

STEP Midwest Town Hall Report



A Sustainable Tools Ecosystem Project Event
Held at Iowa State University, Ames, Iowa
July 11-12, 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 Spatz.

Special thanks are due to Iowa State University and Midwest Town Hall host, Theresa Windus, for generously providing facilities, resources and time for this event.

Finally, the writing team who created this report are appreciated and listed [here](#).

 Terry Jones, STEP PI

TABLE OF CONTENTS

1. Executive Summary & Key Findings.....	3
2. Midwest Town Hall Format & Content	7
2.1 Introduction.....	7
2.2 Plenary Portion	8
2.3 Breakout Portion	8
SESSION 1: Structure & Setting Direction.....	8
Breakout 1.1: Management/governance structure.....	8
Breakout 1.2: Stakeholders & External Advisory Boards.....	9
Breakout 1.3: Processes for Tracking Vision & Gaps Over Time.....	11
SESSION 2: Software Lifecycle.....	14
Breakout 2.1: Bootstrapping Efforts.....	14
Breakout 2.2: Reporting & Maintaining Efforts	15
Breakout 2.3: Reviewing & Sunsetting Efforts.....	18
SESSION 3: Additional Challenges: Governance, Workforce and Organizational Challenges .	19
Breakout 3.1: DEI & Workforce Development	19
Breakout 3.2: Community Activity models & scope	21
Breakout 3.3: Interfacing With Additional Funding Sources.....	22
Breakout 3.4: Facilitating Ongoing Collaboration: What entices Application teams / Facilities to work with STEP Efforts.....	24
3. Next Steps.....	26
3.1 Distribute Findings of Town Hall.....	26
3.2 Refine strategies for understanding current state of tools ecosystem	26
3.3 Design and organize West Coast Town Hall.....	26
3.4 Coordination Activities with other Software Sustainability seedlings	27
4. References.....	28
Appendix 1: Attendees	29
Appendix 1.1 Workshop Organizers.....	29
Appendix 1.2 Writing Leads	30
Appendix 1.3 Workshop Attendees.....	31
Appendix 2: Agenda	33
Day 1 (July 11 th , 2023)	33
Day 2 (July 12 th , 2023)	34
Appendix 3: Glossary	35

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.

The second Midwest town hall was held on July 11-12, 2023, at Iowa State University in Ames, IA. 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 [Appendix 1.3 Workshop Attendees]. This town hall aimed to address a wide range of topics related to the first three tools sustainability challenges outlined in our proposal [1]: exploding hardware complexity, exploding use cases, and coordination.

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.

One of the early plenaries presented a proposal for structuring the STEP Center in terms of four core initiatives (see presentation [here](#)). Unlike the four “challenges” that STEP intends to address

(exploding hardware complexity, exploding use cases, coordination, and management), an initiative is an organizational grouping of different types of activities that collectively will frame the activities that address those challenges. We identified the following organizational initiatives:

- Collaborative teams: This initiative is the most conventional of the four. Its name reflects that most sustainability problems are better addressed through ongoing collaborative effort to address a shared concern rather than support for isolated tool projects. This initiative would support teams (not exclusive to tool developers) who propose to work together to address a long-term sustainability challenge.
- Rapid response: This initiative is driven by a recurring theme in sustainability discussions: technology needs (hardware, use cases and more) often evolve faster than the corresponding tools that support them. We therefore need an agile mechanism to address unanticipated needs as they arise, independent of the conventional funding cycle. Rapid response activities will be characterized by administratively streamlined, short duration projects possibly leading to Collaborative Teams initiative work later. It is somewhat analogous to “contingency funding”, but as a first-class citizen with well-defined and transparent governance.
- Codesign: This initiative is driven by the observation that we need to improve coordination across community stakeholders to truly address sustainability. It consists of explicit cross-ecosystem engagement between tool developers and other parties for codesign to address software tool sustainability problems that span the ecosystem. Tool developers would work directly with applications, facilities, vendors, system software developers, and other parties for joint solutions that require ongoing multi-party contributions to achieve.
- Community activities: This is driven by the need to facilitate community-wide sustainability activities that are beyond the scope (and the burden) of any one HPC tools team. A Community Activity is any center-wide activity that benefits the community at large and is not well-served by implicit manpower contributions from individual teams. Examples (not exhaustive): representation in standards bodies, hosting center-wide meetings, promoting workforce development, and hosting education events. Focus on community sustainability issues that transcend individual projects.

TODO: can we post the presentation about this on the web site and link to it here for more information?

Several key findings and recommendations emerged as recurring themes as we discussed the initiatives and explored more detailed topics during the breakout sessions. These include:

- **Finding:** The STEP Center can enhance sustainability by being as transparent as possible.
 - Recommendation: Establish and publicly publish clear codes of conduct, reporting and review artifacts, management policies, and assessment criteria.
 - Recommendation: Establish a clear plan for how to manage funding in the Rapid Response initiative; there was a consensus that it will be the most difficult to

demonstrate good stewardship for the Rapid Response Initiative. It may initially start with no explicit funding, in which case we will leverage it as a means to adjust or redirect existing effort. Once the Rapid Response initiative is fully enacted, we can ensure that any leftover funding each year is explicitly rolled over into the Community Activities initiative to support high-value community events.

- Recommendation: Decisions regarding funded activities should be made by a diverse leadership team based on clearly documented governance guidelines.
- Finding: Tracking stakeholder needs over time will be crucial to sustainability.
 - Recommendation: Conduct ongoing surveys and interviews, particularly with facility operators and application developers, to maintain an accurate view of what is and isn't working.
 - Recommendation: Create, publish, and continually update a strategic plan that serves as a roadmap for how the STEP Center will focus its activities.
 - Recommendation: The STEP Center should consider methods to support tools that are perceived as crucial to the community even if they are not explicitly funded by STEP. This is an opportunity for other funding sources to influence STEP activities, for example by supporting minimal ongoing maintenance or by Community Activities that support the transition to alternative tools.
- Finding: The STEP Center needs to embrace external review and feedback from a broad perspective.
 - Recommendation: Establish a diverse (across multiple dimensions) external advisory board with at least 7 members whose primary responsibility will be to conduct an annual review of STEP Center activities and provide feedback both to STEP and DOE program management.
 - Recommendation: In addition to annual reporting, the STEP Center web site should be used as a clearinghouse for success stories and a curated list of HPC tools that meet clearly defined criteria for inclusion in STEP.
 - Recommendation: Curate use cases for the community to use for experimentation and evaluation. This may take the form of benchmarks or mini-applications that can be used for standardized evaluation of tools and techniques that are developed and sustained within STEP. Development of these use cases could be pursued within the Codesign initiative.
- Finding: Excessive administrative overhead is a hindrance to productive sustainability efforts. We should streamline the participant experience and make sure that reporting efforts maximize benefit to the community.
 - Recommendation: The proposal process for STEP Center projects should be as lightweight as possible with clear guidance on how to provide concise information

about sustainability needs. We can look to examples such as the nanoscience centers and the Molecular Sciences Software Institute.

- Recommendation: Project reports should take the form of instructions for a reproducible demo of the work along with a concise report. An artifact evaluation team (perhaps funded in part by grants or stipends) will evaluate the demos. Badges will be awarded to projects that meet clear sustainability criteria.
- Recommendation: Semiannual PI meetings will be held with a primary objective of sharing information between project teams.
- Recommendation: Annual STEP meetings will be held to bring together stakeholders and developers to highlight new developments and requirements. Smaller special-interest group meetings may be held at a higher cadence.
- Recommendation: All STEP projects will have clearly defined durations, but the decision making process for reviewing these projects may differ depending on the STEP initiative that it belongs to. PIs may propose additional efforts if needed as projects reach their natural conclusion points.
- Finding: Workforce sustainability is crucial to the long-term viability of the tools community.
 - Recommendation: Community Activities should include outreach to high school and undergraduate students to foster interest. Potential vehicles include internships, assistantships, and REU/SULI type efforts. We should also look to expand outreach through SBIR and Google Summer of Code projects.
 - Recommendation: The STEP Center should coordinate with external groups, such as the Sustainable Horizons Institute, to amplify its impact.
- Finding: The STEP Center should actively seek to engage in additional sources of funding.
 - Recommendation: Act as a clearinghouse for expertise that can be engaged in activities such as:
 - Participation in procurement and evaluation activities (both as a service to the community and to increase visibility)
 - Participate in proposals to advocate for cases in which tools sustainability are needed to support the needs of scientific activities.
 - Recommendation: Establish guidelines to use as a starting point for negotiating with external parties who wish to support sustainability activities within STEP. We cannot establish a complete procedure without working with a third party to determine how to best meet their needs, but we can speed up the process by providing a starting point for discussions about how to leverage the benefit of STEP resources.

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. Midwest 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. As useful, important topics are revisited in subsequent Town Halls to gather additional information, enabling stakeholders who may have missed a previous meeting to provide input.

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

Plenary Session 1: What is STEP and Why Are We Here?

Plenary Session 2: Recap of the East Coast Town Hall and Description of the 4 STEP Initiatives

Plenary Session 3: Primer: Governance, Workforce and Organizational Challenges

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: Structure & Setting Direction

1. Management/governance structure
2. Stakeholders & External Advisory Boards
3. Processes for Tracking Vision & Gaps Over Time

Breakouts Session 2: Software Lifecycle

1. Bootstrapping Efforts
2. Reporting & Maintaining Efforts
3. Reviewing & Sunsetting Efforts

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

1. DEI & Workforce Development
2. Community Activity models & scope
3. Interfacing With Additional Funding Sources
4. Facilitating Ongoing Collaboration: What entices Application teams / Facilities to work with STEP Efforts

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

SESSION 1: Structure & Setting Direction

Breakout 1.1: Management/governance structure

The Management and governance structure breakout session took place on Tuesday, July 11th. The moderator was Terry Jones (ORNL). Additional attendees were James Custer (HPE), Richard Gerber(LBL), Paul Kent (ORNL), David Montoya (Trenza), and Jose Moreira (IBM).

The group discussed what can be done to promote transparency and impartiality in the management of STEP center activities. Here our goals should be to have a code of conduct, clear review, and assessment criteria written down. Since STEP will be new and unproven, we need to plan for plenty of metrics and reporting to establish ourselves; these material should survive audit

The group discussed aspects of STEP's proposed Rapid Response Initiative including ways to fund such activities. ECP maintained a contingency account so there is a DOE/SC precedent of a contingency account. Certainly it should be the goal of STEP to avoid maintaining a large unused contingency funds account if possible. Given the desire to be efficient, we had two strategies to apply to such a need:

- First, whenever appropriate, redirect unused contingency funds.
- Second, we should have Rapid Response option totally without new funds (e.g., rearrange priorities). By redirecting existing efforts to address the new need and allow lower-priority work to slip to revised milestone dates.

Another line of discussion centered on ECP's Continuous Integration (CI): it was noted that such standardized practices should be core to software sustainability.

As the group considered if there other examples in the community that we should look to for guidance, several sources of information were recommended:

- To augment our tool surveys: Determine important applications (perhaps by allocations) and find out what tools those applications use and want to see continued;
- To inform project structure: Learn from Facilities' "small p" projects; talk to their project planning team about what has proven to be useful

It was noted that we got to where we are in the current tools ecosystem organically and that we can do better with an integrated plan. Moreover, since the independent functional teams focus on their topic area (go deep rather than go wide), they avoid a single point of failure.

Breakout 1.2: Stakeholders & External Advisory Boards

The purpose of this breakout discussion was to explore how to facilitate sustainability of the STEP Center as a whole by ensuring accountability to external stakeholders. We discussed key issues, actionable mechanisms, and how the needs of STEP differ from other similar organizations.

There was broad consensus of the need for an external advisory board that is empowered to objectively review STEP activities from a "big picture" view and provide feedback to both STEP leadership and DOE program management. This advisory board must avoid conflicts of interest, be inclusive of different perspectives, and provide earnest feedback to the STEP Center. The breakout participants were able to draw from extensive past experience to provide recommendations on what would make an external advisory board most effective for our purposes.

A second discussion theme was how to ensure that STEP Center reporting activities (both to the advisory board, as well as reporting to the general public) are beneficial to the project and not just a bureaucratic obligation. We identified several ideas for how to provide public reporting that not only share progress on our activities but also directly contribute to community sustainability.

The final discussion theme we covered was how to incorporate public reporting into the proposed STEP Center structure (i.e., the four primary center initiatives.) to make sure that each of them is on track with community expectations. We identified that the "rapid response" initiative is likely to face the biggest reporting and accountability challenges.

Our discussions produced the following actionable sustainability recommendations:

- The STEP Center should assemble an external advisory board to provide broad guidance from an external perspective. We recommend formulating a panel of 7 members who are diverse across the dimensions of DEI (e.g., following SC committee guidelines), community and technical background, agencies (e.g., DOD and NASA), and DOE offices. It should also include a DOE program management representative, though they may elect to participate in an observer role. Terms should be staggered to preserve continuity.
- The primary responsibility of the STEP Center external advisory board is to convene once a year to review the overall STEP Center report and provide guidance, which will then be included in a report submission to the DOE. Advisory board members are likely to also be good candidates for project proposal review panels as well, however.
- STEP Center external reporting will take two forms. The first is a formal annual report presented to the external advisory board and DOE program management. Progress towards the report will be monitored by way of a lightweight quarterly briefing from each of the four STEP initiatives.
- The second form of STEP Center external reporting is ongoing reporting to the general public. This will be done via the STEP Center website. We recommend three such forms of reporting:
 - A synopsis of each year's annual report (with the full report available upon request)

- Ongoing publication of "success stories" that demonstrate STEP's impact on the community. We recommend consideration of small stipend awards to encourage generating these stories.
 - Ongoing curation of a list of high-quality HPC tools. This list must be curated by STEP itself, and all entries must meet criteria that demonstrate sustainable participation in the STEP community.
 - We recommend that reporting for the Rapid Response initiative receive special consideration. If it is an explicitly funded component of STEP, then we must account for any leftover rapid response funding each year. Our recommendation is that any such funding be rolled over into the Community Activities initiative for the purpose of hosting in-depth events.
-

Breakout 1.3: Processes for Tracking Vision & Gaps Over Time

The charge questions for this breakout session explored various critical aspects for tracking STEP's vision and gaps over time. As the tools landscape continually evolves, a key question was how STEP actively tracks sustainability needs and priorities. It involved discussions about identifying priority needs and sustainability problems. Questions also examined the relationship between STEP's vision tracking and applications. Furthermore, a debate on utilizing surveys to understand better the HPC Tool landscape was vital to this breakout. The moderator was Heike Jagode (UTK). Additional attendees were Theresa Windus (AL), Adam Lavelly (LBNL), Shane Snyder (ANL), Chunhua "Leo" Liao (LLNL), and Vitali Morozov (ANL).

Plan Development: To effectively track the sustainability needs and priorities of STEP-funded tools, it's crucial for STEP to develop a strategic plan. This plan should detail the sustainability objectives, priorities, and anticipated outcomes for these tools. The process of developing this plan calls for broad communication with tool developers, vendors, facilities, and developers as well as users of applications. Although surveys (see below) can be useful for collecting information, they often present challenges. Hence, employing other strategies, such as targeted engagement with key stakeholders and discussions with DOE leadership, is vital. Considering the continuously evolving nature of hardware capabilities, software technologies, user and facility requirements, as well as other crucial technical and DOE needs, the plan must undergo regular reviews and updates.

Definition of Metrics and Key Performance Indicators (KPIs): To evaluate priority needs and monitor shifts in the landscape of performance tool usage and development, it's imperative to establish a clear set of metrics and KPIs. This can be a difficult task since one set of metrics may not

be appropriate for all tools and it is often difficult to gather the information for the metrics. Therefore a broad swath of metrics needs to be developed. A few examples to consider include:

- The number of active users, which provides an indication of the tool's reach and usage.
- The frequency of updates or releases, which showcases the responsiveness of the tool's developers to evolving needs and bug fixes.
- Evaluating success in meeting project milestones, which offers insight into the tool's progression and reliability.
- The tool's adaptation to new computing architectures and monitoring capabilities, which is critical to ensuring its future relevance and usability.
- Lastly, user satisfaction levels, which give a direct measure of how well the tool is meeting its users' needs and expectations.

Taken together, these metrics provide a fair view of a tool's performance and potential areas for improvement.

To ensure the successful progression and adaptation of metrics, it's crucial to establish regular reporting, such as an annual review. This review should provide a comprehensive overview of progress made towards the defined metrics and KPIs, any updates or improvements made to the tool, any challenges faced during the course of operation, and an outline of future plans.

User Feedback and Engagement with Applications: Another viable approach to track how the STEP efforts relate to the needs of the scientific applications community is to proactively gather user feedback and promote active engagement with applications. It's crucial that the STEP-funded tools maintain consistent interaction with various applications and incorporate their feedback into regular reports. Furthermore, a fruitful way to connect with the application community may be to require tool developers to interact with a certain number of new or existing applications as a part of their STEP funding requirements. The feedback and success stories resulting from these interactions could serve to streamline STEP's vision tracking relative to applications, confirm that STEP activities meet user needs, and pinpoint areas for tool improvement.

Surveys: To better understand both applications' needs and the HPC tool landscape, STEP could effectively utilize surveys for a variety of insights. These include identifying needs and priorities by comprehending what tool users and developers perceive as essential features. Surveys can also assess user satisfaction by capturing what users appreciate about current tools. Furthermore, they can identify emerging trends or future needs by determining what users think is missing or could be improved. In addition, surveys can help understand user patterns by identifying how, when, and why users utilize certain tools. Lastly, assessing training needs is another crucial aspect, as it helps pinpoint whether users feel they have sufficient knowledge to use the tools effectively.

Nevertheless, as previously mentioned, a significant challenge remains. It can be notably difficult to motivate users to dedicate time to completing these surveys, presenting an obstacle to obtaining the comprehensive insights desired. Several suggested strategies can enhance survey participation and feedback quality:

1. Mandate that users of facilities with INCITE/ALCC or similar allocations provide information related to performance tools within their applications, including all dependencies, or within their annual reports.
2. Leverage special events such as hackathons and workshops to conduct surveys, where attendees are more inclined to participate.
3. Enhance transparency by publishing a roadmap outlining how survey results will shape STEP's activities and make these results publicly accessible.
4. Promote the resolution of problems identified via surveys, as this can incentivize further participation.

Evaluation of Tool Capabilities via Benchmarks: Another effective method to keep track of the proficiency and priorities of performance monitoring features in STEP-funded tools is to assess their capabilities using a standardized set of benchmarks, mini-apps, or proxy apps (in this context, we will use 'benchmarks' as an umbrella term to represent all of these terms). The purpose of recommending the employment of a set of benchmarks on STEP-funded tools is twofold. First, it allows STEP to evaluate the tool's ability to identify specific performance issues on new architectures. Second, it helps in pinpointing areas where the tool may require improvements or enhancements. This approach ensures that the tools are not only efficient but also effective in addressing the continually evolving needs of the hardware landscape.

The selected benchmarks should adequately represent patterns utilized in real-world applications. This can be further supplemented by incorporating full applications into the mix, as smaller benchmarks, while useful, might not fully capture the complexity that arises from the interaction of various hardware components, such as CPU/GPU, memory, and interconnect. Therefore, it's essential to include benchmarks that present significant challenges to ascertain if a tool can effectively identify and handle these issues. This process will also highlight areas needing improvements in the tools. As an added measure, STEP could consider introducing challenge problems that put these tools to the test, fostering a competitive environment for continuous improvement and innovation.

SESSION 2: Software Lifecycle

Breakout 2.1: Bootstrapping Efforts

The Bootstrapping Efforts breakout session took place on Tuesday, July 11th. The moderator was Theresa Windus (Ames National Laboratory). Additional attendees were Adam Lavelly (LBNL), Erdal Mutlu (PNNL), James Custer (HPE), Amedeo Perazzo (SLAC); Terry Jones (ORNL); Paul Kent (ORNL), and Dave Montoya (Trenza). The primary topics of discussion were (1) defining and understanding the four initiatives from the East Coast Town Hall (Collaborative Teams, Rapid Response, Codesign, and Community Activities), (2) the short and long term needs of STEP funding, (3) the process and requirements for teams to apply for funding for sustainability activities in the STEP center, and (4) the decision process for funding projects.

There was considerable discussion around defining and understanding the four initiatives that resulted from the previous Town Hall. Collaborations were defined as more horizontal interactions (within or between different tools), while Codesign was defined by more vertical interactions (between vendors, tools, facilities, and applications). Rapid response was seen as solving a specific problem in a short amount of time or as being a seed to provide initial results or proof of concept for one of the other initiatives. Community activities involve reaching out to the broader community for education, training, workforce development, and visibility of STEP that could be short or longer term activities that are primarily for STEP, broader ASCR initiatives, other community activities, or other sustainability institutions.

The bootstrapping efforts for STEP and STEP projects need to be considered in the context of both the short- and long-term goals of STEP. With the end of ECP in December of 2023, there will be short-term needs to make sure that key tools expertise is not lost due to a delay in funding. In this phase, the group recommended that STEP gather initial data from facilities and other sources to understand what tools are a priority to initially fund some projects to start STEP quickly. However, STEP should transition within 6 months to a proposal based system. In addition, the group recommended that STEP consider delaying the Rapid Response portion for a year or two to allow the STEP organization to obtain a consistent baseline of activities and to allow time for management processes to be developed. The longer-term goals of STEP will require gathering clear data, setting priorities, and making sure that there is broad participation (within funding constraints).

Due to current uncertainties with the duration and funding level of a sustainable software FOA, the duration of the projects within STEP and for STEP itself is an open question. The Collaborations and Codesign projects were seen to likely be longer term activities to encourage sustainability of the software ecosystem as well as the funding sustainability of the tools within STEP. One of the advantages from the ECP project was that longer term funding helped the researchers to focus on

the work to be done without having to constantly write proposals. This was seen as a strong reason to have longer term funding for key components as part of STEP. It was suggested that there could be some minimal “base” funding for tools within STEP with an expectation that projects will search for additional funding through other sources. Rapid response activities should be on the 3-6 month timescale, while Community Activities will likely have a mix of short and long term projects. Within the Rapid and Community Activities, there was significant discussion around whether these should be funded by reprioritizing activities within the Collaboration and Codesign projects (i.e., a pool of funded experts who would address these issues in an “agile” fashion) or whether there should be separate funding for these projects.

All agreed that the proposal process should be as lightweight as possible with explicit criteria and transparent decision making. Potential criteria could include capabilities provided, use cases, value to the user base, other tools that are part of their ecosystem (collaborative elements or dependencies), assessment (reporting, code, etc. that shows the value of what is being created), addressing sustainability (not having to constantly write proposals, justifying resource needs, etc.), software sustainability, and support of the software usage. However, the group recommends that STEP management contact other centers that have lightweight processes such as the nanoscience centers and the Molecular Sciences Software Institute.

The decision making process will be highly dependent on the final management structure of STEP. For example, if STEP has a structure with a STEP Director, an Executive Board (with the Deputy Director as one of the members), and an External Advisory Board (one that has no conflict of interest with STEP), then one possible mechanism is to have two to three reviews from the External Advisory committee that are then reviewed with additional comments by the Executive Board. Based on these inputs, the STEP Director would make a final decision on the approved projects. In any particular decision making scenario, however, the process should be fully transparent with instructive feedback to all proposing teams. It was also noted that the STEP Center has an important opportunity and responsibility to coordinate overlapping efforts between teams that might otherwise compete with each other or simply not be aware of each other.

Breakout 2.2: Reporting & Maintaining Efforts

The Reporting and Maintaining Efforts breakout session took place on Tuesday, July 11th. The moderator was Jose Moreira (IBM Research). Additional attendees were Phil Carns (ANL), Heike Jagode (ICL), Chunhua Liao (LLNL), Bruce Palmer (PNNL), and Andreas Kronfeld (Fermilab). The main topics of this breakout session include (1) process for progress reporting of various projects and (2) mechanisms for adjusting project goals and milestones.

Progress Reporting

STEP Center supervised projects will report their progress using a combination of traditional written reports and live tool demonstrations. Following established DOE practices, each project will submit quarterly progress reports, describing the main accomplishments and milestones achieved in that quarter. Projects will also prepare a yearly summary that will be included in the STEP Center annual report. In addition to these traditional reports, there will be scheduled live sessions to demonstrate the functionality and features of the tools developed by STEP projects.

We believe that the accomplishments and progress of the STEP Center projects should be broadly communicated to all members of the STEP Center and to the HPC community at large. We plan to achieve that by organizing semiannual PI meetings, with open invitation to all STEP members and additional invited guests from the community. Of the two PI meetings in a calendar year, one will be an in-person meeting, to promote closer interaction among STEP members. The other PI meeting will be virtual, to help manage travel demands. At the PI meetings there will be the traditional slide presentations about the various projects, and also simplified demos of the tools being developed. These demos are more for illustration and will be completed with more thorough live demos to be scheduled separately.

We also plan additional public events that will serve both a progress report function as well as a publicity function for STEP projects. These include Birds-of-a-Feather events at the SC conference series and workshop/tutorials at IEEE Cluster and IPDPS conference series. Tutorials can be organized as an overview of various tools in a single tutorial session, thus reducing the burden on any one tool team to prepare a full tutorial.

Live Demonstration of Tools

STEP Center will establish an *Artifact Evaluation Team (AET)*, possibly part of the STEP Codesign Initiative, that will be in charge of testing and verifying the functionality of tools developed by STEP Center projects. The evaluation process would work more or less as follows:

1. The project team, in charge of developing a new tool or a new feature for an existing tool releases the latest version of its work. It also prepares instructions on how to install and use that tool or feature, possibly including a script for a suggested demonstration of the new capabilities of that tool.
2. STEP AET experiments with the tool. Any difficulties in installing or using the tool are reported and discussed with the tool developers, until the AET is satisfied with the solutions. The AET can and should experiment the tool on different use cases, and also investigate its internals to verify if the tool developers followed the best practices and architecture endorsed by STEP Center.

3. STEP AET schedules a live demo of the tool for the broad STEP community, and additional invitees from the community, to showcase the functionality and features of the tool. After successful completion of this live demo, the tool team is considered to have passed the criteria for approval of that live demo milestone.

The evaluation of a new tool or feature concludes with the granting of *badges* (or *tags*) to the tool by the AET. A badge certifies that a tool has a certain capability, tested and verified by the AET. STEP should clearly identify the set of badges to be granted. Possible badges include (1) and Open Source badge to certify that tool is available with a STEP-approved Open Source license; (2) Functionality badges certify that the tool works as advertised; (3) Usability badges certify that independent teams have been able to successfully use the tool without help from the developers.

The AET is expected to consist of a combination of students and more experienced personnel. Students can gain valuable experience by participating in the AET, and this activity could be considered a workforce development initiative. However, supervision and guidance for the more experienced personnel is essential to the success and integrity of the AET.

Use Cases Collection

Effective interaction between the tool project teams and the AET requires establishing a collection of use cases that both teams can experiment with. These use cases are meant to represent relevant situations that the tools are expected to address. They can be used throughout the development of a project to motivate and test features.

STEP Center is in a privileged position to create a collection of relevant use cases. By representing the broad community of HPC tools, it can interact with application teams to identify the important cases, while minimizing the burden on those teams that would otherwise have to interact directly with multiple tool teams.

Working with full-size applications is not always feasible, neither for the tool teams nor the AET. The STEP Center should distill the use cases it collects into a set of *reproducers*. Reproducers are simplified code that represents specific behaviors of full-size applications that are particularly relevant to test a tool.

Adjusting Project Goals and Milestones

As a multi-year project evolves through its lifetime, it is natural to require adjustments of its goals and milestones. The adjustments may be necessary because of both internal and external factors. For example, a project may run into technical challenges that are beyond its capability to solve. Conversely, requirements from applications or systems may change, making certain tool features obsolete and requiring new ones.

Requests for changing a project goal or milestone can be initiated both by the project team itself or by STEP Center, particularly from the Rapid Response Initiative. A *Change Review Board* (CRB) in the STEP Center evaluates those requests, verifying that they are both justified and technically feasible. If the CRB approves the change in milestones/goals, and the tools team agrees with those changes, the changes become official.

Breakout 2.3: Reviewing & Sunsetting Efforts

The Reviewing and Sunsetting Efforts breakout session identified key considerations and mechanisms for addressing the end-of-life for STEP funded initiatives. Shane Snyder led the session and it was attended by John Linford (NVIDIA), Vitali Morozov (ANL), Keita Teranishi (ORNL, S4PST PI), Ryan Adamson (ORNL). The main objectives of this breakout were to determine (i) What criteria should determine when a STEP activity has reached or is about to reach end of life, (ii) How information from STEP activity reports and other sources could influence the decision making process, and (iii) What governance body ultimately makes the final decision regarding sunsetting of STEP funded activities, tools, and projects.

During open discussion, several themes re-emerged during the discussion of these three charge questions, the most important of which is that **STEP funded activities are an investment in software sustainability**. As such, some amount of effort for maintenance and support of these contributions should continue for some period of time after features or tools have been developed. Enterprise operating systems are typically supported by the vendor for several years in order to give stability to consumers – performance tools could enjoy the same benefits. Sunset timelines should most likely be determined at the start of a STEP funded project activity. To reduce potential ‘long tail’ of support funding, maintenance efforts during sunset of a project could be focused solely on ensuring the continued functionality of the capabilities developed by the project. For example, a project that is funded to specifically adapt a tool to measure performance on a MI200 GPU should not fund maintenance/bugfix work for measuring performance on Arm processors unless somehow related to the MI200 GPU scope.

Another key discussion point was that **STEP funded activities may have different decision making processes depending on their type and duration**. For example, short-term *Rapid Response* activities will address disruptive risks in a tactical way and require little to no maintenance funding. These activities can be closed by STEP once the risk to software sustainability has been mitigated. On the other hand, projects that augment existing performance tools (such as adding support for new accelerators) or projects that develop new tools may be proposed to be funded for several years. These projects will have proposed a project plan that was approved by STEP decision makers and

sunset decisions will be made based on the project proposal by STEP. The breakout group realized that the funding profile of these projects would be front loaded: time-consuming feature development takes place at the beginning of a project lifecycle whereas the less-intensive maintenance of that feature continues until closeout. Finally, the longer term, strategic initiatives that STEP may fund are broader in scope and will most likely need to take into account input from scientific application teams, ASCR facilities, and vendor partners. In any case, as a project nears its end of life, PIs may propose additional work on the same tool. When funded, this additional work would be a new project with its own sunset timeline.

Decisions regarding STEP funded activities should be made by a diverse leadership team that is free from conflicts of interest, and in collaborative agreement with the STEP DOE program manager. The leadership team will require information from many sources to make determinations about when to sunset project activities. Representatives on the leadership team or on a STEP external advisory board that have tool development, scientific application, facility, or vendor background is important. Activities should be reviewed in at least two ways: (i) Determining the magnitude of the benefit of activity outcomes on the scientific software community, and (ii) The technical merit of the proposal to determine that the technical plan is indeed appropriate to address the need.

The final discussion of the breakout session focused on minimizing any adverse impacts to the community, software projects, and software project teams when sunsetting. **For software tools or capabilities that STEP will no longer support but are still widely used within a certain community, other funding should be explored with stakeholders within that community.** For tools or capabilities that are no longer in wide use because a similar tool has gained popularity, STEP could utilize outreach and training initiatives to transition developers from one sunset project to another, funded project. Finally, if a tool is being sunset and there is no comparable tool, STEP should ensure that the risk to the scientific software community is indeed low, or consider initiating a *Rapid Response* activity to develop a similar capability within one of the STEP-funded tools.

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

Breakout 3.1: DEI & Workforce Development

This breakout session examined strategies to promote inclusive and equitable research and workforce development in STEP. It primarily focused on identifying the most effective ways to foster

DEI and workforce development through the STEP project. The breakout participants suggested the following approach:

STEP should work to engage high school and undergraduate students to foster interest in computing tools and HPC. Potential strategies include: funding internships, assistantships, or REU / SULI type efforts. Interested students could team up with a scientist participating in STEP or a professor from a collaborating university with the aim of contributing to one of the tools supported by STEP. Some participants also noted that there are existing programs (e.g., mentor / mentee at SC, sustainable research pathways in ECP) that STEP could use as a model or use as a vehicle to implement its own program.

STEP should also aim to integrate supported tools into early undergraduate coursework. Doing so would not only foster interest in younger populations, but it could also increase the exposure and user base of individual tools in STEP. However, multiple participants noted that developing undergraduate coursework based on STEP tools may be very difficult because many of the STEP tools require knowledge and expertise of multiple areas, including architecture, system software stack, and application needs, to operate effectively. However, it was also noted that this process may prove beneficial from a user interface / user experience perspective as making tools accessible to younger / less knowledgeable audiences will also increase the population of professional users that can operate these tools effectively. Additionally, it was noted that STEP could potentially integrate some tools with high school coursework as well, but this work should target platforms and languages that these students may have used and are familiar with (e.g., Python). Participants also agreed that, since course development efforts could be very time consuming, any developed software, projects, and instructional material should be shared with and promoted to the broader community of instructors to facilitate its reuse.

Participants also considered whether there are unique advantages in STEP with regards to DEI and workforce development. One panelist noted the broad reach of the tools community and the popularity of some of the tool products in STEP could be an advantage.

Participants next considered the question of how to organize leadership of DEI activities. Participants agreed that STEP should aim to find groups and individuals with expertise in DEI and workforce development to lead these efforts. For example, STEP could contract with nonprofit groups, such as the Sustainable Horizons Institute, to guide and maximize the impact of these efforts. One participant noted an advisory board consisting of people from universities, labs, and the private sector, with expertise in these different domains would also be beneficial.

Lastly, the participants recommended that STEP should engage with outside groups to promote DEI and workforce development. Some participants noted that STEP should invite students and faculty from MSIs and HBCUs to STEP events, or even consider holding STEP events at a well-located MSI / HBCU to increase engagement with the groups at these schools. The participants also noted several professional and campus organizations with which STEP could collaborate to foster interest and diversity in the HPC tools ecosystem, including: Women in HPC, Alces Flight, Systems (UTK), and many other student and professional organizations with the goal of fostering interest in computing or HPC among individuals within diverse or underrepresented groups.

Breakout 3.2: Community Activity models & scope

The Community Activity models and scope breakout session took place on Wednesday, July 12th. The session lead was Erdal Mutlu (PNNL). Additional attendees were Bruce Palmer (PNNL), Shane Snyder (ANL) and Theresa Windus (AL & ISU). The main topics of this breakout session include (1) describing how the community activities should be structured and (2) what should be their scope to construct a community beyond funding boundaries.

Recurring Meetings

STEP Center aims to follow and build upon the successful community activities from other organizations by having recurring events on different frequencies. One such outreach is Exascale Computing Project (ECP)'s annual meeting where people from various fields come together and observe developments on all of the projects within the ECP portfolio. Such an event will allow a multidisciplinary outreach for expanding the capabilities of the HPC tools as well as increasing the user-space. Following a similar approach, STEP center will plan to have an annual meeting where HPC tool developers and all other stakeholders can meet to highlight developments and new requirements for their ecosystems. This annual meeting can include different types of events such as tutorials, invited talks, and posters sessions. Having such an in-person meeting will help to create new collaborations as well as dissemination of any success stories to the wider community. It will also be beneficial to have some portions of this meeting in a hybrid/virtual setup to increase the accessibility for the newcomers to the community.

Another good example to follow for community activities is having interest group meetings with more frequent cadence. This kind of meeting can easily be structured within the efforts of other initiative (i.e., Collaborative teams, co-design) as well as thorough construction of interest groups for various topics (i.e., performance tools, correctness tools, etc.) where people can meet in-person and/or virtual meetings for discussing gaps and directions for the future. While this might be a more internal

outreach, STEP center also intent to reach-out to a wider audience by proposing tutorials, workshops in other venues such as HPC conferences (i.e., Supercomputing, IPDPS, ... etc.), possible attendance to more application specific meetings and conferences (i.e., ACS, APS, DOE Meetings/Workshops) for disseminating the need for HPC tools, or having seminar series in national laboratories and universities that will allow to extend the users space of these tools as well as possible recruitment as the future workforce for the community.

Other Outreach Activities

The breakout room discussions also focused on the need to have a listing of HPC tools that can provide a source for HPC users/developers. This can also be a STEP certification process that will allow tools to highlight their efforts to make their tools user friendly. This certification process can have different criteria for acceptance and might include tags/badges. Various criteria are discussed throughout the break room discussion the ones that come across are: (1) having extensive user documentation, (2) easy to build on facilities, (3) supports continuous integration and (4) allow interfacing with users code-base. By providing a transparent process that is open to the whole community can make this whole process smooth and crucial for dissemination of the new tools similar to artifact evaluation in conference series (i.e., SC, PLDI etc.).

One other discussion point from the breakout sessions is to have outreach for having internship programs with the national laboratories and the university where different levels of students can interact with experts from the tools community as well as the application community. This will allow both to create new collaborations as well as educating the new workforce. Depending on the policies of the national laboratories, these kinds of efforts can be made virtual by daily meetings for a short period or summer long internships. Another point is having an active communication channel that can allow dynamic remote interactions. Promoting STEP center Slack channel will be useful to allow community level interactions.

Breakout 3.3: Interfacing With Additional Funding Sources

The Interfacing With Additional Funding Sources breakout session took place on Wednesday, July 12th. The session lead was James Custer (HPE). Additional attendees were Phil Carns (ANL), Chunhua Liao (LLNL), and Dave Montoya (Trenza).

The main topic of this breakout session primarily focused on conceptualizing a strategy for interfacing with additional external funding sources to support the STEP project. Anchoring this discussion was

the belief that third-party collaborations could be cultivated with various entities, including the Department of Energy (DOE) facilities, NIH, NASA, vendors, and more.

One significant question that arose concerned the incentives for these organizations to collaborate with STEP. It was theorized that playing a role in solving interoperability issues, contributing to tool enhancements, and expanding research horizon could be beneficial for them.

We explored the potential value STEP could provide to third-party entities. We concluded that elements such as STEP's expert knowledge and base of best practices were considered as critical assets that could attract third-party funding through sponsorships to STEP-led events and initiatives. Additionally, STEP could play a role as a comprehensive platform that could help in the negotiation of contracts, facilitate installs of new machines, and bridge gaps that vendors may avoid.

The feasibility of leveraging existing funding initiatives, such as the DOE Small Business Innovation Research (SBIR) program and Google's Summer of Code, was evaluated. We believe that these avenues can contribute to areas such as internships, workforce development, and diversity, equity, and inclusion (DEI).

We also explored whether STEP could draw funding from the traditional Applications side or whether this would be too challenging. We hypothesize that an already funded applications group would likely not see value in spending money from their already allocated budget, but if STEP were to be involved at the proposal phase, there could be opportunities to include money budgeted for tools to contribute to STEP.

We discussed project scope and structure of deliverables. The need for space to customize collaborations was juxtaposed against the importance of having a consistent way of engaging third-parties. The group largely agreed that a uniform standard for reporting and deliverables is important. The group's recommendation is that traditional quarterly reporting should be replaced by quarterly demos that could be run and are reproducible.

Two crucial unaddressed questions remain unanswered: defining STEP's balance between research and production development and clarifying restrictions on commercial partners collaborating with external entities like SBIR.

The session underscored the complexities involved in effectively collaborating with various funding sources and the importance of maintaining communication and stakeholder engagement throughout the process. It emphasized the need for adaptable strategies and willingness to reassess and refine

these strategies based on their real-world implementations. This adaptability will be critical in shaping successful and sustainable collaborations for the STEP project.

Breakout 3.4: Facilitating Ongoing Collaboration: What entices Application teams / Facilities to work with STEP Efforts

The Facilitating Ongoing Collaboration breakout session took place on Wednesday, July 12th. The co-moderators were Jose Moreira (IBM Research) and Terry Jones (ORNL). Additional attendees were Adam Lavelly (LBNL), John Linford (NVIDIA), Vitali Morozov (ANL), Paul Kent (ORNL), and Richard Gerber (LBL). The overarching goal of this breakout session was to find mechanisms that will foster collaboration between application and tool teams. We have broken down this goal into four separate components: (1) making tools attractive to application developers; (2) successful codesign engagement for sustainability; (3) facilitating collaborative work; and (4) collecting broad feedback from stakeholders.

Making tools attractive to application developers

Successful codesign engagement for sustainability

Facilitating collaborative work

Collecting broad feedback from stakeholders

It is important for STEP Center to continuously collect information and requirements from a broad set of stakeholders. This feedback has to come from the HPC community at large and should not be restricted to DOE labs only. We expect the feedback to come from a combination of surveys and direct engagements with other groups.

For surveys to be effective they must reach out to the broad community and also support useful feedback. Surveys should include free form suggestion boxes so that those completing the surveys can provide as much feedback content as possible.

When engaging with other groups, STEP should look for contacts in Europe (e.g., POP2 and European supercomputing centers) and Japan (e.g., Riken). STEP should also work with NSF centers (e.g., TACC) and interact with the PEARC conference series community. A Birds-of-a-Feather event at PEARC 2024 would be a good way to engage them.

A good way to motivate broad feedback is to report on success stories as soon as possible, ideally in the first year of the Center. The success stories can be about helping applications or facilities to accomplish some goal of importance to DOE. The success does not have to be on massive scale

marquee applications. Success in what would be considered common cases (*everyday* cases) will interest a lot of people.

Attendees also noted that helping new users start using systems more effectively will be increasingly important for sustainability. HPC resources are being used by a growing diversity of application teams, but there is a large gap between the tool proficiency of new users and the tool proficiency of expert users.

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 Midwest Town Hall participants collaboratively explored the current HPC tools space and developed solutions to the sustainability challenges (see [agenda](#)). Capturing the discussion, recommendations and key findings of the Midwest Town Hall held July 11-12th, 2023 at Iowa State University, is a priority outcome from the Midwest Town Hall. When complete, the availability of the Midwest Town Hall Report will be announced via applicable mailing-lists and placed on the STEP website under resources.

3.2 Refine strategies for understanding current state of tools ecosystem

One of the thrusts of STEP is to understand the current state of the Tools Ecosystem as of the Summer of 2023. This includes which tools are in use, which applications use which tools, the extent that the application's success depends on the tool, and how those tools are currently funded. Of particular concern are tools that have a significant dependence on ECP funding as such work will lose that support at ECP's completion on Dec-31-2023.

As of the writing of this report, STEP efforts in this area have relied on two online surveys found on the STEP website ([here](#) and [here](#)), and on personal canvassing of tool developers. While effective, it may be possible to extend these two strategies and STEP will be evaluating new opportunities to capture this information.

3.3 Design and organize West Coast Town Hall

The third and final STEP town hall will be held on August 16-17 at the San Francisco Marriott Fisherman's Wharf in San Francisco, CA. This town hall will shift discussion focus to (a) prioritizing and refining strategies for key sustainability challenges in tools; (b) key components needed for the phase II proposal; and (c) additional challenges including governance, workforce and organizational challenges. As in the first two town halls, it will assemble a collection of experts from the tools community including not only tools developers but also application, facility, and vendor stakeholders to develop actionable collaborative solutions to the challenge of sustainability in the tools ecosystem.

3.4 Coordination Activities with other Software Sustainability seedlings

The STEP Software Sustainability seedling is one of 6 such efforts. Information on the other five Software Sustainability seedlings can be found at their websites:

- [COLABS](#)
- [OSSF](#)
- [PESO](#)
- [S4PST](#)
- [SWAS](#)

The STEP leadership team will continue to communicate with the leaders of these projects to identify collaborative opportunities as we continue to refine our own plan. We will also organize a breakout session at the West Coast Town Hall with the explicit purpose of more carefully exploring potential connection points with not only our peer software sustainability seed projects but also the broader HPC ecosystem.

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 Ex
10. ascale. 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>.
11. 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>
12. NVIDIA Corporation. CUDA Profiler - Release 12.1. April 14, 2023. https://docs.nvidia.com/cuda/pdf/CUDA_Profiler_Users_Guide.pdf
13. IBM Corporation. Power ISA Version 2.07B. April 9, 2015.
14. 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
15. 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>.

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
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
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 edits thereafter.

First Name	Last Name	Affiliation
Ryan	Adamson	Oak Ridge National Laboratory
Phil	Carns	Argonne National Laboratory
James	Custer	Hewlett Packard Enterprise
Heike	Jagode	University of Tennessee
Mike	Jantz	University of Tennessee
Terry	Jones	Oak Ridge National Laboratory
John	Linford	NVIDIA
John	Mellor-Crummey	Rice University
Bart	Miller	University of Wisconsin
José	Moreira	IBM
Erdal	Mutlu	Pacific Northwest National Laboratory
Bruce	Palmer	Pacific Northwest National Laboratory
Shane	Snyder	Argonne National Laboratory
Theresa	Windus	Ames National Laboratory

Appendix 1.3 Workshop Attendees

The following list of people physically attended the East 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
Adam	Lavelly	Lawrence Berkeley National Lab
Ali	Jannesari	Iowa State University
Amadeo	Perazzo	Stanford Linear Accelerator
Andreas	Kronfeld	Fermi National Laboratory
Benoir	Forget	Massachusetts Institute of Technology
Charles	Zaruba	Iowa State University
Bruce	Palmer	Pacific Northwest National Lab
Chunhua “Leo”	Liao	Lawrence Livermore National Lab
David	Montoya	Trenza
Erdal	Mutlu	Pacific Northwest National Lab
Heike	Jagode	University of Tennessee
JaeHyuk	Kwack	Argonne National Laboratory
James	Custer	Hewlett Packard Enterprise
John	Linford	NVIDIA
Jonathan	Graham	Los Alamos National Laboratory
Jordi	Alcaraz	Iowa State University
Jose	Moreira	IBM
Keita	Teranishi	Oak Ridge National Laboratory
Kevin	Huck	University of Oregon
Mike	Jantz	University of Tennessee

Nathan	Tallent	Pacific Northwest National Laboratory
Oscar	Hernandez	Oak Ridge National Laboratory
Paul	Kent	Oak Ridge National Laboratory
Peng	Zu	Iowa State University
Peter	McCorquodale	Lawrence Berkeley National Laboratory
Phil	Carns	Argonne National Laboratory
Richard	Gerber	Lawrence Berkeley National Lab
Ryan	Adamson	Oak Ridge National Laboratory
Sarom	Leang	Iowa State University
Shane	Snyder	Argonne National Laboratory
Suren	Byna	Ohio State University
Terry	Jones	Oak Ridge National Laboratory
Theresa	Windus	Ames National Laboratory
Vitali	Morozov	Argonne National Laboratory
Wei-Keng	Liao	Northwestern University
William	Godoy	Oak Ridge National Laboratory
Wyatt	Spear	University of Oregon

Appendix 2: Agenda

Day 1 (July 11th, 2023)

Time	Topic – Items marked with blue font will be available online via Zoom for remote observers
8:30 – 9:00	Registration & Badging - Howe Hall Atrium
9:00 – 10:00	What is STEP and Why Are We Here? (Terry Jones, ORNL) – Howe Hall Auditorium or zoom link This talk describes STEP, Introduces the need for Initiatives, & introduces the need for a Survey
10:00 - 10:30	Break – Howe Hall Atrium
10:30 - 12:00	Plenary: Recap of East Coast Town Hall – Description of the 4 STEP Initiatives – Howe Hall Auditorium or zoom link <ul style="list-style-type: none">• Recap of East Coast Town Hall (Heike Jagode, University of Tennessee)• Proposed Structure (Phil Carns, Argonne)<ul style="list-style-type: none">• Collaborative Teams• Rapid Response• Codesign• Community Activities
12:00 - 1:00	Working Lunch (Provided) – Discussions on morning sessions – Howe Hall Atrium
1:00 - 1:30	Guidance, Logistics and Desired Outcomes for Breakouts (Terry Jones, ORNL) – Howe Hall Auditorium
1:30 - 3:00	Breakouts Session 1: Structure & Setting Direction <ol style="list-style-type: none">1. Management/governance structure (session lead: Terry Jones, ORNL) – Pearson 21252. Stakeholders & External Advisory Boards (session lead: Phil Carns, Argonne) – Pearson 21313. Processes for Tracking Vision & Gaps Over Time (lead: Heike Jagode, Univ of Tenn) – Pearson 2137
3:00 - 3:30	Break – Pearson 2143
3:30 - 5:00	Breakouts Session 2: Project Lifecycle <ol style="list-style-type: none">1. Bootstrapping Efforts (session lead: Theresa Windus, Ames National Lab) – Pearson 21252. Reporting & Maintaining Efforts (session lead: Jose Moreira, IBM) – Pearson 21313. Reviewing & Sunsetting Efforts (session lead: Shane Snyder, Argonne) – Pearson 2137
5:00 - 5:30	Closing and Guidance for Day 2 – Howe Hall Auditorium

Day 2 (July 12th, 2023)

Time	Topic – Items marked with blue font will be available online via Webex for remote observers
9:00 - 9:10	Introduction and Day 2 Logistics (Terry Jones, ORNL) – Howe Hall Auditorium or zoom link
9:10 - 10:15	<p>Report out summaries from Session 1: Structure & Setting Direction – Howe Hall Auditorium</p> <ol style="list-style-type: none"> 1. Management/governance structure (session lead: Terry Jones, ORNL) 2. Stakeholders & External Advisory Boards (session lead: Phil Carns, Argonne) 3. Processes for Tracking Vision & Gaps Over Time (session lead: Heike Jagode, UTK)
10:15 - 10:45	Break – Howe Hall Atrium
10:45 - noon	<p>Report out summaries from Session 2: Software Lifecycle – Howe Hall Auditorium or zoom link</p> <ol style="list-style-type: none"> 1. Bootstrapping Efforts (session lead: Theresa Windus, Ames National Lab) 2. Reporting & Maintaining Efforts (session lead: Jose Moreira, IBM) 3. Reviewing & Sunsetting Efforts (session lead: Shane Snyder, Argonne)
noon - 1:00	Working Lunch (Provided) – Discussions on morning sessions – Howe Hall Atrium
1:00 - 1:45	<p>Primer: Governance, Workforce and Organizational Challenges (Moderator: James Custer) – Howe Hall Auditorium or zoom link</p> <ul style="list-style-type: none"> ● Perspectives on DEI and Workforce Development Mike Jantz (University of Tennessee) ● Community Activity models & scope Phil Carns (Argonne) ● Interfacing With Additional Funding Sources Terry Jones (ORNL) ● Facilitating Ongoing Collaboration: What entices Application teams / Facilities to work with STEP Efforts Theresa Windus (Ames National Lab / ISU)
1:45 - 3:00	<p>Breakouts Session 3: Additional Challenges: Governance, Workforce and Organizational Challenges</p> <ol style="list-style-type: none"> 1. DEI & Workforce Development (lead: Mike Jantz) VIRTUAL BREAKOUT zoom link 2. Community Activity models & scope (session lead: Erdal Mutlu, PNNL) – Pearson 2125 3. Interfacing With Additional Funding Sources (lead: James Custer, HPE) – Pearson 2131 4. Facilitating Ongoing Collaboration: What entices Application teams / Facilities to work with STEP Efforts (session lead: Jose Moreira, IBM) – Pearson 2137
3:00 - 3:30	Break
3:30 - 4:50	<p>Report out summaries for Session 3: Additional Challenges: Governance, Workforce and Organizational Challenges - Howe Hall Auditorium or zoom link</p> <ol style="list-style-type: none"> 1. DEI & Workforce Development (session lead: Mike Jantz) VIRTUAL PRESENTATION 2. Community Activity models & scope (session lead: Erdal Mutlu, PNNL) 3. Interfacing With Additional Funding Sources (session lead: James Custer, HPE) 4. Facilitating Ongoing Collaboration: What entices Application teams / Facilities to work with STEP Efforts (session lead: Jose Moreira, IBM)
4:50 - 5:00	Closing Remarks / Adjourn (Terry Jones, ORNL) - Howe Hall Auditorium

Appendix 3: 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.