

# Exascale Applications: Opportunities and Challenges

## Update on Activities in the Exascale Computing Project (ECP) Application Development Focus Area

**Douglas B. Kothe**, Oak Ridge National Laboratory

Presentation to the Advanced Scientific Computing Advisory Committee (ASCAC)  
DOE Office of Science Advanced Scientific Computing Research Program  
Washington, DC  
September 20, 2016



EXASCALE COMPUTING PROJECT

# Outline

- ECP Overview
- ECP Application Development (AD): scope and objectives
- Application project selection process
- Application Co-Design
- Application project examples: challenge problems, applications, plans
- Application requirements
- Risks/challenges, merits of projectization, next steps

# DOE's Role in the NSCI

- The National Strategic Computing Initiative (NSCI) that was created by President Obama aims to maximize the benefits of HPC for US economic competitiveness and scientific discovery
  - Leadership in high-performance computing and large-scale data analysis will advance national competitiveness in a wide array of strategic sectors, including basic science, national security, energy technology, and economic prosperity.
  - NSCI Strategic Objective (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 petaflop systems across a range of applications representing government needs.
  - NSCI Strategic Objective (2) Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.
- DOE is a lead agency within NSCI; DOE SC and NNSA will execute a joint effort on advanced simulation through a **capable exascale** computing program emphasizing sustained performance on relevant applications and data analytic computing
- The Exascale Computing Project is the vehicle for that effort

# Capable exascale computing is the gateway to the future


- Support applications solving science problems 50× faster or more complex than today's 20 PF systems
- Operate in a power envelope of 20–30 MW
- Be sufficiently resilient (average fault rate no worse than weekly)
- At least two diverse system architectures
- Possess a software stack that meets the needs of a broad spectrum of applications
- A holistic project approach is needed that uses co-design to develop new platform, software, and computational science capabilities at heretofore unseen scale
  - Essential for tackling much deeper challenges than those that can be solved by hardware scale alone

Builds on 4+ decades of US expertise and leadership in computing and computational science

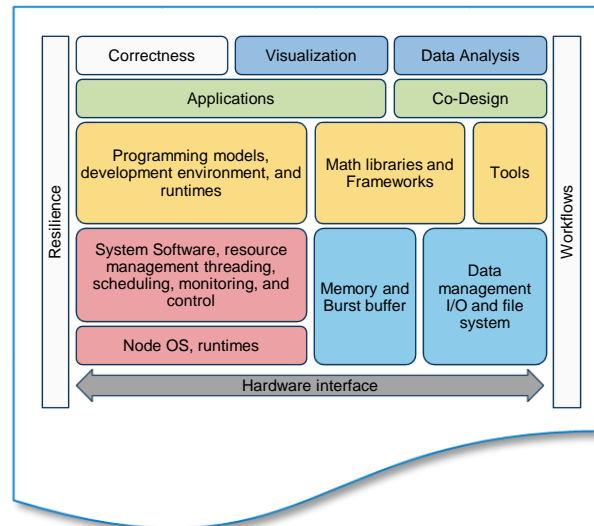
# To achieve capable exascale requires a holistic approach




Science and mission applications



Scalable and productive software stack



Hardware technology elements



Integrated exascale supercomputers



# ECP Mission Need Defines the Application Strategy

Support DOE science <i>and</i> energy missions	Meet national security needs	Key science and technology challenges to be addressed with exascale
<ul style="list-style-type: none"><li>• Discover and characterize next-generation <b>materials</b></li><li>• Systematically understand and improve <b>chemical processes</b></li><li>• Analyze the extremely large datasets resulting from the next generation of <b>particle physics</b> experiments</li><li>• Extract knowledge from <b>systems-biology</b> studies of the microbiome</li><li>• Advance <b>applied energy</b> technologies (e.g., whole-device models of plasma-based fusion systems)</li></ul>	<ul style="list-style-type: none"><li>• Stockpile Stewardship Annual Assessment and Significant Finding Investigations</li><li>• Robust uncertainty quantification (UQ) techniques in support of lifetime extension programs</li><li>• Understanding evolving nuclear threats posed by adversaries and in developing policies to mitigate these threats</li></ul>	<ul style="list-style-type: none"><li>• Materials discovery and design</li><li>• Climate science</li><li>• Nuclear energy</li><li>• Combustion science</li><li>• Large-data applications</li><li>• Fusion energy</li><li>• National security</li><li>• Additive manufacturing</li><li>• Many others!</li></ul>

# Exascale Application Driver

## Wind Energy

### Gaps and Opportunities


- ✓ Wide-scale deployment of unsubsidized wind plants hampered by large plant-level energy losses, currently at ~20%

### Simulation Challenge Problems

- ✓ Wind turbine blade-resolved CFD detached eddy simulation of a ~50-turbine wind plant over a useful operating period (hours)
- ✓ Predict plant flow physics: starting from the atmospheric boundary layer interaction with the wind plant, down to turbine-turbine interactions, the response of individual turbines, and impact of complex terrain

### Prospective Outcomes and Impact

- ✓ Harden wind plant design and layout against energy loss susceptibility
- ✓ 1% wind plant performance improvement translates to >\$100M annual cost savings
- ✓ Higher penetration of wind energy (~20-30%) in U.S. electrical supply supports President's Climate Action Plan



# Exascale Application Driver

## Nuclear Energy

### Gaps and Opportunities


- ✓ Understanding and predicting fuel failure and core damage in severe reactor accidents
- ✓ Near real-time load-following core simulator supporting onsite operating plant decisions
- ✓ Engineering scale predictions of nuclear fuel performance and barriers to higher burnup

### Simulation Challenge Problems

- ✓ Core-wide multi-physics: Monte Carlo neutronics, fluids, fuels, chemistry, material, local wear and contacts, structural response
- ✓ Core-wide multi-scale fluids performance with two-way DNS-to-LES CFD coupling

### Prospective Outcomes and Impact

- ✓ Accelerate and innovate designs for small and advanced reactors
- ✓ Virtual test reactor for advanced fuel designs outside of principal test base
- ✓ Improved efficiency, economics, and safety of existing reactor fleet





# Exascale Application Driver

## Magnetic Fusion Energy

### Gaps and Opportunities


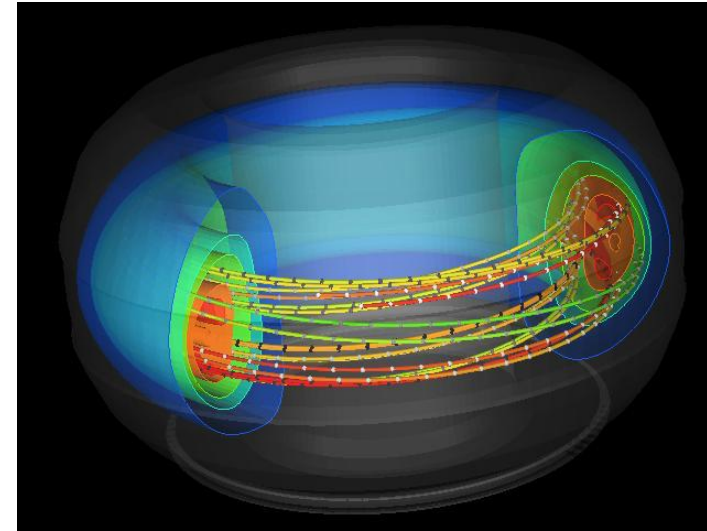
- ✓ Prepare for and exploit ITER and other coming international major experiments such as JET-DT, JT-60SA, and W 7-X

### Simulation Challenge Problems

- ✓ Whole-device fusion reactor simulation: tightly-coupled full-f/delta-f models and loosely coupled source/boundary models, including electron/ion kinetics, MHD, and energetic particles in core and edge regions
- ✓ Simulate and characterize tokamak disruptions and mitigations, incorporating kinetics, MHD, and fast particles
- ✓ Plasma boundary region analysis: edge kinetic effects, material interaction, radiation and detachment, power and particle exhaust

### Prospective Outcomes and Impact

- ✓ Prepare for and fully simulate ITER experiments and increase ROI of validation data and understanding
- ✓ Prepare for next step beyond-ITER devices such as nuclear science facilities and DEMO



# Exascale Application Driver

## Advanced Manufacturing

### Gaps and Opportunities



- ✓ Advance quality, reliability, and application breadth of additive manufacturing (AM)
- ✓ Accelerate innovation in clean energy manufacturing institutes (NNMIs)
- ✓ Capture emerging manufacturing markets

### Simulation Challenge Problems

- ✓ Continuum level predictions of non-uniform microstructure and its relationship to process parameters
- ✓ Predictive mesoscale models for dendritic solidification then scale-bridged to continuum

### Prospective Outcomes and Impact

- ✓ Routine qualification of AM parts via process-aware design specs and reproducibility through process control
- ✓ Fabrication of metal parts with unique properties such as light weight strength and failure-proof joints and welds



# Exascale Application Driver

## Combustion Science and Technology

### Gaps and Opportunities


- ✓ Combustion-based systems will dominate marketplace for decades; must be optimized for energy efficiency and reduced emissions
- ✓ Current cut-and-try approaches for combustion system design take too long; cannot evaluate design parameter space for optimized results

### Simulation Challenge Problems

- ✓ Fully coupled multi-scale and multi-physics LES treatment of complex combustion processes in propulsion and power devices, capturing combined effects of geometry, heat transfer, and multiphase reacting flow
- ✓ DNS of chemical mechanisms for kinetics of complex hydrocarbons and alternative fuels and related turbulence-chemistry interactions

### Prospective Outcomes and Impact

- ✓ Address outstanding challenges for advanced gas turbines and reciprocating engines: flame stabilization, flashback, thermo-acoustics, pollutant formation in gas turbines; effects of fuel composition and spray parameters on ignition and soot formation in engines



# Exascale applications will address key questions of national significance

Some examples . . .

Domain	Key Questions
<b>Climate</b> (BER)	What will the earth's sea level and global temperature be in 2100? How will extreme weather patterns change? How can we best prepare to adapt? What mitigation scenarios are most influential and also still implementable?
<b>Combustion</b> (BES)	Can engine efficiency and pollutant emission be predicted for conventional and new fuels? Can new engines that exhibit a better compromise between maximal energy efficiency and minimal emissions be designed more quickly ?
<b>Nuclear Energy</b> (NE)	What steps are possible to improve the efficiency, economics, and safety of the existing fleet? Can nuclear fuel be safely burned longer while withstanding severe accidents? What advanced nuclear fuel and reactor design concepts hold the most promise for deployment on the power grid?
<b>Carbon Capture and Storage</b> (FE)	What is necessary for safe and permanent carbon capture, storage, and utilization technologies to become technically and commercially feasible? How can the translation of these technologies to market be accelerated?
<b>Wind Energy</b> (EERE)	How can wind plant energy losses be reduced (e.g., by a few percent)? What performance improvements can be identified and implemented that will enable the economic viability of unsubsidized wind plants?
<b>Magnetic Fusion Energy</b> (FES)	Can tokamak plasma disruption physics be understood well enough to predict/avoid/mitigate their occurrence and associated deleterious effects? Can viable candidate plasma-facing components and technologies be identified? What is the optimum magnetic confinement configuration?

# Exascale applications will address key questions of national significance

Some examples . . .

Domain	Key Questions
<b>Stockpile Stewardship</b> (DP)	Will 3D simulations at unprecedented resolution of the multiple coupled physical processes in thermonuclear boost aid predictive capability, just as petascale has helped understand energy balance? Can the effects of unresolved physical processes be adequately captured in sub-grid models informed by UQ analyses resulting from ensembles of 3D simulations at lower resolution?
<b>Manufacturing</b> (EERE)	How can the quality and reproducibility of additively-manufactured (AM) metallic parts be assured? Can the microstructural features of AM parts be predicted, and, if so, can simulation insights guide the real-time monitoring and adjustment of AM processes (heat source spot size, intensity, motion) to yield the desired part?
<b>Astrophysics</b> (NP)	What are the mechanisms and consequences of stellar cataclysms, including “regular” supernovae, “collapsars” (stellar implosions into black holes), neutron star mergers, and thermonuclear X-ray bursts from the surfaces of neutron stars?
<b>QCD</b> (NP/HEP)	Can the structure of the standard model (SM) of particle physics be quantitatively understood, by examining its predictions for the properties of hadrons, nuclei and nuclear matter? Can the precision of properties of quark-anti-quark and three-quark states be increased, including their interactions with the W, Z and Higgs bosons? Can physical quark masses and the effects of electromagnetism be predicted?

# Exascale applications will address key questions of national significance

Some examples . . .

Domain	Key Questions
<b>Chemistry</b> (BES/BER)	Can more robust and selective catalysts and catalytic processes be designed that are orders of magnitude more efficient at temperatures that are hundreds of degrees lower? Can improved methods for sequestration, transport and cycling of chemical constituents of energy byproducts (e.g., carbon) in the environment be developed? Is the rational engineering of enzymes possible in the use of plants and microbes to produce sustainable biofuels and other bioproducts? Can novel and robust branched, hierarchical nanomaterials for energy capture, conversion and storage for our nation's sustainable energy systems be designed in-silico?
<b>Precision Medicine</b> (NIH)	Can predictive models for cancer support treatment choices a physician and patient make? Can promising new treatment options be identified through development of predictive pre-clinical models and can the development of diagnostic and targeted therapy be accelerated based on the ability to predict the fundamental mechanism of RAS/RAF driven cancers?
<b>Power Grid</b> (EDER)	Can we reliably and efficiently predict and optimize the stochastics and dynamics of a grid with intermittent renewable sources, electric vehicles, and smart loads?

# ECP Application Strategy

## Precipitate, foster, and grow a national push in applications

- Deliver to 10 DOE program offices: 5 SC, 4 Energy, NNSA Defense Programs
- Deliver to NSCI “deployment agencies”: NIH, NOAA, NASA, DHS, FBI
- Deliver on next-generation stockpile stewardship applications (3 NNSA Labs)
- Require incubation of “non-traditional” applications (data science, experimental work flow)
- Reassert international leadership and grow the gap between 1st (US) and followers

## Focus applications on specific needs

- Support key national initiatives
  - National security
  - Economic security
  - Energy security
  - Scientific discovery
  - Climate/environmental science
  - Healthcare
- Drive a broad set of requirements for the software ecosystem
- Drive a broad set of key algorithms and methods for applications (mathematical methods, key algorithms, data analytics approaches, ...)
- Integrate co-designed motifs for an essential element of success

## Support application development teams

- Team sizes based on complexity, maturity, strategic importance
- Application teams need not be led by DOE lab personnel (e.g., could be led by other federal agencies, universities, industry)
- Planned data-focused calls will follow the same tiered approach philosophy

# Application Development 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

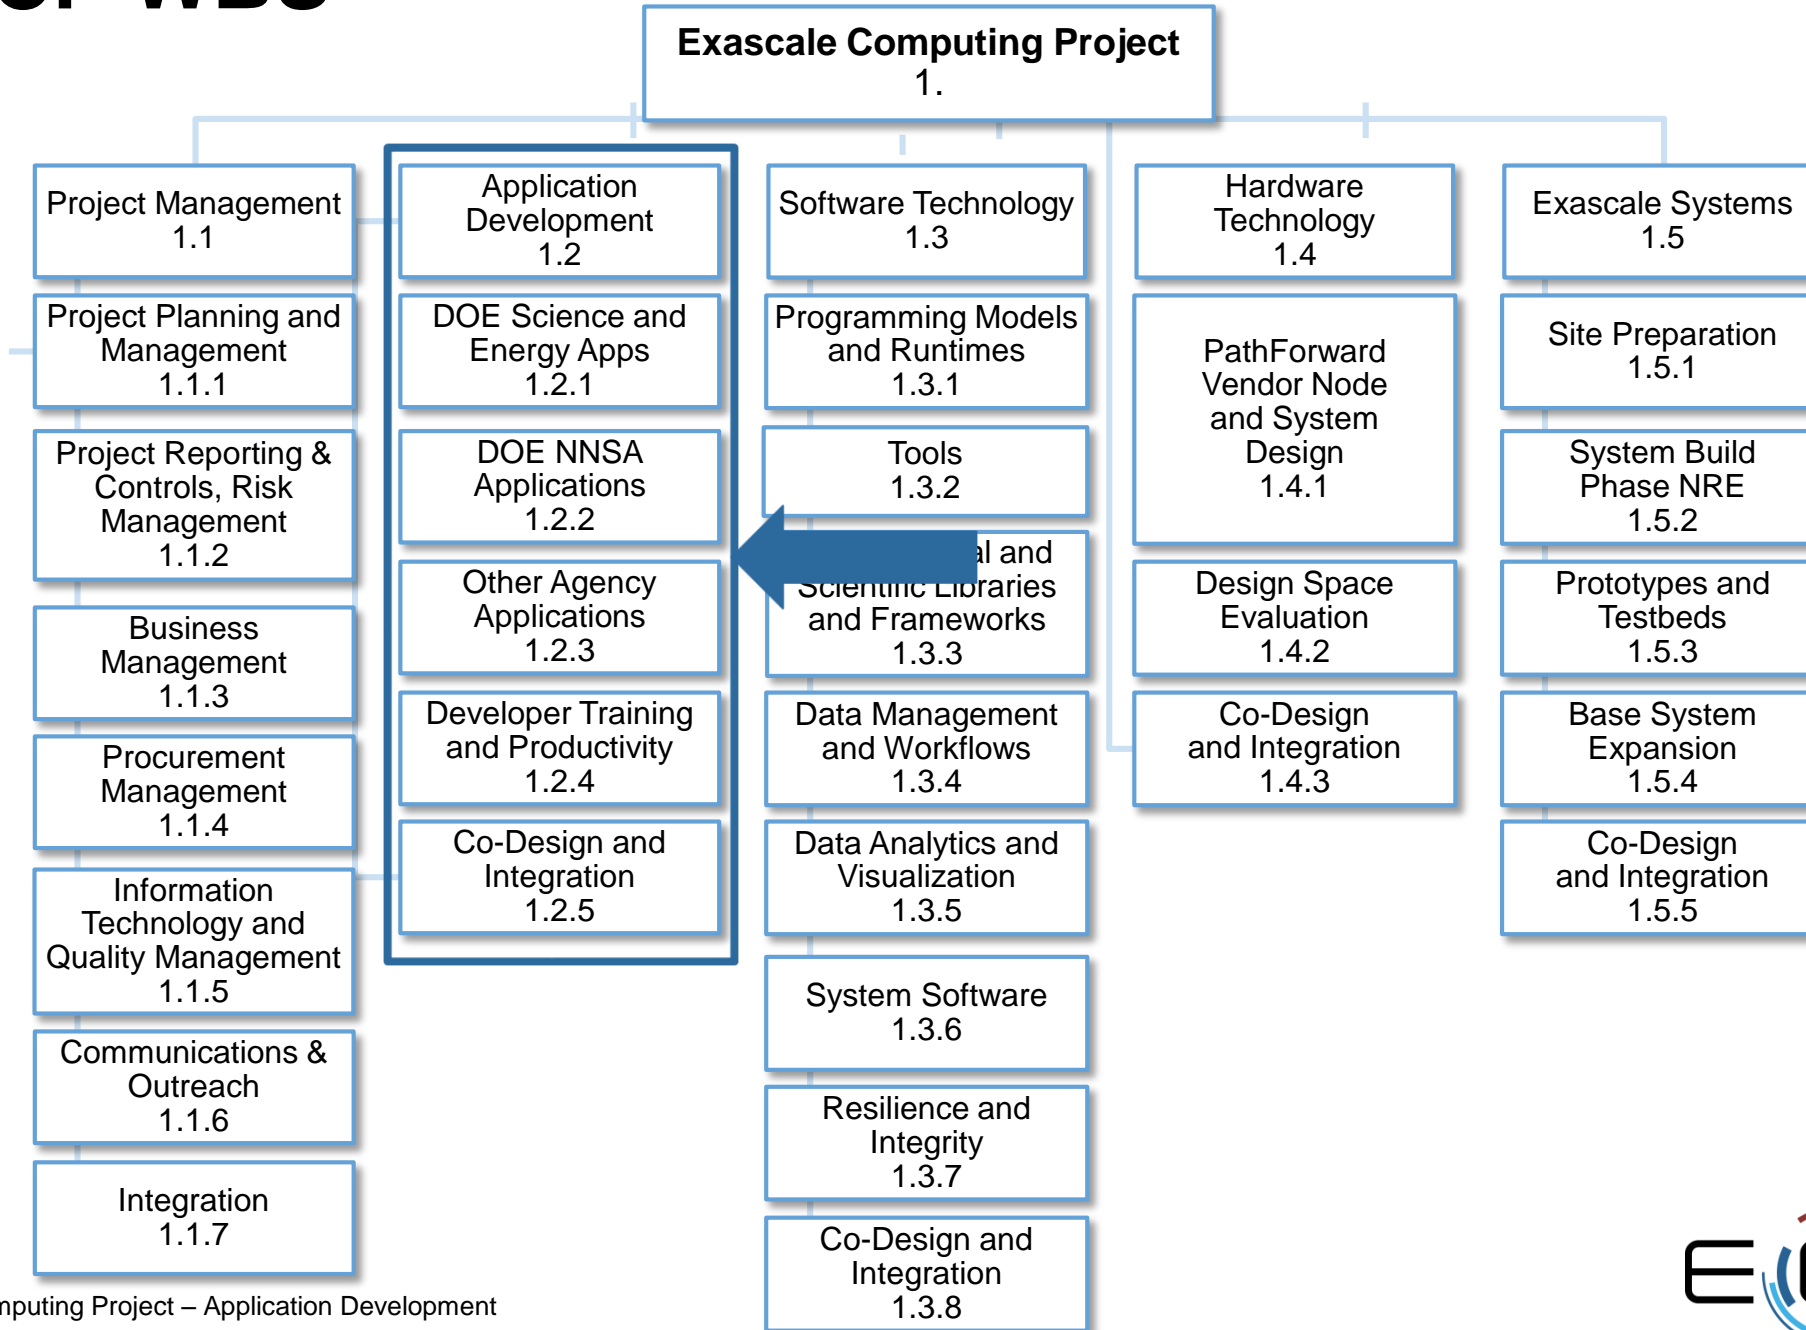
Economic competitiveness

Health enhancement

National security



# ECP WBS



# Selection Process for Application Suite

DOE/NSF/NIH RFI yielded many candidate applications

**First RFI issued to all 17 DOE Laboratories to identify candidate applications proposed to target specific challenge problems of national interest**

- Released May 31, 2015 by DOE SC/ASCR and NNSA/ASC to DOE Laboratories
- Identify key science/engineering areas and corresponding potential exascale applications from DOE National Labs
- 135 responses received in 35+ science and engineering problem areas

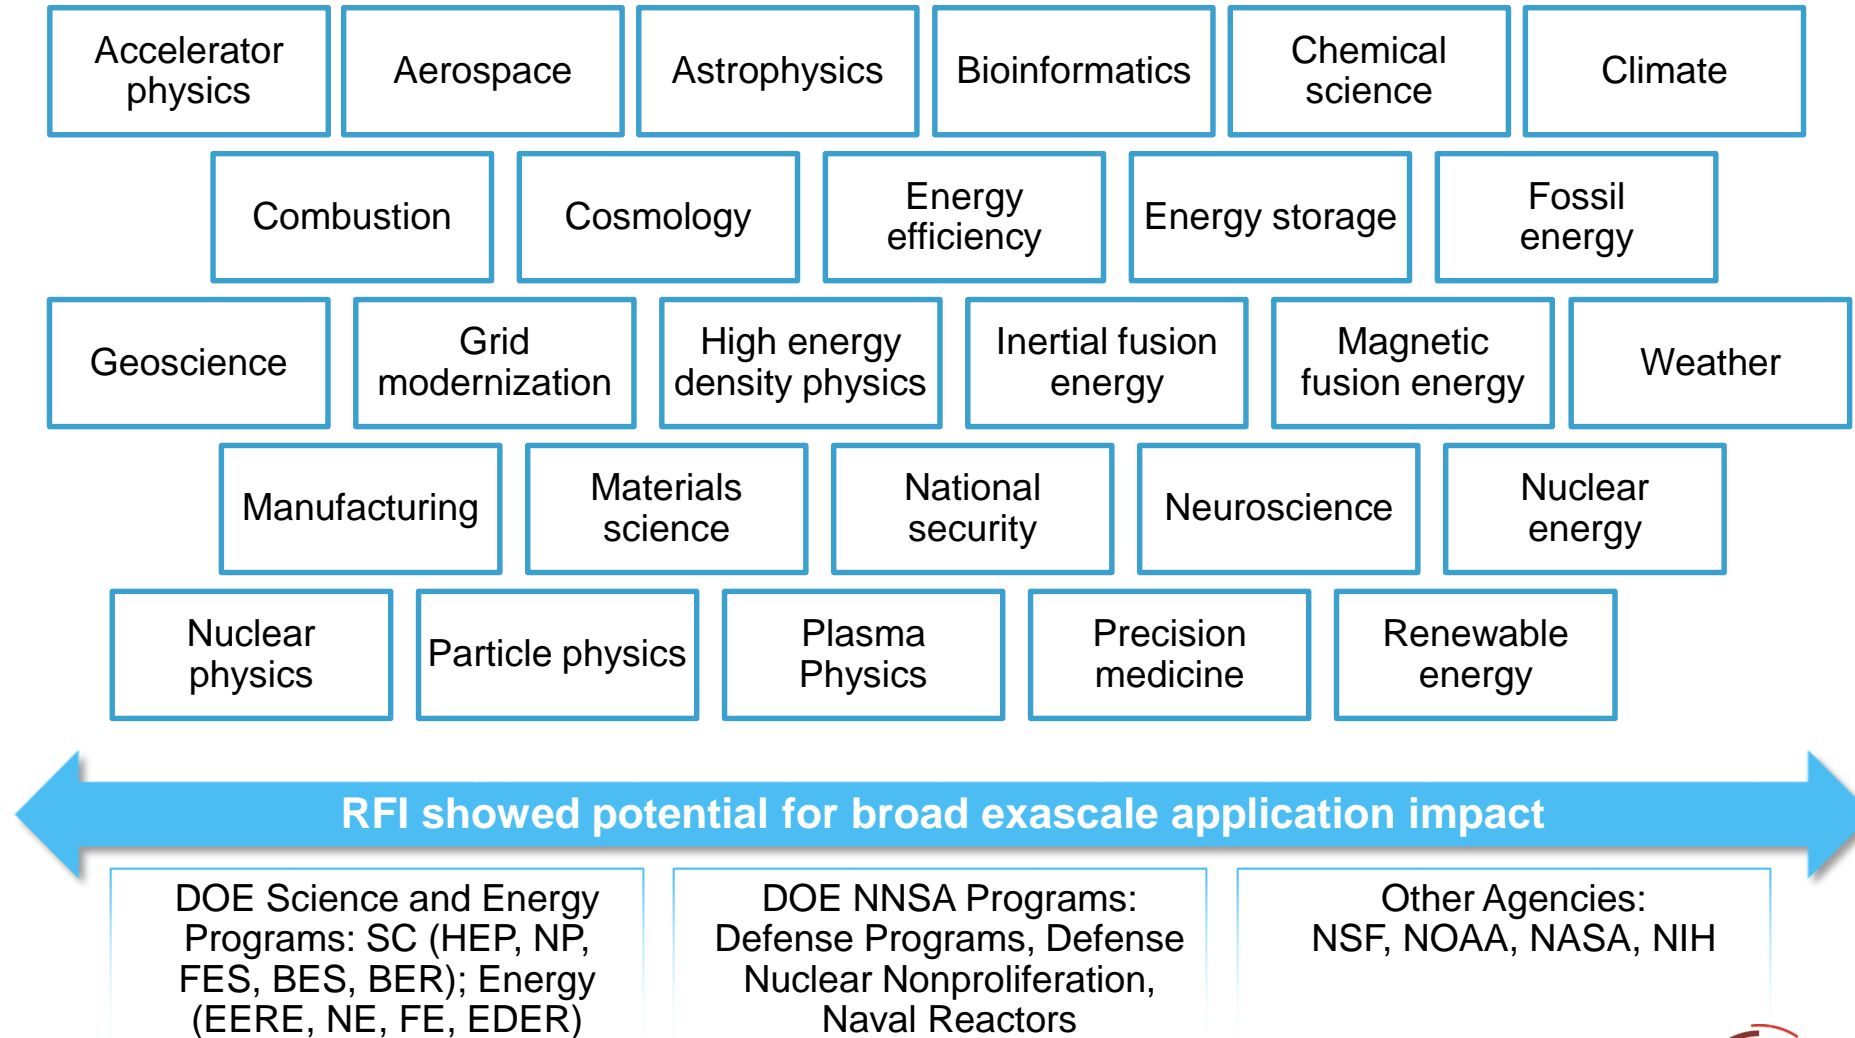
**Second RFI issued to NSF and NIH to identify scientific research topics in need of 100-fold HPC performance on scientific applications**

- Released Sep 15, 2015 by DOE, NIH, and NSF as an Open Solicitation
- Responses will help agencies construct roadmaps, build exascale ecosystems required to support scientific research, and inform the research, engineering and development process
- 114 responses received from NSF and NIH

**249 RFI response provided valuable input for ECP functional requirements**

- Applications can target impactful and mission critical Challenge Problems
- Applications vary in complexity, maturity, programmatic investment
- Application development team will be key to success: leadership, quality, agility, size

# Exascale Application RFI Responses Reflected Broad Domain Coverage of Mission Space in DOE and Other Agencies



# Application Development

## Process for selection of initial ECP applications (FY16)

- Gather input (concept white papers) from the scientific community on candidate exascale applications and the challenge problems targeted
  - Analyze responses to recent Requests for Information (RFI) issued to DOE Labs, NSF, and NIH
  - Work with DOE computing facility leadership in refining earlier (2006-2014) application requirements
    - Specific workshops: HEP (Jun 2015), Turbulence (Aug 2015); BES (Nov 2015), FES (Jan 2016), BER (Mar 2016), NP (Jun 2016), ASCR (Sep 2016)
  - Engage with key application development leaders as needed
- Reduce the candidate set of exascale application concepts (white papers) to a candidate set of ECP application projects (with detailed development plans)
  - Require DOE Labs submitting white papers to prioritize, reduce, and refine their RFI submissions
  - Seek stakeholder (DOE program) feedback on priorities and gaps
  - Form an Exascale Applications Working Group (EAWG) to provide subject matter expertise and facilitate prioritization and multi-institutional collaboration; seek formal review input for selecting application efforts invited to submit full development plans (currently consists of DOE Lab personnel – will likely expand)
  - Solicit detailed development plans from a list of candidate ECP applications
- Select ECP applications from list of candidate application projects
  - Draw upon SMEs as needed for external review of development plans



# Candidate ECP Applications

## Scope of development plans requested

- **4-year projects, with option for 2<sup>nd</sup> 4-year renewal**
- **Development plan template**
  - Challenge problem description
  - Application description & development needs
    - models, algorithms, software, data, testing & assessment
  - Exascale system utilization
  - Expected outcomes and figures of merit
  - Project team and management
  - Challenges and risks
  - Resource allocation
  - Data management
  - Integration

# Candidate ECP Applications

## Domain (number), stakeholder programs, and participating institutions

Domain	Program(s)	Labs	Domain	Program(s)	Labs
Material Science (4)	BES, FES, NE, NNSA	ANL, BNL, LBNL, ORNL, SLAC, SNL, LLNL, LANL	Plasma Physics [MCF, ICF, HEDP] (4)	FES, NNSA	LANL, LLNL, SNL, PPPL, ANL, ORNL
Data Science / Experimental Facility (3)	BES, BER	ANL, BNL, LBNL, SLAC, ORNL, SNL, LANL	Bioinformatics (2)	BER, NIH	LBNL, ORNL
Chemistry (2)	BES, BER	Ames, ANL, PNNL, BNL, ORNL, LBNL	Accelerator Physics (1)	HEP	LBNL, LLNL, SLAC
Combustion (2)	BES, EERE	ANL, LANL, LLNL, SNL, LBNL, NREL	Manufacturing (1)	EERE	ORNL, LANL, LLNL, SNL
QCD / Nuclear Interactions (2)	NP, HEP	BNL, TJNAF, LANL, ANL, LBNL, LLNL, ORNL	Subsurface / Geoscience (1)	BES	LBNL, LLNL, NETL
Cosmology (1)	HEP	ANL, LANL, LBNL	Nuclear Energy (1)	NE, NNSA	ANL, ORNL, INL
Urban Science (1)	EERE, HUD, DOT	ANL, LBNL, NREL, ORNL, PNNL	Seismic (1)	NE, NNSA	LBNL, LLNL
Climate (1)	BER	SNL, ANL, LANL, LBNL, LLNL, ORNL, PNNL	Wind (1)	EERE	NREL, ORNL, SNL
Power Grid (1)	EDER	PNNL, ANL, NREL	Precision Medicine (1)	NIH, NSF	ANL, LANL, LLNL, ORNL
Astrophysics (1)	NP	LBNL, ANL, ORNL	Hypersonic Vehicles (1)	DoD	LLNL
Multiphase Flows (1)	FE, NE	LBNL, NETL			

# Candidate ECP Applications

## Review criteria

- **Quality of the application's *challenge problem(s)***
  - Would the challenge problem, if addressed, represent a major advance in its field? Would its solution address an important and time-urgent national problem?
- **Quality and makeup of project team**
  - Does the team (and particularly the PI) have a historical track record for delivering products and solutions? Does the team consist of respected and leading members of their community in part due to their scientific output? Do all team members add value and/or are any obvious collaborations missing? Will the team size present project management challenges (i.e., too large to function efficiently or too small to execute)?
- **Technical plan**
  - Does the annual progression in capability development, as laid out with milestone targets and progression problems that march toward the challenge problem target, reflect an actionable and well thought-out plan? Are the milestones SMART (specific, measurable, attainable, relevant, and timely)?
- **Need for capable exascale systems**
  - Is the challenge problem uniquely and solely addressed by efficient utilization of an exascale system and accompanying resources?
- **Technology Impact**
  - Does the project have a likelihood of generating modeling and simulation products, technologies, and solutions outside of its domain of focus that are useful for other researchers, users, customers, and/or stakeholders?

# ECP Applications Deliver Broad Coverage of Strategic Pillars

Initial (FY16) selections consist of 15 application projects + 6 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)

## 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



# 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

# 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

# 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													

# Application Co-Design (CD)

## Scope of development plans requested

- Objective
  - CD technologies outputted, impact/urgency, software and application customers and users
- Development plan
  - CD mission/scope, proposed efforts in concurrent algorithms, data architectures, software technologies, hardware technologies
- Application and workflow integration
  - CD plans for integration of products into specific applications and the direct benefit envisioned from this integration
  - CD plans to improve workflows, either by integration of its products into the workflow or by tangible change of the overall workflow itself
- Milestones and expected outcomes
  - Proposed measures for year-by-year evolution of R&D and products
- Project team and management
- Challenges and risks
- Resource allocation

# Application Co-Design (CD)

## Review criteria

- ***Explicitly targets application requirements, with a functionality that helps to enable applications to more efficiently utilize exascale systems***
  - Does the proposed CD Center plan to develop products and technologies that crosscut multiple applications areas?
- ***Quality and makeup of project team***
  - Does the team have a track record for delivering high-quality products and solutions to application developers?
- ***Addresses critical technological needs of exascale applications or significantly increases the potential performance of applications***
- ***Technical Feasibility and Integration***
  - Will the CD Center team assist applications in effecting needed changes to make use of proposed software innovations? Does the CD Center address issues related to the efficient utilization of exascale systems and accompanying resources? Has the CD Center team thought through impacts on applications making use of the software on exascale resources?
- ***Technology Impact***
  - Can the CD Center be expected to generate production-level software products, technologies, and solutions that are useful for basic science, other Federal agencies, and industry?


# Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

## Nuclear Energy (NE)

Accelerate design and commercialization of next-generation small modular reactors\*

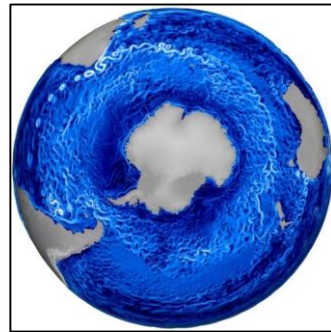
Climate Action Plan;  
SMR licensing support;  
GAIN



## Climate (BER)

Accurate regional impact assessment of climate change\*


Climate Action Plan



## Chemical Science (BES, BER)

Biofuel catalysts design; stress-resistant crops

Climate Action Plan;  
MGI



## Wind Energy (EERE)

Increase efficiency and reduce cost of turbine wind plants sited in complex terrains\*


Climate Action Plan



## Combustion (BES)

Design high-efficiency, low-emission combustion engines and gas turbines\*

2020 greenhouse gas and 2030 carbon emission goals





# Cloud-Resolving Climate Modeling of Earth's Water Cycle

Summary example of an application project development plan

## Exascale challenge problem

- Earth system model (ESM) with throughput needed for multi-decadal coupled high-resolution (~1 km) climate simulations, reducing major systematic errors in precipitation models via explicit treatment of convective storms
- Improve regional impact assessments of climate change on water cycle, e.g., influencing agriculture/energy production
- Integrate cloud-resolving GPU-enabled convective parameterization into ACME ESM using Multiscale Modeling Framework (MMF); refactor key ACME model components for GPU systems
- ACME ESM goal: Fully weather-resolving atmosphere/cloud-resolving superparameterization, eddy-resolving ocean/ice components, throughput (5 SYPD) enabling 10–100 member ensembles of 100 year simulations

## Risks and challenges

- Insufficient LCF allocations
- Obtaining necessary GPU throughput on the cloud-resolving model
- Cloud-resolving convective parameterization via multi-scale modeling framework does not provide expected improvements in water cycle simulation quality
- Global atmospheric model cannot obtain necessary throughput
- MPAS ocean/ice components not amenable to GPU acceleration


PI: Mark Taylor (SNL)

## Applications

- ACME Earth system model: ACME-Atmosphere
- MPAS (Model Prediction Across Scales)-Ocean (ocean)
- MPAS-Seaice (sea ice); MPAS-Landice (land ice)
- SAM (System for Atmospheric Modeling)

## Software technologies

- Fortran, C++, MPI, OpenMP, OpenACC
- Kokkos, Legion
- PIO, Trilinos, PETSc
- ESGF, Globus Online, AKUNA framework



## Development Plan

- **Y1:** Demonstrate ACME-MMF model for Atmospheric Model Intercomparison Project configuration; complete 5 year ACME-MMF simulation with active atmosphere and land components at low resolution and ACME atmosphere diagnostics/ metrics
- **Y2:** Demonstrate ACME-MMF model with active atmosphere, land, ocean and ice; complete 40 year simulation with ACME coupled group water cycle diagnostics/metrics
- **Y3:** Document GPU speedup in performance-critical components: Atmosphere, Ocean and Ice; compare SYPD with and without using the GPU
- **Y4:** ACME-MMF configuration integrated ACME model; document highest resolution able to deliver 5 SYPD; complete 3 member ensemble of 40 year simulations with all active components (atmosphere, ocean, land, ice) with ACME coupled group diagnostics/metrics


# 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\*

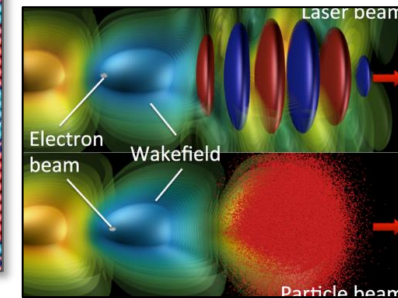
Climate Action Plan; MGI; Light Water Reactor Sustainability; ITER; Stockpile Stewardship Program



## 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



# Virtual Flight Testing for Hypersonic Re-Entry Vehicles\*

## Exascale Challenge Problem

- Virtual flight test simulations of re-entry vehicles from bus separation (exo-atmospheric) to target for normal and hostile environments
- DSMC-based simulation of exo-atmospheric flight regime with hand-off to continuum Navier-Stokes at appropriate altitude
- Time-accurate wall-modeled LES of high Reynolds number (100k-10M) hypersonic gas dynamics
- Fully-coupled simulation of re-entry vehicle ablator/thermal (shape change, ablation products blowing) and structural dynamic (random vibration) response
- DNS and DSMC enhanced reacting gas models and turbulence models via a-priori and on-the-fly model parameter calculations
- Embedded sensitivity analysis, uncertainty quantification and optimization

## Applications & S/W Technologies

### Applications

- SPARC (continuum compressible Navier-Stokes, hypersonic gas dynamics)
- SPARTA (direct simulation Monte-Carlo, rarefied gas dynamics)
- Sierra (Aria – thermal response, Salinas – structural dynamics)

### Software Technologies Cited

- C++
- MPI, Kokkos (OpenMP, Cuda), DARMA (Charm++)
- DataWarehouse, Qthreads, Node-level resource manager
- Trilinos (Belos, MuLue, Tpetra, Sacado, Stokhos, KokkosKernels)
- Percept, IOSS, Exodus, CGNS, netcdf, pnetcdf, HDF5
- In-situ visualization (VTK-M, Catalyst)

## Risks and Challenges

- Scalable solvers for hypersonic gas dynamics (multigrid methods for hyperbolic problems)
- Extreme-scale mesh generation and refinement
- Accurate LES models for hypersonic gas dynamics
- Developing appropriate hypersonic boundary layer/ablator surface interaction models
- Effective task-parallelism and load balancing of heterogeneous physical model workloads

## Development Plan

**FY17:** Demonstrate extreme-scale mesh generation and refinement; continue UQ development efforts; begin research activity on scalable solvers; develop low-dissipation schemes for unsteady turbulent gas dynamics; document KNL performance on ATS-1 (Trinity)

**FY18:** Focus on DARMA task-parallelism implementation; continue physics model development and implementation for hypersonic turbulent flows

**FY19:** Focus on multi-physics coupling and simulation development, including workflows; continue DARMA development activities; document GPU performance on ATS-2 (Sierra)

**FY20:** Full-physics (SPARTA-SPARC coupling, unsteady hypersonic turbulent flows, ablator & structural response coupling) demonstrations with UQ and optimization; document performance

# LANL ASC Advanced Technology Development and Mitigation: Next-Generation Code project (NGC)\*

## Exascale Challenge Problem

- Multi-physics simulations of systems using advanced material modeling for extreme conditions supporting experimental programs at MaRIE
- Multi-physics simulations of high-energy density physics (HEDP) in support of inertial confinement fusion (ICF) experimental programs at NIF
- Routine 3D simulation capabilities to address a variety of new mission spaces of interest to the NNSA complex
- Develop abstraction layer (FleCSI) to separate physics method expression from underlying data and execution model implementations
- Demonstrate the use of advanced programming systems such as Legion for scalable parallel multi-physics, multi-scale code development


## Applications & S/W Technologies

### Applications

- NGC and ASC IC modernization

### Software Technologies Cited

- Legion, MPI
- Kokkos, Thrust, CUDA, OpenMP
- C++17, LLVM/Clang, Python, Lua
- Trilinos, HYPRE
- ParaView, VTK-m, HDF5, Portage, FleCSI, Ingen



## Risks and Challenges

- Immaturity of advanced programming systems such as Legion
- Performance impact of FleCSI abstraction layer may be too great in the context of dynamic multi-physics problem
- Serial nature of existing operator split may impact scalability at exascale and beyond
- Integration of advanced material models in modern unstructured hydrodynamics codes is a research topic
- Scalable storage in support of routine 3D simulations of sufficient resolution is unproven

## Development Plan

- Y1:** Release version 1.0 of a production toolkit for multi-physics application development on advanced architectures. Numerical physics packages that operate atop this foundational toolkit will be employed in a low-energy density multi-physics demonstration problem.
- Y2:** Toolkit release version 2.0. Demonstration of a high-energy density multi-physics problem.
- Y3:** Toolkit release version 3.0. Workflow integration in preparation for Y4 goal.
- Y4:** ASC L1 milestone demonstrating problem of programmatic relevance on ASC Advanced Technology Systems

# 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, 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 Requirements

## Advanced Manufacturing

- **Programming Models and Runtimes**

1. Fortran, C++/C++17, Python, MPI, OpenMP, OpenACC, CUDA, Kokkos, Raja, Boost
2. Legion/Regent, Charm++
3. other asynchronous, task-parallel, programming/execution models and runtime systems

- **Tools**

1. git, CMake, CDash, GitLab
2. Docker, Jira, Travis, PAPI, Oxbow

- **Mathematical Libraries, Scientific Libraries, Frameworks**

1. BLAS/PBLAS, Trilios, PETSc, LAPACK/ScaLAPACK, Hypr, DTK, Chaco, ParMetis, WSMP (direct solver from IBM)
2. HPGMG, MeLu (actually part of Trilinos – replacement for ML), MAGMA, Dakota, SuperLU, AMP

**Requirements Ranking**

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

# Software Technology Requirements

## Advanced Manufacturing

- **Data Management and Workflows**
  1. HDF, netCDF, Exodus
  2. ADIOS
- **Data Analytics and Visualization**
  1. VisIt, ParaView, VTK
- **System Software**

### Requirements Ranking

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

# Merits of Effective Projectization of R&D

- Milestone-based management imparts a sense of urgency & enforces accountability (if you have the “right” milestones)
- Drives setting and adhering to performance metrics that are regularly measured
- Can actually improve breadth and depth of science output
- Forces communication when it’s needed (e.g., when things are not going well)
- Rewards teaming to achieve project goals
- Helps to mentor and train next generation leaders
- Brings helpful process into potentially chaotic situations
- Forces decision points before it’s too late
- Requires active risk management when it’s often an oversight



# 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)

# Next Steps

- Solicit exascale application ideas and project plans from academia, industry
  - Likely target specific application areas representing gaps in current project suite relative to ECP objectives
- Increase application focus on national security and data science problems
- Accelerate current application projects through startup and into a productive cadence (projectization)
- Identify application team “weak spots” for targeted training and productivity improvement
- Publish initial version of exascale application requirements
- Select projects for co-design of key application motifs
- Continue to engage community and stakeholder agencies/programs on priorities and impact
- Proactively engage data science & AI application community (and commercial / open source technologies)
- Develop a sense of community!