LA-UR-23-25337

Approved for public release; distribution is unlimited.

Title:	Final Review of FY23 ASC ATDM L2 Milestone, MRT #8639, SPARC and EMPIRE Readiness for El Capitan
Author(s):	Fung, Jimmy; Francois, Marianne M.; Mosby, Matthew; Castro, Joseph; Draeger, Erik; Heroux, Michael; Mohror, Kathryn; Rockefeller, Gabriel M.; Miller, Douglas; Brunner, Thomas
Intended for:	milestone review memo
Issued:	2023-05-17

National Nuclear Security Administration

Memorandum

date: 5/12/2023

to: Jim Stewart, Rob Hoekstra, David Littlewood

cc: Thuc Hoang, James Peltz, Simon Hammond, Jennifer Gaudioso

from: Jimmy Fung (LANL); Marianne Francois (LANL); Matthew Mosby (SNL), Joseph Castro (SNL), Erik Draeger (LLNL), Michael Heroux (SNL), Kathryn Mohror (LLNL), Gabriel Rockefeller (LANL), Douglas Miller (LLNL), and Thomas Brunner (LLNL)

subject: Final Review of FY23 ASC ATDM L2 Milestone, MRT #8639, SPARC and EMPIRE Readiness for El Capitan

Summary:

- The milestone review meeting was held on May 8-9, 2023 at Sandia National Laboratories, Building 701, Room 2001.

- The ASC-ATDM (Advanced Simulation and Computing – Advanced Technology Development and Mitigation) Subprogram hosted an FY23 Level 2 Milestone Review with ASC-ATDM (MRT #8639), which was held as a classified meeting with representatives from Sandia National Laboratories, Lawrence Livermore National Laboratory, and Los Alamos National Laboratory. Upon hearing the briefing, this committee agreed that **the milestone team has successfully completed and exceeded the milestone requirements**.

Milestone Overview

As displayed in the FY23 ASC Implementation Plan (IP) and the Milestone Reporting Tool (MRT), the milestone description and completion criteria state:

ASC ATDM L2 Milestone, MRT #8639, SPARC and EMPIRE Readiness for El Capitan

Description: The goal of this milestone is to demonstrate readiness of the Sandia Parallel Aerodynamics Reentry Code (SPARC) and the Electro-Magnetic Plasma in Realistic Environments (EMPIRE) code for use on the forthcoming El Capitan exascale platform, which will enable simulation of weapons-relevant problems at unprecedented scale and fidelity.

- A primary challenge is porting SPARC and EMPIRE to AMD CPU/GPU hardware and the corresponding software stack. This will be achieved using performance-portable abstractions.
- The milestone will utilize early-access systems on the LLNL RZ or SCF networks.

Early identification of gaps and development of solution strategies is critical in preparation for El Capitan going into classified production in CY 2024.

Completion Criteria:

1. Identify porting and performance gaps for SPARC and EMPIRE in preparation for the El Capitan system, including the effectiveness of performance-portable abstractions.
2. Assess computational performance and scaling of SPARC on up to 50% of LLNL RZVernal.
3. Implement a strategy for removal of Unified Virtual Memory (UVM) dependencies from EMPIRE and document progress.

Exceeds criteria:

1. Assess computational performance and scaling of EMPIRE on up to 50% of LLNL RZNevada using a simulation of a B-Dot experiment.
2. Assess computational performance and scaling of SPARC on over 50% of LLNL RZVernal.
3. Resolve performance and scaling barriers for SPARC and EMPIRE on AMD CPU/GPU hardware and the corresponding software stack.

The committee found that the team exceeded the milestone requirements for both SPARC and EMPIRE. Along the way to presenting their milestone completion, the SNL team also noted that continuing work on SPARC focuses on solvers and solver performance. For EMPIRE, the use of Relocatable Device Code uncovered clear link-time performance issues in the AMD toolchain.

Scope of Review: For all three laboratory milestone teams, presentations were delivered at the review meeting with associated documents provided prior to the meeting. The presentations were made in three sessions in the following order: SNL, LANL, LLNL. The review committee deliberated after each session and after final deliberations, the committee made its scoring and assessment and presented its findings in an outbrief at the end of the review meeting.

Review Meetings: The overall milestone review schedule was carried out as follows:

- Description of kickoff [11/04/2022]. Introductions were made as well as roles and responsibilities established for milestone team and panel members.
- Mid-year review [12/06/2022]. This meeting included a review of goals and objectives, milestone team presentations, and subsequent feedback from the committee to the teams.
- Final review [05/08-09/2023]. This meeting is formally documented in this memo.

Observations from the committee:

- **General tri-lab comments**
  - Great progress was made towards El Capitan and the outlook is encouraging for production use on El Capitan.
  - The committee appreciates the capture of lessons-learned and collaborative communication across teams and laboratories.
  - Cross-laboratory tools including Mattermost and the Remote Computing Enablement (RCE) are important to effective collaboration and sharing of lessons/experiences and should be encouraged at each laboratory.
  - The El Capitan Center of Excellence (CoE) has done an exemplary job engaging all three labs and vendors.
  - Changes in vendors come with unpredictable costs that are difficult to address and manage and should be recognized in any software porting endeavor. First, the compiler/environment software stack is still very immature and continues to present challenges to software development. Overcoming those challenges has a significant and negative impact on developer productivity.
  - The adoption of software abstractions and components such as RAJA and Kokkos is paying off through reduced costs and effort for porting software.
  - Continued investment in scalable solvers is needed and encouraged, as evidenced by the types of challenges exposed in numerical methods and algorithm development for these simulation capabilities.
  - Software infrastructure and software development/engineering tools have played a significant role in team success.

- **Comments for SNL SPARC**
  - The milestone results clearly demonstrate that Kokkos is an effective portability abstraction layer across GPU architectures.

  - That it took less than 0.1 FTE to compile on EAS-3 is remarkable.
  - Despite memory differences in EAS-3/ATS-4, no code changes are expected to port to El Capitan.
  - Good progress was made between the midterm and final reviews in terms of computational performance, achieving 2X performance improvements from FY23Q1 results.
  - Good performance on EAS-3 was observed compared to ATS-2, achieving 2-4x improvements.
  - The team is commended for its use of Remote Computing Enablement project cross-site testing for Software Quality Assurance activities on EAS-3.
  - The team is commended for gathering and sharing technical details and challenges working on EAS-3.
  - The committee applauds the comparison of computational performance relative to expected hardware peaks.

- **Comments for SNL EMPIRE**
  - Progress on CUDA UVM removal since the mid-year review was impressive.
  - Without UVM, Kokkos enabled straightforward porting to another GPU architecture.
  - Transparent sharing of lessons learned and supplemental slides is valuable.
  - The agile use of cross-platform approaches was very effective in solving or working around issues on EAS-3.
  - The initial performance and scaling on EAS-3 is promising, and continued analysis and optimization is encouraged
  - Close collaboration of Trilinos and EMPIRE developers saw mutual benefits.
  - This effort has prepared Trilinos and other TPLs for use by other applications
  - The team is highly commended for discovering and reporting numerous issues to AMD.

Feedback from the Committee:

Here are recommendations to be considered for future efforts by the teams:

- **Recommendations for SNL SPARC**
  - None.

- **Recommendations for SNL EMPIRE**
  - None.

Conclusions: Following the review, this committee agreed that the milestone team has successfully completed and exceeded the milestone requirements.

Attendees of the final review:

SNL L2 Milestone Team

David Littlewood (lead)
Victor Brunini
Roger Pawlowski
Michael Wolf
Mark Bolstad

LANL L2 Milestone Team

Chris Malone (lead)
Scott Pakin
Nathan Vaughn-Kukura
Jonathan Graham

LLNL L2 Milestone Team

Rob Rieben
Thomas Stitt (lead)
Arturo Vargas
Alejandro Campos
Aaron Skinner
Kenny Weiss
Dave Richards

Review Panel

Joseph Castro, SNL
Matt Mosby, SNL
Gabriel Rockefeller, LANL
Jimmy Fung, LANL (chair)
Douglas Miller, LLNL
Thomas Brunner, LLNL

ECP Reviewers

Erik Draeger, LLNL (ECP AD)
Marianne Francois, LANL (ECP NNSA AD) (co-chair)
Michael Heroux, SNL (ECP ST)
Kathryn Mohror, LLNL (ECP NNSA ST)

ASC Program leadership in attendance

Thuc Hoang, NNSA, ASC program director
Simon Hammond, NNSA, ASC CSSE, ATDM
James Peltz, NNSA, ASC IC
Jennifer Gaudioso, SNL

Jim Stewart, SNL
Rob Hoekstra, SNL
Jason Pruet, LANL
Aimee Hungerford, LANL
Christopher Werner, LANL
Tim Randles, LANL
Rob Neely, LLNL
Teresa Bailey, LLNL

Copy: Thuc Hoang, NNSA, thuc.hoang@nnsa.doe.gov
Simon Hammond, NNSA, simon.hammond@nnsa.doe.gov
James Peltz, NNSA, james.peltz@nnsa.doe.gov
Emily Simpson, NNSA, emily.simpson@nnsa.doe.gov
Jennifer Gaudioso, SNL, jmgaudi@sandia.gov
Jim Stewart, SNL, jrstewa@sandia.gov
Rob Hoekstra, SNL, rjhoeks@sandia.gov
Joseph Castro, SNL, jpcastr@sandia.gov
Matt Mosby, SNL, mdmosby@sandia.gov
Gabriel Rockefeller, LANL, gaber@lanl.gov
Jimmy Fung, LANL, fung@lanl.gov
Douglas Miller, LLNL, miller18@llnl.gov
Thomas Brunner, LLNL, brunner6@llnl.gov
Erik Draeger, LLNL, draeger1@llnl.gov
Marianne Francois, LANL, mmfran@lanl.gov
Michael Heroux, SNL, maherou@sandia.gov
Kathryn Mohror, LLNL, mohror1@llnl.gov
David Littlewood, SNL, djlittl@sandia.gov