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

Paul Messina, ECP Director

NIST

February 14, 2017

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

www.ExascaleProject.org

Exascale Computing Project (ECP) Mission and Scope

- 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 Execution

- 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)

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

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

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

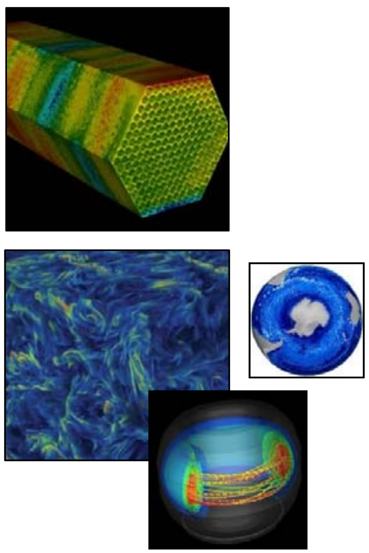

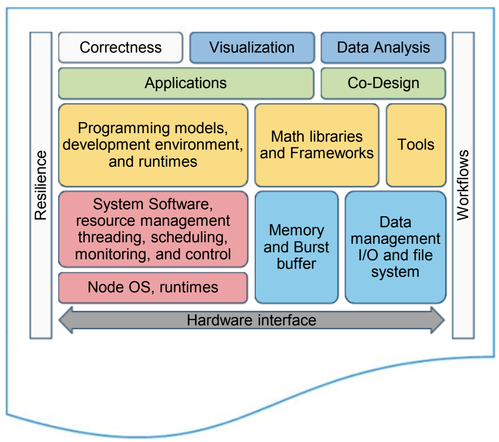

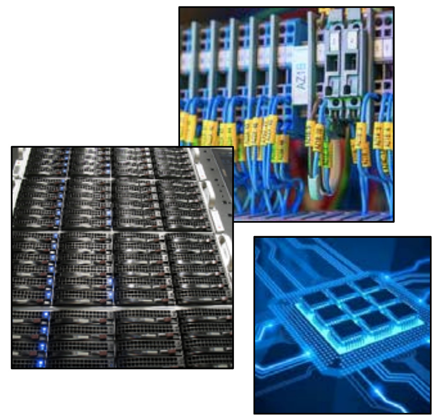

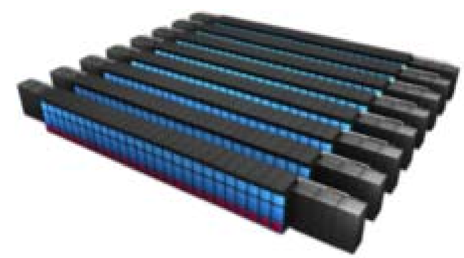

ECP's work encompasses applications, system software, hardware technologies and architectures, and workforce development

What is an exascale *advanced architecture*?

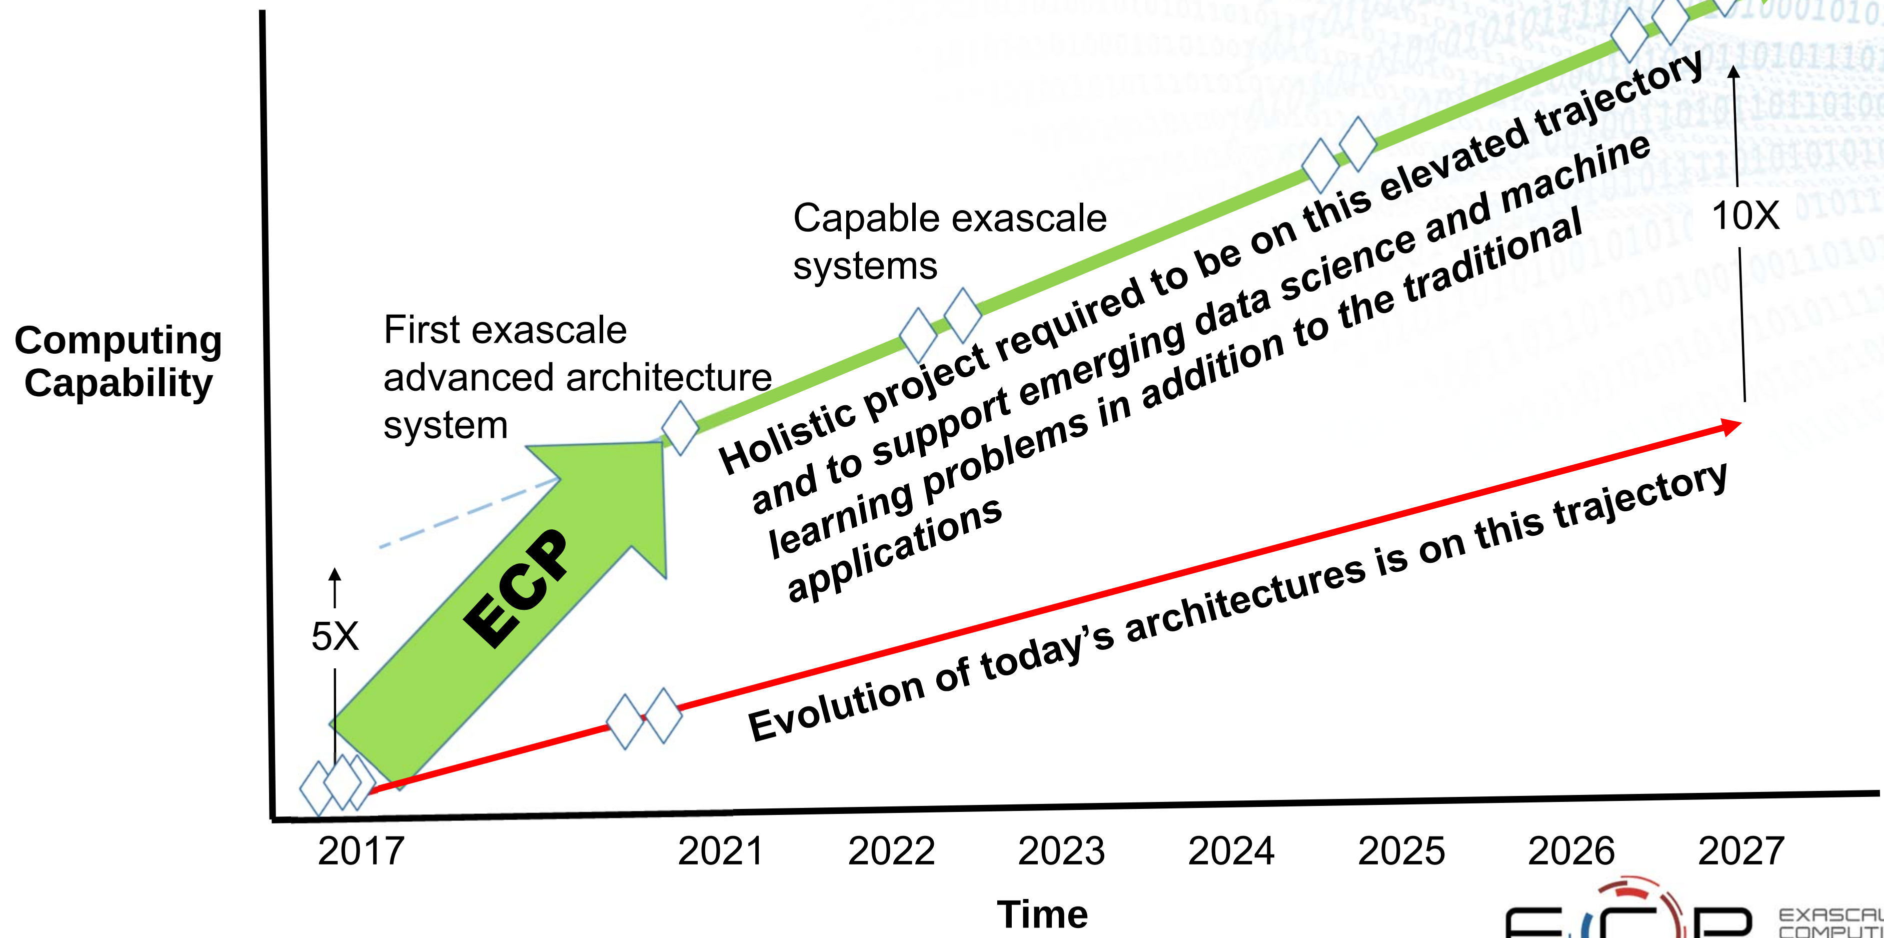

High-level ECP technical project schedule

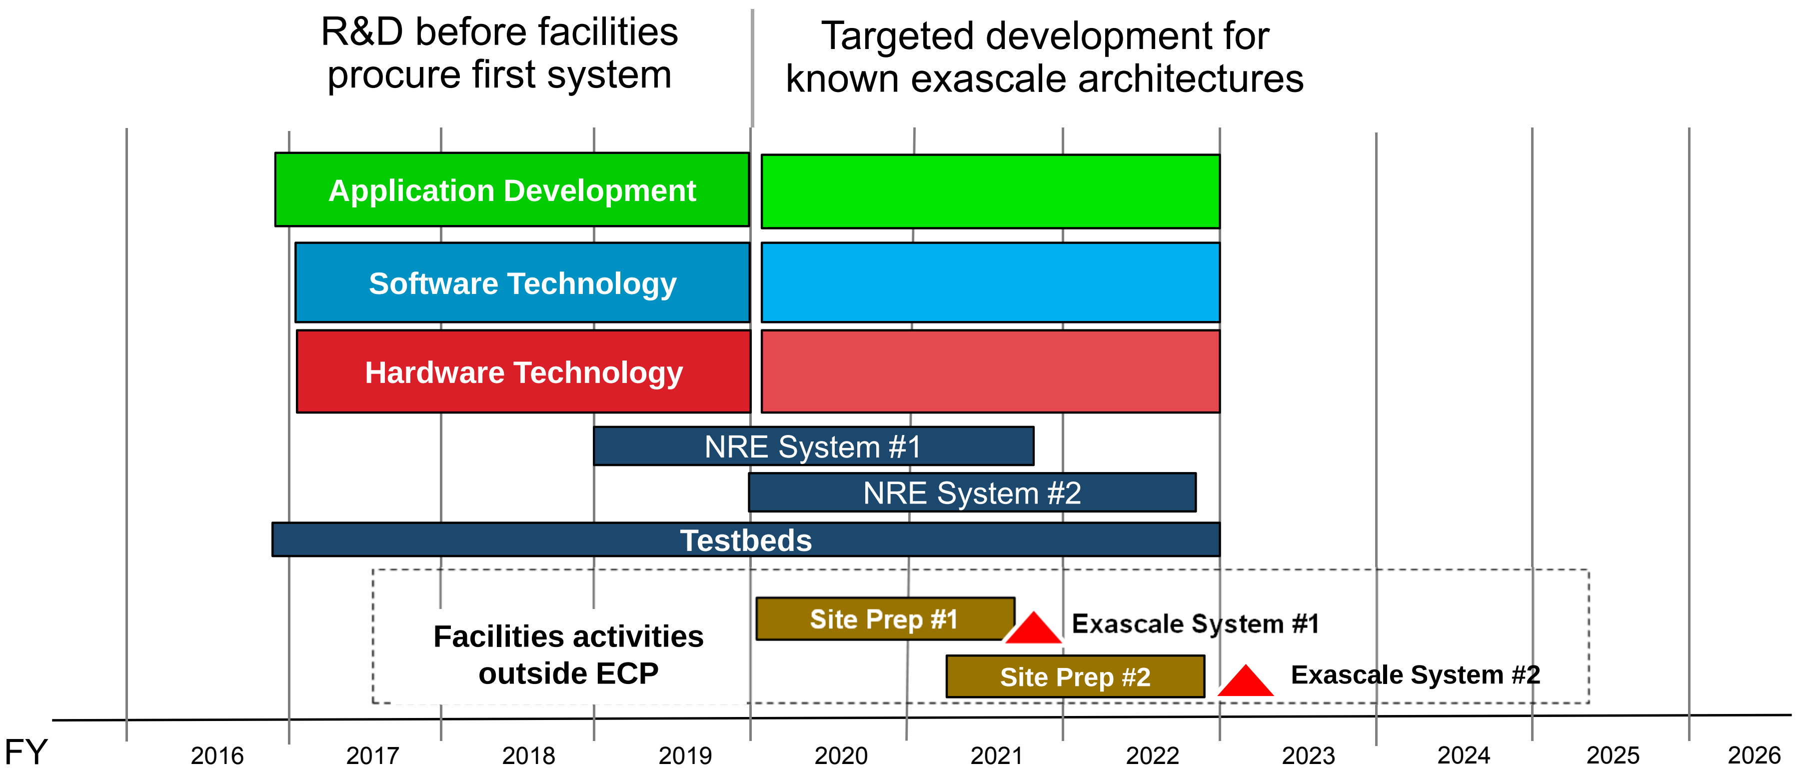

ECP WBS

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
  - PathForward II 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

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

    -- first meeting will be March 6-7

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

ECP Applications Deliver Broad Coverage of Strategic Pillars

Current portfolio consists of 25 application efforts within six pillars

National Security

- Stockpile Stewardship (3)

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)

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 (S)

Health Care

- Accelerate and Translate Cancer Research

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Nuclear Energy (NE)	Climate (BER)	Chemical Science (BES, BER)	Wind Energy (EERE)	Combustion (BES)
Accelerate design and commercialization of next-generation small modular reactors*	Accurate regional impact assessment of climate change*	Biofuel catalysts design; stress-resistant crops	Increase efficiency and reduce cost of turbine wind plants sited in complex terrains*	Design high-efficiency, low-emission combustion engines and gas turbines*
Climate Action Plan; SMR licensing support; GAIN	Climate Action Plan	Climate Action Plan; MGI	Climate Action Plan	2020 greenhouse gas and 2030 carbon emission goals

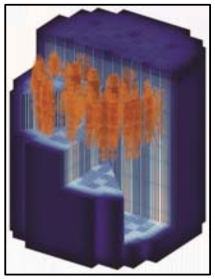

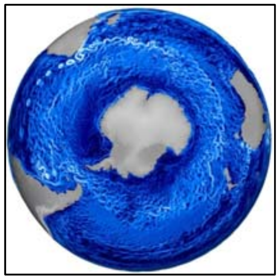

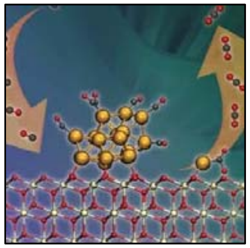

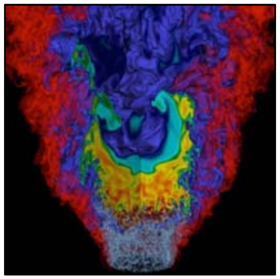

* Scope includes a discernible data science component

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

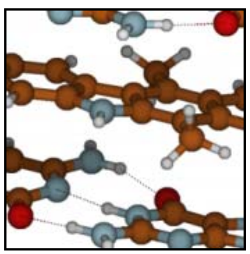

Nuclear Physics (NP)

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

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

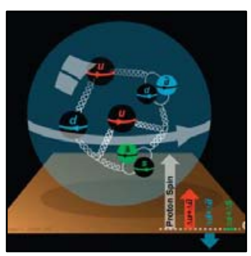

Nuclear Materials (BES, NE, FES)

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

Climate Action Plan; MGI; Light Water Reactor Sustainability; ITER; Stockpile Stewardship Program

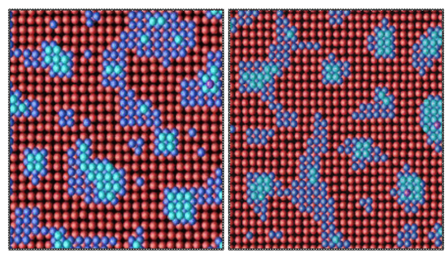

Accelerator Physics (HEP)

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

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

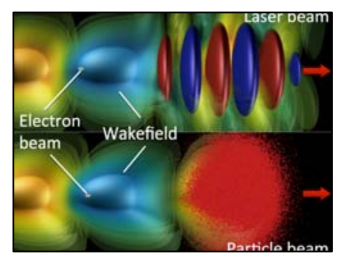

Materials Science (BES)

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

MGI; LCLS-II 2025 Path Forward

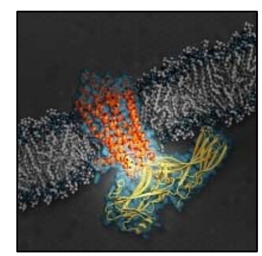

* Scope includes a discernible data science component

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy and Other Agency 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

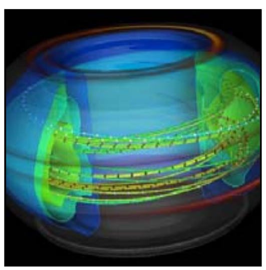

Advanced Manufacturing (EERE)

Additive manufacturing process design for qualifiable metal components*

NNMIs; Clean Energy Manufacturing Initiative

Cosmology (HEP)

Cosmological probe of standard model (SM) of particle physics: Inflation, dark matter, dark energy*

Particle Physics Project Prioritization Panel (P5)

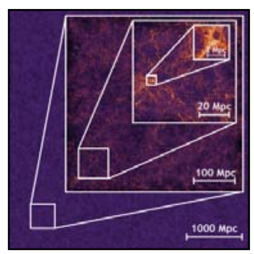

Geoscience (BES, BER, EERE, FE, NE)

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

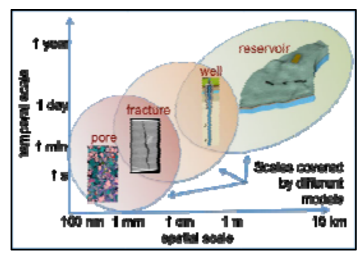

Precision Medicine for Cancer (NIH)

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

Precision Medicine in Oncology; Cancer Moonshot

* Scope includes a discernible data science component

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development seed projects

Seismic (EERE, NE, NNSA)	Carbon Capture and Storage (FE)	Chemical Science (BES)	Urban Systems Science (EERE)
Reliable earthquake hazard and risk assessment in relevant frequency ranges*	Scaling carbon capture/storage laboratory designs of multiphase reactors to industrial size	Design catalysts for conversion of cellulosic-based chemicals into fuels, bioproducts	Retrofit and improve urban districts with new technologies, knowledge, and tools*
DOE Critical Facilities Risk Assessment; urban area risk assessment; treaty verification	Climate Action Plan; SunShot; 2020 greenhouse gas/2030 carbon emission goals	Climate Action Plan; SunShot Initiative; MGI	Energy-Water Nexus; Smart Cities Initiative

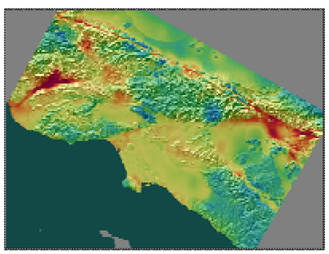

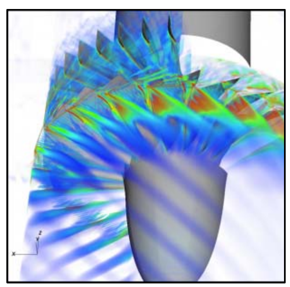

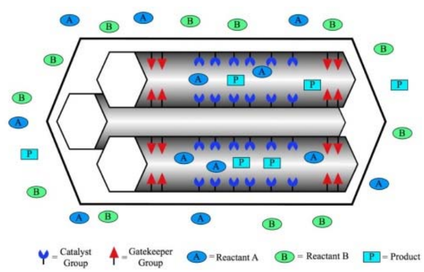

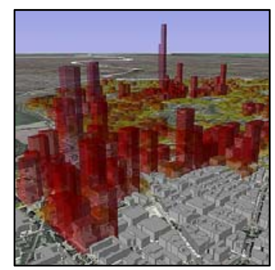

* Scope includes a discernible data science component

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development seed projects

Metagenomics (BER)	Astrophysics (NP)	Power Grid (EERE, OE)
Leveraging microbial diversity in metagenomic datasets for new products and life forms*	Demystify origin of chemical elements (> Fe); confirm LIGO gravitational wave and DUNE neutrino signatures*	Reliably and efficiently planning our nation's grid for societal drivers: rapidly increasing renewable energy penetration, more active consumers*
Climate Action Plan; Human Microbiome Project; Marine Microbiome Initiative	2015 Long Range Plan for Nuclear Science; origin of universe and nuclear matter in universe	Grid Modernization Initiative; Climate Action Plan

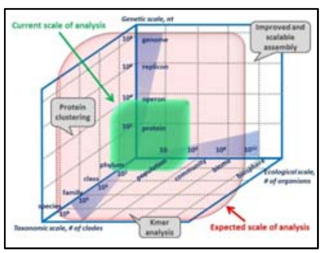

* Scope includes a discernible data science component

ECP EXASCALE COMPUTING PROJECT

Example: NWChemEx: Tackling Chemical, Materials and Biomolecular Challenges in the Exascale Era*

Exascale Challenge Problem

- Aid & accelerate advanced biofuel development by exploring new feedstock for efficient production of biomass for fuels and new catalysts for efficient conversion of biomass derived intermediates into biofuels and bioproducts
- Molecular understanding of how proton transfer controls protein-assisted transport of ions across biomass cellular membranes; often seen as a stress responses in biomass, would lead to more stress-resistant crops thru genetic modifications
- Molecular-level prediction of the chemical processes driving the specific, selective, low-temperature catalytic conversion (e.g., Zeolites such as H-ZSM-5)) of biomass-derived alcohols into fuels and chemicals in constrained environments

Applications & S/W Technologies

Applications

- NWChemEx (evolved from redesigned NWChem)

Software Technologies Cited

- Fortran, C, C++
- Global arrays, TiledArrays, ParSEC, TASCEL
- VisIt, Swift
- TAO, Libint
- Git, svn, JIRA, Travis CI
- Co-Design: CODAR, CE-PSI, GraphEx

Risks and Challenges

- Unknown performance of parallel tools
- Insufficient performance or scalability or large local memory requirements of critical algorithms
- Unavailable tools for hierarchical memory, I/O, and resource management at exascale
- Unknown exascale architectures
- Unknown types of correlation effect for systems with large number of electrons
- Framework cannot support effective development

Development Plan

Y1: Framework with tensor DSL, RTS, APIs, execution state tracking; Operator-level NK-based CCSD with flexible data distributions & symmetry/sparsity exploitation

Y2: Automated compute of CC energies & 1-/2-body CCSD density matrices; HT & DFT compute of >1K atom systems via multi-threading

Y3: Couple embedding with HF & DFT for multilevel memory hierarchies; QMD using HF & DFT for 10K atoms; Scalable R12/F12 for 500 atoms with CCSD energies and gradients using task-based scheduling

Y4: Optimized data distribution & multithreaded implementations for most time-intensive routines in HF, DFT, and CC.

*PI: Thom Dunning (PNNL)

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													

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

ECP Software Technology Overview

- 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
- Accomplished by extending current technologies to exascale where possible, performing R&D required to conceive of new approaches where necessary
  - Coordinate with vendor efforts; i.e., develop software other than what is typically done by vendors, develop common interfaces or services
  - Develop and deploy high-quality and robust software products

Conceptual ECP Software Stack

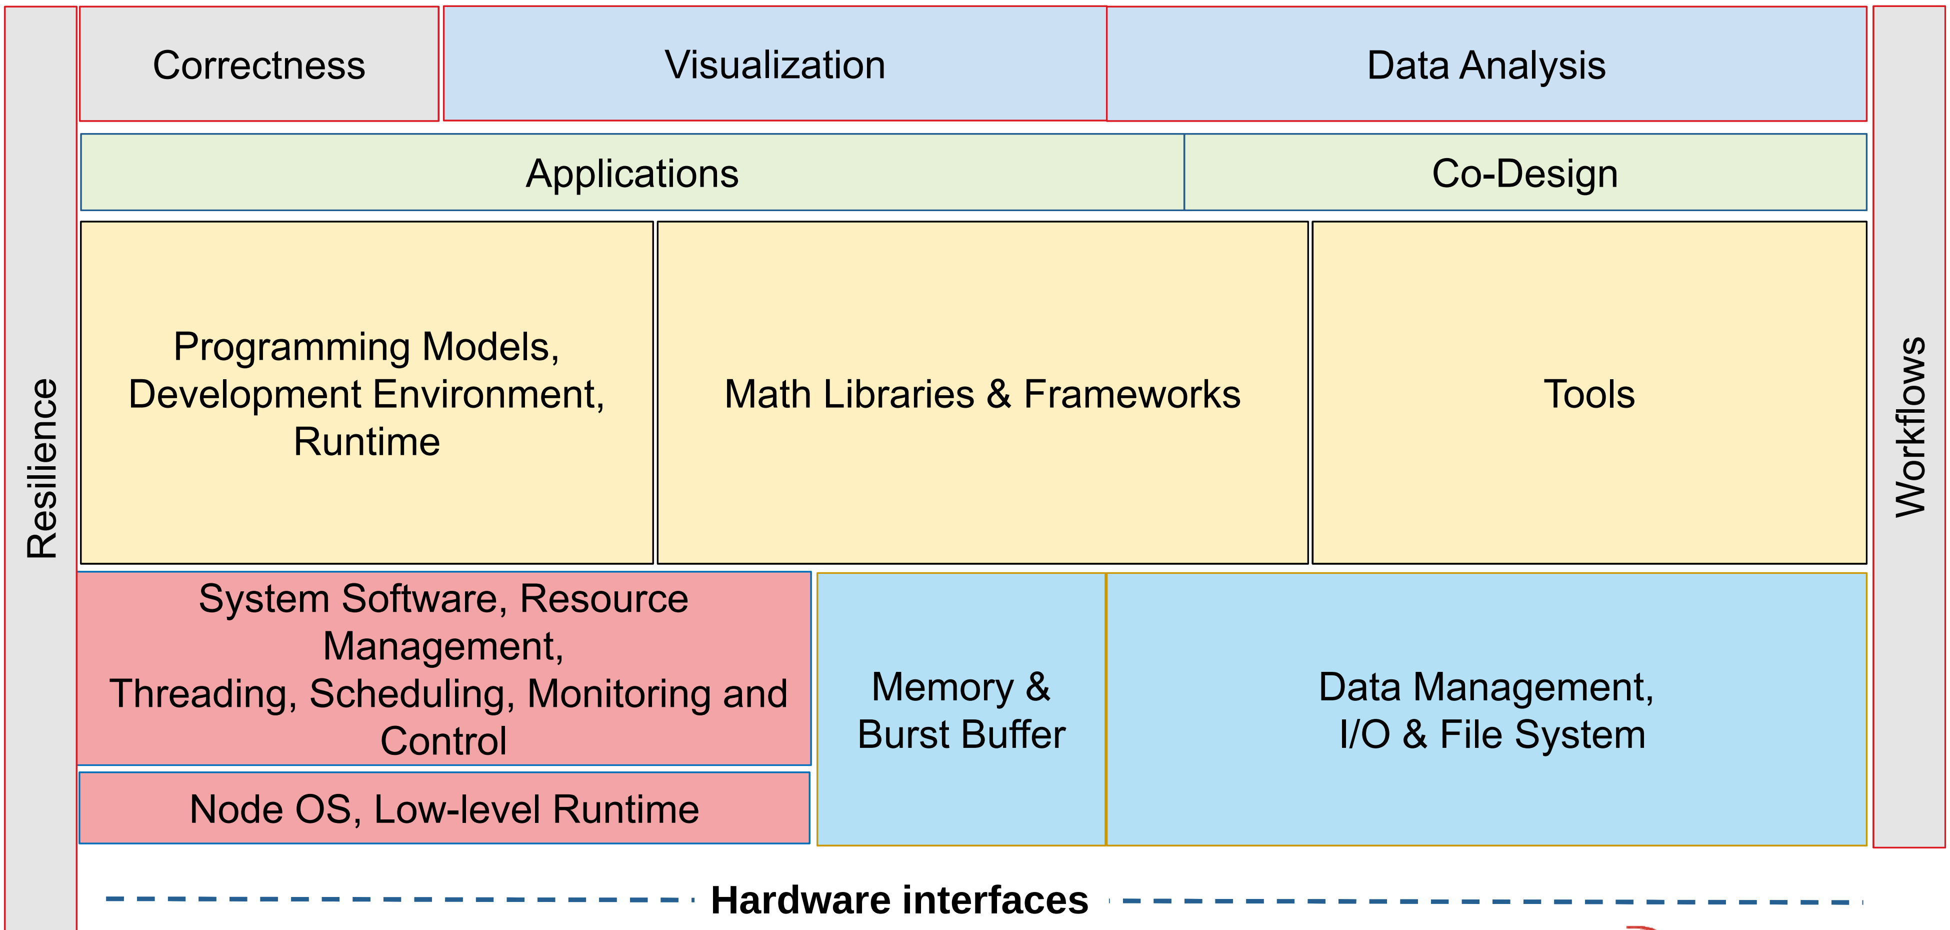

Requirements for Software Technology

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

EXASCALE COMPUTING PROJECT

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**

Software Technology Projects Mapped to Software Stack

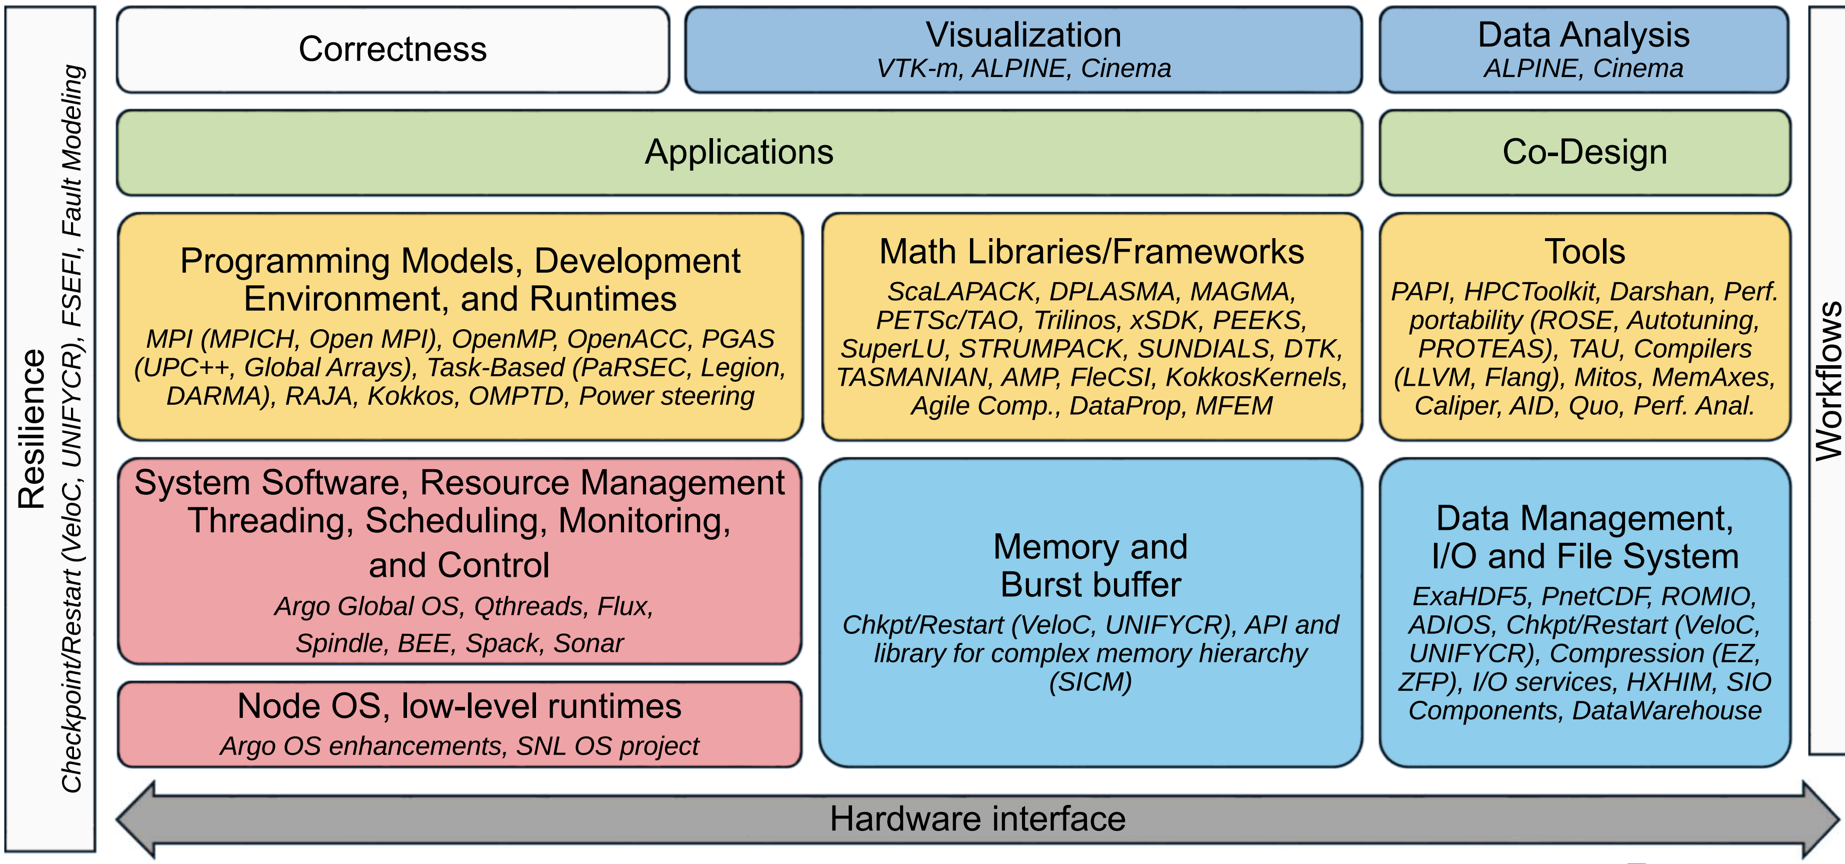

Hardware Technology Overview

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

Issue *PathForward and 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

DOE labs engage to:

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

ECP's plan to accelerate and enhance system capabilities

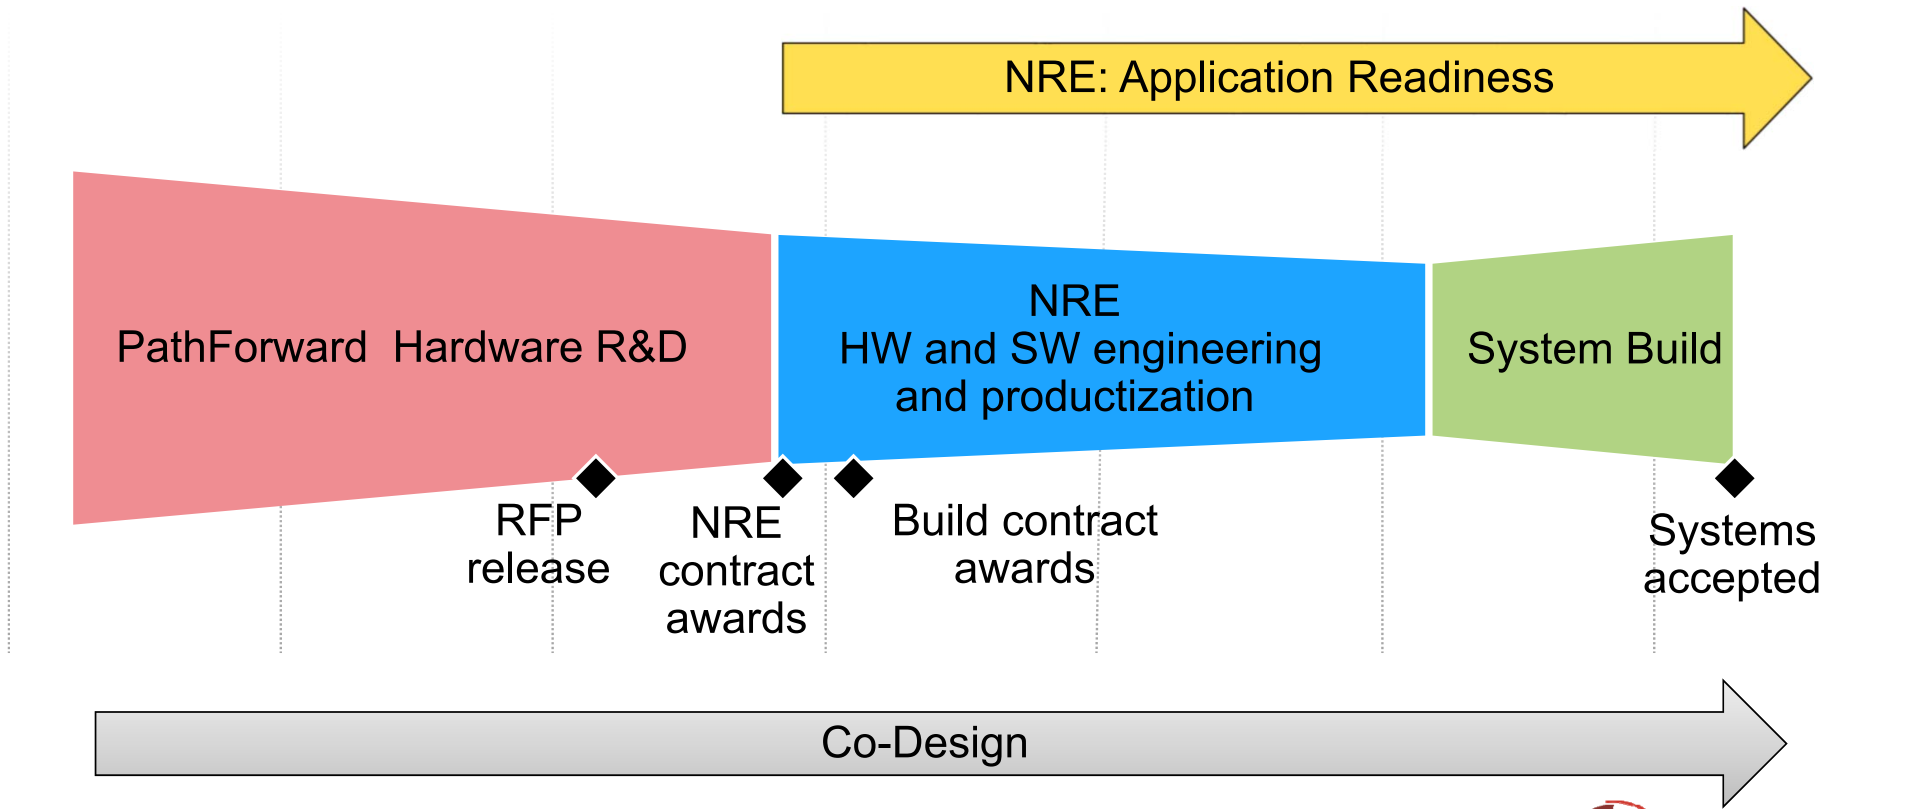

Some Applications Risks and Challenges

- 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

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
- Creating a broadly adopted software stack for HPC
  - *de facto* standards

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

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
- Partner with vendors to develop computer architectures that support exascale applications
- Train a large cadre of computational scientists, engineers, and computer scientists who will be an asset to the nation long after the end of the ECP

Questions?

www.ExascaleProject.org

www.ExascaleProject.org