Toward a Post-ECP Software Sustainability Organization (PESO)

Team: Michael Heroux (Sandia National Laboratories; PI)
James Ahrens (Los Alamos National Laboratory)
Todd Gamblin (Lawrence Livermore National Laboratory)
Timothy Germann (Los Alamos National Laboratory)
Xiaoye Sherry Li (Lawrence Berkeley National Laboratory)
Lois Curfman McInnes (Argonne National Laboratory)
Kathryn Mohror (Lawrence Livermore National Laboratory)
Todd Munson (Argonne National Laboratory)
Sameer Shende (University of Oregon)
Rajeev Thakur (Argonne National Laboratory)
Jeffrey Vetter (Oak Ridge National Laboratory)
James Willenbring (Sandia National Laboratories)

Project Summary: *Toward a Post-ECP Software Sustainability Organization (PESO)* will develop a plan for a software-ecosystem sustainment organization in collaboration with the many stakeholders and partners that engage with the US Department of Energy in the development and use of high-performance computing (HPC) capabilities. The PESO plan will provide a smooth transition of activities from the DOE Exascale Computing Project (ECP) toward the goal of creating a software ecosystem that will deliver high-quality, reliable, and trusted libraries and tools for DOE's mission-critical applications.

PESO efforts will emphasize the broad engagement of the HPC community to assure that the post-ECP organization leverages the knowledge and infrastructure from ECP, while at the same time capturing the contributions and needs of HPC community members who were not engaged in ECP. We believe that an emphasis on software product quality, the continued fostering of software product communities, and the delivery of products via a trusted open-source software stack are key to the success of the PESO plan.

The PESO plan will provide product teams with autonomy and control while leveraging the potential of a team-of-teams approach, managing a fully capable and customizable software stack, and effectively and efficiently using sponsor funding. The PESO plan is guided by principles that place decision-making at the lowest level possible, integrate across collaborators, distribute cost and benefit sharing, create a diverse and inclusive workforce with sustainable career paths, commit to productivity and sustainability improvement, and engage with external community members.

Contents

1 Introduction

The Exascale Computing Project (ECP) is producing a collection of libraries and tools that are essential to running large-scale scientific applications of national interest on production exascale systems[31;46]. As noted in the report *Transitioning ASCR after ECP*[23], it is essential to sustain continued research and development, curation, packaging, testing, and distribution of this software after the completion of ECP, building lasting partnerships with other organizations in the HPC community to improve the cost and benefit sharing of DOE software investments. Likewise, page 175 of the recent *CHIPS and Science Act*[14] specifically calls out the need to sustain the ECP software ecosystem.

In this proposal, leveraging our experience of forming, leading, and executing many of the efforts within ECP, we articulate our long-term vision for sustaining this software ecosystem and describe how we intend to use FY23 funding to flesh out the details of our plan and collaborations. While our ECP experience provides an important foundation for future efforts, we will intentionally step back from our ECP approach to re-baseline. ECP's sharp focus on a specific set of applications, libraries, and tools, and the highly-structured nature of the project necessarily omitted important members of the DOE and HPC community and imposed a strong hierarchical structure on our efforts. Our PESO planning efforts will purposefully include the broad DOE and HPC community by focusing on fundamental goals and relationships, exploring the post-ECP organizational design space on our way to developing a plan. Our FY23 activities will result in a clear strategy for sustaining a broad collection of software and will inform ASCR on how we plan to organize and fund an effective, long-term HPC software sustainability effort in FY24 and beyond.

2 Lessons learned from ECP

ECP has provided a unique opportunity for the DOE scientific libraries and tools communities by sponsoring sustained multi-year funding that is structured to promote cross-institutional collaboration at a scale we have not experienced before. Within this context, we have been able to plan and develop portfolio-based software ecosystems that provide clear value to ECP stakeholders. We have emphasized software quality through the use of software quality community policies, automated testing and continuous integration, and we have fostered product communities through the use of Software Development Kits (SDKs) that provide a common structure for cross-team collaboration. These efforts have been successful in producing a trusted software stack that is ready for production use on DOE exascale systems and that is also generally useful to the broader HPC community.

While high-quality software and collaboration with partners are already central to DOE HPC efforts, ECP has furthered that progress, especially in the development and delivery of reusable libraries and tools. ECP's duration and funding stability have enabled an increased focus on software quality, coordinated delivery of products, and expanded relationships with application teams, leadership facilities, computer system vendors, and more. We believe that ECP efforts to develop and deliver software libraries and tools have provided tangible evidence that an explicit focus on scientific software-ecosystem sustainment makes sense and can work. From our ECP efforts, we plan to carry forward the following lessons learned.

2.1 Fostering a culture of software quality

Every scientific software team strives to produce high-quality software, but other goals compete for time and resources. ECP has provided funding, guidance, recognition and support for teams

to focus on software quality. Among many investments, the IDEAS-ECP project[29;36] within ECP supports product teams with training and resources to improve developer productivity and software sustainability, as key aspects of improving overall scientific productivity and helping to ensure confidence in scientific results. We will continue to support teams in their efforts to improve software quality, and we will continue to emphasize the importance of software quality in future efforts. The PESO plan focuses on how to use our ECP experience to continue fostering a culture of software quality and to continue to improve the quality of the software we deliver.

2.2 Creating intentional product communities

ECP's Software Development Kits (SDKs) establish collaborative structures for product communities. The SDK approach grew out of the original IDEAS xSDK[7;60;66;67] project that started in 2014 to provide cross-team collaboration among math libraries teams by design. The activities conducted within the SDKs have been very effective at accelerating design space exploration and making compatible collections of libraries and tools that benefit users, facilities, and the product development teams themselves. The impact of these multi-team efforts extends far beyond what ECP funds directly, including teams from many different organizations who bring their products and independent funding to the activities. One example of the SDK model's success is the multi-precision effort funded under the xSDK project within ECP. This effort has led to numerous insights and product improvements that would be nearly impossible to accomplish if each independent project conducted the work on their own. The impact of this work is evidenced by the collection of publications[1;2;4;5;12;25;26;32;40;42;68] that have resulted from this funding and by the descriptions in these publications of how math library teams have been able to leverage low-precision hardware capabilities in modern GPUs.

2.3 Delivering a trusted software stack

The goals of ECP required the creation and delivery of software features in dozens of reusable libraries and tools from hundreds of developers over the span of the project. To effectively and efficiently manage this effort, we created the Extreme-scale Scientific Software Stack (E4S)[19;45;48], a curated stack that relies on the Spack software management ecosystem[61]. While E4S is still in its early stages of community use and Spack is still growing rapidly in its adoption, the combination of these capabilities has been very effective at delivering a trusted software stack that is used by many stakeholders. Spack is the preferred software installation tool, and E4S is used at all DOE supercomputing sites[3;24;50;51] and beyond[18;30;33;34;35;43;44;47;53;54;55;56;57;58;59]. We will continue to use E4S and Spack to deliver a trusted software stack that is ready for production use.

2.4 Creating communities

ECP has given DOE scientific software staff an experience of working across institutions and product teams to better engage with each other and with our stakeholders. An essential element of the post-ECP software ecosystem is continuing and growing our ability to work together in these ways. ECP sponsors annual meetings[20], webinars[9;62], web portals and recognition, such as the Better Scientific Software (BSSw) Fellowship Program[10;11]. We will continue to foster collaboration and create intentional communities that are more than the sum of their parts. We will continue to emphasize the importance of software quality and communities, and we will continue to deliver a trusted software stack that is ready for production use.

2.5 PESO team

The PESO team is composed of the level-2 and level-3 leads of the ECP Software Technology (ST) Focus Area, along with the level-3 lead of co-design and the project leads of E4S, SDKs, and Spack: Mike Heroux (ST director), Lois Curfman McInnes (ST deputy director), Rajeev Thakur (programming models and runtimes), Jeff Vetter (development tools), Sherry Li (math libraries), Jim Ahrens (data and visualization), Todd Munson (software ecosystem), Kathryn Mohror (NNSA software), Tim Germann (co-design), Sameer Shende (E4S), Jim Willenbring (SDKs), Todd Gamblin (Spack). Heroux and McInnes also co-lead the IDEAS-ECP project. PESO members have decades of experience and leadership in the collaborative development and delivery of scientific libraries and tools and comprise the team that led the creation of the SDK and E4S efforts within ECP. The PESO team will continue to work together to build on ECP experience and to create a plan that will maximize the impact of DOE software investments on DOE's mission.

3 PESO vision: Better software, better science

Software is a critical component of DOE's mission to advance scientific discovery and innovation. Reusable libraries and tools that underpin and facilitate application codes must be reliable, performant, and easy to use. They must be developed and delivered in a way that is sustainable and responsive to the needs of the DOE HPC community. Finally, and most importantly, they must be available and trusted.

DOE-sponsored software teams have provided tremendous capabilities to the HPC community for decades. DOE's mission to sponsor hardware and software solutions for next-generation HPC computations provides motivation and direction for DOE-funded scientists to produce new capabilities that address DOE mission needs and are also generally useful to the broader HPC community. Even so, the impact of DOE software investments is often limited by the lack of a sustainable software ecosystem that can deliver high-quality software to DOE applications and facilities.

PESO will create a plan that builds on and continues ECP efforts in three areas: **facilitating product improvement while respecting team automony, creating intentional product communities, and delivering a trusted software stack that is ready for production use**. We will also focus on the development of a **sustainable funding model for the software ecosystem that is aligned with the needs of the DOE HPC community.**

The PESO vision is that DOE software investments will have maximum impact on DOE's mission by providing high-quality libraries and tools that deliver the latest in HPC algorithms and capabilities to serve the needs of DOE applications. In addition to doing excellent work and delivering it to our users, we believe that maximum impact also requires broad and deep partnerships with other stakeholders and members of the HPC community. We will build on relationships established in ECP and create new partnerships to maximize the benefit and cost-sharing of DOE software investments.

At the same time, even though ECP provides an example for future software-ecosystem sustainment, ECP is a federally funded construction project with high-consequence Key Performance Parameters (KPPs) and a formal hierarchical work breakdown structure. ECP leadership, including the members of the PESO team, have worked to reduce overhead, but ECP project requirements necessitate top-down planning, execution, tracking, and assessment that we believe can be reconfigured in a post-ECP software sustainability organization to be **lighter weight** and **more consistent with scientific software culture**. Furthermore, post-ECP plans must include a wider set of applications, be agile in meeting changing requirements, and involve many people who were not part of ECP. While many details must be resolved, the plans we create in PESO will

prepare the DOE community for moving toward post-ECP software-ecosystem sustainment.

4 PESO goals

The goal of PESO is to develop a plan for a software sustainability organization that can be instantiated quickly and flexibly as funding resources become available in FY24 and beyond. The plan will leverage the PESO team's ECP experience and emphasize collaboration with DOE stakeholders; other DOE efforts in libraries, tools, applications, and facilities; and numerous external organizations essential to overall DOE software success. **The ultimate goal of PESO efforts is establishing the foundation for a software ecosystem that delivers robust, reliable and trusted libraries and tools for leadership scientific discovery in support of DOE's mission**.

Application success on leadership systems: The fundamental goal of the PESO plan is to enable the success of important application codes by providing high-quality, reusable, and sustainable scientific libraries and tools that provide unambiguous value. Ultimate success hinges on identifying DOE mission needs as determined by national priorities that can be addressed through high-end computing. Identifying and prioritizing our efforts must involve all primary stakeholders, including DOE program managers and research programs, DOE facilities, application developers, and DOE-funded library and tool development teams.

Ecosystem collaboration: Beyond DOE, we collaborate with computer system vendors, commercial scientific software providers, other US agencies, US industry, and international partners. The impact of DOE software investments is greatly amplified through collaboration, sharing the costs and benefits with organizations within and outside of DOE.

End-to-end development and delivery of software capabilities: Scientific software capabilities are typically developed in collaboration with a variety of partners depending on the type of product. One PESO goal is to support product teams by providing the infrastructure and workflows that will enable the sustainable integration of new capabilities into the broader software ecosystem.

Effective and efficient use of sponsor funding: The PESO goals of maximizing the cost and benefit sharing of our products and our desire to deliver high-quality products are driven by our desire to make the best use of sponsor funding.

5 PESO software ecosystem organizational strategy

ECP organizational structures provide significant experiences that inform the PESO plan. We anticipate creating a cross-institutional software organization that will take a fresh look at how we:

1. Provide each product team with autonomy and control over its destiny as it produces the capabilities needed by its funding sponsors.

2. Leverage the potential of a team-of-teams approach to collaborating and coordinating across products in a suite leveraging our SDK experience.

3. Realize the potential of managing a fully capable and customizable software stack that delivers value to many stakeholders within and outside of DOE building on our E4S capabilities.

ECP has enabled the creation of a three-level software ecosystem model that we expect to adapt and carry forward as part of our PESO efforts. While we expect changes from the ECP structure, in particular substantially reducing formality, we believe there is intrinsic value to this model and a leadership structure to guide efforts (Figure 1).

Figure 1: PESO software ecosystem

- **Product team:** Fundamental building block of the PESO plan. In our vision, software product teams set the primary direction for algorithmic, software architecture, and implementation activities for their products in collaboration with their sponsors, stakeholders, and co-design partners. Which teams will participate in the outcome of the PESO plan and how funding will be determined are a central focus for the early phase of our FY23 efforts, as expressed by the governance principles and technical approach outlined below.

- **Product community:** Collection of compatible and complementary products[1] that facilitate interaction among product teams in a variety of activities such as fundamental requirements, design space exploration, training, and the evolution of software practices and tools. The product suite concept was established first in the math libraries community where dozens of library teams (some funded by DOE and others not) collaborate and coordinate on design space exploration (e.g., use of mixed and multi-precision computations), documentation, training, and workshops. Coordination of versioning, vendor interactions, design space exploration, and software delivery at the product suite level helps amortize costs across related products and reduces complexity at the software stack level.

- **Software stack:** Full collection of products that is configurable to provide a subset of products or integrate external and compatible products. Using Spack, we can incorporate external, including proprietary, products. We can also create custom builds that target many environments—from laptop to cluster, leadership, and edge. We also provide binary and containerized environments and manage versioning across products for reproducibility and correctness/security patching. Actively managing a full portfolio of configurable HPC software products that are built and tested on a variety of HPC platforms, including on the leadership platforms under development, provides a rich resource of version-compatible

[1] In ECP we call this collection a Software Development Kit (SDK) but have found the term to be somewhat confusing and may adopt a new and more effective name.

products, built under numerous different parameter settings, and tested on many systems.

6 PESO governance principles

While PESO planning will be shaped by the technical approach described below, we will adhere to the following principles to guide our decisions:

- **Place decision-making ownership at the lowest level possible:** As we determine governance models, we will work to assure that decision-making is placed at the product team level as much as possible, then the product community level, and then the stack level.
- **Integrate across lab, industry, and university collaborators:** Success in HPC requires many diverse skills and knowledge bases. Within a product team, within a product suite, and in the enterprise of providing the software stack, we will promote the integration of all types of collaborators to eliminate organizational barriers and encourage better outcomes.
- **Distribute cost and benefit sharing:** High-quality, widely available software products enable broad contributions across institutional and organizational boundaries and enable the confident use of software from outside organizations. We will improve impact and reduce overall costs, especially through industry and agency collaborations.
- **Create a diverse and inclusive workforce with sustainable career paths:** The PESO PIER plan (see Appendix E) will assure that all members are valued for their contributions and have the opportunity to grow in their position with leadership opportunities. We will collaborate with other organizations to assure that we are an attractive organization for recruiting, on-ramping, and retaining staff.
- **Commit to productivity and sustainability improvement:** We will invest in new ways to improve developer productivity and software sustainability. These investments will be particularly important with the advent of advanced software development environments and the increased success of incorporating modern software practices, processes, and tools into scientific software activities.
- **Engage with external community members:** We will establish external advisory functions wherever guidance and input from the broader stakeholder community are essential. We will seek input at a variety of scope levels to assure we are making good tactical and strategic decisions.

7 PESO FY2023 plan

Given the above vision, goals, and governance principles, we will use FY23 funds to create the PESO plan. The overall plan will emerge from executing four primary planning thrusts that build on ECP experience and the lessons learned from the DOE software ecosystem efforts:

- **Product portfolio management:** How we will initiate and manage the product portfolio. We will leverage the ECP product dependency database and strategy, expanding the set of applications, facilities and user community requirements. We will also establish a transparent and fair assessment process for managing the portfolio with an emphasis on sustainability and transitioning of support to commercial and community partners, to assure the ability to bring in new work each year.

- **Organizational structure and execution:** How we will organize and operate. We will explore lighter-weight approaches relative to ECP, including a more distributed decision-making model.
- **Community engagement:** How we will collaborate with key community partners. We will leverage and expand existing ECP partnerships with DOE facilities, applications, other US agencies, US industry, vendors and commercial partners.
- **Advisor and stakeholder engagement:** How we will inform project efforts to complement and coordinate within the broader ecosystem to better assure success and mitigate risks.

7.1 Product portfolio management strategy

Establishing and managing a product portfolio is essential for overall success. If the management approach is too strong, product teams will be reluctant to participate. However, if the approach is too loose, we risk losing the tremendous value that product communities and robust software stacks provide. While this effort is certainly part of the general organizational structure and execution, we need a portfolio management plan in place as soon as possible to deploy our full plan without disruption as ECP ends.

- **Initial portfolio:** We will define a process to determine which software products will initially be part of the sustainability effort. Decisions will be based on dependency analysis, input from stakeholders (applications, facilities, industry), as well as needs identified by future technology trends. We expect the initial set to have some additions to and deletions from the current set of ECP products.
- **Ongoing portfolio:** In addition to initially identifying and selecting products, we will also need to determine what efforts should continue in the portfolio or be transitioned out. Some efforts could be transitioned to support by a vendor or commercial software company. Others could be off-ramped as a good idea that didn't result in a sufficiently useful capability for continued support. In all cases, we will need to take care that our decisions provide the continuity that product users and developers need.

7.2 Organizational structure and execution strategy

Via meetings and workshops, in collaboration with the broader community, we will explore and define important elements of a sustainable software ecosystem. We will use a mix of approaches, including regular virtual and in-person meetings and workshops, the continuation of the Leadership Scientific Software (`https://lssw.io`) townhalls for exploring key themes in a public forum, and regular briefings to our advisors. Some of the key planning topics are:

- **Funding models** for long-term support of this effort, including a mix of research grants and sustainability funds for individual products, potential contributions from stakeholders (facilities, other DOE programs, applications, industry), as well as subscription models where possible.
- **Approaches for effective and efficient development and delivery** of new software product versions that support teams in sustainably providing new capabilities to their users.
- **Community organization models** that build on the ECP experience of creating the E4S stack, its integration, testing, and support at DOE facilities, and build on the concepts of software development kits (SDKs) to improve collaboration and complementarity.

- **Approaches for improved user support,** including better software usability, cost-effective sustainability, high-quality documentation, and better software availability at all levels of HPC development and use.
- **Development and user partnership models** with stakeholders, especially with facilities, application teams, industry partners, open-source communities (e.g., LLVM development community), and other agencies.
- **Staffing approaches** that improve diversity and inclusion and assure stable career paths.
- **Change management** approaches for regular project reviews and mid-cycle changes.
- **Progress assessment** strategies and cadences for evaluating our software product and collaboration models and processes. Our goal will be to be as effective and efficient as possible.

7.3 Community engagement strategy

We expect the PESO planning effort to be a collaborative process involving other funded FY23 software sustainment projects, the developers of software products, application stakeholders, DOE program managers, DOE lab leadership, industry and agency partners, and the HPC community at large.

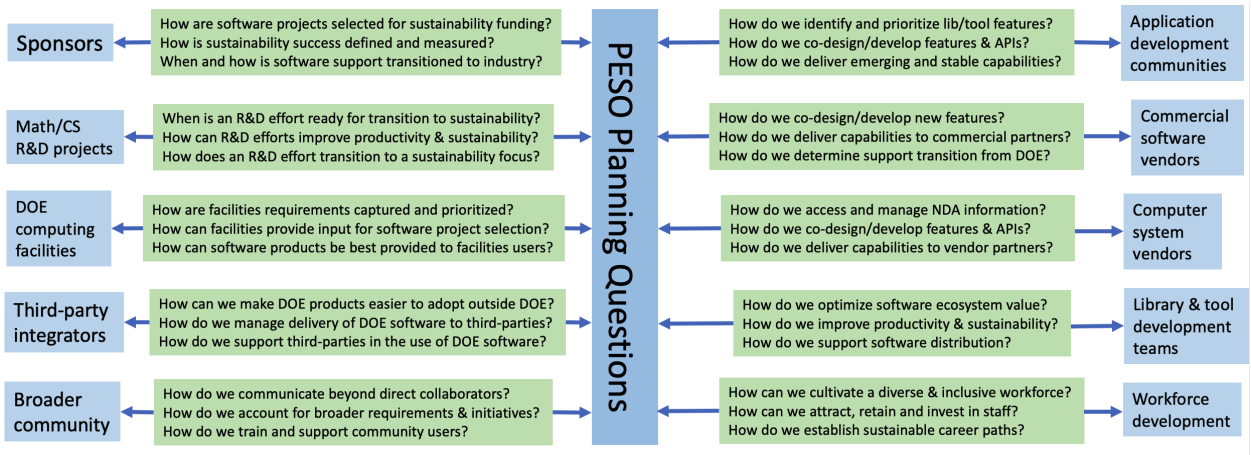

Figure 2: Through townhalls, workshops and ongoing dialogue, PESO will focus on addressing these key questions, building on our ECP experience.

Figure 2 shows some of the key questions we anticipate addressing as part of our PESO community engagement strategy. In each situation, we will look for ways to improve the effectiveness and efficiency of our interactions and build new relationships. The PESO plan will emerge from exploring possible models. Specific engagements include:

- **Sponsors:** DOE sponsors HPC algorithms and software development through multiple programs. DOE ASCR Research, especially SciDAC, and the NNSA ASC program are the most notable, but laboratory LDRD, CRADA, and similar programs also provide significant funds for the development and delivery of new HPC libraries and tools. An essential part of the PESO plan is how to engage with sponsors, especially when developing capabilities that have

the potential for addressing the needs of multiple sponsors, where opportunities for improved cost and benefit sharing exist.

- **Math/CS R&D projects:** Through DOE sponsor funding, algorithms and software R&D projects demonstrate new capabilities that provide the primary input to new library and tool development. An essential part of the PESO plan will be how to effectively and efficiently capture, harden, and deliver new R&D ideas to assure broad scientific community impact.

- **Application development communities:** Application teams are essential partners in all phases of library and tool development. Some application teams work closely with us to identify requirements and design; some teams co-design and co-develop capabilities, while others use the final hardened products. In all phases of design and development, partnering with application communities is essential.

- **DOE computing facilities:** DOE-sponsored libraries and tools play an important role in supporting application codes in the effective use of leadership systems, providing some of the custom capabilities needed to realize the performance potential of these systems. Understanding emerging needs, knowing future system capabilities, and working with facilities and vendor staff are essential to make sure that DOE-sponsored libraries and tools are ready when users need them. In addition, facilities staff are typically the first people that application teams engage when experiencing problems. The PESO plan will build on the E4S facilities support infrastructure established by ECP.

- **Computer system and commercial software vendors:** HPC systems and commercial software vendors provide supported and optimized versions of many products to which DOE contributes. Examples include DOE contributions to the LLVM stack (OpenMP, Fortran, and HPC optimization features), optimized versions of portability layers like Kokkos and RAJA, MPI libraries, performance tools such as PAPI, math libraries, I/O libraries and tools, and more. This collection of relationships will be an important focus of our plan.

- **Third-party software integrators:** E4S is demonstrating that the availability of a stable and robust software stack enables building products on top of what DOE teams provide. For example, application teams outside of DOE increasingly use DOE-sponsored libraries and tools via E4S; the certainty of future availability and commercial support would be beneficial. Furthermore, other language environment communities can make DOE-sponsored libraries and tools available to their user communities. The Julia community is perhaps the most attractive opportunity, and we know that leaders within the Julia community have an interest in providing our software, generally written in C++, to the large Julia community, further expanding the benefit and cost-sharing of DOE software investments.

- **Library and tool development teams:** The first governance principle in Section 6 commits to keeping decision-making at the lowest possible level, assuring that product teams have the necessary autonomy to control their activities. Essential to the PESO plan is optimizing the potential of DOE-sponsored library and tool products through lightweight portfolio management approaches. The PESO plan must assure that being part of a product suite and a software stack has demonstrable value so product teams will engage in portfolio activities.

- **Workforce development and retention:** The PESO plan is committed to cultivating meaningful, satisfying, and stable career paths for all staff who participate in the post-ECP software ecosystem efforts. Our PIER Plan (Appendix E) describes our current and future

Table 1: PESO deliverables

Deliverable	Description
Workshops	Quarterly workshops involving important stakeholders. An initial workshop will focus on defining the overall scope and limits of PESO plans. Later workshops will focus on creating the PESO plan so it is ready to deploy.
Townhalls	Monthly Leadership Scientific Software (LSSw) Townhalls (`https://lssw.io`). Resuming from last year, these virtual events will engage the broad scientific software community to inform our design requirements in developing the PESO plan.
Portfolio management plan	As soon as feasible, we will create a plan to define the initial product portfolio. Knowing which libraries and tools will be part of the post-ECP effort is essential for a smooth transition as ECP ends.
Budget options	Plans at several budget levels that account for possible splits of the budget, the number of products to support, and the possible scope splits between software-ecosystem sustainment and other DOE ASCR funding sources.
Collaboration structure	Using principles from Section 6, depending on the number and kind of post-ECP sustainability efforts, and the relationships defined in our PESO plan, we will establish collaborations that will optimize ecosystem cost and benefit sharing.
Full PESO plan	The final PESO deliverable will be a plan that is ready to deploy without disruption as the Exascale Computing Project ends.

workforce plans that build on our existing commitments to improving diversity, equity, and inclusivity.

- **Broader user and stakeholder communities:** Tutorials, user group meetings, and similar outreach activities are essential for informing the broader community about our capabilities. Furthermore, we need to understand emerging needs and capabilities in the broader community to inform our plans and activities.

7.4 Advisor and stakeholder engagement strategy

Building on DOE, ASCR, facilities, and ECP industry and agency council relationships established during ECP, and incorporating new relationships mentioned in Section 7.3, we will establish an advisor and stakeholder advisory structure. We anticipate having an **advisory council** of distinguished community members who understand the broad context of our work and can provide input to programmatic strategy and a technical advisory committee that can guide technical strategy. Determining the advisory structure will be part of FY23 activities and will depend on the scope of funding expected in FY24 and beyond.

8 Deliverables

The primary deliverable of PESO will be a plan for launching a software-ecosystem sustainment organization in FY24. The plan will be configurable depending on funding availability, the kind of software work that will be directly sponsored by the organization relative to other funding sources, and the number of products we will support. We anticipate that the plan will emerge and stabilize through community engagement activity as described in this proposal. Our specific deliverables are listed in Table 1.

A Bibliography and References Cited

References

[1] Ahmad Abdelfattah, Hartwig Anzt, Alan Ayala, Erik Boman, Erin Carson, Sebastien Cayrols, Terry Cojean, Jack Dongarra, Rob Falgout, Mark Gates, et al. Advances in mixed precision algorithms: 2021 edition. Technical report, Lawrence Livermore National Lab. (LLNL), Livermore, CA (United States), 2021.

[2] Ahmad Abdelfattah, Hartwig Anzt, Erik G Boman, Erin Carson, Terry Cojean, Jack Dongarra, Alyson Fox, Mark Gates, Nicholas J Higham, Xiaoye S Li, et al. A survey of numerical linear algebra methods utilizing mixed-precision arithmetic. *The International Journal of High Performance Computing Applications*, 35(4):344–369, 2021.

[3] E4S at ALCF. `https://www.alcf.anl.gov/events/e4s-hackathon`, 2021.

[4] Hartwig Anzt, Eric G Boman, Marj Gates, Scott Kruger, Sherry Li, Jennifer Loe, Daniel Osei-Kuffuor, Stan Tomov, Yaohung M. Tsai, and Ulrike Meier Yang. Towards use of mixed precision in ECP math libraries. Technical report, Lawrence Livermore National Lab. (LLNL), Livermore, CA (United States), 2021.

[5] Hartwig Anzt, Goran Flegar, Vedran Novaković, Enrique S Quintana-Ortí, and Andrés E Tomás. Residual replacement in mixed-precision iterative refinement for sparse linear systems. In *International Conference on High Performance Computing*, pages 554–561. Springer, Cham, 2018.

[6] Ashley Barker, Dan Martin, Mary Ann Leung, Lois Curfman McInnes, Jini Ramprakash, Julia White, Jim Ahrens, Erik Draeger, Shaun Fomby, Yasaman Ghadar, Rinku Gupta, Mahantesh Halappanavar, Mike Heroux, Christopher Kelly, Mark Miller, Hai Ah Nam, Slaven Peles, Damian Rouson, Dan Turner, Terry Turton, David Brown, and Valerie Taylor. A multipronged approach to building a diverse workforce and cultivating an inclusive professional environment for DOE high-performance computing, 2021. A response to the Stewardship of Software for Scientific and High-Performance Computing Request for Information, `https://doi.org/10.6084/m9.figshare.17192492`.

[7] Roscoe Bartlett, Irina Demeshko, Todd Gamblin, Glenn Hammond, Michael Allen Heroux, Jeffrey Johnson, Alicia Klinvex, Xiaoye Li, Lois Curfman McInnes, J. David Moulton, Daniel Osei-Kuffuor, Jason Sarich, Barry Smith, James Willenbring, and Ulrike Meier Yang. xsdk foundations: Toward an extreme-scale scientific software development kit. *Supercomputing Frontiers and Innovations*, 4(1):69–82, Feb. 2017.

[8] D. Bernholdt, J. Cary, M. Heroux, and L.C. McInnes. The Science of Scientific-Software Development and Use: Workshop Brochure. `https://doi.org/10.2172/1846008`, 2022.

[9] Best Practices for HPC Software Developers: Webinar series. `https://ideas-productivity.org/events/hpc-best-practices-webinars/`.

[10] Better Scientific Software (BSSw) Website. `https://bssw.io`.

[11] Better Scientific Software (BSSw) Fellowship. `https://bssw.io/fellowship`.

[12] Erin Carson, Nicholas J. Higham, and Srikara Pranesh. Three-precision GMRES-based iterative refinement for least squares problems. *SIAM Journal on Scientific Computing*, 42(6):A4063–A4083, 2020.

[13] B. Chapman, H. Calandra, S. Crivelli, J. Dongarra, J. Hittinger, S. Lathrop, V. Sarkar, E. Stahlberg, J. Vetter, and D. Williams. DOE Advanced Scientific Advisory Committee (ASCAC): Workforce Subcommittee Letter. `https://dx.doi.org/10.2172/1222711`, 2014.

[14] CHIPS and Science Act. `https://science.house.gov/chipsandscienceact` and `https://science.house.gov/imo/media/doc/the_chips_and_science_act.pdf`. United States Congress, July 2022.

[15] National Laboratory Directors' Council. Diversity demographics of the 17 DOE National Laboratories. `https://nationallabs.org/staff/diversity`, 2022.

[16] National Science & Technology Council. National Strategic Computing Initiative Update: Pioneering the Future of Computing. `https://www.nitrd.gov/pubs/National-Strategic-Computing-Initiative-Update-2019.pdf`, 2019.

[17] National Science & Technology Council. Pioneering the Future Advanced Computing Ecosystem: A Strategic Plan. `https://www.nitrd.gov/pubs/Future-Advanced-Computing-Ecosystem-Strategic-Plan-Nov-2020.pdf`, 2020.

[18] The Second E4S Extreme-Scale Scientific Software Stack Forum Workshop at EuroMPI/USA 2020. `https://oaciss.uoregon.edu/E4S-Forum20`, 2020.

[19] E4S Web page. `http://e4s.io`.

[20] ECP Annual Meeting Web page. `https://ecpannualmeeting.com`.

[21] ECP Broadening Participation Initiative. `https://www.exascaleproject.org/hpc-workforce`, 2022.

[22] Lisa Arafune et al. CI Workforce Development Workshop 2020. `https://www.rcac.purdue.edu/ciworkforce2020`.

[23] R. Giles et. al. Transitioning ASCR after ECP. `https://science.osti.gov/-/media/ascr/ascac/pdf/meetings/202004/Transition_Report_202004-ASCAC.pdf`, 2020.

[24] Oak Ridge Leadership Computing Facility. Extreme-Scale Scientific Software Stack (E4S). `https://www.olcf.ornl.gov/e4s/`, 2019.

[25] Goran Flegar, Hartwig Anzt, Terry Cojean, and Enrique S. Quintana-Ortí. Adaptive precision block-Jacobi for high performance preconditioning in the Ginkgo linear algebra software. *Proceedings of ISC High Performance*, 47(2), 2021.

[26] Fritz Göbel, Thomas Grützmacher, Tobias Ribizel, and Hartwig Anzt. Mixed precision incomplete and factorized sparse approximate inverse preconditioning on GPUs. In *European Conference on Parallel Processing*, pages 550–564. Springer, Cham, 2021.

[27] William F. Godoy, Ritu Arora, Keith Beattie, David E. Bernholdt, Sarah E. Bratt, Daniel S. Katz, Ignacio Laguna, Amiya K. Maji, Addi Malviya Thakur, Rafael M. Mudafort, Nitin Sukhija, Damian Rouson, Cindy Rubio-Gonzalez, and Karan Vahi. Giving RSEs a larger stage

through the Better Scientific Software Fellowship. https://arxiv.org/pdf/2211.07436.pdf, 2022. Submission to IEEE CiSE.

[28] Evelynn Hammonds, Valerie Taylor, and Editors Rebekah Hutton. Transforming trajectories for women of color in tech. https://doi.org/10.17226/26345, 2022. National Academies of Sciences, Engineering, and Medicine, Washington, DC: The National Academies Press.

[29] M. A. Heroux, L. C. McInnes, D. E. Bernholdt, A. Dubey, E. Gonsiorowski, O. Marques, J. D. Moulton, B. Norris, and E. M. Raybourn et al. Advancing Scientific Productivity through Better Scientific Software: Developer Productivity and Software Sustainability Report. ECP Technical Report ECP-U-RPT-2020-0001, 2020, http://doi.org/10.2172/1606662.

[30] Michael Heroux, Lois Curfman McInnes, Rajeev Thakur, Jeffrey Vetter, Sherry Li, James Ahrens, Todd Munson, Kathryn Mohror, Timothy Germann, Ryan Adamson, Todd Gamblin, Sameer Shende, and James Willenbring. A Response to the Stewardship of Software for Scientific and High-Performance Computing Request for Information. https://github.com/tgamblin/ascr-software-stewardship-rfi-responses, 2022.

[31] Michael A. Heroux, Lois Curfman McInnes, Rajeev Thakur, Jeffrey S. Vetter, Xiaoye Sherry Li, James Ahrens, Todd Munson, Kathryn Mohror, Terece L. Turton, and ECP Software Technology teams. ECP Software Technology Capability Assessment Report, V3.0, June 2022. https://www.osti.gov/biblio/1888898.

[32] Nicholas J. Higham and Theo Mary. Mixed precision algorithms in numerical linear algebra. *Acta Numerica*, 31:347–414, 2022.

[33] HPCWire. Lets Talk Exascale: Helping Users Deploy HPC Tools and Libraries in a Consistent Interoperable Software Stack. https://www.hpcwire.com/off-the-wire/lets-talk-exascale-helping-users-deploy-hpc-tools-and-libraries-in-a-consistent-interoperable-software-stack, 2019.

[34] HPCWire. US DOE Exscale Computing Project Announced Availability of the Extreme-Scale Scientific Software Stack (E4S) v1.2. https://www.hpcwire.com/off-the-wire/us-doe-exascale-computing-project-announced-availability-of-the-extreme-scale-scientific-software-stack-e4s-v1-2, 2019.

[35] HPCWire. Extreme-Scale Scientific Software Stack Releases Version 22.02. https://www.hpcwire.com/off-the-wire/extreme-scale-scientific-software-stack-releases-version-22-02, 2022.

[36] IDEAS-ECP Project Website. https://ideas-productivity.org/ideas-ecp, 2022.

[37] M. A. Leung. Broadening engagement by extending the classroom to conferences, 2018. Society of Industrial and Applied Mathematics Education Conference, Portland, OR.

[38] M. A. Leung, S. Crivelli, and D. Brown. Sustainable Research Pathways: Building connections across communities to diversify the national laboratory workforce, 2019. Collaborative Network for Engineering and Computing Diversity (CoNECD), Washington, D.C.

[39] Mary Ann Leung. Sustainable Horizons Institute. https://shinstitute.org, 2022.

[40] Neil Lindquist, Piotr Luszczek, and Jack Dongarra. Accelerating restarted GMRES with mixed precision arithmetic. *IEEE Transactions on Parallel and Distributed Systems*, 33(4):1027–1037, 2022.

[41] LLNL Computing. Vanessa Sochat is building research software and open-source engagement. `https://computing.llnl.gov/about/people-highlights/vanessa-sochat`, 2021.

[42] Jennifer A. Loe, Christian A. Glusa, Ichitaro Yamazaki, Erik G. Boman, and Sivasankaran Rajamanickam. Experimental evaluation of multiprecision strategies for GMRES on GPUs. In *2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 469–478, 2021.

[43] Lois Curfman McInnes. How a community software ecosystem perspective helps to advance science goals in the Exascale Computing Project. `https://icl.utk.edu/jlesc12/files/jlesc12_McInnes.pdf`, 2021. 12th Joint Laboratory for Extreme Scale Computing (JLESC) Workshop.

[44] Lois Curfman McInnes. Powering HPC discoveries through scientific software ecosystems and communities. `https://www.youtube.com/watch?v=EiG9jlh49J4`, 2021. SC21 invited presentation.

[45] M. Heroux, J. Willenbring, S. Shende, C. Coti, W. Spear, L. Peyralans, J. Skutnik, and E. Keever. E4S: Extreme-scale Scientific Software Stack. White Paper, Collegeville Workshop, July 2020.

[46] Lois Curfman McInnes, Michael A. Heroux, Erik W. Draeger, Andrew Siegel, Susan Coghlan, and Katie Antypas. How community software ecosystems can unlock the potential of exascale computing. *Nature Computational Science*, 1:92–94, 2021. `https://doi.org/10.1038/s43588-021-00033-y`.

[47] AGU Fall Meeting. E4S: Extreme-scale Scientific Software Stack (Invited), IN14A-01: Accelerating Science: Moving from Simply Sharing Data to Reproducibility Through Tools Like Containers Supported by Community. `https://agu.confex.com/agu/fm21/meetingapp.cgi/Paper/992954`, 2021.

[48] Michael Heroux. Toward a Leadership Software Center (LSC): Transforming science R&D into world-class leadership software. `https://www.nitrd.gov/nitrdgroups/images/2022/MAGIC-Michael-Heroux-02022022.pdf`, 2022.

[49] National Science and Technology Council. Best Practices for Diversity and Inclusion in STEM Education and Research: A Guide by and for Federal Agencies. `https://www.whitehouse.gov/wp-content/uploads/2021/09/091621-Best-Practices-for-Diversity-Inclusion-in-STEM.pdf`, 2021.

[50] E4S at NERSC 2022. `https://www.exascaleproject.org/event/e4s_at_nersc_aug2022`, 2022.

[51] NERSC. Extreme-scale Scientific Software Stack (E4S). `https://docs.nersc.gov/applications/e4s`, 2022.

[52] NSF Office of Advanced Cyberinfrastructure, CISE. Transforming Science Through Cyberinfrastructure: NSF's Blueprint for a National Cyberinfrastructure Ecosystem for Science and Engineering in the 21st Century: Blueprint for Cyberinfrastructure Learning and Workforce Development. https://www.nsf.gov/cise/oac/vision/blueprint-2019/CI-LWD.pdf, 2021.

[53] Project Cinema/E4S. http://pantheonscience.org/e4s, 2022.

[54] Rob Farber. The Extreme-Scale Scientific Software Stack (E4S): A new resource for computational and data science research. https://www.exascaleproject.org/the-extreme-scale-scientific-software-stack-e4s-a-new-resource-for-computational-and-data-science-research, 2020.

[55] Sameer Shende. E4S and the TAU Performance System for highly-parallel application profiling on heterogeneous machines, Seminar at CINECA, Italy. https://www.hpc.cineca.it/center_news/seminar-e4s-and-tau-performance-system-highly-parallel-application-profiling, 2022.

[56] Extreme-Scale Scientific Software Stack (E4S), BoF at SC19. https://sc19.supercomputing.org/proceedings/bof/bof_pages/bof193.html, 2019.

[57] Sameer Shende. E4S and TAU Performance System on POWER Systems. https://cfp.openpower.foundation/sc22openpoweradgworkshop/talk/HP3R8E, 2022. ADG Workshop.

[58] Sameer Shende. What is E4S? E4S at NERSC 2022 Workshop. https://www.exascaleproject.org/wp-content/uploads/2022/05/E4S-2022-08-25-Shende_1.pdf, 2022.

[59] Shahzeb Siddiqui and Sameer Shende. Software deployment process at NERSC: Deploying the Extreme-scale Scientific Software Stack (E4S) using Spack at the National Energy Research Scientific Computing Center (NERSC), May 2022. https://dx.doi.org/10.2172/1868332.

[60] B. Smith, R. Bartlett, and xSDK developers. xSDK community package policies, 2016. version 0.3, December 2, 2016, https://dx.doi.org/10.6084/m9.figshare.4495136.

[61] Spack Web page. https://spack.io.

[62] Strategies for Working Remotely Web page. https://ideas-productivity.org/events/strategies-for-working-remotely-panels/.

[63] Sustainable Horizons Institute. Breaking barriers and accelerating particles. https://shinstitute.org/breaking-barriers-and-accelerating-particles, 2022.

[64] The Science of Scientific-Software Development and Use: Workshop Website. https://www.orau.gov/SSSDU2021, 2021.

[65] T. Whitney and V. Taylor. Increasing women and underrepresented minorities in computing: The landscape and what you can do. *Computer*, 51(10):24–31, 2018.

[66] xSDK Community Package Policies. https://github.com/xsdk-project/xsdk-community-policies.

[67] xSDK Web page. http://xsdk.info.

[68] Ichitaro Yamazaki, Erin Carson, and Brian Kelley. Mixed precision s-step conjugate gradient with residual replacement on GPUs. In *2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 886–896, 2022.