

Exceptional service in the national interest



SAND2012-10519P

CSSE L2 Milestone

Data Co-Processing for Extreme Scale Analysis



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP

Multiple Ways for “In Situ” Vis



Embedded

- Direct access to memory.
- Potential no-copy memory sharing.
- Easier to schedule.

In Transit

- Lower (perceived) overhead for simulation.
- Asynchronous computation.
- Smaller vis job size.
- (Potential) process safety (requires better schedulers).

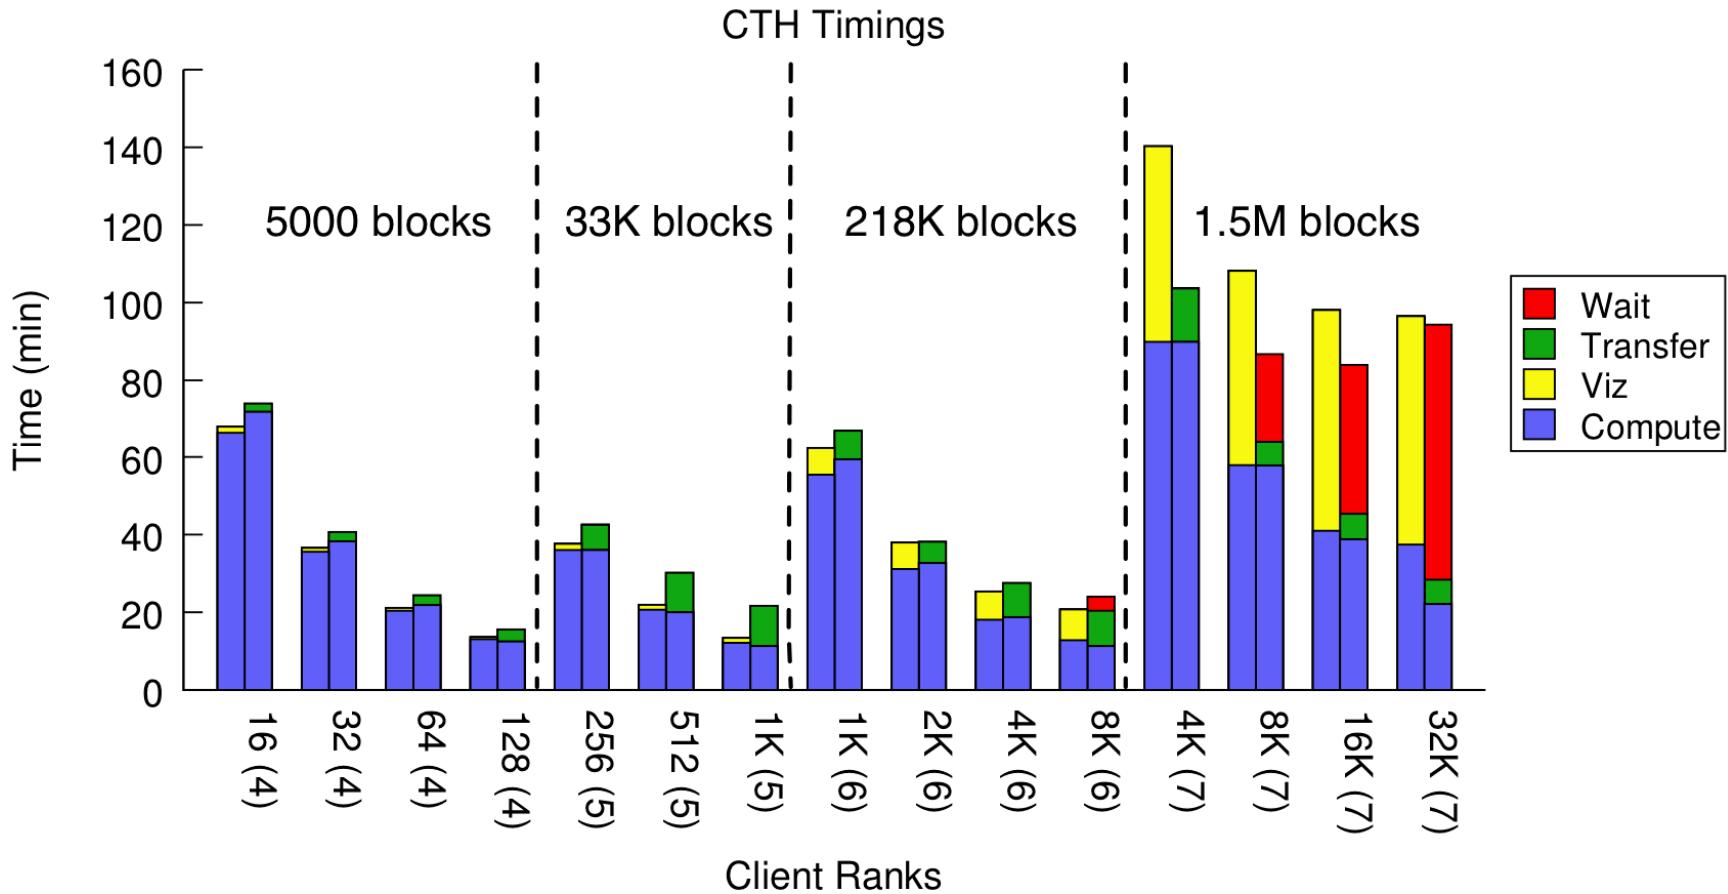
Intention

- Understand performance characteristics of using in situ vis.
 - Specifically, leveraging existing technologies with Sandia simulations.
- Perform metrics of performance.
- Compare embedded in situ with in transit techniques.
- Understand the tradeoffs with the two techniques.
 - What are the overheads?
 - What is the scaling behavior?
 - Which should be used when?

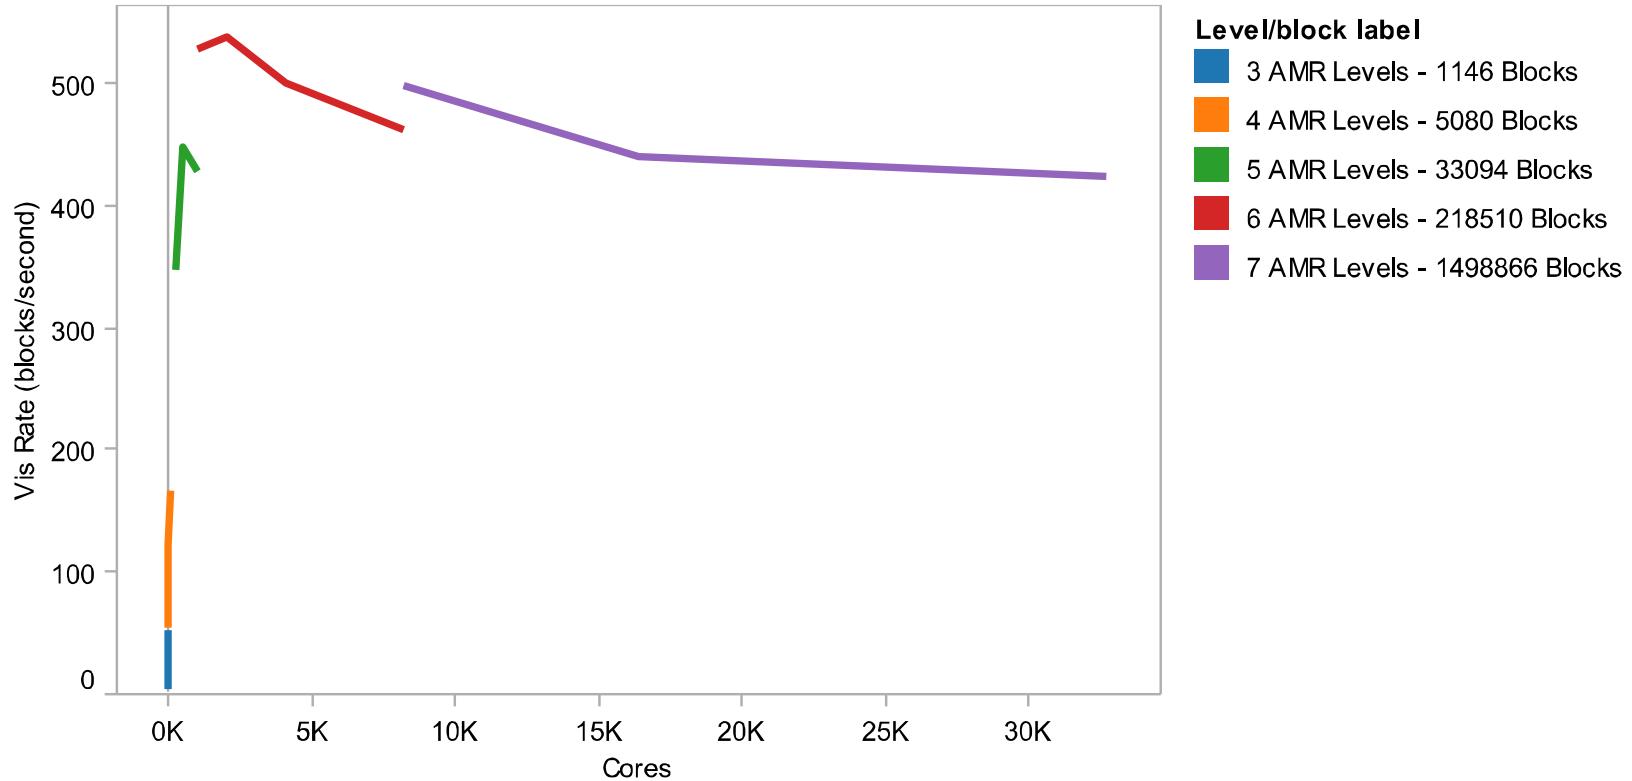
Experiment Setup

- Establish a “proxy” (not sensitive) CTH problem that can be easily scaled up.
 - In CTH, mesh refinement is controlled with a maximum depth of the AMR hierarchy. Given a sufficiently complex geometry, this results in a nonlinearly increasing mesh size.
- Establish in situ and in transit runs with equivalent processing.
 - Watertight contour of volume fraction.
- In Situ: CTH to Catalyst through Spyplot interface.
- In Transit: CTH to Nessie through Spyplot. Nessie transfers to separate job containing 1:16 cores. Nessie provides data to Catalyst through same Spyplot interface.
- Scale jobs to capability sizes.

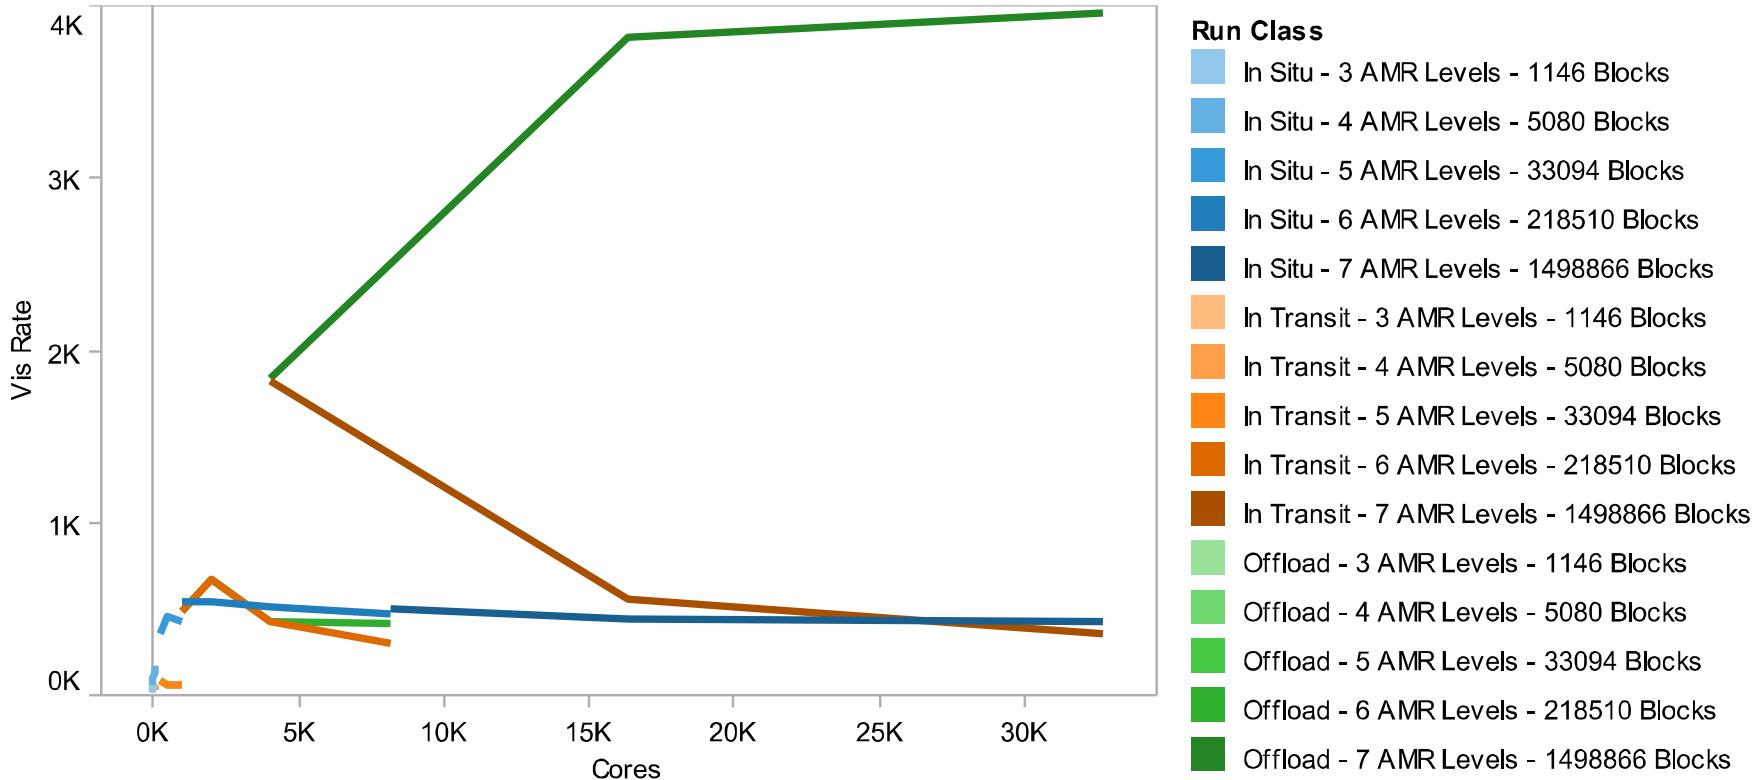
Summary of Data



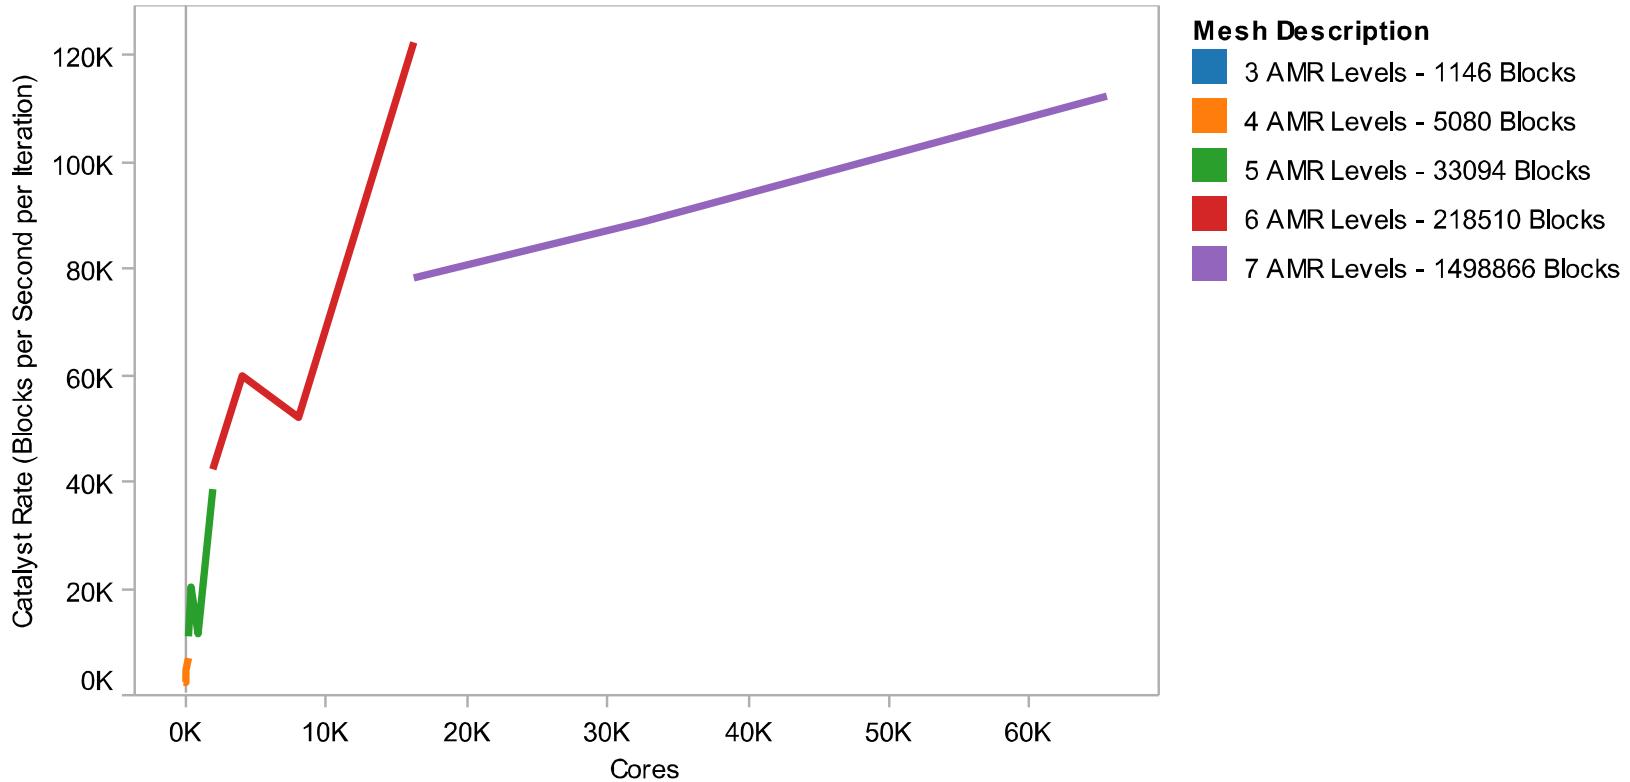
Scaling Behavior (Ron's Measures)



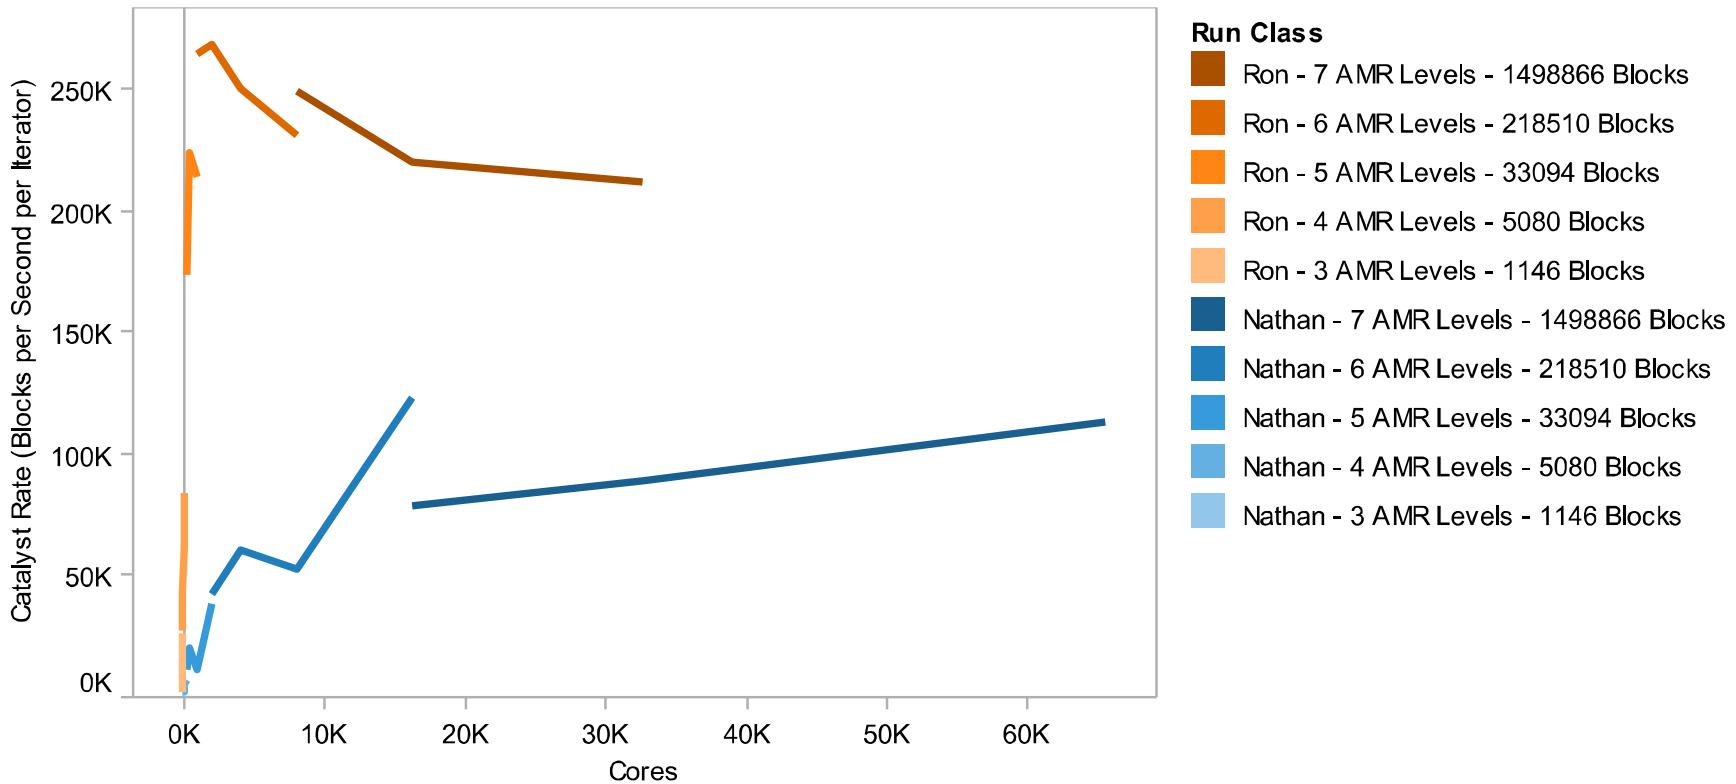
Scaling Behavior (Ron's Measures)



Scaling Behavior (Nathan's Measures)



Scaling Behavior (Measure Compare)



Official Milestone Description

ASC calculations produce complex datasets that are increasingly difficult to explore and understand using traditional post-processing workflows. To advance understanding of underlying physics, uncertainties, and results of ASC codes, SNL must gather as much relevant data as possible from large simulations. This drives SNL to couple data analysis and visualization capability with a running simulation, so that high fidelity data can be extracted and written to disk. This milestone evaluates two approaches for providing such a coupling:

- In-situ processing provides “tightly-coupled” analysis capabilities through libraries linked directly with the simulation. SNL has collaborated on developing an in-situ capability designed for this purpose.
- In-transit processing provides “loosely coupled” analysis capabilities by performing the analysis on separate processing resources. SNL provides this capability through a “data services” capability designed for this purpose.

Official Milestone Description

ASC calculations produce complex datasets that are increasingly difficult to explore and understand using traditional post-processing workflows. To advance understanding of underlying physics, uncertainties, and results of ASC codes, SNL must gather as much relevant data as possible from large simulations. This drives SNL to couple data analysis and visualization capability with a running simulation, so that high fidelity data can be extracted and written to disk. This milestone evaluates two approaches for providing such a coupling:

- In-situ processing provides “tightly-coupled” analysis capabilities through libraries linked directly with the simulation. SNL has collaborated on developing an in-situ capability designed for this purpose.
- In-transit processing provides “loosely coupled” analysis capabilities by performing the analysis on separate processing resources. SNL provides this capability through a “data services” capability designed for this purpose.

Catalyst

Nessie

Official Milestone Completion Criteria



SNL will engineer, test, and evaluate customer-driven data operations on large-scale data created by a running simulation. The data operations will be performed by instrumented versions of both the in-situ and in-transit solutions, with the resulting performance data published and made available to the ASC community.

Official Milestone Completion Criteria



SNL will engineer, test, and evaluate customer-driven operations on large-scale data created by a running simulation.

The data operations will be performed by instrumented versions

Compare performance of equivalent in situ and in transit implementations

CTH Fragment Extraction

Write a technical report

Feedback

- As it stands, have we completed the minimal requirements?
 - If not, what critical work remains to be done?
- Our intention for the next quarter is to scale higher and understand deeper the comparison.
 - What work would be most useful in this regard?