Longhorn Project
TACC's XD Visualization Resource

DOE Computer Graphics Forum

April 14, 2010

Longhorn Visualization and Data Analysis

- In November 2008, NSF accepted proposals for the Extreme Digital Resources for Science and Engineering
- The Longhorn project was proposed as a next generation response to TeraGrid's growing visualization and data analysis needs

TACC

Spur - Visualization System

- 128 cores, 1 TB distributed memory, 32 GPUs
- `spur.tacc.utexas.edu` login node, no GPUs
  don't run apps here!
- `ivisbig.ranger`
  Sun Fire X4600 server
  - 8 AMD Opteron dual-core CPUs @ 3 GHz
  - 256 GB memory
  - 4 NVIDIA FX5600 GPUs
- `ivis[1-7].ranger`
  Sun Fire X4440 server
  - 4 AMD Opteron quad-core CPUs @ 2.3 GHz
  - 128 GB memory
  - 4 NVIDIA FX5600 GPUs

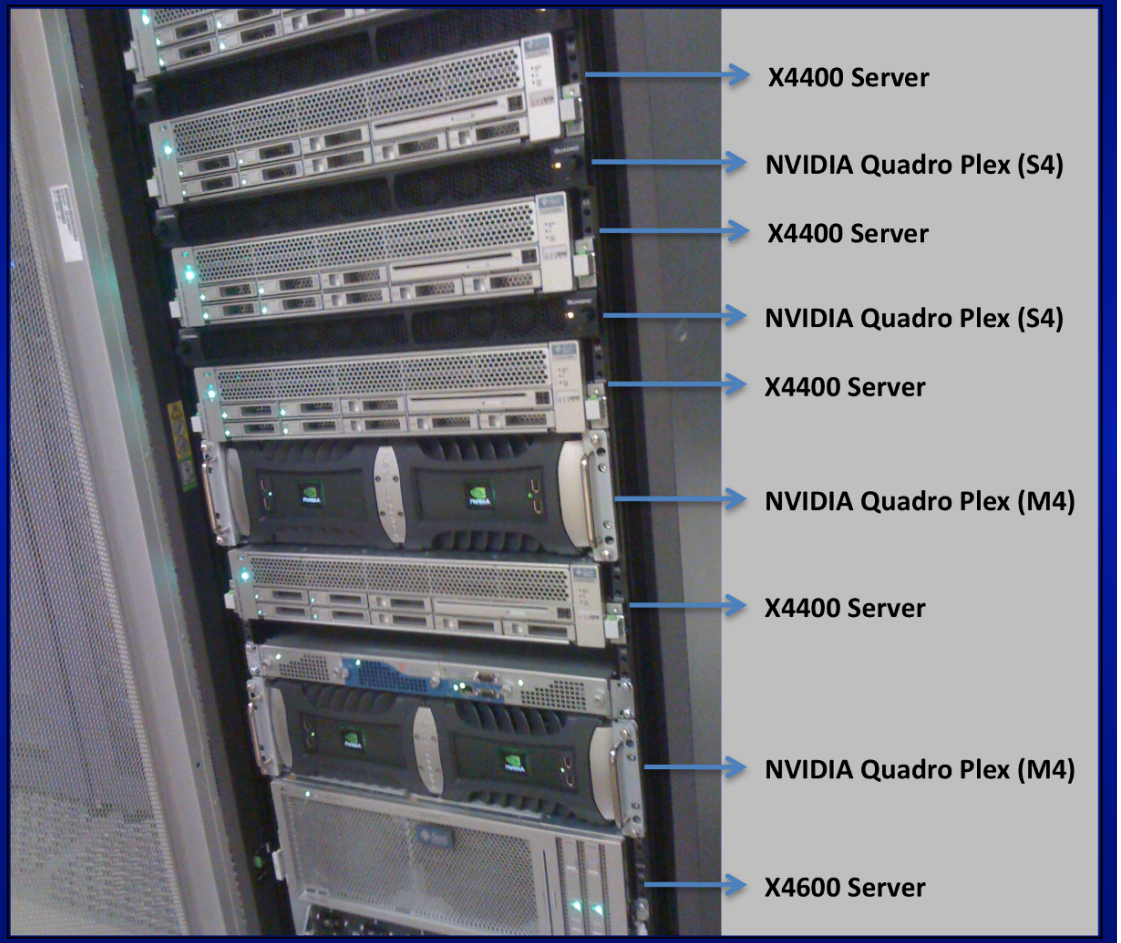

TACC

Spur / Ranger topology

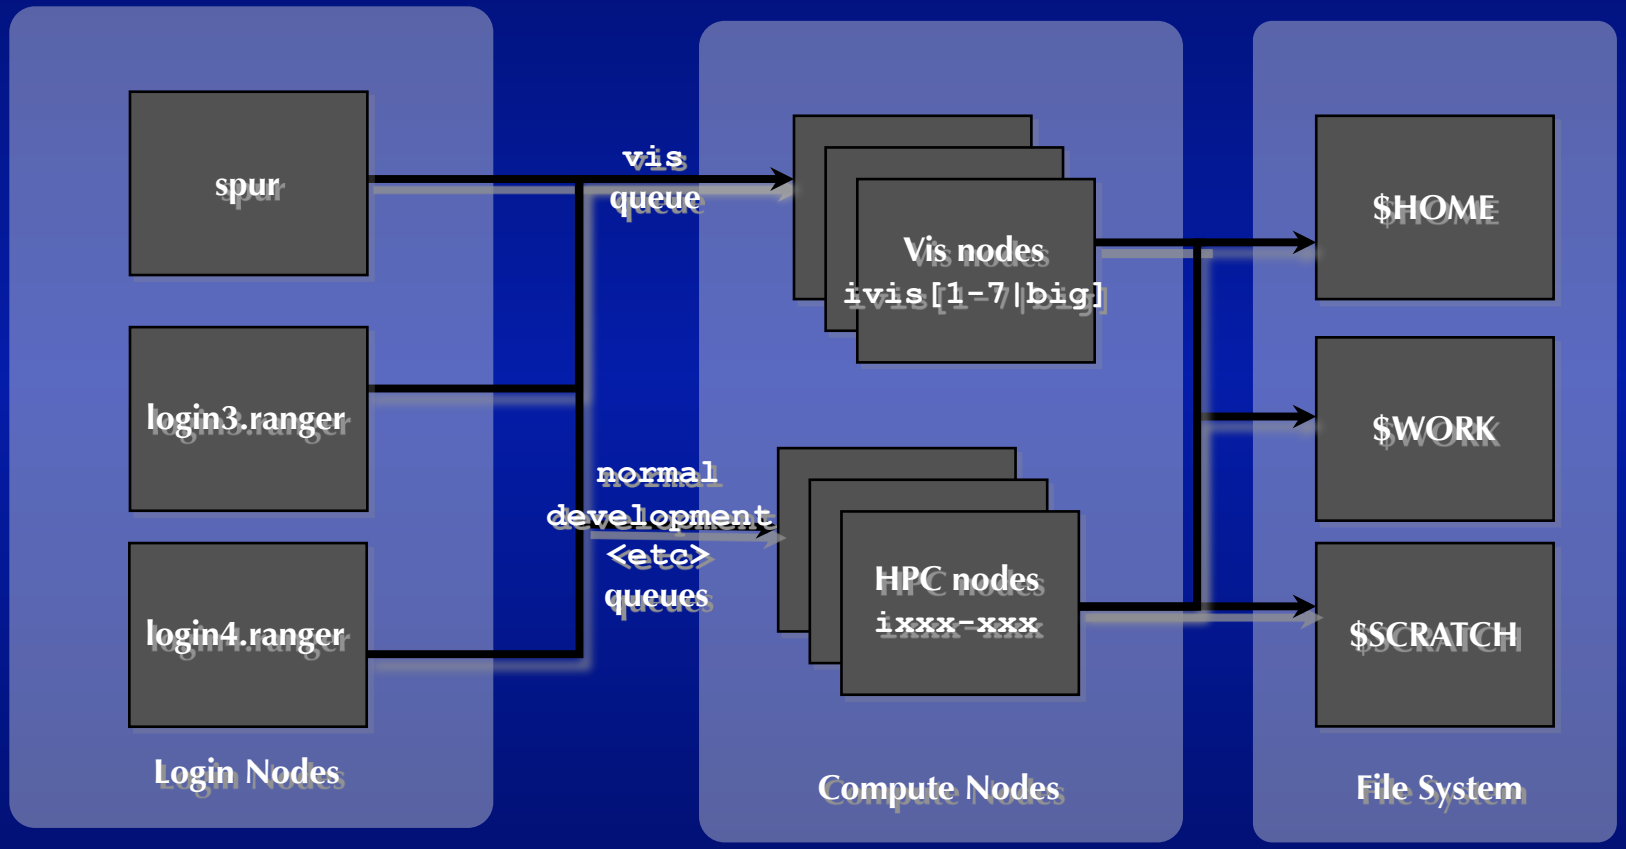

TACC

Connecting to Spur

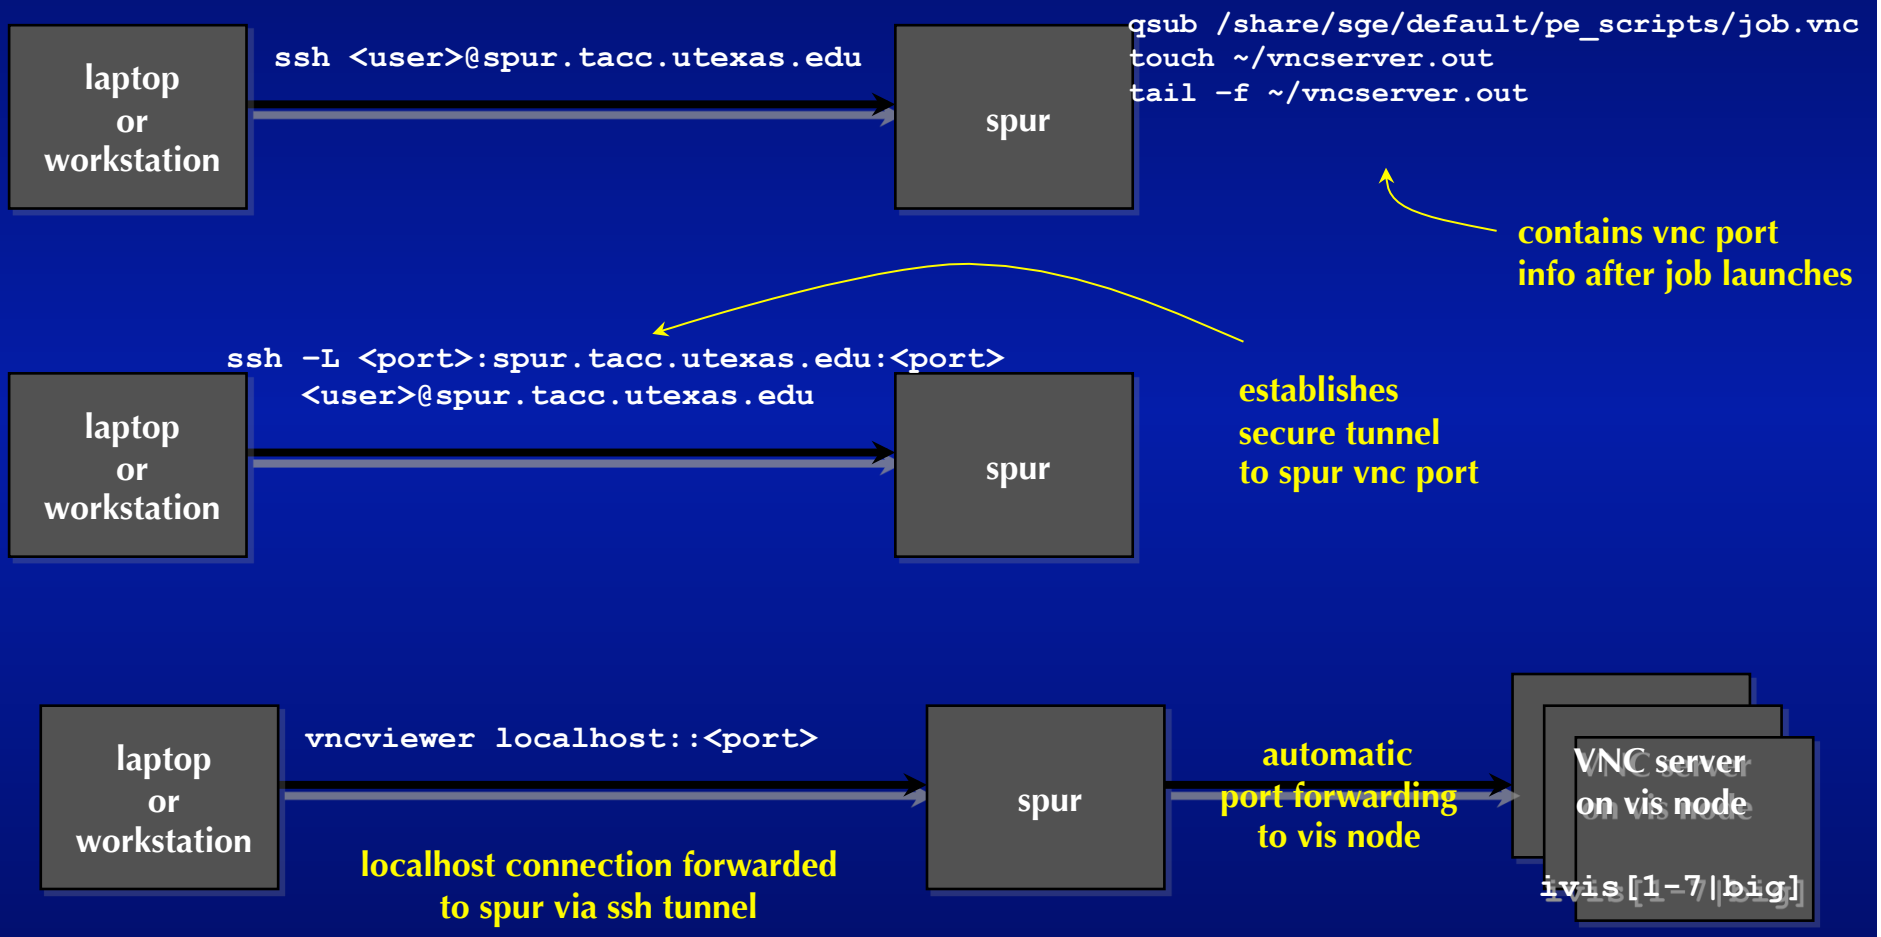

Spur Usage

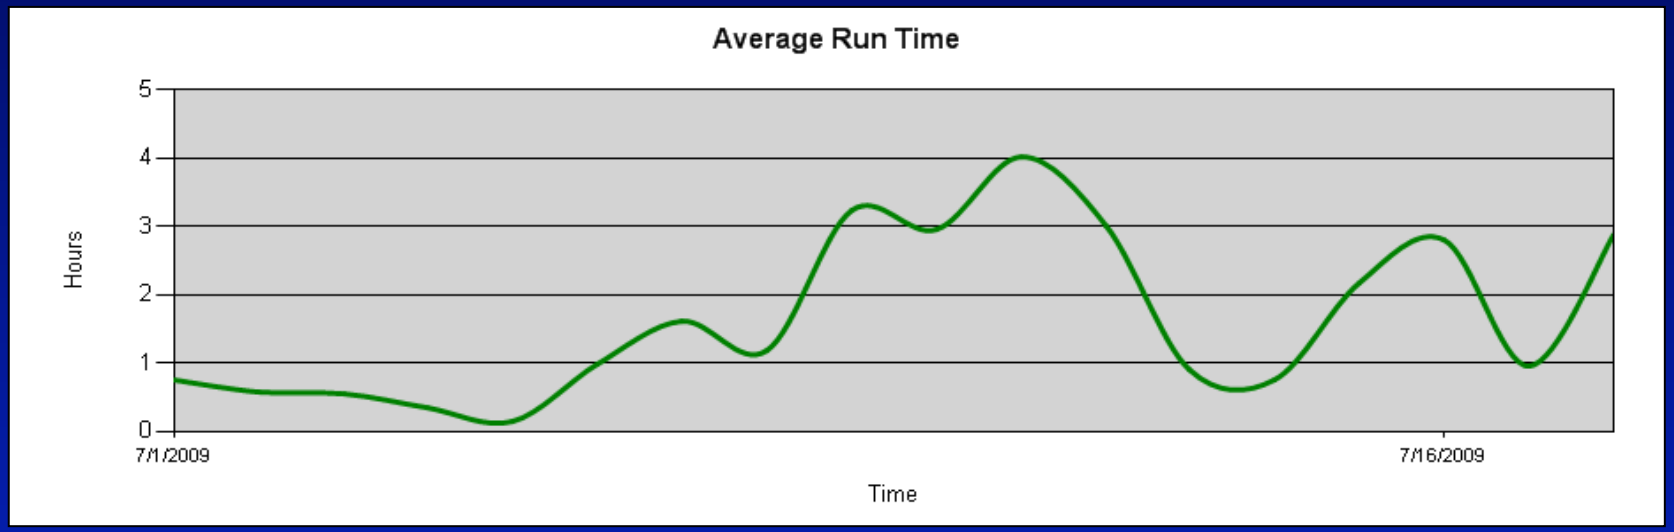

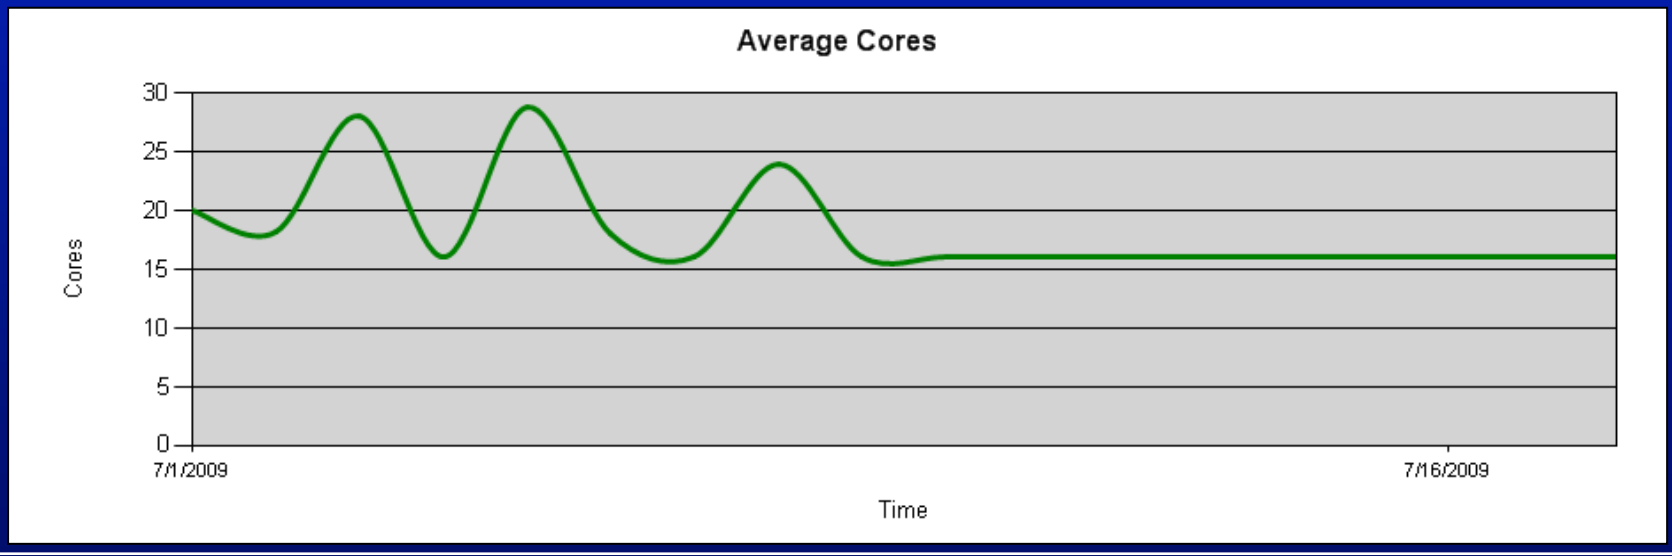

TACC

XD Vis Requirements Analysis

- Surveyed members of the science community via personal interviews and email surveys
- Received ~60 individual responses

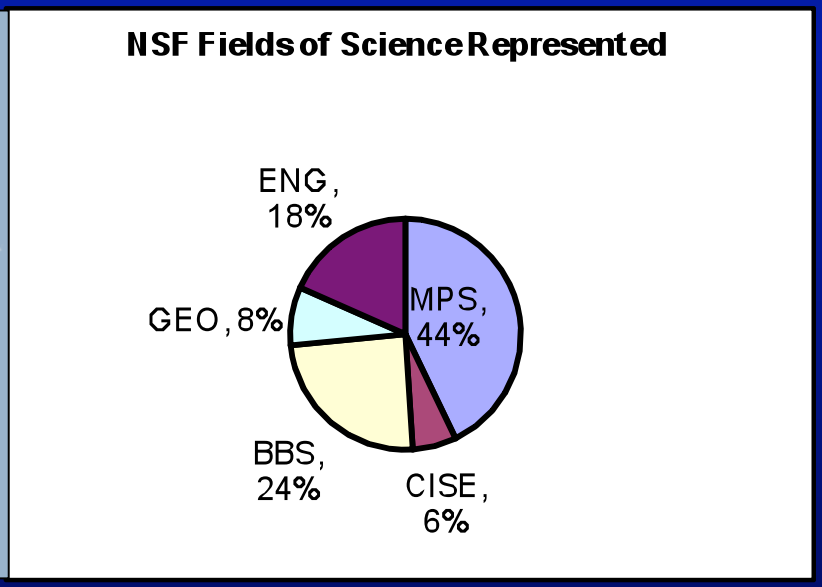

XD Vis Requirements Analysis

Requirement	% Users Requested
User Support and Consulting	96%
Large-Scale DAV Tools/Resources	39%
Remote/Collaborative DAV Services	27%
Computational Steering	10%
In-simulation DAV Tools	6%
Tools for 3D Measurement and Query	6%
Tools for Multiple Length and Time Scales	6%

(DAV = Data Analysis/Visualization)

TACC

Longhorn Configuration

(256 Nodes, 2048 Cores, 512 GPUs, 14.5 TG Aggregate Memory)

- 256 Dell Quad Core Intel Nehalem Nodes
  - 240 Nodes
    - Dual socket, quad core per socket: 8 cores/node
    - 48 GB shared memory/node (6 GB/core)
    - 73 GB Local Disk
    - 2 Nvidia GPUs/node (FX 5800 - 4GB RAM)
  - 16 Nodes
    - Dual socket, quad core per socket: 8 cores/node
    - 144 GB shared memory/node (18 GB/core)
    - 73 GB Local Disk
    - 2 Nvidia GPUs/node (FX 5800 – 4GB RAM)
  - ~14.5 TB aggregate memory
- QDR InfiniBand Interconnect
- Direct Connection to Ranger's Lustre Parallel File System
- 10G Connection to 210 TB Local Lustre Parallel File System
- Jobs launched through SGE

Longhorn's Lustre File System ($SCRATCH)

- OSS's on Longhorn are built on Dell Nehalem Servers Connected to MD10000 Storage Vaults
- 15 Drives Total Configured into 2 Raid5 pairs with a Wandering Spare
- Peak Throughput Speed of the File System is 5.86 GB/sec
- Peak Aggregate Speed of the File System is 5.43 GB/sec

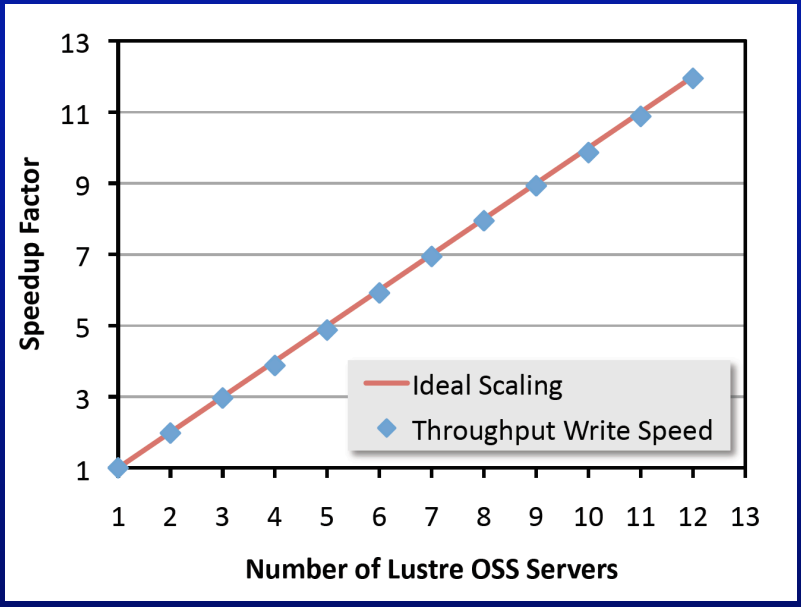

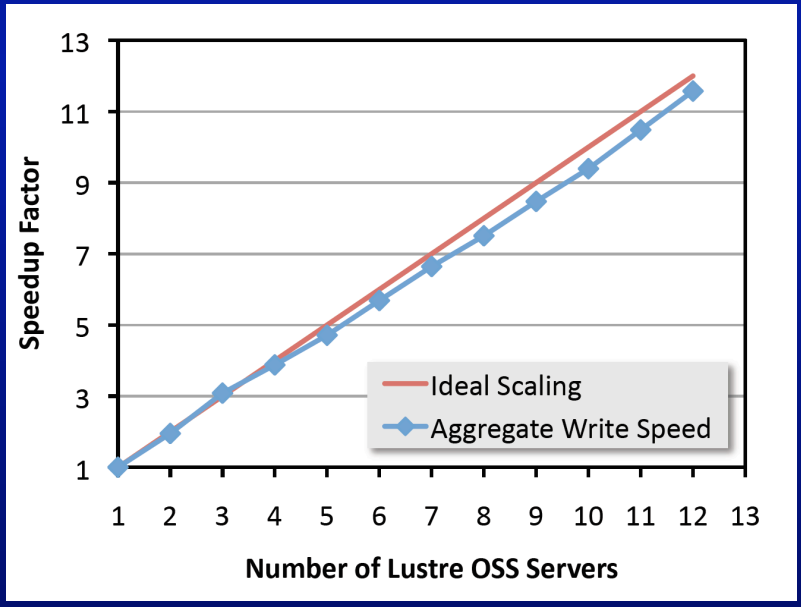

TACC

Longhorn Partners and Roles:

- TACC (Kelly Gaither – PI)
  - **Longhorn** machine deployment
  - User support
  - Visualization and Data Analysis portal development
  - Software/Tool development
- NCAR (John Clyne – CoPI)
  - User support
  - VAPOR Enhancements
- University of Utah (Valerio Pascucci – CoPI, Chuck Hansen)
  - User support
  - Software Integration of RTRT and topological analysis

TACC

Longhorn Partners and Roles:

- Purdue University (David Ebert – CoPI)
  - User support
  - Integration of visual analytics software
- UC Davis (Hank Childs – Chief Software Integration Architect)
  - Directly facilitate tools being integrated into the VisIt software suite
- SURA (Linda Akli – MSI Outreach/Broadening Participation)

Longhorn Usage Modalities:

- Remote/Interactive Visualization
  - Highest priority jobs
  - Remote/Interactive capabilities facilitated through VNC
  - Run on 4 hour time limit
- GPGPU jobs
  - Run on a lower priority than the remote/interactive jobs
  - Run on 12 hour time limit
- CPU jobs with higher memory requirements
  - Run on lowest priority when neither remote/interactive nor GPGPU jobs are waiting in the queue
  - Run 12 hour time limit

Longhorn User Portal

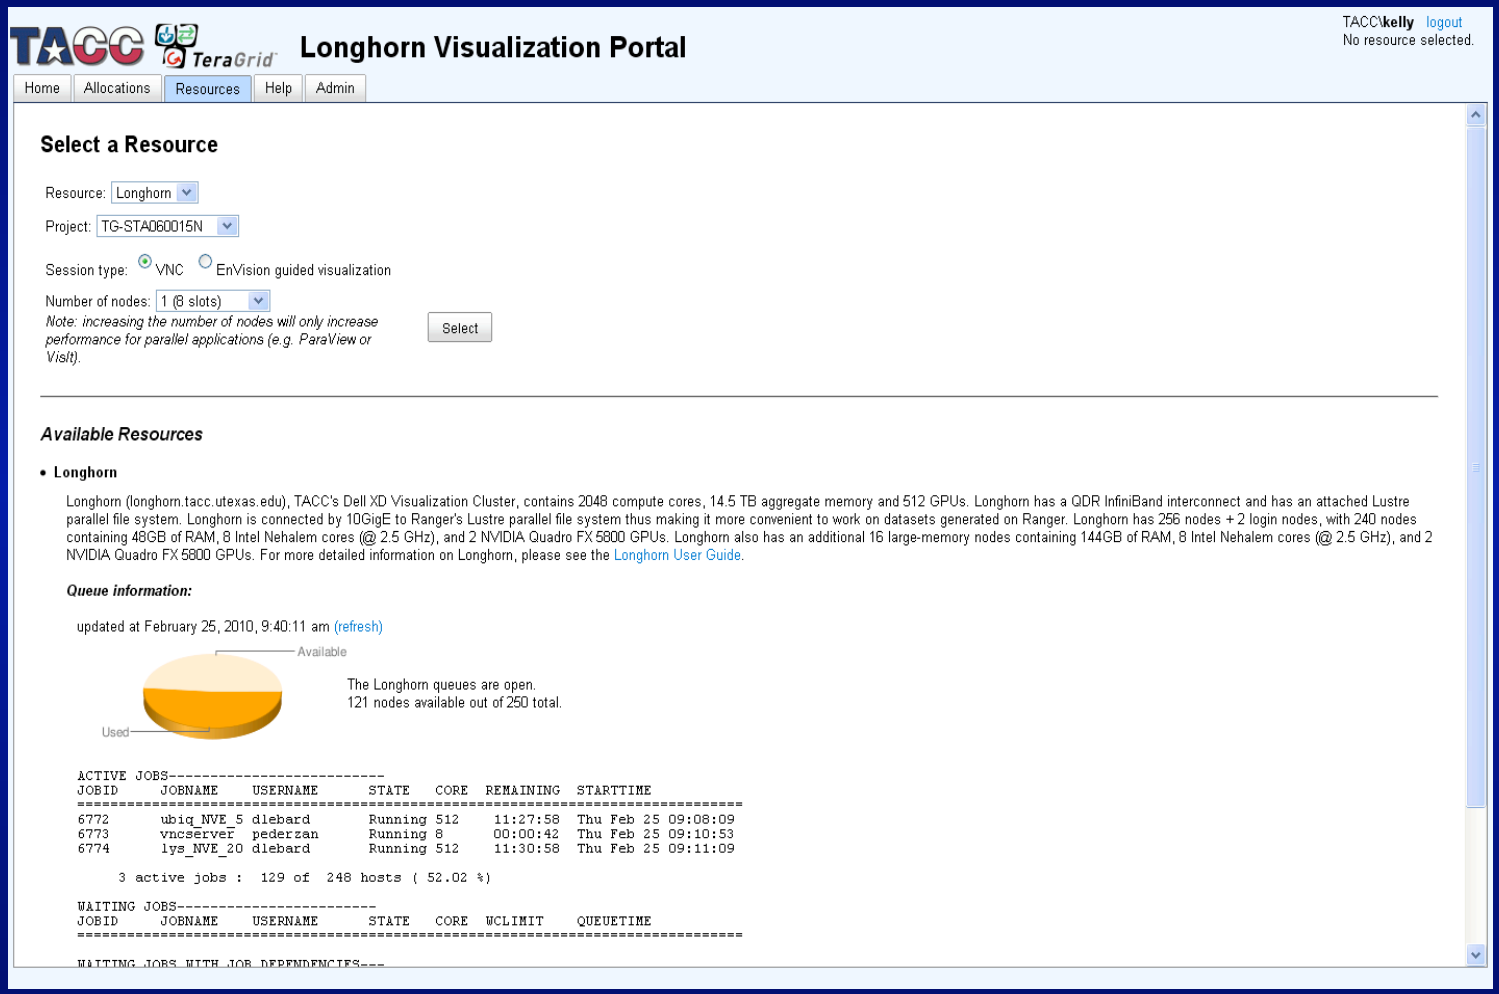

TACC

Longhorn Queue Structure

SGE Batch Environment Queues			
Queue Name	*Max Runtime*	*Max Cores*	*Node Pool*
normal	6 hrs	128	All nodes
long	24 hrs	128	All nodes
largemem	8 hrs	128	16 Large memory nodes
devel	1 hrs	32	8 Nodes
request	---	---	special requests

Project Types		
Type	*Purpose*	*Special Environment Modifications*
vis	Visualization jobs	
data	Data Analysis jobs	
gpgpu	GPGPU jobs	disables X server
hpc	HPC jobs	

qsub -q normal -P vis

TACC

Longhorn Usage

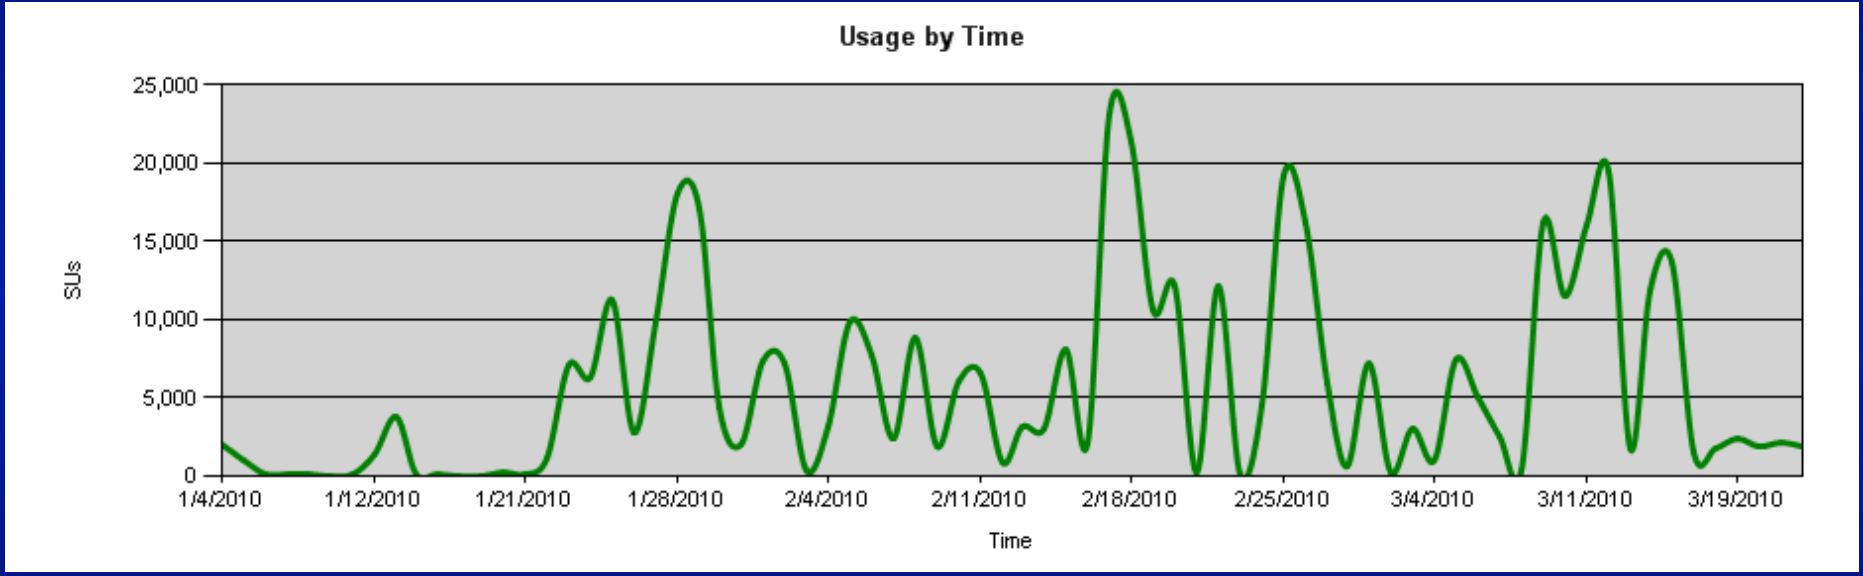

TACC

Longhorn Usage

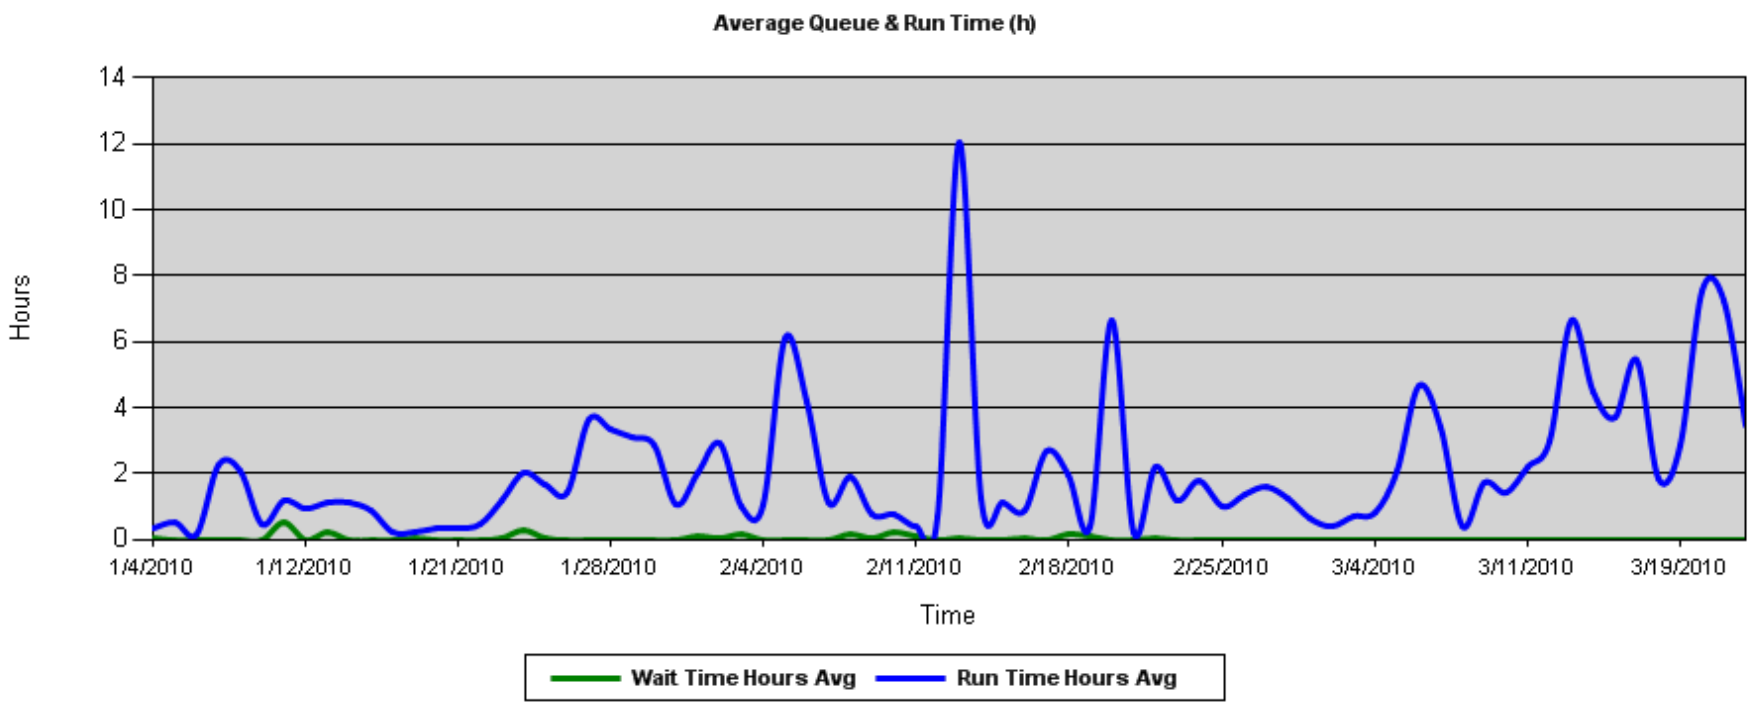

Longhorn Usage

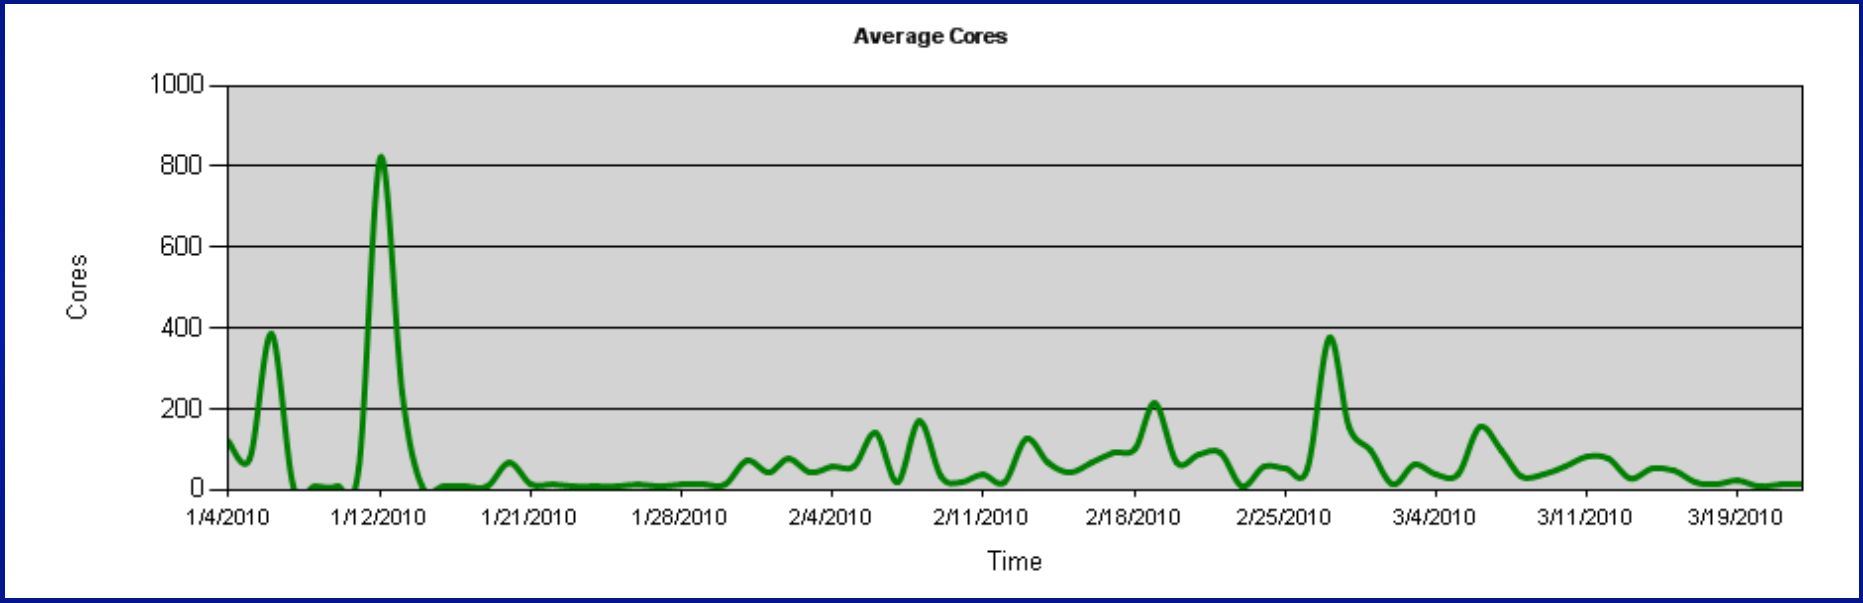

TACC

Longhorn Usage

Field of Science Since Production

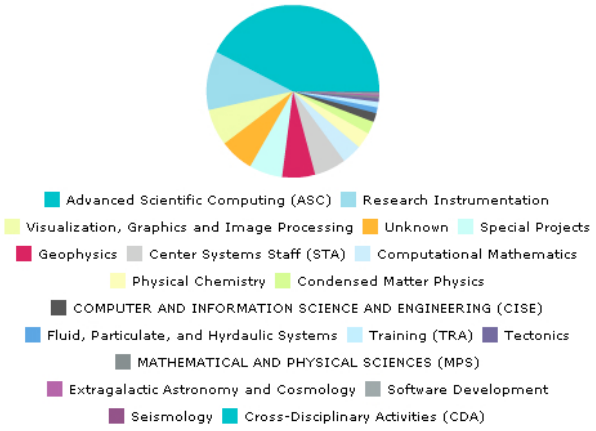

Field of Science Last 30 Days

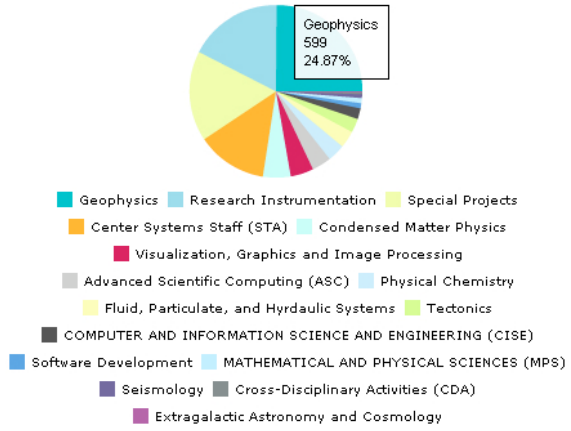

Field of Science Last 7 Days

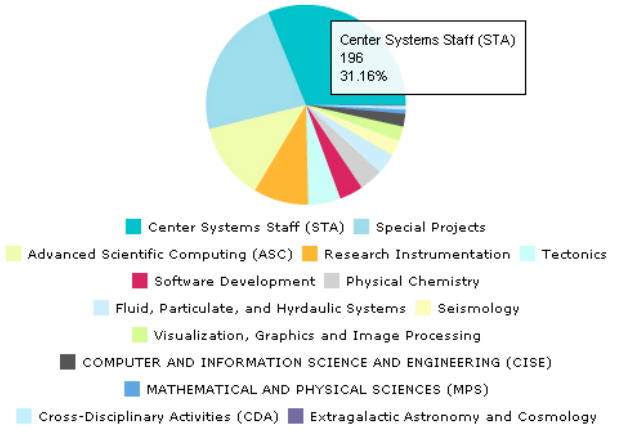

TACC

Sampling of Current Projects

- Computational Study of Earth and Planetary Materials
- Simulation of Quantum Systems
- Visualization and Analysis of Turbulent Flow
- A probabilistic Molecular Dynamics Optimized for the GPU
- Visualization of Nano-Microscopy
- MURI on Biologically-Inspired Autonomous Sea Vehicles: Towards a Mission Configurable Stealth Underwater Batoid
- Adaptive Multiscale Simulations

Visualizing and Analyzing Large-Scale Turbulent Flow

- Detect, track, classify, and visualize features in large-scale turbulent flow.
- Collaborative effort between TACC, VACET and domain scientists.
- Paper submitted to IEEE Vis 2010.
- Thanks to Cyrus Harrison for developing connected components code in VisIt.
- Thanks to Hank Childs for developing chord length distribution code and visualization.
- Thanks to Wes Bethel for facilitating the collaboration!

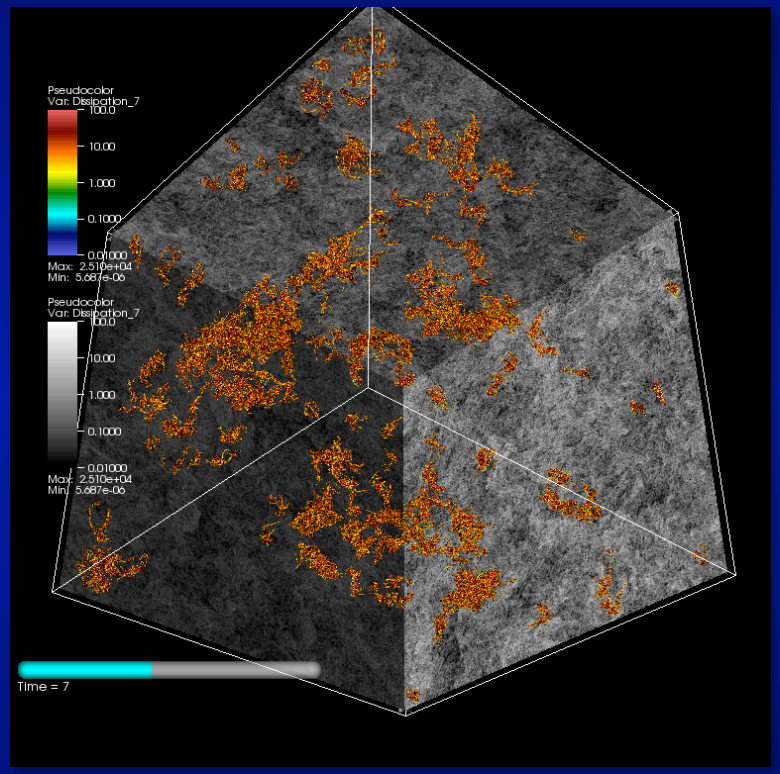

VisIt calculated connected components on a 4K^3 turbulence data in parallel using TACC's Longhorn machine. 2 million components were initially identified and then the map expression was used to select only the components that had total volume greater than 15. Data courtesy of P.K. Yeung & and Diego Donzis

Questions?

Kelly Gaither
kelly@tacc.utexas.edu

TACC