

The Exascale Computing Project

Paul Messina, ECP Project Director

Ad Hoc Big Data Task Force of the
NASA Advisory Council Science Committee

Washington, DC via WebEx

March 6, 2017



EXASCALE COMPUTING PROJECT



U.S. DEPARTMENT OF
ENERGY

Office of
Science

www.ExascaleProject.org

What is the Exascale Computing Project (ECP)?

- As part of the National Strategic Computing initiative, ECP was established to accelerate delivery of a **capable exascale** computing system that integrates hardware and software capability to deliver approximately 50 times more performance than today's 20-petaflops machines on mission critical applications.
 - DOE is a lead agency within NSCI, along with DoD and NSF
 - Deployment agencies: NASA, FBI, NIH, DHS, NOAA
- ECP's work encompasses
 - applications,
 - system software,
 - hardware technologies and architectures, and
 - workforce development to meet scientific and national security mission needs.



What is the Exascale Computing Project?

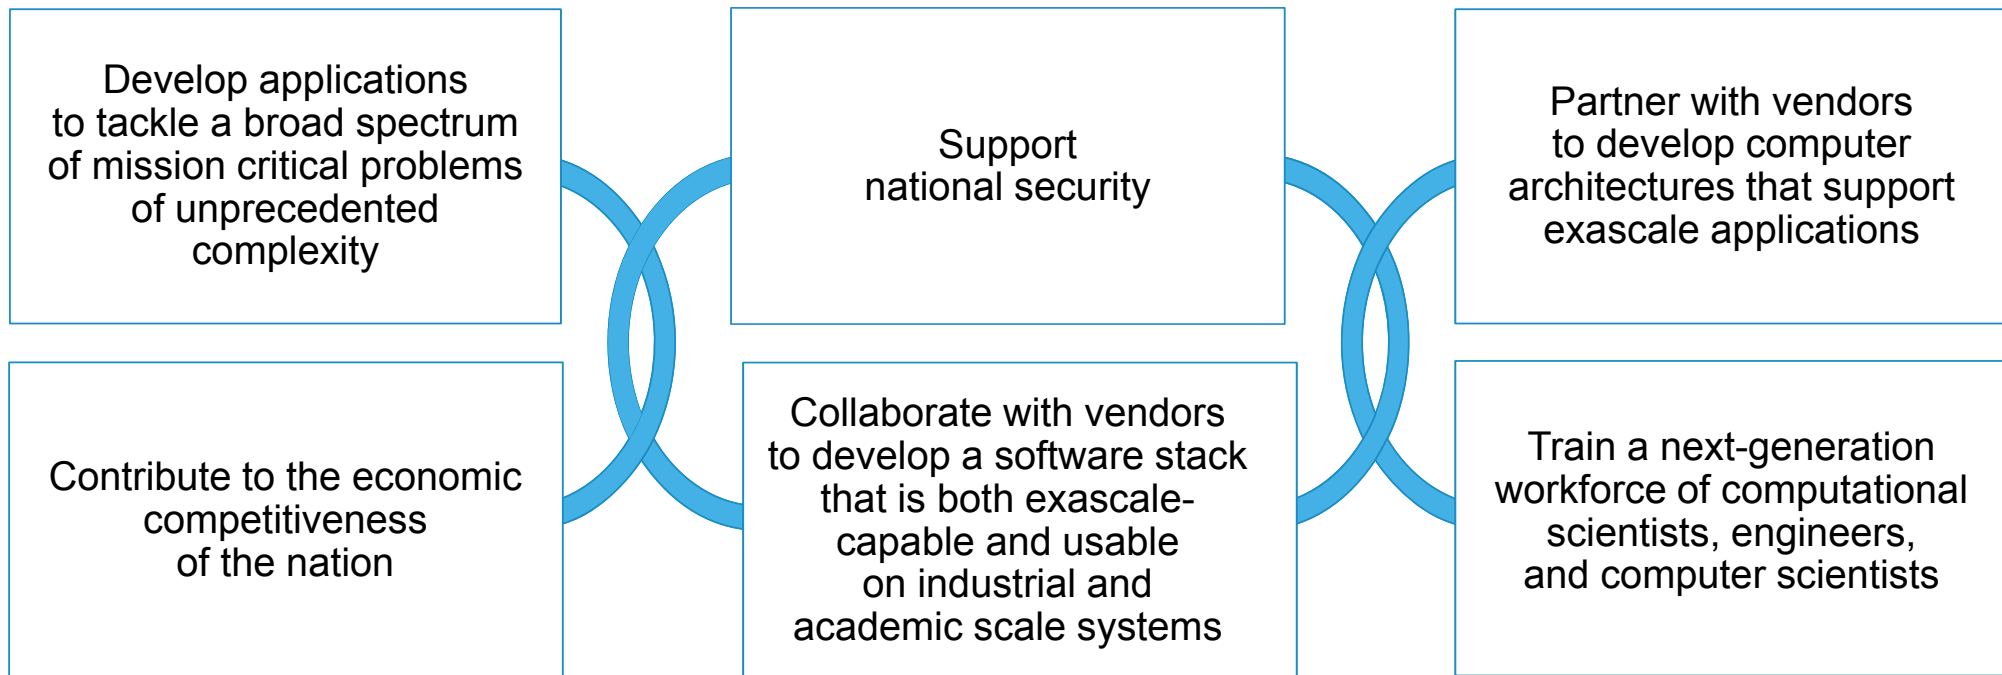
- A collaborative effort of two US Department of Energy (DOE) organizations:
 - Office of Science (DOE-SC)
 - National Nuclear Security Administration (NNSA)
- A 7-year project to accelerate the development of a **capable** exascale ecosystem
 - Led by DOE laboratories
 - Executed in collaboration with academia and industry
 - emphasizing sustained performance on relevant applications



A **capable** exascale computing system will have a well-balanced ecosystem (software, hardware, applications)

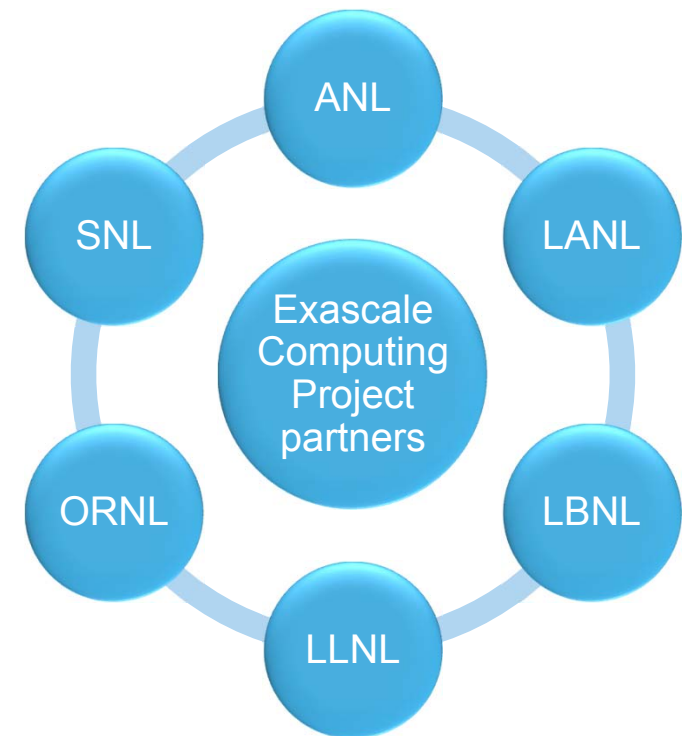


ECP aims to transform the HPC ecosystem and make major contributions to the nation



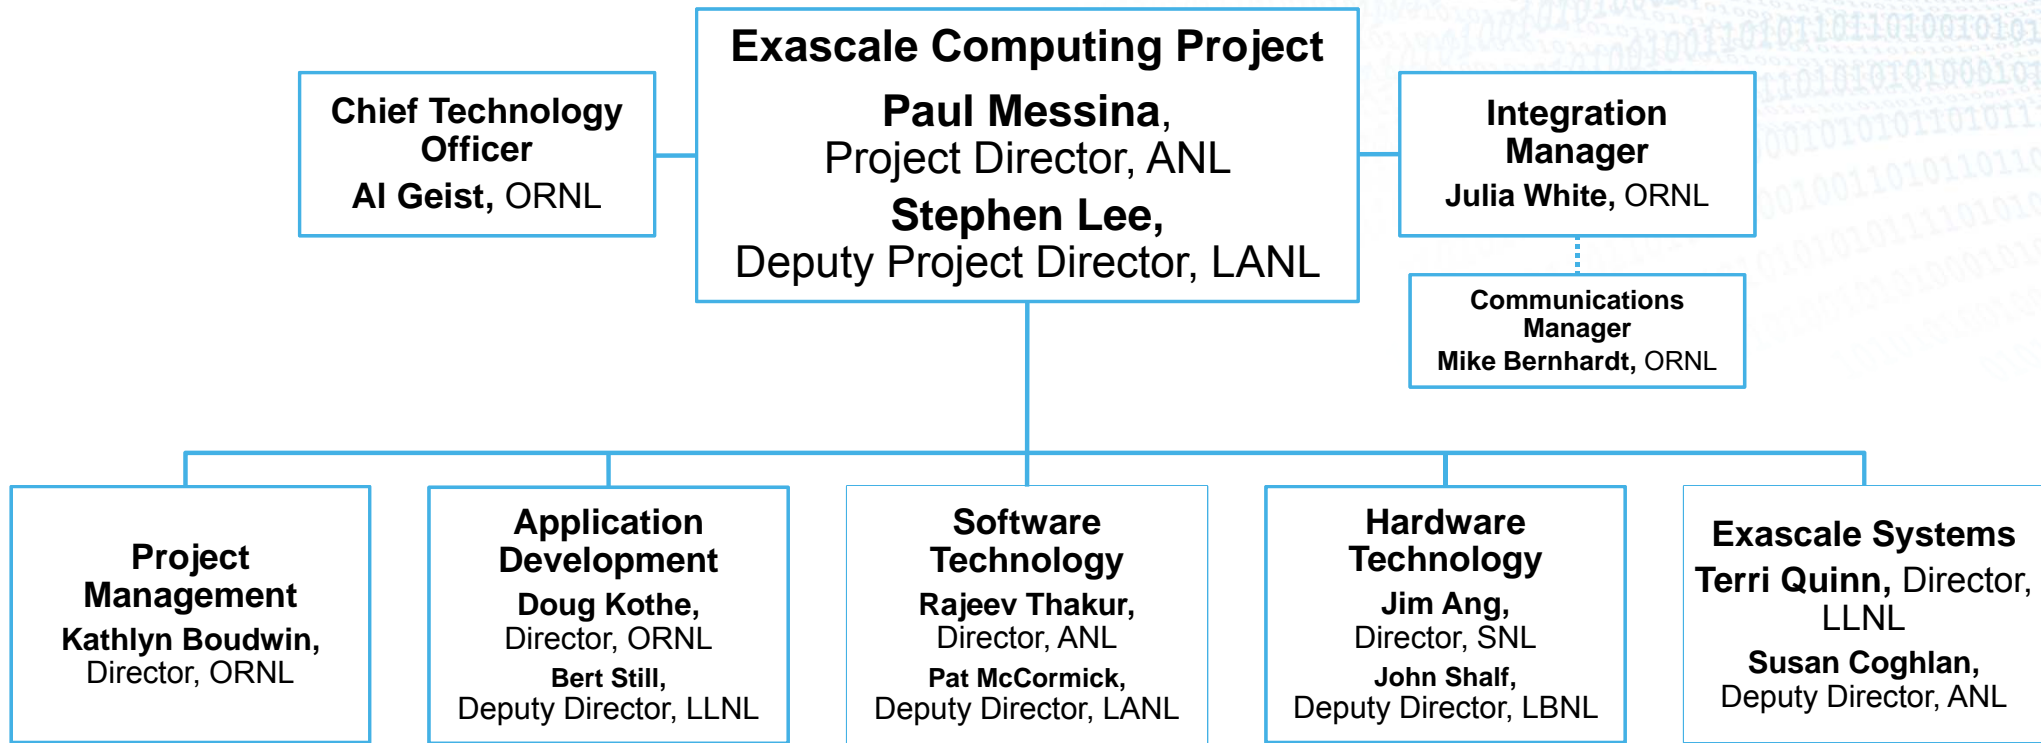
ECP is a collaboration among six labs

- 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 deploying first generation HPC systems
- Leadership team expertise spans all ECP activity areas



ECP leadership team

Staff from 6 national laboratories, with combined experience of >300 years



The 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

Four key challenges that must be addressed to achieve exascale

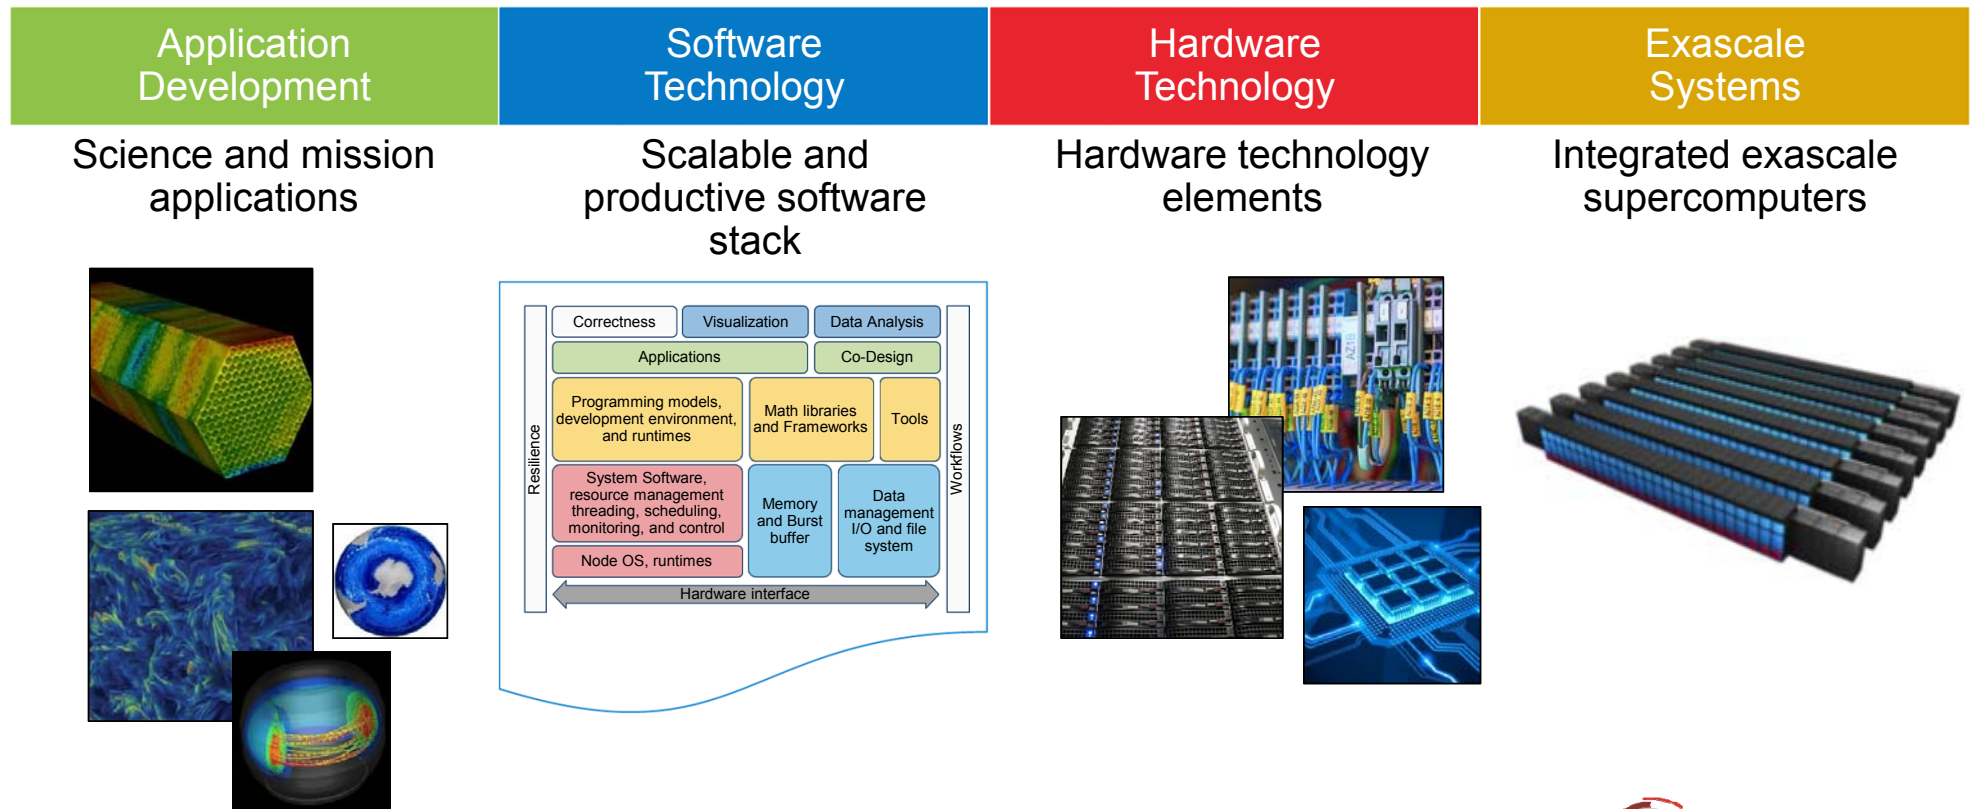
- Parallelism
- Memory and Storage
- Reliability
- Energy Consumption

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

The holistic co-design approach to deliver advanced architecture and capable exascale




Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Nuclear Energy (NE)

Accelerate design and commercialization of next-generation small modular reactors


- ✓Climate Action Plan
- ✓SMR licensing support
- ✓GAIN



Climate (BER)

Accurate regional impact assessment of climate change


- ✓Climate Action Plan



Chemical Science (BES, BER)

Biofuel catalysts design; stress-resistant crops

- ✓Climate Action Plan
- ✓MGI



Wind Energy (EERE)

Increase efficiency and reduce cost of turbine wind plants sited in complex terrains


- ✓Climate Action Plan



Combustion (BES)

Design high-efficiency, low-emission combustion engines and gas turbines

- ✓2020 greenhouse gas and 2030 carbon emission goals




Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Materials Science (BES)

Find, predict, and control materials and properties: property change due to hetero-interfaces and complex structures


✓MGI



Nuclear Physics (NP)

QCD-based elucidation of fundamental laws of nature: SM validation and beyond SM discoveries


✓2015 Long Range Plan for Nuclear Science
✓RHIC, CEBAF, FRIB



Nuclear Materials (BES, NE, FES)

Extend nuclear reactor fuel burnup and develop fusion reactor plasma-facing materials


✓Climate Action Plan
✓MGI
✓LWR Sustainability
✓ITER
✓Stockpile Stewardship Program



Accelerator Physics (HEP)

Practical economic design of 1 TeV electron-positron high-energy collider with plasma wakefield acceleration


✓>30k accelerators in industry, security, energy, environment, medicine



Materials Science (BES)

Protein structure and dynamics; 3D molecular structure design of engineering functional properties

✓MGI
✓LCLS-II 2025 Path Forward




Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Magnetic Fusion Energy (FES)

Predict and guide stable ITER operational performance with an integrated whole device model

- ✓ITER
- ✓Fusion experiments: NSTX, DIII-D, Alcator C-Mod



Advanced Manufacturing (EERE)

Additive manufacturing process design for qualifiable metal components


- ✓NNMIs
- ✓Clean Energy Manufacturing Initiative



Cosmology (HEP)

Cosmological probe of standard model of particle physics: Inflation, dark matter, dark energy


- ✓Particle Physics Project Prioritization Panel (P5)



Geoscience (BES, BER, EERE, FE, NE)

Safe and efficient use of subsurface for carbon capture and storage, petroleum extraction, geothermal energy, nuclear waste

- ✓EERE Forge
- ✓FE NRAP
- ✓Energy-Water Nexus
- ✓SubTER Crosscut



Precision Medicine for Cancer (NIH)

Accelerate and translate cancer research in RAS pathways, drug responses, treatment strategies

- ✓Precision Medicine in Oncology
- ✓Cancer Moonshot




Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Seismic (EERE, NE, NNSA)

Reliable earthquake hazard and risk assessment in relevant frequency ranges


- ✓ DOE Critical Facilities Risk Assessment
- ✓ Urban area risk assessment
- ✓ Treaty verification



Carbon Capture and Storage (FE)

Scaling carbon capture/storage laboratory designs of multiphase reactors to industrial size


- ✓ Climate Action Plan
- ✓ SunShot
- ✓ 2020 greenhouse gas
- ✓ 2030 carbon emission goals



Chemical Science (BES)

Design catalysts for conversion of cellulosic-based chemicals into fuels, bioproducts


- ✓ Climate Action Plan
- ✓ SunShot Initiative
- ✓ MGI



Urban Systems Science (EERE)

Retrofit and improve urban districts with new technologies, knowledge, and tools

- ✓ Energy-Water Nexus
- ✓ Smart Cities Initiative




Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Metagenomics (BER)

Leveraging microbial diversity in metagenomic datasets for new products and life forms

- ✓Climate Action Plan
- ✓Human Microbiome Project
- ✓Marine Microbiome Initiative



Astrophysics (NP)

Demystify origin of chemical elements (> Fe); confirm LIGO gravitational wave and DUNE neutrino signatures

- ✓2015 Long Range Plan for Nuclear Science
- ✓Origin of universe and nuclear matter in universe



Power Grid (EERE, OE)

Reliably and efficiently planning our nation's grid for societal drivers: rapidly increasing renewable energy penetration, more active consumers

- ✓Grid Modernization Initiative
- ✓Climate Action Plan



Application Co-Design (CD)

Application	Monte Carlo	Particles	Sparse Linear Algebra	Dense Linear Algebra	Spectral Methods	Unstructured Grid	Structured Grid	Comb. Logic	Graph Traversal	Dynamical Program	Backtrack & Branch and Bound	Graphical Models	Finite State Machine
Combustion S&T													
Free Electron Laser Data Analytics													
Microbiome Analysis													

Essential to ensure that applications effectively utilize exascale systems

- Pulls ST and HT developments into applications
- Pushes application requirements into ST and HT RD&D
- Evolved from best practice to an essential element of the development cycle

Executed by several CD Centers focusing on a unique collection of algorithmic motifs invoked by ECP applications

- Motif: algorithmic method that drives a common pattern of computation and communication
- CD Centers must address all high priority motifs invoked by ECP applications, including not only the 7 “classical” motifs but also the additional 6 motifs identified to be associated with data science applications

Game-changing mechanism for delivering next-generation community products with broad application impact

- Evaluate, deploy, and integrate exascale hardware-savvy software designs and technologies for key crosscutting algorithmic motifs into applications

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)**

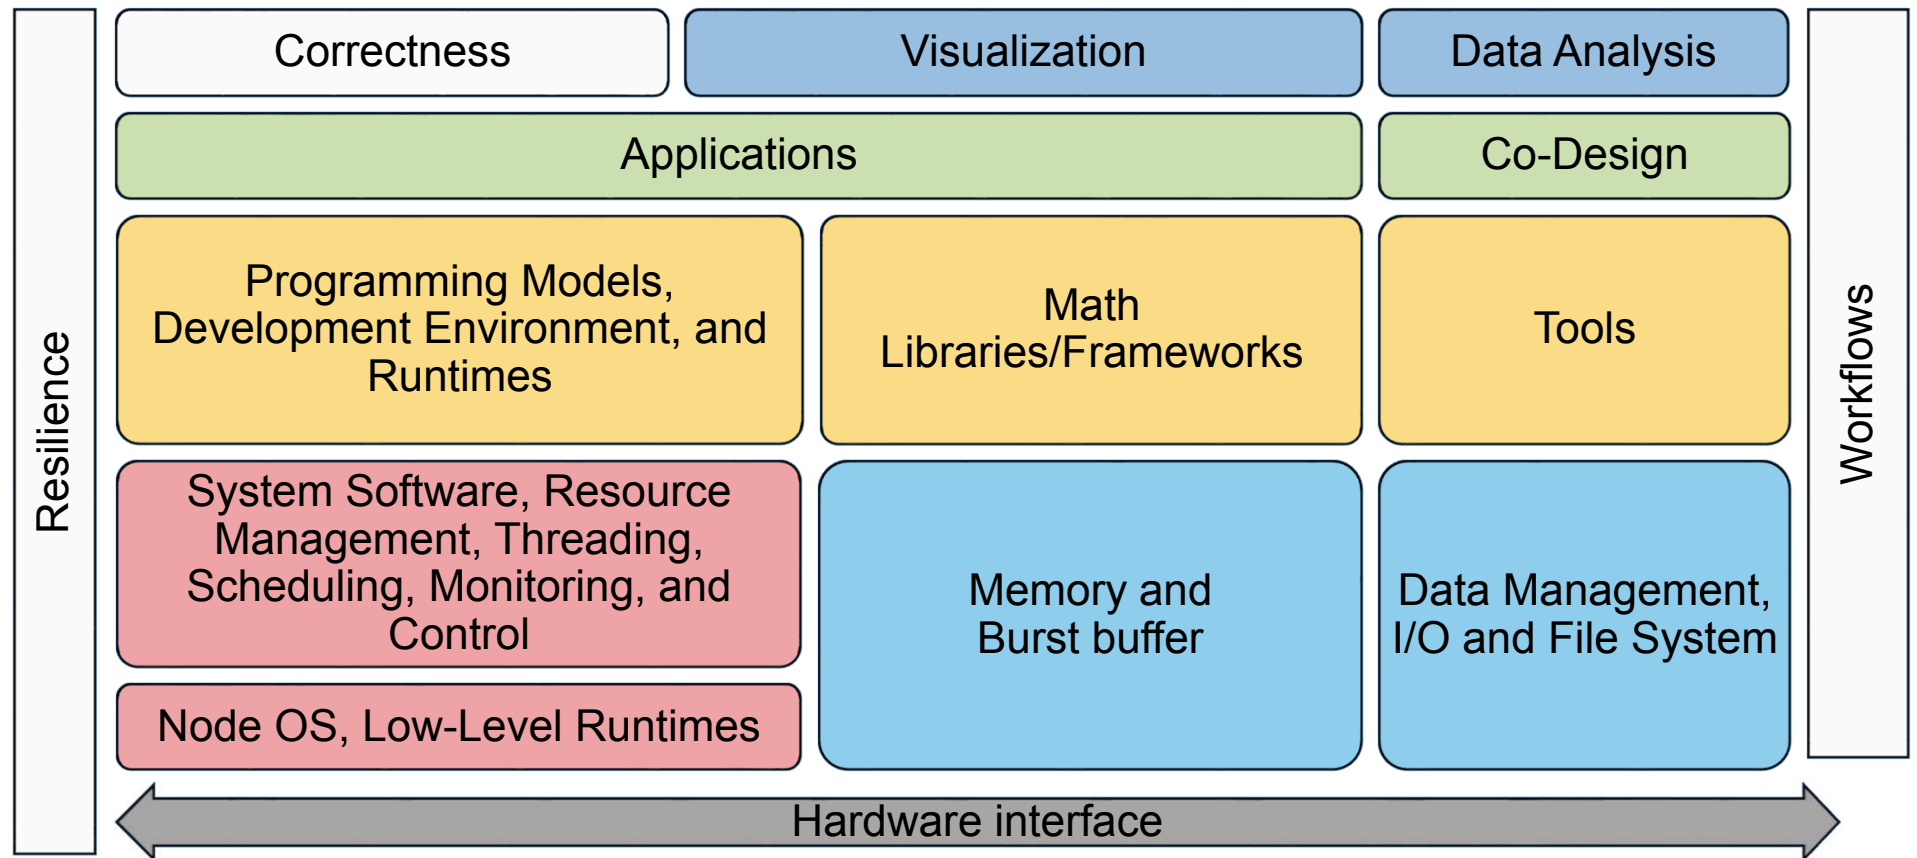
- **Motif(s):** Graph traversals; graph matching; graph coloring; graph clustering, including clique enumeration, parallel branch-and-bound, graph partitioning
- Develop methods and techniques for efficient implementation of key combinatorial (graph) algorithms that play a critical enabling role in numerous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



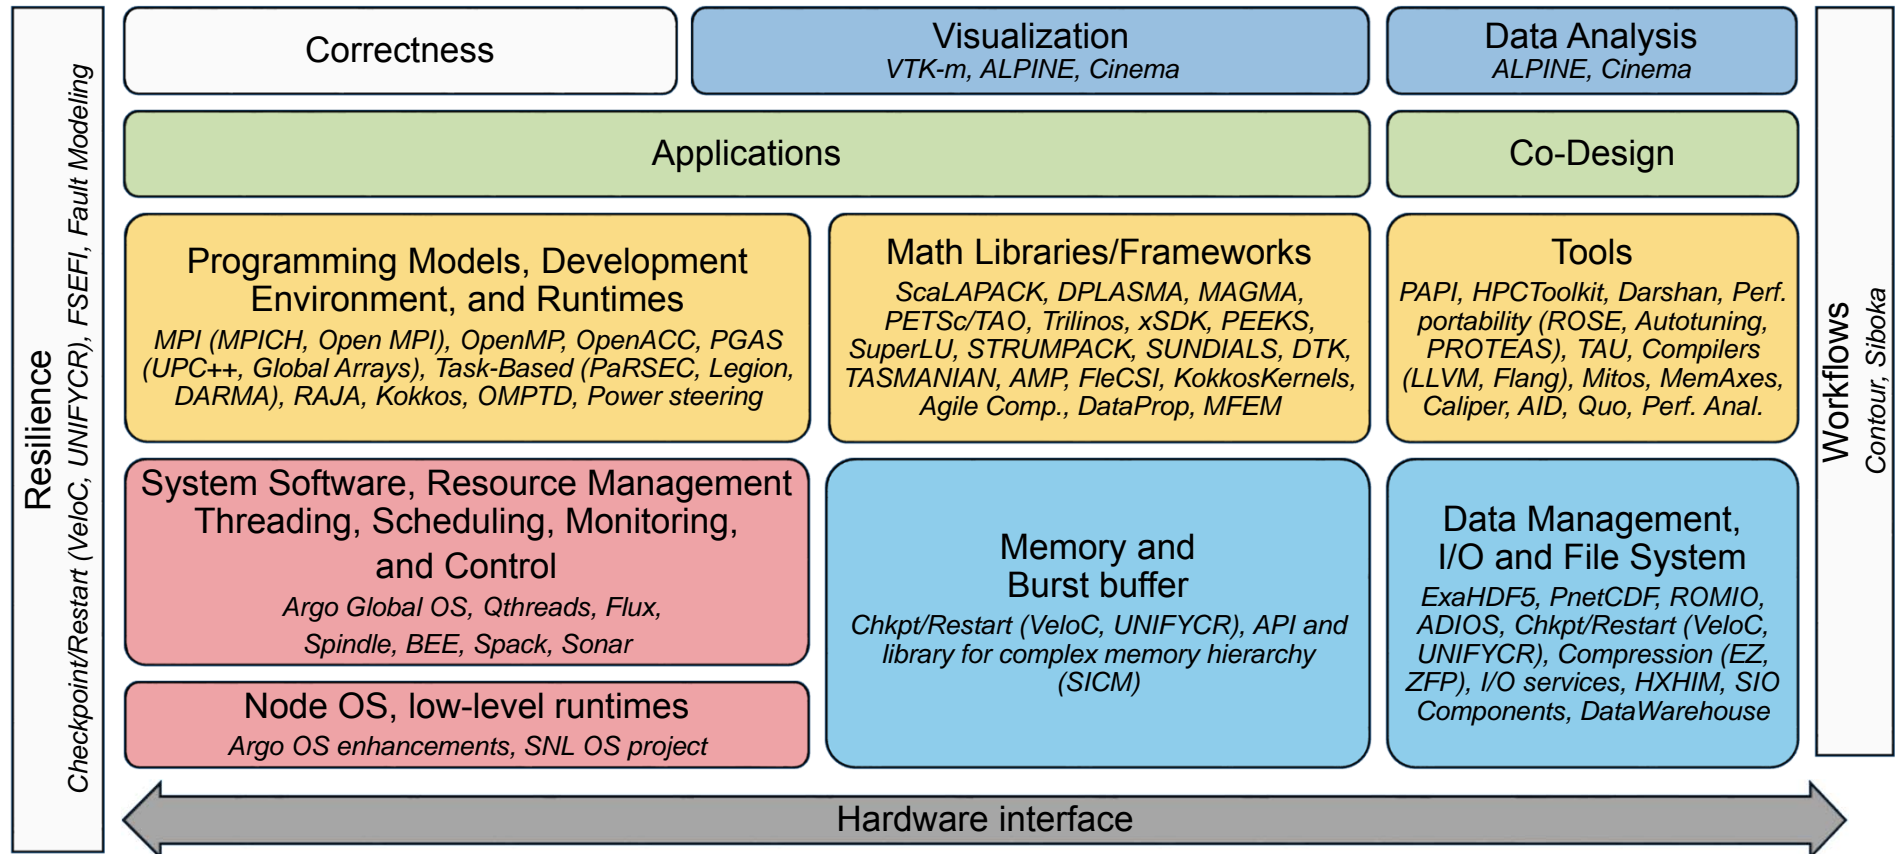
Software Technology Scope

- ECP will build a comprehensive and coherent software stack that will enable application developers to productively write highly parallel applications that can *portably* target diverse exascale architectures
- ECP will accomplish this by
 - extending current technologies to exascale where possible,
 - performing R&D required to conceive of new approaches where necessary,
 - coordinating with vendor efforts, and
 - developing and deploying high-quality and robust software products

Conceptual ECP Software Stack



Current Set of ST Projects Mapped to Software Stack



Current ECP data-intensive applications

- Machine learning for cancer
 - CANDLE project as a good model for frameworks
- Image analysis for light source data
- Several smaller efforts that involve data from
 - urban systems,
 - embedded seismic sensors,
 - cosmological images, and
 - microbial data analysis.

Challenges for Software Technology

- In addition to the usual exascale challenges -- scale, memory hierarchy, power, and performance portability -- the main challenge is the codesign and integration of various components of the software stack with each other, with a broad range of applications, with emerging hardware technologies, and with the software provided by system vendors
- These aspects must all come together to provide application developers with a productive development and execution environment

Next Steps in the Software Stack

- Over the next few months, we will undertake a **gap analysis** to identify what aspects of the software stack are missing in the portfolio, based on requirements of applications and DOE HPC facilities, and discussions with vendors
- Based on the results of the gap analysis, we will issue targeted RFIs/RFPs that will aim to close the identified gaps

Known Gaps

- Our preliminary software stack has been built bottom up, largely based on current usage and plans of the applications teams
- We have few applications that involve big data, large-scale data analytics, deep learning
- Ditto for complex workflows
- Areas for which we deliberately decided to do technology watches before investing in them very much (have not identified enough use cases and what is required to support them)
 - Resilience
 - Workflows

Hardware Technology Overview

Objective: Fund R&D to design hardware that meets ECP's Targets for application performance, programmability, power efficiency, and resilience

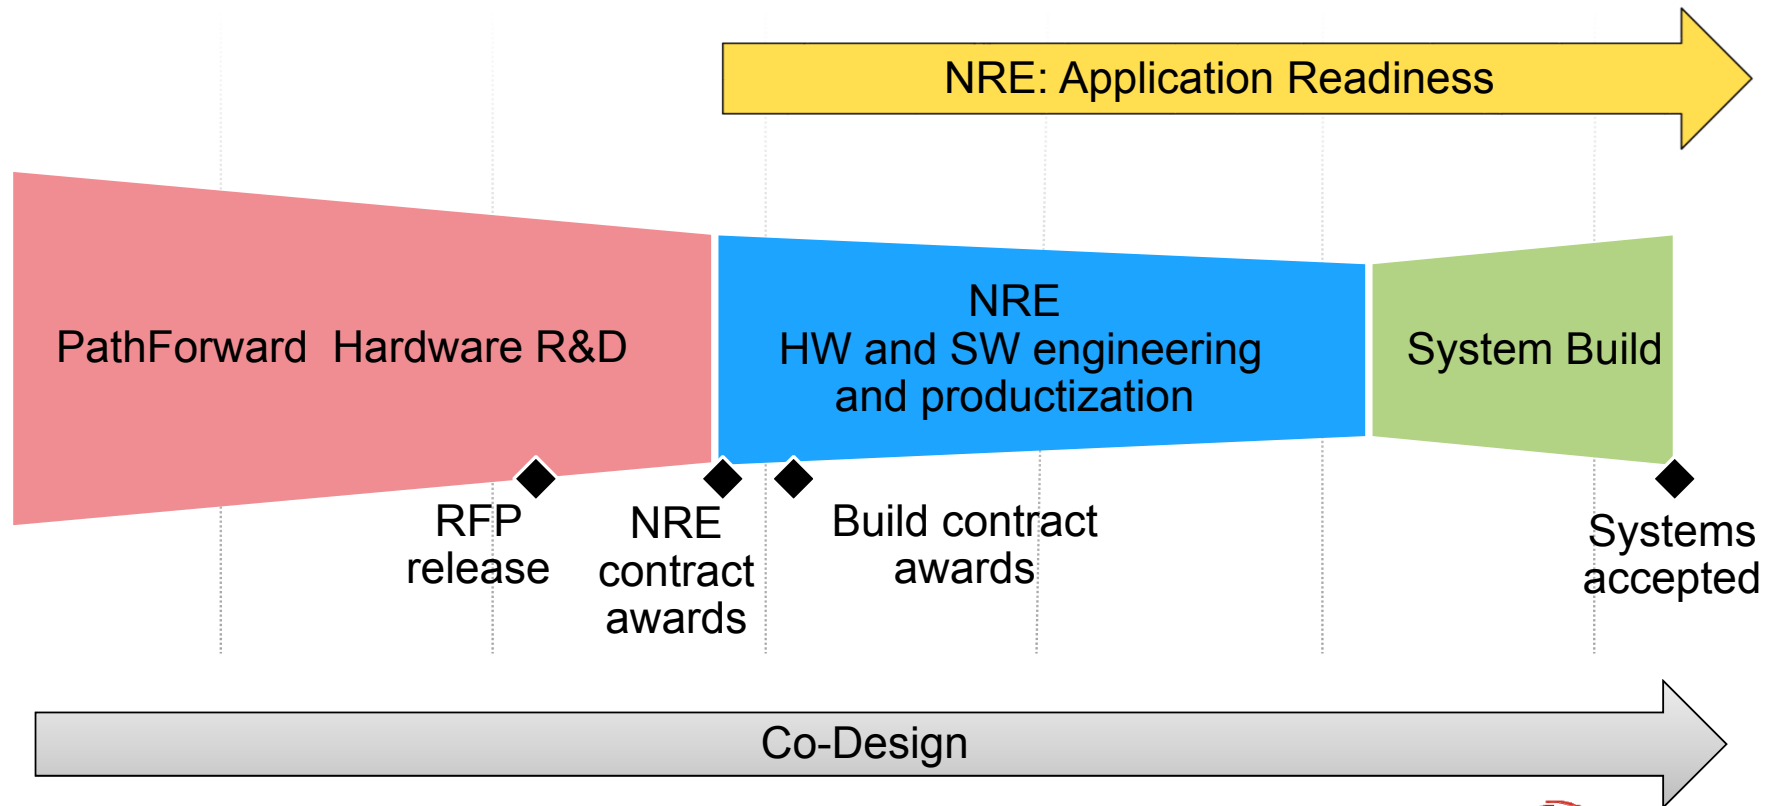
Vendor Partnerships for Hardware Architecture R&D contracts that deliver:

- Conceptual exascale node and system designs
- Analysis of performance improvement on these conceptual system designs
- Technology demonstrators to quantify performance gains over existing roadmaps
- Support for active industry engagement in ECP holistic co-design efforts

DOE Lab Design Space Evaluation Team

- Participate in evaluation and review of Vendor HW Architecture R&D deliverables
- Provide Architectural Analysis, and Abstract Machine Models of Vendor designs to support ECP's holistic co-design

ECP's plan to accelerate and enhance system capabilities



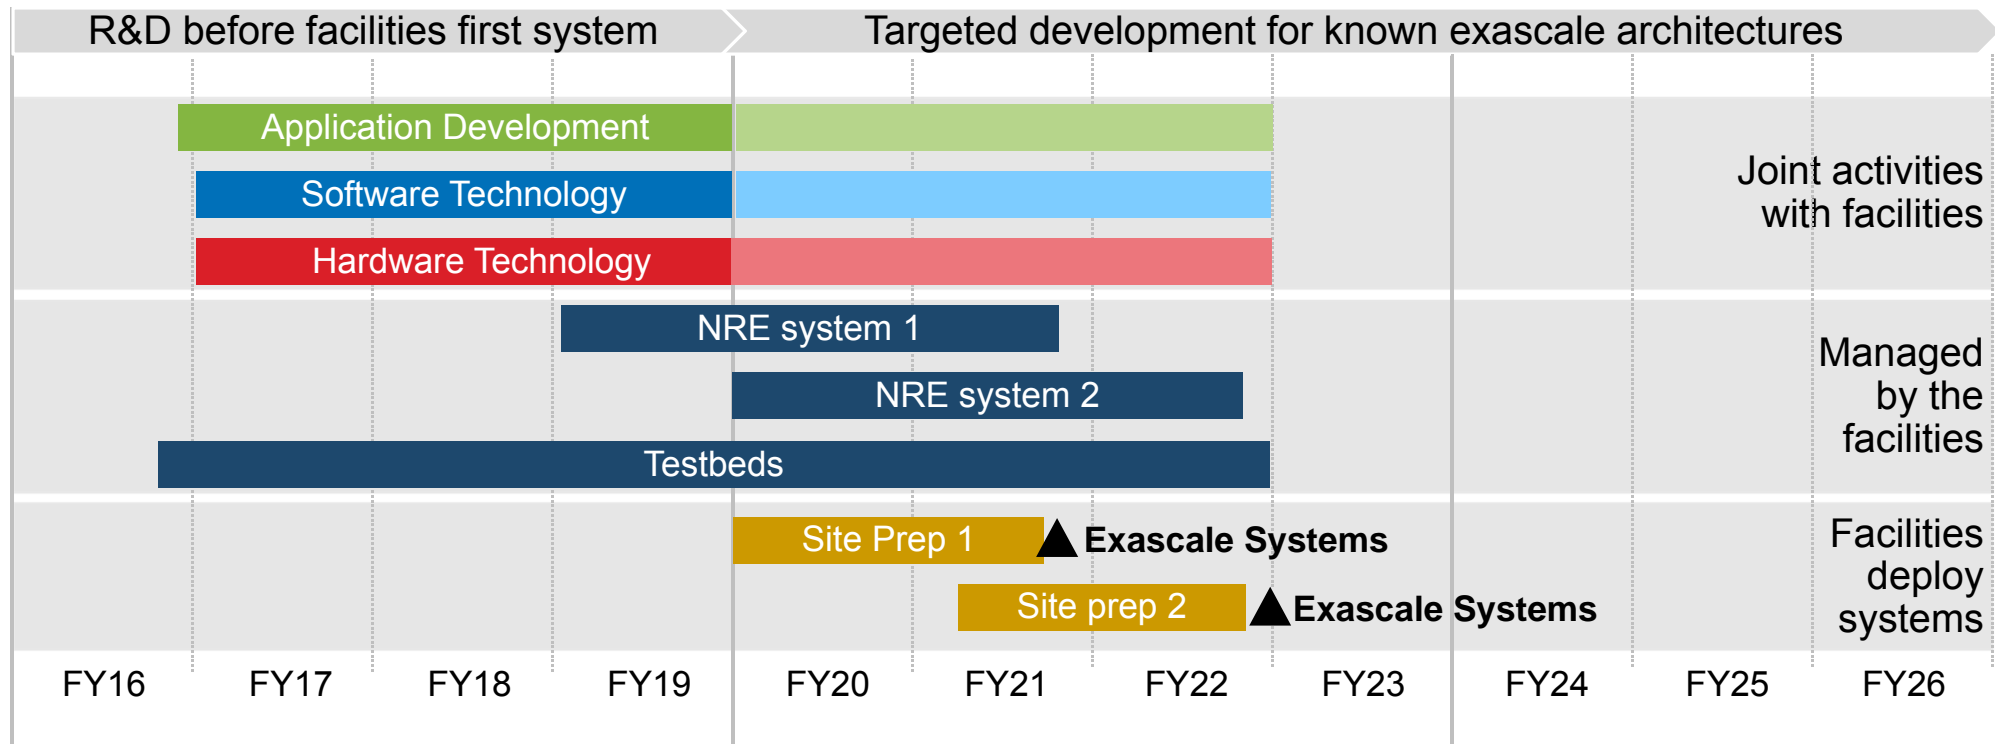
Systems acquisition approach

- DOE-SC and NNSA programs will procure and install the ECP systems
 - ECP's and DOE-SC/NNSA's processes will be tightly coupled and interdependent
 - ECP's requirements will be incorporated into RFP(s)
 - ECP will participate in system selection and co-design
 - ECP will make substantial investments through non-recurring engineering (NRE) contracts coupled to system acquisition contracts

NRE contracts

- Incentivize awardees to address gaps in their system product roadmaps
- Bring to the product stage promising hardware and software research and integrate it into a system
- Accelerate technologies, add capabilities, improve performance, and lower the cost of ownership of system
- Include application readiness R&D efforts
- More than 2 full years of lead time are necessary to maximize impact

High-level ECP technical project schedule





EXASCALE COMPUTING PROJECT



DOE LABORATORIES &
AGENCY PARTNERS

22



PRIVATE SECTOR
PARTNERS

9



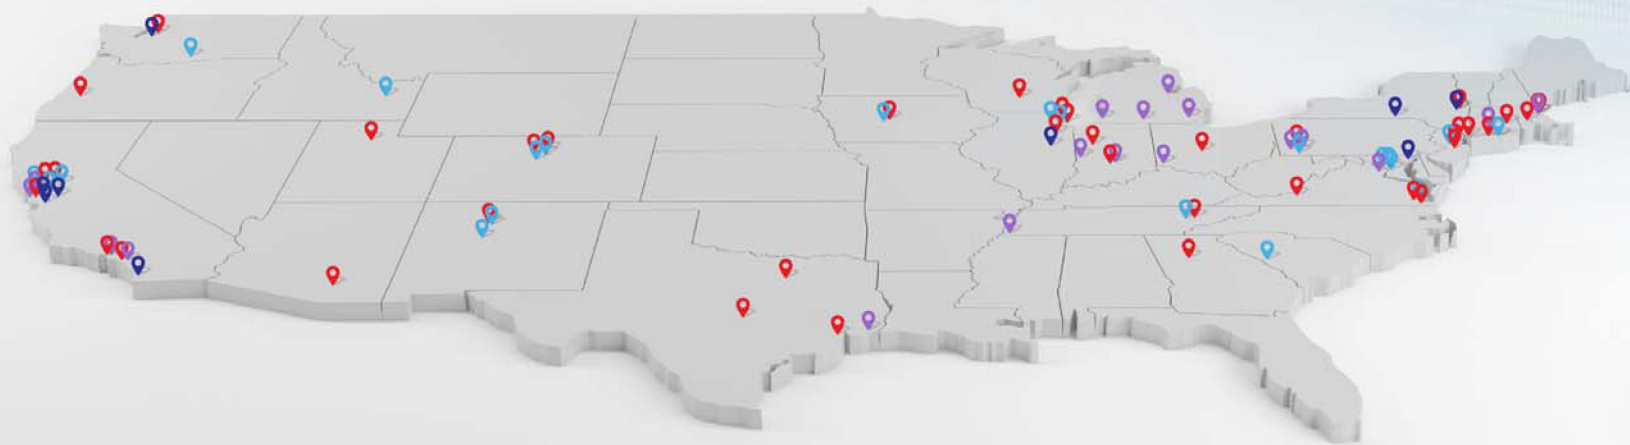
UNIVERSITY RESEARCH
PARTNERS

39



INDUSTRY COUNCIL
MEMBERS

18



Questions?

www.ExascaleProject.org



EXASCALE COMPUTING PROJECT



www.ExascaleProject.org