Exascale Computing Project -- Software

Paul Messina, ECP Director

Stephen Lee, ECP Deputy Director

ASCAC Meeting, Arlington, VA

Crystal City Marriott
April 19, 2017

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

www.ExascaleProject.org

ECP scope and goals

Develop applications to tackle a broad spectrum of mission critical problems of unprecedented complexity

Support national security

Partner with vendors to develop computer architectures that support exascale applications

Contribute to the economic competitiveness of the nation

Develop a software stack that is both exascale-capable and usable on industrial & academic scale systems, in collaboration with vendors

Train a next-generation workforce of computational scientists, engineers, and computer scientists

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 ons	Scalable and productive	Hardware technology elen	Integrated exascale supercomputers

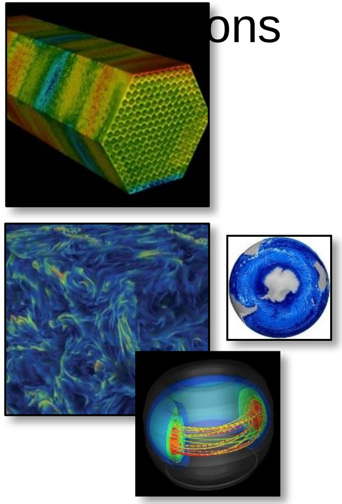

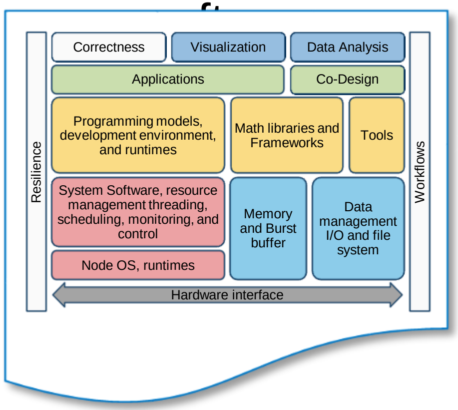

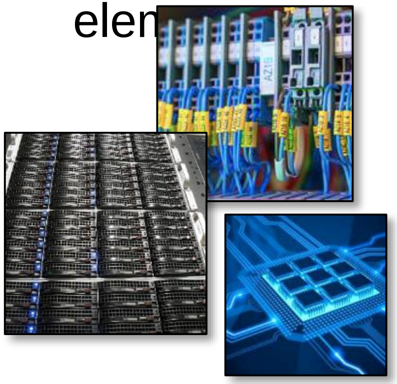

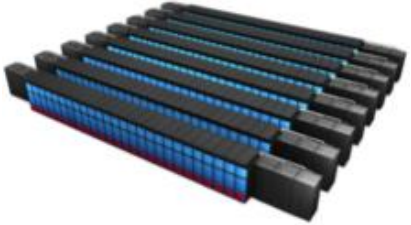

ECP's work encompasses applications, system software, hardware technologies and architectures, and workforce development

EXASCALE COMPUTING PROJECT

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 (perceived fault rate: ≤1/week)
- Includes a software stack that supports a broad spectrum of applications and workloads

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

The ECP Plan of Record

- A 7-year project that follows the *holistic/co-design* approach, which runs through 2023 (including 12 months of schedule contingency)
  - To meet the ECP goals
- Enable an initial exascale system based on *advanced architecture* and 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 facilities

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 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

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

ST 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

EXASCALE COMPUTING PROJECT

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

EXASCALE COMPUTING PROJECT

ECP Requires Strong Integration to Achieve Capable Exascale

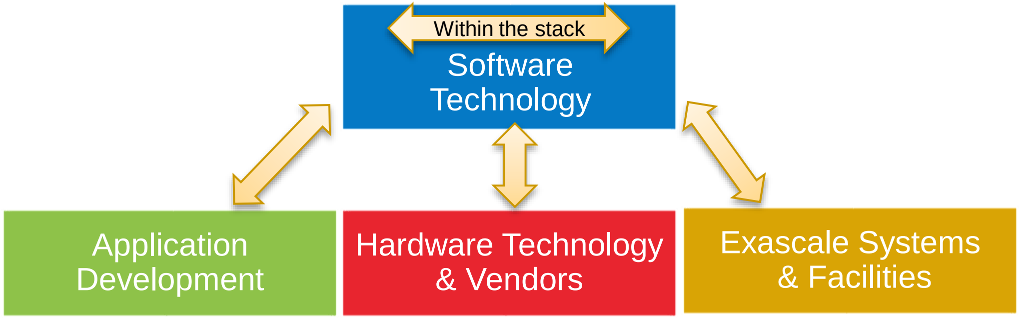

- To achieve a coherent software stack, we must integrate across all the focus areas
  - Understand and respond to the requirements from the apps but also help them understand challenges they may not yet be aware of
  - Understand and repond to the impact of hardware technologies and platform characteristics
  - 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)

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

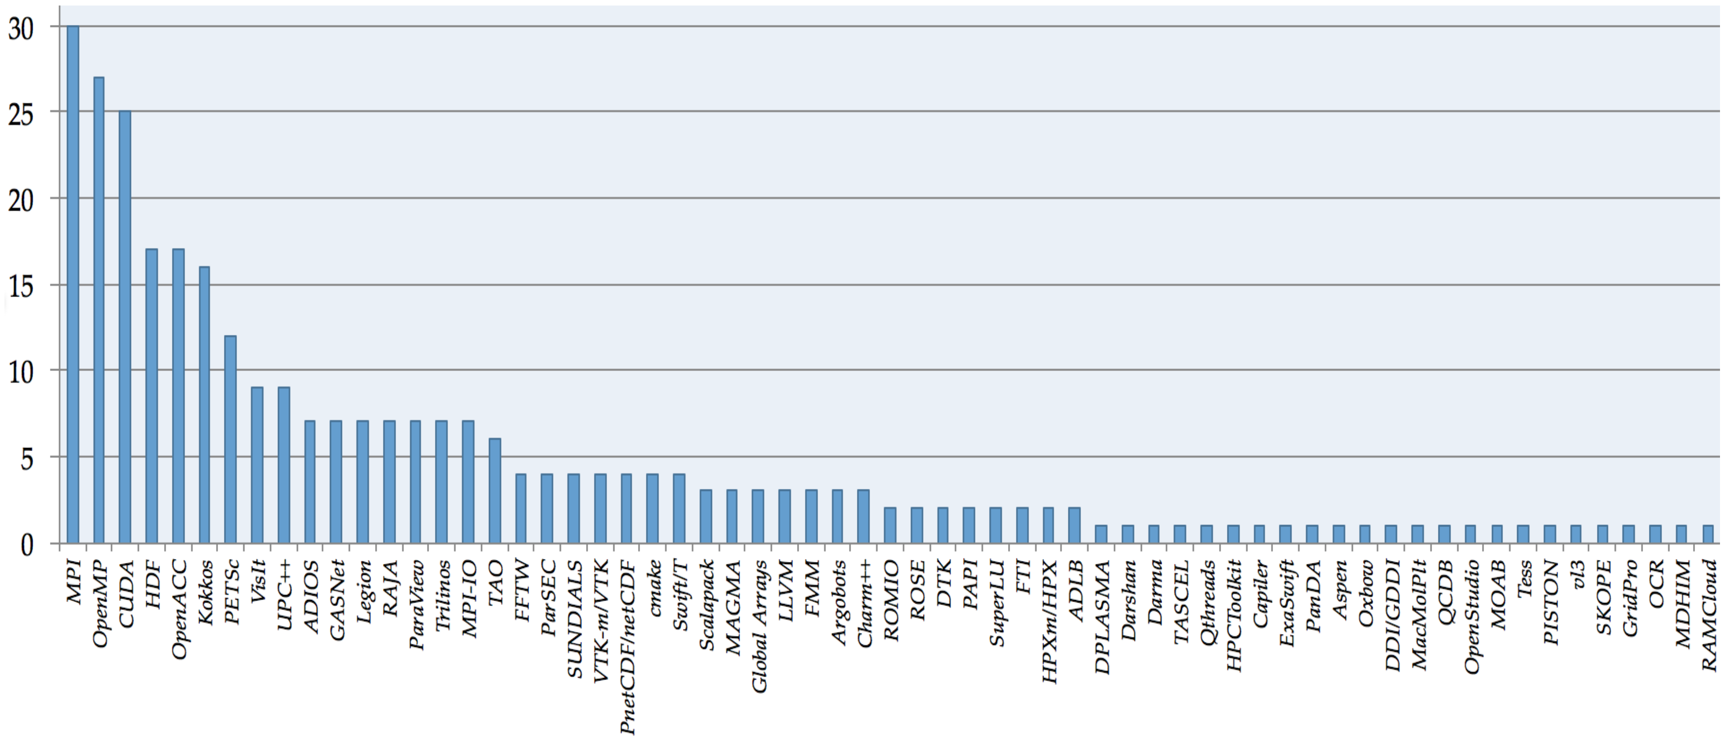

Libraries used at NERSC
(similar data from other facilities)

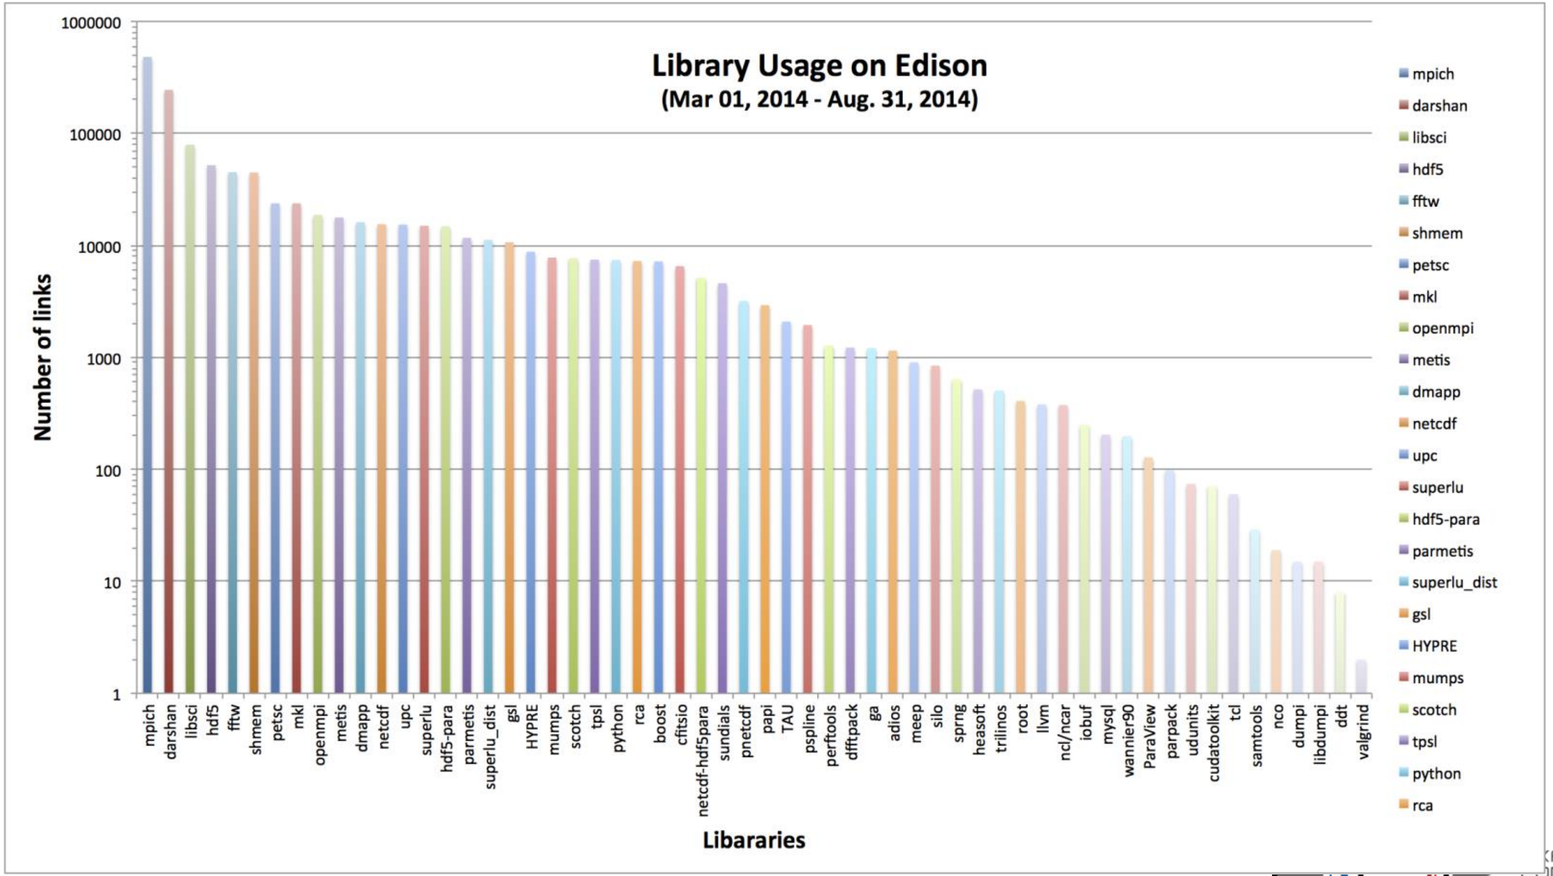

Conceptual ECP Software Stack

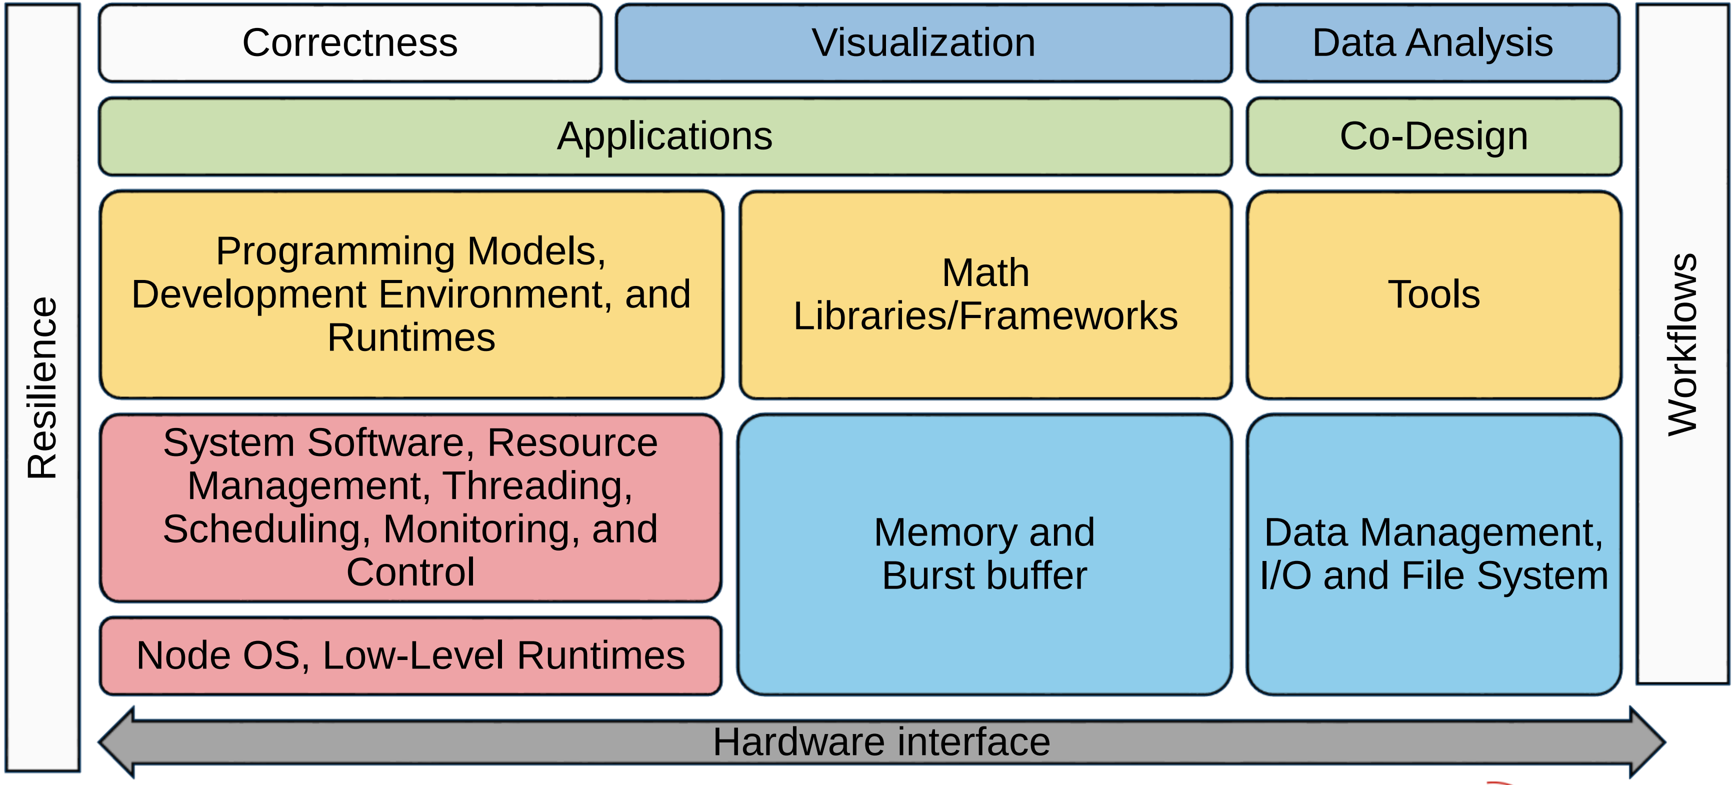

Current Set of ST Projects Mapped to Software Stack

Resilience
Checkpoint/Restart (VeloC, UNIFYCR), FSEFI, Fault Modeling

Correctness

Visualization
VTK-m, ALPINE, Cinema

Data Analysis
ALPINE, Cinema

Applications

Co-Design

Programming Models, Development Environment, and Runtimes
MPI (MPICH, Open MPI), OpenMP, OpenACC, PGAS (UPC++, Global Arrays), Task-Based (PaRSEC, Legion, DARMA), RAJA, Kokkos, OMPTD, Power steering

Math Libraries/Frameworks
ScaLAPACK, DPLASMA, MAGMA, PETSc/TAO, Trilinos, xSDK, PEEKS, SuperLU, STRUMPACK, SUNDIALS, DTK, TASMANIAN, AMP, FleCSI, KokkosKernels, Agile Comp., DataProp, MFEM

Tools
PAPI, HPCToolkit, Darshan, Perf. portability (ROSE, Autotuning, PROTEAS), TAU, Compilers (LLVM, Flang), Mitos, MemAxes, Caliper, AID, Quo, Perf. Anal.

System Software, Resource Management Threading, Scheduling, Monitoring, and Control
Argo Global OS, Qthreads, Flux, Spindle, BEE, Spack, Sonar

Node OS, low-level runtimes
Argo OS enhancements, SNL OS project

Memory and Burst buffer
Chkpt/Restart (VeloC, UNIFYCR), API and library for complex memory hierarchy (SICM)

Data Management, I/O and File System
ExaHDF5, PnetCDF, ROMIO, ADIOS, Chkpt/Restart (VeloC, UNIFYCR), Compression (EZ, ZFP), I/O services, HXHIM, SIO Components, DataWarehouse

Workflows
Contour, Siboka

Hardware interface

EXASCALE COMPUTING PROJECT

Kokkos: On-Node Performance Portable Abstraction Layer 1.3.1

Description and Scope

Project Description:

- Enable application performance portability across next generation platforms (NGPs) on likely paths to exascale; especially DOE's advanced technology systems (ATS)
- Sustain application developer productivity with effective programming model, support, and training that evolves with NGP architectures and sophistication of application algorithms

ECP Scope

- ASC/ATDM funded programming model R&D and ASC application support
- ECP-ST funded DOE-Lab-wide training and application support

Collaborators

ASC and ECP Applications (AD, U, R)

- Sierra Mechanics
- Trilinos algorithms and enabling technologies
- EMPIRE/EMPRESS and SPARC
- Next-Generation Code (NGC) for multi-physics simulations
- Exascale Lattice Gauge Theory Opportunities and Requirements for Nuclear and High Energy Physics
- Molecular Dynamics at the Exascale: Spanning the Accuracy, Length and Time Scales for Critical Problems in Materials Science
- Exascale Predictive Wind Plant Flow Physics Modeling
- QMCPACK: A Framework for Predictive and Systematically Improvable Quantum-Mechanics Based Simulations of Materials
- Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors
- Transforming Additive Manufacturing through Exascale Simulation (TrAMEx)
- Transforming Combustion Science and Technology with Exascale Simulations
- Cloud-Resolving Climate Modeling of the Earth's Water Cycle

Maturity Level and Accessibility

Maturity Indicators and Metrics

- Production since March 2015, support and R&D is ongoing
- Development of first prototyping began in 2010

Accessibility

- github.com/kokkos/kokkos
- BSD 3 clause license

Current users

- numerous @ SNL, LANL, ORNL, U-Utah, ARL, NRL,

FY17 Development Plan

- Support application developers through training, consulting, and enhancements
- Develop initial capability enabling applications to manage and asynchronously utilize NGP node's heterogeneous processing and memory resources
- Develop initial capability to support application algorithms using the on-node parallel pattern of a directed acyclic graph (DAG) of fine grain tasks
- Advocate for improvements to ISO/C++ and other standards for effective use of NGPs

PI: H. Carter Edwards (SNL)

Extending HPCToolkit to Measure and Analyze Code Performance on Exascale Platforms

1.3.2

Description and Scope

Project Description:

Rice University's HPCToolkit performance tools supports measurement, analysis, attribution, and presentation of parallel application performance.

ECP Scope

Extend HPCToolkit with improved capabilities for measurement and analysis of computation, data movement, communication, I/O, power, and energy consumption at extreme scale. Planned enhancements include improved support for lightweight measurement of massive node-level parallelism using hardware counters, measuring OS activity, recording measurement data efficiently, analyzing performance data in parallel, and integrating code-centric, time-centric, data-centric, and resource-centric views for analysis.

Maturity Level and Accessibility

Maturity Indicators and Metrics

- *Released in production*
- *18 years in development*

Accessibility

- *github.com/hpctoolkit, github.com/dyninst*
- *BSD license*

Current users

- *Installed at ANL, LANL, LLNL, NERSC, ORNL, PNNL, SNL*

Collaborators

Potential or likely ties to other ECP Projects

- Engage vendor and co-design efforts to design better instrumentation technologies for performance introspection on future platforms (HT/CD, D, E)
- Partner with developers of programming models and runtimes to add introspection interfaces that deliver insight into application performance (ST, D, E)
- Collaborate with NNSA and ST efforts developing instrumentation technologies for applications and runtimes, as well as measurement and analysis of data movement and network traffic (ST, U, E)
- Help ECP application teams use HPCToolkit to analyze code scalability and performance (AD, U, E)

Synergistic projects

- Exa-PAPI: The exascale performance and application programming interface
- ADTM projects: Caliper, Memtools, NetOpt, OMPT/OMPD

First Year Development Plan

Curate releases of HPCToolkit in Q1 and Q4

Finalize OMPT performance tools interface for OpenMP

Leverage Linux perf_events interface to measure and attribute computation, data movement, OS activity, and blocked threads

Parallelize program structure analysis to aid attribution of performance information

Engage external collaborators and ECP application teams

EXASCALE COMPUTING PROJECT

PI: Mellor-Crummey (Rice University) Partner sites: (University of Wisconsin - Madison)

Exascale MPI

Description and Scope

Project Description:

The Exascale MPI project will design and develop enhancements to MPI that meet the critical needs of MPI applications. The project will focus on three primary capabilities: (1) improvements to the MPI implementation through MPICH; (2) improvements to the MPI standard; and (3) technology transfer to industry to ensure the availability of high-performance MPI implementations on DOE supercomputers.

ECP Scope

Exascale MPI will enable applications to effectively use the latest advances in MPI to scale to the largest supercomputers in the world. It will produce a high performance MPI implementation and affect changes to the MPI standard to meet application and architectural requirements.

Maturity Level and Accessibility

Maturity Indicators and Metrics

- *MPICH is a production software that has been under continuous development for 23 years. MPICH derivatives are used on most supercomputers in the world as the primary MPI implementation.*

Accessibility

- *http://www.mpich.org (download and source links are available from there)*
- *BSD*

Current users

- *Practically all DOE applications and supercomputer centers use MPICH and derivatives. MPICH derivatives include Intel MPI, Cray MPI, IBM BG/Q MPI, Microsoft MPI, MVAPICH.*

Collaborators

Potential or likely ties to other ECP Projects

- *Most ECP applications employ MPI as part of their exascale strategy. We will collaborate with them to better understand the gaps in current implementations or the standard and will work closely with selected teams to develop solutions. Other potential collaborators:*
  - *SOLLVE: Scaling OpenMP with LLVM (ST, E)*
  - *Enhancing Qthreads (ST, E)*
  - *Argo: Operating System and Resource Management (ST, E)*

Synergistic projects

- *OpenMPI for Exascale*

First Year Development Plan

We will

- *survey ECP apps teams to determine needs with respect to MPI implementation changes or standard changes;*
- *select key applications and mini-apps to guide and validate our work;*
- *prototype changes and work with applications to guide future changes;*
- *form teams with external participants to explore selected areas;*
- *integrate changes into releases and work with industry to incorporate first draft of changes for pre-exascale machines (such as CORAL)*

PI: P. Balaji (ANL)

EXASCALE COMPUTING PROJECT

Co-design in action: shared milestone examples

- OMPI-X ST project (Open MPI for exascale) will design, implement, and demonstrate new prototype APIs to facilitate improved coordination between MPI and OpenMP runtimes for use of threads, message delivery, and rank/endpoint placement. (Due 12/31/2018)
- Execution Plan:
  - Interact with ECP Application Development teams to identify requirements about runtime coordination.
  - Interact with appropriate ECP Software Technology teams to define coordination APIs for MPI and OpenMP (and other) runtimes.
  - Experiment and demonstrate the benefit running an ECP application or mini-application running on the ECP testbeds

Shared milestones examples

- SUNDIALS and the AMReX co-design center
  - Enable CVODE to be used from AMReX (Due 6/30/2017)
  - Description: AMReX will provide a build system and interface to CVODE within AMReX SUNDIALS will support AMReX in being able to link to CVODE from within the AMReX build system.
- Evaluation of VeloC (efficient checkpoint/restart) with the QCD application code
- Evaluation of VeloC (efficient checkpoint/restart) with the EXAALT application code (Molecular Dynamics at the Exascale application)

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

Plan to involve additional Universities and Vendors

- The current set of Software Technology projects span 8 DOE national laboratories,15 universities, and 2 independent software vendors
- In FY17-18, budget permitting, we plan to issue additional solicitations for universities and vendors to cover gaps identified by the gap analysis
- We plan to coordinate our activities with NSF to cover the broader university community through NSF
- ECP has formed an Industry Council. This group will also serve as an interface for engaging independent software vendors and identifying industry users' needs.
- Small businesses: DOE SBIR funding could support software of interest to ECP

Update on ECP activities

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

www.ExascaleProject.org

Capable exascale system applications will deliver broad coverage of 6 strategic pillars

National security

Stockpile stewardship

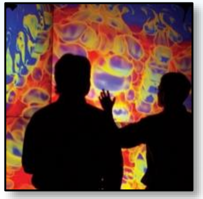

Energy security

Turbine wind plant efficiency

Design and commercialization of SMRs

Nuclear fission and fusion reactor materials design

Subsurface use for carbon capture, petro extraction, waste disposal

High-efficiency, low-emission combustion engine and gas turbine design

Carbon capture and sequestration scaleup

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Urban planning

Reliable and efficient planning of the power grid

Seismic hazard risk assessment

Scientific discovery

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Plasma wakefield accelerator design

Light source-enabled analysis of protein and molecular structure and design

Find, predict, and control materials and properties

Predict and control stable ITER operational performance

Demystify origin of chemical elements

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

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

27 Application Development milestones delivered in Q2FY2017 so far; selected examples

- GAMESS (1.2.1.16)
  - Initiate the development of a GAMESS-QMCPACK interface in collaboration with the QMCPACK group
  - Complete and assess an initial threaded GAMESS RI-MP2 energy + gradient code, conduct benchmarks
- Urban (1.2.1.17)
  - Define data streams between the four proposed models, including what each model is capable of providing and what each model could leverage as input
  - Hold planning/strategy workshop with five-lab team and inviting ECP software teams to assist in evaluating existing coupling/data-exchange systems
- AMReX (1.2.5.3.03)
  - Performance characterization of current AMR application codes

Exascale Proxy Applications Suite

Objectives and Scope

- Assemble and curate a proxy app suite composed of proxies developed by other ECP projects that represent the most important features (especially performance) of exascale applications.
- Improve the quality of proxies created by ECP and maximize the benefit received from their use. Set standards for documentation, build and test systems, performance models and evaluations, etc.
- Collect requirements from app teams. Assess gaps between ECP applications and proxy app suite. Ensure proxy suite covers app motifs/requirements.
- Coordinate use of proxy apps in the co-design process. Connect consumers to producers. Promote success stories and correct misuse of proxies.

Links to Other ECP Projects

- **Application Assessment Project:** Cooperatively assess and quantitatively compare applications and proxy apps.
- **Design Space Evaluation Team:** Will need proxies specially adapted for hardware simulators.
- **Path Forward Vendors:** Evaluate needs and provides proxies & support. Review proxy app usage and results.
- **Application Development Projects & Co-Design Centers:** Producers of proxy apps. Close the loop with lessons learned from proxies.
- **Software Technology Projects:** Consumers of proxy apps. Use proxies to understand app requirements and to test and evaluate proposed ST offerings.

Risks and Challenges

- Proxies are too complex for simulators or HW design activity.
- Proxies fail to accurately represent parent apps.
- Proxy app authors unable or unwilling to meet quality and/or support standards.
- Full coverage of DOE workload/motifs produces a large and unwieldy suite.
- Proxy apps misused by consumers.
- Inability to balance agility and stability/quality of proxies.

Development Plan

- Release updated versions of the proxy app suite every six months. This cadence allows for improved coverage and changing needs while maintaining needed stability.
- Annually update guidance on quality standards. Increase rigor of standards.
- Meet with each application project at least quarterly to maintain a catalog of their requirements, proxies, and key questions for which they are seeking assistance.
- Publish annual proxy app producer report with requirements and assessment of proxies in comparison to parent apps.
- Publish annual proxy app consumer report with success stories, surveys of how proxy consumers are using proxies, and plans to satisfy any unmet needs.

PI: David Richards (LLNL); Institutions: ANL, LANL, LBNL, ORNL, SNL

IDEAS-ECP – Advancing Software Productivity for Exascale Applications

WBS 1.2.4.02

Description and Scope

- **Software Challenges:** Exploit massive on-node concurrency and handle disruptive architectural changes while working toward predictive simulations that couple physics, scales, analytics, and more
- **Goals:** Improve ECP developer productivity and software sustainability, as key aspects of increasing overall scientific productivity
- **Strategy:** In collaboration with ECP community:
  - **Customize and curate methodologies** for ECP app productivity & sustainability
  - **Create an ECP Application Development Kit** of customizable resources for improving scientific software development
  - **Partner with ECP application teams** on software improvements
  - **Training and outreach** in partnership with DOE computing facilities

History and Accessibility

IDEAS productivity

IDEAS: Interoperable Design of Extreme-scale Application Software

- Project began in Sept 2014 as ASCR/BER partnership to improve application software productivity, quality, and sustainability

Resources: https://ideas-productivity.org/resources **Highlights:**

- ***WhatIs and HowTo* docs:** concise characterizations & best practices
  - *What is CSE Software Testing?*
  - *What is Version Control?*
  - *What is Good Documentation?*
  - *How to Write Good Documentation*
  - *How to Add and Improve Testing in a CSE Software Project*
  - *How to do Version Control with Git in your CSE Project*

IDEAS productivity Howto

- **Webinar series, 2016: *Best Practices for HPC Software Developers***
  - *What All Codes Should Do: Overview of Best Practices in HPC Software Development*
  - *Developing, Configuring, Building, and Deploying HPC Software*
  - *Distributed Version Control and Continuous Integration Testing*
  - *Testing and Documenting your Code*
  - more topics …

Slides and videos

Collaborators

IDEAS-ECP team: Catalysts for engaging ECP community on productivity issues

- **Partnerships with ECP applications teams**
  - Understand productivity bottlenecks and improve software practices
- **Collaborate to curate, create, and disseminate software methodologies, processes, and tools** that lead to improved scientific software
  - ECP software technologies projects, applications teams, co-design centers, broader community
  - Software Carpentry-type approach to training for extreme-scale software productivity topics
- **Web-based hub for collaborative content development and delivery**
  - Community-driven collection of resources to help improve software productivity, quality, and sustainability

better scientific software

Phase-1 release: Summer 2017

scientific

First Year Plan

- **MS1/Y1: Templates for Productivity and Sustainability Improvement Plans (PSIPs) and Progress Tracking Cards**
  - Framework for software teams to identify, plan, and track improvements in productivity and sustainability
- **MS2/Y2: Interviews with Phase-1 ECP applications teams**
  - Understand current software practices, productivity challenges, preferred modes of collaboration, and needs for on-line ECP knowledge exchange
  - Determine prioritized needs for productivity improvement; initiate Phase-1 application partnerships
- **MS3/Y1: Training and outreach to ECP community on productivity and sustainability**
  - Develop and deliver tutorials, webinars, and web content

EXASCALE COMPUTING PROJECT

Co-Lead PIs: M. Heroux (SNL) and L.C. McInnes (ANL); Partner sites: LBNL, LLNL, LANL, ORNL, Univ of Oregon

IDEAS-ECP – Highlights of Recent Activities

IDEAS productivity

WBS 1.2.4.02

SIAM CSE17 Conference (Feb 27-Mar 3, 2017)

Premier CSE Conference (largest SIAM meeting, over 1700 attendees)

- **Invited tutorial:** *CSE Collaboration through Software: Improving Software Productivity and Sustainability*
  - Presenters: D. Bernholdt, A. Dubey, M. Heroux, A. Klinvex, L.C. McInnes
    - *Why Effective Software Practices Are Essential for CSE Projects*
    - *An Introduction to Software Licensing*
    - *Better (Small) Scientific Software Teams*
    - *Improving Reproducibility through Better Software Practices*
    - *Testing of HPC Scientific Software*

  Slides and audio: https://doi.org/10.6084/m9.figshare.c.3704287

available via https://ideas-productivity.org/events
Transitioning summer 2017 to
better scientific software

- **Minisymposium and Minisymposterium:** *Software Productivity and Sustainability for CSE and Data Science*
  - Organizers: D. Bernholdt, M. Heroux, D. Katz, A. Logg, L.C. McInnes
  - Slides for presentations: https://doi.org/10.6084/m9.figshare.c.3705946
    - 8 community presentations, including one from the IDEAS team
    - *CSE Software Ecosystems: Critical Instruments of Scientific Discovery*, L.C. McInnes
  - Posters: https://doi.org/10.6084/m9.figshare.c.3703771
    - 29 posters from the community, including several from members of IDEAS-ECP
- **Plenary Presentation:** *Productive and Sustainable: More Effective CSE*, M. Heroux
  - Slides and audio: https://doi.org/10.6084/m9.figshare.4728697

Productivity and Sustainability Improvement Planning Tools

Tools for helping a software team to increase software quality while decreasing the effort, time, and cost to develop, deploy, maintain, and extend software over its intended lifetime.

March 2017: Released PSIP templates & instructions: https://github.com/betterscientificsoftware/PSIP-Tools
Beginning Phase-1 interviews with ECP applications teams: Identify productivity bottlenecks and priorities.

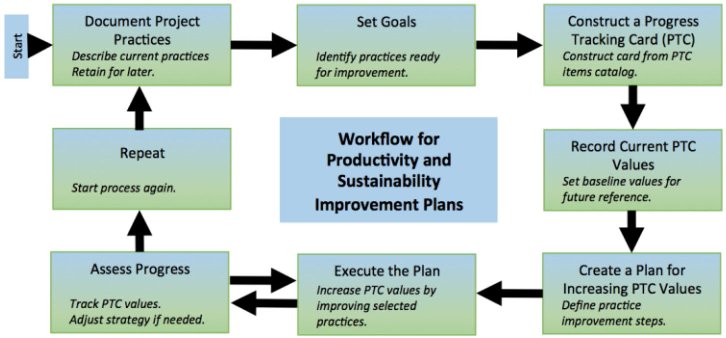

Practice: Test Coverage		Score (0 – 5):
Score Descriptions		
0	No unit testing.	
1	Unit testing used for one example refactored class.	
2	Two developers using unit testing process.	
3	Unit testing developer documentation completed.	
4	All developers trained to write unit tests.	
5	New policy established: All refactored code is covered by unit tests.	

Productivity and Sustainability Improvement Plan (PSIP): a lightweight iterative workflow to identify, plan, and improve selected practices of a software project.

Other practices: Source management system, documentation, software distribution, issue tracking, developer training, etc.

Co-Lead PIs: M. Heroux (SNL) and L.C. McInnes (ANL); Partner sites: LBNL, LLNL, LANL, ORNL, Univ of Oregon

IDEAS-ECP buillds on the foundation of the IDEAS project, funded and supported by Thomas Ndousse-Fettter, Paul Bayer, and David Lesmes.

EXASCALE COMPUTING PROJECT

Argonne Training Program on Extreme-Scale Computing

extremecomputingtraining.anl.gov

Application deadline
March 10

What? An intensive two-week program on HPC methodologies applicable to both current and future supercomputers.

Who? Open to doctoral students, postdocs, and scientists who are conducting CS&E research on large-scale computers.

When? July 30 - August 11, 2017

Where? Q Center, St Charles, IL (USA)

# of Applicants Call closed March 10. Reviewing 167 applicants from 95 institutions worldwide. Expect to select up to 70 participants.

ATPESC 2017

July 30–August 11

PROGRAM CURRICULUM

Renowned scientists, HPC experts, and leaders serve as lecturers and guide hands-on sessions. ATPESC participants will be granted access to DOE Office of Science User Facilities, which are amongst the most powerful supercomputers in the world.

The core curriculum includes:

- Computer architectures and their predicted evolution
- Programming methodologies effective across a variety of today's supercomputers and that are expected to be applicable to exascale systems
- Numerical algorithms and mathematical software
- Approaches to building community codes for HP systems
- Data analysis, visualization, I/O, and methodologies and tools for Big Data applications
- Performance measurement and debugging tools

Argonne Training Program on Extreme-Scale Computing (ATPESC) is an intensive two-week program focused on HPC methodologies that are applicable to both current and future machine architectures, including exascale systems.

ATPESC provides advanced training to 60 participants. Qualified applicants must have:

- Substantial experience in MPI and/or OpenMP programming,
- Used at least one HPC system for a complex application, and
- Plans to conduct CS&E research on large-scale computers.

Admission to the ATPESC program is highly competitive. Participant support is provided, including domestic airfare, meals, and lodging.

ATPESC is part of the Exascale Computing Project, a collaborative effort of the DOE Office of Science and the National Nuclear Security Administration.

APPLICATION DETAILS

Applications for ATPESC 2017 are due by March 10, 2017. All doctoral students, postdocs, and computational scientists are encouraged to submit applications. To apply, visit the ATPESC website at extremecomputingtraining.anl.gov

EXASCALE COMPUTING PROJECT

Hardware Technology Activities

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

PathForward Status

- Five PathForward Projects have received DOE/NNSA HQ's fully signed and executed Coordination and Approval Document (CAP)
  - 4 contracts are fully signed
  - 1 is being routed for signatures by the LLNL Lab Director - Goldstein and his counterparts at the Vendors to fully execute these contracts – anticipated by COB April 14, 2017
- The last PathForward contract is at DOE/NNSA for approval
- The HT leadership team is developing a tiered communication plan to:
  - Introduce the PathForward projects to the rest of ECP, and
  - Establish high-quality co-design collaborations with interested lab-led ECP projects

EXASCALE COMPUTING PROJECT

Design Space Evaluation: Technology Coverage Areas

	DSE (lab POCs)	Memory Technologies	Interconnect/System Simulators	Analytical Models	Node Simulators	Abstract Machine Models and Proxy Architectures
ANL	Ray Bair		ROSS/CODES		Gem5	
LANL	Jeff Kuehn		CoNCEPTuaL	Byfl		
LBNL	David Donofrio			ExaSAT, Roofline Toolkit	Gem5, OpenSOC	Co-lead AMM v3
LLNL	Robin Goldstone	LiME	ROSS/CODES			
ORNL	Jeff Vetter	Blackcomb		ASPEN		
PNNL	Darren Kerbyson			PALM		
SNL	Rob Hoekstra	SST: VaultSimC, MemHierarchy	SST, Merlin, Ember, Firefly		SST: Miranda, Ariel, Gem5	Co-lead AMM v3

ECP teams begin work on Office of Science systems via early access, ALCC awards, and testbeds

Access to Office of Science Systems	Access to testbeds supported with ECP funds

NERSC's Cori II

Cori II early access
- Shared access with other users
- Intel Xeon Phi (KNL) nodes
- Through June 30, 2017

OLCF's Titan

2016 ALCC Reserve Award to ECP
- 13M Titan-core hours
- March 30 to June 30, 2017
- 10 AD, 6 ST, and 1 HT teams are getting access

ALCF's Theta

ECP allocation is 37+M hours on Theta for 2017
- Intel Xeon Phi cpus
- 3,624 nodes
- 64 core processor per node

OLCF's Summitdev

ECP allocation is the equivalent of 18 nodes on Summitdev
- Power 8+
- NVIDIA Pascal GPGPUs
- NVLINK1

ECP submitted an application to the 2017 ASCR Leadership Computing Challenge (ALCC) for time of Titan, Mira, Theta, Cori, and Edison

Communication and Outreach

- We have a website
  - **www.ExascaleProject.org**
- First Annual Meeting January 31 – February 2, 2017
  - 450+ participants
  - Focused on co-design and integration planning
- A newsletter *The ECP Update* has been launched
  - https://exascaleproject.org/newsletter/ecp-update-1/
  - will be published at a minimum on a monthly basis, but we will use this vehicle to send our more frequent updates as we have topics of interest to report on.
  - Mailed to website followers plus Confluence registrants.

ECP Industry Council Charter

- The Exascale Computing Project (ECP) Industry Council is an **advisory body** that serves the **ECP director and leadership team** by providing advice and feedback on project scope and requirements, technical approaches, progress, strategic directions and other issues.
- The ECP Industry Council provides two-way communication and information exchange between the ECP and the high performance computing (HPC) industrial user community as well as the commercial HPC software community.

Industry Council Organization

- ECP IC Chair – Michael McQuade, United Technologies
- ECP Director – Paul Messina, Argonne National Laboratory
- ECP IC Co-Executive Directors – Suzy Tichenor, Oak Ridge National Laboratory and David Martin, Argonne National Laboratory
- ECP IC Members – 18 senior executives, 15 from industry and 3 from ISVs
- Members are appointed by ECP Director for a two-year term
- Members represent their company as well as their industry sector

EXASCALE COMPUTING PROJECT

Industries Represented (Forbes Categories)

- Conglomerates
- Oil & Gas Operations
- Nuclear Power
- Heavy Equipment
- Air Currier
- Household Personal Care
- Pharmaceuticals
- Food, Drink and Tobacco
- Auto & Truck Manufacturers
- Household Appliance
- Diversified Chemicals
- Fusion Energy
- Media (Animation)
- Independent Software Vendors

Wall of Logos

United Technologies

Westinghouse

Lilly

GM

Cummins

ExxonMobil

CASCADE TECHNOLOGIES

FedEx

GE

ANSYS

P&G

MARS

Chevron

Altair

TRI ALPHA ENERGY
THE POWER OF INGENUITY

Dow

Whirlpool CORPORATION

DREAMWORKS

First ECP Industry Council Meeting

- Held March 6-7, 2017 at Argonne National Laboratory
- Tour of the Advanced Photon Source
- Detailed Presentations from ECP Director and Focus Area Leads
- Presentation on International Exascale Competition
- Time for Feedback on Direction and Plans
- Outcome: Exascale is Critical to US, Desire to Provide Detailed Input on Industry Needs

ECP EXASCALE COMPUTING PROJECT

For More Information

- Industry Council Web Site: https://exascaleproject.org/advisory-councils/
- Industry Council Membership List: https://exascaleproject.org/wp-content/uploads/2017/02/ECP-Industry-Council-final-2-23-17-3.pdf
- Industry Council Co-Executive Directors:
  - David Martin <dem@alcf.anl.gov>
  - Suzy Tichenor <tichenorsp@ornl.gov>

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

Thank you!

www.ExascaleProject.org

EXASCALE COMPUTING PROJECT

www.ExascaleProject.org