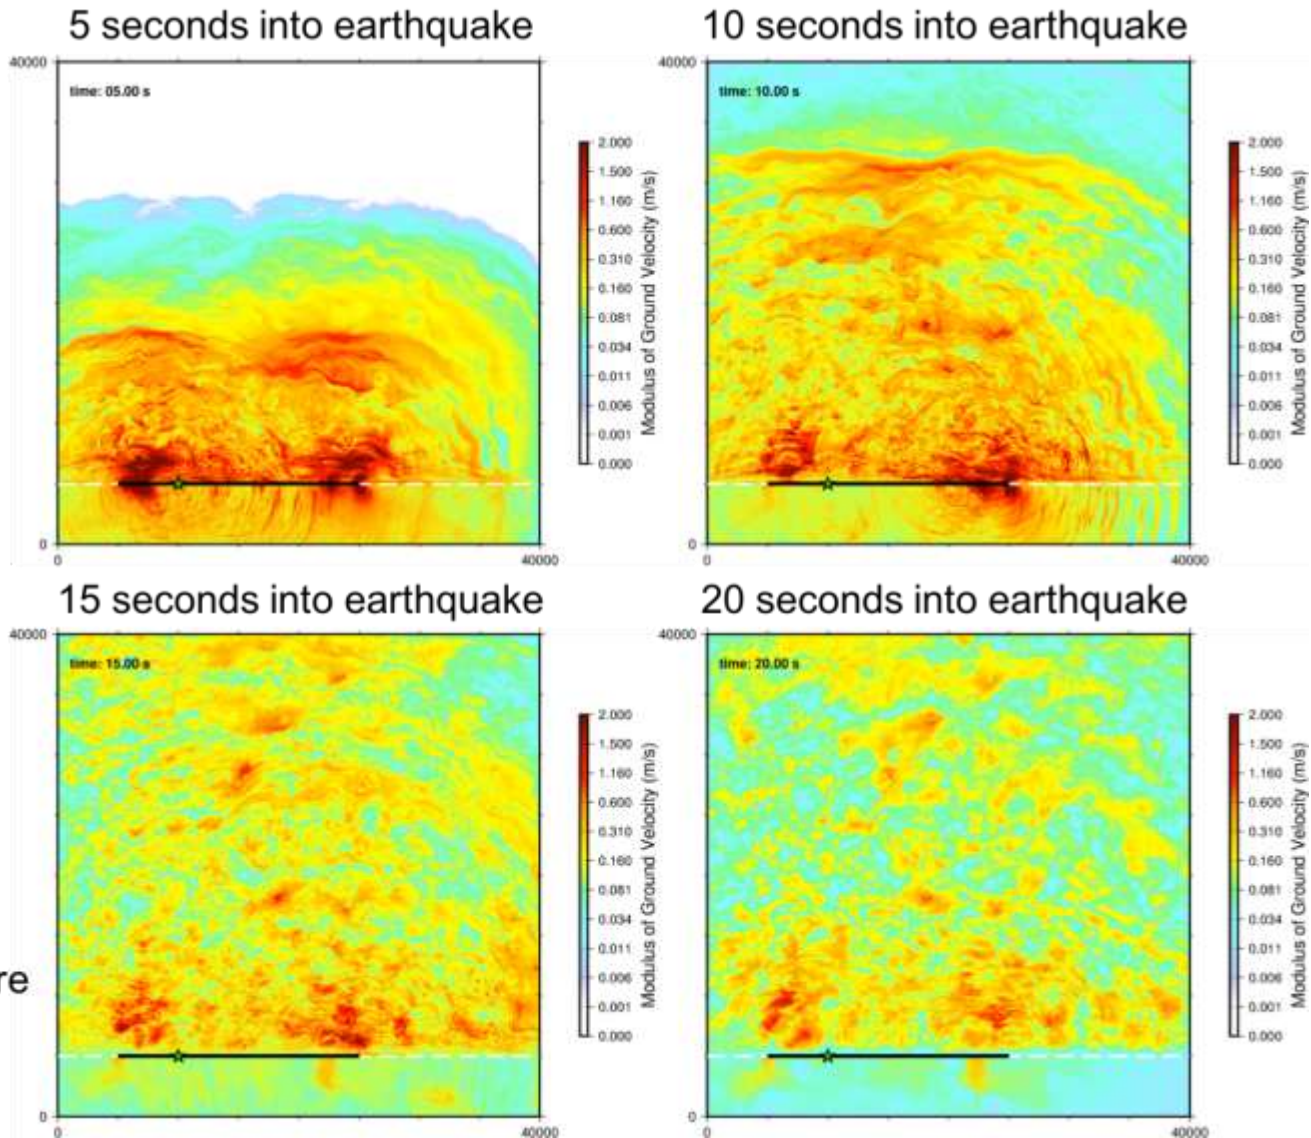
Enabling GAMESS for Exascale Computing in Chemistry & Materials

Heterogeneous Catalysis on Mesoporous Silica Nanoparticles (MSN)



- MSN: highly effective and selective heterogeneous catalysts for a wide variety of important reactions
- MSN selectivity is provided by “gatekeeper” groups (red arrows) that allow only desired reactants A to enter the pore, keeping undesirable species B from entering the pore
- Presence of solvent adds complexity: Accurate electronic structure calculations are needed to deduce the reaction mechanisms, and to design even more effective catalysts
- Narrow pores (3-5 nm) create a diffusion problem that can prevent product molecules from exiting the pore, hence the reaction dynamics must be studied on a sufficiently realistic cross section of the pore
- Adequate representation of the MSN pore requires ~10-100K thousands of atoms with a reasonable basis set; reliably modeling an entire system involves >1M basis functions
- Understanding the reaction mechanism *and* dynamics of the system(s) is beyond the scope of current hardware and software – requiring capable exascale


High Performance, Multidisciplinary Simulations for Regional Scale Earthquake Hazard and Risk Assessments



- Ability to accurately simulate the complex processes associated with major earthquakes will become a reality with capable exascale
- Simulations offer a transformational approach to earthquake hazard and risk assessments
- Dramatically increase our understanding of earthquake processes
- Provide improved estimates of the ground motions that can be expected in future earthquakes
- Time snapshots (map view looking at the surface of the earth) of a simulation of a rupturing earthquake fault and propagation seismic waves

Exascale Predictive Wind Plant Flow Physics Modeling

Understanding Complex Flow Physics of Whole Wind Plants



- Must advance fundamental understanding of flow physics governing whole wind plant performance: wake formation, complex terrain impacts, turbine-turbine interaction effects
- Greater use of U.S. wind resources for electric power generation (~30% of total) will have profound societal and economic impact: strengthening energy security and reducing greenhouse-gas emissions
- Wide-scale deployment of wind energy on the grid without subsidies is hampered by significant plant-level energy losses by turbine-turbine interactions in complex terrains
- Current methods for modeling wind plant performance are not reliable design tools due to insufficient model fidelity and inadequate treatment of key phenomena
- Exascale-enabled predictive simulations of wind plants composed of $O(100)$ multi-MW wind turbines sited within a 10 km x 10 km area with complex terrains will provide validated "ground truth" foundation for new turbine design models, wind plant siting, operational controls and reliably integrating wind energy into the grid

Optimizing Stochastic Grid Dynamics at Exascale

Intermittent renewable sources, electric vehicles, and smart loads will vastly change the behavior of the electric power grid and impose new stochastics and dynamics that the grid is not designed for nor can easily accommodate

- Optimizing such a stochastic and dynamic grid with sufficient reliability and efficiency is a monumental challenge
- Not solving this problem appropriately or accurately could result in either significantly higher energy cost, or decreased reliability inclusive of more blackouts, or both
- Power grid data are clearly showing the trend towards dynamics that cannot be ignored and would invalidate the quasi-steady-state assumption used today for both emergency and normal operation
- The increased uncertainty and dynamics severely strains the analytical workflow that is currently used to obtain the cheapest energy mix at a given level of reliability
- The current practice is to keep the uncertainty, dynamics and optimization analysis separate, and then to make up for the error by allowing for larger operating margins
- The cost of these margins is estimated by various sources to be in \$5-15B per year for the entire United States
- The ECP grid dynamics application can result in the best achievable bounds on these errors and thus resulting in potentially billions of dollars a year in savings.



ECP Co-Design Centers

- **A Co-Design Center for Online Data Analysis and Reduction at the Exascale (CODAR)**

- **Motifs:** Online data analysis and reduction
- Address growing disparity between simulation speeds and I/O rates rendering it infeasible for HPC and data analytic applications to perform offline analysis. Target common data analysis and reduction methods (e.g., feature and outlier detection, compression) and methods specific to particular data types and domains (e.g., particles, FEM)

- **Block-Structured AMR Co-Design Center (AMReX)**

- **Motifs:** Structured Mesh, Block-Structured AMR, Particles
- New block-structured AMR framework (AMReX) for systems of nonlinear PDEs, providing basis for temporal and spatial discretization strategy for DOE applications. Unified infrastructure to effectively utilize exascale and reduce computational cost and memory footprint while preserving local descriptions of physical processes in complex multi-physics algorithms

- **Center for Efficient Exascale Discretizations (CEED)**

- **Motifs:** Unstructured Mesh, Spectral Methods, Finite Element (FE) Methods
- Develop FE discretization libraries to enable unstructured PDE-based applications to take full advantage of exascale resources without the need to “reinvent the wheel” of complicated FE machinery on coming exascale hardware

- **Co-Design Center for Particle Applications (CoPA)**

- **Motif(s):** Particles (involving particle-particle and particle-mesh interactions)
- Focus on four sub-motifs: short-range particle-particle (e.g., MD and SPH), long-range particle-particle (e.g., electrostatic and gravitational), particle-in-cell (PIC), and additional sparse matrix and graph operations of linear-scaling quantum MD

- **Combinatorial Methods for Enabling Exascale Applications (ExaGraph)**

- **Motif(s):** Graph traversals; graph matching; graph coloring; graph clustering, including clique enumeration, parallel branch-and-bound, graph partitioning
- Develop methods and techniques for efficient implementation of key combinatorial (graph) algorithms that play a critical enabling role in numerous scientific applications. The irregular memory access nature of these algorithms makes them difficult algorithmic kernels to implement on parallel systems

Exascale Proxy Applications Suite

Objectives and Scope

- Assemble and curate a proxy app suite composed of proxies developed by other ECP projects that represent the most important features (especially performance) of exascale applications.
- Improve the quality of proxies created by ECP and maximize the benefit received from their use. Set standards for documentation, build and test systems, performance models and evaluations, etc.
- Collect requirements from app teams. Assess gaps between ECP applications and proxy app suite. Ensure proxy suite covers app motifs/requirements.
- Coordinate use of proxy apps in the co-design process. Connect consumers to producers. Promote success stories and correct misuse of proxies.

Links to Other ECP Projects

- **Application Assessment Project:** Cooperatively assess and quantitatively compare applications and proxy apps.
- **Design Space Evaluation Team:** Will need proxies specially adapted for hardware simulators.
- **Path Forward Vendors:** Evaluate needs and provides proxies & support. Review proxy app usage and results.
- **Application Development Projects & Co-Design Centers:** Producers of proxy apps. Close the loop with lessons learned from proxies.
- **Software Technology Projects:** Consumers of proxy apps. Use proxies to understand app requirements and to test and evaluate proposed ST offerings.

Risks and Challenges

- Proxies are too complex for simulators or HW design activity.
- Proxies fail to accurately represent parent apps.
- Proxy app authors unable or unwilling to meet quality and/or support standards.
- Full coverage of DOE workload/motifs produces a large and unwieldy suite.
- Proxy apps misused by consumers.
- Inability to balance agility and stability/quality of proxies.

Development Plan

- Release updated versions of the proxy app suite every six months. This cadence allows for improved coverage and changing needs while maintaining needed stability.
- Annually update guidance on quality standards. Increase rigor of standards.
- Meet with each application project at least quarterly to maintain a catalog of their requirements, proxies, and key questions for which they are seeking assistance.
- Publish annual proxy app producer report with requirements and assessment of proxies in comparison to parent apps.
- Publish annual proxy app consumer report with success stories, surveys of how proxy consumers are using proxies, and plans to satisfy any unmet needs.

Description and Scope

- **Software Challenges:** Exploit massive on-node concurrency and handle disruptive architectural changes while working toward predictive simulations that couple physics, scales, analytics, and more
- **Goals:** Improve ECP developer productivity and software sustainability, as key aspects of increasing overall scientific productivity
- **Strategy:** In collaboration with ECP community:
 - **Customize and curate methodologies** for ECP app productivity & sustainability
 - **Create an ECP Application Development Kit** of customizable resources for improving scientific software development
 - **Partner with ECP application teams** on software improvements
 - **Training and outreach** in partnership with DOE computing facilities

IDEAS-ECP team: Catalysts for engaging ECP community on productivity issues

- **Partnerships with ECP applications teams**
 - Understand productivity bottlenecks and improve software practices
- **Collaborate to curate, create, and disseminate software methodologies, processes, and tools** that lead to improved scientific software
 - ECP software technologies projects, applications teams, co-design centers, broader community
 - Software Carpentry-type approach to training for extreme-scale software productivity topics
- **Web-based hub for collaborative content development and delivery**
 - Community-driven collection of resources to help improve scientific software productivity, quality, and sustainability



Phase-1 release:
Summer 2017

History and Accessibility



IDEAS: Interoperable Design of Extreme-scale Application Software

- Project began in Sept 2014 as ASCR/BER partnership to improve application software productivity, quality, and sustainability

Resources: <https://ideas-productivity.org/resources> **Highlights:**

- **WhatIs and HowTo docs:** concise characterizations & best practices

- *What is CSE Software Testing?*
- *What is Version Control?*
- *What is Good Documentation?*
- *How to Write Good Documentation*
- *How to Add and Improve Testing in a CSE Software Project*
- *How to do Version Control with Git in your CSE Project*



- **Webinar series, 2016: Best Practices for HPC Software Developers**

- *What All Codes Should Do: Overview of Best Practices in HPC Software Development*
- *Developing, Configuring, Building, and Deploying HPC Software*
- *Distributed Version Control and Continuous Integration Testing*
- *Testing and Documenting your Code*
- more topics ...



First Year Plan

- **MS1/Y1: Templates for Productivity and Sustainability Improvement Plans (PSIPs) and Progress Tracking Cards**
 - Framework for software teams to identify, plan, and track improvements in productivity and sustainability
- **MS2/Y2: Interviews with Phase-1 ECP applications teams**
 - Understand current software practices, productivity challenges, preferred modes of collaboration, and needs for on-line ECP knowledge exchange
 - Determine prioritized needs for productivity improvement; initiate Phase-1 application partnerships
- **MS3/Y1: Training and outreach to ECP community on productivity and sustainability**
 - Develop and deliver tutorials, webinars, and web content



SIAM CSE17 Conference (Feb 27-Mar 3, 2017)

Premier CSE Conference (largest SIAM meeting, over 1700 attendees)

Invited tutorial: *CSE Collaboration through Software: Improving Software Productivity and Sustainability*

- Presenters: D. Bernholdt, A. Dubey, M. Heroux, A. Klinvex, L.C. McInnes
 - Why Effective Software Practices Are Essential for CSE Projects*
 - An Introduction to Software Licensing*
 - Better (Small) Scientific Software Teams*
 - Improving Reproducibility through Better Software Practices*
 - Testing of HPC Scientific Software*



Slides and audio: <https://doi.org/10.6084/m9.figshare.c.3704287>

Minisymposium and Minisymposium: *Software Productivity and Sustainability for CSE and Data Science*

- Organizers: D. Bernholdt, M. Heroux, D. Katz, A. Logg, L.C. McInnes
- Slides for presentations: <https://doi.org/10.6084/m9.figshare.c.3705946>
 - 8 community presentations, including one from the IDEAS team
 - CSE Software Ecosystems: Critical Instruments of Scientific Discovery*, L.C. McInnes
- Posters: <https://doi.org/10.6084/m9.figshare.c.3703771>
 - 29 posters from the community, including several from members of IDEAS-ECP

Plenary Presentation: *Productive and Sustainable: More Effective CSE*, M. Heroux

- Slides and audio: <https://doi.org/10.6084/m9.figshare.4728697>

Co-Lead PIs: M. Heroux (SNL) and L.C. McInnes (ANL); Partner sites: LBNL, LLNL, LANL, ORNL, Univ of Oregon

Productivity and Sustainability Improvement Planning Tools

Tools for helping a software team to increase software quality while decreasing the effort, time, and cost to develop, deploy, maintain, and extend software over its intended lifetime.

March 2017: Released PSIP templates & instructions: <https://github.com/betterscientificsoftware/PSIP-Tools>

Beginning Phase-1 interviews with ECP applications teams: Identify productivity bottlenecks and priorities.



Practice: Test Coverage	Score (0 – 5):
Score Descriptions	
0	No unit testing.
1	Unit testing used for one example refactored class.
2	Two developers using unit testing process.
3	Unit testing developer documentation completed.
4	All developers trained to write unit tests.
5	New policy established: All refactored code is covered by unit tests.

Productivity and Sustainability Improvement Plan (PSIP): a lightweight iterative workflow to identify, plan, and improve selected practices of a software project.

Other practices: Source management system, documentation, software distribution, issue tracking, developer training, etc.

Argonne Training Program on Extreme-Scale Computing

extremecomputingtraining.anl.gov

Application
deadline
March 10

What?

An intensive two-week program on HPC methodologies applicable to both current and future supercomputers.

Who?

Open to all doctoral students, postdocs, and scientists interested in conducting CS&E research on large-scale computers.

When?

July 30 - August 11, 2017

Where?

Q Center, St Charles, IL (USA)

of Applicants

Call closed March 10. Reviewing 167 applicants from 95 institutions worldwide. Making final selections now. Expect to select up to 70 participants.

65+ participants	100 h courses & hands-on
\$1.25M 2016-2018	\$0 no cost to attend



Domestic airfare, meals and lodging provided

ATPESC 2017
July 30 - August 11

PROGRAM CURRICULUM
Renowned scientists, HPC experts, and leaders serve as lecturers and guide hands-on sessions. ATPESC participants will be granted access to DOE Office of Science User Facilities, which are amongst the most powerful supercomputers in the world.

The core curriculum includes:

- Computer architectures and their predicted evolution
- Programming methodologies effective across a variety of today's supercomputers and that are expected to be applicable to exascale systems
- Numerical algorithms and mathematical software
- Approaches to building community codes for HPC systems
- Data analysis, visualization, I/O, and workflows and tools for Big Data applications
- Performance measurement and debugging tools

Argonne Training Program on Extreme-Scale Computing (ATPESC) is an intensive two-week program focused on HPC methodologies that are applicable to both current and future machine architectures, including exascale systems.

ATPESC provides advanced training to 60 participants. Qualified applicants must have:

- Substantial experience in MPI and/or OpenMP programming
- Used at least one HPC system for a complex application, and
- Plans to conduct CS&E research on large-scale computers.

Admission to the ATPESC program is highly competitive. Participant support is provided, including domestic airfare, meals, and lodging.

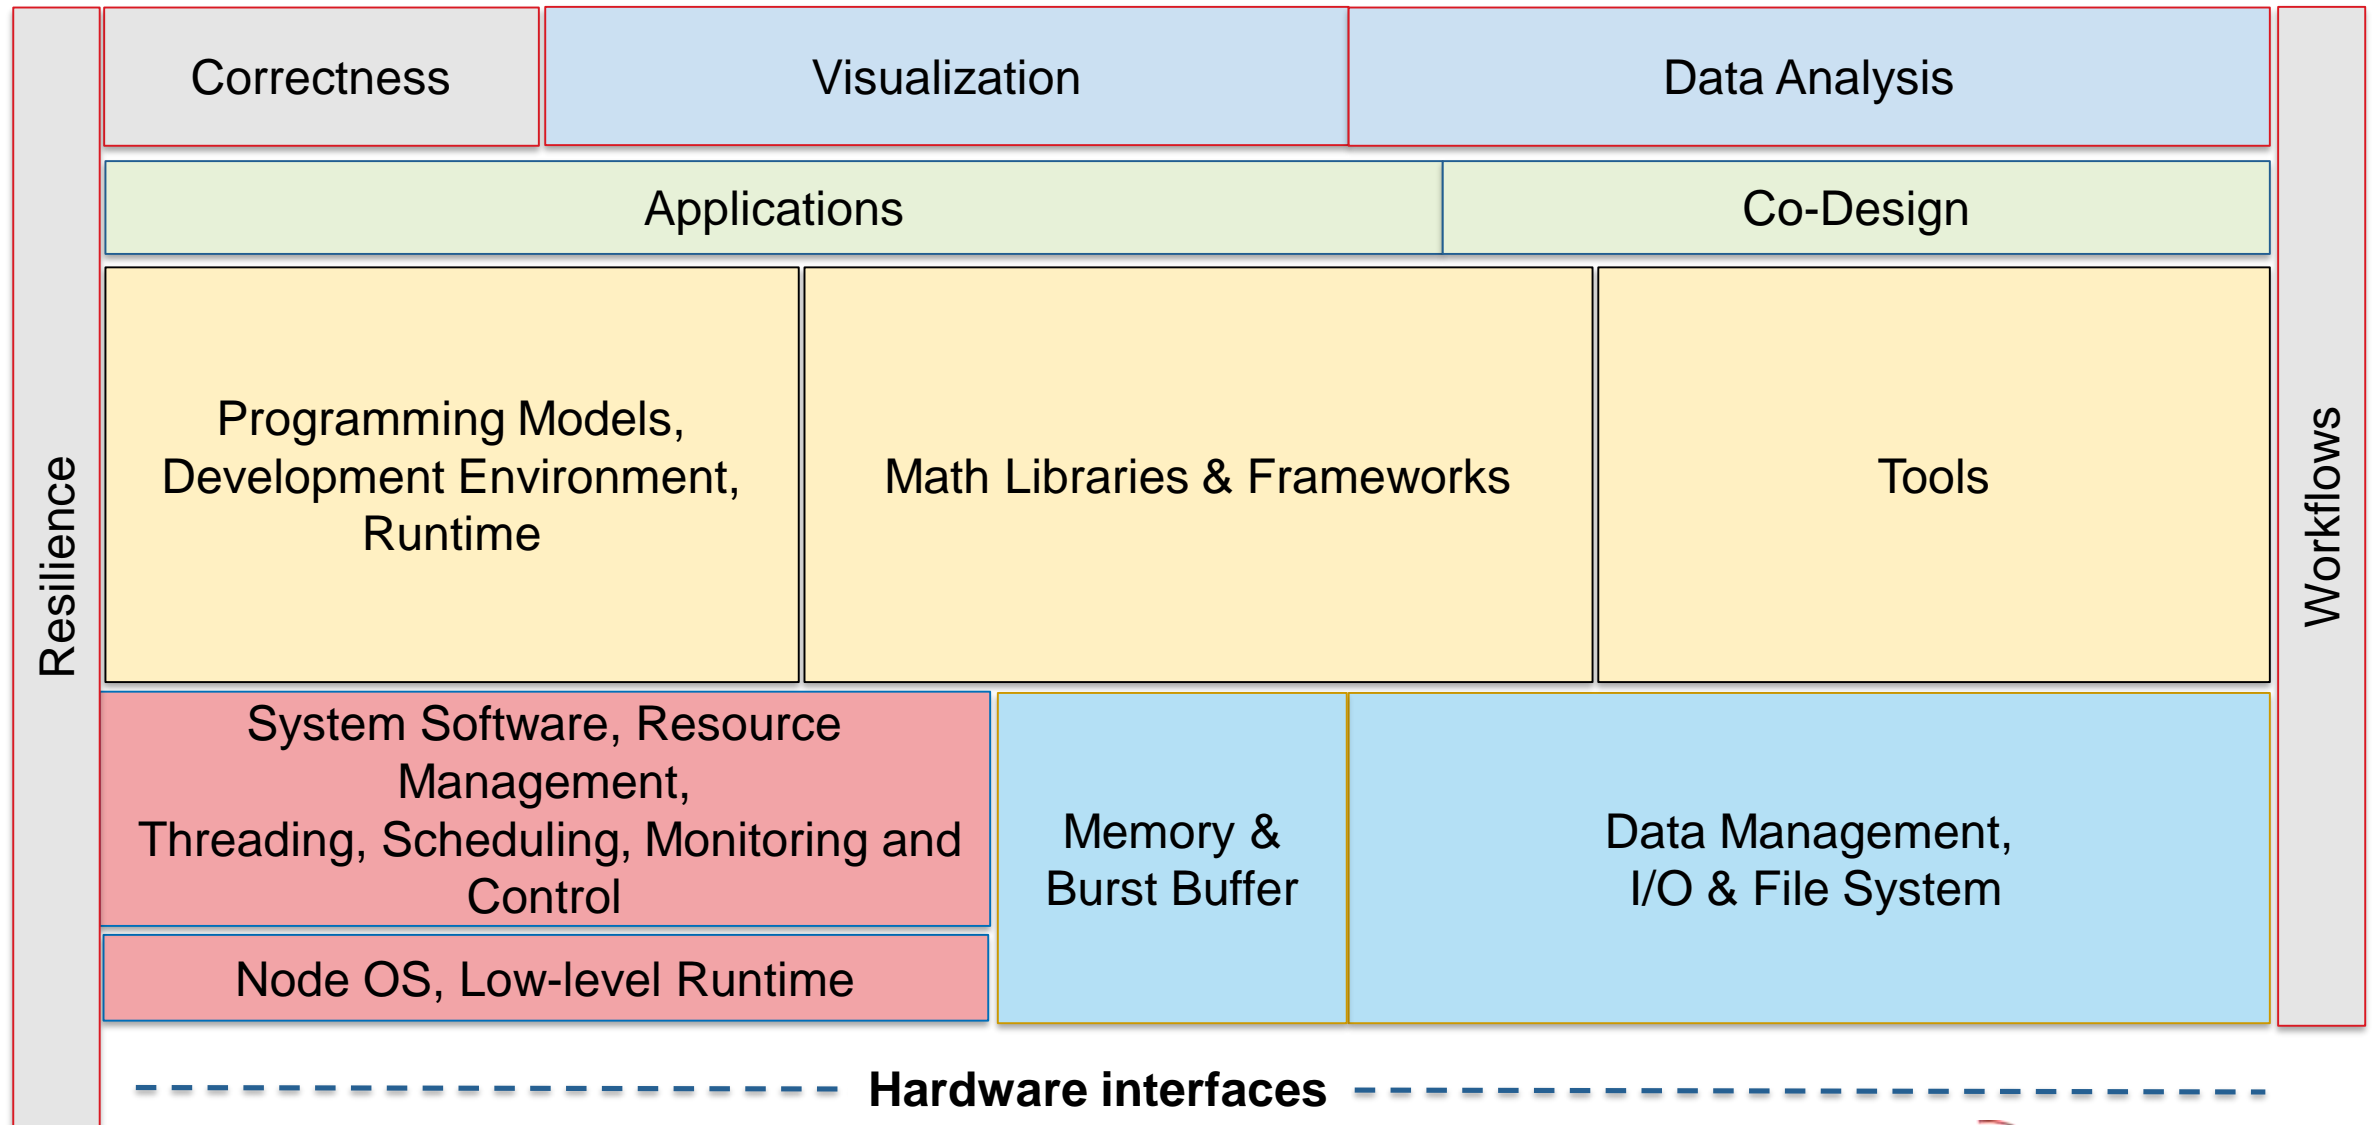
ATPESC is part of the Exascale Computing Project, a collaborative effort of the DOE Office of Science and the National Nuclear Security Administration.

APPLICATION DETAILS
Applications for ATPESC 2017 are due by March 10, 2017. All doctoral students, postdocs, and computational scientists are encouraged to submit applications. To apply, visit the ATPESC website at extremecomputingtraining.anl.gov

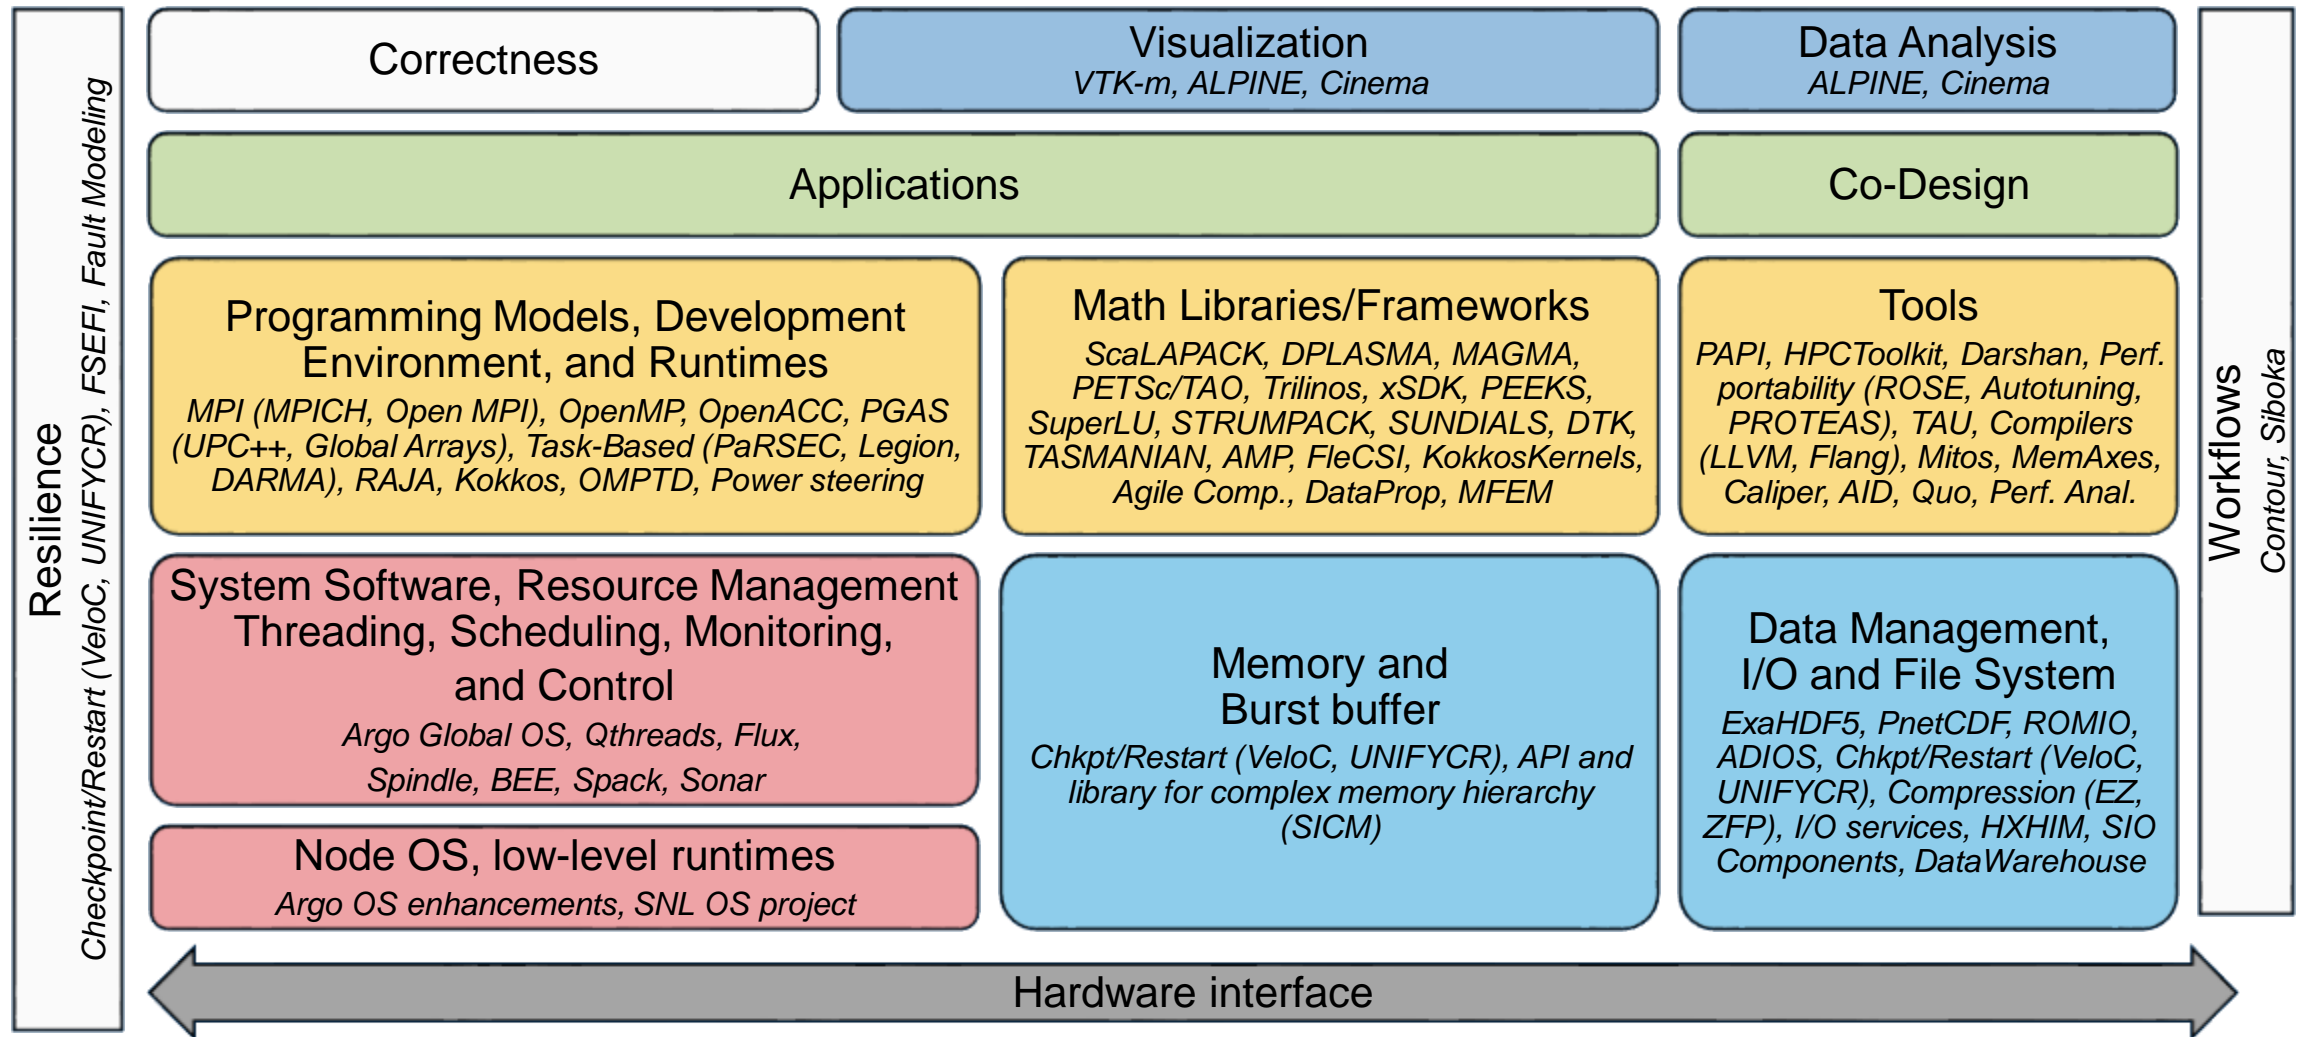
ENERGY NNSA ECOMP Argonne



Conceptual ECP Software Stack



Current Set of ST Projects Mapped to Software Stack



Co-Design in Action

- Design, implement, and demonstrate new prototype APIs to facilitate improved coordination between MPI and OpenMP runtimes for use of threads, message delivery, and rank/endpoint placement. (Due 12/31/2018)
- Execution Plan:
 - Interact with ECP AD teams to identify requirements about runtime coordination.
 - Interact with appropriate ECP ST teams to define coordination APIs for MPI and OpenMP (and other) runtimes.
 - Experiment and demonstrate the benefit running an ECP application or mini-application running on the ECP testbeds

Hardware Technology Activities

- PathForward: support DOE-vendor collaborative R&D activities required to develop exascale systems with at least two diverse architectural features; quote from RFP:
 - PathForward seeks solutions that will improve application performance and developer productivity while maximizing energy efficiency and reliability of exascale systems.
- Design Space Evaluation
 - Apply laboratory architectural analysis capabilities and Abstract Machine Models to PathForward designs to support ECP co-design interactions

PathForward Status

- Five PathForward Projects have received DOE/NNSA HQ's fully signed and executed Coordination and Approval Document (CAP)
 - 4 contracts are fully signed
 - 1 is being routed for signatures by the LLNL Lab Director - Goldstein and his counterparts at the Vendors to fully execute these contracts – anticipated by COB April 14, 2017
- The last PathForward contract is at DOE/NNSA for approval
- The HT leadership team is developing a tiered communication plan to:
 - Introduce the PathForward projects to the rest of ECP, and
 - Establish high-quality co-design collaborations with interested lab-led ECP projects

Design Space Evaluation: Technology Coverage Areas

	DSE (lab POCs)	Memory Technologies	Interconnect/ System Simulators	Analytical Models	Node Simulators	Abstract Machine Models and Proxy Architectures
ANL	Ray Bair		ROSS/CODES		Gem5	
LANL	Jeff Kuehn		CoNCEPTuaL	Byfl		
LBNL	David Donofrio			ExaSAT, Roofline Toolkit	Gem5, OpenSOC	Co-lead AMM v3
LLNL	Robin Goldstone	LiME	ROSS/CODES			
ORNL	Jeff Vetter	Blackcomb		ASPEN		
PNNL	Darren Kerbyson			PALM		
SNL	Rob Hoekstra	SST: VaultSimC, MemHierarchy	SST, Merlin, Ember, Firefly		SST: Miranda, Ariel, Gem5	Co-lead AMM v3

ECP teams begin work on Office of Science systems via early access, ALCC awards, and testbeds

Access to Office of Science Systems

NERSC's Cori II



Cori II early access

- Shared access with other users
- Intel Xeon Phi (KNL) nodes
- Through June 30, 2017

OLCF's Titan



2016 ALCC Reserve Award to ECP

- 13M Titan-core hours
- March 30 to June 30, 2017
- 10 AD, 6 ST, and 1 HT teams are getting access

Access to testbeds supported with ECP funds

ALCF's Theta



ECP allocation is 37+M hours on Theta for 2017

- Intel Xeon Phi cpus
- 3,624 nodes
- 64 core processor per node

OLCF's Summitdev



ECP allocation is the equivalent of 18 nodes on Summitdev

- Power 8+
- NVIDIA Pascal GPGPUs
- NVLINK1

ECP submitted an application to the 2017 ASCR Leadership Computing Challenge (ALCC) for time of Titan, Mira, Theta, Cori, and Edison

Communications and Outreach

- We have a website
 - www.ExascaleProject.org
- A newsletter *The ECP Update* has been launched
 - <https://exascaleproject.org/newsletter/ecp-update-1/>
- First PI meeting November 29 – December 1, 2016
 - 80+ PIs
- First Annual Meeting January 31 – February 2, 2017
 - 450+ participants
 - Focused on co-design and integration planning
- Industry Council established, first meeting March 6-7, 2017



EXASCALE COMPUTING PROJECT



DOE LABORATORIES & AGENCY PARTNERS

22



PRIVATE SECTOR PARTNERS

9



UNIVERSITY RESEARCH PARTNERS

39



INDUSTRY COUNCIL MEMBERS

18



THE ECP ECOSYSTEM

- ◆ 800 Researchers
- ◆ 26 Application Development Projects
- ◆ 66 Software Development Projects
- ◆ 5 Co-Design Centers



U.S. DEPARTMENT OF ENERGY

Office of Science



Thank you!



EXASCALE COMPUTING PROJECT