# Exascale Computing Project: Software Technology Perspective

**Rajeev Thakur, Argonne National Lab.**
ECP Software Technology Director

Charm++ Workshop
Champaign, IL
April 18, 2017

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

www.ExascaleProject.org

# What is the Exascale Computing Project?

- The ECP is a collaborative effort of two US Department of Energy (DOE) organizations – the Office of Science (DOE-SC) and the National Nuclear Security Administration (NNSA).
- As part of the National Strategic Computing initiative, ECP was established to accelerate delivery of a capable exascale computing system that integrates hardware and software capability to deliver 50 times more performance than the nation's most powerful supercomputers in use today.
- 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.

# Approach to executing that DOE role in NSCI

- Starting last 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

# 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: ≤1/week)
- Includes a software stack that meets the needs of 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

# Four key challenges that must be addressed to achieve exascale

- Parallelism
- Memory and Storage
- Reliability
- Energy Consumption

# 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 |

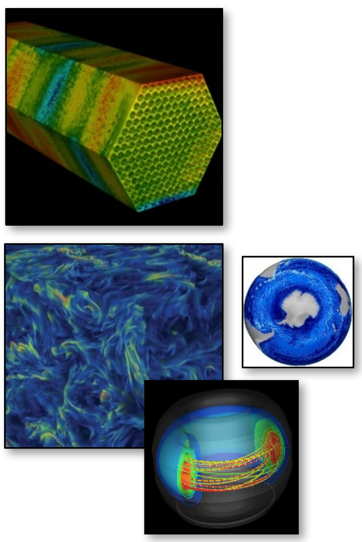

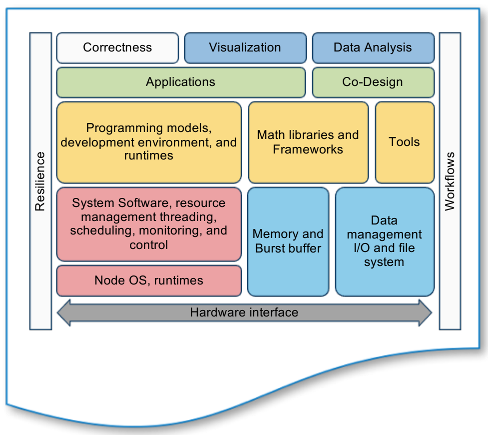


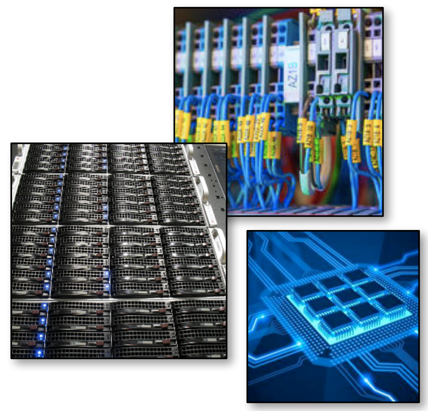

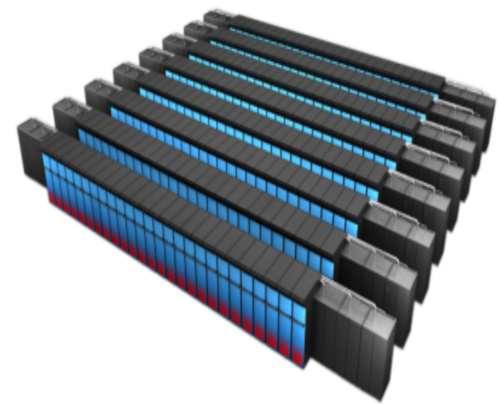



EXASCALE COMPUTING PROJECT

# The ECP Plan of Record

- A 7-year project that follows the **holistic/co-design** approach, that runs through 2023 (including 12 months of schedule contingency)
- Enable an initial exascale system based on **advanced architecture** delivered in 2021
- Enable **capable exascale** systems, based on ECP R&D, delivered in 2022 and deployed in 2023 as part of an NNSA and SC facility upgrades

Acquisition of the exascale systems is outside of the ECP scope, will be carried out by DOE-SC and NNSA-ASC supercomputing facilities



EXASCALE COMPUTING PROJECT

# High-level ECP technical project schedule

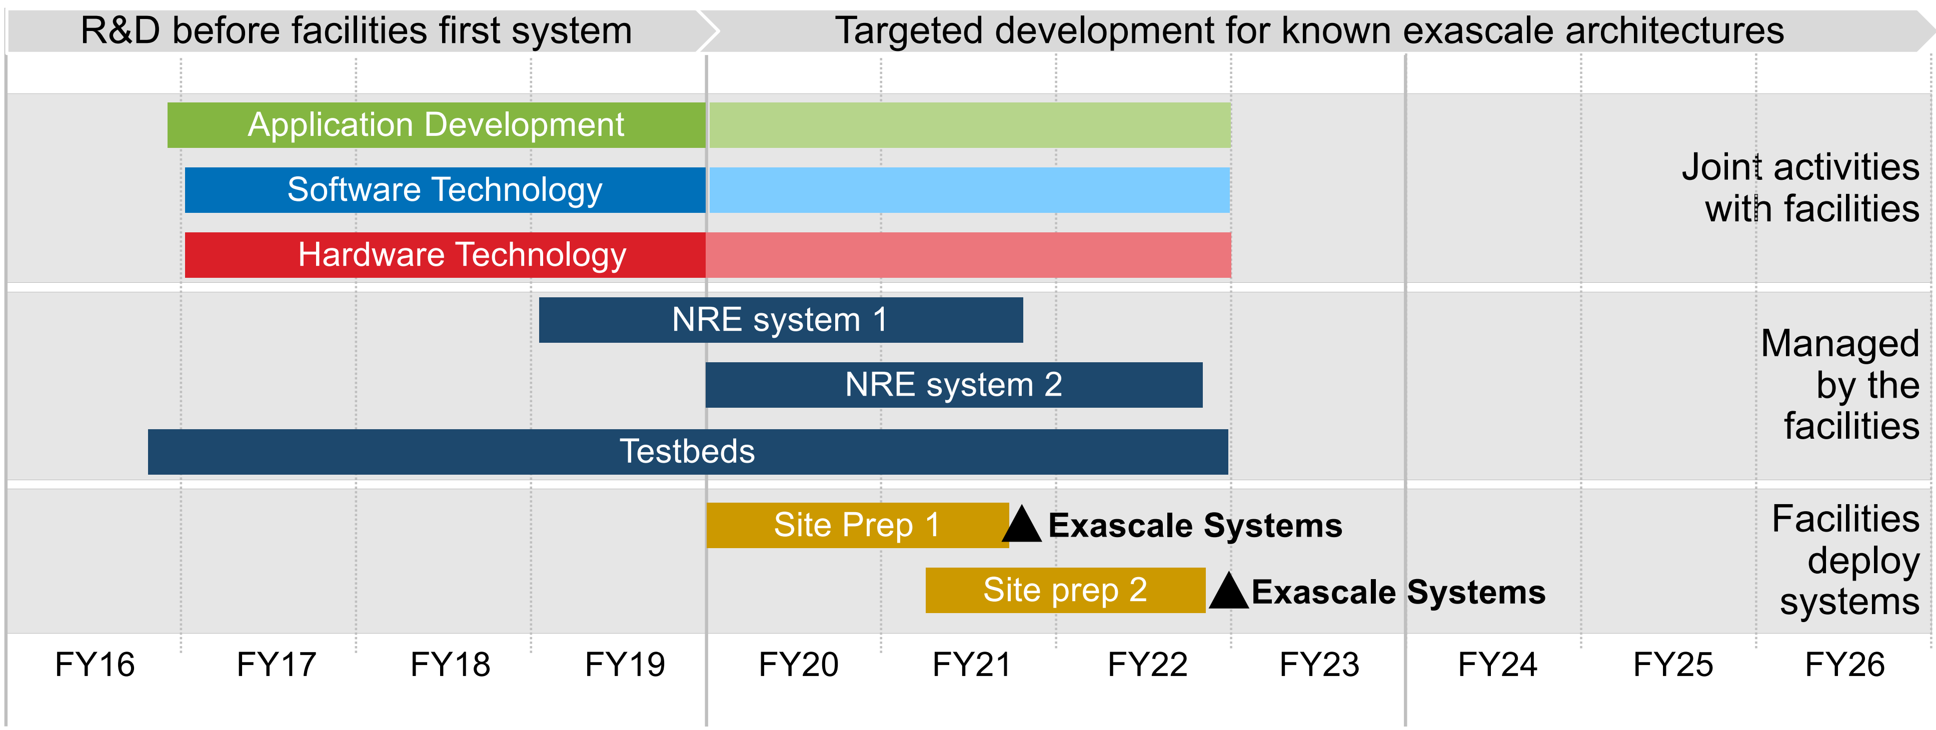

# ECP Projects Status

- 22 application projects have been selected for funding
  - In addition to 4 applications projects already underway at the NNSA labs
- 5 co-design centers have been selected for funding
- 35 software technology projects have been selected for funding
  - In addition to similar number already underway at NNSA labs
- Proposals submitted to PathForward RFP (Hardware Technology R&D by vendors) have been selected for funding
  - Negotiations underway; contracts expected to be signed by May 2017

ECP

EXASCALE COMPUTING PROJECT

DOE LABORATORIES & AGENCY PARTNERS
**22**

PRIVATE SECTOR PARTNERS
**9**

UNIVERSITY RESEARCH PARTNERS
**39**

INDUSTRY COUNCIL MEMBERS
**18**

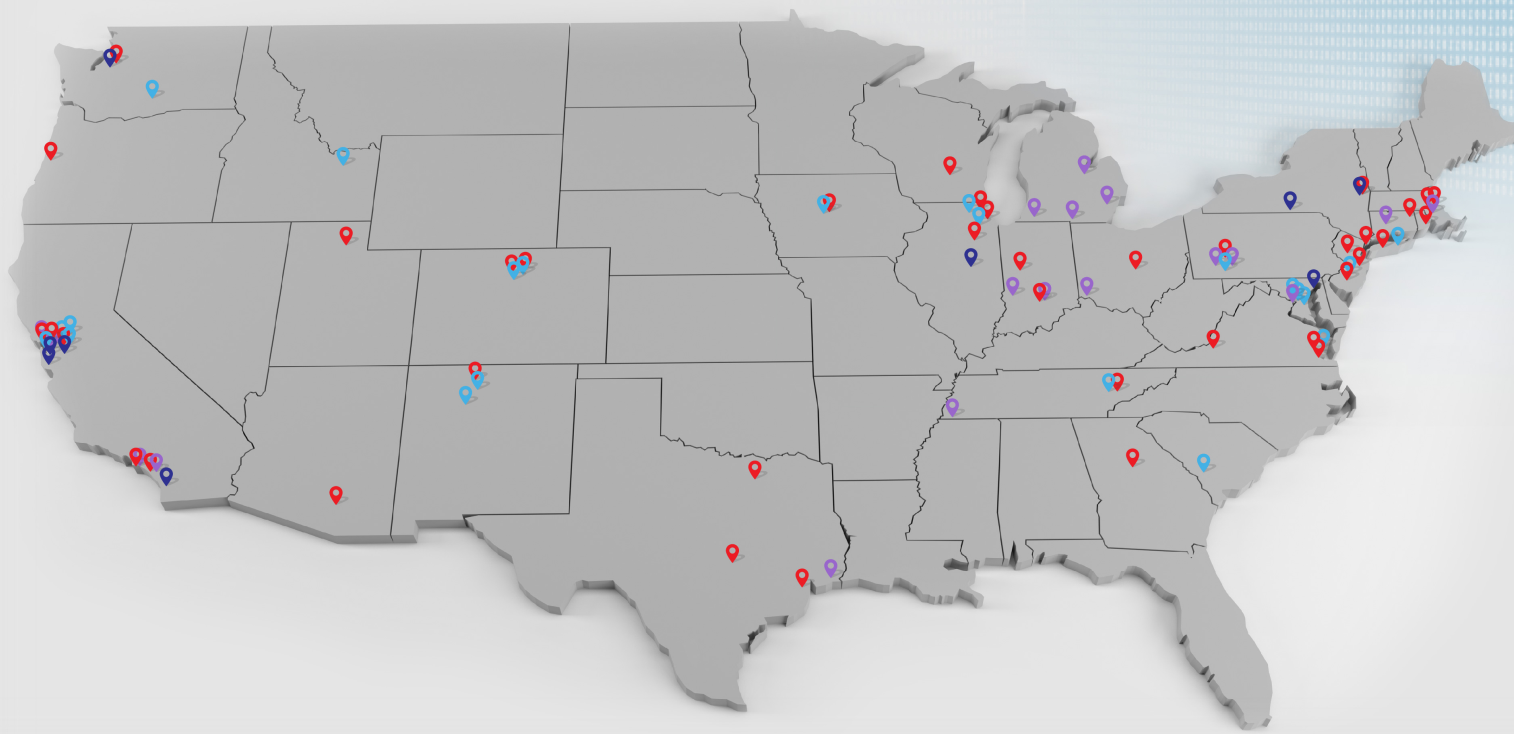

# THE ECP ECOSYSTEM

- 800 Researchers
- 26 Application Development Projects
- 66 Software Development Projects
- 5 Co-Design Centers

U.S. DEPARTMENT OF ENERGY | Office of Science

NNSA National Nuclear Security Administration

# ECP Leadership Team

**Chief Technology Officer**
**Al Geist**, ORNL

**Exascale Computing Project**
**Paul Messina**, Project Director, ANL
**Stephen Lee,** Deputy Project Director, LANL

**Integration Manager**
**Julia White**, ORNL

**Communications Manager**
**Mike Bernhardt,** ORNL

**Project Management**
**Kathlyn Boudwin,**
Director, ORNL

**Application Development**
**Doug Kothe,**
Director, ORNL
**Bert Still,**
Deputy Director, LLNL

**Software Technology**
**Rajeev Thakur,**
Director, ANL
**Pat McCormick,**
Deputy Director, LANL

**Hardware Technology**
**Jim Ang,** Director, SNL
**John Shalf,**
Deputy Director, LBNL

**Exascale Systems**
**Terri Quinn**, Director, LLNL
**Susan Coghlan,**
Deputy Director, ANL

# 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
    - SW PathForward 1.3.9
  - 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
    - PathForward II Vendor Node and System Design 1.4.4
  - Exascale Systems 1.5
    - NRE 1.5.1
    - Testbeds 1.5.2
    - Co-design and Integration 1.5.3

# 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 |

# ECP Software Technology Overview

- Build a comprehensive and coherent software stack that will enable application developers to productively write highly parallel applications that can portably target diverse exascale architectures
- Accomplished by extending current technologies to exascale where possible, performing R&D required to conceive of new approaches where necessary
  - Coordinate with vendor efforts; i.e., develop software other than what is typically done by vendors, develop common interfaces or services
  - Develop and deploy high-quality and robust software products

# Vision and Goals for the ECP Software Stack

| | |
|---|---|
| Deliver and Anticipate | Provide foundational software and infrastructure to applications and facilities necessary for project success in 2021-23, while also pushing to innovate beyond that horizon |
| Collaborate | Encourage and incentivize use of common infrastructure and APIs within the software stack |
| Integrate | Work with vendors to provide a balanced offering between lab/univ developed (open source), vendor-offered (proprietary), and jointly developed solutions |
| Quality | Deploy production-quality software that is easy to build, well tested, documented, and supported |
| Prioritize | Focus on a software stack that addresses the unique requirements of exascale – including extreme scalability, unique requirements of exascale hardware, and performance-critical components |
| Completeness | Perform regular gap analysis and incorporate risk mitigation (including "competing" approaches) in high-risk and broadly impacting areas |

# ECP Requires Strong Integration to Achieve Capable Exascale

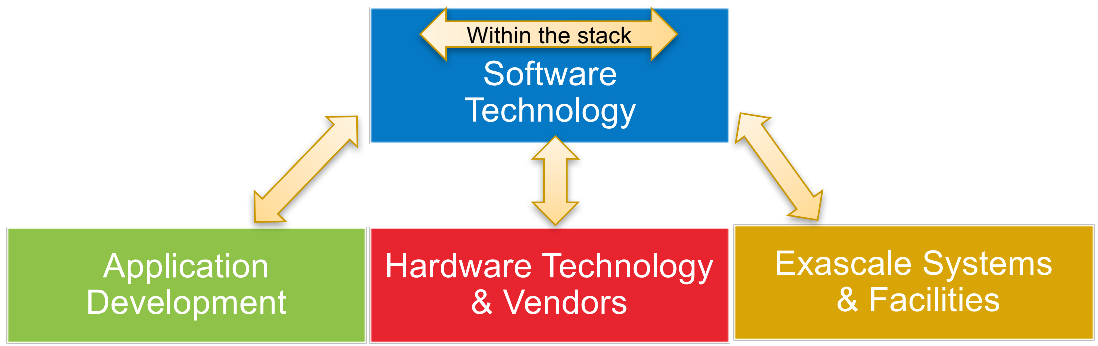


- To achieve a coherent software stack, we must integrate across all the focus areas
  - Understand and respond to the requirements from the apps but also help them understand challenges they may not yet be aware of
  - Understand and repond to the impact of hardware technologies and platform characteristics
  - Work with the facilities and vendors towards a successful stable deployment of our software technologies
  - Understand and respond to dependencies within the stack, avoiding duplication and scope creep
  - This is a comprehensive team effort — not a set of individual projects!

# 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 (CORAL, APEX)
- Expected software environment for an exascale system
- Requirements beyond the software environment provided by vendors of HPC systems
- What is needed to meet ECP's Key Performance Parameters (KPPs)

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

## Exascale Challenge Problem

- Safe and efficient use of the subsurface for geologic $CO_2$ 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

*PI: Carl Steefel (LBNL)

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

## Exascale Challenge Problem

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

## Applications & S/W Technologies

**Applications**

- NWChemEx (evolved from redesigned NWChem)

**Software Technologies Cited**

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

## Risks and Challenges

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

## Development Plan

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

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

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

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

*PI: Thom Dunning (PNNL)



EXASCALE COMPUTING PROJECT

# 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**

# No. of ECP Application Proposals a Software is Mentioned in

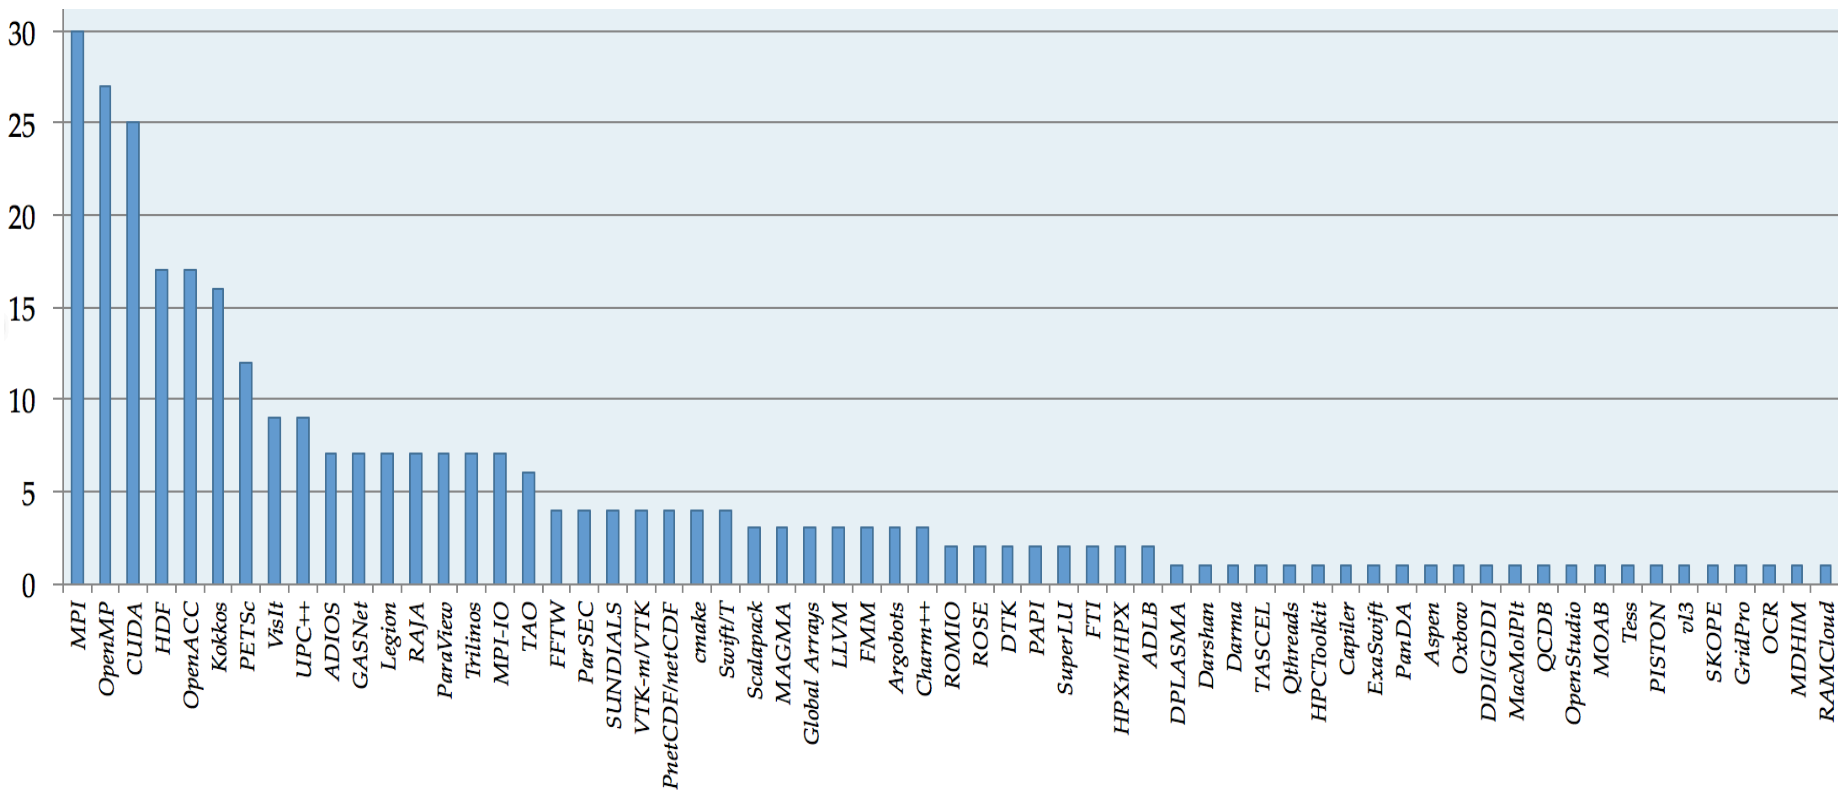

# Libraries used at NERSC
## (similar data from other facilities)

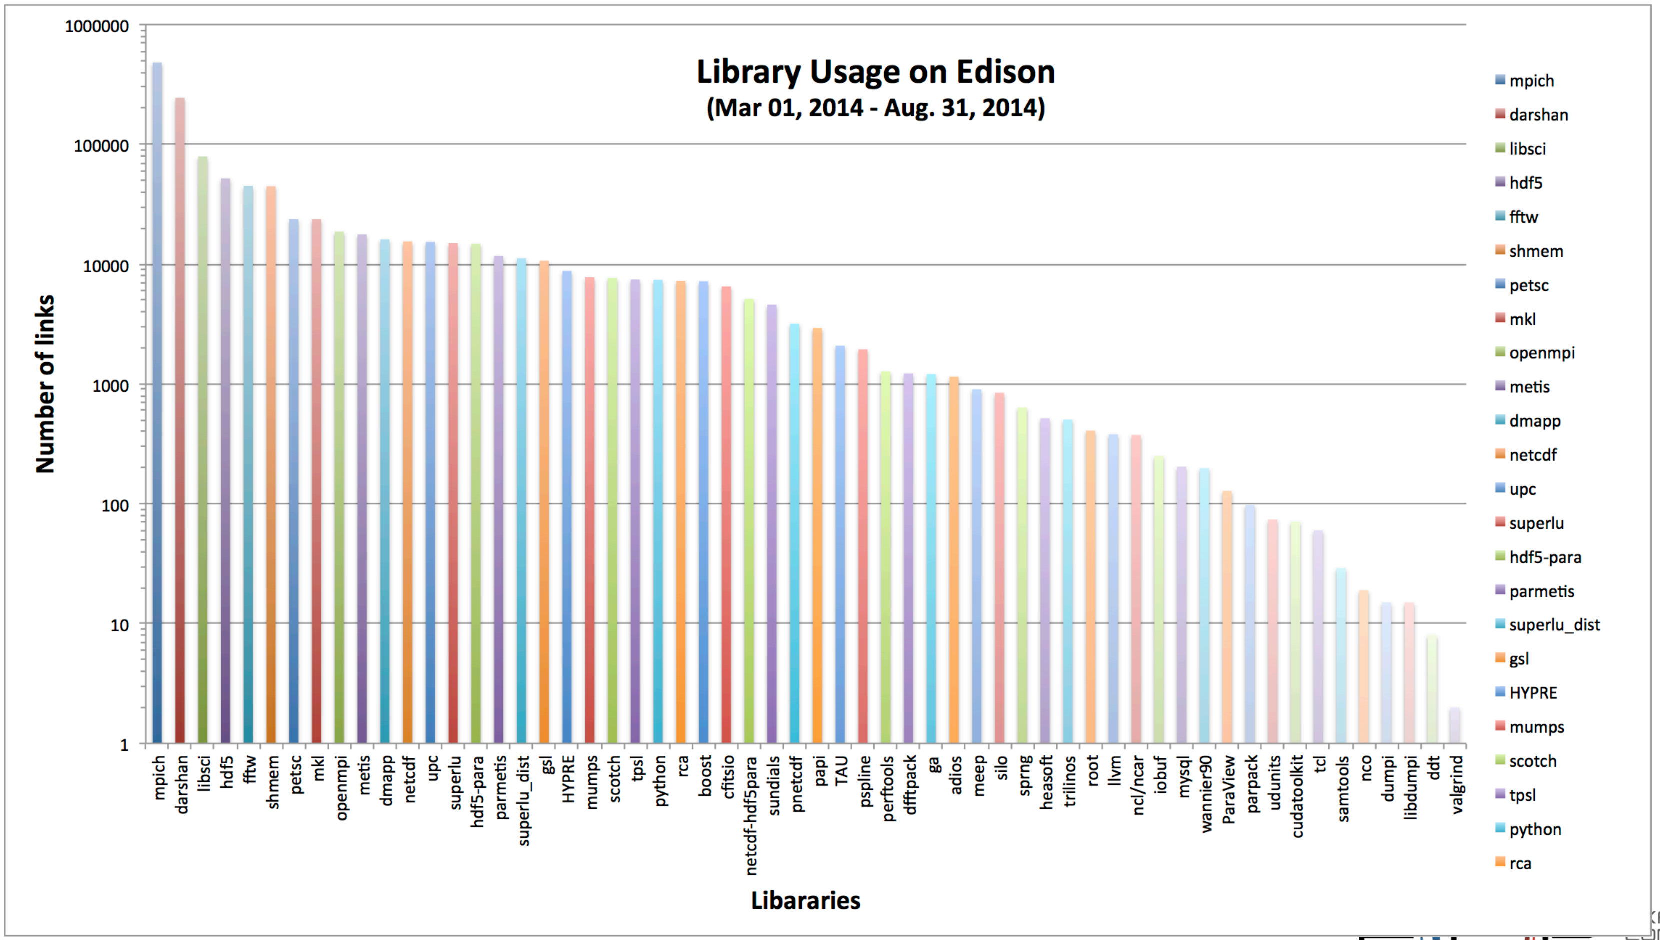

# ECP Projects Status

- 22 application projects have been selected for funding
  - In addition to 3 applications projects already underway at the NNSA labs
- 5 co-design centers have been selected for funding
- 35 software technology projects have been selected for funding
  - In addition to similar number already underway at NNSA labs
- Proposals submitted to PathForward RFP (Hardware Technology R&D by vendors) have been selected for funding
  - Negotiations underway; contracts expected to be signed by April 2017

# Conceptual ECP Software Stack

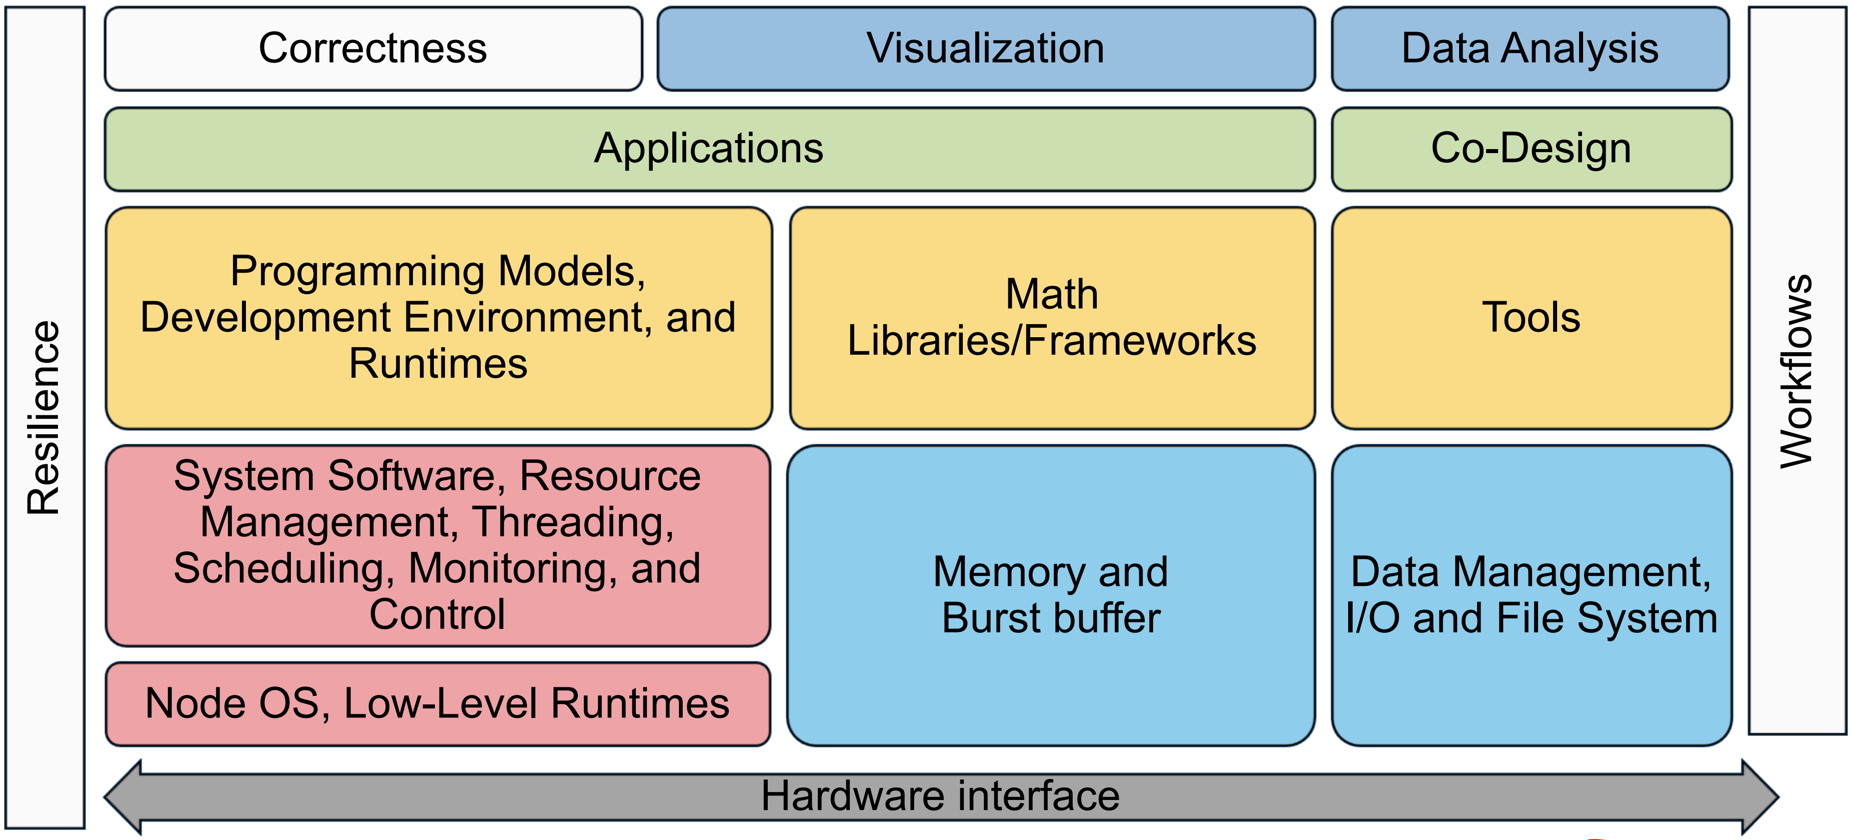

# Current Set of ST Projects (incl. NNSA ATDM)

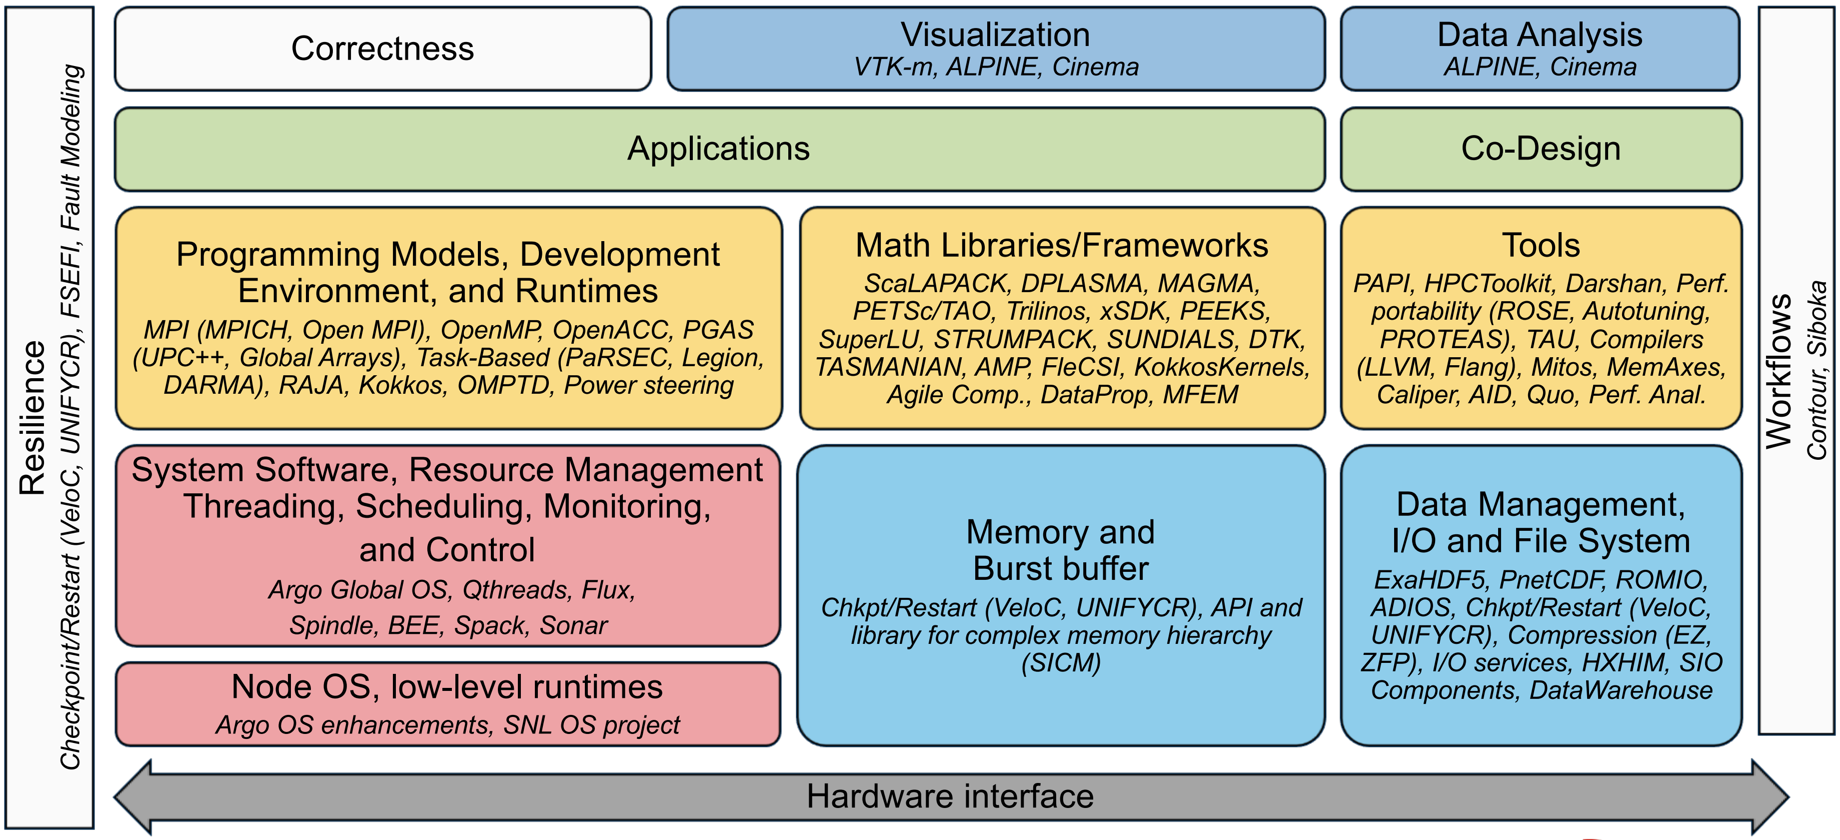

# Software Technology Projects in Programming Models and Runtime

# Programming Models and Runtimes: Overview

- The goal is to support multiple programming models
  - MPI, OpenMP, OpenACC
  - PGAS models
  - Task-based models
  - Abstractions for performance portability

- Applications are encouraged to use some of the newer programming models (beyond MPI and OpenMP), as well as the newer features of MPI and OpenMP

# Exascale MPI (MPICH)

- Lead: Pavan Balaji (ANL)
- Co-PI: Marc Snir (UIUC, ANL)
- Improvements to the MPICH implementation
  - Threads, heterogeneous memory, topology-aware communication, fault tolerance, new features that will be added to the MPI Standard, etc.
- Improvements to the MPI Standard for MPI-4
- Transfer technology to vendors to enable vendor-optimized implementations
- Coordinate with Open MPI and OpenMP efforts

# Open MPI for Exascale (OMPI-X)

- Lead: David Bernholdt (ORNL)
- Co-PIs: Howard Pritchard (LANL), Martin Schulz (LLNL), Ron Brightwell (SNL), George Bosilca (UTK)
- Improvements to the Open MPI implementation
  - Scalability, performance, runtime interoperability for MPI+X, new MPI features, resilience, dynamic execution environments
- Improvements to the MPI Standard for MPI-4
- Transfer technology to vendors to enable vendor-optimized implementations
- Coordinate with MPICH and OpenMP efforts

# SOLLVE: Scaling OpenMP with LLVM for Exascale performance and portability

- Lead: Barbara Chapman (BNL)
- Co-PIs: Pavan Balaji (ANL), Bronis de Supinski (LLNL), David Bernholdt (ORNL), Vivek Sarkar (Rice), Sanjay Kale (UIUC)
- Enhanced OpenMP implementation using LLVM and lightweight threading
- Enhancements to the OpenMP standard
- Facilitate OpenMP-MPI interoperability
- Work with applications to enable them to use advanced OpenMP features for performance portability
- Tech transfer to vendors
- Coordinate with MPI efforts

# Lightweight Communication and Global Address Space Support for Exascale Applications

- Lead: Scott Baden (LBNL); Co-PI: Paul Hargrove (LBNL)
- UPC++ (PGAS abstraction); GASNet communication library
- Effectively utilize RDMA and global-address support in the hardware
- Efficient irregular communication with small/non-contiguous payloads
- Simplify the design of efficient distributed data structures
- Asynchronous remote execution
- Used in Legion runtime, ExaBiome application, AMR codesign center, and others

# xGA: Global Arrays on Extreme Scale Architectures

- Lead: Abhinav Vishnu (PNNL)
- Co-PIs: Bruce Palmer and Jeff Daily (PNNL)
- Extend Global Arrays (GA) for exascale architectures
  - Novel data structures
  - Novel PGAS property types
  - Fine-grained multithreading
  - Interoperability with synergistic programming models
- Used in NWChem, GridPACK, and others

# Distributed Tasking for Exascale (PARSEC)

- Lead: Jack Dongarra (UTK)
- Co-PIs: George Bosilca, Aurelien Boutelier, Thomas Herault (UTK)
- Runtime for dynamic execution on heterogeneous distributed systems
- Domain Specific Languages (DSLs) with debugging, trace collection, and analysis tools
- Used in DPLASMA dense linear algebra package, NWChem (prototype), and others

# Enhancing and Hardening the Legion Programming System for the Exascale Computing Project

- Lead: Galen Shipman (LANL)
- Co-PI: Alex Aiken (Stanford). Partners: ANL, NVIDIA
- High-level task-based programming model
- Runtime system automates the mapping and scheduling of tasks and the movement of data
- Used in S3D (combustion), LANL ATDM next-generation code, and others

# DARMA: Asynchronous Many Task Abstraction Layer

- Lead: Janine Bennett (SNL)
- AMT portability layer and specification that aims to
  - insulate application from runtime system and hardware idiosyncrasies
  - improve AMT runtime programmability by co-designing an API directly with application developers
  - synthesize application co-design activities into meaningful requirements for runtimes
  - facilitate AMT design space characterization and definition, accelerating the development of AMT best practices.
- Used in SNL ATDM next-generation code

# RAJA: On-Node Programming Model Abstractions for C++

- Lead: Richard Hornung (LLNL)
- Abstraction layer that encapsulates architecture and parallel programming model details for C++ application portability
- Core development and outreach to broader set of ECP applications
- Used in production ATDM and ASC applications

# Kokkos: On-Node Performance Portable Abstraction Layer

- Lead: Carter Edwards (SNL)
- Co-PIs: Fernanda Foertter (ORNL), Galen Shipman (LANL)
- C++ template-based interfaces for performance portable node-level parallelism
- Core development and outreach to broader set of ECP applications
- Engagement and advocacy within ISO/C++ Standards committee to introduce Kokkos' abstractions/capabilities into C++20 standard
- Used in several apps at SNL, LANL, ORNL, U-Utah, ARL, NRL

# Runtime Library for Application-Level Power Steering

- Lead: Martin Schulz (LLNL)
- Co-PIs: Barry Rountree (LLNL), David Lowenthal (U Arizona)
- Collaborator: Jonathan Eastep (Intel)
- Job-level power management system that can optimize performance under power and/or energy constraints
- Enhance Intel's GEO (Global Extensible Open Power Manager) library
- Conductor and Adagio power management libraries from LLNL
- Caliper-enabled applications will be able to take greater advantage of the runtime system capabilities

# Tools Interface for OpenMP

- Lead: Martin Schulz (LLNL)
- Collaborators: IBM, RWS TotalView, NVIDIA, U of Utah, Rice U, RWTH Aachen
- Work with the OpenMP standards committee to define
  – A performance tools interface (OMPT)
  – A debugging tools interface (OMPD)
- Will result in better tools support for OpenMP

# ECP RFI to Vendors

- Issued Feb 6, 2017; responses received March 31, 2017
- http://procurement.ornl.gov/rfp/exascale-rfi/
- Asks vendors for information on their hardware and software plans up to 2024 and identify potential gaps

**Exascale Systems 2021-2024**
**Request for Information**
February 6, 2017

**Introduction**

The U.S. Department of Energy (DOE) has a long history of deploying leading-edge computing capability for science and national security. The acquisition plans of the large DOE computer facilities have continued to march forward with new systems being deployed approximately every two years. On this schedule, ANL, ORNL, and LLNL (CORAL) will be acquiring three supercomputers around 2022 and LANL, SNL, and LBNL (APEX) will be acquiring two supercomputers around 2024.

# Next Steps for ST: Gap Analysis

- The results of a number of discussions (with vendors, facilities, applications) and responses to the ECP RFI to vendors are feeding into a gap analysis
- The gap analysis will specify what remaining pieces ECP needs to cover and what will be covered by vendors and other open source efforts
- Examples of known gaps:
  - Workflow infrastructure and tools
  - Software needed for applications in data analytics
  - Integration of ATDM ST efforts into overall architecture
  - Software co-design
  - Commonly used external software not covered in the current portfolio
- Will be reviewed by ECP leadership, senior lab leadership, and others
  - Update continuously based on feedback
- Contingent on available funding, we will issue an RFI/RFP later this year to close the gaps

# ECP Activities in Applications Development

- 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

# ECP Applications Deliver Broad Coverage of Strategic Pillars

## Initial (FY16) selections consist of 15 application projects + 7 seed efforts

| National Security | Energy Security | Economic Security | Scientific Discovery | Climate and Environmental Science | Healthcare |
|---|---|---|---|---|---|
| • Stockpile Stewardship | • Turbine Wind Plant Efficiency<br>• Design/Commercialization of SMRs<br>• Nuclear Fission and Fusion Reactor Materials Design<br>• Subsurface Use for Carbon Capture, Petro Extraction, Waste Disposal<br>• High-Efficiency, Low-Emission Combustion Engine and Gas Turbine Design<br>• Carbon Capture and Sequestration Scaleup (S)<br>• Biofuel Catalyst Design (S) | • Additive Manufacturing of Qualifiable Metal Parts<br>• Urban Planning (S)<br>• Reliable and Efficient Planning of the Power Grid (S)<br>• Seismic Hazard Risk Assessment (S) | • Cosmological Probe of the Standard Model (SM) of Particle Physics<br>• Validate Fundamental Laws of Nature (SM)<br>• Plasma Wakefield Accelerator Design<br>• Light Source-Enabled Analysis of Protein and Molecular Structure and Design<br>• Find, Predict, and Control Materials and Properties<br>• Predict and Control Stable ITER Operational Performance<br>• Demystify Origin of Chemical Elements (S) | • Accurate Regional Impact Assessment of Climate Change<br>• Stress-Resistant Crop Analysis and Catalytic Conversion of Biomass-Derived Alcohols<br>• Metegenomics for Analysis of Biogeochemical Cycles, Climate Change, Environ Remediation (S) | • 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

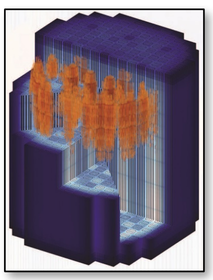

## Climate (BER)

Accurate regional impact assessment of climate change*

Climate Action Plan

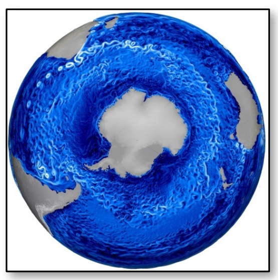

## Chemical Science (BES, BER)

Biofuel catalysts design; stress-resistant crops

Climate Action Plan; MGI

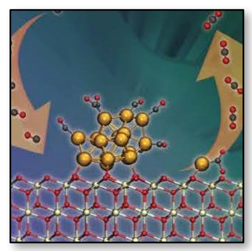

## Wind Energy (EERE)

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

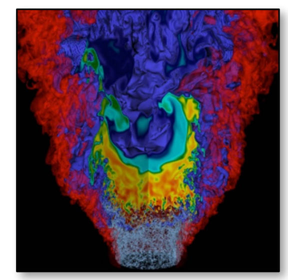

* Scope includes a discernible data science component

# Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development projects

### Materials Science (BES)

Find, predict, and control materials and properties: property change due to hetero-interfaces and complex structures

MGI

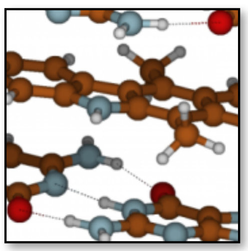

### Nuclear Physics (NP)

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

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

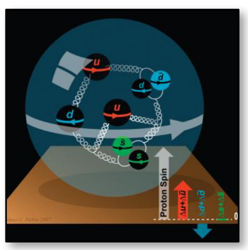

### Nuclear Materials (BES, NE, FES)

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

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

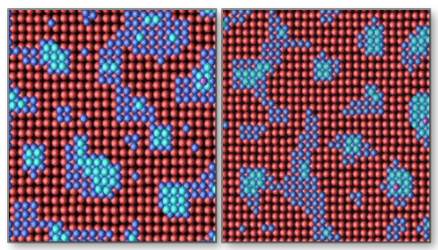

### Accelerator Physics (HEP)

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

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

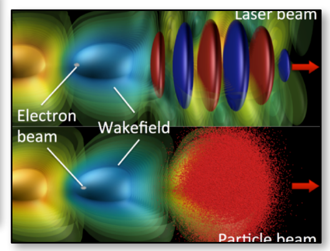


### Materials Science (BES)

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

MGI; LCLS-II 2025 Path Forward

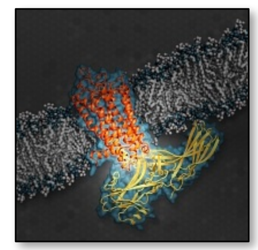

* Scope includes a discernible data science component

# Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy and Other Agency application development projects

### Magnetic Fusion Energy (FES)

Predict and guide stable ITER operational performance with an integrated whole device model*

ITER; fusion experiments: NSTX, DIII-D, Alcator C-Mod

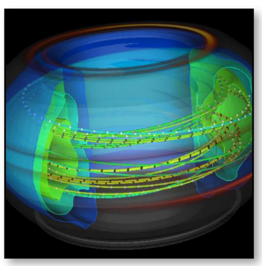

### Advanced Manufacturing (EERE)

Additive manufacturing process design for qualifiable metal components*

NNMIs; Clean Energy Manufacturing Initiative

### Cosmology (HEP)

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

Particle Physics Project Prioritization Panel (P5)

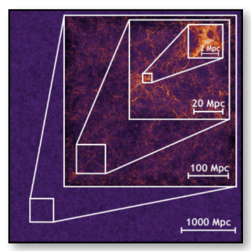


### Geoscience (BES, BER, EERE, FE, NE)

Safe and efficient use of subsurface for carbon capture and storage, petroleum extraction, geothermal energy, nuclear waste*

EERE Forge; FE NRAP; Energy-Water Nexus; SubTER Crosscut

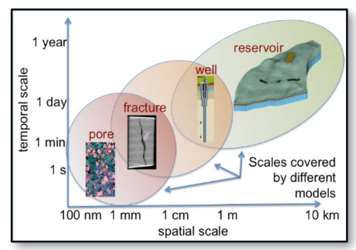


### Precision Medicine for Cancer (NIH)

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

Precision Medicine in Oncology; Cancer Moonshot

* Scope includes a discernible data science component

# Exascale Applications Will Address National Challenges

Summary of current DOE Science & Energy application development seed projects

## Seismic (EERE, NE, NNSA)

Reliable earthquake hazard and risk assessment in relevant frequency ranges*

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

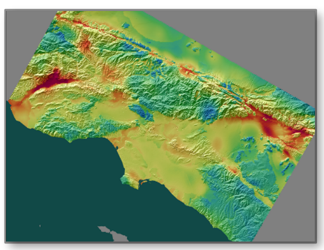

## Carbon Capture and Storage (FE)

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

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

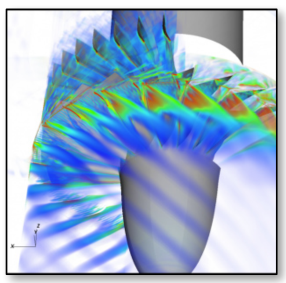

## Chemical Science (BES)

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

Climate Action Plan; SunShot Initiative; MGI

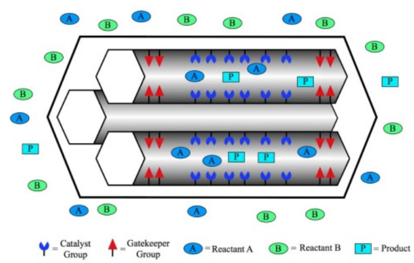


## Urban Systems Science (EERE)

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

Energy-Water Nexus; Smart Cities Initiative

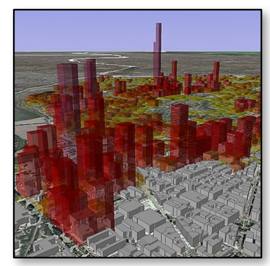

* Scope includes a discernible data science component



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 |

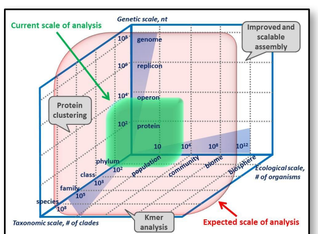




* Scope includes a discernible data science component

# Application Motifs*

Algorithmic methods that capture a common pattern of computation and communication

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



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

# Survey of Application Motifs

| Application | Monte Carlo | Particles | Sparse Linear Algebra | Dense Linear Algebra | Spectral Methods | Unstructured Grid | Structured Grid | Comb. Logic | Graph Traversal | Dynamical Program | Backtrack & Branch and Bound | Graphical Models | Finite State Machine |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Cosmology | | | | | | | | | | | | | |
| Subsurface | | | | | | | | | | | | | |
| Materials (QMC) | | | | | | | | | | | | | |
| Additive Manufacturing | | | | | | | | | | | | | |
| Chemistry for Catalysts & Plants | | | | | | | | | | | | | |
| Climate Science | | | | | | | | | | | | | |
| Precision Medicine Machine Learning | | | | | | | | | | | | | |
| QCD for Standard Model Validation | | | | | | | | | | | | | |
| Accelerator Physics | | | | | | | | | | | | | |
| Nuclear Binding and Heavy Elements | | | | | | | | | | | | | |
| MD for Materials Discovery & Design | | | | | | | | | | | | | |
| Magnetically Confined Fusion | | | | | | | | | | | | | |

# Survey of Application Motifs

| Application | Monte Carlo | Particles | Sparse Linear Algebra | Dense Linear Algebra | Spectral Methods | Unstructured Grid | Structured Grid | Comb. Logic | Graph Traversal | Dynamical Program | Backtrack & Branch and Bound | Graphical Models | Finite State Machine |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Combustion S&T | | ■ | ■ | ■ | | | ■ | | | | | | |
| Free Electron Laser Data Analytics | ■ | ■ | ■ | ■ | ■ | | ■ | | | ■ | ■ | ■ | |
| Microbiome Analysis | | | ■ | | | | | ■ | ■ | ■ | | ■ | |
| Catalyst Design | ■ | ■ | ■ | ■ | ■ | | | | | | | | |
| Wind Plant Flow Physics | | | ■ | ■ | | ■ | | | | | | | |
| SMR Core Physics | ■ | ■ | ■ | ■ | ■ | ■ | ■ | | | | | | |
| Next-Gen Engine Design | | ■ | ■ | | ■ | ■ | | | | | | | |
| Urban Systems | | | ■ | | ■ | ■ | ■ | | | | | | |
| Seismic Hazard Assessment | | | ■ | | ■ | ■ | ■ | | | ■ | | | |
| Systems Biology | ■ | | ■ | ■ | | | | ■ | ■ | ■ | ■ | ■ | |
| Biological Neutron Science | | ■ | ■ | ■ | ■ | | ■ | | ■ | | ■ | | |
| Power Grid Dynamics | ■ | | ■ | ■ | | | | | ■ | ■ | ■ | | |

# Survey of Application Motifs

| Application | Monte Carlo | Particles | Sparse Linear Algebra | Dense Linear Algebra | Spectral Methods | Unstructured Grid | Structured Grid | Comb. Logic | Graph Traversal | Dynamical Program | Backtrack & Branch and Bound | Graphical Models | Finite State Machine |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Stellar Explosions | | | | | | | | | | | | | |
| Excited State Material Properties | | | | | | | | | | | | | |
| Light Sources | | | | | | | | | | | | | |
| Materials for Energy Conversion/Storage | | | | | | | | | | | | | |
| Hypersonic Vehicle Design | | | | | | | | | | | | | |
| Multiphase Energy Conversion Devices | | | | | | | | | | | | | |

# Application Co-Design (CD)

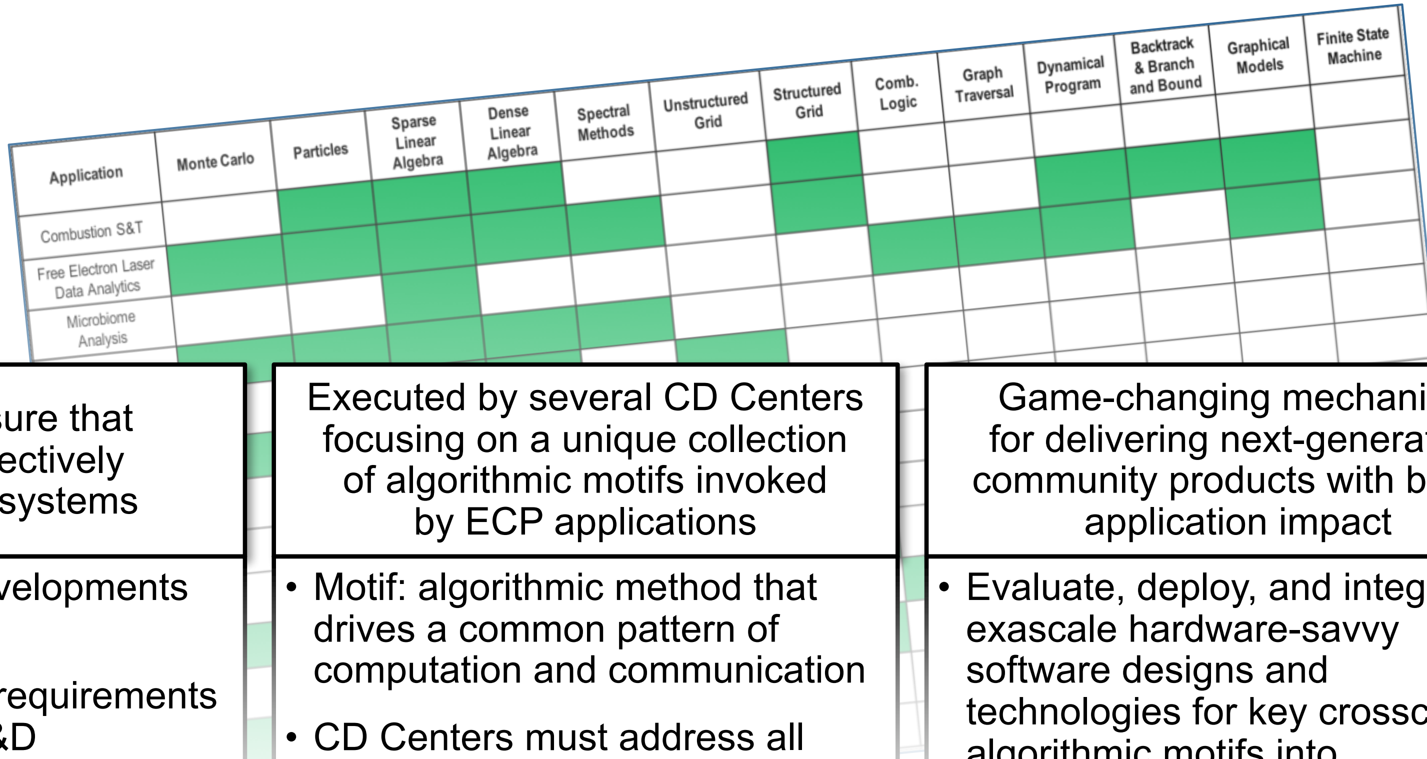


**Essential to ensure that applications effectively utilize exascale systems**

- 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



EXASCALE COMPUTING PROJECT

# ECP Co-Design Centers

- **A Co-Design Center for Online Data Analysis and Reduction at the Exascale (CODAR)**
  - **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 (AMReX)**
  - **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
- **Combinatorial Methods for Enabling Exascale Applications (ExaGraph)**
  - **Motif(s)**: Graph traversals; graph matching; graph coloring; graph clustering, including clique enumeration, parallel branch-and-bound, graph partitioning
  - Develop methods and techniques for efficient implementation of key combinatorial (graph) algorithms that play a critical enabling role in numerous scientific applications. The irregular memory access nature of these algorithms makes them difficult algorithmic kernels to implement on parallel systems

# Some application challenges

- Exploiting on-node memory and compute hierarchies
- Programming models: what to use where and how
- 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)

# Hardware Technology Overview

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

## Vendor Partnerships for Hardware Architecture R&D contracts that deliver:

- Conceptual exascale node and system designs
- Analysis of performance improvement on these conceptual system designs
- Technology demonstrators to quantify performance gains over existing roadmaps
- Support for active industry engagement in ECP holistic co-design efforts

## DOE Lab Design Space Evaluation Team

- Participate in evaluation and review of Vendor HW Architecture R&D deliverables
- Provide Architectural Analysis, and Abstract Machine Models of Vendor designs to support ECP's holistic co-design

# Transition to higher trajectory with advanced architecture

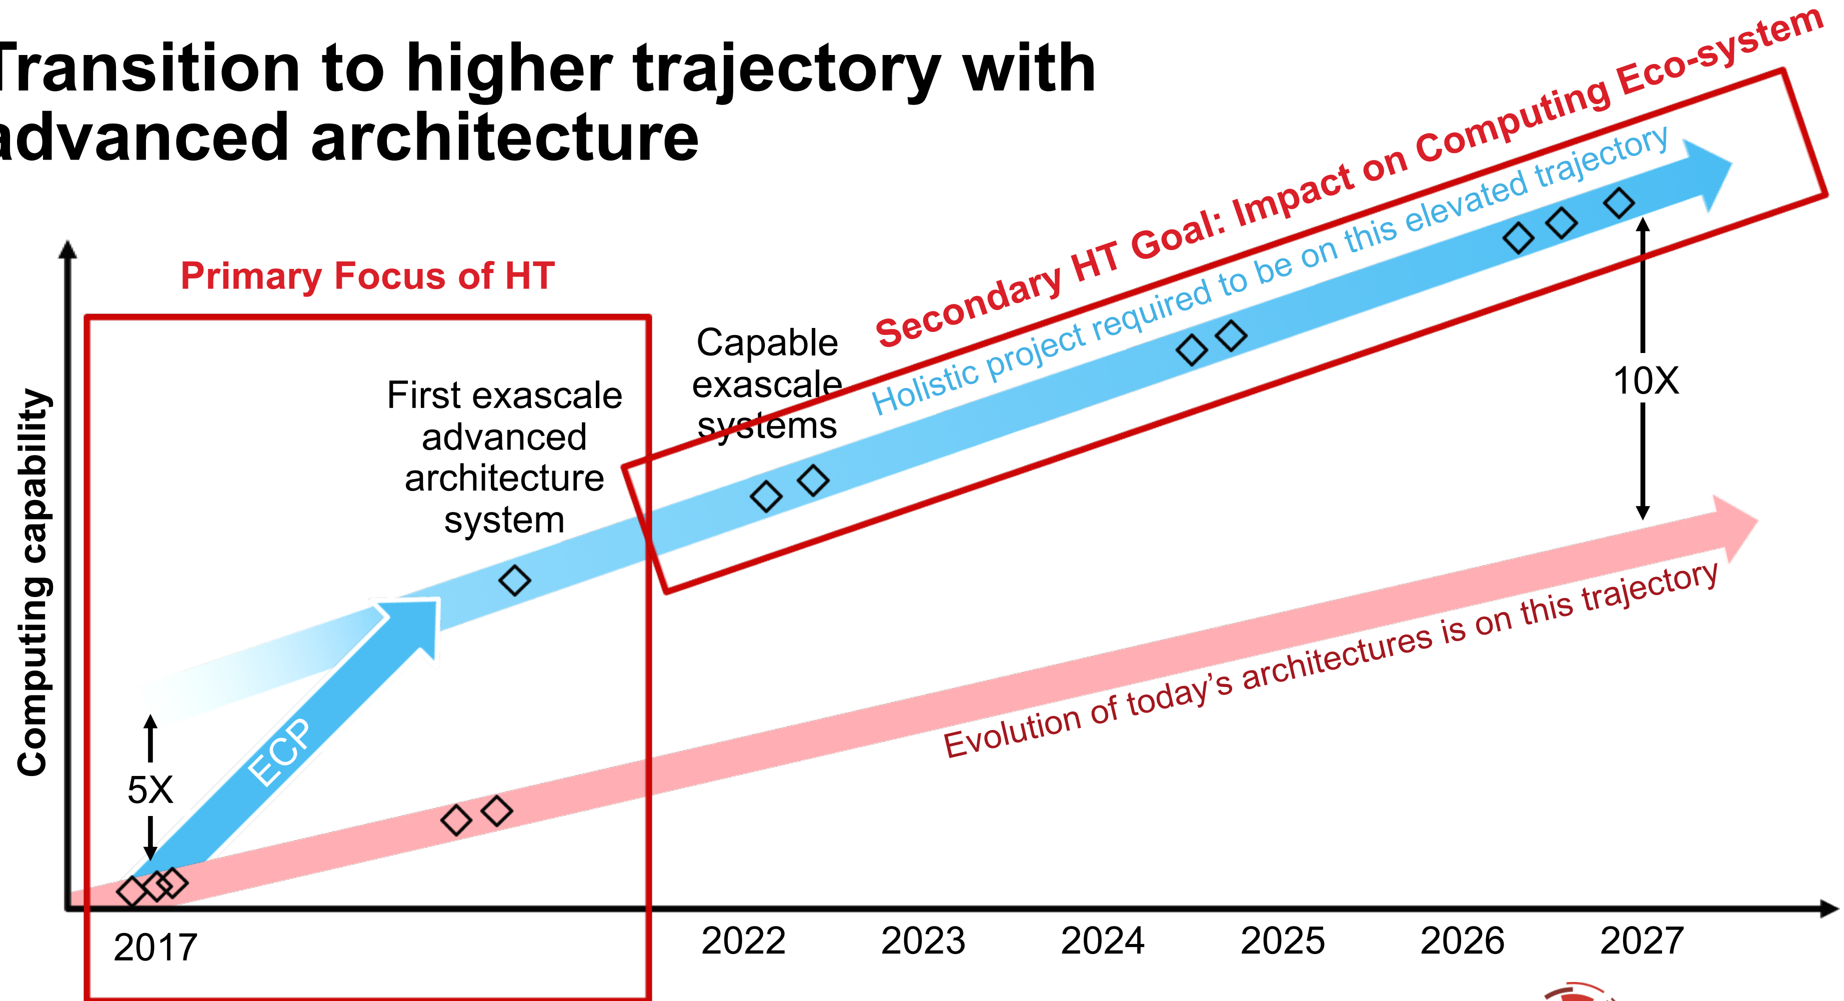

# Systems Acquisition Approach

- DOE-SC and NNSA Facilities will procure and install the ECP systems
  - ECP's and DOE-SC/NNSA's processes will be tightly coupled and interdependent
  - ECP's requirements will be incorporated into RFP(s)
  - ECP will participate in system selection and co-design
  - ECP will make investments through non-recurring engineering (NRE) contracts coupled to system acquisition contracts

**NRE contracts**

- Incentivize awardees to address gaps in their system product roadmaps
- Bring to the product stage promising hardware and software research and integrate it into a system
- Accelerate technologies, add capabilities, improve performance, and lower the cost of ownership of system
- Include application readiness R&D efforts
- More than 2 full years of lead time are necessary to maximize impact

# What the ECP is addressing partially, or not at all

- Only partially tackling convergence of simulation and data analytics
  - Hope to do more, given sufficient funding
  - Deep learning: funding few applications so far, hope to do more but vendors already investing a lot; the number of applications is exploding
    - Do technology watch, find gaps in coverage, be very selective in what we do in ECP
- Post Moore's Law technologies
  - Out of scope for ECP
- Basic research on new programming models
  - Insufficient time to determine their value or deliver production quality implementations

# Planned outcomes of the ECP

- Important applications running at exascale in 2021, producing useful results
- A full suite of mission and science applications ready to run on the 2023 exascale systems
- A large cadre of computational scientists, engineers, and computer scientists who will be an asset to the nation long after the end of ECP
- An integrated software stack that supports exascale applications
- Results of PathForward R&D contract with vendors are integrated into exascale systems and are in vendors' product roadmaps
- Industry and mission critical applications have been prepared for a more diverse and sophisticated set of computing technologies, carrying US supercomputing well into the future

# Participate in EuroMPI/USA 2017


- www.mcs.anl.gov/eurompi2017
- General Chair: Pavan Balaji
- Program Chairs: Bill Gropp, Rajeev Thakur
- Papers Submissions Due: May 1, 2017
- Many good workshops along with the conference, including one **celebrating 25 years of MPI**

# Questions?

www.ExascaleProject.org

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science