An Open Integrated Software Stack for Extreme Scale Computing

Executive Summary

The Department of Energy is preparing to launch an Exascale Initiative (EI), a plan to bring exaFLOPS computing systems online in the 2018 time frame to address critical science questions in its mission areas, e.g., combustion, climate, nuclear energy, renewable energy, and national security. The Initiative is expected to consist of four major components: 1) two platform development partnerships, 2) multiple application-oriented co-design centers, 3) a program to accelerate development of critical technology, and 4) an effort to develop and support the exascale software stack.

Application codes depend on an underlying software stack that is deployed at today's largest HPC Facilities. It is a mix of code from system vendors, independent software vendors (ISVs), high-performance computing (HPC) computer science and math researchers, and the larger open-source software community. However the research and development of these software components is ad-hoc, with no coordination of features or progress tracking across the whole software stack to insure key application-driven features will perform well with new architectures. ***This method of software development injects a great deal of risk into the successful deployment of exascale architectures and increases substantially the effort and elapsed time required to implement applications.*** Without careful design, coordination, integration, and testing for novel software to match the new architectural features and scale of exascale platforms, application codes will not perform adequately to meet their scientific objectives.

This document provides a high-level overview of the planned Exascale Software Center (ESC), a key part of the Exascale Initiative. Carefully applying specific management principles, ESC provides a coordinated and integrated approach to develop the required exascale system software, from the operating system to the math libraries and tools. This plan was developed by experts in the community, including representatives from DOE laboratories, NSF centers, HPC computer scientists and mathematicians at leading universities, and industrial partners. The plan benefits from best practices and many years of community discussion, workshops, and reports on the topic, including Ken Kennedy's effort to develop a Petascale Software Project in 2006. The goal of ESC is straightforward and achievable: **Ensure the timely delivery of a complete, production-quality system software stack for computational science on exascale platforms**.

The **Exascale Software Center** will be involved in the entire software lifecycle:

- Identify required software capabilities in collaboration with application co-design centers and exascale platform partnerships
- Identify gaps in current software capabilities within open source and vendor software packages
- Design and develop open-source software components to address gaps and provide unified, integrated software for exascale platforms
- Test software functionality, stability, and performance in partnership with co-design center applications and exascale platform partnerships to assure quality
- Collaborate with platform vendors to integrate ESC software, streamline system bring-up and acceptance, and partner to provide support
- Apply principles of project management for carefully tracking development progress and milestones
- Coordinate outreach to the broader open source, other government agencies, and the collaboration efforts of the International Exascale Software Project (IESP) to improve planning and collaboration

Software Development within the Exascale Software Center

Considering the exascale software stack as a whole, ESC will be involved in the entire software lifecycle, from working to identify the needs and the gaps in existing software to supporting the deployment and validation of the software on exascale systems. Figure 1 illustrates the tightly coupled nature of co-design and deployment of basic system software.

Co-design requires carefully managed feedback and evaluation processes centered around two key iterative processes. First, the application teams, hardware architects, and software developers must come together to identify the basic requirements of the system software for a given architectural design and determine the key software gaps that must be addressed. The applied research required to create novel solutions must also be started. The process is iterative, with the hardware architects and application co-design centers working together to refine approaches. Second, as prototype solutions are built and show promise, the software matures, with testing and integration becoming a more significant part of the work as well as performance against clearly defined metrics for performance, functionality, scalability, and stability.

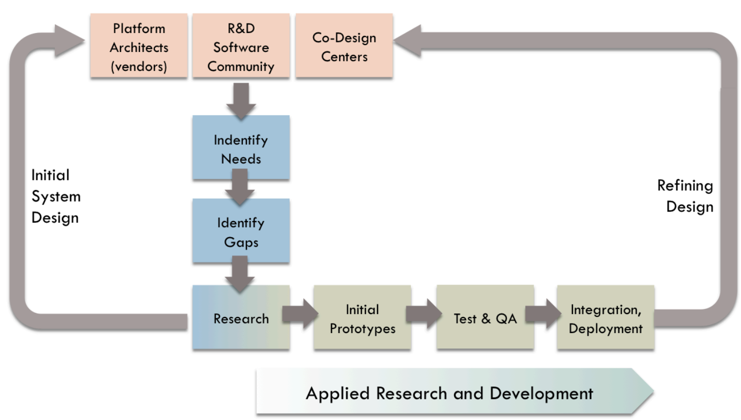

Figure 1: The Exascale Software Development Process

Overall, the Center can be thought of as managing a decentralized software development and integration effort, including stepping in to adjust staffing, priorities, or capabilities to ensure success. Development will be focused on five key areas: programming models, operating systems and run-time systems, application programming tools, numerical libraries and frameworks, data management and analysis, and system management and cybersecurity.

Contributors

Pete Beckman, Jack Dongarra, Jim Ahrens, Pavan Balaji, Ron Brightwell, Barbara Chapman, James Demmel, Al Geist, Bill Gropp, Gary Grider, Paul Hargrove, Michael Heroux, Kamil Iskra, Rusty Lusk, Allen Malony, Arthur Barney Maccabe, Ron Minnich, Pat McCormick, Terry Moore, John Mellor-Crummey, Dan Quinlan, Terri Quinn, Robert Ross, Martin Schulz, Marc Snir, Lee Ward, Bronis de Supinski, Rajeev Thakur, Vinod Tipparajuv, Jeff Vetter, Kathy Yelick, et. al.