

# Examples of In Transit Visualization

## Petascale Data Analytics: Challenges and Opportunities

November 14, 2011

Kenneth Moreland (SNL), Ron Oldfield (SNL), Pat Marion (Kitware),  
Sebastien Jourdain (Kitware), Norbert Podhorszki (ORNL),  
Venkatram Vishwanath (ANL), Nathan Fabian (SNL),  
Ciprian Docan (Rutgers), Manish Parashar (Rutgers),  
Mark Hereld (ANL), Michael E. Papka (ANL), Scott Klasky (ORNL)

SAND 2011-XXXXP

# Collaborators

- **Sandia National Laboratories**
  - Kenneth Moreland
  - Nathan Fabian
  - David Thompson
  - Ron Oldfield
- **Los Alamos National Laboratories**
  - James Ahrens
  - Jonathan Woodring
- **Oak Ridge National Laboratory**
  - Scott Klasky
  - Norbert Podhorszki
- **Argonne National Laboratory**
  - Venkatram Vishwanath
  - Mark Hereld
  - Michael E. Papka

- **Kitware, Inc.**
  - Berk Geveci
  - Utkarsh Ayachit
  - Andrew C. Bauer
  - Pat Marion
  - Sebastien Jourdain
  - David DeMarle
- **University of Colorado at Boulder**
  - Michel Rasquin
  - Kenneth E. Jansen
- **Rutgers University**
  - Ciprian Docan
  - Manish Parashar

# Extreme scale computing

- Trends
  - More FLOPS
  - More concurrency
  - Comparatively less storage, I/O bandwidth
- ASCI purple (49 TB/140 GB/s) –  
JaguarPF (300 TB/200 GB/s)
  - Most people get < 5 GB/sec at scale

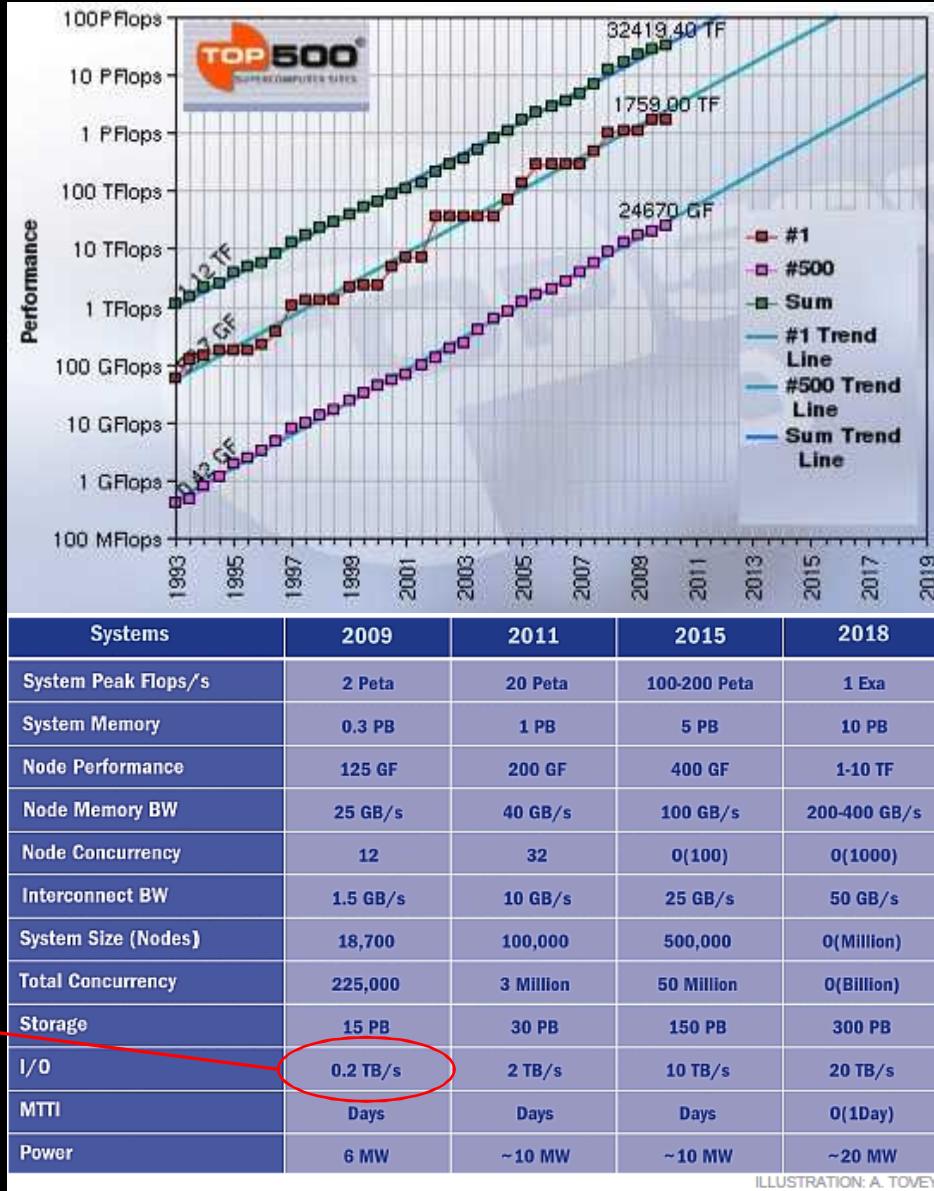
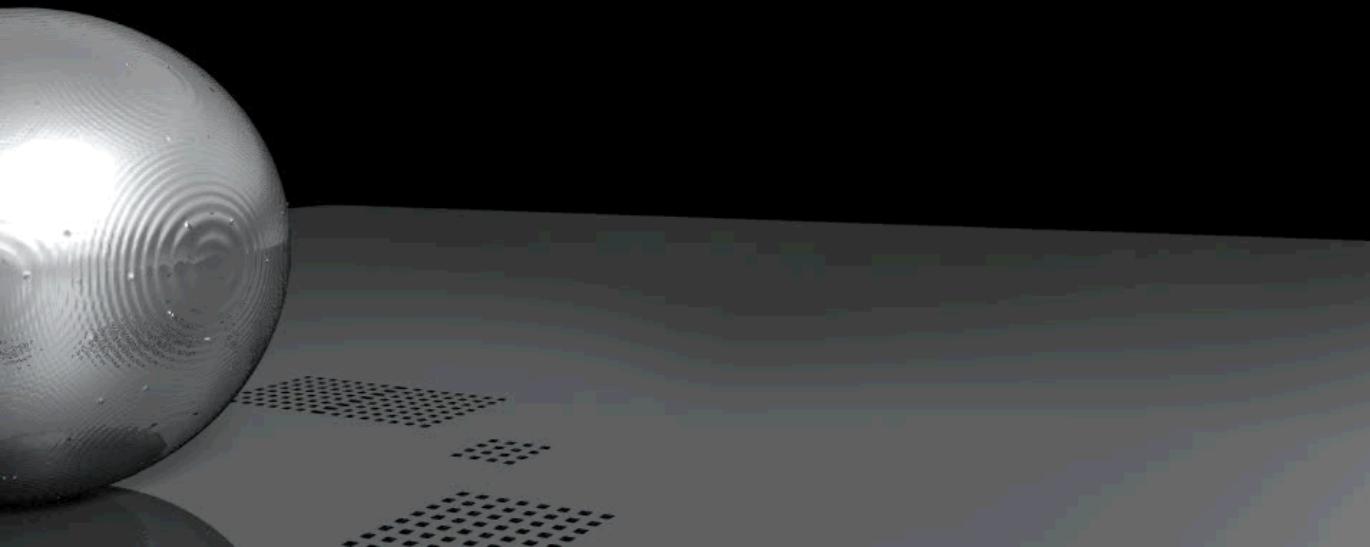
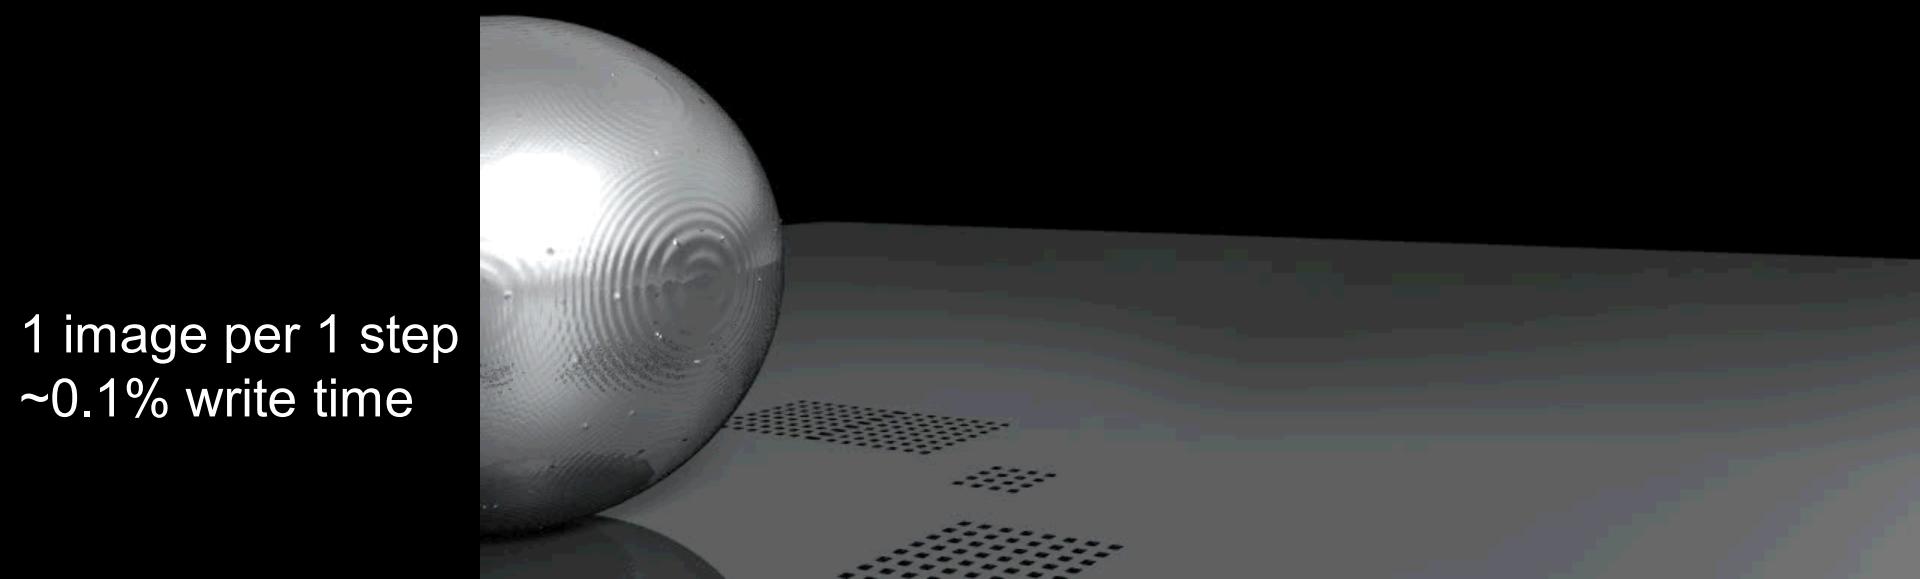


ILLUSTRATION: A. TOVEY

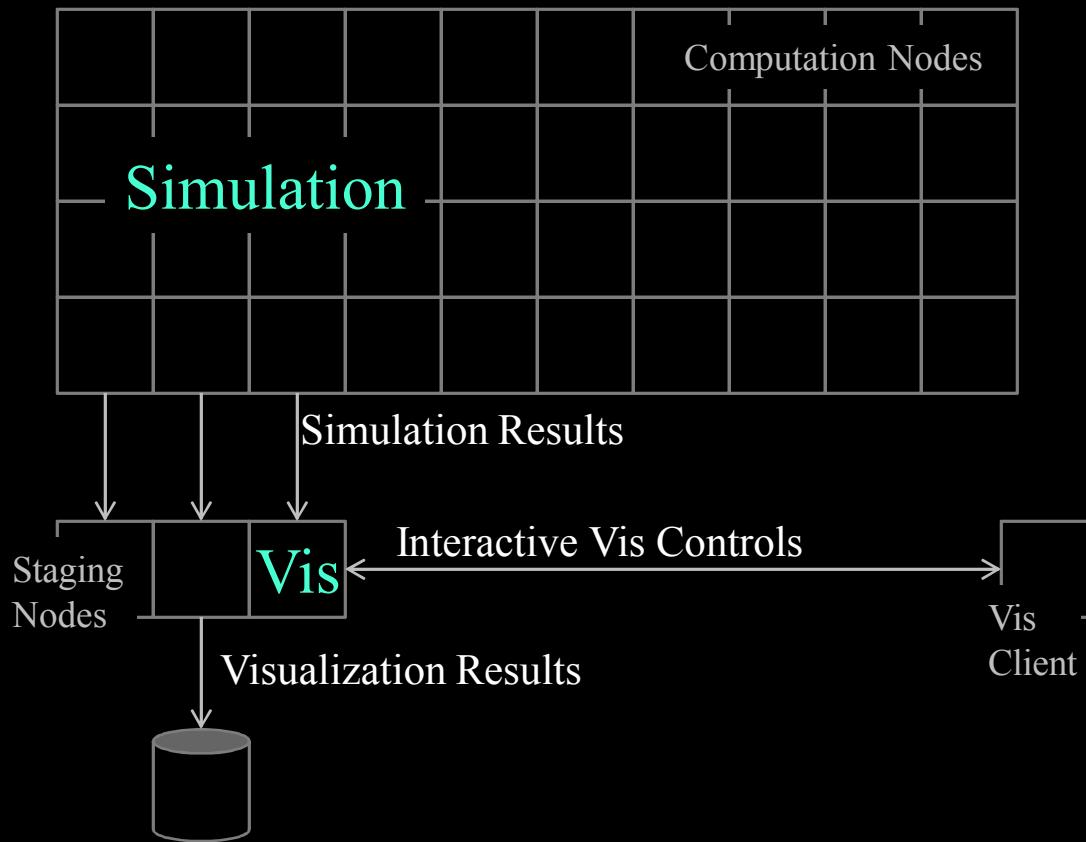


1 mesh per 100 steps  
~1% write time

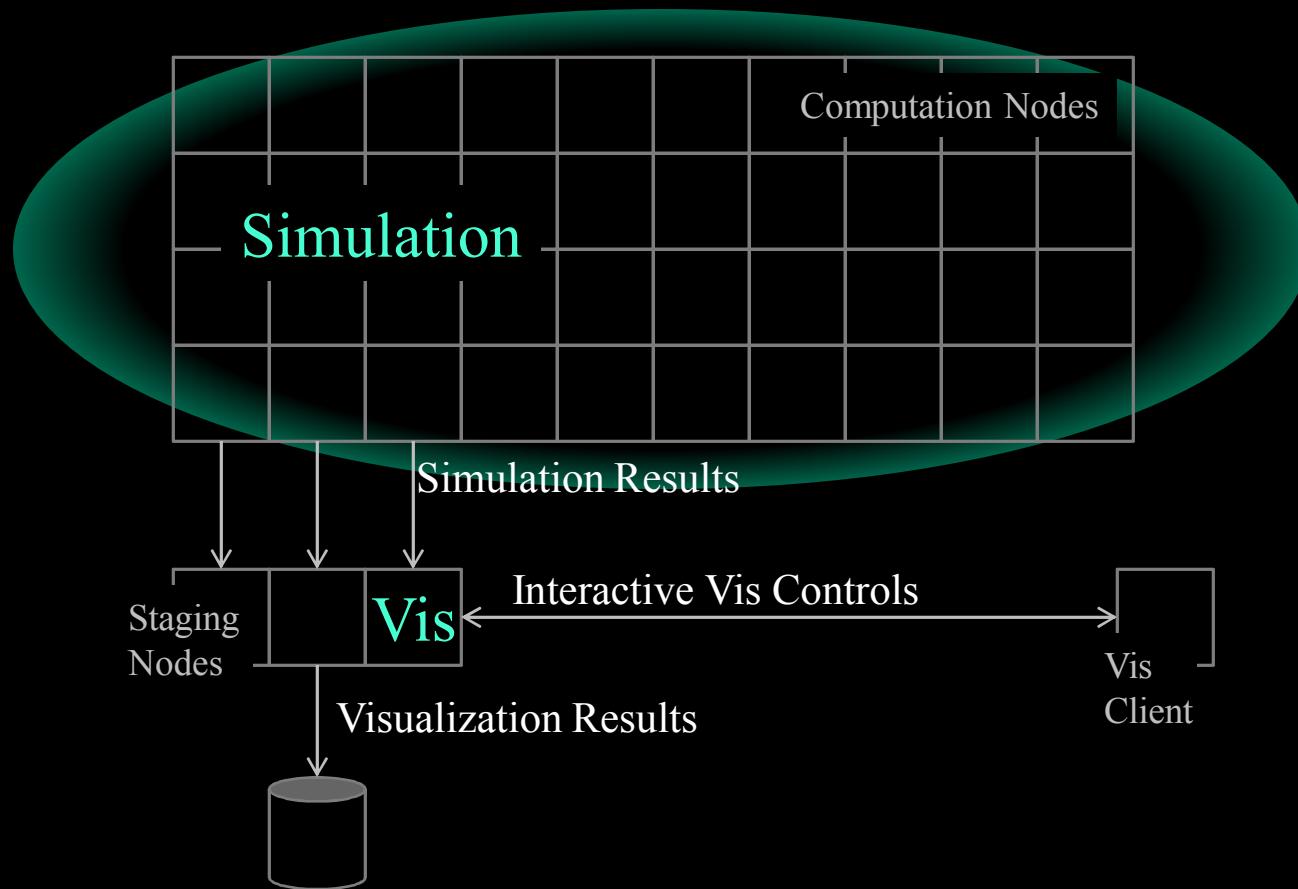


1 image per 1 step  
~0.1% write time

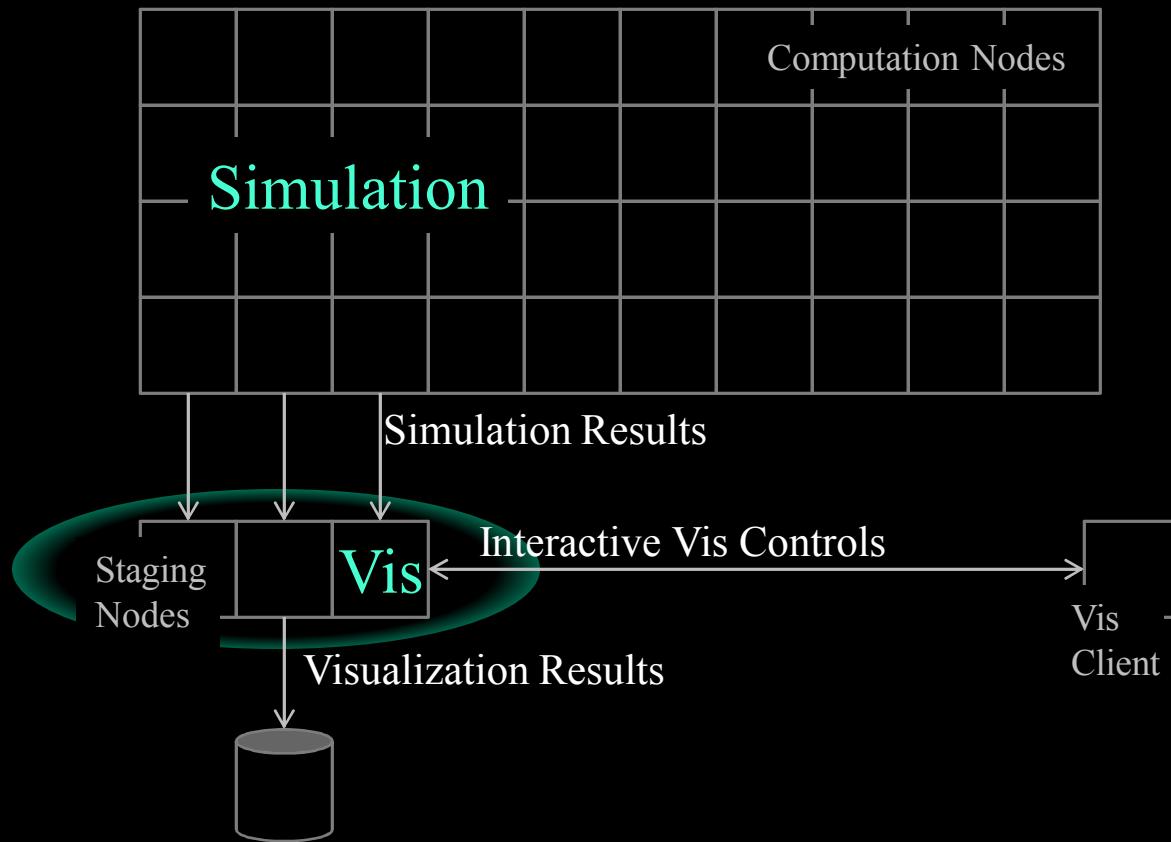
	Capability	Coupling	Footprint	Transfer	Interactive
ParaView In Transit	High	Loose	~5% Extra Nodes	Hi Speed Transfer	Yes



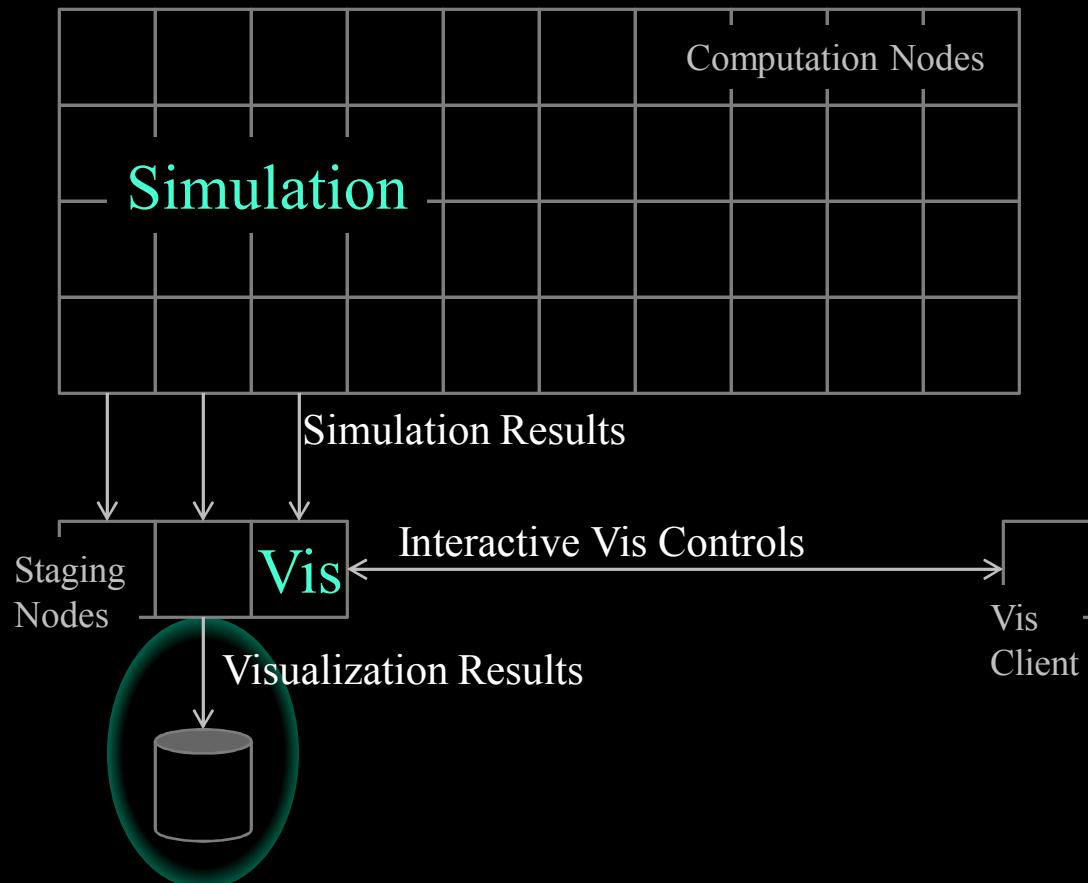
	Capability	Coupling	Footprint	Transfer	Interactive
ParaView In Transit	High	Loose	~5% Extra Nodes	Hi Speed Transfer	Yes



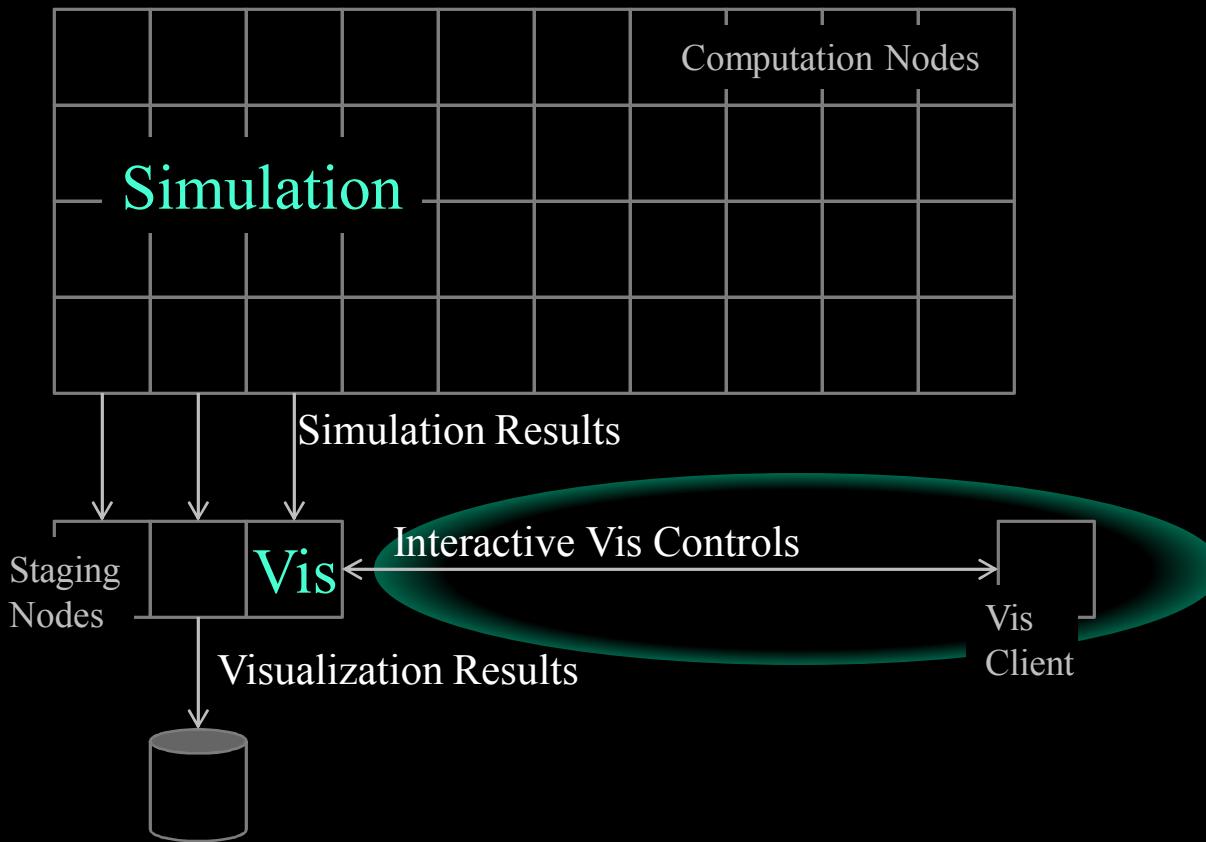
	Capability	Coupling	Footprint	Transfer	Interactive
ParaView In Transit	High	Loose	~5% Extra Nodes	Hi Speed Transfer	Yes

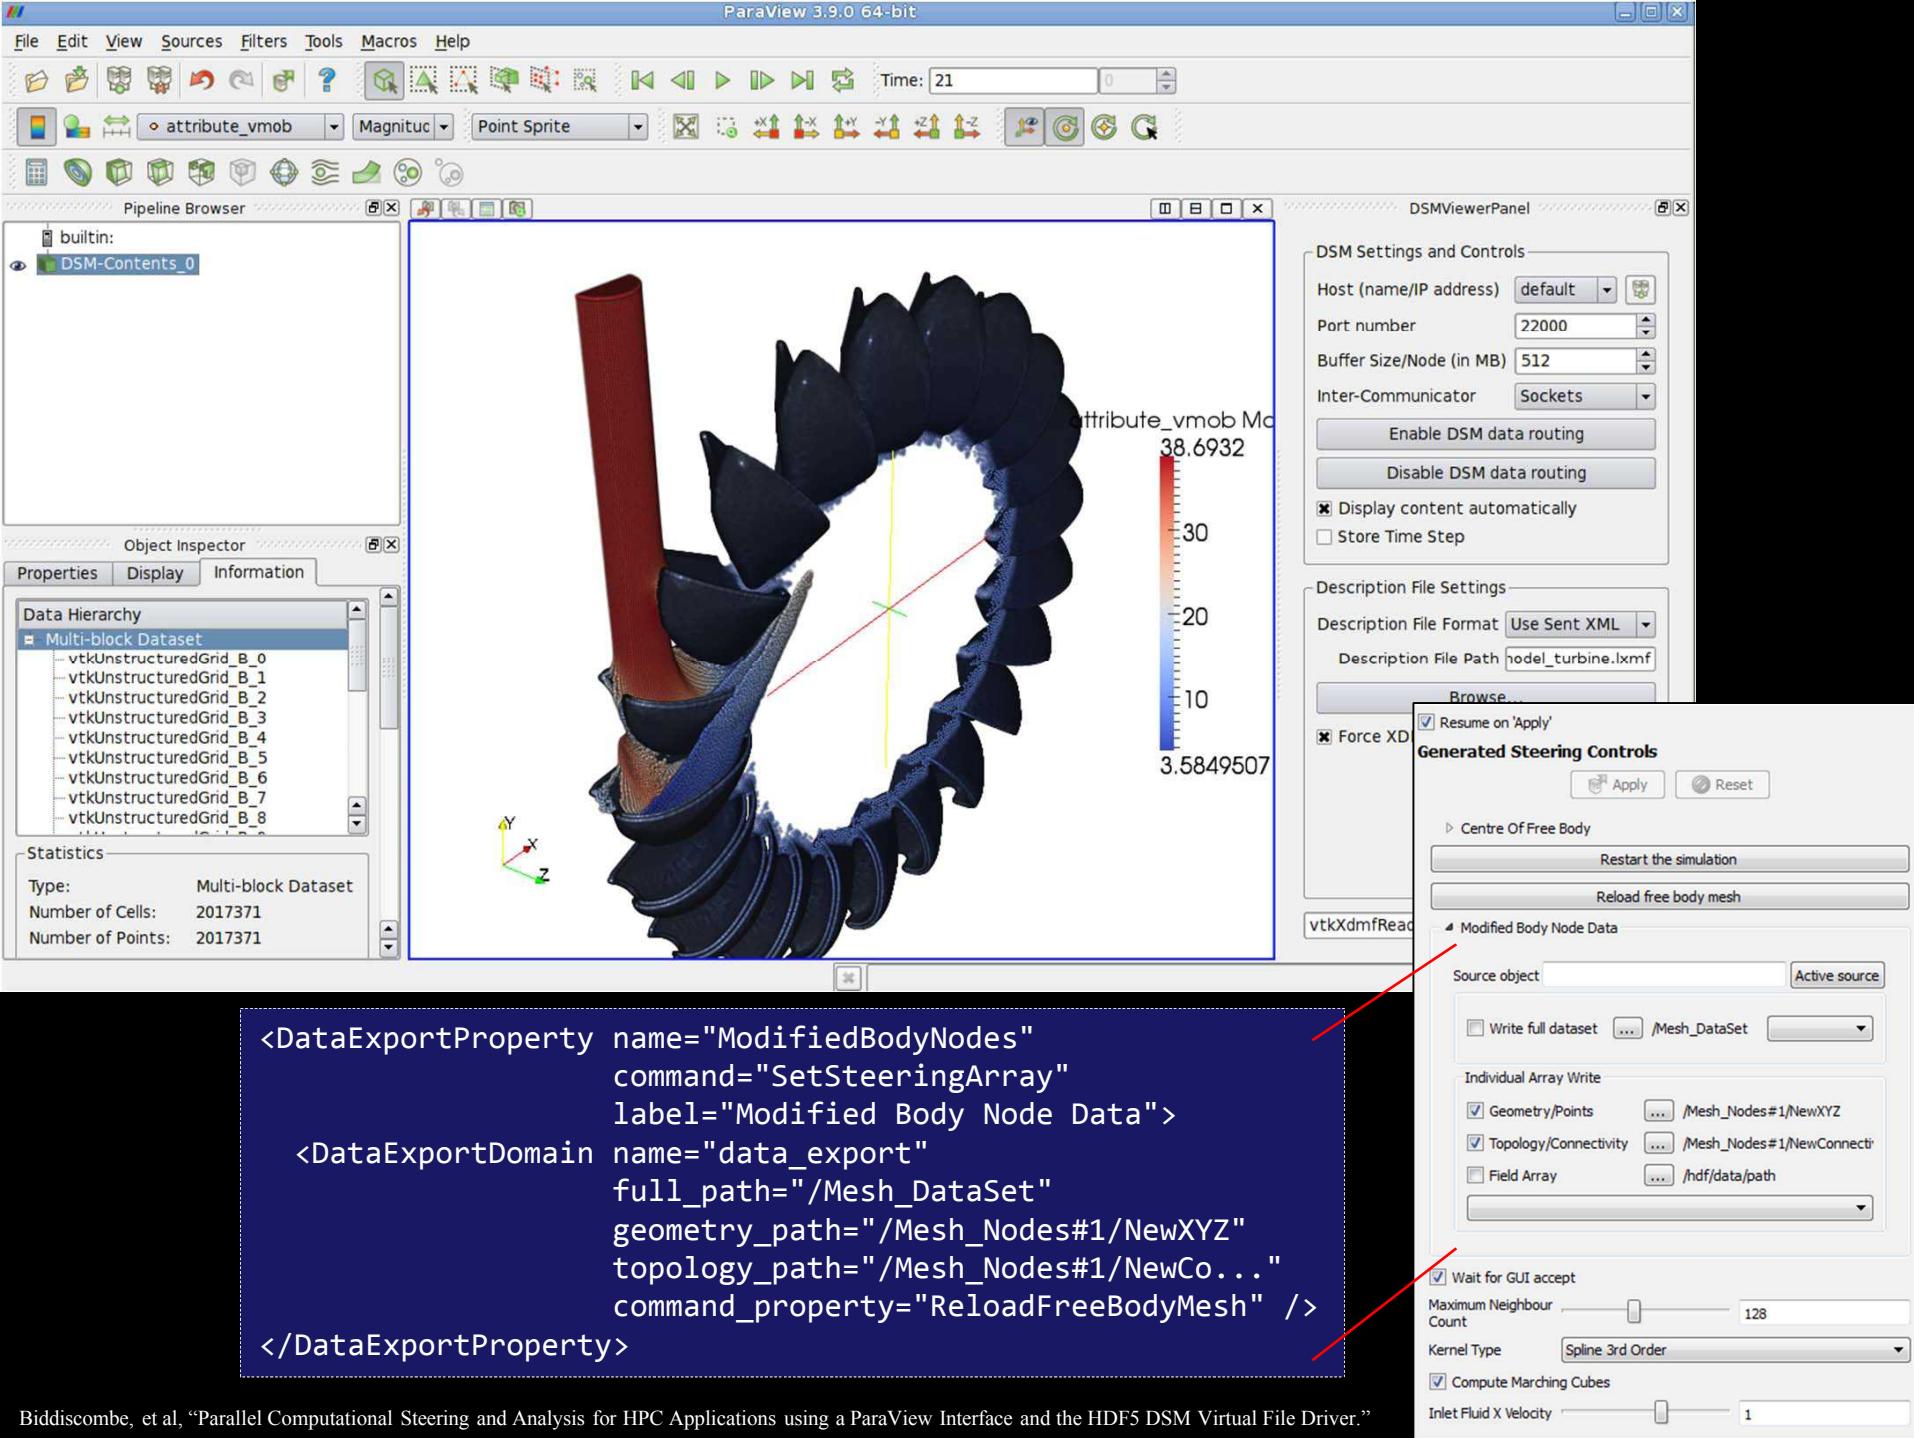


	Capability	Coupling	Footprint	Transfer	Interactive
ParaView In Transit	High	Loose	~5% Extra Nodes	Hi Speed Transfer	Yes



	Capability	Coupling	Footprint	Transfer	Interactive
ParaView In Transit	High	Loose	~5% Extra Nodes	Hi Speed Transfer	Yes







# Network Scalable Service Interface (NESSIE)

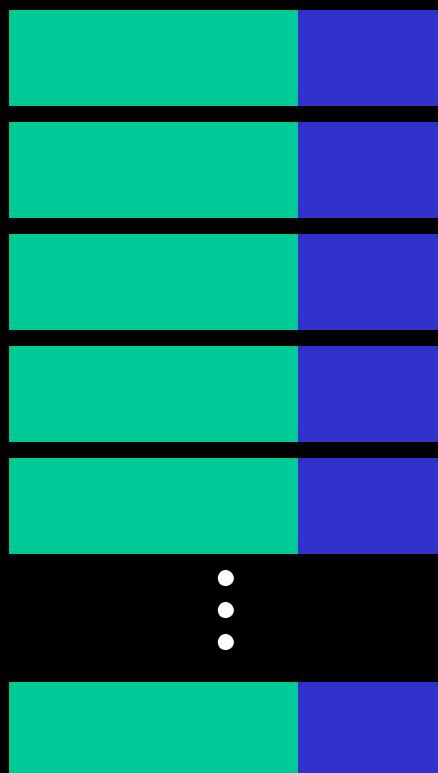
- Provides a remote-procedure call (RPC) abstraction.
- Relies on client/server stubs to marshal/unmarshal procedure call parameters.
  - Also supports data messages for raw data transfers.
- Server uses one-sided API for native remote direct-memory (RDMA).
- Used to implement:
  - File system services
  - HPC database access proxies
  - PnetCDF data staging
- Our work leverages PnetCDF data staging for in transit analysis of CTH.



# NESSIE-CTH-ParaView Integration



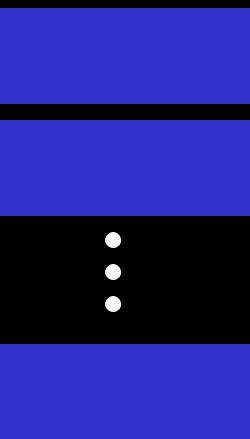
Client Application



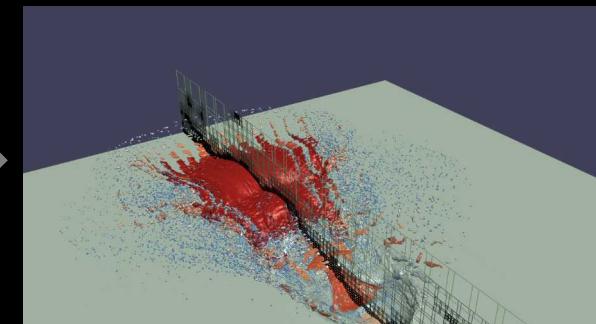
Fragment-Detection  
Service



Raw  
Data



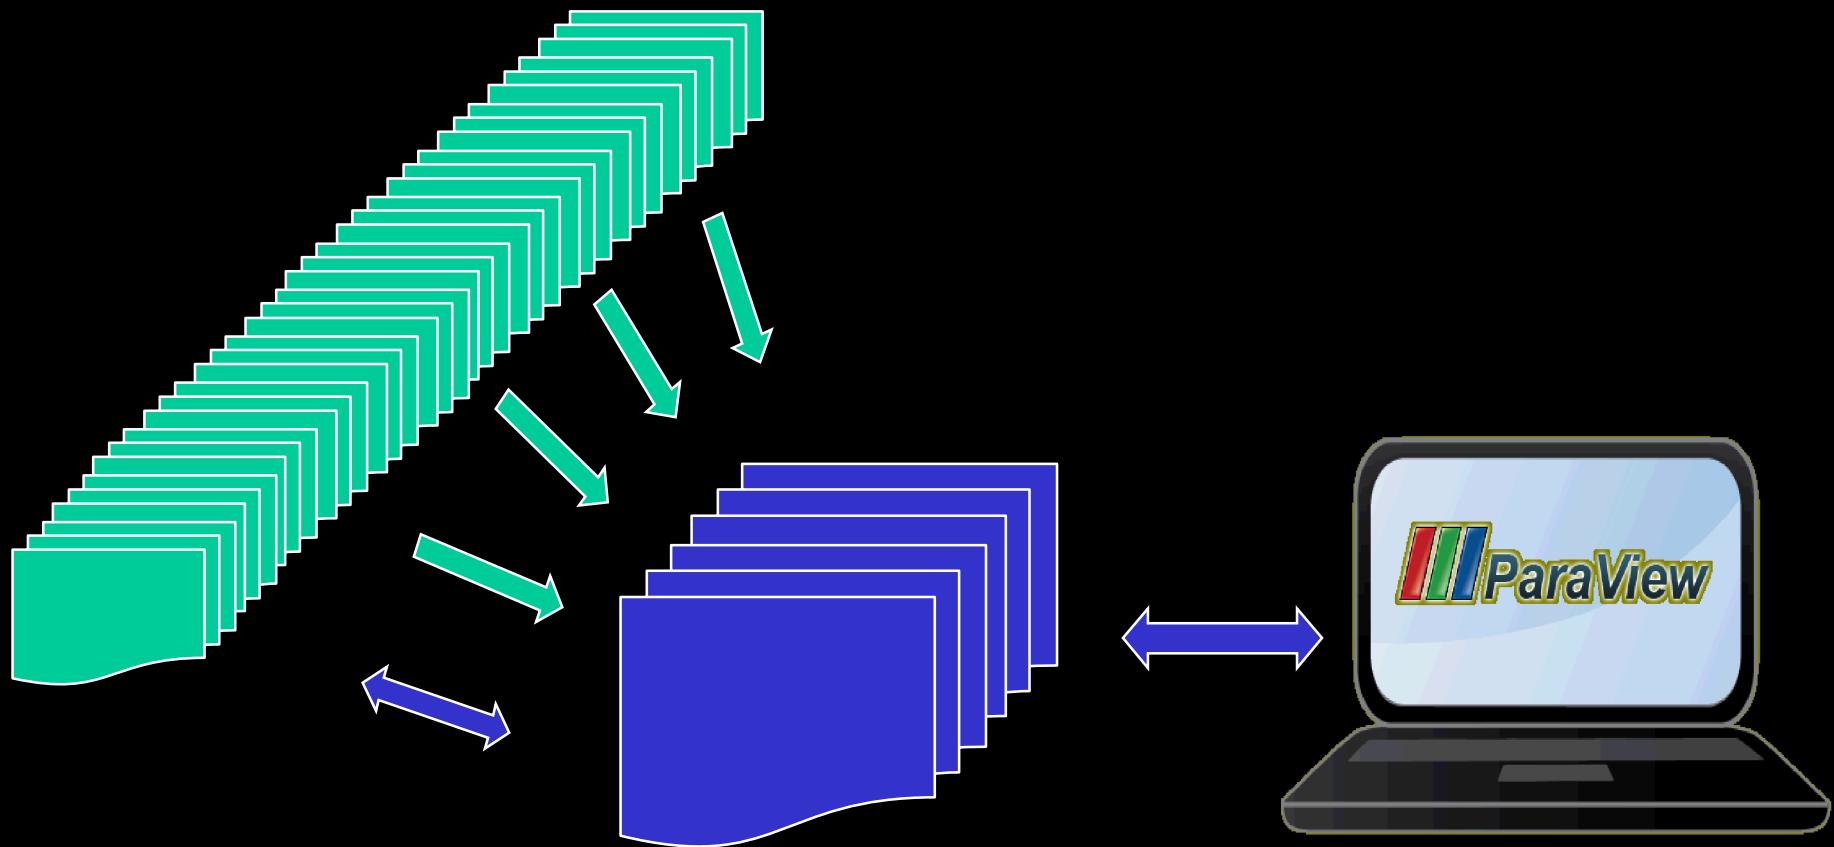
Fragment  
Data

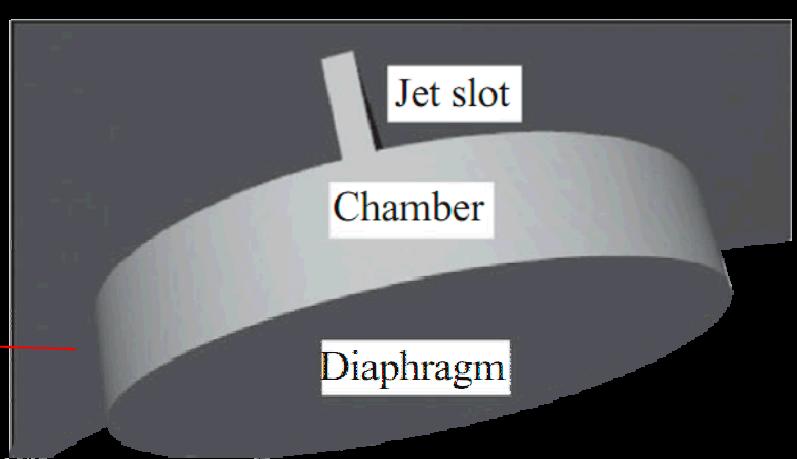
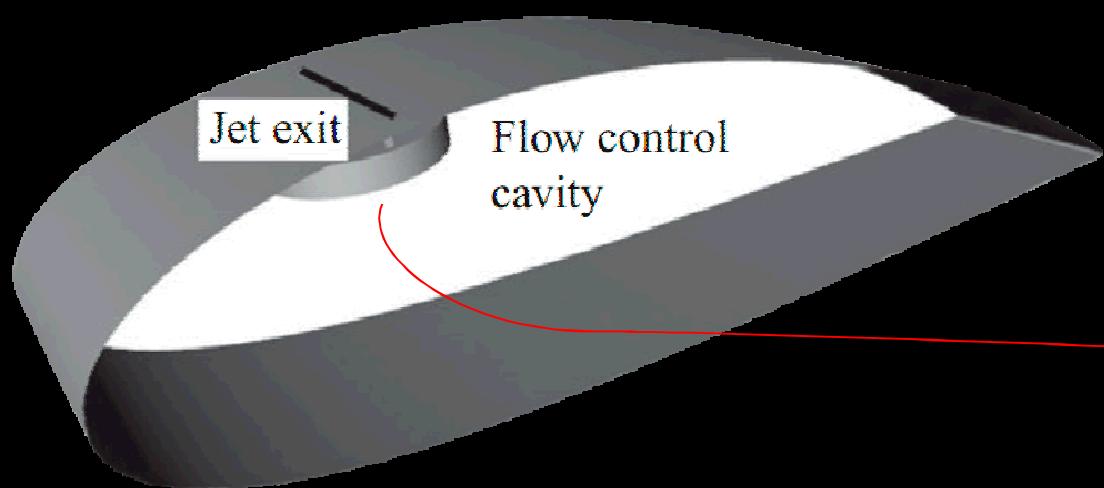
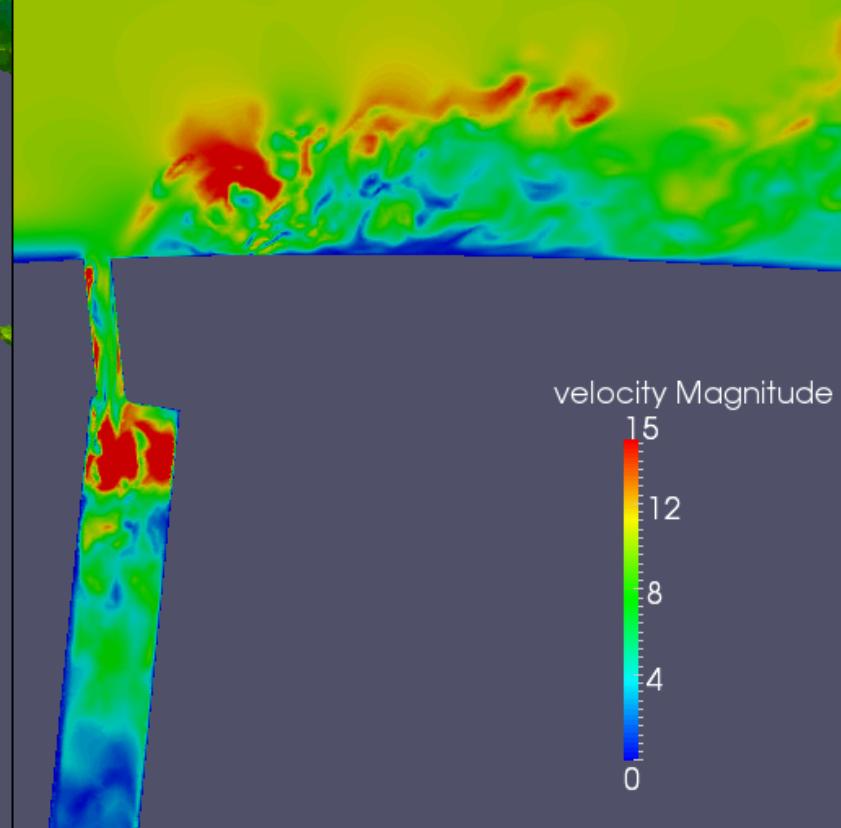
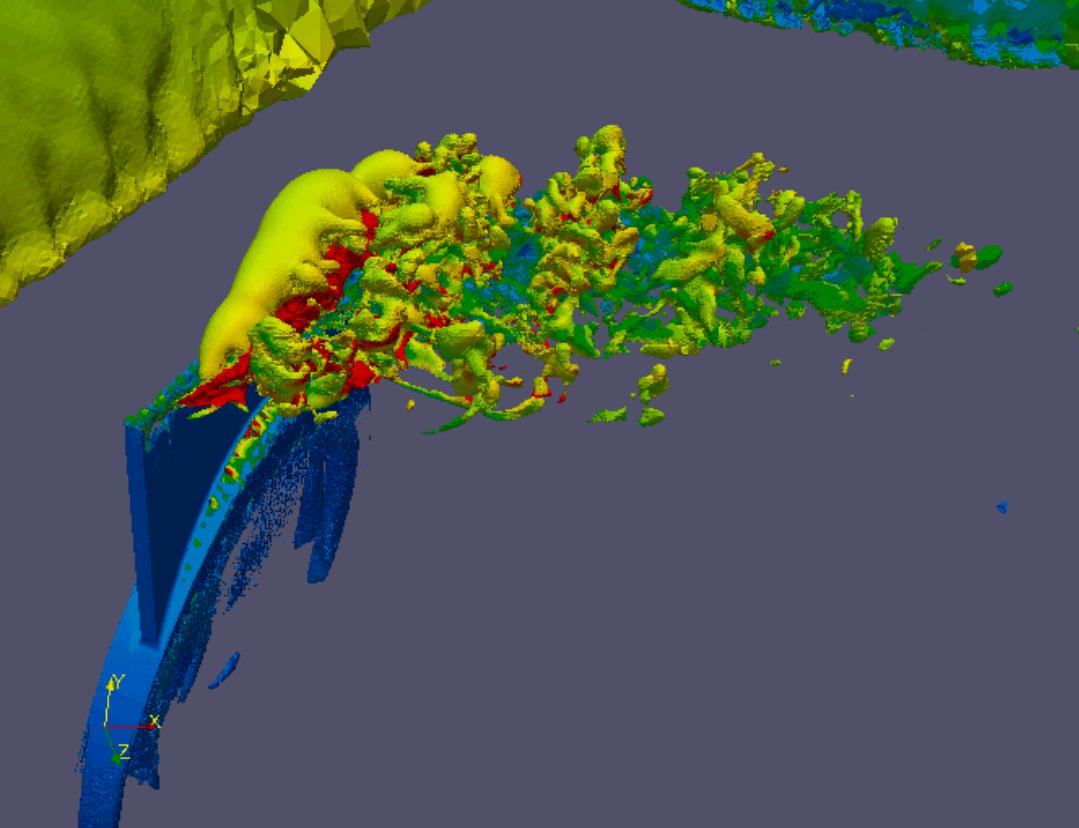


# GLEAN

- Flexible/extensible framework for simulation-time data analysis and I/O acceleration.
- Adjusts to analysis and system characteristics.
- Client/server architecture.
  - Data streams from running solver (client) to dedicated staging nodes (server).
- Integrated in ParaView via VTK reader/writer components.
  - ParaView coprocessing (in situ) library (integrated with server) uses VTK GLEAN writer (“client,” compute nodes).
  - ParaView pvserver (interactive vis service) uses VTK GLEAN reader (“server,” staging nodes).
  - Behaves like file read/write from ParaView perspective.

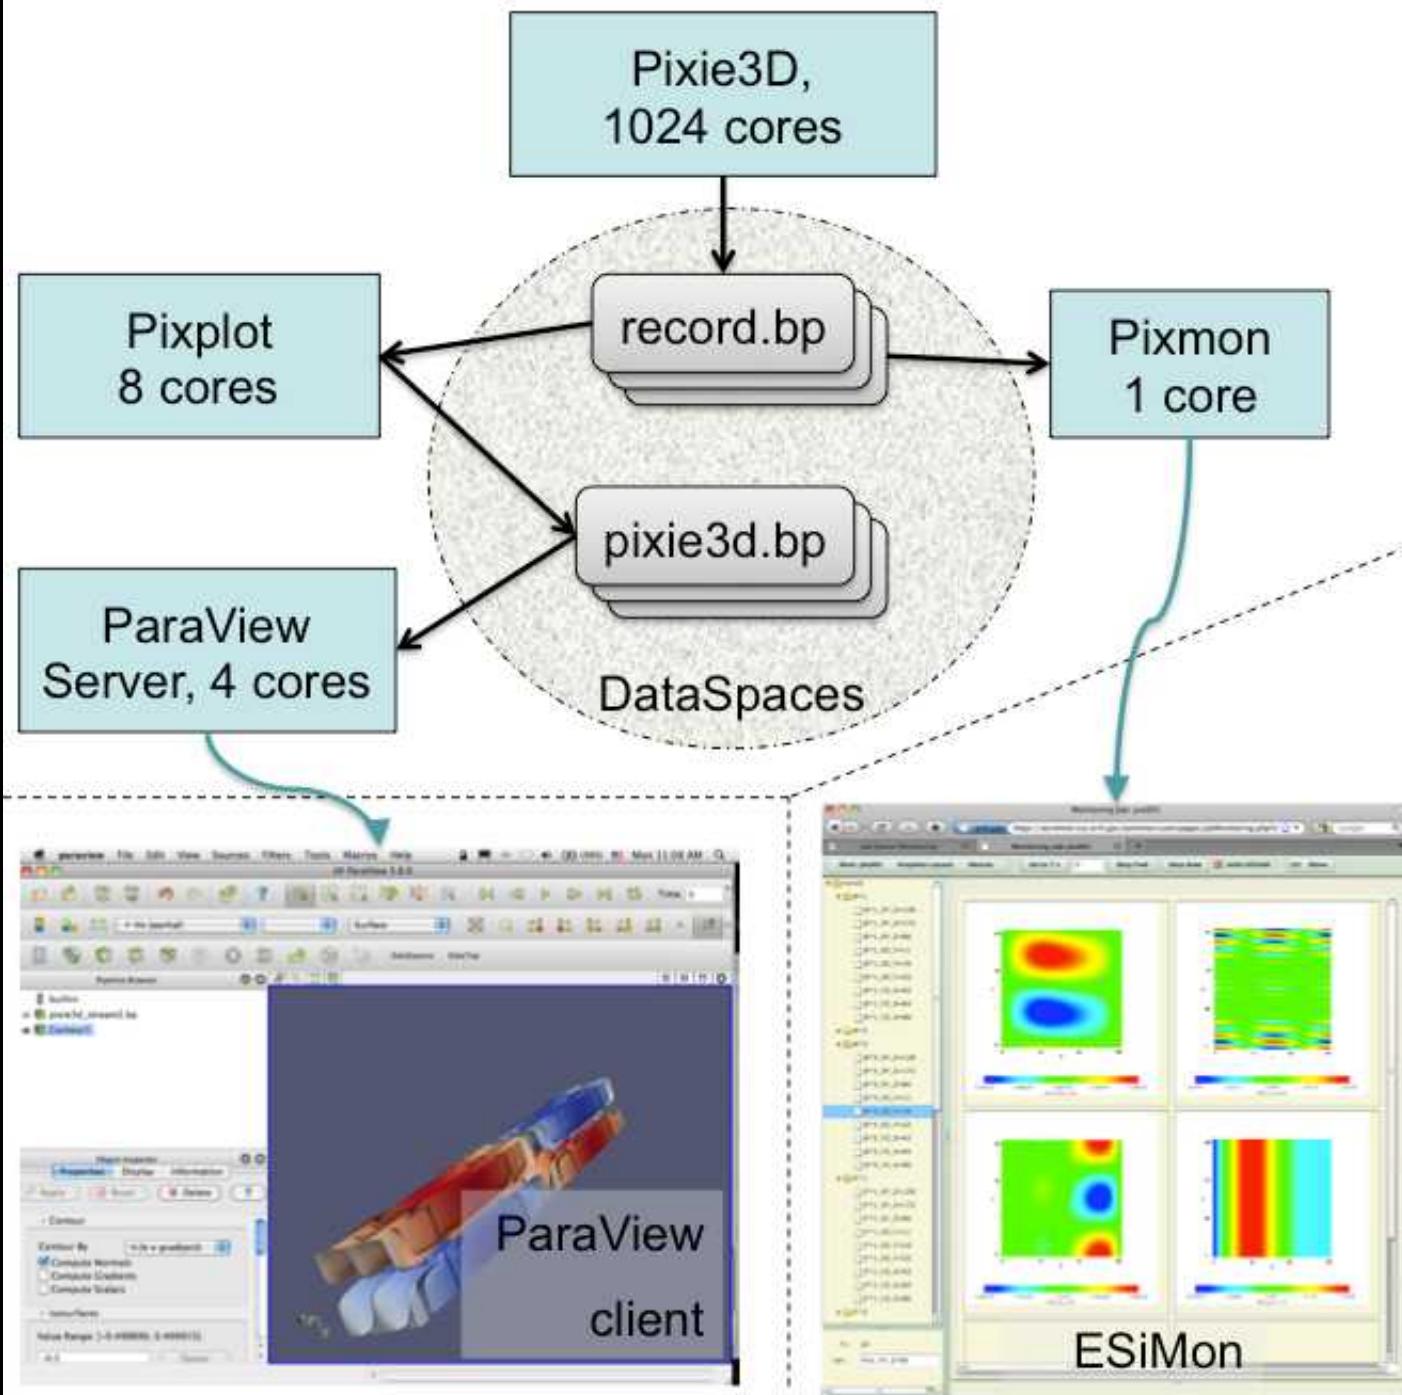
# ParaView In Situ + Covis + Interactive Vis with GLEAN





# Adaptable I/O System (ADIOS)

- Separates I/O API from I/O implementation.
  - Abstracts I/O transport, destination, format.
  - Can be changed at runtime through XML file.
- Uses DataSpaces for virtual semantically-specialized shared space.
  - Location/distribution agnostic.
  - Accessed with asynchronous operations.
  - Exists as separate application – provides fault isolation.
    - Coupled codes do not bring each other down.
  - Manages multiple versions of dataset (e.g. timesteps).
    - Autonomous buffering and eviction.
- Using ADIOS for files automatically enables in transit.
  - But, change read semantics for unavailable timesteps.



# Conclusion

- Potential advantages of in transit visualization.
  - About everything in situ has going for it.
  - Loose coupling between simulation and visualization.
  - Fault isolation.
  - Multiple event loops / asynchronous operation.
  - Federated environment.
- Should stop thinking of I/O in terms of “files.”
  - I/O libraries should (will) progress to more general.
  - Application should stop considering data sources/sinks as static collections.
  - Applications should *not* go low level (POSIX, MPI/IO).

# Shameless Promotion of Other Presentations

Sun	1:30	TCC 205	ParaView In Situ Tutorial
Tue	10:00	ASC Booth	In Situ Visualization and Analysis
Tue	11-1	ASC Booth	ParaView Demonstrations
Tue	5:15	Poster Reception	Co-visualization of full data and in situ data extracts from unstructured grid CFD at 160k cores
Tue	5:30	TCC LL1	BoF: Next Generation I/O and visualization requirements for in situ processing
Thu	10-12	ASC Booth	ParaView Demonstrations
Thu	10:30	ASC Booth	ASC Tri-Lab Visualization on the Cielo Platform

Questions?