“Housekeeping”

Twitter: #ACMLearning

- Welcome to today’s ACM TechTalk, “**The Exascale Computing Project and the Future of HPC**.” The presentation starts at the top of the hour and lasts 60 minutes. Audio and video will automatically play throughout the event. On the bottom panel you’ll find a number of widgets, including Twitter and Sharing apps.

- If you are experiencing any problems/issues, **refresh** your console by pressing the **F5** key on your keyboard in Windows, **Command + R** if on a Mac, or refresh your browser if you’re on a mobile device; or close and re-launch the presentation. You can also view the Webcast Help Guide, by clicking on the “Help” widget in the bottom dock.

- To control volume, adjust the master volume on your computer. If the volume is still too low, use headphones.

- At the end of the presentation, you’ll see a survey open on your screen. Please take a minute to fill it out to help us improve your next webinar experience. You may also open the survey at any time throughout the presentation from the resources window.

- This session is being recorded and will be archived for on-demand viewing in a few days. You will receive an automatic email notification when it is available. See http://learning.acm.org/ for updates. And check out https://learning.acm.org/techtalks for archived recordings of past webcasts.

The U.S. Department of Energy Exascale Computing Project

EXASCALE COMPUTING PROJECT

Douglas B. Kothe (Oak Ridge National Laboratory)
Director, Exascale Computing Project (ECP)
kothe@ornl.gov

Association for Computing Machinery (ACM) Tech Talk
April 30, 2019

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

ACM Highlights

- Learning Center tools for professional development: http://learning.acm.org
  - The Safari Learning Platform featuring the **entire Safari collection of nearly 50,000** technical books, video courses, O'Reilly conference videos, learning paths, tutorials, case studies
  - 1,800+ Skillsoft courses, 4,800+ online books, and 30,000+ task-based short videos for software professionals covering programming, data management, DevOps, cybersecurity, networking, project management, and more; including training toward top vendor certifications such as AWS, CEH, Cisco, CISSP, CompTIA, Oracle, RedHat, PMI.
  - 1,200+ books from Elsevier on the ScienceDirect platform (including Morgan Kaufmann and Syngress titles)
  - TechTalks from thought leaders and top practitioners
  - Podcast interviews with innovators, entrepreneurs, and award winners
- Popular publications:
  - Flagship ***Communications of the ACM (CACM)*** magazine: http://cacm.acm.org
  - *ACM Queue* magazine for practitioners: http://queue.acm.org
- The **ACM Code of Ethics**, a set of principles and guidelines principles and guidelines designed to help computing professionals make ethically responsible decisions in professional practice: https://ethics.acm.org
  ACM Digital Library, the world's most comprehensive database of computing literature: http://dl.acm.org
- International conferences that draw leading experts on a broad spectrum of computing topics: http://www.acm.org/conferences
- Prestigious awards, including the **ACM A.M. Turing** and ACM Prize in Computing: http://awards.acm.org
- And much more… http://www.acm.org.

“Housekeeping”

Twitter: #ACMLearning

- Welcome to today’s ACM TechTalk, “**The Exascale Computing Project and the Future of HPC**.” The presentation starts at the top of the hour and lasts 60 minutes. Audio and video will automatically play throughout the event. On the bottom panel you’ll find a number of widgets, including Twitter and Sharing apps.

- If you are experiencing any problems/issues, **refresh** your console by pressing the **F5** key on your keyboard in Windows, **Command + R** if on a Mac, or refresh your browser if you’re on a mobile device; or close and re-launch the presentation. You can also view the Webcast Help Guide, by clicking on the “Help” widget in the bottom dock.

- To control volume, adjust the master volume on your computer. If the volume is still too low, use headphones.

- At the end of the presentation, you’ll see a survey open on your screen. Please take a minute to fill it out to help us improve your next webinar experience. You may also open the survey at any time throughout the presentation from the resources window.

- This session is being recorded and will be archived for on-demand viewing in a few days. You will receive an automatic email notification when it is available. See http://learning.acm.org/ for updates. And check out https://learning.acm.org/techtalks for archived recordings of past webcasts.

Talk Back

- Tweet your favorite quotes from today's presentation with hashtag #ACMLearning
- Submit questions and comments via Twitter to @acmeducation – we're reading them!
- The ACM Discourse Page is available for post-talk discussion – https://on.acm.org

The U.S. Department of Energy Exascale Computing Project

Douglas B. Kothe (Oak Ridge National Laboratory)
Director, Exascale Computing Project (ECP)
kothe@ornl.gov

Association for Computing Machinery (ACM) Tech Talk
April 30, 2019

The Exascale Computing Project (ECP) enables US revolutions in technology development; scientific discovery; healthcare; energy, economic, and national security

ECP mission

Develop exascale-ready applications and solutions that address currently intractable problems of strategic importance and national interest.

Create and deploy an expanded and vertically integrated software stack on DOE HPC exascale and pre-exascale systems, defining the enduring US exascale ecosystem.

Deliver **US HPC vendor technology advances and deploy ECP products** to DOE HPC pre-exascale and exascale systems.

ECP vision

Deliver **exascale simulation and data science innovations and solutions to national problems** that enhance US economic competitiveness, change our quality of life, and strengthen our national security.

DOE Exascale Program: The Exascale Computing Initiative (ECI)

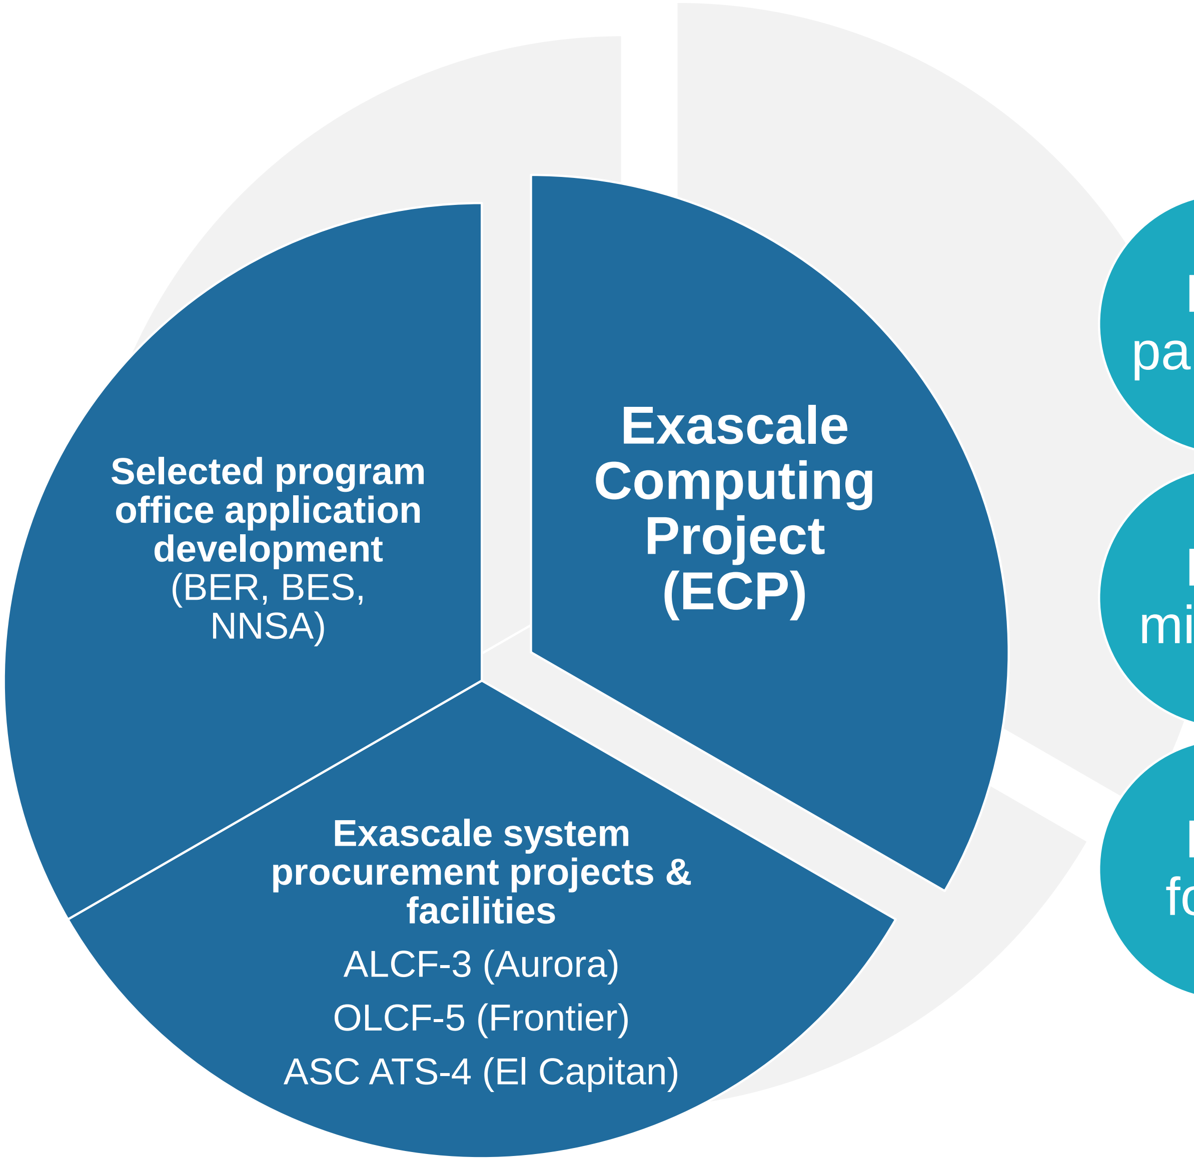

ECI partners — US DOE Office of Science (SC) and National Nuclear Security Administration (NNSA)

ECI mission — Accelerate R&D, acquisition, and deployment to deliver exascale computing capability to DOE national labs by the early- to mid-2020s

ECI focus — Delivery of an *enduring and capable exascale computing capability for use by a wide range of applications* of importance to DOE and the US

Three Major Components of the ECI

What is a "capable" exascale computing ecosystem?

Hardware

At least two diverse system architectures

Delivers 50x the performance of today's 20 petaflop systems and 5x the performance of Summit, Oak Ridge National Laboratory's supercomputer— i.e., allows at least a quintillion floating point operations per second

Functions with sufficient resiliency: an average fault rate of ≤1 per week

Software

Includes a software stack that meets the needs of a broad spectrum of applications and workloads

Applications

Supports a wide range of applications that deliver high-fidelity solutions in less time to problems of greater complexity

Exascale means real capability improvement in the science we can do, and how fast we can do it

Why high performance computing is hard, and getting harder

- Applications need to find more and more concurrency to keep up.
- Moving data becoming increasingly costly relative to computation.
- I/O, vis, analysis becoming major bottlenecks
- Hardware landscape getting more diverse, future architectures have more uncertainty.
- New programming models are being developed to supplement traditional MPI+X approaches:
  - On-node: OCCA, Kokkos, RAJA, OpenACC, OpenCL, Swift
  - Inter-node: Legion, UPC++, Global Arrays

Preparing applications for new architectures can be difficult and time-consuming, working together and learning from each other is crucial

ECP by the Numbers

7
YEARS

$1.7B

A seven-year, $1.7 B R&D effort that launched in 2016

6
CORE DOE
LABS

Six core DOE National Laboratories: Argonne, Lawrence Berkeley, Lawrence Livermore, Oak Ridge, Sandia, Los Alamos

- Staff from most of the 17 DOE national laboratories take part in the project

3
FOCUS
AREAS

Three technical focus areas: Hardware and Integration, Software Technology, Application Development supported by a Project Management Office

100
R&D TEAMS

1000
RESEARCHERS

More than 100 top-notch R&D teams

Hundreds of consequential milestones delivered on schedule and within budget since project inception

EXASCALE COMPUTING PROJECT

Vision: Exascale Computing Project (ECP) Lifts all U.S. High Performance Computing to a New Trajectory

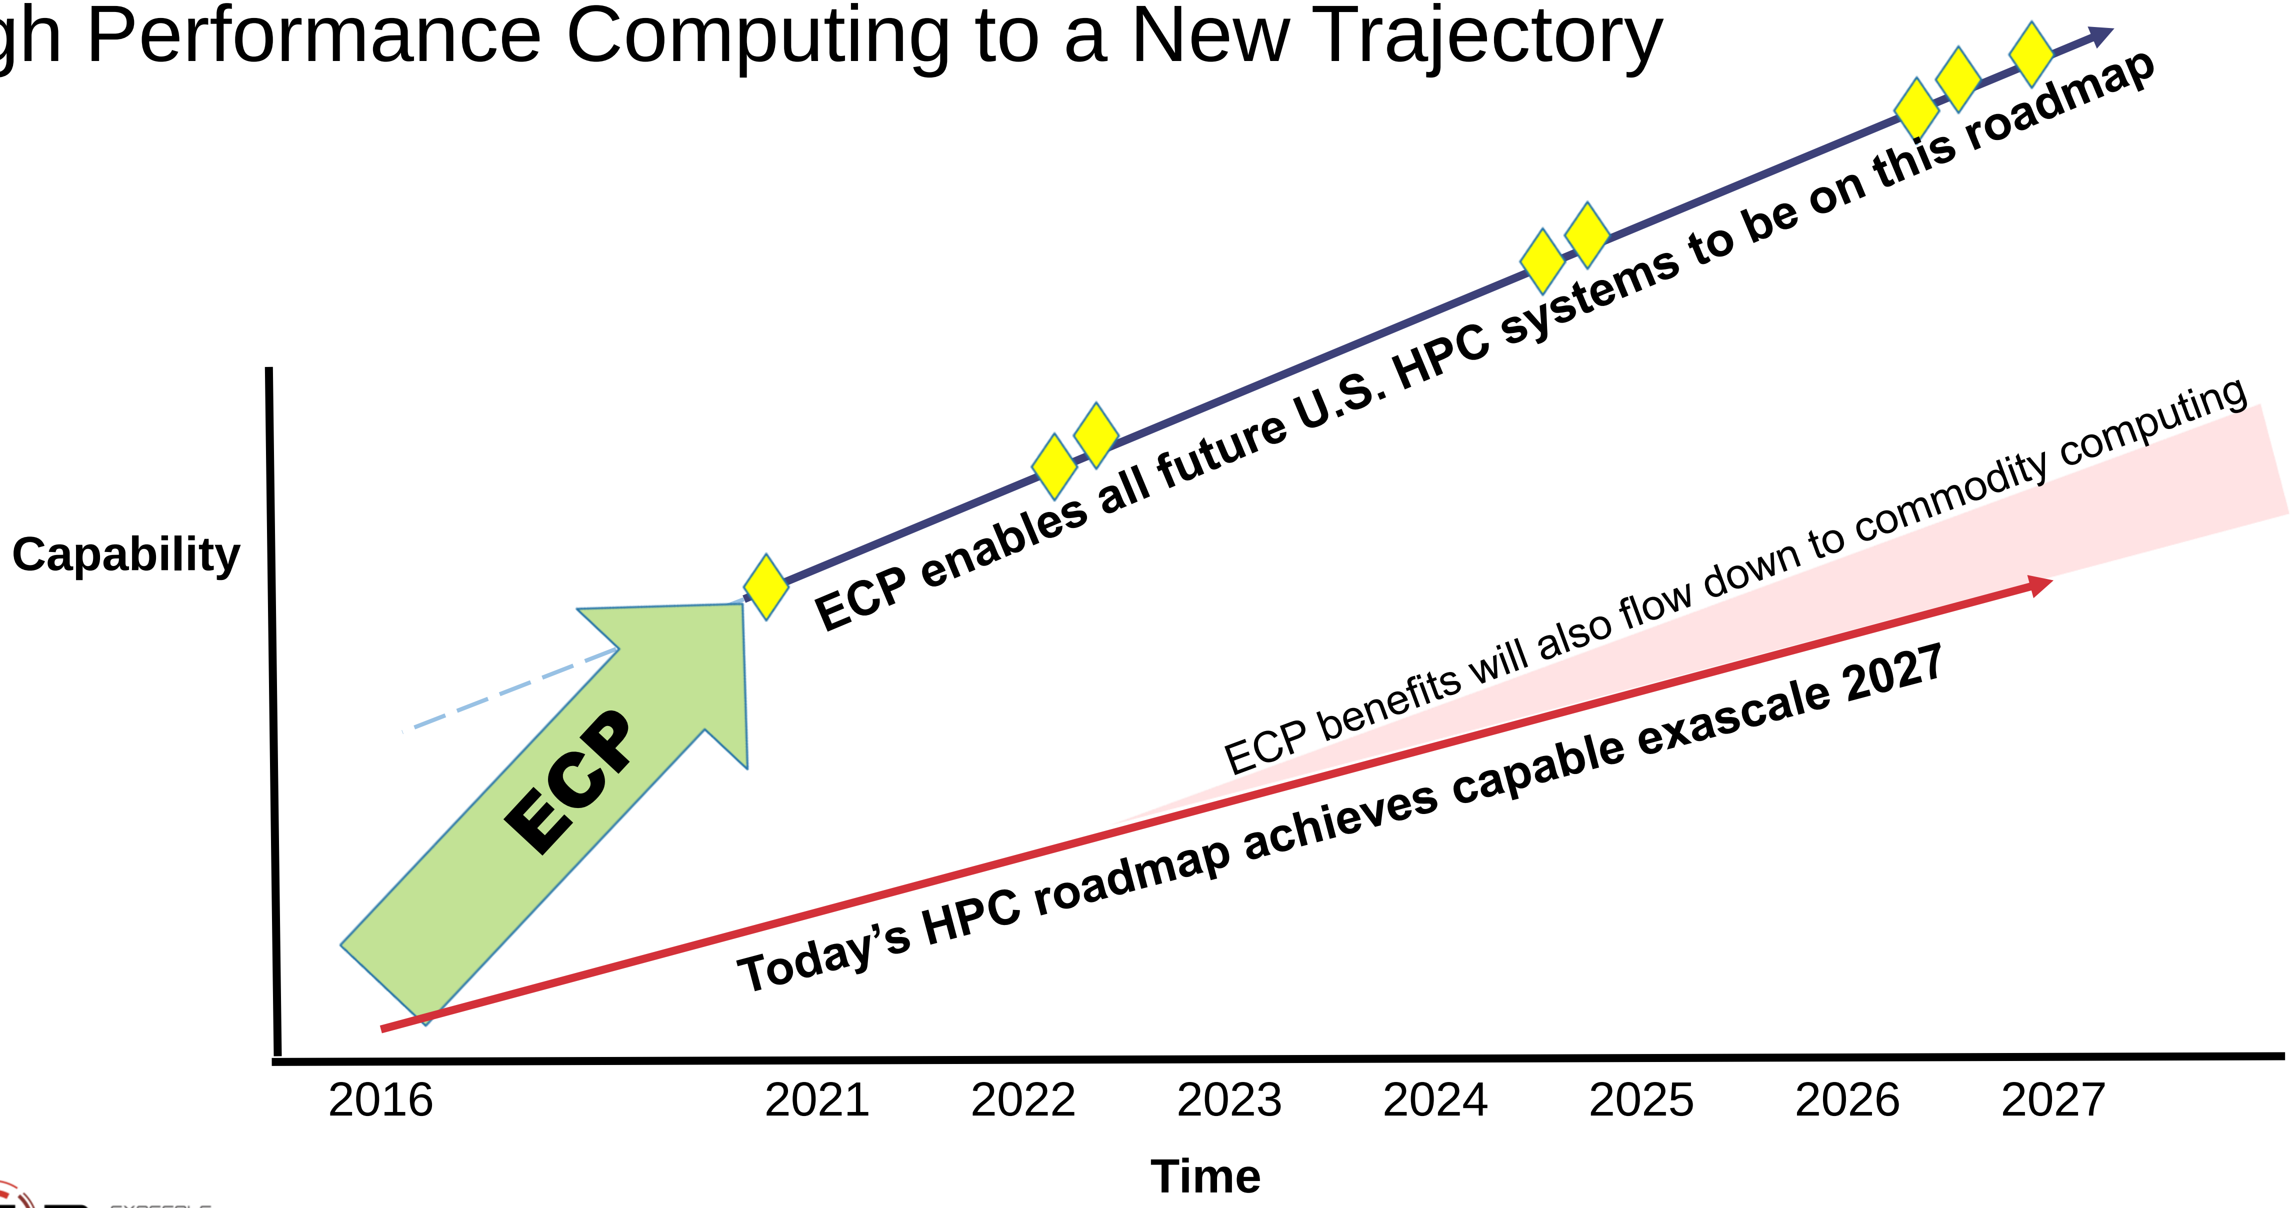

Department of Energy (DOE) Roadmap to Exascale Systems

An impressive, productive lineup of *accelerated node* systems supporting DOE's mission

Pre-Exascale Systems [Aggregate Linpack (Rmax) = 323 PF!]

First U.S. Exascale Systems

2012	2016	2018	2020	2021-2023

2012

ORNL
Cray/AMD/NVIDIA

ANL
IBM BG/Q

LLNL
IBM BG/Q

2016

ANL
Cray/Intel KNL

LBNL
Cray/Intel Xeon/KNL

LANL/SNL
Cray/Intel Xeon/KNL

2018

ORNL
IBM/NVIDIA

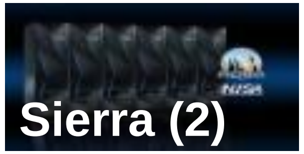

LLNL
IBM/NVIDIA

2020

LBNL
Cray/AMD/NVIDIA

CROSSROADS

LANL/SNL
TBD

2021-2023

FRONTIER

ORNL
TBD

ANL
Intel/Cray

LLNL
TBD

ECP EXASCALE COMPUTING PROJECT

The Summit System @ ORNL

#1 on Top 500

System Performance

- Peak of 200 Petaflops (FP_{64}) for modeling & simulation
- Peak of 3.3 ExaOps (FP_{16}) for data analytics and artificial intelligence
- Max power 13 MW

The system includes

- 4,608 nodes
- Dual-rail Mellanox EDR InfiniBand network
- 250 PB IBM file system transferring data at 2.5 TB/s

Each node has

- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory
  (96 GB HBM2 + 512 GB DDR4)
- 1.6 TB of NV memory

Lawrence Livermore National Laboratory

The Sierra System @ LLNL (#2 on Top 500)

NNSA National Nuclear Security Administration

System specifications

- Peak performance of 125 petaflops for modeling and simulation
- Memory: 1.38 petabytes
- 8,640 Central Processing Units (CPUs)
- 17,280 Graphics Processing Units GPUs)
- Power consumption: 11 megawatts

Each node has

- 2 IBM POWER9 processors
- 4 NVIDIA Tesla V100 GPUs
- 320 GiB of fast memory
  - 256 GiB DDR4
  - 64 GiB HBM2
- 1.6 TB of NVMe memory

The system includes

- 4,320 nodes
- 2:1 tapered Mellanox EDR InfiniBand tree topology (50% global bandwidth) with dual-port HCA per node
- 154 PB IBM Spectrum Scale file system with 1.54 TB/s R/W bandwidth

Sierra supports LLNL's national security mission and ability to advance science in the public interest

Lawrence Livermore National Laboratory
IBM
NVIDIA
Mellanox
redhat

ECP Organization

Board of Directors
Bill Goldstein, Chair (Director, LLNL)
Thomas Zacharia, Vice Chair (Director, ORNL)

Laboratory Operations Task Force (LOTF)

Exascale Computing Project
Doug Kothe, ORNL
Project Director
Lori Diachin, LLNL
Deputy Project Director

Al Geist, ORNL
Chief Technology Officer

Industry Council
Dave Kepczynski, GE, Chair

Julia White, ORNL
Technical Operations

Mike Bernhardt, ORNL
Communications

Project Management
Kathlyn Boudwin, ORNL
Director
Manuel Vigil, LANL
Deputy Director

Monty Middlebrook
Project Controls & Risk

Doug Collins
IT & Quality

Project Office Support
Megan Fielden, Human Resources
Willy Besancenez, Procurement
Sam Howard, Export Control Analyst
Mike Hulsey, Business Management
Kim Milburn, Finance Officer
Susan Ochs, Partnerships
Michael Johnson, Legal
and Points of Contacts at the Core Laboratories

Application Development
Andrew Siegel, ANL
Director
Erik Draeger, LLNL
Deputy Director

Software Technology
Mike Heroux, SNL
Director
Jonathan Carter, LBNL
Deputy Director

Hardware & Integration
Terri Quinn, LLNL
Director
Susan Coghlan, ANL
Deputy Director

DOE HPC Facilities

Core Laboratories

Oak Ridge National Laboratory
Argonne National Laboratory
Berkeley Lab – Lawrence Berkeley National Laboratory
Lawrence Livermore National Laboratory
Los Alamos National Laboratory
Sandia National Laboratories

ECP Exascale Computing Project

ECP Industry Council

Mission

Meet to provide advice and feedback to the ECP Director on:

- ECP project scope and strategic direction
- Technical approaches,
- Progress on milestones and accomplishments
- Industrial requirements,
- Impact on industrial competitiveness

Membership

Executives of U.S. companies (generally VPs of R&D, CTOs, CIOs) for whom HPC is a critical research and production tool

Executives of U.S. independent software vendors (ISVs)

DOE projects and programs historically fall short in executing mutually beneficial collaborations with industry. ECP is committed to working closely with industry: we believe it will make us both better.

EXASCALE COMPUTING PROJECT

ECP Industry Council Member Organizations

ECP Work Breakdown Structure (WBS)

Key leaders at WBS Level 1, 2, 3

Exascale Computing Project 2.0 — Kothe (ORNL)

- **Project Management 2.1 — Boudwin (ORNL)**
  - Project Planning and Management 2.1.1 — Boudwin (ORNL)
  - Project Controls and Risk Management 2.1.2 — Middlebrook (ORNL)
  - Business Management 2.1.3 — Hulsey (ORNL)
  - Procurement Management 2.1.4 — Besancenez (ORNL)
  - Information Technology and Quality Management 2.1.5 — Collins (ORNL)
  - Communications and Outreach 2.1.6 — Bernhardt (ORNL)
- **Application Development 2.2 — Siegel (ANL)**
  - Chemistry and Materials Applications 2.2.1 — Deslippe (LBL)
  - Energy Applications 2.2.2 — Evans (ORNL)
  - Earth and Space Science Applications 2.2.3 — Dubey (ANL)
  - Data Analytics and Optimization Applications 2.2.4 — Hart (SNL)
  - National Security Applications 2.2.5 — Francois (LANL)
  - Co-Design 2.2.6 — Colella (LBL)
- **Software Technology 2.3 — Heroux (SNL)**
  - Programming Models and Runtimes 2.3.1 — Thakur (ANL)
  - Development Tools 2.3.2 — Vetter (ORNL)
  - Mathematical Libraries 2.3.3 — McInnes (ANL)
  - Data and Visualization 2.3.4 — Ahrens (LANL)
  - Software Ecosystem and Delivery 2.3.5 — Munson (ANL)
  - NNSA Software Technologies 2.3.6 — Neely (ANL)
- **Hardware and Integration 2.4 — Quinn (LLNL)**
  - PathForward 2.4.1 — de Supinski (LLNL)
  - Hardware Evaluation 2.4.2 — Hammond (SNL)
  - Application Integration at Facilities 2.4.3 — Hill (ORNL)
  - Software Deployment at Facilities 2.4.4 — Montoya (LANL)
  - Facility Resource Utilization 2.4.5 — White (ORNL)
  - Training and Productivity 2.4.6 — Barker (ORNL)

ECP would fail without its world-class Principal Investigators (PIs) leading the 100+ WBS Level 4 projects responsible for executing ECP's RD&D activities

The three technical areas in ECP have the necessary components to meet national goals

Performant mission and science applications @ scale

Foster application development	Ease of use	Diverse architectures	HPC leadership

Application Development (AD)	Software Technology (ST)	Hardware and Integration (HI)
Develop and enhance the predictive capability of applications critical to the DOE	Produce expanded and vertically integrated software stack to achieve full potential of exascale computing	Integrated delivery of ECP products on targeted systems at leading DOE computing facilities

20+ applications ranging from national security, to energy, earth systems, economic security, materials, and data

80+ unique software products spanning programming models and run times, math libraries, data and visualization

6 vendors supported by PathForward focused on memory, node, connectivity advancements; deployment to facilities

The three technical areas in ECP have the necessary components to address these challenges and meet national goals

Performant mission and science applications @ scale

Foster application development	Ease of use	Diverse architectures	HPC leadership

Application Development (AD)

Develop and enhance the predictive capability of applications critical to the DOE

Software Technology (ST)

Produce expanded and vertically integrated software stack to achieve full potential of exascale computing

Hardware and Integration (HI)

Integrated delivery of ECP products on targeted systems at leading DOE computing facilities

20+ applications ranging from national security, to energy, earth systems, economic security, materials, and data

80+ unique software products spanning programming models and run times, math libraries, data and visualization

6 vendors supported by PathForward focused on memory, node, connectivity advancements; deployment to facilities

EXASCALE COMPUTING PROJECT

ECP Application Development: exascale-capable modeling, simulation, data

Goal

Develop and enhance the predictive capability of applications critical to DOE across science, energy, and national security mission space

- Targeted development of requirements-based methods
- Integration of software and hardware via co-design methodologies
- Systematic improvement of exascale readiness and utilization
- Demonstration and assessment of effective software integration

Chemistry and Materials	Earth and Space Science	Energy	Data Analytics and Optimization	National Security	Co-Design

Hardware realities are forcing new thinking of algorithmic implementations and the move to new algorithms

Algorithmic Implementations

- Reduced communication/data movement
  - Sparse linear algebra, Linpack, etc.
- Much greater locality awareness
  - Likely must be exposed by programming model
- Much higher cost of global synchronization
  - Favor maxim asynchrony where physics allows
- Value to mixed precision where possible
  - Huge role in AI, harder to pin down for PDEs
- Fault resilience?
  - Likely handled outside of applications

New Algorithms

- Adopting Monte Carlo vs. Deterministic approaches
- Exchanging on-the-fly recomputation vs. data table lookup (e.g. neutron cross sections)
- Moving to higher-order methods (e.g. CFD)
- Particle algorithms that favor collecting similar events together rather than parallelism though individual histories

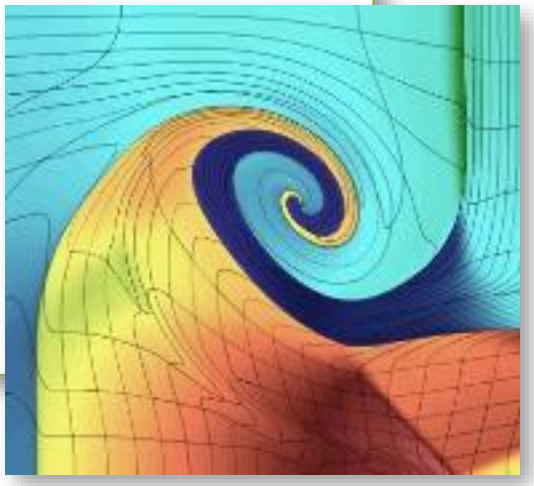

Exascale applications target US national problems in 6 strategic areas

National security

- Stockpile stewardship
- Next-generation electromagnetics simulation of hostile environment and virtual flight testing for hypersonic re-entry vehicles

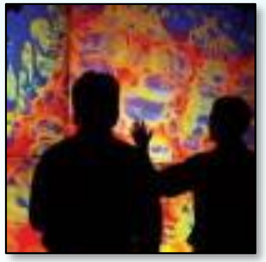

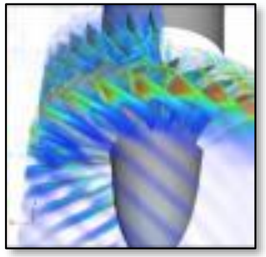

Energy security

- Turbine wind plant efficiency
- High-efficiency, low-emission combustion engine and gas turbine design
- Materials design for extreme environments of nuclear fission and fusion reactors
- Design and commercialization of Small Modular Reactors
- Subsurface use for carbon capture, petroleum extraction, waste disposal
- Scale-up of clean fossil fuel combustion
- Biofuel catalyst design

Economic security

- Additive manufacturing of qualifiable metal parts
- Reliable and efficient planning of the power grid
- Seismic hazard risk assessment

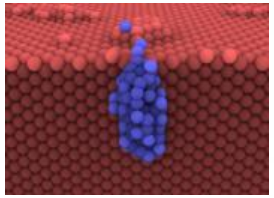

Scientific discovery

- Find, predict, and control materials and properties
- Cosmological probe of the standard model of particle physics
- Validate fundamental laws of nature
- Demystify origin of chemical elements
- Light source-enabled analysis of protein and molecular structure and design
- Whole-device model of magnetically confined fusion plasmas

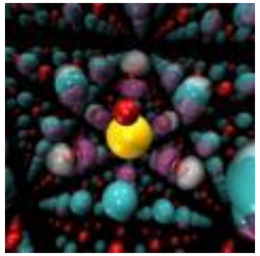

Earth system

- Accurate regional impact assessments in Earth system models
- Stress-resistant crop analysis and catalytic conversion of biomass-derived alcohols
- Metagenomics for analysis of biogeochemical cycles, climate change, environmental remediation

Health care

- Accelerate and translate cancer research

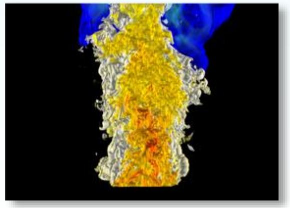

Application Co-Design

Develop efficient exascale libraries that address computational motifs common to multiple application projects

CODAR	Advance understanding of the constraints, mappings, and configuration choices that determine interactions of applications, data analysis and reduction, and exascale platforms
COPA	Create co-designed numerical recipes for particle-based methods that meet application team requirements within design space of STs and subject to constraints of exascale platforms
AMReX	Build framework to support development of block-structured adaptive mesh refinement algorithms for solving systems of partial differential equations on exascale architectures
CEED	Develop next-generation discretization software and algorithms that will enable a wide range of finite element applications to run efficiently on future hardware
ExaGraph	Develop methods and techniques for efficient implementation of key combinatorial (graph) algorithms
ExaLearn	Target learning methods to aid application and experimental facility workflows: deep neural networks (RNNs, CNNs, GANs), kernel & tensor methods, decision trees, ensemble methods, graph models, reinforcement learning
Proxy Apps	Improve the quality of proxy applications created by ECP and maximize the benefit received from their use. Maintain and distribute ECP Proxy App Suite.

Center for Efficient Exascale Discretizations (CEED)

Co-Design of unstructured mesh, FE-based PDE discretizations

Goal

- Develop algorithms and software to enable more efficient HPC simulations in a wide range of PDE-based science applications.
- Focus on next-generation discretization methods: high-order finite elements on general unstructured grids.
- Target high performance on a variety of hardware: CPU, GPU, A21 in a flexible and user-friendly way.

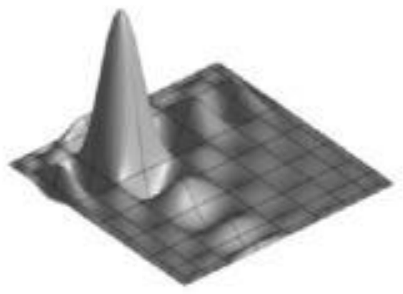

$$A = P^T G^T B^T D B G P$$

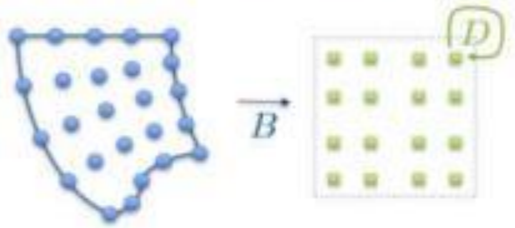

Approach

- Performance-enabling math foundation: high-order operator decomposition
- Fast kernels: CEED benchmarks, combine expertise, engage community
- Library integration: high-level API (MFEM, Nek5000), low-level API (libCEED)
- Application engagement: liaisons, CEED miniapps (Nekbone, Laghos)
- Collaborate with ECP/ST, broader community (SciDAC, xSDK, deal.ii, …)
- High-order software ecosystem: operator format, FMS, matrix-free solvers

PI: Tzanio Kolev (LLNL)

CEED is targeting several ECP applications

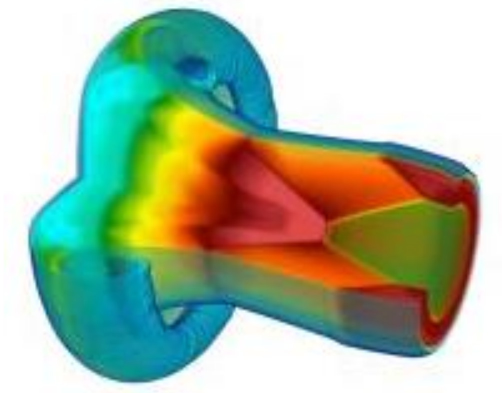

Compressible flow (MARBL)

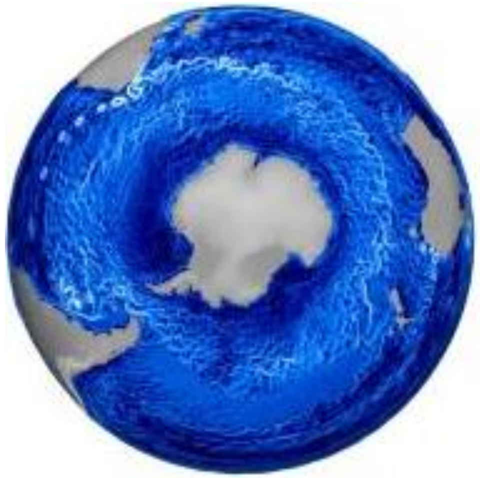

Climate (E3SM)

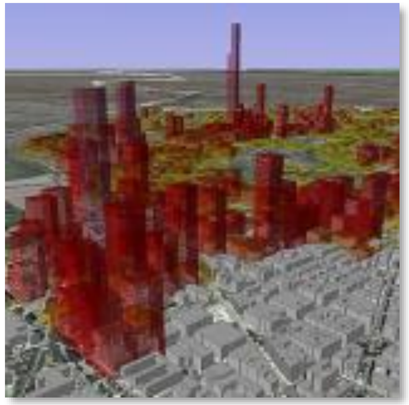

Urban systems (Urban)

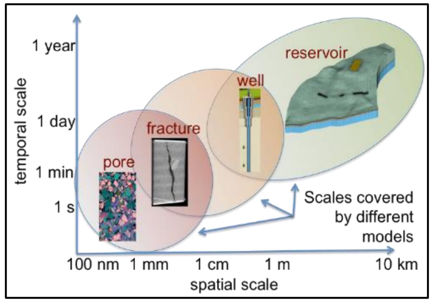

Subsurface (GEOS)

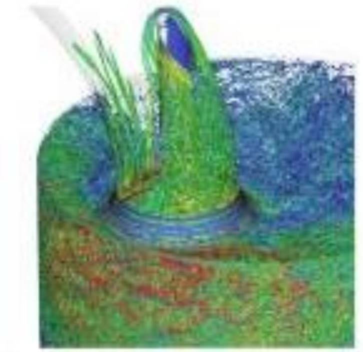

Combustion (Nek5000)

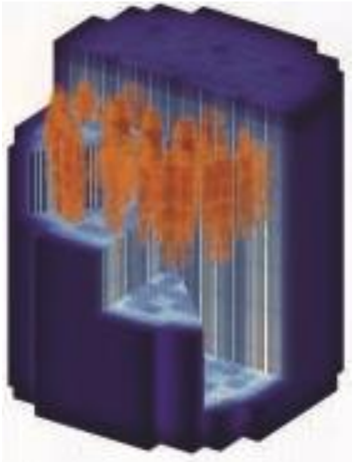

Modular Nuclear Reactors (ExaSMR)

Wind Energy (ExaWind)

Additive Manufacturing (ExaAM)

Magnetic Fusion (WDMApp)

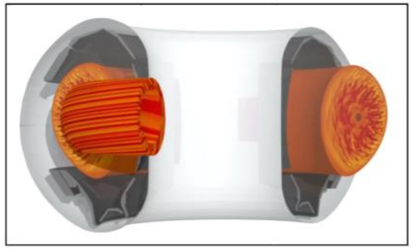

EXASCALE COMPUTING PROJECT

PI: Tzanio Kolev (LLNL)

ECP’s Adaptive Mesh Refinement Co-Design Center: AMReX

- Develop and deploy software to support block-structured adaptive mesh refinement on exascale architectures
  - Core AMR functionality
  - Particles coupled to AMR meshes
  - Embedded boundary (EB) representation of complex geometry
  - Linear solvers
  - Supports two modalities of use
    - Library support for AMR
    - Framework for constructing AMR applications
- Provide direct support to ECP applications that need AMR for their application
- Evaluate software technologies and integrate with AMReXwhen appropriate
- Interact with hardware technologies / vendors

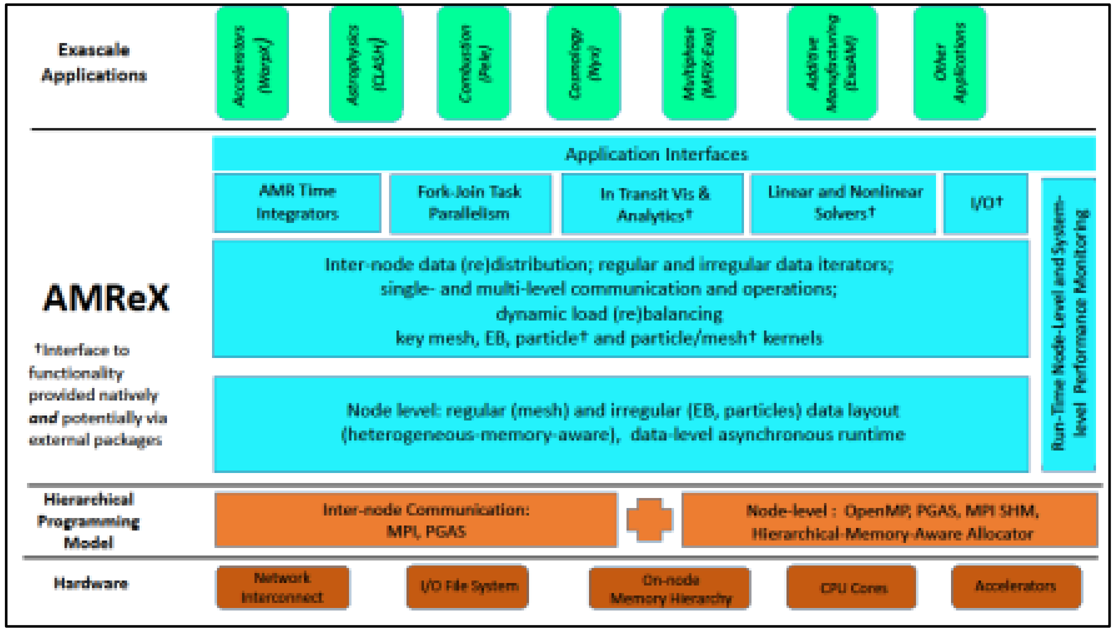

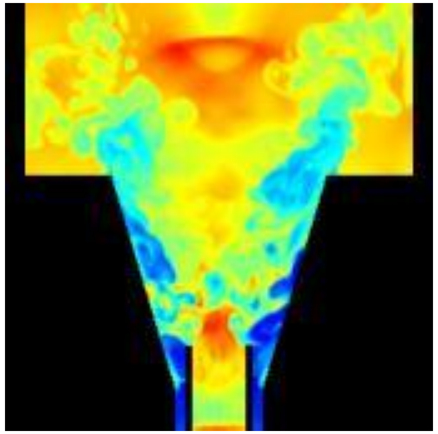

Application	Particles	ODEs	Linear Solvers	EB
Combustion	X	X	X	X
Multiphase	X		X	X
Cosmology	X	X	X	
Astrophysics	X	X	X	
Accelerators	X			

PI: John Bell (LBNL)

ECP's Co-Design Center for Online Data Analysis and Reduction

CODAR

Traditional approach:
Compute…output…analyze [offline]

Write simulation output to secondary storage; read back for analysis

Decimate in time when simulation output rate exceeds output rate of computer

New approach:
Online data analysis and reduction

Co-optimize simulation, analysis, reduction for performance and information output

Substitute CPU cycles for I/O, via online data (de)compression and/or online data analysis

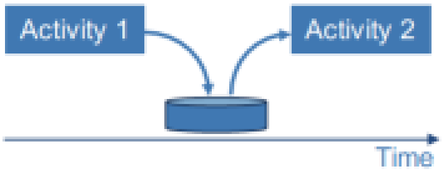

Provide the right information at the right time and place to accelerate discovery!

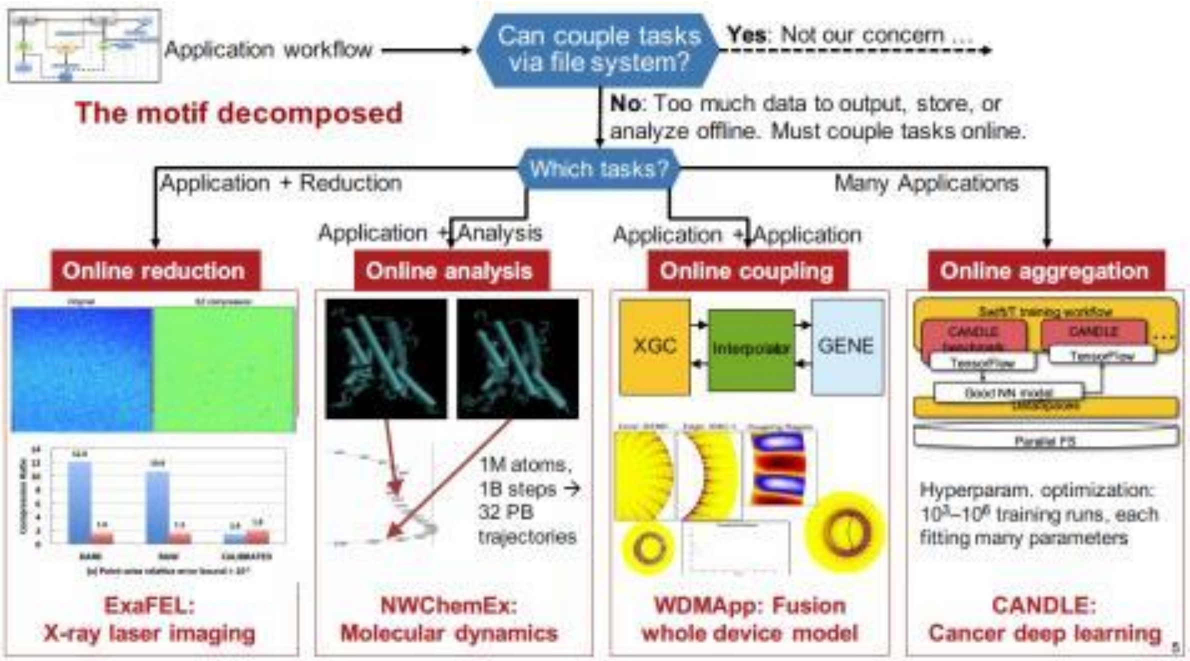

<u>Goal:</u> Replace the activities in HPC workflow that have been mediated through file I/O with in-situ methods / workflows. data reduction, analysis, code coupling, aggregation (e.g. parameter studies).

<u>Cross-cutting tools:</u>

- Workflow setup, manager (Cheetah, Savanna); Data coupler (ADIOS-SST); Compression methods (MGARD, FTK, SZ), compression checker (Z-checker)
- Performance tools (TAU, Chimbuco, SOSFlow)

ECP's Co-Design Center for Particle Applications: CoPA

Goal: Develop algorithms and software for particle methods,

Cross-cutting capabilities:

- Specialized solvers for quantum molecular dynamics (Progress / BML).
- Performance-portable libraries for classical particle methods in MD, PDE (Cabana).
- FFT-based Poisson solvers for long-range forces.

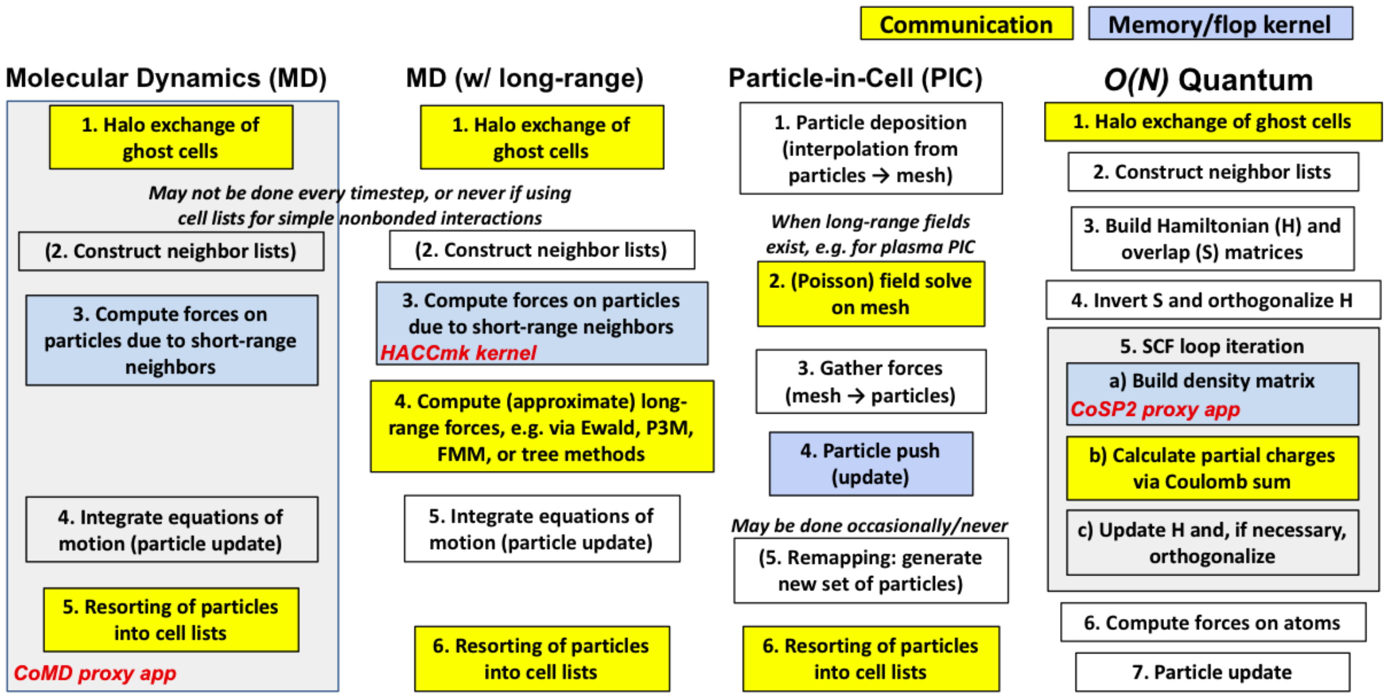

Technical approach:

- High-level C++ APIs, plus a Fortran interface (Cabana).
- Leverage existing / planned FFT software.
- Extensive use of miniapps / proxy apps as part of the development process.

PI: Sue Mniszewski (LANL) recently replacing Tim Germann (LANL), who is taking on a larger role in ECP

ECP's Co-Design Center for Machine Learning: ExaLearn

Bringing together experts from 8 DOE Laboratories

- AI has the potential to accelerate scientific discovery or enable prediction in areas currently too complex for direct simulation (ML for HPC and HPC for ML)
- AI use cases of interest to ECP:
  - *Classification and regression*, including but not limited to image classification and analysis, e.g. scientific data output from DOE experimental facilities or from national security programs.
  - *Surrogate models* in high-fidelity and multiscale simulations, including uncertainty quantification and error estimation.
  - *Structure-to-function relationships*, including genome-to-phenome, the prediction of materials performance based on atomistic structures, or the prediction of performance margins based on manufacturing data.
  - *Control systems*, e.g., for wind plants, nuclear power plants, experimental steering and autonomous vehicles.
  - *Inverse problems* and optimization. This area would include, for example, inverse imaging and materials design.
- Areas in need of research
  - Data quality and statistics
  - Learning algorithms
  - Physics-Informed AI
  - Verification and Validation
  - Performance and scalability
  - Workflow and deployment

Expected Work Product: A Toolset That . . .

- Has a line-of-sight to exascale computing, e.g. through using exascale platforms directly, or providing essential components to an exascale workflow
- Does not replicate capabilities easily obtainable from existing, widely-available packages
- Builds in domain knowledge where possible "Physics"-based ML and AI
- Quantifies uncertainty in predictive capacity
- Is interpretable
- Is reproducible
- Tracks provenance

PI: Frank Alexander (BNL)

Machine Learning in the Light Source Workflow

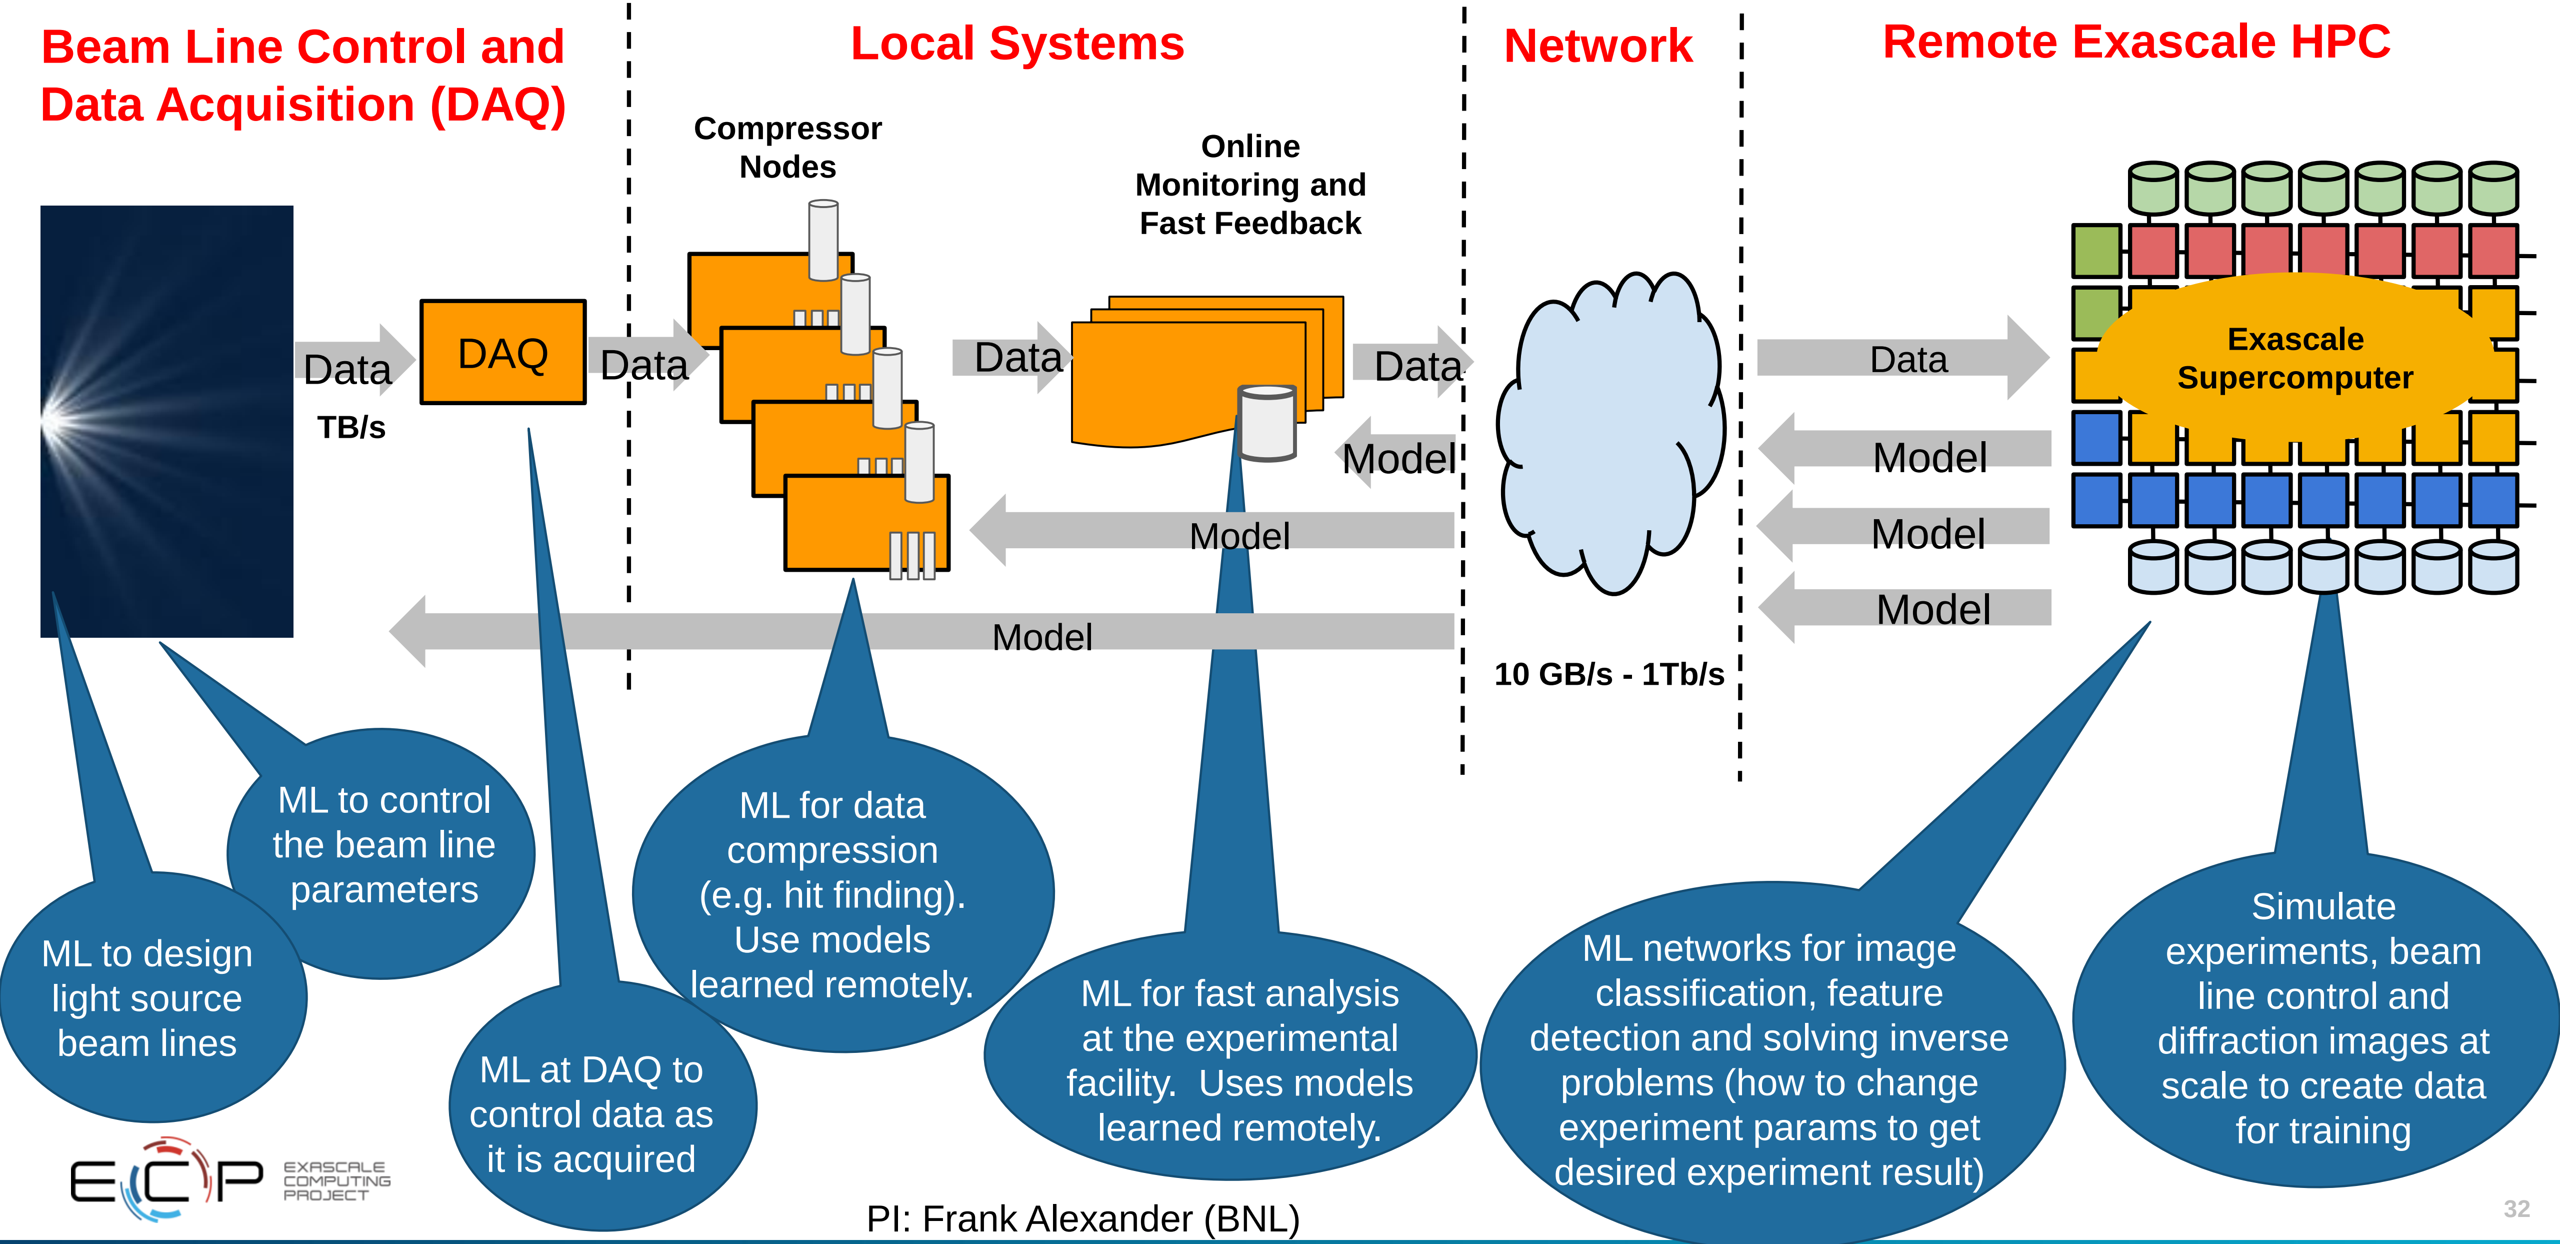

EXASCALE COMPUTING PROJECT

PI: Frank Alexander (BNL)

Exascale apps can deliver transformative products and solutions

ExaWind

Turbine Wind Plant Efficiency
(Mike Sprague, NREL)

- Harden wind plant design and layout against energy loss susceptibility
- Increase penetration of wind energy

Challenges: linear solver perf in strong scale limit; manipulation of large meshes; overset of structured & unstructured grids; communication-avoiding linear solvers

ExaAM

Additive Manufacturing (AM) of Qualifiable Metal Parts
(John Turner, ORNL)

- Accelerate the widespread adoption of AM by enabling routine fabrication of qualifiable metal parts

Challenges: capturing unresolved physics; multi-grid linear solver performance; coupled physics

EQSIM

Earthquake Hazard Risk Assessment
(David McCallen, LBNL)

- Replace conservative and costly earthquake retrofits with safe purpose-fit retrofits and designs

Challenges: full waveform inversion algorithms

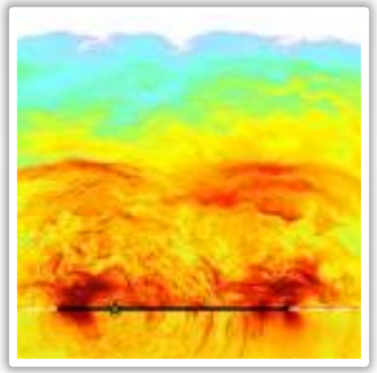

EQSIM: Understanding and predicting earthquake phenomenon

Site ground motions

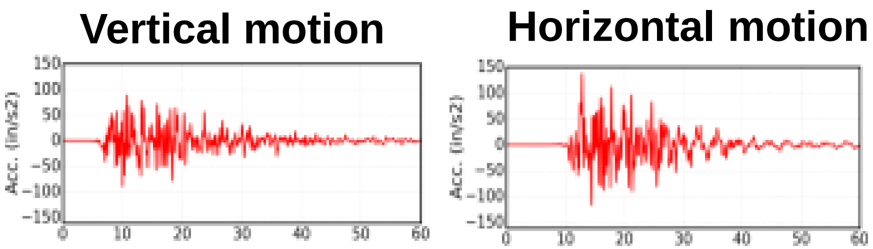

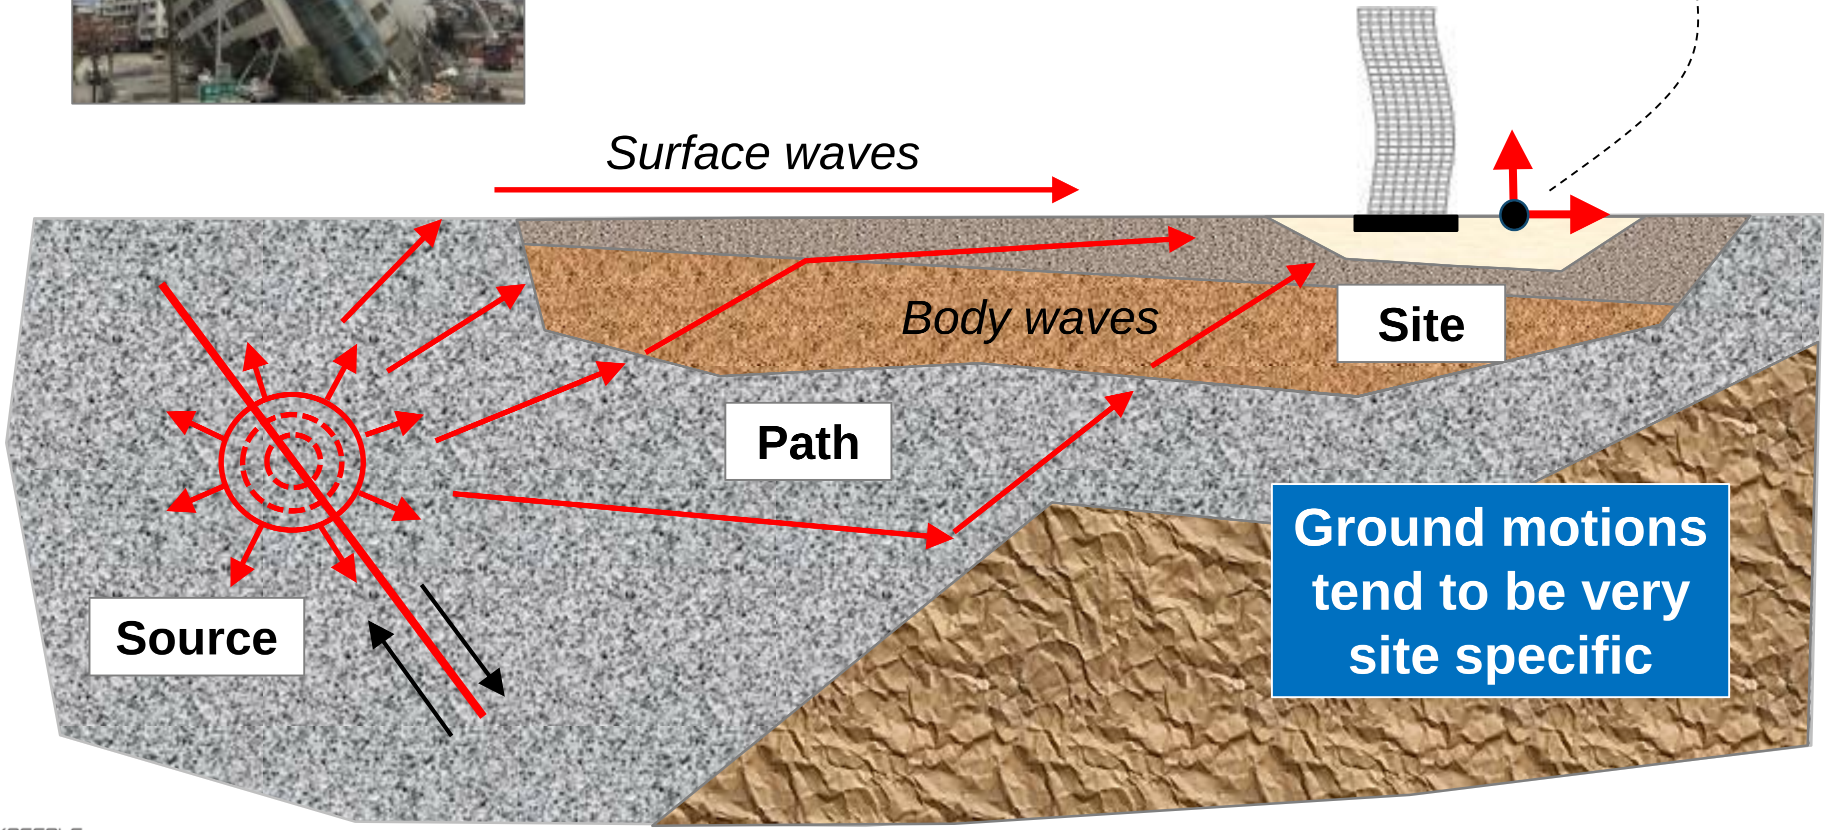

Ground motions tend to be very site specific

PI: David McCallen (LBNL)

EQSIM: The Exascale "Big Lift" – regional ground motion simulations at engineering frequencies

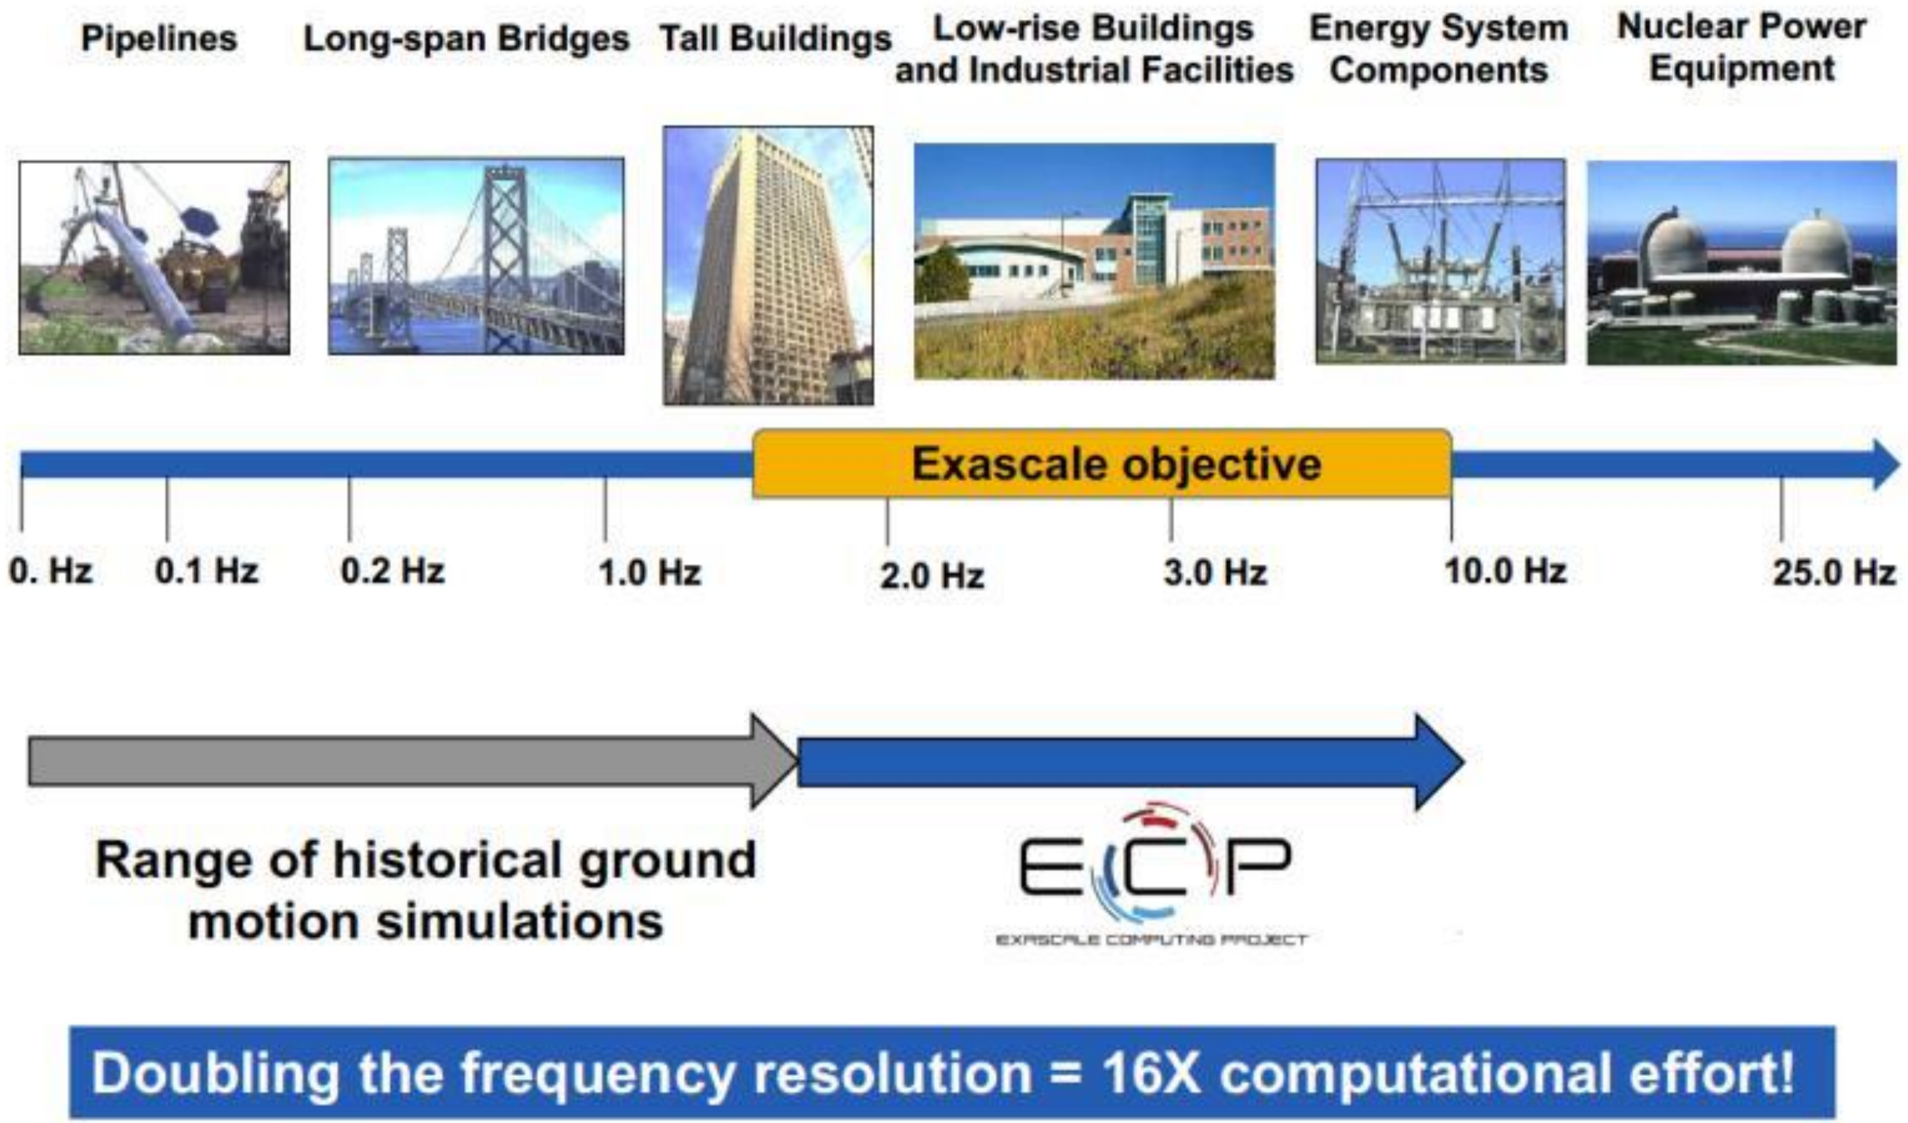

PI: David McCallen (LBNL)

EQSIM: Advancing geophysics and infrastructure applications

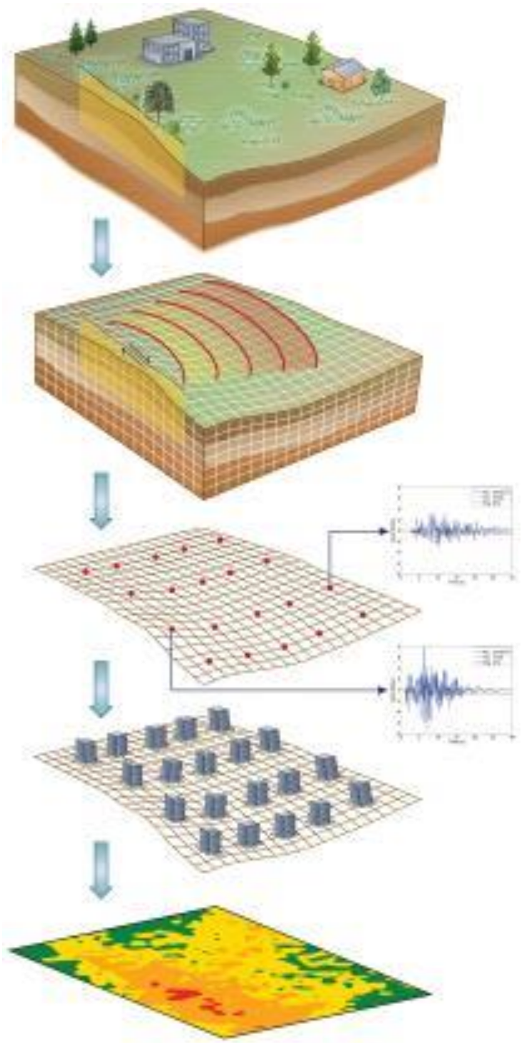

Earthquake Hazard

Earthquake Risk

***SW4* – 4th order finite difference geophysics code for wave propagation**

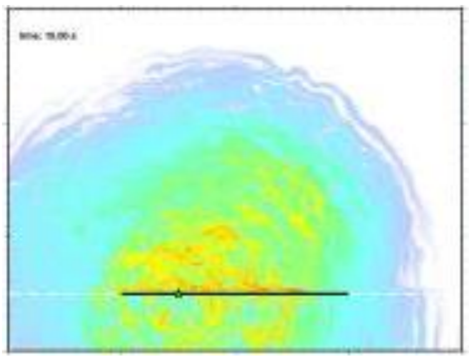

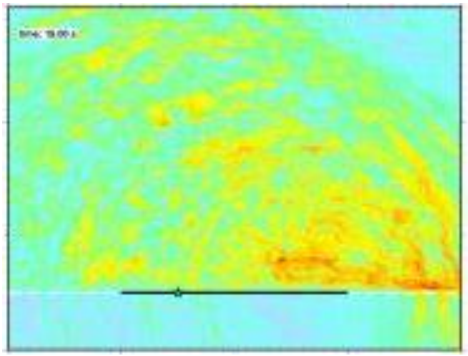

***NEVADA & MSESSI* – finite deformation, inelastic Finite Element codes for structures and soils**

Weak Coupling

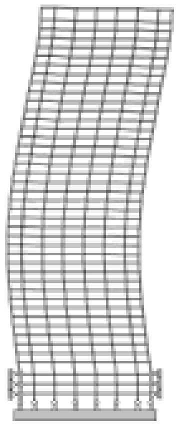

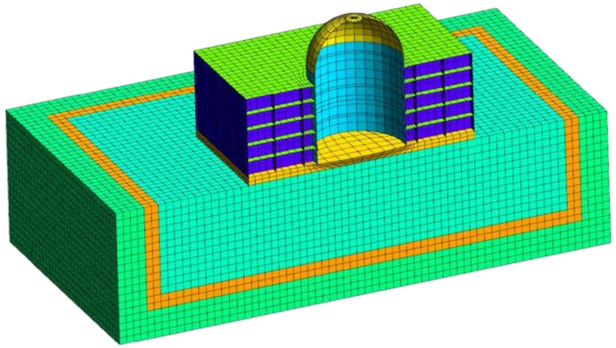

Strong Coupling

EXASCALE COMPUTING PROJECT

PI: David McCallen (LBNL)

EQSIM: Using RAJA to achieve performance portability

- RAJA is a C++ abstraction layer developed at LLNL.
- Same C++ source code for OpenMP and CUDA backends
  - Machine specific options in a policy file
- Coding complexity similar to OpenMP
- Currently running on the Sierra GPU machine at LLNL with low overhead
- August-2018: 1,024 nodes of Sierra, 4,096 GPUs, 6.9 Hz, giving an overall performance (Figure or Merit) improvement of 24.2

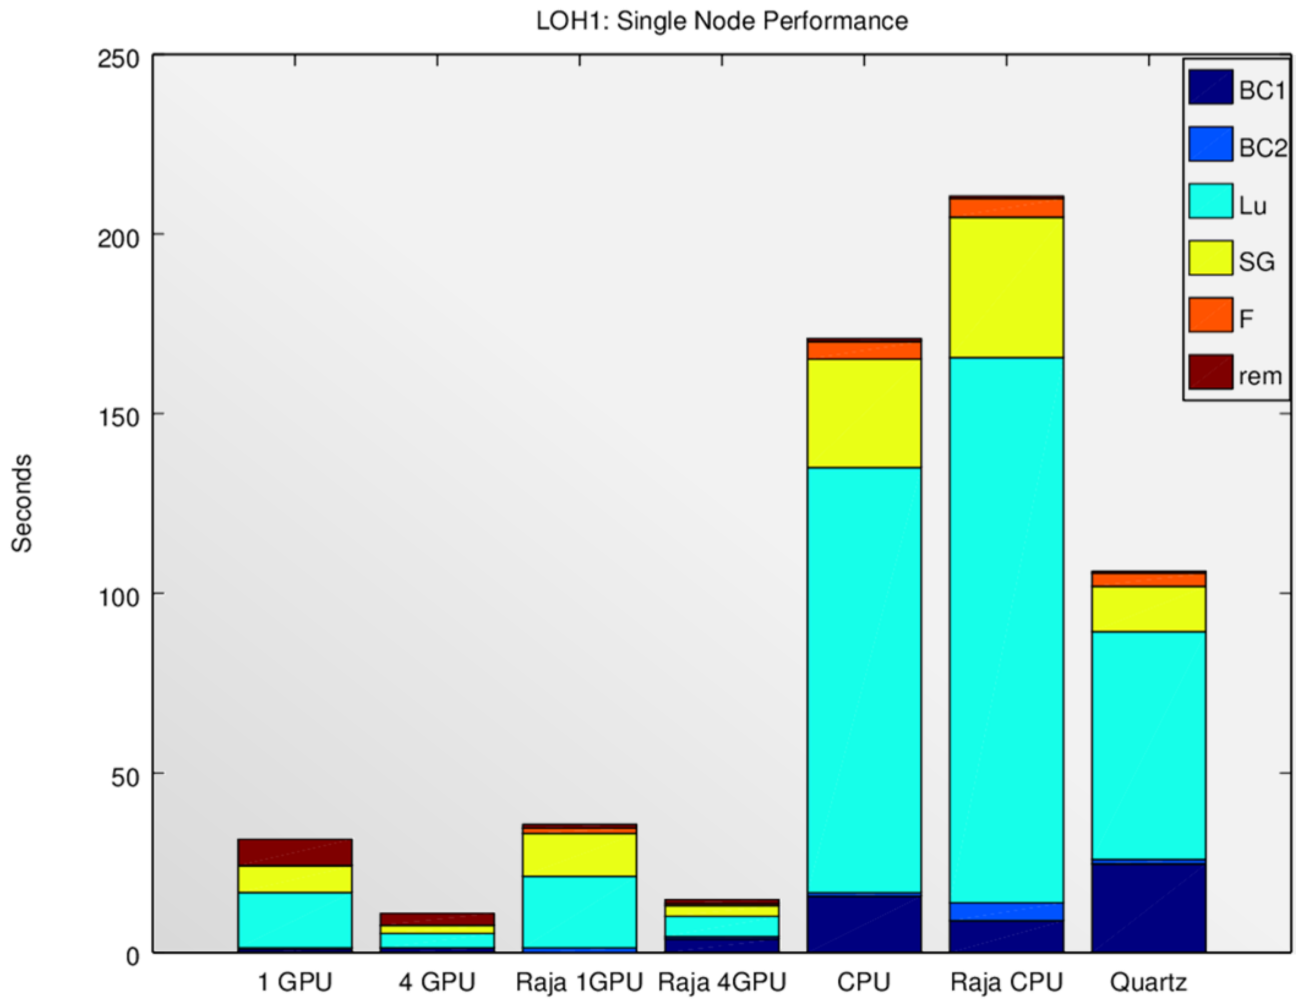

PI: David McCallen (LBNL)

Exascale apps can deliver transformative products and solutions

MFIX-Exa

Scale-up of Clean Fossil Fuel Combustion
(Madhava Syamlal, NETL)

- Commercial-scale demonstration of transformational energy technologies – curbing CO_2 emissions at fossil fuel power plants by 2030

Challenges: load balancing; strong scaling thru transients

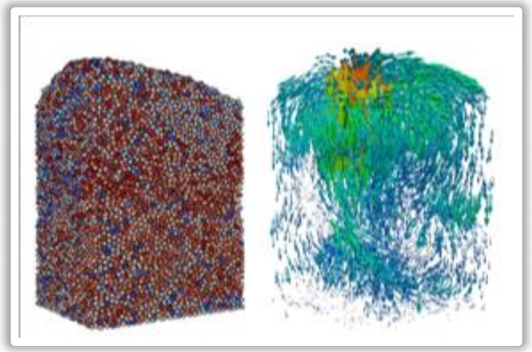

GAMESS

Biofuel Catalyst Design
(Mark Gordon, Ames)

- Design more robust and selective catalysts orders of magnitude more efficient at temperatures hundreds of degrees lower

Challenges: weak scaling of overall problem; on-node performance of molecular fragments

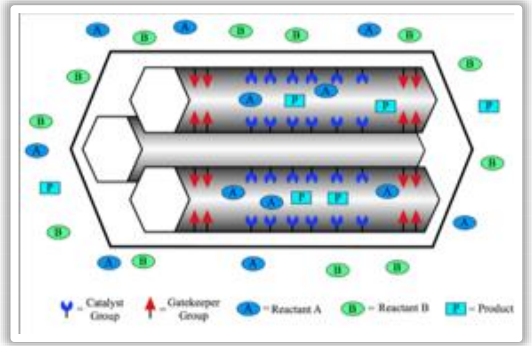

EXAALT

Materials for Extreme Environments
(Danny Perez, LANL)

- Simultaneously address time, length, and accuracy requirements for predictive microstructural evolution of materials

Challenges: SNAP kernel efficiency on accelerators; efficiency of DFTB application on accelerators

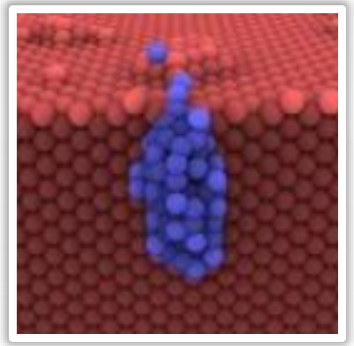

Exascale apps can deliver transformative products and solutions

ExaSMR

Design and Commercialization of Small Modular Reactors
(Steve Hamilton, ORNL)

- Virtual test reactor for advanced designs via experimental-quality simulations of reactor behavior

Challenges: existing GPU-based MC algorithms require rework for hardware less performant for latency-bound algorithms with thread divergence; performance portability with OCCA & OpenACC not achievable; insufficient node memory for adequate CFD + MC coupling

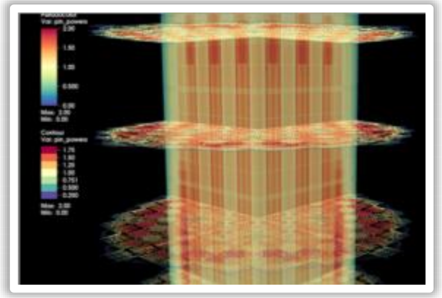

Subsurface

Carbon Capture, Fossil Fuel Extraction, Waste Disposal
(Carl Steefel, LBNL)

- Reliably guide safe long-term consequential decisions about storage, sequestration, and exploration

Challenges: performance of Lagrangian geomechanics; adequacy of Lagrangian crack mechanics) + Eulerian (reaction, advection, diffusion) models; parallel HDF5 for coupling

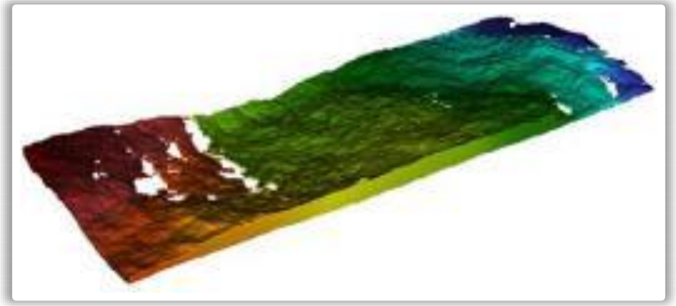

QMCPACK

Materials for Extreme Environments
(Paul Kent, ORNL)

- Find, predict and control materials and properties at the quantum level with unprecedented accuracy for the design novel materials that rely on metal to insulator transitions for high performance electronics, sensing, storage

Challenges: minimizing on-node memory usage; parallel on-node performance of Markov-chain Monte Carlo

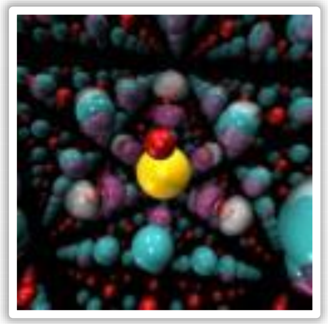

ECP EXASCALE COMPUTING PROJECT

Efficient Monte Carlo on accelerator–based architectures

<u>Challenge</u>: Monte Carlo neutron particle transport is a stochastic method

- Not amenable to single kernel optimization – no "high cost" kernel to optimize
- Independent random walks are not readily amenable to SIMT algorithms
- Sampling data (interaction cross sections) are:
  - randomly accessed
  - characterized by detailed structure
  - in standard applications consist of large point-wise representations (>1 – 5 GB per temperature point)

Distribution of history lengths in SMR core

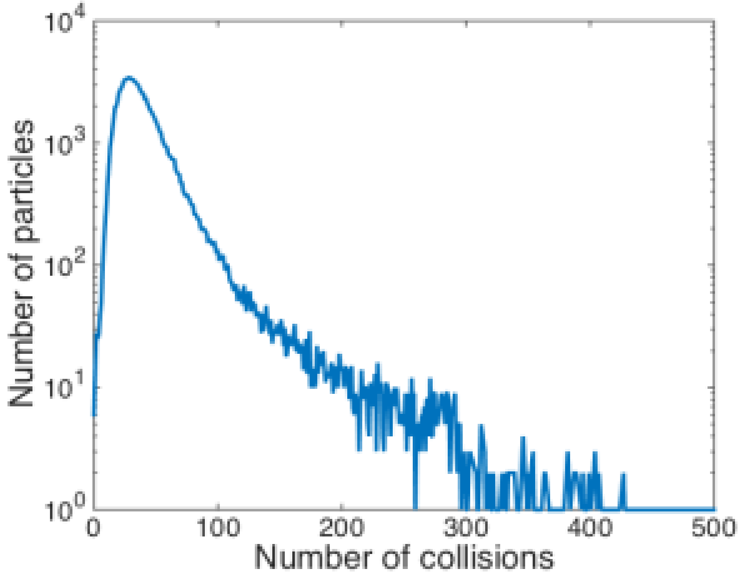

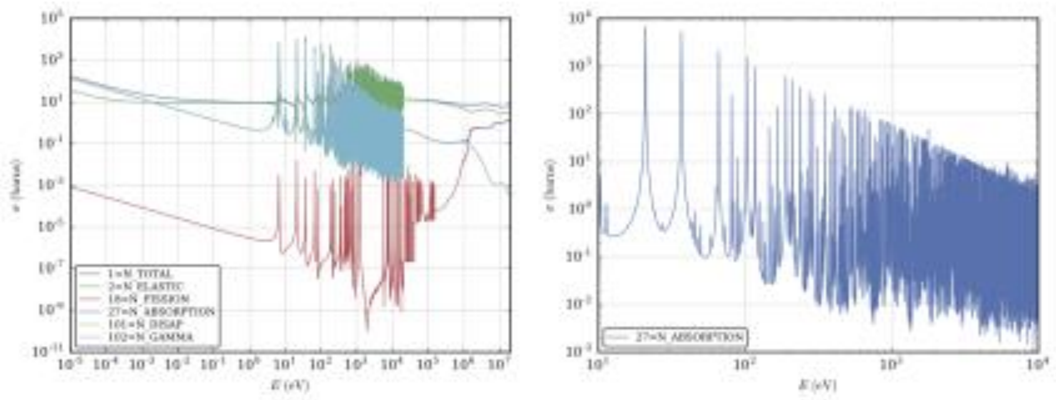

PI: Steve Hamilton (ORNL)

The Monte Carlo algorithm maps well to GPUs after changing from a history-based to an event-based algorithm

- Reduce thread divergence – change from history- to event-based algorithm
- Flatten algorithms to reduce kernel size; smaller kernels = higher occupancy
- Partition events based on fuel and non-fuel regions
- Debuted first comprehensive windowed multipole library for nuclear data (with temp correction)
- MC performance on Summit ~16x that achieved on Titan for the same algorithm
  - Significantly out-pacing gains in machine theoretical peak (7x on LINPACK)
- Overall MC performance has progressed from 15M to 600M particles/s!

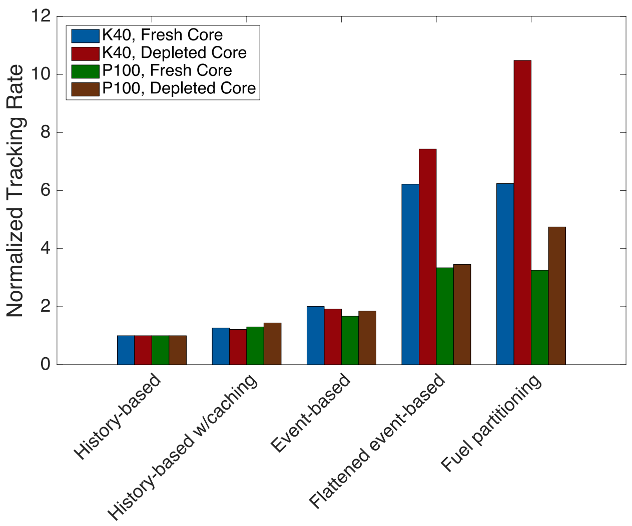

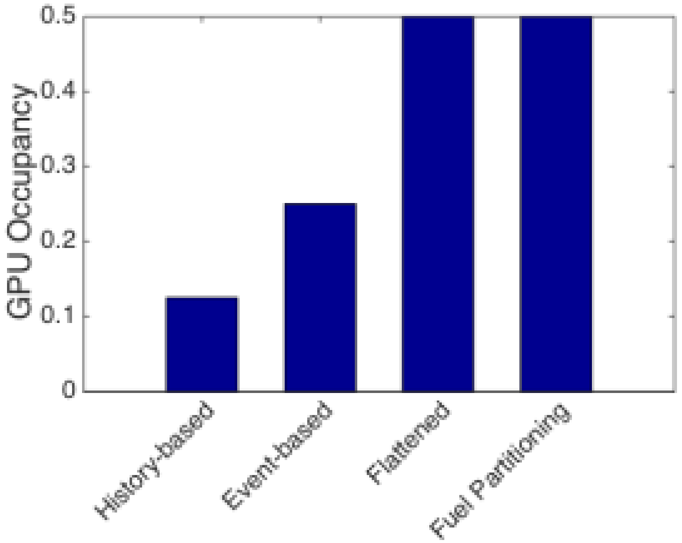

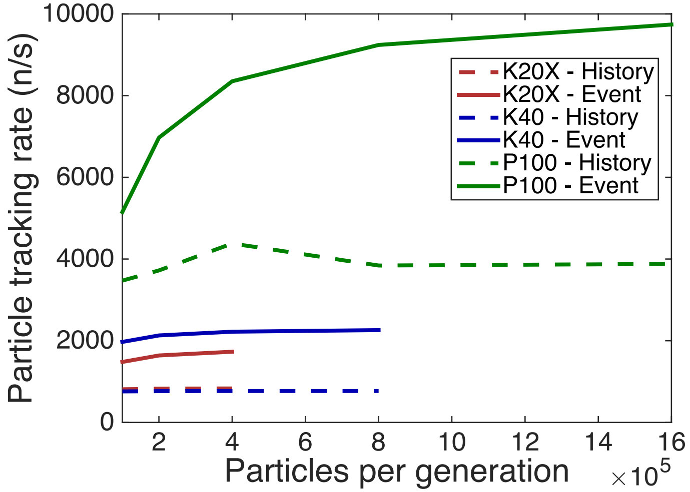

PI: Steve Hamilton (ORNL)

Exascale apps can deliver transformative products and solutions

ExaSGD

Reliable and Efficient Planning of the Power Grid
(Henry Huang, PNNL)

- Optimize power grid planning, operation, control and improve reliability and efficiency

Challenges: parallel performance of nonlinear optimization based on discrete algebraic equations and possible mixed-integer programming

Combustion-PELE

High-Efficiency, Low-Emission Combustion Engine Design
(Jackie Chen, SNL)

- Reduce or eliminate current cut-and-try approaches for combustion system design

Challenges: performance of chemistry ODE integration on accelerated architectures; linear solver performance for low-Mach algorithm; explicit LES/DNS algorithm not stable

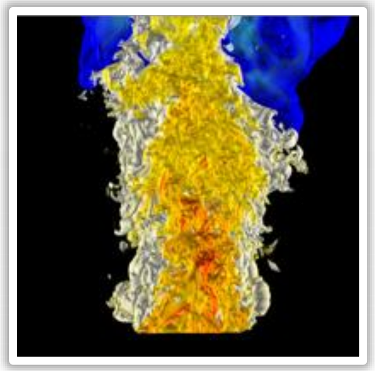

Pele Code Design Overview

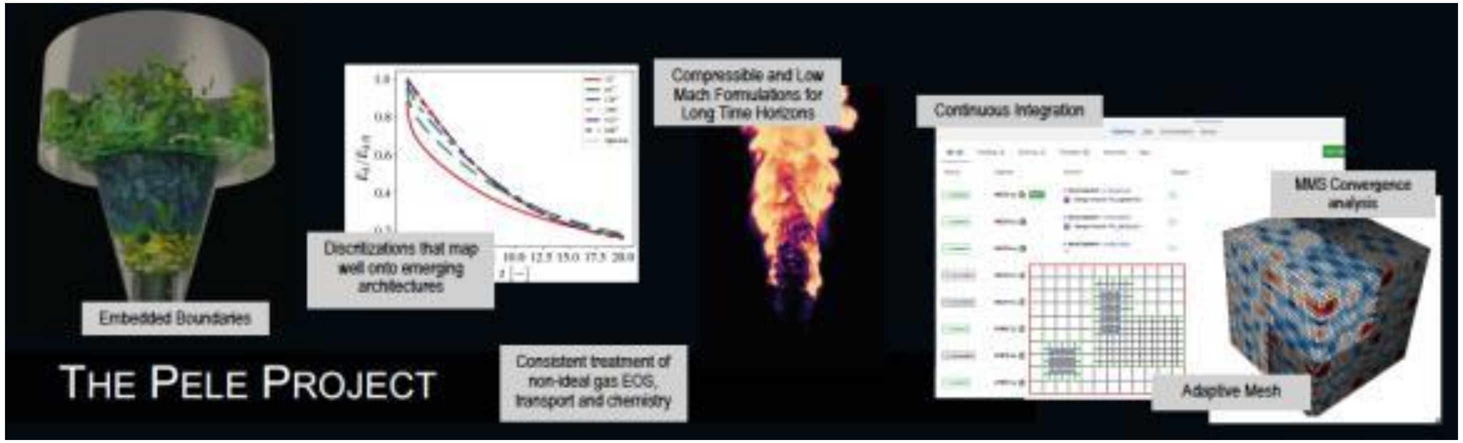

- Baseline algorithm design for multicomponent flowwith stiff reactions, AMR
  - **PeleC**: Comparable advection, diffusion time scales, motivates IMEX-type scheme based on Spectral Deferred Corrections (SDC) with time-implicit chemistry
    - Robust highly efficient time-explicit Godunov-type upwind advection, simple centered diffusion
    - BDF-style implicit chemistry ODE integration, with sources that approximate the other processes
  - **PeleLM**: acoustics filtered away analytically, but still want robust, time-explicit advection
    - Chemistry and diffusion are now time-implicit – iterative timestep simultaneously incorporates flow constraint (constant pressure), mutually coupled species/energy diffusion and chemistry. Entire system evolved stably on slower advection time scales across AMR grid hierarchy
- SDC-based iterative timestep – treats each process essentially independently, with accelerated iteration to couple everything together efficiently
- Robust baseline allows stable, well-behaved extensible time step
  - Switch 2nd order advection scheme with more accurate 4th order algorithm
  - Option for “destiffened” chemistry model that allows highly efficient time-explicit advance
  - Robust to other, potentially stiff, tightly coupled processes, such as sprays, radiation, soot, etc

Exascale apps can deliver transformative products and solutions

E3SM-MMF

Accurate Regional Impact Assessment in Earth Systems
(Mark Taylor, SNL)

- Forecast water resources and severe weather with increased confidence; address food supply changes

Challenges: MMF approach for cloud-resolving model has large biases; adequacy of Fortran MPI+OpenMP for some architectures; Support for OpenMP and OpenACC

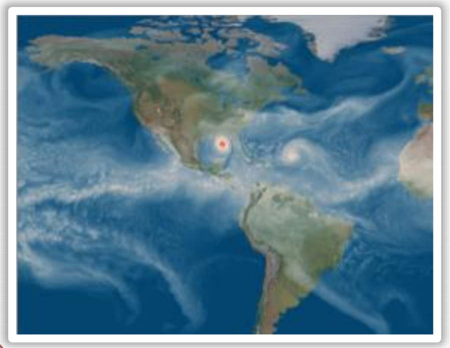

NWChemEx

Catalytic Conversion of Biomass-Derived Alcohols
(Thom Dunning, PNNL)

- Develop new optimal catalysts while changing the current design processes that remain costly, time consuming, and dominated by trial-and-error

Challenges: computation of energy gradients for coupled-cluster implementation; on- and off-node performance

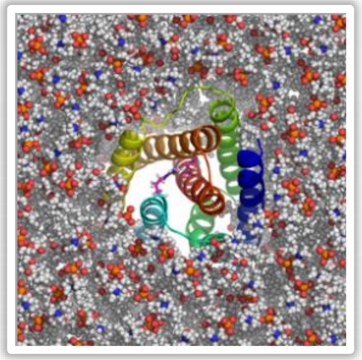

ExaBiome

Metagenomics for Analysis of Biogeochemical Cycles
(Kathy Yelick, LBNL)

- Discover knowledge useful for environmental remediation and the manufacture of novel chemicals and medicines

Challenges: Inability of message injection rates to keep up with core counts; efficient and performant implementation of UPC, UPC++, GASNet; GPU performance; I/O performance

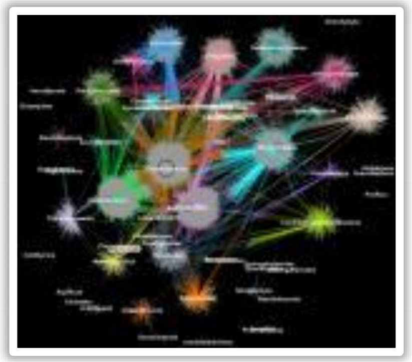

E3SM-Multiscale Modeling Framework (MMF)

Cloud Resolving Climate Model for E3SM

- Develop capability to assess regional impacts of climate change on the water cycle that directly affect the US economy such as agriculture and energy production.
- Cloud resolving climate model is needed to reduce major systematic errors in climate simulations due to structural uncertainty in numerical treatments of convection – such as convective storm systems
- Challenge: cloud resolving climate model using traditional approaches requires *zettascale* resources
- E3SM "conventional" approach:
  - Run the E3SM model with a global cloud resolving atmosphere and eddy resolving ocean.
    - 3 km atmosphere/land (7B grid points) and 15-5 km ocean/ice (1B grid points)
  - Achieve throughput rate of 5 SYPD to perform climate simulation campaigns including a 500 year control simulation
  - Detailed benchmarks on KNL and v100 GPUs show negligible speedups compared to conventional CPUs
    - Low arithmetic intensity of most of the code; throughput requirements lead to strong scaling and low work per node.
- E3SM-MMF: Use a multiscale approach ideal for new architectures to achieve cloud resolving convection on Exascale
  - Exascale will make "conventional" cloud resolving simulations routine for shorter simulations (process studies, weather prediction)
  - For cloud resolving climate simulations, we need fundamentally new approaches to take advantage of exascale resources

Convective storm system nearing the Chicago metropolitan area
http://www.spc.noaa.gov/misc/AbtDerechos/derechofacts.htm

PI: Mark Taylor (SNL)

Exascale apps can deliver transformative products and solutions

ExaSky

Cosmological Probe of the Standard Model of Particle Physics
(Salman Habib, ANL)

- Unravel key unknowns in the dynamics of the Universe: dark energy, dark matter, and inflation

Challenges: subgrid model accuracy; OpenMP performance on GPUs; file system stability and availabilty

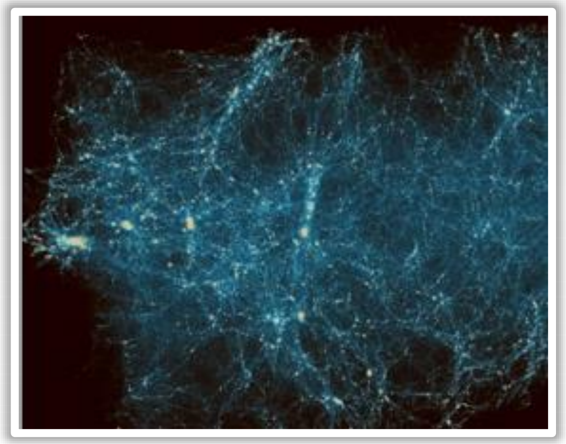

LatticeQCD

Validate Fundamental Laws of Nature
(Andreas Kronfeld, FNAL)

- Correct light quark masses; properties of light nuclei from first principles; <1% uncertainty in simple quantities

Challenges: performance of critical slowing down; reducing network traffic to reduce system interconnect contention; strong scaling performance to mitigate reliance on checkpointing

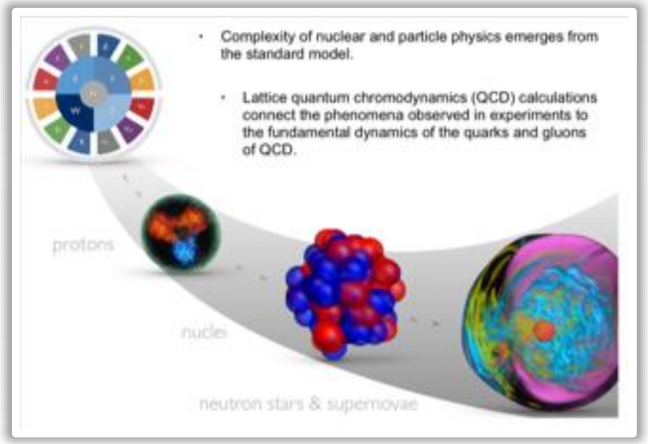

WarpX

Plasma Wakefield Accelerator Design
(Jean-Luc Vay, LBNL)

- Virtual design of 100-stage 1 TeV collider; dramatically cut accelerator size and design cost

Challenges: scaling of Maxwell FFT-based solver; maintaining efficiency of large timestep algorithm; load balancing

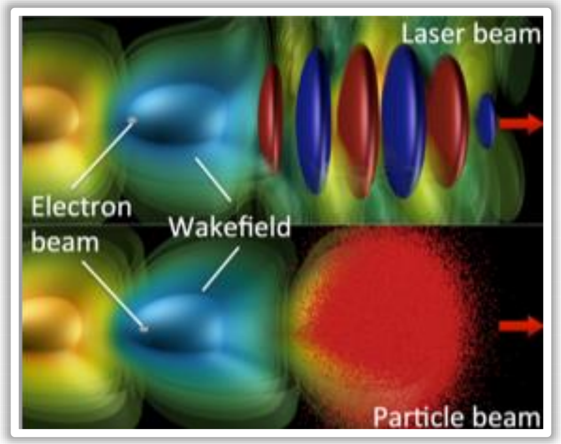

EXASCALE COMPUTING PROJECT

Exascale apps can deliver transformative products and solutions

WDMApp

High-Fidelity Whole Device Modeling of Magnetically Confined Fusion Plasmas
(Amitava Bhattacharjee, PPPL)

- Prepare for ITER exps and increase ROI of validation data and understanding
- Prepare for beyond-ITER devices

Challenges: robust, accurate, and efficient code-coupling algorithm; reduction in memory and I/O usage

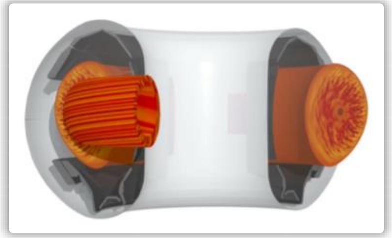

ExaStar

Demystify Origin of Chemical Elements
(Dan Kasen, LBNL)

- What is the origin of the elements?
- How does matter behave at extreme densities?
- What are the sources of gravity waves?

Challenges: delivering performance on accelerators; delivering fidelity for general relativity implementation

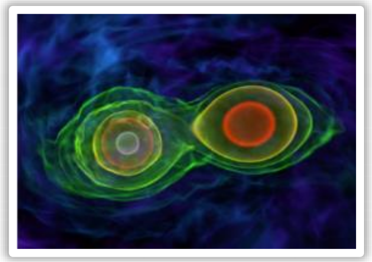

ExaFEL

Light Source-Enabled Analysis of Protein and Molecular Structure and Design
(Amadeo Perazzo, SLAC)

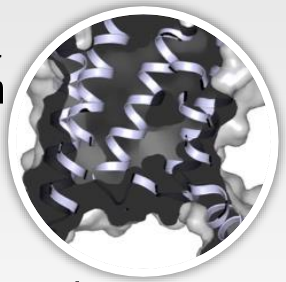

- Process data without beam time loss
- Determine nanoparticle size and shape changes
- Engineer functional properties in biology and materials science

Challenges: improving the strong scaling (one event processed over many cores) of compute-intensive algorithms (ray tracing, M-TIP) on accelerators

CANDLE

Accelerate and Translate Cancer Research
(Rick Stevens, ANL)

- Develop predictive preclinical models and accelerate diagnostic and targeted therapy through predicting mechanisms of RAS/RAF driven cancers

Challenges: increasing accelerator utilization for model search; effectively exploiting HP16; preparing for any data management or communication bottlenecks

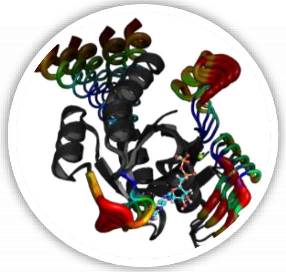

EXASCALE COMPUTING PROJECT

Exascale Application Development Challenges Overall

1) **Porting to accelerator-based architectures**
2) **Exposing additional parallelism**
3) **Coupling codes to create new multiphysics capability**
4) **Adopting new mathematical approaches**
5) **Algorithmic or model improvements**
6) **Leveraging optimized libraries**

ECP will intensify efforts to effectively exploit reduced precision representations of hardware-accelerated operations

The three technical areas in ECP have the necessary components to address these challenges and meet national goals

Performant mission and science applications @ scale

Foster application development	Ease of use	Diverse architectures	HPC leadership

Application Development (AD)	Software Technology (ST)	Hardware and Integration (HI)
Develop and enhance the predictive capability of applications critical to the DOE	Produce expanded and vertically integrated software stack to achieve full potential of exascale computing	Integrated delivery of ECP products on targeted systems at leading DOE computing facilities
20+ applications ranging from national security, to energy, earth systems, economic security, materials, and data	80+ unique software products spanning programming models and run times, math libraries, data and visualization	6 vendors supported by PathForward focused on memory, node, connectivity advancements; deployment to facilities

ECP Software: productive, sustainable ecosystem

Goal

Build a comprehensive, coherent software stack that enables application developers to productively write highly parallel applications that effectively target diverse exascale architectures

- Extend current technologies to exascale where possible
- Perform R&D required for new approaches when necessary
- Guide, and complement, and integrate with vendor efforts
- Develop and deploy high-quality and robust software products

ECP ST Software Ecosystem

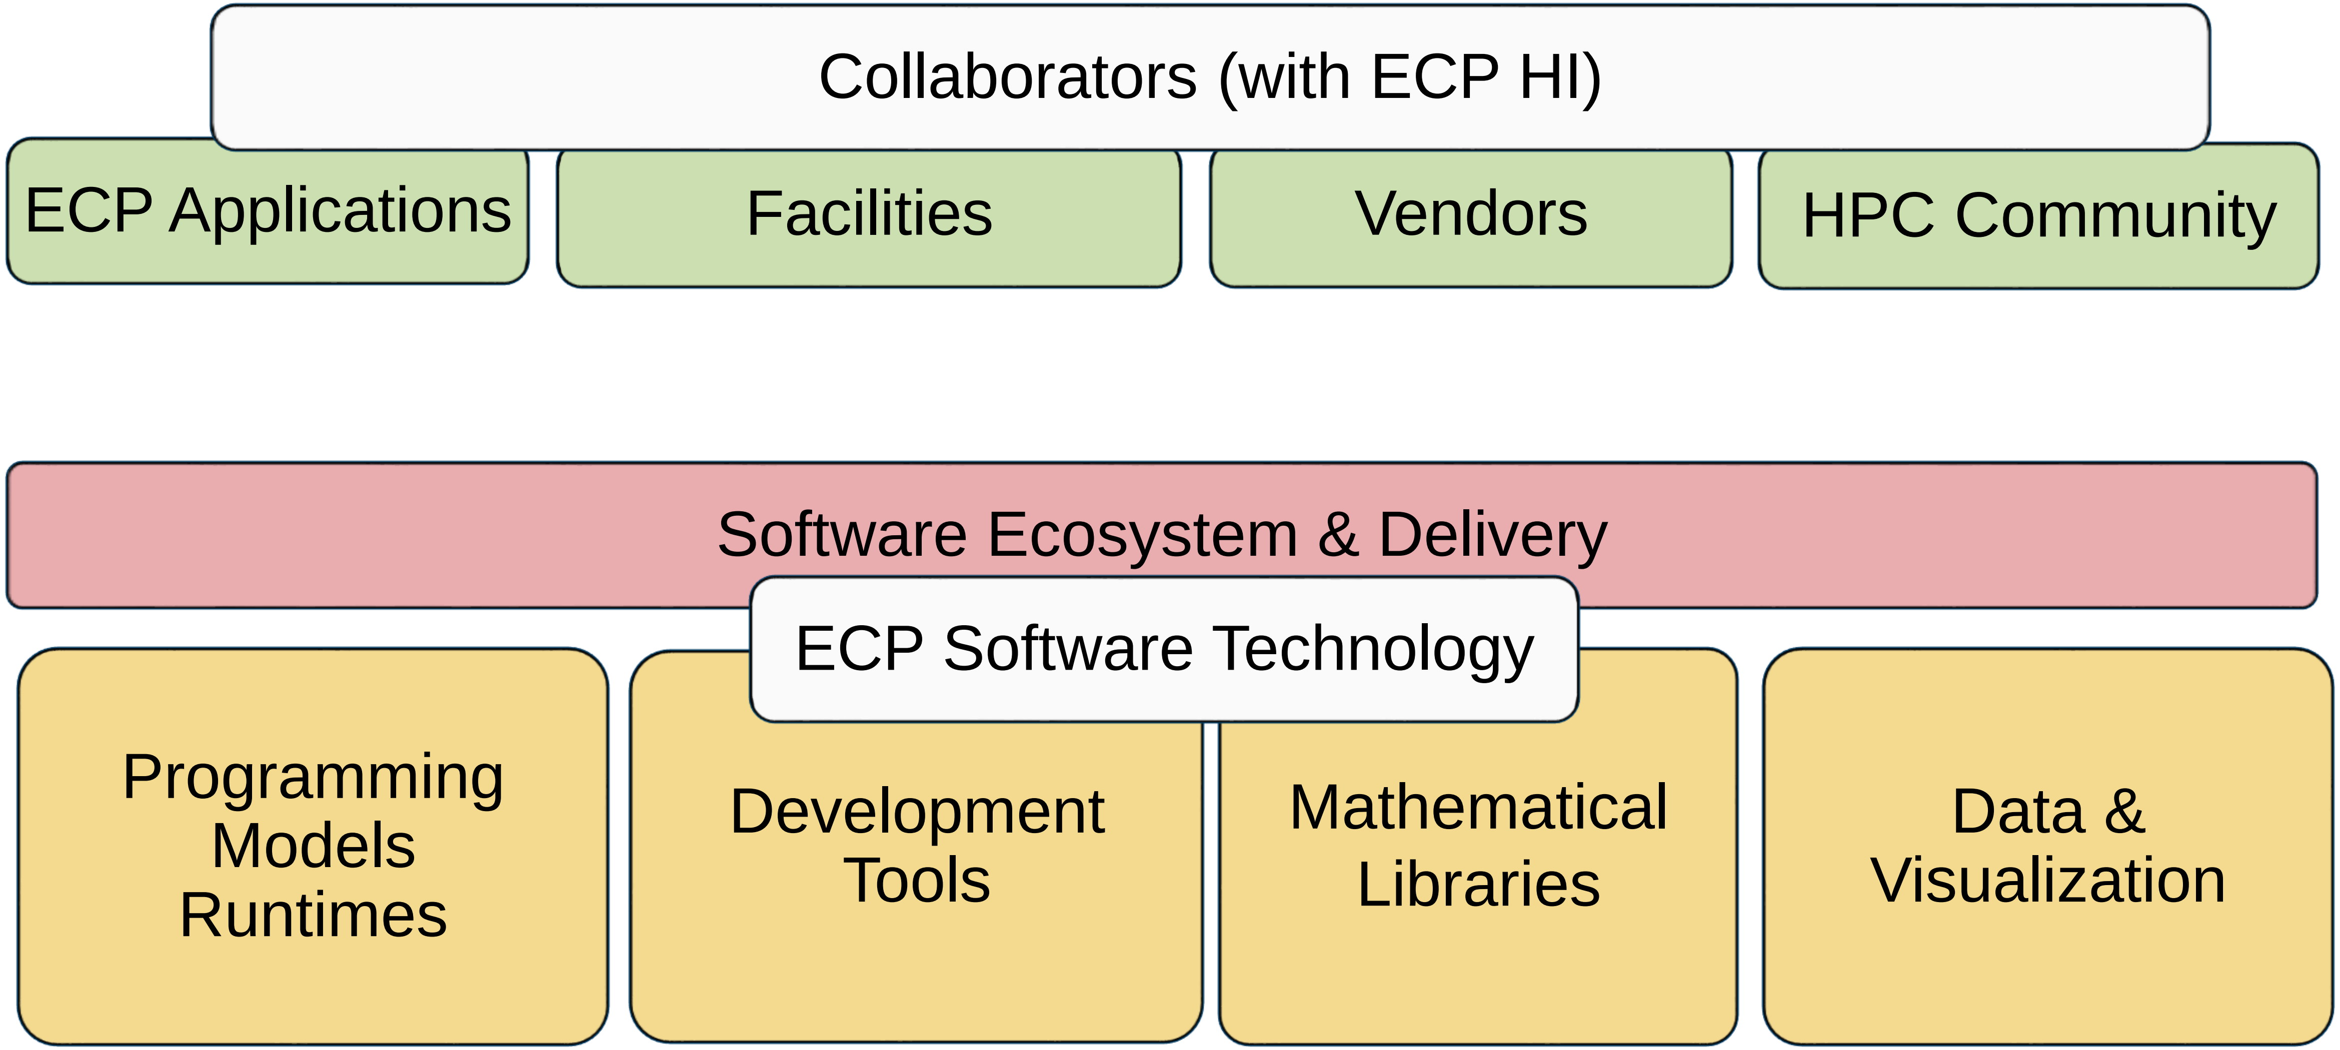

EXASCALE COMPUTING PROJECT

The Bottom Line for ECP Software Technology

- Next-generation **HPC technologies for 90 open source scientific software products**
- The performance potential of leadership computers in preparation for exascale
- **Software development kits (SDKs)** with turnkey installation and interoperability
- The **Extreme-scale Scientific Software Stack (E4S)**:
  - Target: Comprehensive software environment for HPC scientific applications
  - Tested on growing collection of HPC platforms in preparation for Exascale systems
  - Managed complexity using SDKs as components
  - From-source builds for leadership environments
  - Pre-built containers for development, debugging and portability
- A commitment to software quality leveraging industry best practices
- A legacy to build upon for US security, science, industry and technology leadership

ECP investments in software technologies help ensure the exascale computers will be a success

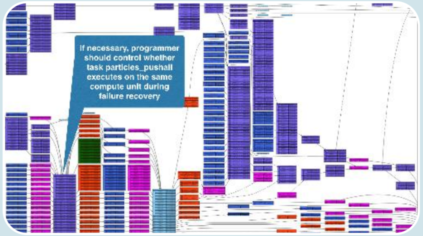

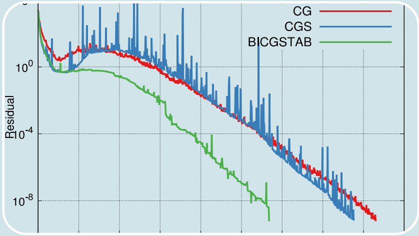

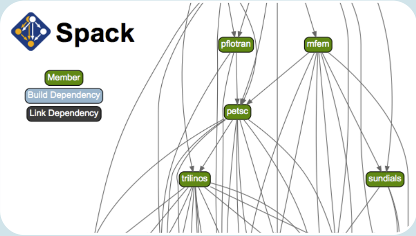

Programming Models & Runtimes

- Enhance and get ready for exascale the widely used MPI and OpenMP programming models (hybrid programming models, deep memory copies)
- Development of performance portability tools (e.g. Kokkos and Raja)
- Support alternate models for potential benefits and risk mitigation: PGAS (UPC++/GASNet) ,task-based models (Legion, PaRSEC)
- Libraries for deep memory hierarchy and power management

Development Tools

- Continued, multifaceted capabilities in portable, open-source LLVM compiler ecosystem to support expected ECP architectures, including support for F18
- Performance analysis tools that accommodate new architectures, programming models, e.g., PAPI, Tau

Math Libraries

- Linear algebra, iterative linear solvers, direct linear solvers, integrators and nonlinear solvers, optimization, FFTs, etc
- Performance on new node architectures; extreme strong scalability
- Advanced algorithms for multi-physics, multiscale simulation and outer-loop analysis
- Increasing quality, interoperability, complementarity of math libraries

Data and Visualization

- I/O via the HDF5 API
- Insightful, memory-efficient in-situ visualization and analysis – Data reduction via scientific data compression
- Checkpoint restart

Software Ecosystem

- Develop features in Spack necessary to support all ST products in E4S, and the AD projects that adopt it
- Development of Spack stacks for reproducible turnkey deployment of large collections of software
- Optimization and interoperability of containers on HPC systems
- Regular E4S releases of the ST software stack and SDKs with regular integration of new ST products

Implementation of a new integer type for global integers in hypre

ECP WBS	WBS 2.3.3.01
PI	Ulrike Meier Yang, LLNL
Members	ANL, LBNL, LLNL, SNL, UC Berkeley, UTK

Scope and objectives

- Determine and document where code changes are required, and implement the new integer type
- Investigate the performance and memory usage of the new mixed-int hypre version in comparison to the 64-bit integer hypre version and the 32-bit version, where possible
- Summarize the results and make the implementation available to application code teams

Performance improvement of new capability

Weak scaling study:
Total runtimes in seconds for AMG-PCG using 1M points/core for 2 different 3D diffusion problems.
The new mixed-int capability performs about 20-25% better than the 64-int version while using less memory and solves larger problems than 32-int.

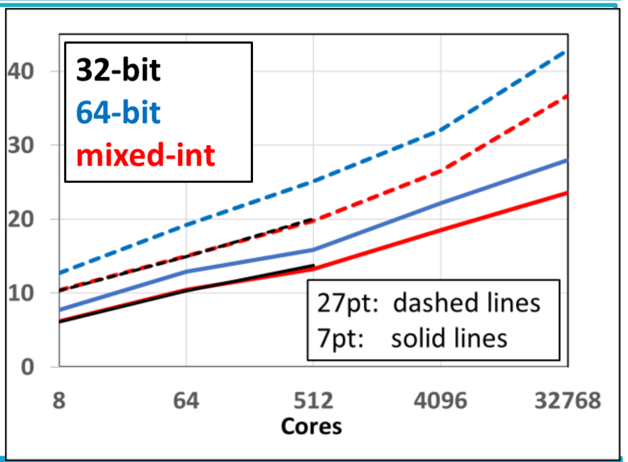

Impact

- Linear systems are an important part of many application codes, and often make up a large portion of their execution times.
- Efficient linear solvers are crucial for ECP applications, and any improvements in performance and memory usage positively impact the applications.

Project accomplishment

- Implemented a new integer type HYPRE_BigInt in hypre to reduce use of 64-bit integers for large problems and improve performance and memory use
- Investigated its performance and memory usage, and summarized the results and code changes in a document
- Generated a new hypre release: v. 2.16.0

Deliverables Hypre release v. 2.16.0 at https://github.com/hypre-space/hypre

ECP ST staff contribute to ISO and *de facto* standards groups: assuring sustainability through standards

Standards Effort	ECP ST Participants
MPI Forum	15
OpenMP	15
BLAS	6
C++	4
Fortran	4
OpenACC	3
LLVM	2
PowerAPI	1
VTK ARB	1

- **MPI/OpenMP:** Several key leadership positions
- Heavy involvement in all aspects.
- **C++:** Getting HPC requirements considered, contributing working code.
- **Fortran:** Flang front end for LLVM.
- ***De facto:*** Specific HPC efforts.
- **ARB*:** Good model for SDKs.

*Architecture Review Board

Many ECP ST products are available for broad community use

For example…

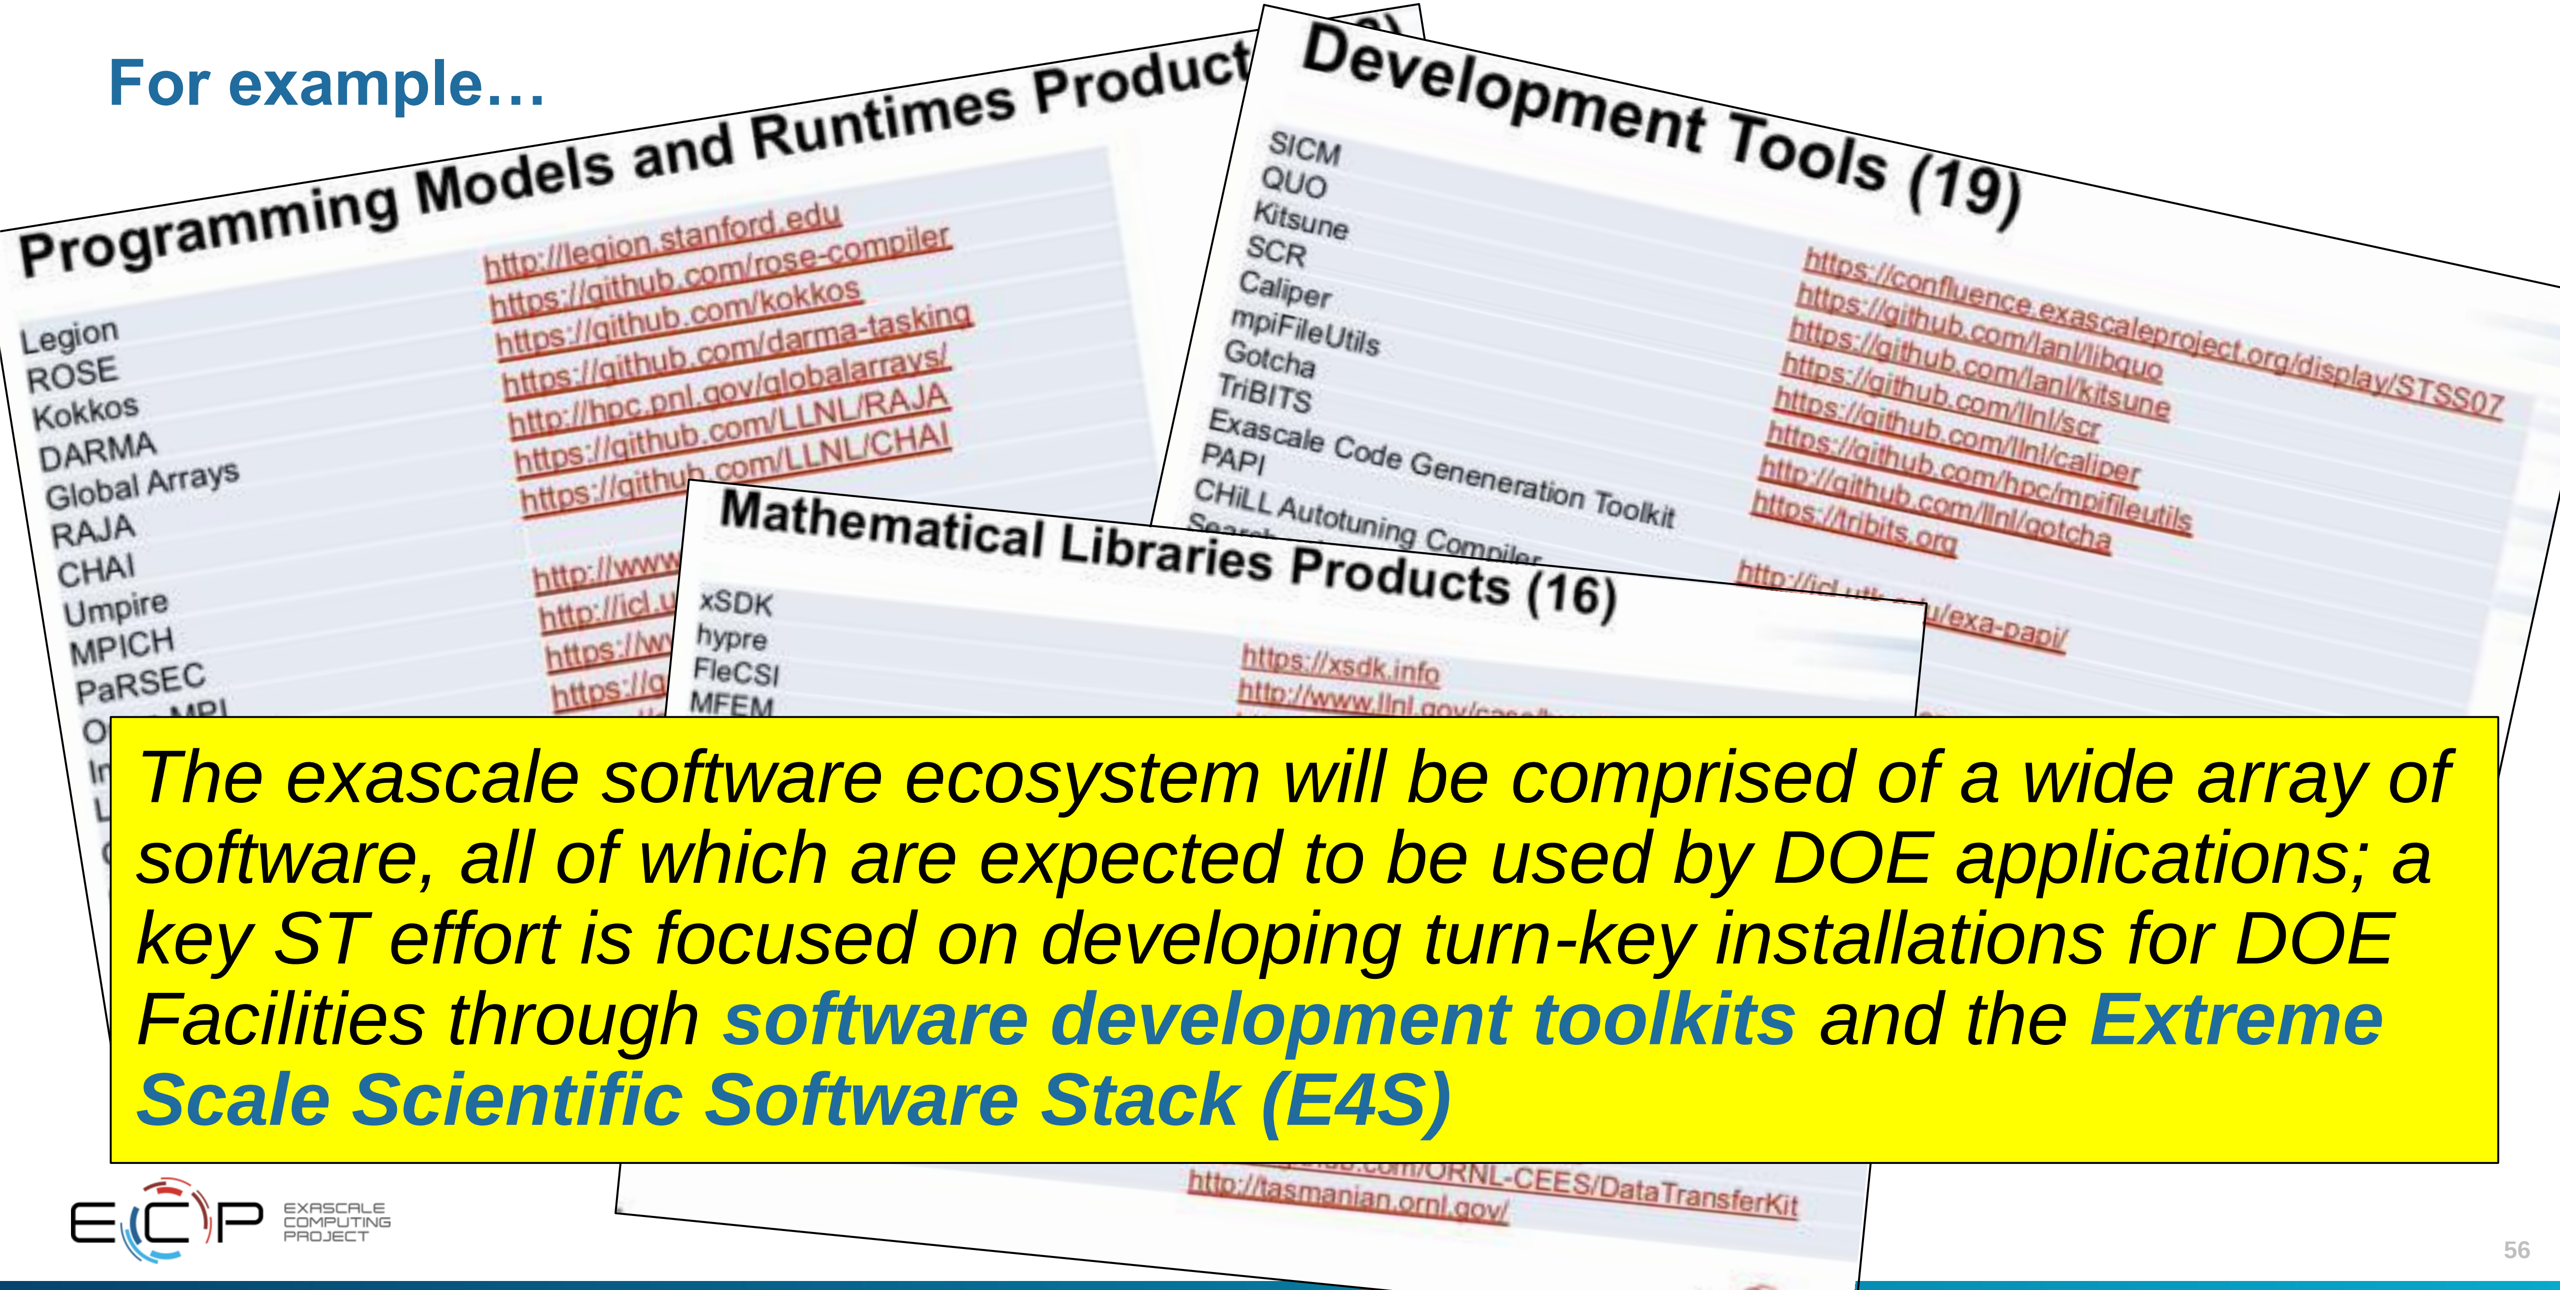

*The exascale software ecosystem will be comprised of a wide array of software, all of which are expected to be used by DOE applications; a key ST effort is focused on developing turn-key installations for DOE Facilities through **software development toolkits** and the **Extreme Scale Scientific Software Stack (E4S)***

EXASCALE COMPUTING PROJECT

Software Development Toolkits are a key delivery vehicle for ECP

- A collection of related software products (called packages) where coordination across package teams will improve usability and practices and foster community growth among teams that develop similar and complementary capabilities
- Attributes
  - Domain scope: Collection makes functional sense
  - Interaction model: How packages interact; compatible, complementary, interoperable
  - **Community policies**: Value statements; serve as criteria for membership
  - Meta-infrastructure: Encapsulates, invokes build of all packages (Spack), shared test suites
  - Coordinated plans: Inter-package planning. Does not replace autonomous package planning
  - Community outreach: Coordinated, combined tutorials, documentation, best practices
- Overarching goal: Unity in essentials, otherwise diversity

xSDK-0.3.0: Dec 2017… (that was then..)

https://xsdk.info

Notation: A → B:
A can use **B** to provide functionality on behalf of **A**

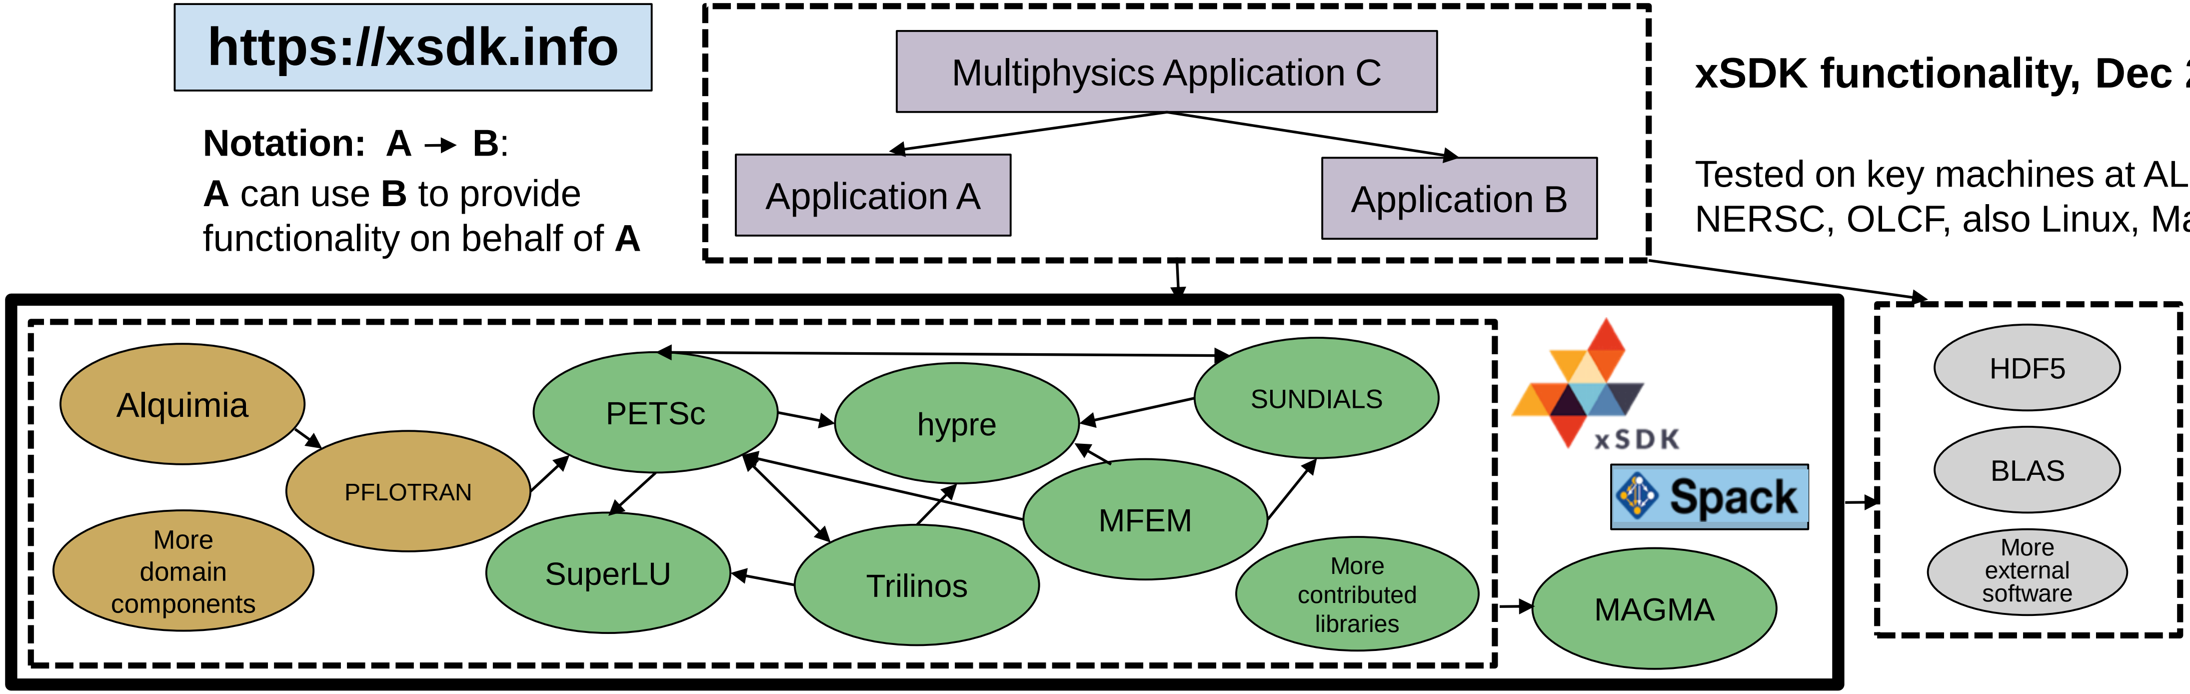

xSDK functionality, Dec 2017

Tested on key machines at ALCF, NERSC, OLCF, also Linux, Mac OS X

July 2018: Revisions of xSDK Community Policies

https://xsdk.info/policies

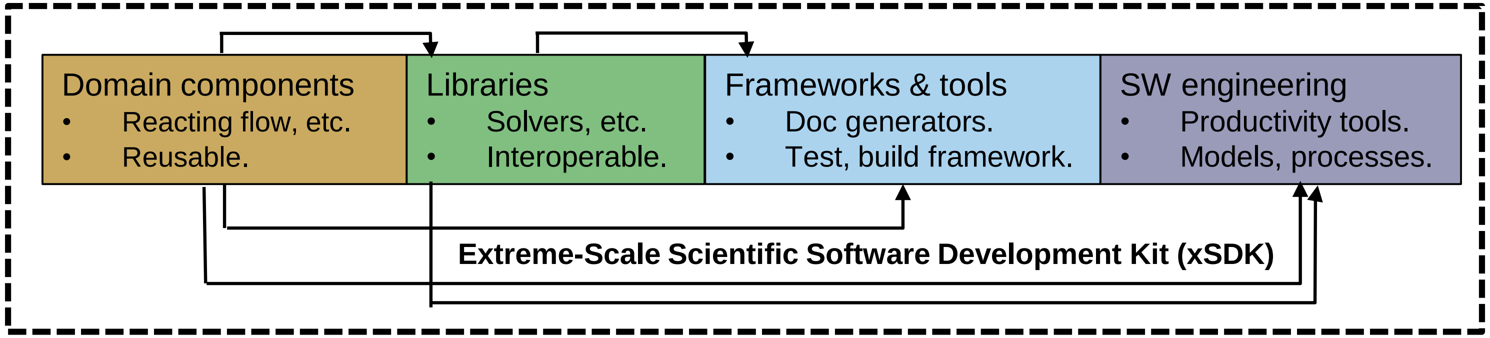

EXASCALE COMPUTING PROJECT

xSDK Version 0.4.0: December 2018 (this is now)

https://xsdk.info

Each xSDK member package uses or can be used with one or more xSDK packages, and the connecting interface is regularly tested for regressions.

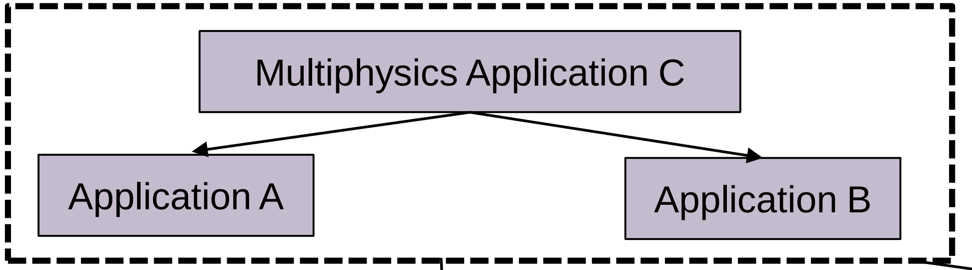

xSDK functionality, Dec 2018

Tested on key machines at ALCF, NERSC, OLCF, also Linux, Mac OS X

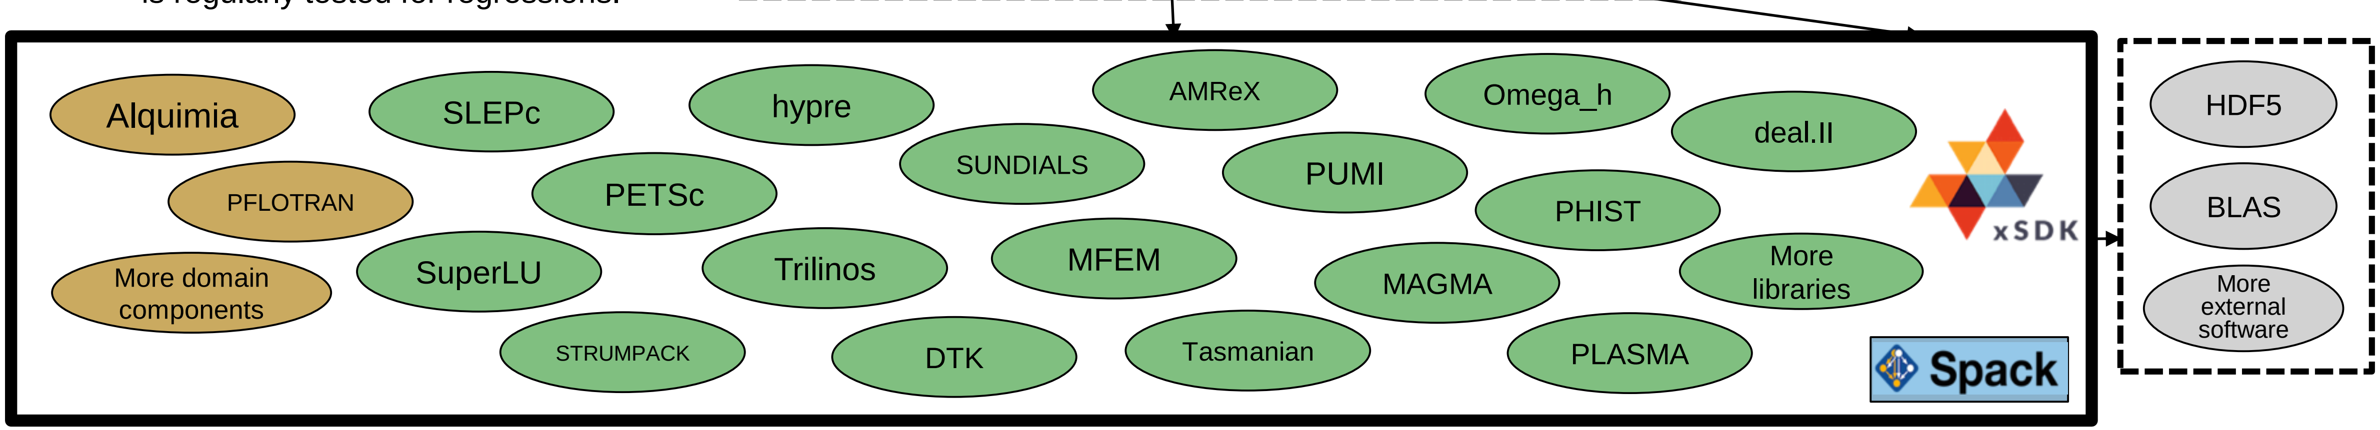

December 2018

- 17 math libraries
- 2 domain components
- 16 mandatory xSDK community policies
- Spack xSDK installer

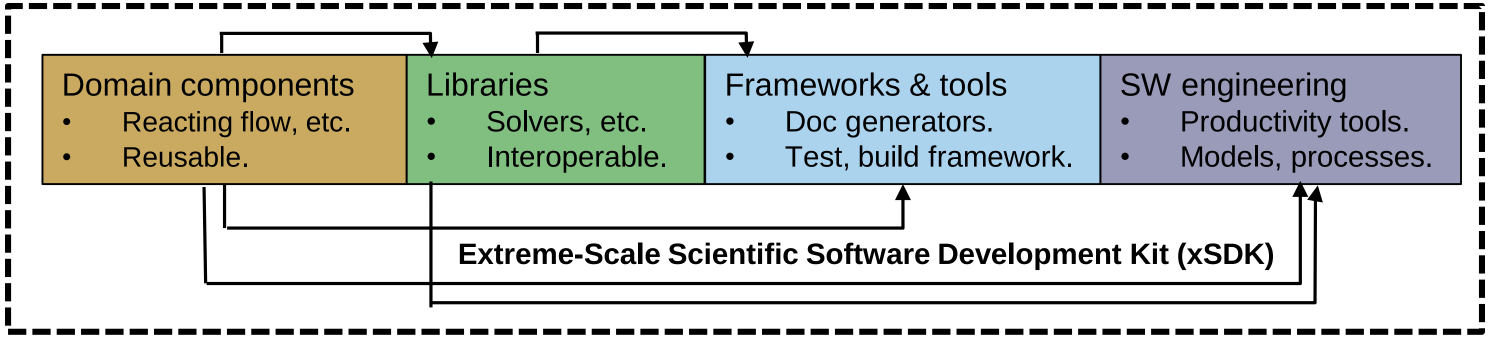

Impact: Improved code quality, usability, access, sustainability

Foundation for work on performance portability, deeper levels of package interoperability

EXASCALE COMPUTING PROJECT

The planned ECP ST SDKs will span all technology areas

xSDK (16)	PMR Core (17)	Tools and Technology (11)	Compilers & Support (7)	Visualization Analysis & Reduction (9)	Data Mgmt, I/O Services, & Checkpoint restart (12)	Ecosystem/E4S at-large (12)
hypre	Legion	TAU	openarc	ParaView	FAODEL	BEE
FleSCI	Kokkos (Support)	HPCToolkit	Kitsune	Catalyst	ROMIO	FSEFI
MFEM	RAJA	Dyninst Binary Tools	LLVM	VTK-m	Mercury (part of Mochi suite)	Kitten Lightweight Kernel
Kokkoskernels	CHAI	Gotcha	CHiLL Autotuning Compiler	SZ	HDF5	COOLR
Trilinos	PaRSEC*	Caliper	LLVM OpenMP compiler	zfp	Parallel netCDF	NRM
SUNDIALS	DARMA	PAPI	OpenMP V & V	VisIt	ADIOS	ArgoContainers
PETSc/TAO	GASNet-EX	Program Database Toolkit	Flang/LLVM Fortran compiler	ASCENT	Darshan	Spack
libEnsemble	Qthreads	Search using Random Forests		Cinema	UnifyCR	MarFS
STRUMPACK	BOLT	Siboka		ROVER	VeloC	GUFI
SuperLU	UPC++	C2C			IOSS	Intel GEOPM
ForTrilinos	MPICH	Sonar			HXHIM	mpiFileUtils
SLATE	Open MPI				SCR	TriBITS
MAGMA	Umpire					
DTK	QUO					
Tasmanian	Papyrus					
TuckerMPI	SICM					
	AML					

Key
PMR
Tools
Math Libraries
Data and Vis
Ecosystems and Delivery

Extreme-Scale Scientific Software Stack – E4S

- E4S: A Spack-based distribution of ECP ST and related and dependent software tested for interoperability and portability to multiple architectures
- Provides distinction between SDK usability / general quality / community and deployment / testing goals
- Will leverage and enhance SDK interoperability thrust
- Oct: E4S 0.1 - 24 full, 24 partial release products
- Jan: E4S 0.2 - 37 full, 10 partial release products
- Current primary focus: Facilities deployment

E4S

e4s.io

Lead: Sameer Shende
(U Oregon)

E4S Full Release and Installed Packages

- Adios
- Bolt
- Caliper
- Darshan
- Gasnet
- GEOPM
- GlobalArrays
- Gotcha
- HDF5
- HPCToolkit
- Hypre
- Jupyter
- Kokkos
- Legion
- Libquo
- Magma
- MFEM
- MPICH
- OpenMPI
- PAPI
- Papyrus
- Parallel netCDF
- ParaView
- PETSc/TAO
- Program Database Toolkit (PDT)
- Qthreads
- Raja
- SCR
- Spack
- Strumpack
- Sundials
- SuperLU
- Swift/T
- SZ
- Tasmanian
- TAU
- Trilinos
- VTKm
- Umpire

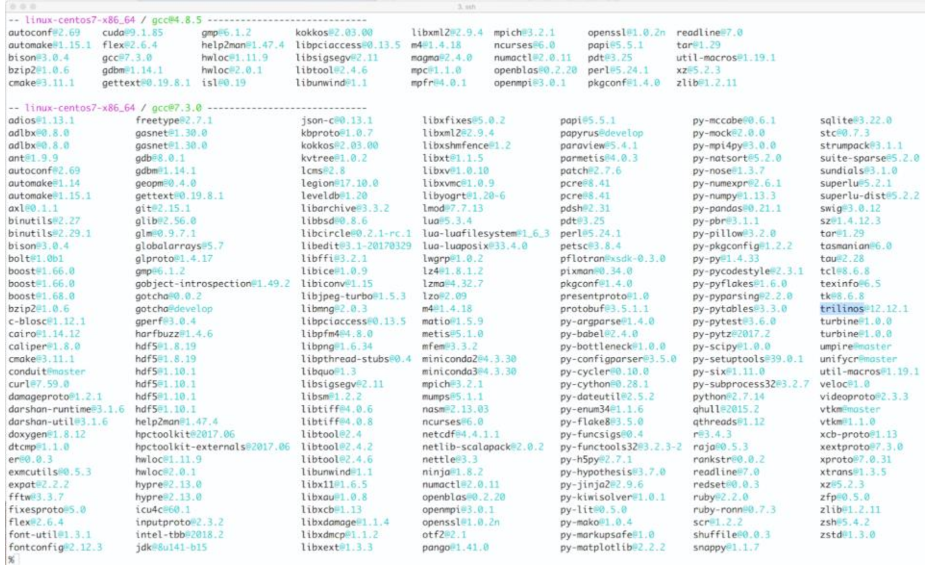

Packages installed using Spack

- UnifyCR
- Veloc
- xSDK
- Zfp

Detailed Information about the software technology projects is available in the ECP ST Capability Assessment Report

- Products discussed here are presented with more detail and further citations.
- We classify ECP ST Products deployment as Broad, Moderate or Experimental.
  - Broad and Moderate Deployment is typical suitable for collaboration.
  - Web links are available for almost all products.
  - About 1/3 of ECP ST Products are available as part of the Extreme-scale Scientific Software Stack (E4S) http://e4s.io.

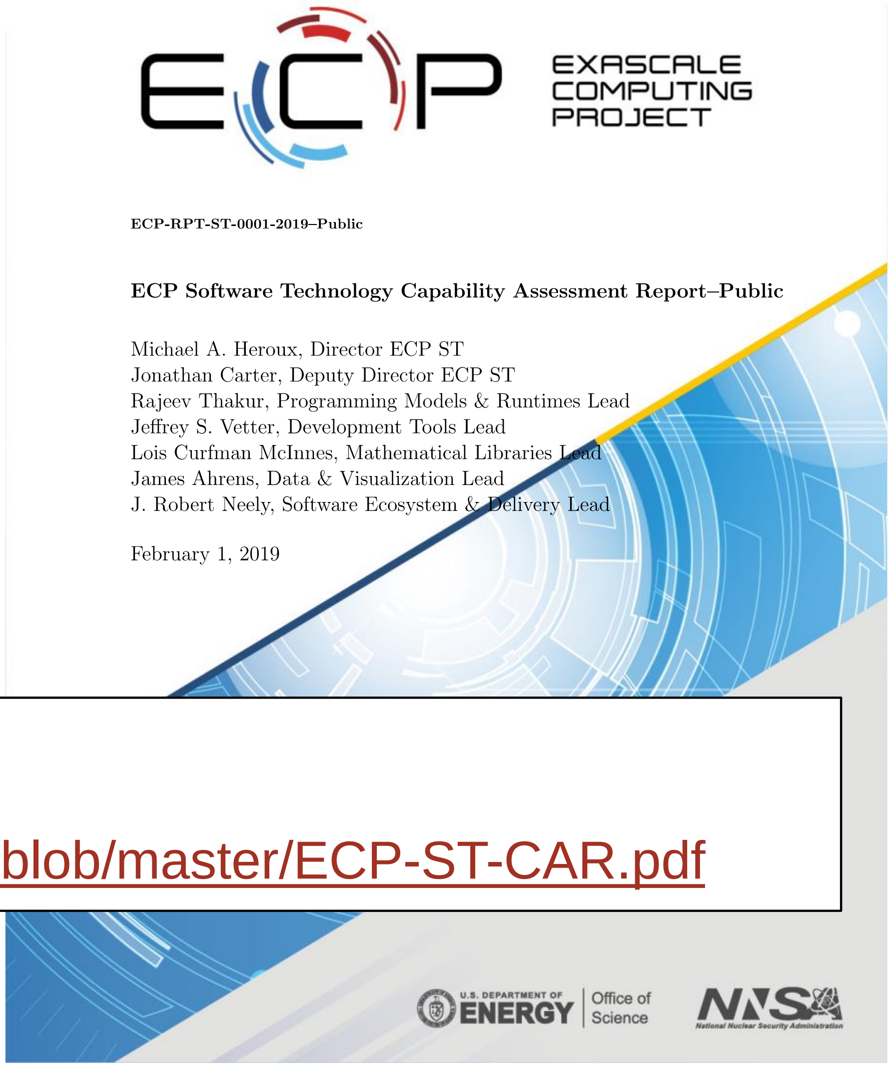
EXASCALE COMPUTING PROJECT

ECP-RPT-ST-0001-2019–Public

ECP Software Technology Capability Assessment Report–Public

Michael A. Heroux, Director ECP ST
Jonathan Carter, Deputy Director ECP ST
Rajeev Thakur, Programming Models & Runtimes Lead
Jeffrey S. Vetter, Development Tools Lead
Lois Curfman McInnes, Mathematical Libraries Lead
James Ahrens, Data & Visualization Lead
J. Robert Neely, Software Ecosystem & Delivery Lead

February 1, 2019

U.S. DEPARTMENT OF ENERGY | Office of Science
National Nuclear Security Administration

V 1.0 https://www.exascaleproject.org

V 1.5 https://github.com/E4S-Project/ECP-ST-CAR-PUBLIC/blob/master/ECP-ST-CAR.pdf

The three technical areas in ECP have the necessary components to address these challenges and meet national goals

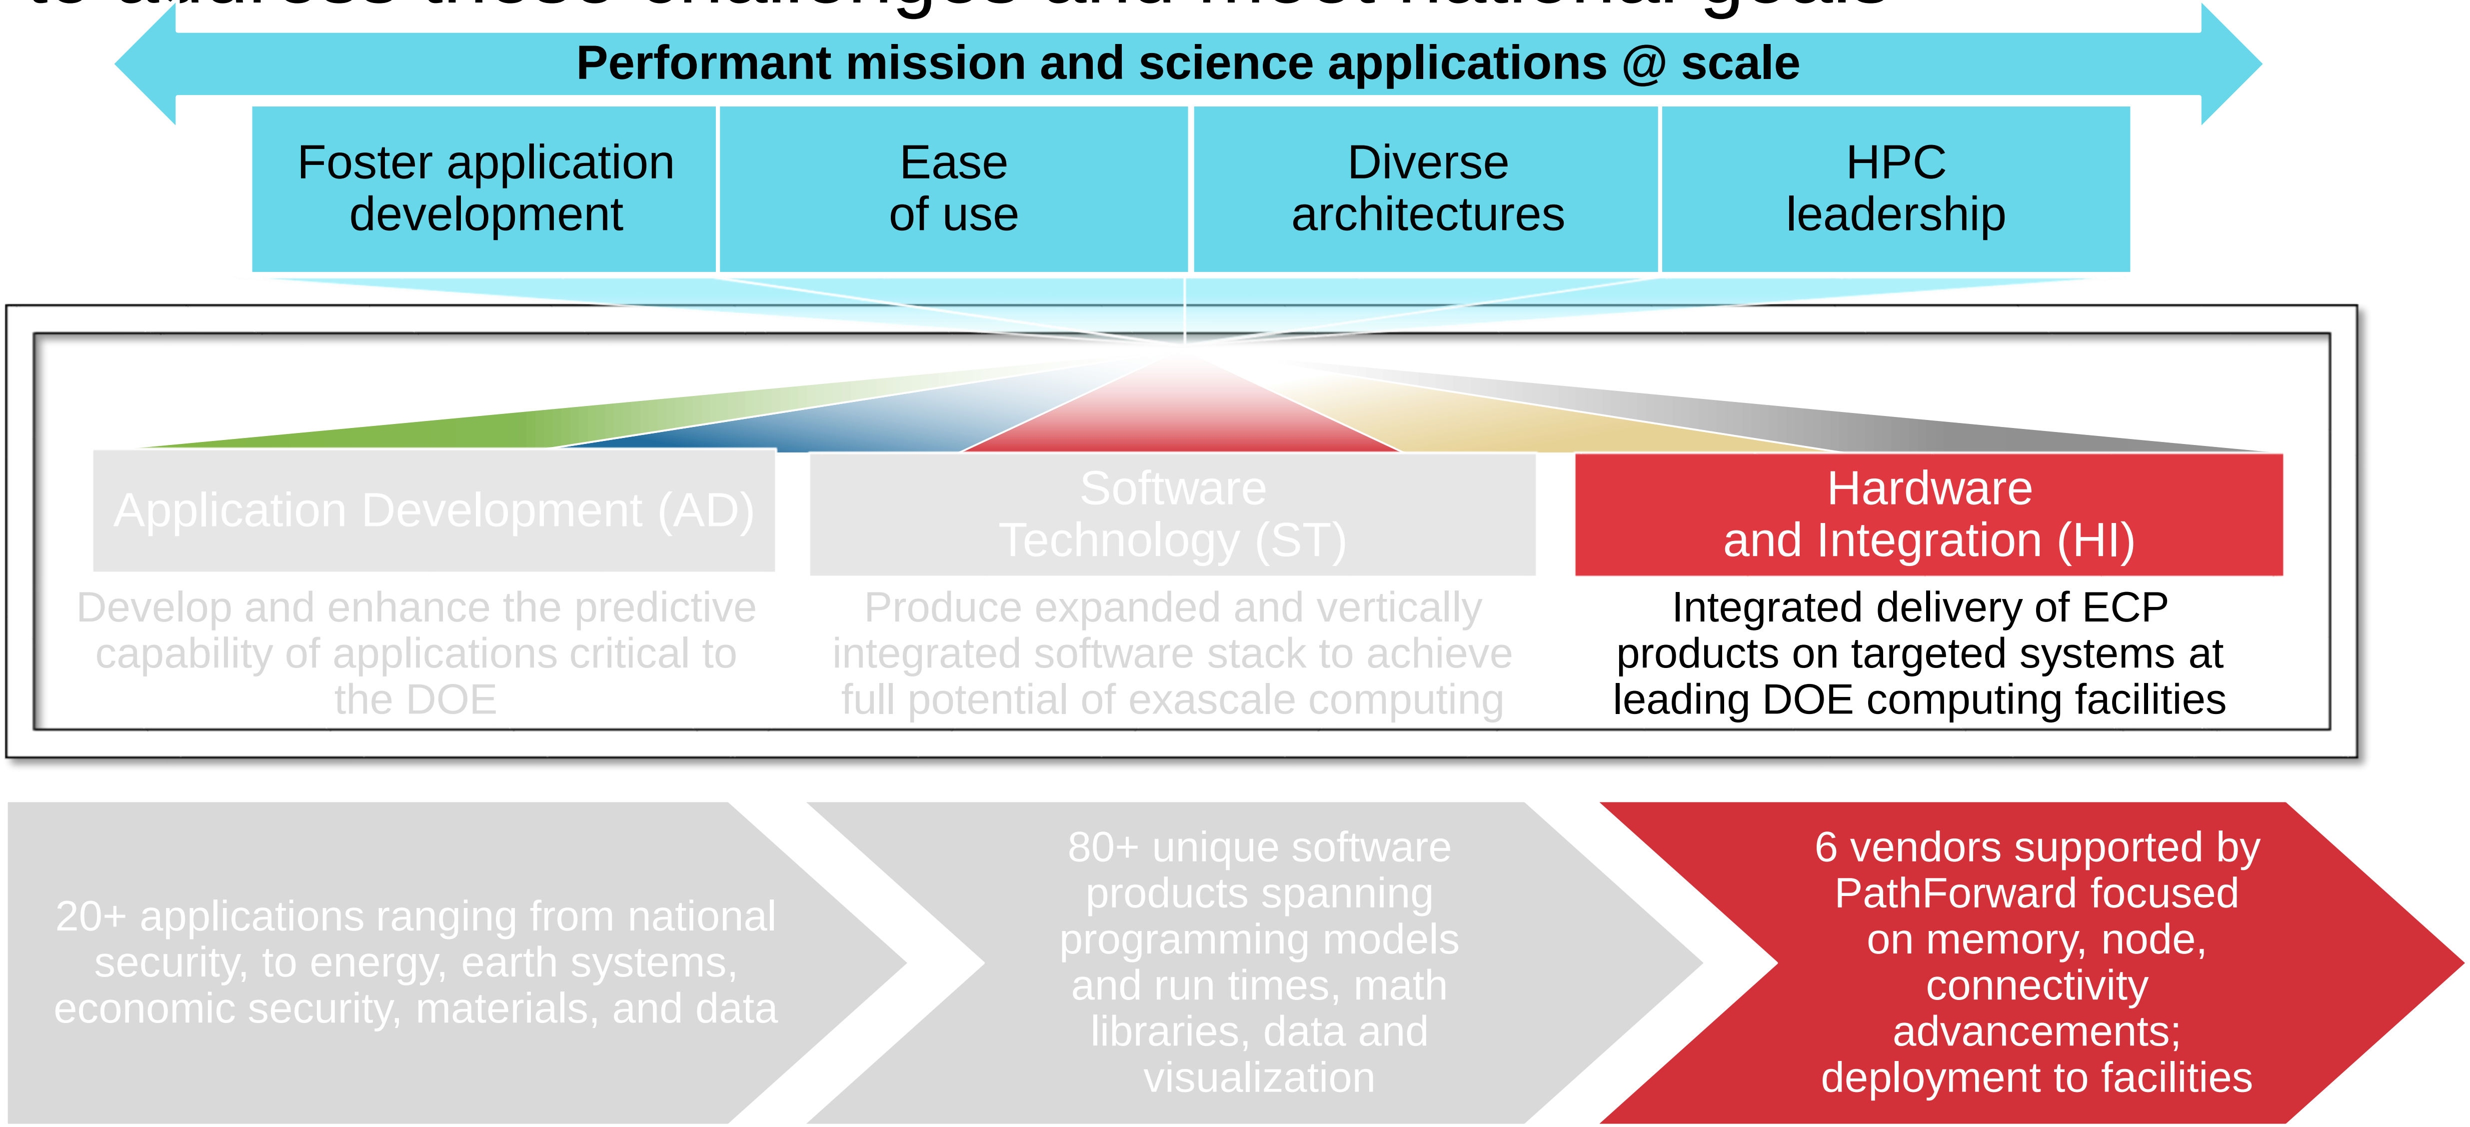

ECP Hardware and Integration: delivery of integrated ECP/DOE Facility products

Goal

A capable exascale computing ecosystem made possible by integrating ECP applications, software and hardware innovations within DOE facilities

- Innovative supercomputer architectures for competitive exascale system designs
- Accelerated application readiness through collaboration with the facilities
- A well integrated and continuously tested exascale software ecosystem deployed at DOE facilities through collaboration with facilities
- Training on key ECP technologies, help in accelerating the software development cycle and in optimizing the productivity of application and software developers

ECP's PathForward Vendor Hardware R&D Efforts Accelerate Hardware Technologies for Exascale Systems

PathForward began in 2017; builds upon the FastForward I & II and DesignForward I & II efforts

Total value of the work is ~$430M; DOE paid 60% of the price or $250+M

Examples of work funded include:

a) innovative memory architectures

b) higher-speed interconnects

c) improved system reliability

d) Innovations for increased parallelism - approaches for increasing computing power without prohibitive increases in energy demand

- Advanced Micro Devices (AMD)
- Cray Inc. (CRAY)
- Hewlett Packard Enterprise (HPE)
- International Business Machines (IBM)
- Intel Corp. (Intel)
- NVIDIA Corp. (NVIDIA)

ECP's PathForward project and ASCR Facility system acquisition projects work together to accelerate the delivery of exascale systems

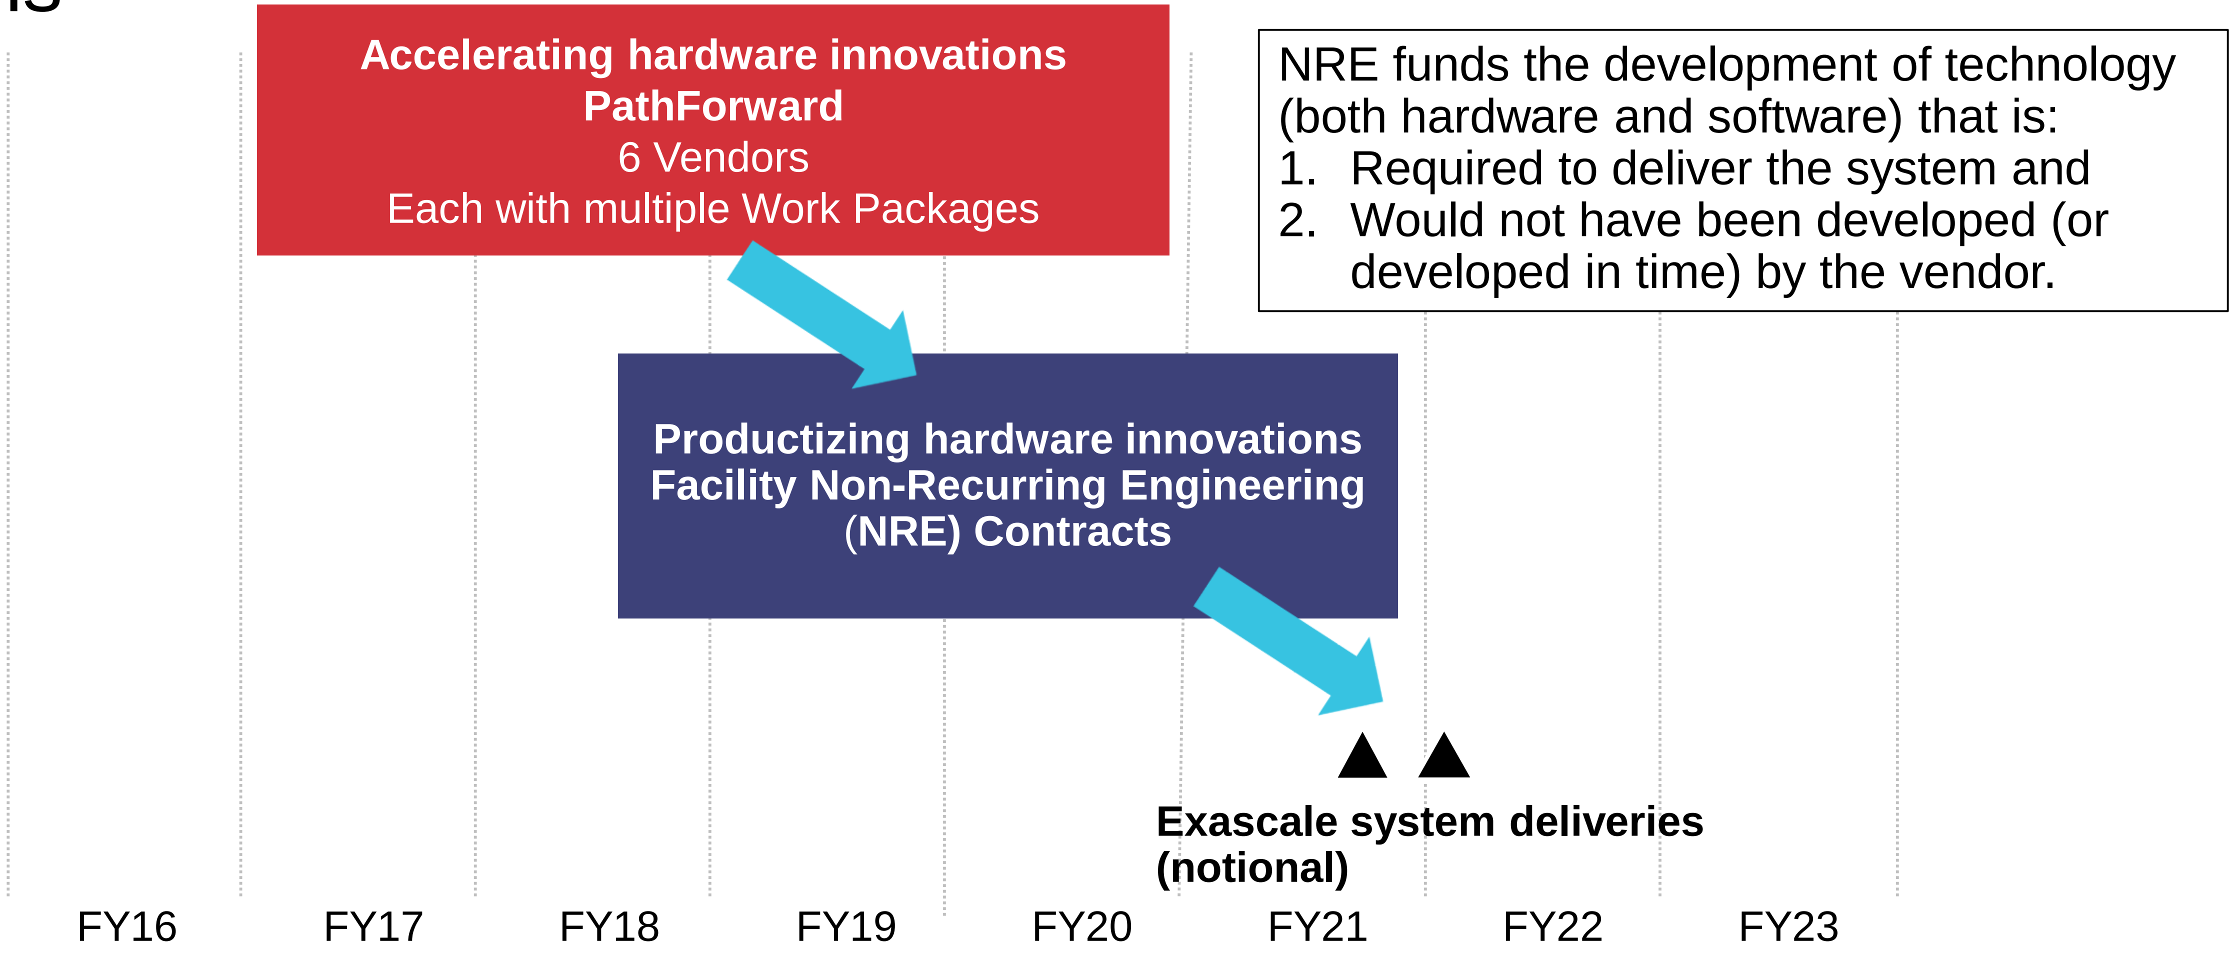

EXASCALE COMPUTING PROJECT

Application Integration at the Facilities Accelerates ECP AD's Application Readiness on ASCR Facilities Exascale Architectures

Approach: Leverage DOE Facilities Application Readiness expertise (from OLCF CAAR, ALCF ESP, NERSC NESAP programs) by providing:

- Facility computational scientists and performance engineering expertise to AD teams
- Access to Facilities Vendor Centers of Excellences

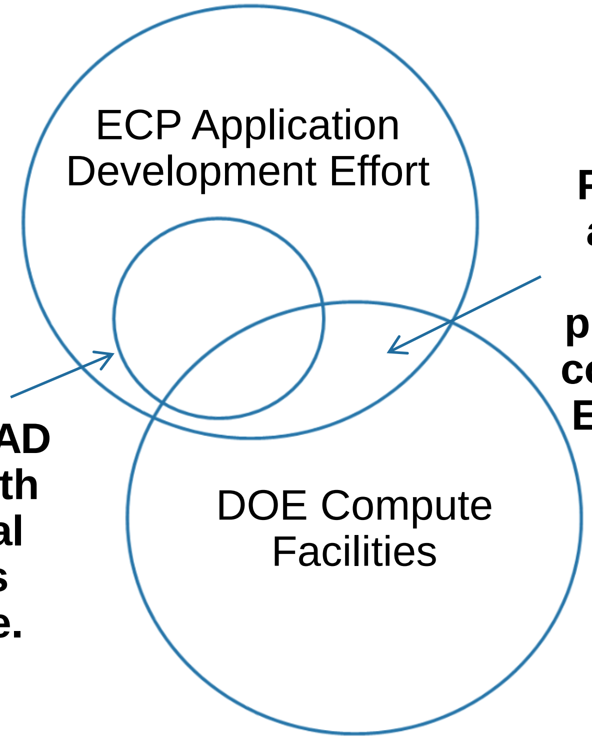

Augment AD efforts with additional facilities expertise.

Provide AD efforts with access to (1) Facilities Vendor COEs and (2) production/development computing resources via ECP allocation program (WBS 2.4.5)

Benefits of this Approach

- Facilities benefit by having more applications prepared to run on their systems as soon as the system is ready
- ECP benefits by leveraging the Facilities' successful application readiness efforts

Software Deployment at Facilities deploys ECP's ST products to meet ECP application needs

- Project Goal: ECP software integrated with facility software and vendor software targeting application needs
- Establish and operate a Continuous Integration Testing infrastructure for automated testing across HPC sites
- Develop a Software Deployment pipeline that supports packaging, efficient deployment at multiple facilities, and allows for container deployment approaches

Approach

Establish ongoing collaborations (via funded efforts) across Facilities that have historically been ad-hoc

Define infrastructure and production-quality processes that will live beyond the lifetime of ECP and establish a long-lasting DOE software sustainability model

Address unique site-specific deployment models while attempting to drive more commonality where it benefits the ECP user community

Maximize cross-fertilization of ST technology across multiple sites, vendors, and open source offerings (e.g. OpenHPC)

ECP affords a first-ever opportunity to drive the major DOE HPC Facilities and software developers to establish, share, and leverage common practices that will be critical for post-ECP software sustainability

The central DOE GitLab managed by OSTI for the CI process will provide software centralization with cross-site build and run capabilities

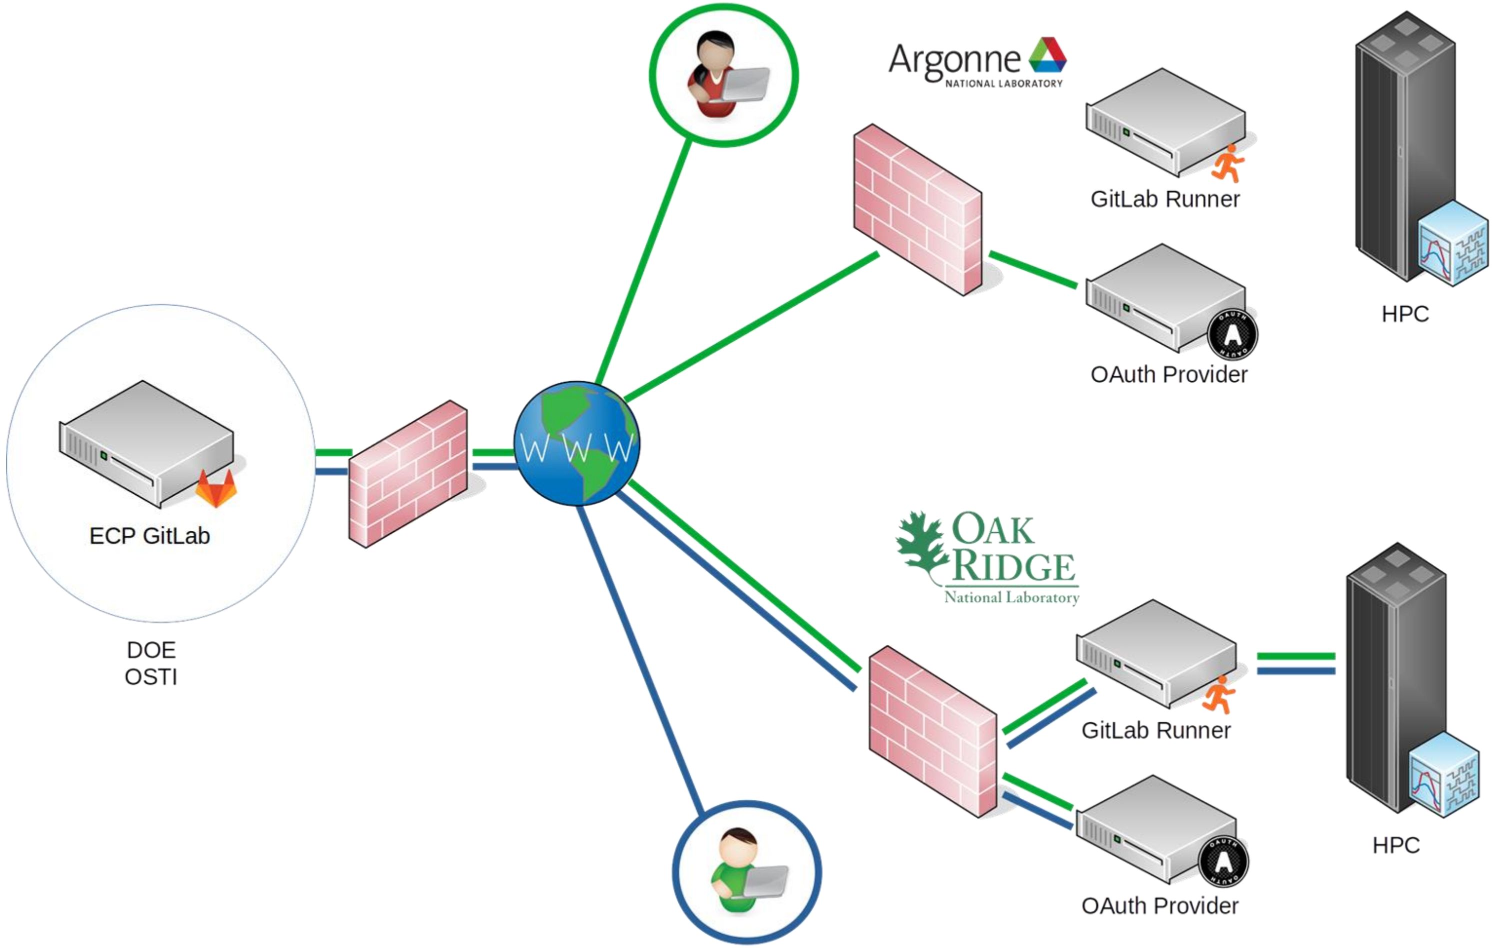

EXASCALE COMPUTING PROJECT

ECP's Flow of Software and Application Delivery and Deployment

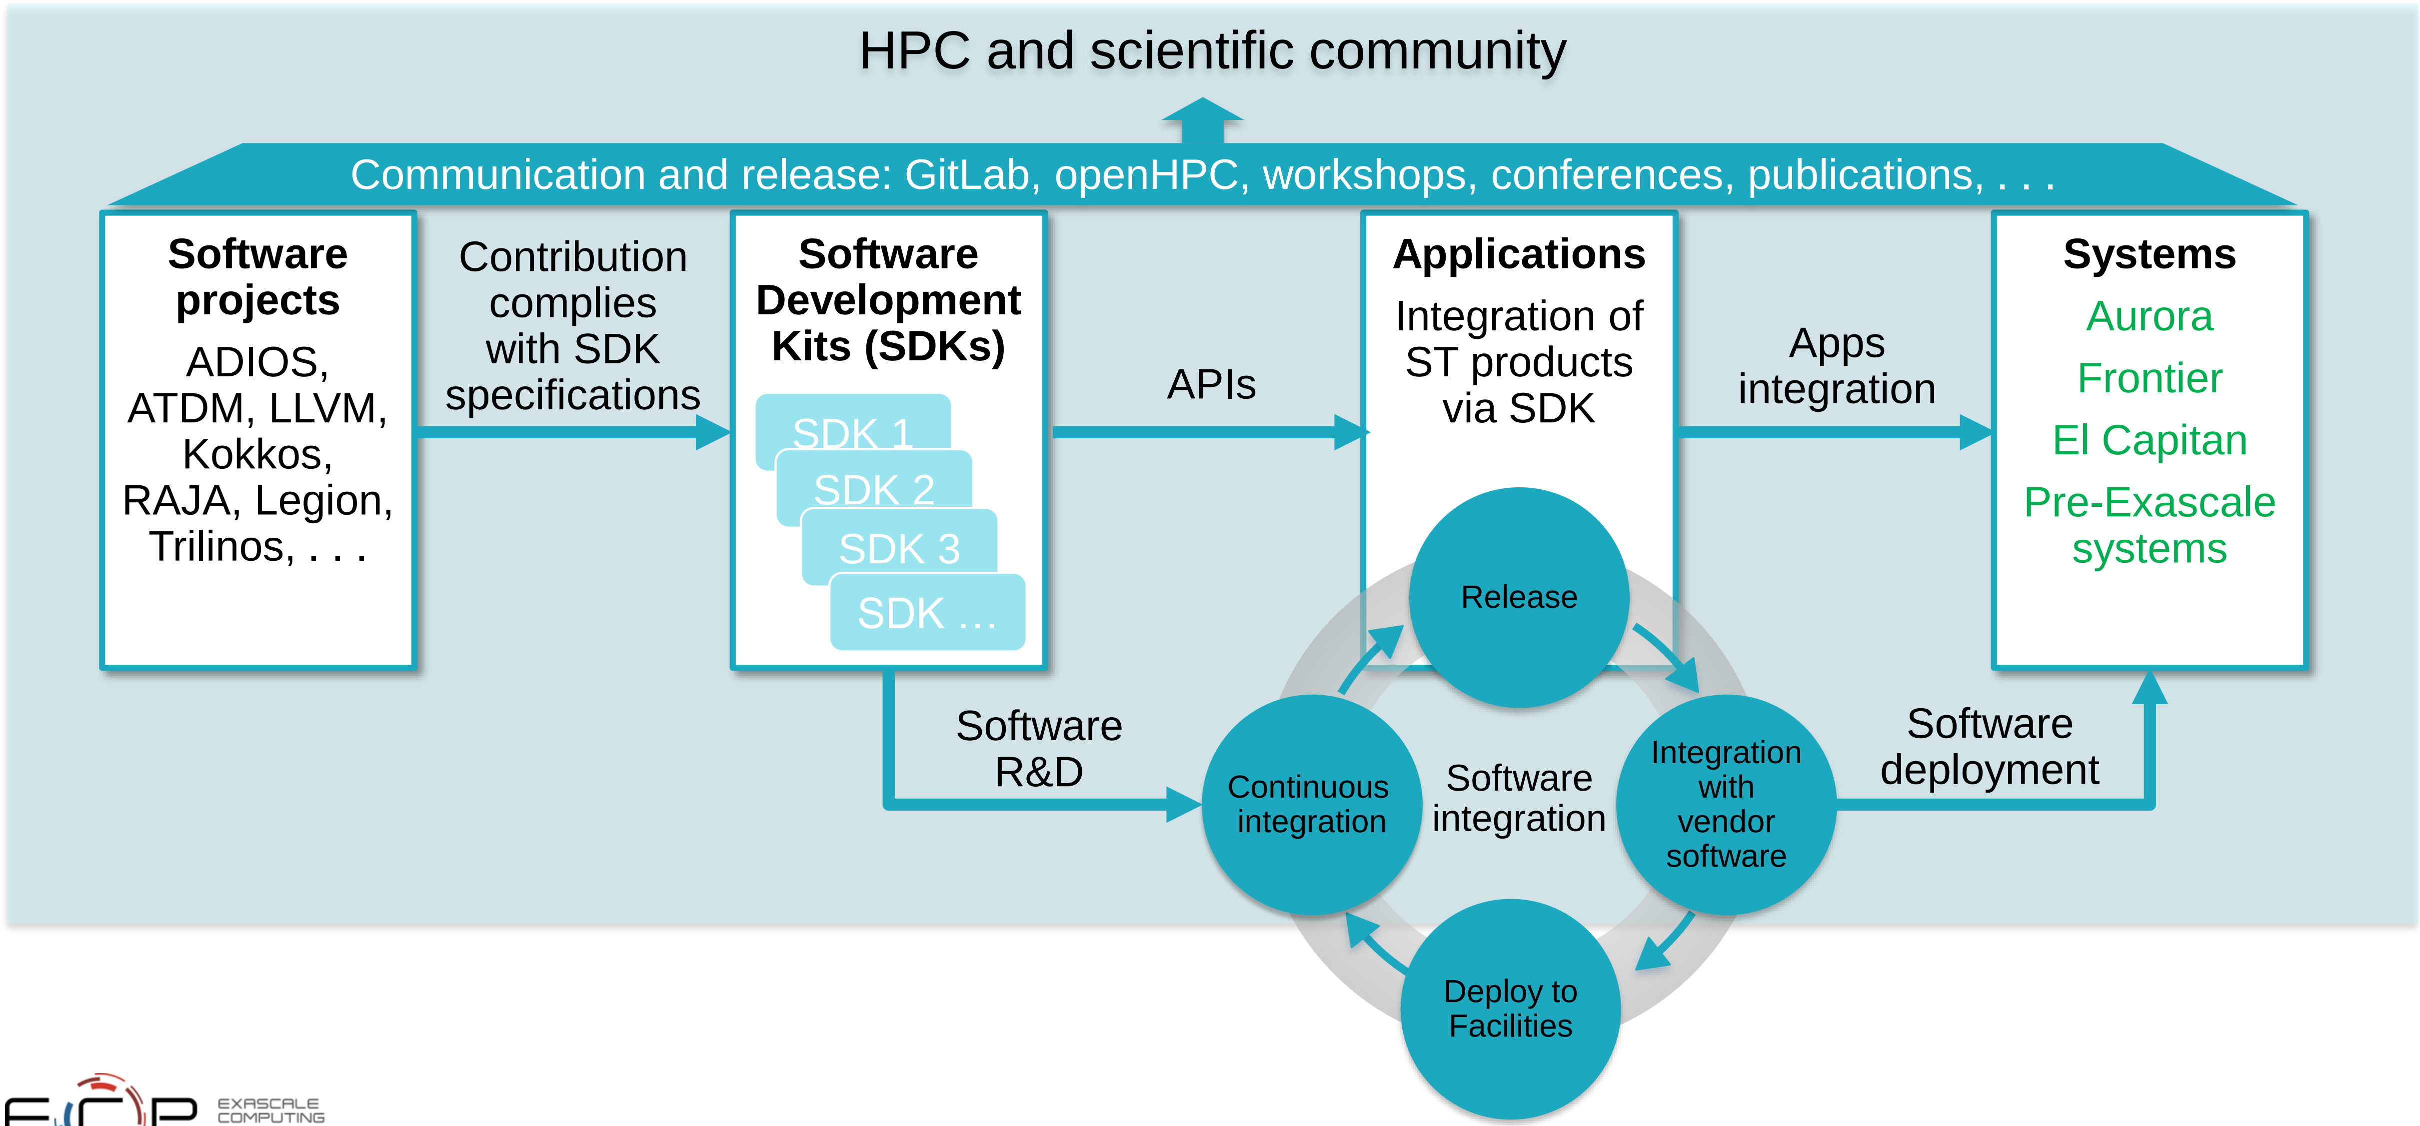

EXASCALE COMPUTING PROJECT

The ECP is on track to deliver a capable exascale computing ecosystem

Applications

- 20+ application teams actively engaged in targeted development and capability enablement for 2+ years
- Apps have well-defined exascale challenge problem targets with associated “science work rate” goals
- Initial performance experiences on pre-exascale systems (Summit, Sierra) exceeding expectations

Software Stack

- Over 80 software technology products being actively developed for next generation architectures
- Regular assessment of software stack products ensures line-of-sight to apps and HPC Facilities
- Plans for broad containerized delivery of products via SDKs and the E4S being executed

Hardware & Integration

- Return on PathForward vendor hardware R&D element evident in recent exascale RFP responses
- Plans for deployment and continuous integration of SDKs into DOE HPC Facilities being executed
- Prioritized performance engineering of applications targeting first three exascale systems underway

ECP Acknowledgments
Department of Energy (DOE) Support and Leadership

- **DOE Office of Science Advanced Scientific Computing Research (ASCR) Program**
  - *Barb Helland*
    - ASCR Associate Director and ECP Program Manager
- **DOE National Nuclear Security Administration (NNSA) Advanced Simulation and Computing (ASC) Program**
  - *Thuc Hoang*
    - ASC Program Manager for Computational Systems and Software Environment & Facility Operations and User Support and ECP Program Manager
- **Oak Ridge National Laboratory (ORNL) Site Office (OSO)**
  - *Dan Hoag*
    - OSO Senior Technical Advisor and ECP Federal Project Director
- **Lawrence Livermore National Laboratory (LLNL) Site Office (LSO)**
  - *Sam Brinker*
    - LSO Assistant Manager for National Security Implementation and ECP Deputy Federal Project Director

EXASCALE COMPUTING PROJECT

For more information…

https://www.exascaleproject.org

or reach out to the leadership team in the areas that interest you..

Doug Kothe, ECP Director
kothe@ornl.gov

Lori Diachin, ECP Deputy Director
diachin2@llnl.gov

Andrew Siegel, Application Development Director
seigela@uchicago.edu.

Erik Draeger, Application Development Deputy Director
draeger1@llnl.gov

Mike Heroux, Software Technology Director
maherou@sandia.gov

Jonathan Carter, Software Technology Deputy Director
jtcarter@lbl.gov

Terri Quinn, Hardware and Integration Director
quinn1@llnl.gov

Susan Coghlan, Hardware and Integration Deputy Director
smc@alcf.anl.gov

Questions?

NVIDIA's tesla v100

- 5,120 CUDA cores (64 on each of 80 SMs)
- 640 NEW Tensor cores (8 on each of 80 SMs)
- 20MB SM RF | 16MB Cache | 16GB HBM2 @ 900 GB/s
- 300 GB/s NVLink
- 7.5 FP64 TFLOPS | 15 FP32 TFLOPS | 120 Tensor TFLOPS
- >27K of these on ORNL's Summit system!
- Mixed precision matrix math 4x4 matrices

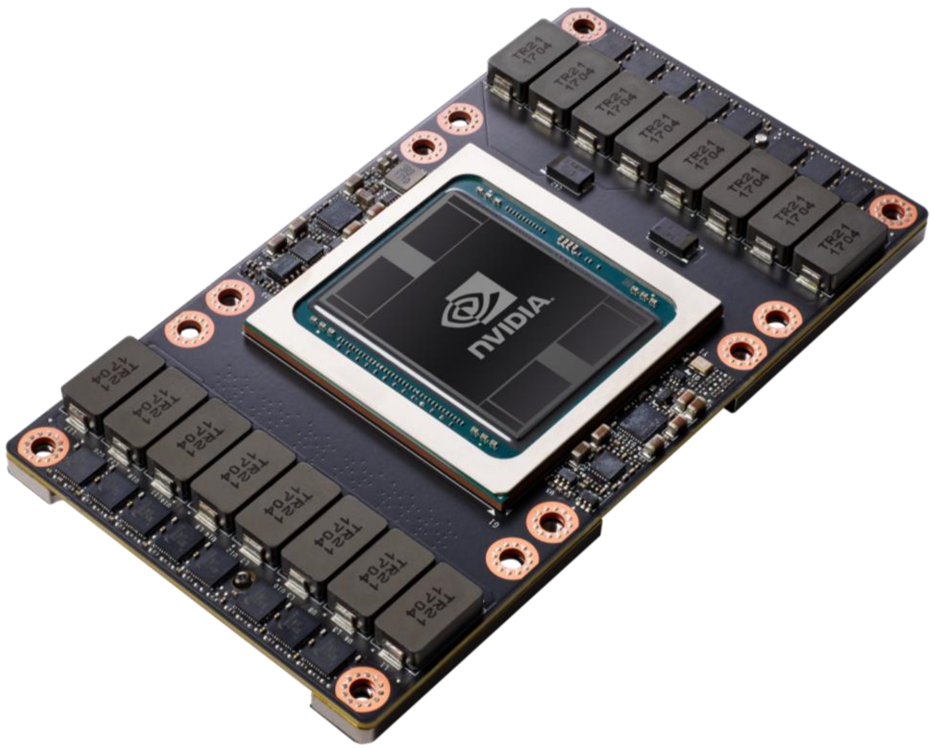

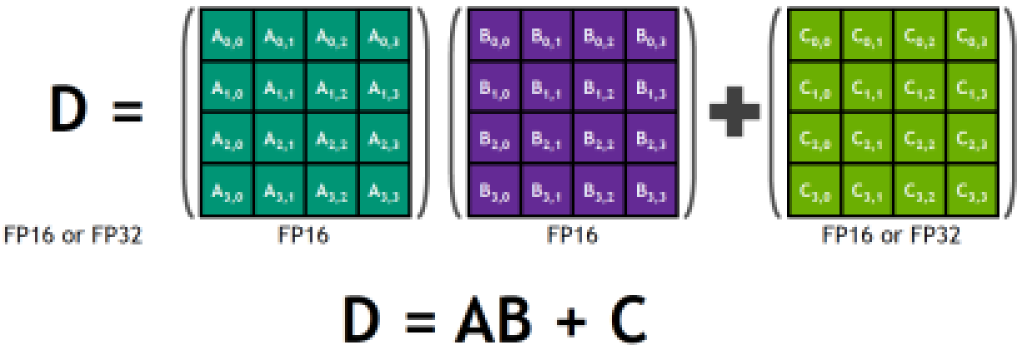

Type	Size	Range	$u = 2^{-t}$
half	16 bits	$10^{\pm 5}$	$2^{-11} \approx 4.9 \times 10^{-4}$
single	32 bits	$10^{\pm 38}$	$2^{-24} \approx 6.0 \times 10^{-8}$
double	64 bits	$10^{\pm 308}$	$2^{-53} \approx 1.1 \times 10^{-16}$
quadruple	128 bits	$10^{\pm 4932}$	$2^{-113} \approx 9.6 \times 10^{-35}$

EXASCALE COMPUTING PROJECT

- The M&S community must figure how out to "cheat" and utilize mixed / reduced precisions
- Ex: Jack Dongarra shows he can get 4x FP64 peak for 64bit LU on V100 with iterative mixed precision (using GMRES!)

Principles for a Healthy ECP and Facilities Partnership

Make it a win-win partnership

- ECP outputs are usable by and meet the needs of the Facilities
- Facilities make available the expertise and resources needed by ECP

Leverage each others capacities

Two examples are:

- ECP makes use of the Facilities' application preparation programs for ECP applications
- ECP's early hardware R&D investments improve the systems DOE Facilities are acquiring

Align our plans to the extent that makes sense

- Two examples are:
- Both are interested in improving software quality and ease of deployment within the Facilities' HCP centers
- Both would like ECP applications to be ready to run on their exascale system

Application Co-Design (CD)

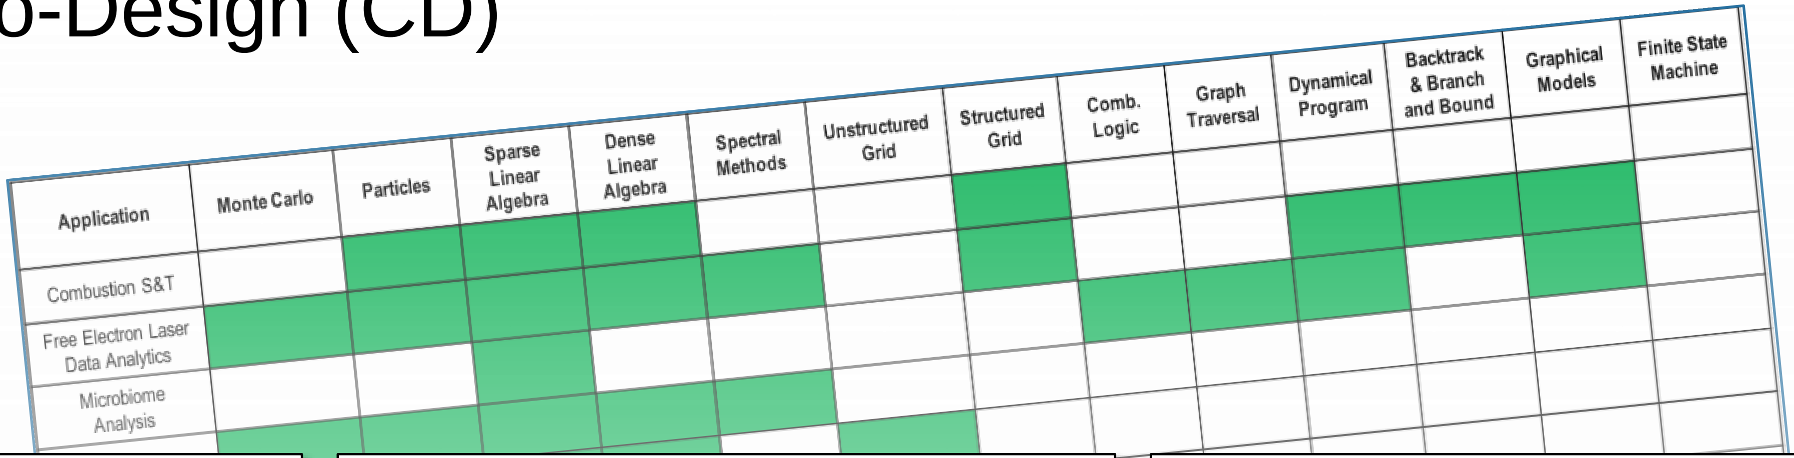

Essential to **ensure that applications effectively utilize** exascale systems

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

CD Centers **focus 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
- An appropriate nexus for reduced precision?

EXASCALE COMPUTING PROJECT

E3SM: Energy Exascale Earth System Model

- Global Earth System Model
- Atmosphere, Land, Ocean and Ice component models
- 8 DOE labs, 12 university subcontracts, 53 FTEs spread over 87 individuals
- Development driven by DOE Office of Scienc mission interests: Energy/water issues looking out 40 years
- Key computational goal: Ensure E3SM will run well on upcoming DOE pre-exascale and exascale computers
- E3SM is open source / open development
  - Website: www.e3sm.org
  - Github: https://github.com/E3SM-Project
  - DOE Science youtube channel: https://www.youtube.com/channel/UC_rhpi0lBeD1U-6nD2zvlBA

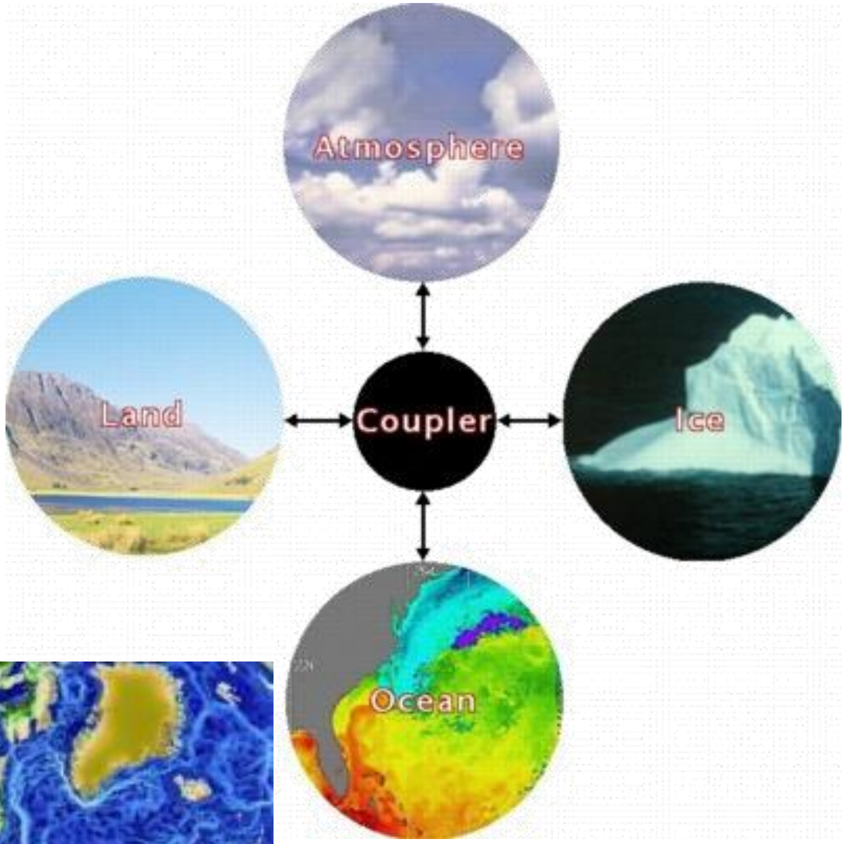

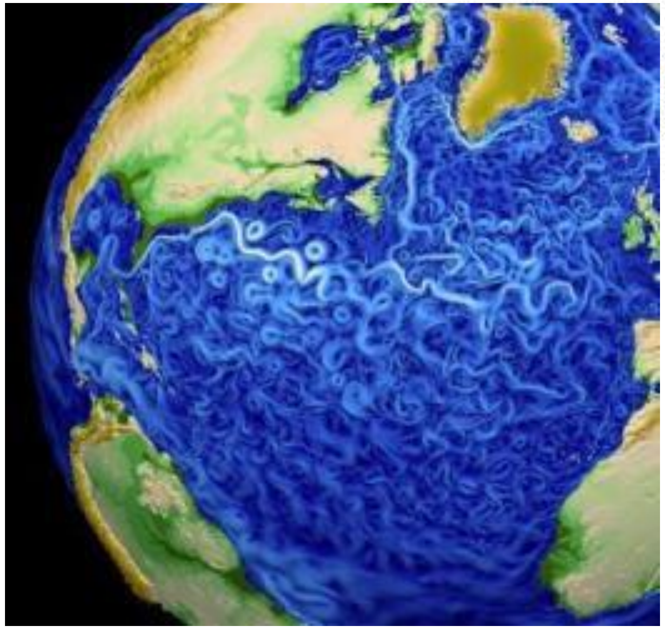

PI: David Bader (LLNL)

We work on products applications need now and into the future

Key themes:
- Exploration/development of new algorithms/software for emerging HPC capabilities:
- High-concurrency node architectures and advanced memory & storage technologies.
- Enabling access and use via standard APIs.

Software categories:
- The next generation of well-known and widely used HPC products (e.g., MPICH, OpenMPI, PETSc)
- Some lesser used but known products that address key new requirements (e.g., Kokkos, RAJA, Spack)
- New products that enable exploration of emerging HPC requirements (e.g., SICM, zfp, UnifyCR)

Example Products	Engagement
MPI – Backbone of HPC apps	Explore/develop MPICH and OpenMPI new features & standards.
OpenMP/OpenACC –On-node parallelism	Explore/develop new features and standards.
Performance Portability Libraries	Lightweight APIs for compile-time polymorphisms.
LLVM/Vendor compilers	Injecting HPC features, testing/feedback to vendors.
Perf Tools - PAPI, TAU, HPCToolkit	Explore/develop new features.
Math Libraries: BLAS, sparse solvers, etc.	Scalable algorithms and software, critical enabling technologies.
IO: HDF5, MPI-IO, ADIOS	Standard and next-gen IO, leveraging non-volatile storage.
Viz/Data Analysis	ParaView-related product development, node concurrency.

Software Development Toolkit Motivation

- The exascale software ecosystem will be comprised of a wide array of software, all of which are expected to be used by DOE applications.
- The software must be:
  - interoperable
  - sustainable
  - maintainable
  - adaptable
  - portable
  - scalable
  - deployed at DOE computing facilities
- Without these qualities:
  - Value will be diminished
  - Scientific productivity will suffer
- Provides intermediate coordination points to better manage complexity

xSDK

Spack

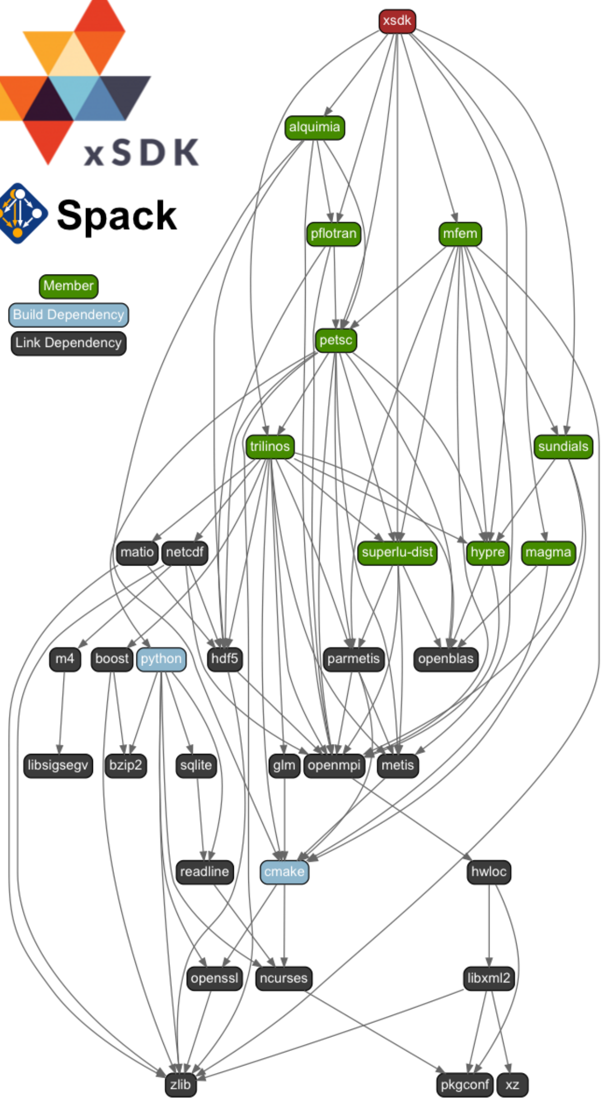

ECP ST SDK community policies: Important team building, quality improvement, membership criteria.

SDK Community Policy Strategy

- Review and revise xSDK community policies and categorize
  - Generally applicable
  - In what context the policy is applicable
- Allow each SDK latitude in customizing appropriate community policies
- Establish baseline policies in FY19 Q2, continually refine

xSDK compatible package: Must satisfy mandatory xSDK policies:

M1. Support xSDK community GNU Autoconf or CMake options.

M2. Provide a comprehensive test suite.

M3. Employ user-provided MPI communicator.

M4. Give best effort at portability to key architectures.

M5. Provide a documented, reliable way to contact the development team.

…

Prior to defining and complying with these policies, a user could not correctly, much less easily, build hypre, PETSc, SuperLU and Trilinos in a single executable: a basic requirement for some ECP app multi-scale/multi-physics efforts.

Recommended policies: encouraged, not required:

R1. Have a public repository.

R2. Possible to run test suite under valgrind in order to test for memory corruption issues.

R3. Adopt and document consistent system for error conditions/exceptions.

R4. Free all system resources it has acquired as soon as they are no longer needed.

R5. Provide a mechanism to export ordered list of library dependencies.

xSDK member package: An xSDK-compatible package, *that* uses or can be used by another package in the xSDK, and the connecting interface is regularly tested for regressions.

https://xsdk.info/policies

Initially the xSDK team did not have sufficient common understanding to jointly define community policies.

ECP ST Technologies that may be particularly suited to industry interactions

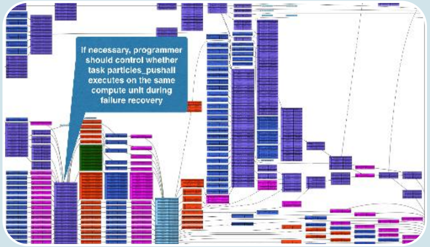

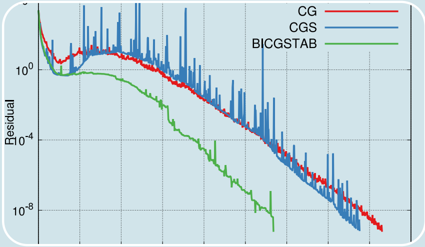

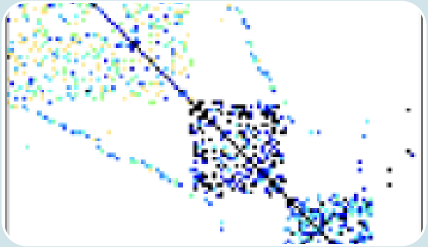

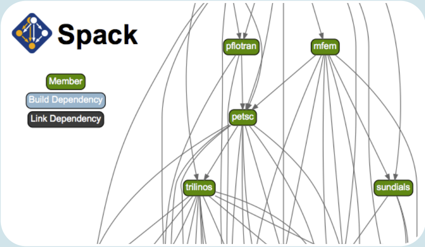

Programming Models & Runtimes

- Leverage new features in MPICH, OpenMP libraries
- Use C++ compile-time polymorphism to generate node-specific code from common source code (e.g., Kokkos, RAJA)
- Experiment with alternative programming models (Legion, UPC++/GASNet)

Development Tools

- Tools for performance analysis:
  - PAPI, TAU, HPC Toolkit, Dyninst:
  - Widely used in HPC community
- Portable, open-source LLVM compiler ecosystem to support expected ECP architectures, including support for F18

Math Libraries

- Use hypre, PETSc, SuperLU, Trilinos, others: All widely used parallel solvers being adapted for massive on-node concurrency.
  - APIs are largely unchanged
  - Provides performance portability across platforms
- Try STRUMPACK
  - Suitable SuperLU replacement
  - Highly scalable (for a direct solver).
  - Turnkey solver (easy to install and use)

Data and Visualization

- New storage software and workflows associated with non-volatile memory
  - Fundamental I/O game-changer
  - Examples: Fast offload of checkpoints, all-flash storage system
- Data compression tools: Same impact as increasing memory and storage size and bandwidth.
- In situ workflows: Increased opportunities to analyze and transform data as part of the workflow.

Software Ecosystem

- Advanced resource management:
  - Fast, scalable checkpoint/restart (leverage NVRAM).
  - Resource managers, e.g., Flux.
- SDKs and Spack are emerging as attractive combination for managing software components:
  - Involvement and input from industry can be beneficial both ways

Hardware and integration

Develop technology advances for exascale and deploy ECP products

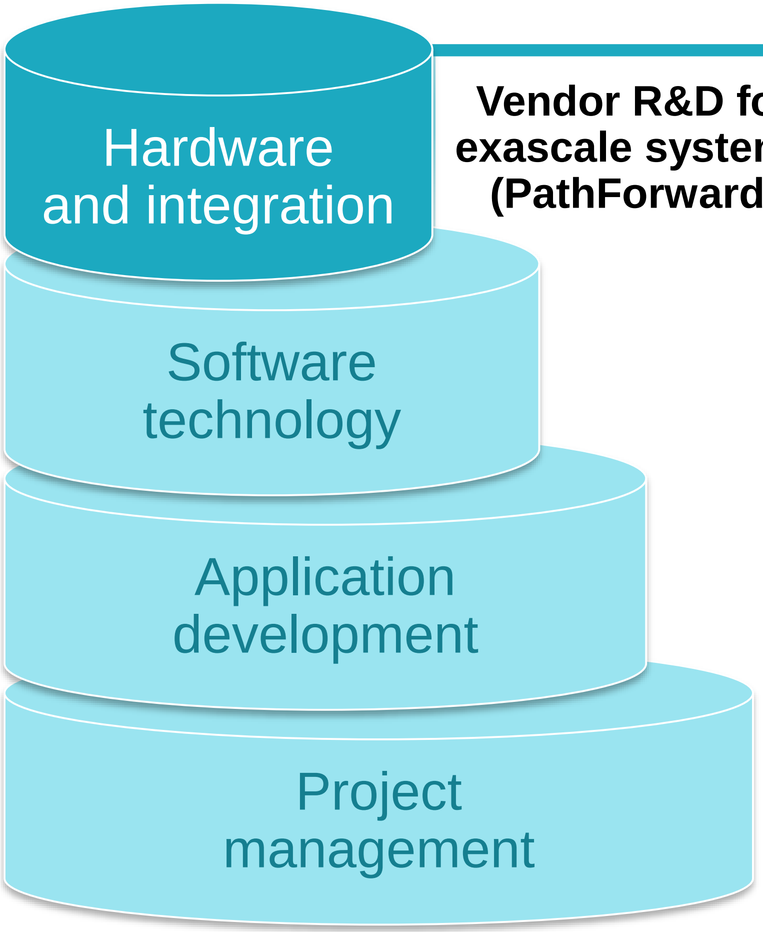

- **Vendor R&D for exascale systems (PathForward)**
- Evaluation of hardware technology and performance
- **Software deployment and application integration at HPC facilities**
- Pre-exascale and exascale system utilization measurement and tracking
- Community training and productivity

HI has primary ECP responsibility for the partnership with the Facilities and each HI project is collaborating with the Facilities.

EXASCALE COMPUTING PROJECT

PathForward vendors, objectives, and R&D

Company (HPC Vendor)	Objective	R&D thrusts
AMD	Develop innovative technologies to enable our system integration partners to design and assemble a variety of world-class solutions for exascale computing	• Innovations in memory interfaces • CPU and GPU microarchitecture • Component integration • High-speed interconnects
Cray Inc.	Improvements in sustained performance, power efficiency, scalability, and reliability	• Arm processor enhancements • Network enhancements • Memory architectures
Hewlett Packard Enterprise (HPE)	A system architecture and the integration technologies that can deliver an exascale system in a node-agnostic way	• System, node, and I/O design • Advance the Gen-Z interconnect • Optical interconnects
IBM	Optimize an exascale system for improved application performance and developer productivity while maximizing energy efficiency and reliability	• Architectural and system component innovations for a system combining processors, GPUs, high-performance networks, and high-performance storage
Intel Corp.	Energy efficiency, fabric costs, memory and high-speed IO, performance, scalability and usability	• Energy efficiency • Reduced fabric costs and power • Scalable storage and memory arch. • Optimized communication characteristics
NVIDIA Corp.	Accelerate efforts to develop highly efficient throughput computing technologies	• Energy-efficient GPU architectures • Resilience

Application Matching to Facilities Plan and Status

Strategy: Match applications with existing facility readiness efforts

Progress Assessment: Progress towards technical execution plans measured quarterly; annual external assessment.

Goal: 22 performant exascale applications that run on Aurora and/or Frontier

Application Development 2.2

- Chemistry and Materials Applications 2.2.1
- Energy Applications 2.2.2
- Earth and Space Science Applications 2.2.3
- Data Analytics and Optimization Applications 2.2.4
- National Security Applications 2.2.5
- Co-Design 2.2.6

ALCF

12 initial applications engaged by ALCF for Aurora. Other teams can follow best practices for Aurora readiness, and will be engaged as staffing allows.

OLCF

An initial set of ~10 ECP applications will be identified to participate in CAAR-ECP in FY19. Applications may transition in and out of the program as progress is made.

NERSC

5 ECP AD applications participating in NESAP for NERSC-9 . Additional applications may participate with NERSC funding.

Goal: Progress towards exascale readiness develops, and NESAP-ECP apps transition to LCF facilities

Delivering a secure, easy-to-use Continuous Integration (CI) solution to support Software Product testing on DOE HPC environments

- Allows for the verification of development efforts through automated building & testing across sites to better identify errors and improve code efficiency
- Continuous Integration/Continuous Delivery pipelines enabled through a combination of the web application and project configuration files and are executed by a selected runners

OnyxPoint + GitLab were chosen based upon a request for proposal with participation from the E6

Use cases being addressed and the solution being worked

Use Cases

- Build, test, package, and deploy project code in a streamlined, repeatable workflow
- Develop CI pipelines that can directly leverage and submit jobs to HPC resources
- Target wider range of testing and deployment configurations by utilizing resources across facilities

Solution

- Use GitLab for the entire DevOps lifecycle (plan, create, verify, release) coupled with runner components located across all sites
- Fund Facility staff to define the solution and to deploy and operate the service at their site
- Onyx Point developed runner improvements for DOE secure HPC data centers (identifying users uniquely and enabling submitting jobs directly to resource job management system)
- Centralized core repository managed by DOE Office of Scientific and Technical Information (OSTI) accessible by ECP and E6 users
- Shared runner strategy that offers security tied to user authentication while minimizing system administration overhead – A runner that serves all projects is called a shared Runner

Facilities and ECP Partnering is Key to ECI

- Leverage ECP capabilities
- Provides expertise, staff, and computer resources
- Acquires, deploys, and operates exascale computers
- Shares in the deployment of exascale computing capabilities

- Mission is enduring
- ASCR has been delivering on its mission for decades

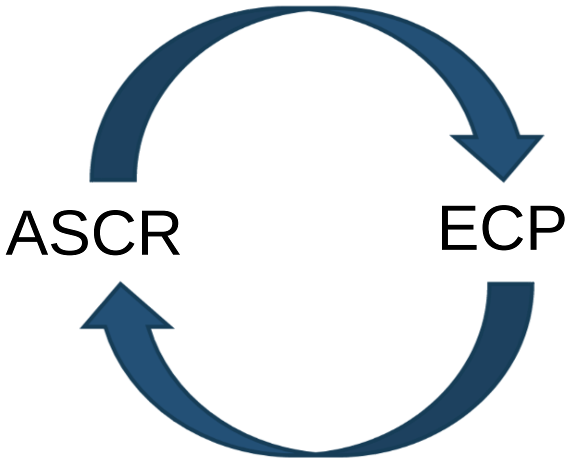

- Mission ends with the project
- ECP is a few years old

- Leverage ASCR capabilities
- Products delivered are accelerated apps, sw, and hw
- Shares in the deployment of exascale computing capabilities

EXASCALE COMPUTING PROJECT

More detailed information about the applications will be available in the ECP Application Assessment Report

- Comprehensive assessment of ECP's application projects and 7 co-design centers
- Primary document elements:
  - Completion requirements, including detailed description of exascale challenge problem and figure of merit (FOM) formula
  - FY18 milestone execution
  - DOE stakeholders
  - Software details, integration with ST projects
- Teams and AD leadership already acting on these findings
- Public version available May 2019


EXASCALE COMPUTING PROJECT

Status and Progress of ECP Applications
Application Development

Andrew Siegel[1], Erik W. Draeger[2], Jack Deslippe[3], Tom Evans[4], Anshu Dubey[1], William Hart[5], and Phillip Colella[3]

December 31, 2018

[1] Argonne National Laboratory (ANL)
[2] Lawrence Livermore National Laboratory (LLNL)
[3] Lawrence Berkeley National Laboratory
[4] Oak Ridge National Laboratory
[5] Sandia National Laboratory

EXASCALE COMPUTING PROJECT

ACM: The Learning Continues…

- Questions/comments about this webcast? learning@acm.org
- ACM Code of Ethics: https://ethics.acm.org
- ACM's Discourse Page: http://on.acm.org
- ACM TechTalks (on-demand archive): https://learning.acm.org/techtalks
- ACM Learning Center: http://learning.acm.org
- ACM SIGHPC: https://www.sighpc.org/
- ACM Queue: http://queue.acm.org