# The U.S. Exascale Computing Project

What is it, why do we have this, and how is it structured?

**Paul Messina**
Director, Exascale Computing Project

**Stephen Lee**
Deputy Director, Exascale Computing Project

Presented to the
**ECP Industry Council**

Argonne National Laboratory
March 7, 2017

EXASCALE COMPUTING PROJECT

# 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.

# Exascale meeting science and national security needs

- High Performance Computing (HPC) has become an indispensable tool for fundamental understanding and for prediction of properties and behaviors of materials and entire systems
- HPC based modeling and simulation is used extensively in the advancement of DOE missions in science and in the national security space, including stewardship of the nation's nuclear stockpile
- Without a sustained commitment for HPC advancement, the US will lose its competitive edge with attendant, adverse consequences to scientific discovery and economic competitiveness
- The Exascale Computing Project (ECP) is part of the Exascale Computing Initiative (ECI), a coordinated effort to support US leadership in achieving next-generation HPC
  - DOE is a lead agency within NSCI, along with DoD and NSF
  - DOE-SC and NNSA are executing a joint effort on advanced performance on relevant applications and data analytic computing
  - The ECP is a 7-year project with a cost range of \$3.5B–\$5.7B

## 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 adaption 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

# From Giga to Exa, via Tera & Peta*

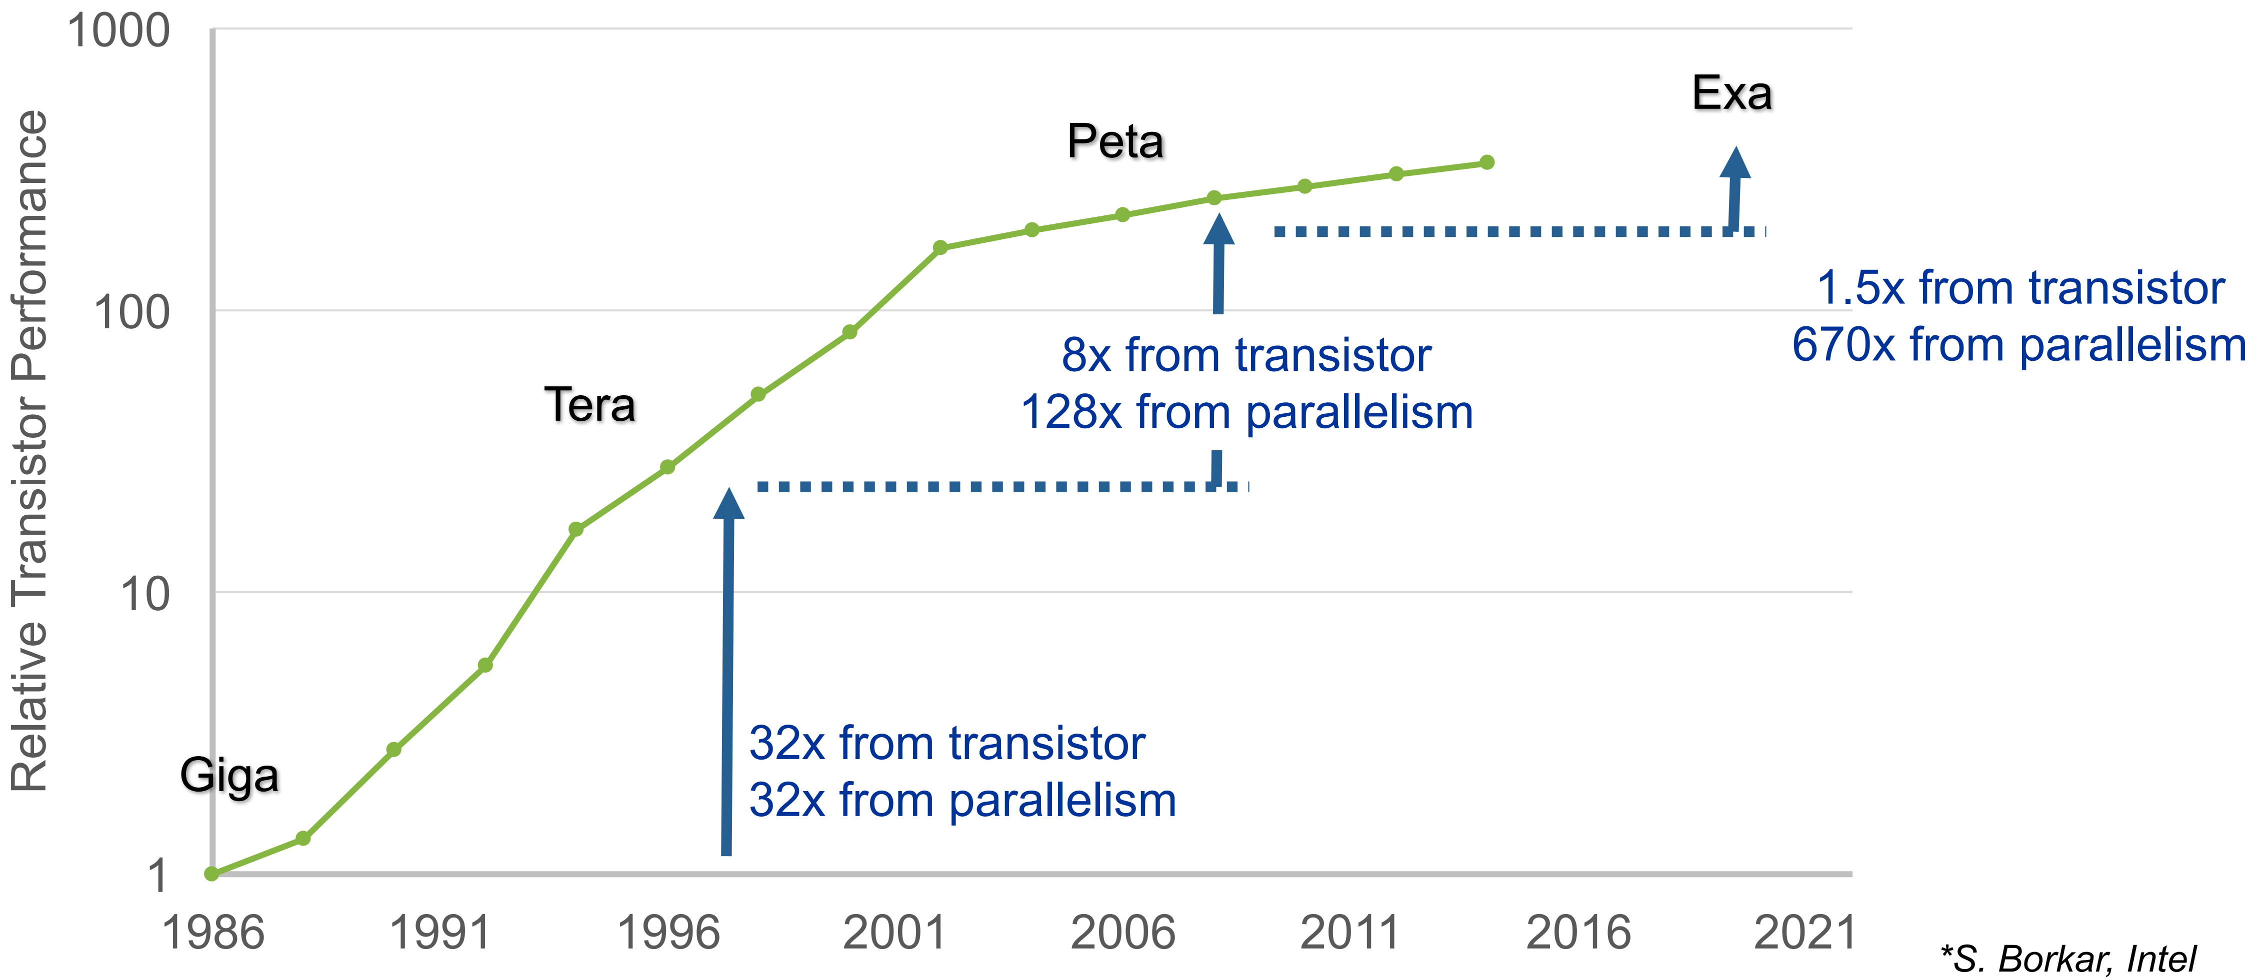


*S. Borkar, Intel

**Performance from parallelism**

EXASCALE COMPUTING PROJECT

# Why the Department of Energy?

- The DOE national labs have been among the leaders in HPC since the early days of digital (and analog) computers
- Starting in the 1950s Argonne, Los Alamos, Livermore designed and built computers, often in collaboration with vendors
  - Commercial products were usually the result
  - This has continued to the present, with additional labs as well, e.g., Sandia
- Software developed at the DOE labs is widely used
  - Mathematical software libraries, programming models, I/O, visualization, …
  - Application software, e.g., KIVA, LS-DYNA

# AVIDAC: Argonne's Version of the Institute's Digital Arithmetic Computer: 1949-1953

“Moll” Flanders, Director
Jeffrey Chu, Chief Engineer

**Margaret Butler wrote AVIDAC’s interpretive floating-point arithmetic system**

- Memory access time: 15 microsec
- Addition: 10 microsec
- Multiplication: 1 millisec

**AVIDAC press: @ 100,000 times as fast as a trained “Computer” using a desk calculator**

EXASCALE COMPUTING PROJECT

# Early work on computer architecture

Margaret Butler helped assemble the ORACLE computer with ORNL Engineer Rudolph Klein. In 1953, ORACLE was the world's fastest computer, multiplying 12-digit numbers in
.0005 seconds. Designed at Argonne, it was constructed at Oak Ridge.

EXASCALE COMPUTING PROJECT

# Advancements in (High Performance) Computing Have Occurred in Several Distinct “Eras”

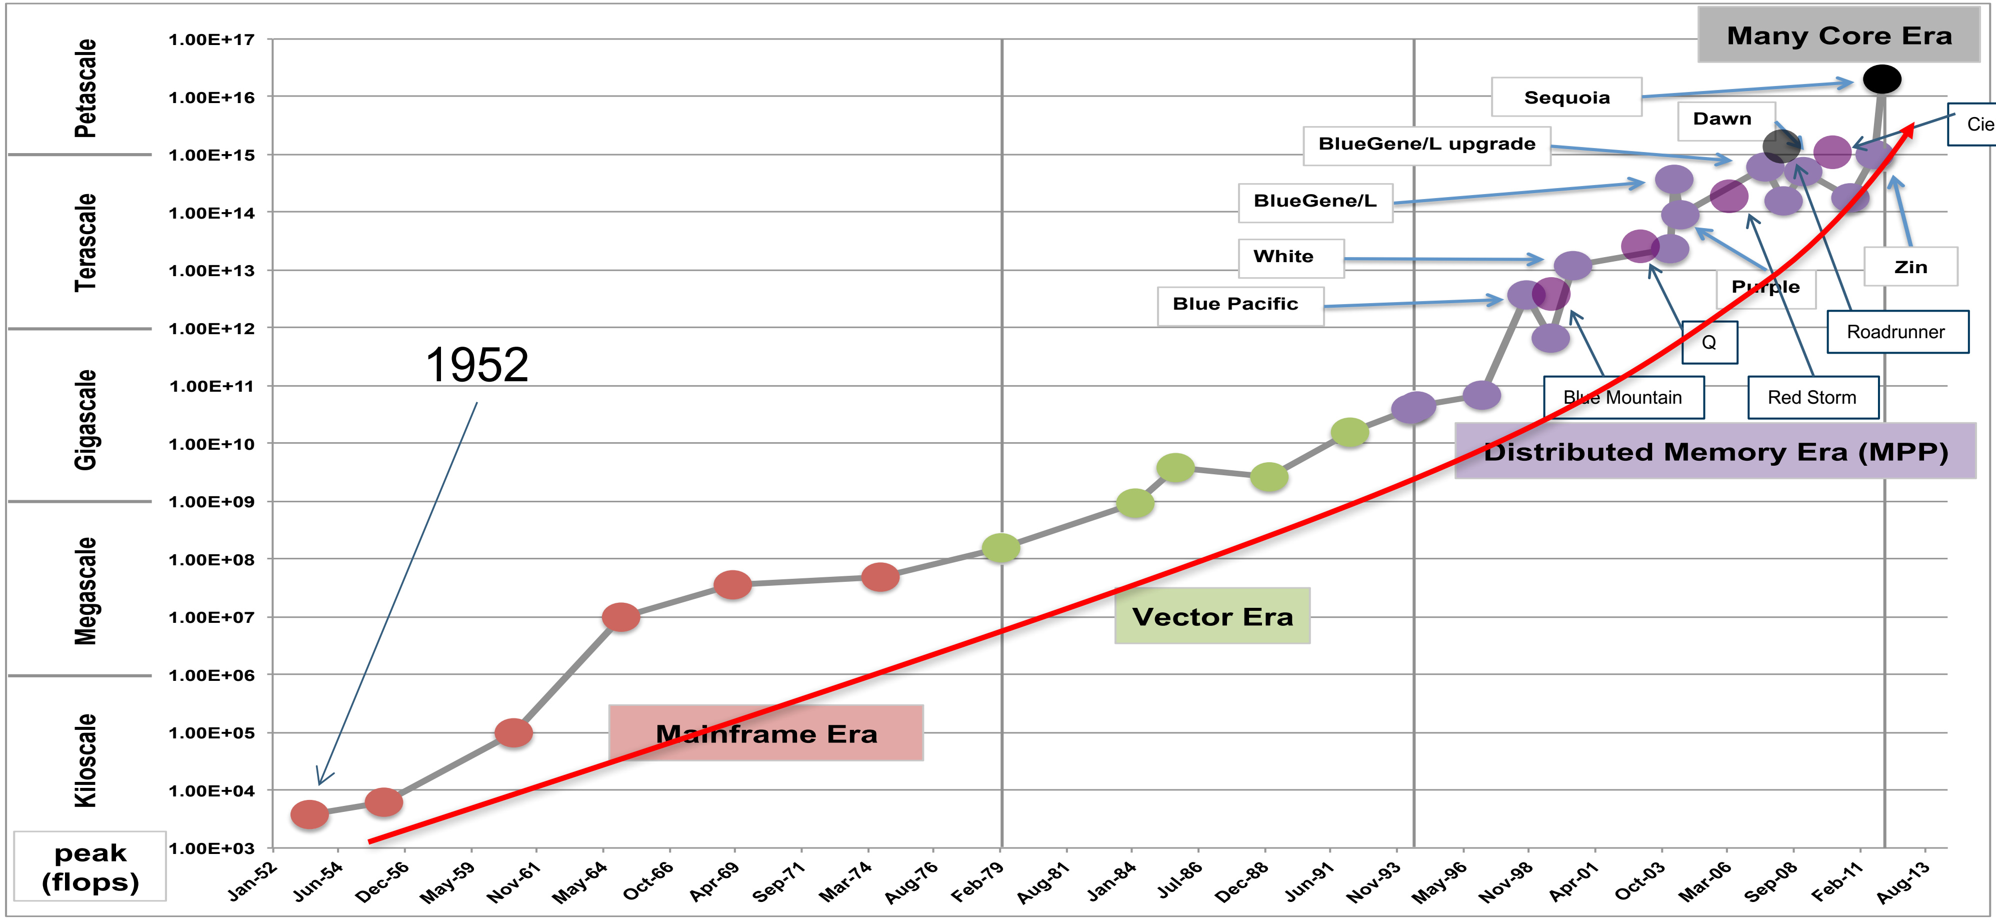


Truth is: we don’t what to call this next era. It’s currently defined by its inability to be defined!

Each of these eras define not so much a common hardware architecture, but a common <u>programming model</u>

*Slide credit: Rob Neely, LLNL, used with permission*

# The “mainframe” era – general purpose computers designed for scientific computing

1953
1954
1956
1958
1960

- Univac 1
  - First machine installed at LLNL in 1953
- IBM 701
  - Installed at Los Alamos also in 1953
  - Williams tubes (fast memory, but unreliable)
- IBM 704
  - Core memory, floating point arithmetic, CRT
  - Commercially successful
- IBM 709
  - Seamless porting
- IBM 7090
  - Transistor-based
  - Large speed increases
- Univac LARC
  - Co-designed with, and for, LLNL
  - One of the first transistor-based machines

*Slide credit: Rob Neely, LLNL, used with permission*



EXASCALE COMPUTING PROJECT

10

# The "mainframe" era – continued

1961
1962
1964
1969

- IBM 7030 (Stretch)
  - Competitor to LARC
  - Considered a failure at the time (only achieved 50% of performance goals)
  - Introduced many concepts that went into the IBM System/360
- CDC 1604
  - First designed by Seymour Cray
- CDC 3600
  - 48-bit words (higher precision)
- CDC 6600
  - Considered the first real "supercomputer"
  - Full separation of input/output from computing
- CDC 7600
  - Hierarchical memory design
  - Fastest computer in world from 1969-1975

*Slide credit: Rob Neely, LLNL, used with permission*

# ASCI has led the US to world leadership in high performance simulation [*this slide is from a presentation I gave in August 2000*]

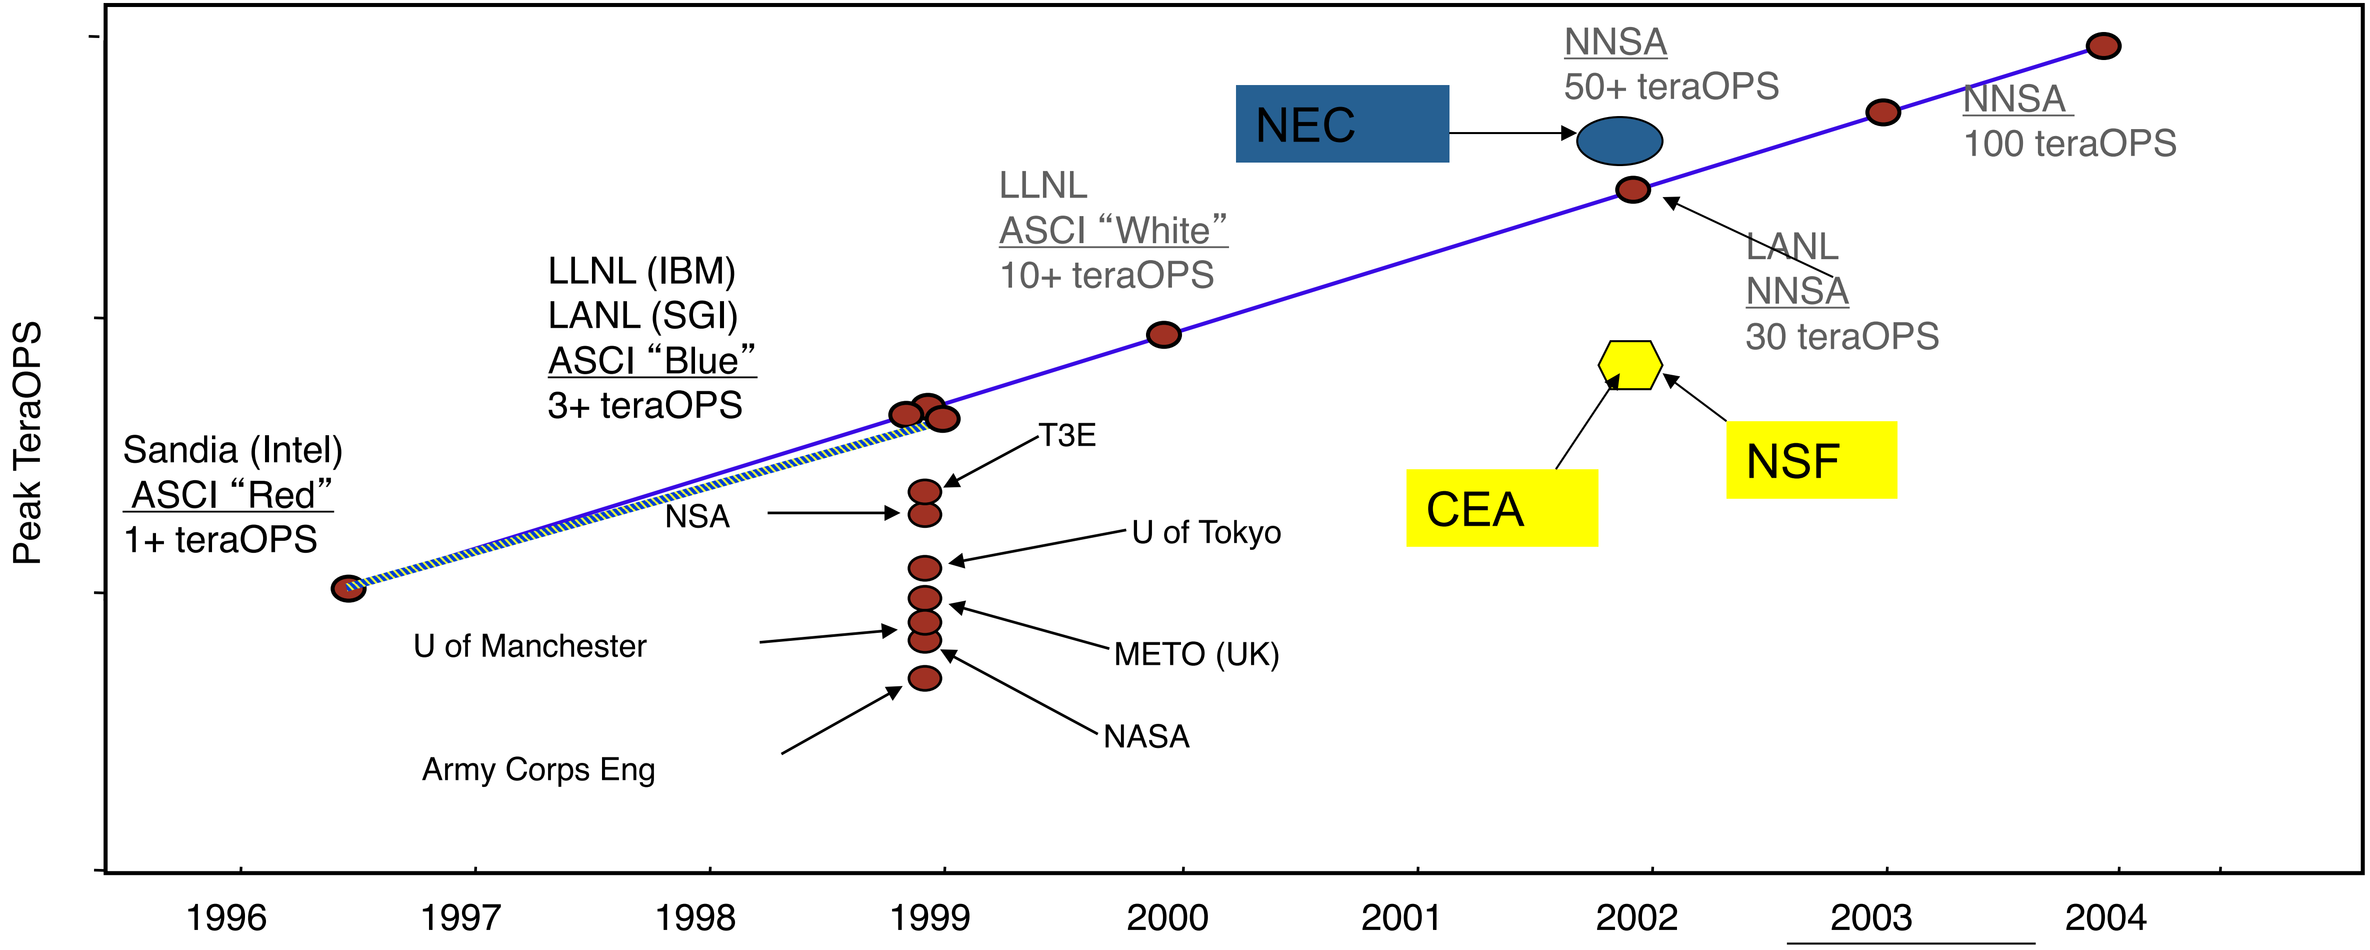


The NEC Earth Simulator @ 40TF& 5000 Processors is projected to be faster than the most capable NNSA platforms in 2002.



EXASCALE COMPUTING PROJECT

# International competition in HPC continues to intensify

- China has had the #1 spot on the Top500 list since June 2013 and the top two spots since November 2016
  - The new #1 machine is a Chinese machine that uses indigenously designed and manufactured processors and software stack
  - Chinese researchers using the new machine won November 2016 Gordon Bell prize recognizing outstanding achievement in HPC
- China vastly outspends the US on exascale
  - Has 3 potential exascale architectures in development and plans to deploy prototypes by the end of 2017
- China is growing a skilled national workforce, partially through repatriation of experts trained abroad

**China's goals:**

Economic competitiveness and national security

Export sales of HPC to other countries

Asserting national dominance

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

Develop applications to tackle a broad spectrum of mission critical problems of unprecedented complexity

Support national security

Partner with vendors to develop computer architectures that support exascale applications

Contribute to the economic competitiveness of the nation

Collaborate with vendors to develop a software stack that is both exascale-capable and usable on industrial and academic scale systems

Train a next-generation workforce of computational scientists, engineers, and computer scientists

# 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

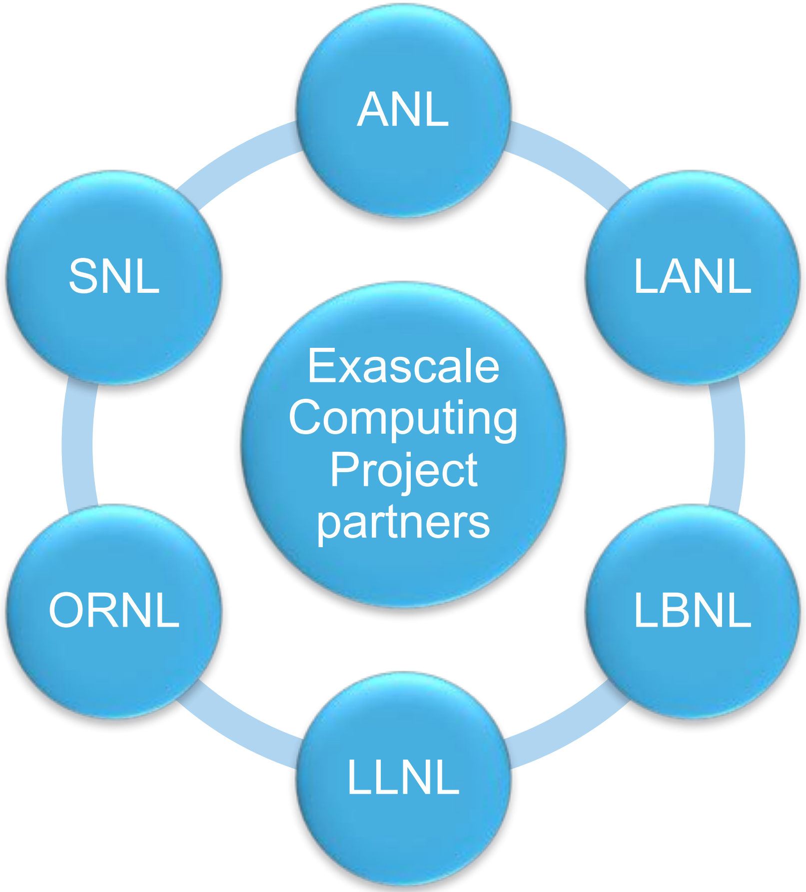


EXASCALE COMPUTING PROJECT

# 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 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



ECP EXASCALE COMPUTING PROJECT

# 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: ≤1/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

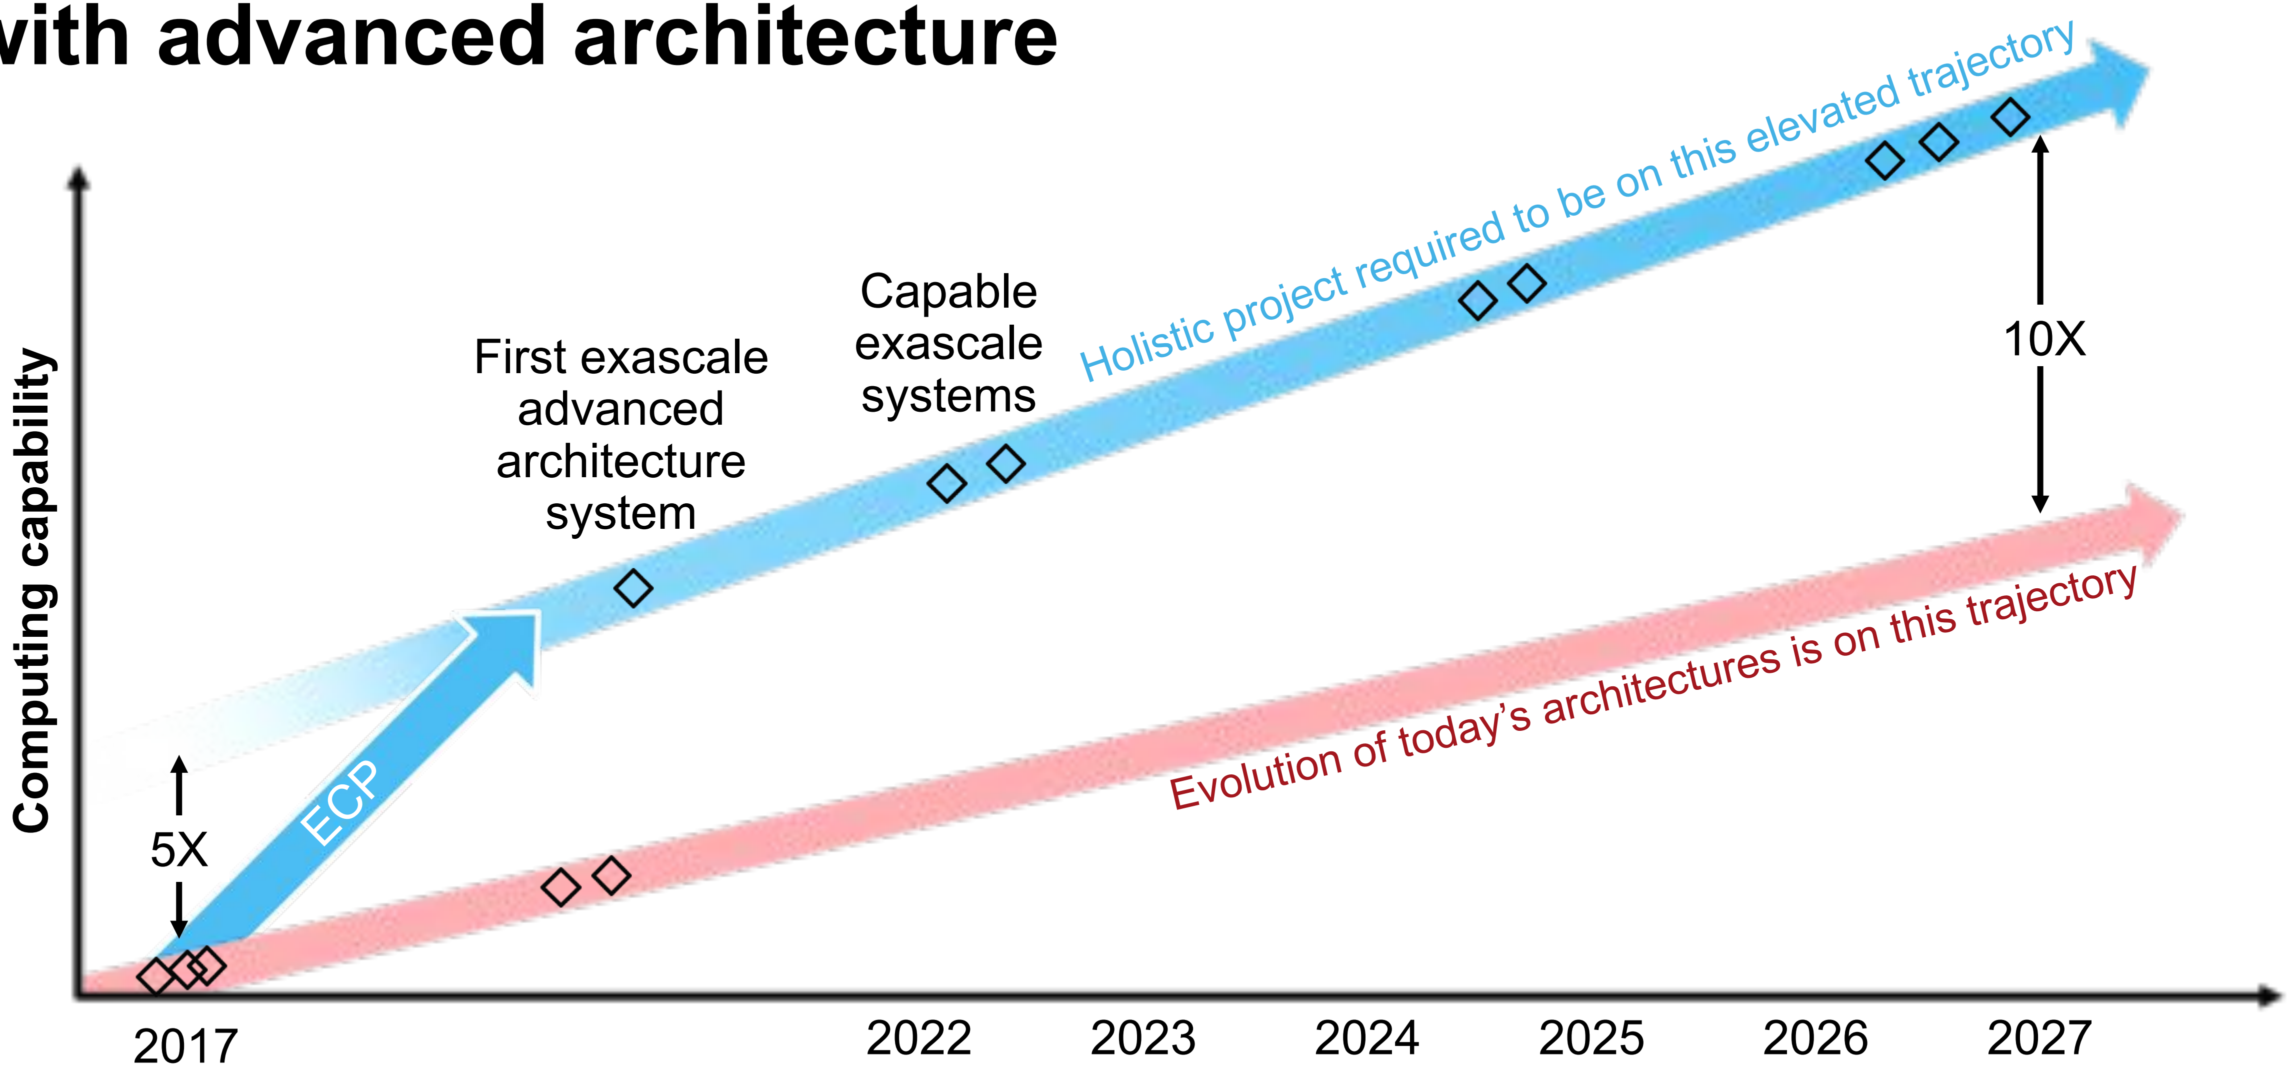




EXASCALE COMPUTING PROJECT

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

| Application Development | Software Technology | Hardware Technology | Exascale Systems |
|---|---|---|---|
| Science and mission applications | Scalable and productive software stack | Hardware technology elements | Integrated exascale supercomputers |

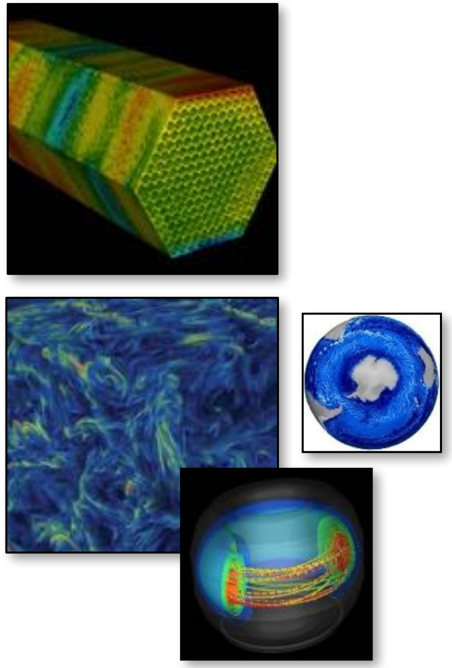

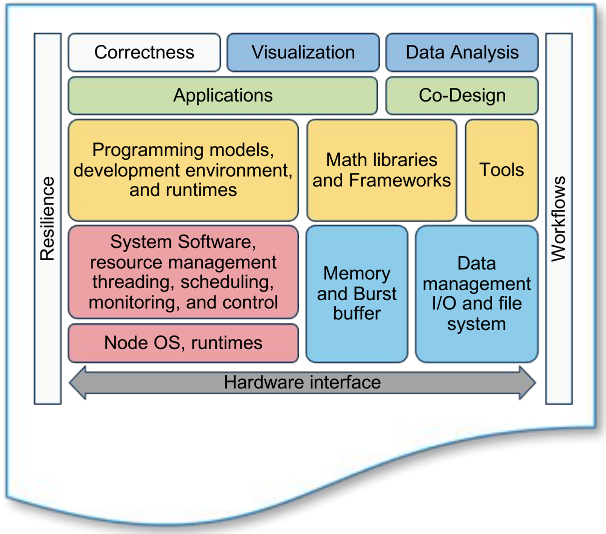


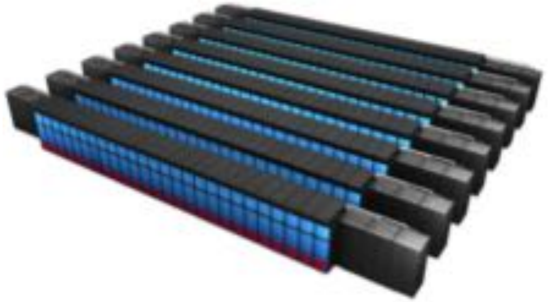

EXASCALE COMPUTING PROJECT

# ECP leadership team

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

**Chief Technology Officer**
**Al Geist**, ORNL

**Exascale Computing Project**
**Paul Messina**, Project Director, ANL
**Stephen Lee**, Deputy Project Director, LANL

**Integration Manager**
**Julia White**, ORNL

**Communications Manager**
**Mike Bernhardt**, ORNL

**Project Management**
**Kathlyn Boudwin**, Director, ORNL

**Application Development**
**Doug Kothe**, Director, ORNL
**Bert Still**, Deputy Director, LLNL

**Software Technology**
**Rajeev Thakur**, Director, ANL
**Pat McCormick**, Deputy Director, LANL

**Hardware Technology**
**Jim Ang**, Director, SNL
**John Shalf**, Deputy Director, LBNL

**Exascale Systems**
**Terri Quinn**, Director, LLNL
**Susan Coghlan**, Deputy Director, ANL

# ECP WBS

- **Exascale Computing Project 1.**
  - Project Management 1.1
    - Project Planning and Management 1.1.1
    - Project Controls & Risk Management 1.1.2
    - Business Management 1.1.3
    - Procurement Management 1.1.4
    - Information Technology and Quality Management 1.1.5
    - Communications & Outreach 1.1.6
    - Integration 1.1.7
  - Application Development 1.2
    - DOE Science and Energy Apps 1.2.1
    - DOE NNSA Applications 1.2.2
    - Other Agency Applications 1.2.3
    - Developer Training and Productivity 1.2.4
    - Co-Design and Integration 1.2.5
  - Software Technology 1.3
    - Programming Models and Runtimes 1.3.1
    - Tools 1.3.2
    - Mathematical and Scientific Libraries and Frameworks 1.3.3
    - Data Management and Workflows 1.3.4
    - Data Analytics and Visualization 1.3.5
    - System Software 1.3.6
    - Resilience and Integrity 1.3.7
    - Co-Design and Integration 1.3.8
    - SW PathForward 1.3.9
  - Hardware Technology 1.4
    - PathForward Vendor Node and System Design 1.4.1
    - Design Space Evaluation 1.4.2
    - Co-Design and Integration 1.4.3
    - PathForward II Vendor Node and System Design 1.4.4
  - Exascale Systems 1.5
    - NRE 1.5.1
    - Testbeds 1.5.2
    - Co-design and Integration 1.5.3

EXASCALE COMPUTING PROJECT

# ECP Project organization

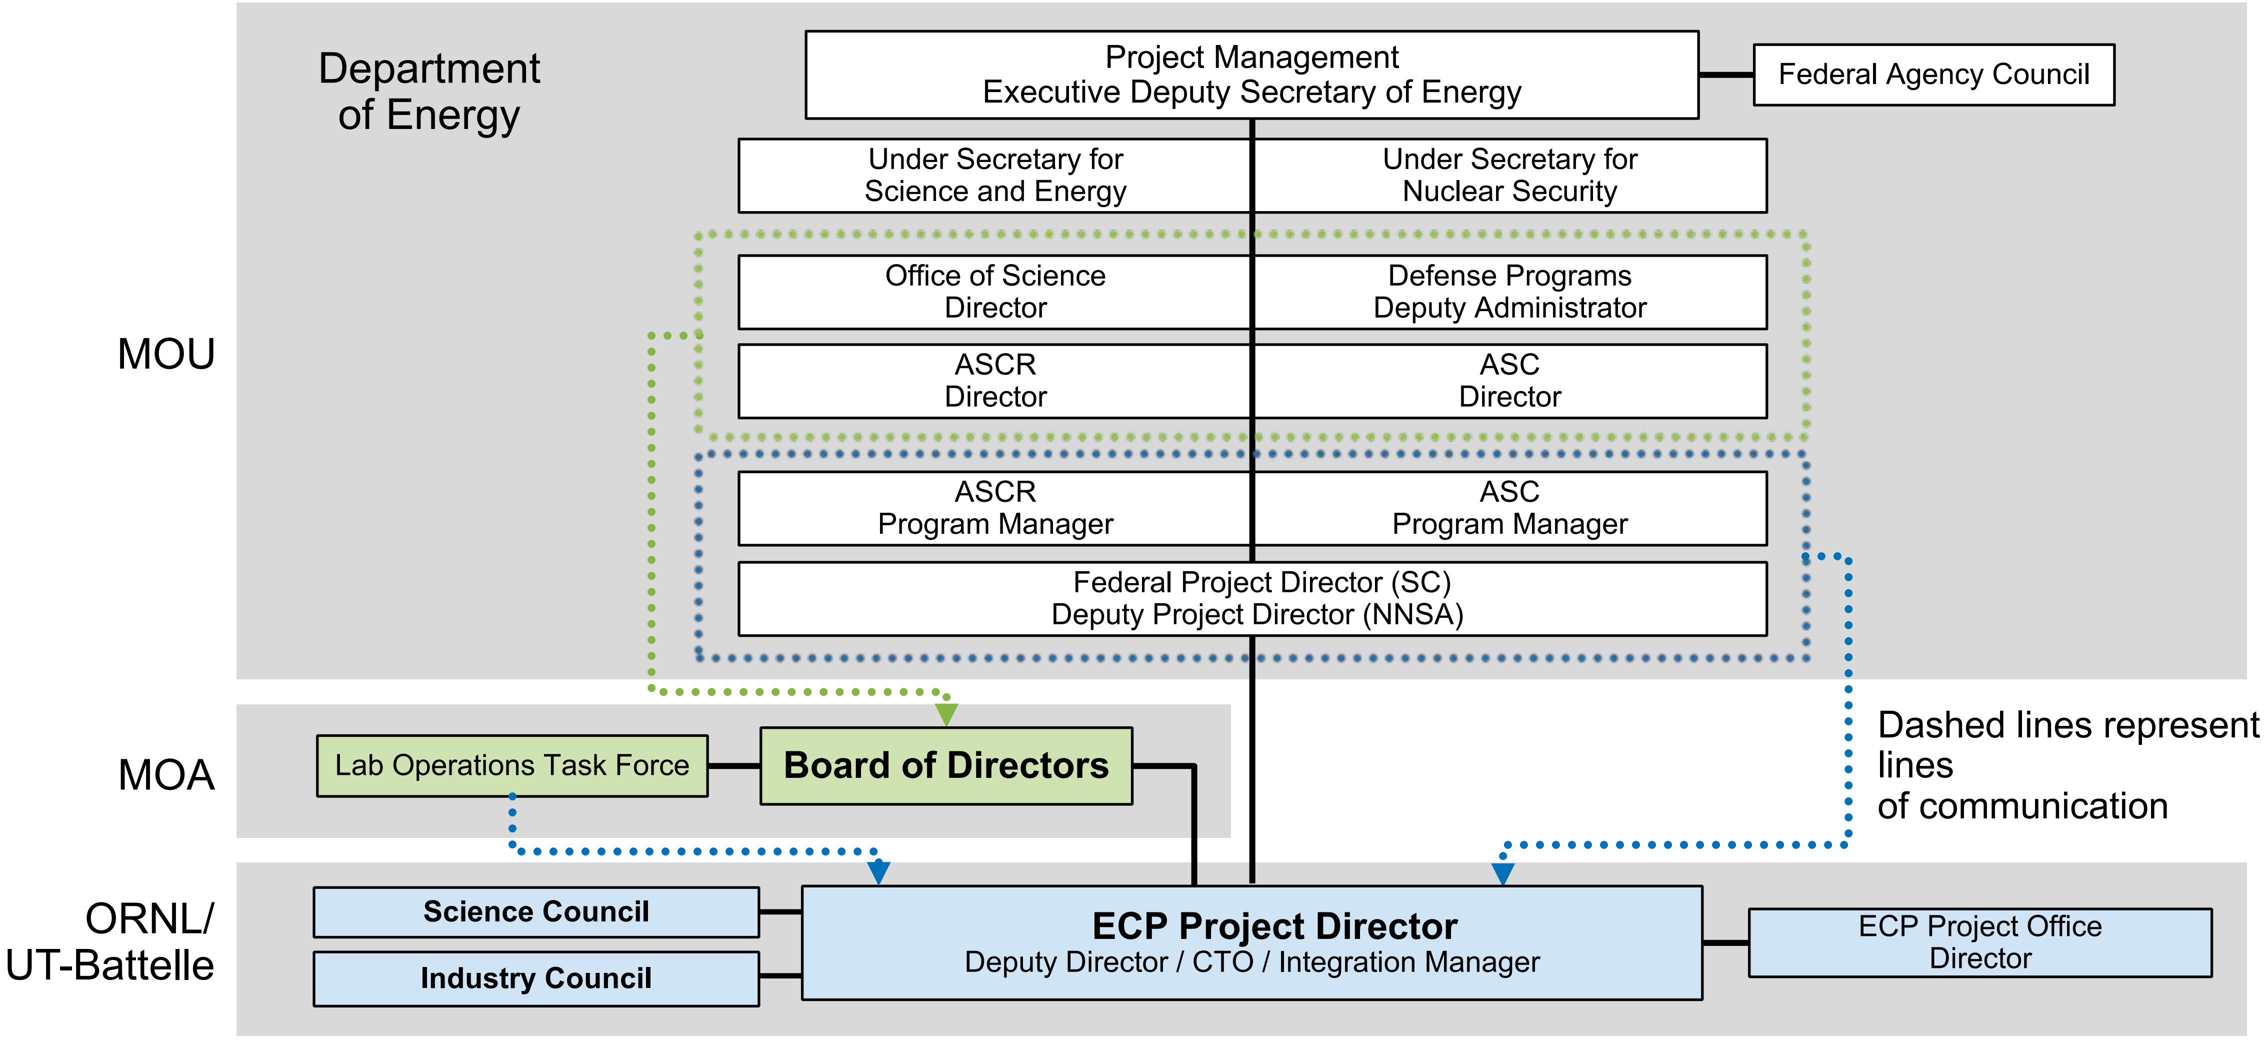

# Capable exascale system applications will deliver broad coverage of 6 strategic pillars

## National security

Stockpile stewardship

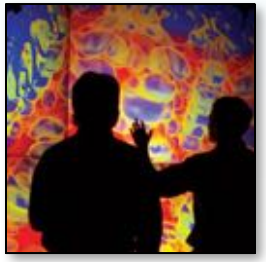

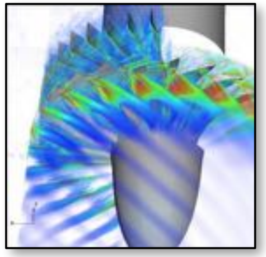

## Energy security

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

EXASCALE COMPUTING PROJECT

# Technology Overview

## Software

- Build a comprehensive and coherent software stack
  - Enable application developers to productively write highly parallel applications that can portably target diverse exascale architectures
- Extending current technologies to exascale where possible, performing R&D required to conceive of new approaches where necessary
  - Coordinate with vendor efforts to develop software other than what is typically done by vendors, develop common interfaces or services
  - Develop and deploy high-quality and robust software products

## Hardware

- Fund R&D to design hardware meeting ECP targets for application performance, power efficiency, and resilience
- Issue **PathForward and PathForward II** Hardware Architecture R&D contracts that deliver:
  - Conceptual exascale node and system designs
  - Analysis of performance improvement on conceptual system design
  - Technology demonstrators to quantify performance gains over existing roadmaps
  - Support for active industry engagement in ECP holistic co-design efforts

# 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

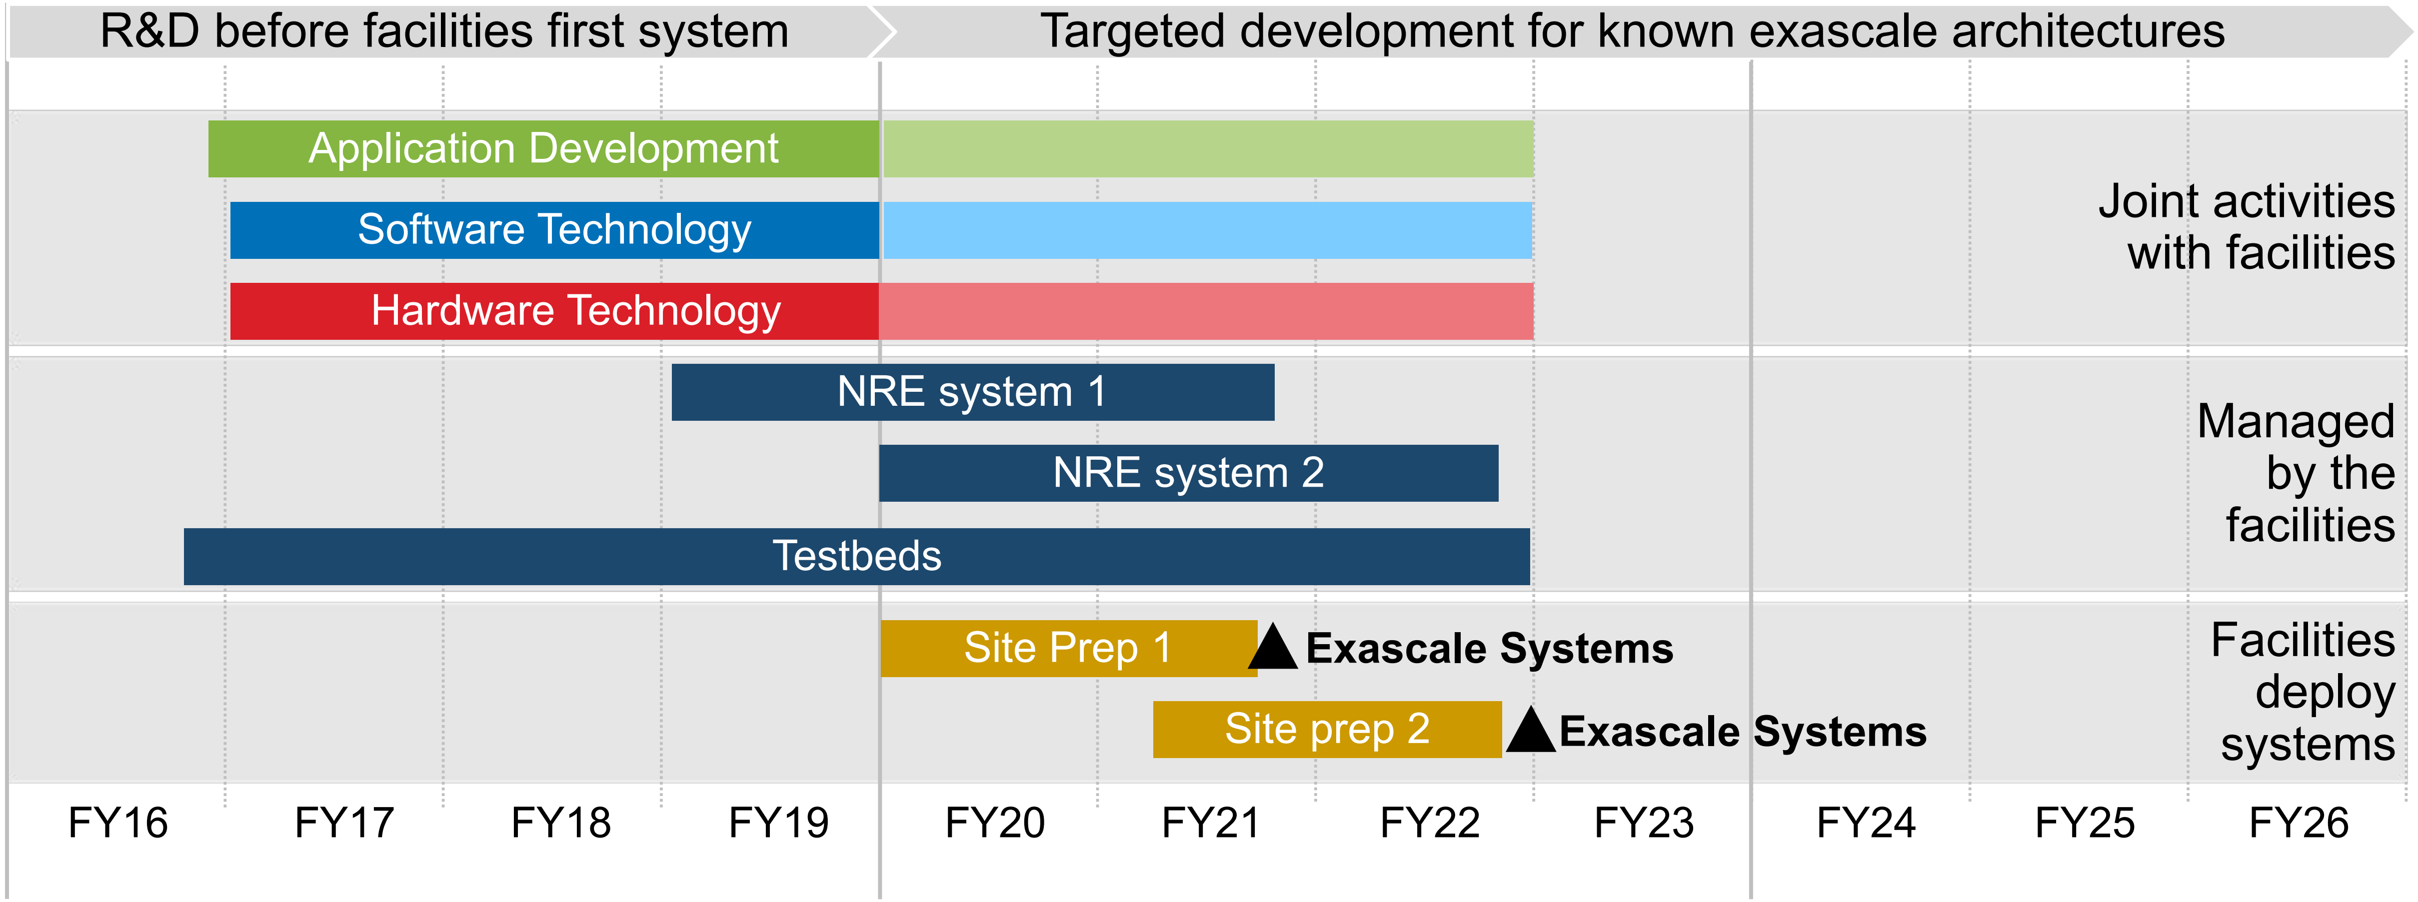




EXASCALE COMPUTING PROJECT

# The top ECP risks include the following

- Risk ID 1001: Insufficient funding from SC and NNSA for the life of the project
- Risk 2001: Unable to recruit and/or retain qualified staff needed to execute R&D
- Risk 2023: Availability and enforcement of programming model standards insufficient for portable application development and performance - split, application developers not willing to change models; ST: no adequate models to move to.
- Risk 4010: PathForward designs fail to meet ECP requirements

# Current ECP Status

The Mission Need Statement was jointly approved by the Office of Science and NNSA on April 14, 2016

Critical Decision-0, Approve Mission Need, was approved by the Project Management Executive (PME) on July 28, 2016

Critical Decision 1, Alternative Selection and Cost Range for Exascale Computing Project, was signed by the PME on January 3, 2017

Cost Range: $3.5B–$5.7B

Critical Decision 3A, Long Lead Procurements, was signed by the PME on January 3, 2017

Cost: $694M which includes funding for vendor hardware and software partnerships and prototype testbeds

# Planned outcomes of the ECP

- Important applications running at exascale in 2021, producing useful results
- A full suite of mission and science applications ready to run on the 2023 exascale systems
- A large cadre of computational scientists, engineers, and computer scientists who will be an asset to the nation long after the end of ECP
- An integrated software stack that supports exascale applications
- Results of PathForward R&D contract with vendors are integrated into exascale systems and are in vendors' product roadmaps
- Industry and mission critical applications have been prepared for a more diverse and sophisticated set of computing technologies, carrying US supercomputing well into the future


ECP
DOE LABORATORIES & AGENCY PARTNERS
22
PRIVATE SECTOR PARTNERS
9
UNIVERSITY RESEARCH PARTNERS
39
INDUSTRY COUNCIL MEMBERS
18

# Discussion

EXASCALE COMPUTING PROJECT

# Back-up slides

EXASCALE COMPUTING PROJECT

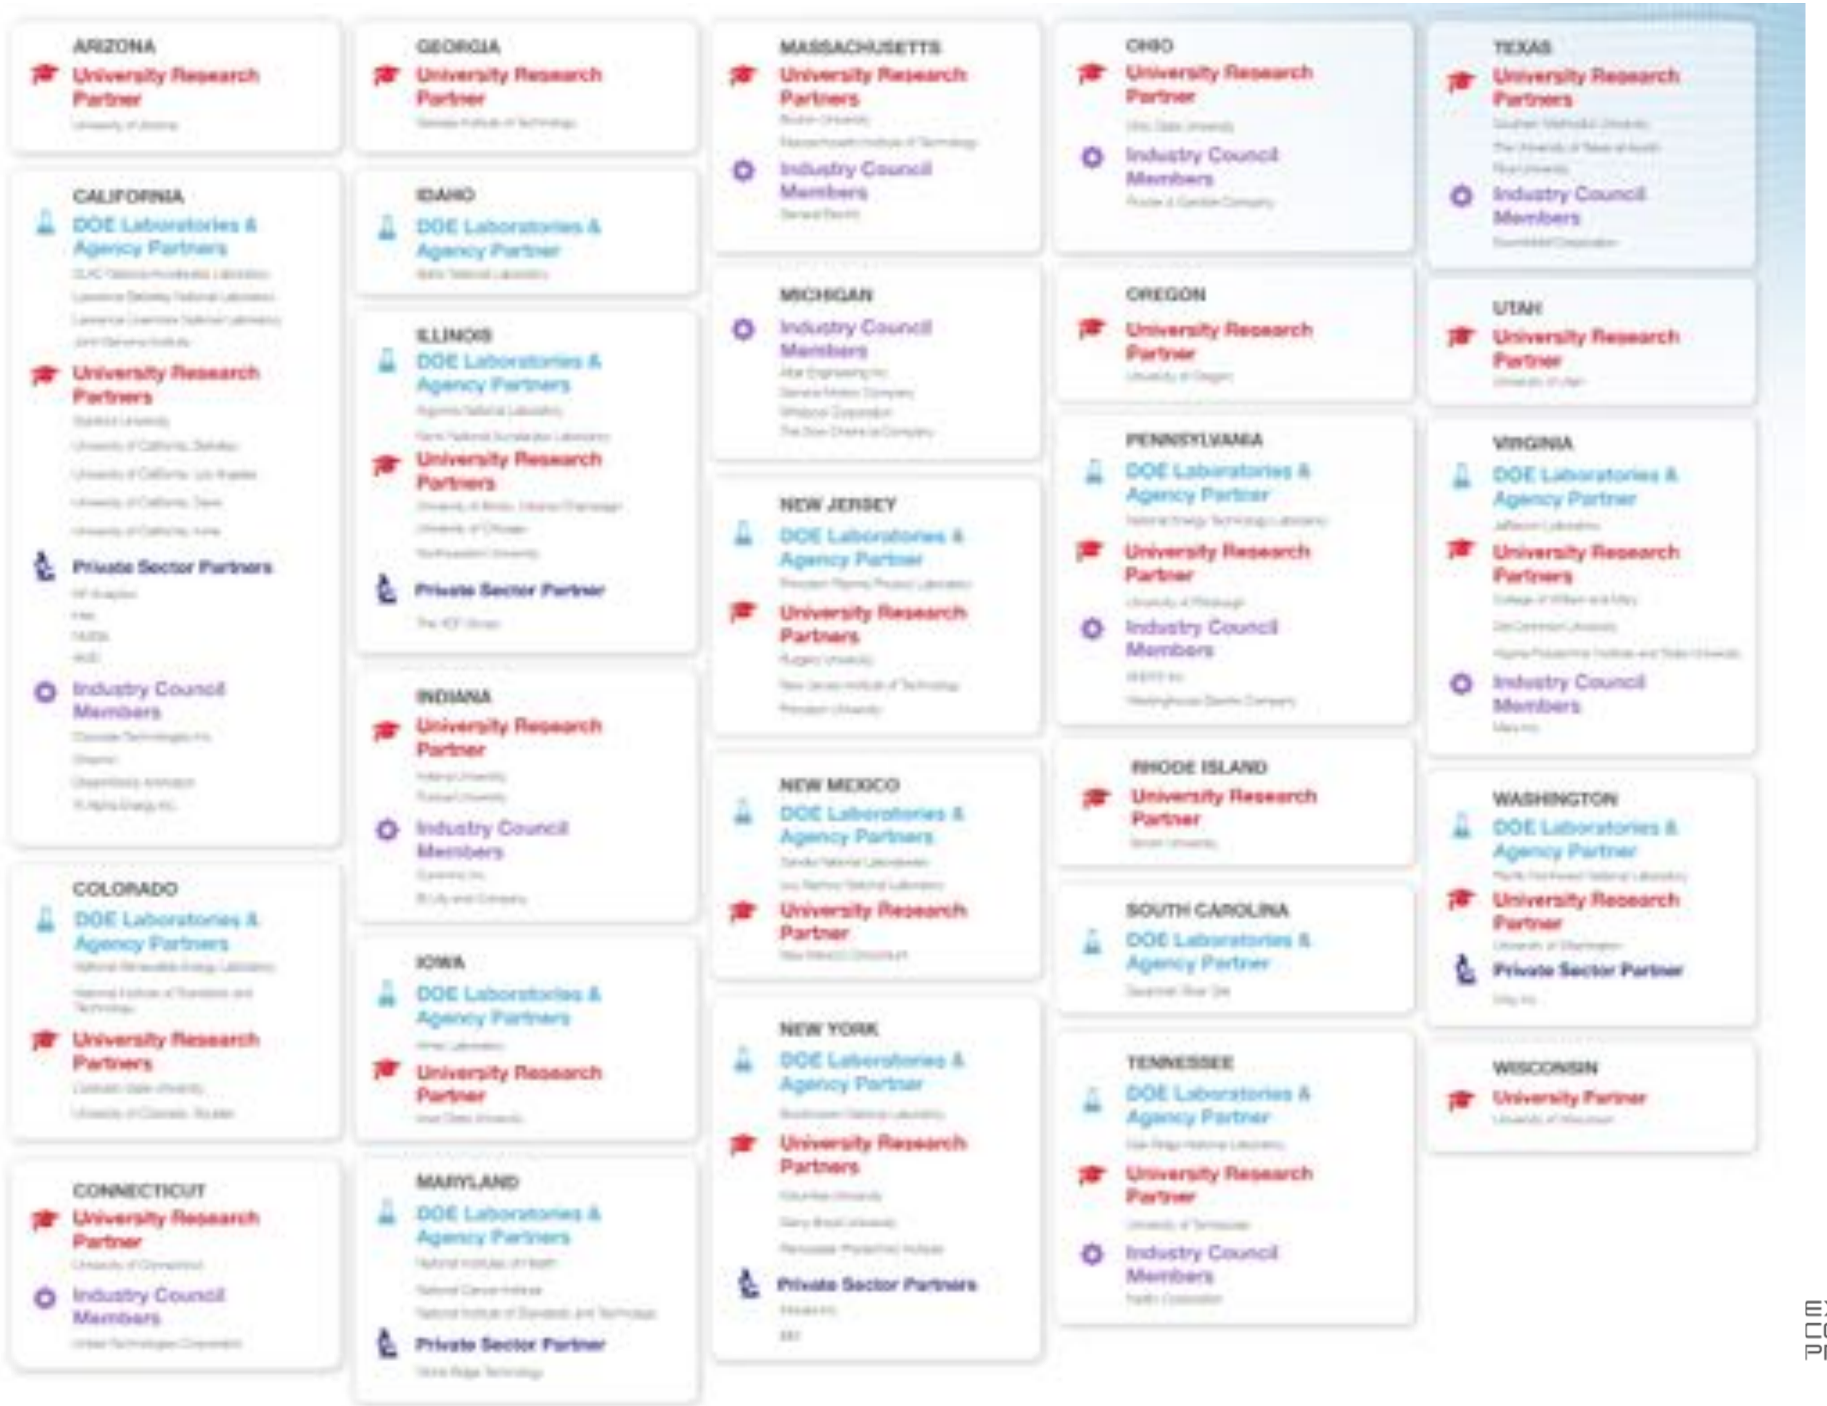


EXASCALE COMPUTING PROJECT

# What the ECP is addressing partially, or not at all

- Only partially tackling convergence of simulation and data analytics
  - Hope to do more, given sufficient funding
  - Deep learning: funding few applications so far, hope to do more but vendors already investing a lot; the number of applications is exploding
    - Do technology watch, find gaps in coverage, be very selective in what we do in ECP
  - Would be good to develop motifs along the lines of Colella's motifs
- Post Moore's Law technologies
  - out of scope for ECP
- Basic research on new programming models
  - Insufficient time to determine their value or deliver production quality implementations