Partnering with the Office of Science Computing Facilities

Terri Quinn
Director, Hardware and Integration, ECP

Presentation to the Advanced Scientific Computing Advisory Committee
March 26, 2019
Cambria Hotel Rockville, Maryland

Outline

- Introduce the ASCR computing facility and ECP partnership
- Introduce the Hardware and Integration Focus Area
- Partnering Example #1 - Hardware innovations leading to exascale systems
- Partnering Example #2 - Application preparation for exascale systems
- Partnering Example #3 - Software deployment and testing capabilities
- Engagement Plans

DOE Exascale Program: The Exascale Computing Initiative (ECI)

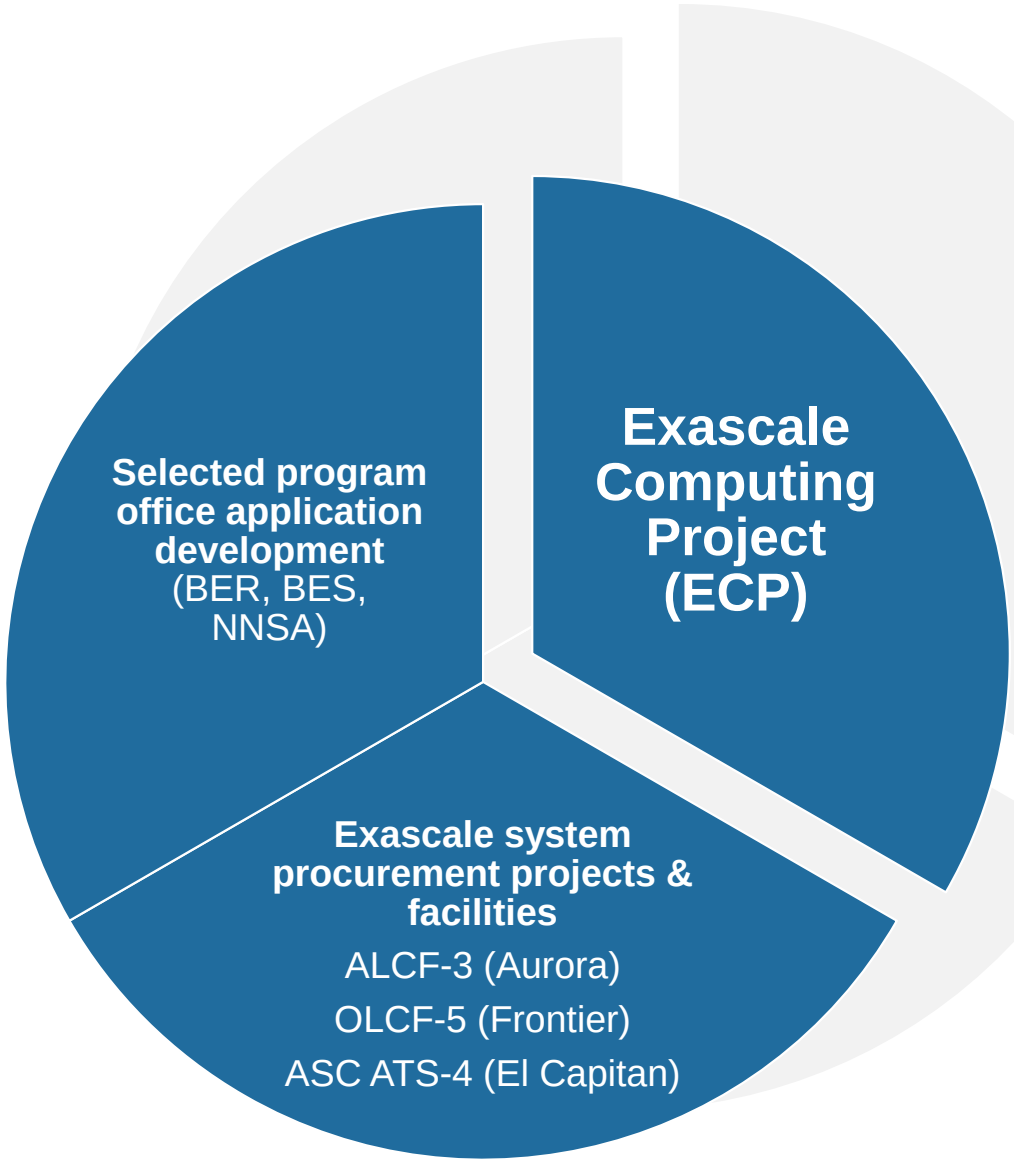

ECI partners

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

ECI mission

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

ECI focus

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

ECI Roles for ASCR Computing Facilities (Facilities) and ECP

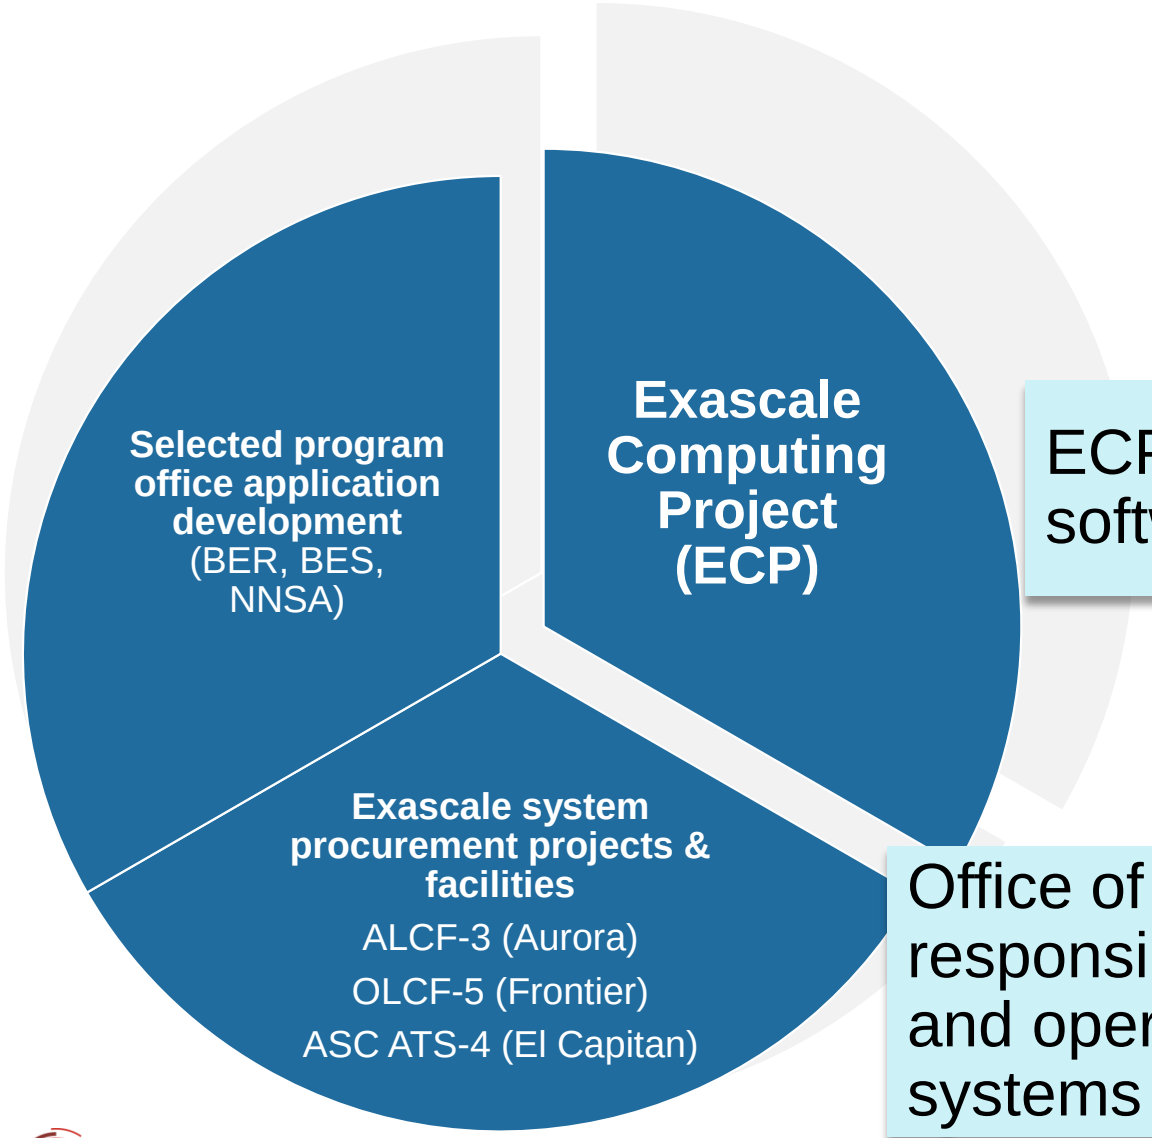

ECP is responsible for advances in software, hardware, and applications

Office of Science Computing Facilities are responsible for the acquisition, installation, and operation of computers – advances in systems

Facilities and ECP Partnering

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

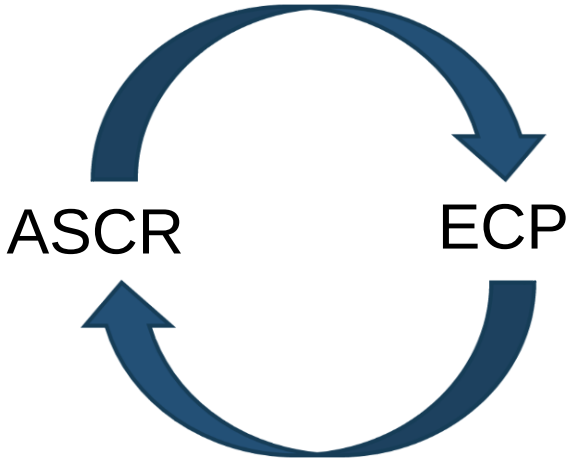

- Mission ends with the project
- ECP is a few years old

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

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

Principles for a Healthy ECP and Facilities Partnership

Make it a win-win partnership

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

Leverage each others capacities

Two examples are:

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

Align our plans to the extent that makes sense

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

Introducing the Hardware and Integration Focus Area

Hardware and integration

Develop technology advances for exascale and deploy ECP products

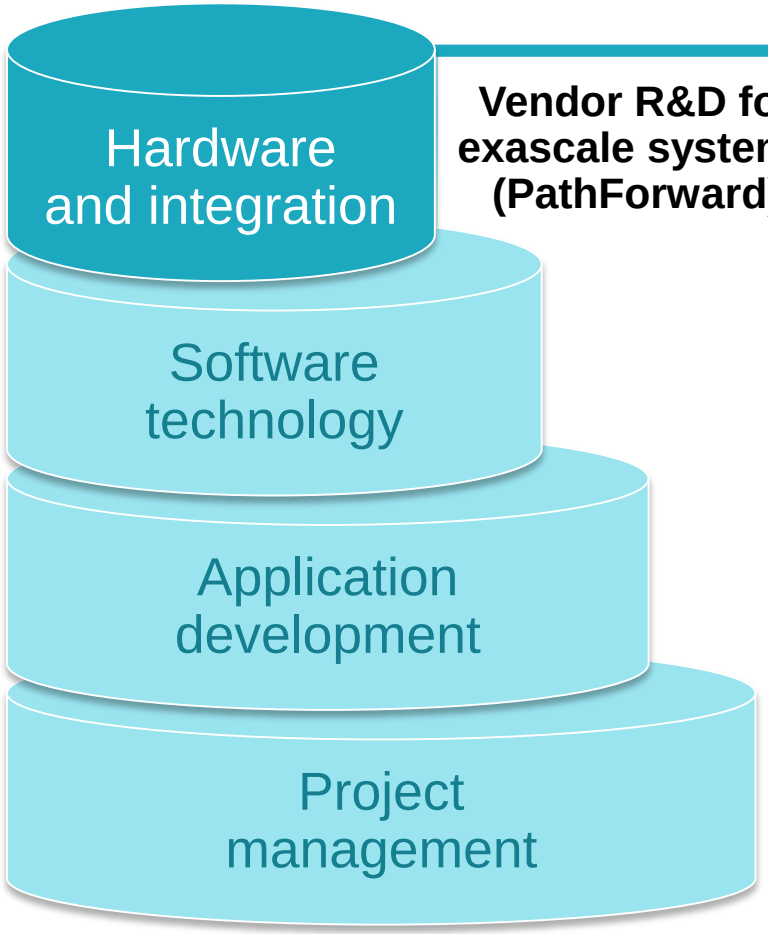

Vendor R&D for exascale systems (PathForward)

Evaluation of hardware technology and performance

Software deployment and application integration at HPC facilities

Pre-exascale and exascale system utilization measurement and tracking

Community training and productivity

Within ECP HI has the primary responsibility for the partnership with the Facilities. Each HI project is collaborating with the Facilities.

ECP Hardware and Integration (HI) Leadership Team

Terri Quinn, HI Director (2.4)
Associate Program Director for Livermore Computing Systems & Environment. (LLNL)

Susan Coghlan, HI Deputy Director (2.4)
Extreme-scale supercomputer and systems integration expert. Currently the Argonne Future Systems Project Director, responsible for acquisition and deployment of the Argonne supercomputers. (Argonne)

Bronis de Supinski, PathForward (2.4.1)
Chief Technology Officer and head of the Advanced Technology Office for Livermore Computing (LLNL)

Simon Hammond, Hardware (HW) Evaluation (2.4.2)
Scalable computer architecture and app performance (Sandia)

Judy Hill, Application Integration at Facilities (2.4.3)
Computational scientist and currently managing a user program for the ASCR Leadership Computing Facilities (ORNL)

Dave Montoya, Software Deployment at Facilities (2.4.4)
HPC software development and deployment expert (LANL)

Julia White, Facility Resource Utilization (2.4.5)
Experience as the INCITE program manager. Currently Manager of ECP Technical Operations (ORNL)

Ashley Barker, Training and Productivity (2.4.6)
Group leader of user assistance and outreach (ORNL)

Partnering Example #1
Hardware innovations leading to exascale systems

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

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

Total value of the work is $431M; DOE paid 60% of the price or $259M

Examples of work funded include:

a) innovative memory architectures

b) higher-speed interconnects

c) improved system reliability

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

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

PathForward R&D Thrusts

Company (HPC Vendor)	R&D thrusts
AMD	• Innovations in memory interfaces • CPU and GPU microarchitecture • Component integration • High-speed interconnects
Cray Inc.	• Arm processor enhancements • Network enhancements • Memory architectures
Hewlett Packard Enterprise (HPE)	• System, node, and I/O design • Advance the Gen-Z interconnect • Optical interconnects

Company (HPC Vendor)	R&D thrusts
IBM	• Architectural and system component innovations for a system combining processors, GPUs, high-performance networks, and high-performance storage
Intel Corp.	• Energy efficiency • Reduced fabric costs and power • Scalable storage and memory arch. • Optimized communication characteristics
NVIDIA Corp.	• Energy-efficient GPU architectures • Resilience

The PathForward project and ASCR facility system acquisition projects work together to accelerate the delivery of exascale systems

FastForward I & II
DesignForward I & II
First contract awarded 2012

Accelerating hardware innovations
PathForward
6 Vendors
Each with multiple Work Packages

Non-recurring Engineering (NRE) contracts fund the development of technology (both hardware and software) that is:

1. Required to deliver the system and
2. Would not have been developed (or developed in time) by the vendor.

First FF/DF contract was awarded in 2012

Productizing hardware innovations
Facility Non-Recurring Engineering (NRE) Contracts

Exascale system deliveries (notional)

FY16 FY17 FY18 FY19 FY20 FY21 FY22 FY23

Partnering Example #2
Application preparation for exascale systems

ECP's Flow of Software and Application Delivery and Deployment

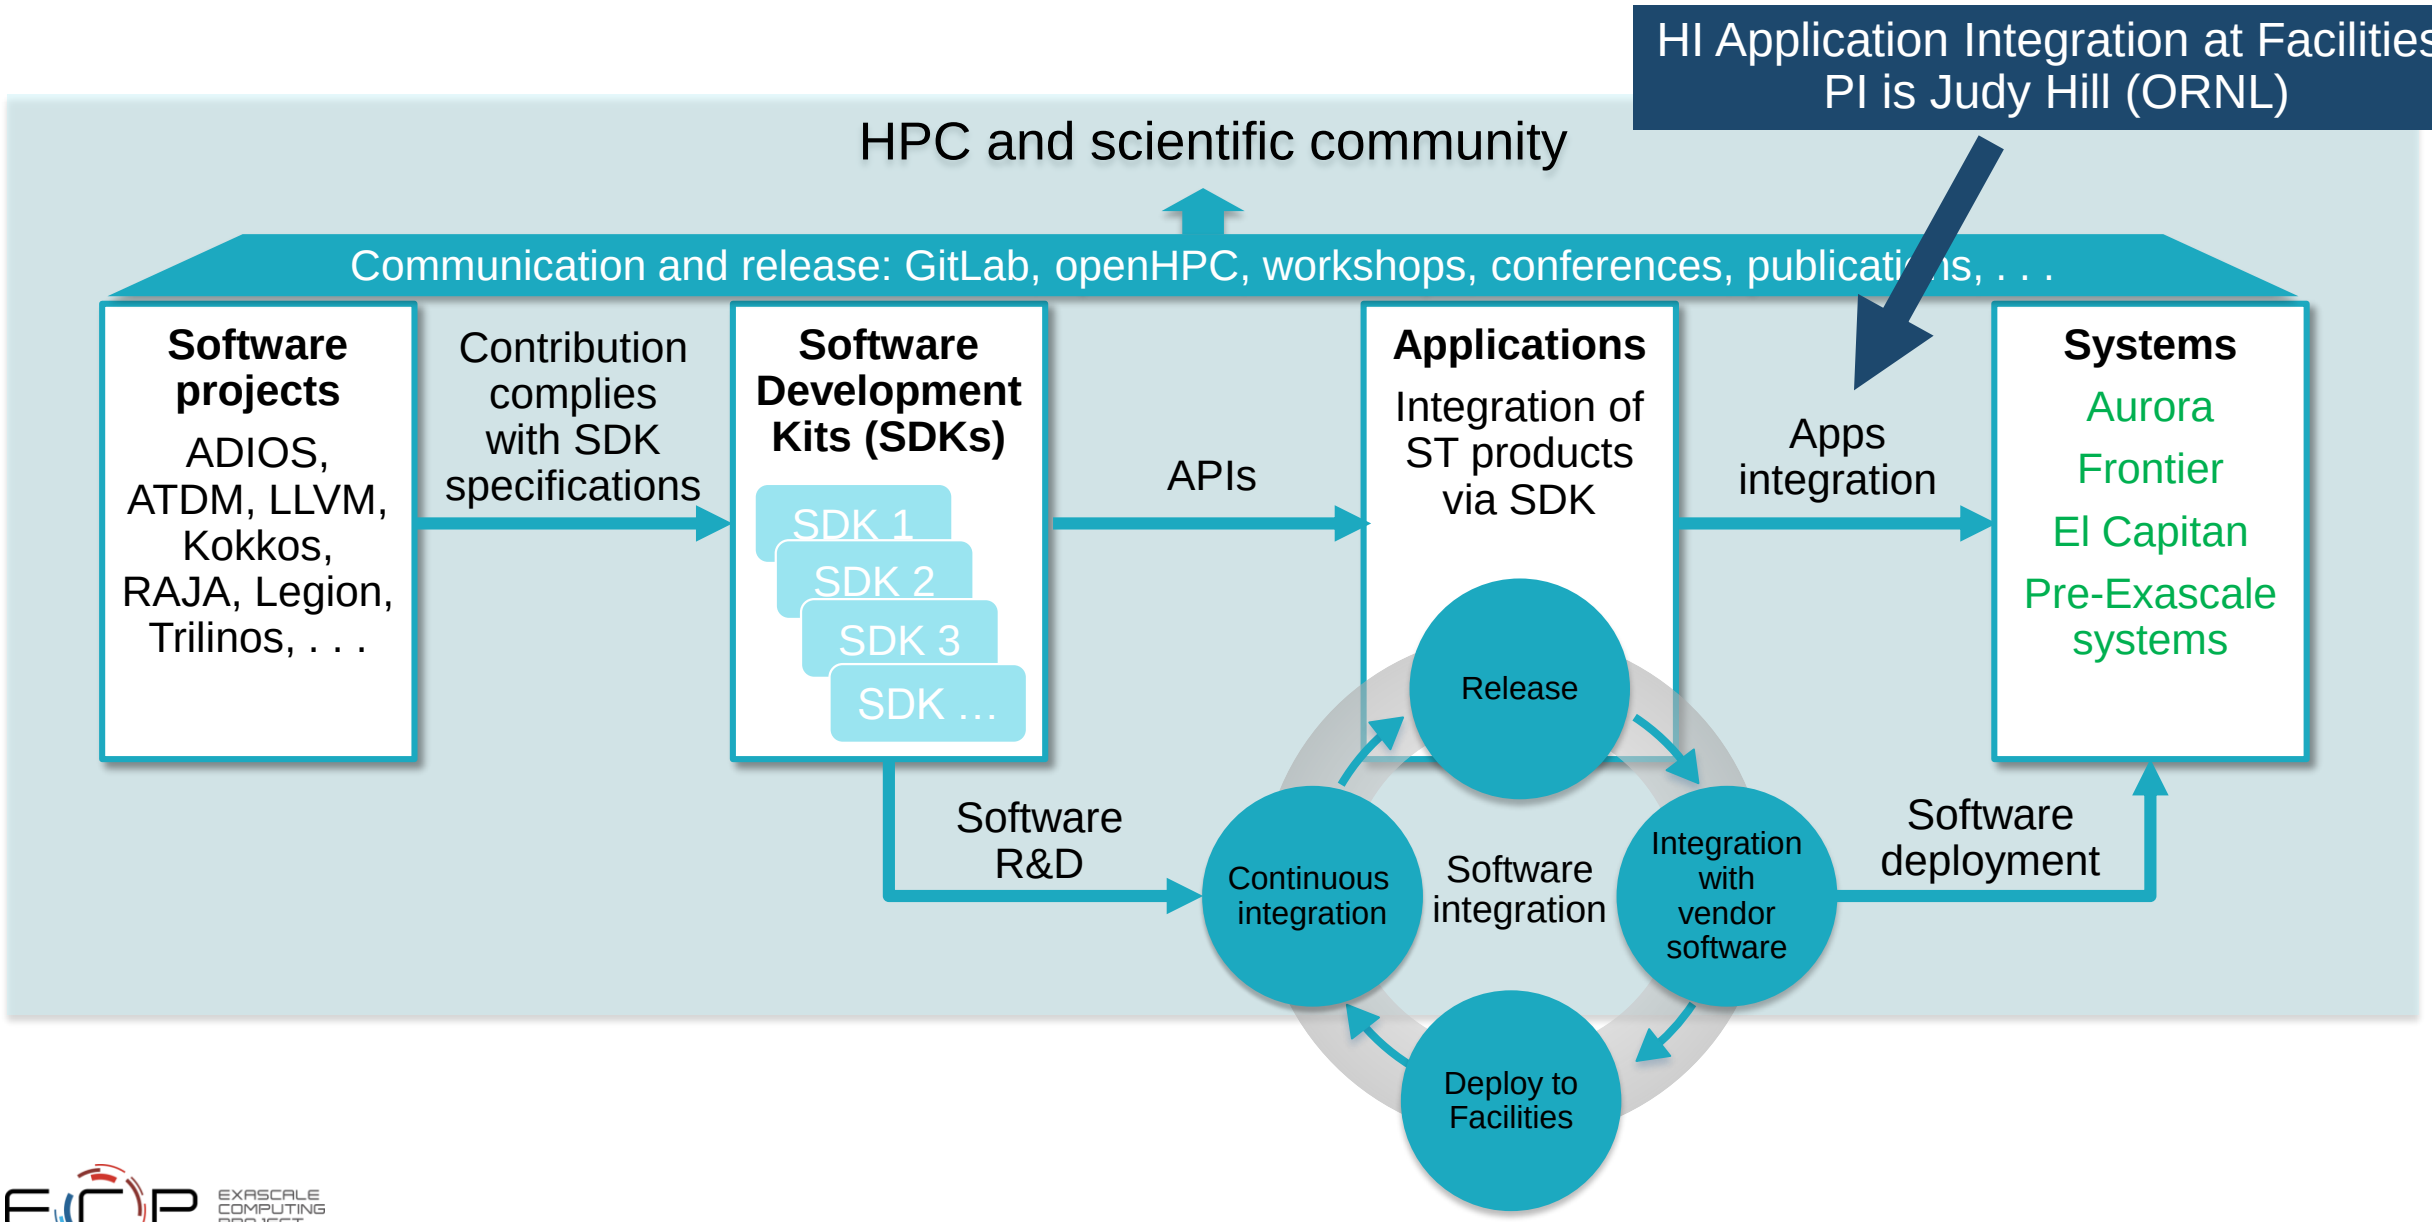

HI's Application Integration Project Accelerates ECP AD's Application Readiness on ASCR Facilities Exascale Architectures

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

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

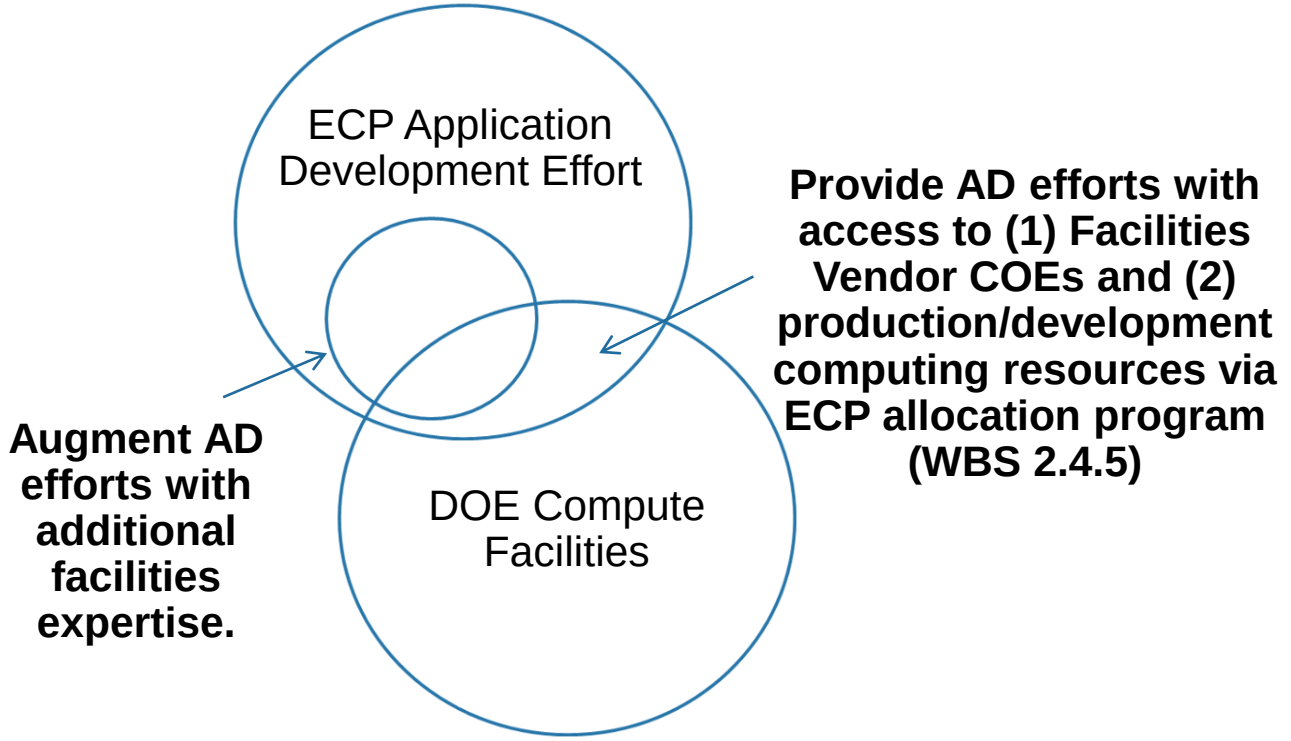

Benefits of this Approach

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

Application Matching to Facilities Plan and Status

Strategy: Match applications with existing facility readiness efforts

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

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

- **Application Development 2.2**
  - Chemistry and Materials Applications 2.2.1
  - Energy Applications 2.2.2
  - Earth and Space Science Applications 2.2.3
  - Data Analytics and Optimization Applications 2.2.4
  - National Security Applications 2.2.5
  - Co-Design 2.2.6

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

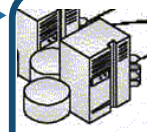

An initial set of ~10 ECP applications will be identified to participate in CAAR-ECP in FY19. Applications may transition in and out of the program as progress is made.

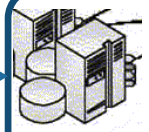

5 ECP AD applications participating in NESAP for NERSC-9 . Additional applications may participate with NERSC funding.

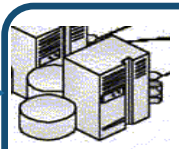

Goal: Progress towards exascale readiness develops, and NESAP-ECP apps transition to LCF facilities

Application Matching to Facilities Key Points

- ECP applications may participate in more than one ECP-funded facilities readiness effort
- ECP applications may participate directly in a facilities-funded readiness effort outside of ECP (scope may not overlap)
- Leadership Computing Facilities will be agile and ECP applications may transition into and out of the readiness efforts as progress is made and achieved
- Application matching and the ECP-funded support is a collaboration between ECP and the Facilities
  - Facilities have input into which applications and how many work with their facility
  - Facilities have input into the process by which applications are identified.

Partnering Example #3
Software deployment and testing capabilities

ECP's Flow of Software and Application Delivery and Deployment

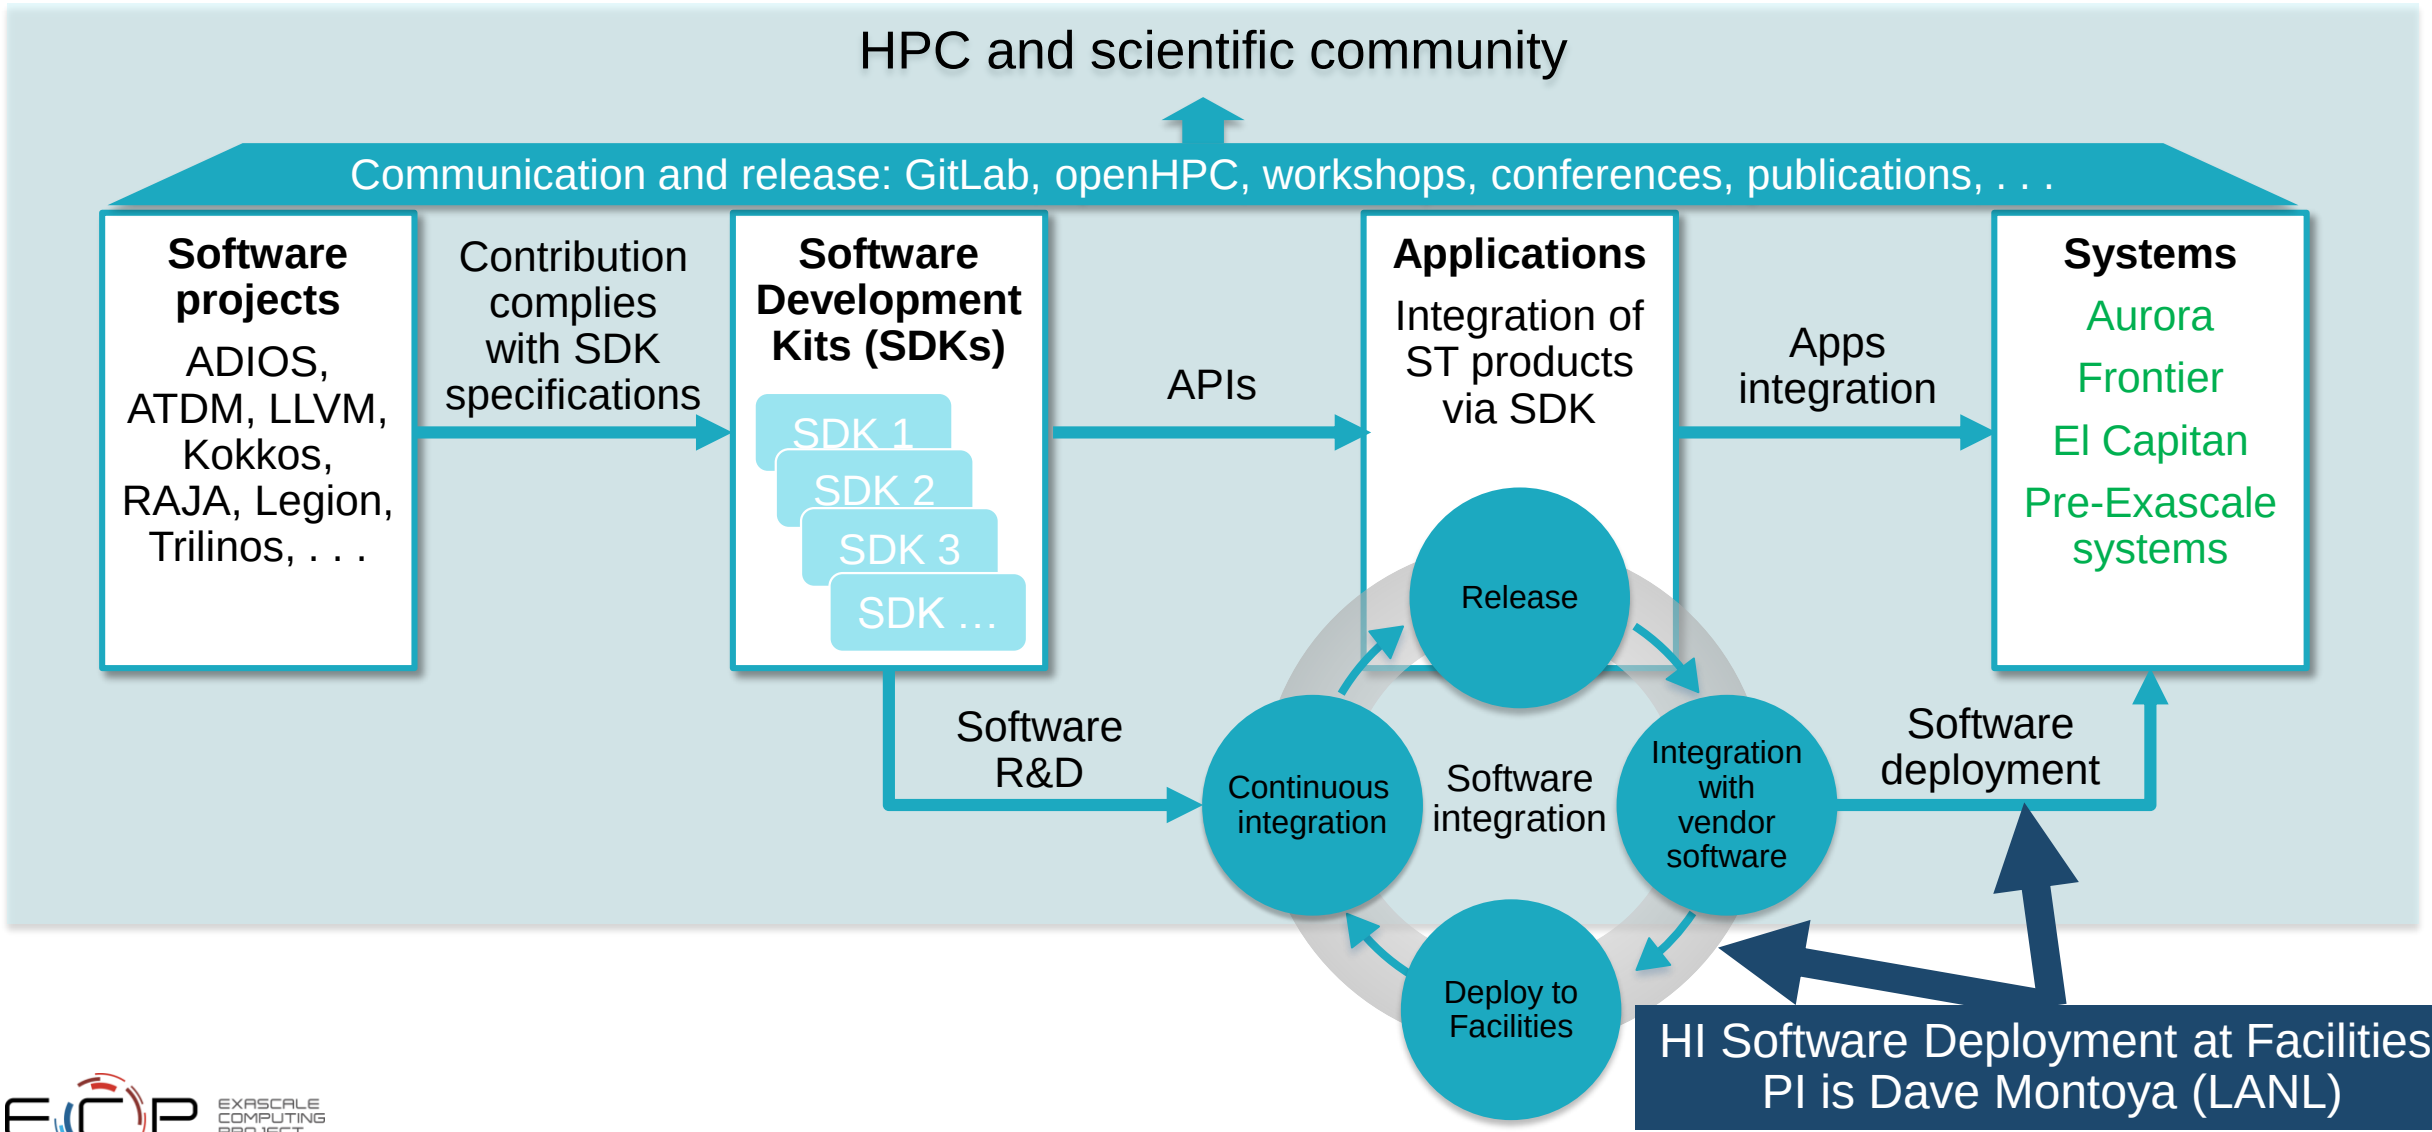

HI's Software Deployment Project Deploys ECP's ST Products to Meet ECP Application and Facility Needs

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

Approach

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

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

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

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

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

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

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

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

Use cases being addressed and the solution being worked

Use Cases

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

Solution

- Use GitLab for the entire DevOps lifecycle (plan, create, verify, release) coupled with runner components located across all sites
- Fund facility staff to define the solution and to deploy and operate the service at their site
- Onyx Point developed runner improvements for DOE secure HPC data centers (identifying users uniquely and enabling submitting jobs directly to resource job management system)
- Centralized core repository managed by DOE Office of Scientific and Technical Information (OSTI) accessible by ECP and E6 users

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

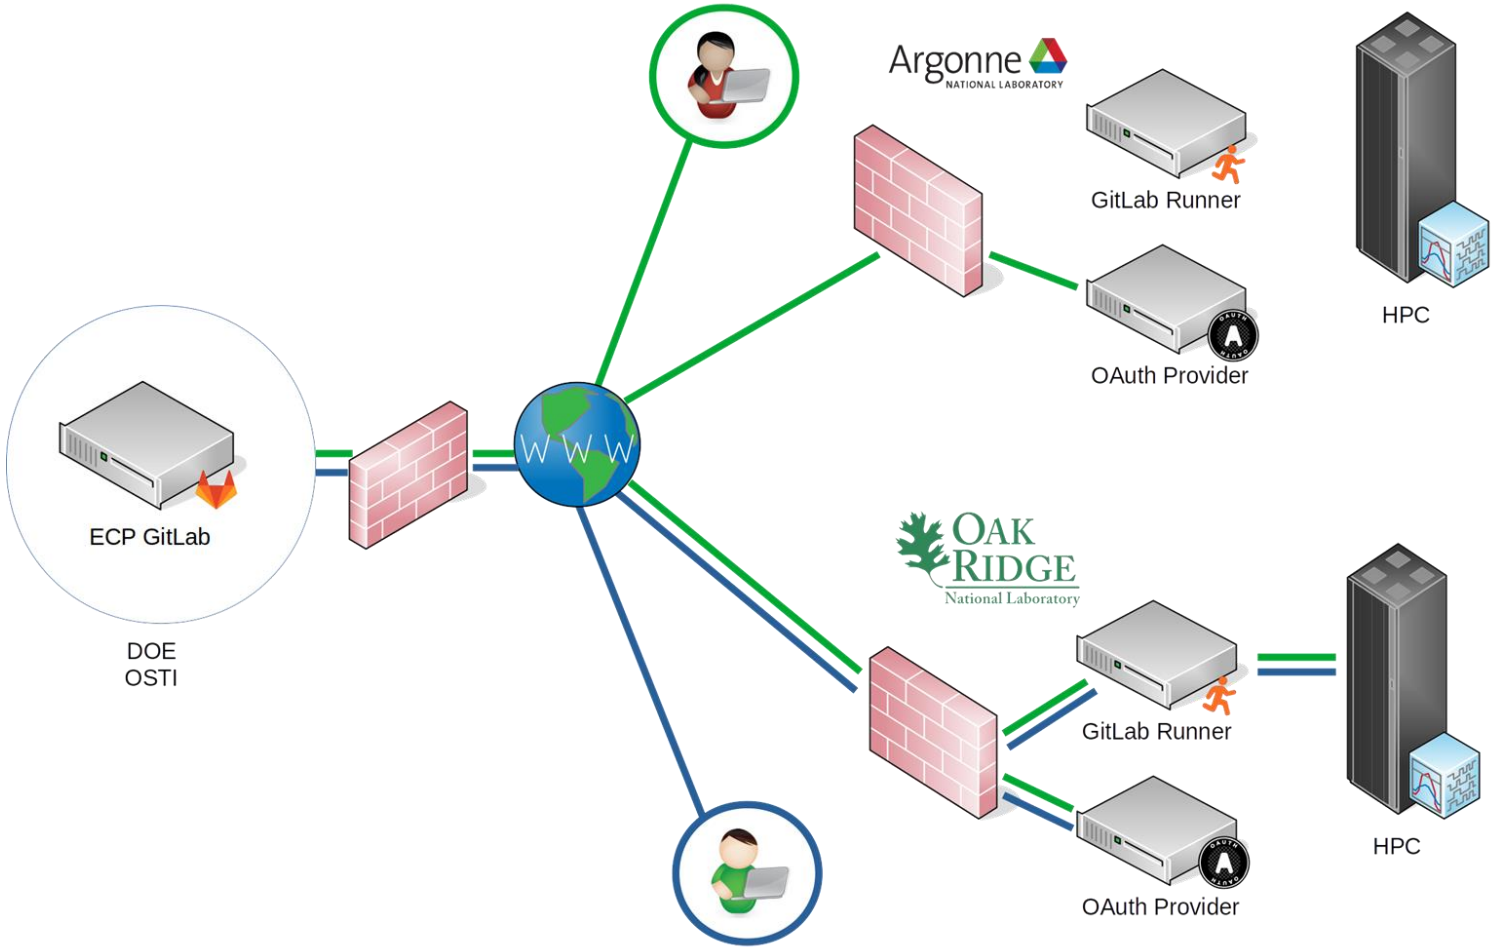

Accomplishments and plan for the next 18 months

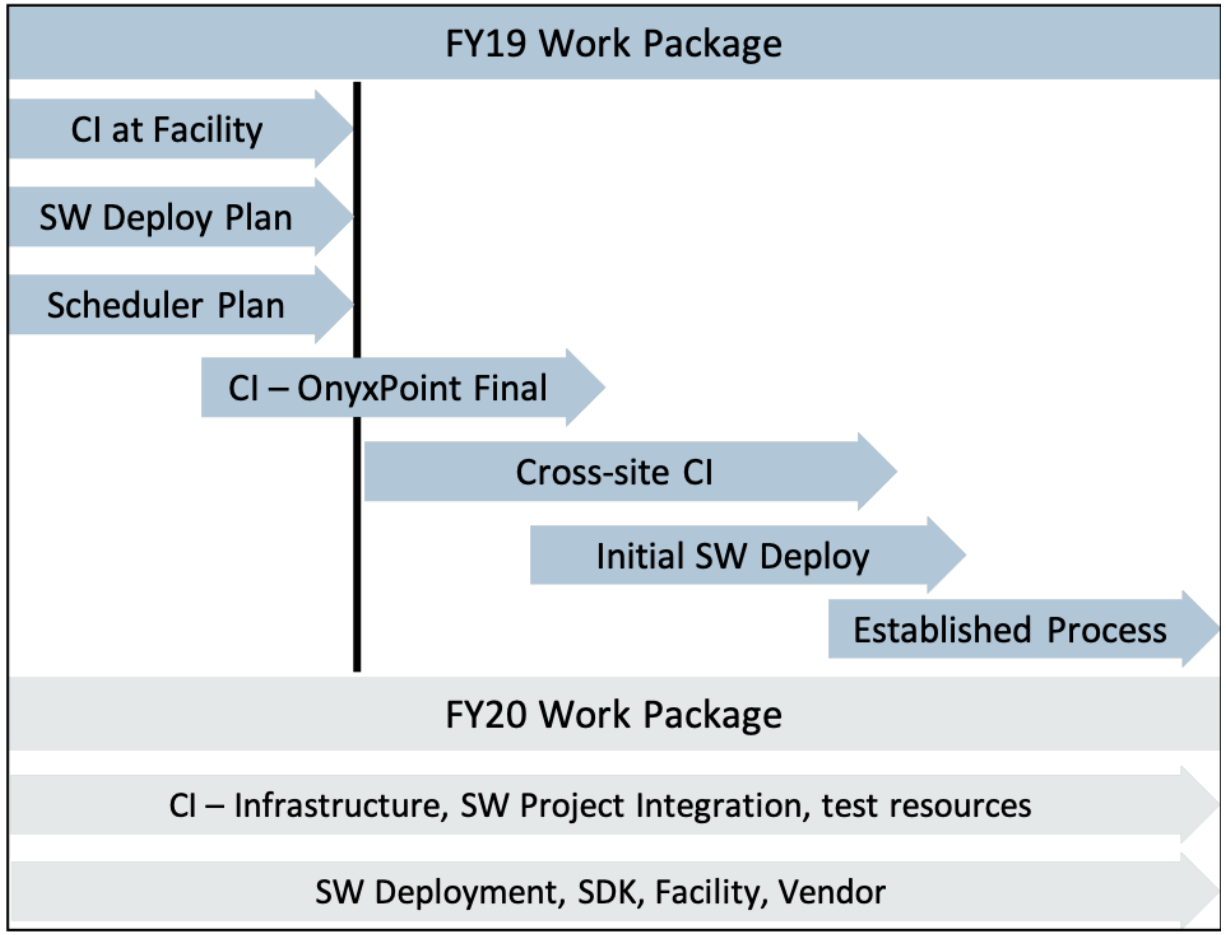

- Onyx Point progress
  - Security plan delivered and approved by 6 DOE labs for the design and implementation of a secure mechanism for launching runners
  - Above enhancement delivered and tested
  - Batch submission runner for HPC specific job management
  - In progress is a web user interface for the above capability
- An initial Continuous Integration capability (GitLab instance including the Onyx Point enhancement) was brought up and tested at two labs
- ECP ST SW projects were used to test the capability.

State of Deployments and Integration at DOE HPC Facilities

Facility	CI Resources	Teams/Products Integrating into CI
OLCF	Open science Gitlab has associated runners connected to the login node of the open HPC cluster Ascent, which is a power9 “summit-like” cluster.	ciminiqmc, ecp-copa/cabana, proxy-ci, laristra/flecsi, and datatransferkit
ALCF	Local GitLab server VM, has associated runner installed on test hpc service node for submitting jobs to the iota testing hpc.	Argo-AML, Datalib/Darshan, ECP-Proxy, and HDF5
NERSC	Local GitLab server VM, has assisted runner installed on a test cluster on Edison/Cori	HDF5 and ECP-Proxy
LLNL	Gitlab server running in production environment accessible to all Livermore Computing users (currently using administrative controls within Gitlab to restrict full access to a limited number of test users).	• Onyx Point runners with both "batch" and "shell" executors running on 14 login nodes of 2600 compute-node production cluster "quartz" using slurm scheduler. • Onyx Point runner with "batch" executor running on single login node of 18 compute-node Power9 test cluster "butte" using lsf scheduler.

Engagement Plans to Identify Formal Commitments

Facility Engagement Plans Define and Formalize the Partnership

- Plans were signed with Individual appendices for each SC facility (ALCF, NERSC, and OLCF)
- ECP's commitments are mapped to ECP milestones and are tracked
- Early March 2019, the facilities and ECP met for a one and a half day meeting for a check on the partnership and to update the appendices
- An additional goal was to work on a shared vision – my assessment is that much progress was made in answering these questions:
  - How will ECP's ST software be selected for deployment at a facility and how will the software be supported
  - How AD applications will demonstrate their final objectives on the exascale computers (how many runs, what types of runs, and when is the access needed)
  - When and how can AD and ST teams get access to early hardware for their development
- Outcomes were: draft undated appendices, plans to increase communications, and notes on the shared vision

Facilities and ECP Partnering is Key to ECI

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

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

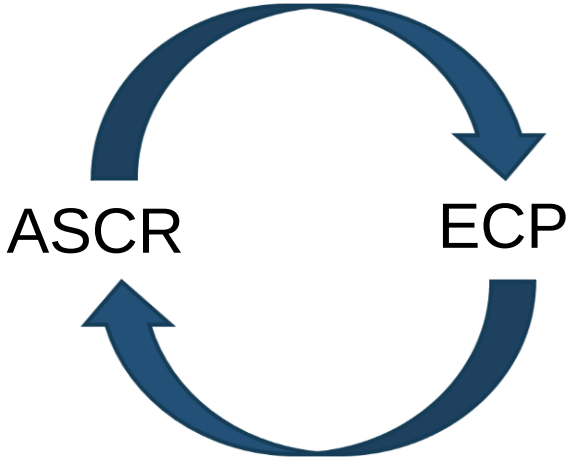

- Mission ends with the project
- ECP is a few years old

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