STEP West Coast Town Hall Report

A Sustainable Tools Ecosystem Project Event
Held in San Francisco, California
August 16-17, 2023

Recognition & Acknowledgements

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research program under Award Number DE-SC-0002844, program managers Hal Finkel and William Spotz.

Special thanks are due to West Coast Town Hall host, Anita Benn, for generously providing support and time for this event.

Finally, the writing team who created this report are appreciated and listed here.

 , STEP PI

TABLE OF CONTENTS

1. Executive Summary & Key Findings

The Sustainable Tools Ecosystem Project (STEP) brings together a diverse community of High Performance Computing (HPC) tools developers and stakeholders to develop plans for the sustainability of the HPC tools ecosystem. We define HPC tools as the collection of tools and utilities for analyzing and optimizing application performance, identifying correctness problems, and debugging. These tools interact with hardware features, compilers, communication libraries, programming model runtime systems, and operating systems capabilities that support HPC Tools, as well as the many applications that use these tools. We define the HPC tools ecosystem as the broader ecosystem encompassing the collection of stakeholders, platform dependencies, and interactions that influence those tools. This ecosystem is dependent on such diverse elements as hardware features, compilers, communication libraries, programming model runtime systems, and operating systems capabilities that support HPC Tools. Given the nature of HPC tool development, it requires extensive hardware and application interaction and understanding, and it must adapt to evolving technology. The utilization of tools can greatly enhance the performance and effectiveness of supercomputers in advancing scientific discovery.

STEP is organizing a series of three town hall meetings during the summer of 2023. The main objective of these meetings is to develop a strategic action plan for DOE/ASCR, outlining a recommended approach to ensure the long-term sustainability of an HPC tools ecosystem. Participants will collaboratively explore the current HPC tools space and develop solutions to the sustainability challenges.

This report documents the proceedings and findings for the third STEP Town Hall. Reports from the first and second town halls are available here: East Coast Town Hall Report and Midwest Town Hall Report. The third town hall was held August 16-17, 2023 at the San Francisco Mariott Fisherman's Wharf in San Francisco, CA. It brought together 41 stakeholders, 31 people on-site and 10 remotely attending, from diverse communities including tool developers, vendors, facility operators, and application developers [see Appendix 1.3 Workshop Attendees]. This town hall aimed to prioritize strategies and thrusts of the upcoming STEP Phase II proposal, and to refine certain topics related to the first two town hall findings.

Expert plenary talks and panel discussions were organized to provide a comprehensive understanding of the topics at hand. This was followed by breakout room discussions, where participants had the opportunity to delve deeper into sustainability challenges and explore actionable solutions.

The goal of this specific town hall was to assess the provisional STEP plan to identify remaining gaps, refine actionable recommendations, and prioritize initial STEP Center activities. It also served to affirm the overarching division of STEP responsibilities in both the topical dimension and management structure dimension. Breakout discussions reviewed and assessed priorities across topics that were split up into **four key challenges to HPC tools sustainability: exploding hardware complexity, exploding use cases, coordination, and management**. We also reviewed actionable recommendations across **four initiatives for STEP Center management: collaborative teams, community activities, codesign, and rapid response**. Finally, we discussed contingency to execute portions of the STEP plan under different timing or funding constraints and resolved elements that we anticipate to be key components of a phase 2 proposal.

The town hall attendees identified the following as the highest-priority high-level STEP Center responsibilities:

- **Priority STEP Center responsibility**: establish clear and productive guidelines for STEP Center management and project lifecycle management.
    - **Recommendation**: Plan for minimal management infrastructure with the ability to expand as needed according to the size of the center. The minimum viable level of management personnel should include a director, deputy director, external review board, and internal management board. The internal management board will be further divided into specific community responsibilities as needed.
    - **Recommendation**: Establish and publish clear and transparent guidance that includes core principles, program scope, expectations and assessment criteria. Provide models for now new participants can engage in the community.
    - **Recommendation**: Rely on an external panel to objectively resolve difficult issues within the community (i.e. unproductive competition between overlapping tools or changes in funding priorities).
- **Priority STEP Center responsibility**: Improve communication across the tools ecosystem to engage the full breadth of stakeholders including application users and developers, vendors, and facility operators.
    - **Recommendation**: Engage and coordinate with vendors via organized vendor briefings and a public vendor-tool alliance portal to consolidate and disseminate information.
    - **Recommendation**: Engage with application users and developers via hackathons, training events, and surveys.
    - **Recommendation**: Engage with facility operators by soliciting their feedback on priorities, collecting usage data, and providing procurement that will help enable high-quality tools support.

    - **Recommendation**: Designate specific personnel, or small groups of personnel, to act as liaisons to different stakeholder groups.
    - **Recommendation**: Where possible, leverage the resources of the nascent COLABS seed project to execute community, outreach, and educational activities.
    - **Recommendation**: Coordinate compute time allocations and early system access across DOE facilities for STEP participants.
    - **Recommendation**: Establish outreach programs to academic communities to foster workforce development. Examples include curriculum development, recruiting, engagement with under-represented groups and institutions, and development of content to increase awareness and interest in HPC optimization.
- **Priority STEP Center responsibility**: Proactively track technology trends and develop a strategic plan for addressing those trends and promoting interoperability across the tools ecosystem.
    - **Recommendation**: Work with vendors to improve documentation of hardware features and counters and act as a conduit for communication as early as possible in the product design process.
    - **Recommendation**: Coordinate the construction of a community-driven tools portability layer to share capabilities across tools and use cases. This effort should begin with simple examples to gain momentum and should learn from experience in other communities such as LLVM and Kokkos.
    - **Recommendation**: Leverage STEP's unique view across tool projects to identify opportunities to share components and increase modularity across tools. This will serve not just to strengthen existing tools, but also facilitate adopting existing tools to new use cases.
    - **Recommendation**: Improve the connection between existing HPC tool technology and the Python profiling ecosystem to help bridge the gap between the agile data analytics and scientific computing communities.
    - **Recommendation**: Improve the connection between existing HPC tool technology and the cloud and workflows community. Coordinate with the nascent SWAS seedling in this endeavor where possible.
- **Priority STEP Center responsibility**: Develop metrics and evaluation methods that ensure accountability while also maximizing community sustainability impact.
    - **Recommendation**: curate a collection of mini-apps representing a broad cross section of use cases to serve as a proving ground for sustainable tool technology. These mini-apps should be developed and maintained as a codesign activity with application stakeholders.

    - **Recommendation**: Make sure that the mini-app collection evolves to include emerging use cases. STEP should monitor emerging benchmarks from broader communities to identify trends in fields such as machine learning (ML).
    - **Recommendation**: Maximize the value of project reports by delivering them in the form of instructions for a reproducible demo of capabilities along with a concise written report.
    - **Recommendation**: Ensure that there is a low-friction mechanism for teams to adjust planned milestones in response to unanticipated needs.
- **Priority STEP Center responsibility**: improve security practices across the tools community as a whole.
    - **Recommendation**: leverage existing resources including the OWASP Software Assurance Maturity Model, NIST SP 800-218, and the Trusted CI (Cybersecurity Center of Excellence) as a foundation for improving security practice.
    - **Recommendation**: In addition to the above, distill secure development guidance into more concise actionable best practice, tailored to the tools community, to provide an accessible entry point for tools developers.
    - **Recommendation**: Leverage shared resources for security expertise; this can range from training resources to software engineers with security expertise who can be assigned to assist projects. Leverage nascent COLABS seedling resources for this purpose where available.
- **Priority STEP Center responsibility**: Coordinate with entities beyond STEP to maximize impact and make effective use of community resources.
    - **Recommendation**: Coordinate with the nascent SWAS and S4PST seedlings to identify the tool needs of the workflow and runtime communities and develop shared solutions.
    - **Recommendation**: Coordinate with the nascent PESO and SRSI seedlings to leverage potential shared management infrastructure and resources for coordination with vendor and facility teams.
    - **Recommendation**: Coordinate with the nascent COLABS seedling to deliver community activities, provide software engineering expertise across tools projects, and establish DEI and workforce development initiatives.
    - **Recommendation**: Partner with cloud and other large data service providers to amplify our influence when our needs overlap.
    - **Recommendation**: Coordinate shared interests and resources with other DOE entities, including but not limited to NNSA laboratories, SciDAC, the IRI initiative, and additional DOE programs.

- **Recommendation**: engage with open-source software (OSS) ecosystem creation efforts, like NSF POSE, to track momentum on OSS efforts and facilitate ongoing software maintenance and support.

The subsequent sections of this document provide a detailed account of the outcomes from each plenary, breakout, and group discussion. Section 3 highlights the next steps based on these findings, along with key insights on enhancing the effectiveness of the upcoming two town halls in the series.

2. West Coast Town Hall Format & Content

2.1 Introduction

Each of the STEP Town Halls is designed to focus discussion on certain *broad themes/objectives* and gather information about them in their entirety. For example, the East Coast Town Hall focused on challenges related to exploding hardware complexity, increasing diversity of use cases, and coordination among stakeholders, while the Midwest Town Hall focused on challenges related to project management and identifying existing tool capabilities and users. The West Coast Town Hall aimed to prioritize strategies and thrusts of the upcoming STEP Phase II proposal, and to refine certain topics related to the first two town hall findings.

Each Town Hall consists of a variety of sessions that seek to clarify the important challenges and urgent gaps in the HPC tools ecosystem and generate concrete actions to move this area forward in the near and long-term future.

Breakout sessions explore critical challenges and opportunities in building a sustainable tools ecosystem relevant to application software. Dedicated session leads facilitate productive conversations that yield actionable outcomes. Session leads document key discussion topics and proposed actions. In addition, the STEP team prepared a comprehensive summary of results obtained from the *tool survey* and *tool sustainability* sessions. The purpose of these town hall reports is to provide actionable tasks for immediate implementation as well as long-term strategies for achieving broader and sustained impact. We are designing the town halls to focus on a subset of the sustainability challenges, with some topics appearing in at least two of the three town halls to ensure coverage and diversity of viewpoints.

Collaborative efforts often face obstacles in engaging a core constituency, which leads to limited progress. In contrast, the STEP effort involves support from a substantial number of the HPC tool and national labs communities, demonstrating a shared recognition of the need to address these sustainability challenges. Further, many collaborative efforts lose momentum when internal priorities outweigh the benefits of collaboration. Our approach increases the collaborative gain by bringing together communities to collectively address challenges caused by their dependencies. As part of the objectives addressing the coordination challenge, we are committed to establishing mechanisms for ongoing communications among these communities even after the Town Halls have concluded.

Participation in the Town Halls is by invitation. Invitations were sent to senior HPC professionals selected from several categories, including: (a) tool developers; (b) facilities staff; (c) HPC vendors; (d) application developers; (e) individuals knowledgeable about diversity, equity and inclusion concerns.

2.2 Plenary Portion

The following presentations were helped to provide context and foster discussion:

Plenary 1

What is STEP and Why Are We Here? [45 mins] (Terry Jones, ORNL) [available online]

Plenary 2

Provisional STEP Structure [10 mins] (Phil Carns, Argonne) [available online]
Challenge 1: Exploding HW Complexity [20 mins] (Mike Jantz, UTK) [available online]
Challenge 2: Exploding Use Cases [20 mins] (James Custer, HPE) [available online]
Challenge 3: Coordination [20 mins] (Kevin Harms, Argonne) [available online]
Challenge 4: Management [20 mins] (Phil Carns, Argonne) [available online]

Plenary 3: Workforce, Security & Survey – Golden Gate Ballroom or zoom link

Establishing a Workforce [30 mins] (Dorian Arnold, Emory Univ.)
Practices Related to the Security & Integrity of Software & Data [30 mins] (Ryan Adamson, ORNL)
Survey Results and Primer for Breakout Sessions [15 mins]

2.3 Breakout Portion

This report section provides summaries for each breakout. The breakouts were grouped into 3 topic-based sessions – each session was subdivided into four facets of the session topic:

Breakouts Session 1: Prioritizing & Refining Strategies for Key Sustainability Challenges in Tools

1. Challenge 1: Exploding HW Complexity (breakout lead: Mike Jantz)
2. Challenge 2: Exploding Use Cases (breakout lead: James Custer/Phil Carns)
3. Challenge 3: Coordination (breakout lead: Kevin Harms)
4. Challenge 4: Management (breakout lead: Phil Carns

Breakouts Session 2: Key Components Needed for Phase II Proposal

1. Short-term Objectives, Priorities & Risks - What we want to do (breakout lead: John Mellor-Crummey)
2. STEP Roles & responsibilities, processes to staff them (breakout lead: David Montoya)
3. STEP Integration in whole Software Sustainability Landscape (breakout lead: Kshitij Doshi)
4. Charting Technical Directions (lead: John Linford)

Breakouts Session 3: Additional Challenges: Governance, Workforce and Organizational Challenges

1. DEI & Workforce Development (breakout lead: Mike Jantz)
2. Security and Integrity of Software and Data (breakout lead: Barton Miller)
3. Exercise: What would we want new HPC Users to see from an aspirational viewpoint? (breakout lead: Matt Legendre)
4. Bootstrapping STEP in 2024 - How to "Bridge" Tasks in a Flexible way (breakout lead: Kevin Harms)

In the summaries below we provide a synopsis of the breakouts along with their key takeaways and findings.

Session 1: Prioritizing & Refining Strategies for Key Sustainability Challenges in Tools

Breakout 1.1: Exploding HW Complexity

The objectives of this breakout were: 1) to summarize key issues, tasks, and recommendations from earlier breakouts on the exploding hardware challenge, and 2) to rank and refine these items based on their importance for sustainability as well as their likelihood of success. The ultimate goal is to distill and organize the material from the earlier town halls in preparation for a full proposal of the STEP vision.

The breakout panelists considered issues and recommendations identified in four breakout sessions from the STEP East Coast Town Hall held earlier in summer 2023. The discussion and recommendations for the material from each breakout are described next.

Breakout #1: Support Obstacles

The objective of this breakout was to describe the most significant obstacles to supporting new hardware or platforms, as well as mechanisms that could potentially address these obstacles. It identified four primary obstacles that can prevent support of new systems:

1. Developers only have hands-on access to a machine after it has already been accepted at a facility (the so-called "reactive cycle" problem)
2. There is a tension between wanting to provide general availability of a machine as quickly as possible and need for coordination between key personnel at facilities and vendors

3. Lack of coordinated forum for communication among stakeholders
4. Challenges with familiar human traits (e.g., denial, resistance to change)

Of these obstacles, the panelists agreed that, while human traits, such as resistance to change, can be a barrier to support, they are often not the primary reason for failure to support a particular technology. Additionally, while machine access is a problem in some cases, access is not really blocked in many cases, as developers often have early access to machines through NDAs and other early access programs. However, it is important to note that access and influence over the design process can hinder tool developers' willingness and ability to support a particular platform.

There was consensus in the panel that the most important factor in blocking support for new platforms is the lack of coordinated communication among stakeholders. Hence, it is important to emphasize and pursue solutions that address this lack of communication. Indeed, almost all of the solutions described in the original breakout panel, including vendor to community briefings, open communication channels, training for end users, and hackathons, aim to increase communication among stakeholders.

In addition to the solutions described by the earlier breakout, this breakout panel discussed several other mechanisms STEP should pursue to address the lack of communication among stakeholders. In particular, STEP should encourage and/or support better documentation of hardware features and counters. Better standardization, including through well-defined interfaces and APIs, would also enable vendors and tool developers to collaborate more closely without needing to consider issues related to intellectual property. Additionally, STEP should aim to establish communication channels early in the design process. Simply inviting vendors to participate in tool design meetings could reduce many of the challenges that arise from supporting new hardware and platforms.

Breakout #2: Hardware Coverage

This breakout aimed to describe challenges and possible solutions for how to achieve tools sustainability in terms of hardware coverage. The panel from the earlier town hall meeting identified two significant gaps in hardware coverage that will need to be addressed in the immediate and near term: 1) tool support for GPUs, and 2) tool support for accelerators and FPGAs throughout the system architecture and along the data path. The original panel also described some cross-cutting gaps related to both hardware and system software: 1) lack of support for power / energy management, and 2) lack of support for load balancing across heterogeneous architectures.

In this meeting's breakout, the panel discussed tool support that already exists for these technologies. Many GPUs include counters and other architectural features that enable tools to collect hardware events, but more research is needed to convert this information into actionable intelligence. Some other efforts, such as Microsoft's Project Brainwave, have already begun building tool support for FPGAs. In addition to these more urgent coverage gaps, the panelists also noted that there is a gap in tool support for quantum simulation and computing, which will need to be addressed in the medium to long term future.

The panelists in the original breakout made several recommendations for addressing these coverage gaps, which are listed below:

- Contingency funding (e.g., to help the community react more quickly to new HW)
- Standardization of counters (impractical broadly, but would be useful if targeted)
- Considering alternatives to counters for HW instrumentation
- Availability of “mini-apps for tooling”
- Early access to hardware
- Open source hardware designs
- FOA mandates that support HPC tools
- Better community coordination to ensure coverage (top down documentation, sharing profile data)

While the panelists agreed that all of these recommendations could potentially be useful, several ideas stood out as especially important for STEP. In particular, mini-apps that are able to demonstrate a gap or show that a tool can effectively identify a specific problem could be enormously helpful to many stakeholders across the community, including developers, users, vendors, and academics.

Other recommendations that were highlighted as particularly important include: standardization of counters as well as better community coordination. With regards to standardization, the panel notes that, while standardizing hardware counters across vendors and devices is impractical, standardizing capabilities can be practical and also very useful.

Breakout #3: Vendor Engagement

The vendor engagement breakout aimed to explore how vendors can contribute to tool sustainability and determine how tools can support vendors in this initiative.

The original breakout made several key recommendations to improve vendor engagement, including: 1) vendors should proactively disclose forthcoming vendor-specific software and hardware to tool developers, and 2) tool developers should engage in proactive communication with vendors prior to

releasing vendor-specific support within their toolsets. They also described several potential benefits for more proactive disclosures from vendors. However, this panel felt that it is unlikely that vendors would be more willing to proactively disclose hardware or software to tool developers because of these benefits. For instance, some identified benefits (e.g., proactive disclosure could potentially relieve the vendor of the burdens of tool development) are not realistic for many vendors.

Additionally, the original breakout recommended that STEP establish a "Vendor-Tools Alliance" workshop and website, where the workshop could serve to facilitate communication between vendors and tools, and the website could serve as a hub for collecting and consolidating issues / feedback from various tools. This panel agreed that both aspects of the Vendor-Tools Alliance could potentially help address issues related to vendor engagement. However, they note that vendors are not likely to participate in public websites that document problems with their products. They recommend that the website should be a way for people to ask questions about how to use tools and to learn about how they work on different platforms.

Breakout #4: Event Correlation

This breakout aimed to describe key use cases for correlating hardware events with program execution and/or application source code. It also identified technologies and capabilities that are needed to improve correlation between applications and hardware and ensure it can be achieved sustainably.

The original breakout described several use cases, including identifying and understanding correctness and performance problems, where event correlation is useful and necessary. The panelists agreed that STEP should aim to sustain event correlation capabilities so that these use cases will continue to be supported.

The original breakout also noted some key gaps in event correlation, which are listed below:

- New architectures often lack necessary HW and SW mechanisms for monitoring hardware events and correlating them with source code positions
- New architectures often lack documentation, such as top-down models, that inform how to use the available HW mechanisms for measurement and event correlation
- GPU architectures have under-developed correlation mechanisms
- Mechanisms to correlate network performance with network and I/O operations are also lacking

The panelists in this breakout agreed that it will be important to address these gaps in order to achieve sustainable event correlation. The first two gaps could potentially be addressed through more open and active communication channels among stakeholders. Additionally, the panelists note that there do exist some monitoring and correlation mechanisms on GPUs, but these are often difficult to use effectively. Many users desire events or metrics that have been developed on CPUs, even if these metrics are not meaningful when collected on a GPU or accelerator. The panelists agreed that more research is needed to develop better root cause models so that the community can use event correlation effectively on devices that are not CPUs.

Lastly, the original breakout made multiple recommendations to improve event correlation. Specifically, they suggest that subject matter experts should work independently from the procurement process to identify best practices and gaps in event correlation and publish these in a white paper (or platform capability standard). In this way, vendors will be able to work to address these needs before offering them in a procurement. Additionally, they recommend that the HPC community should join forces with other relevant stakeholders (including cloud providers and other large data service providers: 1) to identify common needs and communicate these needs to technology developers, and 2) to identify gaps where new methods for better event correlation are needed. This breakout's panelists agree that these suggestions are both useful and could potentially help address the gaps in event correlation. They also note that collaborating or partnering with cloud and other large data service providers will enable the HPC community to speak with a larger voice when making requests for capabilities or technologies to vendors.

Breakout 1.2: Exploding Use Cases

The purpose of the Exploding Use Case Challenge breakout was to review, refine, and prioritize the STEP project's findings and recommendations to date with respect to the Exploding Use Case challenge. This challenge was defined as follows in the initial STEP seed proposal: "Exploding use cases: New and emerging application paradigms, including AI/ML, edge, and embedded instrumentation are shifting the usages that tools need to support. Additionally, there are new opportunities for tools in traditional HPC areas, such as feedback-driven dynamic resource management." The provisional STEP plan intended to address this challenge was based on the first two STEP workshops; it is summarized in a presentation available at the STEP website here.

The attendees of this breakout discussed how to prioritize recommendations for how to address exploding use cases in a way that will maximize the impact of the STEP Center. The highest priority recommendations can be broadly characterized as either a) being proactive with respect to application trends, technology trends, and program mandates or b) promoting interoperability and

portability of tools across use cases. The latter category of recommendations was further refined to focus initially on topics that attendees believe that STEP can address quickly.

The breakout attendees recommended the following actions as the highest priority, highest impact in the category of ensuring responsiveness to application trends, technology trends, and program mandates:

- Improve the connection between Python profiling infrastructure and tools. The Python ecosystem has grown to be a widely used and integral part of data analytics and agile development in scientific computing. However, the most widely used HPC tools do not presently integrate well with Python performance instrumentation. This is possibly an opportunity for a STEP tools portability layer to have a high degree of production impact.
- Improve performance analysis of workflows and cloud computing. Most scientific campaigns, but especially those that require distributed resources or integration of observational data, are actually made up of explicit or implicit multi-step workflows, many of which may involve cloud resource components as well as conventional on-premise resources. This calls for improved instrumentation of virtual machines, improved instrumentation of high-level data movement, and improved workflow awareness.
- Performance tools must incorporate more stringent security principles. Security of HPC resources is a growing concern within the DOE and other agencies, but to date has not been a top priority for tool developers due to limited resources and expertise. Addressing this recommendation requires developer education, clear guidelines, and availability of security expertise so that tools can correctly incorporate encryption, role based access, and secure programming principles where needed.
- Hackathons and training events are crucial to addressing the exploding use case challenge. In addition to the first-order benefits to the community, it was noted that these also serve as an important vehicle for user engagement to confirm that STEP is being responsive to the community. It is also an opportunity for STEP to educate the community on how to interpret instrumentation data, as well as for the community to educate STEP on how to improve our data interpretation strategies.

The breakout attendees recommended the following actions as the highest priority, highest impact in the category of promoting interoperability and portability of tools across use cases:

- Foster a modular HPC tool ecosystem. It is clear that many of the underlying capabilities of existing HPC tools could be leveraged in additional use cases if capabilities were more modular. This would also promote sharing of functionality across tool developer communities. The attendees noted that modular, cross-tool architectures also pose a risk of being too restrictive (by generalizing functionality that was previously specialized for a specific purpose), but the potential reward for executing it well compensates for this risk.

- Create a community-driven STEP tool portability layer. A STEP tool portability layer would be a STEP sanctioned, community-driven approach to creating explicit common instrumentation interfaces across tools. The purpose is to maximize portability and interoperability and simplify the task of porting broad collections of tools to new platforms. Attendees noted that this task is more difficult in some contexts than others, and advised that STEP should begin with "low-hanging fruit" in order to foster early adoption.
- Incentivize vendors to expose and document performance interfaces. STEP has an important opportunity to amplify influence by communicating needs as an entire community rather than as individual developers or teams. It should use this opportunity to work with vendors to expose more performance interfaces and clearly document their use and semantics for the broader community.

The breakout attendees notably perceived AI methods for enabling new performance insights as a promising direction, though it does not at present reach the bar for sustainability by the definitions of the STEP project. During group discussions we did note, however, that there may be opportunities for tool developers to prepare for this possibility in the future by ensuring that performance data can be made available in clean, machine-readable formats that facilitate AI methods.

Breakout 1.3: Coordination

The Coordination Challenge breakout was charged with assessing the findings and recommendations developed during the East Coast and Midwest Town Hall sessions. The STEP seed proposal defines the coordination challenge as follows, "Tools themselves are uniquely and closely tied to design decisions across different layers of the execution stack, including: hardware, system software, middleware, and applications." The current siloed nature of HPC tool development precludes large community coordination, which in turn limits the breadth of domain expertise that can be brought to bear. The ultimate impact of this coordination challenge is that there is limited interoperability between key tools and piecemeal response to technology trends. At present there is no venue facilitating communication between the HPC tool community at large and key stakeholders such as vendors or facility operators. There is also no ongoing forum for promoting interoperability.

The breakout examined six of the recommendations generated during the previous East Coast and Midwest town halls. Of these six recommendations, one was considered high priority, two were considered moderate to high priority and three were moderate priority. The impacts of these six recommendations were split with three being high impact and the other three being moderate impact. There was discussion in the breakout on the importance of evaluating the level of effort required to accomplish each of these recommendations and this should be considered as part of the analysis when selecting recommendations to prioritize within the STEP center.

The following lists each of the six recommendations with commentary and evaluation of each one.

The STEP center should establish continuous communication channels (options include alliance organizations, recurring meetings, and public forums) with the explicit goal of gathering and formulating requirements, engaging diverse expertise, and exchanging information for the benefit of the entire ecosystem. ***This may require new strategies to address intellectual property (IP) and non-disclosure agreement (NDA) concerns.*** This recommendation has a broad scope and the breakout considered that it is prudent to separate the strategy for addressing intellectual property and non-disclosure agreement concerns as a separate recommendation. The setup of the communications channels is considered a moderate priority activity that could begin at launch which can be accomplished by different mechanisms with different levels of effort. The level of engagement from the tool developers will determine the overall impact for this. The activity for addressing IP and NDA concerns is considered a high priority with a high impact for reducing communication difficulties for tool developers within the STEP center. The breakout thought these activities would be best carried out within the Collaboration teams and Community activities.

The STEP center should actively engage vendors, facilities, applications, hardware designers, and tool developers to help establish desired requirements, guidelines, testing methods, and early system access. These collaborative efforts will benefit all parties involved, fostering a culture of proactive collaboration and continuous improvement. The recommendation carries both high priority and high impact. One of the biggest struggles for tool developers is getting access to vendor hardware and design early enough in the life cycle to provide feedback to address limitations and shortcomings. This activity within the Codesign initiative would be mutually beneficial to both tool developers and hardware vendors alike and should be started as soon as possible. The effort to create the initial engagement is a minimal technical hurdle, instead we need to overcome the social engineering for these groups to meet.

Produce and broadly adapt a tool portability layer (or layers) that enables tool developers to share common functionality, drawing inspiration from successful examples from other communities such as LLVM and Kokkos. Additionally, establish community guidelines to promote interoperability among stakeholders. This recommendation is considered to have a high impact which would reduce effort and redundancy for long term tool development. However, this was rated as a medium priority due to the uncertainty of success and the high effort required to accomplish this. For the portability layer to be a success, establishing buy-in from a wide portion of the tool development community is necessary. Which initiative this work fits best in is unclear, but the breakout determined the Codesign initiative is likely best as it provides a mechanism to engage the vendor population.

Project reports should take the form of instructions for a reproducible demo (including quick start) of the work along with a concise report. An artifact evaluation team (perhaps funded in part by grants or stipends will evaluate the demos. This Midwest Town Hall recommendation was relevant to coordination challenges as it provides a mechanism to increase the accessibility of tools for the wider application developer community providing a significant impact to increase the appropriate usage of tools. However, to accomplish this work has a high level of effort due to the inability to fully automate the generation of materials and evaluate their correctness. The Collaboration initiative would be the appropriate place to enact this activity though with the main beneficiaries being the Applications and Facilities.

The final two recommendations were both rated as medium priority with medium impact. There was a concern that the *Establish a collaborative platform that is independent of any one facility to foster the exchange of best practices, mini-apps, case studies, and profiling data among various deployment sites*, would have difficulty remaining relevant over time as hardware and software changed with an unclear payoff to the tool developers. The *Develop methods to effectively document and streamline tool dependencies, accounting for subtle interactions such as required build options for underlying libraries, implicit dependencies, and performance dependencies. Utilize existing tools like Spack and E4S, where appropriate recommendation* was considered reasonably low effort and cited some existing work done during ECP could be leveraged to initiate some of this work. A dependence on Spack was noted as a potential concern due to Spack being outside of the STEP community. The STEP center would need assurance that there would be continued development resources dedicated to Spack as well as some commitment from the Spack development team to engage with the STEP center to prioritize some of the needed improvements for tools build and deployment.

The breakout also reviewed the remaining recommendations from the East Coast Town Hall which were not prioritized and agreed that none of these recommendations were high priority. A few gaps were discussed as well including the lack of static analysis and code transformation tools within STEP. No recommendation related to this is forthcoming but perhaps STEP can solicit these type of tools in the future as its scope broadens. Finally, two new recommendations were developed during the breakout.

The STEP Center should work with ASCR DOE User Facilities management to ensure that compute allocations are provided for the tool developers within STEP. Currently, tool developers must compete with science proposals for allocation time within INCITE or ALCC calls with no guarantee of success. Otherwise, tool developers are left with using smaller discretionary allocations

which require the developer to pick and choose what to test and limits any ability to test at large scale.

The STEP Center should develop a set of requirements and specifications that can be utilized by facilities during the system acquisition process to inform vendors of the needs for the tool community. This allows the STEP community to have an impact on future ASCR user facilities without needing to resolve all constraints with a multi-institution association getting access to the procurement process.

Breakout 1.4: Management

The purpose of the Management Challenge breakout was to review, refine, and prioritize the STEP project's findings and recommendations to date with respect to Management. The management challenge was defined as follows in the original STEP seed proposal: "The management challenge: Building a sustainable tools ecosystem will require plans for organizing, operating budgets, community standards, technology tracking, workforce development with particular attention to promoting inclusive and equitable research (PIER)." We continued to use this as the basis for framing the topic. Management is a cross-cutting challenge; as such, many of the issues and findings are closely tied to the other three challenges defined by STEP. As earlier referenced for the Exploding Use Cases challenge, STEP seeks to address these issues with a provisional plan based on the first two STEP workshops was summarized in a presentation available at the STEP website here.

The attendees of this breakout had considerable previous experience with the ECP program. The general consensus is that ECP participants appreciated the stability and predictability of the program, and hope that we can replicate those aspects in the STEP Center. However, we also aspire to do so with lower management overhead so that we can maximize productivity and flexibility for STEP participants.

In addition to the above observation, the attendees identified three management risks that the STEP center should be mindful of:

- Risk 1: Overplanning management mechanisms: Because the funding parameters for STEP are not known at this time, there is a possibility that we could inadvertently mandate a level of management infrastructure that is disproportionately large for the amount of technical work performed. The mitigation for this is to plan for a lean management structure with the expectation that it will be scaled out quickly as needed.
- Risk 2: Locking out new participants in STEP: Conventional review, reporting, and continuation policies for projects are necessary for oversite, but if care is not taken they can create a closed loop where it is easy for existing participants to continue but difficult for new participants to gain entry. This was identified as a shortcoming of the ECP program for some

potential participants. The mitigation for this risk is to clearly communicate expectations for how important new tools in need of sustainability support can gain entry to the program.

- Risk 3: Unhealthy competition between overlapping tools: The STEP Center will bring together a broad collection of tools under one funding vehicle, and some of these tools provide overlapping capabilities. This can be a positive situation that fosters friendly competition and risk mitigation. There is also a risk that it can duplicate effort or foster a negative culture among participants. The mitigation for this to ensure that resolution of competitive issues between teams comes from an authoritative, credible, objective party (i.e., an external advisory board) with as much transparency as possible.

Overall, the breakout produced four primary recommendations to refine the existing plan for addressing the management challenge:

- The STEP Center, regardless of funding level, should prioritize establishing and publishing initial transparent documentation and governance for STEP. This includes sustainability criteria, stakeholder needs, governance methods, principles, scope, and expectations. These documents are equally important across any range of funding scope, and are crucial to avoid misinterpretation of expectations for potential participants. Even the word "sustainability" can be interpreted in a number of different ways by different stakeholders and could lead to confusion as to how to best participate in STEP if we are not perfectly clear.
- The STEP Center should proceed with the "milestones are demos" concept for tracking the progress of funded projects. This concept was first proposed at the STEP Midwest Town Hall and has continued to resonate among town hall attendees. In this breakout we noted that this method is a good mechanism to ensure that STEP activities are well-integrated with real-world applications and responsive to application needs. They also serve to document best practice for tool usage.
- STEP should provide a low-friction mechanism for altering milestones over time. This was based on reflection and assessment of ECP policies. The consensus of the breakout participants was that planning milestones (i.e., quarterly demos in STEP) well in advance is a good exercise for ensuring that funded projects have a clear vision, but the nature of HPC tools and their dependencies means that plans are subject to change as unexpected problems arise. We must be respectful of this and allow projects to be agile with their milestone and demo plans.

One issue that remains unresolved is the exact mechanism for distributing funding from ASCR to project participants in STEP. This requires additional negotiation with program management and laboratories to clarify what is possible. We anticipate that the STEP project will propose a mechanism and be prepared for flexibility as needed.

Finally we observed a clear opportunity for collaboration across seedlings: many of the outreach, education, and workforce development activities that we have identified as crucial for the Community Activities initiative within STEP align well with services that could be provided by the COLABS seed project. We recommend that the STEP Center coordinate with the COLABS effort, particularly with respect to management of community outreach activities, to minimize duplicating management effort. However, we will also be prepared to execute these activities autonomously as a risk-mitigation strategy if these services are not available in the form that we need.

Session 2: Key Components Needed for Phase II Proposal

Breakout 2.1: Short-term Objectives, Priorities & Risks

Breakout 2.1 discussed risks and impacts associated with tools. The breakout considered impacts to the tools landscape under the scenario that STEP ***is*** funded, and the scenario that STEP ***is not*** funded.

For the scenario where STEP is not funded, the breakout group listed risks associated with the ecosystem (open source tools will degrade as vendors change interfaces, older interfaces will be deprecated and then discontinued), the integration across stakeholders, the higher reliance upon vendor tools by stakeholders, and the implications of a high vendor tool reliance (likely a higher focus on node level performance). The impact for these risks apply to application scientists (who will need to develop ad-hoc tools to meet their needs, and who may need to learn to learn more platform dependent tools), facilities (who also may need to write their own tools: partial duplication of effort, and who will use tools less efficiently), and tool developers (who will struggle to fund tool development with patchwork of grants, and who will be unable to support all HW and use cases).

For the scenario where STEP is funded, the breakout group listed risks associated with shutting out participation of people outside the center, STEP management making poor decisions (Inappropriate allocation of resources, stovepiped solutions, premature standardization on fewer solutions, pick the wrong platforms as targets), underestimating the effort to re-architect the software ecosystem, and slow to adapt to changing needs of stakeholders.

The breakout group was able to suggest a number of measures that STEP could employ to mitigate these risks including carefully chosen leadership committee; outside board of stakeholders & annual reviews; coordination to avoid duplication of effort; coalesce around best of breed solutions for foundational capabilities (for example, stack unwinding, function interposition); rapid response efforts (extend a tool with a new capability that addresses an urgent need); an open proposal process for new efforts (aided by understood criteria for what will be considered, and possibly by allowing research prototype demonstrations); and sunsetting projects that are no longer important.

Finally, the breakout discussed the largest non-monetary impediments facing the tools ecosystem. These were listed as stable staffing (retaining well-qualified staff and attracting new staff with the right skills), inertia of the status quo (Insufficient commitment to software integration among competing collaborators), Lack of sufficient vendor engagement (reverse engineering is no substitute for vendor documentation of invariants), and lack of standardization.

Breakout 2.2: STEP Roles & responsibilities, processes to staff them

The STEP center will require various roles to support the tools ecosystem and interface with collaborative programs and interface with target users, facilities and vendors.

The core organization should be lean but allow for fluid creation of specific domain boards and forums. The coordination layer should include executive and deputy directors (part-time), an external review board to assess STEP progress and be the decision body for acceptance of projects and funding distribution, an internal board made up of tool community members, and various liaison and interface coordination roles, and administrative support.

The director role will also serve as funding coordinators based on agreed on project distribution from the project acceptance process and review. It is expected that funding distribution across the labs and contractors will have challenges. They will also oversee the outcomes and be involved in the decision-making process regarding the boards and forums formed within STEP.

The external review board is made up of individuals external to the tools projects that are within STEP and should include DoE executives, and other domain experts. It should include 4 or 5 individuals (? On how long and staggered commitment). The board is involved in initial project acceptance and funding distribution, quarterly (or bi-annual) STEP center review, and new project assessment and acceptance.

The internal board is similar in structure to the external board and is made up of tool community members. Its role is to support the executive directors in STEP center planning, board and forum development, coordination support and decision making. They will perform project reviews/demos, introduce new projects for review, support assessment and creation of rapid response projects, and coordinate and staff sub-committees focused on applications, facilities, and vendor interface.

Communication and collaboration boards and forums. These are groups that can form organically or be developed for specific needs. There are unique tool needs and collaborations that can be developed as well as community wide or vendor focused topics that need to be addressed. These could include efforts such as a focus on standards committees or vendor NDA concerns and interaction.

STEP center members should include a broad spectrum of the tools community needed to support HPC. This includes projects that are funded, projects that are identified as key contributors but not funded, vendors, application teams, and other stakeholders. Engagement and communication should be seen as a key to continued success and growth.

Breakout 2.3: STEP Integration in whole Software Sustainability Landscape

Software sustainability has many themes that cut across diverse HPC community endeavors worldwide. In this light, several notable opportunities for coupling STEP goals and initiatives with those of other software sustainability projects, and with other communities in DoE and beyond DoE are identified at an outline level next.

The mission of sustainable tools cannot be met in its own silo – it requires high energy collaboration with other software efforts and domain experts, and needs to be responsive to continuous assessment and feedback from stakeholders, facilities operations, and vendors. STEP must rapidly identify dependencies that crisscross between tools, software stacks, operation and management planes, security frameworks, and so on, and prioritize them for resourcing and execution. Specifically, the following are major points of collaboration and coordination:

- With S4PST: STEP can collaborate on compiler and runtime tools, defining common tooling APIs and standards, scoping static analysis and debugging tools.
- With SWAS: STEP can collaborate on workflows, structuring network I/O caching as a unified API behind which different usages and hardware assemblies are supported with different parameter settings, leveraging IRI and using it for better cross-coordination among people involved in different parts of workflows, and recommendations for separately funded areas of research on addressing tools gaps and developing performance insights across end to end workflows.
- With PESO: STEP can identify areas of overlap with PESO's tools spoke and how to best identify opportunities for division of labor and collaboration, and on bringing in a CI/CD focus, especially on performing CI/CD automation for performance debugging and regression analysis.
- With OSSF: Many critical observation and control tools may need to be vendor proprietary in some respects. STEP can work with OSSF to create a foundation that vendors can join to work with STEP for customizing the tools' functions and interfaces responsively to community needs.
- With COLABS: STEP can capitalize on the professional SW engineering focus of COLABS, securing from COLABS any specialized expertise needed in various SW engineering tasks

while transmitting guidance, documentation, and training on tools via COLABS to other stakeholders and user communities. COLABS can also help STEP with hardening of tools.

- *STEP would collaborate closely with other DoE communities including but not limited to:* Facilities, Laboratories, IRI, and NNSA labs.

Tools that are not designed with some important applications in mind risk falling into disuse and abandonment, while application impact provides a clear metric of success and an impetus for rapid improvements in tools. With this in mind, STEP will focus on enabling early interactions with application teams. With key application-oriented tools developed in close concert with major applications teams, STEP will seek to circumvent a common problem with tools feedback: that it is very difficult to elicit constructive feedback on a tool when application developers are either not aware of the tool or are only using it infrequently. The difficulty is compounded by information and expertise silos, and challenges for developers in keeping up with what, if any, is the right tool for each given circumstance. STEP will create hackathons and work with other sustainability projects to organize outreach and stand up a tools consultation body where rapid help/guidance/bridging-to-experts can be provided. COLABS and IRI can facilitate the structuring, resourcing, and coordinating these activities.

Breakout 2.4: Charting Technical Directions

The Charting Technical Directions breakout discussion focused on identifying (1) processes to track technology trends, DOE mission needs, and user communities; (2) interfaces and interactions with other communities (e.g., system software) to foster sustainability and innovation; and (3) methods for conflict resolution for disagreements on technical directions.

The following themes were discussed. Technology trends need to be tracked. For example, a lot of tools missed the emergence of AI/ML applications and there is a general lack of tools that drive actionable information for AI/ML frameworks. There is also a severe lack of tooling for interpreted languages. What the community does today (trend location) needs to be tracked as well, such as using (a) center support tickets and user requests, (b) system-wide lightweight monitoring, and (c) reproducibility artifacts. Where the community is going (trend velocity/trajectory) also needs to be tracked, such as with (a) center request tickets (e.g., software install requests), (b) generating an understanding of emerging benchmark usage (e.g., MLPerf as a bellwether for AI/Python), and (c) observing user communities that disrupt the tools landscape. The NSF Pathways to Enable Open-Source Ecosystems (POSE) program may serve as an example for community interaction as it helps to form new managing organizations for existing open-source products.

The following recommendations were made during the breakout discussion in the areas of reporting, benchmarks, and community:

- ***Reporting:*** When feasible, HPC centers participating in STEP should share software usage statistics to tool developers to help track and intercept trends. Provide regular summary reports of user requests, including frequent requests for popular packages. Highlight unusual/bleeding edge user requests. STEP participants should deploy system-wide monitoring, e.g. XALT, and track job-level usage of libraries and executables.
- ***Benchmarks:*** STEP should track bleeding edge benchmarks as potential bellwethers, including benchmarks used or debated by startups. Such benchmarks should be used to test tools and evaluate their readiness for emerging technology, e.g. quantum simulation and quantum compute. Commercial entities should be leveraged, as they are highly motivated to optimize these benchmarks. Vendors will eagerly engage a tool demonstrating value-add to a key bleeding edge benchmark.
- ***Community:*** STEP should engage with Open Source Software (OSS) ecosystem creation efforts, like NSF POSE, to track momentum on individual OSS efforts and facilitate ongoing software maintenance and support. Reproducibility initiatives should be engaged as they can show trends by community. Reproducibility artifacts show exactly what software is being used by conference and even by conference track.

Session 3: Additional Challenges: Governance, Workforce and Organizational Challenges

Breakout 3.1: DEI & Workforce Development

The objectives of this breakout were to summarize and refine DEI proposals collected during the previous town hall as well as the Establishing a Workforce plenary session to create a plan for promoting DEI and workforce development through the STEP project.

The breakout panel included Anshu Dubey, a Senior Staff Scientist at Argonne National Lab, as well as PI of the COLABS seedling project. COLABS already has a comprehensive and detailed DEI & Workforce Development plan, which is directly in line with the goals of the STEP project. The COLABS vision includes engagement and recruitment from HBCUs and MSIs, an internship and training program, as well as partnerships with external organizations, such as the Sustainable Horizons Institute, which have a proven track record of developing programs to promote and increase diversity and inclusion in scientific organizations. The panel agreed that STEP, COLABS, and the other seedlings, should work together to build and implement DEI and workforce development plans to reach their common goal of promoting equitable and inclusive research through the software sustainability programs.

The panel discussed several other specific ideas that the STEP program, in collaboration with the other seedlings, could leverage to achieve PIER objectives, as described next.

First, students in many large universities, including schools with well established Computer Science and engineering programs, are not aware of DOE national laboratories as career opportunities. STEP should increase engagement with these schools, including through faculty connections, career fairs, and other university programs, to engage and recruit a new pool of students and young professionals.

Next, STEP should develop new programs and content to increase interest in scientific software and tools among young people. For example, some projects have increased interest and engagement in robotics through guided "boot camp"-style programs. In the tools context, such a boot camp could potentially challenge students to learn and use software tools to solve a puzzle, complete a game, or participate in some other competitive design activities. Additionally, STEP should create educational materials, including video and/or short course modules, that can be distributed quickly and widely to help teach novices about scientific tools and their capabilities. All of these developed programs and materials should aim to highlight the sense of wonder and purpose that comes with a career in science.

Finally, STEP should try to recruit students and populations that are not typically engaged with DOE national laboratories, including community college and night school students, adults considering a career change, and younger people, including high school students. STEP should consider modeling these efforts on other successful DEI and workforce development programs, including those that have been deployed by the US Army. For example, STEP could establish a program that pays for a student's education in exchange for an agreement that requires the student to work on a particular project or in a particular lab for some number of years after their graduation.

Breakout 3.2: Security and Integrity of Software and Data

Software and data security are an area of concern for HPC (“the S in HPC stands for security”), as unfortunately, attention to security is often eclipsed by performance, cost, or efficiency considerations. These security concerns come in two forms. First, are we designing and developing code that conforms to best (or even reasonable) practices that avoid flaws that can introduce vulnerabilities and security threats. Certainly, any team that develops code that is widely shared or is deployed as a service has a great need for a software security plan and a maturity model on which to base that plan

Second, are the evolving hardware, operating system, runtime, and networking security features meant to increase the security of applications and services. These security measures and policies are becoming a hurdle to collecting, disseminating, and correlating performance data from several important sources (e.g., uncore events and counters from modern processors). The presence of such measures require the vendors to enable secure and filtered access at fine granularity to such data. In contrast, HPC tools are vulnerable to security flaws in vendor resources such as GPU drivers.

Software Security Plans and Reference Models

Several security resources and best practices can be applied with current HPC tools to evolve the tools securely and to ensure their continued integrity. The goal is to create an HPC-relevant software security maturity model and to have a guide to developing a secure software development plan. STEP can be the focal point for such efforts.

A maturity model outlines levels of conformance to security concerns. It provides a staged approach to developing a fully effective ("mature") security program. A good starting place is the OWASP Software Assurance Maturity Model[1]. While SAMM is an excellent starting point, it has many many different aspects and choices to be made. A focused maturity model for HPC projects could greatly accelerate their adoption.

A maturity model is descriptive, i.e., it tells you what you want to accomplish. A companion to a maturity model is a reference guide, which is prescriptive, i.e., it tells you how to reach the different levels in the maturity model. As a HPC software security maturity model is developed, a companion guide should also be developed.

As part of the model and guide development, several resources should be leveraged:

- NIST SP 800-218 for overall requirements.
- Threat modeling methodologies and Secure Software Design Lifecycle Methodologies
- Static analysis and dependency tools (free and open source)
- Free training resources from the Trusted CI (Cybersecurity Center of Excellence)[2]. These resources can be supplemented by training resources from CMU SEI and other, albeit expensive, commercial training resources and testing resources.
- The Trusted CI Guide to Securing Scientific Software, Security Framework,

STEP can take a lead in providing an HPC software maturity model and guidance on using the model and the level of standardization that should be targeted for sustainability. STEP will need staffing to evaluate, advise, and train tool teams on security, and can leverage expertise and staffing at the computing facilities.

The rich detail in resources like NIST SP 800-218 ironically makes it an uphill task for the HPC domain experts to comprehend and apply rigorously to their coding and evaluation practices. One important role for STEP would be to create prescriptive guidance (short and to-the-point) that the developers can incorporate seamlessly into their coding, reviewing, and testing processes. STEP can also be a catalyst in the process by creating online training resources and in-person and online

[1] https://owasp.org/www-project-samm/

[2] https://research.cs.wisc.edu/mist/SoftwareSecurityCourse/

workshops for accelerating the adoption of best security processes throughout the software and systems development lifecycle.

System Software and Hardware Privacy Issues

STEP also needs to address the key question of how to provide sufficiently detailed performance and execution data from an HPC application without exposing facilities and users to security and privacy risks.

On the user side, the key issue is whether runtime tools (such as for profiling and debugging) have access to hardware and software resources to provide information that is sufficiently detailed to understand the behavior of the program. Providing complete access to hardware performance counters and operating system performance statistics are important examples of this type of information. STEP can be the advocate and advisor to the vendors and system software producers to help insure that such information is available

On the facilities side, the key concern is privacy and safety: does the data expose information about proprietary aspects of the system software or other users. There has been a long history of drivers for accelerators containing flaws that could allow one user's data to be exposed to another user. STEP provides a resource to work with the vendors to help detect such security flaws and advise on the design of solutions to these flaws to ensure that the fixes are not so conservative that they block useful functionality.

Breakout 3.3: Exercise: What would new HPC Users want from an aspirational viewpoint?

The STEP center would like to bring a new focus to interaction and value of tools to new and existing users. To accomplish this, we need to understand the various user communities, their workload, their knowledge of the environment, and what would benefit their effective use of the HPC environment. This does not provide a one-size-fits-all solution, but starts with engaging users with high level information to educate them on resource use and capability, how their jobs are interacting with the HPC environment, and possibly provide guidance on better effective use. We identify three categories of users:

1. Application authors are code focused on code development and benefit from Continuous Integration (CI) performance analysis. They track performance during development and are concerned with the architecture and networking capabilities of the environments for which they target their application. Simple workflows enable these users to run tools.
2. Simulation scientists set up simulations and are the main users of compute cycles). There may be a wide span of knowledge within this group regarding how to configure jobs for effective

use of the system environment: they are more likely to mis-configure rather than mistune an application. A monitoring-focused performance analysis tool provides these users with better insight into their resource use and effective performance. A characteristic of these tools is that they are always on, accessible to users and sysadmins.

3. Data analysis (ML/AI/data-centric) users are a growing user base. They are more familiar with tools related to data analysis such as Jupyter notebooks and analytics libraries (such as Pandas). As part of their workflow, they may need to perform large data analysis and run into similar configuration issues that the simulation scientists noted above have.

Regardless of the user category, there are tooling and insight directions that can be taken. Some thoughts include tool integration with familiar IDEs, such as VSCode. Some possible integrations include annotating code views with performance data in these IDEs; provide users with guidance that tells them what tool they need; and provide guidance as to what resources are available to help the programmer deal with their performance problem? Additionally, we would like to see that users have access to training and documentation that would help them to understand what the tools are telling them. Some of these mechanisms already exist but are not integrated in a way that supports a given user base; how to take it to the next level is part of the effort.

From a STEP perspective, we need quick ways to short circuit performance analysis ("are they even doing much IO?") with intuitive feedback. In the real world this must include, easy to run: "Module load <tool>", run job normally, look at dashboard. We must have portability across many architectures, but also have the ability to go deep on each architecture. We need more memory tools (example - memaxis using hwloc), to help users understand memory. There needs to be tools that understand whole-system context for IO, network, etc.. This is where we look toward integration with system monitoring. Monitoring infrastructures also need integration from many sources: darshan, network, etc... We know that this is not an easy integration as most monitoring infrastructure are very facility based and it is difficult to have a general interface.

There are many opportunities in this space, but the viewpoint has to be holistic. Users are not the same, HPC environments are not the same, but the approach should be similar. The STEP center can help coordinate toward this vision.

Breakout 3.4: Bootstrapping STEP in 2024 - How to "Bridge" Tasks in a Flexible way

The Bootstrapping STEP in 2024 - How to "Bridge" Tasks in a Flexible Way breakout was dedicated to looking at the scenario where the STEP Center is funded but is required to distribute funding prior to setting up the complete leadership and governance structures including the complete set of policies and procedures. The breakout considered methods to accomplish a faster evaluation of

potential projects and provide the needed funding according to the requirements set forth by the DOE directions. In order to do this, the breakout considered the following charge questions.

If the STEP Center is tasked with distributing funding before the full governance structure can be established:

- *What is the most equitable way to sustain existing activities?*
- *What logistical infrastructure must be in place, and at what institutions?*
- *What time horizon is appropriate for this bridge period, and what conditions must be met before transitioning to the proposed STEP operational model?*

The breakout developed three high level methods to address an equitable distribution for sustaining the existing activities. These methods all share a set of common requirements that will be needed to evaluate and execute the decisions as part of sustaining an activity.

First a list of potential target projects will be required. This list will come from the interested parties involved in STEP, other Seedlings and existing ECP funded projects. Given a tight deadline, it will not be possible to execute a full open call for projects. Once a list of potential projects is created, a set of guidelines to evaluate these projects is needed. The STEP leadership will be responsible for creating these initial guidelines based on the overall goals of the STEP Center and any guidance provided by DOE. It is critical the selection criteria and policy should not go against any of the planned STEP center goals. A few example of what might be contained in these guidened are:

- What project is considered a Tool?
- What is the scope of work defined for the bridge period?
- What should be milestones or reportable products/outcomes that fit within the period?
    - Projects should be able to demonstrate some outcome during the bridge period.
- What is equitable? There are many ways to define equitable within the STEP context.

The breakout discussed various ways to select candidate projects in an equitable manner. There was no clear and obvious choice for this, but here we list the ideas that were discussed in the breakout.

- Projects that were involved in the STEP Center proposal and processes (such as attending Town Halls)
- Projects that were funded as part of ECP
- Projects that are most in-need of funding that will immediately loose staff
- Projects ability to achieve a sustainment activity within the defined bridge period
- Projects which form some interconnected collection which can work together
- Based on the needs from scientists, applications and facilities
- Distribution between well known and well established projects and newer less mature projects
- Based on an objective ranking system (most users, most downloads, …)

The selection process will also need to consider the ability to distribute funds to the project. Any project should have a Principal Investigator with their associated existing financial infrastructure in place and able to receive funding from the STEP Center having all of the existing DOE and government paperwork and processes already in place, as these can take months to establish.

Projects awarded funding during this bridge period should be provided the funding for at least one year, with six months being the minimum time. A year will provide adequate time for the STEP Center to put into place all of the governance structures, policies and procedures for running STEP as well as selection and onboarding of projects. This year will allow the initial bridge projects reasonable time to execute on sustainment activities and produce reportable outcomes that can be assessed within STEP. Once the year period has been reached, each bridge project should create a formal proposal to submit to STEP and be subject to the STEP project selection process. This year of funding also provides a "workforce" for STEP, the members will be required to serve on review panels for new proposals and artifacts produced.

Finally, the breakout developed the following three high level methods for selection of projects in an equitable manner. The methods are ordered by what was perceived as the most equitable method to the lesser equitable method. The order is indicative of the level of effort and time to achieve, so the most equitable method will require the most effort and time while the last method requires the least effort and least amount of time.

Selection by External Review

A panel of reviewers will be selected from tool developers inside and outside of the STEP center, leadership from other Seedling projects and specifically precluding the STEP leadership. The panel will review the projects list provided by the STEP leadership and select the best candidates based on the guidelines provided by the STEP leadership including the analysis of the selections.

Selection by Review of Applications Developers and Facility Experts

A panel of reviewers composed of scientific application developers and facility experts. The panel will review the projects list provided by the STEP leadership and select the best candidates based on the guidelines provided by the STEP leadership including the analysis of the selections.

Selection by STEP Leadership

The STEP leadership will select the best candidates from the list of projects and use the guidelines for selection that were developed by the STEP leadership. The selections report will include an analysis of the selections.

For all the methods, once the selection process is complete the STEP leadership will determine the exact award amount based on the request and the results of the survey results presented in section Plenary 3.

3. Next Steps

3.1 Distribute Findings of Town Hall

STEP uses a grass roots approach based on Town Halls and community input to develop a plan for sustaining a tools ecosystem for HPC over the long term. Each Town Hall meeting is designed around invitations to experts representing the HPC tools developers, vendors, HPC facilities and application teams communities. These disparate communities have not typically or regularly interacted as a group, but have significant inter-dependencies. The West Coast Town Hall participants collaboratively explored the current HPC tools space and developed solutions to the sustainability challenges (see Appendix 2: Agenda). Capturing the discussion, recommendations and key findings of the West Coast Town Hall held August 16-17th, 2023 at the San Francisco Marriott Fisherman's Wharf, is a priority outcome from the town hall. When complete, the availability of the West Coast Town Hall Report will be announced via applicable mailing-lists and placed on the STEP website under resources.

3.2 Begin STEP Center Proposal planning

Prior to calendar year 2024, we anticipate an opportunity to submit a proposal to DOE/ASCR for how STEP would sustain the tools ecosystem. We will develop this plan with the full input from our 3 town hall events:

- The East Coast was structured around the theme of 3 challenges facing the tools ecosystem: exploding hardware complexity, exploding use cases, and coordination. It featured plenary talks by Terry Jones (STEP Overview), Hal Finkel (View from DOE/HQ), John Mellor-Crummey (the Exploding Hardware Complexity problem), a panel of experts who discussed the Use Case Problem (Amadeo Perazzo of SLAC, Kerstin Kleese Van Dam of BNL, and Sam Reeve of ORNL), and a panel of experts who discussed the Coordination Challenge (Matthew Legendre of Lawrence Livermore National Laboratory, JaeHyuk Kwack of Argonne National Laboratory, and Dave Montoya of Trenza). It featured 44 stakeholders from diverse communities including tool developers, vendors, facility operators, and application developers.

- The Midwest Town Hall focused on various aspects of the Management Challenge. It featured plenary talks from Terry Jones (STEP Overview); Phil Carns (STEP Initiatives); Mike Jantz (Perspectives on DEI and Workforce Development); and Theresa Windus (Interfacing with Application Teams). It brought together 37 stakeholders, 20 people on-site and 17 remotely attending, from diverse communities including tool developers, vendors, facility operators, and application developers
- The West Coast Town Hall aimed to prioritize strategies and thrusts of the upcoming STEP Phase II proposal, and to refine certain topics related to the first two town hall findings. It featured plenary talks from Terry Jones (STEP Overview); Bill Spotz (View from DOE/HQ); Phil Carns (STEP Initiatives); Dorian and Theresa Windus (Interfacing with Application Teams). It brought together 41 stakeholders, 31 people on-site and 10 remotely attending, from diverse communities including tool developers, vendors, facility operators, and application developers.

Using the wealth of information from these three town halls, and the collective experience of our leadership team of 23 HPC professionals, we are working diligently to assess and address the tools community concerns. Our plans, which will feature extensive coordination and collaboration with other DOE Software Sustainability efforts, will be documented in detail in our forthcoming proposal.

4. References

1. Sustainable Tools Ecosystem Project (STEP): Available Online: https://ascr-step.org/wp-content/uploads/2023/07/STEP-Proposal-Abbreviated.pdf
2. David Boehme, Todd Gamblin, David Beckingsale, Peer-Timo Bremer, Alfredo Gimenez, Matthew LeGendre, Olga Pearce, and Martin Schulz. 2016. Caliper: performance introspection for HPC software stacks. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '16). IEEE Press, Article 47, 1–11.
3. DWARF Version 5. https://dwarfstd.org/dwarf5std.html.
4. GNU Binutils. https://www.gnu.org/software/binutils.
5. Brinkley Sprunt. Pentium 4 performance-monitoring features. IEEE Micro, 22(4):72-82, 2002.
6. M. Srinivas, B. Sinharoy, R. J. Eickemeyer, R. Raghavan, S. Kunkel, T. Chen, W. Maron, D. Flemming, A. Blanchard, P. Seshadri, J. W. Kellington, A. Mericas, A. E. Petruski, V. R. Indukuru, and S. Reyes. IBM POWER7 performance modeling, verification, and evaluation. IBM Journal of Research and Development, 55(3):4:1-4:19, May 2011.
7. A. Yasim. A top-down method for performance analysis and counters architecture. In IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), pages 35-44, March 2014.
8. A. Nowak, D. Levinthal and W. Zwaenepoel, "Hierarchical cycle accounting: a new method for application performance tuning," *2015 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Philadelphia, PA, USA, 2015, pp. 112-123, doi: 10.1109/ISPASS.2015.7095790.
9. Keith Underwood - Cassini: The Slingshot NIC Hardware and Software Enabling Exascale. Eighth International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale (Exacomm2023), May 25, 2023. https://youtu.be/yxPrDBSgK1Y.
10. Samuel Williams, Andrew Waterman, and David Patterson. 2009. Roofline: an insightful visual performance model for multicore architectures. Commun. ACM 52, 4 (April 2009), 65–76. https://doi.org/10.1145/1498765.1498785
11. NVIDIA Corporation. CUDA Profiler - Release 12.1. April 14, 2023. https://docs.nvidia.com/cuda/pdf/CUDA_Profiler_Users_Guide.pdf
12. IBM Corporation. Power ISA Version 2.07B. April 9, 2015.
13. Paul J. Drongowski. Instruction-Based Sampling: A New Performance Analysis Technique for AMD Family 10h Processors. November 16, 2007. https://www.amd.com/system/files/TechDocs/AMD_IBS_paper_EN.pdf
14. Intel 64 and IA-32 Architectures Software Developer's Manual. Document Number: 252046-043. June 2014. http://kib.kiev.ua/x86docs/Intel/SDMs/252046-043.pdf.
15. Barton Miller: Add referenced for tool compatibility - Barton P. Miller, Ana Cortés, Miquel A. Senar, and Miron Livny, "The Tool Daemon Protocol (TDP)", SC2003, Phoenix, Arizona, November 2003.

Appendix 1: Attendees

Appendix 1.1 Workshop Organizers

First Name	Last Name	Affiliation
Jim	Brandt	Sandia National Laboratories
Phil	Carns	Argonne National Laboratoy
James	Custer	Hewlett Packard Enterprise
Kshitij	Doshi	Intel
Ann	Gentile	Sandia National Laboratories
Richard	Gerber	Lawrence Berkeley National Laboratory
Tim	Haines	University of Wisconsin
Heike	Jagode	University of Tennessee
Mike	Jantz	University of Tennessee
Terry	Jones	Oak Ridge National Laboratory
Matt	Legendre	Los Alamos National Laboratory
John	Linford	NVIDIA
Keith	Lowery	Advanced Micro Devices
John	Mellor-Crummey	Rice University
Bart	Miller	University of Wisconsin
José	Moreira	IBM
Erdal	Mutlu	Pacific Northwest National Laboratory
Phil	Roth	Oak Ridge National Laboratory
Sameer	Shende	University of Oregon
Galen	Shipman	Los Alamos National Laboratory
Shane	Snyder	Argonne National Laboratory
Devesh	Tiwari	Northeastern University
Theresa	Windus	Ames National Laboratory

Appendix 1.2 Writing Leads

This report was created at an accelerated pace with the help and diligence of the following people. Each person led a breakout summary and provided information for the Executive Summary and findings, or remained at the event for an additional day to help expedite this writing of the report, or provided significant written materials before the meeting and/or edits thereafter.

First Name	Last Name	Affiliation
Dorian	Arnold	Emory University
Phil	Carns	Argonne National Laboratory
James	Custer	Hewlett Packard Enterprise
Kshitij	Doshi	Intel Corporation
Kevin	Harms	Argonne National Laboratory
Mike	Jantz	University of Tennessee
Terry	Jones	Oak Ridge National Laboratory
John	Linford	NVIDIA
John	Mellor-Crummey	Rice University
Bart	Miller	University of Wisconsin
David	Montoya	Trenza
Theresa	Windus	Ames National Laboratory

Appendix 1.3 Workshop Attendees

The following list of people physically attended the West Coast Town Hall. This report would not be possible without the time and commitment from this distinguished group of HPC experts and tools ecosystem veterans. Their insights and input affirm that a sustainable tools ecosystem founded on openness and coordination can be realized. People listed in **bold font** were on-site during the event while people listed in regular font attended the available virtual portion remotely.

First Name	Last Name	Affiliation
Ryan	Adamson	Oak Ridge National Laboratory
Jordi	Alcaraz	Iowa State University
Dorian	**Arnold**	Emory University
Anita	**Benn**	Oak Ridge National Laboratory
Phil	**Carns**	Argonne National Laboratory
John	**Cazes**	University of Texas (TACC)
Piercy	Clark	Oak Ridge National Laboratory
James	**Custer**	Hewlett Packard Enterprise
Kshitij	**Doshi**	Intel Corporation
Anshu	**Dubey**	Argonne National Laboratory
Christian	**Engelmann**	Oak Ridge National Laboratory
Kevin	**Harms**	Argonne National Laboratory
Oscar	Hernandez	Oak Ridge National Laboratory
Kevin	**Huck**	University of Oregon
Clayton	**Hughes**	Sandia National Laboratories
Tanzima	Islam	Texas State University
Michael	**Jantz**	University of Tennessee, Knoxville
Terry	**Jones**	Oak Ridge National Laboratory
Mariam	**Kiran**	Oak Ridge National Laboratory
Jae Hyuk	**Kwack**	Argonne National Laboratory

Matt	**Legendre**	Lawrence Livermore National Laboratory
Ang	**Li**	Northwest National Laboratory
John	**Linford**	Nvidia
Jaelyn	**Litzinger**	Pacific Northwest National Laboratory
Peter	**McCorquodale**	Lawrence Berkeley National Laboratory
John	**Mellor-Crummey**	Rice University
Bart	**Miller**	University of Wisconsin
Dave	**Montoya**	Trenza
Vitali	**Morozov**	Argonne National Laboratory
Swann	**Perarnau**	Argonne National Laboratory
Amedeo	**Perazzo**	Stanford SLAC
Scott	**Richmond**	ESNet
Phil	Roth	Oak Ridge National Laboratory
Justin	Simpson	Oak Ridge National Laboratory
Bill	Spotz	DOE/ASCR
Nathan	**Tallent**	Pacific Northwest National Laboratory
Eric	Van Hensbergen	ARM
Brice	**Videau**	Argonne National Laboratory
Sam	Williams	Lawrence Berkeley National Laboratory
John	**Wu**	Lawrence Berkeley National Laboratory
Weiqun	**Zhang**	Lawrence Berkeley National Laboratory

Appendix 2: Agenda

Day 1 (August 16, 2023) – San Francisco Marriott Fisherman's Wharf

Time (PT)	Topic – Items marked with blue font will be available online via Zoom for remote observers
8:00 – 9:00	Registration & Badging - <outside Golden Gate Ballroom>
8:00 – 9:00	Working Breakfast – St. Helena Room
9:00 – 10:00	Plenary – Opening Comments — Golden Gate Ballroom or <zoom link> • What is STEP and Why Are We Here? [45 mins] (Terry Jones, ORNL) • Remarks and Q&A from DOE Headquarters [15 mins] (Bill Spotz, DOE/ASCR)
10:00 - 10:30	Break – Golden Gate Ballroom
10:30 - 12:00	Plenary – Review of STEP & Our Goals — Golden Gate Ballroom or <zoom link> ● Provisional STEP Structure [10 mins] (Phil Carns, Argonne) ● Challenge 1: Exploding HW Complexity [20 mins] (breakout lead: Mike Jantz) ● Challenge 2: Exploding Use Cases [20 mins] (breakout lead: James Custer) ● Challenge 3: Coordination [20 mins] (breakout lead: Kevin Harms) ● Challenge 4: Management [20 mins] (breakout lead: Phil Carns)
12:00 - 1:00	Working Lunch (Provided) – Discussions on morning sessions – St. Helena Room
1:00 - 1:30	Primer and Desired Outcomes for Breakouts (Terry Jones, ORNL) – Golden Gate Ballroom
1:30 - 3:00	Breakouts Session 1: Prioritizing & Refining Strategies for Key Sustainability Challenges in Tools 1. Challenge 1: Exploding HW Complexity (breakout lead: Mike Jantz) – Mendocino Rm 2. Challenge 2: Exploding Use Cases (breakout lead: James Custer) – Sonoma Rm 3. Challenge 3: Coordination (breakout lead:Kevin Harms) – Monterey Rm 4. Challenge 4: Management (breakout lead: Phil Carns) – Boardroom Rm 5. Virtual Breakout: What Questions Do You Have About STEP? (lead Terry Jones, ORNL)
3:00 - 3:20	Break – Golden Gate Ballroom
3:20 - 3:35	Primer and Desired Outcomes for Breakouts (Terry Jones, ORNL) – Golden Gate Ballroom
3:35 - 5:00	Breakouts Session 2: Key Components Needed for **Phase II Proposal** 1. Short-term Objectives, Priorities & Risks (lead: John Mellor-Crummey) – Mendocino Rm 2. STEP Roles & responsibilities, processes to staff them (lead: David Montoya) – Sonoma Rm 3. STEP Integration in whole Software Sustainability Landscape (lead: Doshi) – Monterey Rm 4. Charting Technical Directions (lead: John Linford) – Boardroom Rm
5:00 - 5:30	Closing Guidance / Adjourn day (Terry Jones, ORNL) – Golden Gate Ballroom

Day 2 (August 17, 2023) – San Francisco Marriott Fisherman's Wharf

Time (PT)	Topic – Items marked with blue font will be available online via Webex for remote observers
8:00 – 9:00	Registration & Badging - <outside Golden Gate Ballroom>
8:00 – 9:00	Working Breakfast - St. Helena Room
9:00 - 9:10	Introduction and Day 2 Logistics (Terry Jones, ORNL) – Golden Gate Ballroom or zoom link
9:10 - 10:30	Report out summaries from Breakout Session 1: Prioritizing & Refining Strategies for Key Sustainability Challenges in Tools – Golden Gate Ballroom or zoom link 1. Challenge 1: Exploding HW Complexity (breakout lead: Mike Jantz) 2. Challenge 2: Exploding Use Cases (breakout lead: James Custer/Phil Carns) 3. Challenge 3: Coordination (breakout lead: Kevin Harms) 4. Challenge 4: Management (breakout lead: Phil Carns)
10:30 - 10:50	Break – Golden Gate Ballroom
10:50 - 12:10	Report out summaries from Breakout Session 2: Key Components of Phase II Proposal – Golden Gate Ballroom or zoom link 1. Short-term Objectives, Priorities & Risks (breakout lead: John Mellor-Crummey) 2. STEP Roles & responsibilities, processes to staff them (breakout lead: David Montoya) 3. STEP Integration in whole Software Sustainability Landscape (breakout lead: Kshitij Doshi) 4. Charting Technical Directions (lead: John Linford)
12:10 - 1:10	Working Lunch (Provided) – Discussions on morning sessions – St. Helena Room
1:10 - 2:25	Plenary 3: Workforce, Security & Survey – Golden Gate Ballroom or zoom link 1. Establishing a Workforce [30 mins] (Dorian Arnold, Emory Univ.) 2. Practices Related to the Security & Integrity of Software & Data [30 mins] (Ryan Adamson, ORN 3. Survey Results and Primer for Breakout Sessions [15 mins] (Terry Jones, ORNL)
2:25 - 3:45	Breakouts Session 3: Additional Challenges: Governance, Workforce and Organizational Challenges 1. DEI & Workforce Development (breakout lead: Mike Jantz) – Mendocino Rm 2. Security and Integrity of Software and Data (breakout lead: Barton Miller) – Sonoma Rm 3. Exercise: What would we want new HPC Users to see from an aspirational viewpoint? (breakout lead: Matt Legendre) – Monterey Rm 4. Bootstrapping STEP in 2024 - How to "Bridge" Tasks in a Flexible way (breakout lead: Kevin Harms) – Boardroom Rm
3:45 - 4:15	Break – Golden Gate Ballroom
4:15 - 5:35	Report out summaries for Breakout Session 3: Golden Gate Ballroom or zoom link 1. DEI & Workforce Development (breakout lead: Mike Jantz) 2. Security and Integrity of Software and Data (breakout lead: Bart Miller) 3. Exercise: What would we want new HPC Users to see from an aspirational viewpoint? (lead: Matt Legendre) 4. Bootstrapping STEP in 2024 - How to "Bridge" Tasks in a Flexible way (lead: Kevin Harms)
5:35 - 5:40	Closing Remarks / Adjourn Town Hall (Terry Jones, ORNL) – Golden Gate Ballroom or zoom link

Appendix 3: Breakout Charge Questions

Breakout 1.1: Exploding HW Complexity

- Rank the collected issues/tasks based on priorities.
- Rank the gathered recommendations based on "most valuable or highest success rate, etc.", and also with the goal in mind to trim down the list of ideas.
- Are there any recommendations to date that seem impractical?
- Should any of the recommendations be phased-in over time rather than enacted at initial launch?
- Are there any remaining gaps to be addressed that you have not seen presented yet?
- Do all of the recommendations clearly address sustainability?
- Are all of the recommendations clearly responsive to the following stakeholders? Poll the room to see what perspective each participant most closely identifies with and make sure that their opinions are heard:
    - Developers
    - Facilities
    - Vendors
    - Applications
- Do the activities for this challenge clearly address the needs of the broader community, whether they were a part of ECP or not?
- Which of the four initiatives most closely align with this challenge?

Breakout 1.2: Exploding Use Cases

- Rank the collected issues/tasks based on priorities.
- Rank the gathered recommendations based on "most valuable or highest success rate, etc.", and also with the goal in mind to trim down the list of ideas.
- Are there any recommendations to date that seem impractical?
- Should any of the recommendations be phased-in over time rather than enacted at initial launch?
- Are there any remaining gaps to be addressed that you have not seen presented yet?
- Do all of the recommendations clearly address sustainability?
- Are all of the recommendations clearly responsive to the following stakeholders? Poll the room to see what perspective each participant most closely identifies with and make sure that their opinions are heard:
    - Developers
    - Facilities
    - Vendors
    - Applications
- Do the activities for this challenge clearly address the needs of the broader community, whether they were a part of ECP or not?
- Which of the four initiatives most closely align with this challenge?

Breakout 1.3: Coordination

- Rank the collected issues/tasks based on priorities.

- Rank the gathered recommendations based on "most valuable or highest success rate, etc.", and also with the goal in mind to trim down the list of ideas.
- Are there any recommendations to date that seem impractical?
- Should any of the recommendations be phased-in over time rather than enacted at initial launch?
- Are there any remaining gaps to be addressed that you have not seen presented yet?
- Do all of the recommendations clearly address sustainability?
- Are all of the recommendations clearly responsive to the following stakeholders? Poll the room to see what perspective each participant most closely identifies with and make sure that their opinions are heard:
    - Developers
    - Facilities
    - Vendors
    - Applications
- Do the activities for this challenge clearly address the needs of the broader community, whether they were a part of ECP or not?
- Which of the four initiatives most closely align with this challenge?
    - Collaboration teams
    - Rapid response
    - Codesign
    - Community activities

Breakout 1.4: Management

- Rank the collected issues/tasks based on priorities.
- Rank the gathered recommendations based on "most valuable or highest success rate, etc.", and also with the goal in mind to trim down the list of ideas.
- Are there any recommendations to date that seem impractical?
- Should any of the recommendations be phased-in over time rather than enacted at initial launch?
- Are there any remaining gaps to be addressed that you have not seen presented yet?
- Do all of the recommendations clearly address sustainability?
- Are all of the recommendations clearly responsive to the following stakeholders? Poll the room to see what perspective each participant most closely identifies with and make sure that their opinions are heard:
    - Developers
    - Facilities
    - Vendors
    - Applications
- Do the activities for this challenge clearly address the needs of the broader community, whether they were a part of ECP or not?
- Which of the four initiatives most closely align with this challenge?

Breakout 2.1: Short-term Objectives, Priorities & Risks

- What are the risks/costs for the DOE if the STEP Center is not funded.
    - What would this scenario look like for stakeholders?
    - Who would be responsible for sustaining this ecosystem?
- What are the most important foreseeable risks if the STEP Center is funded?
    - What strategies can be employed to mitigate them?
- What are the largest non-monetary impediments to performing this work?
- Of the four initiatives (collaborative teams, rapid response, codesign, community activities):
    - Which is most crucial to spin up in year 1?
    - What should be the schedule for phasing in additional initiative efforts?
- What should be the earliest metrics for STEP Center success?
- How do we validate that we are on track with bootstrapping STEP?
    - E.g. do we need early engagement with an external advisory board?

Breakout 2.2: STEP Roles & responsibilities, processes to staff them

- Create a draft summary document for the governance roles in the STEP Center
    - What would this look like as an org chart?
- How large are the boards (both internal and external advisory)?
- What is the initial nomination process to fill these roles?
- How do you stagger terms to keep the positions fresh without losing too much institutional knowledge at once?
- What additional administrative/executive roles are needed for day-to-day operations?
- Is it clear what roles would be responsible for quarterly/annual reporting of activities in each initiative?

Breakout 2.3: STEP Integration in whole Software Sustainability Landscape

- What are the biggest opportunities for interaction with other software sustainability seedlings?
    - What services should we leverage from other seedlings?
    - What can we offer to other seedlings?
    - Which initiatives/activities would be the best contact points?
- What other entities with the DOE should we closely collaborate with, and how?
- What other entities beyond the DOE should we closely collaborate with, and how?
- How do we assess the impact of our work in the community?
- How do we gather feedback on our priorities from stakeholders?

Breakout 2.4: Charting Technical Directions

- What processes should we use to track:
    - technology (i.e., emerging devices that will impact tool sustainability)
    - DOE mission trends (i.e., what are the prominent DOE focus areas)
    - who's using DOE facilities (i.e., what communities we should serve)
- How do we interface with other communities?

    - For example, what can be done to make interactions between HPC tools and system software more sustainable?
    - What can we do to continue fostering research innovation while serving sustainability needs?
- How do we handle conflict resolution if two teams disagree on technical directions
    - For example: if two or more teams should work together, but disagree on interfaces, conventions, or file formats, how do we reach a resolution?

Breakout 3.1: DEI & Workforce Development

- Of the plans / strategies for promoting DEI that were discussed in the previous town hall, which are most important? Which are most likely to succeed?
- Are these strategies sufficient for our DEI goals? What are we missing?
- What potential workforce development efforts should be prioritized first?
- What entities in the governance structure are responsible for tracking our DEI progress?
- What existing organizations/efforts/conferences should we align ourselves with (i.e., by sponsoring/organizing/hosting events)?
- Some extra questions after Dorian's talk:
    - What can we do to measure diversity in our current groups / organizations?
    - Best ways to recruit cohorts?
    - How do we go about measuring impact?

Breakout 3.2: Security and Integrity of Software and Data

- What present day security concerns are hindering the sustainability of HPC tools?
- What are the best practices for ensuring the ongoing integrity of HPC tools that have already been developed?
- Should we prescribe specific guidelines for software maintained within the STEP Center?
    - If so, what?
    - What additional resources outside of STEP can we leverage?
- How do we share expertise (either in the form of manpower or knowledge) on this topic in the STEP Center?
- Is it plausible to make more HPC tool data open to the community without exposing users and facilities to security or privacy risks?

Breakout 3.3: Exercise: What would we want new HPC Users to see from an aspirational viewpoint?

- Imagine a clean slate design and a new user trying to leverage HPC resources for a scientific campaign.
    - What is their "on ramp" for HPC tool usage?
    - How do you minimize jargon?
    - How would you identify what kinds of resources are most crucial to their use case (i.e., is it using GPUs, is it I/O intensive, is it memory bound, etc.)?
    - How do you identify what kinds of tools may be relevant for their needs?
    - How do you branch off to different specialized tools?
    - What kind of visual representations or motifs would be most productive?

Breakout 3.4: Bootstrapping STEP in 2024 - How to “Bridge” Tasks in a Flexible way

- If the STEP Center is tasked with distributing funding before the full governance structure can be established:
    - What is the most equitable way to sustain existing activities?
    - What logistical infrastructure must be in place, and at what institutions?
    - What time horizon is appropriate for this bridge period, and what conditions must be met before transitioning to the proposed STEP operational model?

Appendix 4: Glossary

API	Application programming interface. Syntax and semantics for invoking services from within an executing application.
ALCF	Argonne Leadership Computing Facility
ASCR	The Advanced Scientific Computing Research (ASCR) Program within the Department of Energy Office of Science is a program with the mission to discover, develop, and deploy computational and networking capability to analyze, model, simulate and predict complex phenomena important to the Department of Energy and the advancement of science.
BER	Biological and Environmental Research (DOE-SC)
BES	Basic Energy Sciences (DOE-SC)
BRN	Basic Research Needs
Co-design	Co-design refers to a computer system design process where scientific problem requirements influence architecture design and technology and constraints inform formulation and design of algorithms and software.
Continuous Integration	A software engineering practice during the build and unit testing stages of the software release process. Every revision that is committed triggers an automated build and test.
DOE	The United States Department of Energy.
DOI	A digital object identifier is a persistent identifier or handle used to uniquely identify various objects.

ECP	Exascale Computing Project
ESnet	Energy Sciences Network
FAIR	Findable, Accessible, Interoperable, Reusable
FOA	Funding Opportunity Announcement
GPU	Graphics Processing Unit. A GPU may be used together with a CPU to accelerate scientific and analytical workloads.
HPC	High Performance Computing.
LLNL	Lawrence Livermore National Laboratory
LLVM	Low Level Virtual Machine
ML	Machine Learning. Machine Learning is a field of science devoted to understanding and building methods that 'learn', that is, methods that leverage data to improve performance on some set of tasks. It is seen as a part of artificial intelligence.
NERSC	National Energy Research Scientific Computing Center
Open source	Software that is available to users in source form and can be used and modified freely.
ORNL	Oak Ridge National Laboratory
OS	Operating System
SC	Office of Science (DOE)
STEP	Sustainable Tools Ecosystem Project.
Tool	The collection of tools and utilities for analyzing and optimizing application performance, identifying correctness problems, and debugging.
UQ	Uncertainty quantification.
V&V	Validation and verification.