

ECP Software Technology Overview



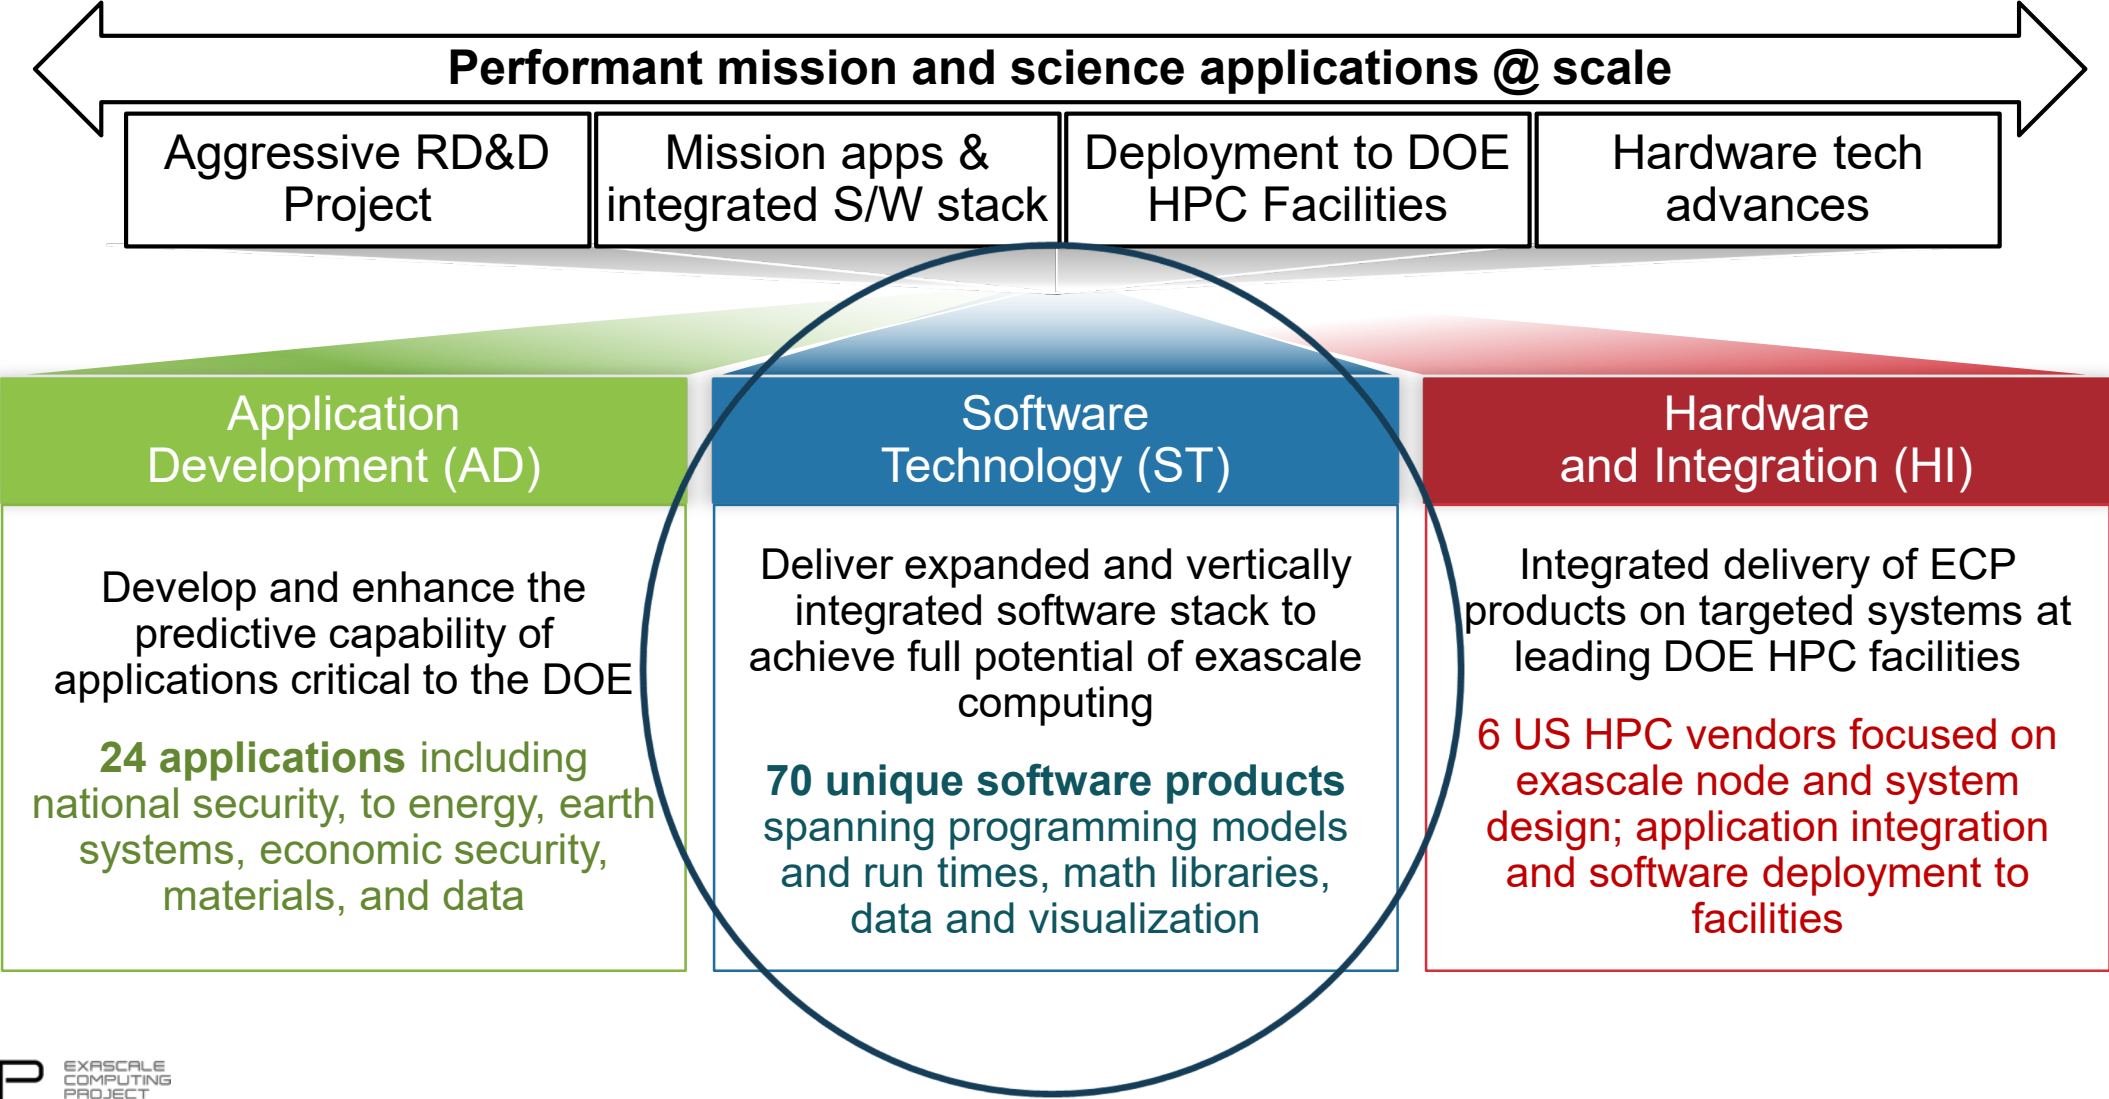
Michael A. Heroux, Sandia National Laboratories
Director of Software Technology

E4S at DOE Facilities with Deep Dive at NERSC, October 4, 2021




ECP Organizational Sketch



ECP Software Technology (ST) is one of three focus areas



ECP ST has six technical areas



Programming Models & Runtimes

- Enhance and get ready for exascale the widely used MPI and OpenMP programming models (hybrid programming models, deep memory copies)
- Development of performance portability tools (e.g. Kokkos and Raja)
- Support alternate models for potential benefits and risk mitigation: PGAS (UPC++/GASNet), task-based models (Legion, PaRSEC)
- Libraries for deep memory hierarchy and power management

Development Tools

- Continued, multifaceted capabilities in portable, open-source LLVM compiler ecosystem to support expected ECP architectures, including support for F18
- Performance analysis tools that accommodate new architectures, programming models, e.g., PAPI, Tau

Math Libraries

- Linear algebra, iterative linear solvers, direct linear solvers, integrators and nonlinear solvers, optimization, FFTs, etc
- Performance on new node architectures; extreme strong scalability
- Advanced algorithms for multi-physics, multiscale simulation and outer-loop analysis
- Increasing quality, interoperability, complementarity of math libraries

Data and Visualization

- I/O via the HDF5 API
- Insightful, memory-efficient in-situ visualization and analysis – Data reduction via scientific data compression
- Checkpoint restart

Software Ecosystem

- Develop features in Spack necessary to support all ST products in E4S, and the AD projects that adopt it
- Development of Spack stacks for reproducible turnkey deployment of large collections of software
- Optimization and interoperability of containers on HPC systems
- Regular E4S releases of the ST software stack and SDKs with regular integration of new ST products

NNSA ST

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

ECP Software Technology 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.



Lois Curfman McInnes, Software Technology Deputy Director

Lois is a senior computational scientist in the Mathematics and Computer Science Division of ANL. She has over 20 years of experience in HPC numerical software, including development of PETSc and leadership of multi-institutional work toward sustainable scientific software ecosystems.



Rajeev Thakur, Programming Models and Runtimes

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

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, such as heterogeneous computing and nonvolatile memory, for HPC for over 15 years.



Xaioye (Sherry) Li, Math Libraries

Sherry is a senior scientist at Berkeley Lab. She has over 20 years of experience in high-performance numerical software, including development of SuperLU and related linear algebra algorithms and software.



Jim Ahrens, Data and Visualization

Jim is a senior research scientist at the Los Alamos National Laboratory (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.



Todd Munson, Software Ecosystem and Delivery

Todd is a computational scientist in the Math and Computer Science Division of ANL. He has nearly 20 years of experience in high-performance numerical software, including development of PETSc/TAO and project management leadership in the ECP CODAR project.



Kathryn Mohror, NNSA ST

Kathryn is Group Leader for the CASC Data Analysis Group at LLNL. Her work focuses on I/O for extreme scale systems, scalable performance analysis and tuning, fault tolerance, and parallel programming paradigms. She is a 2019 recipient of the DOE Early Career Award.

ST L4 Teams

- WBS
- Name
- PIs
- PCs - Project Coordinators

ECP ST Stats

- 35 L4 subprojects
- 11 PI/PC same
- 24 PI/PC different
- ~27% ECP budget



WBS	WBS Name	CAM/PI	PC
2.3	Software Technology	Heroux, Mike, McInnes, Lois	
2.3.1	Programming Models & Runtimes	Thakur, Rajeev	
2.3.1.01	PMR SDK	Shende, Sameer	Shende, Sameer
2.3.1.07	Exascale MPI (MPICH)	Balaji, Pavan	Guo, Yanfei
2.3.1.08	Legion	McCormick, Pat	McCormick, Pat
2.3.1.09	PaRSEC	Bosilica, George	Carr, Earl
2.3.1.14	Pagoda: UPC++/GASNet for Lightweight Communication and Global Address Space Support	Hargrove, Paul	Hargrove, Paul
2.3.1.16	SICM	Lang, Michael	Vigil, Brittney
2.3.1.17	OMPI-X	Bernholdt, David	Grundhoffer, Alicia
2.3.1.18	RAJA/Kokkos	Trott, Christian Robert	Trujillo, Gabrielle
2.3.1.19	Argo: Low-level resource management for the OS and runtime	Beckman, Pete	Gupta, Rinku
2.3.2	Development Tools	Vetter, Jeff	
2.3.2.01	Development Tools Software Development Kit	Vetter, Barton	Tim Haines
2.3.2.06	Exa-PAPI++: The Exascale Performance Application Programming Interface with Modern C++	Dongarra, Jack	Jagode, Heike
2.3.2.08	Extending HPCToolkit to Measure and Analyze Code Performance on Exascale Platforms	Mellor-Crummey, John	Meng, Xiaozhu
2.3.2.10	PROTEAS-TUNE	Vetter, Jeff	Glassbrook, Dick
2.3.2.11	SOLLVE: Scaling OpenMP with LLVM for Exascale	Wachsmuth, Heesara	Kale, Vivek
2.3.2.12	FLANG	McCormick, Pat	Perry-Holby, Alexis
2.3.3	Mathematical Libraries	Li, Sherry	
2.3.3.01	Extreme-scale Scientific xSDK for ECP	Yang, Ulrike	Yang, Ulrike
2.3.3.06	Preparing PETSc/TAO for Exascale	Munson, Todd	Munson, Todd
2.3.3.07	STRUMPACK/SuperLU/FFTX: sparse direct solvers, preconditioners, and FFT libraries	Li, Sherry	Li, Sherry
2.3.3.12	Enabling Time Integrators for Exascale Through Sundials/ Hypr	Woodward, Carol	Woodward, Carol
2.3.3.13	CLOVER: Computational Libraries Optimized Via Exascale Research	Dongarra, Jack	Carr, Earl
2.3.3.14	ALExa: Accelerated Libraries for Exascale/ForTrilinos	Trujillo, Gabrielle	Grundhoffer, Alicia
2.3.3.15	Sake: Scalable Algorithms and Kernels for Exascale	Rajamanickam, Siva	Trujillo, Gabrielle
2.3.4	Data and Visualization	Ahrens, James	
2.3.4.01	Data and Visualization Software Development Kit	Ahrens, James	Bagha, Neelam
2.3.4.09	ADIOS Framework for Scientific Data on Exascale Systems	Grundhoffer, Alicia	Grundhoffer, Alicia
2.3.4.10	DataLib: Data Libraries and Services Enabling Exascale Science	Ross, Rob	Ross, Rob
2.3.4.13	ECP/VTK-m	Moreland, Kenneth	Moreland, Kenneth
2.3.4.14	VeloC: Very Low Overhead Transparent Multilevel Checkpoint/Restart	Edging, Scott	Edging, Scott
2.3.4.15	ExaIO - Delivering Efficient Parallel I/O on Exascale Computing Systems with HDF5 and Chimera	Bagha, Neelam	Bagha, Neelam
2.3.4.16	ALPINE: Algorithms and Infrastructure for In Situ Visualization and Analysis/ZFP	Turton, Terry	Turton, Terry
2.3.5	Software Ecosystem and Delivery	Munson, Todd	
2.3.5.01	Software Ecosystem and Delivery Software Development Kit	Munson, Todd	Munson, Todd
2.3.5.09	SW Packaging Technologies	Gamblin, Todd	Gamblin, Todd
2.3.5.10	ExaWorks	Laney, Dan	Laney, Dan
2.3.6	NNSA ST	Mohror, Kathryn	
2.3.6.01	LANL ATDM	Mike Lang	Vandenbusch, Tanya Marie
2.3.6.02	LLNL ATDM	Becky Springmeyer	Gamblin, Todd
2.3.6.03	SNL ATDM	Jim Stewart	Trujillo, Gabrielle

• ~250 staff
• ~70 products
• 34 teams
• ~30 universities
• ~9 DOE labs
• 6 technical areas
• 1 focus area of 3 in ECP

We work on products applications need now and into the future

Key themes:

- Focus: GPU node architectures and advanced memory & storage technologies
- Create: New high-concurrency, latency tolerant algorithms
- Develop: New portable (Nvidia, Intel, AMD GPUs) software product
- Enable: Access and use via standard APIs

Software categories:

- **Next generation established products:** Widely used HPC products (e.g., MPICH, OpenMPI, PETSc)
- **Robust emerging products:** Address key new requirements (e.g., Kokkos, RAJA, Spack)
- **New products:** Enable exploration of emerging HPC requirements (e.g., SICM, zfp, UnifyCR)

Example Products	Engagement
MPI – Backbone of HPC apps	Explore/develop MPICH and OpenMPI new features & standards
OpenMP/OpenACC –On-node parallelism	Explore/develop new features and standards
Performance Portability Libraries	Lightweight APIs for compile-time polymorphisms
LLVM/Vendor compilers	Injecting HPC features, testing/feedback to vendors
Perf Tools - PAPI, TAU, HPCToolkit	Explore/develop new features
Math Libraries: BLAS, sparse solvers, etc.	Scalable algorithms and software, critical enabling technologies
IO: HDF5, MPI-IO, ADIOS	Standard and next-gen IO, leveraging non-volatile storage
Viz/Data Analysis	ParaView-related product development, node concurrency

One example: SLATE port to AMD and Intel platforms

Scope and objectives

- SLATE is a distributed, GPU-accelerated, dense linear algebra library, intended to replace ScaLAPACK
- SLATE covers parallel BLAS, linear system solvers, least squares, eigensolvers, and the SVD

Impact

- Initially supported NVIDIA's cuBLAS for use on current machines like Summit
- Can now use AMD's rocBLAS in preparation for Frontier, and Intel's oneMKL in preparation for Aurora
- Other projects can also leverage BLAS++ for portability

Deliverables Report: <https://www.icl.utk.edu/publications/swan-016>
Code in git repos: bitbucket.org/icl/slate/ and bitbucket.org/icl/blaspp/

Port to AMD and Intel

- SLATE and BLAS++ now support all three major GPU platforms



Accomplishment

- Refactored SLATE to use BLAS++ as portability layer
- Ported BLAS++ to AMD rocBLAS and Intel oneMKL

Key ECP Software Stack Legacy:

- Portable execution on:
 - CPUs
 - 3 different GPUs
- A bridge from CPUs to GPUs

E4S Planning, Executing, Delivering



ECP ST Planning Process: Hierarchical, three-phase, cyclical


Baseline



FY20–23 Baseline Plan High level Definitions

- Q2 FY19 start
- FY20 Base plan
- FY21–23 planning packages

Annual Refinement



FY Refine Baseline Plan As Needed Basic activity definitions

- 6 months prior to FY
- 4–6 P6 Activities/year
- Each activity:
 - % annual budget
 - Baseline start/end
 - High level description

Per Activity



Detailed Plan Complete activity definitions

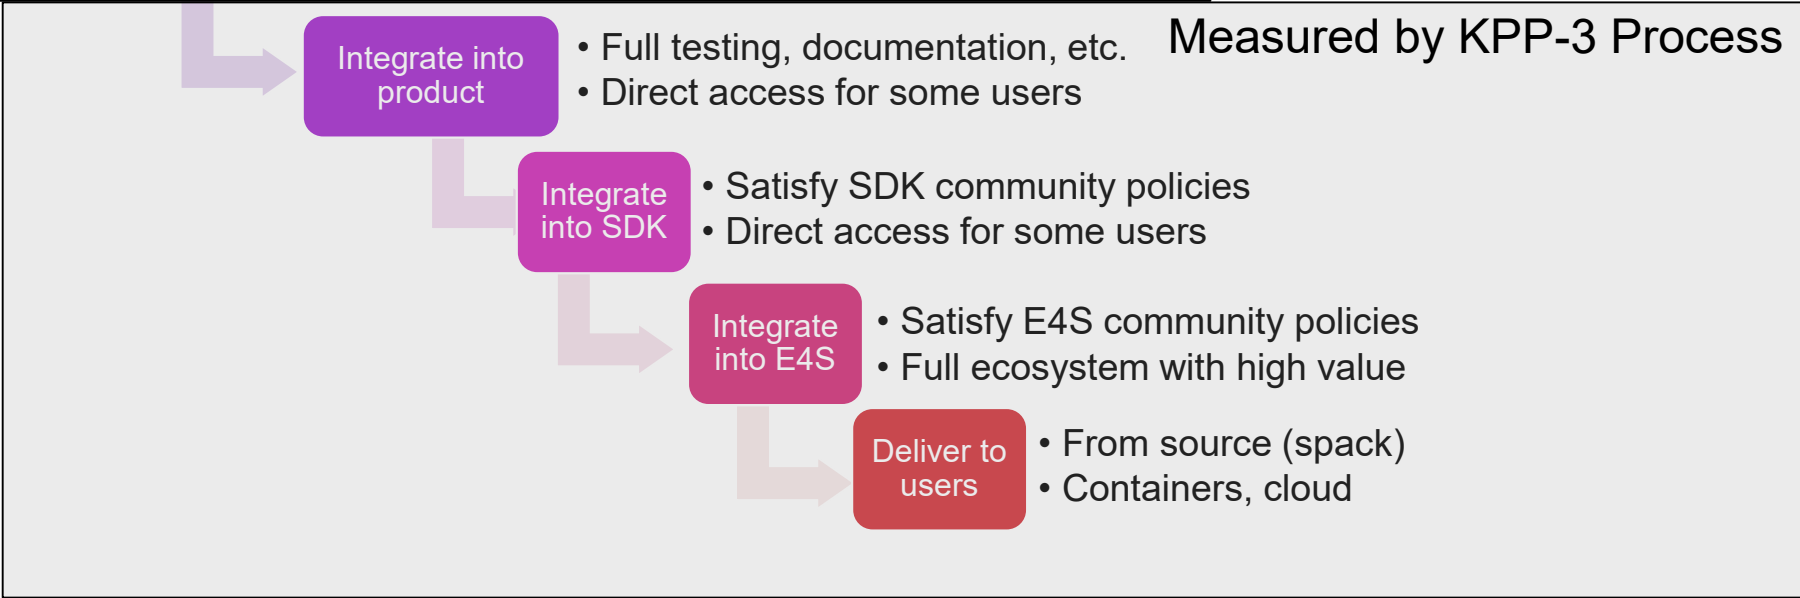
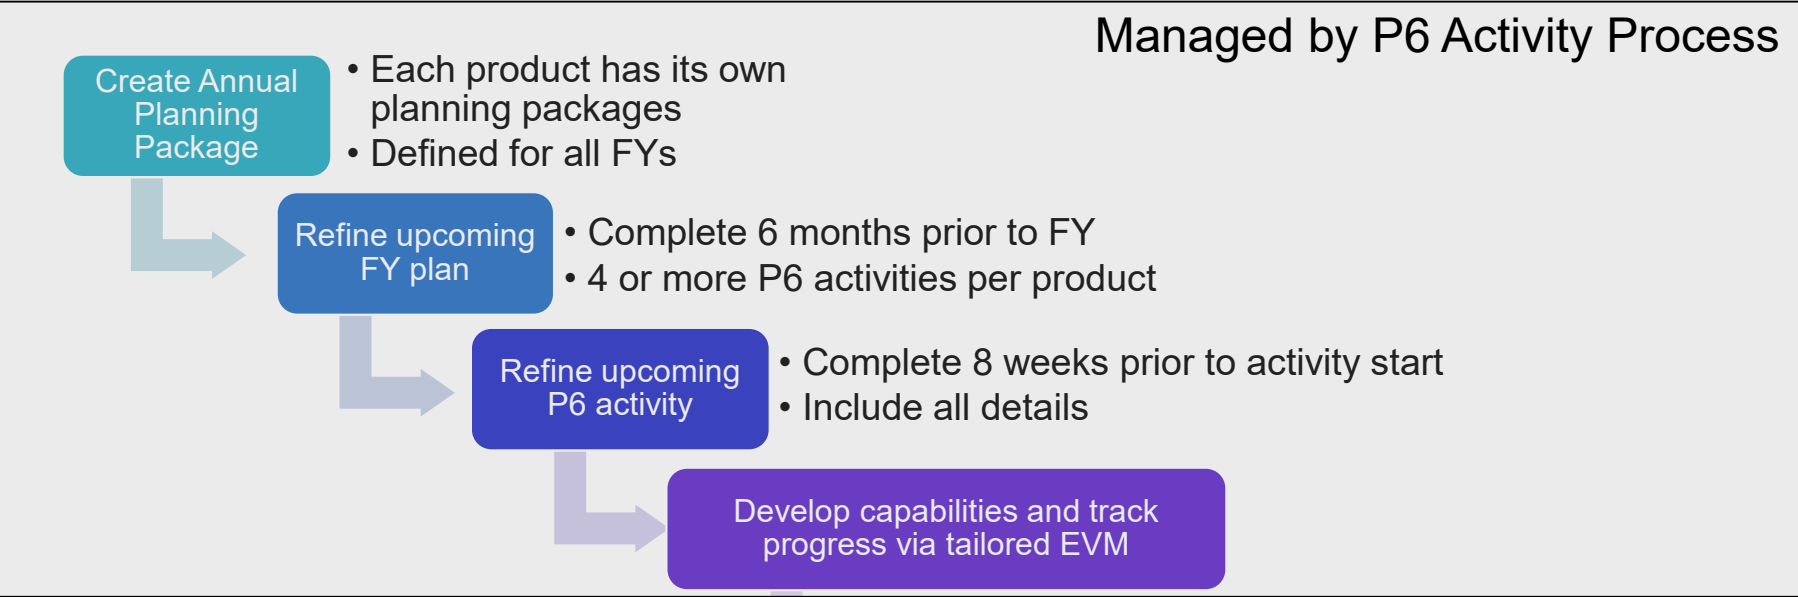
- 8 weeks prior to start
- High-fidelity description
- Execution strategy
- Completion criteria
- Personnel details

Two-level Review Process	
Changes to Cost, Scope, and Schedule	
Minor	Major
Lightweight Review in Jira, L3 and L2 leads	Change Control Board Review, ECP leadership
Variance Recorded in Jira Proceed with Execution	

KPP-3: Focus on capability integration

- **Capability:** Any significant product functionality, including existing features adapted to the pre-exascale and exascale environments, that can be integrated into a client environment.
- **Capability Integration:** Complete, sustainable integration of a significant product capability into a client environment in a pre-exascale environment (tentative score) and in an exascale environment (confirmed score).

ECP ST Lifecycle summary



The Extreme-Scale Scientific Software Stack (E4S) and Software Development Kits (SDKs)



Extreme-scale Scientific Software Stack (E4S)



- E4S: HPC Software Ecosystem – a curated software portfolio
- A **Spack-based** distribution of software tested for interoperability and portability to multiple architectures
- Available from **source, containers, cloud, binary caches**
- Leverages and enhances SDK interoperability thrust
- Not a commercial product – an open resource for all
- Oct 2018: E4S 0.1 - 24 full, 24 partial release products
- Jan 2019: E4S 0.2 - 37 full, 10 partial release products
- Nov 2019: E4S 1.0 - 50 full, 5 partial release products
- Feb 2020: E4S 1.1 - 61 full release products
- Nov 2020: E4S 1.2 (aka, 20.10) - 67 full release products
- Feb 2021: E4S 21.02 - 67 full release, 4 partial release
- May 2021: E4S 21.05 - 76 full release products
- August 2021: E4S 21.08 - 88 full release products



<https://e4s.io>

Lead: Sameer Shende
(U Oregon)



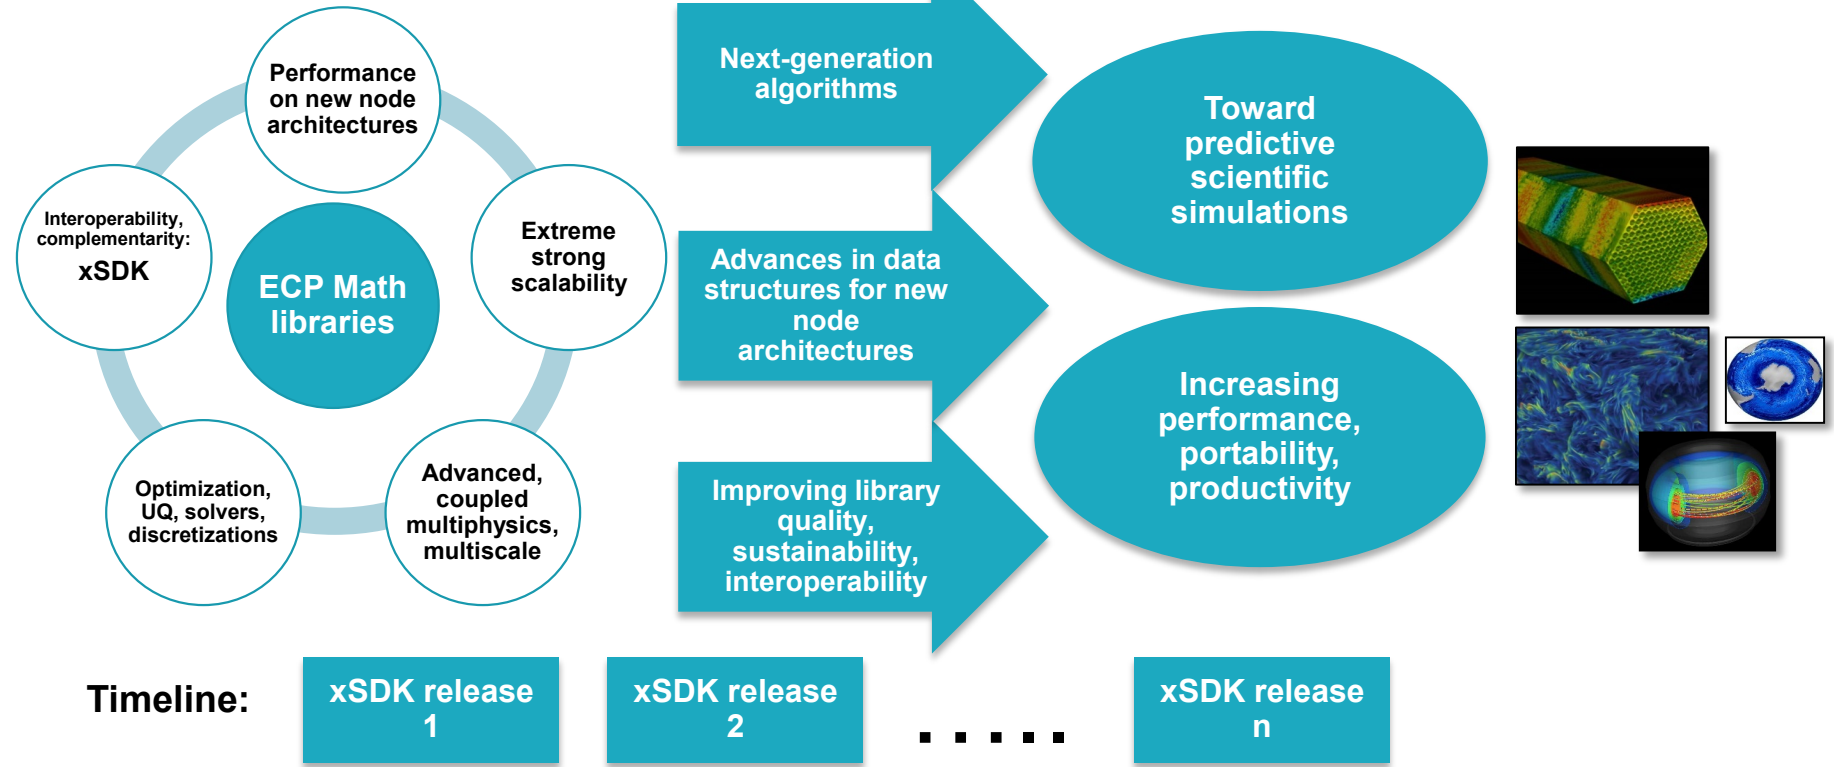
xSDK: Primary delivery mechanism for ECP math libraries' continual advancements toward predictive science

xSDK release 0.6.0 (Nov 2020)

hypr
PETSc/TAO
SuperLU
Trilinos
AMReX
ButterflyPACK
DTK
Ginkgo
heFFTe
libEnsemble
MAGMA
MFEM
Omega_h
PLASMA
PUMI
SLATE
Tasmanian
SUNDIALS
Strumpack
Alquimia
PFLOTRAN
deal.II
preCICE
PHIST
SLEPc

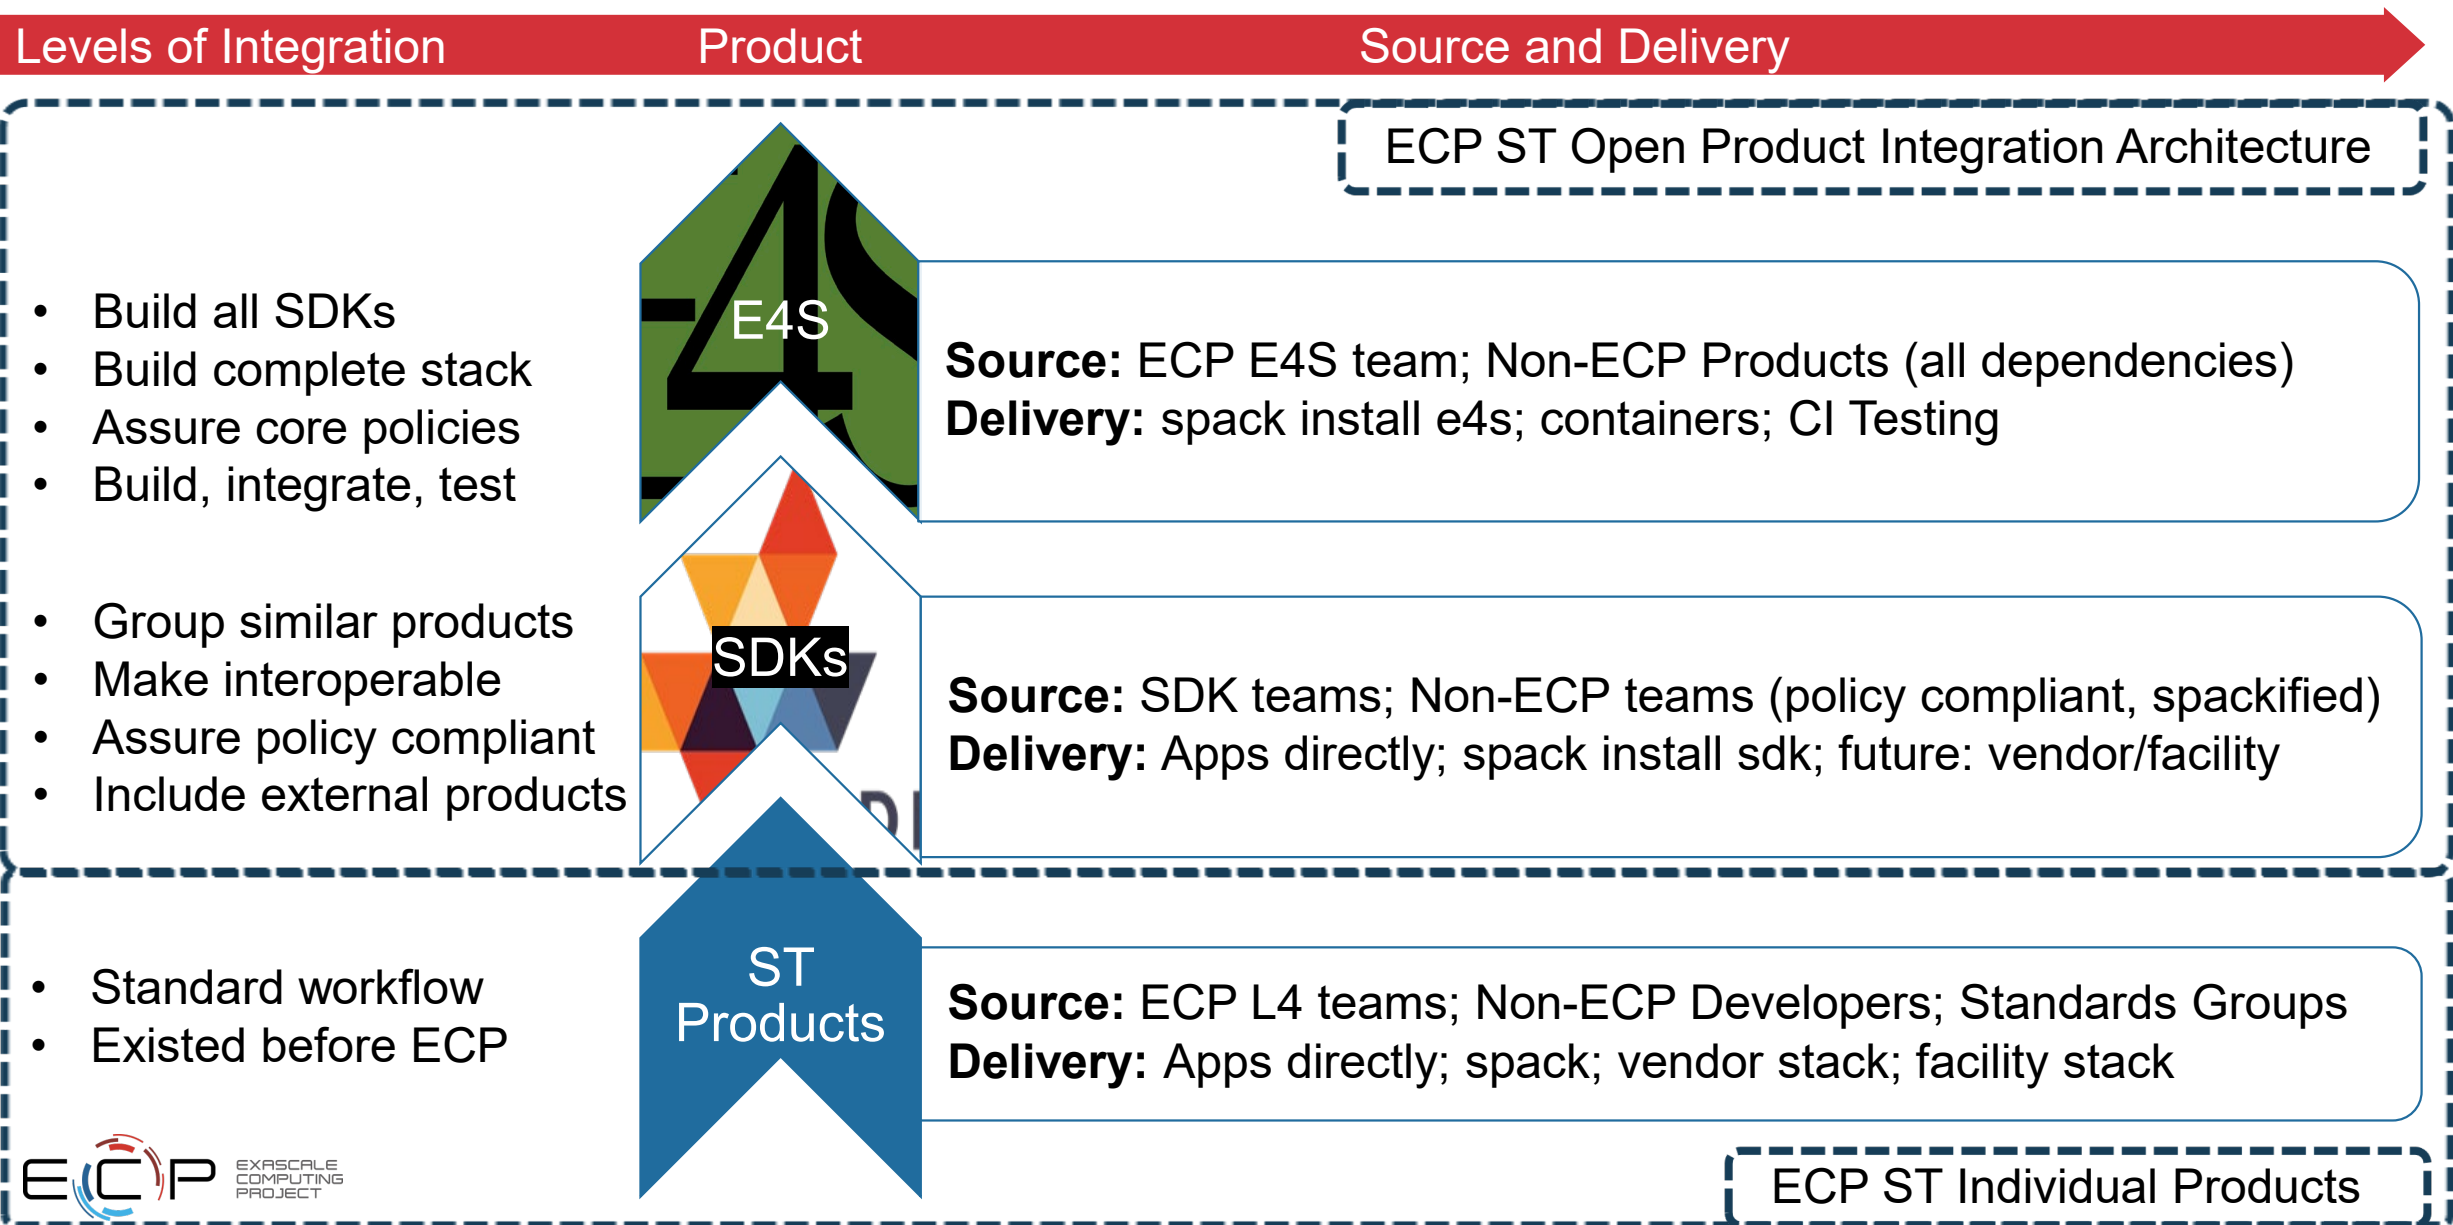
} from the
broader
community

As motivated and validated by
the needs of ECP applications:



Delivering an open, hierarchical software ecosystem


More than a collection of individual products



E4S Community Policies



E4S Community Policies V1.0 Released



What is E4S?

The Extreme-scale Scientific Software Stack (E4S) is a community effort to provide open source software packages for developing, deploying and running scientific applications on high-performance computing (HPC) platforms. E4S provides from-source builds and containers of a **broad collection of HPC software packages**.



Purpose

E4S exists to accelerate the development, deployment and use of HPC software, lowering the barriers for HPC users. E4S provides containers and turn-key, from-source builds of more than 80 popular HPC products in programming models, such as MPI; development tools such as HPCToolkit, TAU and PAPI; math libraries such as PETSc and Trilinos; and Data and Viz tools such as HDF5 and Paraview.



Approach

By using Spack as the meta-build tool and providing containers of pre-built binaries for Docker, Singularity, Shifter and CharlieCloud, E4S enables the flexible use and testing of a **large collection of reusable HPC software packages**.

E4S Community Policies Version 1

A Commitment to Quality Improvement

- Will serve as membership criteria for E4S
 - Membership is not required for *inclusion* in E4S
 - Also includes forward-looking draft policies
- Purpose: enhance sustainability and interoperability
- Topics cover building, testing, documentation, accessibility, error handling and more
- Multi-year effort led by SDK team
 - Included representation from across ST
 - Multiple rounds of feedback incorporated from ST leadership and membership
- Modeled after xSDK Community Policies
- <https://e4s-project.github.io/policies.html>

P1 *Spack-based Build and Installation* Each E4S member package supports a scriptable *Spack* build and production-quality installation in a way that is compatible with other E4S member packages in the same environment. When E4S build, test, or installation issues arise, there is an expectation that teams will collaboratively resolve those issues.

P2 *Minimal Validation Testing* Each E4S member package has at least one test that is executable through the E4S validation test suite (<https://github.com/E4S-Project/testsuite>). This will be a post-installation test that validates the usability of the package. The E4S validation test suite provides basic confidence that a user can compile, install and run every E4S member package. The E4S team can actively participate in the addition of new packages to the suite upon request.

P3 *Sustainability* All E4S compatibility changes will be sustainable in that the changes go into the regular development and release versions of the package and should not be in a private release/branch that is provided only for E4S releases.

P4 *Documentation* Each E4S member package should have sufficient documentation to support installation and use.

P5 *Product Metadata* Each E4S member package team will provide key product information via metadata that is organized in the *E4S DocPortal* format. Depending on the filenames where the metadata is located, this may require *minimal setup*.

P6 *Public Repository* Each E4S member package will have a public repository, for example at GitHub or Bitbucket, where the development version of the package is available and pull requests can be submitted.

P7 *Imported Software* If an E4S member package imports software that is externally developed and maintained, then it must allow installing, building, and linking against a functionally equivalent outside copy of that software. Acceptable ways to accomplish this include (1) forsaking the internal copied version and using an externally-provided implementation or (2) changing the file names and namespaces of all global symbols to allow the internal copy and the external copy to coexist in the same downstream libraries and programs. This pertains primarily to third party support libraries and does not apply to key components of the package that may be independent packages but are also integral components to the package itself.

P8 *Error Handling* Each E4S member package will adopt and document a consistent system for signifying error conditions as appropriate for the language and application. For e.g., returning an error condition or throwing an exception. In the case of a command line tool, it should return a sensible exit status on success/failure, so the package can be safely run from within a script.

P9 *Test Suite* Each E4S member package will provide a test suite that does not require special system privileges or the purchase of commercial software. This test suite should grow in its comprehensiveness over time. That is, new and modified features should be included in the suite.

E4S DocPortal



E4S DocPortal

- Single point of access
- All E4S products
- Summary Info
 - Name
 - Functional Area
 - Description
 - License
- Searchable
- Sortable
- Rendered daily from repos

E4S Products

*: Member Product

Show 10 entries

Search:

	Name	Area	Description	
+	ADIOS2	Data & Viz	I/O and data management library for storage I/O, in-memory code coupling and online data analysis and visualization workflows.	2021-03-10 16:45:25
+	AML	PMR	Hierarchical memory management library from Argo.	2019-04-25 13:03:01
+	AMREX	PMR	A framework designed for building massively parallel block- structured adaptive mesh refinement applications.	2021-05-02 17:26:43
+	ARBORX	Math libraries	Performance-portable geometric search library	2021-01-05 15:39:55
+	ARCHER			
+	ASCENT			
+	BEE	Software Ecosystem	Container-based solution for portable build and execution across HPC systems and cloud resources	2018-08-22 22:26:19
+	BOLT	Development Tools	OpenMP over lightweight threads.	2020-05-04 11:24:57
+	CALIPER	Development tools	Performance analysis library.	2020-11-04 23:53:07
+	CHAI	PMR	A library that handles automatic data migration to different memory spaces behind an array-style interface.	2020-11-02 19:58:24

Name

<https://e4s-project.github.io/DocPortal.html>

Latest Doc Update

Showing 1 to 10 of 76 entries

Previous

1

2

3

4

5

...

8

Next

All we need from the software team is a repo URL + up-to-date meta-data files

Goal: All E4S product documentation accessible from single portal on E4S.io (working mock webpage below)

CS373 Issues

Database Projects

Virtual Events

Reproducibility

Save to Instapaper

2019 Spring CS317

Monte Carlo

Design for Manufacturing

Save to Mendeley

GitHub policies

Machine Learning for PDEs

E4S

HOME

EVENTS

ABOUT

DOCPORTAL

CONTACT US

FAQ

DOWNLOAD

* Member Product

Show 10 entries

Name	Area
ADIOS2	Data & Viz
AML	PMR
ARCHER	Tools
ASCENT	Data & Viz
BEE	Software ecosystem
BOLT	Development tools
CALIPER	Development tools
CHAI	PMR
CINEMA	Data & Viz
DARSHAN	Data & Viz

Showing 1 to 10 of 75 entries

E4S

HOME

EVENTS

ABOUT

DOCPORTAL

CONTACT US

FAQ

DOWNLOAD

E4S Products

* Member Product

Show 10 entries

Search:

Name	Area	Description
ADIOS2	Data & Viz	I/O and data management library for storage I/O, in-memory code coupling and online data analysis and visualization workflows.

Description: The Adaptable Input Output System version 2, developed in the Exascale Computing Program

Homepage: <https://csmd.ornl.gov/software/adios>

Document Summaries

ReadMe.md

License: Apache 2.0

docs: passing

release: v2.6.0

build: v2.6.0

More...

LICENSE

Apache License
Version 2.0, January 2004
<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

More...

CS373 Issues

Database Projects

Virtual Events

Reproducibility

Save to Instapaper

2019 Spring CS317

Monte Carlo

Design for Manufacturing

Save to Mendeley

GitHub policies

Machine Learning for PDEs

OAK RIDGE
National Laboratory

Home

About

Research

Software


Groups

Jobs

Computer Science and Mathematics

Software

ADIOS2




ADIOS 2: The Adaptable Input Output (I/O) System version 2 is an open-source framework that addresses scientific data management challenges, e.g. scalable parallel I/O, as we approach the exascale era in high-performance computing (HPC). ADIOS 2 bindings are available in C++, C, Fortran, Python and can be used on supercomputers, personal computers, and cloud systems running on Linux, macOS and Windows. ADIOS 2 has out-of-the-box support for MPI and serial environments.


ADIOS 2 unified application programming interface (API) focuses on what scientific applications produce and consume in terms of n-dimensional Variables, Attributes, and Steps, while hiding the low-level details of how the data byte streams are transported as efficiently as possible from application memory to HPC networks, files, wide-area-networks, and direct memory access media. Typical use cases include file storage for checkpoint-restart and analysis, data streaming for code-coupling, and in situ analysis and visualization workflows. ADIOS 2 also provides high-level APIs that resemble native I/O libraries in Python (file) and C++ (fstream) for easy integration with their rich data analysis ecosystems. In addition, XML and YAML runtime configuration files are provided so users can fine tune available parameters to enable efficient data movements without recompiling their codes. ADIOS 2 also supports data compression via third party libraries for lossy: zfp, SZ, MGARD, and lossless: blosc, bzip2, png operations.

The ADIOS 2 development process adopts modern software engineering practices such as unit testing, continuous integration, and documentation to make the final product accessible to the scientific community. Our commitment is to release a new version every 6 months. Distributions are currently available via modern package management systems: conda, spack, homebrew (and more to come). Overall, applications using ADIOS 2 do not need to dramatically modify their source code to evaluate I/O performance trade-offs, thus reducing integration and maintenance costs in their development process. For those coming from ADIOS 1.x, ADIOS 2


ORNL Researchers




Nicholas Thompson




Norbert Podhorszki




William Godoy



Scott Klasky



Lipeng Wan



Jeremy Logan

1 2 Last

Other Researchers

Kesheng (John) Wu
Lawrence Berkeley National Laboratory
0000-0002-6907-3393

Greg Eisenhauer
Georgia Institute of Technology
0000-0002-2070-043X

Manish Parashar
Rutgers University
0000-0003-0983-7408

Philip Davis
Rutgers University
0000-0002-2205-8268


Group
Scientific Data Group

E4S: Better quality, documentation, testing, integration, delivery, building & use


Delivering HPC software to facilities, vendors, agencies, industry, international partners in a brand-new way



Community Policies
Commitment to software quality



DocPortal
Single portal to all E4S product info



Portfolio testing
Especially leadership platforms



Curated collection
The end of dependency hell



Quarterly releases
Release 1.2 – November



Build caches
10X build time improvement



Turnkey stack
A new user experience



<https://e4s.io>



Community & LSSw
US agencies, industry, international

Growing and Sustaining the Software Community



IDEAS-ECP team works with the ECP community to improve developer productivity and software sustainability as key aspects of increasing overall scientific productivity.

- 1 Customize and curate methodologies**
 - Target scientific software productivity and sustainability
 - Use workflow for best practices content development

- 2 Incrementally and iteratively improve software practices**
 - Determine high-priority topics for improvement and track progress
 - *Productivity and Sustainability Improvement Planning (PSIP)*



- 3 Establish software communities**
 - Determine community policies to improve software quality and compatibility
 - Create Software Development Kits (SDKs) to facilitate the combined use of complementary libraries and tools

- 4 Engage in community outreach**
 - Broad community partnerships
 - Collaboration with computing facilities
 - Webinars, tutorials, events
 - *WhatIs* and *HowTo* docs
 - Better Scientific Software site (<https://bssw.io>)

BSSw Fellowship: Meet the Fellows

<https://bssw.io/fellowship>

Meet Our Fellows

The BSSw Fellowship program gives recognition and funding to leaders and advocates of high-quality scientific software. Meet the Fellows and Honorable Mentions and learn more about how they impact Better Scientific Software.

[Fellowships Overview](#)[Apply](#)[Meet Our Fellows](#)[BSSw Fellowship FAQ](#)

Community Growth

2018 - 2021

2018 Class

Fellows



Jeffrey Carver
University of Alabama
Improving code quality through modern peer code review



Ivo Jimenez
University of California, Santa Cruz
Enabling reproducible research through automated computational experimentation



Daniel S. Katz
University of Illinois at Urbana-Champaign, National Center for Supercomputing Applications
Giving software developers long-overdue credit through principles for software citation



Andrew Lumadine
Pacific Northwest National Laboratory, University of Washington, Northwest Institute for Advanced Computing
Building efficient use of modern C++ for high-performance computing

Honorable Mentions



Neal Davis
University of Illinois at Urbana-Champaign
Teaching Assistant Professor, Computer Science



Marc Henry de Frahan
National Renewable Energy Laboratory
Postdoctoral Researcher



Elsa Gonsiorowski
Lawrence Livermore National Laboratory
HPC VLSI Specialist, Livermore Computing



Ying Li
Argonne National Laboratory
Argonne Scholar, Argonne Leadership Computing Facility

2019 Class

Fellows



Rene Gasmoeiller
University of California, Davis
Guiding your scientific software project from inception to long-term sustainability



Ignacio Laguna
Lawrence Livermore National Laboratory
Improving the reliability of scientific applications by analyzing and debugging floating-point software



Tanu Malik
DePaul University
Reducing technical debt in scientific software through reproducible containers



Kyle Niemeyer
Oregon State University
Educating scientists on best practices for developing research software

Honorable Mentions



Stephen Andrews
Los Alamos National Laboratory
Staff Scientist, XCP-R Verification and Analysis



Nasir Eisty
University of Alabama
Ph.D. Student, Computer Science



Benjamin Pritchard
Virginia Tech
Software Scientist, Molecular Sciences Software Institute



Vanessa Sochat
Stanford University
Research Software Engineer, Stanford Research Computing Center

2020 Class

Fellows



Nasir Eisty
University of Alabama
Automating testing in scientific software



Damian Rouson
Sustainable Horizons Institute, Sorceury Institute
Introducing agile scientific software development to underrepresented groups



Cindy Rubio-Gonzalez
University of California, Davis
Improving the reliability and performance of numerical software

Honorable Mentions



David Boehme
Lawrence Livermore National Laboratory
Research Staff, Center for Applied Scientific Computing




Sumana Harihareswara
Changest Consulting
Founder and Principal, Open source software management and collaboration




David Rogers
National Center for Computational Sciences, Oak Ridge National Lab
Computational Scientist

2021 Class


Fellows




Marisol García-Reyes
Farallon Institute
Increasing accessibility of data & cloud technologies



Mary Ann Leung
Sustainable Horizons Institute
Increasing developer productivity and innovation through diversity




Chase Million
Million Concepts
Project management best practices for research software




Amy Roberts
University of Colorado Denver
Enabling collaboration through version control user stories


Honorable Mentions




Keith Beattie
Lawrence Berkeley National Laboratory
Computational Research Division, Computer Systems Engineer



Julia Stewart Lowndes
National Center for Ecological Analysis and Synthesis (NCEAS), UC Santa Barbara
Openscapes Director



Jonathan Madsen
Lawrence Berkeley National Laboratory
NERSC, Application Performance Specialist



Addi Thakur Malviya
Oak Ridge National Laboratory
Software Engineering Group, Group Leader

Advancing Scientific Productivity through Better Scientific Software: Developer Productivity & Software Sustainability Report

Disruptive changes in computer architectures and the complexities of tackling new frontiers in extreme-scale modeling, simulation, and analysis present daunting challenges to software productivity and sustainability.

This report explains the IDEAS approach, outcomes, and impact of work (in partnership with the ECP and broader computational science community).

Target readers are all those who care about the quality and integrity of scientific discoveries based on simulation and analysis. While the difficulties of extreme-scale computing intensify software challenges, issues are relevant across all computing scales, given universal increases in complexity and the need to ensure the trustworthiness of computational results.



<https://exascaleproject.org/better-scientific-productivity-through-better-scientific-software-the-ideas-report>

Preparing for Sustainable Software Efforts after ECP: Leadership Scientific Software (LSSw.io)



Background

- US Department of Energy (DOE) Exascale Computing Project ([ECP](#))
 - Developing enabling technologies for upcoming exascale computers
 - ECP Software Technology (ST) focus area:
 - Uses a macro-engineering software lifecycle to
 - Plan, execute, track, and assess product development toward the
 - Delivery of a curated portfolio of reusable, open-source software products called
 - The Extreme-scale Scientific Software Stack or E4S (<https://e4s.io>)
- During the final years of ECP, one key objective is to:
 - Transition our efforts to a sustainable organization and model for
 - Continued development and delivery of future capabilities, including
 - Incorporation of new scientific software domains, and
 - Expansion of the contributor and user communities
- LSSw is key component toward sustainability

LSSw Mission

- LSSw is dedicated to
 - Building community and understanding around the
 - Development and sustainable delivery of
 - Leadership scientific software
- Development
 - Portfolio-driven approach
 - Co-design with hardware, system software, applications
- Sustainable
 - Organizational stability
 - Emphasis on quality
 - Broad accessibility

Leadership Scientific Software (defn)

- Libraries, tools and environments that
 - Contribute to scientific discovery and insight in
 - New and emerging computing environments
- Are end-user applications within scope?
 - Yes, as stakeholders in the effort
 - Goal is to provide
 - High-priority functionality not available elsewhere
 - Portable performance on leading edge and emerging platforms
 - A sustainable turnkey software ecosystem

Leadership Scientific Software (defn)

- Push the boundary of feasibility
 - Enabling
 - Larger scale, higher fidelity and greater integration of
 - Advanced computing ecosystems
- Does “leadership” limit the scope of discussion?
 - Yes, we are directly focused on non-commodity environments, but:
 - Still use laptops, desktops, CPU clusters as part of our development efforts
 - Many of our tools and libraries need to be available everywhere
 - Non-commodity focus does not mean we work only on non-commodity systems
- Focus is on efforts that include co-design of
 - **Computing platforms:** Modeling & simulation, AI/ML, edge: at scale
 - **System software:** Collaborative co-design with vendors
 - **Science-specific tools and libraries:** What we are developing for users

ECP Efforts

- ECP is a notable project:
 - Stable, sustained funding of a national project with clear goals
 - Infrastructure to innovate and establish new collaborative work
- ECP enables tremendous opportunities to:
 - Create a new generation of scientific software
 - Provide a curated portfolio of reusable software products for apps
 - Qualitatively change how we plan, develop and deliver leadership SW

Join the conversation

- <https://lssw.io>: Main portal for the LSSw community
- LSSw Town Hall Meetings:
 - 3rd Thursday each month, 3 – 4:30 pm Eastern US time
- Slack: Share your ideas interactively
- White Papers: Written content for LSSw conversations
 - We need your ideas
 - 2 – 4 page white paper
 - Submit via GitHub PR or attachment to contribute@lssw.io
- References:
 - Help us build a reading list
 - Submit via GitHub PR or email to contribute@lssw.io


Summary & Next Steps

- Scientific software capabilities and complexity are increasing
- Computing systems are becoming more diverse
- A portfolio approach to planning and delivering is attractive
- ECP provides a working example to address complexity:
 - ECP ST lifecycle enables coordinated planning, executing, tracking and assessing
 - E4S and SDKs provide a scalable software architecture and portfolio for “turnkey” software stack
 - The IDEAS project and BSSw provide community building for scientific software developers
 - Goal: Better, faster and cheaper
- We believe the next steps require broad community engagement:
 - What are other fundamental requirements for improving leadership scientific software?
 - How can we collaborate as a broad community in development and use?
 - Are there other working software ecosystems we should learn from?
 - What topics are missing from the conversation?
- We need your engagement in this effort!

Thank you

<https://www.exascaleproject.org>

This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.



EXASCALE COMPUTING PROJECT

Thank you to all collaborators in the ECP and broader computational science communities. The work discussed in this presentation represents creative contributions of many people who are passionately working toward next-generation computational science.