Flexible In-Situ Analytics

Karsten Schwan, Greg Eisenhauer, Matt Wolf, and Ph.D. students
Georgia Institute of Technology, CERCS Research Center

Applications:

- GTC, GTS, Pixie3D Fusion applications
- LAMMPS materials modeling code
- S3D combustion code

Problems:

- Rapid output processing for timely science insights
- Large I/O output data volume
- Coupling to science users

Technology Basis

- ADIOS I/O interface
- EVPath data streaming middleware as ADIOS transport
- Location options for locating analytics processing: compute nodes, staging, remote, offline
- NNTI (Sandia) efficient transport for RDMA

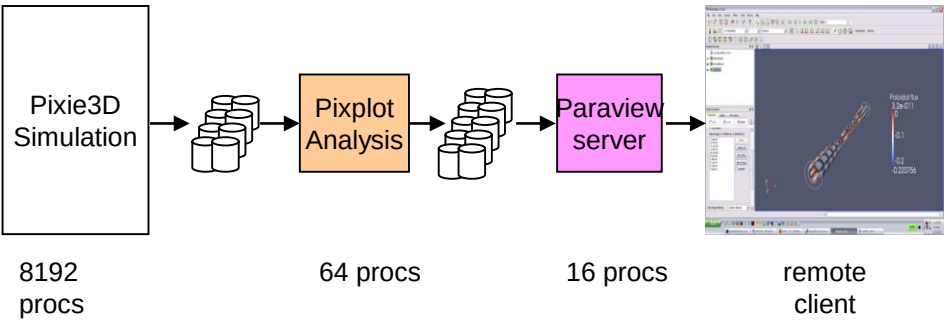

Pixie3D I/O Processing Pipeline

Challenges

- Limited resources for I/O and analytics
- High I/O performance with additional online analytics
  - Require online data reduction
  - Require limiting use of disk subsystem
  - Require judicious data movement, analytics placement, and analytics scheduling

Flexible Placement and Execution for In-Situ Analytics

Technology Contributions

- ADIOS/EVPath I/O middleware
  - High performance data movement on IB and UGNI
  - Support diverse in situ analytics placement options
  - Higher-level API: meta-data rich, easy-to-use
- Flexible Placement
  - Metric-driven optimization, including for end-to-end performance/cost objectives
- Resource Containers:
  - Resource provisioning for analytics components

Result/Impact

- Extended ADIOS with new transport to support location-flexible in situ analytics
- Implemented in situ analytics for GTS, LAMMPS, Pixie3D, S3D
- Up to 30% end-to-end performance improvement of those applications through flexible placement
- Utilized DOE-provided NNTI RDMA transport for support of data staging

Accelerating Pixie3D I/O Pipeline via Flexible Placement

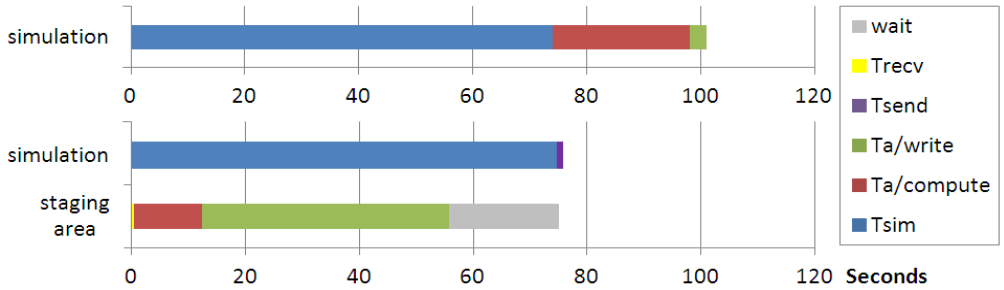

Using 0.78% additional nodes offloading Pixplot and I/O to staging area increases performance by 33% in comparison to inline placement at the scale of 8192 cores

Managing I/O Resources with I/O Containers

Applications:

- LAMMPS materials modeling code
- DOE Sandia applications
- SmartPointer Scientific Annotation Toolkit

Problems:

- Poor staging resource allocations can cause dataflow bottleneck
- Complex computational models for analytics execution

Technology Basis

- ADIOS I/O interface
- EVPath data streaming for monitoring and control
- Multilevel management hierarchy
- Runtime resource management for I/O pipelines
- Scalable transactions for resilience (with DOE Sandia)

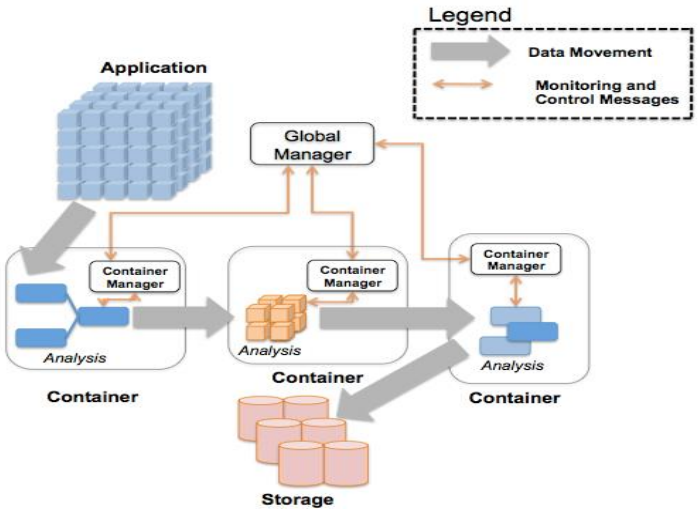

I/O Containers Overview

Challenges

- Limited resources for I/O and analytics
- Move offline analysis workflows online
  - Must support multiple computational models
  - Provide scalability for non-scalable analysis codes
  - Controlled data movement
  - Provide fault and performance isolation for analysis components and scientific applications

Managing I/O Resources with I/O Containers

Technology

- I/O Containers
  - Move offline workflows online to operate on data in-transit
  - Runtime resource management to balance resource usage amongst online analysis codes
- Doubly Distributed Transactions
  - Provide resilience for data movement and control operations in HPC environments

Result/Impact

- Extended ADIOS and DataTap to use I/O Containers middleware
- Increased end to end performance for online analysis pipelines
- Measured performance impact of implementing transactions in HPC environments

Improved end to end performance though I/O Containers management

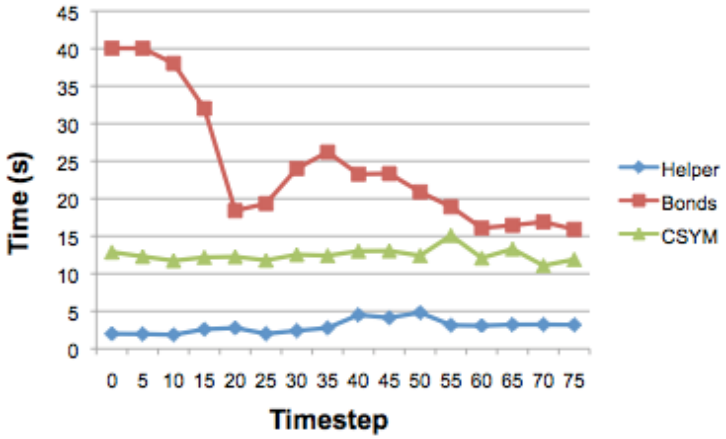

256 simulation nodes; 13 staging nodes

Performance impact of transactions for control operations