EXASCALE COMPUTING PROJECT

Exascale Computing Project Update

Douglas B. Kothe (ORNL), Director
Stephen Lee (LANL), Deputy Director

Presentation to the DOE Advanced Scientific Computing Advisory Committee (ASCAC)

December 20, 2017

The Exascale Computing Project in Context

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 is focused on significant mission-critical applications, including:
  - Stewardship of today's stockpile and the growing threat of potential adversaries defeating U.S. weapons
  - Development of clean energy systems, improving our infrastructure resilience, designing new materials, adapting to regional water cycle changes, demystifying the origin of elements in the universe, developing smaller and more powerful accelerators for use in medicine and industry, …
  - Discover new knowledge from data in powerful DOE experimental facilities, environmental genomes, cancer research, high-fidelity simulations
- ECP is necessary, but not sufficient – the broader necessary ECI elements include:
  - 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 Strategic Goals and Outcomes

- ***Goal: Applications***
  - Delivery of mission and science results using exascale systems enabled by ECP, each addressing an exascale challenge problem (problems of strategic importance and national interest, intractable without at least 50x computational power of today's systems)
- ***Goal: Software Technologies***
  - Deliver next generation DOE software capabilities for exascale applications and platforms, targeted to specific exascale systems as a high quality, sustainable product suite
- ***Goal: Hardware and Integration***
  - Integrated deployment of specific outcomes and products on targeted systems at computing Facilities, including the completion of PathForward projects transitioning to Facility NRE (where appropriate), and the integration of software and applications on pre-exascale and exascale system resources at Facilities
- ***Outcome: Accelerated delivery of a capable exascale computing ecosystem*** providing breakthrough solutions addressing our most critical challenges in scientific discovery, energy assurance, economic competitiveness, and national security
  - *Capable*: a wide range of applications can effectively use the systems enabled by ECP, ensuring that both science and security needs will be addressed (affordable, usable, useful)
  - *Exascale*: application performance at least 50x of today's systems
  - *Ecosystem*: all methods and tools needed for effective use of ECP-enabled exascale systems to be acquired by DOE labs (not just more powerful systems)

Activities Since Oct 1 2017

- Overall ECP design re-formulated and documented
  - Revised work breakdown structure (WBS)
  - New extended leadership team
    - 33 of the most experienced computational scientists, leaders, and project managers from your laboratories
  - Revised Key Performance Parameters (KPPs)
  - Revised Key Milestones (in Preliminary Project Execution Plan)
- Re-tooled our approach to project management processes and tracking
  - Devised a new approach to tracking progress (cost/schedule performance; dashboard)
  - Revamped resource loaded schedule
  - Fresh and close analysis of risks and risk management
  - Critical path analysis
- Plans for closer integration with HPC Facilities
  - A much more direct, mutually beneficial, and necessary tie for success

ECP has multiple necessary components to meet national goals

Application Development	Software Technology	Hardware and Integration	Project Management
Develop and enhance the predictive capability of applications critical to the DOE	Produce expanded and vertically integrated software stack to achieve full potential of exascale computing	Integrated delivery of ECP products on targeted systems at leading DOE computing facilities	Project scheduling, tracking cost tracking, analysis, and accountability across the entire project spread across the country
Andrew Siegel (ANL) **Erik Draeger (LLNL)**	**Mike Heroux (SNL)** **Jonathan Carter (LBNL)**	**Terri Quinn (LLNL)** **Susan Coghlan (ANL)**	**Kathlyn Boudwin (ORNL)** **Manuel Vigil (LANL)**

EXASCALE COMPUTING PROJECT

ECP WBS

Exascale Computing Project
2.0

- **Project Management 2.1**
  - Project Planning and Management 2.1.1
  - Project Controls and Risk Management 2.1.2
  - Business Management 2.1.3
  - Procurement Management 2.1.4
  - Information Technology and Quality Management 2.1.5
  - Communications and 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
  - Math 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 Structure and Leadership Addresses Evolving Requirements

Possessing the requisite skills, experience and resources for success

- Structure: changes to the WBS structure (Oct 2017) needed to increase ECP focus on *product delivery* and "steady state" execution post-startup
  - **AD** changed from programmatic (e.g., "DOE Science & Energy") to domain-based structure (e.g., "Chemistry and Materials") to enable closer and more effective leadership and project management by domain science experts.
  - **ST** structure consolidated and streamlined, with aggressive movement into a critical *product development* stage after its initial R&D stage and line of sight of ST products to applications.
  - **HI** [formerly Hardware Technology (HT)] required expanded scope to more proactively and directly integrate - *including a formal handoff of ECP products and technologies*- with DOE HPC facilities.
- Leadership: size and complexity of ECP warranted a more empowered extended leadership team (ELT) with explicit roles, responsibilities, authorities, accountabilities (ECP Management Plan)
  - **L2 leads**: strategic thinkers (integrating their technical area within the larger US exascale ecosystem); tactical planners (relative to project management): influential (relative to the technical community and also within ECP); able to represent the full ECP scope in a compelling manner
  - **L3 leads**: technical leaders for sub-projects within their element, providing technical oversight and management; also serve as Control Account Managers (CAM) for their element
  - **ELT** (SLT+L2s+L3s): now involved in most leadership discussions, with SLT decisions based on informed-input from ELT

ECP Senior Leadership Team

Doug Kothe, Director
Doug is a leader in the development of computational science applications and formerly the Director of the CASL DOE Hub. He most recently led the ECP Application Development focus area, and is a veteran of DOE with 33 years of service at LANL (20), LLNL (1), and presently ORNL (12).

Stephen Lee, Deputy Director
Stephen is a computational physicist and a long time leader at LANL, with 30 years of experience in computational physics and code development. He was most recently the Division Leader for the Computer, Computational, and Statistical Sciences Division at LANL.

Kathlyn Boudwin, Project Management Director
Kathlyn has 25 years of experience working on large scientific and HPC projects for DOE and NSF. She was previously the Deputy Project Director for OLCF, Associate Project Director at SNS, and Project Controller for APS and Argonne Guest House at ANL. She holds both an MBA and a degree in economics.

Manuel Vigil, Project Management Deputy Director
Manuel most recently served as the Trinity Project Director for the acquisition and project management of the Trinity Advanced Technology System and as the Platforms Program Manager for LANL's Advanced Simulation and Computing Program. Manuel has over 40 years of experience in the acquisition and management of HPC systems for scientific and national security missions.

Andrew Siegel, Application Development Director
Andrew is a Senior Scientist at ANL with appointments in the Mathematics and Computer Science and Nuclear Engineering Divisions. For the past decade he has led ANL's program in simulation of advanced nuclear reactors. He has over 20 years of experience developing applications and using HPC systems to study problems in a broad range of disciplines including nuclear engineering, atmospheric science, astrophysics, and oceanography.

Erik Draeger, Application Development Deputy Director
Erik is the group leader of the High Performance Computing group in the Center for Applied Scientific Computing at LLNL, with over 15 years of experience developing and optimizing highly-scalable applications for physics and biology.

ECP Senior Leadership Team

Mike Heroux, Software Technology Director
Mike has been involved in scientific software R&D for 30 years. His first 10 were at Cray in the LIBSCI and scalable apps groups. At Sandia he started the Trilinos and Mantevo projects, is author of the HPCG benchmark for TOP500, and leads productivity and sustainability efforts for DOE.

Jonathan Carter, Software Technology Deputy Director
Jonathan has been involved in the support and development of HPC applications for chemistry, the procurement of HPC systems, and the evaluation of novel computing hardware for over 25 years. He currently a senior manager in Computing Sciences at Berkeley Lab.

Terri Quinn, Hardware and Integration Director
Terri is a long-time leader in high performance computing at LLNL, with over 20 years as a program manager for the Advanced Simulation and Computing program. She is the Associate Program Director for Livermore Computing Systems and Environment.

Susan Coghlan, Hardware and Integration Deputy Director
Susan is a respected computer scientist and leader in high performance computing at ANL, with over 20 years as an extreme-scale supercomputer and systems integration expert. She is currently the Argonne Project Director for CORAL.

Al Geist, Chief Technology Officer
Al is a well known computer scientist and an ORNL Corporate Fellow. He serves as CTO for the LCF at ORNL and Chief Technologist of the Computer Science and Mathematics Division at ORNL.

Julia White, Technical Operations Manager
Julia holds a Ph.D. in chemistry as well as an MBA, and has a wealth of experience in computational science. She previously held management roles at ORNL, PNNL, and Physical Review B and was most recently the program manager for INCITE, which allocates computer time on DOE leadership class computers.

ECP AD WBS

Exascale Computing Project 2.0

- **Project Management 2.1**
  - Project Planning and Management 2.1.1
  - Project Controls and Risk Management 2.1.2
  - Business Management 2.1.3
  - Procurement Management 2.1.4
  - Information Technology and Quality Management 2.1.5
  - Communications and 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
  - Math 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 Application Development Leadership Team

WBS 2.2: Diverse team bringing new ideas & approaches to results-driven application development

Jack Deslippe, Chemistry and Materials Applications (2.2.1)
Jack is the acting group lead of the Application Performance Group NERSC and leads the NERSC Exascale Science Applications Program (NESAP). Jack has a history of support and development of HPC apps in Materials Science and Ph.D. in Condensed Matter Physics from UC Berkeley..

Tom Evans, Energy Applications (2.2.2)
Tom is a Distinguished R&D Staff member and Team Leader of the HPC Methods and Applications Team in Reactor and Nuclear Systems Division at ORNL. He over 20 years of experience in the development of single- and coupled-physics development on HPC hardware spanning Office of Science and NNSA applications.

Anshu Dubey, Earth and Space Science Applications (2.2.3)
Anshu is a Computer Scientist in the Mathematics and Computer Science Division at ANL. She has a long history of working with all aspects of PDE based scientific software.

Bill Hart, Data Analytics and Optimization Applications (2.2.4)
Bill is a senior research in the Center for Computing Research (CCR) at SNL. As a researcher and former manager in CCR, Bill has led and championed research in scalable optimization, data science and cybersecurity. Bill has a Ph.D. in Computer Science from UC San Diego.

Marianne Francois, National Security Applications (2.2.5)
Marianne is the group leader of the Methods and Algorithms Group in the Computational Physics Division in the Weapons Physics Directorate at LANL. She has a history of developing methods for complex multi-material flow and Ph.D. in Aerospace Engineering from University of Florida.

Phil Colella, Co-Design (2.2.6)
Phil is Senior Scientist at in the Computational Research Division at LBNL, and a Professor in Residence in the EECS Department at UC Berkeley. His research has been in the area of numerical methods and HPC software development for partial differential equations.

ECP EXASCALE COMPUTING PROJECT

ECP Applications: Exascale-capable modeling, simulation, data

Goal

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

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

32 WBS L4 subprojects executing RD&D

134 L4 subproject (P6) milestones delivered in FY17

318 L4 subproject (P6) milestones planned in FY18-19

Selected FY17 AD Highlights

Release

- CoPA co-design center prepared the HACC parallel FFT algorithm for broad distribution (milestone ADCD02-3)

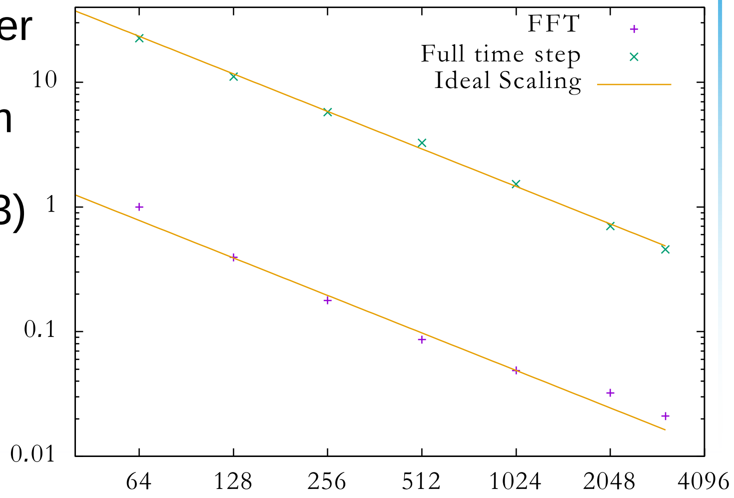

Strong scaling of the 3D FFT on Intel KNL

Science Capability

- ExaSky performed large-scale Nyx simulation to interpret Lyman alpha observations, important for understanding the properties of dark matter and neutrinos (milestone ADSE01-26)

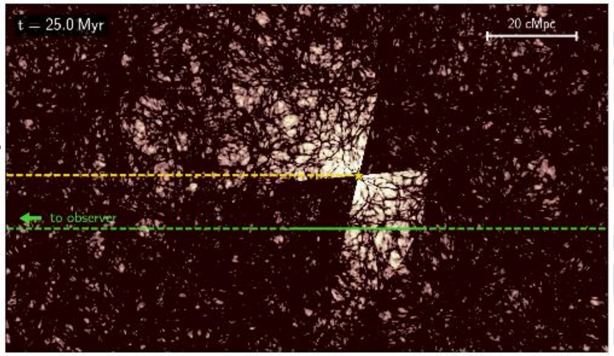

Helium II cosmic web in a Nyx simulation lit by a bi-conical quasar emission showing the parabolic shaped region that can be reached for the given quasar age (25 Myr).

Algorithm or Physical Model Enhancement

- Combustion-PELE developed the capability to automatically predict the thermochemical properties for the combustion of syngas, natural gas, and butane (milestone ADSE14-57)

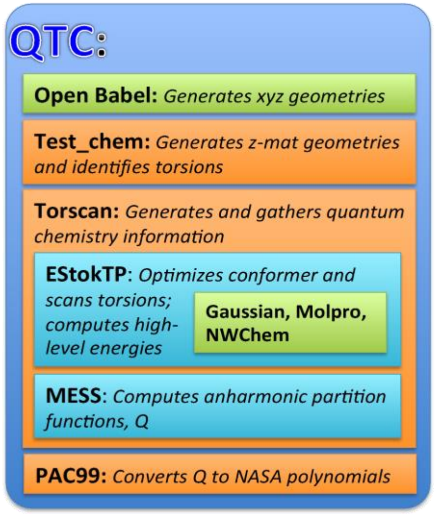

Performance

- ExaSMR demonstrated first ever full-core reactor simulation with continuous-energy MC on GPUs (milestone ADSE08-36)

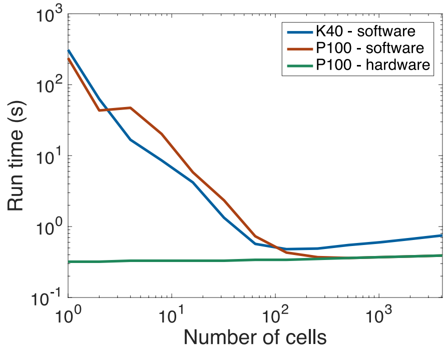

EXASCALE COMPUTING PROJECT

AD Challenges

Challenges

Adopting and co-designing programming models that portably expose critical system features to attain portable, efficient performance. Dealing with extreme parallelism, increased faults, and complex node architectures.

- Exposing billion way concurrency in applications: instruction-level, SIMD, MIMD, shared and distributed memory.
- Optimizing applications for deeper and more heterogeneous memory hierarchies.
- Efficiently analyzing exabyte-sized data sets for scientific discovery.
- Avoiding bulk-synchronization and dynamically handling faults within applications.

ECP ST WBS

Exascale Computing Project 2.0

- **Project Management 2.1**
  - Project Planning and Management 2.1.1
  - Project Controls and Risk Management 2.1.2
  - Business Management 2.1.3
  - Procurement Management 2.1.4
  - Information Technology and Quality Management 2.1.5
  - Communications and 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
  - Math 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 EXASCALE COMPUTING PROJECT

ECP Software Technology Leadership Team

Rajeev Thakur, Programming Models and Runtimes (2.3.1)
Rajeev is a senior computer scientist at ANL and most recently led the ECP Software Technology focus area. His research interests are in parallel programming models, runtime systems, communication libraries, and scalable parallel I/O. He has been involved in the development of open source software for large-scale HPC systems for over 20 years.

Jeff Vetter, Development Tools (2.3.2)
Jeff is a computer scientist at ORNL, where he leads the Future Technologies Group. He has been involved in research and development of architectures and software for emerging technologies in HPC, such as heterogeneous computing and nonvolatile memory, for over 15 years.

Lois Curfman McInnes, Mathematical Libraries (2.3.3)
Lois is a senior computational scientist in the Mathematics and Computer Science Division of ANL. She has over 20 years of experience in high-performance numerical software, including development of PETSc and leadership of multi-institutional work toward sustainable scientific software ecosystems.

Jim Ahrens, Data and Visualization (2.3.4)
Jim is a senior research scientist at LANL and an expert in data science at scale. He started and actively contributes to many open-source data science packages, including ParaView and Cinema.

Rob Neely, Software Ecosystem and Delivery (2.3.5)
Rob has several leadership roles at LLNL spanning applications, CS research, platforms, and vendor interactions. He is an Associate Division Leader in the Center for Applied Scientific Computing, chair of the Weapons Simulation and Computing Research Council, and the lead for the Sierra Center of Excellence.

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
185 L4 subproject (P6) milestones delivered in FY17
426 L4 subproject (P6) milestones planned in FY18-19

Software Development Kits are a key design feature of ST and are an important delivery vehicle for software products

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

Math SDK

Tools SDK

PM&RT SDK

DataViz SDK

Facility SDK

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

Assume all releases are delivered as “build from source” via Spack – at least initially

Focus on ensuring that software compiles robustly on all platforms of interest to ECP (including testbeds)

xSDK release 0.3.0 for ECP

ECP WBS	WBS 2.3.3.01
PI	Lois Curfman McInnes, ANL
Members	LBNL, LLNL, SNL, UC Berkeley, UTK

Milestone lead: Jim Willenbring, SNL

Scope and objectives

- Demonstrate the impact of community policies to simplify the combined use and portability of independently developed software packages.
- Increase formality of xSDK release process.
- Expand xSDK package members to include additional key ECP numerical libraries.

Spack install graph for xsdk-0.3.0

Legend:
xSDK
xSDK packages
Build dependencies
Link dependencies

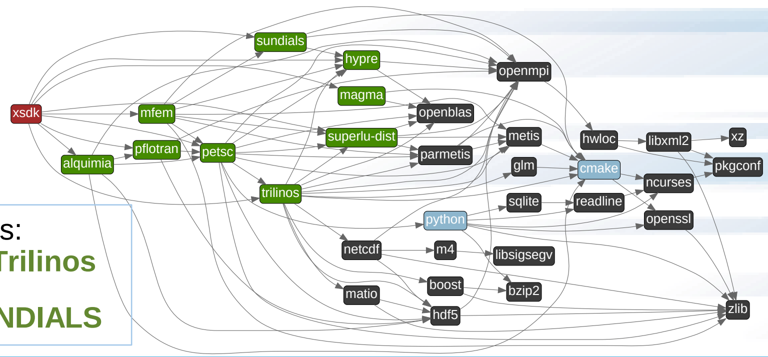

Original xSDK math libraries:
hypre, PETSc, SuperLU, Trilinos

New: **MAGMA, MFEM, SUNDIALS**

Impact

- Lay the groundwork for addressing broader issues in software interoperability and performance portability.
- Improve access to numerical libraries for ECP apps.
- Provide ECP SDKs with an example release process.
- Build roles and experience for future exascale platform porting.

Project accomplishment

- xSDK 0.3.0 released December 2017.
- xSDK 0.3.0 tested on key platforms at ANL, NERSC, and ORNL (Cori, Titan, Theta), also Linux and Mac OS X.
- Documented package compatibility with xSDK community policies.

Deliverables https://xsdk.info/download https://xsdk.info/release-0-3-0 https://xsdk.info/installing-the-software https://github.com/xsdk-project/installxSDK https://xsdk.info/packages https://xsdk.info/policies

EXASCALE COMPUTING PROJECT

ST Challenges

Challenges

Qualitative changes: Massive concurrency; Multi-scale, multi-physics, data-driven science; Ecosystem integration

- Billion way concurrency: Several novel compute nodes.
- Coupled apps: Physics, scales, in situ data, more.
- Data-driven: New software HPC environments, containers.
- Ecosystem: Part of a large, complex, evolving SW environment.

ECP HI WBS

Exascale Computing Project
2.0

- **Project Management 2.1**
  - Project Planning and Management 2.1.1
  - Project Controls and Risk Management 2.1.2
  - Business Management 2.1.3
  - Procurement Management 2.1.4
  - Information Technology and Quality Management 2.1.5
  - Communications and 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
  - Math 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

EXASCALE COMPUTING PROJECT

ECP Hardware and Integration Leadership Team

WBS 2.4: Possessing breadth and depth of knowledge and experience well-matched to objectives

Bronis de Supinski, PathForward (2.4.1)
Bronis is the CTO for Livermore Computing at LLNL, where he formulates LLNL's large-scale computing strategy and oversees its implementation. He earned his Ph.D. in Computer Science from the University of Virginia. Bronis is also a Professor of Exascale Computing at Queen's University of Belfast and an Adjunct Associate Professor in the Department of Computer Science and Engineering at Texas A&M University.

Simon Hammond, Hardware Evaluation (2.4.2)
Si has worked for over a decade in application performance optimization and hardware design. He received his PhD. from the University of Warwick where he worked extensively with the UK's Atomic Weapons Establishment. He now works as a Research Scientist at Sandia's CSRI on hardware/software co-design and next-generation computing architectures.

Judy Hill, Application Integration at Facilities (2.4.3)
Judy is currently the INCITE Program Manager for the LCFs and the Task Lead for INCITE Liaison Activities at the OLCF. She is heavily involved in application readiness activities at the OLCF. Judy holds a PhD in computational science with a subsequent 10+ years of experience bringing numerical methods and HPC expertise to a variety of application domains.

Dave Montoya, Software Deployment at Facilities (2.4.4)
David, the Technology Futures Lead for Software Environments in the HPC organization at LANL, has experience in workflow characterization, integrated application performance metrics, application/system workflow assessment, and coordinating of the NNSA tri-lab Common Computing Environment Project. He holds a BA in Information Systems Analysis from New Mexico State University and an MBA from the University of Rochester.

Julia White, Facility Resource Utilization (2.4.5)
Holds a Ph.D. in chemistry as well as an MBA, and has a wealth of experience in computational science. She previously held management roles at ORNL, PNNL, and Physical Review B and was most recently the program manager for INCITE, which allocates computer time on DOE leadership class computers.

Ashley Barker, Training and Productivity (2.4.6)
Ashley is the Group Leader for the User Assistance and Outreach team at the OLCF, where her group is responsible for facilitating access to OLCF resources, providing training, documentation, and technical support to users, collecting and reporting on user facility data, and acquainting the public with the work conducted at the OLCF through scientific highlights. Ashley earned her MS in information sciences from the University of Tennessee.

ECP Hardware and Integration: Delivery of integrated ECP/DOE facility products

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

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

20 WBS L4 subprojects executing RD&D
38 L4 subproject (P6) milestones delivered in FY17
266 L4 subproject (P6) milestones planned in FY18-19

HI: Designed to enable integration of ECP's products into HPC environments at the Facilities

ECP will demonstrate meeting objectives on Facility resources

ECP products

- Applications
- Software
- Early HW R&D

Hardware and Integration

- Facility resource utilization
- Developer training and productivity
- Application integration at Facilities
- SW deployment at Facilities
- PathForward
- HW evaluation

Facilities

- ASCR and NNSA HPC Facilities
- US vendor system offerings

HI is the result of a maturing of ECP's thinking about the end game

EXASCALE COMPUTING PROJECT

Relevant Pre-Exascale and Exascale Systems for ECP

Pre-Exascale Systems

2013

Argonne
IBM BG/Q
Open

ORNL
Cray/NVidia K20
Open

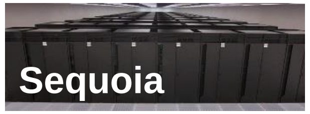

LLNL
IBM BG/Q
Secure

2016

Argonne
Intel/Cray KNL
Open

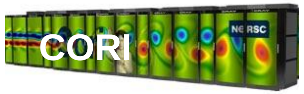

LBNL
Cray/Intel Xeon/KNL
Unclassified

LANL/SNL
Cray/Intel Xeon/KNL
Secure

2018

ORNL
IBM/NVidia
P9/Volta
Open

LLNL
IBM/NVidia
P9/Volta
Secure

2020

NERSC-9

LBNL
TBD
Open

Crossroads

LANL/SNL
TBD
Secure

Exascale Systems

2021-2022

Argonne
Intel/Cray TBD
Open

Frontier

ORNL
TBD
Open

El Capitan

LLNL
TBD
Secure

Selected HI FY17 Highlights: Training and Productivity Activities

Date	Event	# Registered	# Attended
FY17 Q2	Introduction to ECP Project Mgmt Tools (Confluence & JIRA)	N/A	N/A
FY17 Q3	Python in HPC	210	132
	OpenMP Tutorial	113	60
FY17 Q4	Argonne Training Program on Extreme-Scale Computing 2017	168 (applied)	71 (selected)
	Intermediate Git	149	78
	Using the Roofline Model and Intel Advisor	113	71
	Barely Sufficient Project Management	190	122
	Scalable Node Programming with OpenACC	75	46
FY18 Q1	HipChat Webinar	N/A	N/A
	Managing Defects in HPC Software Development	135	63
	TAU Performance System Webinar	53	31

- Over 650 people have participated in ECP Training activities to-date
- ECP team members have led additional tutorials and workshops at events such as SC, ISC, and PASC.
- Have created documentation and best practices covering topics like ECP project management tools
- Launched Better Scientific Software website – resource for developer productivity and sustainability

Public training materials are archived on the ECP External Website:
http://exascaleproject.org
Internal training materials are found on the ECP Confluence Knowledgebase site:
https://confluence.exascaleproject.org/display/KB

Collaboration between ECP and Facilities is Critical for Success

Engagement Plan starts with existing touch points and expands into formal activities of mutual benefit:

- **Application Development**
  - ECP will collaborate and share experiences in user programs, facility testbeds, and development work on pre-exascale and exascale systems
  - ECP will work with facilities to obtain access to vendor staff and system expertise for help in preparing applications for the exascale and pre-exascale systems

- **Software Deployment and Support**
  - ECP will collaborate with facilities staff to explore deployment and integration of software on production systems
  - ECP will closely collaborate with facilities staff on strategies for continuous testing and hardening of advanced software components on facilities architectures and within the facilities existing software stacks
  - Facilities will be actively involved in reviewing software projects and gap analysis

- **Hardware and Integration**
  - ECP and Facilities will hold joint training and workshops, with invitations to the trainings held by either organization to be open to the other organization
  - DOE pre-exascale and exascale systems will benefit from vendor R&D efforts supported by the DOE Pathforward program as will the Facilities and ECP applications
  - SC Facilities and ECP will collaborate in the allocation of facility resources to ECP to most effectively utilize the resources for science and exascale development

RFP for Continuous Integration Solution

- Released ECP RFP for a Continuous Integration Solution on 12/12/17.
- Selected vendor will deliver a secure, easy-to-use Continuous Integration (CI) solution for deployment at large-scale HPC facilities that will run *on the real hardware* at each site and *in that site's environment* to test for hard-to-find bugs related to the site's deployment environment. The solution should be deployable on-premises at the HPC facilities or as a trusted cloud application that can run jobs at HPC facilities.
- RFP closes 01/05/2018.
- RFP will be evaluated by members of ECP ST, AD, and HI teams along with designated members of the facilities.
- Activity will provide benefit to both ECP and the ASCR/NNSA computing facilities.
- RFP URL: http://procurement.ornl.gov/rfp/6400015283/

OAK RIDGE National Laboratory

This website should be monitored routinely for notification of posted amendments.

Solicitation No. 6400015283

CONTINUOUS INTEGRATION SOLUTION FOR DEPLOYMENT AT LARGE-SCALE HIGH PERFORMANCE COMPUTING FACILITIES

- Vendor Instruction Letter ***(Added 12/18/2017)***
- Solicitation *(Added 12/12/2017)*
- Attachments
  - Attachment 1 - Statement of Work *(Added 12/12/2017)*
  - Attachment 2 - ECP CI Use Cases *(Added 12/12/2017)*
  - Attachment 3 - Intro to CI for HPC Centers *(Added 12/12/2017)*
- Exhibits
- Special Articles and Forms
- Terms and Conditions

EXASCALE COMPUTING PROJECT

HI Challenges

Challenges

Time pressure on HPC ecosystem and DOE facilities; Developing and maintaining strong collaborations with many facilities, AD, and ST; 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

ECP Preliminary Design Report (PDR)

Present comprehensive technical design and approach to meeting ECP's objectives (mission need, KPPs)

Accelerate and Translate Cancer Research (CANDLE)

The DOE has entered into a partnership with the National Cancer Institute (NCI) of the National Institutes of Health (NIH). This partnership has identified three key challenges that the combined resources of DOE and NCI can accelerate. The first challenge (called the "RAS pathway problem") is to understand the molecular basis of key protein interactions in the RAS/RAF pathway that is present in 30% of cancers. The second challenge (called the "drug response problem") is to develop predictive models for drug response that can be used to optimize pre-clinical drug screening and drive precision medicine based treatments for cancer patients. The third challenge (called the "treatment strategy problem") is to automate the analysis and extraction of information from millions of cancer patient records to determine optimal cancer treatment strategies across a range of patient lifestyles, environmental exposures, cancer types, and healthcare systems. While each of these three challenges are at different scales and have specific scientific teams collaborating on the data acquisition, data analysis, model formulation, and scientific runs of simulations, they also share several common threads. The Exascale Deep Learning and Simulation Enabled Precision Medicine for Cancer focuses on the machine learning aspect of the three challenges and, in particular, builds on a single scalable deep neural network code called CANDLE (CANcer Distributed Learning Environment).

CANDLE Application Overview

Code Name(s)	Computational Motifs	Figure of Merit	Software Integration
CANDLE	Deep learning training, machine learning, neural networks, dense matrix multiply, sparse matrix multiply, FFT	Training rate: number of networks trained/day + number of "weight" updates/sec	CODAR, ADIOS, ExaHDF5, Swift, SPINDLE, Sonar, OpenMP, PROTEAS, MPI

CANDLE FY18 Milestones

Milestone ID	Milestone Title
ADOA01-13	Demonstrate a prototype DNN that includes drug descriptors, drug fingerprints and at least 3 cell line properties
ADOA01-18	Demonstrate a prototype DNN that performs unsupervised feature learning on MD simulation neighborhoods
ADOA01-22	Scalable, distributed, and resilient toolkit
ADOA01-26	CANDLE version 1.0
ADOA01-32	Demonstrate a prototype multi-task learning DNN that extracts simultaneously multiple targets and determine optimum network architecture for maximizing capacity FOM
ADOA01-35	Demonstrate an improved DNN that adds drug target descriptions to the input and improves cell line properties set

CANDLE FY19 Milestones

Milestone ID	Milestone Title
ADOA01-41	Demonstrate scaling of unsupervised feature learning to full-size MD neighborhoods
ADOA01-46	CANDLE workflow and high-level support library version 0.1
ADOA01-52	Provide support for multi-scale multi-task learning in internal network layers
ADOA01-55	CANDLE version 2.0
ADOA01-60	Demonstrate ladder networks to extract O(100)s bio-marker and other relevant fields from pathology clinical reports
ADOA01-62	Demonstrate a local receptive field layer that captures the locality of biological features
ADOA01-65	Demonstrate a DNN that can predict PDX drug response from cell line data
ADOA01-70	Demonstrate protein configuration HMM learning in a learned MD feature space

The CANDLE project has defined three specific challenge problems. For the RAS pathway problem, multi-scale molecular dynamics (MD) runs will be guided through a large-scale state-space search using unsupervised learning to determine the scope and scale of the next series of simulations based on the history of previous simulations. This will be used to develop a predictive capability for modeling the behavior of proteins on membranes and to apply that capability to RAS and effector proteins along the primary RAS signaling pathways. For the drug response problem, unsupervised machine learning methods will be used to capture the complex, nonlinear relationships between the properties of drugs and the properties of the tumors to predict response to treatment (and therefore develop a model that can provide treatment recommendations for a given tumor). For the treatment strategy problem, semi-supervised machine learning will be used to automatically read and encode millions of clinical reports into a form that can be used to develop the predictive models necessary for population-wide cancer surveillance that extends beyond the clinical trial setting.

AD WBS L4 subproject example (cancer application)

xGA: Global Arrays on Extreme Scale Architectures

The xGA project is focused on improving the performance, scalability, and user productivity of the Global Arrays (GA) library for exascale systems. This is essential to the ECP applications that already depend on GA such as NWChemEx, GAMESS and GridPACK (used as part of the Stochastic Grid Dynamics project). In addition, GA is being considered for use in the ECP application QMCPACK. GA supports a shared memory-like programming model on distributed memory platforms. This allows users to create distributed multidimensional arrays that can be accessed from any processor using simple one-sided put/get/accumulate communication primitives. We are extending the GA library to take advantage of exascale architecture features including deep memory hierarchies and accelerators, while continuing to tune and improve the performance of the GA runtime.

xGA Project Overview

Scope & Intent	R&D Themes	Delivery Process	Target ECP Users	Support Model
Improved performance and scalability though advantageous use of exascale hardware features.	User productivity and application performance via scalable programming models.	Regular releases of software and documentation via the GitHub project page, with interim availability of new features via git branch management.	Many apps, including NWChemEx, GAMESS, QMCPACK, and ExaSGD. Early users of ExaMPI, OMPI-X, SCIM.	Ongoing developer support. Dedicated email list. GitHub issue tracking. Web-based documentation.

xGA FY18 Milestones

Milestone ID	Milestone Title	ECP Users
STPM07-4	Initial xGA Property Types	Apps already using xGA, e.g., NWChemEX, and those with large read-only dense array data structures.
STPM07-6	Sparse Data Support in xGA	Apps with sparse data requirements such as ExaSGD.
STPM07-7	Advanced Multi-threaded xGA	Apps already using xGA, e.g., NWChemEX, that require GA to be called within threaded loops and remain performant.
STPM07-9	Integration with ECP Applications and Public Release	Ensure Apps already using xGA, e.g., NWChemEX, remain supported and tested. Demonstrate new app integration, e.g., QMCPACK.

xGA FY19 Milestones

Milestone ID	Milestone Title	ECP Users
STPM07-8	xGA Read Only Property	Apps already using xGA, e.g., NWChemEX, early adopters of the initial xGA property types, and those with large read-only dense array data structures.
STPM07-10	xGA Hash-table support	Apps requiring shared distributed hash tables such as QMCPACK and ExaSGD.
STPM07-11	User-supplied Memory Support for PGAS view in xGA	Improved interoperability with existing app data structures that require temporary global shared access to data.
STPM07-12	Non-algebraic Collectives	

xGA Impact Goals & Metrics

Goal	Metric
xGA is used as the primary source for scalable dense arrays in ECP applications.	The number of ECP applications that have integrated the xGA product into their production build environment and rely on xGA distributed dense arrays for scalable, robust execution on exascale-ready problems.
xGA facilitates the use of exascale hardware features.	The number of exascale hardware features utilized by xGA including memory hierarchies and accelerators.

ST WBS L4 subproject example (global arrays)

ECP Preliminary Design Review Feedback

From Dec 5, 2017 Independent Design Review (IDR)

- **Selected Observations**
  – Review committee was impressed with the quality of the Preliminary Design Report (PDR) and its interaction with the ECP leadership team. The leadership team is fully staffed down through level 3 of the WBS with highly competent leaders
  – ECP has restructured and realigned their technical approach to reflect the changes to the project's scope, schedule and budget in order to meet the overall project goals and KPPs
  – The project has engaged with the SC and NNSA facilities to begin planning for the successful deployment of exascale applications
- **Selected Recommendations**
  – Review ECP milestones and Figures of Merit (FOM) for level consistency
  – Make sure each of the Software Technology (ST) projects have at leaset one impact goal which directly connects to one or more Application Development (AD) efforts.
  – Clearly define how the risk register is developed and used (define risk categories, cost estimates, contingency, proposed mitigations)
  – Once the technical and contractual details for the first two exascale systems are known, the Hardware and Integration (HI) focus area should review and update their preliminary planning, schedule, and risks to identify potential barriers to successfully achieving KPPs.

ECP: Summary

- We believe we have a good team and plan, executing full bore now
- Enjoying DOE stakeholder support (e.g., SC and Applied Offices for ECP applications)
- Still have work to do: tracking progress (dashboard), maintaining agility for highest ROI on RD&D, more proactive risk management, better integration with facilities
- Working to actively manage our highest risks
- Moving into a product development stage with more quantitative metrics
- Near term radar
  - ECP Independent Project Review (IPR): Jan 9-11, 2018 (ORNL)
  - ECP 2nd Annual Meeting: Feb 5-9, 2018 (Knoxville, TN)