

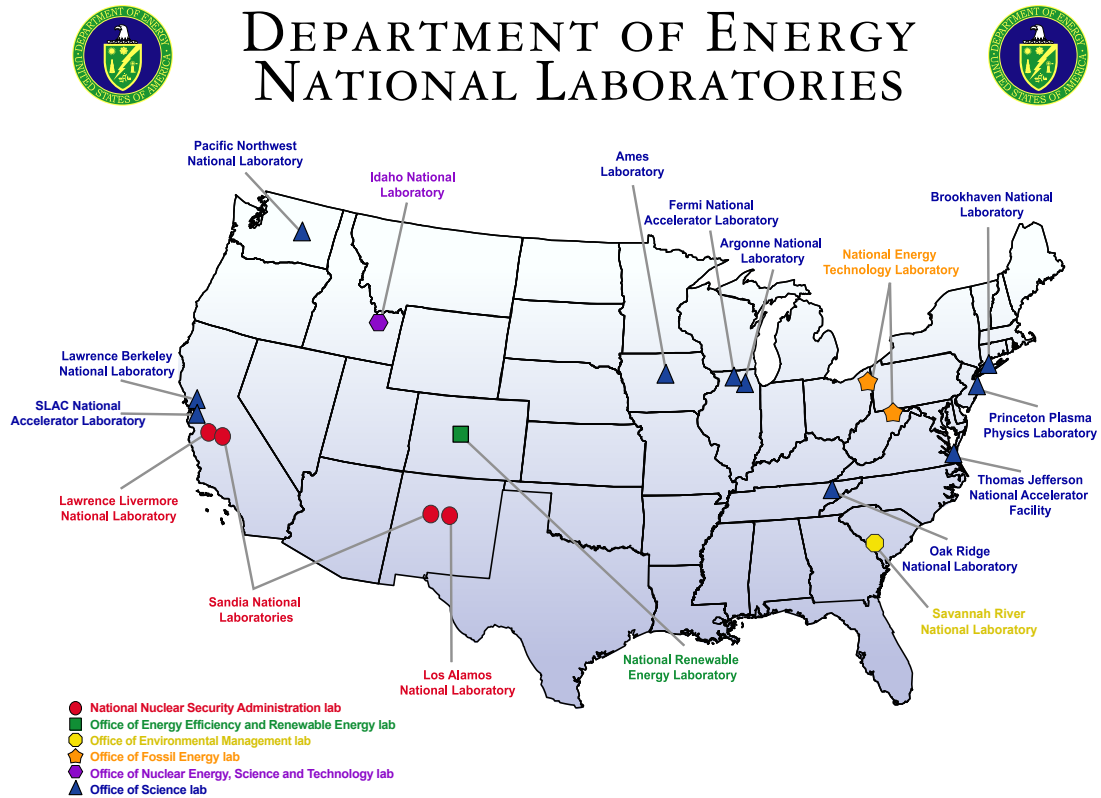
LANL Site Report DOE CGF 2011

Jon Woodring, LANL

Who We Are

- **CCS-7 group**
 - Jim Ahrens, team lead
 - R&D Visualization team
 - “Data Science at Scale”
 - ASC, OSC ASCR and BER, and LDRD
- **HPC-5 group**
 - Laura Monroe, team lead
 - Production Visualization team
 - Hardware, facilities, software and support
 - ASC and institutional funding

ASC core funding



HPC-5 Activities in the Past Year

- **Redesigning the Vis Corridor**
 - Viewmaster2
 - New video extenders
- **TLCC visualization**
- **Visualization Projects**
 - Thermonuclear Burn Initiative
 - ASC
 - ASCEM with LBL and Savannah River
- **Blue Room redesign (LANL classified museum)**
- **Interface investigation (touch table, haptic, and pen interfaces)**
- **ASC Exascale planning**

HPC-5 Viewmaster2

■ Viewmaster2

- 183-node rendering cluster
 - Dual Intel Westmere-EP, X5650, 6 cores each 2.66 GHz
 - Mellanox IBA 4x QDR fat tree interconnect
 - 200 PB Panasas dedicated file system, also mounts global Panasas file system
 - 60 Back-end Rendering Nodes
 - 2 NVIDIA Quadro 6000s (6 GB GDDR5 on-board memory), 96 GB DDR3-1333
 - 10 Advanced User Nodes
 - 1 NVIDIA Quadro 6000, 96 GB DDR3-1333 memory
 - 50 Standard User Nodes
 - 1 NVIDIA Quadro 5000 (2 GB GDDR5 on-board memory), 24 GB DDR3-1333
 - 37 Facilities Nodes to support CAVE, PowerWall, etc
 - 1 NVIDIA Quadro 6000, with g-sync daughterboard for framelock, 96 GB DDR3-1333
 - 12 I/O, 2 RPS, 2 LSM support nodes

HPC-5 Visualization Support

- **Conejo**
 - SGI XE 1300
- **Mapache**
 - SGI XE 1300
- **Lobo**
 - 4x4 AMD Opteron cluster
- **Cerrillos (Roadrunner for IC)**
 - AMD Opteron + IBM Cell blade cluster
- **Cielo**
 - Cray XE6
- **Tri-Lab Vis Software Production: EnSight, ParaView, and VisIt**

CCS-7 Activities in the Past Year


- **ParaViewWeb**
- **ASC Level 2 Milestone with Sandia**
- **New CCS-7 research cluster (replaces DQ, the Viewmaster prototype)**
 - Multiple mixed GPUs (Fermis & AMD), large mechanical drives and SSDs
 - Connected to 6x3 powerwall
- **We like “fat nodes” too, 32+ core machines – 4GB per core (128 GB)**
- **IKS Review**
 - Much more: Switch to poster used at IKS review

Today

VPIC

- Magnetic Reconnection
- Extremely Large Data (~5 billion cells)
- Parallel Analysis
- Visualization performed by end user
- ParaView framework used for Remote Visualization

W. Daughton et al. "Role of Electron Physics in the Development of Turbulent Magnetic Reconnection in Collisionless Plasmas" Accepted to *Nature Physics*.




Visualization of VPIC data

Visualization of VPIC data

Multiphase Flow with Interface eXchange (MFIx)


- Carbon Capture Simulation Initiative (CCSI)
- Parallel model
- Reactive Flow



Visualization of MFIx dataset

Cosmology

- Verification and Validation
- Large Data
- Particles
- In-situ analysis




Dark matter cluster visualization

J. Ahrens et al. "Verification of the Scientific Simulations via Hypothesis-Driven Comparative and Quantitative Visualization". *IEEE Computer Graphics and Applications*, 30, 6, November/December 2010.

Cosmology visualization in ParaView

Climate

- Ultra-scale Visualization Climate Data Analysis Tools (UVCDAT)
- Earth Systems Grid (ESG)
- Ocean Modeling (POP)



Let p be the number of processors available, s the total number of time steps, m and n be factors of p such that $n \cdot m = p$, n is the number of processor groups and m is the size of each group.


Time parallelism for UVCDAT

UVCDAT work to be submitted to LDAV 2011

Windblade

- Alternative Energy
- Computational Data
- Experimental Data
- Turbine-turbine interference
- Feature detection and visualization

Work being submitted to *IEEE Visualization 2011*



Lambda 2 vs. Absolute vorticity for identifying vortices

Smoke tracer in wind turbine simulation

Current Research

Analysis for Emerging Architectures

- Heterogeneous Systems
 - GPU, CPUs, etc.
- Homogeneous Systems
 - Multi-core general purpose CPUs
- Thrust Library
- Manta in ParaView

"A Report Documenting the Completion of the Los Alamos National Laboratory Portion of the ASC Level II Milestone "Visualization on the Supercomputing Platform", *ASC Level II Milestone Meeting*, August 2010.

Analysis for Complex Data

- Feature Identification
- Feature Extraction
- Qualitative
- Quantitative

S. Williams, et al., "Global Eddy Analysis and Visualization". Accepted to *Eurographics/IEEE-VGTTC Symposium on Visualization* (2011).

Data Intensive Visualization

- Large data storage and retrieval
- Time Series
- Distributed file system support
- Increase the aggregate I/O bandwidth

C. Mitchell, et al., "VisIO: Enabling Interactive Visualization of Ultra-Scale, Time Series Data via High-Bandwidth Distributed I/O Systems". *IEEE International Parallel and Distributed Processing Symposium*, May 2011.

Reducing Data at Run-time


- Sampling
- In-situ analysis
- Run-time data reduction
- Statistical unbiased estimators
- Quantifiable

J. Woodring et al., "In-situ Sampling of a Large-Scale Particle Simulation for Interactive Visualization and Analysis". Accepted to *Eurovis 2011*.


Real-time interactivity

- Multi-Resolution Streaming
- Level of Detail Scaling
- Streaming
- Prioritization
- Builds on Sampling and Compression work


J. Ahrens et al. "Interactive Remote Large-Scale Data Visualization via Prioritized Multi-resolution Streaming", *Workshop on Ultrascale Visualization*, Nov. 2009.




Rendering with Manta raytracer with many-core CPU





Simulation with contour calculation using Thrust on GPU





Global eddy census, showing the average depth of eddies averaged per 1 degree of latitude and longitude, averaged across 350 daily snapshots



Metrics for finding eddies in ocean



Visualization pre- and post-reduction by statistical sampling



Multi-resolution, streaming, prioritization, high resolution where needed