

PETROGLYPHS TO PETAFLIPS
— CUG 2005 —



Red Storm


Data Analysis & Visualization Environment

Constantine "Dino" Pavlakos
David White



Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy under contract DE-AC04-94AL85000.






We are building an environment that addresses high performance application needs



With the following strategies:

- *Integrated end-to-end infrastructure*
- *Scalability (hardware and software tools)*
- *Leveraging use of commodity-based graphics HW (clusters and office)*
- *Accessibility from the Desktop/Office*
- *Ease of use*




The Data Analysis / Visualization Process ... a lot more than just rendering



PE

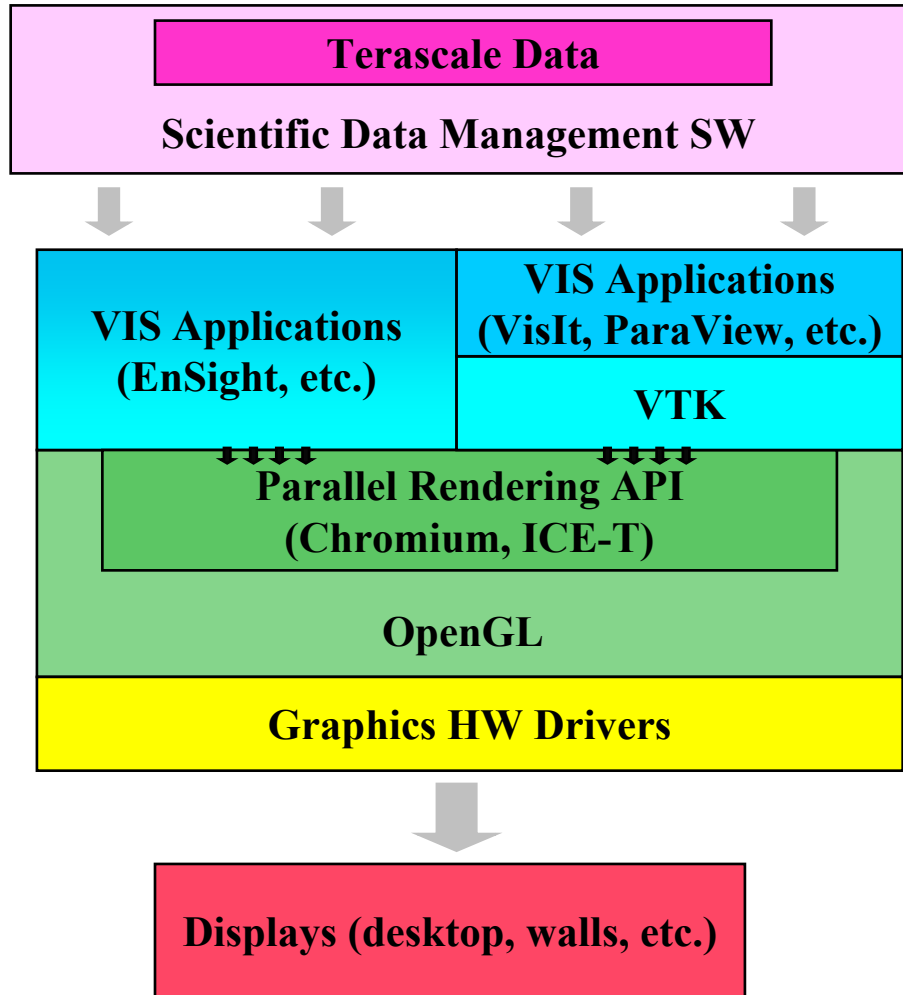
Data Sources: Simulations, Archives, Experiments					User Services: Navigation Rendering Control Advanced User Interface Collaborative Control Display Control	
Data Services:	Data Access	Filtering	Data Query	Subsetting Resampling		Data Algebra $x,y,z \Rightarrow \text{mag}/\Phi$
	Format/Representation Conversion (e.g., multi-res)		Data Simplification			Re-partitioning $M \Rightarrow N$
	Data Mining	Feature Detection, Extraction, Tracking		Data Fusion & Comparison		Time History Generation
Visualization Services:	Renderable-Object Generation (eg., surface extraction)		Attribute Specification (eg., Volume Transfer-function)			
		Multi-Visualization Technique Combine	Time Sequence Generation			
	Surface rendering	Volume rendering	Plotting	Image-based Rendering		
Display Modalities:	Desktop Display	Theater Display	Powerwalls	Immersive Stereoscopic		



The end-to-end environment we are deploying has benefited from a cooperative development effort



PETROGLYPHS TO PETAFLOPS
CUG 2005




ASCI Tri-labs:

- Sandia National Laboratories
- Lawrence Livermore National Laboratory
- Los Alamos National Laboratory


Partners (past and present):

- NCSA (HDF5)
- CEI (Parallel **EnSight**)
- Kitware (Parallel VTK, **ParaView**)
- Stanford University (see “The Chromium Project” at sourceforge.net)
- Princeton University (Scalable Displays)
- NVIDIA (Linux graphics drivers)
- RedHat / Tungsten Graphics (Chromium and Distributed X Server)
- IBM (Bertha displays)

The Big Picture (Tri-lab architecture)



PETROGLYPHS TO PETAFLORS
— CUG 200



Really Big Data Handling

- Data Manipulation
 - Data Subsetting
 - Geometry Extraction
 - Feature Detection / Extraction (Data Mining)
 - Data Preparation for
 - Hierarchical / Multi-res
 - Out-of-core
- Some Full-featured Vis and Rendering (maybe SW only)

- EnSight Server of Servers
- ParaView
- VisIt
- Scalable Data Services

Big Data Handling


- Data Manipulation
- Full-featured Visualization
 - Raw data
 - Hierarchical / Multi-res
 - Out-of-core
- Parallel HW Rendering

- EnSight Server of Servers
- EnSight Client
- ParaView
- VisIt
- Scalable Data Services

Smaller Data

- Remote Image Display
- Visualization
 - Hierarchical
 - Out-of-core
 - Anything with smaller data
- HW Rendering

- EnSight Client
- EnSight
- ParaView
- VisIt
- Access to Data Services




PETROGLYPHS TO PETAFLIPS


CUG 2005



Red Storm basics



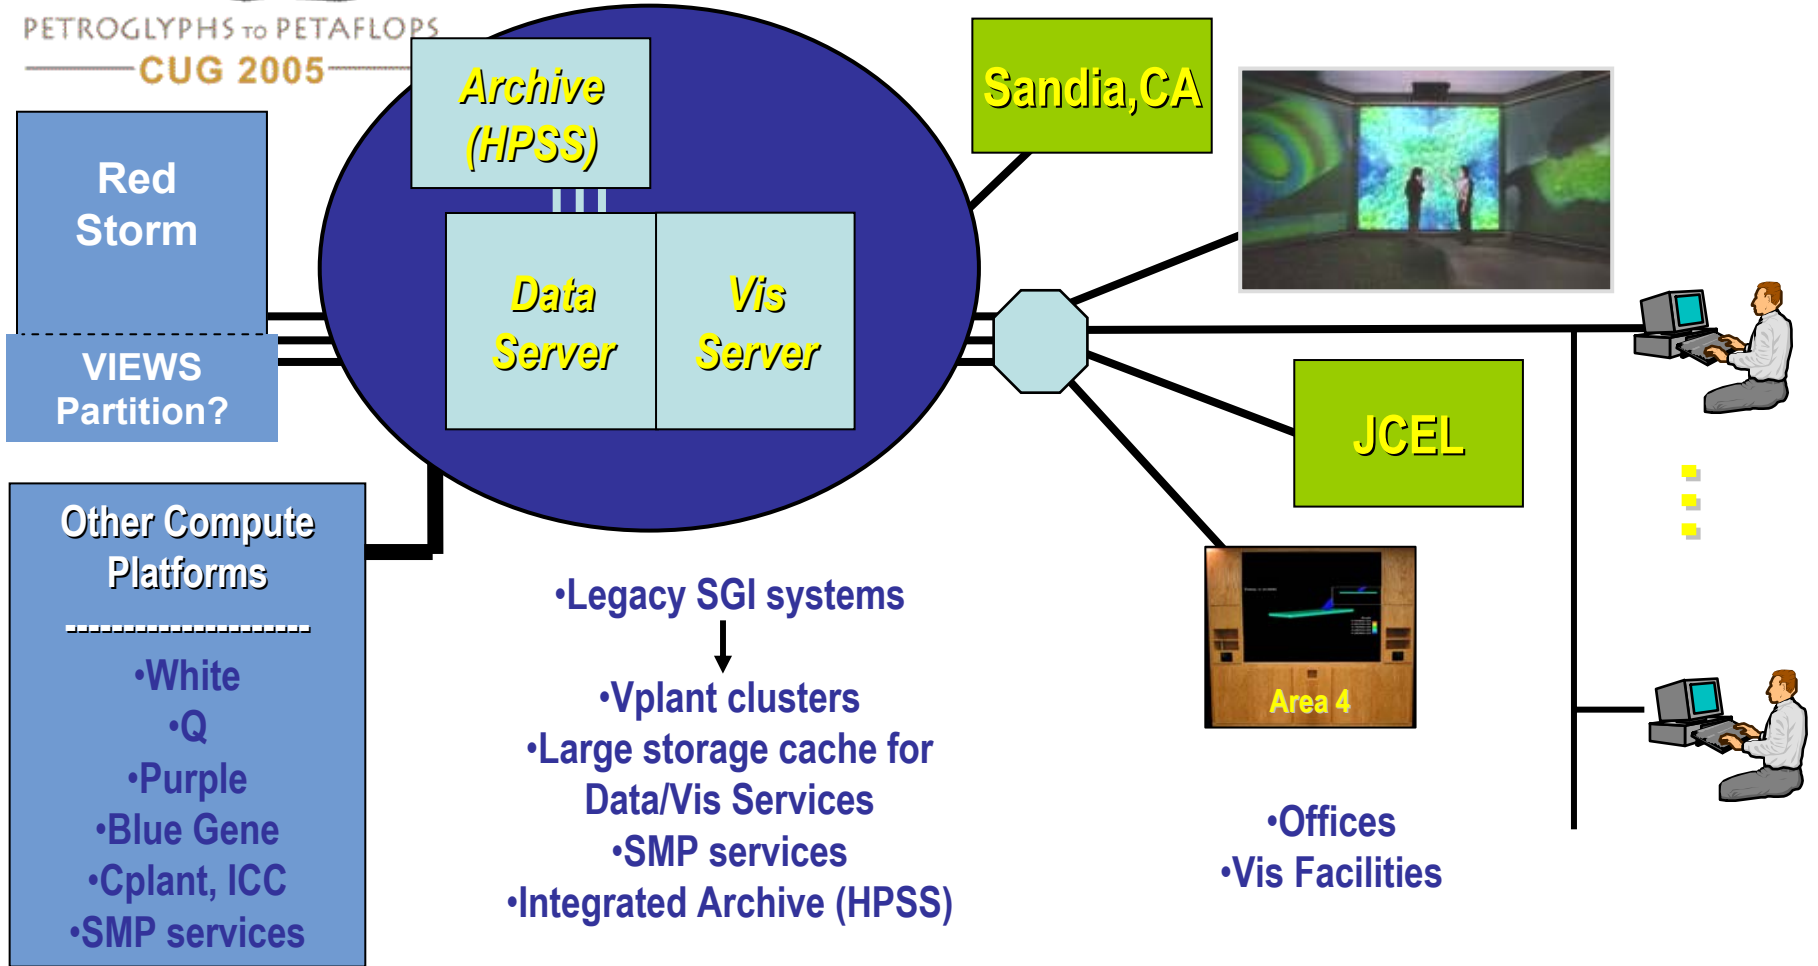
- Sustained PFS I/O at 50 GB/s
- Sustained external I/O at 25 GB/s (10 TB @ 10 GB/s takes 17 min)
- Plan has been for 256 IEWS nodes on each end of Red Storm (uncertain)
 - With PCI-Express graphics
 - “Full-service OS” ... will support EnSight SOS



We are building a Data/Vis Center that provides a breadth of services for our complete high performance computing environment




PETROGLYPHS TO PETAFLIPS
CUG 2005

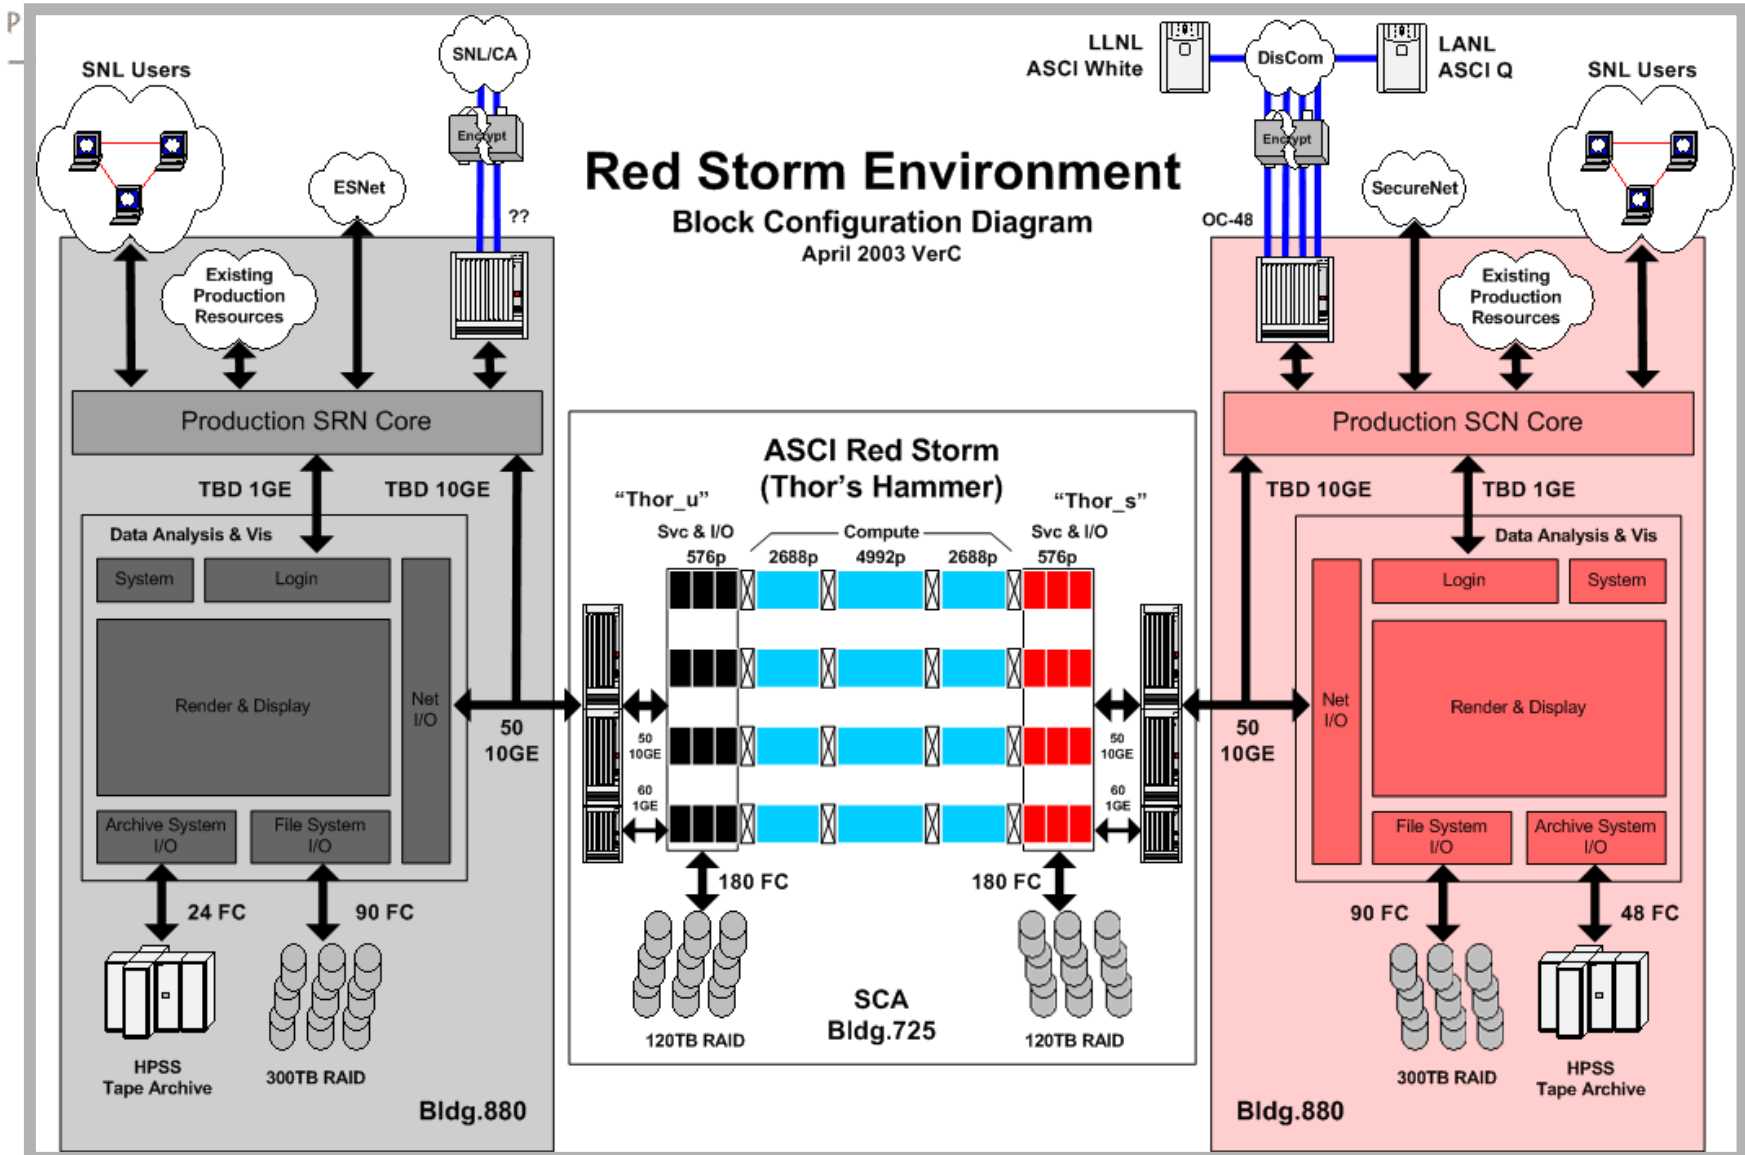



... for pre & post processing, data manipulation, archival services, etc.

... Red Storm is a primary driver, but not the only one



Red Storm + Vplant Clusters





PETROGLYPHS TO PETAFLOPS
CUG 2005

“Vplant” – We are making substantial use of clusters for data analysis and visualization



Unclassified (Feynman)

- Demonstrate technology scalability
- Development platform for Scalable Vis and Data Services
- Archiving into production
- Includes 128 graphics nodes for parallel rendering, 80 compute nodes.
- 15 TB of RAID storage, 2.3 GB/s RAID bandwidth, ~1.5 GB/s PVFS bandwidth
- 0.2 PB of HPSS tape storage, 0.5 GB/s

Production Classified (Da Vinci)


- Core systems to become part of RedRoSE
- 8 TB of RAID storage, 2 GB/s RAID bandwidth, ~1 GB/s PVFS bandwidth
- 0.2 PB of HPSS tape storage, 0.6 GB/s


Production Classified for VIEWS Corridor (Wilson)

- 64 render nodes
- 2 TB of distributed storage

JCEL cluster (RoSEbud)

- 12 nodes each with graphics & dual-ported display outputs (drives 6x4 powerwall)





PETROGLYPHS TO PETAFLIPS
CUG 2005

Building toward ...


Red Storm Environment (RoSE) Clusters

– Current RoSE HW on hand



- 136 Dual Xeon server nodes
 - 3.06 GHz, 4 GB RAM
 - Storage and 10GigE nodes
- 264 Dual Xeon graphics nodes
 - 3.6 GHz EM64T, 4 GB RAM
 - Visualization nodes
- 4x InfiniBand Interconnect
 - ~800 Mbytes/s measured bandwidth per port
- 480 TB SATA RAID, ~ 20 Gbyte/s bandwidth






Delivering the Red Storm Environment (configurations targeted for spring/summer 2005)



PETROGLYPHS TO PETAFLIPS
— CUG 2005 —

- **Classified: RedRoSE (InfiniBand)**
 - 8-12 GB/s file system bandwidth, ~200 TB capacity
 - 8 GB/s network bandwidth to Red Storm
 - 264 visualization nodes
 - 20 Tape drives for DSA, Silo has ~2 PB capacity
- **Unclassified: Feynman (Myrinet)**
 - 4-6 GB/s file system bandwidth, ~100 TB capacity
 - 8 GB/s network bandwidth to Red Storm
 - 128 visualization nodes, 80 compute nodes
 - 16 Tape drives for DSA, Silo ~2 PB capacity



RoSE Targeted Capabilities


(if/when fully deployed, Red & Black)



PETROGLYPHS TO PETAFLOPS
CUG 2005

		RoSE
Total Processors		1024 (512 nodes)
Interconnect Bandwidth	Line speed:	8 Gb/s
Visualization Rate	triangles/sec:	4×10^9
Storage Rate	Storage system:	25 GB/s
	parallel file system:	25 GB/s
Storage Capacity		300 TB

Scalability of parallel file system is a key challenge




PETROGLYPHS TO PETAFLUPS
CUG 2005

The VPlant clusters must support a large community of end users and facilities for a range of capability data / visualization needs.



- The “Wylie Theorem” ... the sweet spot for scaling interactive visualization of unstructured mesh data is at about 1 million elements per graphics node.
- 256 graphics nodes will support:
 - 1 user at 250 million elements
 - 1 user at 100 million elements + 3 at 50 million elements (4 simultaneous users)
 - 6 simultaneous users at 40 million elements
 - 12 simultaneous users at 20 million elements




Enabling the desktop is a fundamental part of the end-to-end strategy




PETROGLYPHS - PETROGLYPHS

DESKTOP REQUIREMENTS

- Technical Network Bandwidth (1 Gbit/s)
- Secure Environment (Need to Know/SRD)
- COTS-based high performance desktop visualization
- COTS and ASCI Modeling & simulation tools supported
- Seamless access to tools & high-end compute, data & visualization services





Advanced Visualization Tool Development




PETROGLYPHS TO PETAFLIPS
CUG 2005


Remember then ...

- Vis the old fashioned way

And now, as a result of our scalable visualization efforts ...

- ParaView examples
 - Parallel rendering (with desktop delivery)
 - Cone with rendering mode comparison
 - 470 million polygons the new way





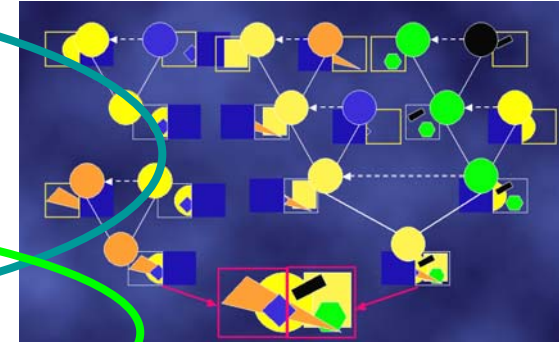
PETROGLYPHS TO PETAFLOPS
CUG 2005

Sandia Contributions to ParaView



Scalability




- Load Balancing and Data Distribution (D3)
- Tiled Display Rendering (ICE-T)
- Distributed Parallel Rendering (ICE-T)
- Image compression and Delivery (Squirt)
- Quadratic Cell Support (SHOE)
- Unstructured Volume Rendering (GATOR)



Desktop Delivery





PETROGLYPHS TO PETAFLOPS
CUG 2005

We expect to be ready for capability applications that use Red Storm



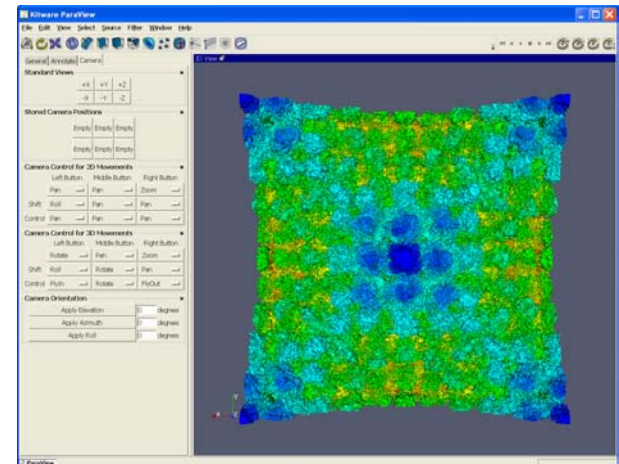
Recent rendering test using RoSE hardware


- 128 graphics nodes, Nvidia Quadro FX 3400 cards, PCI-Express
- “Stock” ParaView software (including IceT rendering and Squirt image delivery)
- 100 Million triangle test data (subdivided sphere)
- 100T network connection to the desktop

Results:

- 12-15 frames per second
- 1.5 Billion triangles/sec aggregate rendering performance

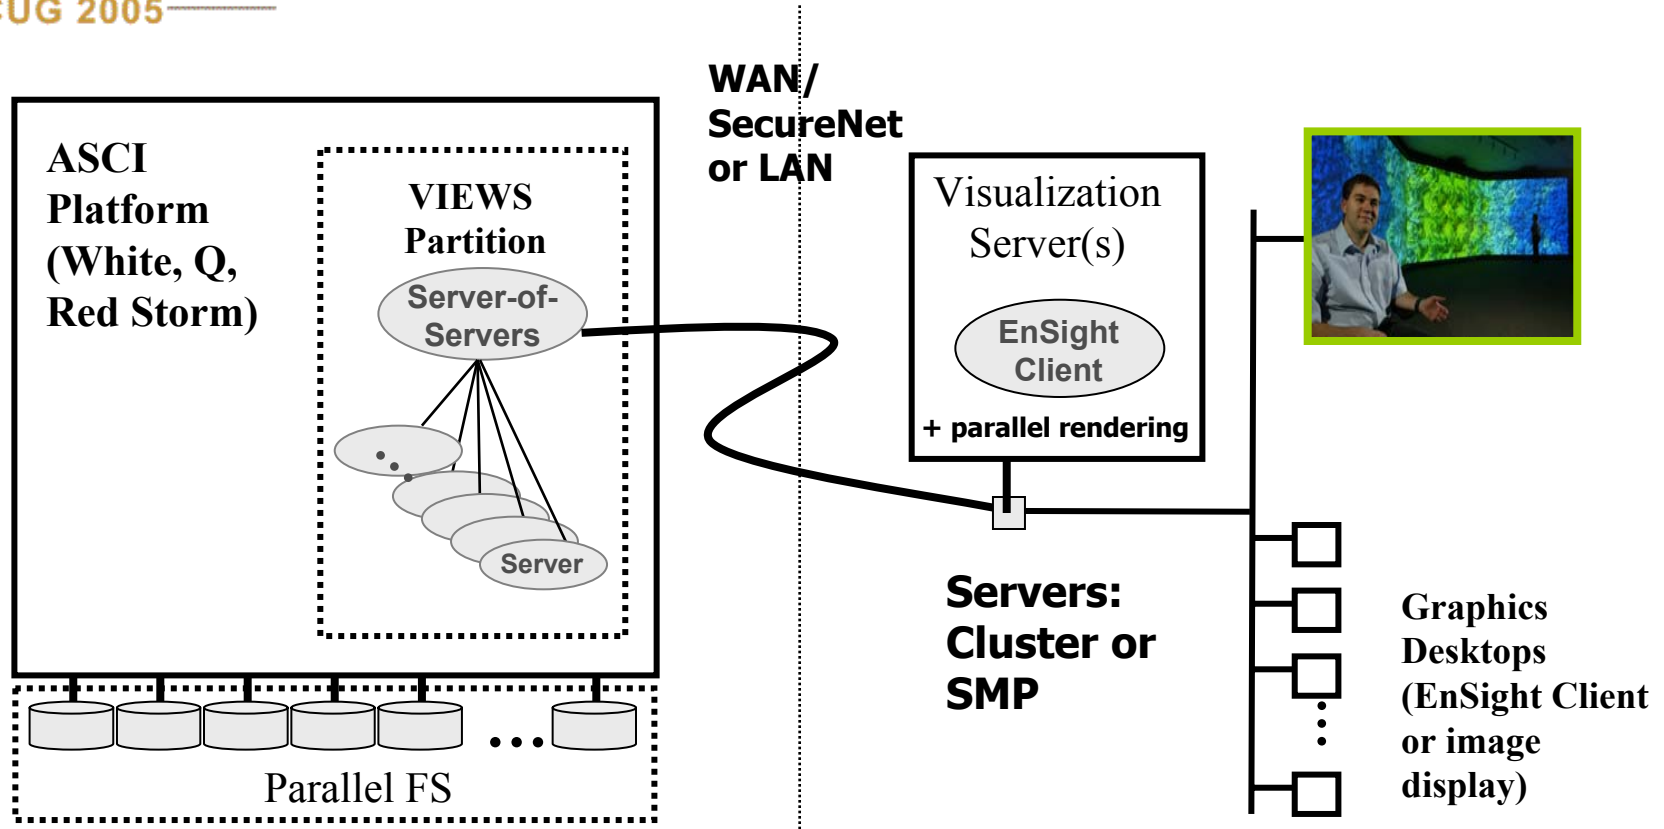
Aggregate rendering performance results reproduced for 470M triangle turbulence isosurface (~3fps)






PETROGLYPHS TO PETAFLUPS
CUG 2005

EnSight Gold will be provided as a core visualization tool ... some details still TBD



- EnSight Server-of-Servers now
- Parallel client (cluster-enabled) under development
- [This would also enable EnSight rendering on Red Storm VIEWS nodes]




Data Services / Tools



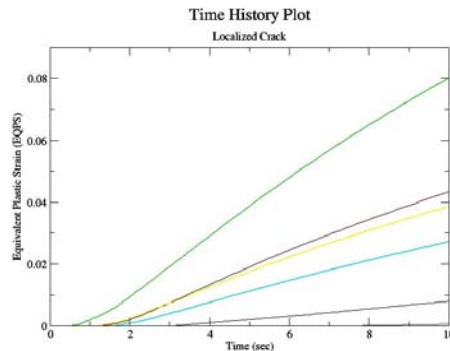
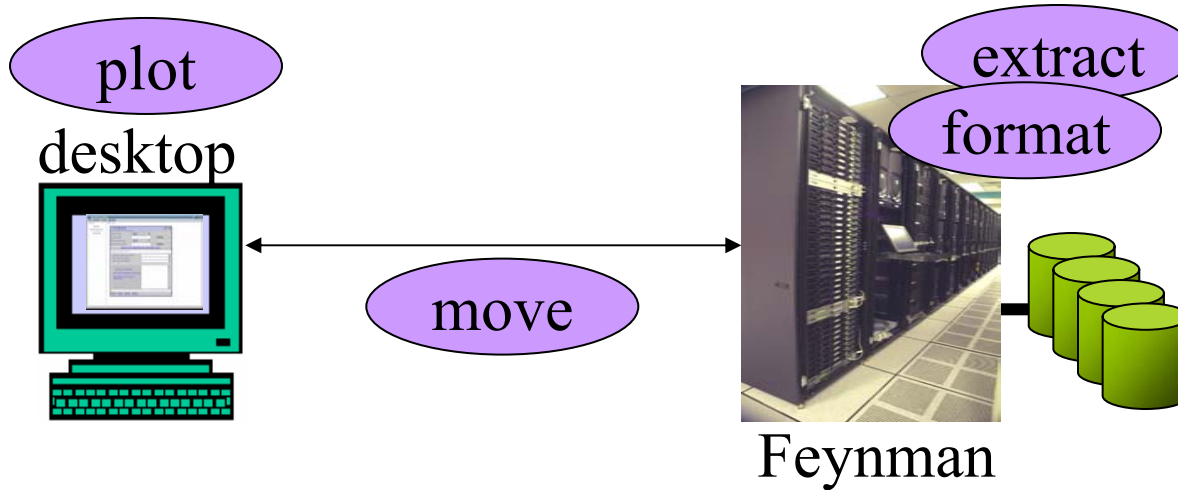
PETROGLYPHS TO PETAFLOPS
— CUG 2005 —

- **Support Post-Processing of Large / Complex Data Sets**
- **Parallel tools (leverage clusters)**
 - **Data movement and management**
 - **Data query**
 - **Data reduction**
 - **Data extraction**
 - **Data transformation**
 - **Data manipulation**
 - **Data mining**
 - **Data derivation (e.g., algebra)**
- **Access to high performance resources from the office/desktop**
- **Manipulation / migration of data objects, rather than complete data sets**
 - **10 TB @ 1Gb/s (peak) takes > 22 hours**
 - **1 GB @ 1Gb/s (peak) takes 8 seconds**




PETROGLYPHS TO PETAFLIPS
CUG 2005

Data Extraction and Plot Demo



- Extract Time History Data across Dataset Timeslices
- Format X,Y Data
- Move to desktop
- Plot
- ... all in a few seconds


[Demo Video](#)



Data Services ToolKit (DSTK) – Beta Release



- Python scriptable interface
- Functionality
 - Data extraction, output, math, filtering, selection, mesh modification, query, subsetting
- Integrated plotting
- Matlab output
- Exodus 2 and SAF 1.6 data formats
- Current release not fully parallel
 - internals are designed for parallelism
 - initial parallelism through parallel Python

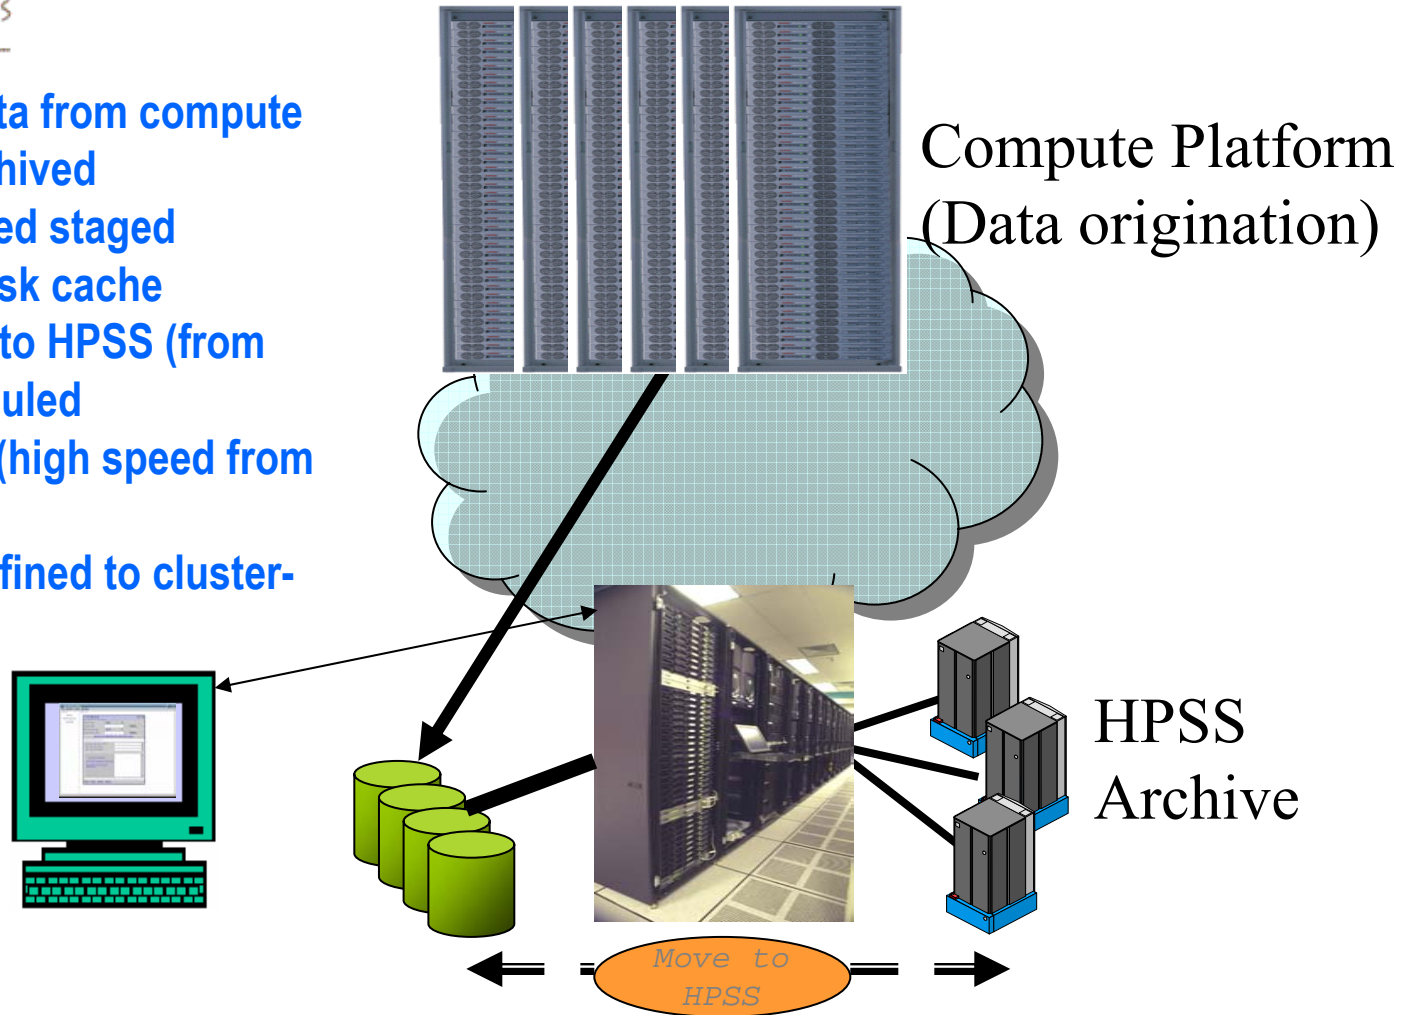



Data Service Archive (DSA)



PETROGLYPHS TO PETAFLOPS
— CUG 2005 —

- User requests data from compute platform to be archived
- Data to be archived staged through cluster disk cache
- Data transfers onto HPSS (from disk cache) scheduled
- Reliable transfer (high speed from Red Storm)
- HPSS clients confined to cluster-data-server





We are also working to integrate Scientific Data Management (SDM) tools



PETROGLYPHS TO PETAFLOPS
CUG 2005

SimTracker [Preferences](#) | [Help](#)

[Table of Contents](#) | [Summary Manager](#) | [PSG Manager](#) | **Summary Page**

[Previous](#) | [Next](#) | [Up](#) | [Delete](#) | [Customize Layout](#)

Calculation Summary [Edit](#)

Kernel: mmm Run Date: 05-02-2000 13:58
 Code Sys: Sample Run Status: Done
 Project: Viper Run Host: Blue
 Title: Viper Jet Formation
 Author: Wimer
 Comment: calculation of the viper jet formation problem in preparation for performing a hydro experiment.

Associated Files [Add](#) | **Thumbnail Snapshots** [Edit](#)

Input Files

- simt.in (.49KB)
- mminput (.73KB Compressed)

Log Files

- xxxlogmmm (.43KB)
- mmmbatch (.07KB)

Edit Files

- hcmmm100 (2.6KB)
- hcmmm101 (2.6KB)
- hcmmm200 (2.6KB)
- hcmmm201 (2.6KB)

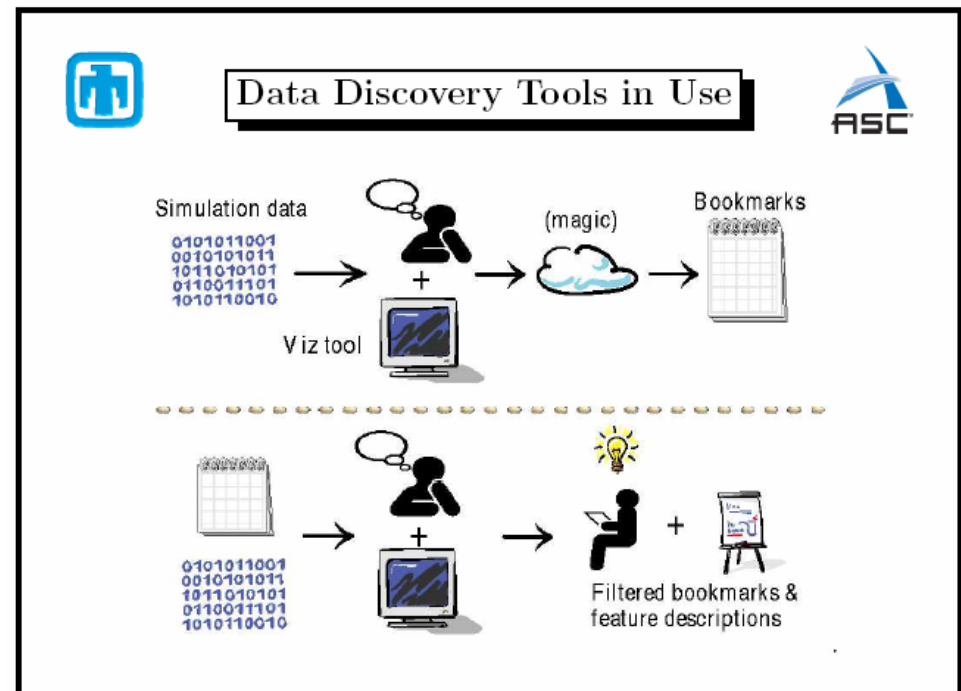
Science Files


- shcmmm000 (2.6KB)
- shcmmm100 (2.6KB)
- shcmmm200 (2.6KB)
- shcmmm300 (2.6KB)

[Previous](#) | [Next](#) | [Up](#) | [Delete](#) | [Customize Layout](#)

Tools for metadata-based management of simulation results (“SimTracker”) – currently undergoing “DART” (Design through Analysis Realization Team) assimilation at Sandia

Research in “Data Discovery” – e.g., prototype saliency-based feature identification and “Lookmarks”






PETROGLYPHS TO PETAFLOPS
— CUG 2005 —

When fully deployed, the Red Storm Environment will be new and improved



...

- All the production data analysis already accustomed to
- Plus
 - Improved I/O environment
 - Scalable data/vis infrastructure
 - Scalable visualization tools
 - Scalable data service tools



PETROGLYPHS TO PETAFLIPS
CUG 2005



END