




NCAR

# Overview of NCAR Cyberinfrastructure Support for the Atmospheric and Related Sciences



Tom Bettge

Director of Operations and Services

Computational and Information Systems Laboratory

# The NCAR Mission

- To understand the behavior of the atmosphere and related physical, biological, and social systems
- To support, enhance and extend the capabilities of the university community and the broader scientific community, nationally and internationally
  - *Provide robust, accessible, innovative and advanced services and tools*
  - *Capacity and capability supercomputing*
  - *Create an Earth System Knowledge Environment*

**From 1959 Blue Book: “.....requires facilities and technological assistance beyond those that can be made available at individual universities.”**

- To foster the transfer of knowledge for the betterment of life on Earth




# Outline

- High Performance Computing
- Mass Storage
- Research Data Archive (RDA) Stewardship
- Data Analysis and Visualization Services
- Earth System Grid
- Computational Science – Toward Petascale Computing

# Scientific Computing Division

## History of Supercomputing at the National Center for Atmospheric Research




# Communities Served



NCAR

## *Universities and Affiliates*



***NCAR***



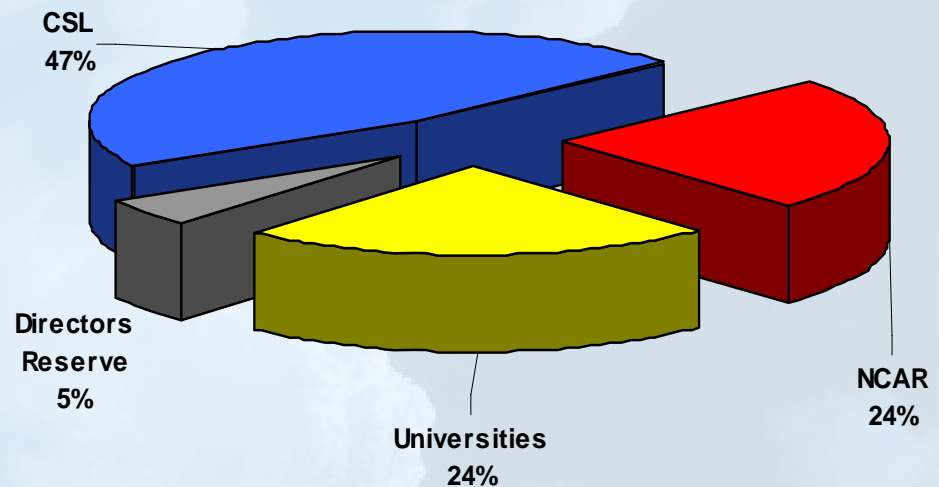
## *CSL - The Climate System Laboratory at NCAR*



# The Computing Resource Pie

- NCAR Computing Resources
  - *Climate Simulation Laboratory*
    - multi-agency access for large climate-related projects
  - *UCAR Community*
    - University
    - NCAR
    - NCAR Director's Reserve


- Facts and Figures
  - *Active Users:* **1345**
  - *Projects:*
    - CSL 47
    - Univ 412
    - NCAR 283





# The Science Pie

- Near the end of FY2006, the combined peak production supercomputing capacity at NCAR was over 15 TFLOPs
- Roughly 64% of that capacity was used for climate simulation and analysis





# Managing NCAR HPC Resources

- **Goals:**

- *Fulfill mission*
- *Ensure users/groups/projects/proposals get fair share*
- *Ensure equitable access*
- *Resource utilization is maximized*
- *Flexibility exists to give important projects priority access*
- *Flexibility exists to allocate unanticipated scientific opportunities*

- **Methods:**

- *Allocation Process*
  - Peer Review
- *Resource Scheduling*
  - Recognizes Desired Distribution
- *Queue Priority Structure*
  - Premium – Regular – Economy – Standby
- *Resource Limits*
  - Wallclock Time
  - Job Size



# Allocation Process

- **University Community**
  - *Review Panel of peers from scientific community*
  - *Atmospheric and related sciences, large allocation request requires associated NSF grant*
  - *“seed” allocations to unsponsored graduate students, postdocs, and new faculty are available*
  - *Life of grant allocation period (term allocation)*



# Allocation Process

- **NCAR Community**
  - *Support NCAR scientific mission*
  - *Allocated annually by the NCAR Director*
  - *12 month allocation period (monthly allocation)*
- **Climate System Laboratory (CSL)**
  - *Review Panel from Climate Programs (multi-agency)*
  - *Climate Change Science Program (multi-agency)*
  - *Large-scale, long running simulations of the Earth's climate system*



# HPC Production Platforms



## IBM Power5

- IBM Power5 p575
- 78 8-way nodes
- 624 single core 1.9 GHz
- Federation Switch
- 4.7 Tflop Peak
- 55 TB Disk Storage
- 210 kW



## IBM Power4

- IBM Power4 p690
- 76 8-way nodes
- 31 32-way nodes
- 1600 1.3 GHz pes
- Colony Switch
- 8.3 Tflop Peak
- 29 TB Disk Storage
- 400 kW

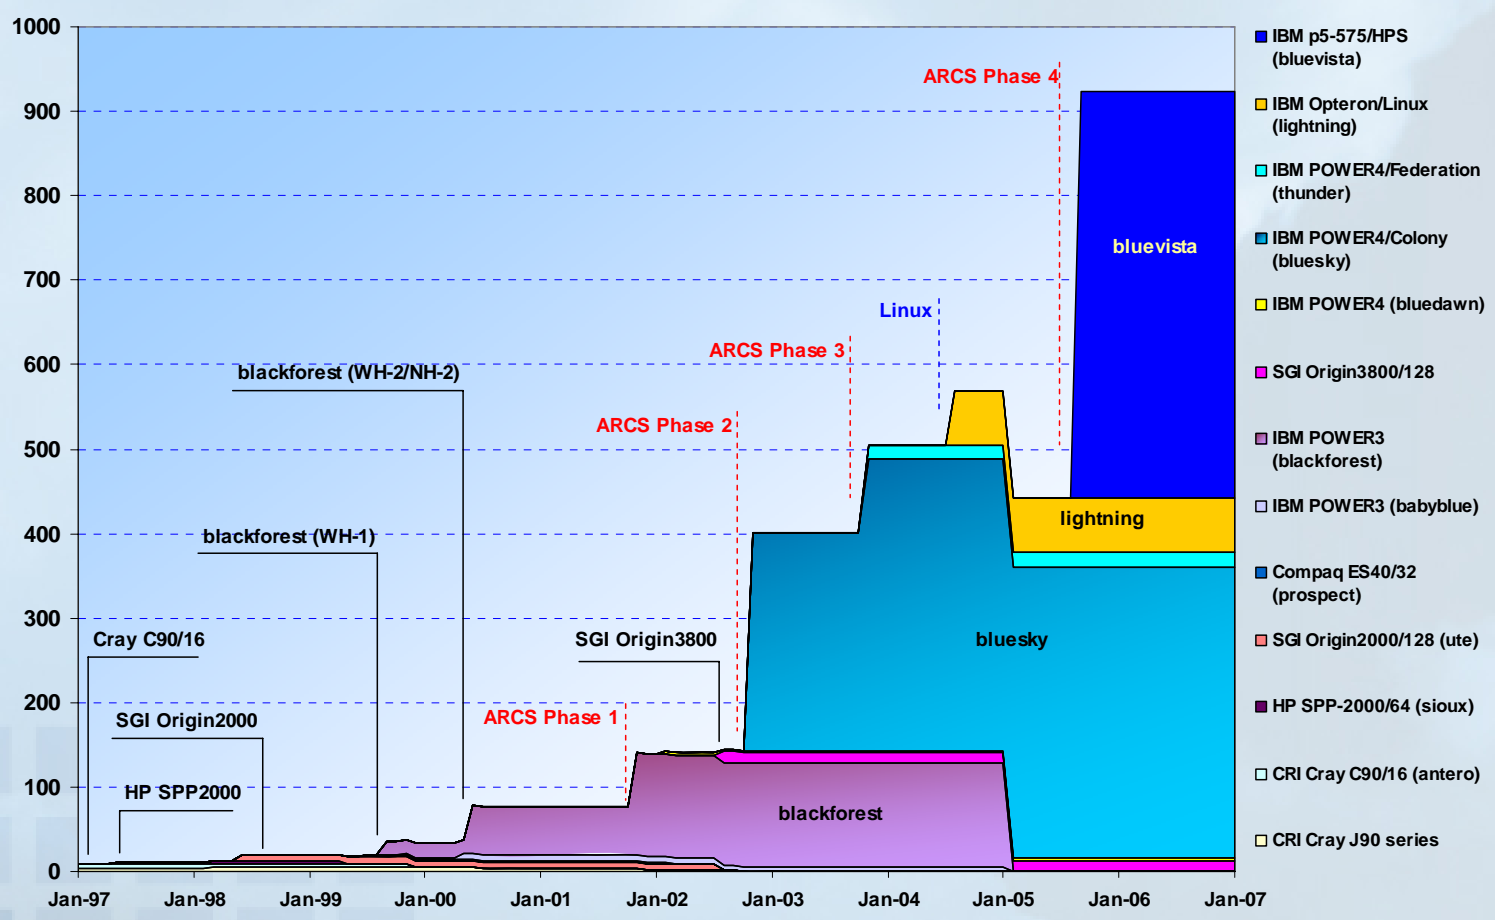


## Linux Cluster

- IBM Cluster e1350
- SuSE Linux
- 132 2-way nodes
- 264 2.2 GHz Opteron pes
- Myrinet Switch
- 1.2 Tflop Peak
- 6 TB Disk Storage
- 48 kW

# Sustained Production GFLOPs

## Estimated Sustained GFLOPs at NCAR (Production Systems)

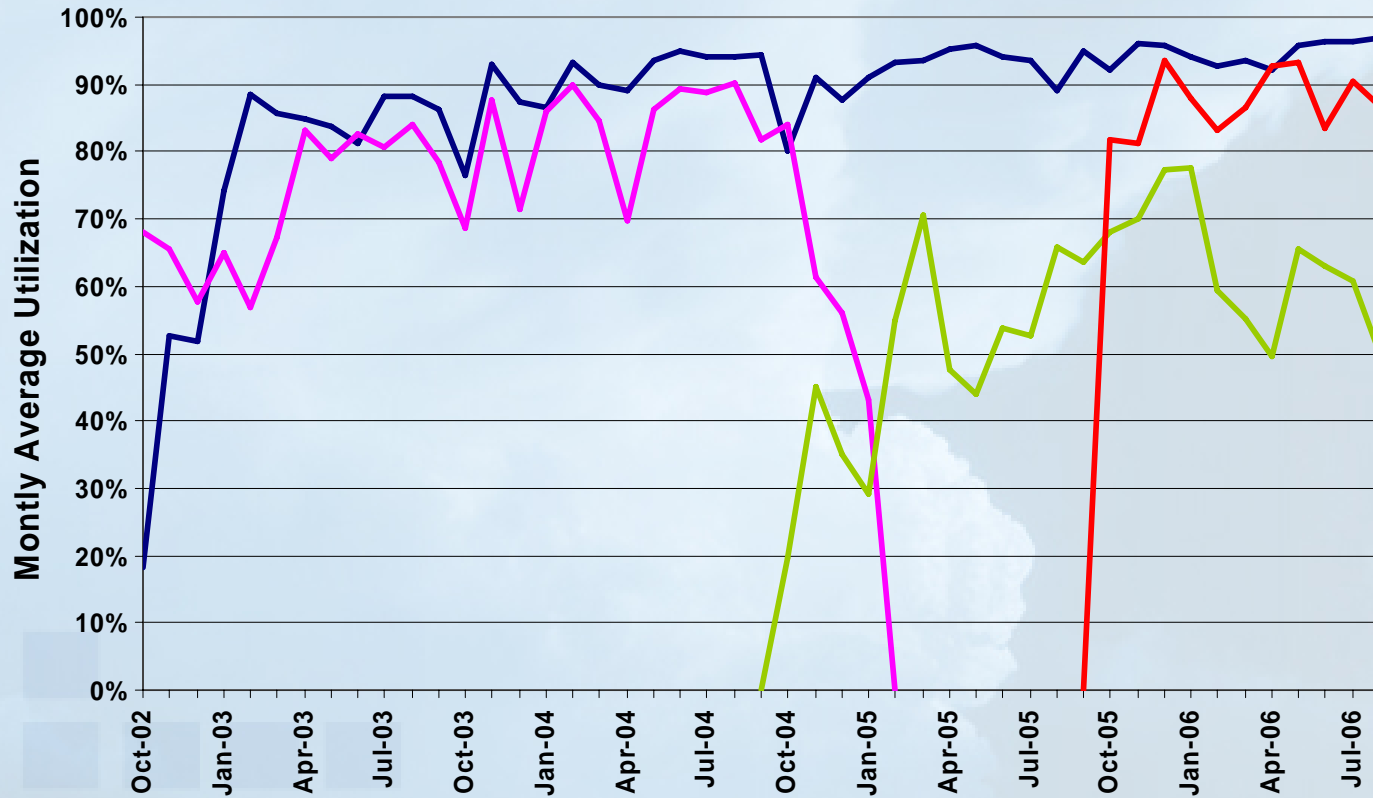




# High Utilization

## Average % Utilization of Production Systems

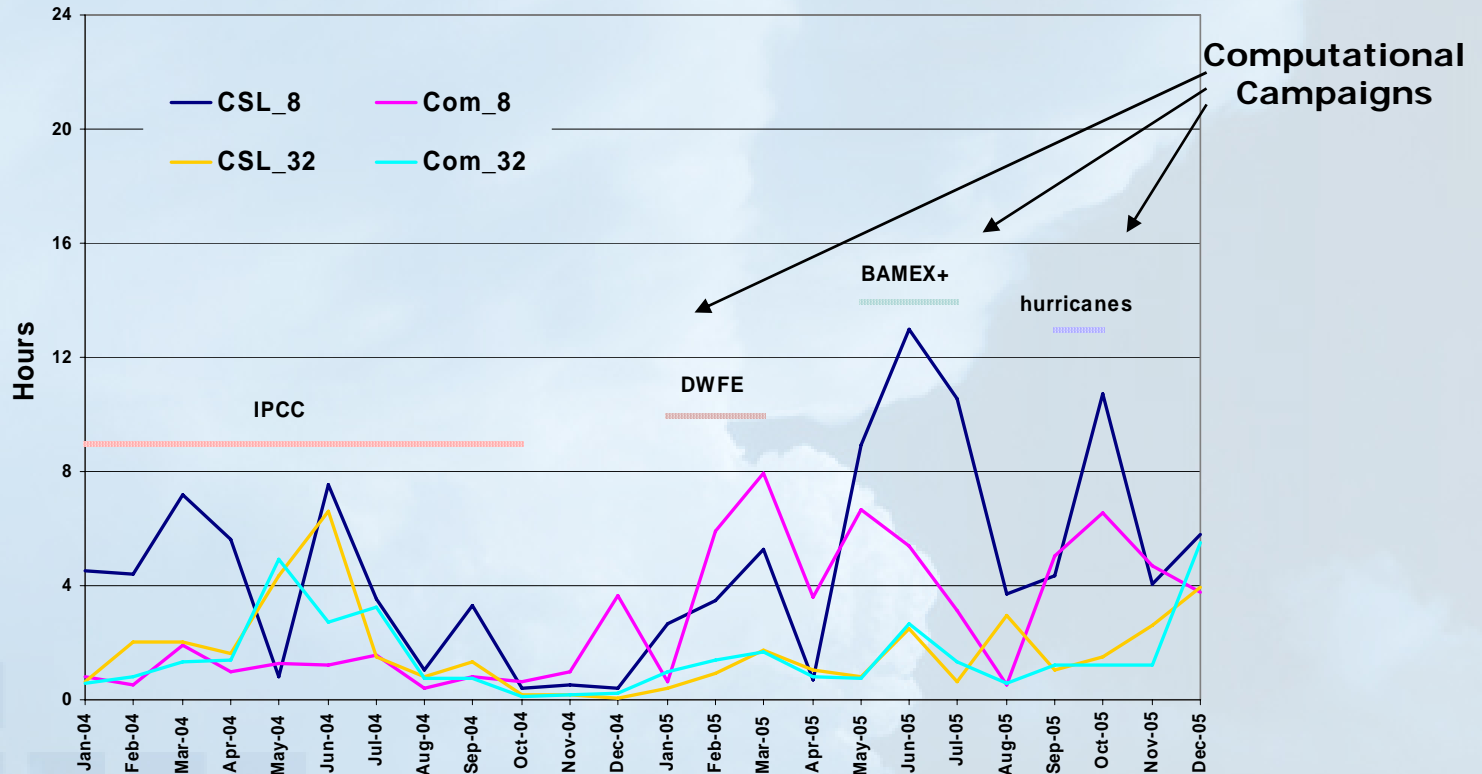
bluesky      blackforest  
lightning      bluevista





# Short Queue Wait Times.....

## Monthly Average Queue Wait Time Regular Queue Jobs on IBM Power4



# HPC Computational Campaigns



- **Definition**

- *Address capability or capacity computation which is otherwise difficult to obtain within NCAR's resource scheduling mechanisms*
- *Support near real-time computational need*

- **Recent Examples**

- *CCSM climate simulations for participation in IPCC (24x7 use of 576 processors over 1 year)*
- *MHD Turbulence simulation with  $1536^3$  grid points requiring 100,000 pe hours*
- *March 2006 real-time weather forecasts in support of field program for the Megacity Impacts on Regional and Global Environments (MIRAGE) in Mexico City*
- *Spring 2004-2006 experimental, real-time convective storm forecasts with WRF*
- *Summer-Fall 2005-2006 Atlantic hurricane forecasts with WRF*
- *Summer 2006 special access by students in the ASP Colloquium "The Art of Climate Modeling"*

# IBM BlueGene/L at NCAR

- **Blue Gene/L is managed by a partnership:**
  - *NCAR*
  - *University of Colorado – Boulder*
  - *University of Colorado – Denver*
- **Scalable platform for prototyping petascale applications.**
- **Blue Gene/L will soon be deployed as a TeraGrid resource.**



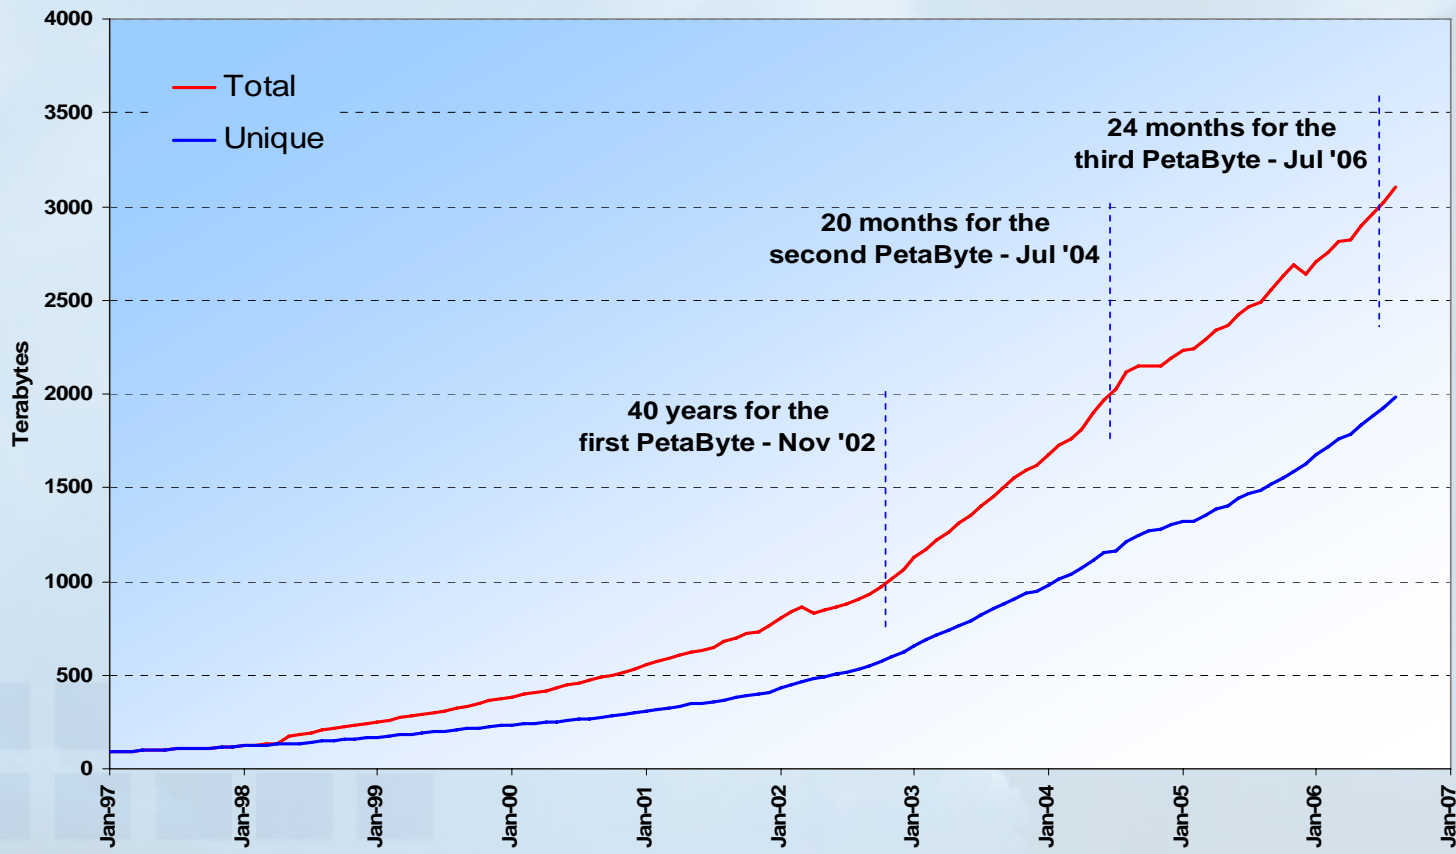
## BlueGene/L

- IBM BlueGene/L
- 1024 2-way nodes
- 2048 0.7 GHz pes
- 5.73 Tflop Peak
- 25 kW



# NCAR Mass Storage System

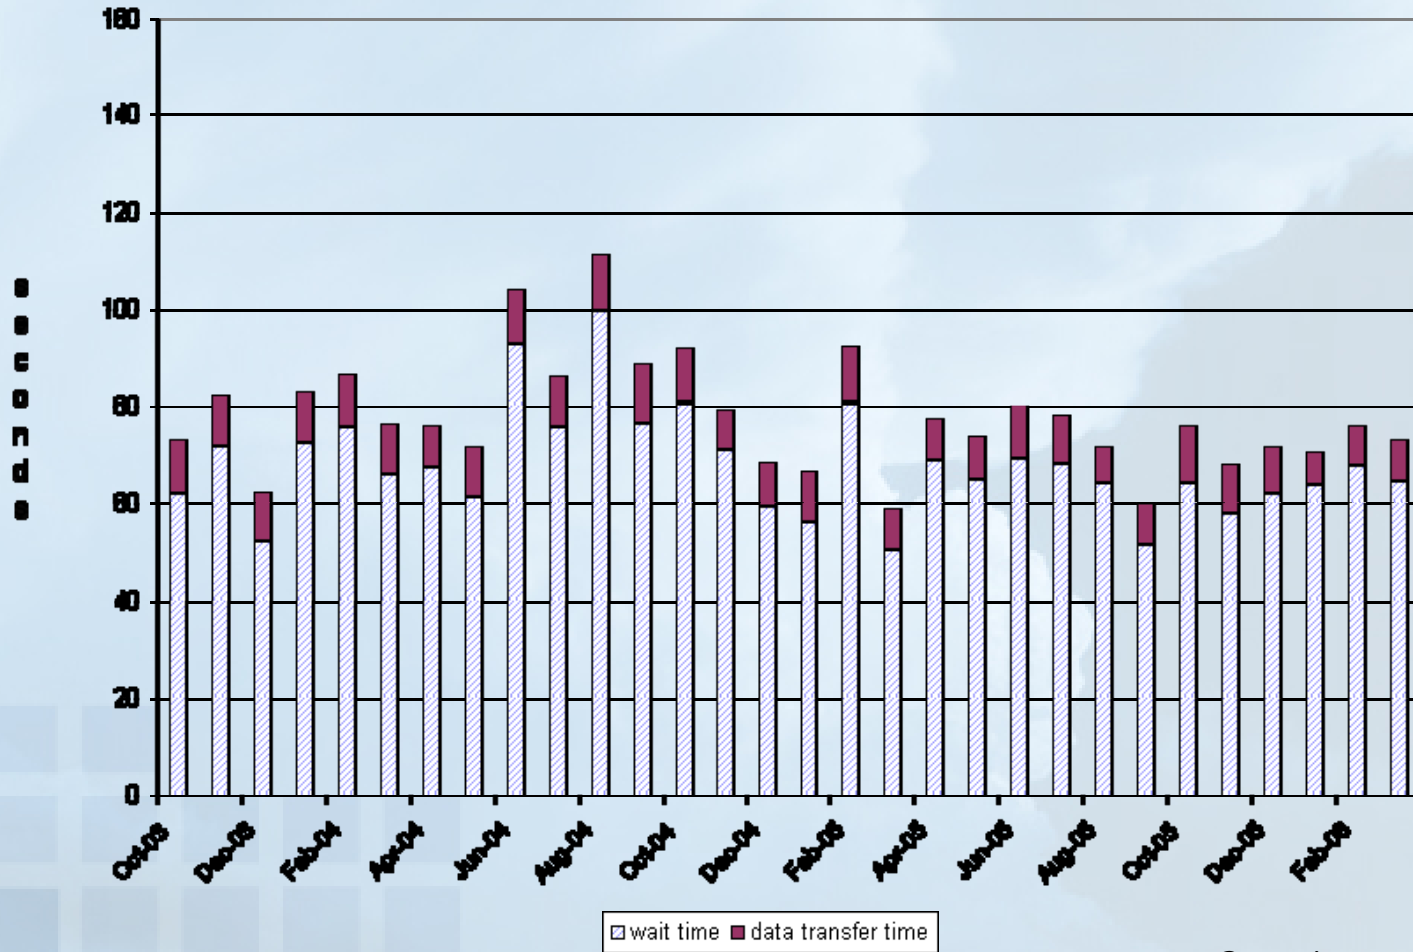
## NCAR MSS - Data Holdings





# MSS Read Response Time

Monthly Average "response" times (reads, all devices)





# Research Data Archives at NCAR


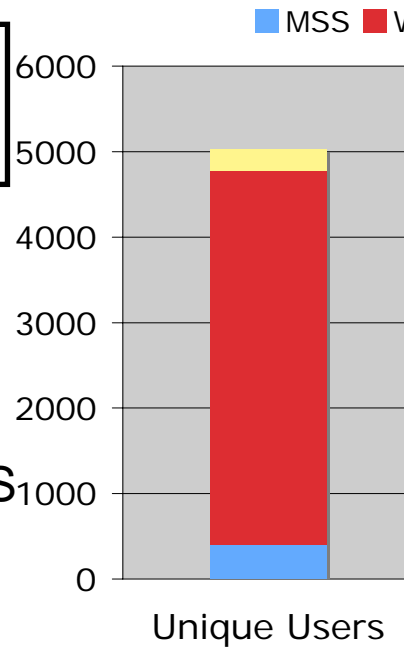
- Collection of reference datasets used in atmospheric and related sciences
- Over 600 datasets currently archived
- 10-20 new datasets added annually
- The RDA is maintained by 9 staff
- First established about 40 years ago
- All data on the MSS
  - *Over 550K files*
  - *Over 110 TB of data*
- Selected data is online.....

**dss.ucar.edu**

# Research Data Archive, Service and Growth - FY06

## Service:


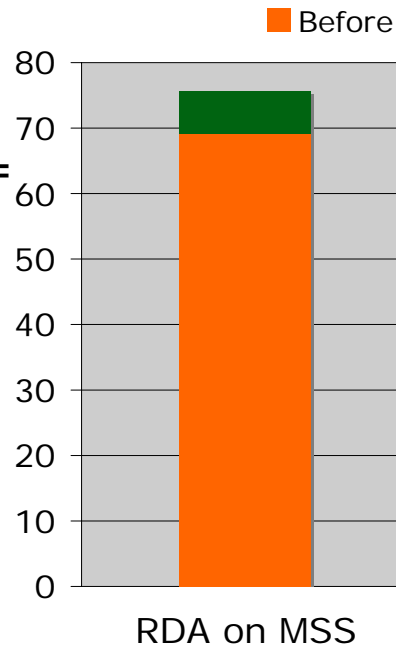
- Unique users topped 5000
  - majority (4.4K) on the public web servers (DSS server + CDP)
- Data provided is 93 TB
  - 2:1 ratio, MSS to Web



## Growth:

- Added 6.5 TB (9%) for total = 76 TB on MSS
- Added 3.8 TB (39%) for total = 13.6 TB on web servers

9/21/06, final FY06 metrics tbd 10/06




# Visualization and Analysis Platform for oceanic, atmospheric, and solar Research (VAPoR)



- **Key concepts**

1. A visual data discovery environment tailored toward the specialized needs of the geosciences CFD community
2. Desktop viz and analysis of terascale datasets in both large spatial and temporal scales
3. Interactively steer the analysis
4. Employ multiresolution data representation as a data reduction technique

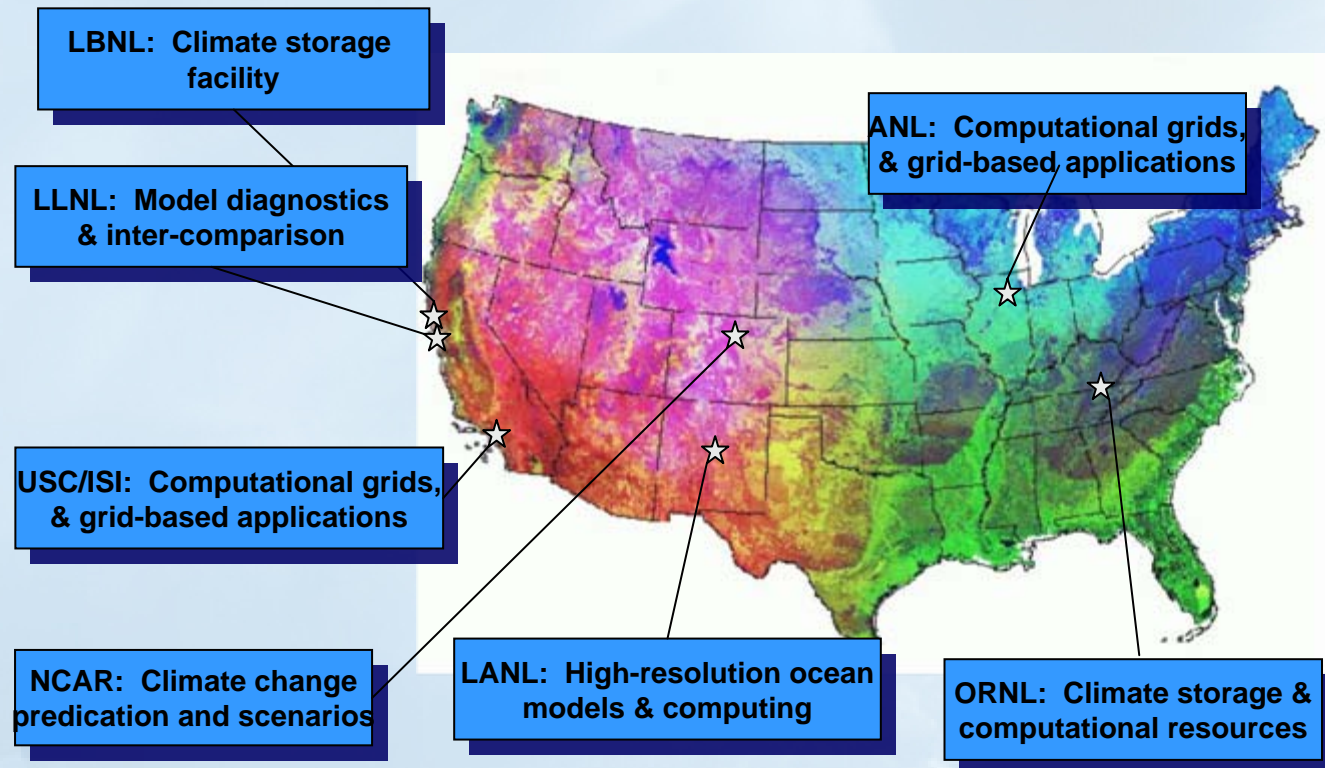


Size: Full    1/8    1/64    1/512

[www.vapor.ucar.edu](http://www.vapor.ucar.edu)



# The Earth System Grid (ESG)





# ESG: An Operational DataGrid for Climate Research

The screenshot shows the Earth System Grid website in a browser window. The browser address bar shows <https://www.earthsystemgrid.org/>. The website has a navigation menu with links for Home, Data, About ESG, and Login. The main content area is divided into several sections:

- ESG News:** Contains registration information and several news items regarding IPCC Working Group 1 data, NCAR MSS downtime, NERSC HPSS maintenance, and ORNL HPSS downtime.
- Data Search:** A search box for dataset metadata with a search button and examples like 'c02, B06.77'.
- Browse Dataset Catalogs:** Lists 'CCSM (Community Climate System Model)' and 'PCM (Parallel Climate Model)' with folder icons.
- Welcome to ESG:** A text block explaining the ESG's role in integrating supercomputers and servers for climate research, accompanied by a colorful map of the United States.
- ESG Collaborators:** A list of institutions including Argonne National Laboratory, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, National Center for Atmospheric Research, Oak Ridge National Laboratory, and University of Southern California/Information Sciences Institute.
- ESG Current Status:** A table showing the status of various services across different labs.

At the bottom of the page, there is a footer with 'Web Portal software version 3.0 © 2004, UCAR. All rights reserved.', 'Login Status: Not logged in.', and 'Problem using site? Contact ESG'.

# NCAR and the TeraGrid



NCAR



Announced at TG2006 in Indianapolis, June 2006

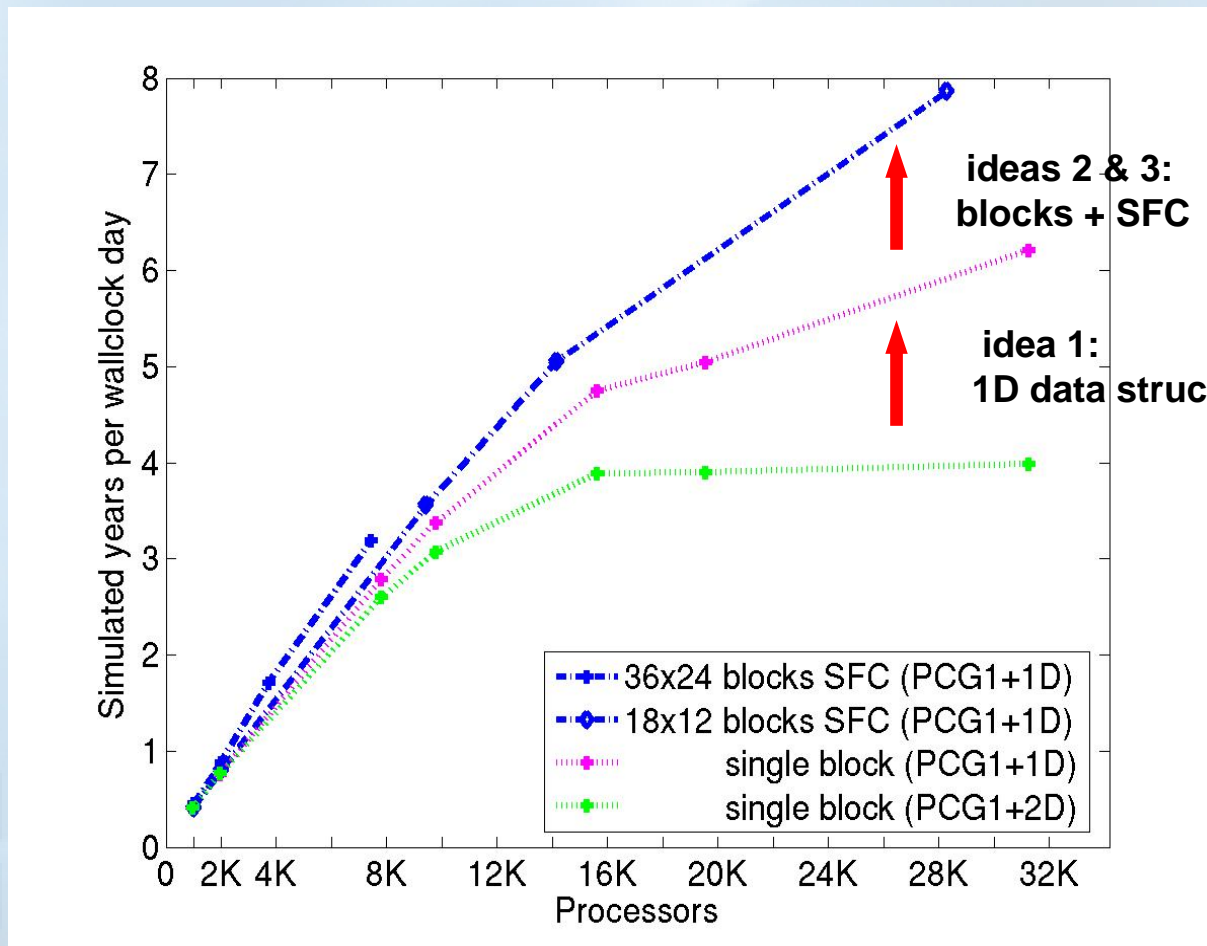
# Preparing POP for the Petascale



NCAR

- Parallel Ocean Program (POP)
  - *Developed at Los Alamos National Laboratory*
  - *Used with CCSM*
  - *Used to simulate basin flows and general oceanic circulation at very high resolution (0.1 degree) – “eddy resolving”*
- Performance Enhancements
  1. *Eliminate land points (from 2D to 1D data structure).*
  2. *Redesign internal data structure (fixed block sizes, reduce load imbalance).*
  3. *Apply Space Filling Curves to improve scalability (enhance load balance).*

# POP 0.1 Degree BG/L Improvement



Courtesy of John Dennis, NCAR



NCAR

# Questions and Discussion