

The Exascale Computing Project (ECP)

Paul Messina, ECP Director

Stephen Lee, ECP Deputy Director

SC16 Birds of a Feather, "The U.S. Exascale Computing Project"

November 16, 2016

Salt Lake City, Utah



EXASCALE COMPUTING PROJECT



U.S. DEPARTMENT OF
ENERGY

Office of
Science

www.ExascaleProject.org

What is the Exascale Computing Project?

- As part of President Obama's 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 petaflop machines.
- ECP's work encompasses
 - applications,
 - system software,
 - hardware technologies and architectures, and
 - workforce development to meet the scientific and national security mission needs of DOE.

Four key challenges that must be addressed to achieve exascale

- Parallelism
- Memory and Storage
- Reliability
- Energy Consumption

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 (average fault rate: $\leq 1/\text{week}$)
- Includes a software stack that supports a broad spectrum of applications and workloads

This ecosystem will be developed using a co-design approach to deliver new software, applications, platforms, and computational science capabilities at heretofore unseen scale

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

Foster application development

Create software that makes exascale systems usable by a wide variety of scientists and engineers across a range of applications

Ease of use

Enable by 2023 \geq two diverse computing platforms with up to 50 \times more computational capability than today's 20 PF systems, within a similar size, cost, and power footprint


\geq Two diverse architectures

Help ensure continued American leadership in architecture, software and applications to support scientific discovery, energy assurance, stockpile stewardship, and nonproliferation programs and policies

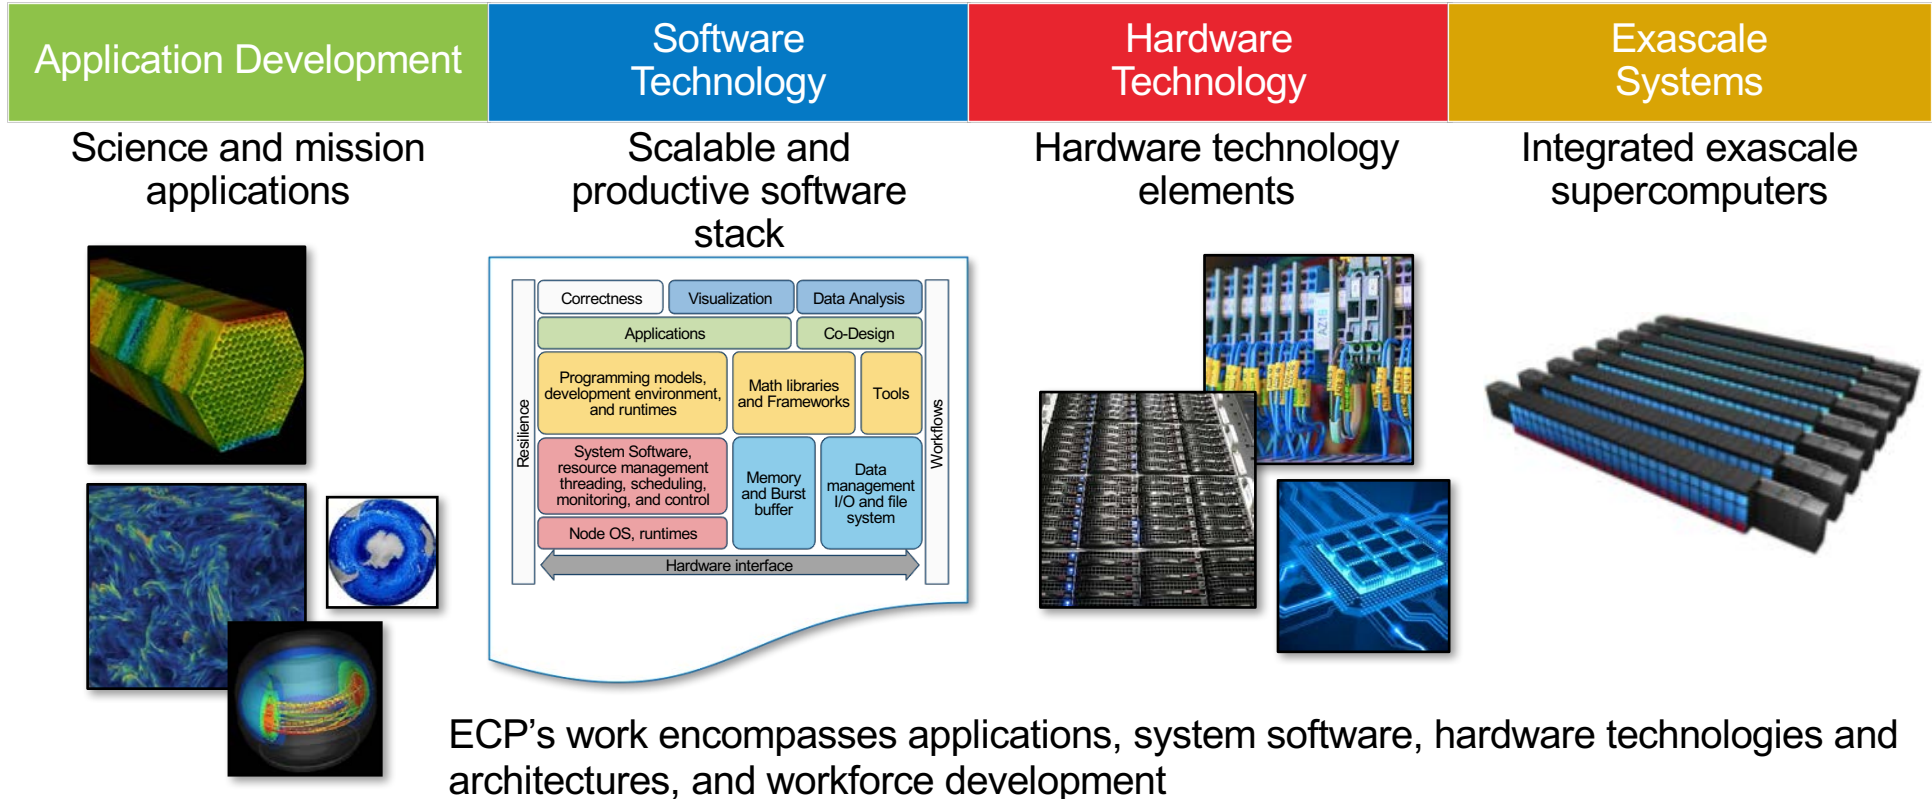
US HPC leadership

ECP leadership team

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



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



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



NEWS RELEASE

The Exascale Computing Project Awards \$34 Million for Software Development

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

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



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 Computing Project (ECP) today announced its first round of funding with the selection of 15 application development proposals for full funding and seven proposals for seed funding, representing teams from 45 research and academic organizations.

The awards, totaling \$39.8 million, target advanced modeling and simulation solutions for specific challenges supporting key DOE missions in science, clean energy and national security, as well as collaborations such as the Precision Medicine Initiative with the National Institutes of Health's National Cancer Institute.



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.

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



ECP Application Development (AD) Focus Area

Douglas B. Kothe, ECP AD Director

Charles H. Still, ECP AD Deputy Director

SC16 Birds of a Feather, "The U.S. Exascale Computing Project"

November 16, 2016

Salt Lake City, UT



EXASCALE COMPUTING PROJECT



U.S. DEPARTMENT OF
ENERGY

Office of
Science

www.ExascaleProject.org

Summary

- Applications are the tool for delivering on Mission Need
 - Vehicle for high-confidence insights and answers to national science, energy, and national security Challenge Problems
 - Necessary for all KPPs; on point for Scalable Science Performance, Application Readiness, and Productive Software Ecosystem metrics
- Mission Need requirements will be met only through broad coverage of DOE programs
 - 10 program offices are targeted
 - Each office has multiple high-priority strategic goals addressable with exascale applications
- Application Co-Design is an essential element of success
- Application challenges can be met with efficient and productive development teams sharing lessons learned and best practices

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
<ul style="list-style-type: none">• 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)	<ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• Materials discovery and design• Climate science• Nuclear energy• Combustion science• Large-data applications• Fusion energy• National security• Additive manufacturing• Many others!

AD Scope

Deliver science-based applications able to exploit exascale for high-confidence insights and answers to problems of National importance

Mission need

Create and enhance applications through:

- Development of models, algorithms, and methods
- Integration of software and hardware using co-design methodologies
- Improvement of exascale system readiness and utilization
- Demonstration and assessment of challenge problem capabilities

Objective

Deliver a broad array of comprehensive science-based computational applications that effectively exploit exascale HPC technology to provide breakthrough modeling and simulation solutions for National challenges:

- Scientific discovery
- Energy assurance
- Economic competitiveness
- Health enhancement
- National security

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 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)
<p>Accelerate design and commercialization of next-generation small modular reactors*</p> <p>Climate Action Plan; SMR licensing support; GAIN</p>	<p>Accurate regional impact assessment of climate change*</p> <p>Climate Action Plan</p>	<p>Biofuel catalysts design; stress-resistant crops</p> <p>Climate Action Plan; MGI</p>	<p>Increase efficiency and reduce cost of turbine wind plants sited in complex terrains*</p> <p>Climate Action Plan</p>	<p>Design high-efficiency, low-emission combustion engines and gas turbines*</p> <p>2020 greenhouse gas and 2030 carbon emission goals</p>
				


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




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



Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development seed projects

Seismic (EERE, NE, NNSA)

Reliable earthquake hazard and risk assessment in relevant frequency ranges*


DOE Critical Facilities Risk Assessment; urban area risk assessment; treaty verification



Carbon Capture and Storage (FE)

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


Climate Action Plan; SunShot; 2020 greenhouse gas/2030 carbon emission goals



Chemical Science (BES)

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


Climate Action Plan; SunShot Initiative; MGI



Urban Systems Science (EERE)

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

Energy-Water Nexus; Smart Cities Initiative




Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development seed projects

Metagenomics (BER)

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

Climate Action Plan; Human Microbiome Project; Marine Microbiome Initiative



Astrophysics (NP)

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

2015 Long Range Plan for Nuclear Science; origin of universe and nuclear matter in universe



Power Grid (EERE, OE)

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

Grid Modernization Initiative; Climate Action Plan



Application Co-Design (CD)

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

Essential to ensure that applications effectively utilize exascale systems

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

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

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

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

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

ECP Co-Design Centers

- **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

Some 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
- Mapping “traditional” HPC applications to current and inbound data hardware
- Infusing data science apps and components into current workflows (e.g., ML for OTF subgrid models)
- Achieving portable performance (without “if-def’ing” 2 different code bases)
- Multi-physics coupling: both algorithms (Picard, JFNK, Anderson Acceleration, HOLO, ...) and S/W (e.g., DTK, ADIOS, ...); what to use where and how
- Integrating sensitivity analysis, data assimilation, and uncertainty quantification technologies
- Staffing (recruitment & retention)

ECP Software Technology (ST) Focus Area

Rajeev Thakur, ECP ST Director

Pat McCormick, ECP ST Deputy Director

SC16 Birds of a Feather, "The U.S. Exascale Computing Project"

November 16, 2016

Salt Lake City, UT



EXASCALE COMPUTING PROJECT



U.S. DEPARTMENT OF
ENERGY


Office of
Science

www.ExascaleProject.org

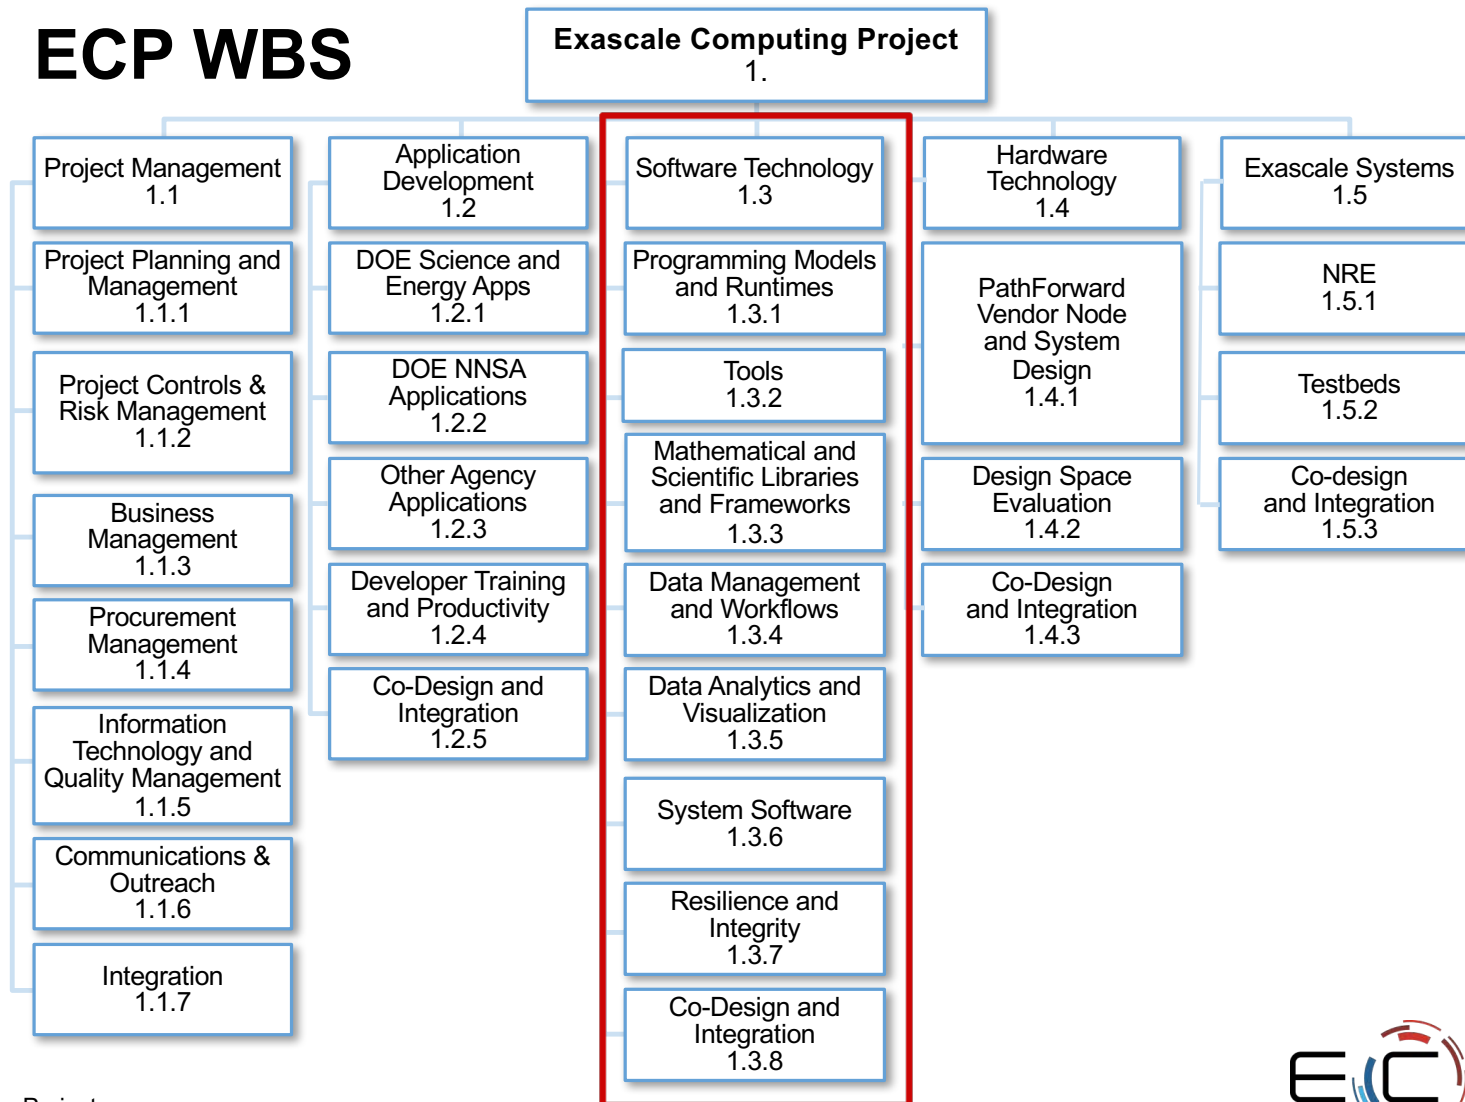
Summary

- 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

ECP leadership team



ECP WBS



Software Technology Level 3 WBS Leads

Programming Models and Runtimes 1.3.1	Rajeev Thakur, ANL
Tools 1.3.2	Jeff Vetter, ORNL
Mathematical and Scientific Libraries and Frameworks 1.3.3	Mike Heroux, SNL
Data Management and Workflows 1.3.4	Rob Ross, ANL
Data Analytics and Visualization 1.3.5	Jim Ahrens, LANL
System Software 1.3.6	Martin Schulz, LLNL
Resilience and Integrity 1.3.7	Al Geist, ORNL
Co-Design and Integration 1.3.8	Rob Neely, LLNL

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

Example: An Exascale Subsurface Simulator of Coupled Flow, Transport, Reactions and Mechanics*

Exascale Challenge Problem

- Safe and efficient use of the subsurface for geologic CO₂ sequestration, petroleum extraction, geothermal energy and nuclear waste isolation
- Predict reservoir-scale behavior as affected by the long-term integrity of hundreds of thousands deep wells that penetrate the subsurface for resource utilization
- Resolve pore-scale (0.1-10 μm) physical and geochemical heterogeneities in wellbores and fractures to predict evolution of these features when subjected to geomechanical and geochemical stressors
- Integrate multi-scale (μm to km), multi-physics in a reservoir simulator: non-isothermal multiphase fluid flow and reactive transport, chemical and mechanical effects on formation properties, induced seismicity and reservoir performance
- Century-long simulation of a field of wellbores and their interaction in the reservoir

Applications & S/W Technologies

Applications

- Chombo-Crunch, GEOS

Software Technologies Cited

- C++, Fortran, LLVM/Clang
- MPI, OpenMP, CUDA
- Raja, CHAI
- Chombo AMR, PETSc
- ADIOS, HDF5, Silo, ASCTK
- VisIt

Risks and Challenges

- Porting to exascale results in suboptimal usage across platforms
- No file abstraction API that can meet coupling requirements
- Batch scripting interface incapable of expressing simulation workflow semantics
- Scalable AMG solver in PETSc
- Physics coupling stability issues
- Fully overlapping coupling approach results inefficient.

Development Plan

Y1: Evolve GEOS and Chombo-Crunch; Coupling framework v1.0; Large scale (100 m) mechanics test (GEOS); Fine scale (1 cm) reactive transport test (Chombo-Crunch)

Y2: GEOS+Chombo-Crunch coupling for single phase; Coupling framework w/ physics; Multiphase flow for Darcy & pore scale; GEOS large strain deformation conveyed to Chombo-Crunch surfaces; Chombo-Crunch precip/dissolution conveyed to GEOS surfaces

Y3: Full demo of fracture asperity evolution-coupled flow, chemistry, and mechanics

Y4: Full demo of km-scale wellbore problem with reactive flow and geomechanical deformation, from pore scale to resolve the geomechanical and geochemical modifications to the thin interface between cement and subsurface materials in the wellbore and to asperities in fractures and fracture networks

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.

Software Technologies

Aggregate of technologies cited in all 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, Hypr, 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**

No. of ECP Application Proposals a Software is Mentioned in



Conceptual ECP Software Stack



Selection Process for Software Technology Projects

- RFI (Request for Information) sent on Feb 26, 2016, to selected PIs from DOE labs and universities
- PIs selected based on their history of developing software that runs on large-scale HPC systems.
- Total of 81 recipients of the RFI
 - They could include others (from labs or universities) as collaborators
- Received 109 3-page preproposals on March 14
- Preproposals were reviewed by subject matter experts, DOE lab leadership, and ECP team based on published selection criteria

Selection Process for Software Technology Projects

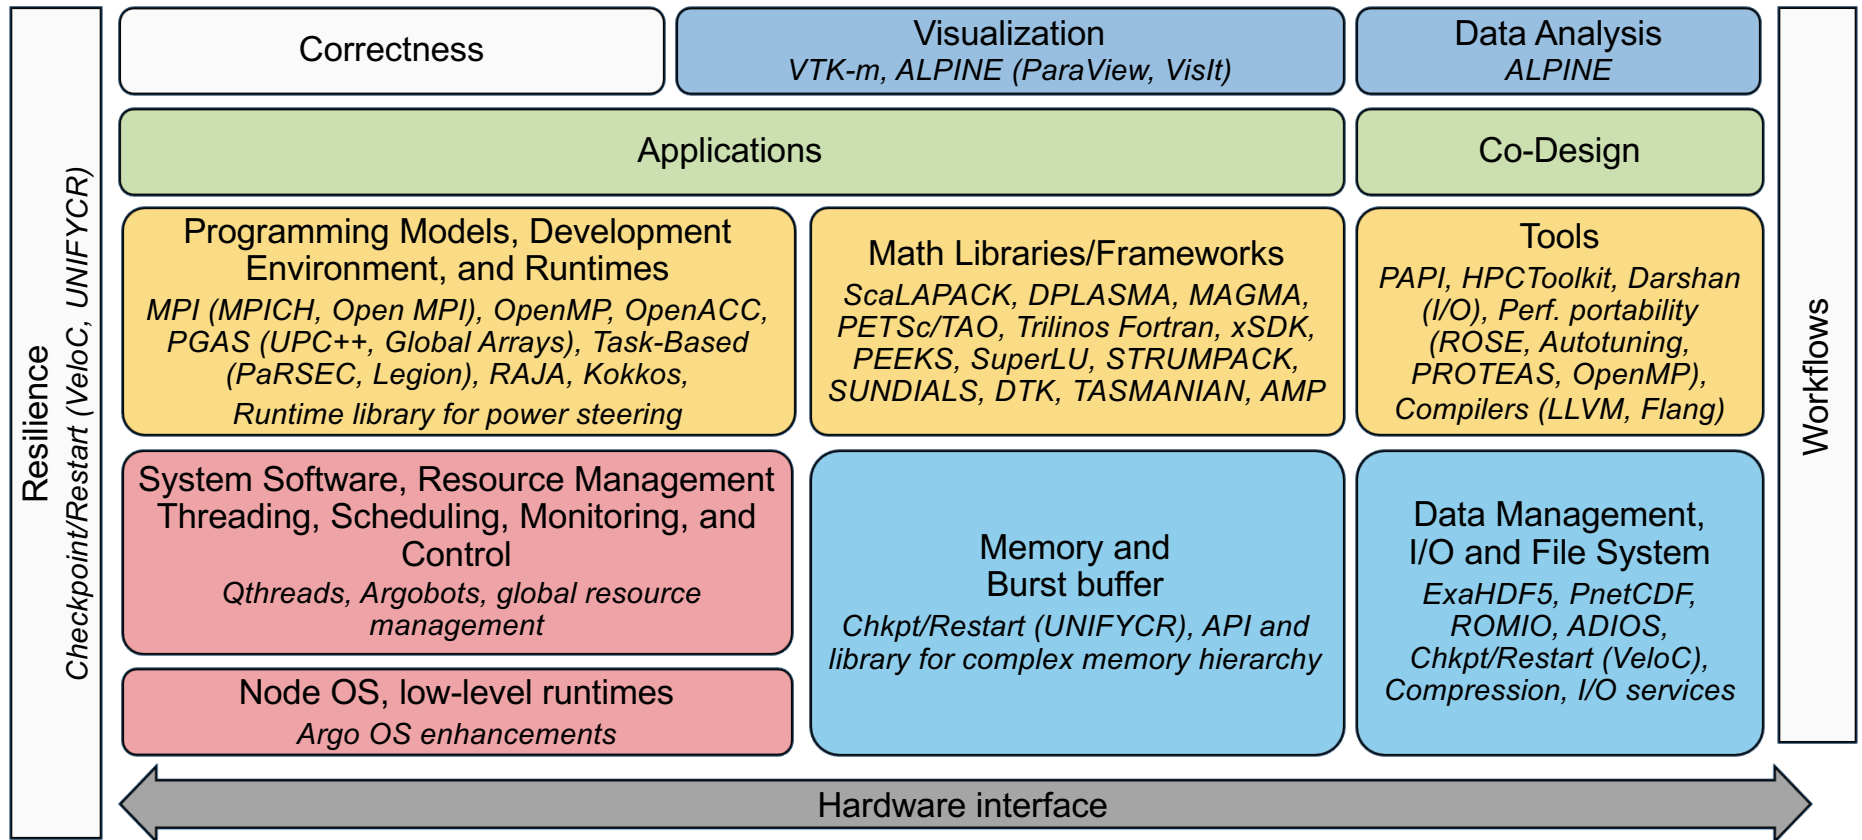
- RFP (Request for Proposals) issued on July 5 to 50 of the preproposals
- Some related preproposals were asked to merge, bringing the number down to 43
- 43 full proposals were received by the deadline of August 10
- Proposals reviewed by independent experts from universities, labs, industry, and abroad, based on published evaluation criteria
- Based on the reviews, requirements analysis, coverage analysis, and needs of the project, 35 of the proposals were selected for funding

The Exascale Computing Project Awards \$34 Million for Software Development

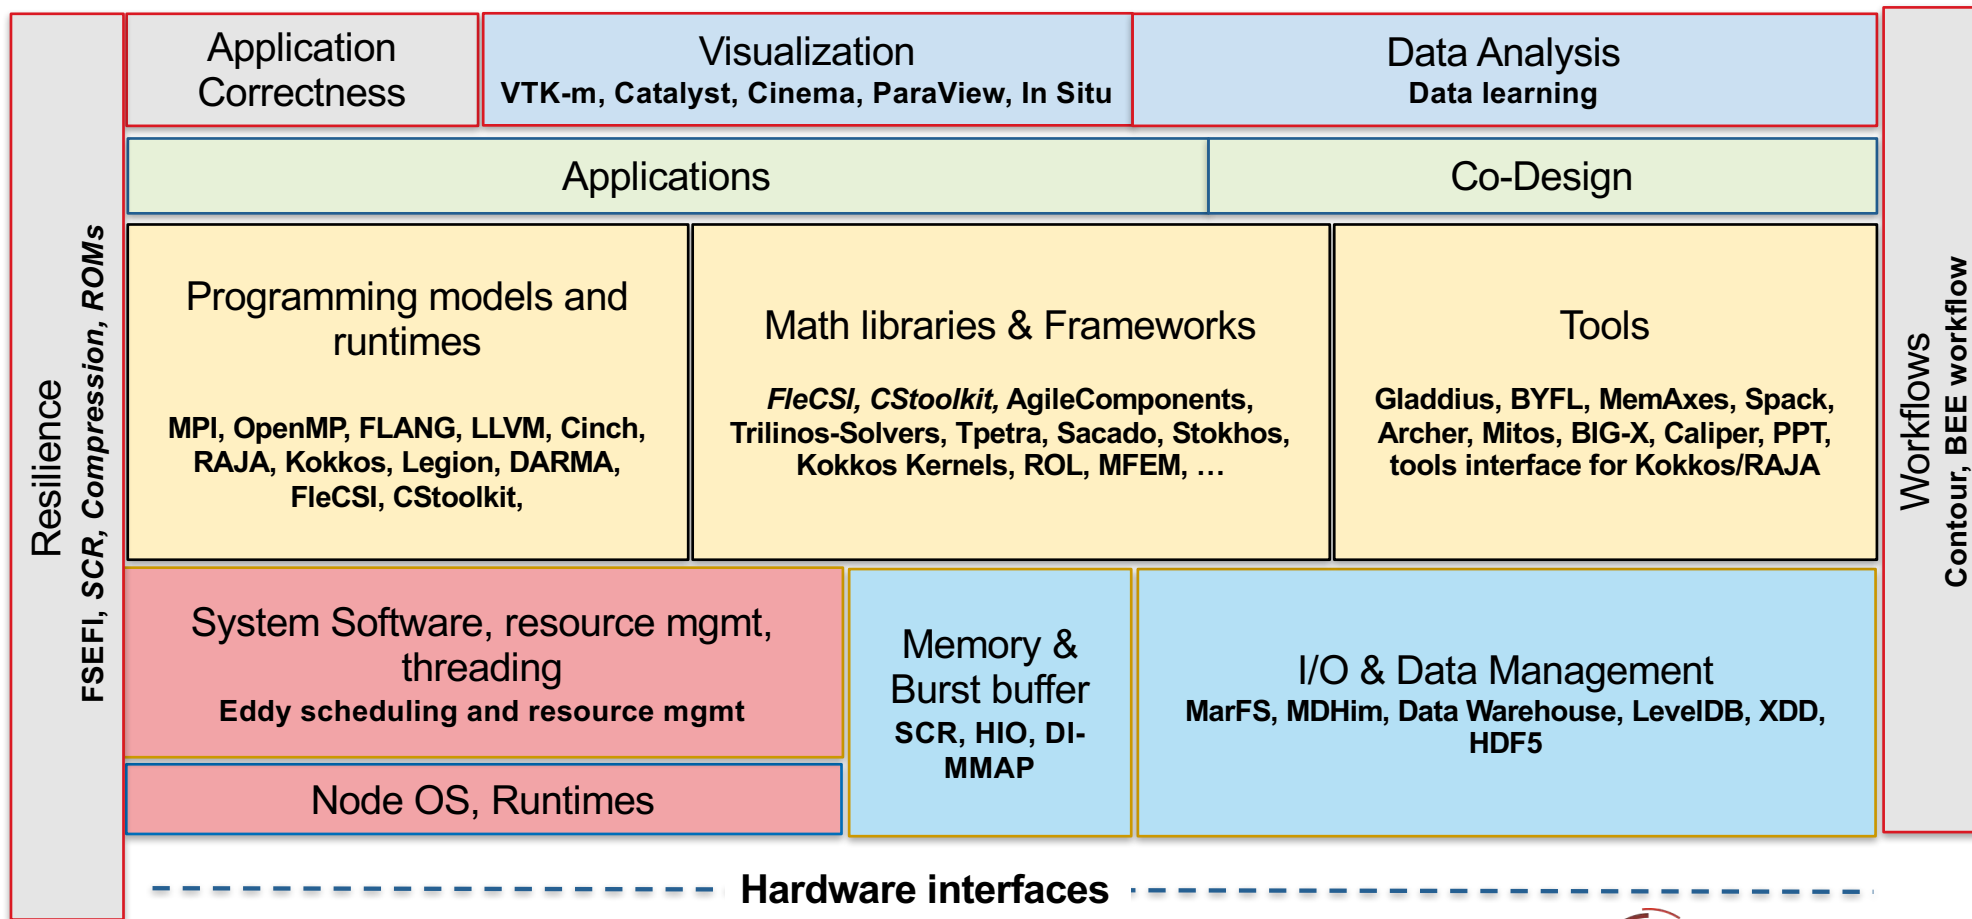
OAK RIDGE, Tenn., Nov. 10, 2016 – The Department of Energy’s Exascale Computing Project (ECP) today announced the selection of 35 software development proposals representing 25 research and academic organizations.

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

Recent ST Selections Mapped to Software Stack



NNSA ATDM Projects in ECP Software Technology



Challenges for Software Technology

- In addition to the usual exascale challenges of 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

- Over the next few months, we plan to 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 Hardware Technology (HT) Focus Area

Jim Ang, ECP HT Director

John Shalf, ECP HT Deputy Director

SC16 Birds of a Feather, "The U.S. Exascale Computing Project"

November 16, 2016

Salt Lake City, UT



EXASCALE COMPUTING PROJECT



U.S. DEPARTMENT OF
ENERGY

Office of
Science

www.ExascaleProject.org

ECP HT 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

- Reduces the Technical Risk for NRE investments in Exascale Systems (ES)
- 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

Scope

Develop the technology needed to build and support the Exascale systems

Mission need

The Exascale Computing Project requires Hardware Technology R&D to enhance application and system performance for science, engineering and data-analytics applications on exascale systems

Objective

Support hardware architecture R&D at both the node and system architecture levels

Prioritize R&D activities that address ECP performance objectives for the initial Exascale System RFPs

Enable Application Development, Software Technology, and Exascale Systems to improve the performance and usability of future HPC hardware platforms (holistic codesign)

Hardware Technology Focus Area

- Leverage our window of time to support advances in both system and node architectures
- Close gaps in vendor's technology roadmaps or accelerate time to market to address ECP performance targets *while* affecting and intercepting the 2019 Exascale System RFP
- Provide an opportunity for ECP Application Development and Software Technology efforts to *influence* the design of future node and system architecture designs

Hardware Technology Overview

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

Issue *PathForward* 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 deliverables
- Lead Design Space Evaluation through Architectural Analysis, and Abstract Machine Models of PathForward designs for co-design

Overarching Goals for PathForward

- Improve the quality and number of competitive offeror responses to the 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 Exascale Systems RFP
- Improve DOE confidence in technology performance benefit, programmability and ability to integrate into a credible system platform acquisition

PathForward addresses the disruptive trends in computing due to the power challenge

Power challenge

Processor/node trends

System trends

Disruptive changes

4 Challenges: power, memory, parallelism, and resiliency

- End of Dennard scaling
- Today's technology: ~50MW to 100 MW to **power** the largest systems


- GPUs/accelerators
- Simple in order cores
- Unreliability at near-threshold voltages
- Lack of large-scale cache coherency
- Massive on-node **parallelism**

- Complex hardware
- Massive numbers of nodes
- Low bandwidth to **memory**
- Drop in platform **resiliency**

- New algorithms
- New programming models
- Less "fast" memory
- Managing for increasing system disruptions
- High power costs

PathForward will drive improvements in vendor offerings that address ECP's needs for scale, parallel simulations, and large scientific data analytics

Integration and Co-Design is key



Capable exascale computing requires close coupling and coordination of key development and technology R&D areas

Holistic Co-Design

- ECP is a very large DOE Project, composed of over 80 separate projects
 - Many organizations: National Labs, Vendors, Universities
 - Many technologies
 - At least two diverse system architectures
 - Different timeframes (Three phases)
- For ECP to be successful,
the whole must be more than the sum of the parts

Co-Design requires Culture Change

- AD and ST teams cannot assume that the node and system architectures are firmly defined as inputs to develop their project plans
- HT PathForward projects also cannot assume that applications, benchmarks and the software stack are fixed inputs for their project plans
- Initial assumptions about inputs lead to preliminary project plans with associated deliverables, but there needs to be flexibility
- Each ECP project needs to understand that they do not operate in a vacuum
- *In Holistic Co-Design, each project's output can be another project's input*

Co-Design and ECP Challenges

- Multi-disciplinary Co-Design teams
 - ECP project funding arrives in technology-centric bins, e.g. the focus areas
 - ECP leadership must foster integration of projects into collaborative Co-Design teams
 - Every ECP project's performance evaluation will include: *how well they play with others*

Co-Design and ECP Challenges

- With ~25 AD teams, ~50 ST teams, ~5 PathForward teams:
All-to-all communication is impractical
- The Co-Design Centers and the HT Design Space Evaluation team will provide capabilities to help manage some of the communication workload
 - Proxy Applications and Benchmarks
 - Abstract Machine Models and Proxy Architectures
- The ECP Leadership team will be actively working to identify:
 - Alignments of cross-cutting projects that will form natural Co-Design collaborations
 - Orthogonal projects that do not need to expend time and energy trying to force an integration – this will need to be monitored to ensure a new enabling technology does not change this assessment

ECP Exascale Systems (ES) Focus Area

Terri Quinn, ECP ES Director

Susan Coghlan, ECP ES Deputy Director

SC16 Birds of a Feather, "The U.S. Exascale Computing Project"

November 16, 2016

Salt Lake City, UT



EXASCALE COMPUTING PROJECT



U.S. DEPARTMENT OF
ENERGY

Office of
Science

www.ExascaleProject.org

Exascale Systems: Capable exascale systems by 2023

Systems must meet ECP's essential performance parameters

Four key challenges will be addressed through targeted R&D investments to bridge the capability gap

“Non-recurring engineering” (NRE) activities will be integral to next-generation computing hardware and software.



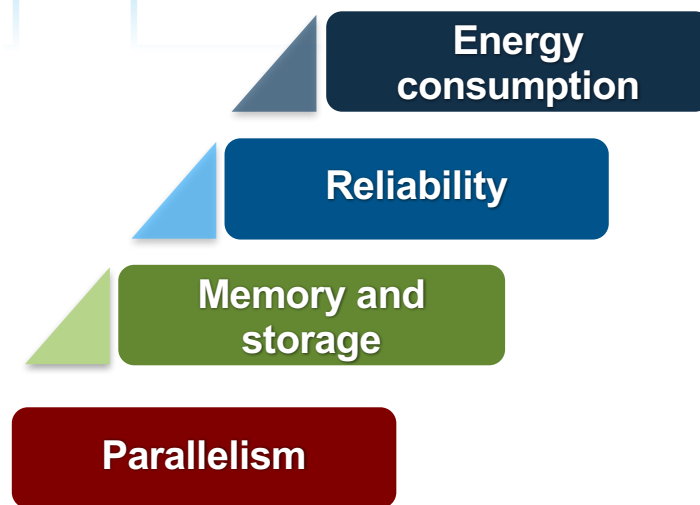
50 times
the current
performance



10 times
reduction in power
consumption



System resilience:
6 days
without app failure



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

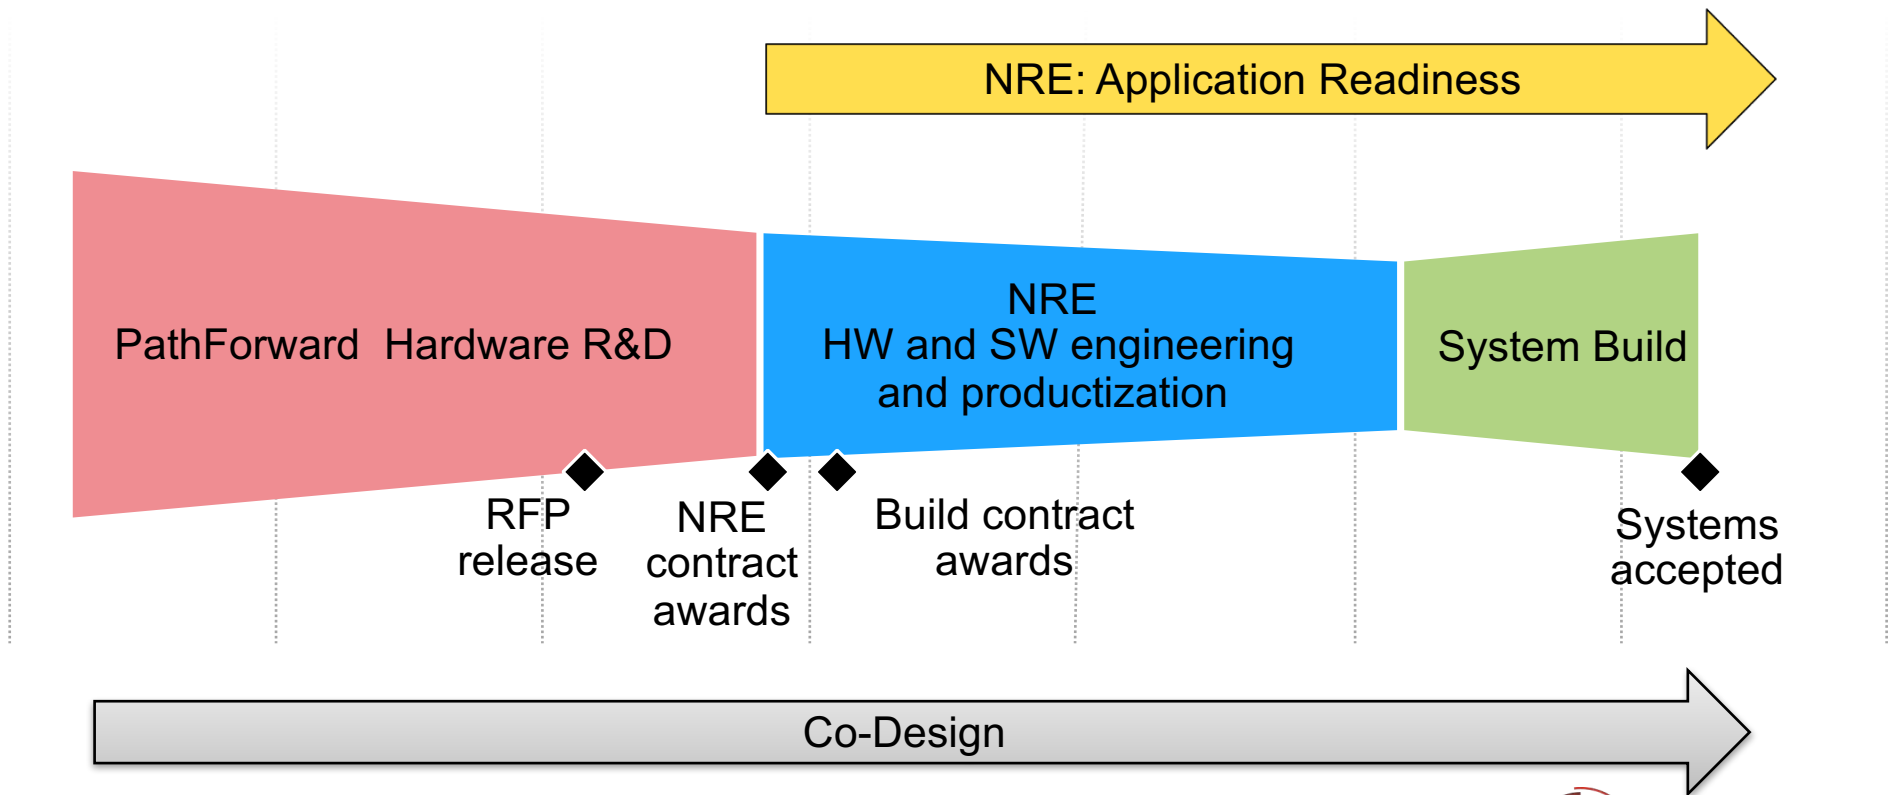
ECP's and SC/NNSA's processes will be tightly coupled and interdependent

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



Summary

- ES ensures at least two systems are accepted by no later than 2023 that meet ECP's requirements
- SC and NNSA will acquire these systems in collaboration with ECP
- ECP will make substantial investments in the systems through NRE contracts
- ES will acquire testbeds for application and software development projects and for hardware investigations

Thank you!

www.ExascaleProject.org



EXASCALE COMPUTING PROJECT



U.S. DEPARTMENT OF
ENERGY

Office of
Science