Update on the Exascale Computing Project

Paul Messina, ECP Director

Stephen Lee, ECP Deputy Director

ASCAC Meeting, Washington, DC

American Geophysical Union
December 20, 2016

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

www.ExascaleProject.org

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

- Develop applications that will tackle a broad spectrum of mission critical problems of unprecedented complexity with unprecedented performance
- Contribute to the economic competitiveness of the nation
- Support national security
- Develop a software stack, in collaboration with vendors, that is exascale-capable and is usable on smaller systems by industry and academia
- Train a large cadre of computational scientists, engineers, and computer scientists who will be an asset to the nation long after the end of ECP
- Partner with vendors to develop computer architectures that support exascale applications
- Revitalize the US HPC vendor industry
- Demonstrate the value of comprehensive co-design

Exascale Computing Project Goals

Develop scientific, engineering, and large-data applications that exploit the emerging, exascale-era computational trends caused by the end of Dennard scaling and Moore's law	Create software that makes exascale systems usable by a wide variety of scientists and engineers across a range of applications	Enable by 2021 and 2023 at least two diverse computing platforms with up to 50× more computational capability than today's 20 PF systems, within a similar size, cost, and power footprint	Help ensure continued American leadership in architecture, software and applications to support scientific discovery, energy assurance, stockpile stewardship, and nonproliferation programs and policies
Foster application development	Ease of use	Rich exascale ecosystem	US HPC leadership

ECP EXASCALE COMPUTING PROJECT

ECP has formulated a holistic approach that uses co-design and integration to achieve 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

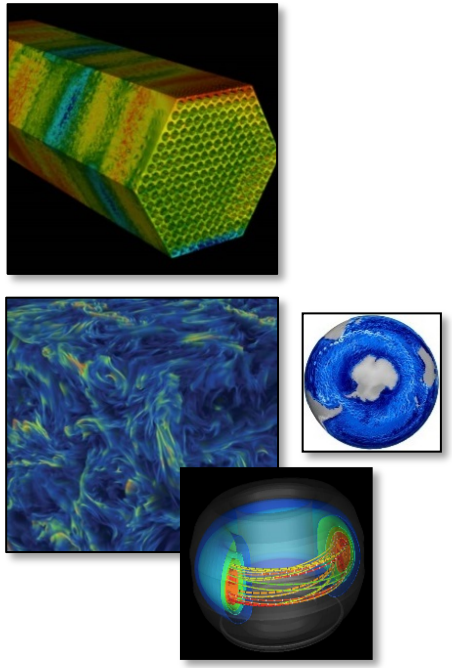

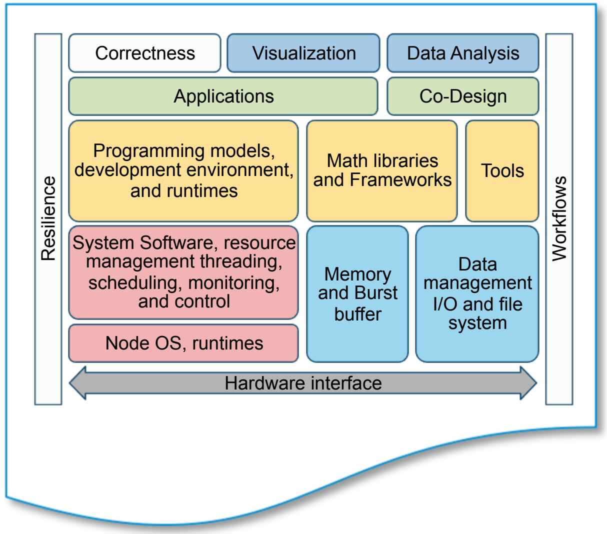

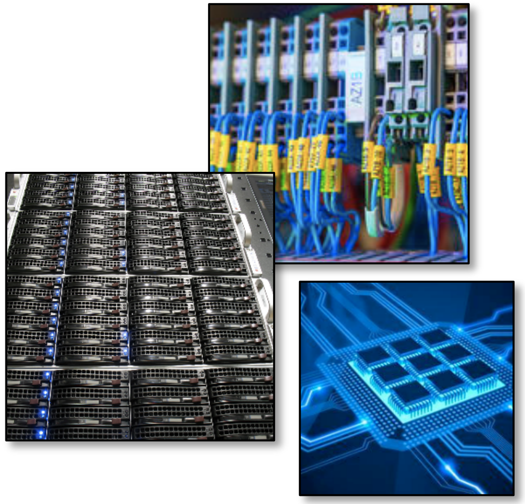

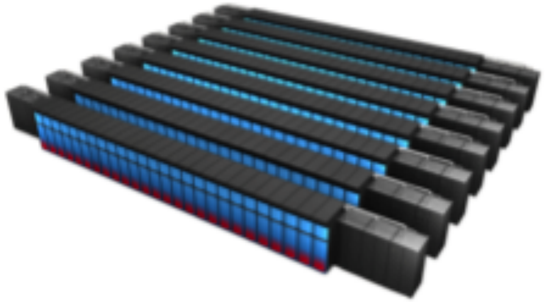

ECP's work encompasses applications, system software, hardware technologies and architectures, and workforce development

ECP EXASCALE COMPUTING PROJECT

What is a *capable* exascale computing system?

A capable exascale computing system requires an entire computational ecosystem that:

- 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

The ECP Plan of Record

- A 7-year project that follows the *holistic/co-design* approach, which runs through 2023 (including 12 months of schedule contingency)
  - To meet the ECP goals
- Enable an initial exascale system based on *advanced architecture* and 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 facilities

What is an exascale *advanced architecture*?

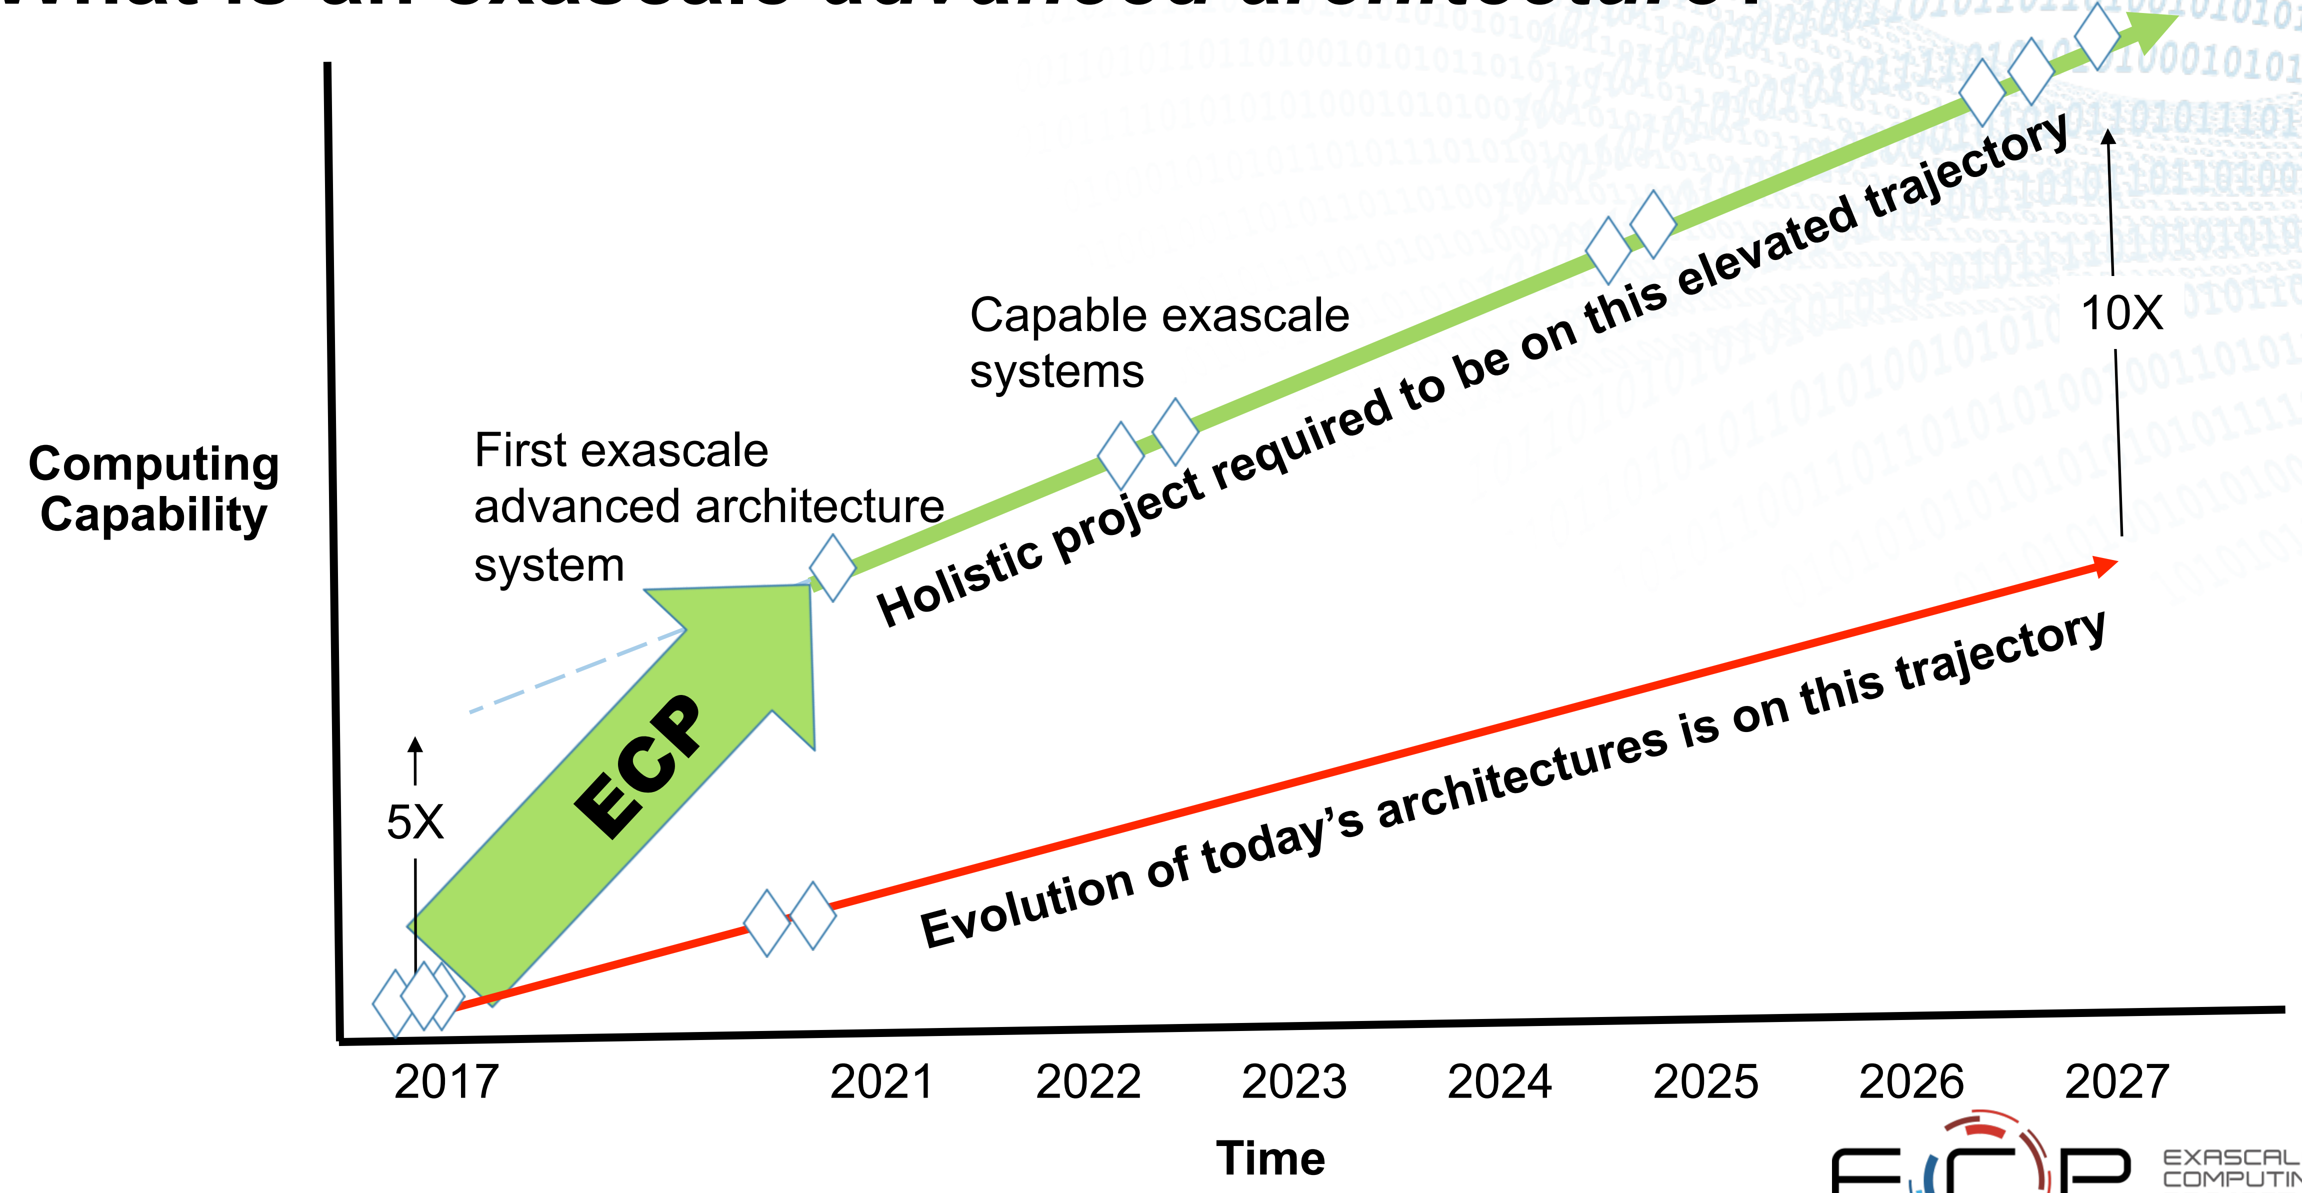

Reaching the Elevated Trajectory will require Advanced and Innovative Architectures

In order to reach the elevated trajectory, advanced architectures must be developed that make a big leap in:

- Parallelism
- Memory and Storage
- Reliability
- Energy Consumption

> The exascale advanced architecture developments benefit all future U.S. systems on the higher trajectory

In addition, the exascale advanced architecture will need to solve emerging data science and machine learning problems in addition to the traditional modeling and simulations applications.

High-level ECP technical project schedule

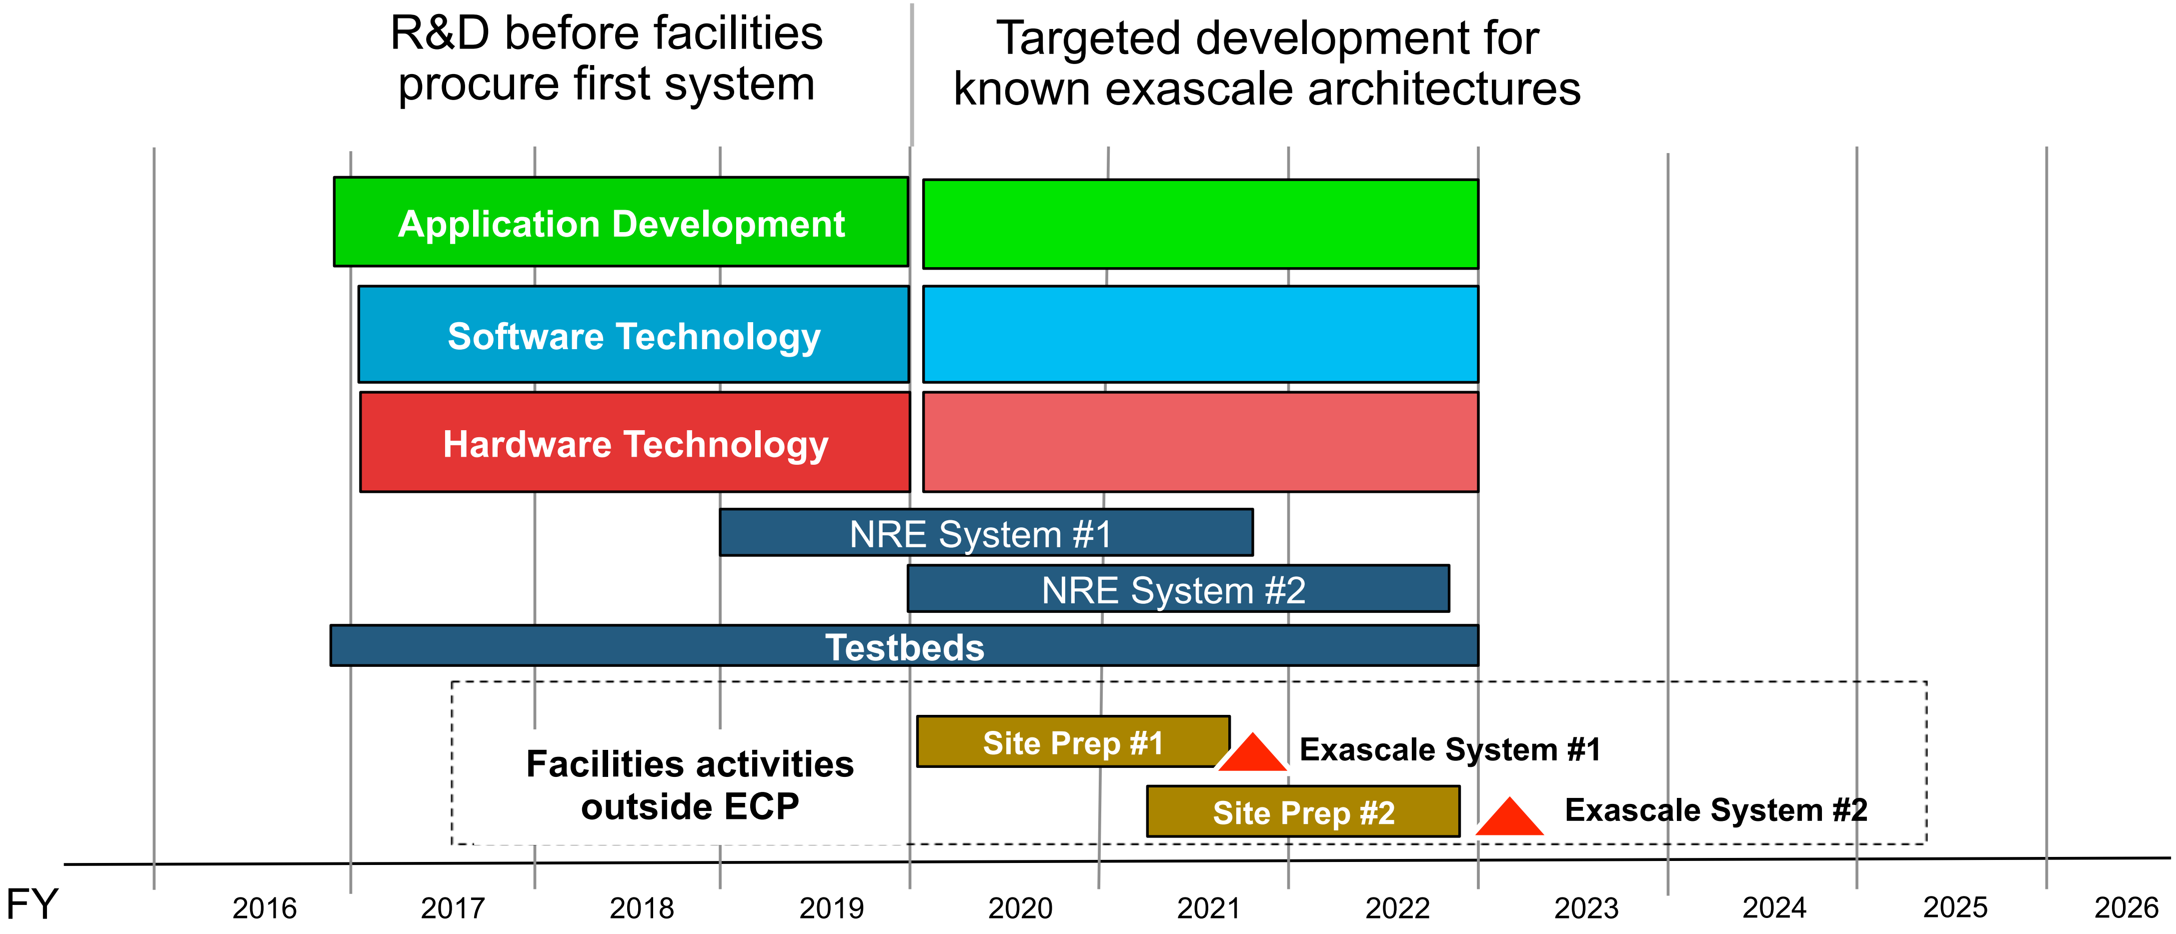

ECP leadership team

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

EXASCALE COMPUTING PROJECT

ECP WBS

Exascale Computing Project
1. Paul Messina

- **Project Management** – 1.1 Kathlyn Boudwin
  - Project Planning and Management – 1.1.1 Kathlyn Boudwin
  - Project Controls & Risk Management – 1.1.2 Monty Middlebrook
  - Business Management – 1.1.3 Dennis Parton
  - Procurement Management – 1.1.4 Willy Besancenez
  - Information Technology and Quality Management – 1.1.5 Doug Collins
  - Communications & Outreach – 1.1.6 Mike Bernhardt
  - Integration – 1.1.7 Julia White
- **Application Development** – 1.2 Doug Kothe
  - DOE Science and Energy Apps – 1.2.1 Andrew Siegel
  - DOE NNSA Applications – 1.2.2 Bert Still
  - Other Agency Applications – 1.2.3 Doug Kothe
  - Developer Training and Productivity – 1.2.4 Ashley Barker
  - Co-Design and Integration – 1.2.5 Phil Colella
- **Software Technology** – 1.3 Rajeev Thakur
  - Programming Models and Runtimes – 1.3.1 Rajeev Thakur
  - Tools – 1.3.2 Jeff Vetter
  - Mathematical and Scientific Libraries and Frameworks – 1.3.3 Mike Heroux
  - Data Management and Workflows – 1.3.4 Rob Ross
  - Data Analytics and Visualization – 1.3.5 Jim Ahrens
  - System Software – 1.3.6 Martin Schulz
  - Resilience and Integrity – 1.3.7 Al Geist
  - Co-Design and Integration – 1.3.8 Rob Neely
- **Hardware Technology** – 1.4 Jim Ang
  - PathForward Vendor Node and System Design – 1.4.1 Bronis de Supinski
  - Design Space Evaluation – 1.4.2 John Shalf
  - Co-Design and Integration – 1.4.3 Jim Ang
  - LeapForward Vendor Node And System Design – 1.4.4 TBD
- **Exascale Systems** – 1.5 Terri Quinn
  - NRE – 1.5.1 Terri Quinn
  - Testbeds – 1.5.2 Terri Quinn
  - Co-design and Integration – 1.5.3 Susan Coghlan

ECP application, co-design center, and software project awards

EXASCALE COMPUTING PROJECT

NEWS RELEASE

The Exascale Computing Project Awards $34 Million for Software Development

OAK RIDGE, Tenn., Nov. 10, 2016 – The Department of Energy's Exascale Co[...] (ECP) today announced the selection of 35 software development proposals [...] research and academic organizations.

The awards for the first year of funding total $34 million and cover many co[...] software stack for exascale systems, including programming models and run[...] mathematical libraries and frameworks, tools, lower-level system software, [...] and I/O, as well as in situ visualization and data analysis.

EXASCALE COMPUTING PROJECT

[...]EWS RELEASE

The Exascale Computing Project Announces $48 Million to Establish Four Exascale Co-Design Centers

OAK RIDGE, Tenn., Nov. 11, 2016 – The Department of Energy's Exascale Computing Project (ECP) today announced that it has selected four co-design centers as part of a 4 year, $48 million funding award. The first year is funded at $12 million, and is to be allocated evenly among the four award recipients.

[...]he ECP is responsible for the planning, execution, and delivery of technologies necessary for a capable [...]xascale ecosystem to support the nation's exascale imperative including software, [...] [...]nd early testbed platforms.

EXASCALE COMPUTING PROJECT

NEWS RELEASE

For Immediate Distribution

The Exascale Computing Project (ECP) Announces $39.8 million in First-Round Application Development Award[...]

OAK RIDGE, Tenn., Sept. 07, 2016 – The Department of Energy's Exascale Computi[...] (ECP) today announced its first round of funding with the selection of 15 applicatio[...] development proposals for full funding and seven proposals for seed funding, repr[...] teams from 45 research and academic organizations.

The awards, totaling $39.8 million, target advanced modeling and simulation solu[...] specific challenges supporting key DOE missions in science, clean energy and natio[...] as well as collaborations such as the Precision Medicine Initiative with the Nation[...] Health's National Cancer Institute.

EXASCALE COMPUTING PROJECT

ECP Mission Need Defines the Application Strategy

Support DOE science *and* energy missions	Meet national security needs	Key science and technology challenges to be addressed with exascale
• Discover and characterize next-generation **materials** • Systematically understand and improve **chemical processes** • Analyze the extremely large datasets resulting from the next generation of **particle physics** experiments • Extract knowledge from **systems-biology** studies of the microbiome • Advance **applied energy** technologies (e.g., whole-device models of plasma-based fusion systems)	• Stockpile Stewardship Annual Assessment and Significant Finding Investigations • Robust uncertainty quantification (UQ) techniques in support of lifetime extension programs • Understanding evolving nuclear threats posed by adversaries and in developing policies to mitigate these threats	• Materials discovery and design • Climate science • Nuclear energy • Combustion science • Large-data applications • Fusion energy • National security • Additive manufacturing • Many others!

EXASCALE COMPUTING PROJECT

ECP Applications Deliver Broad Coverage of Strategic Pillars

Initial selections consist of 15 application projects + 7 seed efforts

National Security	Energy Security	Economic Security	Scientific Discovery	Climate and Environmental Science	Healthcare
• Stockpile Stewardship	• Turbine Wind Plant Efficiency • Design/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 (S) • Biofuel Catalyst Design (S)	• Additive Manufacturing of Qualifiable Metal Parts • Urban Planning (S) • Reliable and Efficient Planning of the Power Grid (S) • Seismic Hazard Risk Assessment (S)	• Cosmological Probe of the Standard Model (SM) of Particle Physics • Validate Fundamental Laws of Nature (SM) • 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 (S)	• Accurate Regional Impact Assessment of Climate Change • Stress-Resistant Crop Analysis and Catalytic Conversion of Biomass-Derived Alcohols • Metegenomics for Analysis of Biogeochemical Cycles, Climate Change, Environ Remediation (S)	• Accelerate and Translate Cancer Research

EXASCALE COMPUTING PROJECT

Application Development: Partnerships

INDUSTRY
Intel
NVIDIA
Stone Ridge Technology

UNIVERSITY
College of William & Mary
Columbia University
University of Indiana
University of Utah
University of Illinois
University of Tennessee
Rutgers University
University of California, Irvine & Los Angeles
Colorado State University
University of Chicago
Virginia Tech.
University of Colorado

LABORATORY PARTNERSHIPS	
Ames	Argonne
Brookhaven	Fermi National Accelerator
Idaho	Oak Ridge
Los Alamos	Lawrence Berkeley
Lawrence Livermore	Pacific Northwest
National Renewable Energy Laboratory	Thomas Jefferson National Accelerator Facility
National Energy Technology Laboratory	Princeton Plasma Physics Laboratory
SLAC National Accelerator	Sandia

Meeting ECP Mission Need Application Requirements

- Current ECP application portfolio appropriately targets high-priority strategic problems in key DOE science, energy, national security programs
- Some application gaps remain relative to ECP Mission Need requirements
  - Data analytic computing (DAC) applications per the NSCI Strategic Plan
  - High priority applications for some NSCI deployment agencies (NASA, NOAA, DHS, FBI)
- ECP has not yet adequately targeted the talented U.S. computational science leadership in academia and industry
- A Request for Information (RFI) for candidate ECP application projects led by academia or industry is currently being drafted
  - Tentative release date of Jan 2017 with an informational videoconference scheduled shortly thereafter
  - Industry participation will adhere to standard DOE cost share requirements and IP provisions
  - A subset of RFI responses, determined by DOE, ECP, and external peer review, will be asked to submit full proposals, with final selection of 3-5 projects (sized initially with teams of 2-4 staff) by April 2017
- New Data Analytic Computing methods and applications are of particular interest, in addition to proposals to enhance the current ECP portfolio
  - Critical infrastructure, superfacility, supply chain, image/signal processing, in situ analytics
  - Machine/statistical learning, classification, streaming/graph analytics, discrete event, combinatorial optimization

Application Co-Design (CD)

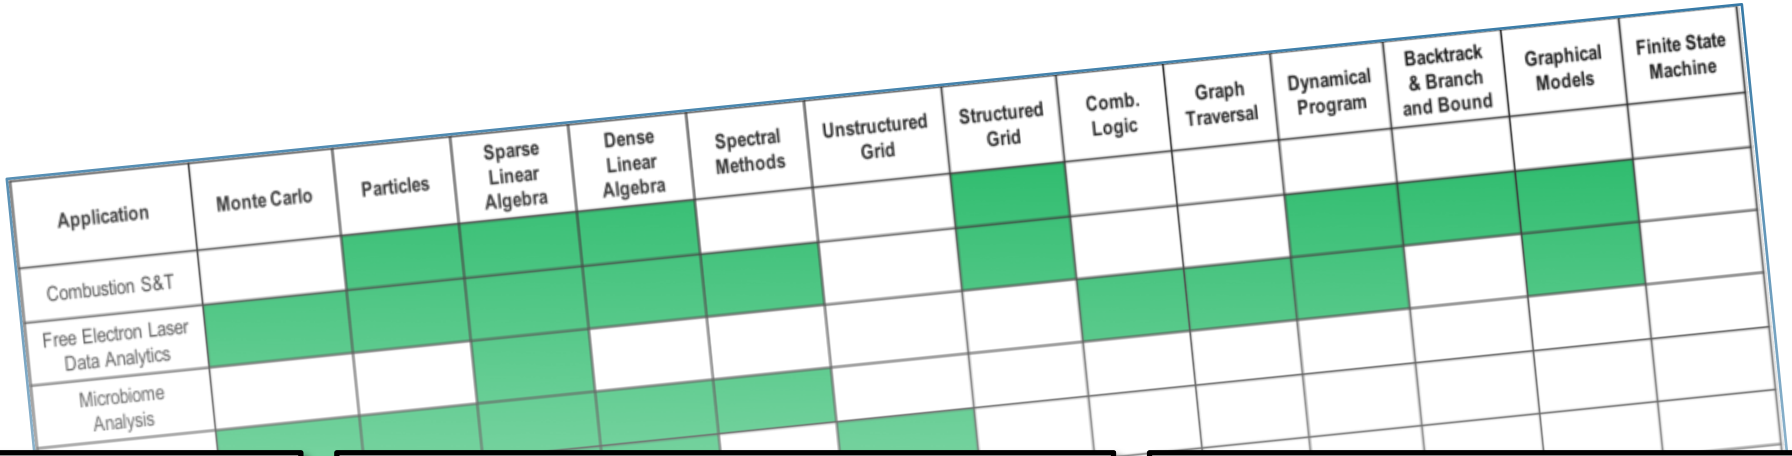

Essential to ensure that applications effectively utilize exascale systems

- 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

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

Ongoing Training: Important for ECP Development Teams

- Training for ECP AD and ST project teams is crucial to keep them abreast of key emerging exascale technologies and productive in integrating them
  - Latest algorithms and methods, high performance libraries, memory and storage hierarchies, on-node and task-based parallelism, application portability, and software engineering design principles and best practices.
- ECP training project will offer both generic and focused training activities through topical workshops, deep-dives, hands-on hackathons, seminars, webinars, videos, and documentation
  - Leverage partnerships with the ASCR and NNSA facilities and complement their existing training programs
  - Model training events on previous facility events such as the ATPESC
  - Disseminate lessons learned, best practices, and other T&P materials to the ECP teams and to the general HPC community through the use of the ECP website.
- Early training activities have been focused on developing training and best practices for Agile software development tools and methodologies
- Near term training will focus on using the newly provisioned ECP testbed systems
- Training materials (documents, slides, videos) will be disseminated broadly whenever possible

Ensuring ECP Development Teams are Productive

- ECP must assess, recommend, develop and/or deploy software engineering tools, methodologies, and/or processes for software development teams and to cultivate and disseminate software engineering best practices across the teams for improved scientific software development
- ECP is currently standing up a Productivity Project modeled in part on the recent ASCR IDEAS project
  - Includes participation from six DOE Labs and one University partner
- The productivity project team will first assess ECP AD and ST productivity needs and then address these needs through a combination of technical deep dives, implementation of software engineering tools, the development of "how to" documents, training, and one-on-one assistance.
- The productivity work will kick-off in January, 2017

EXASCALE COMPUTING PROJECT

Software Technology Summary

- 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

ST Projects (incl. ATDM) Mapped to Software Stack

Resilience
Checkpoint/Restart (VeloC, UNIFYCR), FSEFI, Fault Modeling

Correctness

Visualization
VTK-m, ALPINE, Cinema

Data Analysis
ALPINE, Cinema

Applications

Co-Design

Programming Models, Development Environment, and Runtimes
MPI (MPICH, Open MPI), OpenMP, OpenACC, PGAS (UPC++, Global Arrays), Task-Based (PaRSEC, Legion, DARMA), RAJA, Kokkos, OMPTD, Power steering

Math Libraries/Frameworks
ScaLAPACK, DPLASMA, MAGMA, PETSc/TAO, Trilinos, xSDK, PEEKS, SuperLU, STRUMPACK, SUNDIALS, DTK, TASMANIAN, AMP, FleCSI, KokkosKernels, Agile Comp., DataProp, MFEM

Tools
PAPI, HPCToolkit, Darshan, Perf. portability (ROSE, Autotuning, PROTEAS), TAU, Compilers (LLVM, Flang), Mitos, MemAxes, Caliper, AID, Quo, Perf. Anal.

System Software, Resource Management Threading, Scheduling, Monitoring, and Control
Argo Global OS, Qthreads, Flux, Spindle, BEE, Spack, Sonar

Node OS, low-level runtimes
Argo OS enhancements, SNL OS project

Memory and Burst buffer
Chkpt/Restart (VeloC, UNIFYCR), API and library for complex memory hierarchy (SICM)

Data Management, I/O and File System
ExaHDF5, PnetCDF, ROMIO, ADIOS, Chkpt/Restart (VeloC, UNIFYCR), Compression (EZ, ZFP), I/O services, HXHIM, SIO Components, DataWarehouse

Workflows
Contour, Siboka

Hardware interface

Plan to involve additional Universities and Vendors

- The current set of Software Technology projects span 8 DOE national laboratories,15 universities, and 2 independent software vendors
- In FY17-18, we plan to issue additional solicitations for universities and vendors
- We plan to coordinate our activities with NSF to cover the broader university community through NSF
- ECP has formed an Industry Council. This group will also serve as an interface for engaging industrial users and software vendors.
- Small businesses: The solicitations in FY17-18 will include independent software vendors. DOE SBIR funding could also support software of interest to ECP.

Next Steps

- 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

Hardware Technology Overview

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

Issue *PathForward and LeapForward* 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

DOE labs engage to:

- Participate in evaluation and review of PathForward and LeapForward deliverables
- Lead Design Space Evaluation through Architectural Analysis, and Abstract Machine Models of PathForward/LeapForward designs for ECP's holistic co-design

ECP EXASCALE COMPUTING PROJECT

Overarching Goals for PathForward

- Improve the quality and number of competitive offeror responses to the Capable Exascale Systems RFP
- Improve the offeror's confidence in the value and feasibility of aggressive advanced technology options that would be bid in response to the Capable Exascale Systems RFP
- Improve DOE confidence in technology performance benefit, programmability and ability to integrate into a credible system platform acquisition

EXASCALE COMPUTING PROJECT

PathForward Status

- LLNL is managing the procurement
- Competitive RFP released in June
- 14 PathForward proposals received in August
- 6 proposals selected for SOW negotiations of final Work Packages
  - As of Early December: The First Vendor SOW has completed the LLNL Procurement Review and is on to Livermore Field Office review, will need final review by NNSA/HQ
  - Two more Vendor SOWs are entering the LLNL Procurement Review
  - Final three are pending
- Firm Fixed Price contracts with milestone deliverables and payments
- DOE Advance IP Waivers for vendors that provide ≥ 40% cost share
- Project duration: 3 years
- Targeting Contract Awards in January 2017

Overarching Goals for LeapForward

- Improve the quality and number of competitive offeror responses to the Advanced Architecture Exascale Systems RFP
- Develop Advanced Architecture concepts that arise from Holistic Co-design to impact ECP applications that are addressing:
  – Evolutionary Challenge Problems
  – Revolutionary Challenge Problems
- The Advanced Architecture Exascale System will maintain application performance requirements for a system that is delivered in 2021, but there would need to be tradeoffs in other performance parameters

Hardware Technology Summary

Accelerate innovative hardware technology options that create a rich, competitive HPC ecosystem that supports at least two diverse Capable Exascale Systems, and enhance system and application performance for traditional science and engineering applications, as well as data-intensive and data-analytics applications

- PathForward reduces the Technical Risk for NRE investments in the 2023 Exascale System (ES)
- LeapForward accepts Technical Risk in return for game-changing impact on the U.S. computing eco-system and potential for early deployment in 2021 exascale system
- Establishes a foundation for architectural diversity in the HPC eco-system
- Provides hardware technology expertise and analysis
- Provides an opportunity for inter-agency collaboration under NSCI

ECP EXASCALE COMPUTING PROJECT

ECP's systems acquisition approach

- DOE's Office of Science (SC) and National Nuclear Security Administration (NNSA) programs will procure and install the systems, not ECP
- ECP's requirements will be incorporated into RFP(s)
- ECP will participate in system selection and co-design
- ECP will make substantial investments through NRE contracts to accelerate technologies, add capabilities, improve performance, and lower the cost of ownership of systems
- NRE contracts are coupled to system acquisition contracts

ECP's and SC/NNSA's processes will be tightly coupled and interdependent

ECP EXASCALE COMPUTING PROJECT

Non-Recurring Engineering (NRE) incentivizes awardees to address gaps in their system product roadmaps

- Brings to the product stage promising hardware and software research (ECP, vendor, Lab, etc.) and integrates it into a system
- Includes application readiness R&D efforts
- Must start early enough to impact the system - more than two full years of lead time are necessary to maximize impact

Experience has shown that NRE can substantially improve the delivered system

ECP's plan to accelerate and enhance system capabilities

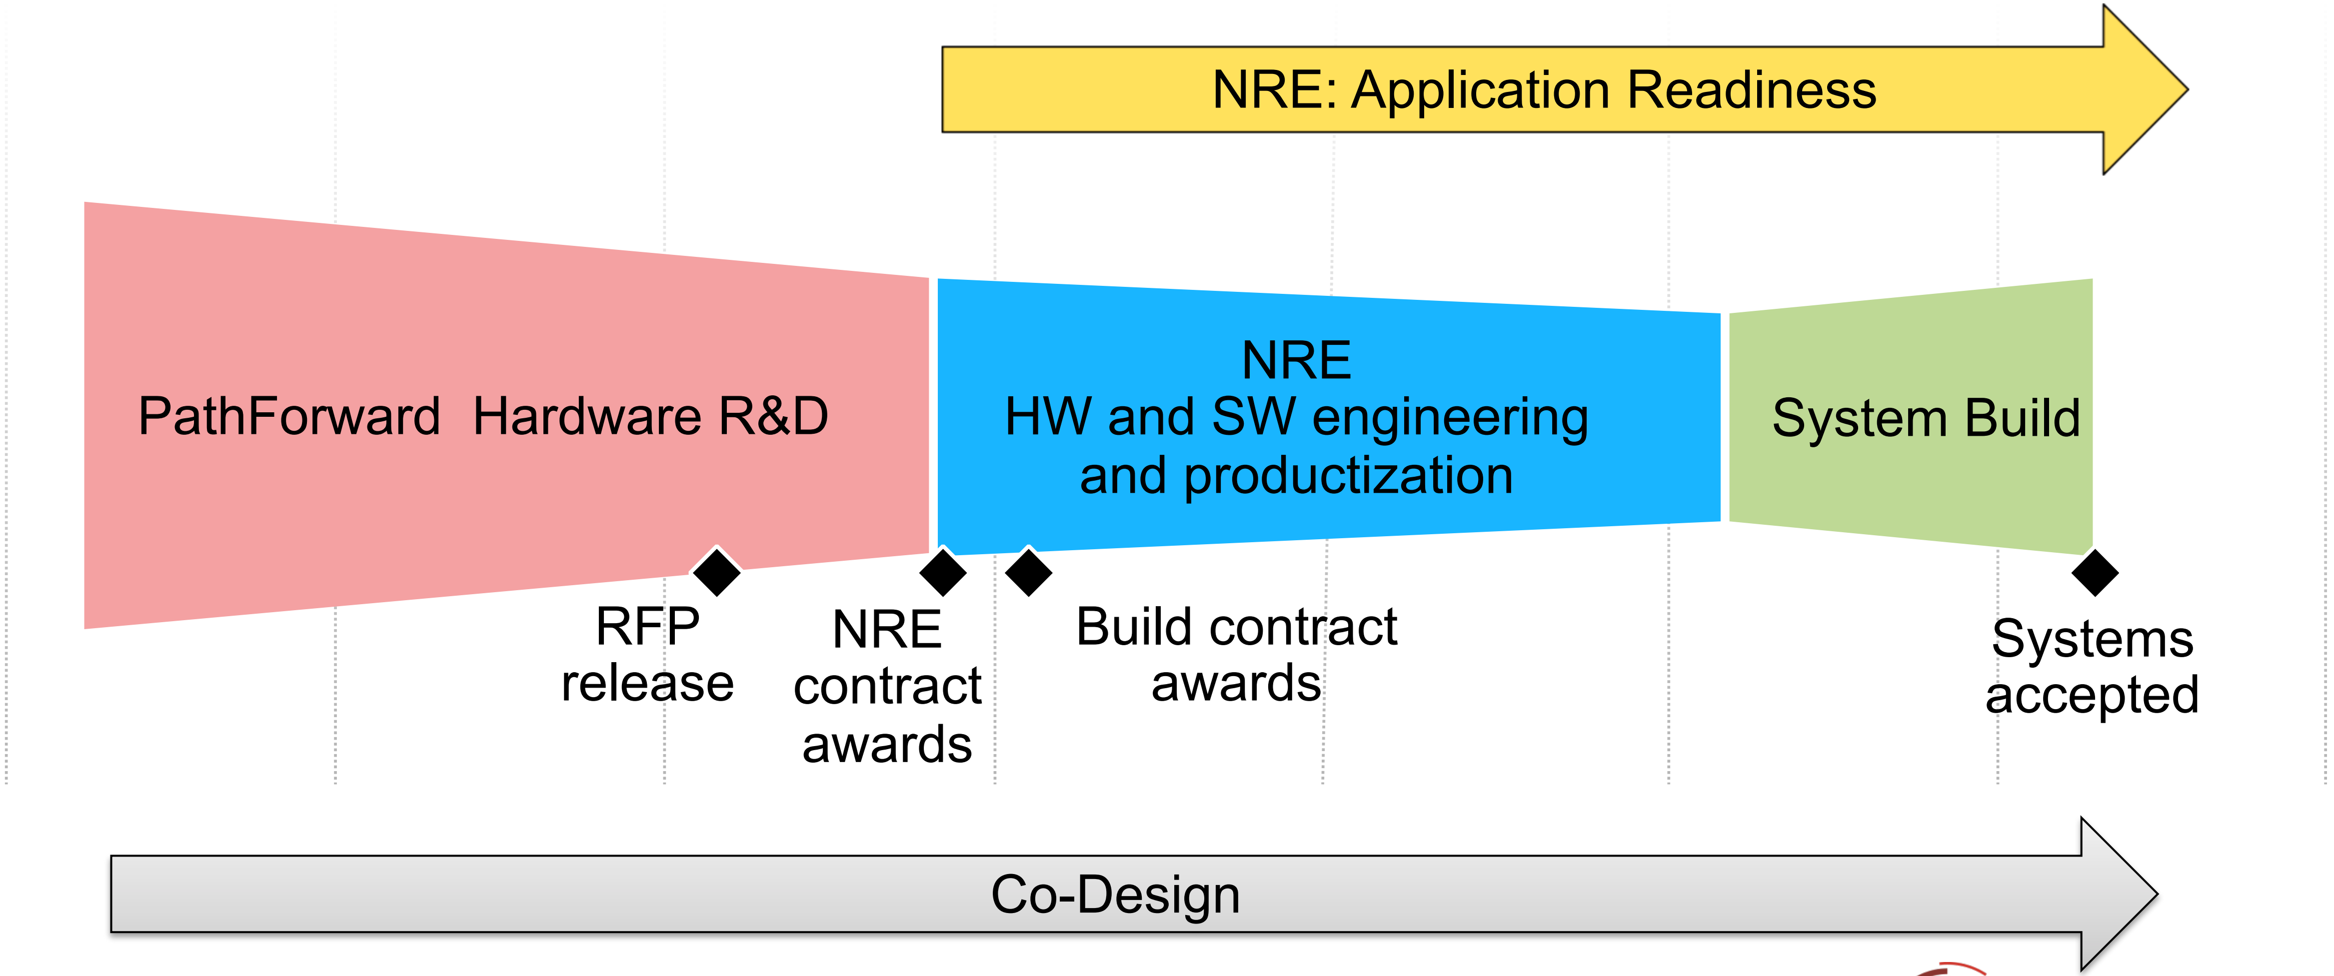

ECP will acquire and operate testbeds for ECP users

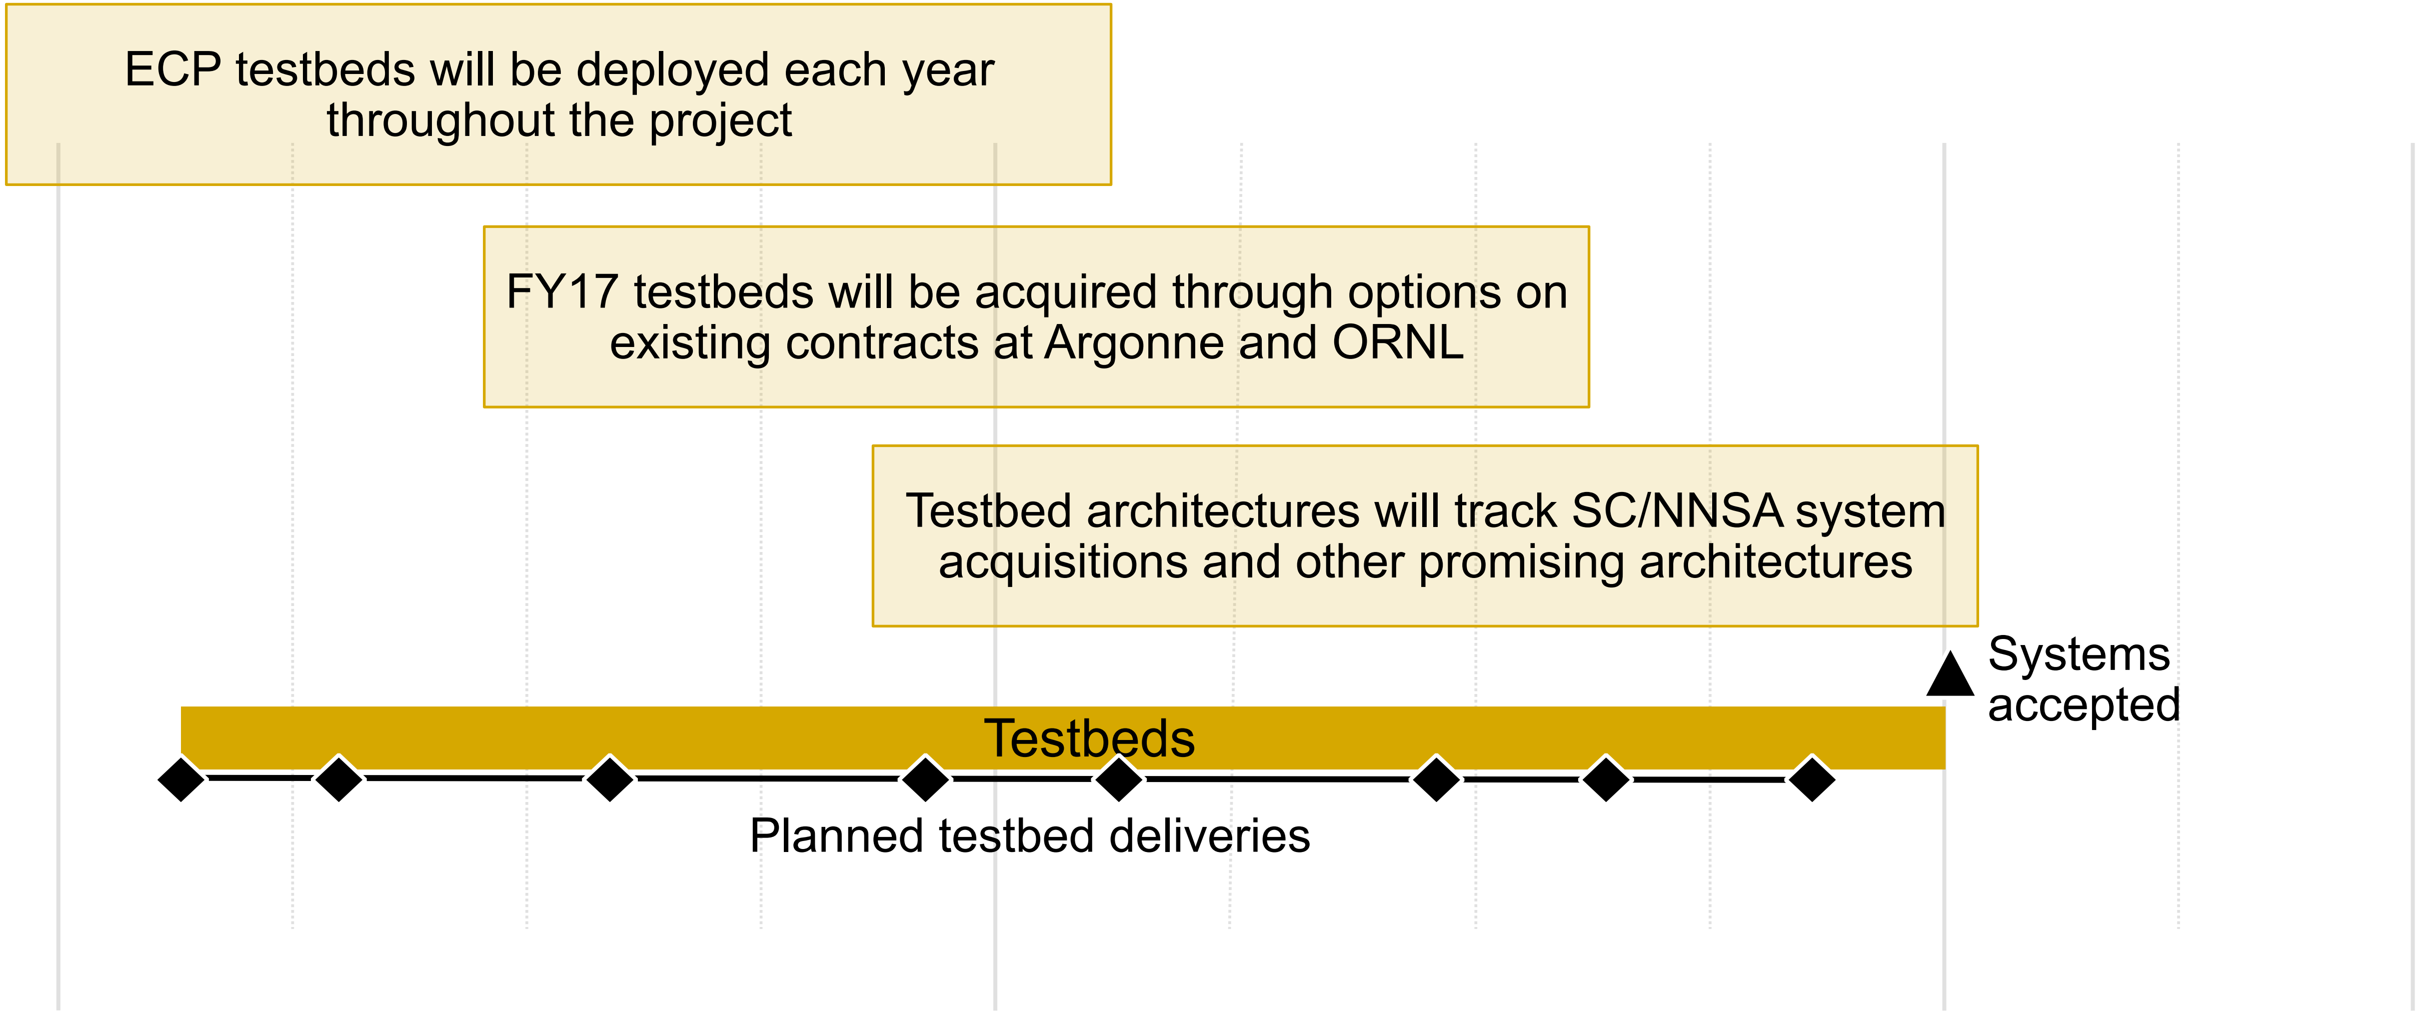

ECP PI and Integration Meeting
November 29 – December 1, 2016

Meeting Goals

Foster ECP Community	Integrate Across Projects	Promote Holistic Co-Design
Exchanged funded project summaries via quad charts and presentations Held parallel sessions for focus areas to meet and strategize	Arranged more than two dozen topic-focused breakout sessions to promote discussion between current and future collaborators	Plenary talks given by ECP leadership and focus area representatives on current and future activities Outlined the objective for integration

- Representatives for all ECP projects; a total of 125 attendees
  - 84% of the thirty-four respondents to a meeting survey Strongly Agreed or Agreed that they “found this meeting to be **worthwhile to attend**” and that “**the topics covered at this meeting were relevant**” to them.
- ECP Annual Meeting: January 31 – February 2, 2017

ECP Extended Project Leadership Team

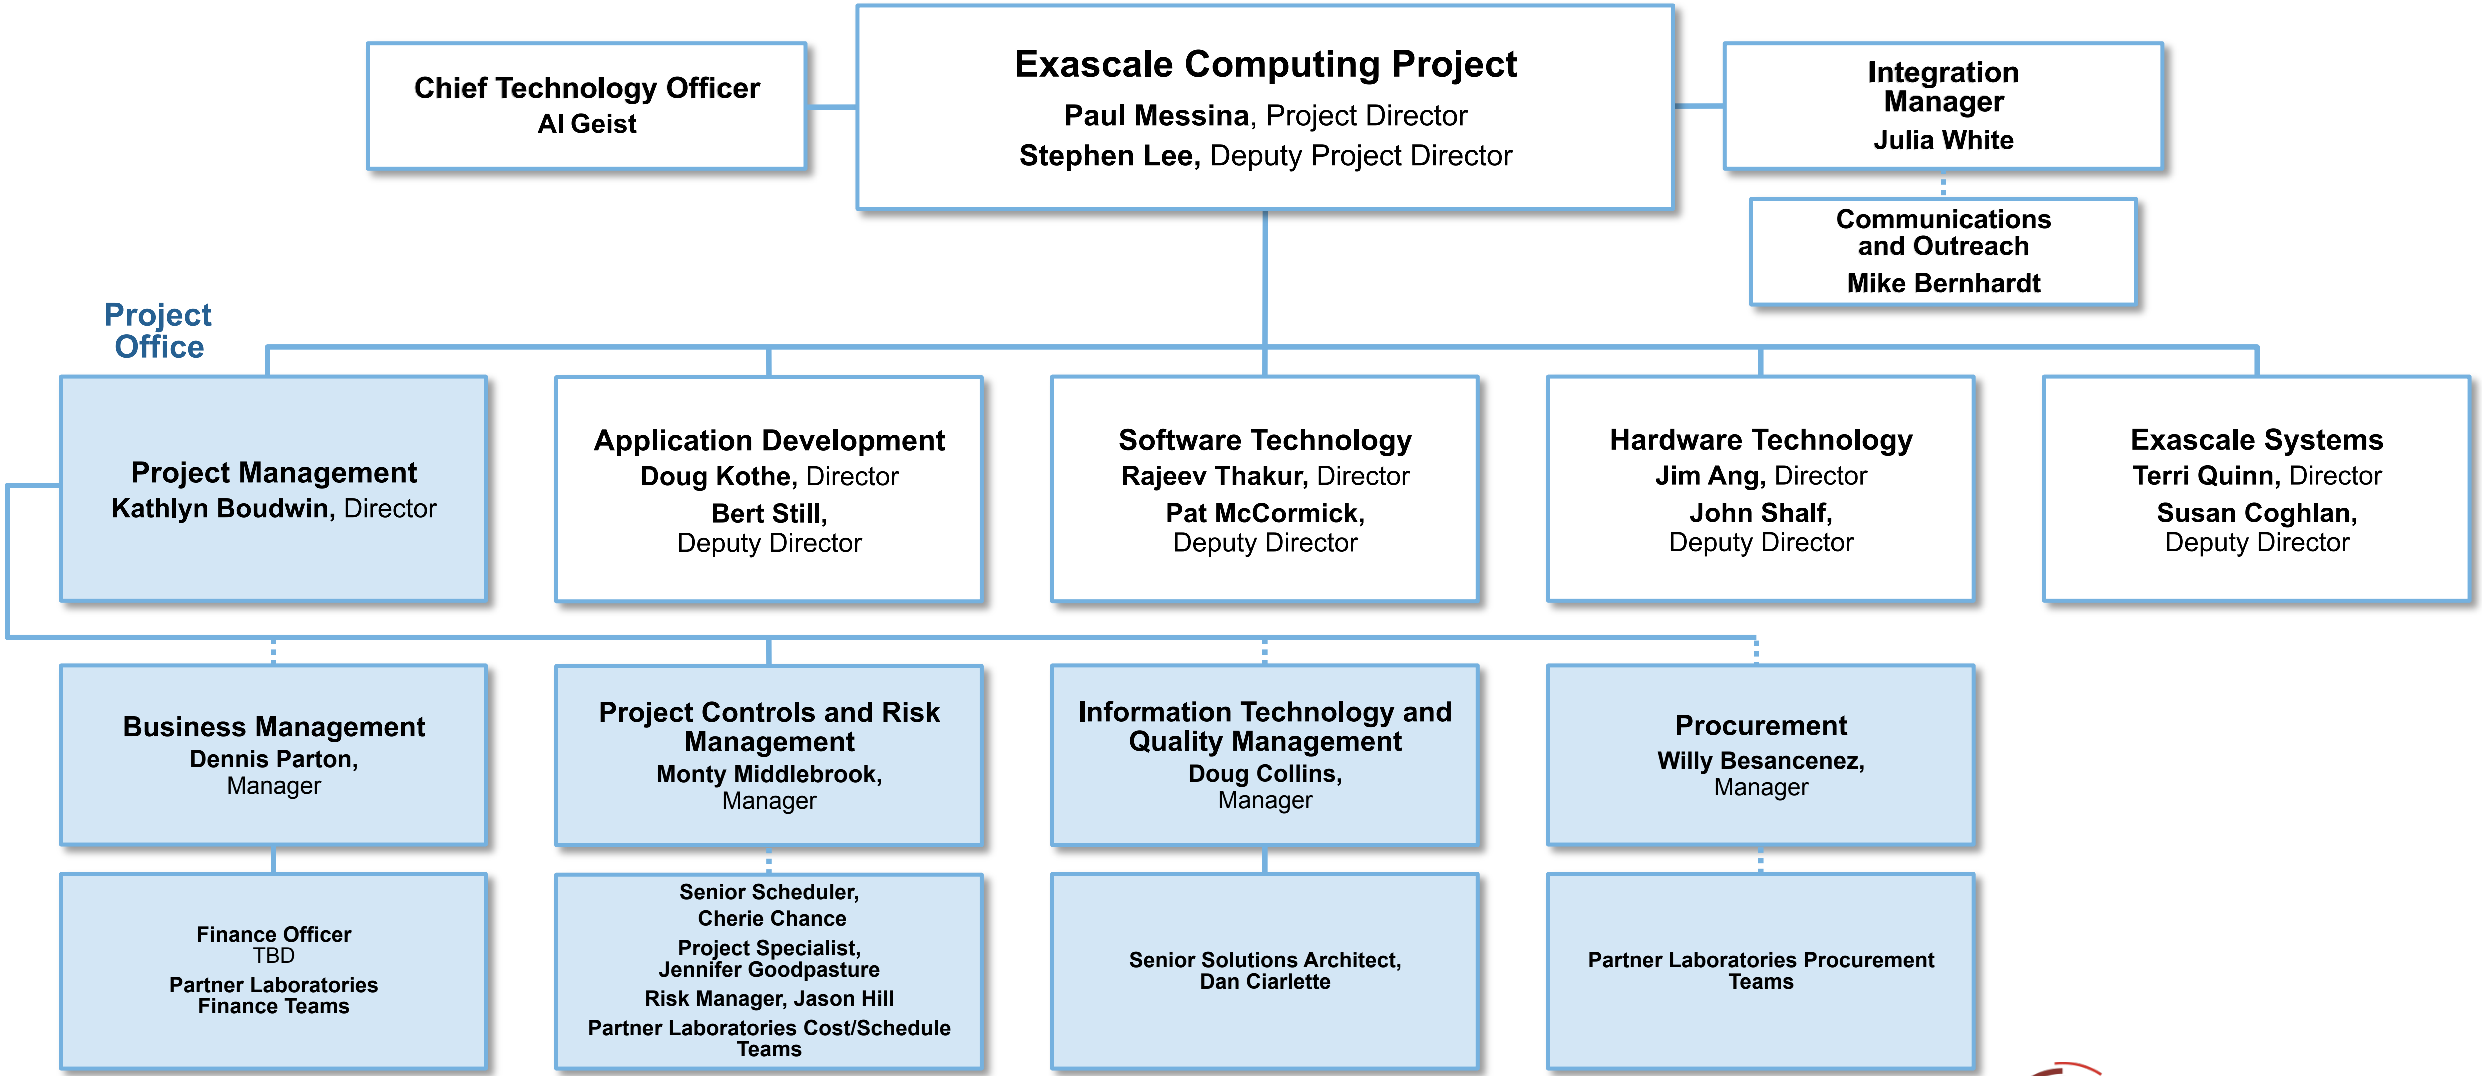

EXASCALE COMPUTING PROJECT

High-level project management milestones (past and future)

- ✓ July 28, 2016: CD-0 achieved
  - Approved mission need
- ✓ November 17, 2016: CD-1 and CD-3a achieved
  - Approved alternative selection and cost range
  - Shortened timeline includes work to support availability of an exascale advanced architecture system in 2021
- ❑ September 2019: CD-2 and CD-3 achieved (planned)
  - Approve project performance baseline (final KPPs defined and approved)
  - Award remainder of R&D procurements
- ❑ September 2023: CD-4 achieved (planned project closeout)
  - Approve project completion
  - 2023 includes 12 months of schedule contingency

This is a very exciting time for computing in the US

- Unique opportunity to do something special for the nation on a rapid time scale
- The advanced architecture system in 2021 affords the opportunity for
  - More rapid advancement and scaling of mission and science applications
  - More rapid advancement and scaling of an exascale software stack
  - Rapid investments in vendor technologies and software needed for 2021 and 2023 systems
  - More rapid progress in numerical methods and algorithms for advanced architectures
  - Strong leveraging of and broader engagement with US computing capability
- When ECP ends, we will have
  - Run meaningful applications at exascale in 2021, producing useful results
  - Prepared a full suite of mission and science applications for 2023 capable exascale systems
  - Demonstrated integrated software stack components at exascale
  - Invested in the engineering and development, and participated in acquisition and testing of 2023 capable exascale systems
  - Prepared industry and critical applications for a more diverse and sophisticated set of computing technologies, carrying US supercomputing well into the future

ECP EXASCALE COMPUTING PROJECT

Thank you!

www.ExascaleProject.org

EXASCALE COMPUTING PROJECT