ECP Update

EXASCALE COMPUTING PROJECT

Douglas B. Kothe (ORNL), ECP Director
Lori A. Diachin (LLNL), ECP Deputy Director

ASCAC Meeting
September 24, 2019

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

ECP BLUF (Bottom Line Up Front)

- ***Technical Highlights***
  - **Hardware and Integration** (HI): PathForward element paying dividends; turning focus to deploying ECP's E4S (Extreme-Scale Scientific Software Stack); tuning applications to exascale systems
  - **Software Technology** (ST): Deploying E4S now; defining ST product integration metrics; increasing focus on software abstraction layers, hardware-driven algorithms for math libs (mixed precision), programming models
  - **Application Development** (AD): Get skin in the game (quantitative criteria for challenge problems); refine plans for exploitation of accelerators; what are the performance bottlenecks to delivering on the challenge problems?
  - ECP's **AI/ML scope** is cutting-edge & impactful (CANDLE, ExaLearn) - **will expand as risks are retired**
- ***External engagements across international borders (UK, Japan) and with other US Gov't agencies (NOAA, NSF, NASA, DoD)***
- ***DOE Facility Engagement remains very active and is retiring many unknown unknown risks***
  - First-mover exascale system (Aurora, Frontier, El Capitan) **schedules and technology targets set**
  - Formal relationships and **shared-milestone plans with DOE HPC Facilities** defined for mutual success
  - Known unknowns remain, e.g., those related to robust, portable, and performant accelerator programming model
- ***Recent and upcoming reviews***
  - Recent (Jun 2019) external review of ECP's "**Final Design**" reaffirmed that ECP is on track
  - CD 2/3 preparations are well underway; extensive revision of project documentation; baseline technical scope

Technical Highlights

ECP's three technical areas have the necessary components to meet national goals

Performant mission and science applications @ scale

Aggressive RD&D Project	Mission apps & integrated S/W stack	Deployment to DOE HPC Facilities	Hardware tech advances

Application Development (AD)

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

24 applications including national security, to energy, earth systems, economic security, materials, and data

Software Technology (ST)

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

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

Hardware and Integration (HI)

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

6 US HPC vendors focused on exascale node and system design; application integration and software deployment to facilities

ECP **Application Development (AD)**

Goal

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

- High impact science or engineering exascale *challenge problem*
- Detailed criteria for assessing successful completion of challenge problem
- A Figure of Merit (FOM) formula quantifying performance of challenge problem
- Demonstration and assessment of effective software integration

Chemistry and Materials

Earth and Space Science

Energy

Data Analytics and Optimization

National Security

Co-Design

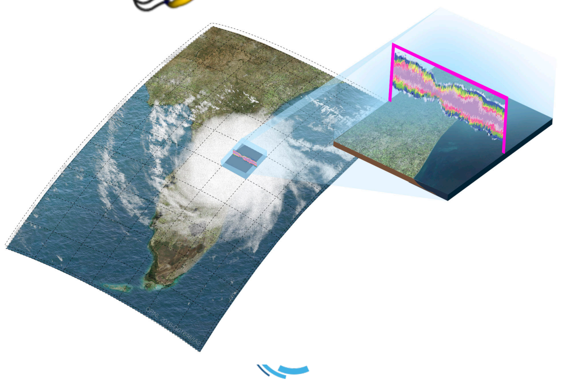

Exascale Computing Project: Application Development

Goal: Ensure that exascale hardware impacts DOE science/engineering mission

Approach: Significant investment in scientific applications well in advance of exascale machines

- ~~Code Porting~~
- Algorithmic Restructuring
- New Numerical Approaches
- Alternate choice of Physical Models

Hardware has significant impact on all aspects of simulation strategy

Portfolio of ECP Applications

Application Categories	Projects
Chemistry and Materials	6
Energy (generation)	5
Earth and Space Sciences	5
Data Analytics and Optimization	4
National Security	4

24 Domain Science/Engineering Simulation Projects

50+ separate codes

2/3 C/C++ ; 1/3 Fortran

Most pure MPI, or MPI+OpenMP at outset

Well defined, evolving dependencies on ECP software technology projects

ECP Apps: Delivering on Challenge Problems

Requires Overcoming Computational Hurdles

Domain	Challenge Problem	Computational Hurdles
Wind Energy	Optimize 50-100 turbine wind farms	Linear solvers; structured / unstructured overset meshes
Nuclear Energy	Virtualize small & micro reactors	Coupled CFD + Monte Carlo neutronics; MC on GPUs
Fossil Energy	Burn fossil fuels cleanly with CLRs	AMR + EB + DEM + multiphase incompressible CFD
Combustion	Reactivity controlled compression ignition	AMR + EB + CFD + LES/DNS + reactive chemistry
Accelerator Design	TeV-class 100-1000X cheaper & smaller	AMR on Maxwell's equations + FFT linear solvers + PIC
Magnetic Fusion	Coupled gyrokinetics for ITER in H-mode	Coupled continuum delta-F + stochastic full-F gyrokinetics
Nuclear Physics: Lattice QCD	Use correct light quark masses for first principle light nuclei properties	Critical slowing down; strong scaling performance of MG-preconditioned Krylov solvers
Chemistry	Heterogeneous catalysis: MSN reactions	HF + DFT + coupled cluster (CC) + fragmentation methods
Chemistry	Catalytic conversion of biomass	Hybrid DFT + CC; CC energy gradients
Extreme Materials	Microstructure evolution in nuclear matls	AMD via replica dynamics; OTF quantum-based potentials
Additive Manufacturing	Born-qualified 3D printed metal alloys	Coupled micro + meso + continuum; linear solvers
Quantum Materials	Predict & control matls @ quantum level	Parallel on-node performance of Markov-chain Monte Carlo
Astrophysics	Supernovae explosions & neutron star mergers	AMR + nucleosynthesis + GR + neutrino transport

ECP Apps: Delivering on Challenge Problems

Requires Overcoming Computational Hurdles

Domain	Challenge Problem	Computational Hurdles
Cosmology	Extract "dark sector" physics from upcoming cosmological surveys	AMR or particles (PIC & SPH); subgrid model accuracy; insitu data analytics
Earthquakes	Regional hazard and risk assessment	Seismic wave propagation coupled to structural mechanics
Geoscience	Geomechanical and geochemical evolution of a wellbore system at near-reservoir scale	Coupled AMR flow + transport + reactions to Lagrangian mechanics and fracture
Earth System	Assess regional impacts of climate change on the water cycle @ 5 SYPD	Viability of Multiscale Modeling Framework (MMF) approach for cloud-resolving model; GPU port of radiation and ocean
Power Grid	Efficient planning; underfrequency response	Parallel performance of nonlinear optimization based on discrete algebraic equations and MIP
Cancer Research	Predictive preclinical models and accelerate diagnostic and targeted therapy	Increasing accelerator utilization for model search; exploiting reduced/mixed precision; preparing for any data management or communication bottlenecks
Metagenomics	Discover, understand (find genes) and control species in microbial communities	Efficient and performant implementation of UPC, UPC++, GASNet; graph algorithms; SpGEMM performance
FEL Light Source	Light source-enabled analysis of protein and molecular structure and design	Strong scaling (one event processed over many cores) of compute-intensive algorithms (ray tracing, M-TIP) on accelerators

Applications Face Common Challenges

1) **Flat performance profiles**
2) **Strong Scaling**
3) **Understanding/analyzing accelerator performance**
4) **Choice of programming model**
5) **Selecting mathematical models that fit architecture**
6) **Software dependencies**

Despite these challenges, recent (Sep 2019) external SME reviews indicate that almost all application projects are on or ahead of schedule. Corrective actions are being implemented for only a few projects.

ECP's Co-design Centers Target Key Motifs

Address computational motifs common to multiple application projects

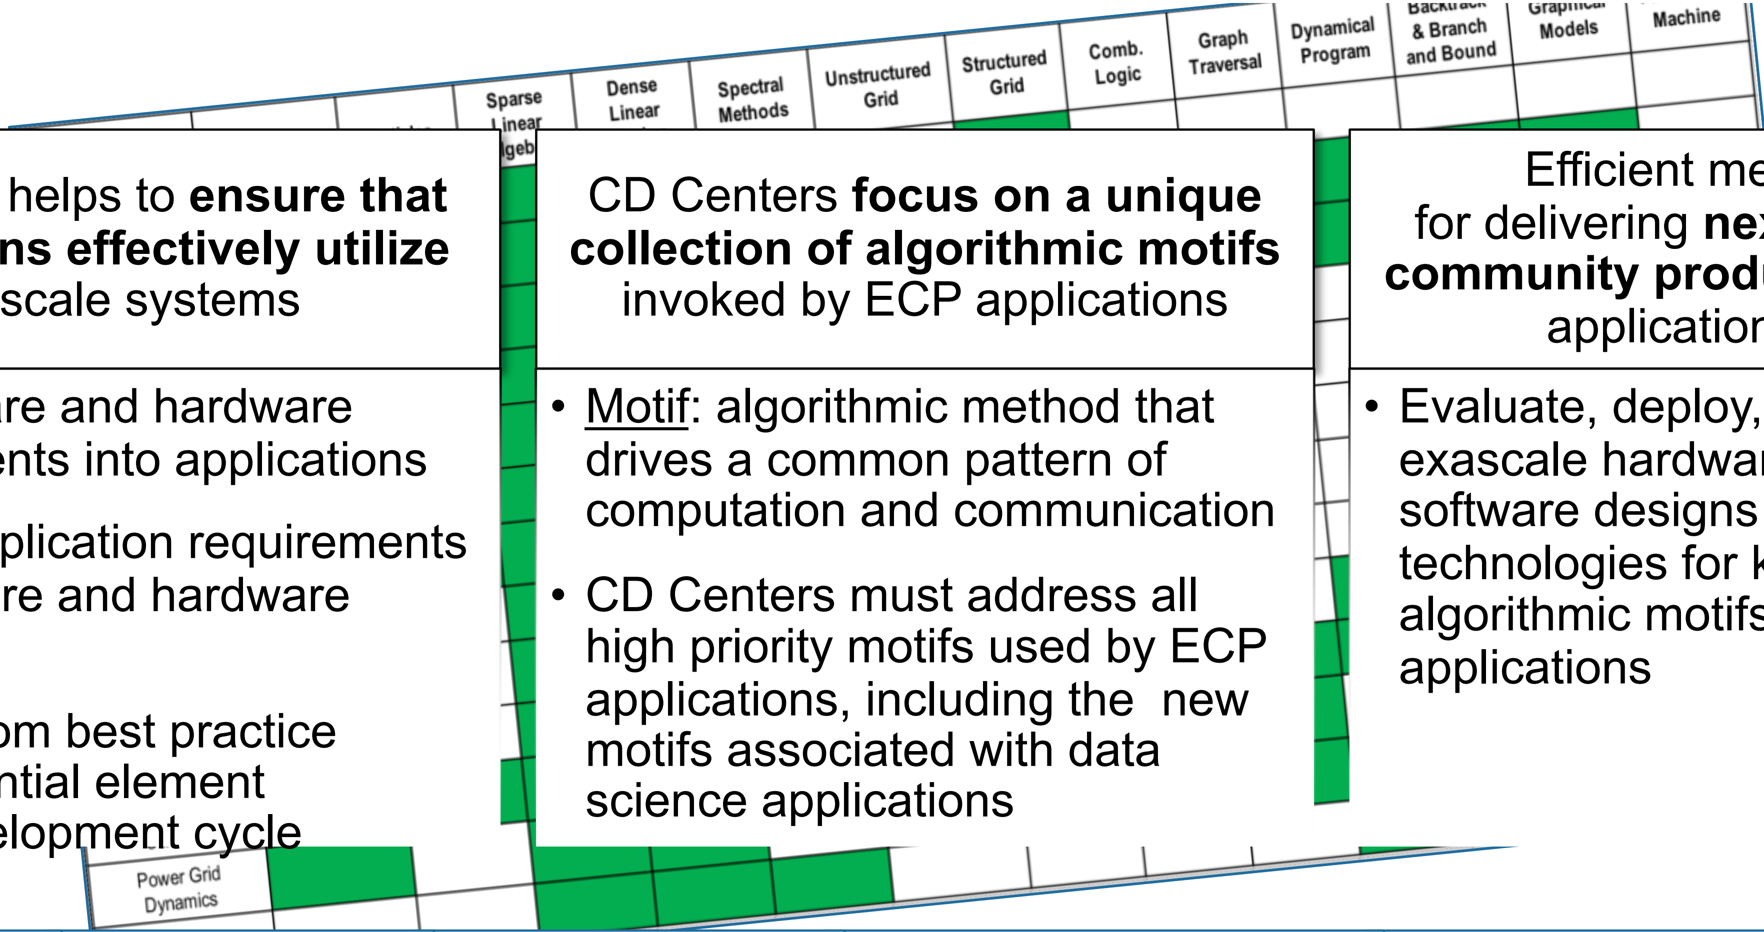

Co-design helps to **ensure that applications effectively utilize** exascale systems

- Pull 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 used by ECP applications, including the new motifs associated with data science applications

Efficient mechanism for delivering **next-generation community products** with broad application impact

- Evaluate, deploy, and integrate exascale hardware-aware software designs and technologies for key crosscutting algorithmic motifs into applications

CODAR *Data and workflows*	**COPA** *Particles/mesh methods*	**AMReX** *Block structured AMR*	**CEED** *Finite element discretization*	**ExaGraph** *Graph-based algorithms*	**ExaLearn** *Machine Learning*

ECP's Co-Design Centers: Impacting Multiple Applications

Co-design	Application
CEED	ExaSMR, LLNL NNSA App, *ExaAM, ExaWind, Combustion-PELE, Subsurface, E3SM, SNL NNSA App*
AMReX	ExaAM, Combustion-PELE, MFIX-Exa, WarpX, ExaStar, ExaSky
CoPA	EXAALT, ExaAM, WDMApp, *MFIX-Exa, WarpX, ExaSky, AMReX*
CODAR	WDMApp, CANDLE, NWChemEx, *EXAALT, Combustion-PELE, ExaSky*
ExaGraph	ExaWind, ExaBiome, *ExaSGD, SNL NNSA App*
ExaLearn	ExaSky, *CANDLE, NWChemEx, ExaAM*

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

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.

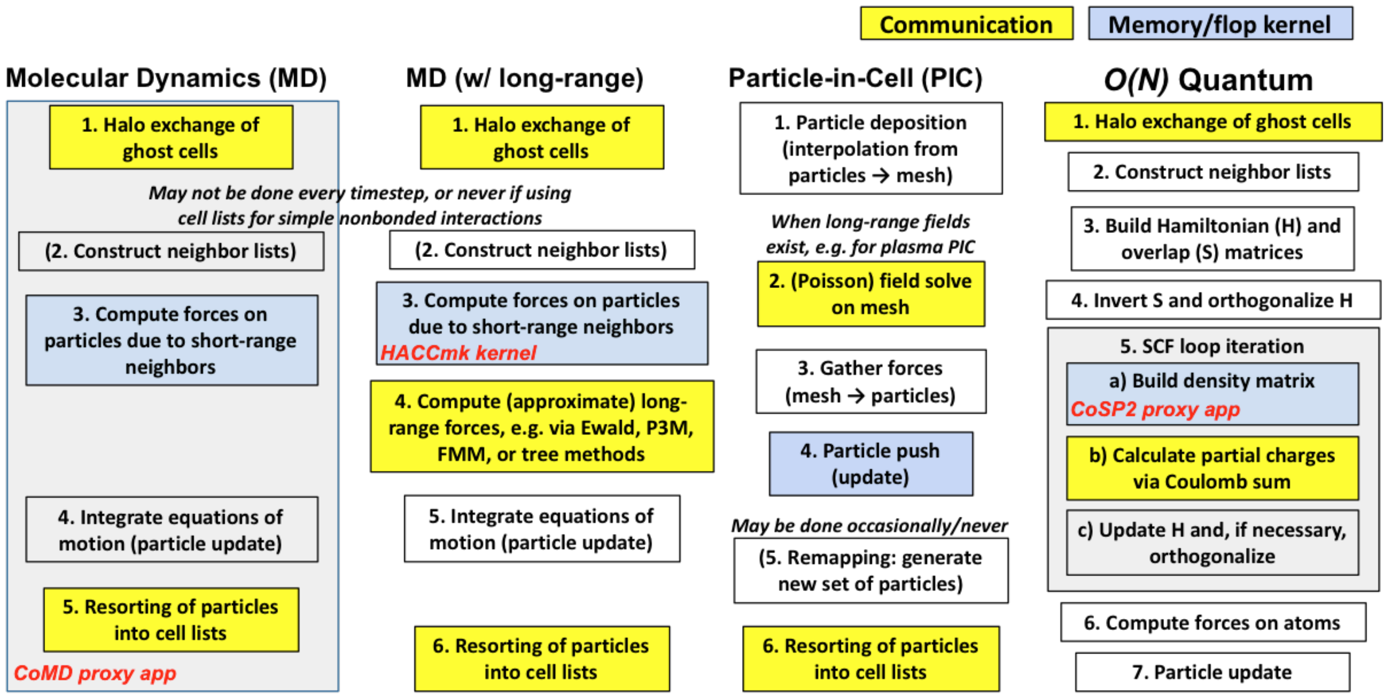

Technical approach:

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

CoPA Cabana: Co-Designed Numerical Recipes for Particles

Cabana:

- is a software library for developing exascale applications that use particle algorithms
- contains general particle data structures and algorithms implemented with those data structures
- provides a platform to develop and deploy advanced scalable and portable methods for particle-based physics algorithms
- is designed for modern DOE HPC architectures and **builds directly on Kokkos**
- is open source and distributed on GitHub

Core ECP stakeholders include projects with codes for molecular dynamics (MD), N-body and smoothed particle hydrodynamics (SPH), and various particle-in-cell (PIC) derivatives.

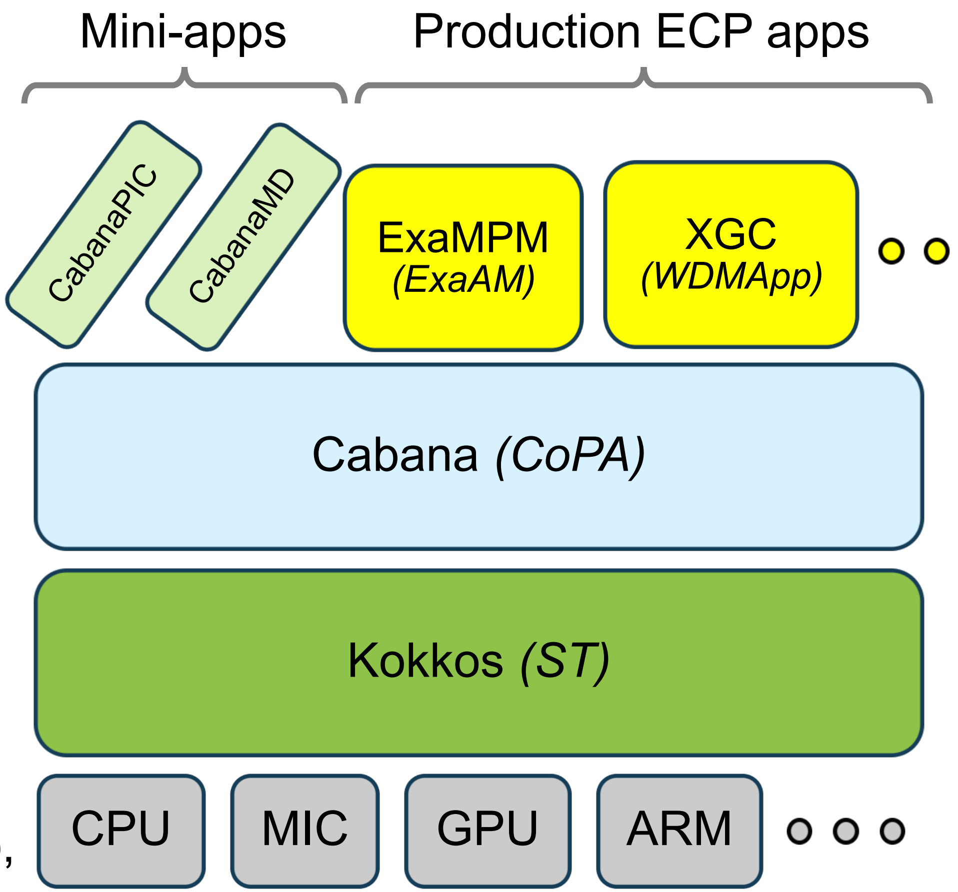

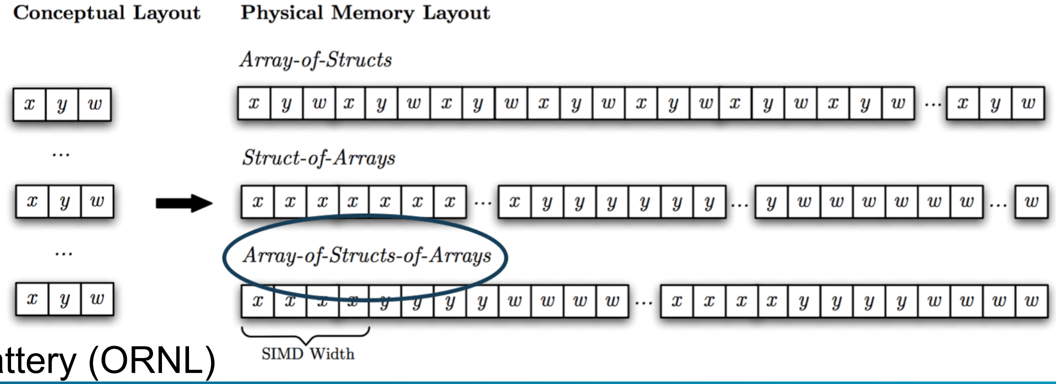

ECP's Co-Designed Motif Approach is Working

ECP's CoPA is ensuring portable performance of the XGC fusion application

- XGC utilizes Cabana/Kokkos for portable platform performance: Summit, Perlmutter, Aurora, Frontier
- Fortran interface has been developed for XGC via Cabana particle library in ECP-CoPA
- Similar performance on Summit and KNL

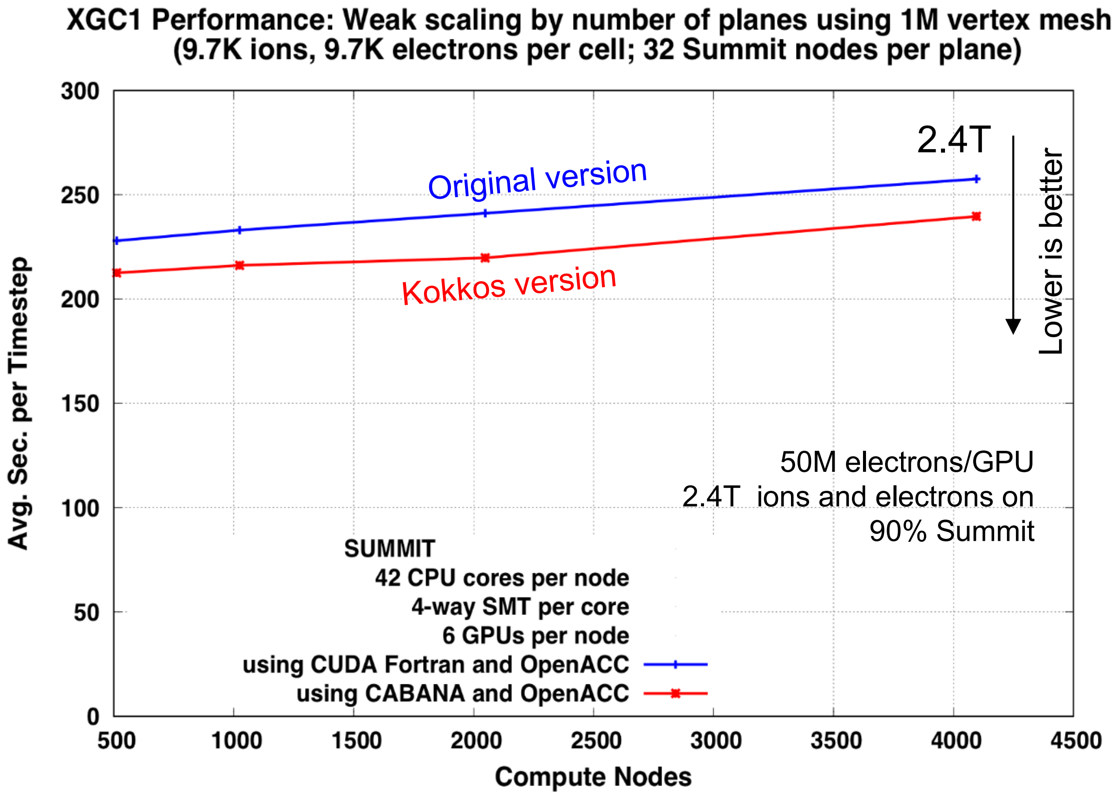

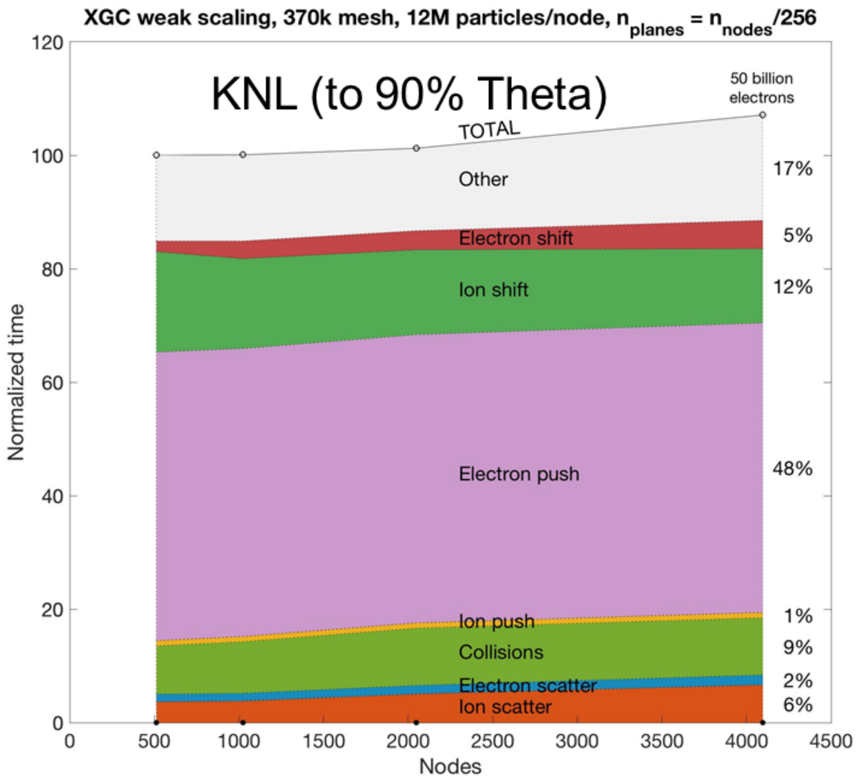

ExaLearn: Machine Learning for Inverse Problems in Materials

- **Impact**
  - Enable learning from vast troves of materials-related empirical and simulation data available within the DOE complex, as well as in various repositories around the world ***quickly*** and ***more accurately***.
  - Core ML technologies will be built based on ***domain-agnostic design principles***, but deployment will target domain-specific benchmarks in *neutron scattering* and *X-ray crystallography*.
  - Of particular interest are studies of an important class of materials called ***perovskites*** (see chart).

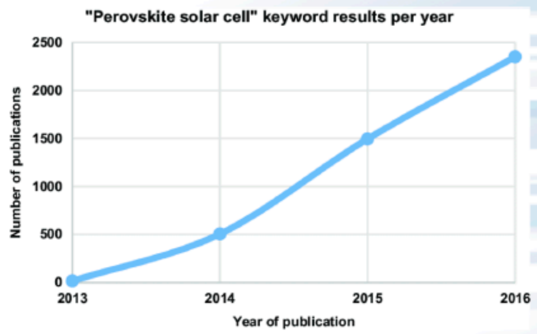

ExaLearn Pipeline for Material Structure Determination from Neutron Scattering/Diffraction Data

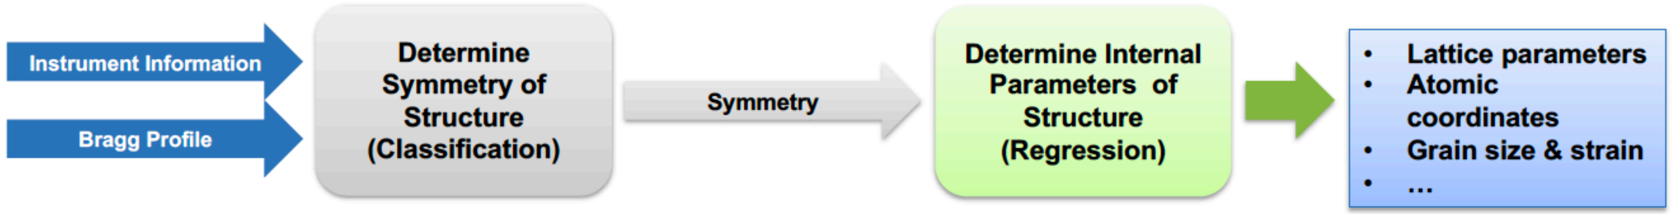

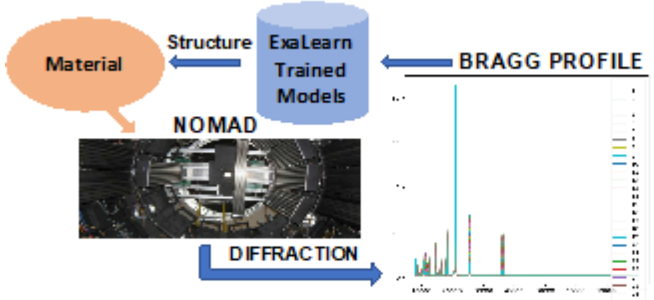

Figure 1: The learning framework ingests experimentally acquired neutron diffraction patterns of target samples and predicts their structural parameters, such as, inter-ionic spacings, bond angles and other thermal parameters.

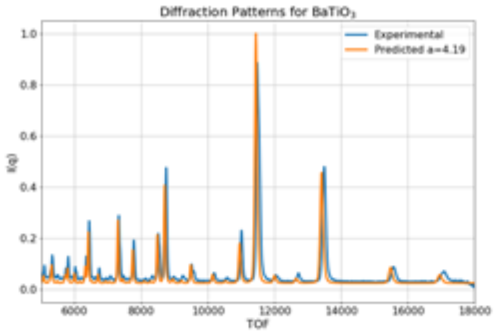

Figure 2: Experimental diffraction profile (blue) compared with a diffraction profile simulated using the structural parameters predicted by the learning framework (orange).

Machine Learning in the Light Source Workflow

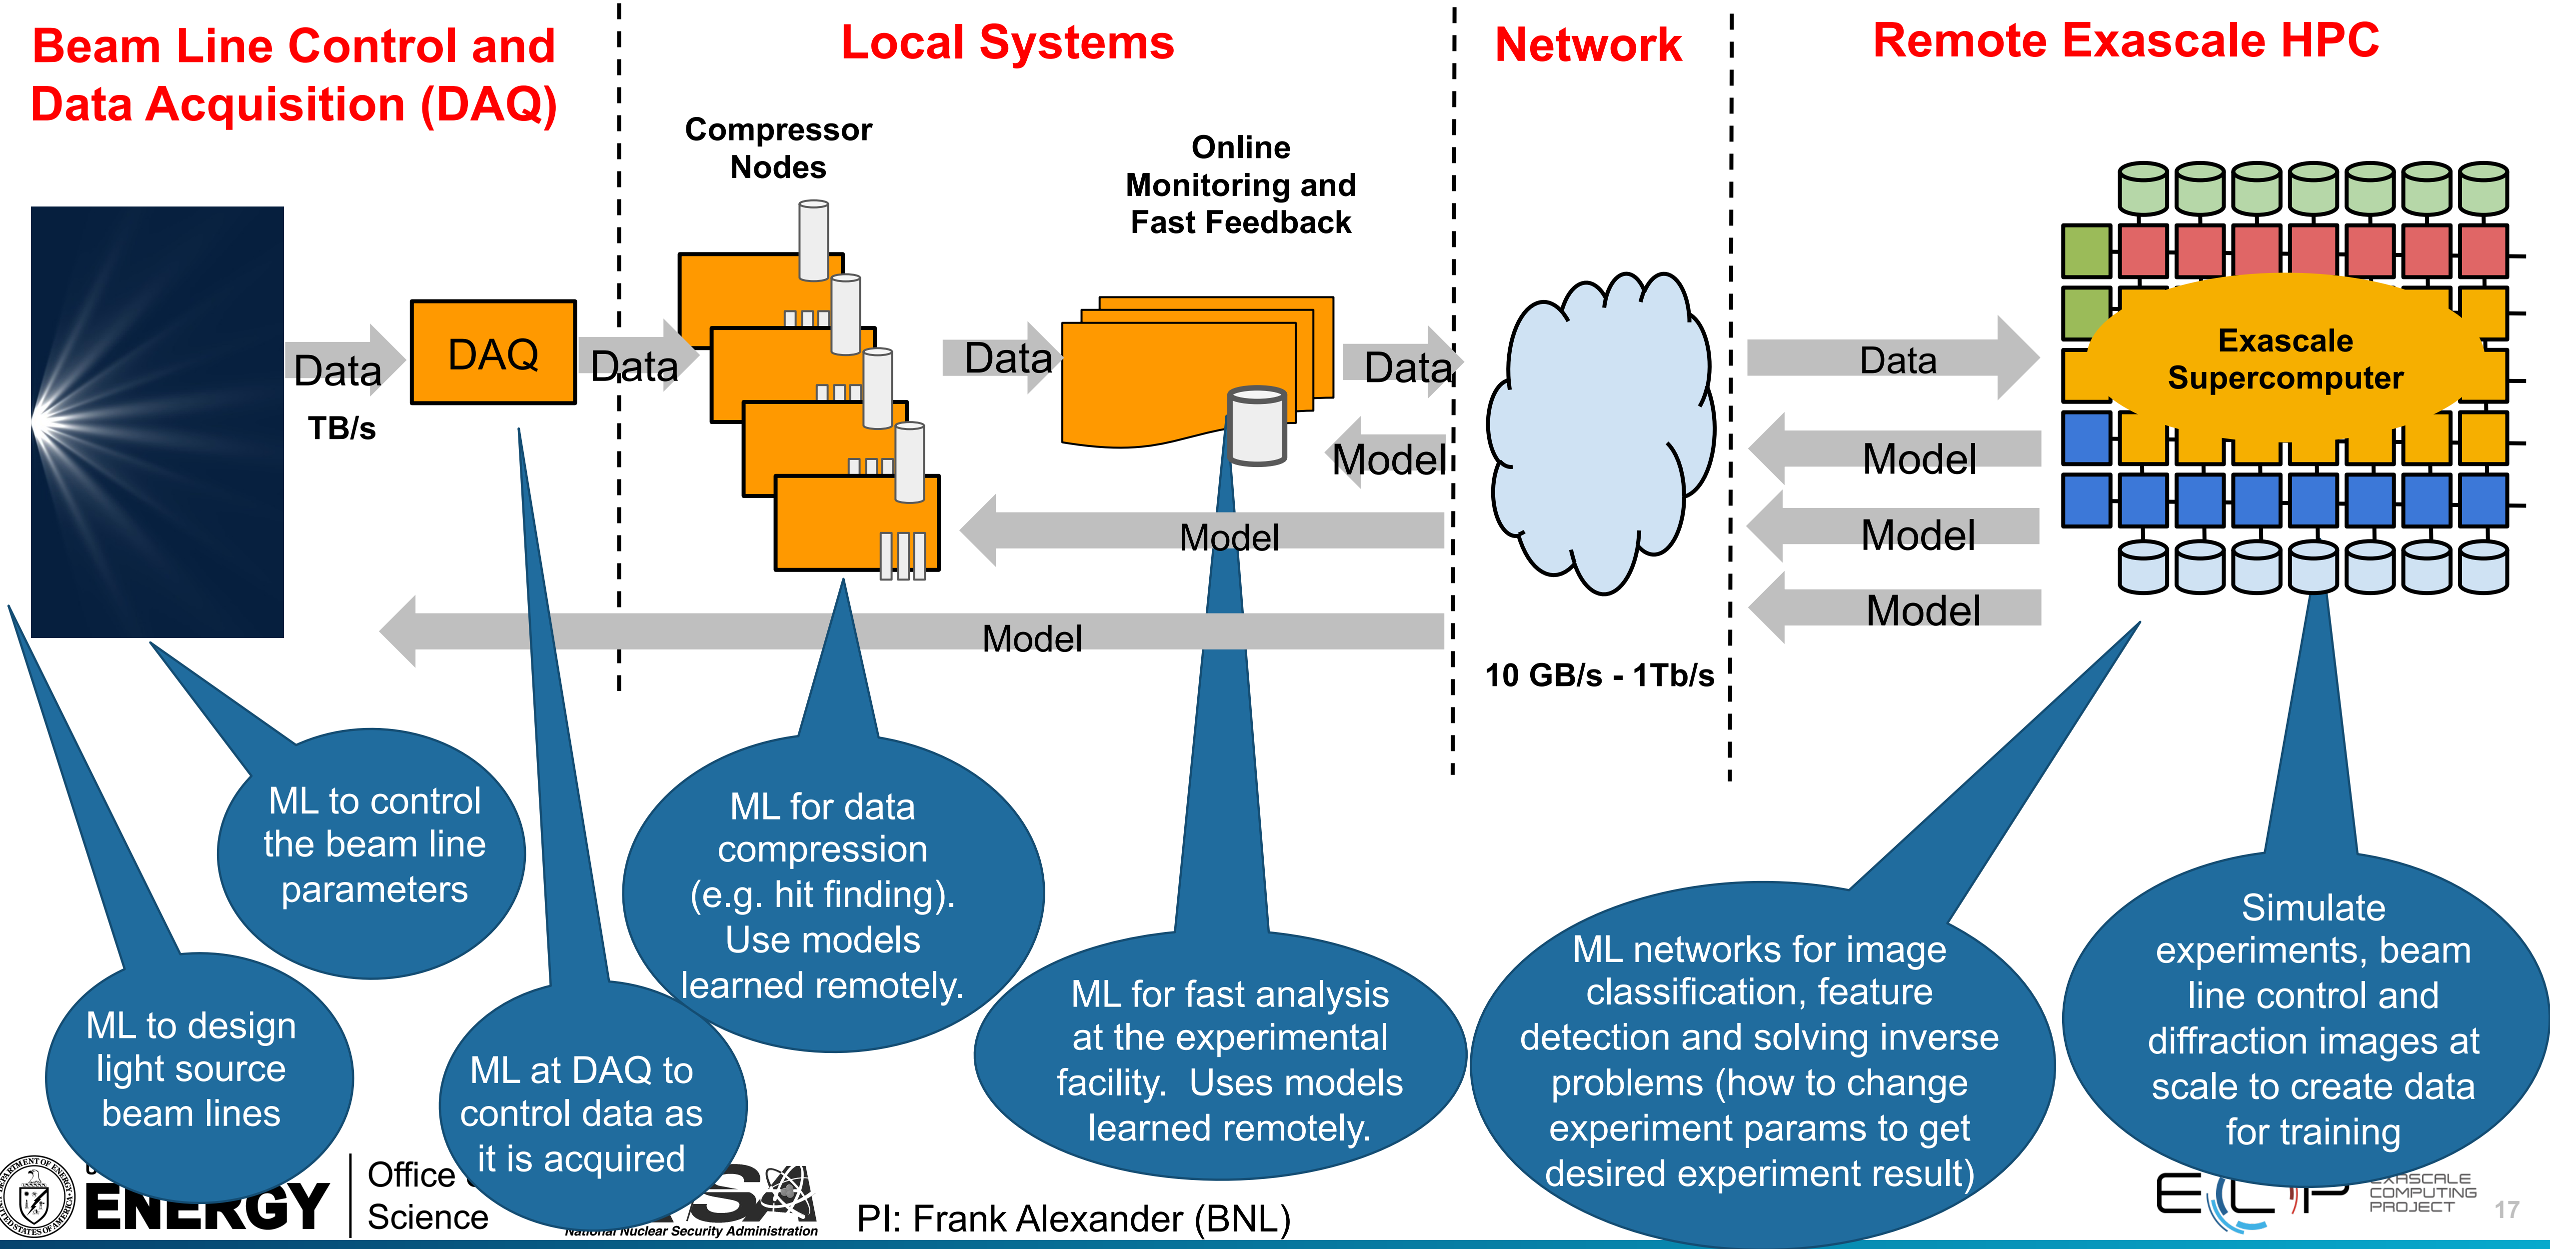

PI: Frank Alexander (BNL)

ECP **Software Technology (ST)**

Goal

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

- Prepare SW stack for scalability with massive on-node parallelism
- Extend existing capabilities when possible, develop new when not
- Guide, and complement, and integrate with vendor efforts
- Develop and deliver high-quality and robust software products

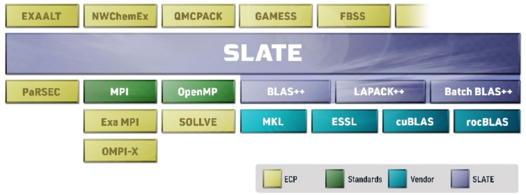

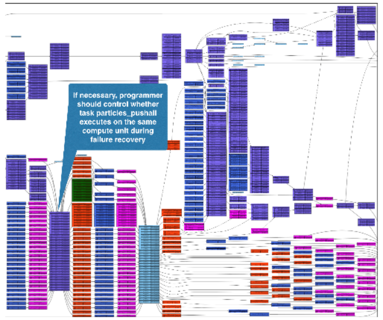

ECP ST Software Ecosystem

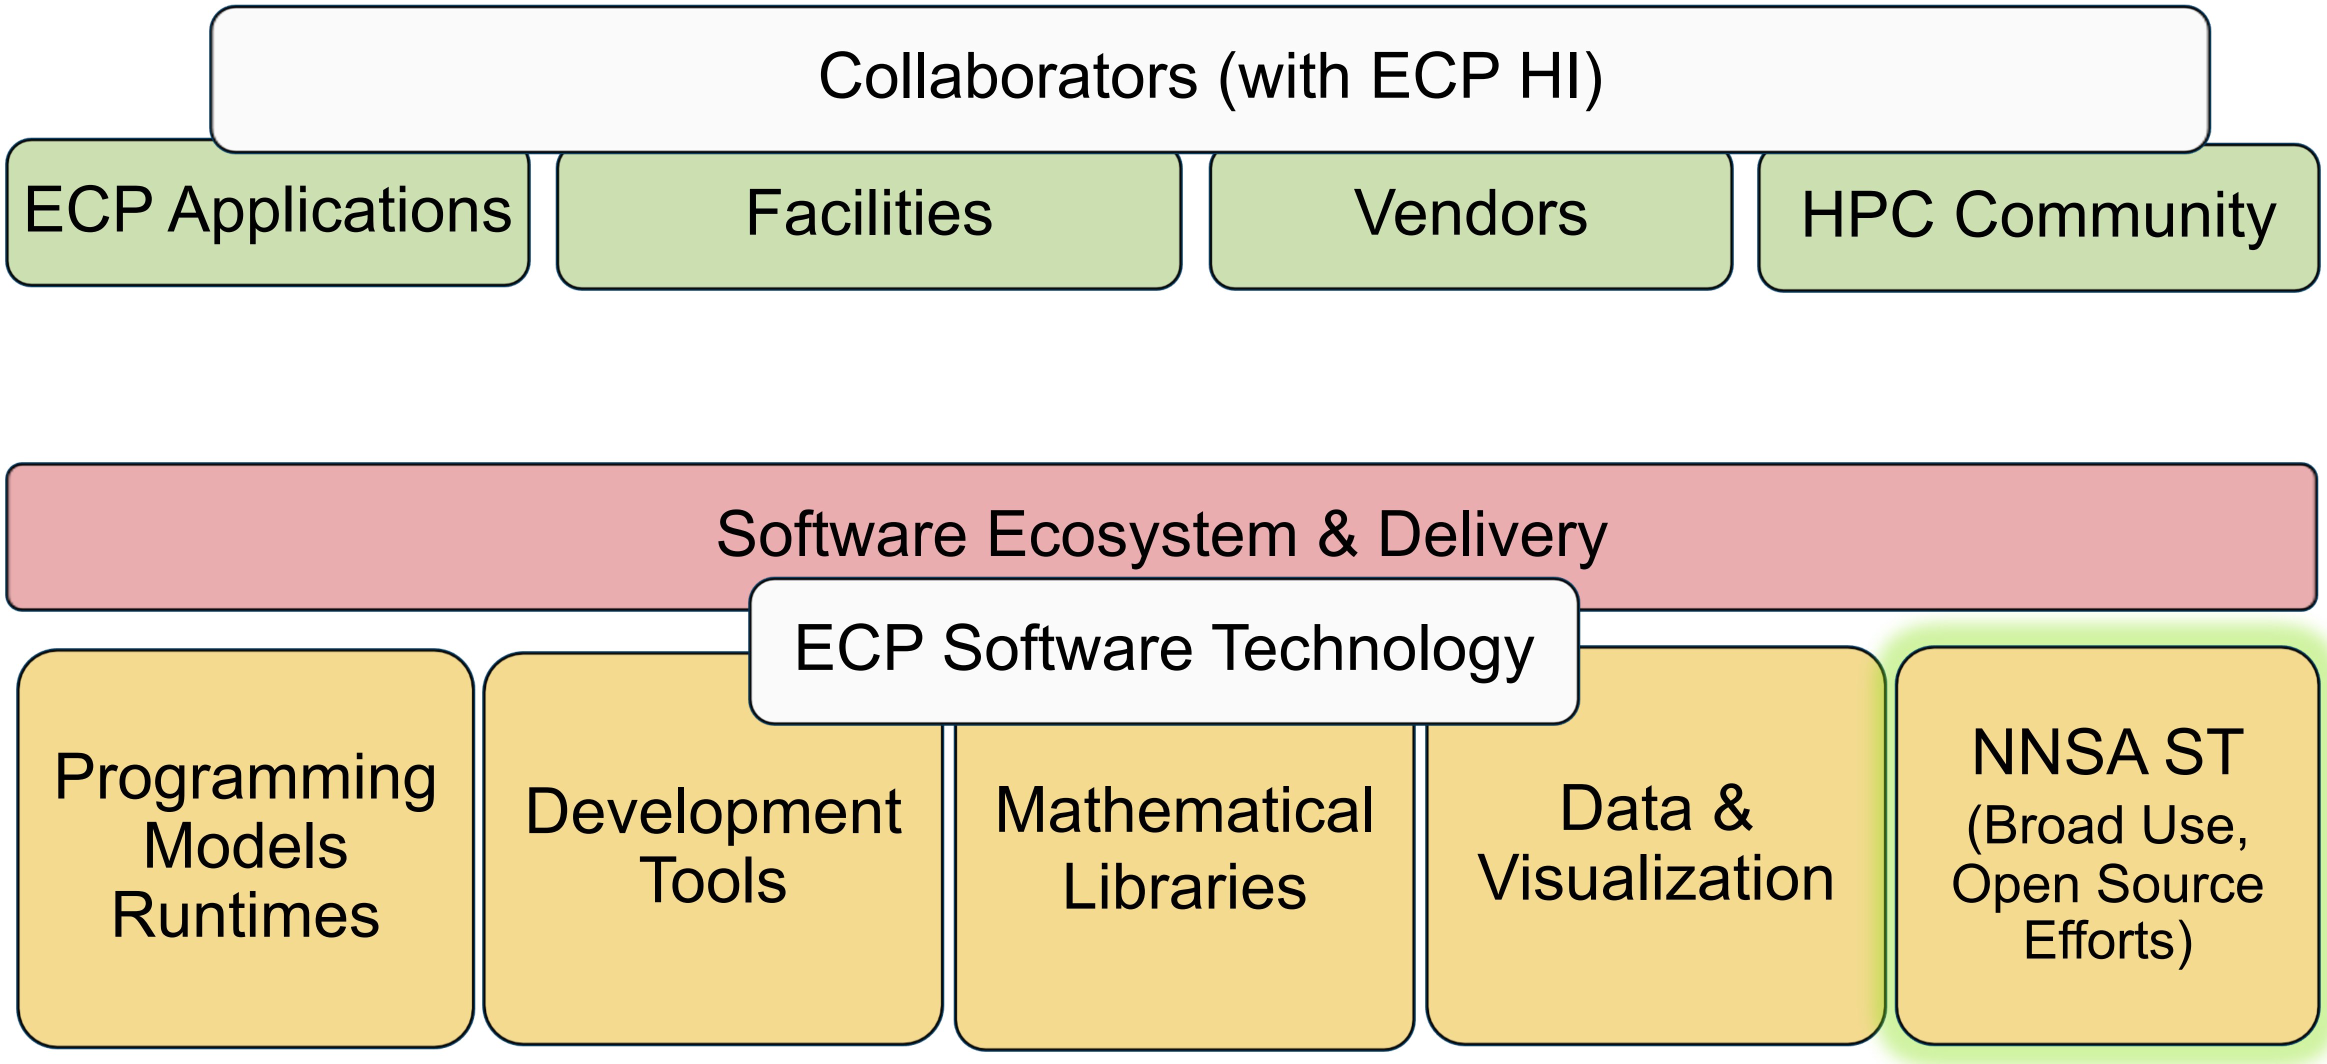

ECP software technologies overview

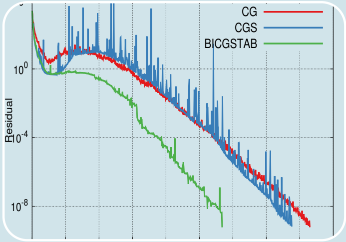

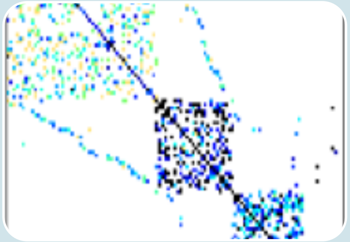

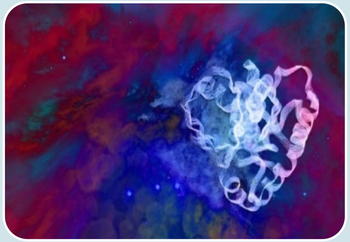

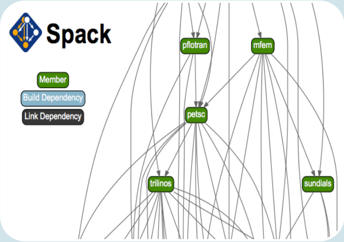

National Nuclear Security Administration

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,
- 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
- Exploit reduced and mixed precision hardware operations

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

NNSA ST

- Open source NNSA Software projects
- Projects that have both mission role and open science role
- Major technical areas: New programming abstractions, math libraries, data and viz libraries
- Cover most ST technology areas
- Subject to the same planning, reporting and review processes

Software Technology Ecosystem

Levels of Integration | **Product** | **Source and Delivery**

ECP ST Individual Products

- Standard workflow
- Existed before ECP

ST Products

Source: ECP L4 teams; Non-ECP Developers; Standards Groups
Delivery: Apps directly; spack; vendor stack; facility stack

ECP ST Open Product Integration Architecture

- Group similar products
- Make interoperable
- Assure policy compliant
- Include external products

SDKs

Source: ECP SDK teams; Non-ECP Products (policy compliant, spackified)
Delivery: Apps directly; spack install sdk; future: vendor/facility

- Build all SDKs
- Build complete stack
- Containerize binaries

E4S

Source: ECP E4S team; Non-ECP Products (all dependencies)
Delivery: spack install e4s; containers; CI Testing

There are many examples of AD/ST/CD interactions resulting in significant progress

QMCPACK / SOLLVE+Kokkos

- Portable programming required for CPU, GPU and other accelerated systems
- High priority kernels ported to both OpenMP and Kokkos and their preliminary performance has be assessed.
- OpenMP GPU branch provides current best FOM on Summit

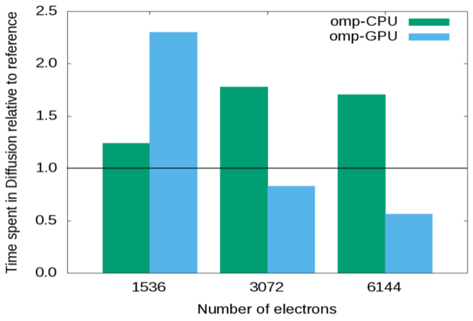

WDMApp / ADIOS

- Provide high-performance IO and coupling framework for Gyrokinetic codes
- ADIOS is the backbone of the KITTIE framework that allows coupling between GENE (core) and XGC (edge) gyrokinetic codes
- Will allow incorporation of community-fusion models into WDMApp for exascale whole device modeling

MFIX-Exa / AMReX

- Provide adaptive mesh refinement, field solvers for implicit projection Navier-Stokes, and embedded boundaries for non-orthogonal geometries
- Embedded boundaries and implicit projection solver integrated and tested; performance optimization and scaling ongoing
- Allow performance-portable DEM mechanics on combinatorial geometry reactor models

ST research teams are advancing the state-of-the-art in preparation for exascale computing

GASnet-EX RMA matches or exceeds MPI RMA

- Three different MPI implementations; Two distinct network hardware types
- On four systems the performance of GASNet-EX matches or exceeds that of MPI RMA :
  - 8-byte Put latency 6% to 55% better
  - 8-byte Get latency 5% to 45% better
  - Better flood bandwidth efficiency, typically saturating at ½ or ¼ the transfer size

8-Byte RMA Operation Latency (one-at-a-time)

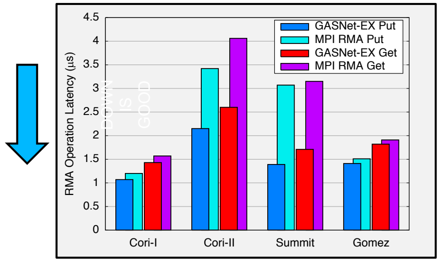

Added PAPI support for TESLA V100 GPUs and NVLINK

- Production interface of PAPI to CUPTI and nVidia Management Library (NVML)
- Demonstrated PAPI NVIDIA GPU power reading and control, and the use of performance counters across multiple GPUs
- Allows developers to change run profiles to reduce energy cost (NVML)
- Aids developers in producing more efficient code by profiling the utilization of the latest GPU resources and diagnosing performance bottlenecks.

Easing ECP software stack deployment via containers

- Defn: OVA file is a virtual appliance used by virtualization applications
- Provides a platform for easy deployment and use of HPC container runtimes and access to the Extreme Scale Scientific Software Stack container
- Created an OVA file that includes Docker, Singularity, Shifter, and Charliecloud runtimes that contains Spack based packages as well as the Singularity E4S image

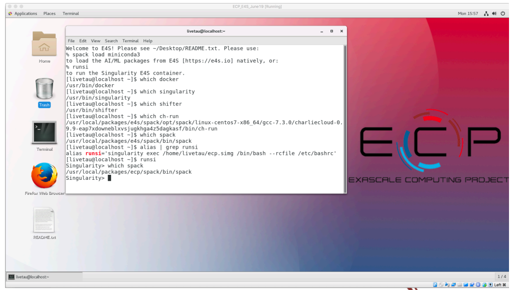

ECP **Hardware and Integration (HI)**

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

Access to the computer resources at facilities: early access, test and dev. systems, and pre-exascale and exascale systems

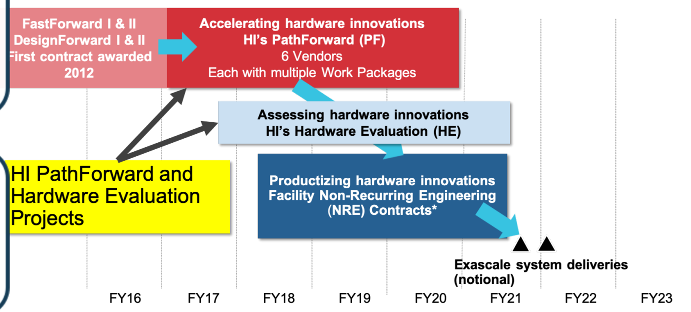

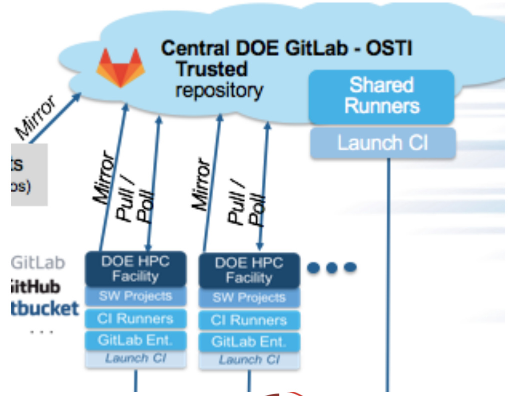

HI six L3 technical projects and their scope

PathForward (PF)	Critical early vendor HW R&D for multiple exascale-capable system designs
Hardware Evaluation (HE)	HW evaluations to influence system designs and to inform Facilities, AD, and ST
Application Integration (AI)	Facility support for ECP application development efforts to port and optimize for exascale or pre-exascale systems
Software Deployment (SD)	Facility support for deploying ECP SW at the Facilities and integrating with each Facility's exascale SW ecosystem
Facility Resource Utilization (FRU)	Access to compute resources made available to ECP through the Facilities
Training and Productivity (T&P)	Disseminated development knowledge, lessons learned, and best practices to AD and ST teams in collaboration with AD, ST, and the Facilities

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: 21 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 12 ECP applications 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

Final Design Review

ECP: The Road Ahead

FY19

- **Final Design Review**
- Establish performance baseline
- AD KPP completion criteria and ST integration goals set
- *Access to pre-exascale systems*

FY20

- **CD-2/3 Review and Approval**
- Did PathForward deliver? Are AD and ST performance and integration projections on track?
- *Access to Aurora and Frontier early hardware*

FY21

- **Status Independent Project Review (IPR)**
- AD application projections firm for target system
- ST integration goals assessed
- *Access to Aurora and Frontier Test and Development Systems (TDS)*

FY22

- **Status Independent Project Review (IPR)**
- AD and ST readiness demonstrated
- *Access to earliest Aurora and Frontier full system*

FY23

- **CD-4 Review and Approval (project completion)**
- Deliver KPP completion evidence
- *Access to Aurora and Frontier full system*

The ECP Final Design Review was held June 25-26, 2019 at Argonne National Laboratory

Charge Questions

- Is the ECP project plan and structure adequate to deliver exascale-capable applications and software that meet the KPPs?
- Is the final design sufficiently detailed and mature to establish a reliable baseline cost and schedule for the project?
- Is the ECP adequately managing complex interdependencies across the project activities?
- Is the ECP satisfactorily managing the engagements and synergistic activities with the DOE Facilities that are critical to project success?
- Have the appropriate technical risks and mitigation strategies been identified and addressed?

Review Panel

- Dan Stanzione, TACC
- Mike Norman, UCSD
- Gianluca Iaccarino, Stanford Ed Seidel, UIUC
- Bill Carlson, IDA
- Sadaf Alam, CSCS
- Fred Johnson, Retired, DOE
- Edmond Chow, GA Tech
- Chip Watson, TJNAF
- Keith Obenschain, NRL Christine Cuicchi, DoD

ECP's final design consists of three primary components

Project Structure

- Three technical focus areas teamed with project management expertise
- Hierarchical break down of work scope with strong technical leadership at each level
- Key Performance Parameters (KPPs) to measure success in meeting project objectives
- Critical dependencies
  - Integration within the project
  - Integration with DOE Facilities

Technical Plans

- Detailed definition of KPPs for each project with verifiable completion criteria
- Capability development plans for each subproject including scope and schedule
  - Mileposts, milestones
- Technical risks and mitigation strategies identified
- Key integration points and dependencies identified

Management Processes

- Project planning
  - Activity/milestone development
  - Maintaining agility
- Project tracking
  - Technical leaders and supporting tools (Jira, Confluence, Primavera); Dashboards; Milestone reports; Monthly reports
- Project assessment
  - External reviews; Milestone review and approval; Stakeholder discussions
- Dependency management
- Risk management
- Change management

The Final Design Report provides a comprehensive description of the technical scope of the ECP

ECP overview:

- Technical focus areas, KPPs
- Integration among project elements and with DOE Facilities
- ECP schedule and funding
- Project planning, tracking, and assessment
- Risk and change management

L2 and L3 areas:

- In depth description of each KPP, verification procedures
- High level risks for each focus area
- Planning, assessment and prioritization
- Engagement with other focus areas, facilities

Each L4 project:

- Project overview, description of challenge problems and FOMs or impact goals and metrics
- FY20–FY23 development plans
- Accomplishments
- Major integration points
- Technical risks and mitigation strategies
- Activities/Milestones


ECP

EXASCALE COMPUTING PROJECT

Final Design Report

May 2019

Review Recommendations have helped to refine ECP's Final Design

Recommendations

- The review team recommends the project proceed to CD-2 at the earliest opportunity
- Transition to Operations - the PEP requires a Transition to Operations Plan, and the project should contribute a written recommendation to assist DOE sponsors. *(Considered a risk reduction for staff retention near the end of the ECP project)*

  Transition to Operations
- Keep the KPPs as they are, but consider some fine-tuning to criteria particularly in KPP-2 (AD) and KPP-3 (ST). *(Refine KPP-2 completion criterion; add base and stretch goals. For KPP-3, complete conversion to new scoring system)*

  KPP Update
- The project should further mitigate architecture risks. *(Now that architectures are known, develop early access plans and schedule, consider how contingency can help)*

  Early Access

Cross cutting findings requiring action:

- Refine the precision of information used to track AD dependencies on ST products

  AD/ST Dependencies
- Continue training efforts to the end of the project

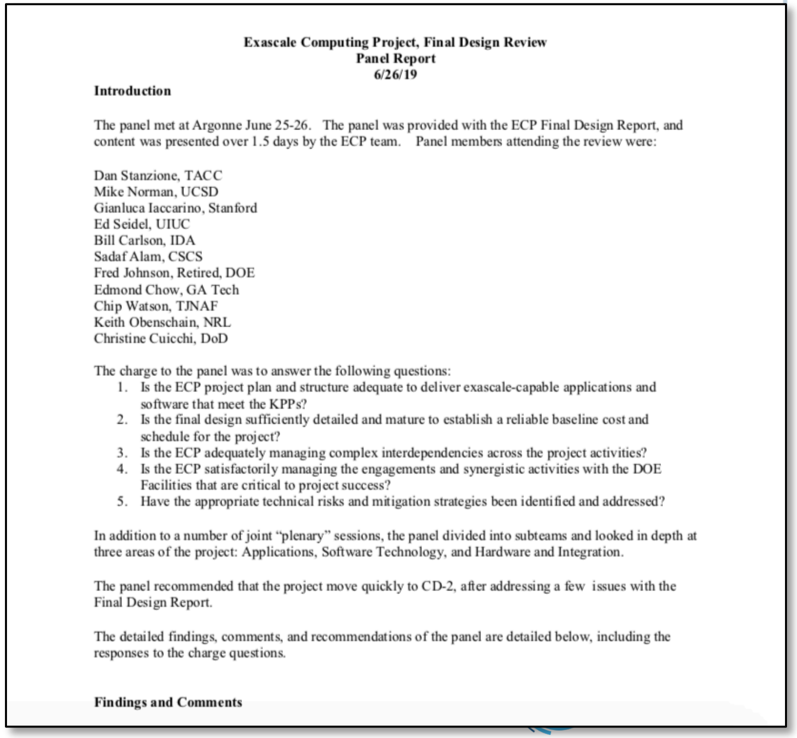

Exascale Computing Project, Final Design Review
Panel Report
6/26/19

Introduction

The panel met at Argonne June 25-26. The panel was provided with the ECP Final Design Report, and content was presented over 1.5 days by the ECP team. Panel members attending the review were:

Dan Stanzione, TACC
Mike Norman, UCSD
Gianluca Iaccarino, Stanford
Ed Seidel, UIUC
Bill Carlson, IDA
Sadaf Alam, CSCS
Fred Johnson, Retired, DOE
Edmond Chow, GA Tech
Chip Watson, TJNAF
Keith Obenschain, NRL
Christine Cuicchi, DoD

The charge to the panel was to answer the following questions:

1. Is the ECP project plan and structure adequate to deliver exascale-capable applications and software that meet the KPPs?
2. Is the final design sufficiently detailed and mature to establish a reliable baseline cost and schedule for the project?
3. Is the ECP adequately managing complex interdependencies across the project activities?
4. Is the ECP satisfactorily managing the engagements and synergistic activities with the DOE Facilities that are critical to project success?
5. Have the appropriate technical risks and mitigation strategies been identified and addressed?

In addition to a number of joint "plenary" sessions, the panel divided into subteams and looked in depth at three areas of the project: Applications, Software Technology, and Hardware and Integration.

The panel recommended that the project move quickly to CD-2, after addressing a few issues with the Final Design Report.

The detailed findings, comments, and recommendations of the panel are detailed below, including the responses to the charge questions.

Findings and Comments

U.S. DEPARTMENT OF ENERGY | Office of Science

NNSA National Nuclear Security Administration

ECP is Formulating its “Transition to Operations” Plan

- Document goals:
  - Highlight the key artifacts ECP is leaving behind that should be sustained post-ECP
  - Suggest strategies for that sustainment
- Key Artifacts
  - Application codes that can be used for scientific exploration on exascale machines; require ongoing development, maintenance and support
  - Ensure that application workflows are able to be augmented with AI / ML
  - Software Products that are an integral part of the exascale ecosystem
  - New strategies across the DOE complex for testing and deploying software (continuous integration at Facilities)
  - Highly trained workforce that works collaboratively across the DOE complex; a key consideration for ECP is helping ensure they can fully leverage exascale computers (including ML and AI strategies for scientific computing)
  - Sophisticated project management tools and best practices for large-scale, distributed R D & D projects
  - External engagements in HPC with international collaborators, other US gov’t agencies, and industry

The ECP Key Performance Parameters (KPPs)

KPP definition: A vital characteristic, function, requirement or design basis that if changed, would have a major impact on the facility or system performance, scope, schedule, cost and/or risk or the ability of an interfacing project to meet its mission requirements

ECP KPPs…

- are driven by the scope of the project and the mission needs statement, are not required to encompass the full scope of the project
- have an associated minimum threshold value (required for project completion) and a desired objective value
- provide a measurable benchmark can be tracked to measure progress during project execution
- drive integration among ECP project focus areas and technical teams
- have evolved significantly over time based on gained experience, feedback from independent project reviews, and sponsors
- have been approved by DOE

Based on feedback from the FDR we have revised KPP-1/2/3 to include stretch goals

KPP ID	Description of Scope	Threshold KPP	Objective KPP	Verification Action/Evidence
KPP-1	Performance of scientific and national security applications relative to today's performance	50% of selected applications achieve Figure of merit improvement ≥50	100% of selected applications achieve Figure of merit improvement stretch goal	Independent assessment of measured results and report that threshold goal is met
KPP-2	Broaden the reach of exascale science and mission capability	50% of selected applications can execute their challenge problem	100% of selected applications can execute their challenge problem stretch goal	Independent assessment of mission application readiness
KPP-3	Productive and Sustainable Software Ecosystem	Software teams meet 50% of their weighted impact goals	Software teams meet 100% of their weighted impact stretch goals	Independent assessment verifying threshold goal is met
KPP-4	Enrich the HPC Hardware Ecosystem	Vendors meet 80% of all the PathForward milestones	Vendors meet 100% of all the PathForward milestones	Independent assessment of the impact and timeliness of PathForward milestones

AD (KPP-1, KPP-2)

ST/CD (KPP-3)

HI (KPP-4)

KPP-1 Definition: Apps Performance

- KPP-1 is based on a Figure of Merit (FOM) defined individually for each project to capture the relevant **scientific work rate** for an application.
- Goal of KPP-1 is to measure the overall impact of ECP project, including both hardware-driven and algorithmic improvement.
- Each application measured a **baseline FOM value** at the inception of ECP.
- KPP-1 is calculated as the ratio of the FOM **on the exascale challenge problem** to the baseline

$$KPP\text{-}1 = \frac{FOM_{exascale}}{FOM_{baseline}}$$

- The FOM ratio is measured throughout the project to track progress.

KPP-1 Threshold

50% of KPP-1 applications achieve Figure of merit* improvement ≥50

KPP-1 Objective

100% of KPP-1 applications achieve Figure of merit improvement stretch goal

Example Base and Stretch Goal: ExaSky

- Enable extraction of fundamental physics from upcoming cosmological surveys
- FOM: number of particles and time to solution as measured per time step for gravity and hydro solvers
- Base goal: Volume of 3000^3 Mpc/h, # particles: 23,040^3, particle mass: ~2 x 10^8
- Stretch goal: Volume of 3000^3 Mpc/h, # particles: 30,720^3, particle mass: ~8 x 10^7

Current Figure of Merit Improvements on Summit/Sierra

AD KPP-1 FOM Status: Measured and Extrapolated FOM Increase

Measured KPP-1 values are the ratio of the highest reported FOM to the baseline FOM.
Extrapolated values assume perfect scaling to full machine size.
The Y-Axis default is limited to a maximum of 50 to ensure smaller FOM increases are shown.

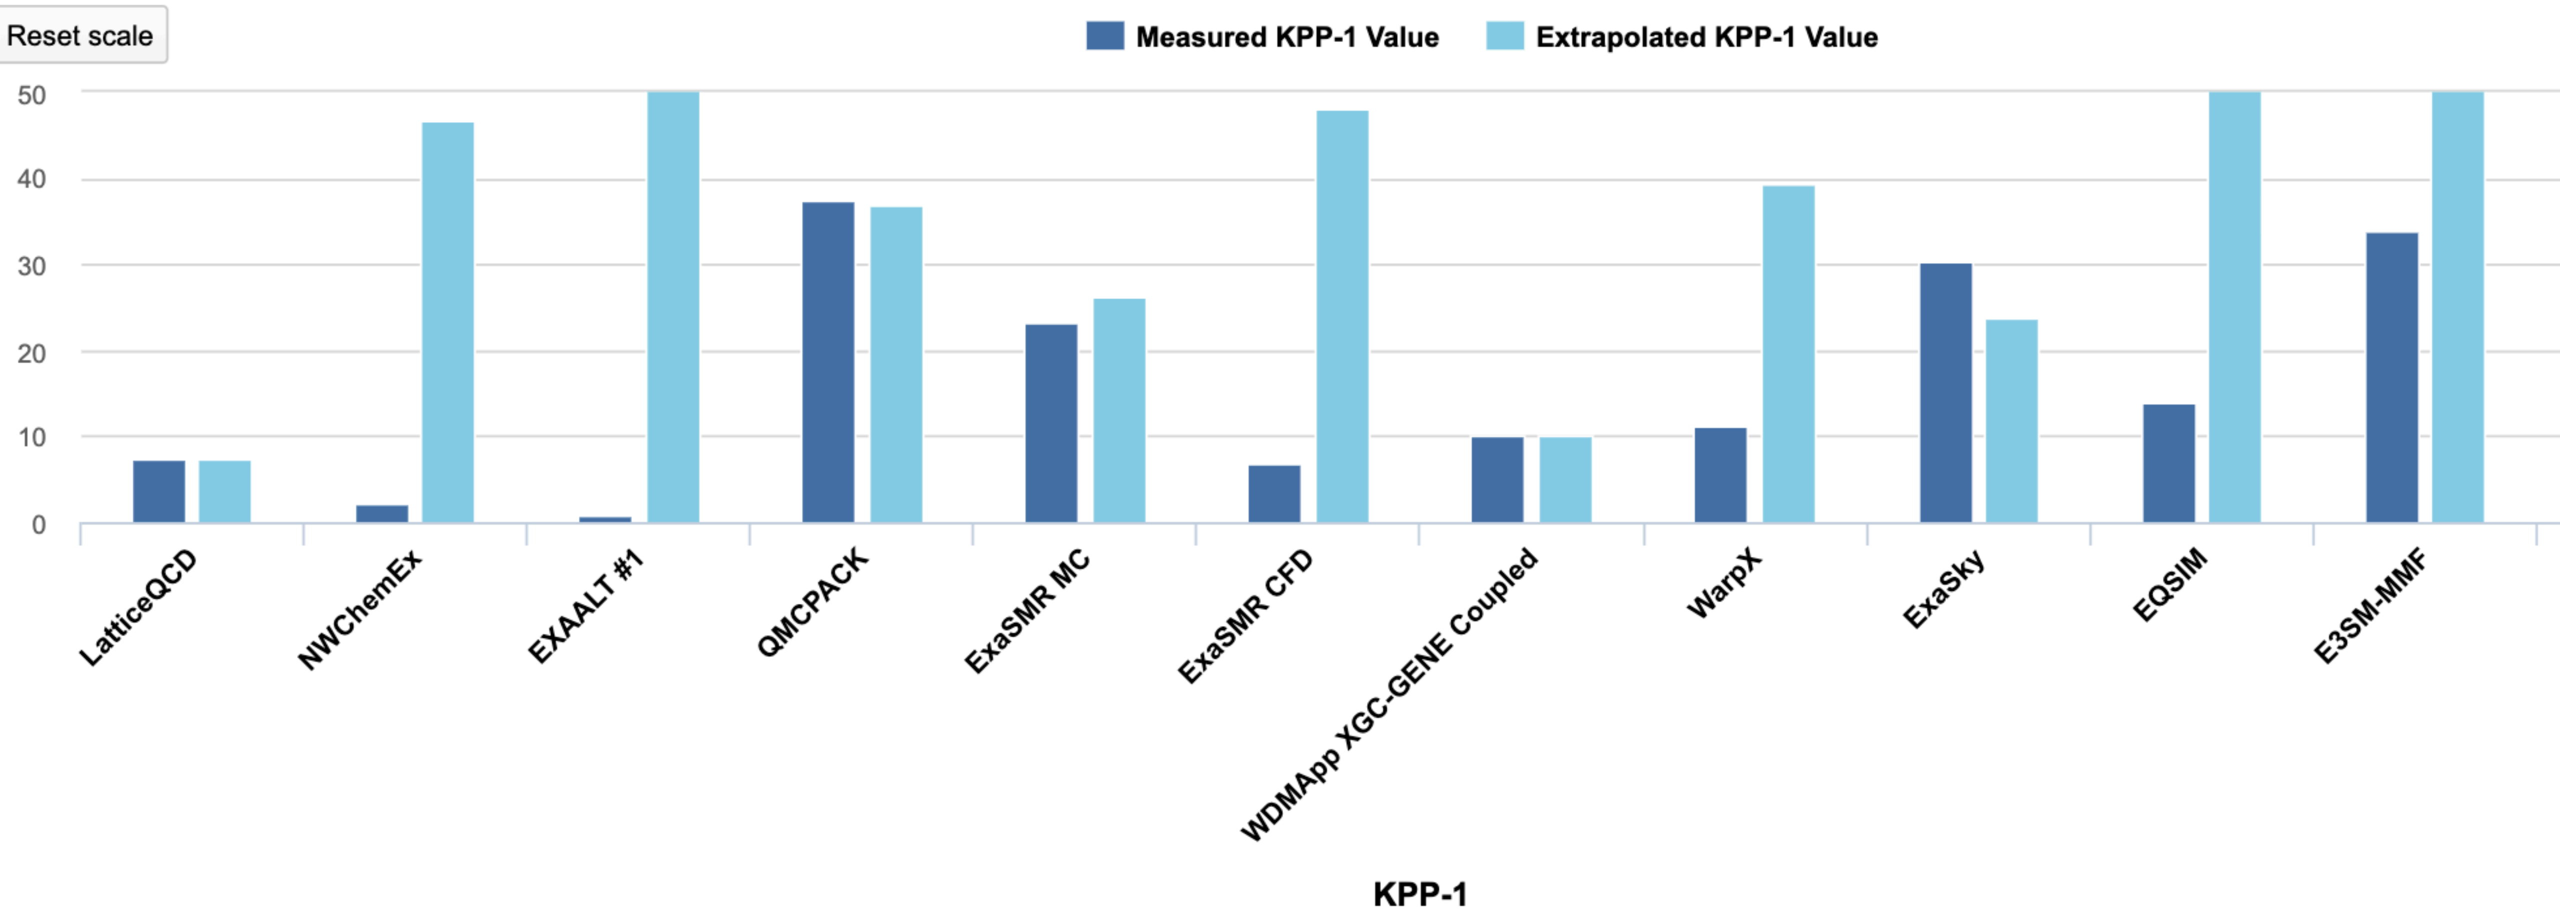

KPP-2 Definition: Apps Capability

- KPP-2 is **based on developing new mission-critical capabilities at exascale** per the ECP mission needs statement to broaden the reach of exascale computing.
- Unlike KPP-1 applications, a well-defined baseline was not available at the inception of ECP.
- To meet KPP-2 an application must successfully execute a capability demonstration of the challenge problem on an exascale platform.

- Performance requirements for KPP-2
  - Must demonstrate parallel scalability on the exascale systems
  - Must sufficiently utilize hardware accelerators on a node
  - Must execute simulation using all necessary physics and algorithmic capabilities of the challenge problem

KPP-2 Threshold

50% of KPP-2 applications can execute their exascale challenge problem

KPP-2 Objective

100% of KPP-2 applications can execute their challenge problem stretch goal

Example Base and Stretch Goal: ExaWind

- Predictive simulation of wind farms
- Base Goal: 3x3 array of wind turbines in a 4x4x1 km^3 domain; 5-MW wind turbine (126 m rotor), at rated wind speed (11.4 m/s); run in the strong scaling limit
- Stretch Goal: O(100) multi-MW wind turbines in a 10x10x1 km^3 complex domain; first target is to increase turbines and domain to demonstrate weak scaling

Verifying KPP-1 and KPP-2 completion

- Compute time for dedicated KPP-1 and KPP-2 demonstration calculations pre-negotiated with ALCF and OLCF
- Projects demonstrate KPPs on rolling basis after machine deployment
- KPP-1 and KPP-2 success is verified by an external SME review at end of project
  - KPP-1 run must be fully documented and reproducible, including any caveats
  - KPP-1 must include full documentation of baseline calculation
  - KPP-2 must demonstrate all new capability in place to execute challenge problem
  - KPP-2 must demonstrate reasonably efficient port to exascale machine
    - Make effective use of the accelerators on a node
    - Scale up to a significant fraction of the exascale machine
- Each team will be asked to provide a short report that describes the challenge problem, FOM, key steps needed to get performance

ST and co-design projects use KPP-3 to measure integration and drive creation of a productive and sustainable ecosystem

KPP-3 Basics

- **Integration Goal:**
  A statement of impact on the ECP ecosystem, consequential and sustainable use by client.
- **Metric:**
  Capability integration – Use of the product for the first time or a significant feature set recently developed representing an FTE or more worth of effort.
- **Threshold/Objective:**
  50%/100% of the weighted impact goals are met.

KPP-3 Details

- Weights correlate with scope of impact. Examples:
  - OpenMP, MPICH – Weight of 2.
  - Most – Weight of 1.
  - Legion, ParSEC – Weight of 0.5.
- Integration must represent sustainable progress, not just “tried it” or “considering it”.
- Not looking for hero-level integration score counts. Integration is hard work.

KPP-3 Scoring Summary

- **Individual** KPP-3 Goal/Metric scoring:
  - Passing – Minimum acceptable success value.
  - Stretch – Maximum reasonable achievable value, used for normalizing.
  - Tentative Present – Current score on existing platforms.
  - Confirmed Present – Score after success migrating to exascale environment.
- **Cumulative** KPP-3 scoring:
  - Weighted sum of individual KPP-3 scores that have reached minimum value.
  - Currently 82 total KPP-3 Jira issues for ST.
    - Impact levels: 6 High, 68 Normal, 8 Risk-Mitigating.
    - The objective value for the cumulative KPP-3 would be 2*6 + 1*68 + ½*8 = 84.
    - The threshold value is 42 (half of 84).
  - Scenario: 2, 12 and 0 of the high, normal, and risk-mitigating KPP-3 issues have achieved minimum.
    - Present value of the cumulative KPP-3 would be 2*2 + 1*12 + ½*0 = 16.
- Note: All KPP-3 issue scores are final when they are confirmed on exascale environments.

Early Access Project will help ensure adequate access to Facility resources

Goal	Method of Demonstration
• Acquire access to Aurora and Frontier "N-1" and final resources architectures • Successful execution of the Early Access project element ensures ECP subprojects have sufficient resources to carry out research, development, and deployment activities to meet their respective FOM and KPP targets.	• Negotiations with ALCF and OLCF for dedicated access to 0.5 to 1.0 rack of "N-1" and Aurora and Frontier resources is underway. This will determine the mode of access and agreed-upon deliverables to the DOE HPC facilities.

The ECP is actively managing several dependencies both within the project and with the DOE Facilities

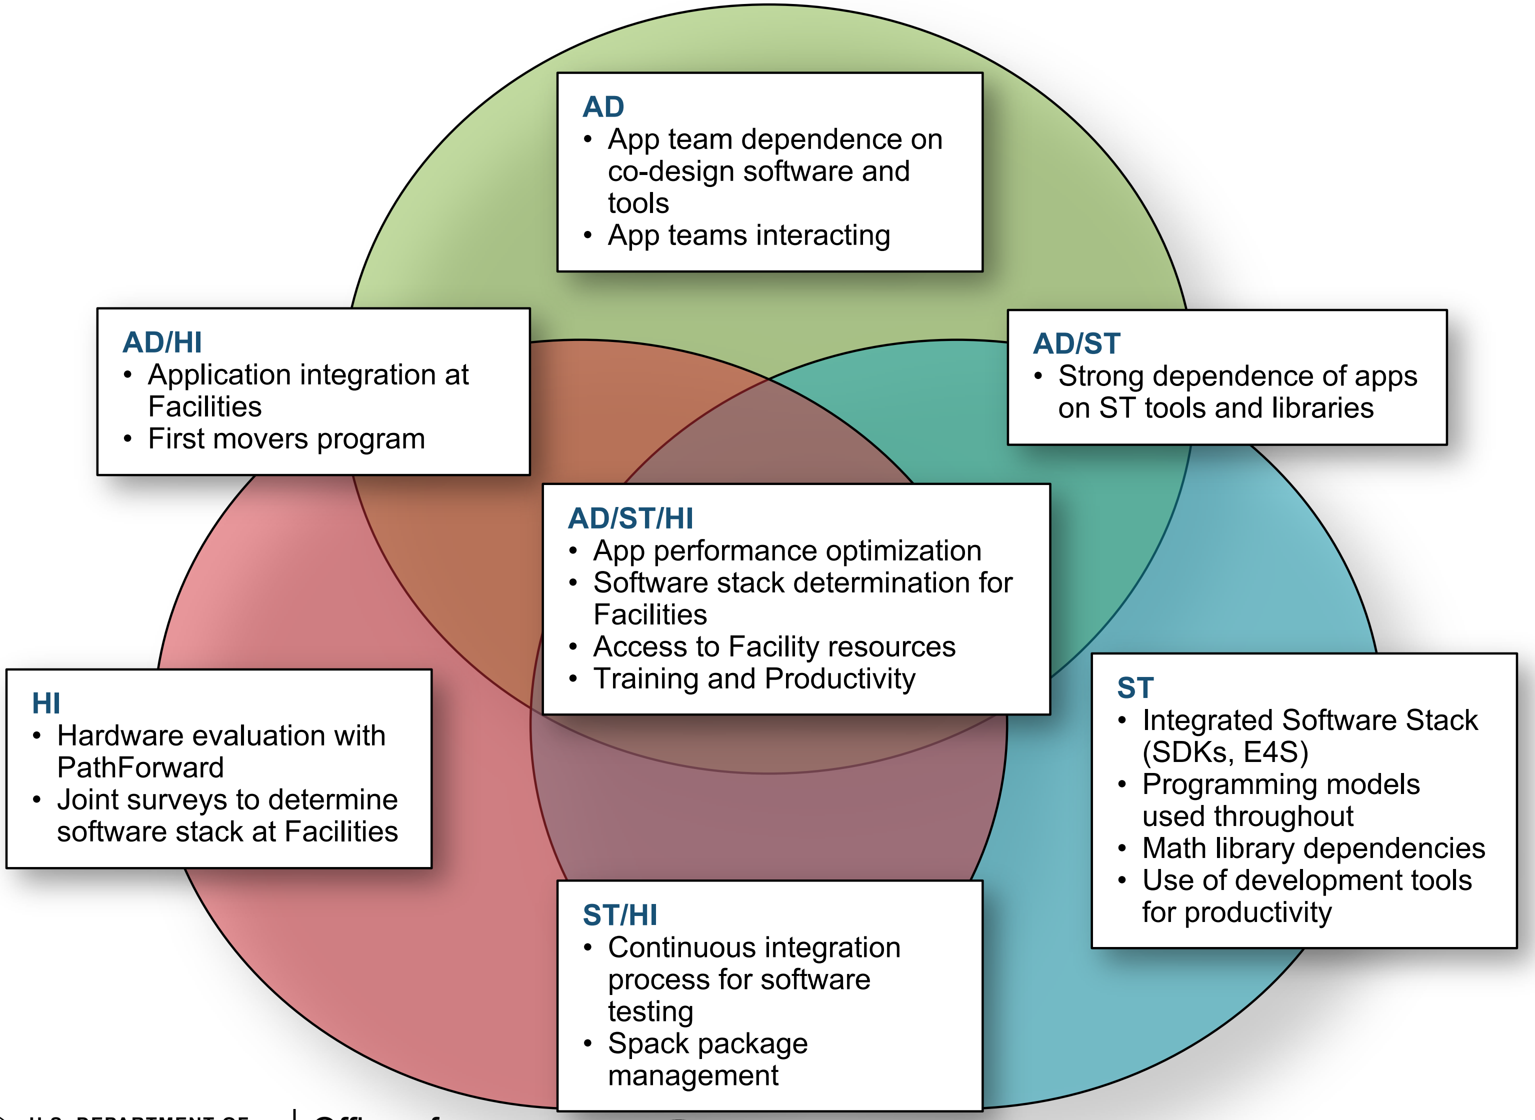

DOE Facility Dependencies

- ECP requires access to Facility resources to develop, test, and demonstrate KPPs
- ECP software stack must leverage and complement vendor and Facility software stack
- PathForward program designed to keep US industry healthy and feed into Facility procurements

Managing AD-ST complexity, i.e. “taming the hairball”

- We currently have significant usage of ST and co-design products by AD application teams.
- To manage dependencies, it was necessary to first gather accurate data:
  - AD applications filled out detailed tables of software specs and dependencies on Confluence
  - ST teams reported application dependencies
  - HI interviews with application teams
- Data was not initially fully consistent.
  - ST teams reported working with applications who didn't list them as dependencies
  - Applications reported depending on ST projects who didn't list them as customers
- Consistent interdependency data now being imported into ECP’s database for configuration control, analysis and planning
  - **Has this ever been done? We don’t think so . . .**

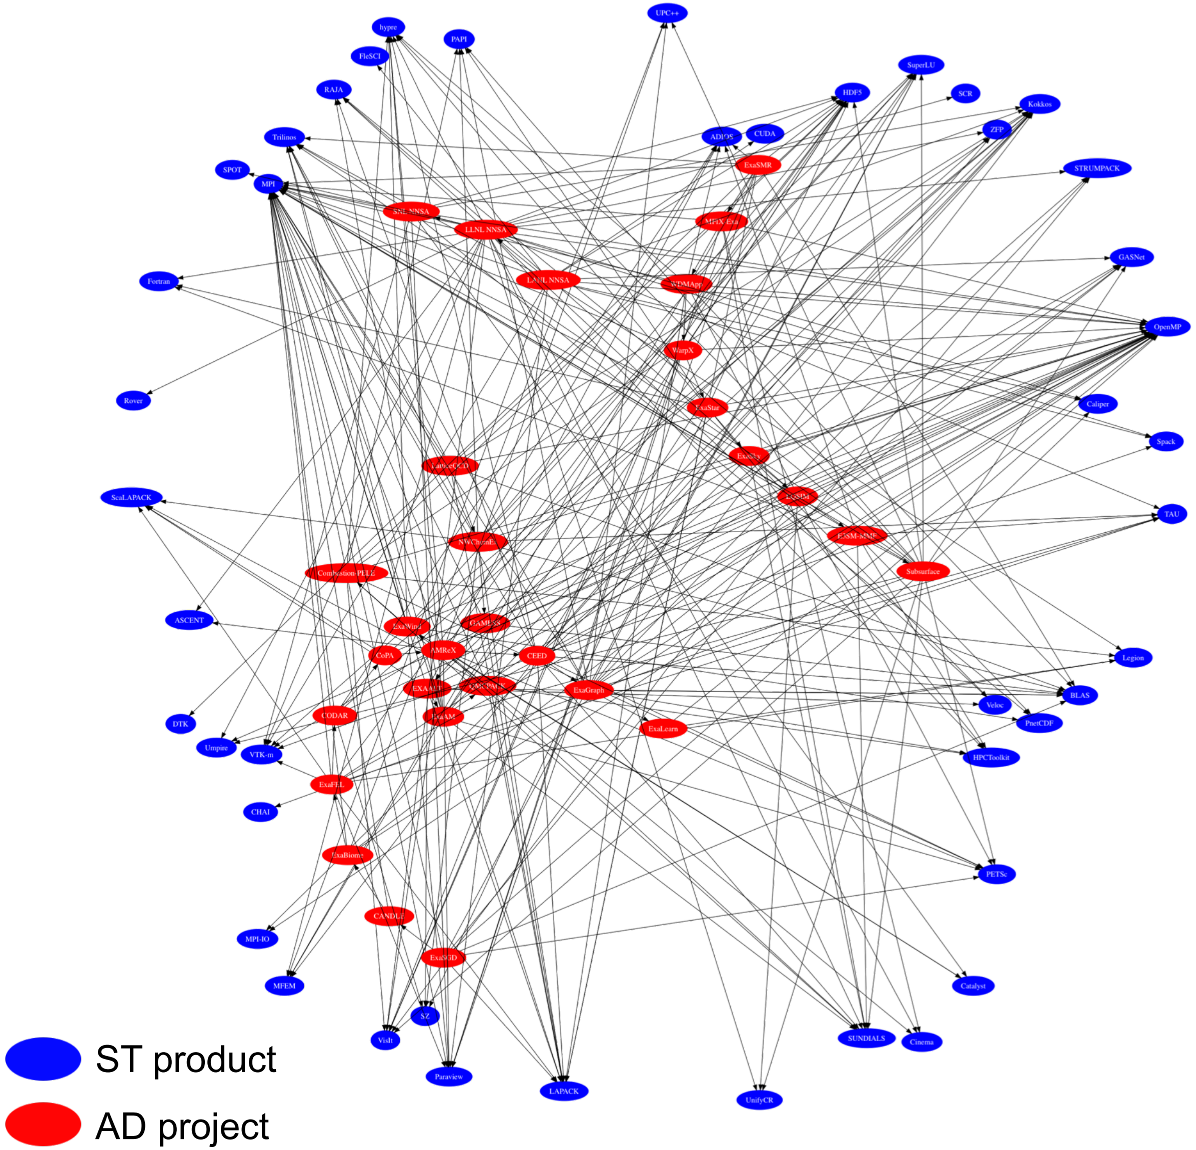

AD/ST dependencies varied based on how data was collected and lacked a time dimension

ST Project Owner	ST Column AD Row Proje	ST?	2.2.1.01 - LatticeQCD Chroma, MILC, CPS	2.2.1.02 - NW ChemEX	2.2.1.03 - GAMESS GAMESS, LIBCCHEM	2.2.1.04 EXAALT PARSPLICE, LAMMPS, LATTE	2.2.1.05 ExaAM Diablo, TruchasPBF, ExaCA, MEUMAPPS-SS,	2.2.1.06 QMCPACK	2.2.2.01 ExaWind Nalu-Wind	2.2.2.02 Pele, PeleC, PelePhysics, PeleLM	2.2.2.03 ExaSMR Shift, OpenMC, Nek5000	2.2.2.04 MFIX-Exa Cabana	2.2.2.05 WDMApp GENE, XGC	2.2.2.06 WarpX fftMPI	2.2.3.01 ExaStar Castro, FLaSH	2.2.3.02 ExaSky Nyx, HACC	2.2.3.03 EQSIM SW4	2.2.3.04 Subsurface chombo-Crunch, GEOSX	2.2.3.05 E3SM-MMF	2.2.4.02 ExaSGD GOSS, GridPACK, PIPS, StructjuMP	2.2.4.03 CANDLE CANDLE	2.2.4.04 ExaBiome MetaHipMer	2.2.4.05 ExaFEL psana, CCTBX, LUNUS, M-TIP	2.2.5.01 ATDM LANL Ristra	2.2.5.02 ATDM LLNL MARBL	2.2.5.03 ATDM SNL SPARC, EMPIRE	2.2.6.03 CODAR	2.2.6.04 CoPA	2.2.6.05 AMReX	2.2.6.06 CEED	2.2.6.07 ExaGraph	2.2.6.08 ExaLearn
2.3.4.09	ADIOS	ST		T, DM - Y		O	T, DM - I	T	T		O		R						R		T						R	T, DM - i		T, DM - I		
2.3.4.16	ALPINE	ST					DM - Y		DM - E	T, DM - Y	DM - Y	T, DM - Y		O	DM - Y	T, DM - I	DM - Y	DM - E					DM - Y	R	DM - E	DM - E			R	R		
2.2.6.05	AMREX	ST					R		O			R		R	R	R							O									
2.3.4.16	Ascent - in ALPINE	ST								DM - I		DM - I		DM - I		DM - I	DM - I							DM - E	R				DM - Y	DM - I		
2,3,3,13	BLAS	ST	O	R	R			R					R							R											R	
2.3.4.14	blosc - Veloc	ST														O																
2.3.6.02	Caliper	ST																DM - E					DM - I	R	R					DM - I		
2.3.4.12	Catalyst - ALPINE	ST							O			DM - I		DM - I		DM - I								DM - E		DM - Y			DM - Y			
2.3.1.18	CHAI	ST																R														
2.3.4.16	Cinema	ST										DM - Y		DM - I	DM - I	DM - I								DM - E	O			DM - I	DM - Y	DM - I		
2.3.5.01	cmake	ST																							R							
2.3.3.12	conduit	ST																							R							
2.3.4.06	EZ	ST		O																												
2.3.2.12	FLANG	ST					T																						T			
2.3.6.01	FleCSI	ST																						R								
2.3.4.15	HDF5	ST	O	O	R		R	R	DM - I	T, DM - Y	R	T, DM - I		T, DM - Y	O	T, DM - I	R	R					R	R	R	R			R	DM - I		
2.3.2.08	HPC Toolkit	ST		T	T		DM - I	T, DM - Y	DM - I	DM - I				T, DM - I	DM - Y	T, DM - I		T	DM - Y				T, DM - I		O				DM - I		DM - Y	
2.3.6.01	HXHIM	ST																						O								
2.3.3.12	Hypre	ST					R		R	O			R		O			R							R				T, DM - Y	T, DM - E		
2.3.1.18	Kokkos	ST	DM - Y	T		T, DM - E	DM - Y	R	R	T, DM - Y				DM - I					DM - Y			DM - I		R		R		R		R	T, DM - E	T
2.3.3.13	LAPACK	ST	O	R	R	R	R	R					R		O					R			R									
2.3.1.08	Legion	ST								DM - Y	DM - I										T		T, DM - E	R								R
OSS	LLVM	ST	T					T											T	R									T	T		T
2.3.3.13	MAGMA - Dongarra	ST	DM - Y	DM - E		O		DM - i															DM - I	DM - I	O			R		DM - I		
2.3.6.02	MFEM 3.4	ST					R																	DM - Y	R					R		
hese MPI referer	MPI	ST	R	R	R	R	R	R									R			R		R	O	R	R	R		R			O	
2.3.1.07	Mpich	ST	DM - E	T, DM - E		DM - I			DM - E	DM - Y	DM - Y			DM - Y	DM - Y	DM - Y	DM - Y	DM - E	DM - Y	DM - E	R	DM - Y	DM - E			DM - Y				R	DM - Y	
2.3.1.13	OpenMP	ST	R	DM - I	R	DM - E	Dm - Y	R	DM - Y	DM - E	DM - Y	DM - E	DM - E	DM - E	DM - E	DM - E	R		DM - E	DM - Y		DM - E	DM - E	O	R	R		DM - E	DM - E	R	DM - E	R
2.3.1.13	OpenMP + LLVM (SOLLV	ST																														
2.3.1.17	OpenMPI, PMIX, Qthrea	ST	DM - E	DM - E		T, DM - Y		T	DM - E	DM - Y	DM - Y			DM - Y	DM - Y	DM - Y	DM - Y	DM - Y	DM - E	DM - I		DM - Y	DM - Y	T	T	DM -Y			O		O	
2.3.2.06	PAPI	ST		O											DM - Y	DM - I		DM - Y													DM - Y	
2.3.4.12	Paraview - ALPINE	ST							O															O					O			
2.3.3.06	PETSc	ST					DM - E			Dm - Y			R		T, DM - I			R		R			DM - I		DM - I				T, DM - Y	R		
2.3.3.06	Petsc - SLEPc	ST											R																			
2.3.4.10	PnetCDF / NetCDF	ST							DM - E										R						DM - I	R				DM - I		
2.3.1.17	Q-threads	ST																						DM - I		R						
2.3.1.18	RAJA	ST					DM - Y										R	R		DM- I				DM - I	R					T, DM - Y		
2.3.3.13	SCALAPACK	ST																					R									
2.3.4.08	SCR	ST																							O							
2.3.3.13	SLATE/MAGMA	ST		DM - E	T	O		R												T			DM - I					DM - Y				
2.3.5.09	Spack	ST																						R	R					R		
2.3.6.02	SPOT	ST																							R							
2.3.3.07	STRUMPACK	ST											T							T, DM - Y					O					R	R	
2.3.3.12	SUNDIALS	ST					T, DM - Y			T, DM - E		DM - E			T, DM - I	R		DM - Y	DM - Y						O				R	R		
2.3.3.12	SUNDIALS _ CVODE	ST														R																
2.3.3.07	SuperLu	ST		DM - E			DM - Y		DM - Y				O					DM - Y		O			DM - I		DM - I	DM - E				T, DM - Y	R	
2.3.3.08	Swig	ST																			R											
2.3.4.14	SZ	ST		T, DM - Y	T	O								DM - I	DM - Y	T, DM - I							T				R					
2.3.2.10	Tau	ST	DM - I	O			DM - Y		DM - I	DM - I				DM - I					DM - Y				DM - Y		O		R		DM - I		DM - Y	
2.3.3.08	Trillinos	ST	DM - Y				DM - Y		R		R							O	T				DM - I	R		R					DM - Y	
2.3.1.18	UMPIRE	ST															R	R						R	R					DM - I		
2.3.4.15	UNIFY	ST		DM - I														T, DM - Y							O							
2.3.1.14	UPC++ & GASNet	ST		T														O	T			R		O					DM - Y		O	
2.3.4.14	Veloc	ST	T, DM - I	O		O										R									O							
2.3.4.16	Visit	ST									O			O											R							
2.3.4.13	VTK-m	ST							DM - Y					DM - I	DM - Y	O		DM - I					DM - I	R	R	R		DM - I	DM - Y	R		
2.3.3.01	xSDK	ST																											R	R		
2.3.4.16	zfp	ST		DM - Y				T					T	DM - I			T, DM - I							T	O		R			R		

ECP's agile project management database allows more rigorous tracking of AD/ST/CD dependencies

Official ST and AD Product lists enable rigorous dependency management

ST Product List:

Use widely-recognized product names. Enables mapping between AD & Facilities dependencies and ST development efforts.

- MPI – MPICH, OpenMP
- C++/C/Fortran - LLVM
- Fortran – Flang
- hypre – hypre

67 Products with descriptions and POC

AD Code List:

Currently creating a complementary AD code dictionary list to facilitate interactions among teams

ADIOS	Flux
AML	Fortran
Ascent	GASNet
BLAS	Ginkgo
C	HDF5
C++	HPCToolkit
Caliper	hypre
Catalyst	Kokkos
CHAI	KokkosKernels
Cinema	LAPACK

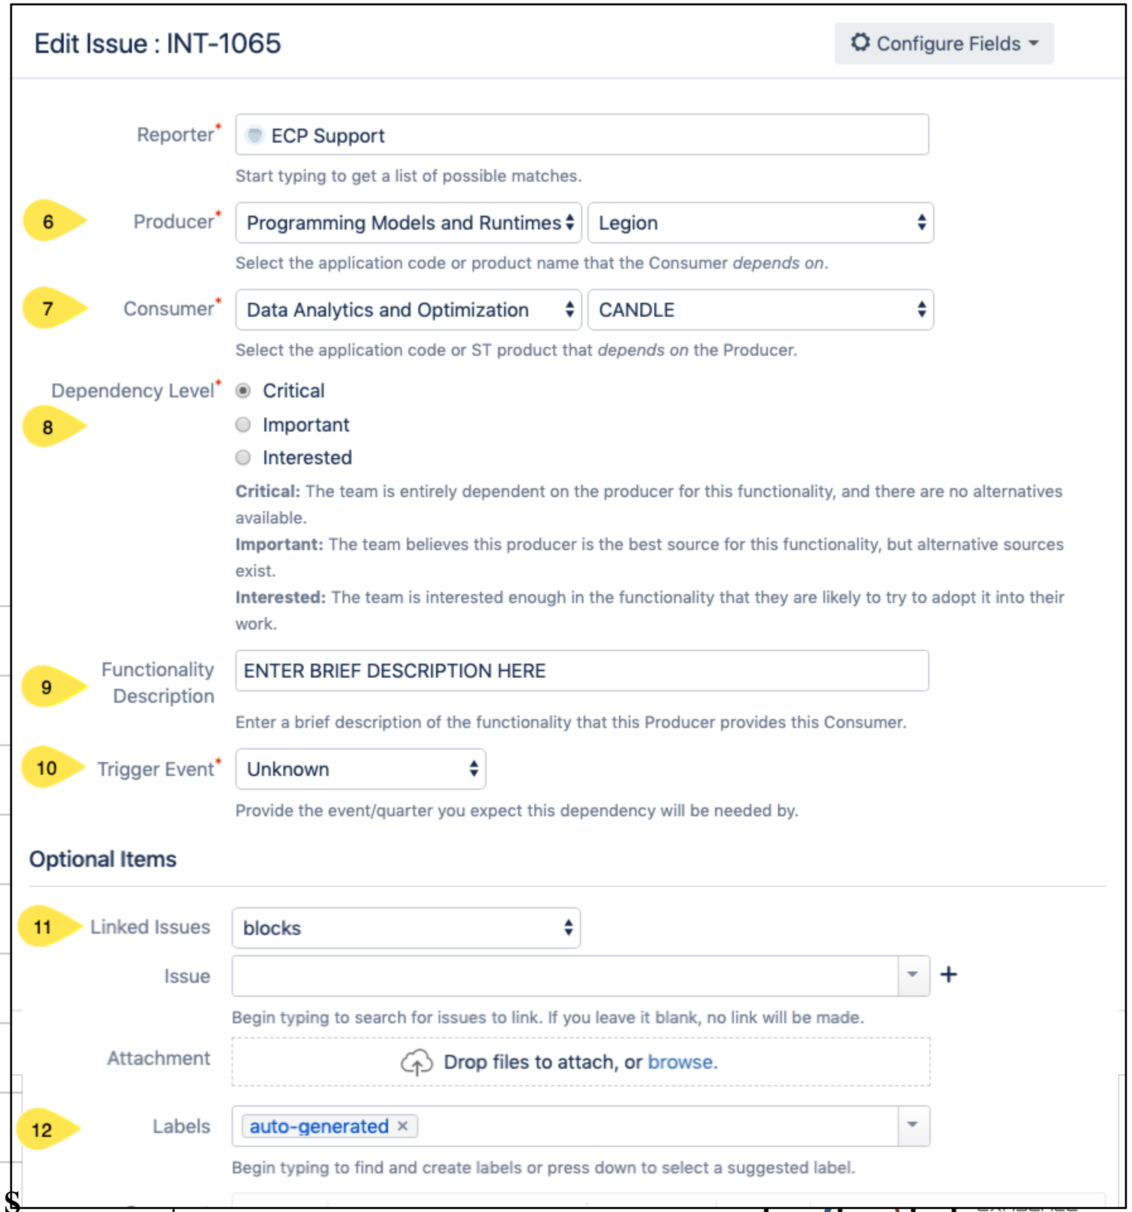

AD/ST Dependencies

ECP's use of agile PM tools (JIRA) to track these dependencies allows for significant real time data analytics

Dependency Matrix

Note: By default, this chart only shows ST -> AD dependencies. To show other kinds of dependencies, change the second and third dropdowns.

Critical, Important, Interested ▾ | ST Producers ▾ | AD Consumers ▾ | Draft, Approved ▾

	+ Chemistry and Materials	+ Co-Design	+ Data Analytics and Optimization	+ Earth and Space Science	+ Energy	+ Nati Sec
	Dependencies	Dependencies	Dependencies	Dependencies	Dependencies	Deper
+ Data & Visualization	42	22	17	35	22	
+ Development Tools	42	19	25	23		
+ Math Libraries	57	23	44	23		
+ Programming Models and Runtimes	39	26	20	22		
+ Software Ecosystem	5	9	1			
Total	185	99	107	103		

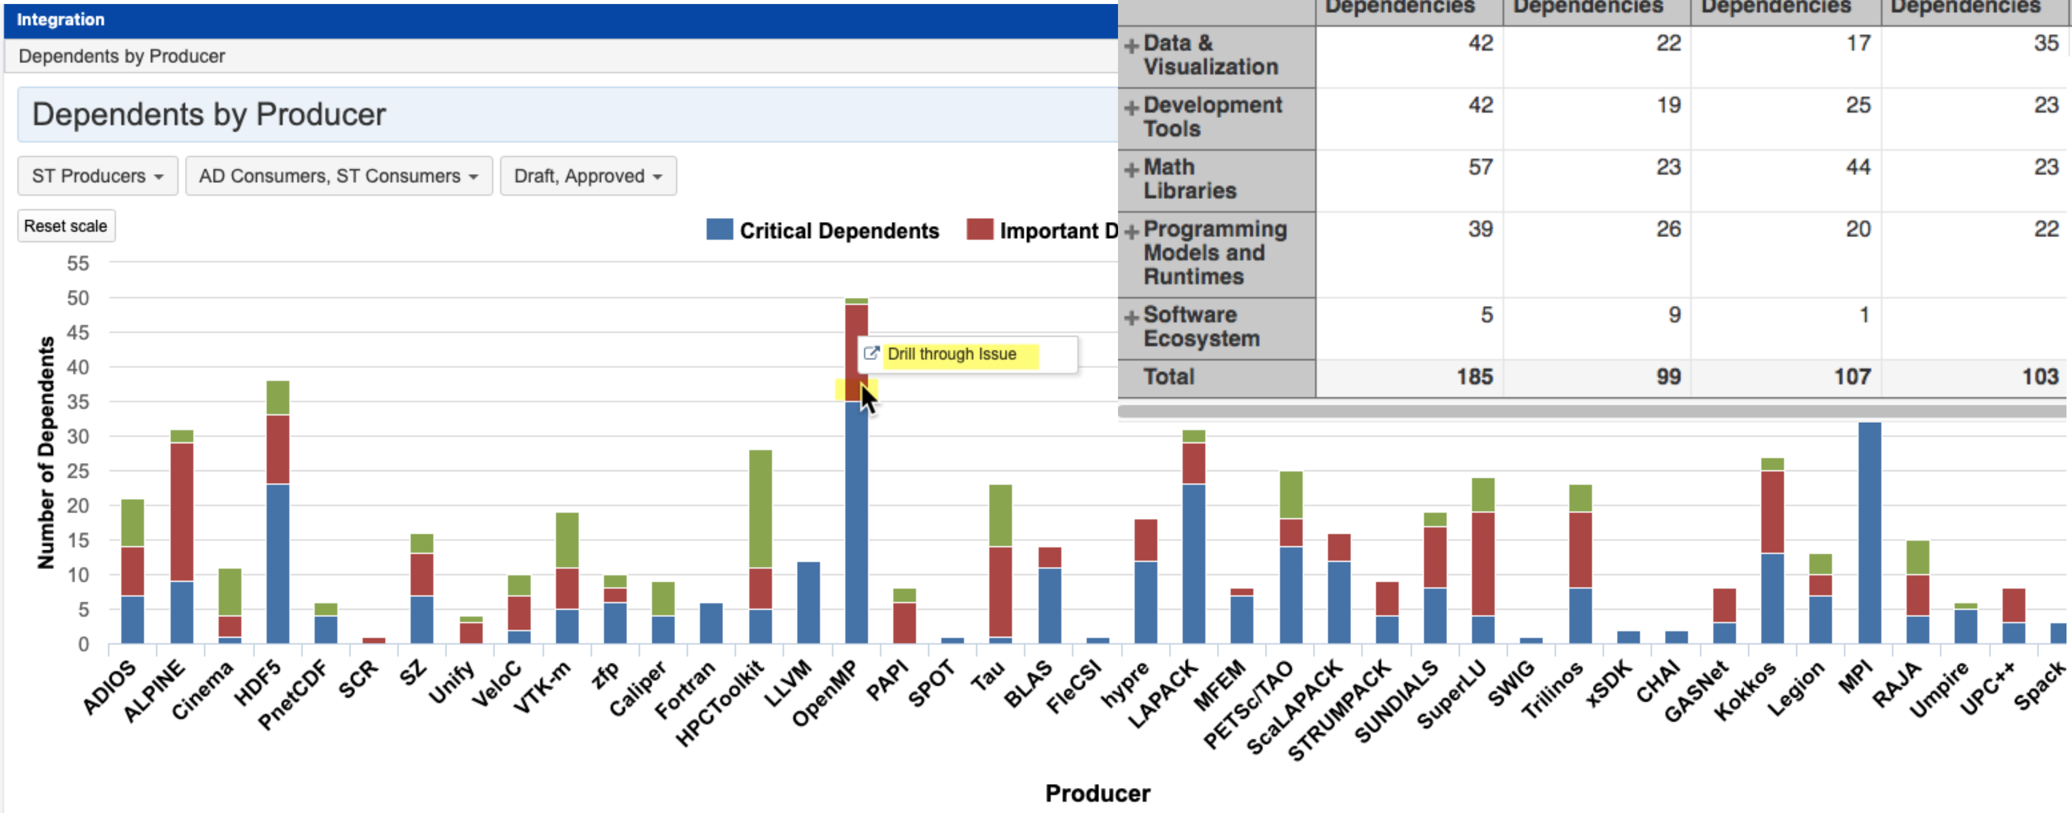

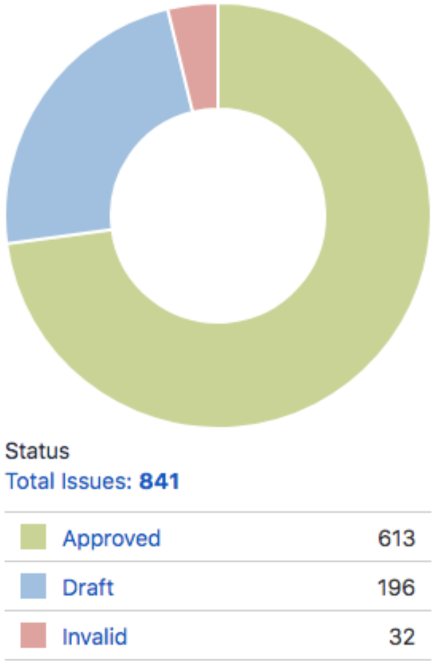

The ECP is ready to set its performance baseline

CD-2/3 review on December 3-5, 2019

Charge questions

- 1. Have the recommendations from the Oct 2018 IPR and Final Design Review been addressed?
- 2. Are the proposed cost and schedule and scope baselines sufficient to meet the KPPs and complete the project?
- 3. Have the risks been adequately identified and have appropriate risk responses been developed for this phase of the project? Is there adequate contingency?
- 4. Is the management of the ECP appropriately structured and empowered to ensure the project's success? Has ECP accounted for the critical external dependencies required for ECP success?
- 5. Has the project met all the requirements for a CD-2/3 and is the project ready for CD-2/3 approval?

Major Preparation Activities

- Baseline Scope set in the Final Design Review; revised to reflect refined KPP definitions and new scope associated with recent DOE ASCR budget increase
- Baseline Schedule: L4 Technical planning for FY20-FY23 complete; significant details in FY20, planning packages in FY21-FY23. Resource loaded schedule entered into Primavera
- Baseline budget: Assigned to L4 projects for the remainder of the project
- FDR review responses including AD/ST dependency
- Risk register significantly updated to reflect latest information from Facilities
- Document Scrub
- Red team review scheduled Oct 15-16, 2019

Questions?