The University of Tennessee Center for Remote Data Analysis and Visualization (RDAV)

Sean Ahern, PI, Director – University of Tennessee
Jian Huang, co-PI, Associate Director – University of Tennessee
Wes Bethel, co-PI – Lawrence Berkeley National Laboratory
Scott Klasky, co-PI – Oak Ridge National Laboratory
Dave Semeraro, co-PI – National Center for Supercomputing Applications (NCSA)
George Ostrouchov, Senior staff – Oak Ridge National Laboratory
Miron Livny, Senior staff – University of Wisconsin

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation

Co-located at National Institute for Computational Sciences (NICS)

- NICS is a collaboration between UT and ORNL
- Awarded the NSF Track 2B ($65M)
- Phased deployment of Cray XT systems
- Home of Kraken, used to be #3 on Top 500

NSF Teragrid

- Under the auspices of NSF OCI (Office of Cyberinfrastructure)
- The world's largest distributed cyberinfrastructure for open science research
- 11 partner sites of integrated, persistent computational resources
- 2.5 petaflops, 50 petabytes storage, 100+ discipline specific databases -- connected through high speed network
- Entering its new era of XD: eXtreme Digital (2011 - 2016)

RDAV - Eyes of the Teragrid

- Provide remote and shared resources for the purpose of exploring/analyzing/ visualizing large scale data.
- Provide the ability to easily take advantage of remote and shared computing/ data storage infrastructure.
- Provide unique architecture for data analysis and visualization
- Leverage large amount of existing experience in deploying similar capabilities.
- Allocated through TRAC

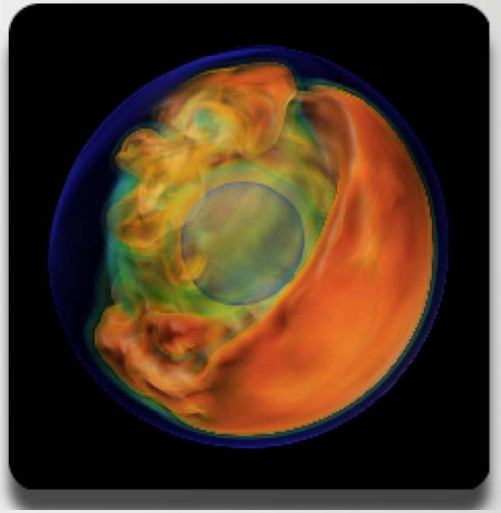

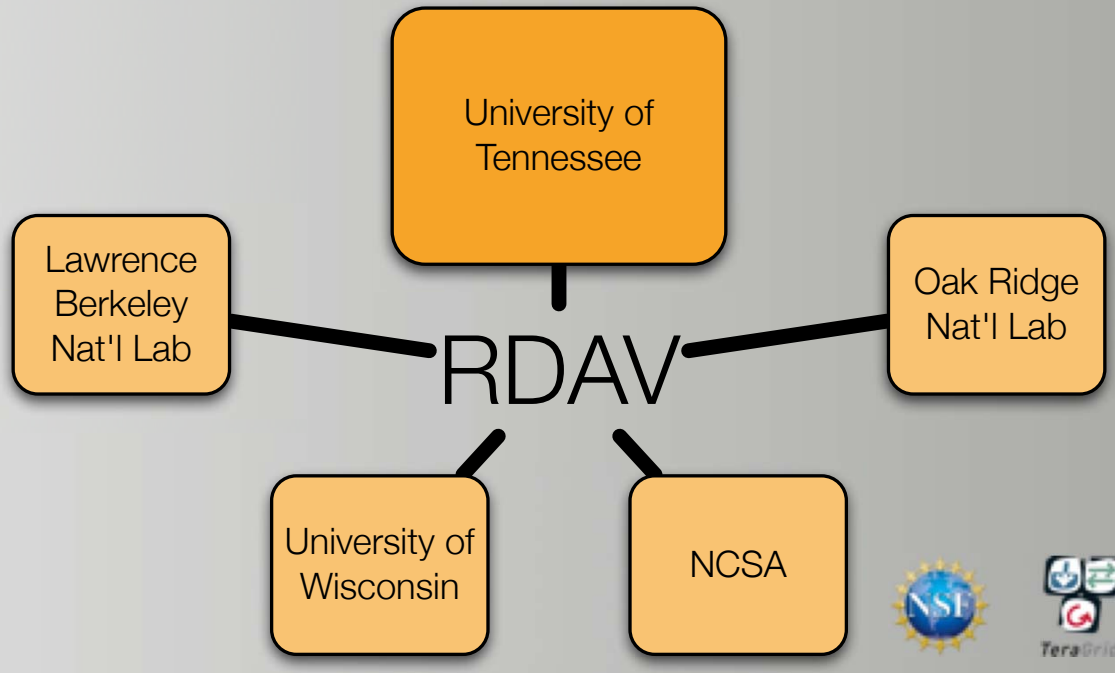

RDAV's Central Hardware: Nautilus - SGI Ultraviolet SMP

- Nautilus in stable production
- System configuration:
  - 1024 Nehalem cores
  - 2x racks, 64x blades/sockets
    - 4 TB memory
    - SLES 11
  - Infiniband 24x QDR
  - 4x 10 Gigabit Ethernet
  - 8-16x Nvidia Fermi Tesla
  - ~1 PB parallel file system

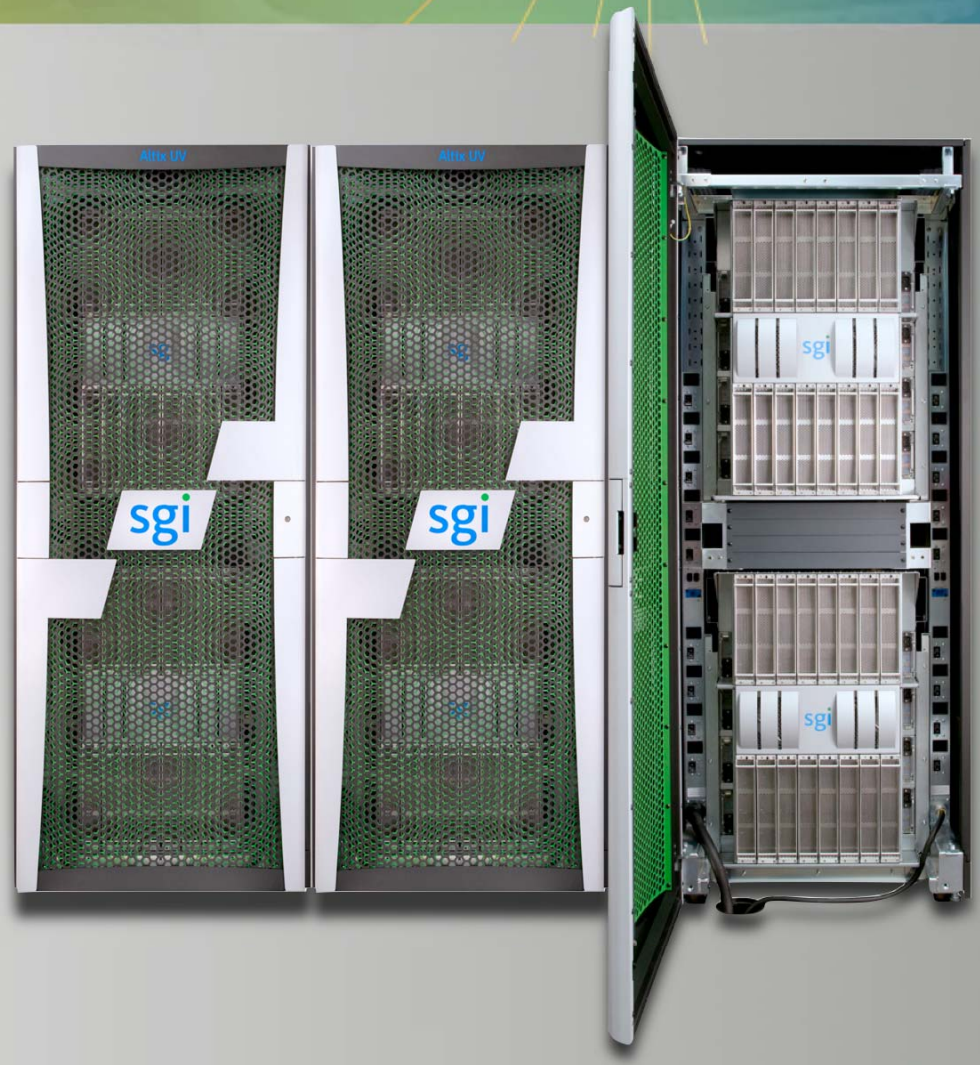

What's on each blade?

Each blade:
16 Intel® X7500 "Nehalem-EX" Cores & up to 128GB DDR3

I/O risers provide choice of expansion slot capabilities

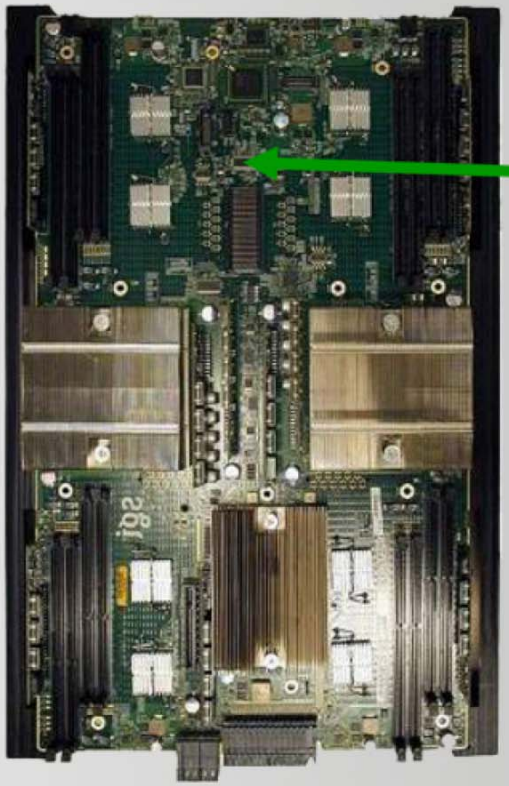

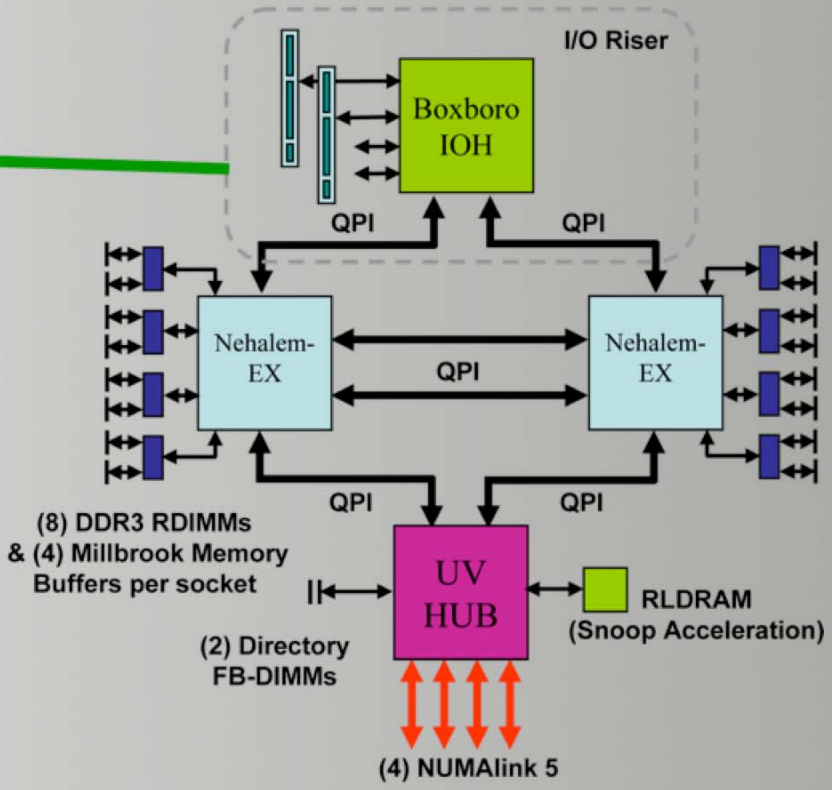

- SGI® NUMAlink® 5 = 15.0 GB/s bidirectional (7.65 GB/s payload)
- Intel® Quick Path Interconnect (QPI) = 23.4 GB/s aggregate (5.86 GT/s)
- Directory FBD1 = 6.4GB/s Read + 3.2GB/s Write (800MHz DIMMs)
- Millbrook Memory Buffers with 4 channels of DDR3 DIMMs
- Intel® Scalable Memory Interconnect (SMI)

How Nautilus Fits Into NICS

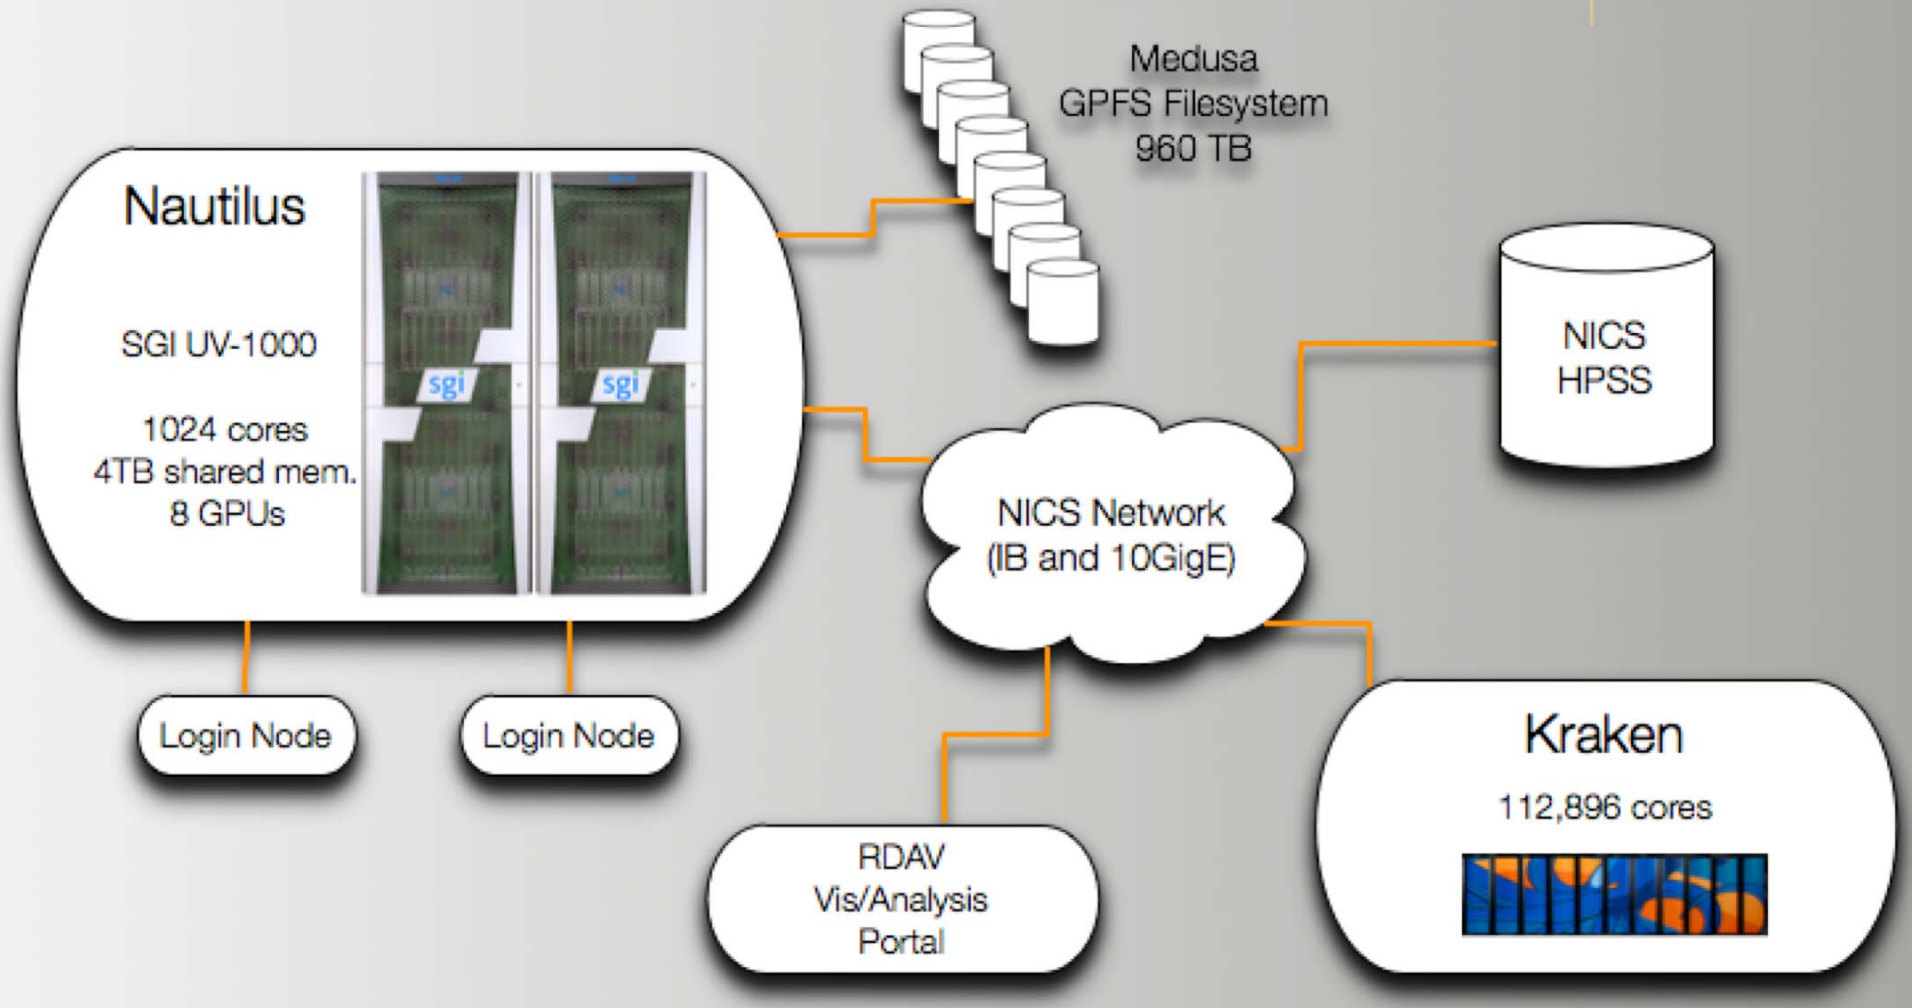

Diverse use cases dictate unique architecture

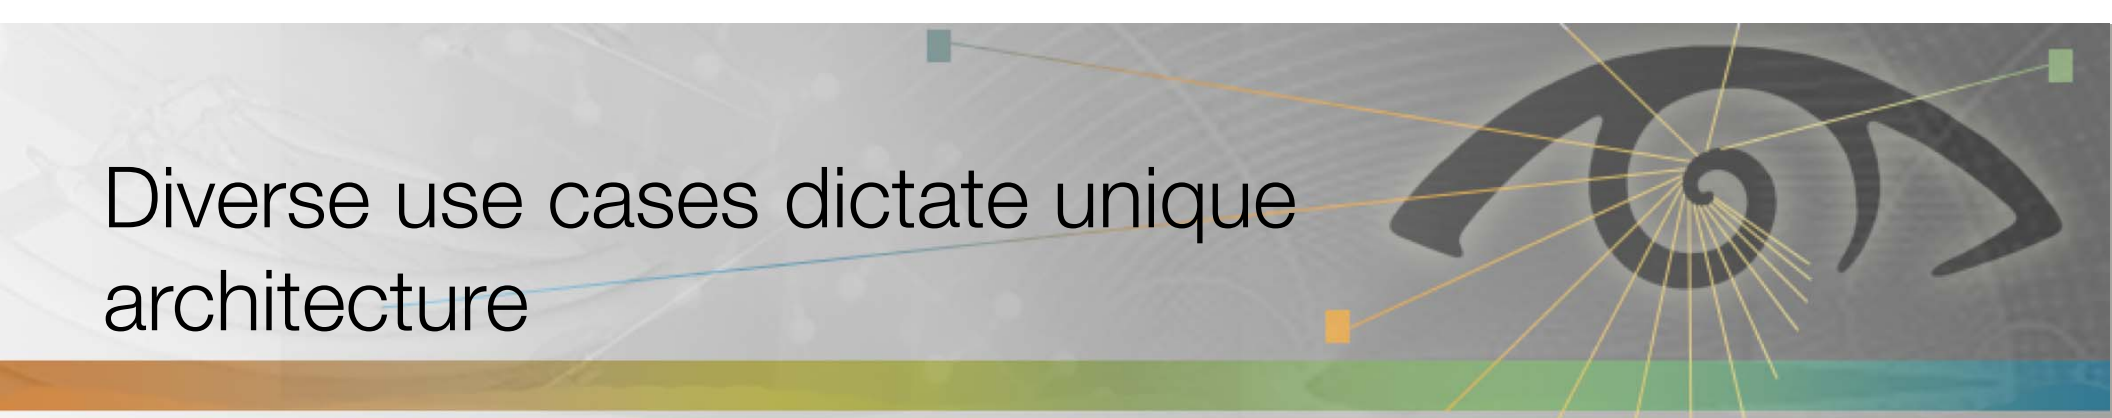

- Many HPC users can use distributed memory analysis
  - data parallel, time parallel
- However, many general and statistical analysis algorithms favor large shared memory
  - Document clustering/searching
  - Generalized graph structures
  - Bioinformatics, genomics
  - …
- *Large shared memory* is the only reasonable way to address all of these needs
- SGI UltraViolet architecture provides:
  - Large memory single-system image through NUMA
  - A “better” cluster architecture, accelerating distributed memory MPI

RDAV provides User Services

- Routine user services staff – Connected to NICS
  - Routine user services
  - Ticket triage and routing to specialist
- Specialized staff for Advanced Support for TeraGrid Applications (ASTA)
  - Specific for remote visualization, data analysis, workflow services, portal
  - Educates on effective use of existing tools or on custom development
  - Provides individualized assistance for center-wide software
- Dedicated staff for education, outreach, and training
- Dedicated staff for tool discovery and certification

RDAV provides a range of software services

- Analysis applications: to be dictated by user needs and technology needed to solve user problems.
  "Whatever it takes!"
- Write any glue software needed
  - Eden
  - Custom scripts in python, etc.
- Remote visualization and image generation
  - Provide interactive and batch image generation tools. (gnuplot, ImageMagick, etc.)
  - Remote parallel visualization (VisIt, ParaView, etc.)
  - Tools for custom application development
- Data analysis and statistical analysis
  - Octave, Parallel R, Matlab, etc.
- Workflow systems
  - DAGMan system automates batch actions on behalf of users
  - Infrequent current use, however, value is increasing and many users wish to explore.
- Dashboard delivery
  - Leverage DoE funding for eSimMon dashboard system.
- Portal system
  - Builds upon standard Liferay platform
  - Provides SAS services for analysis and visualization