Red Storm Infrastructure
Cray User's Group, May 2005

Robert A. Ballance
Dept 9328

Red Storm Computer facility
Sandia National Laboratories
Albuquerque, NM

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

Advanced Simulation & Computing™

Co-Authors

- John Noe – Mgr, 9328
- Milt Clauser – RoSE
- Martha Ernest – CEM
- Barbara Jennings – User Support
- David Logsted – Viz
- David Martinez – Facilities
- John Naegle – Networking
- Len Stans – Facilities

Sandia National Laboratories

Red Storm Infrastructure

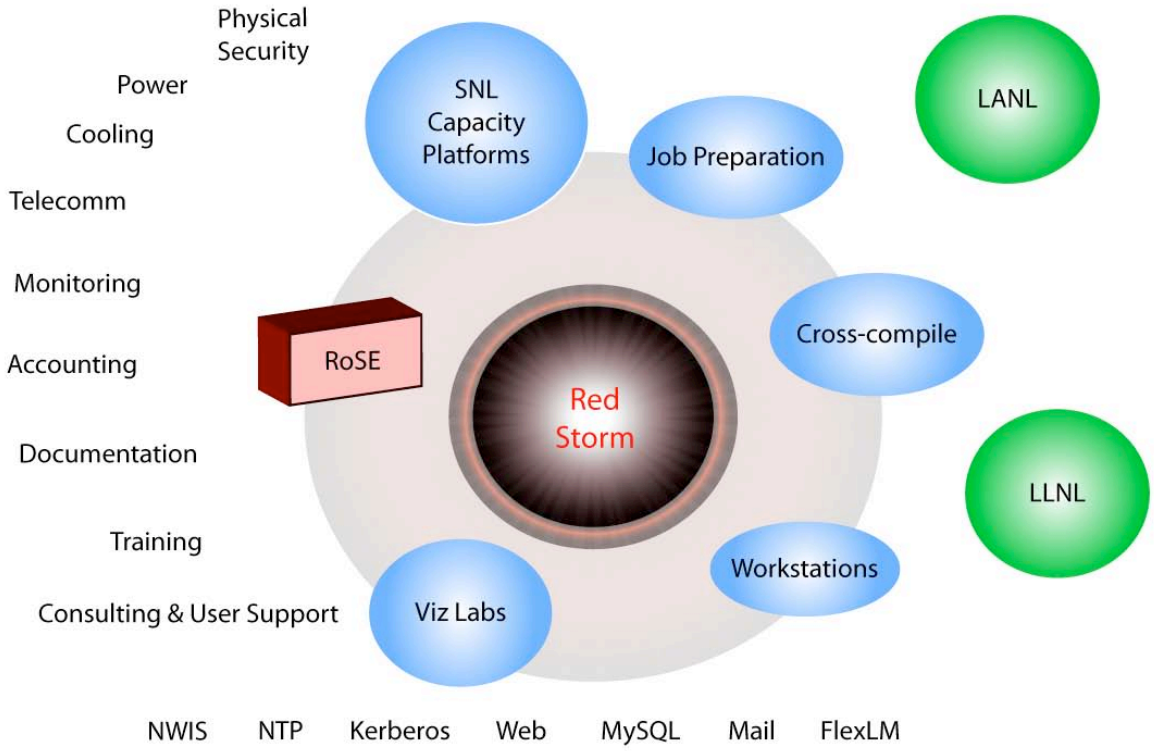

Advanced Simulation & Computing™

Facilities

Sandia National Laboratories

Light, Power, Cooling!

- Super Computer Annex (SCA)
  - 20,250 Sq. Ft. Facility
  - Clear Span of 150'
  - Additional 4,500 Sq. Ft for Office and Support Equipment
  - 3.5 Megawatts
- Computer Utility Building (CUB)
  - Supply Chilled Water for Bldg. 880 and/or the Super Computer Annex (SCA)
  - 2000 Tons of Chilled Water for Red Storm
  - Expandable to 4000 Tons

A “Big” Empty Floor

Floorplan

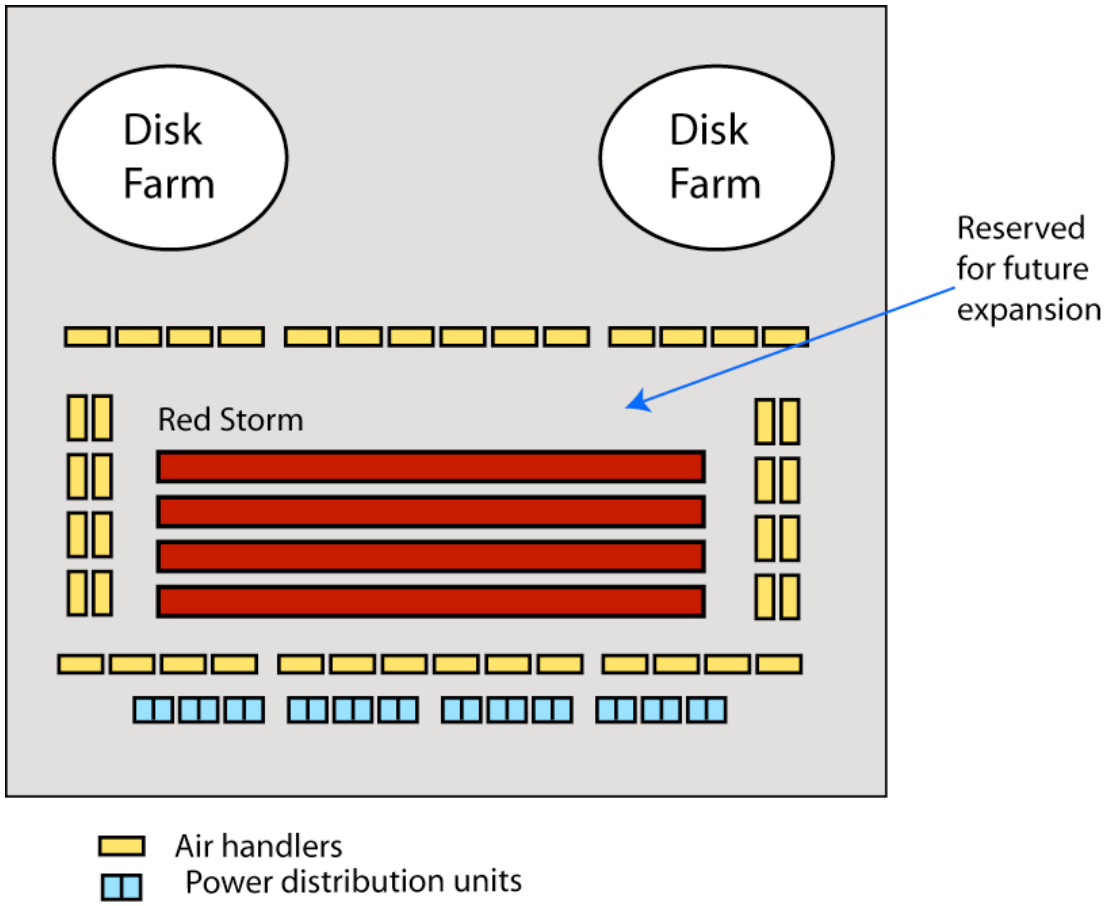
Disk Farm
Disk Farm
Reserved for future expansion
Red Storm
Air handlers
Power distribution units

Advanced Simulation & Computing™

Computational Setting

Sandia National Laboratories

A Data and Visualization Center to provide interactive & batch services for a complete high-performance computing environment

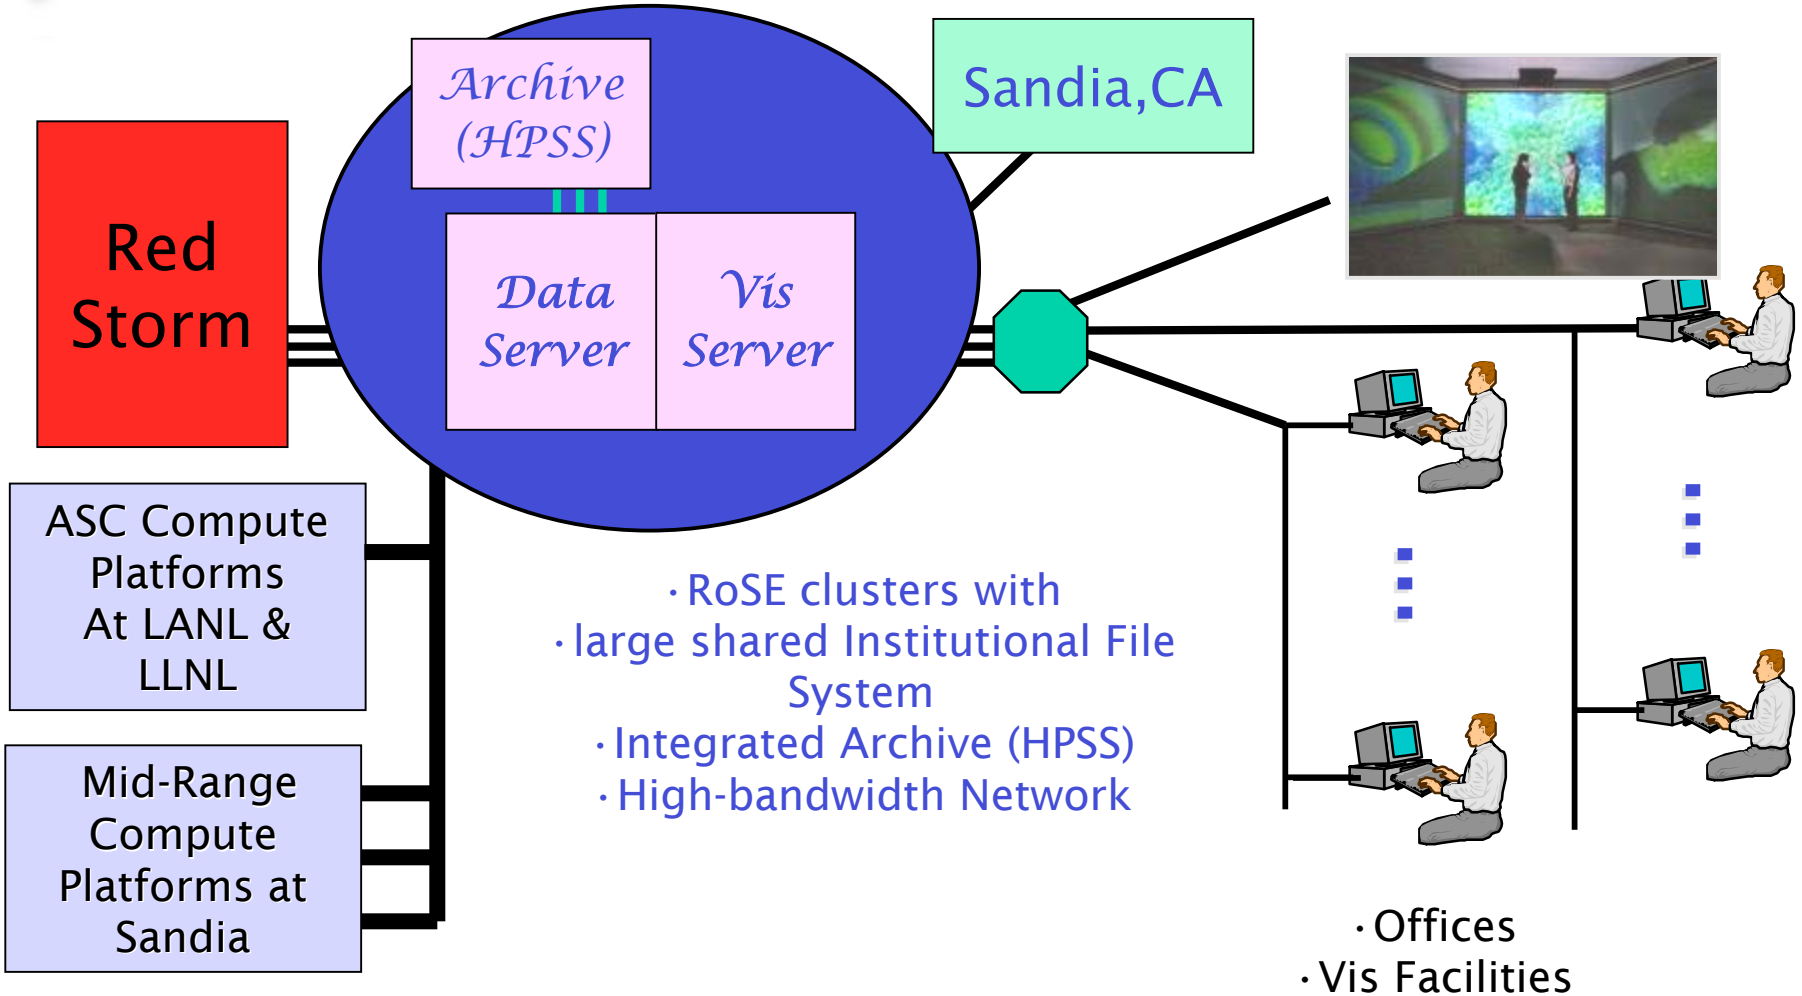

… for pre & post processing, visualization, data manipulation, archival services.

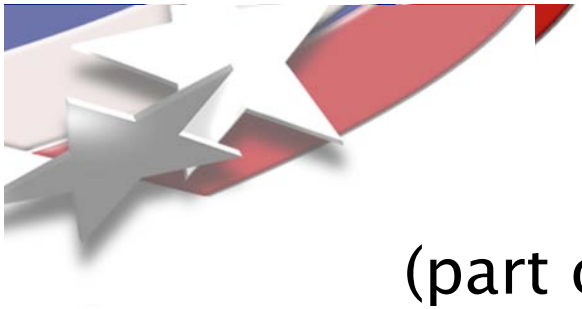

RoSE Clusters

(part of the Red Storm Environment)

Two terascale clusters for highly interactive visualization, data analysis and archiving of output from Red Storm

- Red RoSE cluster acquired: Summer 2004
  - initial unclassified operation: Fall 2004
  - initial classified operation: June 2005
- Black RoSE cluster acquisition in FY 06
- Initial unclassified support for Red Storm: May 2005
  Feynman cluster's vis nodes + part of Red RoSE storage

Two Key Performance Requirements for RoSE:

Vis Power

- for highly interactive visualization and analysis of large (300 mega-cell and larger) datasets
- At least 4 X present capability of Feynman Vis nodes
  - 10^{10} (peak) triangles/sec render power
  - 4 Teraflops (peak) compute power (~1/10 Red Storm)

I/O Power

- for accessing terascale data within RoSE cluster at interactive rates
  - 25 GB/s parallel file system
  - 1 second to access one time-step of a 300-mega-cell calculation
- for moving terascale data from Red Storm
  - 25 GB/s (90 TB/hour) parallel file transfer, to minimize impact on Red Storm file system

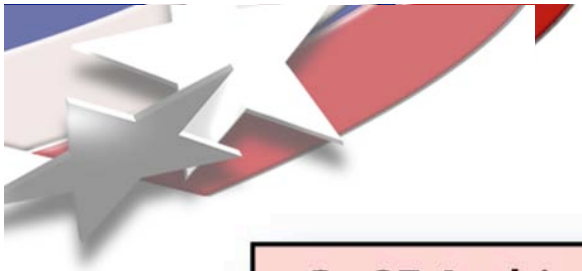

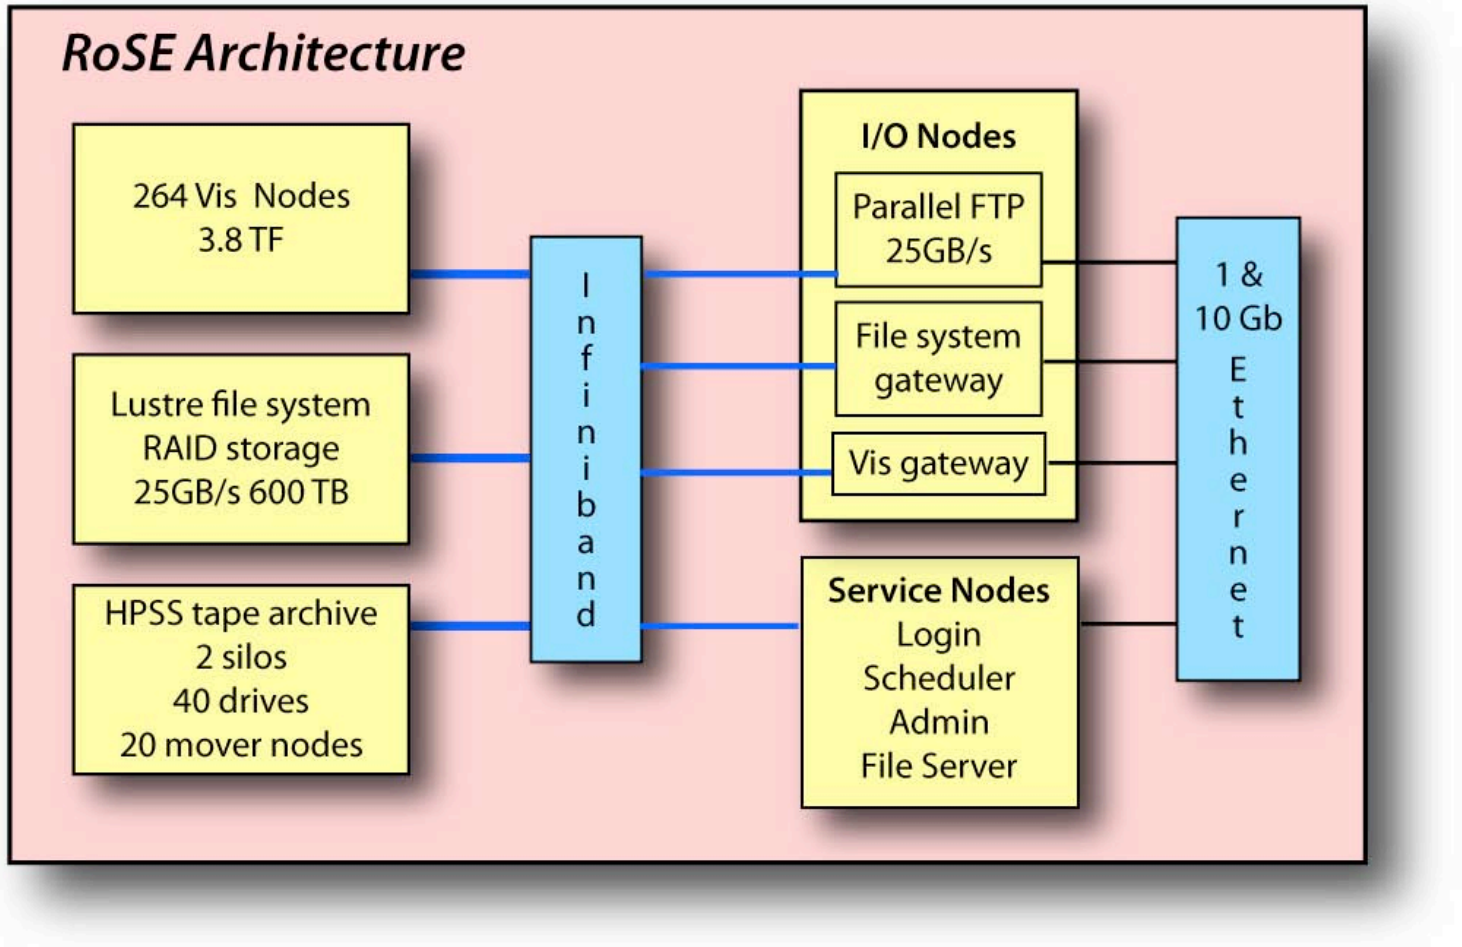
RoSE Architecture
264 Vis Nodes
3.8 TF
Lustre file system
RAID storage
25GB/s 600 TB
HPSS tape archive
2 silos
40 drives
20 mover nodes
Infiniband
I/O Nodes
Parallel FTP
25GB/s
File system
gateway
Vis gateway
Service Nodes
Login
Scheduler
Admin
File Server
1 &
10 Gb
Ethernet

Red Storm/RoSE Connections

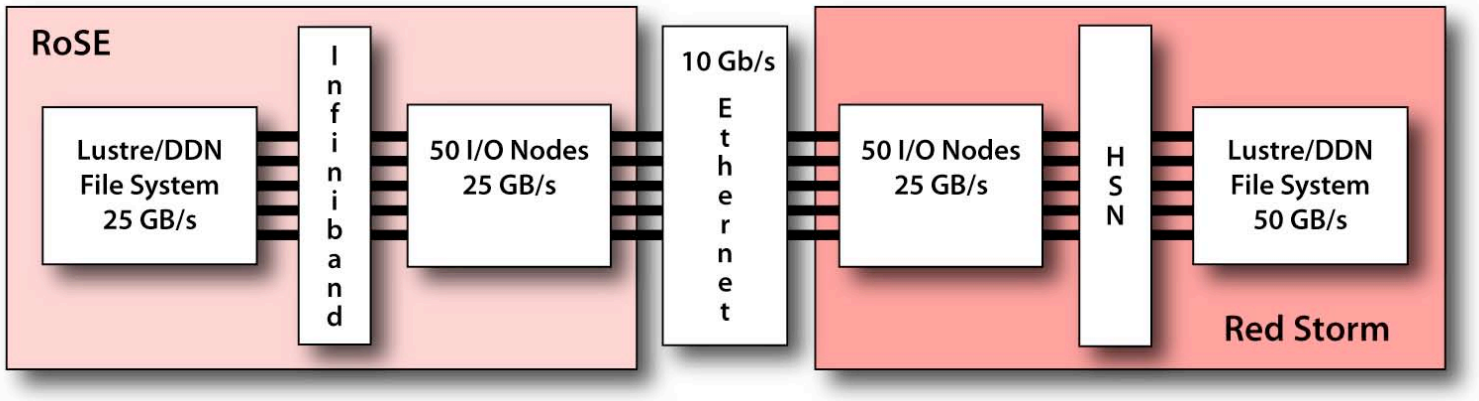

Advanced Simulation & Computing™

Vis Labs

Sandia National Laboratories

Sandia Vis Facilities

- Views Corridor
- DISL Interactive Design Center
- JCEL
- MESA Design and Education Center
- Mesa/WIF Weapons Integration Facility

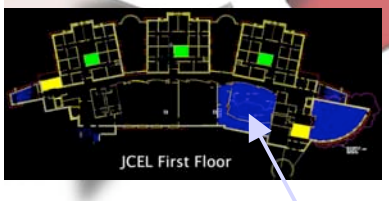

JCEL

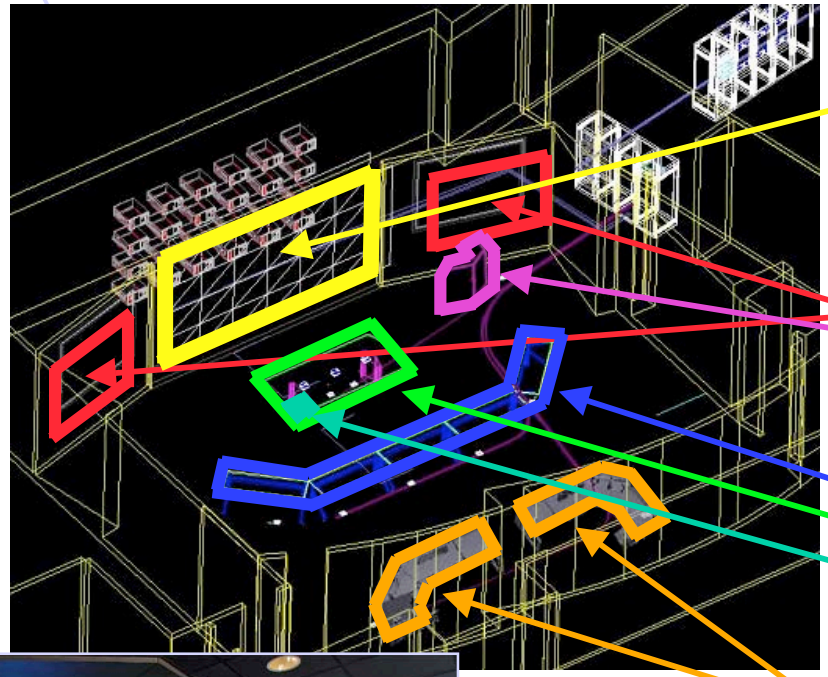

Design Features:

- ✓ High-Resolution Display
  - 31.5 Mega-pixel
  - 6-horizontal x 4-vertical
  - 24 DPI HIGHlite 4000Dsx projectors
- ✓ 24 node display cluster driven
- ✓ 2-SXGA side displays
- ✓ Presentation Lectern
- ✓ Tablet Display Interface
- ✓ Videoconferencing (Red & Black)
- ✓ AV interface boxes
- ✓ Training Tables
- ✓ Conference Table
- ✓ Document Cameras
- ✓ Intuitive Design
- ✓ Crestron Room Automation
- ✓ 10-Gigabit Ethernet Networking
- ✓ Red & Black Command Stations

Advanced Simulation & Computing™

Networking

Sandia National Laboratories

Network Core

- Full scale switched 10G infrastructure for Red Storm
  - High speed network required to support Red Storm connectivity to external resources (users, storage, visualization, etc.)
- High speed connectivity for Red Storm user community
  - High speed network support required for effective user access to Red Storm
  - Multiple locations across SNLA

Unclassified Core Network

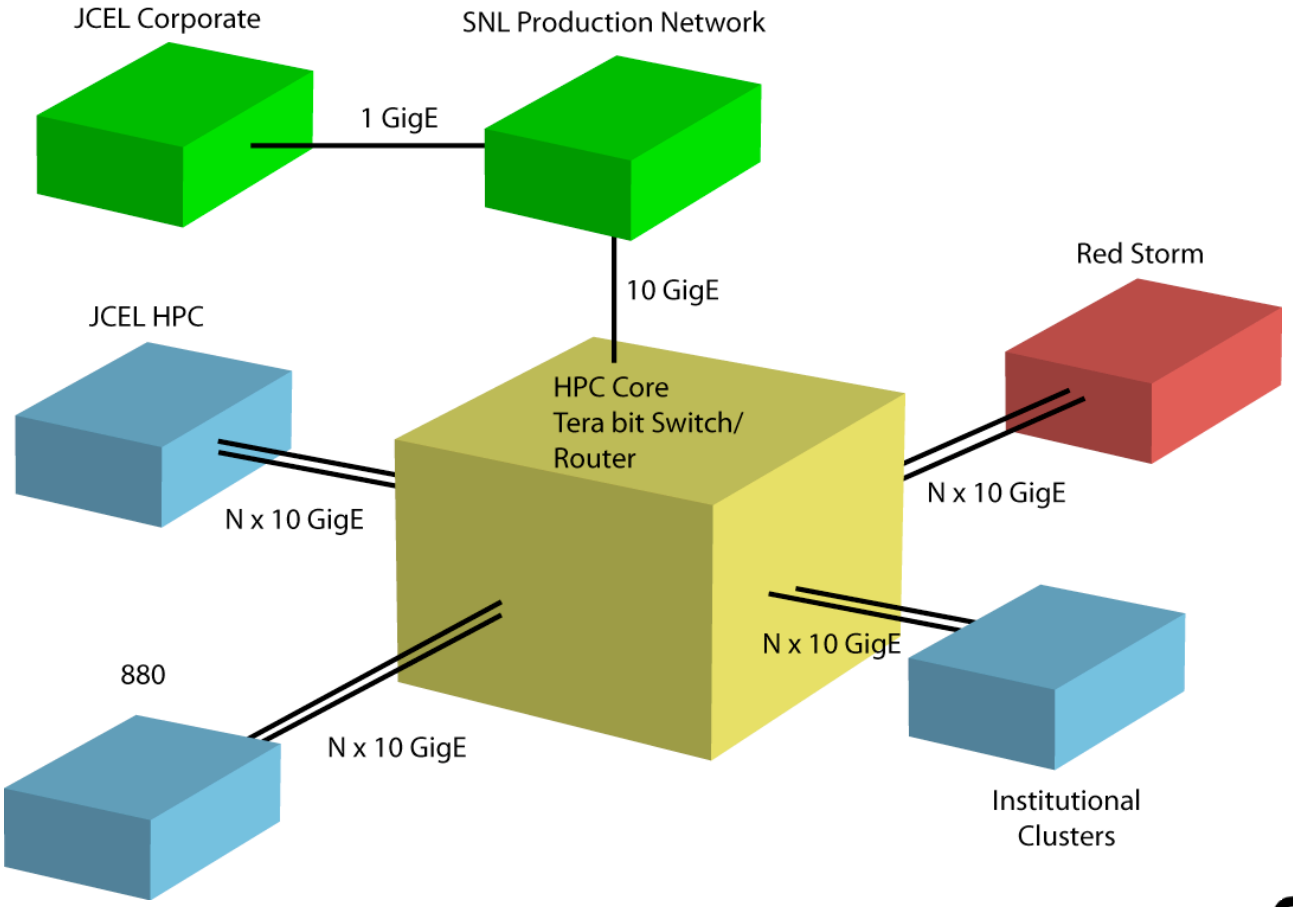

Red Storm Connectivity

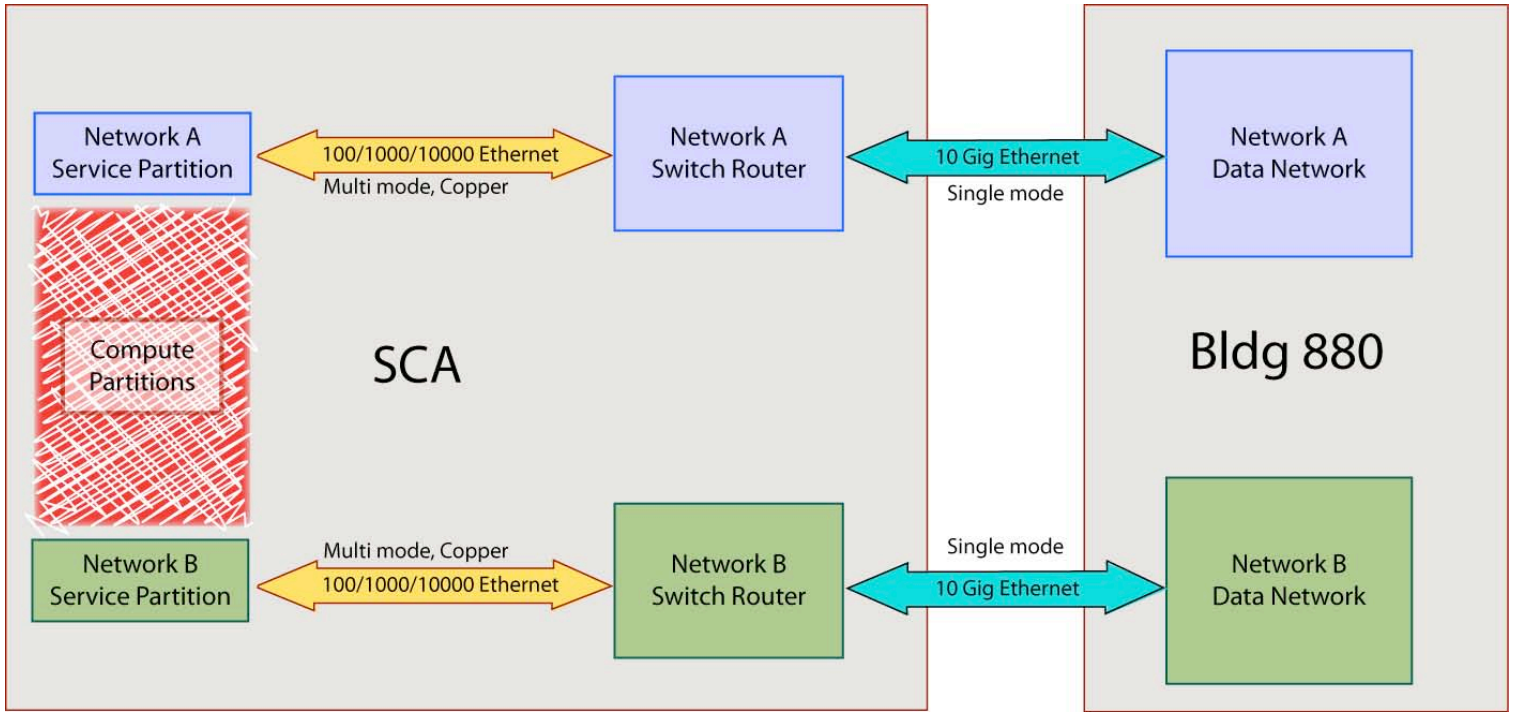

Advanced Simulation & Computing™

Operational Setting

Sandia National Laboratories

Network Services

- DHCP
  - Can't be used everywhere
  - Required local fixes for 10G nodes
- DNS
  - Lbnamed?
- Keberized ssh for external access
  - Passwordless ssh on inside
  - Can we use the same server?
- FlexLM
  - On compile servers for now

Job Accounting

- XT3 provides basic job accounting
  - CPA, PBS
  - CPA-based scripts
  - CUP measures
- SNL has centralized accounting server
  - AIRS
  - Nightly inhales
  - Superset of CPA info
- CPA also keeps I/O statistics not in AIRS
- 2 sources of accounting data?

Monitoring

- Cyber-Enterprise Management (CEM)
  - Business Service View
  - 24x7 Staffing
- RAS and CEM
  - How do you leverage CRMS to fit with CEM?
    - Events?
    - Probes?
    - Extensions
  - Example: Rolling service failure

User Support

- Web
- Email
- Phone
- Collaborative Learning Environment (CLIK!)
  - Come to Barbara Jennings' talk later today!