Lessons Learned by the DOE Exascale Computing Project

EXASCALE COMPUTING PROJECT

Lois Curfman McInnes, Argonne National Laboratory

ISC: First International Workshop on Readiness of HPC Extreme-scale Applications (RHEA)
Hamburg, Germany
May 16, 2024

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

The Exascale Computing Project was designed to launch the exascale era

The ECP funded teams that spanned national labs, universities, and industry

$1.8B over 7 years in a formal DOE 413.3B project

Funded by DOE SC and NNSA programs

81 **Total research, development, and deployment projects**

6 Co-Design Centers
25 Application Development
35 Software Technology
17 Hardware and Integration

15	57	36
US DOE Labs, 6 of which are Core Partners	University Partners	Industry Organizations

- A **unique collaboration** brought together some of the brightest application, software, and computational experts from coast to coast
- **Best practices and lessons learned** for how to program GPUs – moving the nation forward
- **1000+ researchers** trained and ready for accelerator-based computing
- **1000+ students** introduced to HPC and Exascale computing through ECP's outreach, training, and workforce development initiatives

Technical work was largely complete as of Dec 31, 2023

Project leadership team is now working to close out the formal DOE 413.3B project

Science and beyond: Applications and discovery in ECP

And more:
- 25 applications
- 6 co-design centers
- 70 software products

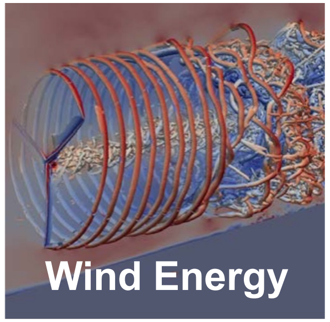

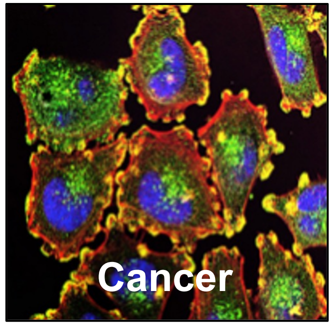

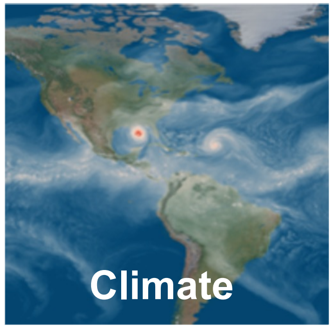

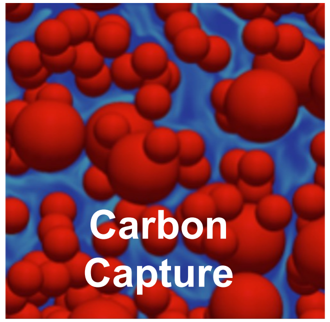

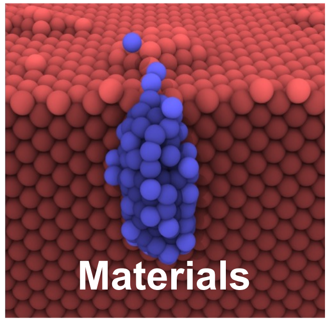

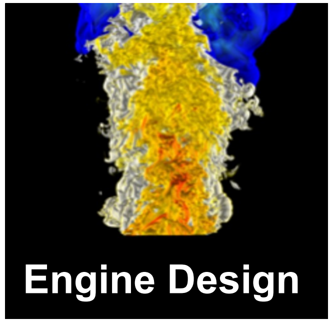

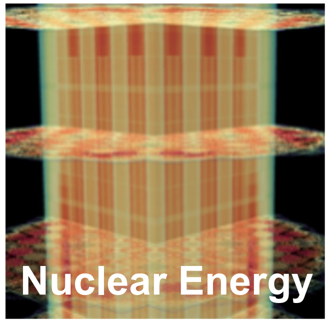

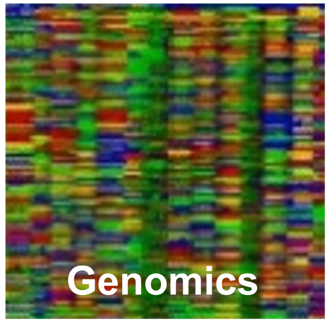

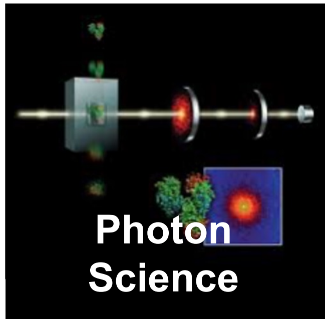

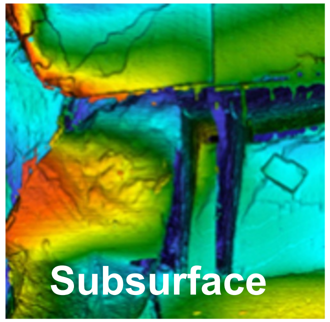

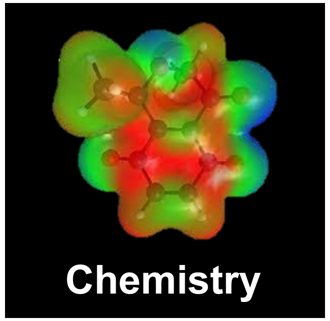

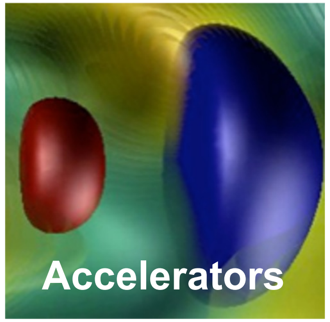

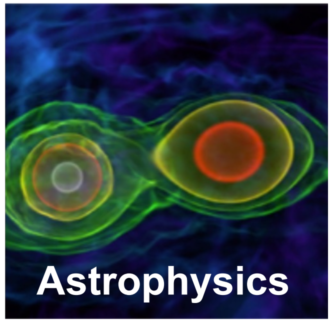

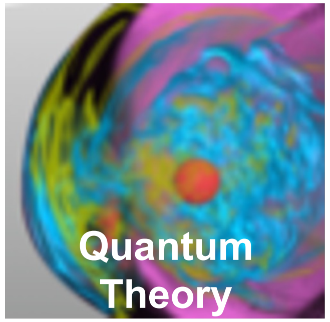

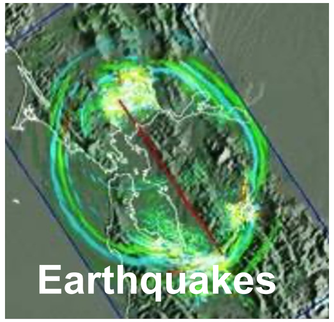

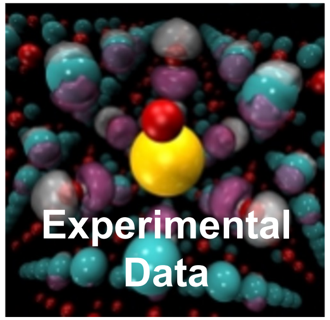

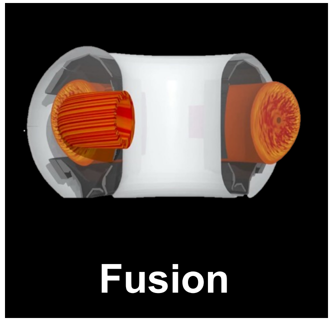

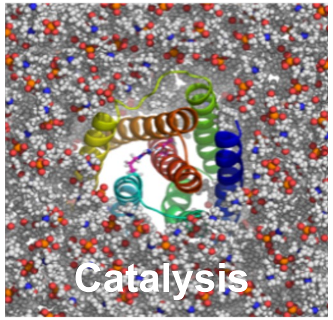

ECP EXASCALE COMPUTING PROJECT

Thank you to ECP teams: applications, software technologies, hardware and integration

ECP software technologies power diverse applications across multiple architectures

Project/PI	EXAALT: Molecular Dynamics Danny Perez
Challenge Problem	Damaged surface of Tungsten in conditions relevant to plasma facing materials in fusion reactors • 100,000 atoms • T=1200K
FOM Speedup	398.5
Nodes Used	7000
ST/CD Tools	Used in KPP Demo: **Kokkos, CoPa**

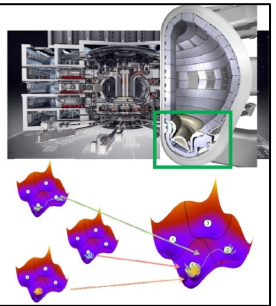

Project/PI	ExaSMR: Small Modular Reactors Steve Hamilton
Challenge Problem	NuScale-style Small Module Reactor (SMR) with depleted fuel and natural circulation • 213,860 Monte Carlo tally cells/6 reactions • 5.12×10^{12} particle histories/cycle, 40 cycles • 1098×10^{6} CFD spatial elements • 376×10^{9} CFD degrees of freedom • 1500 CFD timesteps
FOM Speedup	70
Nodes Used	6400
ST/CD Tools	Used in KPP Demo: **CEED** Additional: Trilinos

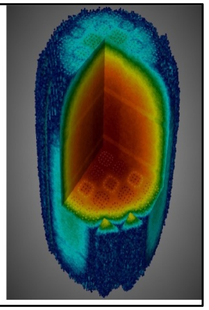

Project/PI	ExaSky: Cosmology Salman Habib
Challenge Problem	Two large cosmology simulations • gravity-only • hydrodynamics
FOM Speedup	271.65
Nodes Used	8192
ST/CD Tools	Used in KPP demo: **none** Additional: CoPa, VTK-m, CINEMA, HDF5.0

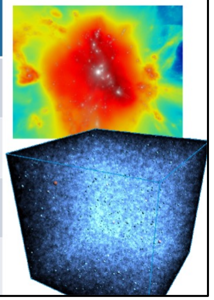

Project/PI	WarpX: Plasma Wakefield Accelerators Jean-Luc Vay
Challenge Problem	Wakefield plasma accelerator with a 1PW laser drive • 6.9×10^{12} grid cells • 14×10^{12} macroparticles • 1000 timesteps/1 stage
FOM Speedup	500
Nodes Used	8576
ST/CD Tools	Used in KPP Demo: **AMReX**, **libEnsemble** Additional: ADIOS, HDF5, VTK-m, ALPINE

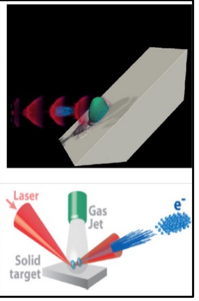

Project/PI	WDMApp: Fusion Tokamaks Amitava Bhatacharjee
Challenge Problem	Gyrokinetic simulation of the full ITER plasma to predict the height and width of the edge pedestal
FOM Speedup	150
Nodes Used	6156
ST/CD Tools	Used in KPP Demo: **CODAR, CoPA, PETSc, ADIOS** Additional: VTK-m

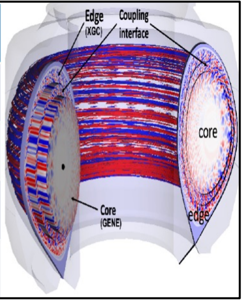

Project/PI	EQSIM: Earthquake Modeling and Risk Dave McCallen
Challenge Problem	Impacts of Mag 7 rupture on the Hayward Fault on the bay area.
FOM Speedup	3467
Nodes Used	5088
ST/CD Tools	Used in KPP Demo: **RAJA, HDF5**

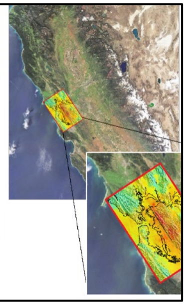

ECP Software Technologies

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

70 software products across 6 technical areas

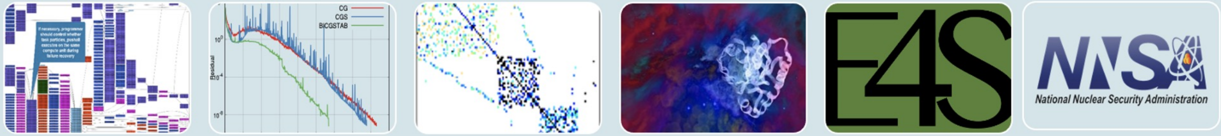

Programming Models & Runtimes
•Enhance and get

Development Tools
•Continued, multifaceted

Math Libraries
•Linear algebra, iterative linear solvers, direct linear solvers, integrators

Data and Visualization
• I/O via the HDF5 API

Software Ecosystem
•Develop features in Spack necessary to

NNSA ST
•Open source NNSA Software projects
•Projects that have

ECP EXASCALE COMPUTING PROJECT

Thank you

https://www.exascaleproject.org

<u>Thank you</u> 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.**

EXASCALE COMPUTING PROJECT

ECP Director: Lori Diachin
ECP Deputy Director: Ashley Barker
Former ECP Director: Doug Kothe



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.

ECP's Technical Focus Areas

Providing the necessary components to meet national goals

Performant mission and science applications at scale

Aggressive RD&D project	Mission apps; integrated S/W stack	Deployment to DOE HPC Facilities	Hardware technology advances

Application Development (AD)

Develop and enhance the predictive capability of applications critical to DOE

25 applications
National security, energy, Earth systems, economic security, materials, data

6 co-design centers
ML, graph analytics, mesh refinement, PDE discretization, particles, online data analytics


Andrew Siegel, AD Director
Erik Draeger, AD Deputy Director

Software Technology (ST)

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

70 unique software products spanning programming models and runtimes, math libraries, data and visualization, development tools


Mike Heroux, ST Director
Lois Curfman McInnes, ST Deputy Director

Hardware and Integration (HI)

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

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


Richard Gerber, HI Director
Susan Coghlan, HI Deputy Director

EXASCALE COMPUTING PROJECT

HPC systems have come a long way since ECP's inception

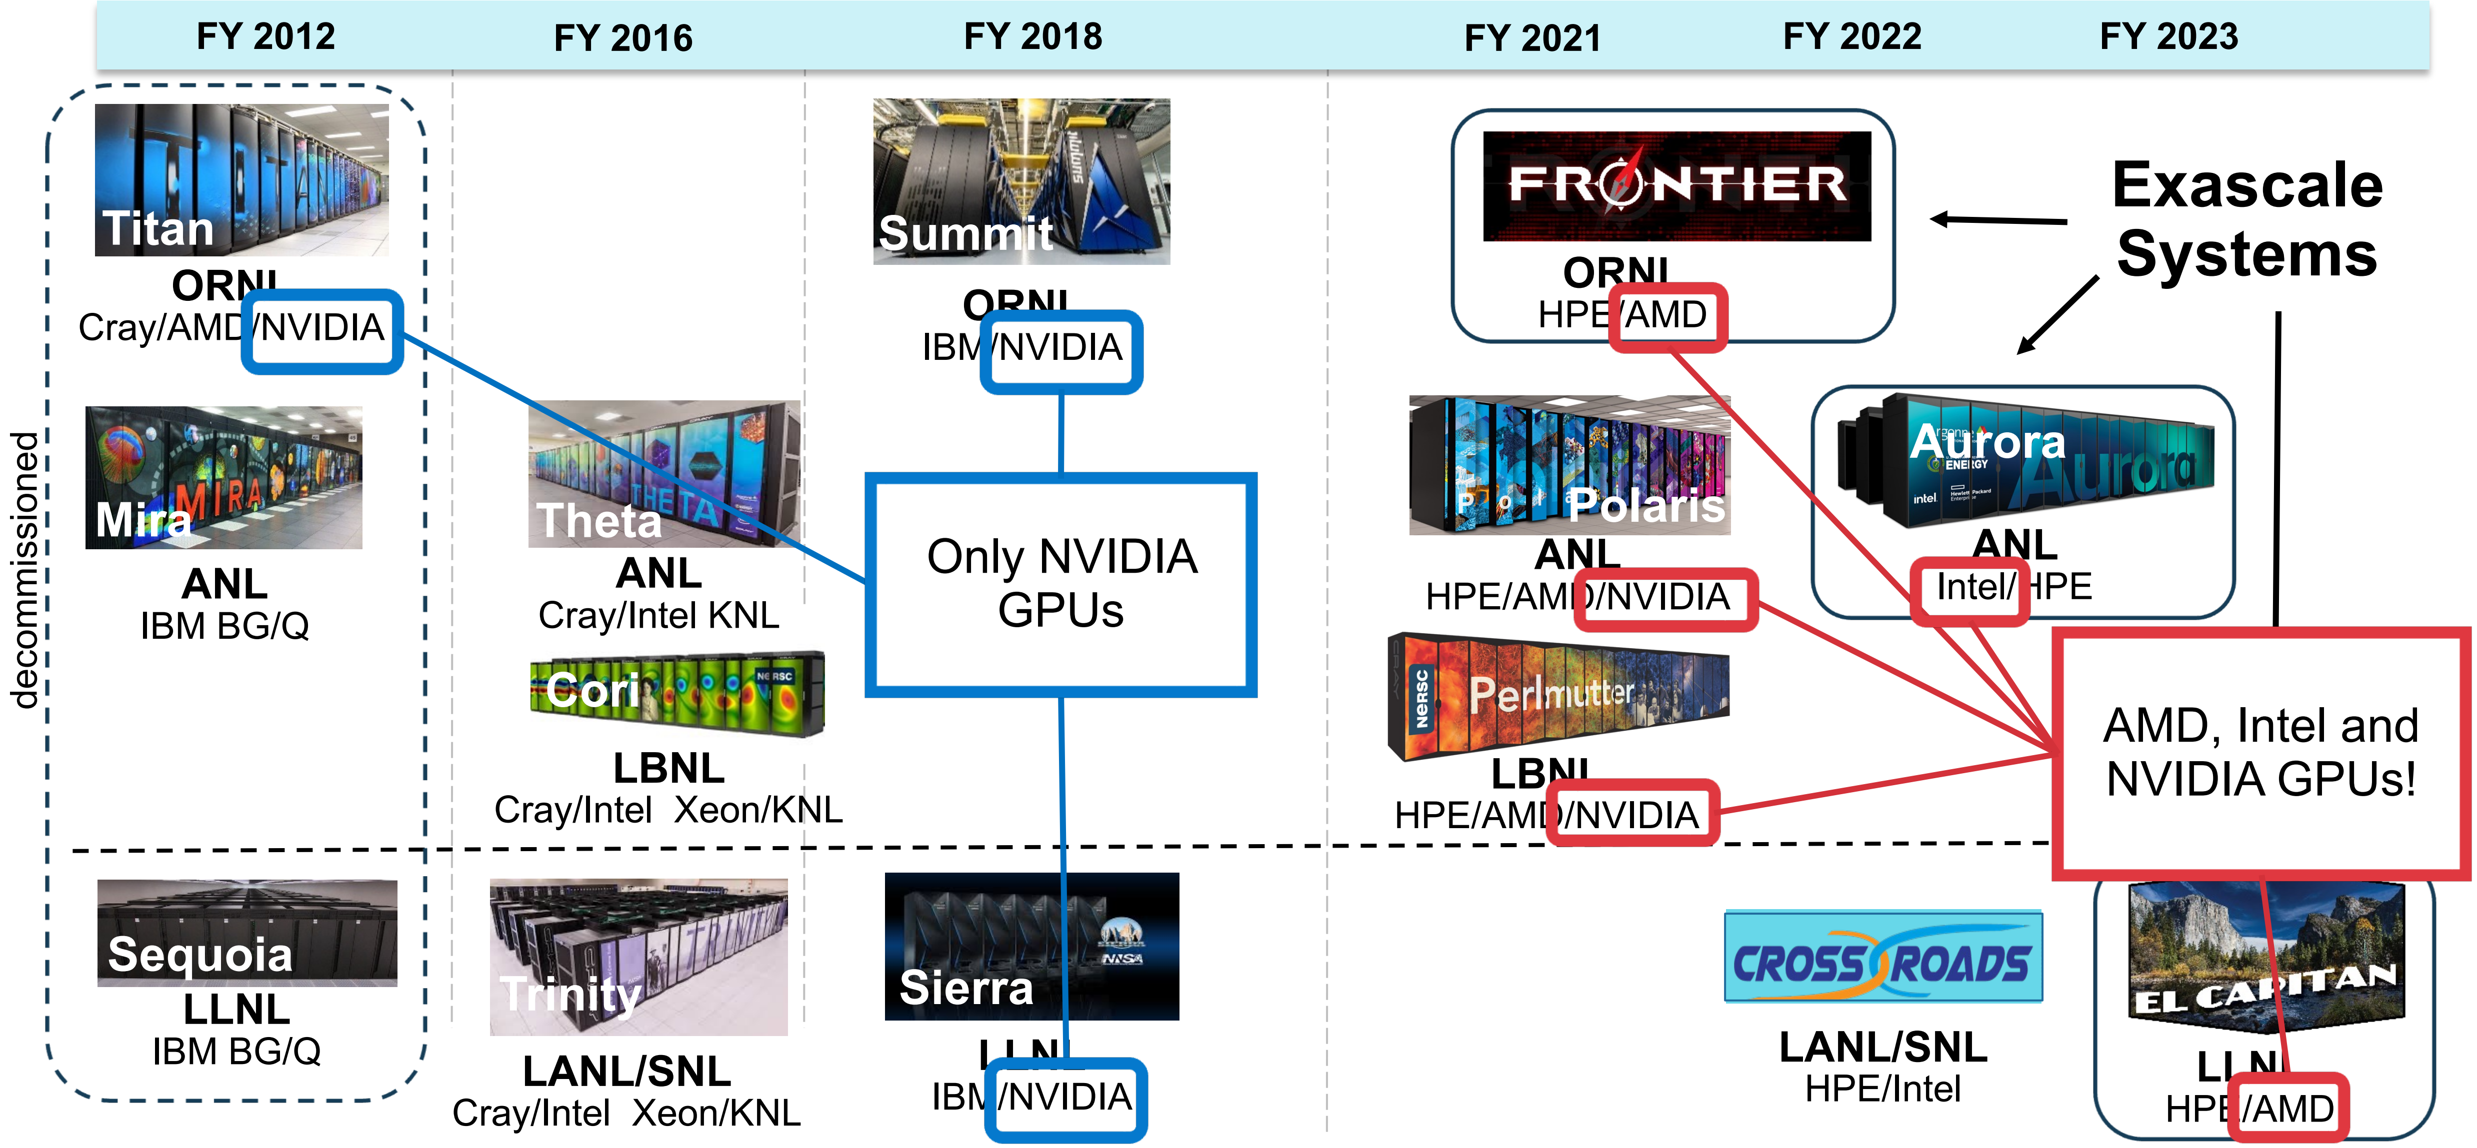

ECP EXASCALE COMPUTING PROJECT

Version 2.0

ECP set ambitious capability and performance goals across a diverse application space

slide credit: A. Siegel and E. Draeger

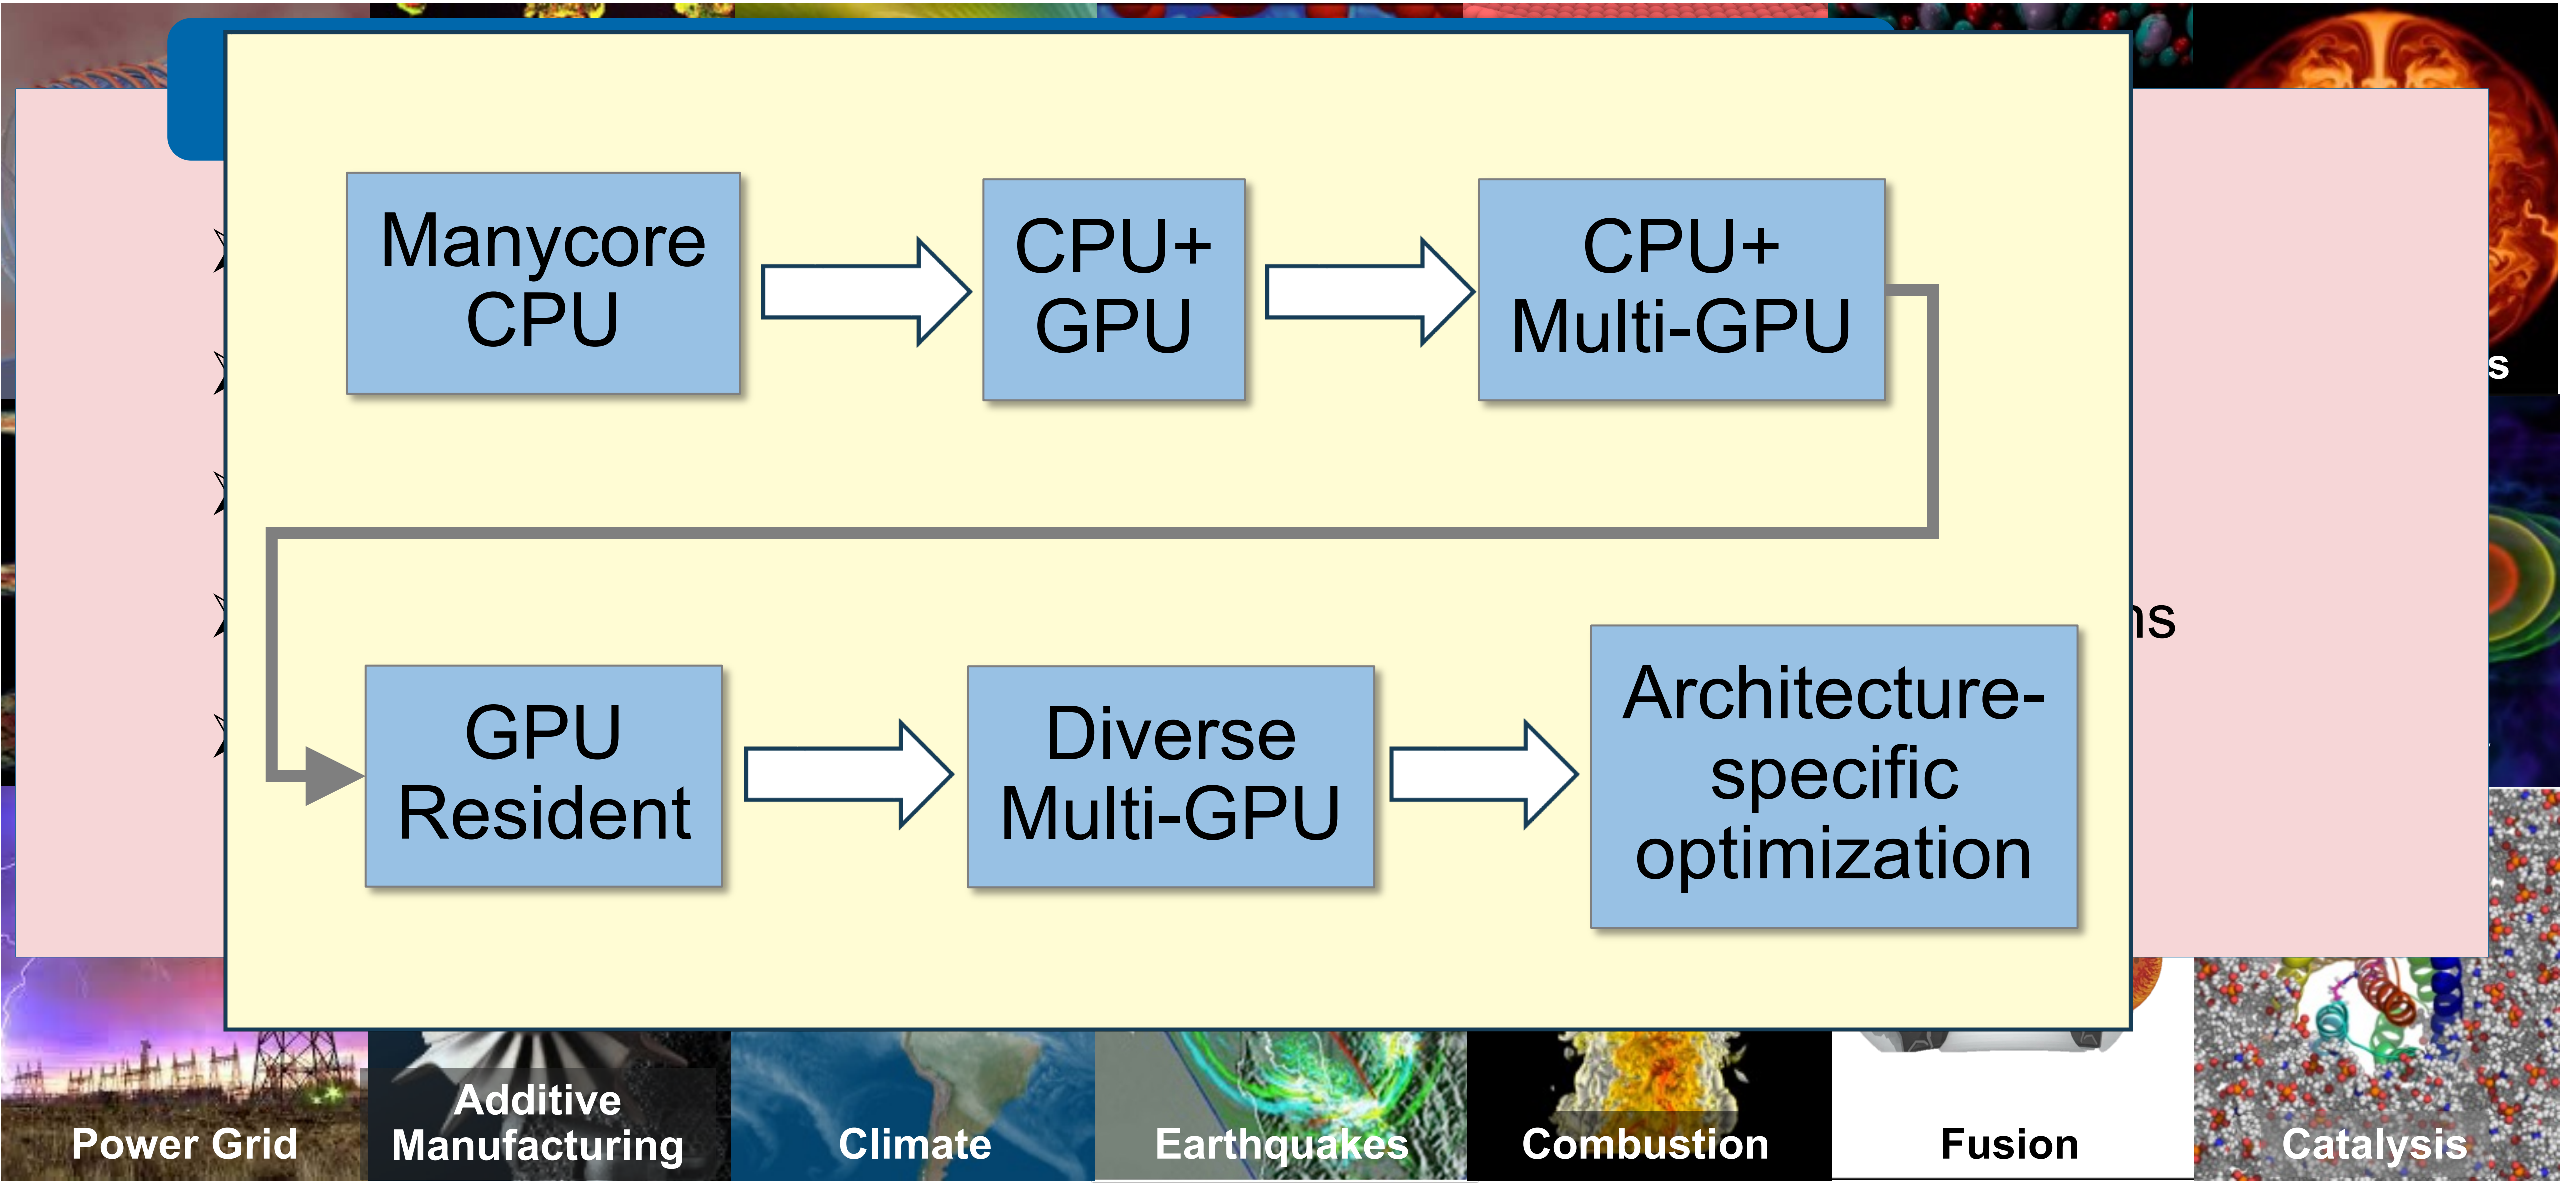

ECP application results exceeded expectations

10 out of 11 KPP-1 projects surpassed ambitious 50x performance target

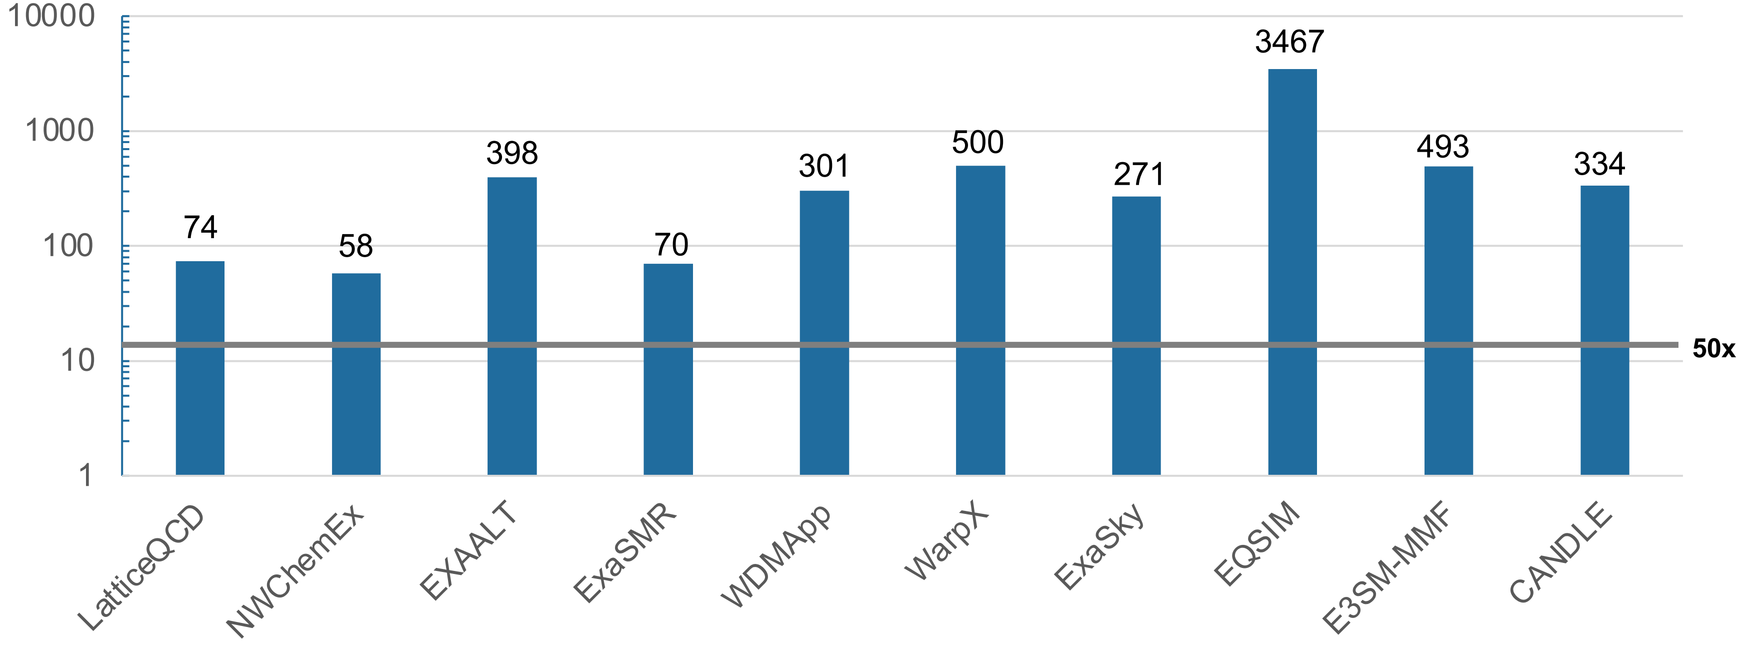

Co-design played a critical role

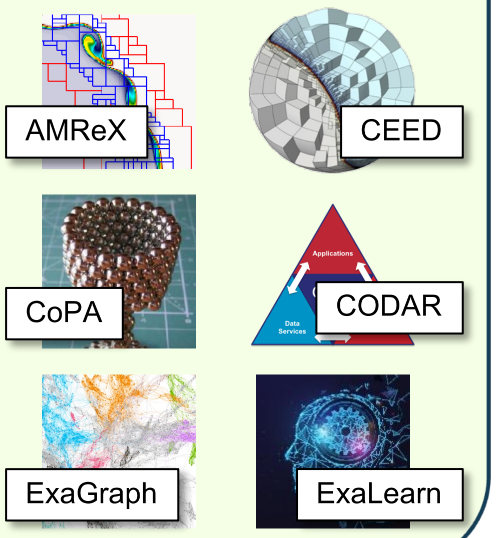

9 out of 10 science KPP-2 projects completed exascale capability demonstrations

3 out of 4 NNSA KPP-2 applications demonstrated exascale readiness

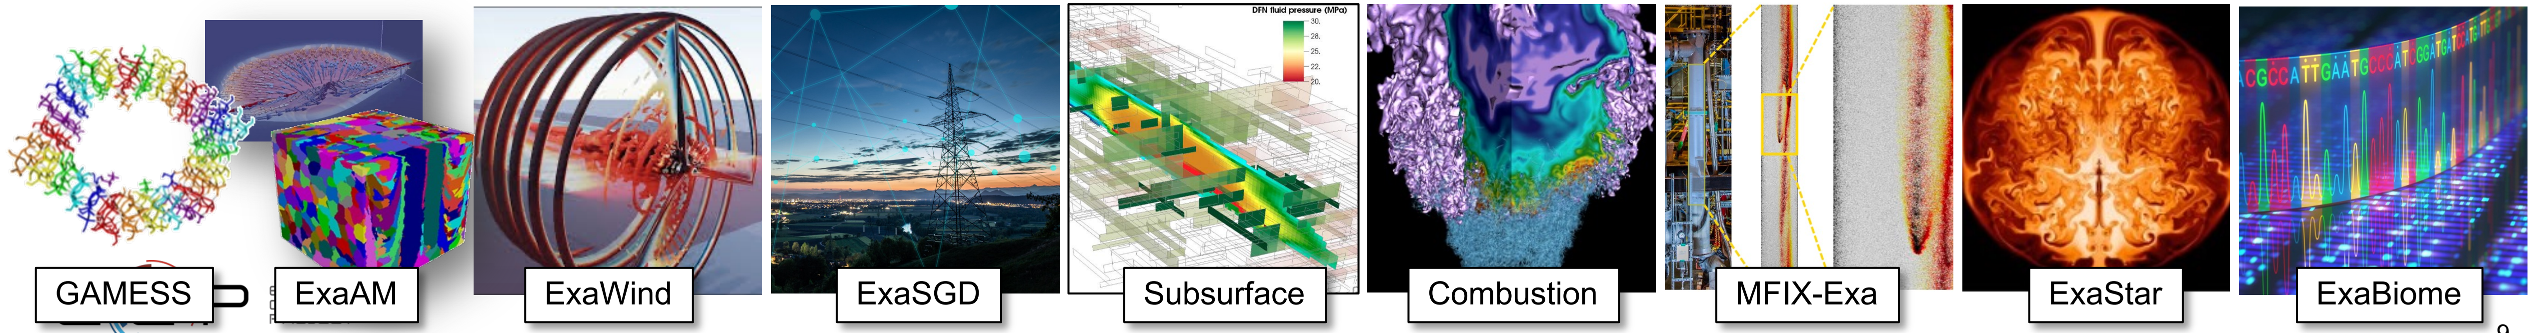

slide credit: A. Siegel and E. Draeger

ExaSMR developed first-of-a-kind simulations of advanced nuclear reactors such as small modular reactors

- **Objective**: Help DOE meet its goal of an operational small modular reactor (SMR) in 10 years, a key part of the Department's goal to develop safe, clean, and affordable nuclear power options
- Accomplishments:
  - First fully coupled, fully resolved simulation of nuclear reactor core with coupled Monte Carlo neutron transport and CFD; achieved nearly 100X overall performance improvements in the science work rate
  - Allows study of evolution of nuclear fuel for the first time
  - **Gordon Bell Finalist – SC23!**
- Deliver experiment-quality simulations of reactor behavior to enable the design and commercialization of advanced nuclear reactors and fuels at significant savings in time (from years to months) and money

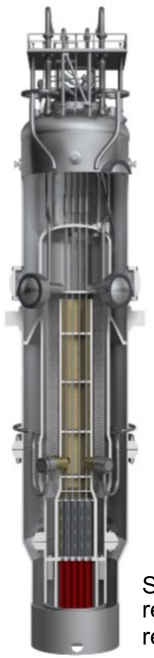

Small modular reactor. Image reproduced with permission of NuScale

EQSIM developed a framework that advances the ability to predict and use site-specific ground motions in earthquake risk assessment

- **Objective**: Create high-performance simulation tools that establish a coupled assessment of earthquake hazard (ground motion) and earthquake risk (infrastructure demands) at regional scale
- Accomplishments:
  - Increased the resolution of earthquake ground motion frequencies from 2Hz to 10Hz
  - Simulations now include strong coupling between regional geophysics and local soil/building models
  - Achieved more than 1000X improvement in computational performance compared to all previous San Francisco Bay Area simulations
- This is a transformational tool for addressing questions of earthquake risk to buildings, energy systems, and other critical infrastructure spanning an entire region

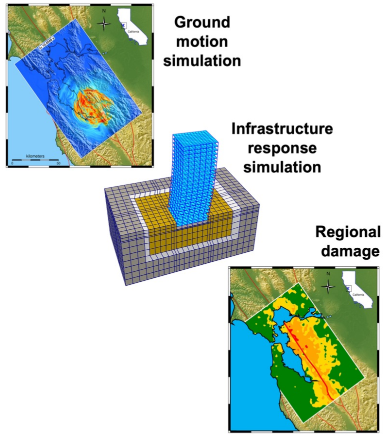

Exascale applications require significant software infrastructure

Programming abstractions

Memory / heterogeneous data management

In situ visualization and analysis

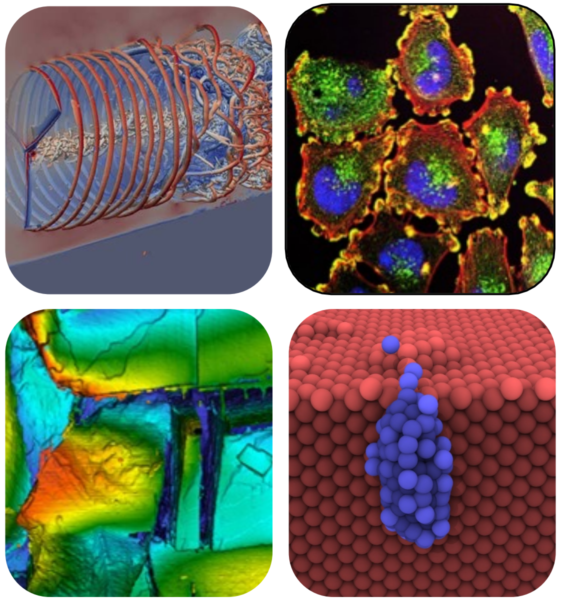

Solvers / preconditioners

Numerical infrastructure / discretizations

Build and deployment

Performance analysis and monitoring

Achievements in new application science would not have been possible without a robust software stack

- E4S: HPC software ecosystem – a curated software portfolio built on software development toolkits (SDKs)
- A **Spack-based** distribution of software tested for interoperability and portability to multiple architectures
- Available from **source**, **containers**, **cloud, binary caches**
- Not a commercial product – an open resource for all
- Supported by DOE and commercial entities (ParaTools)
- Growing functionality: February 2024: E4S 24.02 – 120+ full release products

https://spack.io

Spack lead: Todd Gamblin (LLNL)

Spack

E4S

https://e4s.io

E4S lead: Sameer Shende (U Oregon)

Community Policies Commitment to SW quality	DocPortal Single portal to all E4S product info	Portfolio testing Especially leadership platforms
Curated collection The end of dependency hell	Quarterly releases	Build caches 10X build time improvement
Turnkey stack A new user experience	https://e4s.io	Post-ECP Strategy LSSw, ASCR Task Force

One of the largest impacts of ECP

- Deployed on Frontier, Aurora, and Perlmutter
- Used on DoD, NOAA, NSF, and many other systems

ECP EXASCALE COMPUTING PROJECT

ECP **Software Technology (ST)**

Goal

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

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

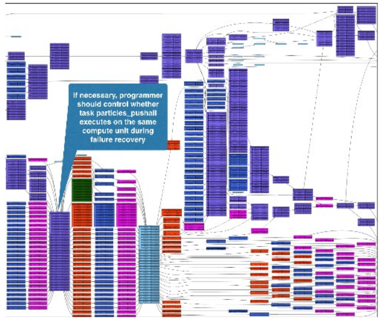

ECP ST six technical areas

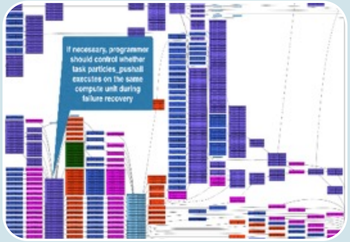

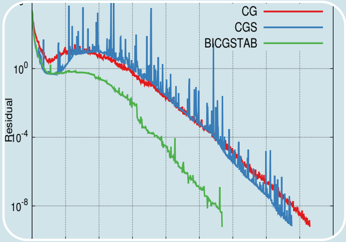

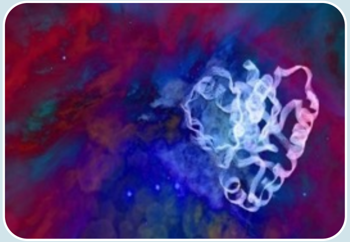

E4S

NNSA
National Nuclear Security Administration

Programming Models & Runtimes

- Enhance and get ready for exascale the MPI and OpenMP programming models (hybrid programming models, deep memory copies)
- Develop 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 ST products in E4S, and the AD projects that adopt it
- Develop Spack stacks for reproducible turnkey software deployment
- Optimization and interoperability of containers for HPC
- 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

Area Leads:	Rajeev Thakur	Jeff Vetter	Sherry Li	Jim Ahrens	Todd Munson	Kathryn Mohror

ST L4 Teams

- WBS
- Name
- PIs
- PCs – Project Coordinators

ECP ST Stats

- 250 staff
- 70 products
- 35 L4 subprojects
- 30 universities
- 9 DOE labs
- 6 technical areas
- 1 of 3 ECP focus areas

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	
2.3.1.07	Exascale MPI (MPICH)	Guo, Yanfei	
2.3.1.08	Legion	McCormick, Pat	
2.3.1.09	PaRSEC	Antz, Hartwig	
2.3.1.14	Pagoda: UPC++/GASNet for Lightweight Communication and Global Address Space Support	Hargrove, Paul	Hargrove, Paul
2.3.1.16	SICM	Pakin, Scott	Turton, Terry
2.3.1.17	OMPI-X	Bernholdt, David	
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	Miller, Barton	Tim Haines
2.3.2.06	Exa-PAPI++: The Exascale Performance Application Programming Interface with Modern C++	Anzt, Hartwig	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	
2.3.2.11	SOLLVE: Scaling OpenMP with LLVm for Exascale	Chandrasekaran, Sunita	Vergara, Veronica
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	
2.3.3.06	Preparing PETSc/TAO for Exascale	Munson, Todd	
2.3.3.07	STRUMPACK/SuperLU/FFTX: sparse direct solvers, preconditioners, and FFT libraries	Li, Sherry	
2.3.3.12	Enabling Time Integrators for Exascale Through SUNDIALS/ Hypre	Woodward, Carol	
2.3.3.13	CLOVER: Computational Libraries Optimized Via Exascale Research	Anzt, Hartwig	
2.3.3.14	ALExa: Accelerated Libraries for Exascale/ForTrilinos	Prokopenko, Andrey	
2.3.3.15	Sake: Solvers 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	O'Leary, Patrick	Bagha, Neelam
2.3.4.09	ADIOS Framework for Scientific Data on Exascale Systems	Klasky, Scott	
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/Sz	Cappello, Franck	Ehling, Scott
2.3.4.15	ExaIO - Delivering Efficient Parallel I/O on Exascale Computing Systems with HDF5 and Unify	Wu, John	Bagha, Neelam
2.3.4.16	ALPINE: Algorithms and Infrastructure for In Situ Visualization and Analysis/ZFP	Ahrens, James	Turton, Terry
2.3.5	**Software Ecosystem and Delivery**	**Munson, Todd**	
2.3.5.01	Software Ecosystem and Delivery Software Development Kit	Willenbring, James M	Willenbring, James M
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	Randles, Tim	Turton, Terry
2.3.6.02	LLNL ATDM	Springmeyer, Becky	Gamblin, Todd
2.3.6.03	SNL ATDM	Stewart, Jim	Trujillo, Gabrielle

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

EXASCALE COMPUTING PROJECT

A Sampler of Products

- No two project alike
- Some personality driven
- Some community driven
- Small, medium, large

ECP Software Technology works on products that apps need now and in 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, Ginkgo, Spack**)
- **New products:** Enable exploration of emerging HPC requirements (e.g., **Variorum, zfp**)

Legacy: A stack that enables performance portable application development on leadership platforms

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
C++ Performance Portability Abstractions	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

ECP applications would not have been successful without close collaboration with software technology developers

ExaWind: advanced wind farm modeling	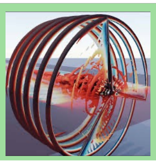	Sparse linear solvers optimized for strong scaling and GPU performance **from hypre and Trilinos**
QMCPACK: quantum Monte Carlo for materials	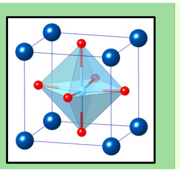	Batched dense linear algebra kernels significantly improved GPU performance **from SLATE**
ExaSGD: power grid optimization		Optimize sparse indefinite solvers developed and optimized for large-scale grid problems **from SuperLU and Ginkgo**
ExaSMR: small modular reactor modeling	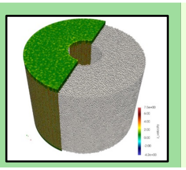	Custom discretization designed and tuned for specific reactor assembly **from CEED**
WarpX: plasma wakefield accelerator design	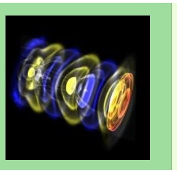	Adaptive mesh data structures and solvers highly optimized for GPU performance **from AMReX**
ExaFEL: real-time light source analysis and reconstruction	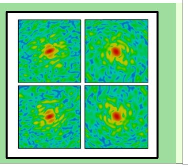	Non-uniform FFTs designed to minimize data motion **from FFTX**

ECP EXASCALE COMPUTING PROJECT

Delivering an open, hierarchical software ecosystem

More than a collection of individual products

Levels of Integration	Product	Source and Delivery
• Build all SDKs • Build complete stack • Assure core policies • Build, integrate, test	E4S	**Source:** ECP E4S team; Non-ECP Products (all dependencies) **Delivery:** spack install e4s; containers; CI Testing
• Group similar products • Make interoperable • Assure policy compliant • Include external products	SDKs	**Source:** SDK teams; Non-ECP teams (policy compliant, spackified) **Delivery:** Apps directly; spack install sdk; future: vendor/facility
• Standard workflow • Existed before ECP	ST Products	**Source:** ECP L4 teams; Non-ECP Developers; Standards Groups **Delivery:** Apps directly; spack; vendor stack; facility stack

ECP ST Open Product Integration Architecture (E4S and SDKs levels)

ECP ST Individual Products (ST Products level)

ECP EXASCALE COMPUTING PROJECT

E4S and SDKs as platforms are providing tremendous value

Activity	SDKs	E4S
Planning	Transparent and collaborative requirements, analysis and design, delivery – better plans, less effort, improved complementarity	Campaign-based portfolio planning coordinated with Facilities, vendors, community ecosystem, non-DOE partners
Implementation	Leverage shared knowledge, infrastructure, best practices	ID and assist product teams with cross-cutting issues
Cultivating community	Within a specific technical domain: Portability layers, LLVM coordination, sparse solvers, etc.	Across delivery and deployment, with software teams, facilities' staff, with non-DOE users in industry, US agencies
Resolving issues, sharing solutions	Performance bottlenecks and tricks, coordinated packaging and use of substrate, e.g., Desul for RAJA and Kokkos	Build system bugs and enhancements, protocols for triage, tracking & resolution, leverage across & beyond DOE
Improving quality	Shared practice improvement, domain-specific quality policies, reduced incidental differences and redundancies, per-commit CI testing of portfolio	Portfolio-wide quality policies with assessment process and quality improvement efforts, documentation portal, portfolio testing on many platforms not available to developers. Address supply chain needs
Path-finding	Collaborative exploration and development of leading-edge tools and processes	Exploration and development of leading-edge packaging and distribution tools and workflows that provide capabilities and guidance for others
Training	Collaborative content creation and curation, coordinated training events for domain users, deep, problem-focused solutions using multiple products	Portfolio installation and use, set up of build caches, turnkey and portable installations, container and cloud instances
Developer experience	Increased community interaction, increased overhead (some devs question value), improved R&D exploration, e.g., variable precision	Low-cost product visibility via doc portal, wide distribution via E4S as from-source/pre-installed/container environment
User experience	Improve multi-product use, better APIs through improved design, easier understanding of what to use when	Rapid access to latest stable feature sets, installation on almost any HPC system, leadership to laptop
Scientific Software R&D	Shared knowledge of new algorithmic advances, licensing, build tools, and more	Programmatic cultivation of scientific software R&D not possible at smaller scales
Community development	Attractive and collaborative community that attracts junior members to join, establishes multi-institutional friendships & careers	Programmatic cultivation of community through outreach and funded opportunities that expand the sustainable membership possibilities

The SDK and E4S platforms provide compelling value for modest cost in ways that become more important going forward

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

xSDK lead: Ulrike Meier Yang (LLNL)
xSDK release lead: Satish Balay (ANL)

xSDK release 1.0.0 (Nov 2023)

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

(deal.II, preCICE, PHIST, SLEPc: from the broader community)

As motivated and validated by the needs of ECP applications:

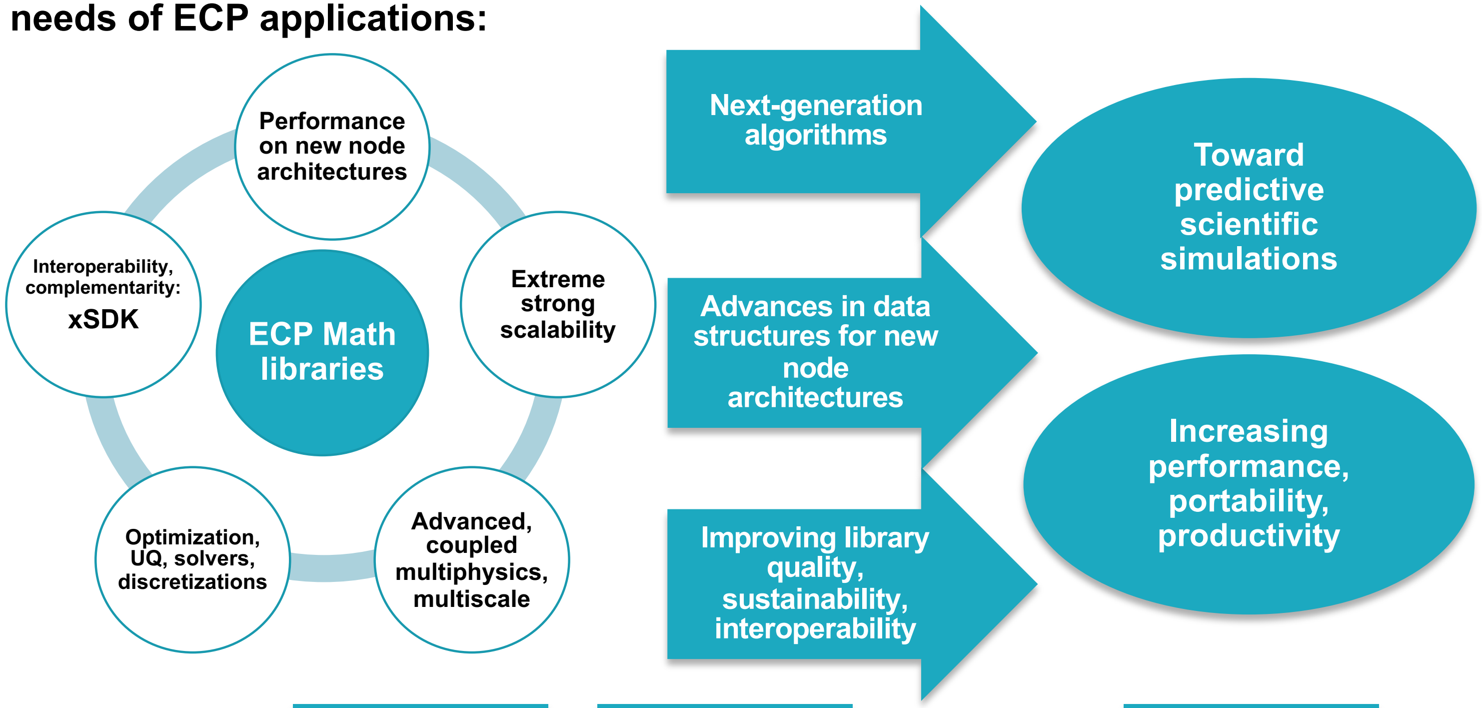

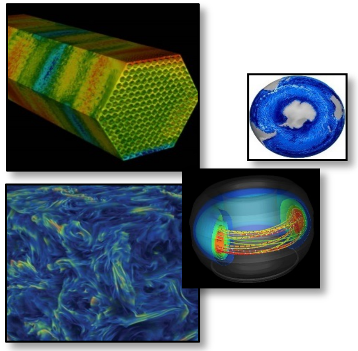

Timeline: xSDK release 1 | xSDK release 2 | ..... | xSDK release n

Ref: xSDK: Building an Ecosystem of Highly Efficient Math Libraries for Exascale, **SIAM News**, Jan 2021

Macro-engineering lifecycle summary

Managed by Milestone Process

Create Annual Planning Package
- Each product has its own planning packages
- Defined for all FYs

Refine upcoming FY plan
- Complete months prior to FY
- 4 or more milestones per product

Refine upcoming milestone
- Complete weeks prior to activity start
- Include all details

Develop capabilities and track progress via tailored agile

Measured by Capability Integration Process

Integrate into product
- Full testing, documentation, etc.
- Direct access for some users

Integrate into SDK
- Satisfy SDK community policies
- Direct access for some users

Integrate into E4S
- Satisfy E4S community policies
- Full ecosystem with high value

Deliver to users
- From source (spack)
- Containers, cloud

E4S Business Model: Optimize Cost & Benefit Sharing

DOE E4S Team enables a portfolio approach
- Integrated delivery/support of libs/tools
- Single POC for planning and issues

New

HIGH PERFORMANCE SOFTWARE FOUNDATION

https://hpsf.io

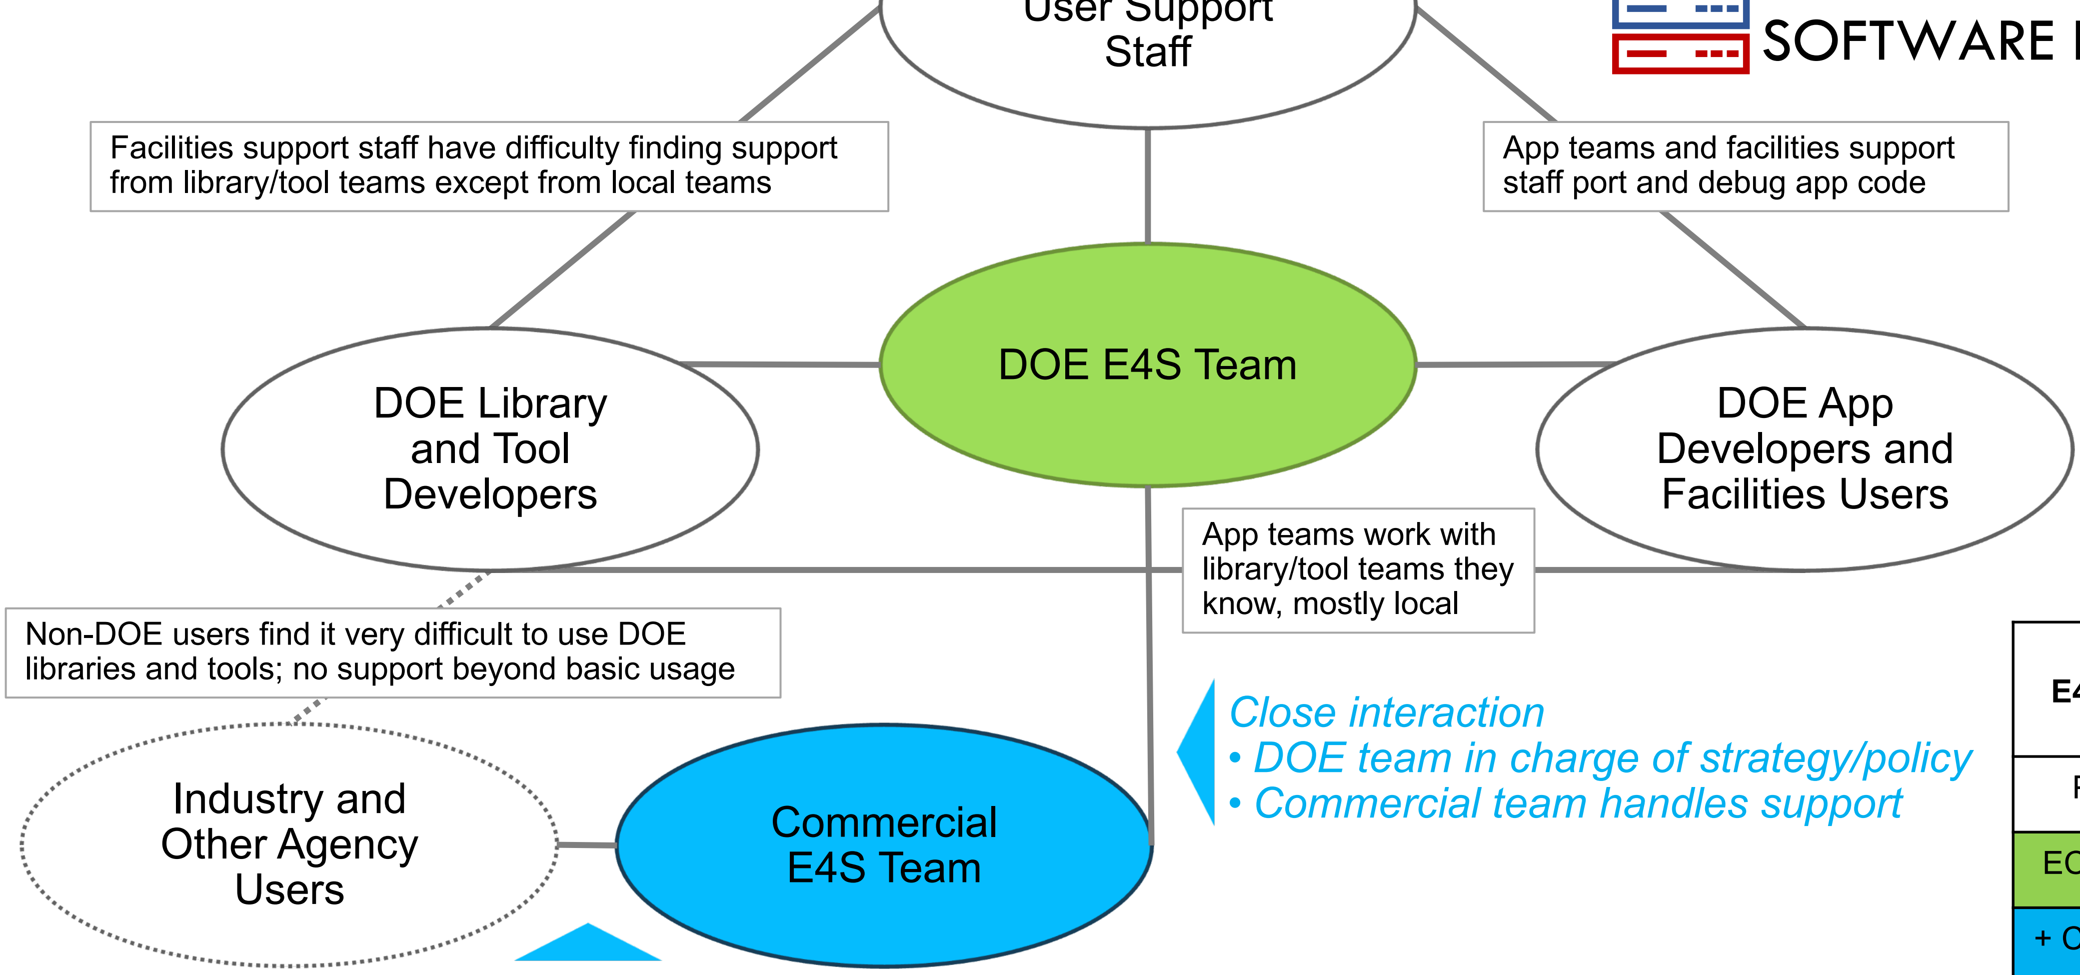

Close interaction
- *DOE team in charge of strategy/policy*
- *Commercial team handles support*

First of a kind interactions
- *Industry/agencies can acquire support*
- *Shared costs and benefits with DOE*

E4S Phase	**Cost & Benefit Scope**
Pre-E4S	Local Facility
ECP support	DOE complex
+ Commercial support	Universal

100X Demonstrated: ECP-sponsored application FOMs

Project/PI	EXAALT: Molecular Dynamics Danny Perez
Challenge Problem	Damaged surface of Tungsten in conditions relevant to plasma facing materials in fusion reactors • 100,000 atoms • T=1200K
FOM Speedup	398.5
Nodes Used	7000
ST/CD Tools	Used in KPP Demo: **Kokkos, CoPa**

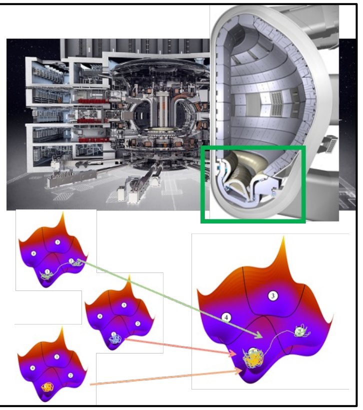

Project/PI	ExaSMR: Small Modular Reactors Steve Hamilton
Challenge Problem	NuScale-style Small Module Reactor (SMR) with depleted fuel and natural circulation • 213,860 Monte Carlo tally cells/6 reactions • 5.12×10^{12} particle histories/cycle, 40 cycles • 1098×10^{6} CFD spatial elements • 376×10^{9} CFD degrees of freedom • 1500 CFD timesteps
FOM Speedup	70
Nodes Used	6400
ST/CD Tools	Used in KPP Demo: **CEED** Additional: Trilinos

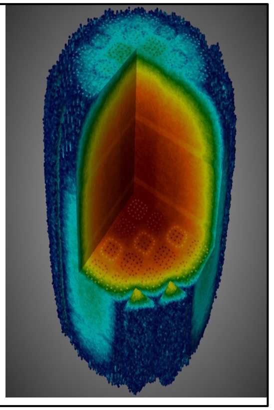

Project/PI	ExaSky: Cosmology Salman Habib
Challenge Problem	Two large cosmology simulations • gravity-only • hydrodynamics
FOM Speedup	271.65
Nodes Used	8192
ST/CD Tools	Used in KPP demo: **none** Additional: CoPa, VTK-m, CINEMA, HDF5.0

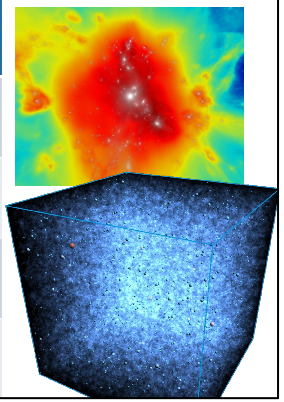

Project/PI	WarpX: Plasma Wakefield Accelerators Jean-Luc Vay
Challenge Problem	Wakefield plasma accelerator with a 1PW laser drive • 6.9×10^{12} grid cells • 14×10^{12} macroparticles • 1000 timesteps/1 stage
FOM Speedup	500
Nodes Used	8576
ST/CD Tools	Used in KPP Demo: **AMReX**, **libEnsemble** Additional: ADIOS, HDF5, VTK-m, ALPINE

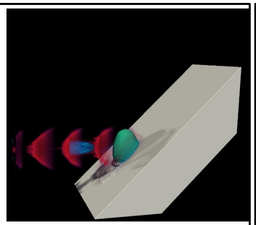

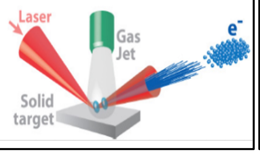

Project/PI	WDMApp: Fusion Tokamaks Amitava Bhatacharjee
Challenge Problem	Gyrokinetic simulation of the full ITER plasma to predict the height and width of the edge pedestal
FOM Speedup	150
Nodes Used	6156
ST/CD Tools	Used in KPP Demo: **CODAR**, **CoPA**, **PETSc**, **ADIOS** Additional: VTK-m

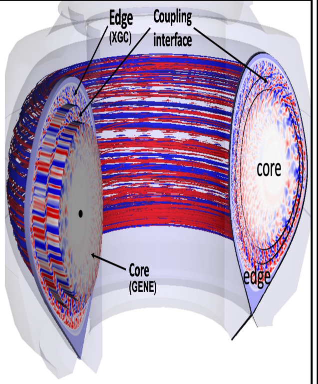

Project/PI	EQSIM: Earthquake Modeling and Risk Dave McCallen
Challenge Problem	Impacts of Mag 7 rupture on the Hayward Fault on the bay area.
FOM Speedup	3467
Nodes Used	5088
ST/CD Tools	Used in KPP Demo: **RAJA, HDF5**

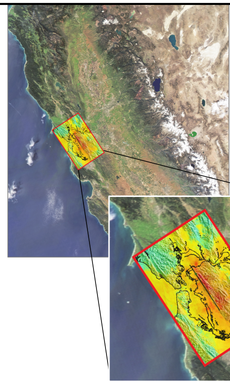

FOM = Figure of Merit: Speedup in a scientifically meaningful way.
50X or more improvement for each code was required for success.

Advancing scientific productivity through better scientific software

Reducing technical risk by building a firmer foundation for computational science

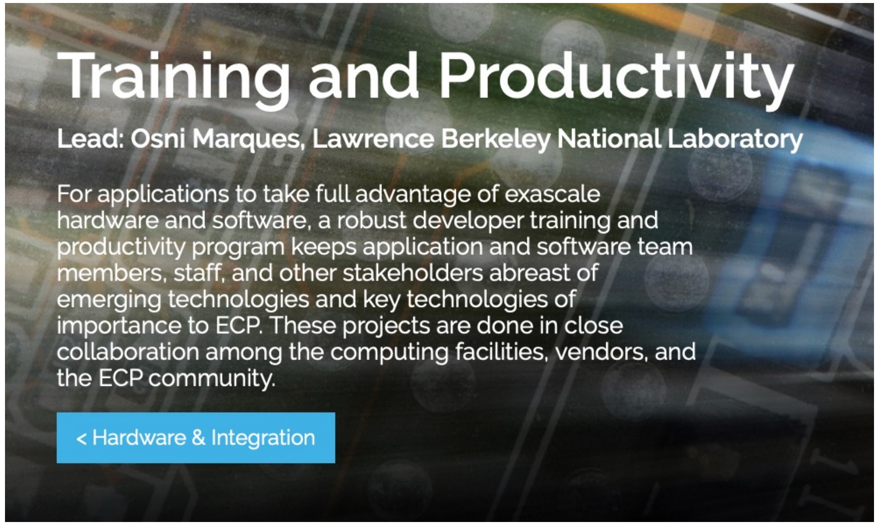

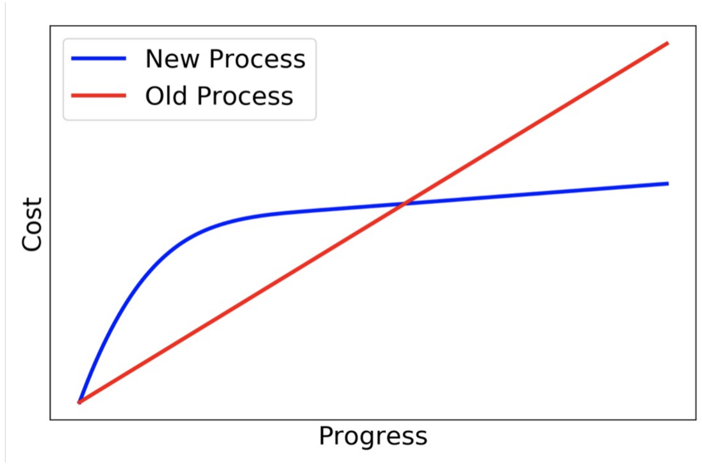

Improving developer productivity and software sustainability:
Nurturing a culture of continual improvement in software practices

Recognizing that change requires investment but pays off over time

Impact: Helping ECP teams to achieve:

- **Better:** Science, portability, robustness, composability
- **Faster:** Execution, development, dissemination
- **Cheaper:** Fewer staff hours and lines of code

IDEAS productivity

better scientific software

ECP EXASCALE COMPUTING PROJECT

IDEAS productivity

https://ideas-productivity.org

Advancing scientific productivity through better scientific software

Science through computing is only as good as the software that produces it.

- Accelerating design space exploration
- Advancing quality, usability and interoperability, while respecting team autonomy
- Software community policies
- Software Development Kits (SDKs) and E4S
  - xSDK, CAT-SDK, DAV-SDK, SWAS, etc

- Webinar Series:
  - HPC Best Practices
  - HPC Workforce Development and Retention
- Tutorials on Practices for Better Scientific Software

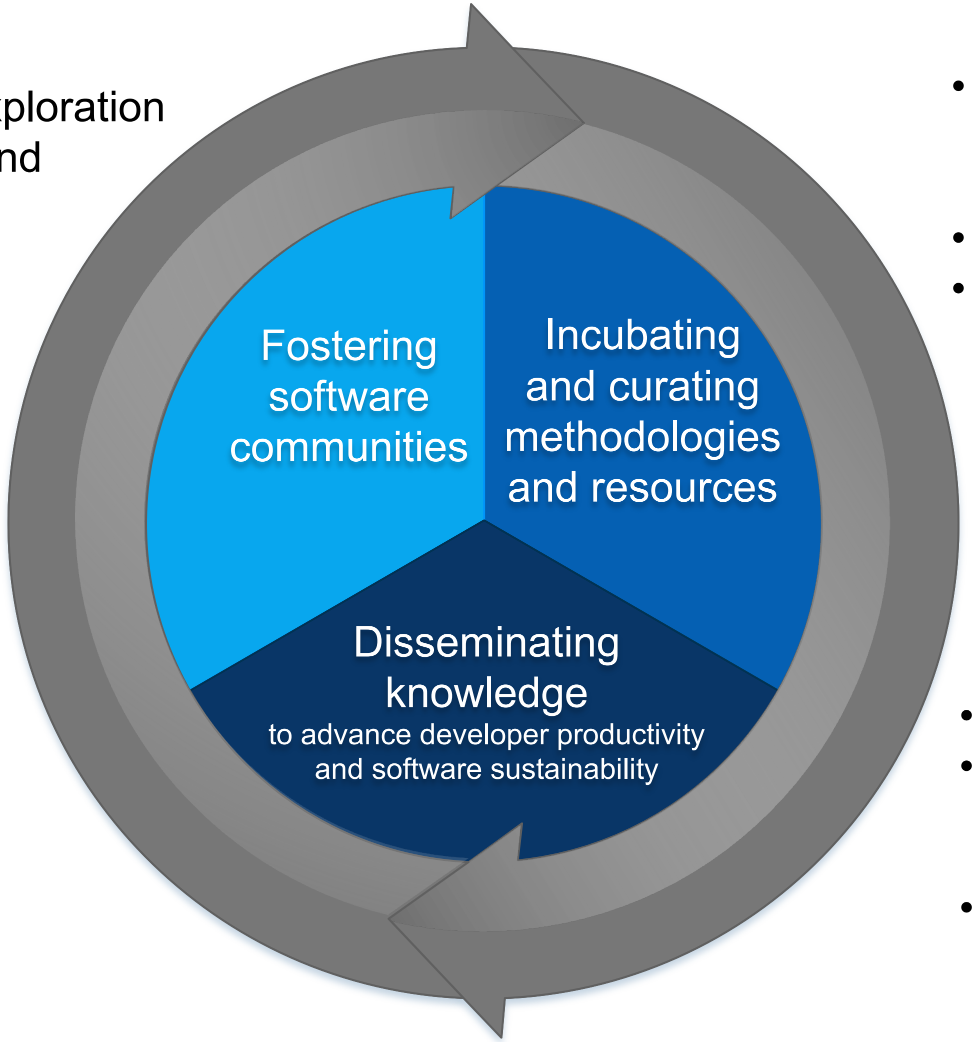

- Productivity and Sustainability Improvement Planning (PSIP)
- *Team of teams* concepts
- Better Scientific Software (BSSw.io) website

- BSSw Fellowship Program
- Panel Series:
  - Strategies for Working Remotely
  - Performance Portability
- Events: BOFs, workshops, and more

EXASCALE COMPUTING PROJECT

More info: ***A cast of thousands: How the IDEAS Productivity project has advanced software productivity and sustainability,*** arXiv, Nov 2023, https://doi.org/10.48550/arXiv.2311.02010 (to appear in IEEE CiSE).

Moving Forward

doi:10.1109/MCSE.2023.3260475

doi:10.2172/1846009

We believe that IDEAS has been an important catalyst for the scientific software community in ECP and beyond, but two further elements are essential for continued qualitative growth...

- Increasing focus on **research software science**
  - Applying our experience with hypothesis-driven science to improve our understanding of how scientific software is developed and used
  - Social-, cognitive-, hard-science, and engineering viewpoints needed
- Changing the prevailing attitude that scientific software productivity, sustainability, and trustworthiness are not just “nice to have” but **“must-have”**
  - Innovators and early adopters are there
  - But many still prioritize more scientific results at the expense of beneficial investments in the software behind them

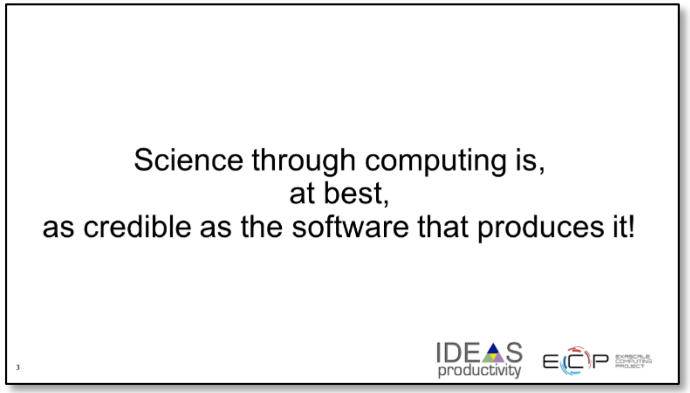

A key message of the BSSw tutorials

EXASCALE COMPUTING PROJECT

Consortium for the Advancement of Scientific Software

DOE Office of Advanced Scientific Computing Research (ASCR)
Next-Generation Scientific Software Technologies
Post-ECP: Collaboration among Software Stewardship Organizations (SSOs)

https://cass.community

- **PESO:** Stewarding, evolving and integrating a cohesive ecosystem for DOE software
- **RAPIDS & FASTMath:** Stewardship, advancement, integration for math, data/vis, and ML/AI packages.
- **SWAS:** Stewardship and project support for scientific workflow software and its community
- **S4PST:** Stewardship, advancement and engagement for programming systems.
- **STEP:** Stewardship, advancement of software tools for understanding performance and behavior.
- **COLABS:** Training, workforce development, and building the RSE community.
- **CORSA:** Partnering with foundations to provide onboarding paths for DOE-funded software.

Addressing the stewardship needs of the ASCR scientific software ecosystem

Member organizations have distinct and complementary foci with the goal of ensuring the long-term viability of the ASCR software ecosystem so that it may continue to serve as the base for future DOE-funded research.

https://pesoproject.org

PESO: Partnering for Scientific Software Ecosystem Stewardship Opportunities

PESO Partnerships

Stakeholder Engagement and Consortium Partnerships

Consortium
- Spack support
- Products in E4S
- SW practices
- PIER activities

Applications Community
- Spack support
- E4S support

Computing Facilities
- Spack support
- E4S integration
- E4S user support

Commercial HPC Companies
- Engage in joint activities to advance the scientific software ecosystem
- Develop business models to further partnerships
- Increase release coordination activities

US Agencies
- Engage in business model discussions & plans for use of E4S
- Work with commercial providers to establish a support model
- Explore opportunities for joint product development

Industrial Users
- Engage in business model discussions & plans for use of E4S
- Work with commercial providers to establish a support model
- Explore models for mixed open-proprietary software stacks

Community Development

Broadening Participation of Underrepresented Groups in DOE Computing Sciences
- Coordinate consortium crosscutting-layer PIER planning
- Seek support for Sustainable Research Pathways Program
- Lead HPC Workforce Development and Retention Action Group

Better Scientific Software (BSSw) Fellowship Program
- Coordinate BSSw Fellowship Program – which gives recognition and funding to leaders and advocates of high-quality scientific software
- Seek sustainable support for BSSw Fellows, Honorable Mentions, travel, and program management for 2025 and beyond

PESO Services

Integration Partnerships

Provide resources and support for portfolio build, integration, and testing capabilities
- Spack integration
- CI testing
- Portfolio support & management

in collaboration with & co-funded by SSOs
- On-node & inter-node programming systems (w. S4PST)
- Math libraries, Data & viz, ML/AI (w. OASIS)
- Tools (w. STEP), Workflows (w. SWAS)
- NNSA software (funded by NNSA)

SQA & Security

Provide infrastructure to support and leverage product team SQA efforts
- Supply chain, Product quality
- Testing, Documentation

PESO Products

E4S
- Support for product integration
- E4S website
- Documentation, Training

Spack
- Features for consortium products
- Documentation, Training

Port & Test Platforms
- Frank test & development system
- Cloud resources
- Documentation, training

BSSw.io Content (w. COLABS)
- Short articles on topics related to scientific software productivity and sustainability (recruit, write, review, & edit)

EXASCALE COMPUTING PROJECT

PESO

Stakeholders:
Applications Community
Commercial HPC Companies
Industrial Users
US Agencies

DOE Computing Facilities:
ALCF NERSC OLCF

CRLC: Computational Research Leadership Council: ANL, BNL, LBNL, LLNL, LANL, ORNL, PNNL, SNL

PESO Advisory Board
Reps from ANL, LBNL, LLNL, LANL, ORNL, SNL

S3C Consortium
PESO, COLABS, CORSA, OASIS, STEP, SWAS, S4PST

DOE Program Managers
ASCR: Hal Finkel, Ben Brown, Saswata Hier-Majumder, Robinson Pino, Bill Spotz, David Rabson
NNSA: Si Hammond

PESO: Partnering for Scientific Software Ecosystem Stewardship Opportunities

Mike Heroux, SNL - PI **Lois Curfman McInnes, ANL - Co-PI**

PESO Partnerships

Stakeholder Engagement
(Mike Heroux, SNL)

Partnerships Coordinator
(Terece Turton, LANL)

Strategic engagement with consortium partners, applications, facilities, industry and agencies (in collaboration with and co-funded by SSOs)

- William Godoy, ORNL, On-node programming systems (w. S4PST)
- Rajeev Thakur, ANL, Inter-node programming systems (w. S4PST)
- Sameer Shende, Univ of Oregon, Tools (w. STEP)
- Sherry Li, LBNL, Math libraries (w. OASIS)
- Berk Geveci, Kitware, Data and viz (w. OASIS)
- Lavanya Ramakrishnan, LBNL, Workflows (w. SWAS)
- Mahantesh Halappanavar, PNNL, AI/ML (w. OASIS)

Unfunded partners: Strategic engagement with NNSA, communities of practice, applications, facilities, industry, agencies

- David Bernholdt, ORNL, RSE engagement (funded by COLABS)
- Addi Malviya-Thakur, ORNL, Foundation engagement (funded by CORSA)
- Elaine Raybourn, SNL, Consortium-wide community development (funded by CORSA)
- Ulrike Yang, LLNL, NNSA software (funded by NNSA)
- Partners at ALCF, NERSC, OLCF (funded by facilities, SW integration)

Community Development
(Lois Curfman McInnes, ANL)

Broadening Participation Initiative

- Mary Ann Leung, Sustainable Horizons Institute, PIER planning, lead of Sustainable Research Pathways (SRP)
- Daniel Martin, LBNL, lab lead of Sustainable Research Pathways
- Suzanne Parete-Koon, ORNL, lead of HPC Workforce Development and Retention Action Group

Better Scientific Software (BSSw) Fellowship Program

- Elsa Gonsiorowski, LLNL, Coordinator of BSSw Fellowship Program
- Erik Palmer, LBNL, Deputy Coordinator of BSSw Fellowship Program

PESO Services

Integration Coordinator
(Jim Willenbring, SNL)

Software portfolio management and integration (in collaboration with and co-funded by SSOs)

- Damien Lebrun-Grandie, ORNL, On-node prog systems (w. S4PST)
- Hui Zhou, ANL, Inter-node programming systems (w. S4PST)
- Bill Hoffman, Kitware, Tools (w. STEP)
- Satish Balay, ANL, Math libs (w. OASIS)
- Patrick O'Leary, Kitware, Data & viz (w. OASIS)
- Matteo Turilli, BNL, Workflows (w. SWAS)
- Sam Browne, SNL, NNSA software (funded by NNSA)

SQA & Security
(David Bernholdt, ORNL)

- Ross Bartlett (SNL)
- Berk Geveci (Kitware)
- Jim Willenbring (SNL)

PESO Products

E4S
(Sameer Shende, U Oregon)

- Luke Peyralans, Erik Keever, Wyatt Spear, Jordi Rodriguez

Spack (Todd Gamblin, LLNL)

- Greg Becker, LLNL
- Tammy Dahlgren, LLNL

Port & Test Platforms
(Gamblin & Shende)

- In partnership with Univ of Oregon, Cloud, etc.

BSSw.io Content (w. COLABS)

- Ross Bartlett, SNL
- Keith Beattie, LBNL
- Patricia Grubel, LANL
- Mark Miller, LLNL

Strategy & Integration – Members are part of other SSO teams, NNSA, for tight collaboration

ECP has been very active in outreach to industry and other US Government agencies

NOAA collaboration points

- E4S use on the cloud
- Allowed first successful run of GFLD Earth system model on the cloud
- Exploring Spack-stack builds

DoD collaboration points

- E4S deployment on five DoD systems (Cray EX, HPE SGI 8600)
- Spack tools extensively used
- Deep dive evaluation of 16 E4S packages (Ascent, Exaworks, Flux, Kokkos, RAJA, PETSc, LAMMPS, etc.)

TAE Technologies collaboration points

- Adopted WarpX, AMReX, Adios, Ascent, for fusion framework
- Extending WarpX with new physics models
- Went from ALCF Theta CPUs to Perlmutter GPUs

Date	ECP/IAC Events
April 11, 2023	NASA/AD deep dive
July 18, 2023	NASA/ST deep dive
August 17, 2023	IAC Quarterly Call
October 3, 2023	Continuous Integration Workshop
October 25, 2023	Final in-person meeting of the IAC members
November 16, 2023	Impact of Exascale on Industry Panel at SC23
January 31-Feb 1	Final in-person meeting of the IAC tech reps

ECP created programs to train the next generation pipeline for the DOE HPC workforce

Sustainable Research Pathways

https://shinstitute.org/sustainable-research-pathways

Summer 2022

- 61 participants at 9 labs: 13 student track, 16 faculty track (+29 students), 3 self-funded students
- 82% of overall participants represent at least 1 element of diversity

Summer 2023

- Multi-lab CRLC program spanning ECP and other computational and data science projects
- 189 faculty & students at 10 labs (120 funded thru ECP, 69 funded thru labs

Summer 2024

- ECP launch via SRP Matching Workshop, Oct 31-Nov 3, 2023
- Targeting @ 100 participants

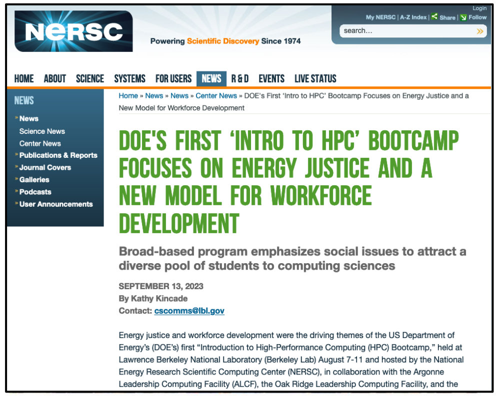
NERSC
Powering Scientific Discovery Since 1974
Login
My NERSC | A-Z Index | Share | Follow
search...
HOME ABOUT SCIENCE SYSTEMS FOR USERS NEWS R & D EVENTS LIVE STATUS
NEWS
- News
- Science News
- Center News
- Publications & Reports
- Journal Covers
- Galleries
- Podcasts
- User Announcements

Home » News » News » Center News » DOE's First 'Intro to HPC' Bootcamp Focuses on Energy Justice and a New Model for Workforce Development

DOE'S FIRST 'INTRO TO HPC' BOOTCAMP FOCUSES ON ENERGY JUSTICE AND A NEW MODEL FOR WORKFORCE DEVELOPMENT

Broad-based program emphasizes social issues to attract a diverse pool of students to computing sciences

SEPTEMBER 13, 2023
By Kathy Kincade
Contact: cscomms@lbl.gov

Energy justice and workforce development were the driving themes of the US Department of Energy's (DOE's) first "Introduction to High-Performance Computing (HPC) Bootcamp," held at Lawrence Berkeley National Laboratory (Berkeley Lab) August 7-11 and hosted by the National Energy Research Scientific Computing Center (NERSC), in collaboration with the Argonne Leadership Computing Facility (ALCF), the Oak Ridge Leadership Computing Facility, and the


SRP-HPC students and faculty, summer 2022

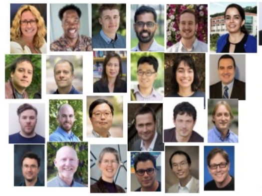
SRP-HPC mentors/co-mentors


Dan Martin, LBL
SRP thrust lead
for ECP


Keisha Moore, SHI
SRP Program
Coordinator

Intro to HPC Bootcamp

- Hosted at NERSC, August 7-11, 2023
- 60 students from 22 states
  - 48% first-generation scholars
- 14 project leaders & trainers (ANL, LBL, ORNL)
- 7 group projects: HPC topics related to energy justice
- Many have now applied to be part of SRP matching workshop for 2024
- **Website:** https://shinstitute.org/introduction-to-high-performance-computing-bootcamp

Argonne Training Program on Extreme-Scale Computing (ATPESC)

- Incorporated into ECP at project inception; 2-week training program at ANL focused on late graduate career, postdocs, early career scientists
- "*There are numerous HPC-related training offerings available across the community, but ATPESC is truly the standard by which all others are measured.*" -- Eric Neilsen, NASA Langley, Lead of FUN3D

EXASCALE COMPUTING PROJECT

ECP investments enabled a 100X improvement in capabilities

- **7 years** building an **accelerated**, **cloud-ready** software ecosystem
- Positioned to utilize **accelerators from multiple vendors** that others cannot
- **Emphasized software quality**: testing, documentation, design, and more
- Prioritized **community engagement**: Webinars, BOFs, tutorials, and more
- DOE **portability layers** are the credible way to
  - Build codes that are sustainable **across multiple GPUs** and
  - **Avoid vendor lock-in**
  - **Avoid growing divergence** and hand tuning in your code base
- ECP software can **lower costs** and **increase performance** for **accelerated** platforms
- Outside of AI, industry has not caught up
  - DOE enables an entirely different class of applications and capabilities to use accelerated nodes
  - In addition to AI
- **ECP legacy: A path and software ecosystem for others to leverage**

Software Ecosystem Lessons Learned

- **A hierarchical, process-driven approach to scientific library & tool development can work**
  - ECP produced a large collection of products that provide portable access to GPU-based systems
  - Working across labs, facilities, universities, vendors was extremely success, never done before for software
- **Software as a facility has promise**
  - ECP enabled us to operate a software organization as a facility for several years
  - A single organization with connections to all other HPC entities enabled DOE software to be a first-order entity in the organizational ecosystem
  - Many fruitful relationships grew out of this model


EXASCALE COMPUTING PROJECT

ECP-U-ST-RPT_2022_00285

ECP Software Technology Capability Assessment Report V3.0

Michael A. Heroux, Director ECP ST
Lois Curfman McInnes, Deputy Director ECP ST
Rajeev Thakur, Programming Models & Runtimes Lead
Jeffrey S. Vetter, Development Tools Lead
Sherry Li, Mathematical Libraries Lead
James Ahrens, Data & Visualization Lead
Todd Munson, Software Ecosystem & Delivery Lead
Kathryn Mohror, NNSA ST Lead
Terece L. Turton, Integration Lead

June 1, 2022

Performance Portability Lessons Learned

- **Node-level performance portability is possible**
  - Library approaches like Kokkos, RAJA, others can work well
    - Offer the first real possibility that apps can have one source code base, with perhaps some device-specific kernels
  - Markup approaches (OpenMP, OpenACC), vendor-specific (CUDA, HIP, SYCL) needed
    - But are best used underneath Kokkos and RAJA (so far)
- **Library support for diverse application node programming choices is possible**
  - Regardless of which choice the apps make (Kokkos, RAJA, OpenMP, OpenACC, CUDA, HIP, SYCL)
    - ECP libraries and tools can be compiled and configured to support that choice
- **ECP investments in node-parallel algorithms, libraries, and tools have broad impact**
  - All mainstream energy-efficient computing requires use of accelerator devices
  - ECP investments provide capabilities for NVIDIA, AMD, and Intel GPUs
  - ECP software strategies also greatly reduce the time and effort for porting to future accelerators

Software Governance Lessons Learned

- **ECP as a 413.3b project with tailored EVM provided a clear process for software governance**
  - **Hierarchical organizational structure with deep domain leadership enabled robust decision making**
    - **Structure gave ECP the ability to proactively manage cost, scope, schedule:**
      - Manage and adjust budgets: Annual reviews, plus ups, reductions
      - Coach underperforming teams: Better ECP alignment, better software practices
      - Off-ramp chronically mis-aligned or underperforming projects
      - Identify and address functionality gaps:
      - Iteratively plan, execute, track, and assess toward project goals
    - **Holistic scope established software libraries and tools as first-class entity in the HPC ecosystem**
      - SDK/E4S communities and products
      - Coordinated planning with computing facilities
      - Collaboration with vendor software teams, US agencies, international partners
- Post-ECP governance is still emerging but learning from ECP experiences

Concluding Remarks:

- **Scientific software ecosystems** are critical for the future of computational science – at all scales of computing
- **ECP legacy: A path and software ecosystem for others to leverage**
- **More than one way to leverage 100X**
  - 100X can be realized as exciting new science capabilities at the high end
    - Fundamental new science on leadership platforms
    - New opportunities on affordable machines that fit in current data centers
  - **But can also reduce costs**
  - Migration to accelerated platforms can be used to
    - Migrate a problem from an HPC cluster to a deskside or laptop systems
    - Lower your AWS monthly charges – E4S is available for container/cloud
    - Keep energy costs in check while still growing computing capabilities
  - **Biggest ECP impact will be accelerating GPU transition – at all levels**
  - **Transitioning software stacks to GPUs is essential**
    - CPU-based HPC systems realize only modest energy efficiency improvements
    - Migrating to GPUs is key to improving HPC environmental impact

We are capturing and sharing our lessons learned in technical, collaboration, and project management areas

slide credit: L. Diachin

Lessons Learned: Technical

- Performance portability, programming models, strategies for increasing arithmetic intensity, refactoring code, new algorithm design, etc..
- Strategies for moving from GPU-accelerated to GPU-resident
- When facing an inflection point in the HW, S/W investment must be a first class citizen
- Node-level solutions apply at all levels of computing

Lessons Learned: Collaboration

- The value of diverse, multi-disciplinary teams
- Collaborative solutions can't be dictated but they can be incentivized
- Build integration into project structure and measures of success
- High-quality software is the foundation for collaboration in scientific computing
- Open, frequent communication (good/bad/ugly) is imperative with sponsors, stakeholders, staff

Lessons Learned: Project Management

- Projectizing R&D works if agile project management and aggressive change control are in place
- Empower the leadership team then hold them accountable
- Understand and manage external dependencies
- Highly functioning diverse leadership team are a must
- Good centralized project management tools do not guarantee success but bad ones can sure impede progress
- Stability and long term planning results in remarkable innovation

Thank you

Abstract

The U.S. Department of Energy's (DOE) Exascale Computing Project (ECP) recently successfully completed its work in developing a capable exascale computing ecosystem comprising applications, software technologies, and deployment and integration capabilities. We discuss major accomplishments and lessons learned by the ECP community over the course of seven years on the development of an integrated scientific computing software stack (which enables and fosters success on a wide variety and scales of computers) and the demonstration of new physics capabilities in a wide variety of scientific applications. We emphasize issues in creating the exascale ecosystem, particularly in algorithm design and implementation for accelerator-based compute nodes, performance portability across a range of platforms, fostering strong collaborations across multidisciplinary teams, and managing and measuring the success of a computational science project of this scale.

ISC 2024 Workshop:

First International Workshop on Readiness of HPC Extreme-scale Applications

- https://pop-coe.eu/news/events/readiness-of-hpc-extreme-scale-applications