The Exascale Computing Project

Paul Messina, ECP Director

November 3, 2016

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

What is the Exascale Computing Project?

- Who in this room has heard of the Exascale Computing Project?

- When we say the Exascale Computing Project – what comes to mind?
  - Hardware / systems / platforms?
  - Software / software stack?
  - Applications?

If you were thinking ‘all the above’ – you were right.

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 to 100 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.

The Role of ECP within NSCI

- DOE is a lead agency within NSCI, along with DoD and NSF
- In particular: DOE SC and NNSA will execute a joint effort on advanced simulation through a **capable exascale** computing program emphasizing sustained performance on relevant applications
  - This is ECP's role
- Deployment agencies: NASA, FBI, NIH, DHS, NOAA

Approach to executing that DOE role in NSCI

- Starting this year, the Exascale Computing Project (ECP) was initiated as a DOE-SC/NNSA-ASC partnership, using DOE's formal project management processes
- The ECP is a project led by DOE laboratories and executed in collaboration with academia and industry
- The ECP leadership team has staff from six U.S. DOE labs
  - Staff from most of the 17 DOE national laboratories will take part in the project
- The ECP collaborates with the facilities that operate DOE's most powerful computers

From Giga to Exa, via Tera & Peta*

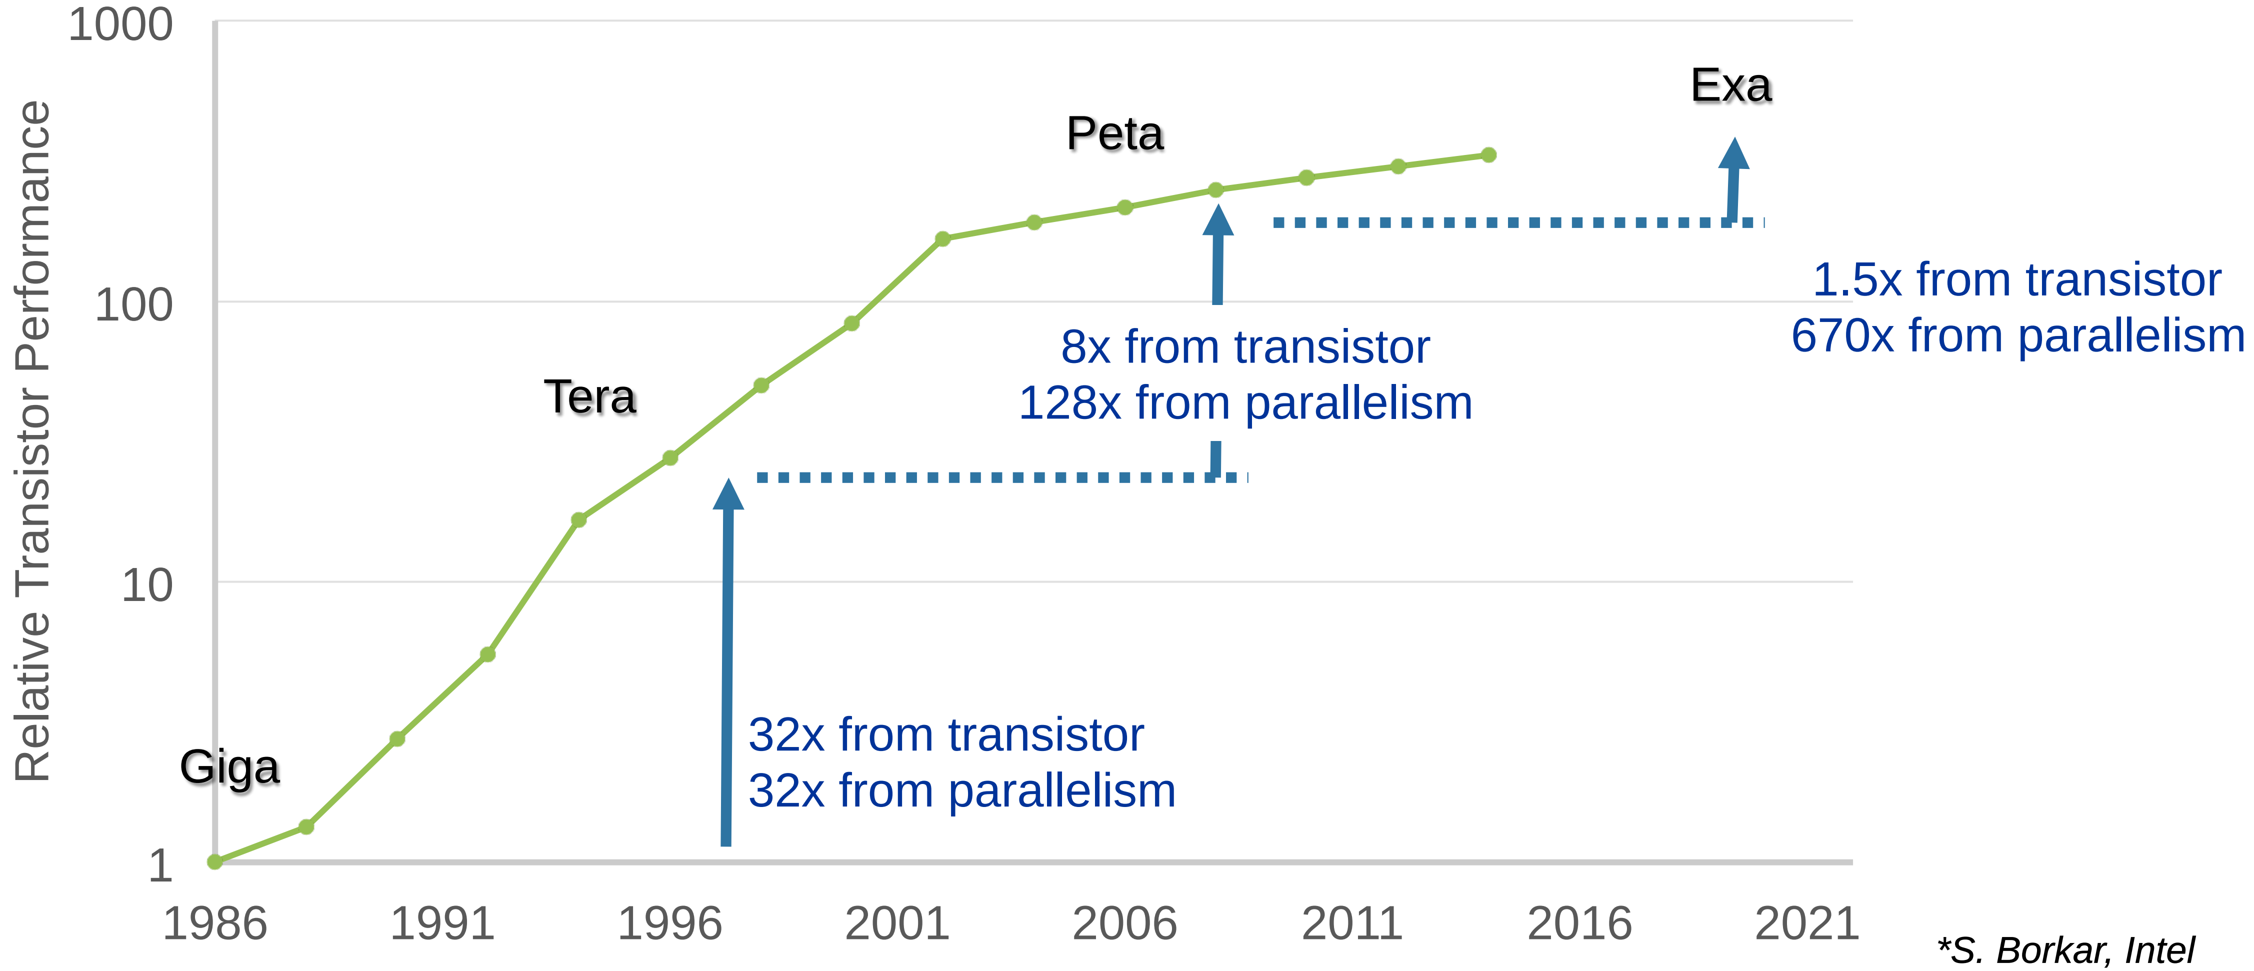

*S. Borkar, Intel

Performance from parallelism

Basic Energy Sciences Advisory Committee Briefing 2.11.2016

Achieving *capable* exascale computing

- 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

Four key challenges that must be addressed to achieve exascale

- Parallelism
- Memory and Storage
- Reliability
- Energy Consumption

Exascale Computing Project Goals

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

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 software stack	Hardware technology elements	Integrated exascale supercomputers

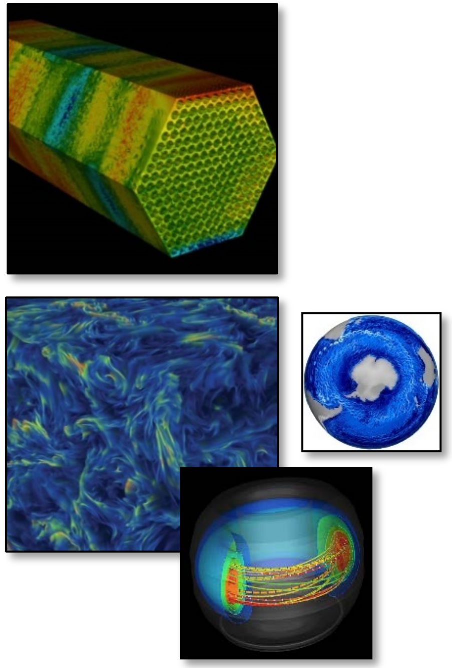

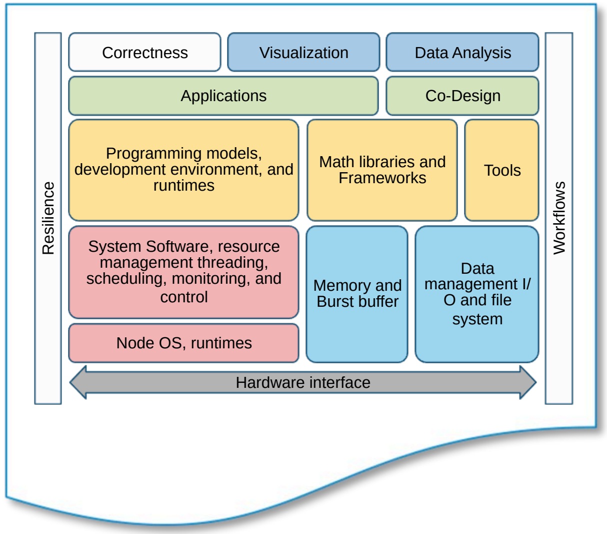

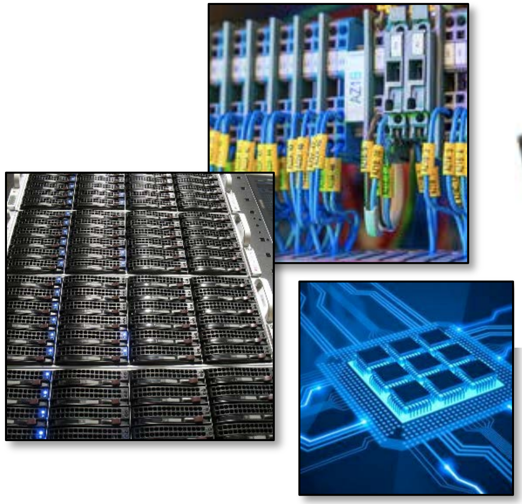

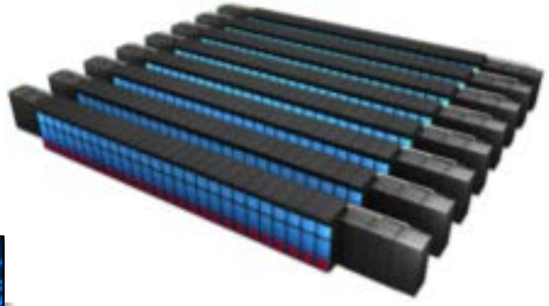

Applications Development activities

- Fund applications development teams
  - Each aiming at capability and specific challenge problems
  - Following software engineering practices
  - Tasked to provide software and hardware requirements
  - Execute milestones jointly with software activities
- Establish co-design centers for commonly used methods
    - E.g., Adaptive Mesh Refinement, Particle-in-Cell
- Developer training

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

- Accelerator physics
- Aerospace
- Astrophysics
- Bioinformatics
- Chemical science
- Climate
- Combustion
- Cosmology
- Energy efficiency
- Energy storage
- Fossil energy
- Geoscience
- Grid modernization
- High energy density physics
- Inertial fusion energy
- Magnetic fusion energy
- Weather
- Manufacturing
- Materials science
- National security
- Neuroscience
- Nuclear energy
- Nuclear physics
- Particle physics
- Plasma Physics
- Precision medicine
- Renewable energy

RFI showed potential for broad exascale application impact

DOE Science and Energy Programs: SC (HEP, NP, FES, BES, BER); Energy (EERE, NE, FE, EDER)	DOE NNSA Programs: Defense Programs, Defense Nuclear Nonproliferation, Naval Reactors	Other Agencies: NSF, NOAA, NASA, NIH

ECP EXASCALE COMPUTING PROJECT

ECP Applications Deliver Broad Coverage of Strategic Pillars

Initial (FY16) 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
- Metegenomics 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)

Accelerate design and commercialization of next-generation small modular reactors

Climate Action Plan; SMR licensing support; GAIN

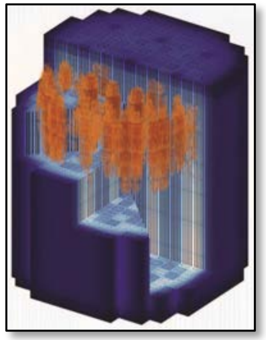

Climate (BER)

Accurate regional impact assessment of climate change

Climate Action Plan

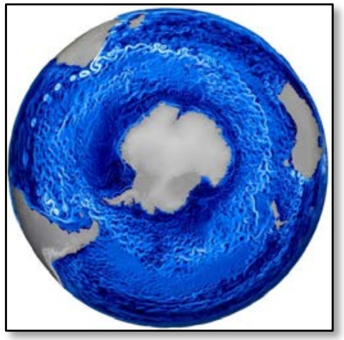

Chemical Science (BES, BER)

Biofuel catalysts design; stress-resistant crops

Climate Action Plan; MGI

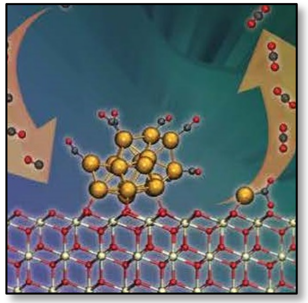

Wind Energy (EERE)

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

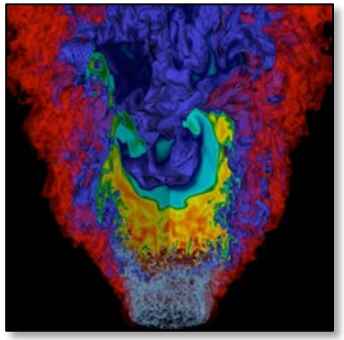

Exascale Applications Will Address National Challenges

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

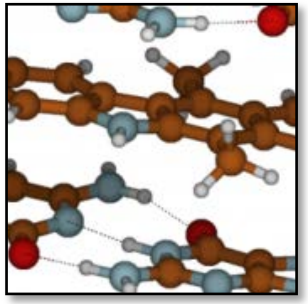

Nuclear Physics (NP)

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

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

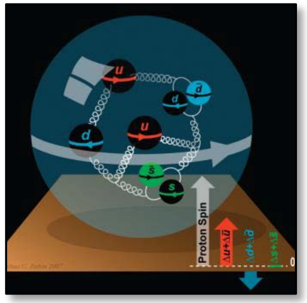

Nuclear Materials (BES, NE, FES)

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

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

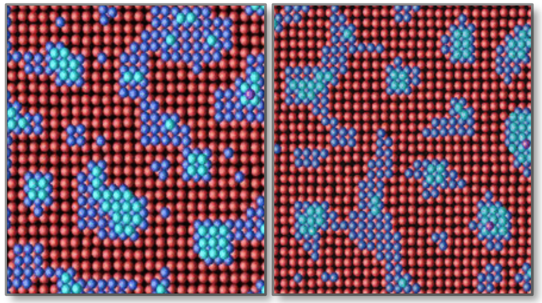

Accelerator Physics (HEP)

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

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

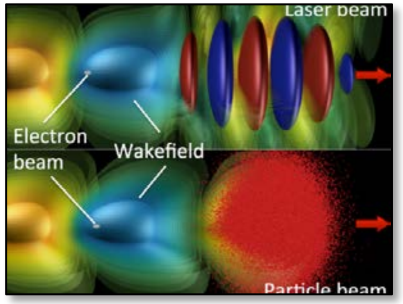

Materials Science (BES)

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

MGI; LCLS-II 2025 Path Forward

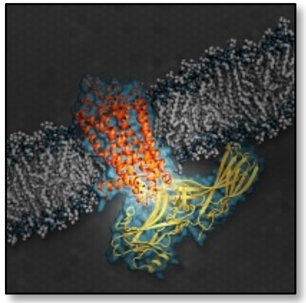

EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

Magnetic Fusion Energy (FES)	Advanced Manufacturing (EERE)	Cosmology (HEP)	Geoscience (BES, BER, EERE, FE, NE)
Predict and guide stable ITER operational performance with an integrated whole device model	Additive manufacturing process design for qualifiable metal components	Cosmological probe of standard model (SM) of particle physics: Inflation, dark matter, dark energy	Safe and efficient use of subsurface for carbon capture and storage, petroleum extraction, geothermal energy, nuclear waste
ITER; fusion experiments: NSTX, DIII-D, Alcator C-Mod	NNMIs; Clean Energy Manufacturing Initiative	Particle Physics Project Prioritization Panel (P5)	EERE Forge; FE NRAP; Energy-Water Nexus; SubTER Crosscut

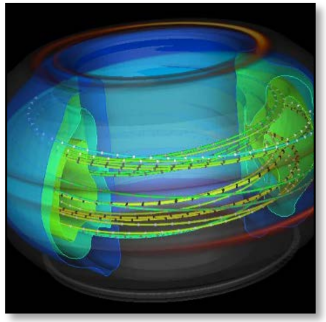

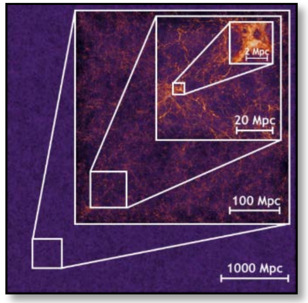

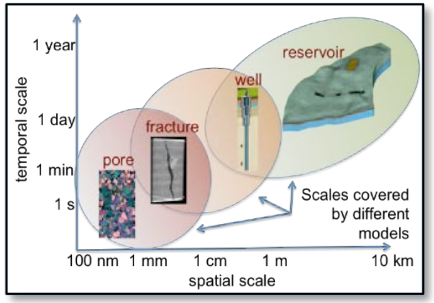

EXASCALE COMPUTING PROJECT

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

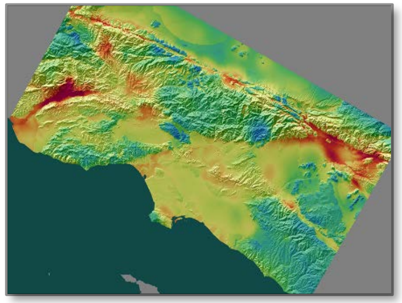

Carbon Capture and Storage (FE)

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

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

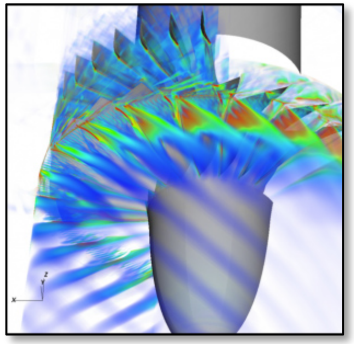

Chemical Science (BES)

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

Climate Action Plan; SunShot Initiative; MGI

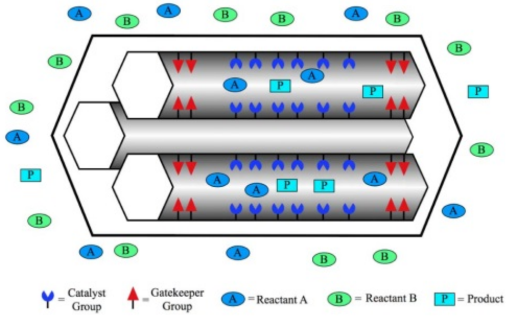

Urban Systems Science (EERE)

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

Energy-Water Nexus; Smart Cities Initiative

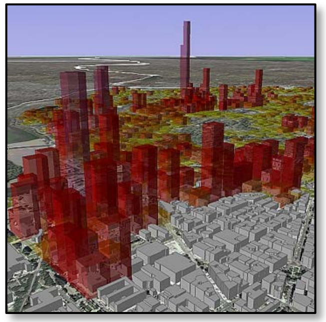

ECP EXASCALE COMPUTING PROJECT

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

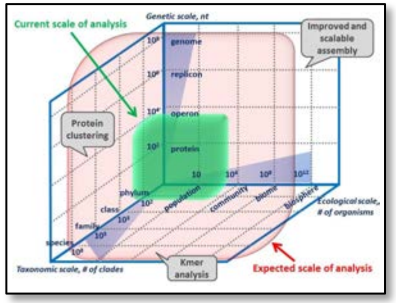

EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Summary of current DOE NNSA application development projects

Stockpile Stewardship

Gaps and Opportunities

- Complete understanding of thermonuclear boost
- Resolution of important length scales with appropriate fidelity

Simulation Challenge Problems

- 3D boost simulations with multiple coupled physical processes at unprecedented resolution
- Detailed highly resolved 3D nuclear safety simulations
- UQ performed in 3D at lower resolution with sub-grid models to capture unresolved physics

Prospective Outcomes and Impact

- Simulation of appropriately complex material at engineering scale through formal and rigorous validation of sub-grid models
- Improved interpretation and understanding of nuclear test data
- High-confidence predictions of thermonuclear boost less dependent upon 2D calibrations

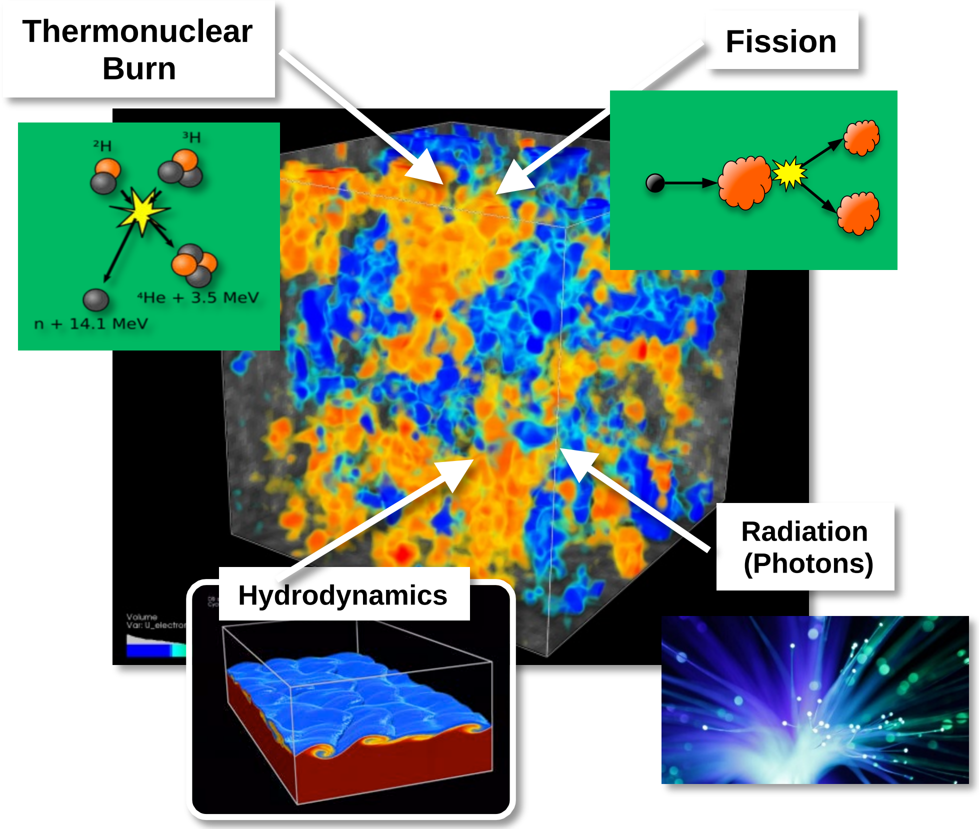

Exascale Applications Will Address National Challenges

Summary of current Other Agency application development projects

Precision Medicine for Cancer (NIH)

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

Precision Medicine in Oncology; Cancer Moonshot

Exascale Predictive Wind Plant Flow Physics Modeling*

Exascale Challenge Problem

- A key challenge to wide-scale deployment of wind energy, without subsidy, in the utility grid is predicting and minimizing plant-level energy losses. Current methods lack model fidelity and inadequately treat key phenomena.
- Deliver predictive simulation of a wind plant composed of O(100) multi-MW wind turbines sited within 10km x 10km area, with complex terrain (O(10e11 grid points).
- Predictive physics-based high-fidelity models validated with target experiments, provide fundamental understanding of wind plant flow physics, and will drive blade, turbine, and wind plant design innovation.
- This work will play a vital role in addressing urgent national need to dramatically increase the percentage of electricity produced from wind power, without subsidy.

Applications & S/W Technologies

Applications

- Nalu, FAST

Software Technologies Cited

- C++, MPI, OpenMP (via Kokkos), CUDA (via Kokkos)
- Trilinos (Tpetra), Muelu, Sierra Toolkit (STK), Kokkos
- Spack, Docker
- DHARMA (Distributed asynchronous Adaptive Resilient Management of Applications)

Risks and Challenges

- Transition to next generation platforms
- Robustness of high-order schemes on turbulence model equations
- Sliding-mesh algorithm scalability
- Kokkos support

Development Plan

Y1: Baseline run for canonical ABL simulation with MPI; single-blade-resolved sim in non-rotating turbulent flow; incorporate Kokkos and demonstrate faster ABL run; demonstrate single-blade-resolved simulation with rotating blades

Y2: Baseline single-blade-resolved capability (SBR) run; demonstrate mixed-order run with overset or sliding mesh algorithm; demonstrate faster SBR run; demonstrate single-turbine blade-resolved simulation

Y3: Demonstrate simulation of several turbines operating in flat terrain

Y4: Demonstrate simulation of O(10) turbines operating in complex terrain

*PI: Steve Hammond (NREL)

ECP WBS

- **Exascale Computing Project** 1.
  - Project Management 1.1
    - Project Planning and Management 1.1.1
    - Project Controls & Risk Management 1.1.2
    - Business Management 1.1.3
    - Procurement Management 1.1.4
    - Information Technology and Quality Management 1.1.5
    - Communications & Outreach 1.1.6
    - Integration 1.1.7
  - Application Development 1.2
    - DOE Science and Energy Apps 1.2.1
    - DOE NNSA Applications 1.2.2
    - Other Agency Applications 1.2.3
    - Developer Training and Productivity 1.2.4
    - Co-Design and Integration 1.2.5
  - Software Technology 1.3
    - Programming Models and Runtimes 1.3.1
    - Tools 1.3.2
    - Mathematical and Scientific Libraries and Frameworks 1.3.3
    - Data Management and Workflows 1.3.4
    - Data Analytics and Visualization 1.3.5
    - System Software 1.3.6
    - Resilience and Integrity 1.3.7
    - Co-Design and Integration 1.3.8
  - Hardware Technology 1.4
    - PathForward Vendor Node and System Design 1.4.1
    - Design Space Evaluation 1.4.2
    - Co-Design and Integration 1.4.3
  - Exascale Systems 1.5
    - Site Preparation 1.5.1
    - NRE 1.5.2
    - Testbeds 1.5.3
    - Base System Expansion 1.5.4
    - Co-Design and Integration 1.5.5

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).*

Application Co-Design (CD)

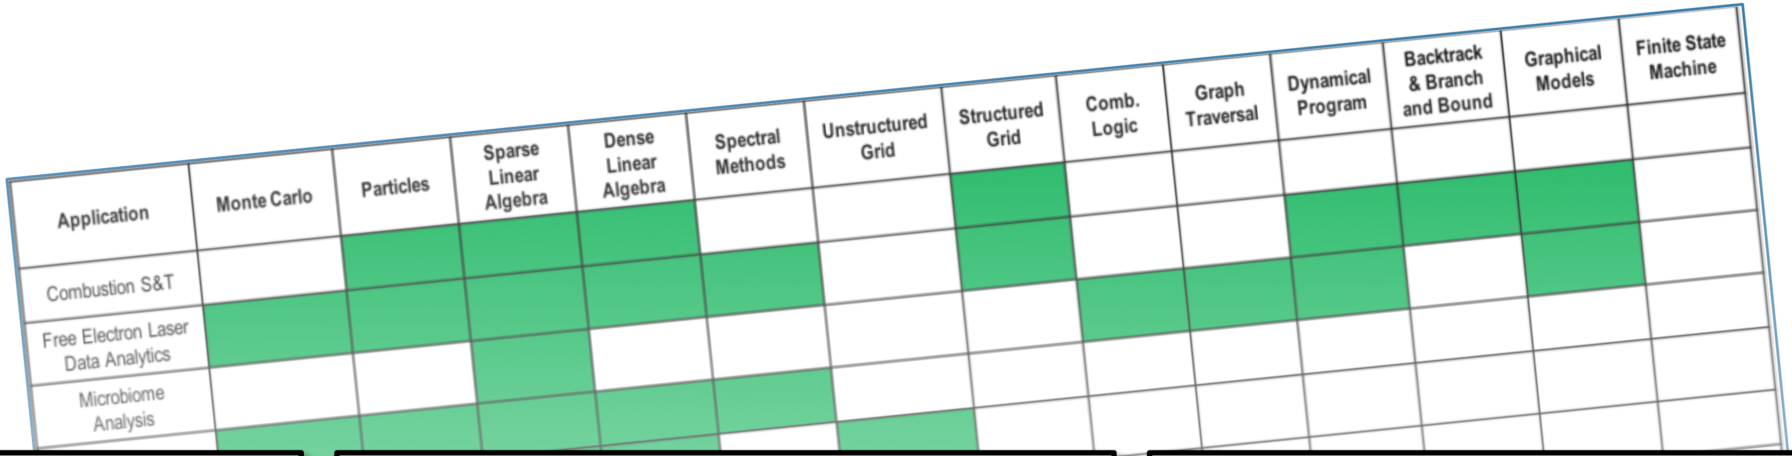

Essential to ensure that applications effectively utilize exascale systems	Executed by several CD Centers focusing on a unique collection of algorithmic motifs invoked by ECP applications	Game-changing mechanism for delivering next-generation community products with broad application impact
• Pulls ST 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	• 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	• Evaluate, deploy, and integrate exascale hardware-savvy software designs and technologies for key crosscutting algorithmic motifs into applications

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		■	■			■							

EXASCALE COMPUTING PROJECT

Initial Co-design Centers

- Co-Design Center for Particle-Based Methods: From Quantum to Classical, Molecular to Cosmological
- CODAR: A Co-Design Center for Online Data Analysis and Reduction at the Exascale
- Center for Efficient Exascale Discretizations (CEED)
- Block-Structured AMR Co-Design Center

ECP EXASCALE COMPUTING PROJECT

Block-Structured AMR Co-Design Center

Center Objectives (Co-Designed Motifs)

- **Motif(s)**: 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
- Hierarchical solution at multiple levels of resolution, with each level of refinement being the union of data containers at that resolution, each of which represents the solution over a logically rectangular subregion of the domain
- AMReX supports conventional representation of field variables on a mesh as well as particle data and embedded boundary representations of complex geometries

Applications and S/W & H/W Technologies

Applications Targeted

- CLASH, WarpX, Pele, Nyx, MFIX-Exa, ChomboCrunch, SW4

Software Technologies Cited

- MPI, OpenMP, OpenACC, PGAS, GASNet, UPC++, C++, Fortran
- BoxLib, Chombo, FLASH, Perilla, ExaSAT, HDF5, HPGMG

Hardware Technologies Addressed

- Inter-node communication (topology, bandwidth, latency), on-node design (number of compute cores, compute heterogeneity, deeper and more non-uniform, on-node memory hierarchies) increase in the relative cost of data movement and synchronization (especially off-node)

Risks and Challenges

- GASNet not available on future machines
- OpenMP not performant on GPU's
- External AMG solver fails to have either good performance or low set-up time
- Prototype hardware not available for performance testing
- Loss of key personnel
- HDF5 I/O not performant
- Resilience issues lead to high node failure rate

Development Plan

Y1: Mesh/particle data structures; MPI+OpenMP for synchronous iterators of mesh/particle data; Characterize communication models & non-EB load balancing; Implement MG solvers & multi-level MG solvers with HPGMG agglomeration strategy; AMReX release (mesh/particle functionality)

Y2: EB data structures & particle containers; MPI+OpenMP for iterators of EB data structures; Asynchronous iterators for mesh/particle data; MPI+OpenMP alternatives (PGAS); EB & non-EB load balancing issues & strategies; Single-level EB MG solver; AMReX release (single-level EB functionality)

Y3: non-EB data structures; Scheduling of EB stencil operations; Iterators & MPI+OpenMP alternatives for GPUs; Load balancing strategies for EB apps; Multi-level, cell-centered EB multigrid solve; Optimized MG solvers; AMReX release (multi-level EB functionality)

Y4: EB data structures; Real-time load balancing for partner ECP apps; Single and multi-level nodal EB multigrid solvers; Full suite of optimized MG solvers; AMReX release (supports all partner ECP apps)

PI: John Bell (LBNL); Institutions: ANL, LBNL, NREL

Conceptual ECP Software Stack

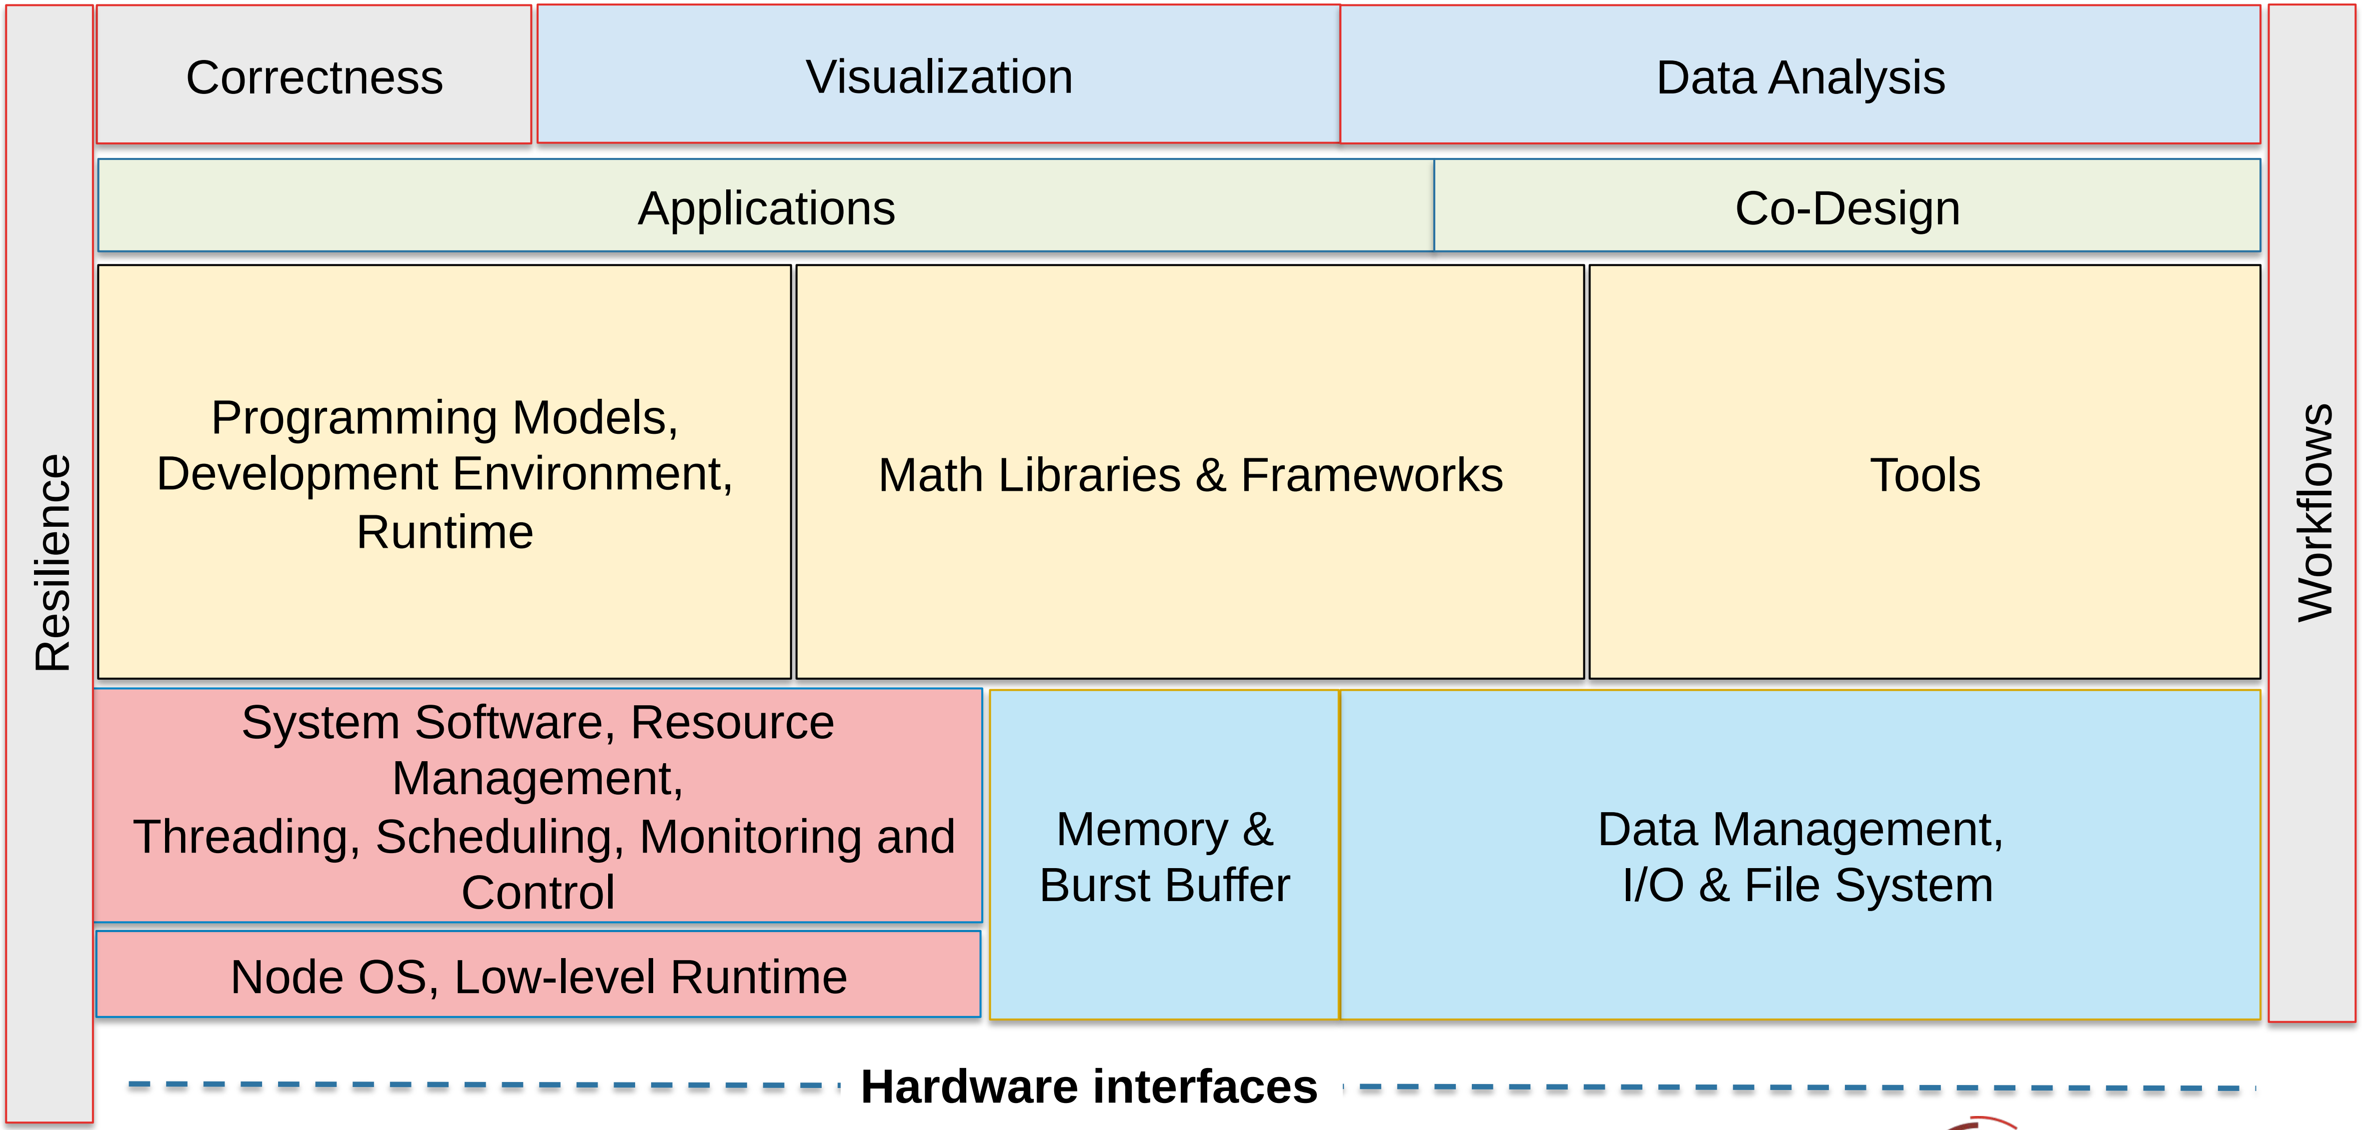

Requirements for Software Technology

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**

No. of ECP Application Proposals a Software is Mentioned in

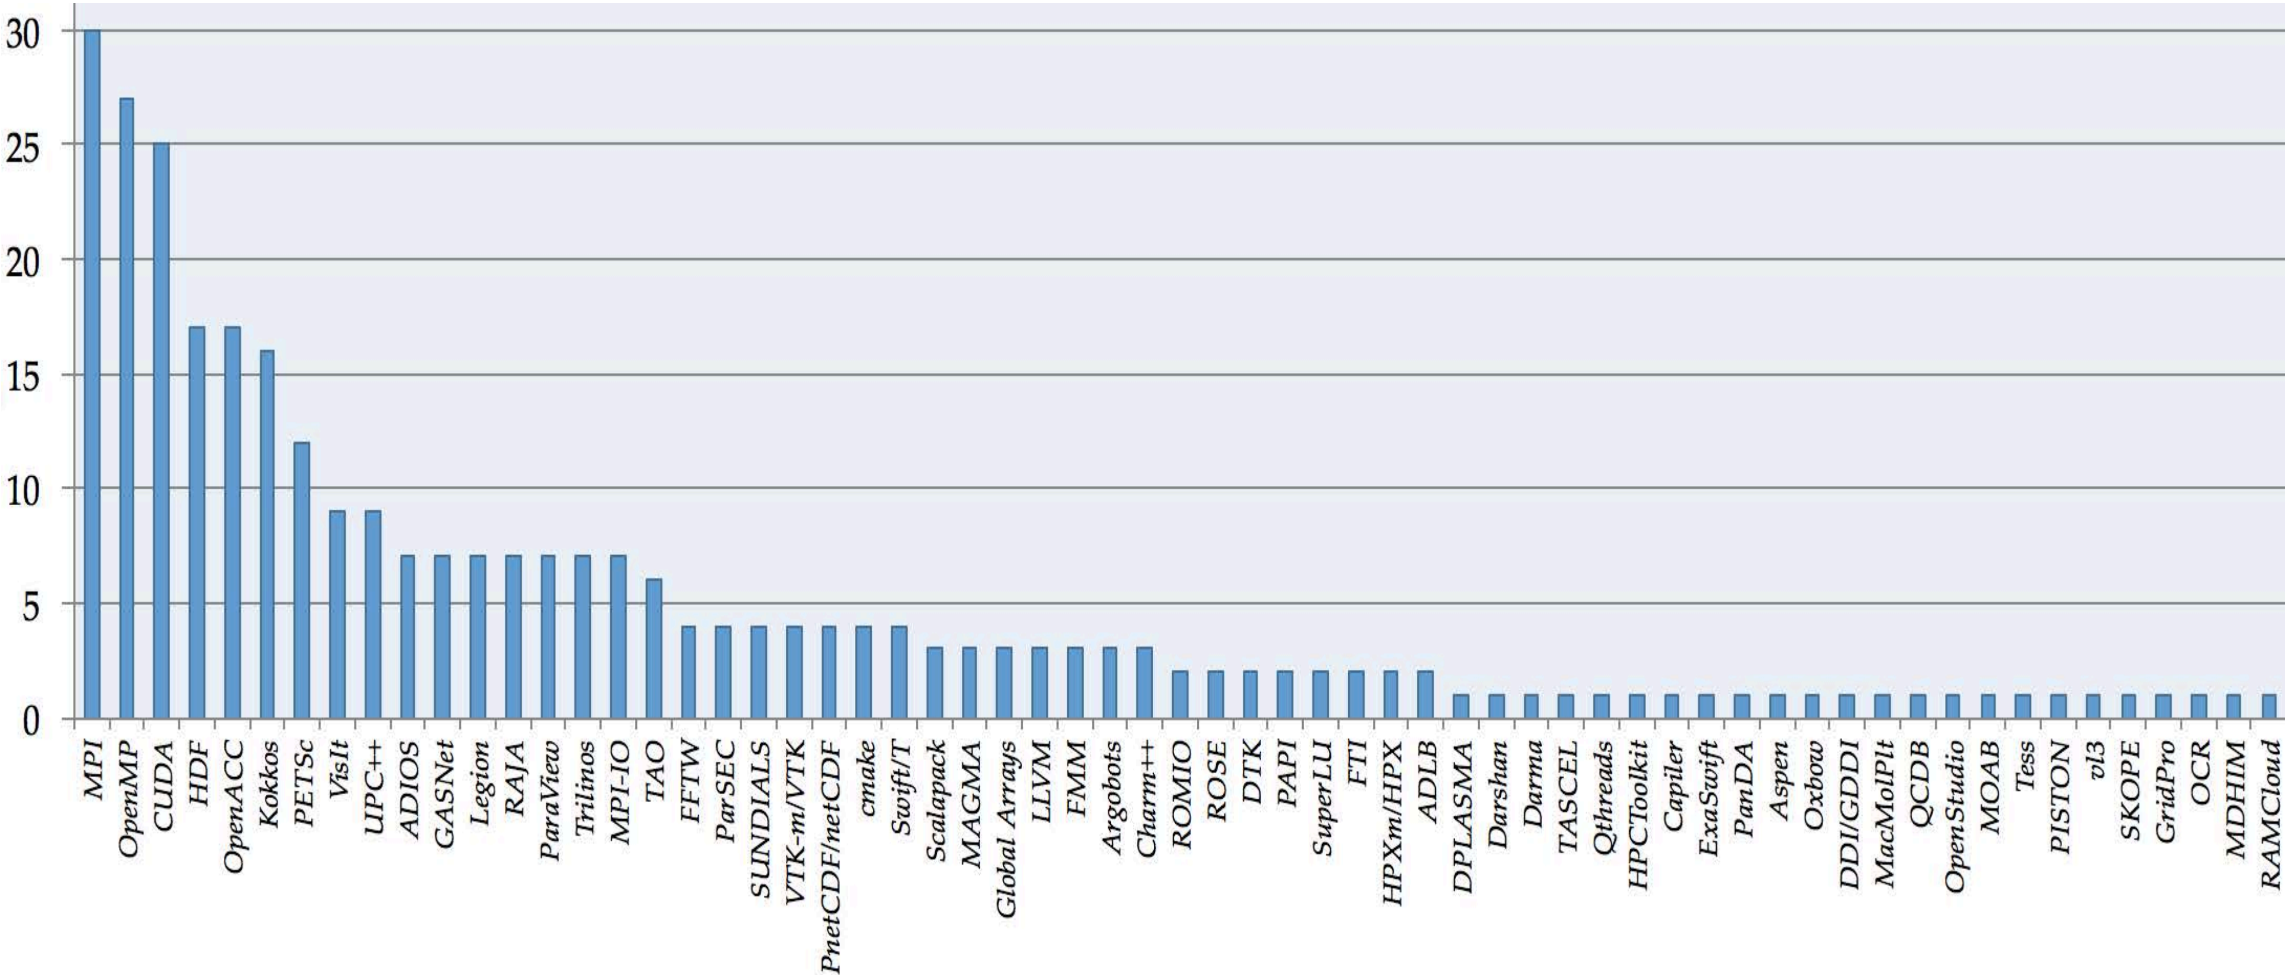

Libraries used at NERSC
(similar data from other facilities)

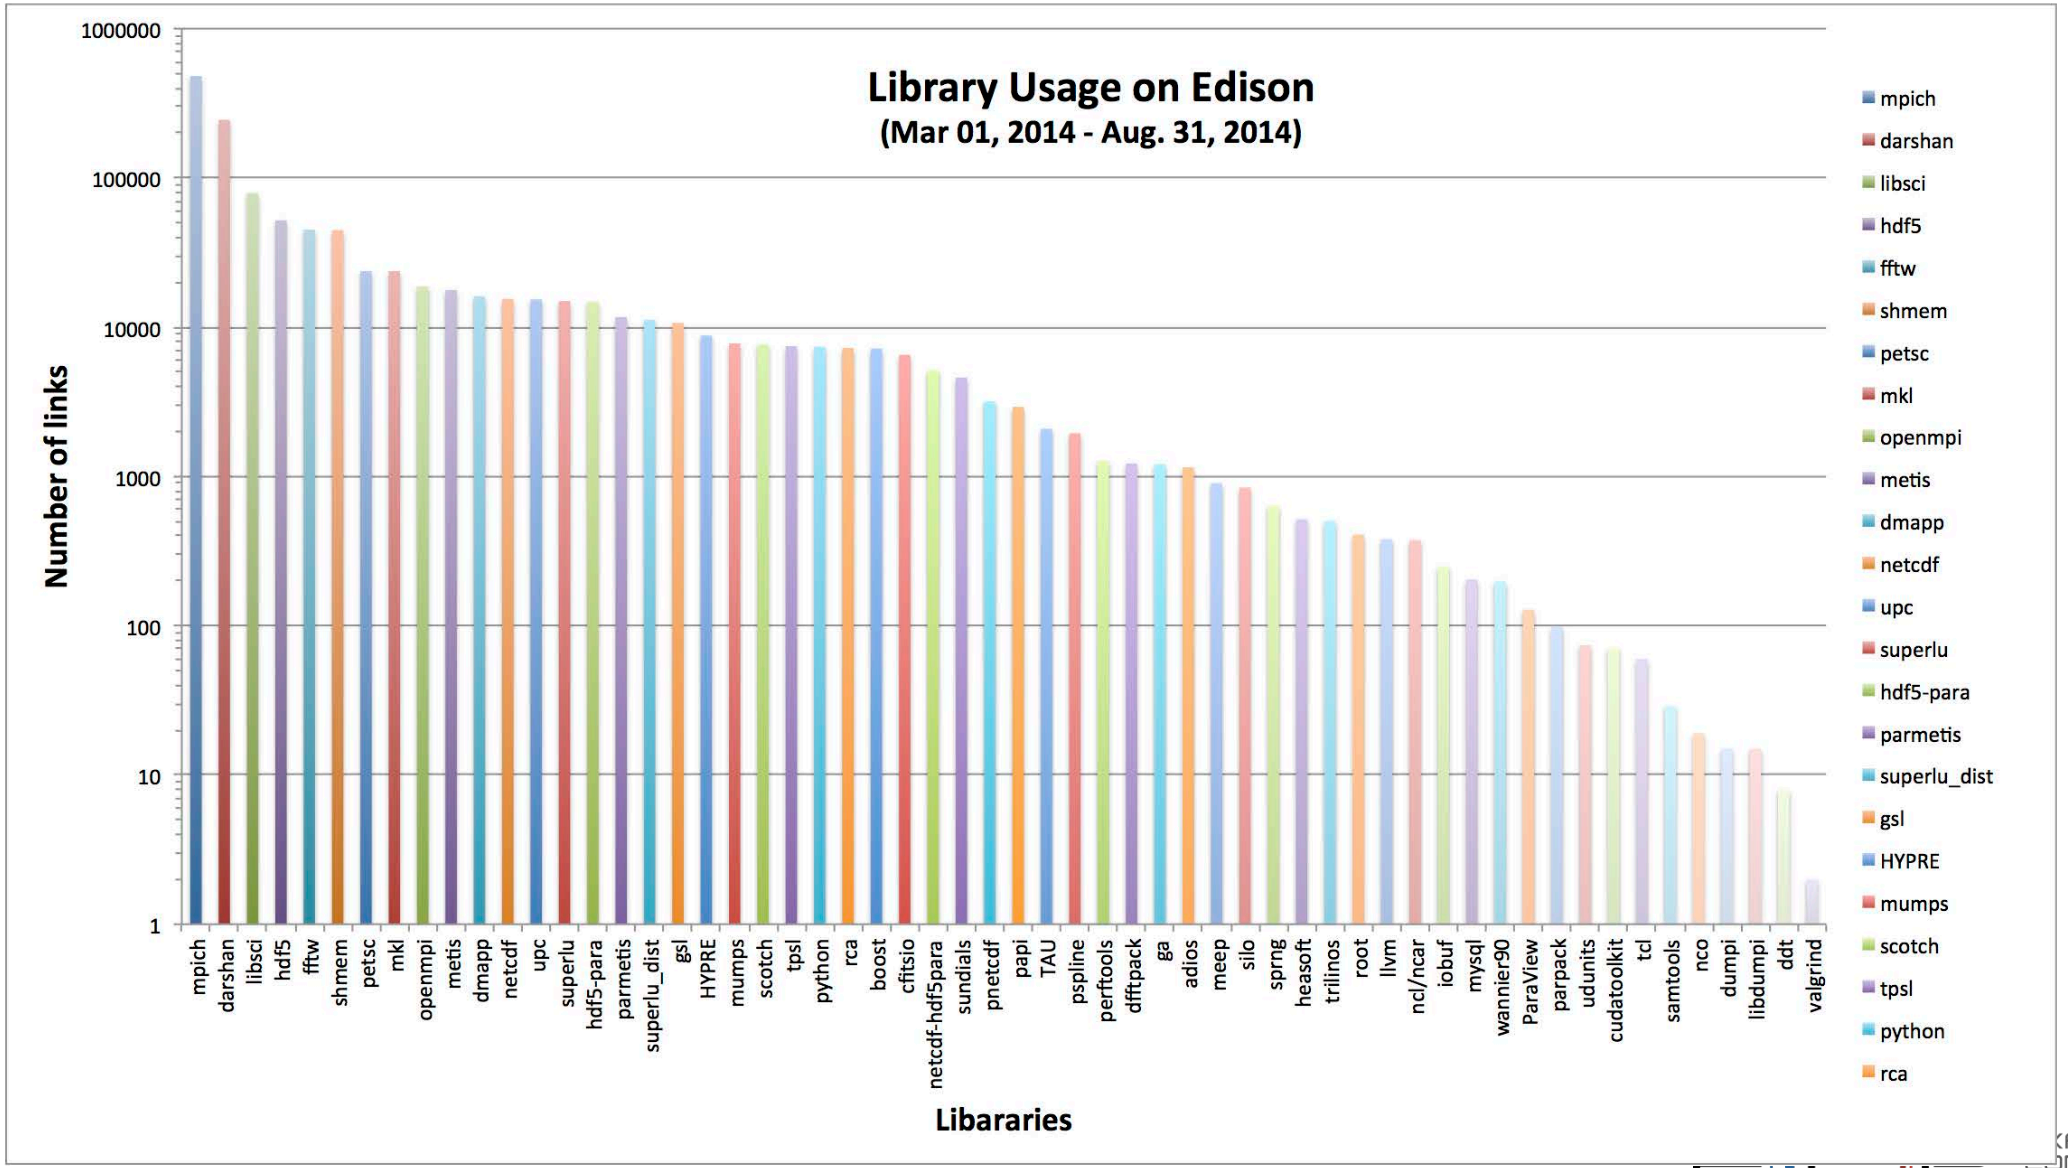

Software Technology Selections Mapped to Software Stack

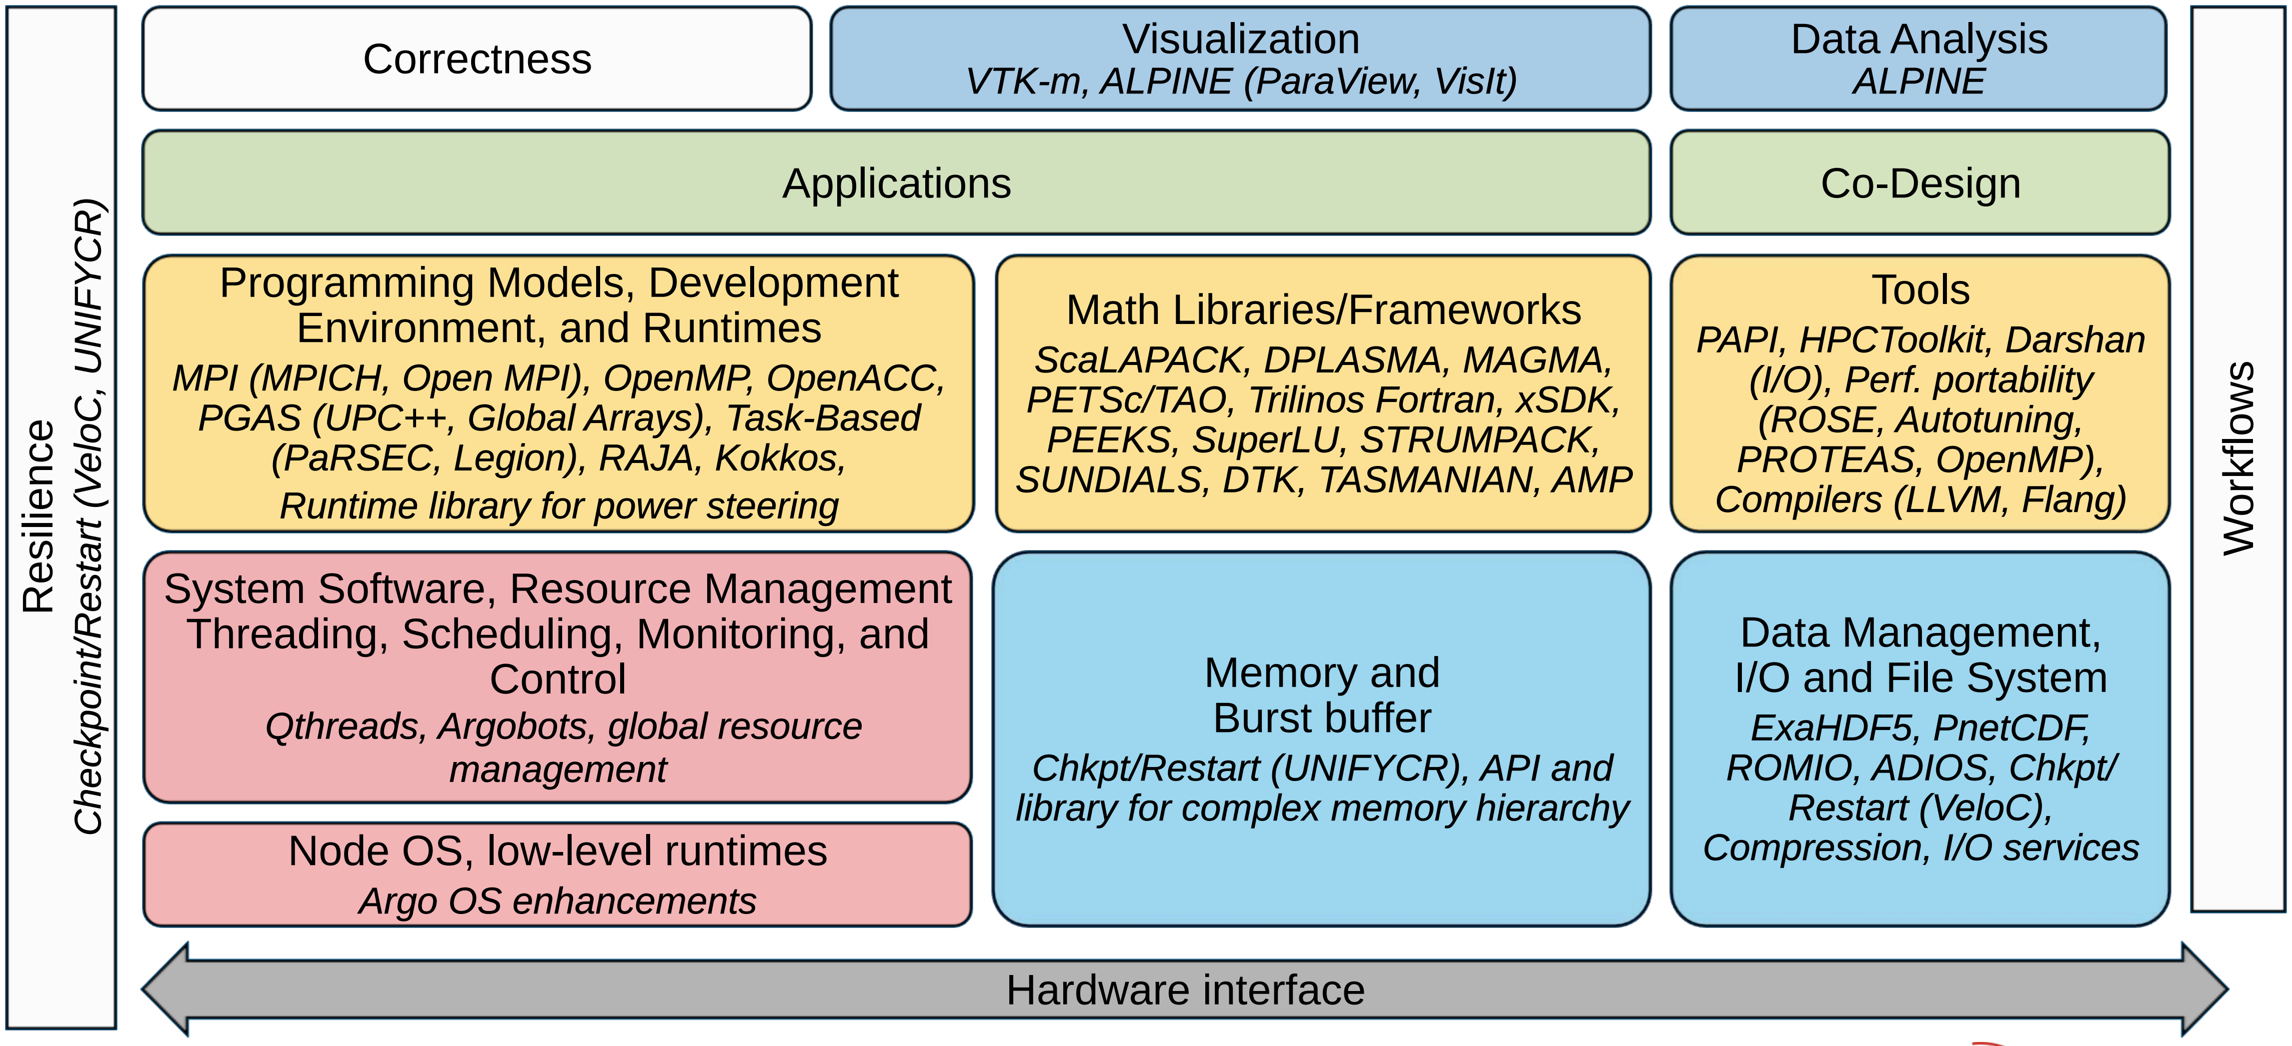

Hardware Technology Activities

- PathForward: support DOE-vendor collaborative R&D activities required to develop exascale systems with at least two diverse architectural features; quote from RFP:
  - PathForward seeks solutions that will improve application performance and developer productivity while maximizing energy efficiency and reliability of exascale systems.
- Design Space Evaluation
  - Apply laboratory architectural analysis capabilities and Abstract Machine Models to PathForward designs to support ECP co-design interactions

Exascale Systems Activities

- Ensure at least two exascale-class systems are accepted no later than 2023 and the systems are diverse, affordable, production-ready, and capable
  - NRE contracts
  - Convey results of ECP R&D to RFP for exascale systems procurement that facilities will issue
- Provide requirements from facilities viewpoint
- Acquire and oversee operation of testbeds for application and software development projects and for hardware investigations

ECP phases

- 2016 – 2019
  - Develop applications, conduct R&D&D on software technologies
    - Use current systems, CORAL systems as testbeds
  - Vendor R&D on node and system designs that are better suited for HPC applications
- 2019
  - **ECP insights are used in formulation of RFP for exascale systems**
  - DOE and NNSA laboratories issue RFP for exascale systems, select offers, award build and NRE contracts
- 2019-2023
  - ECP Applications and software technologies are modified with knowledge of systems
  - Software technologies are “productized”
- 2023-2025
  - Exascale systems are in production, applications and software deal with actual system behavior

EXASCALE COMPUTING PROJECT

ECP status

- 22 application proposals have been selected for funding
- 4 co-design centers have been selected for funding
- 35 software technology proposals have been selected for funding
  - In addition to similar number already underway at NNSA labs
- Responses to PathForward RFP (Hardware Technology R&D by vendors) have been evaluated and proposals selected for funding
  - Negotiations underway, contracts expected to be put in place Q1/2 FY2017

To achieve capable exascale requires a holistic approach

Application Development	Software Technology	Hardware Technology	Exascale Systems
Science and mission applications	Scalable and productive software stack	Hardware technology elements	Integrated exascale supercomputers

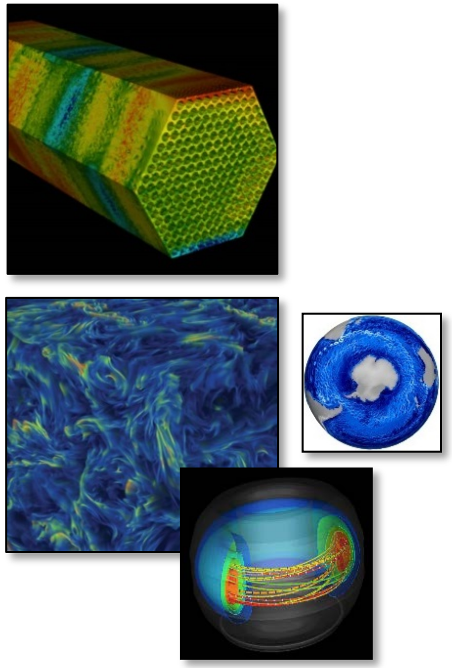

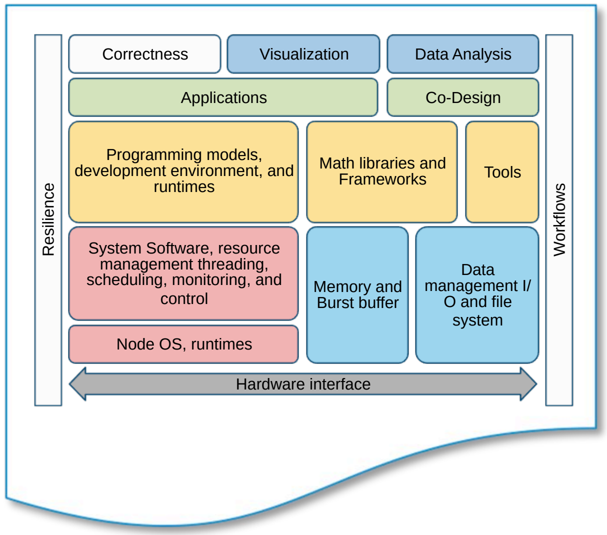

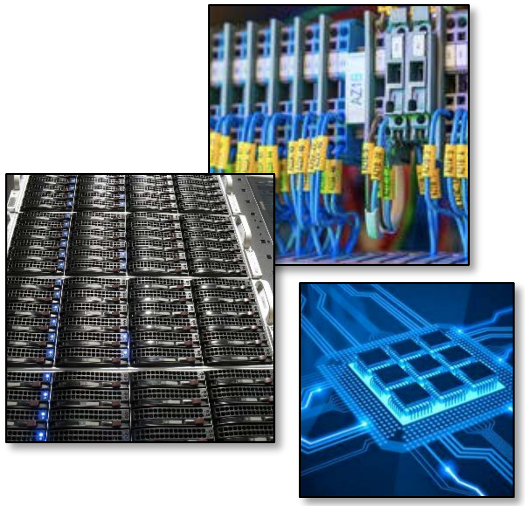

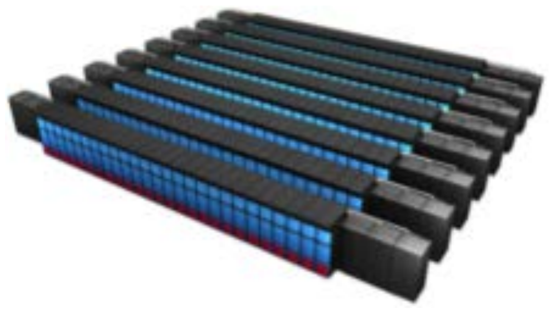

ECP's work encompasses applications, system software, hardware technologies and architectures, and workforce development

EXASCALE COMPUTING PROJECT

Thank you!

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science