RAPIDS

Scientific Data Management and Understanding:
Accelerating Data Operations to achieve Science Insights

RAPIDS

Ana Gainaru, Rob Latham, Scott Klasky, Lenny Oliker, Robert Ross

Scientific Data Management

Data management remains a challenge as simulations continue to grow, evolve to tackle new domains, and adopt new technologies. Our goal is to optimize the I/O process for SciDAC applications in three directions:

Enable faster, more scalable I/O
across the diverse resources at the DOE and enable efficient I/O on different storage architectures

Effectively parse large amounts of data
extracting the relevant information by querying the primary data along with the quantities of interest for local and federated access

I/O monitoring and analysis capabilities
that facilitate a simple way to monitor I/O performance on LCFs and NERSC

Stability-preserving lossy compression for checkpoints

Checkpointing large PDE simulations stresses file systems, but lossy compression risks corrupting restarts. A stability-preserving compressor bounds physical error across restart cycles (relative L2 ≈ 2e-6) while accelerating checkpoint writes **on Frontier by up to 8.4× and reads by 6.3×.**

Gong et al., "Stability-preserving Lossy Compression for Large-scale Partial Differential Equations", SC '25: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (2025)

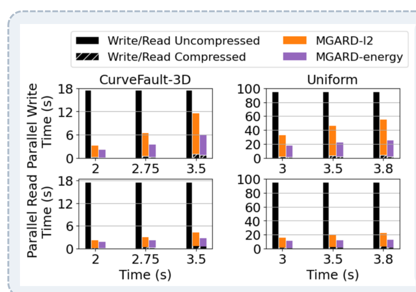

Parallel I/O with lossy checkpointing across 1024 nodes on Frontier at the OLCF

Parallel data object creation

- Simulation campaigns now create millions of data objects, making metadata handling — not bandwidth — the bottleneck.
- Partitioned metadata blocks in PnetCDF/HDF5 decouple per-object consistency checks, achieving **up to 196× faster parallel object creation** for 5.68M objects across 4,096 MPI processes.

Li et al., "Parallel Data Object Creation: Scalable Metadata Management in Parallel I/O Library", Proceedings of the SC '25 Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis (2025)

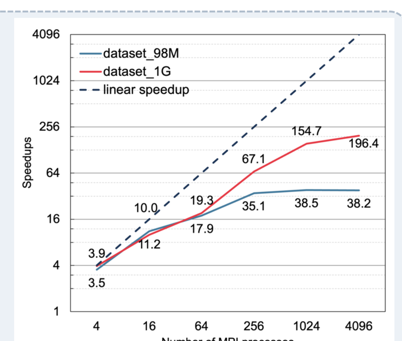

Speedups of the header format approach.

Data Orchestration

Data orchestration focuses on developing and enhancing workflow management systems to handle the growing complexity of tools that enable scientific understanding, providing more automation through traditional and AI techniques:

Monitor the services during execution
by establishing a multi-agent workflow with verified memory and planning modules and offering command and control capabilities

Novel ways to connect services
through data streaming and efficient near-real time querying and compression techniques

Integrate intelligent agents
at experimental facilities for adaptive experiment optimization

Command-and-control for in-situ workflows

- **Automated feedback, graceful checkpointing, dynamic reduction, and event-triggered analysis**, demonstrated across three use cases.
- The framework sequence relies on EFFIS and ADIOS to monitor tasks reading data to scan for pre-defined events, signal the command-and-control system when an event is detected and launche an external application at runtime.

Mehta et al., "Enabling Command-and-Control in Advanced In Situ Workflows," ICPP Workshops Proceedings of the 54th International Conference on Parallel Processing (2025), doi:10.1145/3750720.3757296

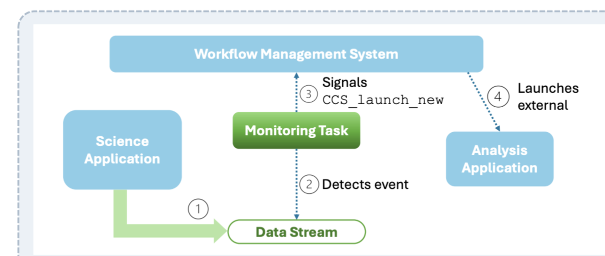

Sequence of interactions between components of the event-triggered workflow

I/O orchestration through Data-Size Based aggregation

- Data-Size Based (DSB) aggregation partitions ranks dynamically through a greedy algorithm that groups them into chains that write roughly equal data.
- DBS shows ~3-4x improvement in I/O time when extracts are severely imbalanced (Cartiso miniapp, 1024–8192 MPI ranks, at low (≥90% of ranks write) and severe (only 6–9% write) imbalance) and is on par with the best existing scheme when data is balanced

Ross, Wittenburg, Eisenhauer & Wetterer-Nelson, A Data-Size Adaptive Approach to I/O of Poorly Load Balanced In Situ Data Extracts, SC '25 Workshops (2025)

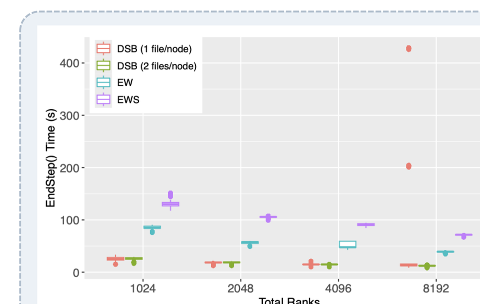

Performance for the severe load imbalance workload

Agentic AI Orchestration for Nuclear Fuel Qualification

- Agents receive a plain-language task, determine tool calls autonomously, and save successful runs as reusable skills for new fuel systems.
- Workflow management systems are used to run DFT calculations for the materials not existing in the campaign **reproducing studies that took months in hours.**

D.A. Lopes, A, Gainaru et al., "Agentic AI for accelerated fuel qualification framework: Demonstrating workflows with the AERIS enthalpy reasoning model," Journal of Nuclear Materials 631 (2026)

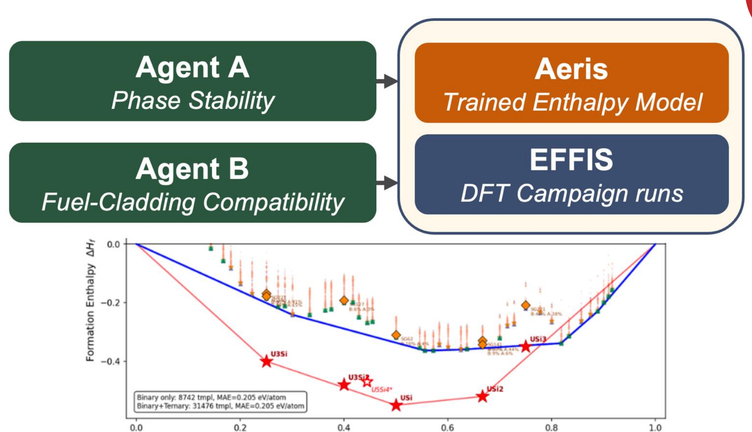

Agent A recovered the full U–Si convex hull, including U_3Si_2, in ~1.5 hour versus weeks to months for the equivalent evolutionary DFT search (bottom)

Agent B identified Mo and W as the most stable UN claddings and V, Ta as the least, consistent with experimentally reported nitride-layer formation.

Accelerate all the ways applications use data

Through efficient storage access, wrangling workflows, and insightful analysis of data we keep science moving.

AI threaded through everything

ChatVis: an LLM agent that writes visualization scripts

- Writing ParaView scripts is a barrier for many scientists and LLMs still struggle with highly specialized programming tasks such as scientific visualization.
- ChatVis pairs a large-language-model agent with retrieval and iterative correction, reaching a 100% compilation rate on a 100-task benchmark while running up to **6× faster** and **6× cheaper** than baseline language models.

Peterka et al., "ChatVis: LLM Agent for Generating Scientific Visualizations," IEEE Symposium on Large Data Analysis and Visualization (2025)

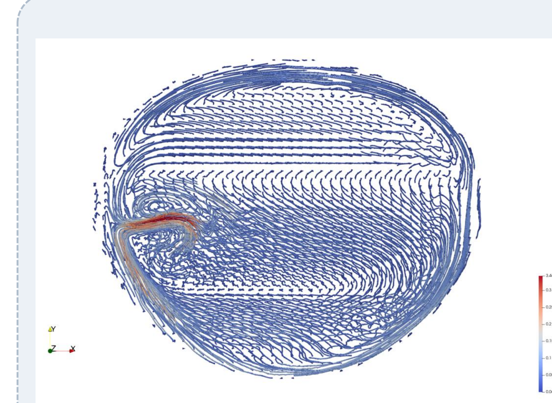

ChatVis-generated visualization of pathlines traced from a vector field simulated by an ocean climate model.

Scientific Visualization and Understanding

As simulations generate increasingly large datasets, scientists need more advanced visualization technologies to explore them across the diverse range of hardware at different facilities. Our goal is to drive improved understanding of data in three ways:

Enabling high-performance viz
through scalable algorithms across a diversity of architectures

In-situ visualization
to create new services for production visualization systems, integrated into data orchestration workflow technologies

Production visualization
to integrate campaign-based data management solutions into visualization software

Uncertainty-aware visualization at interactive speed

Quantifying uncertainty in surfaces extracted from simulations was too costly with Monte Carlo sampling. MAGIC's closed-form **formulation accelerates isosurface uncertainty by up to 585x** — and a companion vector-field method reaches ≈19,698× on GPUs — making uncertainty-aware analysis practical for meteorology, ocean, urban-flow, and astrophysics data.

Athawale et al., "MAGIC: Marching Cubes Isosurface Uncertainty Visualization for Gaussian Uncertain Data With Spatial Correlation," IEEE Transactions on Visualization and Computer Graphics(2026), Ouermi et al., "Efficient Probabilistic Visualization" (2025)

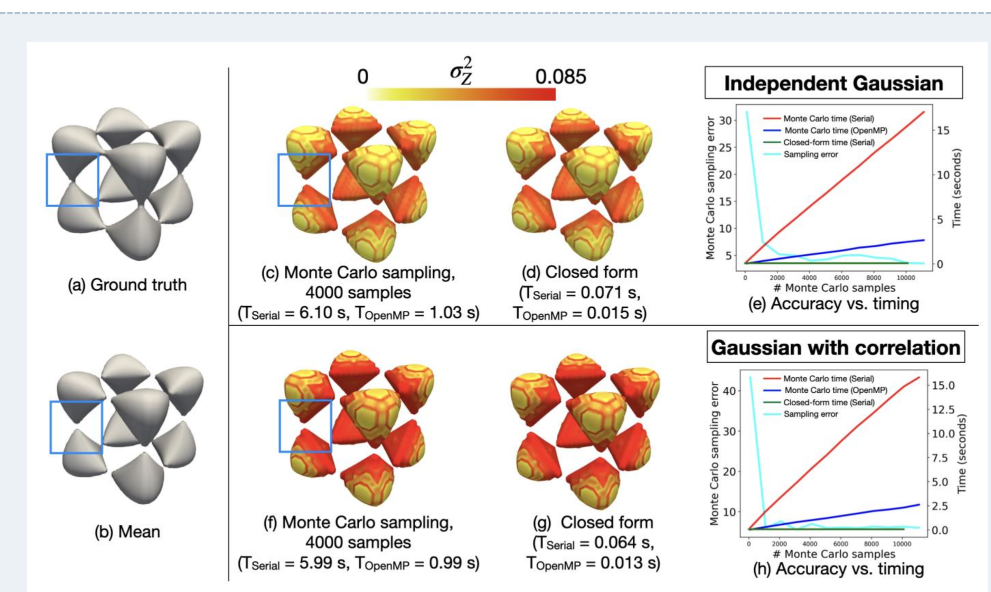

MAGIC framework for isosurface uncertainty visualization of the tangle dataset. Image (a) visualizes the ground truth isosurface. The ensemble with correlated noise is generated from the ground truth, and the isosurface from the mean of the ensemble is visualized in image (b). The meanfield isosurface extracted with MCA breaks in the area enclosed by the blue box and does not convey its spatial uncertainty visually. Images (c)-(d) visualize the isosurface uncertainty computed by ignoring spatial correlation. As seen from image (e), the runtime of the Monte Carlo solution (red and blue curves) increases with the number of samples, whereas the closed-form solution (green curve) is much faster and independent of the number of samples.

In-situ visualization in production: ParaView Catalyst

Exascale simulations cannot afford to write everything to disk before looking at it. ParaView/Catalyst 2.1 was deployed on three leadership-class systems, with an end-to-end in-situ workflow demonstrated with AMReX and WarpX and an initial integration with the MFEM finite-element library. Three new capabilities:

- Asynchronous execution.
- GPU-resident workflows.
- Closed-loop steering.

ParaView/Catalyst 6 deployment, RAPIDS3 software stewardship (2026)

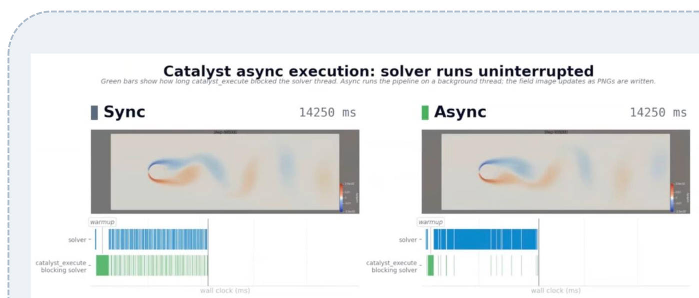

2D LBM cylinder flow on GU, identical pipeline run both ways. Solver blocked 33% of wall clock under sync, 6% under async. Async runs 27% more iterations per second in the same wall clock.

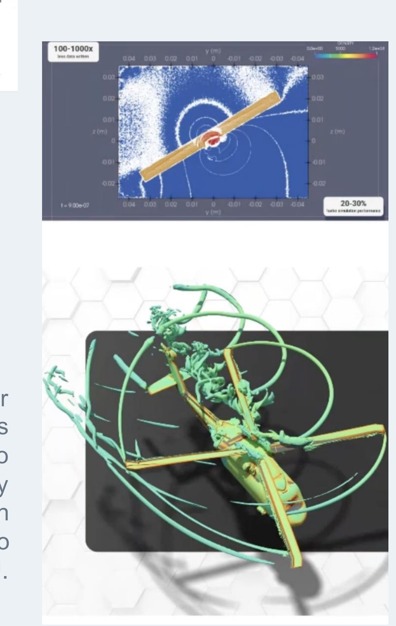

On-device Viskores reduction computes a scalar signal from the flow field. That single value is returned to the driver, which uses it as the input to a controller adjusting the solver's boundary condition on the next step. The round-trip stays on device except for eight bytes of scalar crossing to the driver: field data never leaves the GPU.

Argonne National Laboratory · Brookhaven National Laboratory · Berkeley Lab · Oak Ridge National Laboratory · Lawrence Livermore National Laboratory · kitware · Northwestern University · Georgia Tech · The Ohio State University · University of Florida · The HDF Group · U.S. Department of Energy Office of Science