Exascale Computing Project: Status and Next Steps

John Sarrao (PADSTE, LANL)
Chair, ECP Laboratory Operations Task Force (LOTF)

DOE Office of Science Advanced Scientific Computing Advisory Committee (ASCAC)
April 17, 2018

ECP is a critical component of the broader ECI strategy

Computational and data science solutions for the Nation

- Broader ECI elements are essential for success
  - Deploying exascale systems quickly enough to impact schedule-sensitive mission problems
  - Maintaining and advancing the “HPC ecosystem” after ECP
  - Developing U.S. industry and academia partnerships to ensure the benefits of advanced computing have broad and enduring impacts
- ECP Vision (where ECP is headed)
  - *Accelerating innovation with exascale simulation and data science solutions that enhance US economic competitiveness, improve our quality of life, and strengthen our national security.*
- ECP Mission (purpose; reason for existence)
  - *Deliver e*xascale-ready applications and solutions that address currently intractable problems of strategic importance and national interest;
  - *Create and deploy* an expanded and vertically integrated software stack on DOE HPC exascale and pre-exascale systems, defining the enduring US exascale ecosystem
  - *Leverage* US HPC vendor R&D activities and products into DOE HPC exascale systems.

ECP enables the continued US global strategic advantage in science and technology which is the foundation of future revolutions in technology development, energy security, scientific discovery, and national security.

ECP History: from the Start of Project to Present

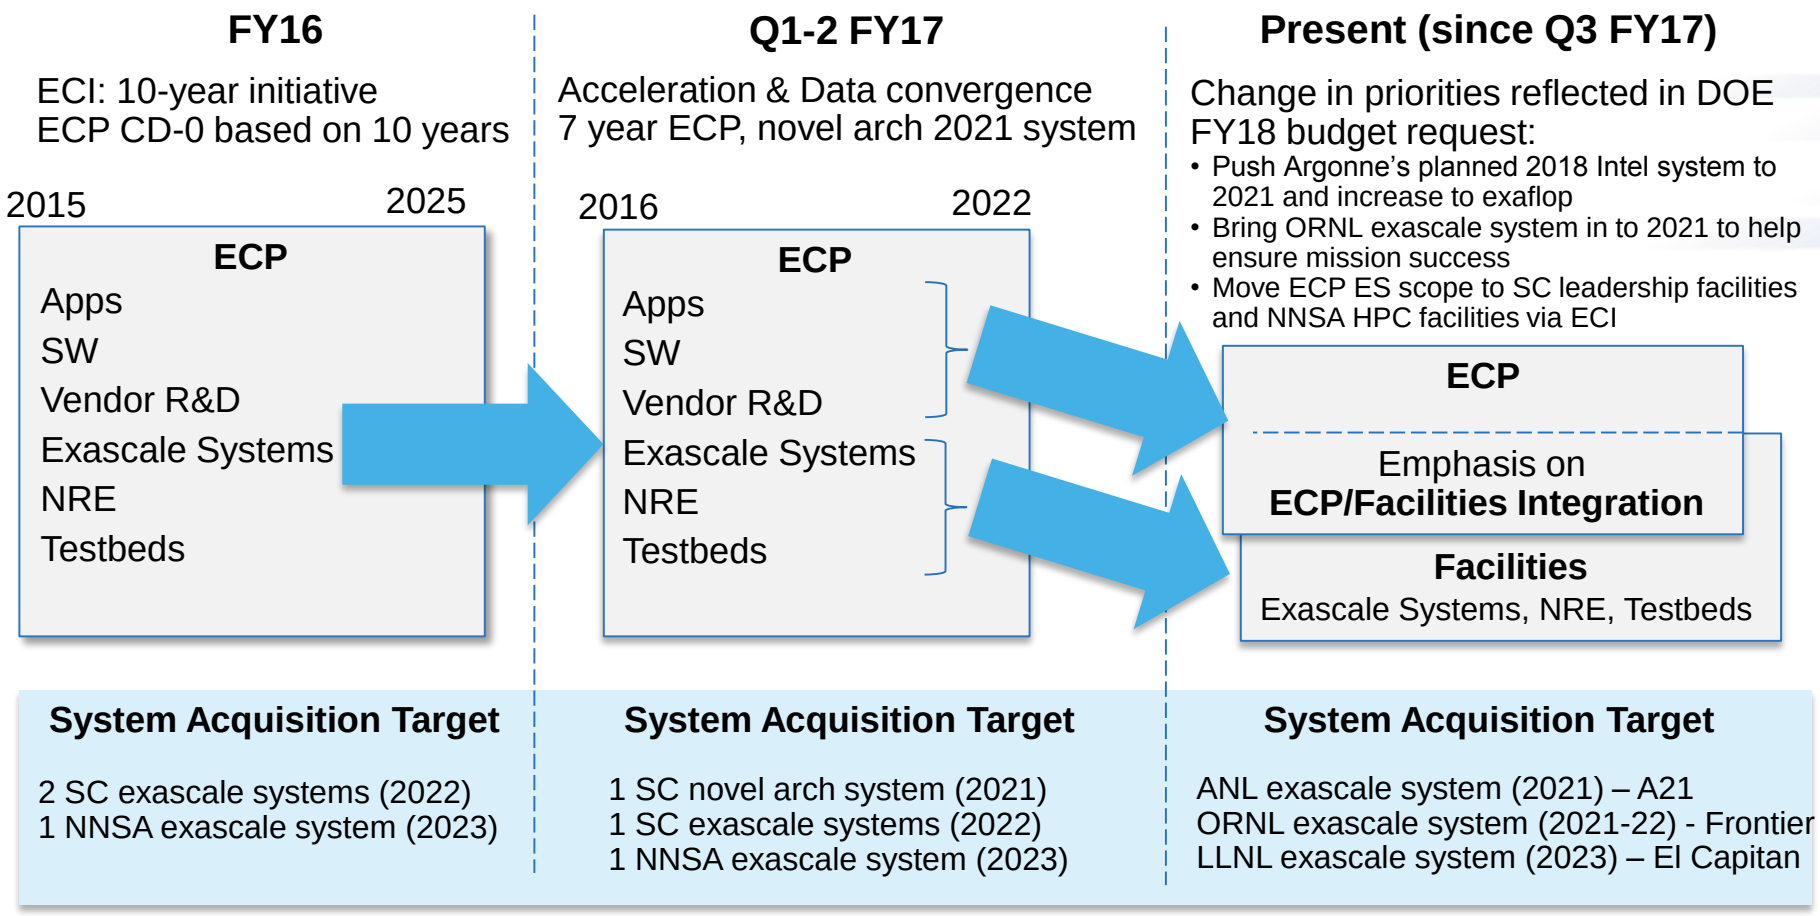

ECP Areas of Technical Focus

Application Development

The Application Development effort develops and enhances the predictive capability of applications critical to the DOE, including the science, energy, and national security mission space. The scope of the AD focus area includes

- targeted development of requirements-based models, algorithms, and methods,
- integration of appropriate software and hardware via co-design methodologies,
- systematic improvement of exascale system readiness and utilization, and
- demonstration and assessment of effective software integration.

Software Technology

Software Technology spans low-level operational software to high-level applications software development environments, including the software infrastructure to support large data management and data science for the DOE SC and NNSA computational science and national security activities at exascale. Projects will have:

- line of sight to application's efforts
- inclusion of a Software Development Kit to enhance the drive for collaboration, and
- delivery of specific software products across this focus area.

Hardware and Integration

This focus area is centered on the integrated delivery of specific outcomes (ECP Key Performance Parameters, or KPPs) and products (e.g., science as enabled by applications, software, and hardware innovations) on targeted systems at leading DOE computing facilities. Areas include:

- PathForward
- Hardware Evaluation
- Application Integration at Facilities
- Software Deployment at Facilities
- Facility Resource Utilization
- Training and Productivity

DOE Program and core laboratory agreements provide effective oversight, advisory bodies provide effective advice

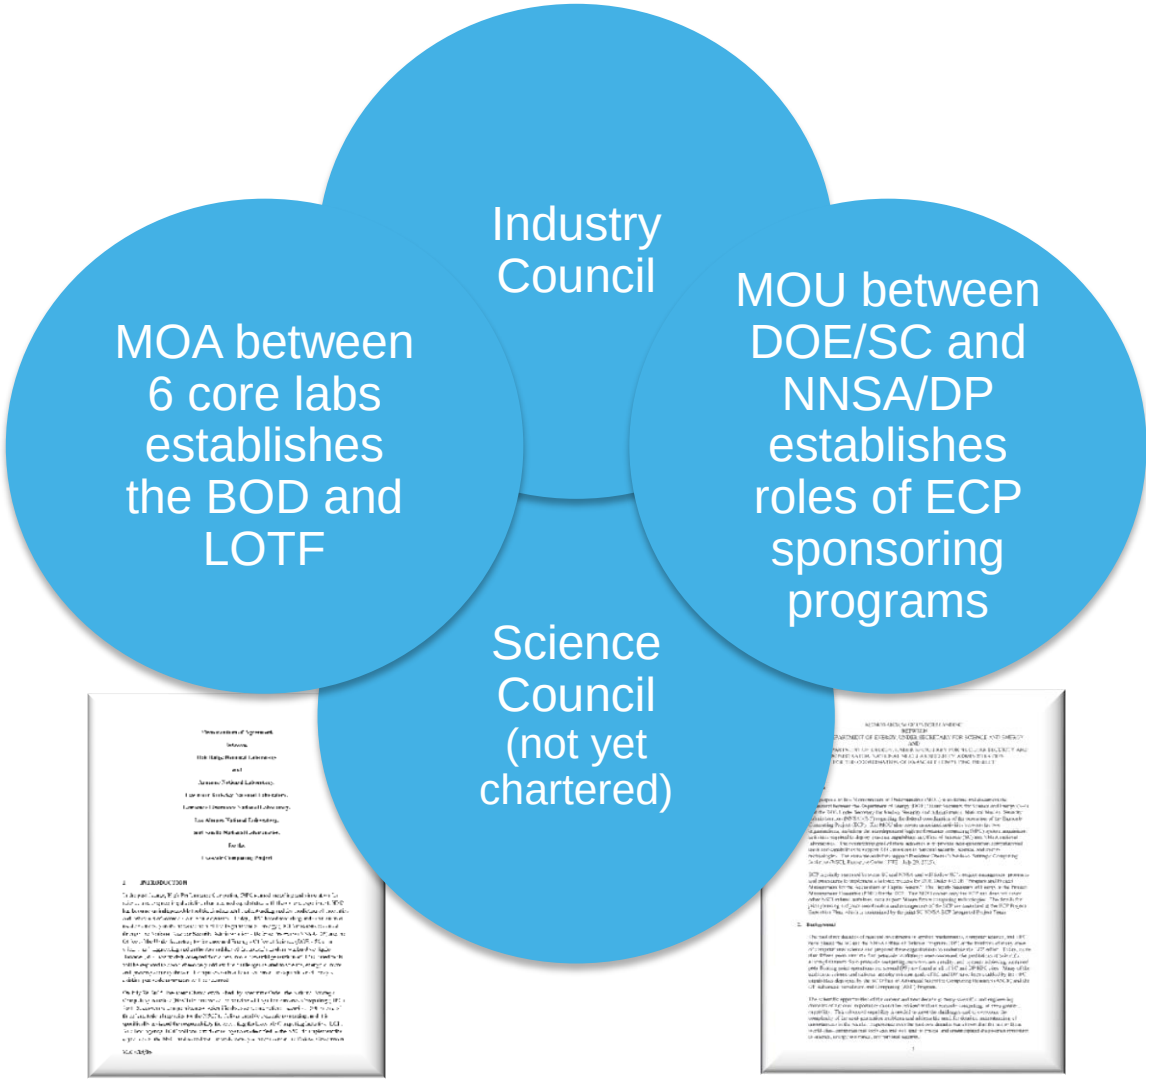

GOVERNANCE

ECP Board of Directors (BOD):

- Laboratory directors of 6 core national labs
- ECP Project Director reports to the BOD

Laboratory Operations Task Force (LOTF):

- ALDs from each of the 6 core laboratories
- Meets with ECP project director and deputy weekly to discuss staffing, operational, and management issues

Governance model has proven to be effective:

- BOD provides oversight and advice to ECP Director, and high-level interface with stakeholders
- LOTF provides engaged and effective leadership, ensuring core partner labs work together and cooperatively
- DOE/NNSA provide clear programmatic guidance across office boundaries

ADVICE

- **Industry Council** provides communication and information exchange between ECP and industries reliant on computer and computational science
- **Science Council** will examine ECP scientific output and ECP plans from a science and engineering quality perspective
- Both bodies will examine ECP results, processes, and activities and provide respected and useful experience-based advice to ECP management

ECP Integrated Project Team (IPT)

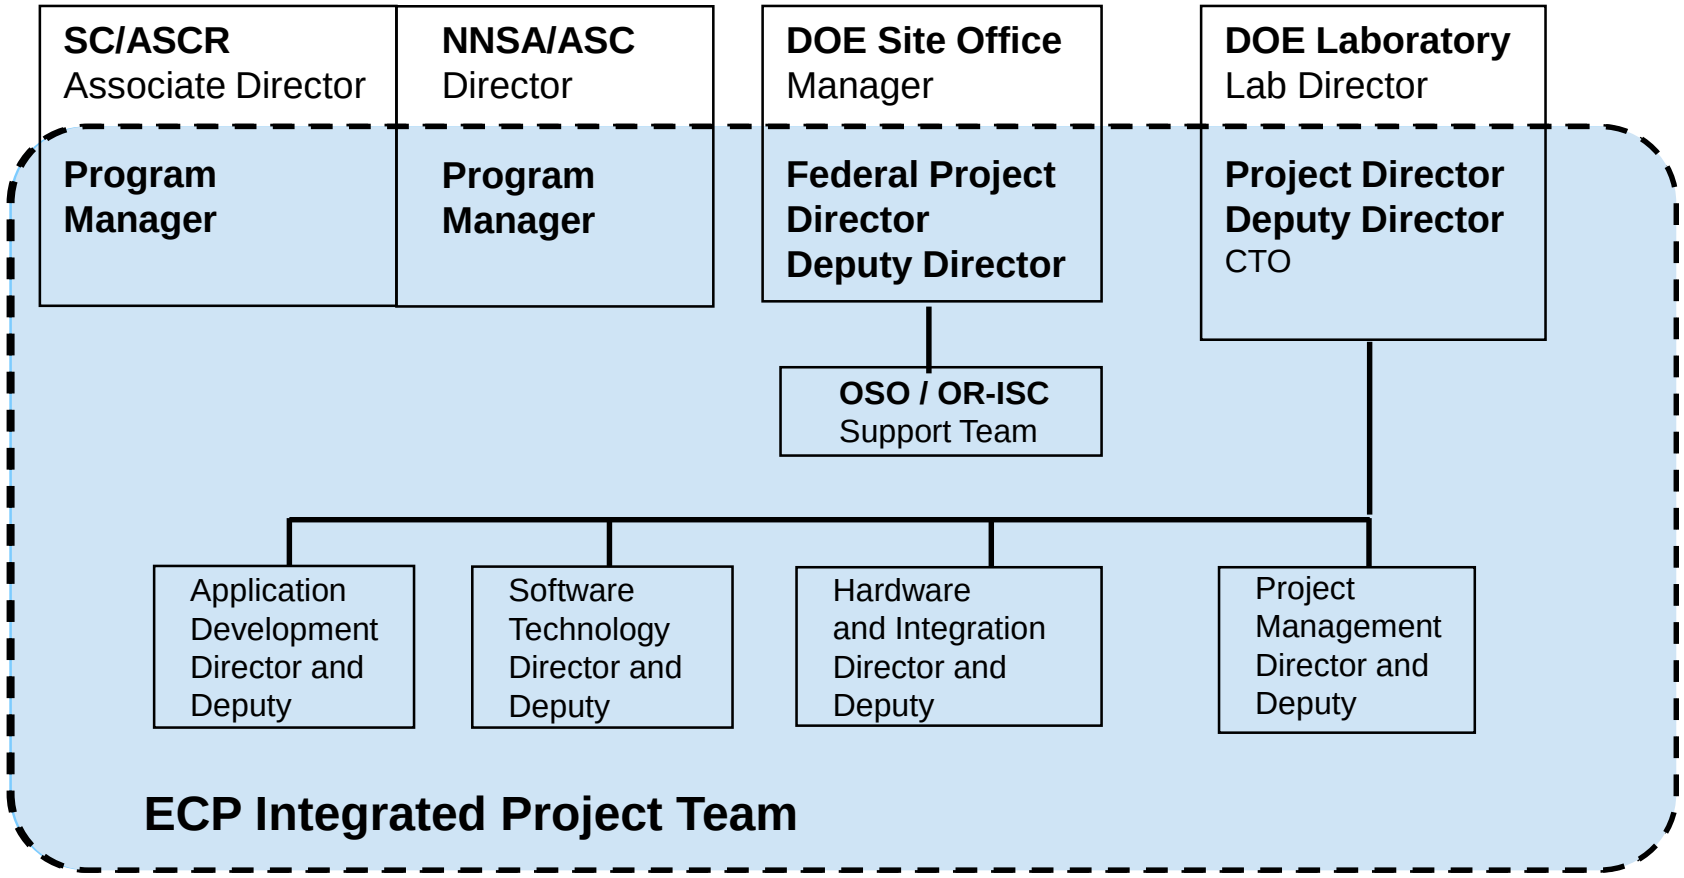

- **IPT Members**: DOE program managers, Federal Project Director (FPD) and FPD Deputy, members from the ORNL site office, Oak Ridge Integrated Support Center (OR-ISC), Project Director (PD) and PD Deputy, CTO, focus area directors and deputies, and project office staff (when needed)
- **Mission**: Provide strategic planning, coordination, and communication for the ECP to ensure that the project's objectives are achieved on schedule, within budget, and consistent with quality standards and DOE guidance

ECP EXASCALE COMPUTING PROJECT

ECP WBS

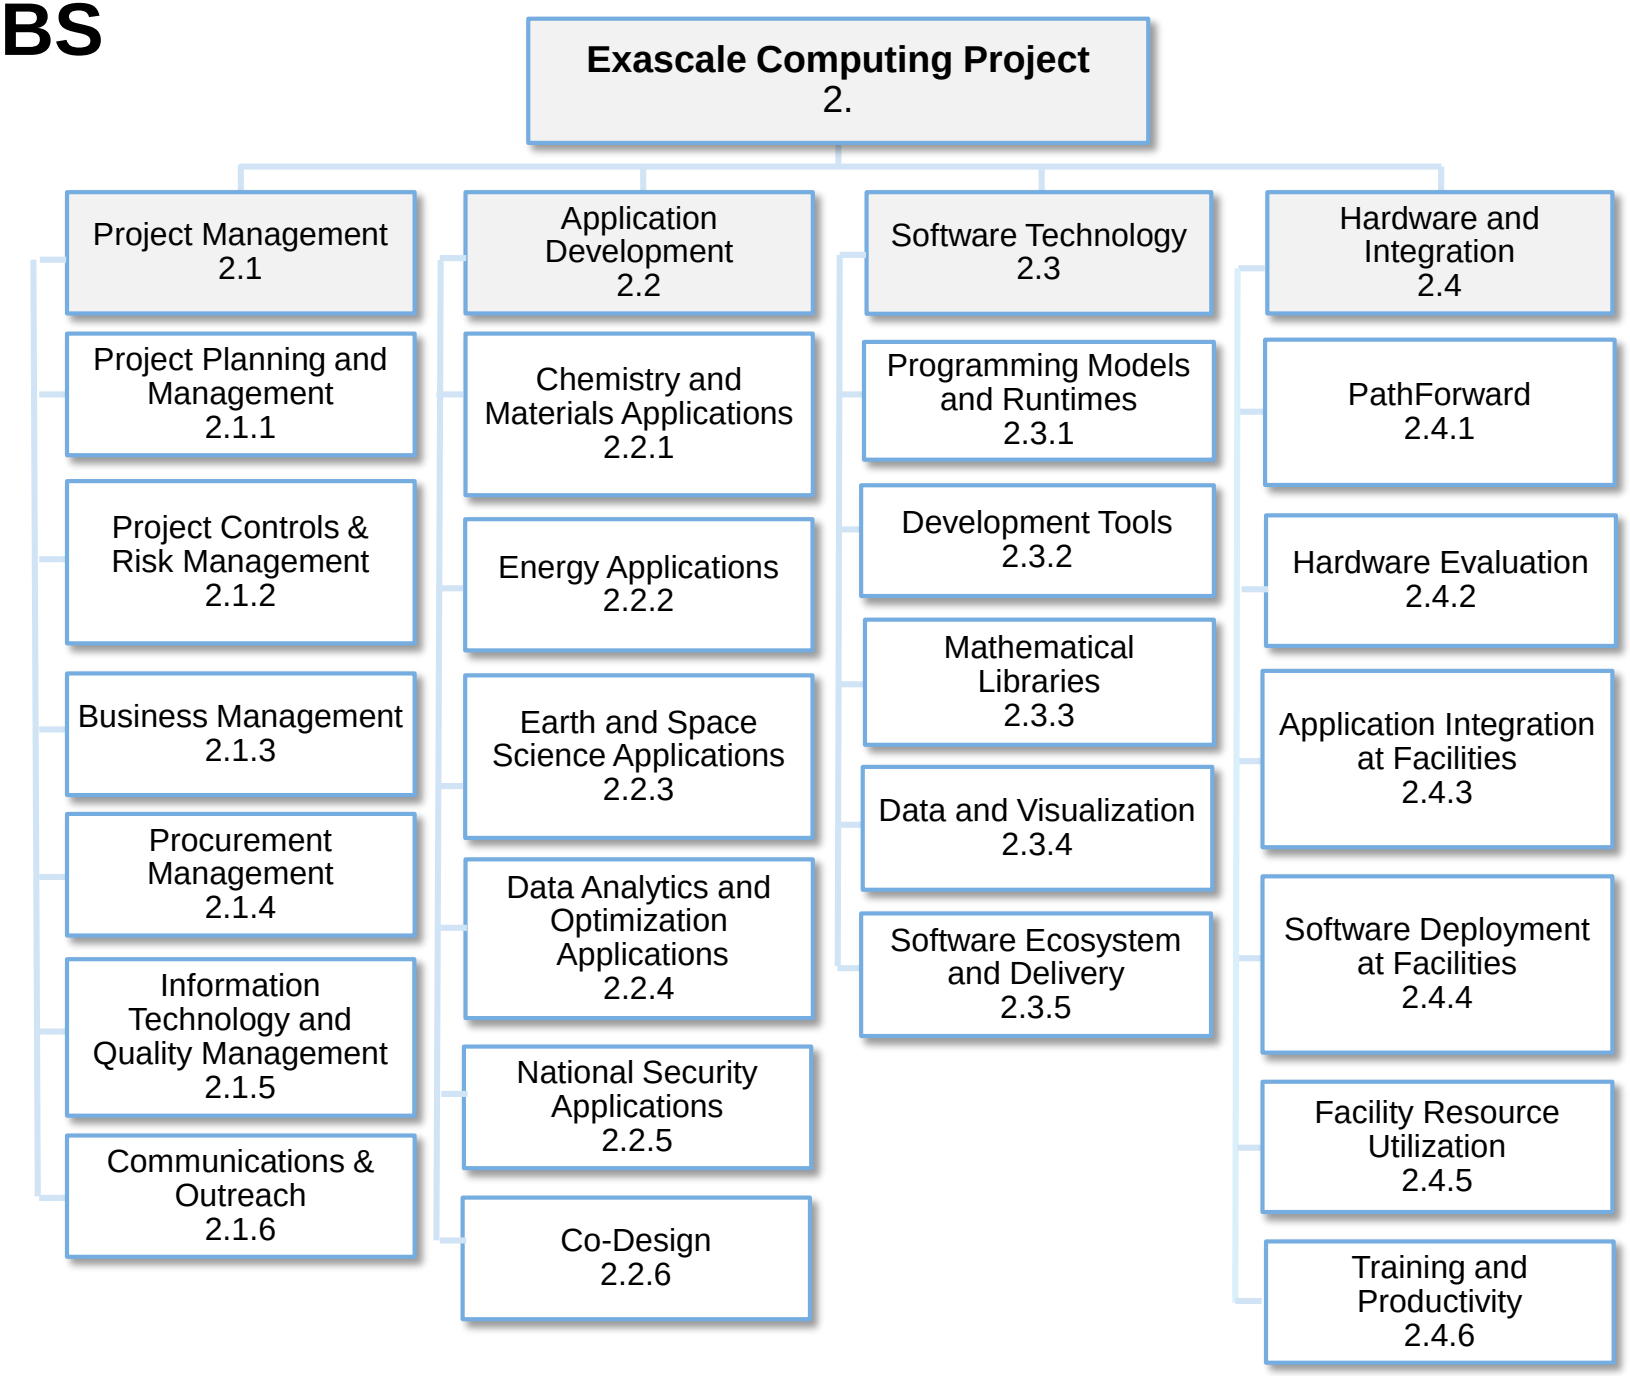
Exascale Computing Project
2.
Project Management
2.1
Project Planning and Management
2.1.1
Project Controls & Risk Management
2.1.2
Business Management
2.1.3
Procurement Management
2.1.4
Information Technology and Quality Management
2.1.5
Communications & Outreach
2.1.6
Application Development
2.2
Chemistry 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
Software Technology
2.3
Programming Models and Runtimes
2.3.1
Development Tools
2.3.2
Mathematical Libraries
2.3.3
Data and Visualization
2.3.4
Software Ecosystem and Delivery
2.3.5
Hardware and Integration
2.4
PathForward
2.4.1
Hardware Evaluation
2.4.2
Application Integration at Facilities
2.4.3
Software Deployment at Facilities
2.4.4
Facility Resource Utilization
2.4.5
Training and Productivity
2.4.6

ECP's Independent Project Review: Key Recommendations

From Jan 9-11, 2018 review

- Revise the ST gap analysis document regularly, at least annually.
- Revise initial ST Level 4 milestones by third quarter FY 2018 to remove vagueness and improve relevance.
- ECP and facilities need to work on an ongoing basis to achieve earliest possible access to test-beds, timelines, and information on chip/system architecture features.
- Add a HI and facilities breakout session to future independent project reviews including members of all the facilities in use by ECP.
- Expedite NDAs for the projects so AD and ST understand the hardware challenges and implications they will be facing early on.
- Refine and finalize the KPPs (including Figures of Merit and impact metrics) as soon as possible prior to CD-2.

ECP Funding thru FY18 and Plans for FY19

The FY18 budget was approved by Congress in March.

The President's Budget Request (PBR) for FY19 is public.

	FY-16 actual	FY-17 actual	FY-18 actual	FY-19 PBR
Total ECP Funding	**221**	**249**	**334**	**300**
ASCR contribution	157	164	205	232
ASC contribution	64	85	129	68

FY-18 looks good with a slight increase over ASCR's planned $197M and ASC contributing the planned amounts to AD and ST plus $54M rather than the planned $40M to PathForward.

FY-19 contains challenges, most notable is a potential large drop in ASC support.

The PBR specifies that ASC contribute $30M to AD, $35M to ST, and $0 to HI. The drop significantly affects PathForward and also reduces ST ATDM projects

Execution is ECP's highest FY18 priority

- Execution of our plan of record – FY18 budget enables the vigorous execution of our plan
  - The ECP plan of record includes additional scope and funding for some applications (seed and non-seed)
- Complete detailed joint engagement and execution plans with HPC Facilities
  - e.g., alignment of PathForward projects with NRE contracts once placed
- Evaluate additional strategic scope to enhance ECP's long-term impact
  - Possible addition of data science scope (e.g., cyber) to the electrical grid application
  - Possible addition of a data science/artificial intelligence co-design center
- Advancing project management toward CD-2
  - Complete implementation of project dashboards for efficient project status and management
  - Complete clear plan to achieve CD-2 (in FY19) – Next IPR late fall, 2018

In addition to execution, our focus is on rapid evaluation, decision, and pushing out increased funding for added scope

ECP Application Development (AD)

Chemistry and Materials Applications	Describes underlying properties of matter needed to optimize and control the design of new materials and energy technologies
Energy Applications	Model and simulation of existing and future technologies for the efficient and responsible production of energy to meet the growing needs of the U.S.
Earth and Space Science Applications	Spans fundamental scientific questions from the origin of the universe and chemical elements to planetary processes and interactions affecting life and longevity
Data Analytics and Optimization Applications	Applications partially based on modern data analysis and machine learning techniques rather than strictly on approximate solutions to equations that state fundamental physical principles or reduced semi-empirical models
National Security Applications	Stewardship of the US nuclear stockpile and assessment of future threats; related physics and engineering modeling and scientific inquiries consistent with that mission space
Co-Design	Focused on crosscutting algorithmic methods that capture the most common patterns of computation and communication in ECP applications

ECP Applications must Deliver Useful & Impactful Science

1. **Deliver improved and impactful science & engineering (*performance*)**
   - New or enhanced predictability on a problem of national importance (the “challenge problem”) measured in part by the application’s rate of executing useful work (W/t) on the CP
2. **As performance portable as possible and reasonable (*portability*)**
   - No “boutique” one-off applications able to only execute on one (and likely ephemeral) system
3. **Able to make effective use of a capable system (*readiness*)**
   - *Effective* is app specific (weak, strong, ensembles, single-node performance)
4. **Able to integrate latest relevant software technologies (*modern*)**
   - Needed to demonstrate agility, flexibility, modern architecture; overall app portfolio must apply pressure to all key attributes of the system design characteristics
5. **High priority (*strategic*)**
   - Key stakeholders care about using application to make consequential decisions

Selected FY18 AD Highlights

Release

- CEED team released first software distribution (CEED 1.0), consisting of 12 integrated Spack packages for libCEED, mfem, nek5000, nekcem, laghos, nekbone, hpgmg, occa, magma, gslib, petsc and pumi, plus a new CEED “meta-package” (Milestone CEED-13).

CEED
EXASCALE DISCRETIZATIONS

Performance

- ExaSky ported the HACC code to Summit, and has already achieved excellent strong scaling and a 30-35X speedup over the CPU.

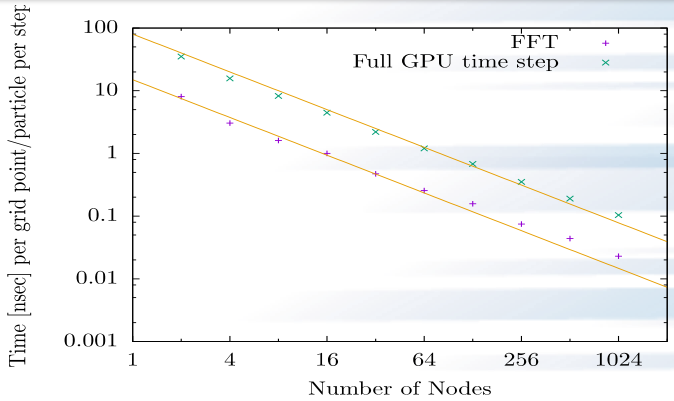

Weak scaling results for HACC on Summit Phase I up to 1024 nodes; straight lines represent ideal scaling. FFT optimization for Summit has not been carried out as yet.

Algorithm or Physical Model Enhancement

- WDMApp team completed coupling of GENE and XGC, first time a continuum core gyro-kinetic code (GENE) coupled to an edge gyro-kinetic code (XGC) based on the particle-in-cell (PIC) method, enabling a complete kinetic description of the entire plasma.

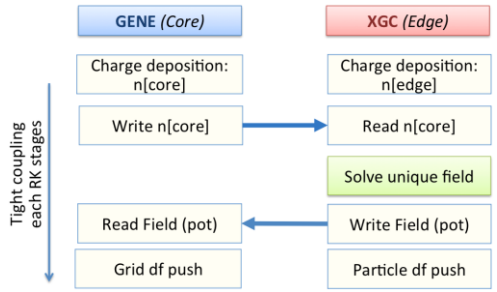

Science Capability

- E3SM-MMF successfully completed the milestone on GPU acceleration of ocean, and also completed the initial port of CRM to GPU.

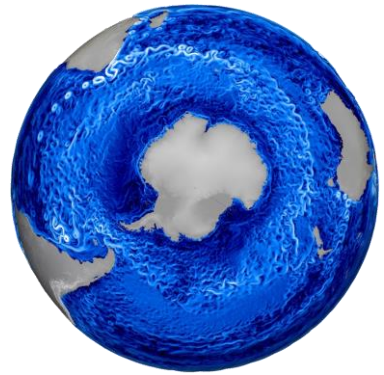

ECP EXASCALE COMPUTING PROJECT

ECP Software: Productive, sustainable ecosystem

Goal

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

- Extend current technologies to exascale where possible

- Perform R&D required for new approaches when necessary

- Coordinate with and complement vendor efforts

- Develop and deploy high-quality and robust software products

56 WBS L4 subprojects executing RD&D

ECP SW Stack: Strategic Alignment & Synergies

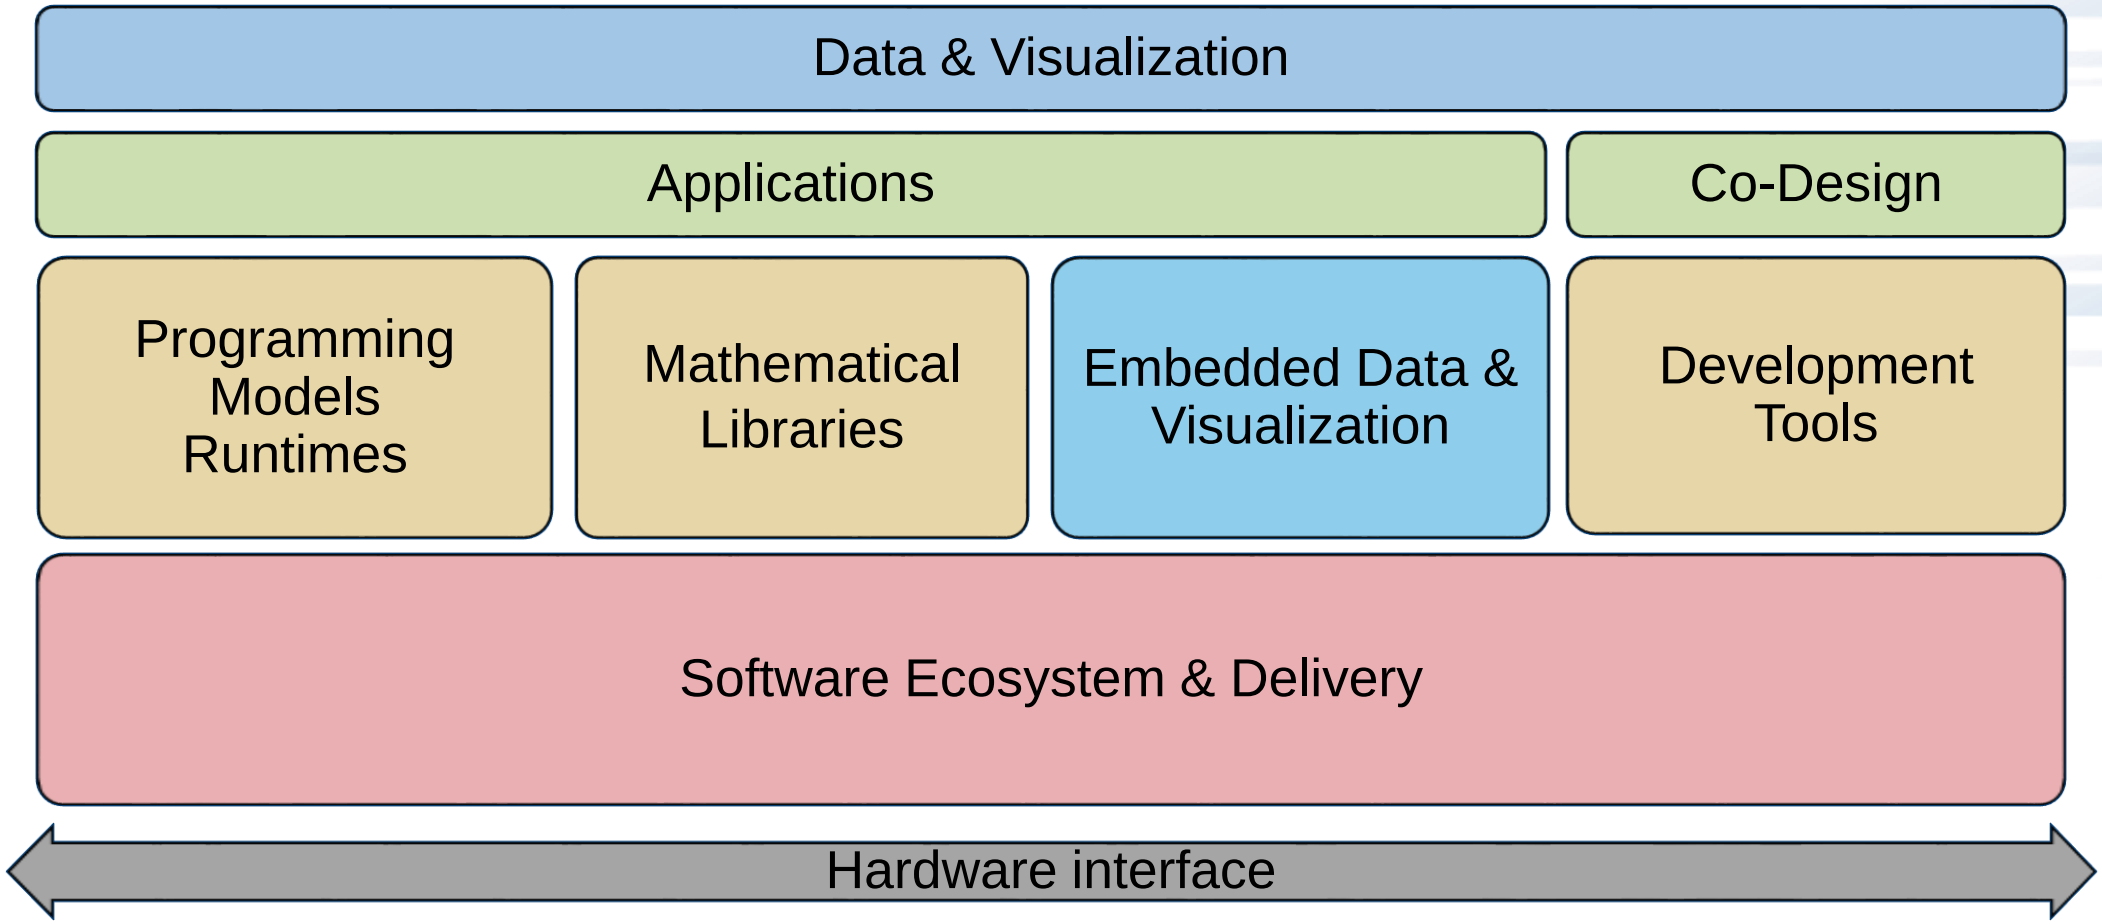

ECP EXASCALE COMPUTING PROJECT

Programming Models and Runtimes

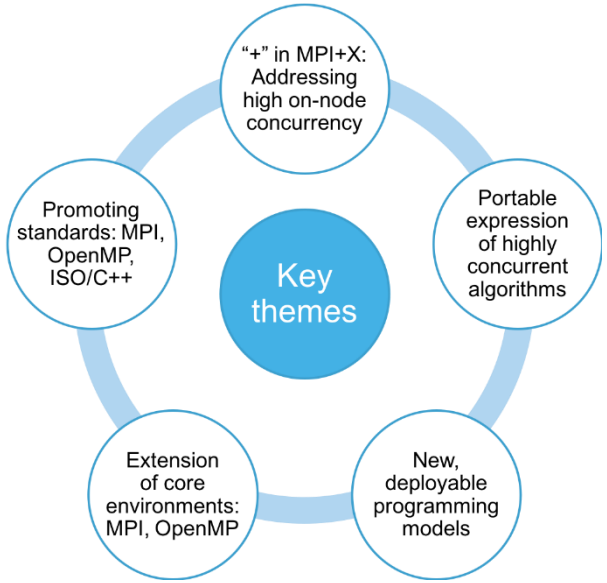

Development Tools

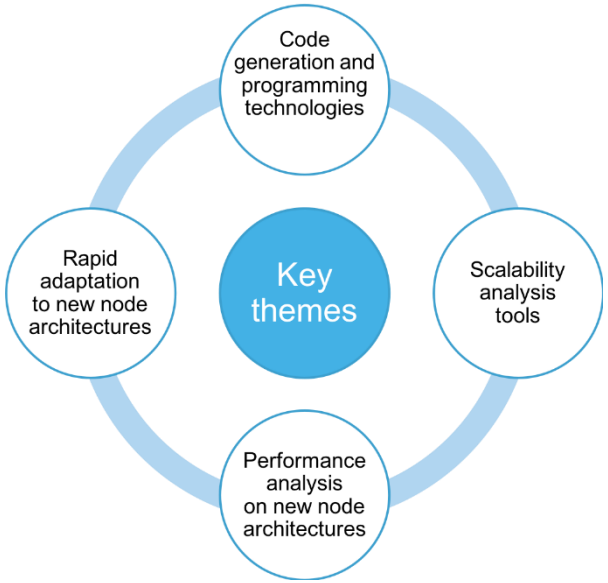

Math Libraries

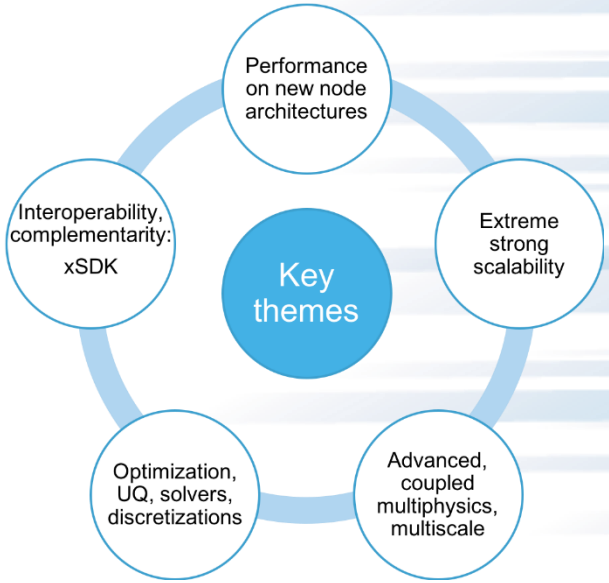

Data and Visualization

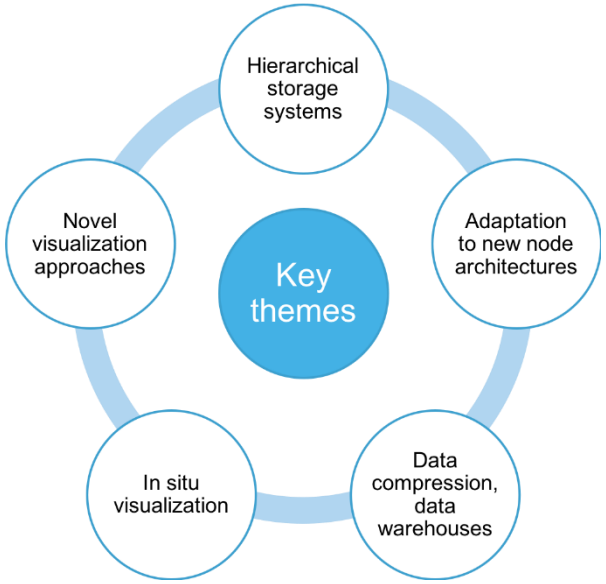

Software Ecosystem and Delivery

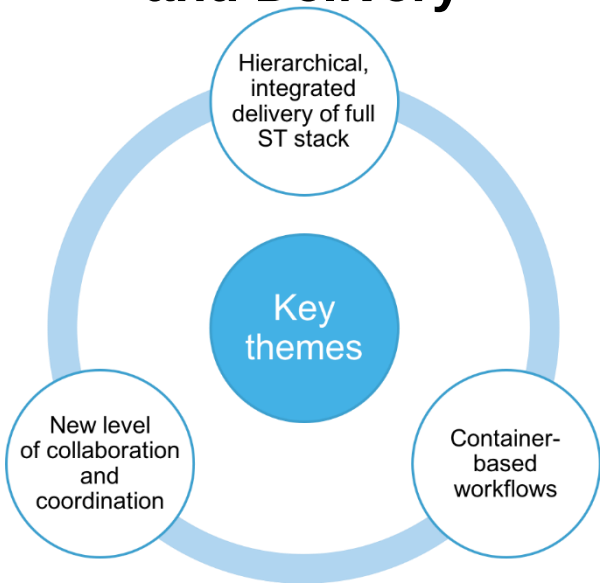

ECP EXASCALE COMPUTING PROJECT

Software Development Kits Progress: Leadership in place, Spack packaging making rapid progress

ECP software projects

Each project to define (potentially ≥2) release vectors

More projects ← → Fewer projects

SDKs

Reusable software libraries embedded in applications; cohesive/interdependent libraries released as sets modeled on xSDK

- Regular coordinated releases
- Hierarchical collection built on Spack
- Products may belong to >1 SDK based on dependences
- Establish community policies for library development
- Apply Continuous Integration and other robust testing practices

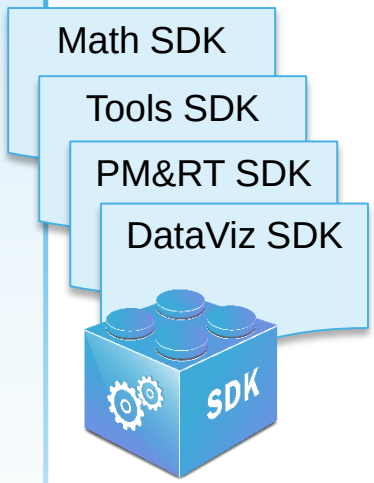

OpenHPC

Potential exit strategy for binary distributions

- Target similar software to existing OpenHPC stack
- Develop super-scalable release targeting higher end systems

Direct2Facility

Platform-specific software in support of a specified 2021–2023 exascale system

- Software **exclusively** supporting a specific platform
- System software, some tools and runtimes

94 – Number of Products ECP ST contributes to

(46, 23) – Number of Spack-ready products (Done, In progress)

SDK Leadership Team: Decades of Software Experience

- **Jim Willenbring – SDK Coordinator and Release Manager**
- **Sameer Shende – Programming Models & Runtimes**
- **Bart Miller – Development Tools**
- **Lois McInnes – Math Libraries**
- **Chuck Atkins - Data & Viz**

ECP EXASCALE COMPUTING PROJECT

Q2FY18 ECP ST Highlights

ECP ST Q2 Progress: 3 Examples

- Threaded execution of data compression: 16x time improvement.
  - Compression and threaded execution are essential.
- New OpenMP Runtime: 70% time overhead reduction.
  - Addresses the '+' of MPI+X, essential for Exascale.
- New Checkpoint/Restart: 3s offload (purple), scale independent.
  - Rapid CPR is the ECP resilience strategy

Data Compression Speedups

Speedup
32 cores: 14~16
64 cores: 15~17
Breakdown performance
Speedup: up to about 40

Overall performance

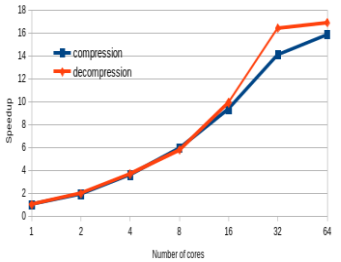

Performance breakdown

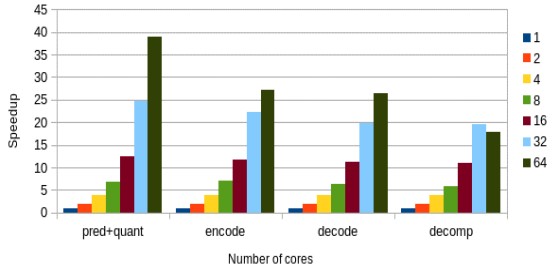

Checkpoint/Restart Progress

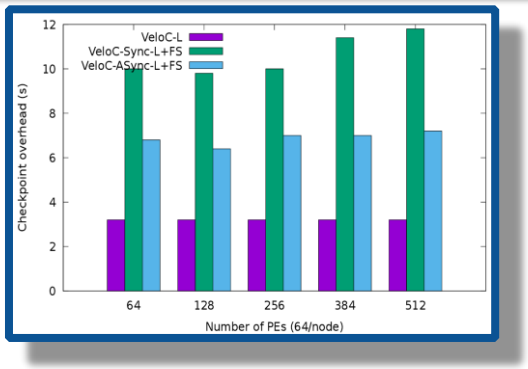

Preliminary experiments on Theta: Better scalability and performance for asynchronous vs synchronous mode

Joint ASCR-NNSA effort, bringing together capabilities from two existing projects for better collaboration and products.

MPI+OpenMP Overhead Progress

- H: # of threads for horizontal parallelism
- V: # of threads for vertical parallelism
- Evaluated on a 16-core machine
- 2 x Intel Xeon CPU E5-2670
- Demonstration for vendor adoption

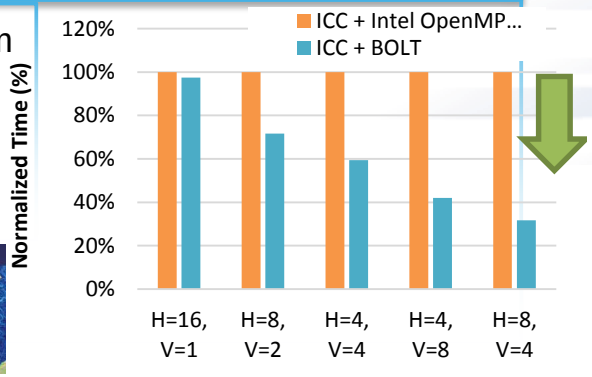

Software Available from all projects:
SZ Compression: New release (1.4.13.0) can be downloaded from github (https://github.com/disheng222/SZ/releases/tag/v1.4.13.0).
OpenMP: BOLT v1.0b1 and v1.0a1 (http://www.bolt-omp.org/downloads) and Argobots v1.0b1 and v1.0a1 (http://www.argobots.org/downloads)
CPR: Veloc Backend software (xgitlab.cels.anl.gov/ecp-veloc)

ECP EXASCALE COMPUTING PROJECT

Hardware Integration: Designed to Enable Integration of ECP's products into HPC Environments at the Facilities

ECP will meet its objectives on Facility resources

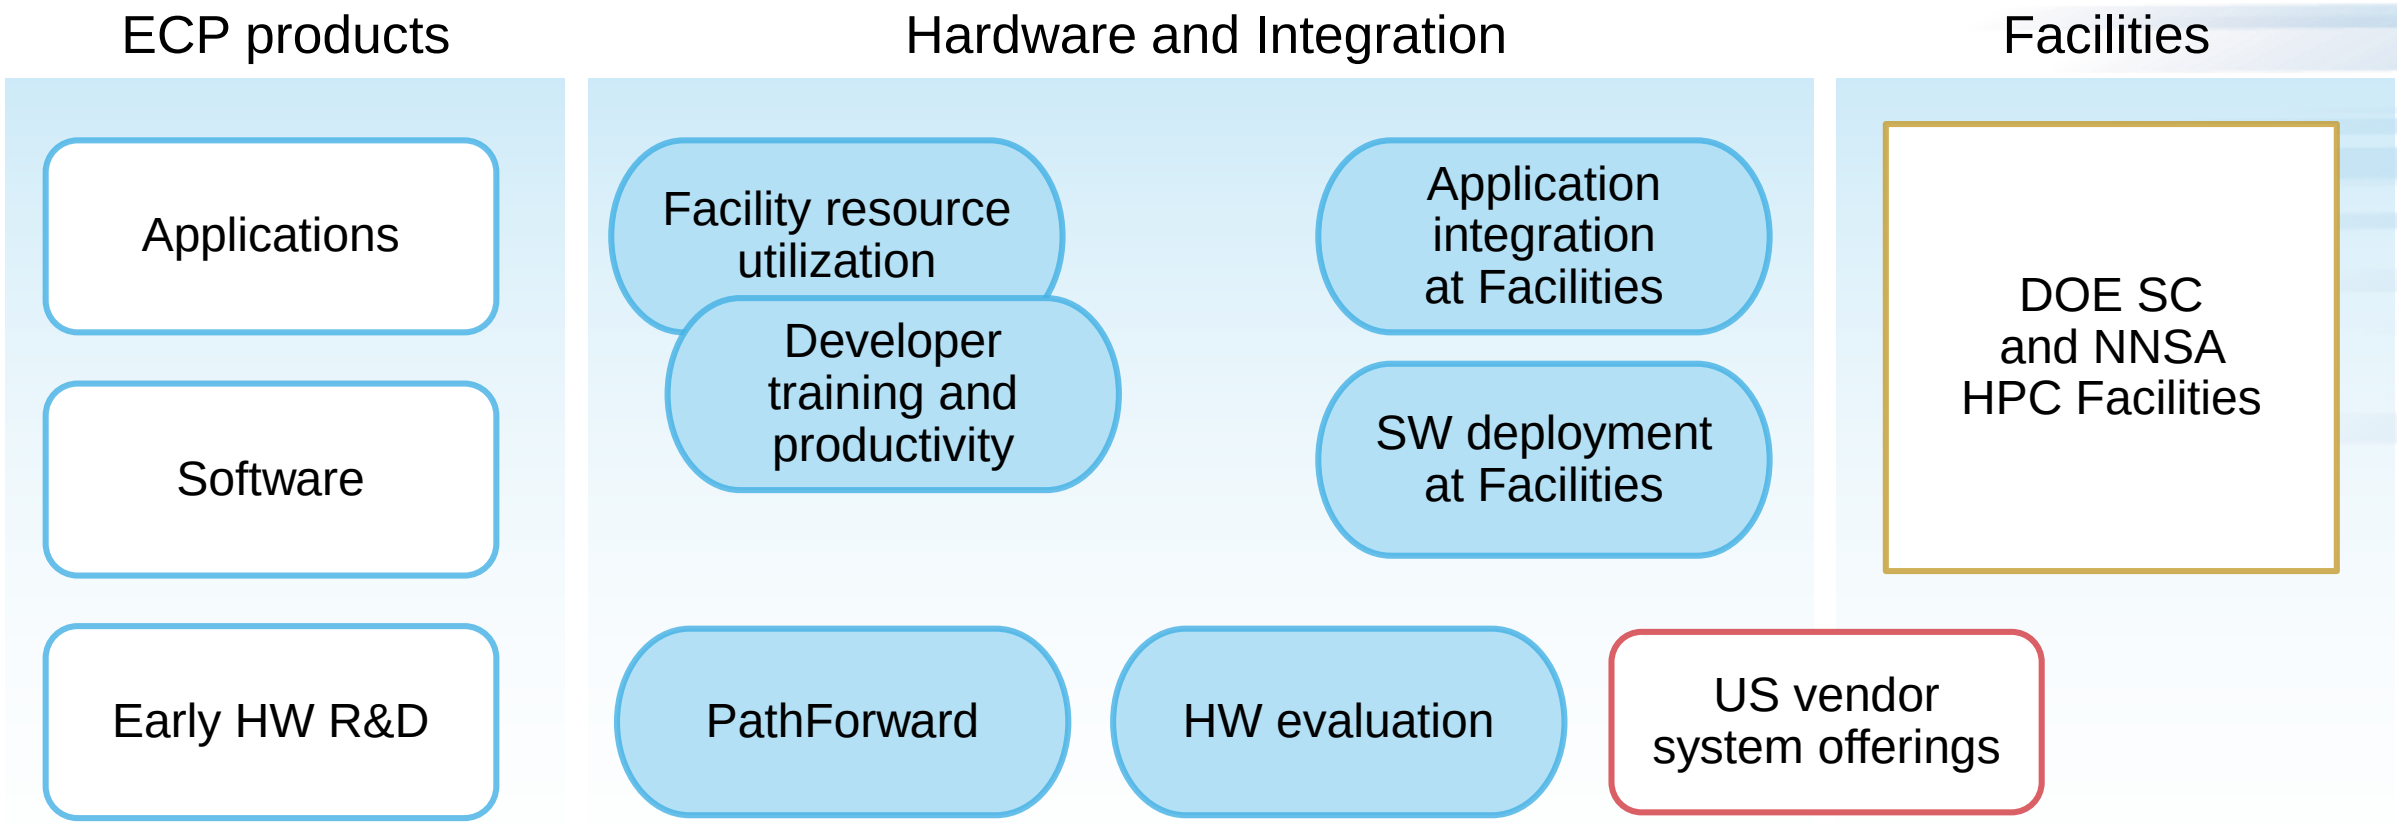

HI project	• FY18 Q1 and Q2 Accomplishments
Application Integration at Facilities	• Developed a plan jointly with the Facilities to prepare and deploy ECP applications on the Facilities' pre-exascale and exascale systems • Prepared paperwork to send funds to the Facilities for this project
Software Deployment at Facilities	• Developed a plan jointly with the Facilities to deploy ECP software on the Facilities' pre-exascale and exascale systems • Prepared paperwork to send funds to the Facilities for this project • Collaborated with ST on awarding and managing a contract to develop a secure automated testing capability
Training and Productivity	• Developed and hosted multiple training and educational events • Awarded the 2018 Better Scientific Software (BSSw) Fellows • Made available on ECP's public website recordings of training events and material

HI project	• FY18 Q1 and Q2 Accomplishments
Facility Resource Allocation	• Stood up the ECP Resource Allocation Council (RAS) • Allocated, managed, and reported on ECP's initial computer allocations • Drafted a user program for managing future computer allocations
PathForward	• 2nd review held March 2018 – all 6 contractors reported progress to be reviewed by a team of Lab/Gov't experts • ~25% of milestones completed; another ~25% in progress • PF research efforts are expected to significantly enhance the CORAL2 bids
Hardware Evaluation	• Published first version of an Abstract Machine Model document to be used as a aid for understanding key architectural features without revealing proprietary information • Completed a memory performance study

HI has three main challenges: 1) time pressures 2) developing and maintaining strong partnerships, and 3) standing up HI two years into ECP

- Shortened timeline: less time for HW innovations, app and SW development, facilities preparation and acquisition
- Integration with 6 facilities: building strong collaborations, aligning priorities and schedules
- HI late start: reduced time to achieve objectives, budgetary constraints
- Integration with AD and ST: many projects, connecting them with the facilities, building collaboration between many stakeholders

HPC and Scientific Community

- GitLab
- openHPC

Communication and Release

- Workshops
- Conferences
- Publications
- …

Communication and Release

Software Projects

ADIOS
ATDM
LLVM
Kokkos
RAJA
Legion
Trilinos
.
.
.

Contribution Complies w/ SDK Specifications

SDKs

SDK 1
SDK 2
SDK 3
SDK …

Software Development Kits (SDKs)
(delivery specs and mechanism)

APIs

Applications

Integration of ST products via SDK

Apps integration

Facilities

A21
Frontier
El Capitan
Pre-Exascale

S/W R&D

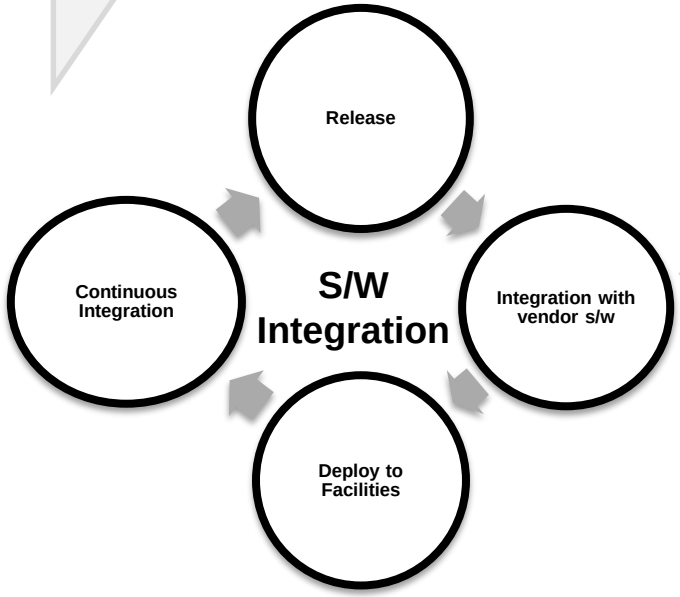

S/W deployment

The ECP Exascale Computing Ecosystem

ECP EXASCALE COMPUTING PROJECT

Dashboard Being Developed to Track Progress

Attributes currently being prototyped – ECP will evolve what we learn is most useful

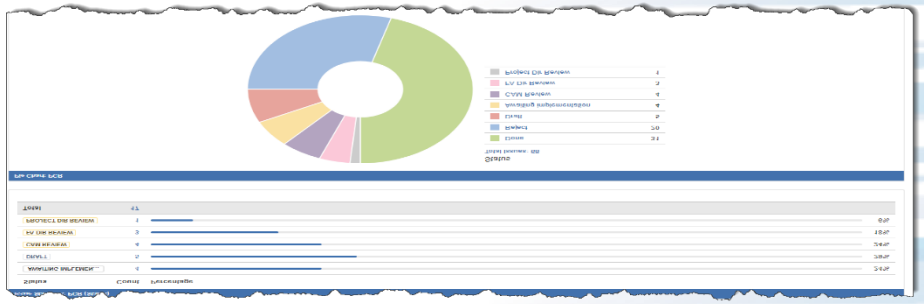

- **Applications**
  - Mission (priority, uniqueness)
  - Exascale challenge problem status (on track, caution, trouble)
  - Figure of Merit (FOM) status (ratio of required challenge problem work rate / current work rate)
  - Node- and system-level performance and scalability (relative to baseline)
  - Readiness and portability (facility platforms supported & specs associated with successful execution)
  - Facility commitment (readiness, acceptance, early science)
  - Ecosystem Requirements (coverage of motifs & ST products)
- **Software Technologies**
  - Utilization: by applications, facilities, and overall users
  - Continuous integration results: regression tests, performance, integrity, scaling
  - Impact goal status (on track, caution, trouble)
- **Hardware and Integration**
  - PathForward milestone deliverables (how many, mapping to exascale hardware challenges)
  - Status of application and software deployment at DOE HPC Facilities
  - EPC application and software HPC resource usage at DOE HPC Facilities

ECP Risks and Challenges

- **Cost/funding issues**
  - Accept (track funding issues; implement flexible contracts; design for agility)
- **Exascale platform meeting requirements unavailable to ECP**
  - Mitigate (impart 12 months of schedule contingency, re-baseline)
- **Unclear messaging to stakeholders**
  - Mitigate (develop common talking points and training)
- **Standard programming models are insufficient for applications**
  - Mitigate (formal constraints on s/w ecosystem delivered by vendors, regular assessments)
- **PathForward funding does not align with costs**
  - Mitigate (realign funding commitments to profiles, re-negotiate vendor milestones)
- **Relationship between ECP and Facilities**
  - Mitigate (facility engagement plans, track interactions and facility utilization)

ECP Status: Summary

- Great team, great plan, executing full bore
- Great stakeholder support and high priority to DOE
- Moving into a product development stage with more quantitative metrics
- Actively managing our highest risks
- Putting facility engagement plans into action
- Still have work to do: tracking progress (dashboard), maintaining agility for highest ROI on RD&D, adjusting plan to ever-changing budget and cost profiles

Thank you!

ECP Governance: DOE Program & Lab Agreements

- MOU between DOE/SC and NNSA/DP to define roles & responsibilities of Programs responsible for ECI & ECP
- MOA between six national labs (“core partners”) responsible for leading and managing ECP: ANL, LANL, LBNL, LLNL, ORNL, SNL
- ECP Board of Directors serves an advisory, oversight, line-management role for ECP
  - Advise Project Director (PD) on strategic direction and performance goals; evaluate PD performance on an annual or as-needed basis; support ECP SLT in managing interfaces with key stakeholders
  - Comprised of the laboratory directors from each core partner laboratory and 2-3 external members
  - Operations: establish a standing Laboratory Operations Task Force (LOTF) comprised of one delegate from each core partner laboratory with management responsibility at their laboratory for the resources needed for ECP to be successful
    - Rick Stevens (ANL); John Sarrao (LANL); Kathy Yelick (LBNL)
    - Mike McCoy (LLNL); Jeff Nichols (ORNL); Scott Collis (SNL)
  - LOTF currently meets weekly with the PD to discuss staffing, operational and management issues

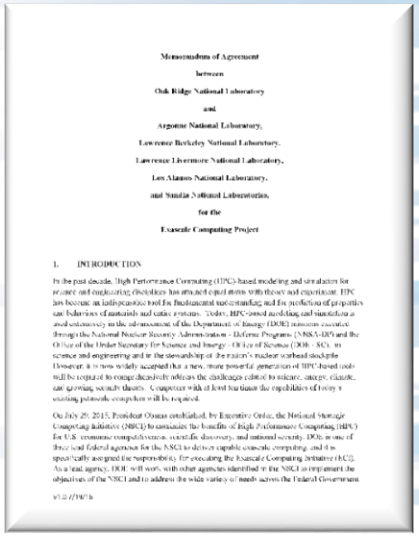
Memorandum of Agreement
between
Oak Ridge National Laboratory
and
Argonne National Laboratory,
Lawrence Berkeley National Laboratory,
Lawrence Livermore National Laboratory,
Los Alamos National Laboratory,
and Sandia National Laboratories,
for the
Exascale Computing Project

1. INTRODUCTION

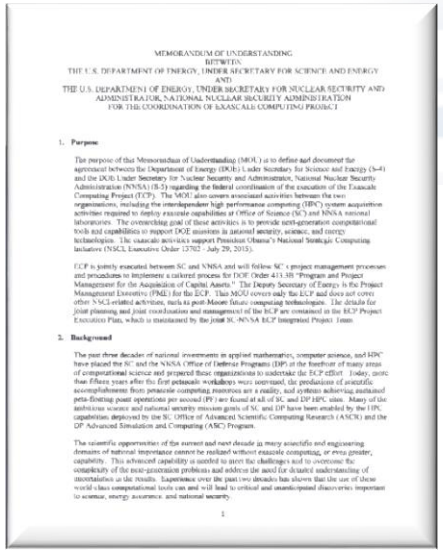
MEMORANDUM OF UNDERSTANDING
BETWEEN
THE U.S. DEPARTMENT OF ENERGY, UNDER SECRETARY FOR SCIENCE AND ENERGY
AND
THE U.S. DEPARTMENT OF ENERGY, UNDER SECRETARY FOR NUCLEAR SECURITY AND ADMINISTRATOR, NATIONAL NUCLEAR SECURITY ADMINISTRATION
FOR THE COORDINATION OF EXASCALE COMPUTING PROJECT

1. Purpose

2. Background

Governance model has proven to be effective. LOTF provides engaged and effective leadership to ECP and in keeping ECP core partner labs working together consistently and cooperatively.

ECP Advisory Bodies: Industry and Science Councils

- Industry Council (meetings Mar & Oct 2017, teleconference Jun 2017)
  - Provides two-way communication and information exchange between ECP and private sector software and hardware companies and broader engagement with the computational science and HPC user community from an industry perspective
  - Member companies represent computer and computing technology, design and engineering, consumers of scientific computing resources as part of their business activities, and ISVs
  - Current activity: provide ECP with industry exascale challenge problem drivers and requirements
  - Future activities: participate in red team reviews, provide feedback on milestone reports and accomplishments, provide feedback on project office's approach to incorporating agile with formal PM
- Science Council (not yet chartered)
  - SC will be asked to examine and provide technical feedback and advice to ECP leadership on the plan for and quality of the science, engineering, and technology achievements and milestone deliverables
  - SC will be composed of leading computer scientists, computational scientists, and applied mathematicians (e.g., NAS, NAE members)

ECP will benefit in FY18 and beyond from increased external SC/IC advice on its plans and milestone deliverables