The U.S. D.O.E. Exascale Computing Project – Goals and Challenges

Paul Messina, ECP Director

Big Simulation and Big Data Workshop

Indiana University

January 9, 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 COMPUTING PROJECT

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)

EXASCALE COMPUTING PROJECT

Exascale Computing Project Goals

Foster application development	Ease of use	Rich exascale ecosystem	US HPC leadership
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

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

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

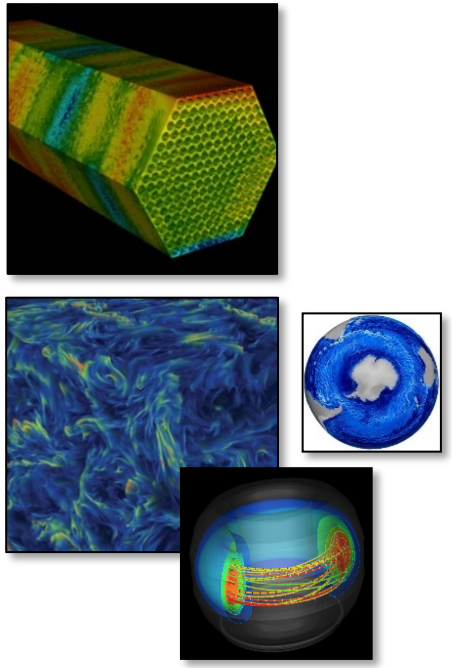

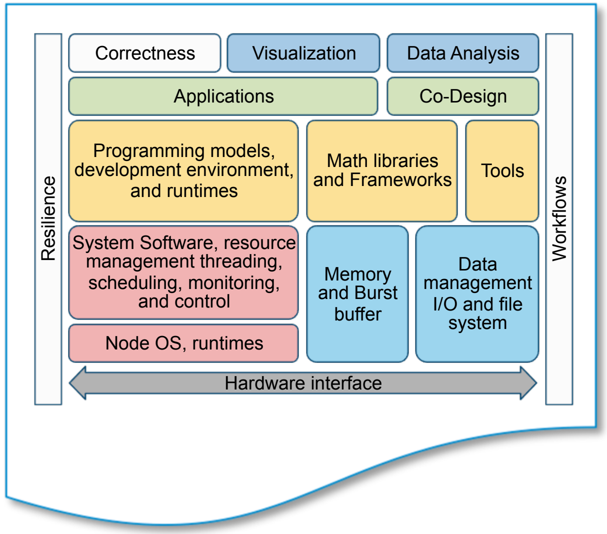

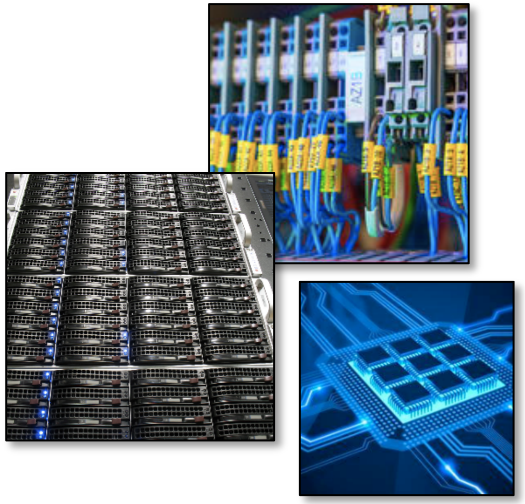

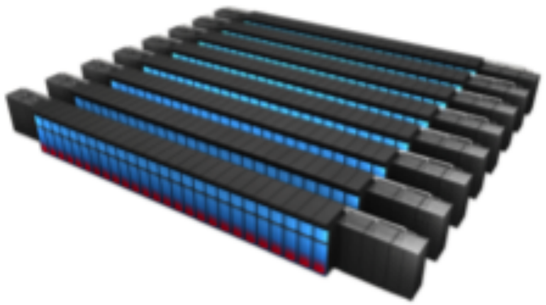

ECP’s work encompasses applications, system software, hardware technologies and architectures, and workforce development

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)
- 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*?

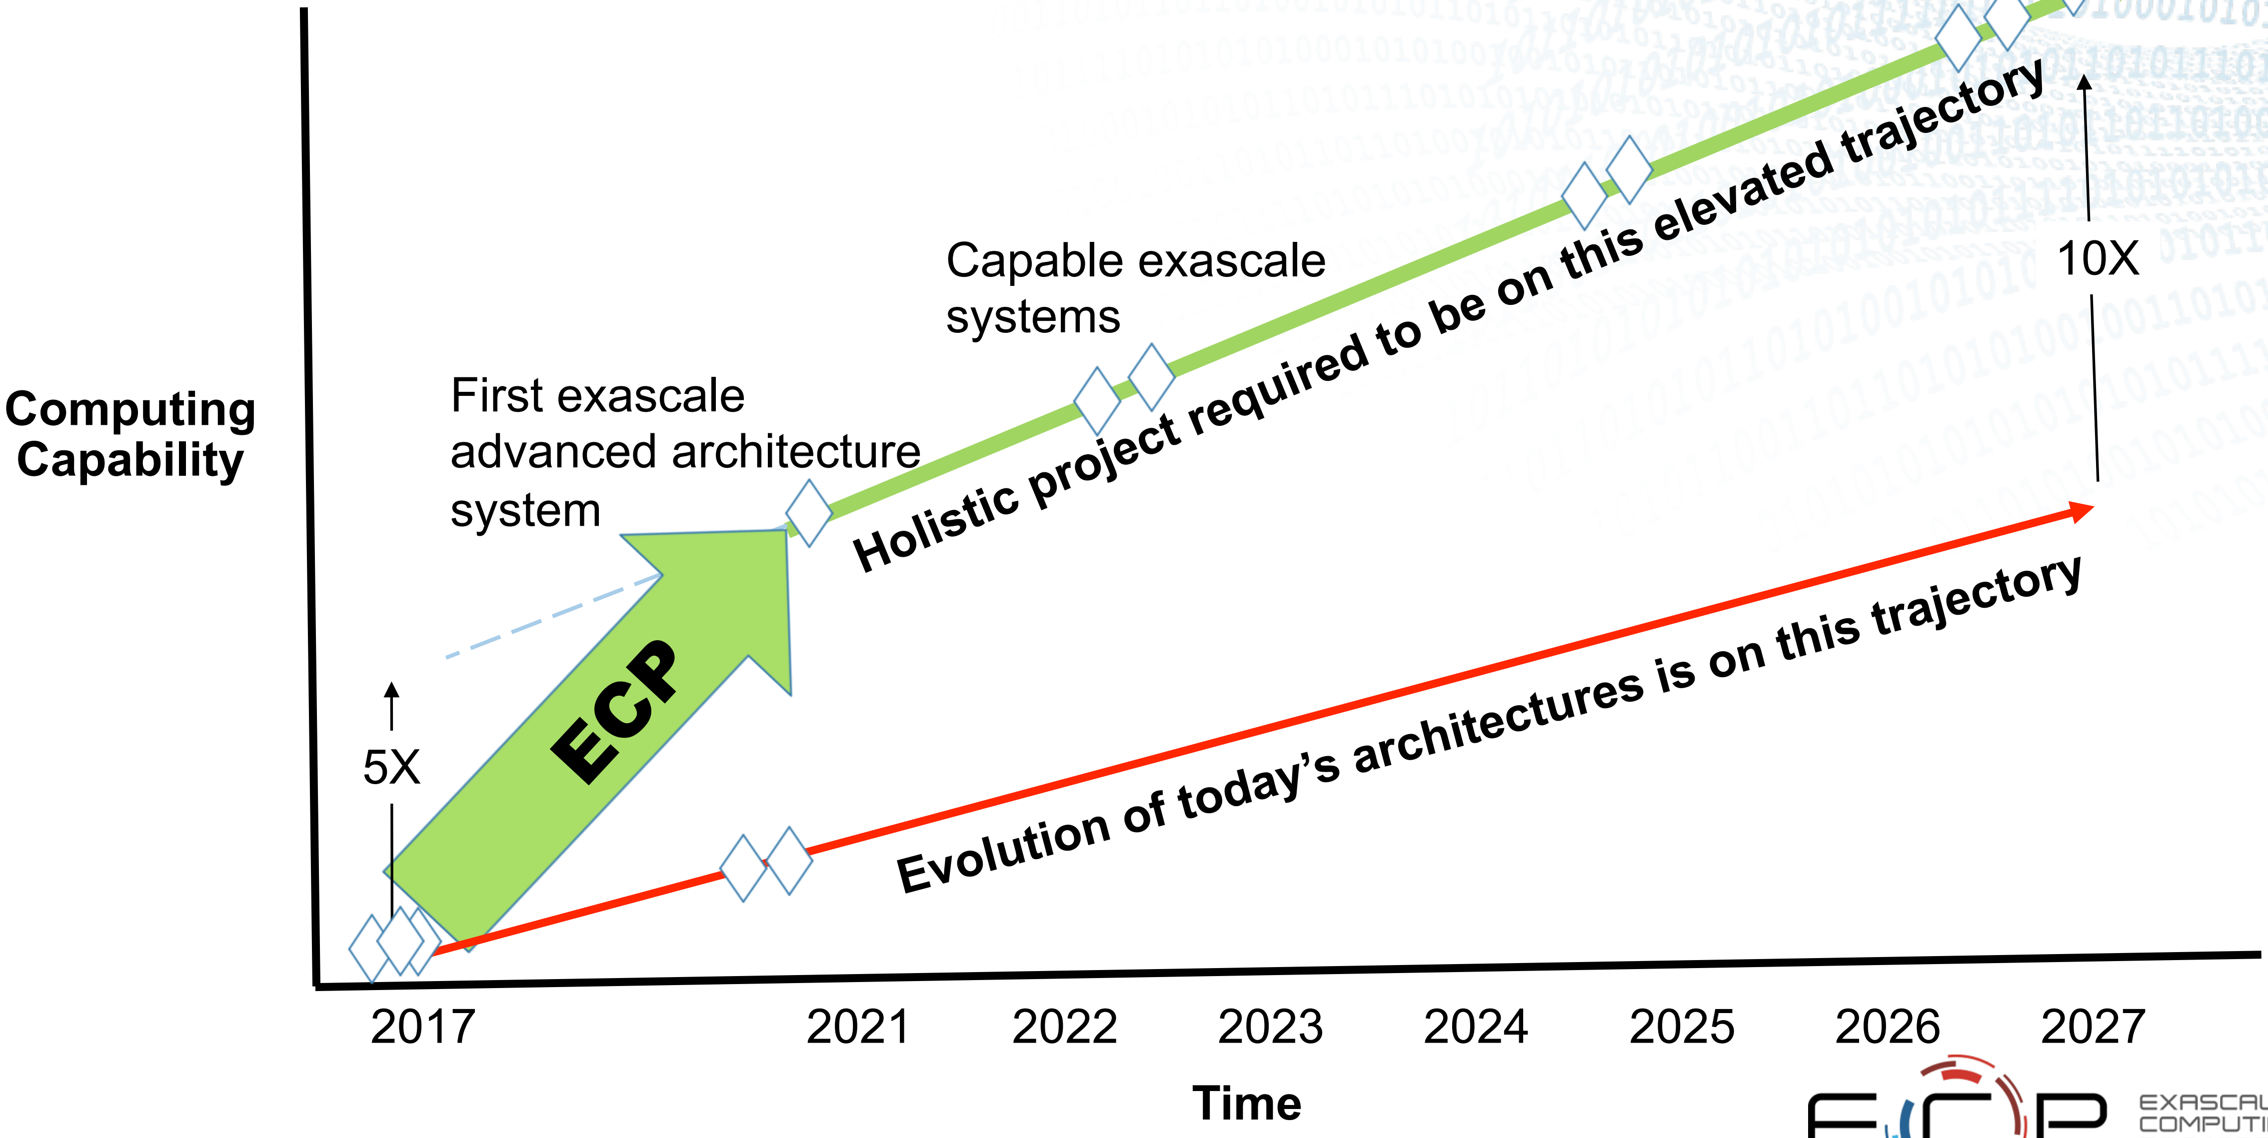

Reaching the Elevated Trajectory will require Advanced and Innovative Architectures

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.

EXASCALE COMPUTING PROJECT

High-level ECP technical project schedule

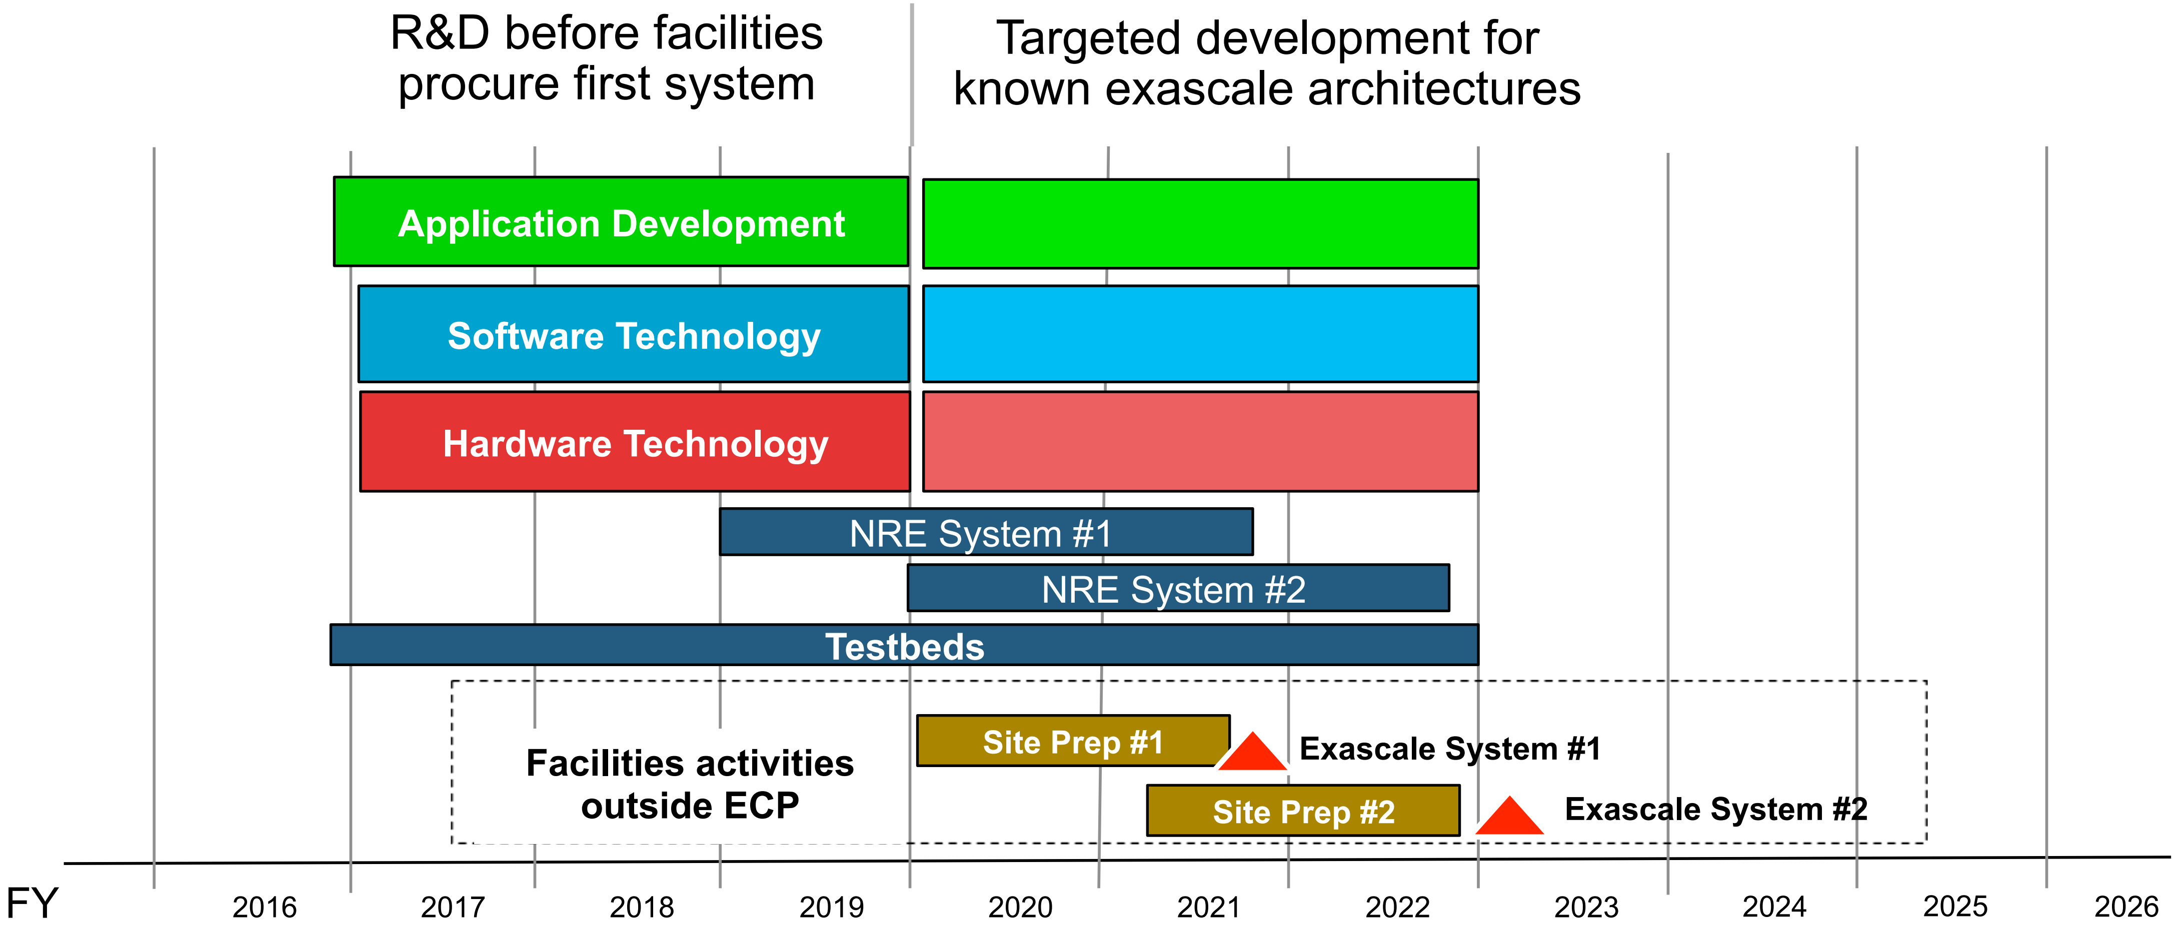

EXASCALE COMPUTING PROJECT

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

EXASCALE COMPUTING PROJECT

Science and Industry Councils

- The ECP is in the process of establishing two advisory bodies:

  An Industry Council composed of ~20 representatives from end-user industries and software vendors

  A Science Council composed of computer scientists, applied mathematicians, and computational scientists

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

NEWS RELEASE

EXASCALE COMPUTING PROJECT

The Exascale Computing Project Awards $34 Million for Software Development

OAK RIDGE, Tenn., Nov. 10, 2016 – The Department of Energy's Exascale Cor
(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 cor
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

NEWS 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 softwar
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 solut
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 Applications Deliver Broad Coverage of Strategic Pillars

Initial selections consist of 15 application projects + 7 seed efforts

National Security

- Stockpile Stewardship

Energy Security

- 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)

Economic Security

- Additive Manufacturing of Qualifiable Metal Parts
- Urban Planning (S)
- Reliable and Efficient Planning of the Power Grid (S)
- Seismic Hazard Risk Assessment (S)

Scientific Discovery

- 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)

Climate and Environmental Science

- Accurate Regional Impact Assessment of Climate Change
- Stress-Resistant Crop Analysis and Catalytic Conversion of Biomass-Derived Alcohols
- Metagenomics for Analysis of Biogeochemical Cycles, Climate Change, Environ Remediation (S)

Healthcare

- Accelerate and Translate Cancer Research

EXASCALE COMPUTING PROJECT

Application Motifs*

Algorithmic methods that capture a common pattern of computation and communication

1. **Dense Linear Algebra**
   - Dense matrices or vectors (e.g., BLAS Level 1/2/3)
2. **Sparse Linear Algebra**
   - Many zeros, usually stored in compressed matrices to access nonzero values (e.g., Krylov solvers)
3. **Spectral Methods**
   - Frequency domain, combining multiply-add with specific patterns of data permutation with all-to-all for some stages (e.g., 3D FFT)
4. **N-Body Methods (Particles)**
   - Interaction between many discrete points, with variations being particle-particle or hierarchical particle methods (e.g., PIC, SPH, PME)
5. **Structured Grids**
   - Regular grid with points on a grid conceptually updated together with high spatial locality (e.g., FDM-based PDE solvers)
6. **Unstructured Grids**
   - Irregular grid with data locations determined by app and connectivity to neighboring points provided (e.g., FEM-based PDE solvers)
7. **Monte Carlo**
   - Calculations depend upon statistical results of repeated random trials
8. **Combinational Logic**
   - Simple operations on large amounts of data, often exploiting bit-level parallelism (e.g., Cyclic Redundancy Codes or RSA encryption)
9. **Graph Traversal**
   - Traversing objects and examining their characteristics, e.g., for searches, often with indirect table lookups and little computation
10. **Graphical Models**
    - Graphs representing random variables as nodes and dependencies as edges (e.g., Bayesian networks, Hidden Markov Models)
11. **Finite State Machines**
    - Interconnected set of states (e.g., for parsing); often decomposed into multiple simultaneously active state machines that can act in parallel
12. **Dynamic Programming**
    - Computes solutions by solving simpler overlapping subproblems, e.g., for optimization solutions derived from optimal subproblem results
13. **Backtrack and Branch-and-Bound**
    - Solving search and global optimization problems for intractably large spaces where regions of the search space with no interesting solutions are ruled out. Use the divide and conquer principle: subdivide the search space into smaller subregions ("branching"), and bounds are found on solutions contained in each subregion under consideration

The Landscape of Parallel Computing Research: A View from Berkeley, Technical Report No. UCB/EECS-2006-183 (Dec 2006).

Survey of Application Motifs

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
Cosmology													
Subsurface													
Materials (QMC)													
Additive Manufacturing													
Chemistry for Catalysts & Plants													
Climate Science													
Precision Medicine Machine Learning													
QCD for Standard Model Validation													
Accelerator Physics													
Nuclear Binding and Heavy Elements													
MD for Materials Discovery & Design													
Magnetically Confined Fusion													

Survey of Application Motifs

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			■					■	■	■		■	
Catalyst Design	■	■	■	■	■								
Wind Plant Flow Physics			■	■		■							
SMR Core Physics	■	■	■	■	■	■	■						
Next-Gen Engine Design		■	■		■	■							
Urban Systems			■		■	■	■						
Seismic Hazard Assessment			■		■	■	■			■			
Systems Biology	■		■	■				■	■	■	■	■	
Biological Neutron Science		■	■	■	■		■		■		■		
Power Grid Dynamics	■		■	■					■	■	■		

EXASCALE COMPUTING PROJECT

Survey of Application Motifs

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
Stellar Explosions													
Excited State Material Properties													
Light Sources													
Materials for Energy Conversion/Storage													
Hypersonic Vehicle Design													
Multiphase Energy Conversion Devices													

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 Application Development and Software Technology 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

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

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

Conceptual ECP Software Stack

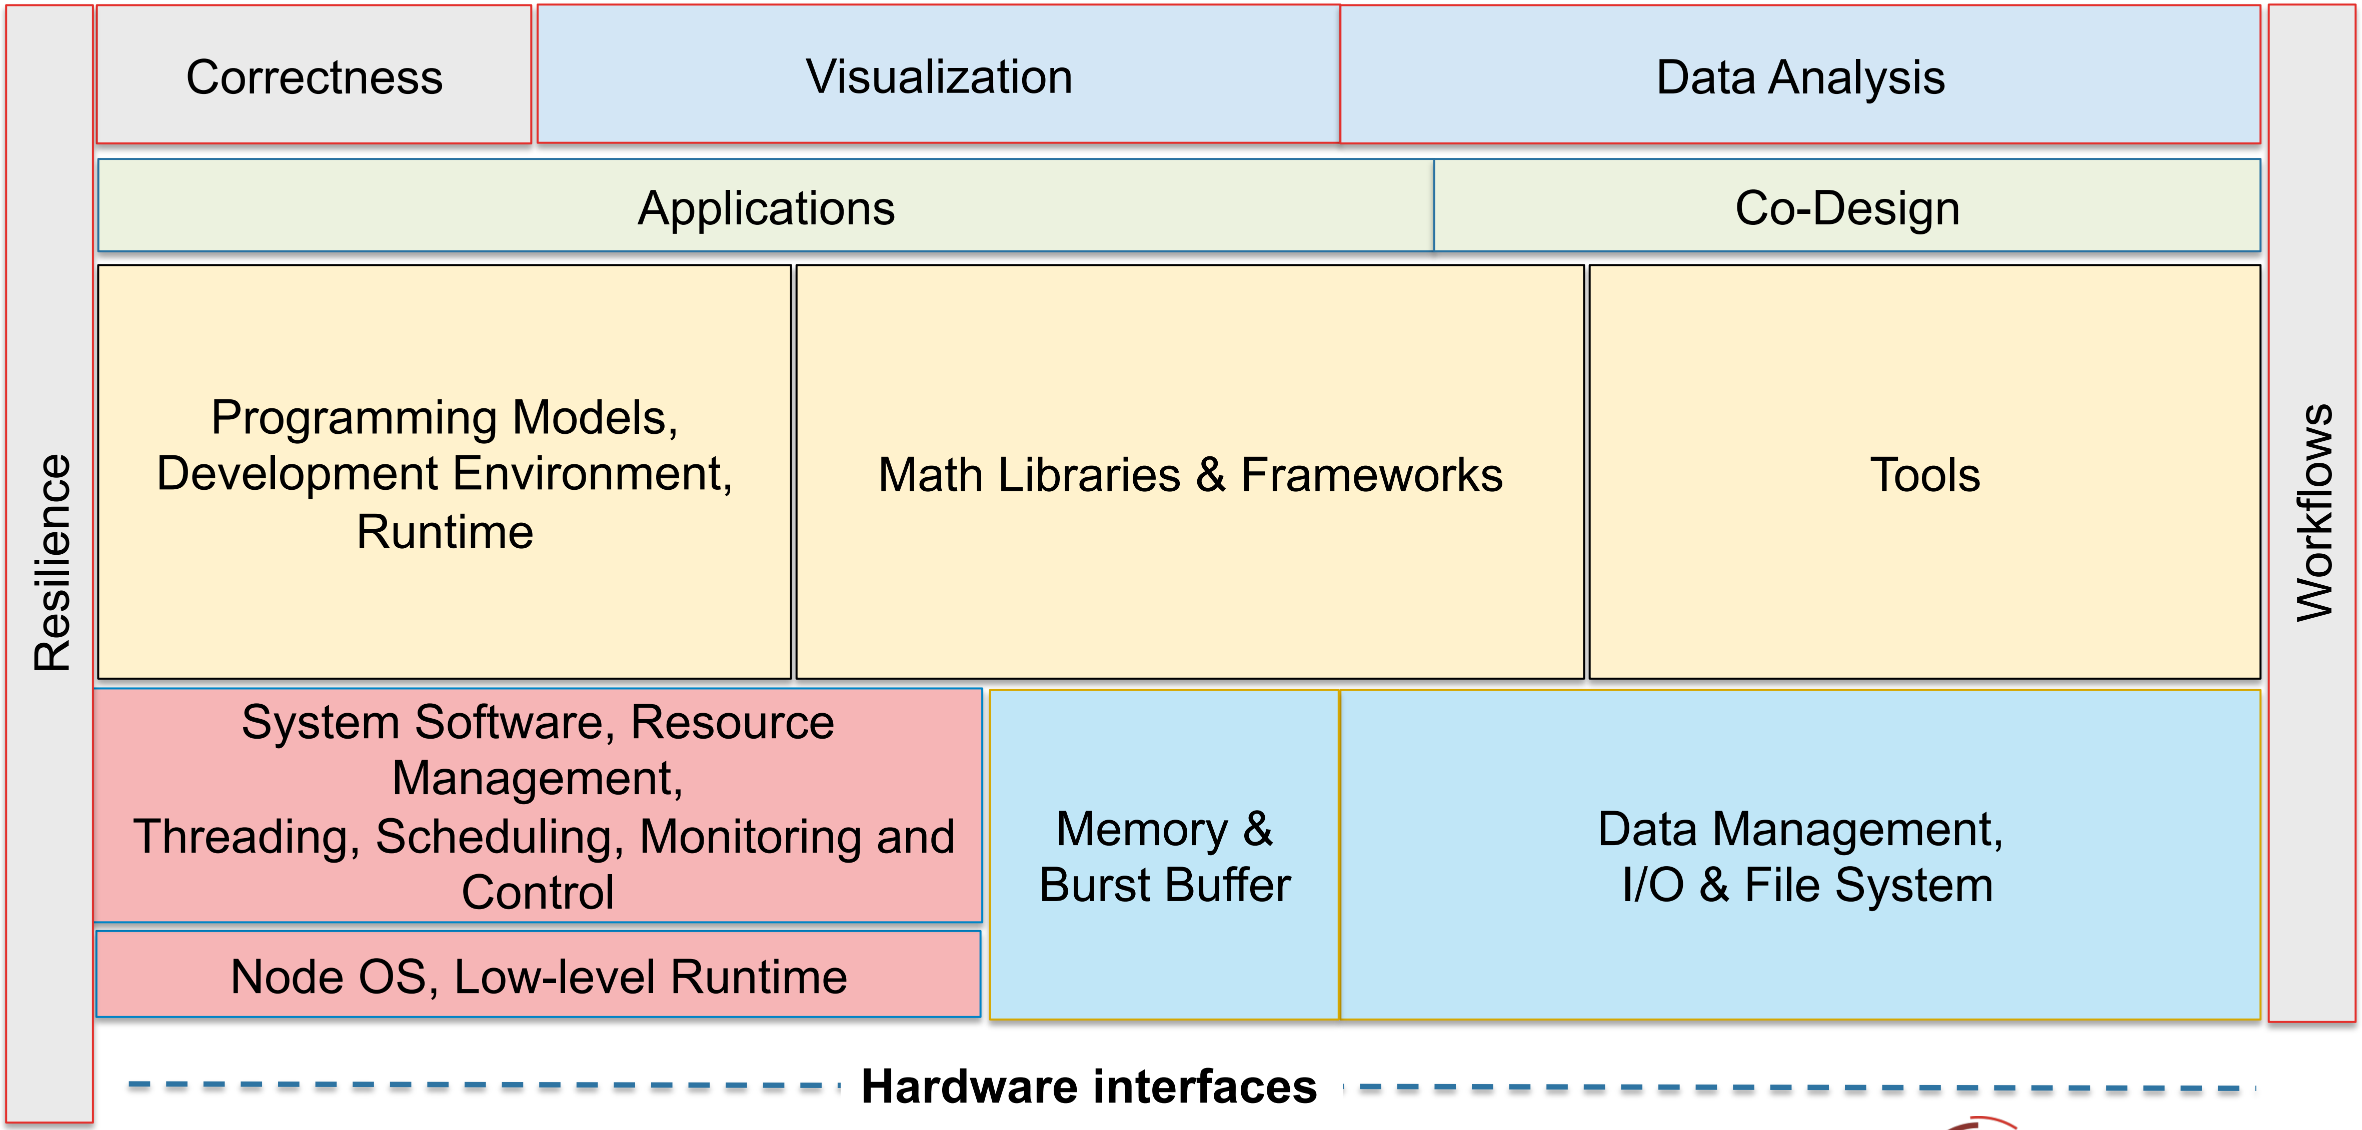

Requirements for Software Technology

Derived from

- Analysis of the software needs of exascale applications
- Inventory of software environments at major DOE HPC facilities (ALCF, OLCF, NERSC, LLNL, LANL, SNL)
  - For current systems and the next acquisition in 2–3 years
- Expected software environment for an exascale system
- Requirements beyond the software environment provided by vendors of HPC systems

Software Technology Requirements
Nuclear Reactors

- **Programming Models and Runtimes**
  1. C++/C++-17, C, Fortran, MPI, OpenMP, Thrust, CUDA, Python
  2. Kokkos, OpenACC, NVL-C
  3. Raja, Legion/Regent, HPX
- **Tools**
  1. LLVM/Clang, PAPI, Cmake, git, CDash, gitlab, Oxbow
  2. Docker, Aspen
  3. TAU
- **Mathematical Libraries, Scientific Libraries, Frameworks**
  1. BLAS/PBLAS, Trilinos, LAPACK
  2. Metis/ParMETIS, SuperLU, PETSc
  3. Hypre

Requirements Ranking
1. Definitely plan to use
2. Will explore as an option
3. Might be useful but no concrete plans

Software Technology Requirements
Nuclear Reactors

- **Data Management and Workflows**
  1. MPI-IO, HDF, Silo, DTK
  2. ADIOS
- **Data Analytics and Visualization**
  1. VisIt
  2. Paraview
- **System Software**

Requirements Ranking
1. Definitely plan to use
2. Will explore as an option
3. Might be useful but no concrete plans

Software Technologies
Aggregate of technologies cited in candidate ECP Applications

- **Programming Models and Runtimes**
  - Fortran, C++/C++17, Python, C, Javascript, C#, R, Ruby
  - MPI, OpenMP, OpenACC, CUDA, Global Arrays, TiledArrays, Argobots, HPX, OpenCL, Charm++
  - UPC/UPC++, Co-Array FORTRAN, CHAPEL, Julia, GDDI, DASK-Parallel, PYBIND11
  - PGAS, GASNetEX, Kokkos, Raja, Legion/Regent, OpenShmem, Thrust
  - PARSEC, Panda, Sycl, Perilla, Globus Online, ZeroMQ, ParSEC, TASCEL, Boost
- **Tools** (debuggers, profilers, software development, compilers)
  - LLVM/Clang,HPCToolkit, PAPI, ROSE, Oxbow (performance analysis), JIRA (software development tool), Travis (testing),
  - ASPEN (machine modeling), CMake, git, TAU, Caliper, , GitLab, CDash (testing), Flux, Spack, Docker, Shifter, ESGF, Gerrit
  - GDB, Valgrind, GitHub, Jenkins (testing), DDT (debugger)
- **Mathematical Libraries, Scientific Libraries, Frameworks**
  - BLAS/PBLAS, MOAB, Trilios, PETSc, BoxLib, LAPACK/ScaLAPACK, Hypre, Chombo, SAMRAI, Metis/ParMETIS, SLEPc
  - SuperLU, Repast HPC (agent-based model toolkit), APOSMM (optimization solver), HPGMG (multigrid), FFTW, Dakota, Zero-RK
  - cuDNN, DAAL, P3DFFT, QUDA (QCD on GPUs), QPhiX (QCD on Phi), ArPack (Arnoldi), ADLB, DMEM, MKL, Sundials, Muelu
  - DPLASMA, MAGMA,PEBBL, pbdR, FMM, DASHMM, Chaco (partitioning), libint (gaussian integrals)
  - Smith-Waterman, NumPy, libcchem

Software Technologies
Cited in Candidate ECP Applications

- **Data Management and Workflows**
  - Swift, MPI-IO, HDF, ADIOS, XTC (extended tag container), Decaf, PDACS, GridPro (meshing), Fireworks, NEDB, BlitzDB, CouchDB
  - Bellerophon, Sidre, Silo, ZFP, ASCTK, SCR, Sierra, DHARMA, DTK, PIO, Akuna, GridOPTICS software system (GOSS), DisPy, Luigi
  - CityGML, SIGMA (meshing), OpenStudio, Landscan USA
  - IMG/KBase, SRA, Globus, Python-PANDAS
- **Data Analytics and Visualization**
  - VisIt, VTK, Paraview, netCDF, CESIUM, Pymatgen, MacMolPlt, Yt
  - CombBLAS, Elviz, GAGE, MetaQuast
- **System Software**

EXASCALE COMPUTING PROJECT

Software Technology Projects 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

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

EXASCALE COMPUTING PROJECT

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

ECP's plan to accelerate and enhance system capabilities

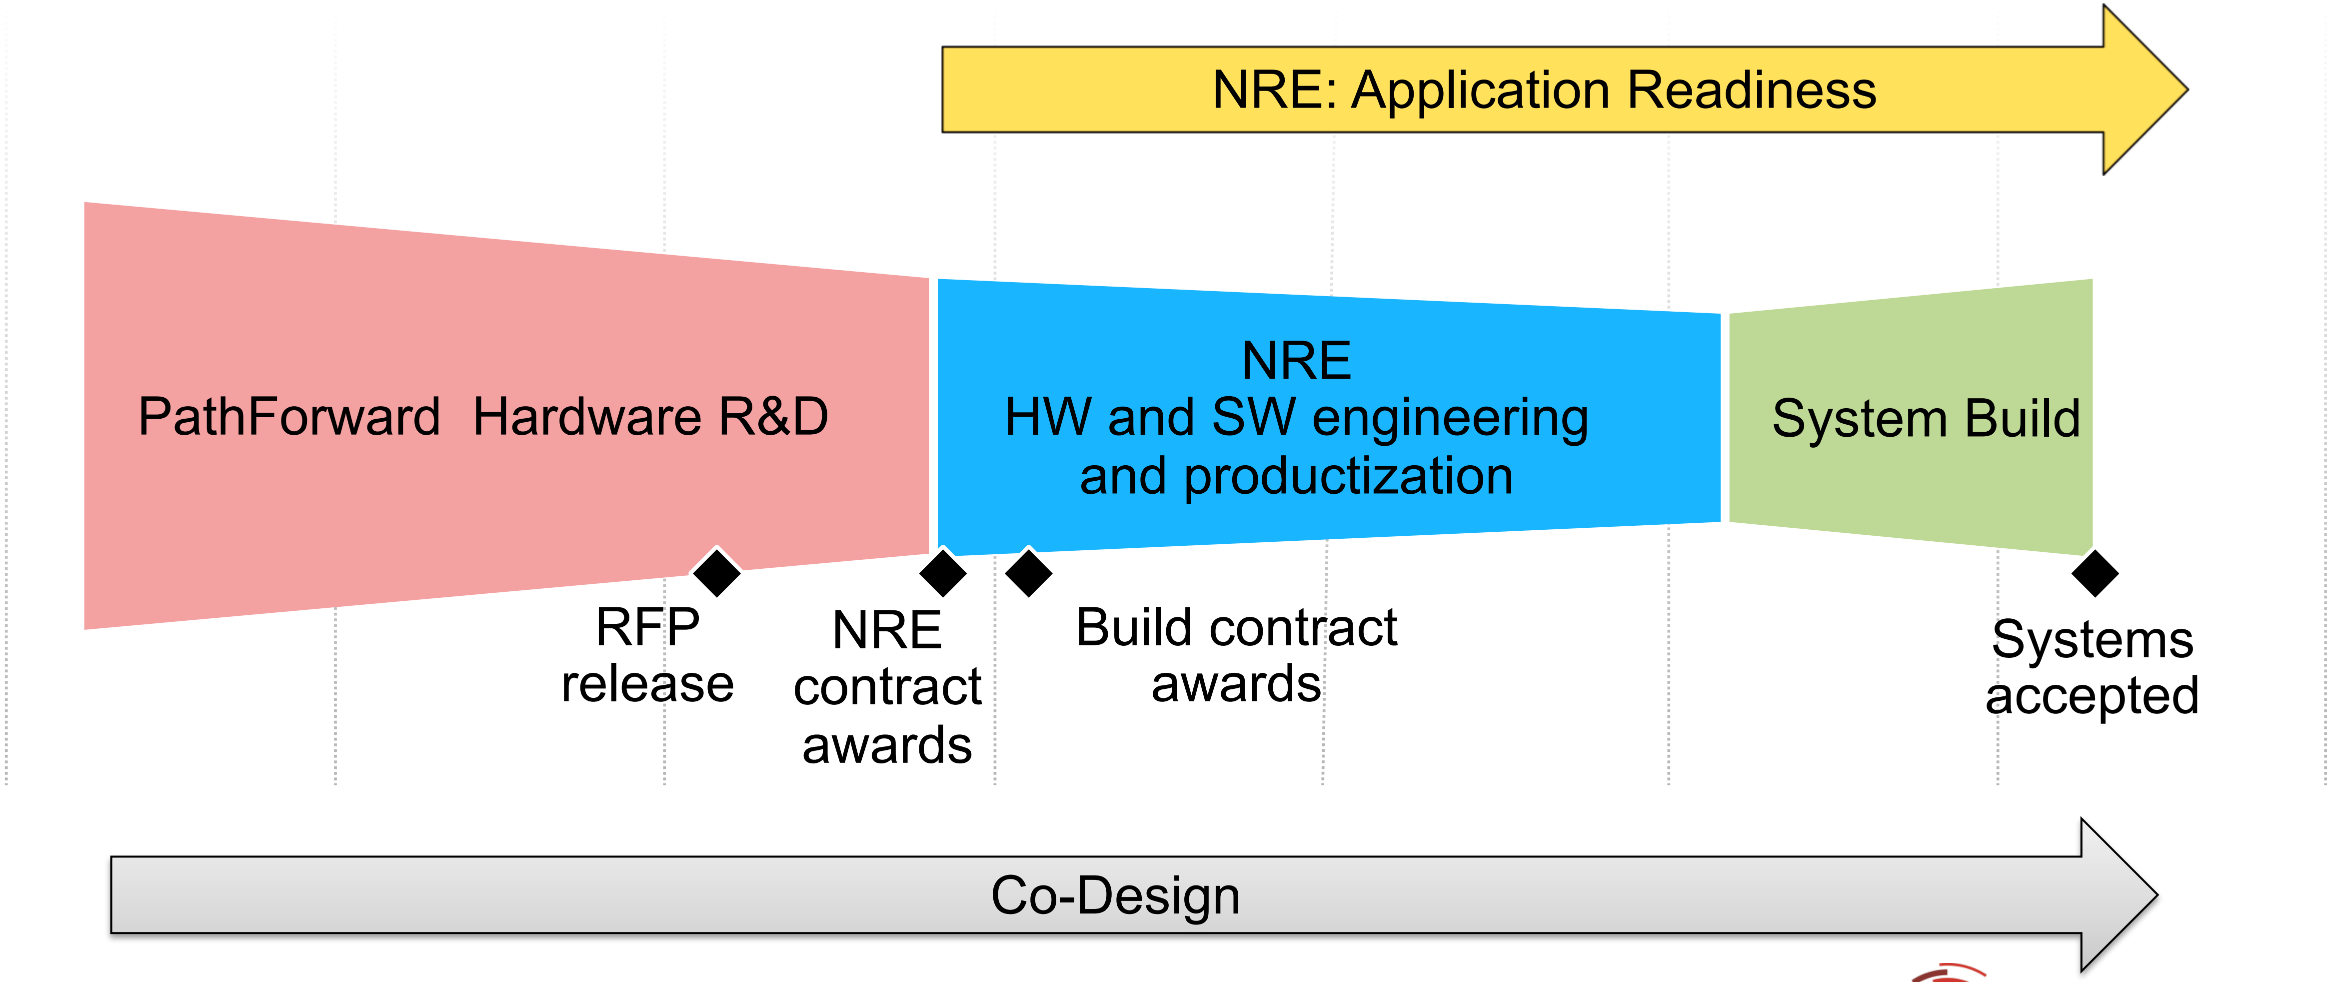

NSCI Objectives

Executive departments, agencies, and offices participating in the NSCI shall pursue five strategic objectives:

1) Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflops systems across a range of applications representing government needs.

2) Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.

3) Establishing over the next 15 years, a viable path forward for future HPC systems in the Post-Moore's-Law Era to advance beyond traditional lithographic scaling of devices.

4) Increasing the capacity and capability of an enduring national HPC ecosystem, employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, and workforce development.

5) Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the U.S. commercial, government, and academic sectors.

EXASCALE COMPUTING PROJECT

What the ECP is not addressing, fully or at all

- Only partially tackling convergence of simulation and data analytics
  - Hope to do more
- Post Moore's Law issues – out of scope for ECP

Some Applications Risks and Challenges

- Exploiting on-node memory and compute hierarchies
- Programming models: what to use where and how (e.g., task-based RTS)
- Integrating S/W components that use disparate approaches (e.g., on-node parallelism)
- Developing and integrating co-designed motif-based community components
- Achieving portable performance (without “if-def’ing” 2 different code bases)
- Multi-physics coupling: both algorithms and software
- Integrating sensitivity analysis, data assimilation, and uncertainty quantification technologies
- Understanding requirements of Data Analytic Computing methods and applications
  - Critical infrastructure, superfacility, supply chain, image/signal processing, in situ analytics
  - Machine/statistical learning, classification, streaming/graph analytics, discrete event, combinatorial optimization

EXASCALE COMPUTING PROJECT

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

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 or large-scale data analytics
- Ditto for complex workflows
- Areas for which we deliberately decided to do technology watches before investing in them very much (don’t understand the use cases, the requirements)
  – Resilience
  – Workflows

EXASCALE COMPUTING PROJECT

Questions?

www.ExascaleProject.org

EXASCALE COMPUTING PROJECT

www.ExascaleProject.org