LA-UR-23-25448

Approved for public release; distribution is unlimited.

Title: Final Review of FY23 ASC ATDM L2 Milestone, MRT #8541, ATDM Multiphysics Scaling on EAS-3

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: 2023-05-19

Los Alamos
NATIONAL LABORATORY

UNCLASSIFIED

Memorandum

date: 5/12/2023

to: Chris Werner, Tim Randles, Chris Malone

cc: Thuc Hoang, James Peltz, Simon Hammond, Jason Pruet, Aimee Hungerford

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 #8541, ATDM Multiphysics
Scaling on EAS-3

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 #8541), 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 has been successfully completed**.

UNCLASSIFIED

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 #8541, ATDM Multiphysics Scaling on EAS-3

Description: Demonstrate the capability to execute mission-relevant calculations with a multiphysics code from the Ristra project at scale on the pre-El Capitan Early Access System 3 (EAS-3). These calculations must use the GPU hardware available on EAS-3, the AMD MI-250s, to achieve any reasonable performance.

Completion Criteria:
1. A 3d shape charge is run on at least 50% of rzvernal, the RZ component of EAS-3
2. A scaling study is performed and documented up to the scale achieved in item 1
3. A report is delivered outlining the Ristra project's experiences with the HPE/AMD software stack in anticipation of El Capitan

The committee found that the team met their milestone requirements. Along the way, the team ran a 3D shaped charge simulation in pure Lagrangian mode until the initial movement of the lining (e.g. early simulation times) on up to 16 nodes (50% of the Confluence published node count on RZVernal). In addition, Moya code performance on CTS-1 and EAS-3 was compared to FLAG code performance on CTS-1.

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 trilab comments**
  - Great progress was made towards El Capitan and the outlook is encouraging for production use on El Capitan.

o The committee appreciates the capture of lessons-learned and collaborative communication across teams and laboratories.
o 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.
o The El Capitan Center of Excellence (CoE) has done an exemplary job engaging all three labs and vendors.
o 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.
o The adoption of software abstractions and components such as RAJA and Kokkos is paying off through reduced costs and effort for porting software.
o 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.
o Software infrastructure and software development/engineering tools have played a significant role in team success.

- **Comments for Ristra**
  o Over the past 6+ years, the Ristra project has taken high risk approaches with many course corrections and is now far behind on the path toward production simulation on GPU architectures. The committee commends the movement to a new API (FleCSI 2.0) that is more maintainable. The committee also believes that the team is now on a promising path and that, following a goal of production capability on El Capitan, the team should continue with a narrower focus and less risk.
  o Legion has thus far has not been demonstrated to be a viable performance solution. In addition, the dependency on vendors for Legion support is high risk.
  o Moya has come a long way since its recent start (10/22), but is still immature in capability. An ALE scheme should be the top priority as a next step towards Moya as a production capability.
  o The complication of many backends is proving a challenge and focus is needed for success. More than one backend is needed to maintain and develop an abstraction, but depth-first for a single backend provides more value. Demonstrating value of Moya to users should be prioritized over broadening tasking models.
  o The committee was surprised that the Ristra team did not know the accurate size of RZVernal.

Feedback from the Committee:

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

- **Recommendations for LANL Ristra**
    - o Consistent with a goal of production capability on El Capitan, the Ristra project should focus more on meeting programmatic application needs. The committee recommends the development of a detailed plan or roadmap to guide this effort. A differentiating application needs to be found and selected where Moya can excel relative to current capabilities. In addition, the committee recommends deemphasizing alternative paths such as tasking backends and CartaBlanca.
    - o The committee recommends that the Ristra team meets with LANL line and program management to review their plan/roadmap and to agree on strategy and scope moving forward.

Conclusions: Following the review, this committee agreed that the milestone has been successfully completed.

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
     Jason Pruet, LANL, pruet@lanl.gov
     Aimee Hungerford, LANL, aimee@lanl.gov
     Christopher Werner, LANL, cwerner@lanl.gov
     Tim Randles, LANL, trandles@lanl.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
Chris Malone, LANL, cmalone@lanl.gov