# High Performance Computing at the National Center for Computational Sciences at ORNL

**Go to Menu**

# Outline


- Our Mission
- Computer Systems: Present, Past, Future
- Challenges Along the Way
- Resources for Users


**Go to Menu**

# Our Mission

**Go to Menu**

# ORNL is the U.S. Department of Energy's largest science and energy laboratory

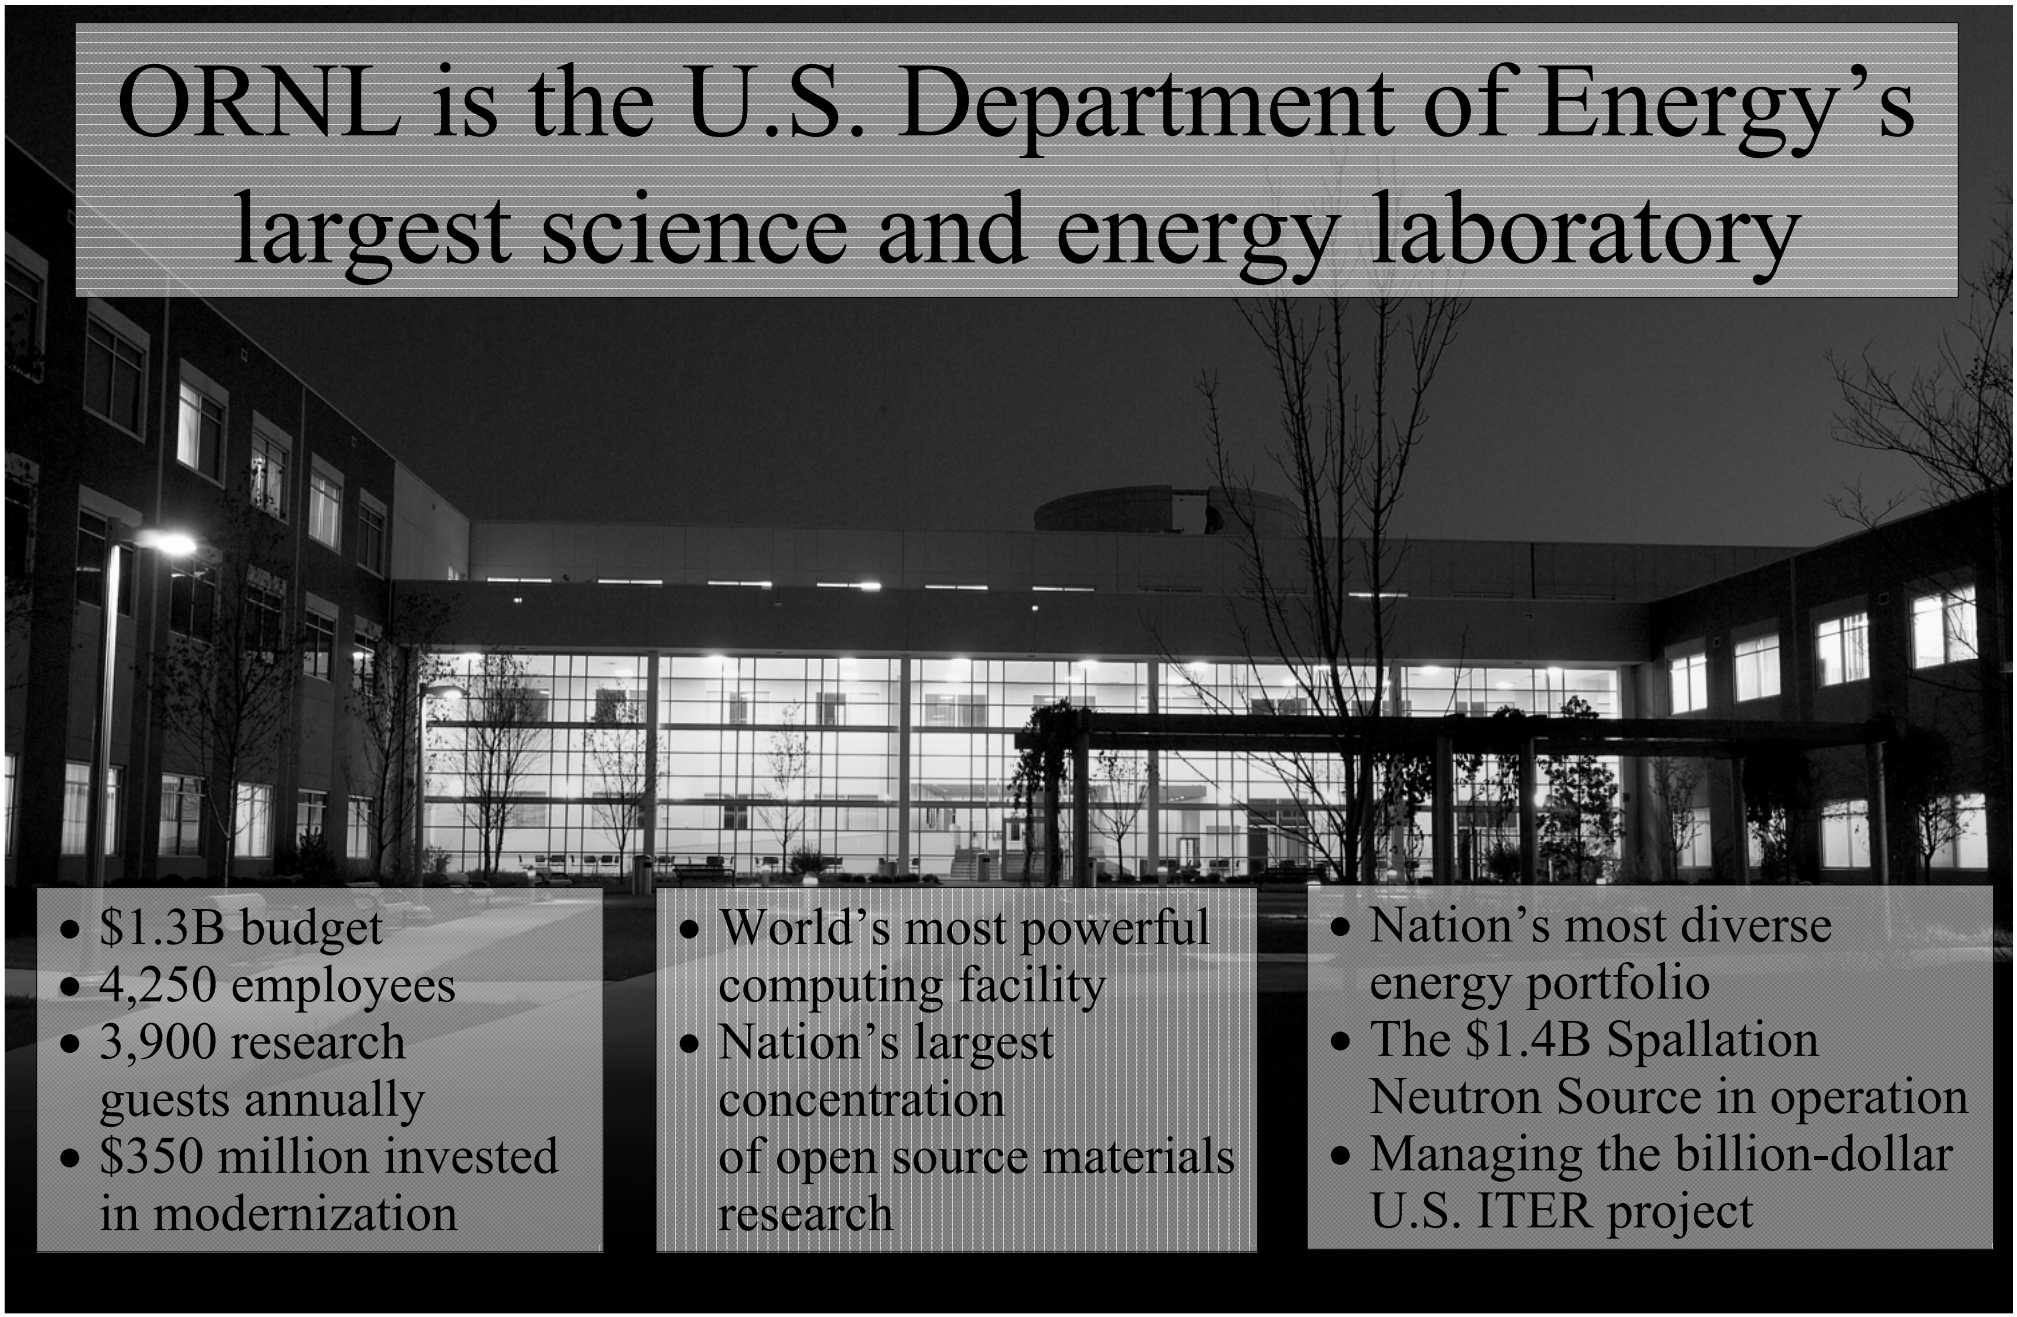

- $1.3B budget
- 4,250 employees
- 3,900 research guests annually
- $350 million invested in modernization

- World's most powerful computing facility
- Nation's largest concentration of open source materials research

- Nation's most diverse energy portfolio
- The $1.4B Spallation Neutron Source in operation
- Managing the billion-dollar U.S. ITER project

**Go to Menu**

# Computing Complex @ ORNL

**$70M Operating budget to deploy and operate the computational resources required to tackle global challenges**

- Providing world-leading computational resources and specialized services for the most computationally intensive problems
- Providing stable hardware/software path of increasing scale to maximize productive applications development
- Delivering transforming discoveries in materials, biology, climate, energy technologies, etc.


World's most powerful computer for open science

Go to Menu

# Computer Systems: Present, Past, Future

**Go to Menu**

# NCCS systems, 1991–2008

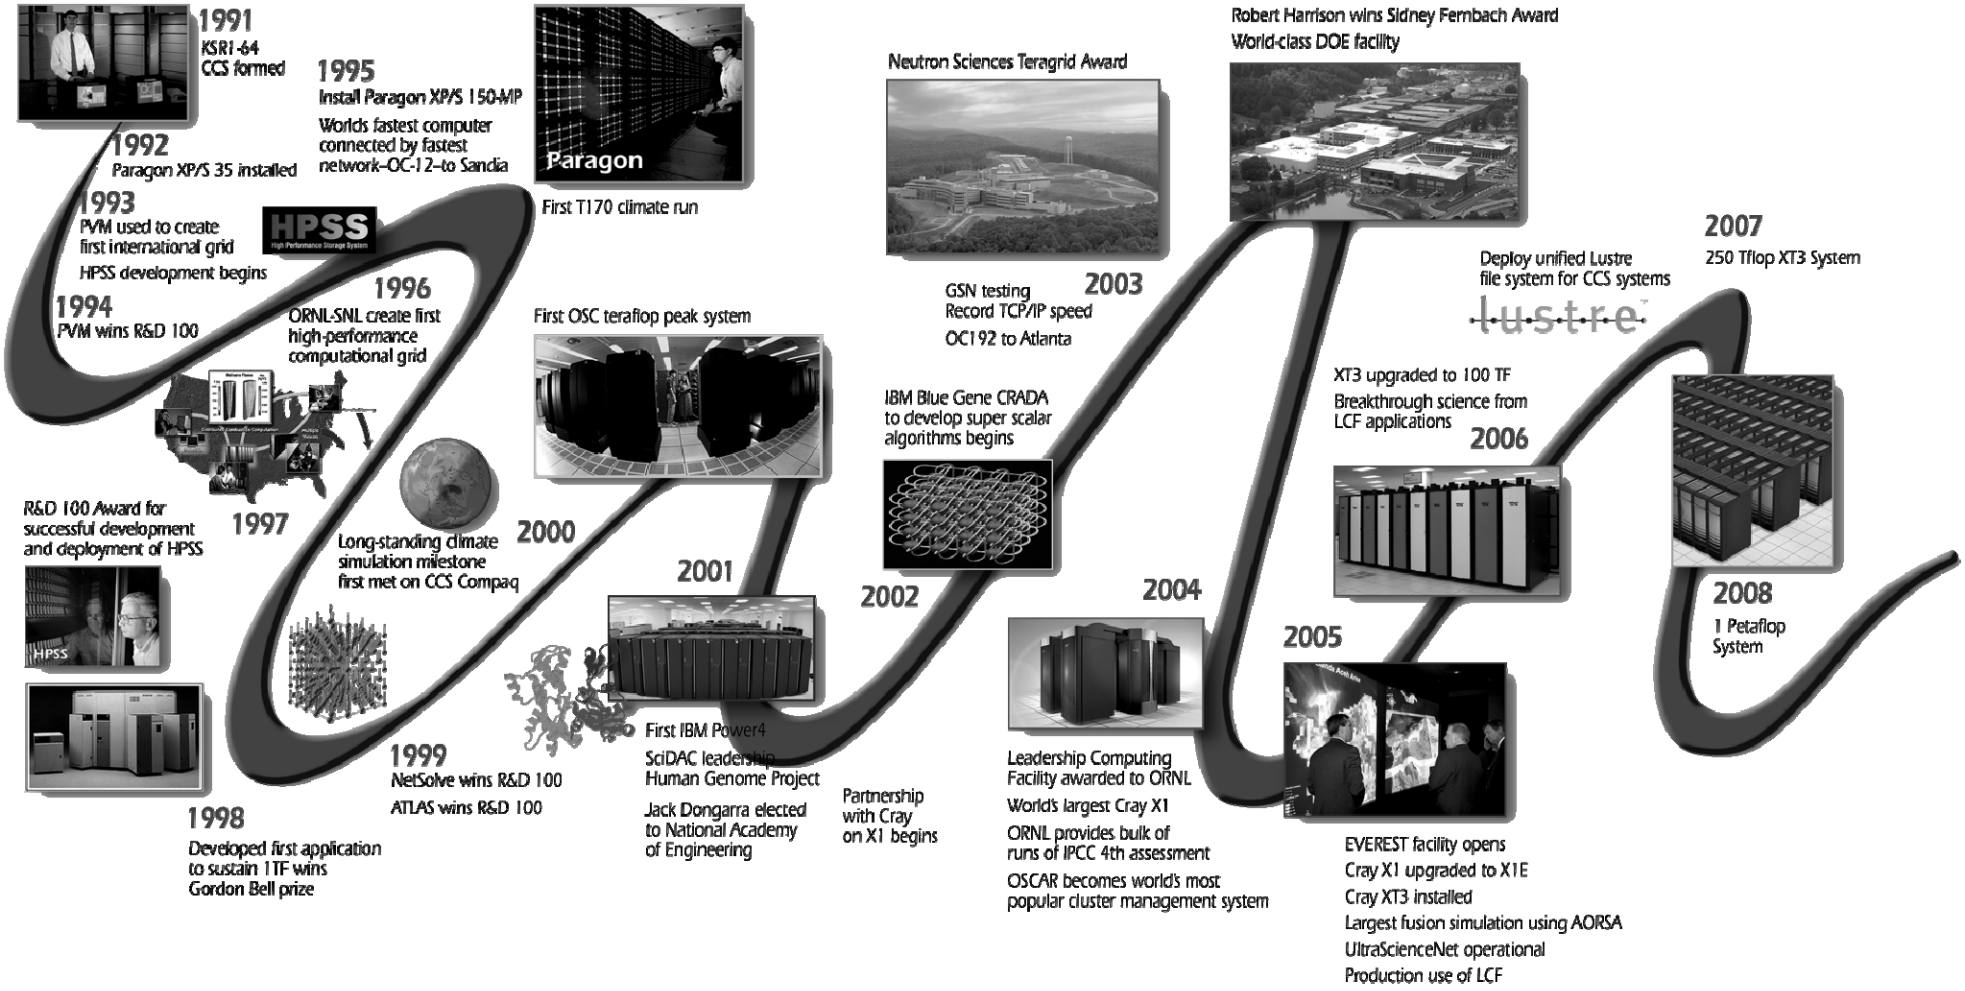


Go to Menu

# Million-fold increase in computing and data capabilities

Cray X1
3 TF

2004

2005

Cray XT3
Single-core
26 TF

2006

Cray XT3
Dual-core
54 TF

2007

Cray XT4
119 TF

2008

Cray XT4
Quad-core
166 TF
18TB,
0.3 PB

Cray XT4
Quad-core
263 TF
62 TB,
1 PB

2009

Cray XT5
8/12-core,
dual-socket
SMP ~1 PF
100TB,
3.3 PB

Jaguar
8-core,
dual-socket
SMP 1.64PF
362 TB,
10 PB

2009

Jaguar
8/12-core,
dual-socket
SMP 2.60PF
362 TB,
10 PB

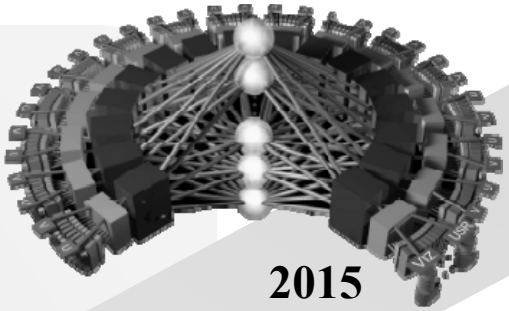

2015

Future
system
100–250
PF

2018

Future
system
1000 PF
(1 EF)

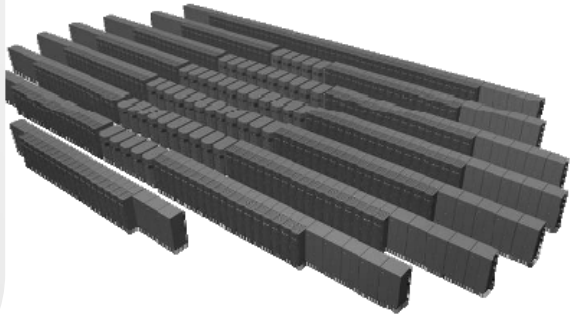

Go to Menu

# A Growth-Factor of a Billion in Performance in a Single Career




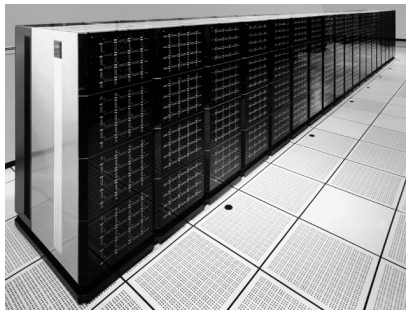



1949 Edsac
1954 ORACLE
1959 IBM 7094
1976 Cray 1
1991 Intel Delta
1995 Intel Paragon
1996 Cray T3E
2003 Cray X1
2008 Cray XT5

1
$10^3$
$10^6$
$10^9$
$10^{12}$
$10^{15}$

One OPS
KiloOPS
MegaOPS
GigaOPS
TeraOPS
PetaOPS

1823 Babbage Difference Engine
1943 Harvard Mark 1
1951 Univac 1
1964 CDC 6600
1982 Cray XMP
1988 Cray YMP
1997 ASCI Red
2001 Earth Simulator
2005 Cray XT3

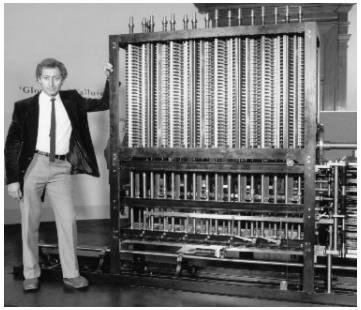



Go to Menu

# NCCS resources

**October 2009 summary**

**6 Systems**

**Supercomputers**

**>265,000 cores**
**>360 TB Memory**

**>2,6 PFLOPS**

**CRAY XT5 JAGUAR**

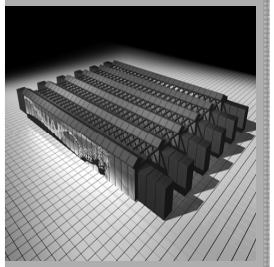

(224,256)
2.6GHz
292 TB Memory

10,000 TB

**CRAY XT4 JAGUAR**

(31,328)
2.1GHz
61 TB Memory

750 TB

**IBM BLUE/GENE P**

(8192)
850 MHz
4 TB Memory

60 TB

**LINUX CLUSTER SMOKY**


(1280)
2.4 GHz
2.5 TB

4.5 TB

**LINUX CLUSTER LENS**

(128)
2.2 GHz
128 GB

9 TB

**IBM HPSS**

Many storage devices supported

Max. 30 PB

**Scientific visualization lab EVEREST**

27-projector PowerWall
35 million pixels

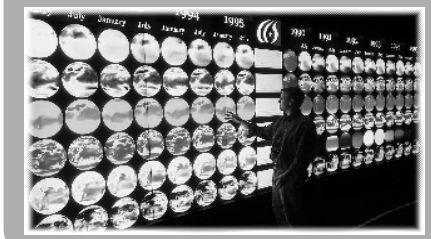

Go to Menu

# NCCS resources: Jaguar Supercomputer

Jaguar: World's Most Powerful Computer Designed for Science from the Ground Up


| | jaguar XT4 | jaguarpf XT5 |
|---|---|---|
| Peak Performance | 263.16 TFLOPS | 2.33 PFLOPS |
| System Memory | 61 TB | 292 TB |
| Disk Space | 750 TB | 10,000 TB |
| Disk Bandwidth | 44 GB/s | 240 GB/s |
| Interconnect Bandwidth | 157 TB/s | 374 TB/s |

Go to Menu

# NCCS resources: BG/P


- 27 TFlop System
  - 2048 850Mhz IBM quad core 450d PowerPC
  - 2 GB/node
  - 64 I/O nodes
  - 10 Terabyte GPFS parallel file system
- Available to ORNL + core university partner members

Go to Menu

# NCCS resources: Smoky


- Resource provided to users needing system comparable to major NCCS resources for application porting, development
- 80 node Linux cluster
- Four quad-core 2.0GHz AMD Opteron processors per node
- 32 GB of memory (2GB per core)
- Gigabit Ethernet network with infiniband interconnect

**Go to Menu**

# NCCS resources: Visualization Facilities

- The visualization capabilities of NCCS include:
  - visualization/data analysis cluster called Lens
  - large PowerWall display called EVEREST

- Scientists can make use of the EVEREST facility by contacting any member of the visualization team and booking a time.

# NCCS resources: Lens

- Resource for data visualization
- 32 node Linux cluster dedicated to data analysis and high-end visualization
- Each node:
  - Four quad-core 2.3 GHz AMD Opteron processors
  - 64 GB memory
  - 2 NVIDIA 8800 GTX GPUs.

Go to Menu

# NCCS resources: Visualization PowerWall (EVEREST)

EVEREST - Exploratory Visualization Environment for REsearch in Science and Technology

- 27-projector PowerWall
- Viewing at a 30 feet by 8 feet
- 11,520 by 3,072 pixels, or a total of 35 million pixels
- The wall is integrated with the rest of the computing center, creating a high-bandwidth data path between large-scale high-performance computing and large-scale data visualization.

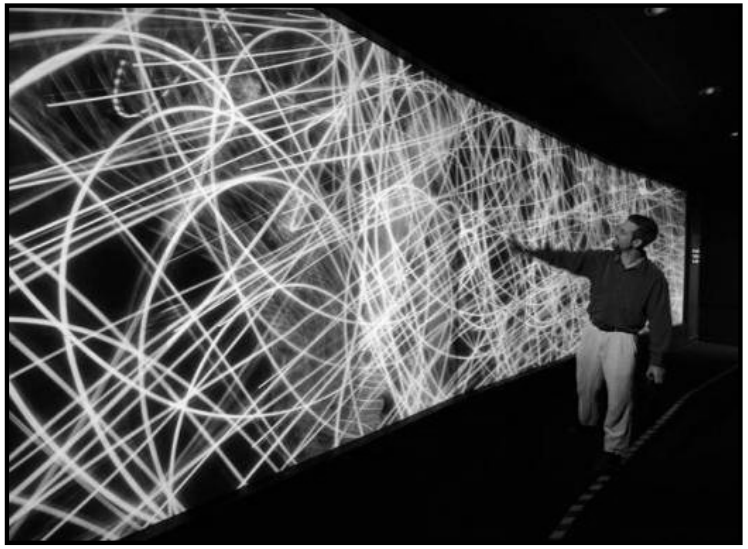

- EVEREST is controlled by a 14 node cluster with GPUs for remote visualization.
- Each node contains four dual-core AMD Opteron processors.
- These 14 nodes have nVidia QuadroFX 3000G graphics cards connected to the projectors, providing a very-high-throughput visualization capability.

**Go to Menu**

# NCCS resources: High Performance Storage System (HPSS)

- HPSS is an archival Back-up system which consists of
  - two types of storage technology:
    - disk – “on-line” for frequently/recently accessed files
    - tape – “off-line” for very large or infrequently accessed files
  - Linux servers
  - High Performance Storage System software


- Tape storage is provided by robotic tape libraries.
- HPSS has three SL8500 tape libraries. Each can hold up to 10,000 cartridges.
- The StorageTek SL8500 libraries house a total of
  - twenty-four T10000A tape drives (500 gigabyte cartridges, uncompressed)
  - thirty-six T10000B tape drives (1 terabyte cartridges, uncompressed).
- Each drive has a bandwidth of 120 MB/s
- As of October, 2009, HPSS has 7.2 PB stored in over 16.1 million files.

**Go to Menu**

# NCCS resources: Center-Wide File System (SPIDER)

"Spider" provides a shared, parallel file system for all LCF systems and based on Lustre file system

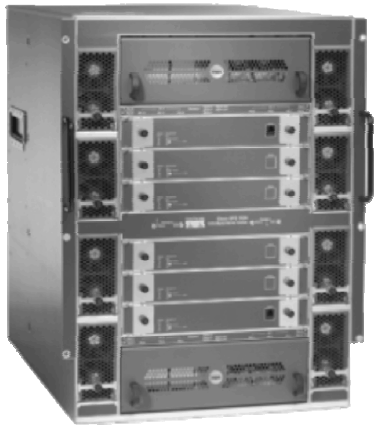

- Over 10 PB of RAID-6 Capacity
  - 13,440 1Gb SATA Drives (33 tons of discs)
  - 192 OSSs and 1344 OSTs (7 OSTs/OSS)
  - 3 Terabytes of memory
- Demonstrated bandwidth of over 200 GB/s
  - 30,000 files created per second
- Demonstrated stability on a number of LCF Systems
  - Over 26,000 lustre clients at NCCS mounting the file system and performing I/O
- Available from all systems via our high performance scalable I/O network
  - 4 InfiniBand core switches
  - Over 3,000 InfiniBand ports
  - Over 3 miles of cables


Go to Menu

# Challenges Along the Way

**Go to Menu**

# High Bandwidth Connectivity

- High Bandwidth Connectivity to NCCS Enables Efficient Remote User Access
- Connected to Major Science Networks

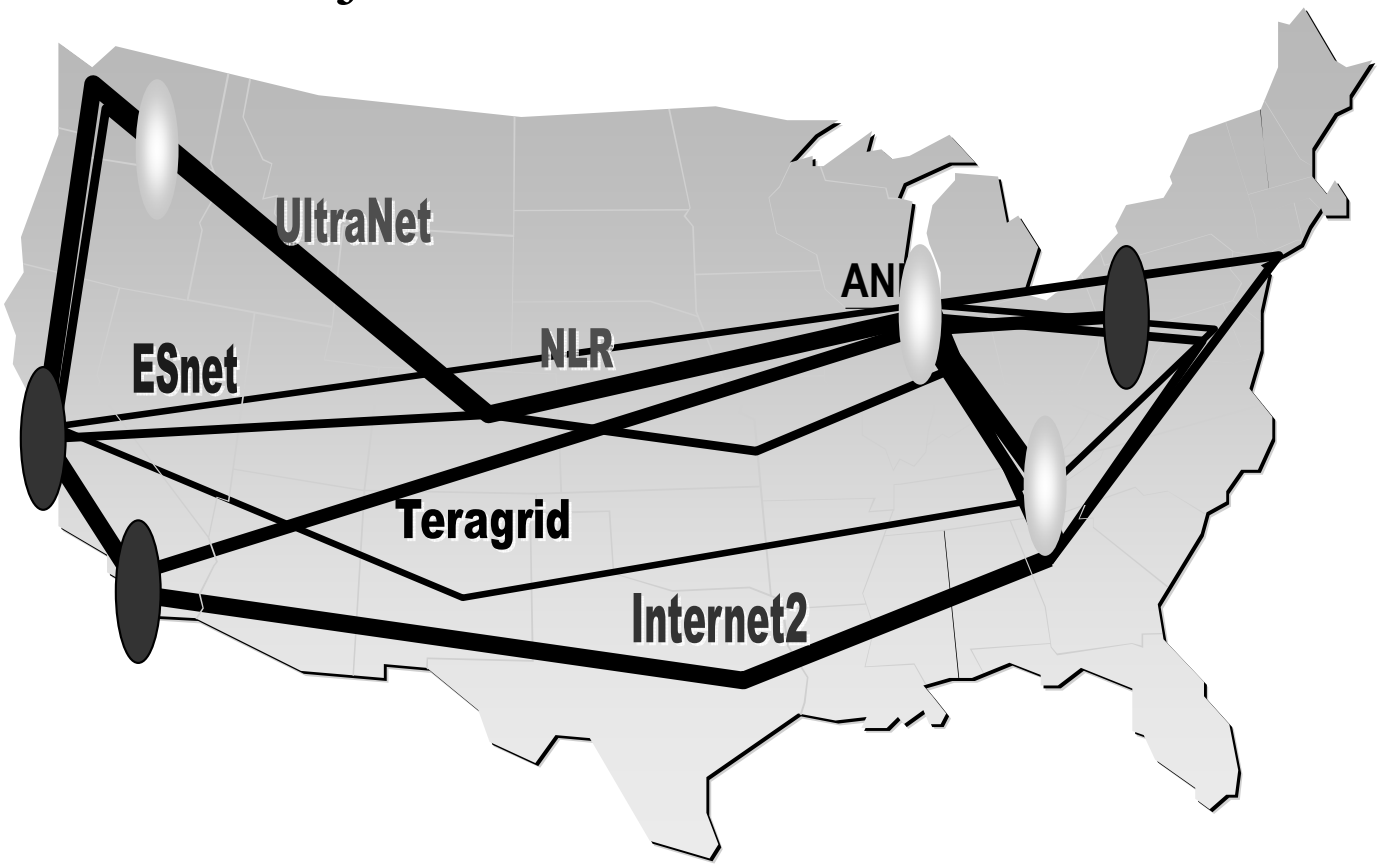


Target

- 100Gb/lamda (NRL+Vendor+ORNL)

Go to Menu

# We are Addressing a Broad Range of Science Challenges

| Projects | 2006 | 2007 | 2008 | 2009 |
|---|---|---|---|---|
| Accelerator physics | 1 | 1 | 1 | 1 |
| Astrophysics | 3 | 4 | 5 | 5 |
| Chemistry | 1 | 1 | 2 | 4 |
| Climate change | 3 | 3 | 4 | 5 |
| Combustion | 1 | 1 | 2 | 2 |
| Computer science | 1 | 1 | 1 | 1 |
| Fluid Dynamics | | | 1 | 1 |
| Fusion | 4 | 5 | 3 | 5 |
| Geosciences | | 1 | 1 | 1 |
| High energy physics | | 1 | 1 | |
| Life sciences | 2 | 2 | 2 | 4 |
| Materials science | 2 | 3 | 3 | 4 |
| Nuclear physics | 2 | 2 | 1 | 2 |
| Industry | 2 | 3 | 3 | 3 |
| Total Projects: | 22 | 28 | 30 | 38 |
| CPU Hours: | 36,156,000 | 75,495,000 | 145,387,000 | 469,683,000 |

Go to Menu

# Science Application Development & Readiness

Ensuring Application Codes Can Effectively Utilize HPC Systems on "Day One"

- Model coupling
- Hierarchical algorithms
- Solver technology and innovative solution techniques
- Accelerated time integration
- Parallel programming models
- Maintaining application libraries
- Software and algorithm strategies to mitigate high hardware latencies
- Automated diagnostics

**Go to Menu**

# Resources for Users

**Go to Menu**

# Resources for Users: Getting Started

- About Jaguar

  http://www.nccs.gov/computing-resources/jaguar/

- Quad Core AMD Opteron Processor Overview

  http://www.nccs.gov/wp-content/uploads/2008/04/amd_craywkshp_apr2008.pdf

- PGI Compilers for XT5

  http://www.nccs.gov/wp-content/uploads/2008/04/compilers.ppt

- NCCS Training & Education – archives of NCCS workshops and seminar series, HPC/parallel computing references

  http://www.nccs.gov/user-support/training-education/

- 2009 Cray XT5 Quad-core Workshop

  http://www.nccs.gov/user-support/training-education/workshops/2008-cray-xt5-quad-core-workshop/

**Go to Menu**

# Resources for Users: Advanced Topics

- Debugging Applications Using TotalView

  http://www.nccs.gov/user-support/general-support/software/totalview

- Using Cray Performance Tools - CrayPat

  http://www.nccs.gov/computing-resources/jaguar/debugging-optimization/cray-pat/

- I/O Tips for Cray XT4

  http://www.nccs.gov/computing-resources/jaguar/debugging-optimization/io-tips/

- NCCS Software

  http://www.nccs.gov/computing-resources/jaguar/software/

- Cray Documentation

  http://docs.cray.com/

Go to Menu

# Resources for Users: More Information

- NCCS website

  http://www.nccs.gov/

- How to obtain Access to NCCS Resources

  http://www.nccs.gov/user-support/access/

- Contact us

  help@nccs.gov

**Go to Menu**