



# **ORNL's Mass Storage System**

**Stan White**

**May 5, 2004**

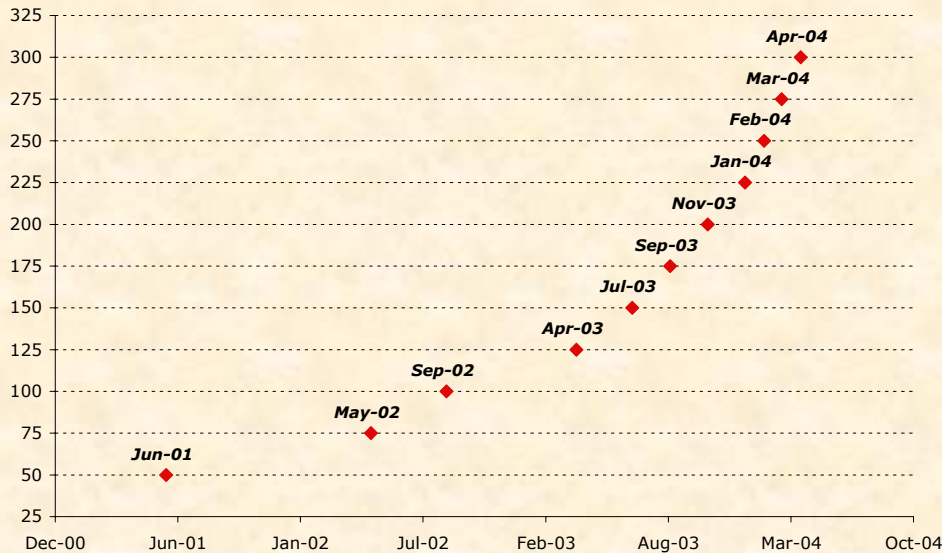
**Storage Systems Analyst  
Storage Networking Group  
whitesr@ornl.gov**

The submitted manuscript has been authored by a contractor of the U.S. Government under Contract No. DE-AC05-00OR22725. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

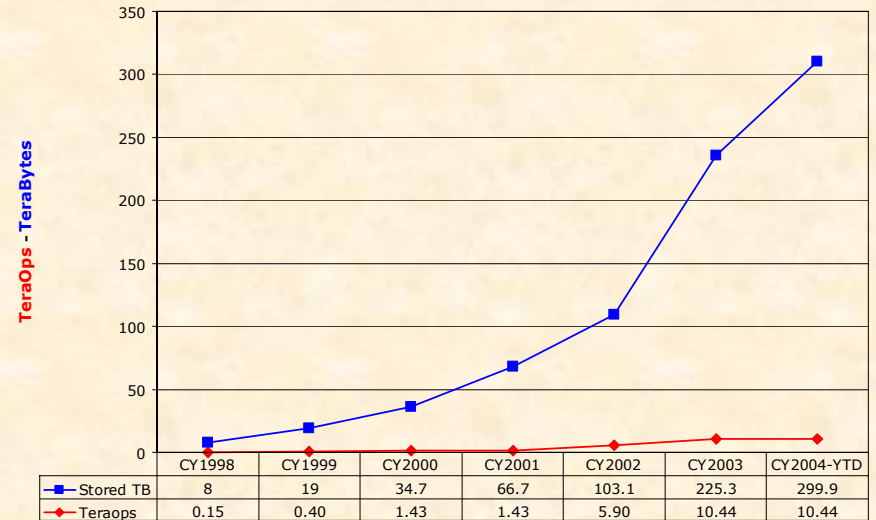
## Production System

- Over 5 Million files
- ~300 Terabytes stored

HPSS Data Growth over Time  
(25 TB Increments)



ORNL CCS Data & Compute Capacity Growth



- Growing ~ 6 TB/week
- Moving as much as 17 TB per day

- **ORNL is responsible for the software to manage HPSS – the Storage System Management (SSM) components.**
  - Several elements
    - System Manager (SM)
    - Data Server (DS)
    - GUI and command-line interfaces (for operators and admins)
- **Developers – Deryl Steinert, Vicky White and Dan Million together with Kathleen Tinch, IBM and Debbie Morford, LLNL**
- **Huge effort**
  - 310 Java source files 265,000 source lines
  - 170 C source files 230,000 source lines
  - 272 auto-generated files 173,000 generated lines  
121,000 executable statements
- **100+ screens, 1000+ variables**
- **ORNL also part of the “HSI shadow” development team.**

**ORNL supports research in a number of scientific disciplines with the HPCC and Storage resources located in the Center for Computational Sciences (CCS).**

- **Astrophysics** -- <http://www.ccs.ornl.gov/astro/>
- **Atmospheric Radiation Measurement** -- <http://www.arm.gov/docs/>
- **Climate and Carbon research** -- <http://www.ccs.ornl.gov/CCR/>
- **Computational Biology** -- <http://www.ccs.ornl.gov/cbi/>
- **Computational Materials** -- <http://www.ccs.ornl.gov/mri/>
- **Computational Sciences** -- <http://www.csm.ornl.gov/>
- **Fusion Simulation** -- <http://www.ccs.ornl.gov/fsi/>
- **Nanoscience** -- <http://www.cnms.ornl.gov/>
- **Neutron Sciences** -- <http://www.sns.gov/>
- **Transportation** -- [http://www.ornl.gov/info/ornlreview/v33\\_3\\_00/features.htm](http://www.ornl.gov/info/ornlreview/v33_3_00/features.htm)

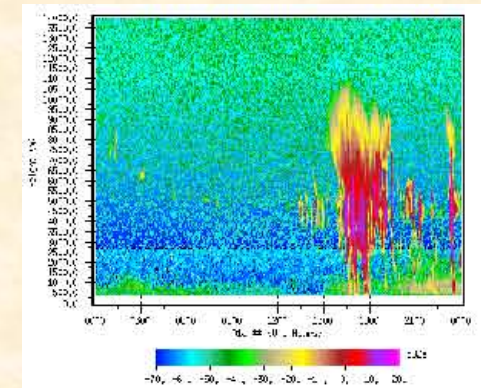
# Atmospheric Radiation Measurement Program (ARM)

The Atmospheric Radiation Measurement Program is the largest global change research program supported by the U.S. Department of Energy (DOE).

- The ARM Program's primary objective is to develop and test parameterizations of clouds and their effect on the radiative energy balance, with the ultimate goal of improving general circulation models used for climate research and prediction.
- To meet this goal, data is gathered at a number of sites around the globe to measure key aspects of the radiation field under vastly different climate conditions.
- ARM data stored in HPSS at ORNL accounts for ~25% of the data stored, and ~half of the total number of files.

**For more information see:**

- <http://www.arm.gov/docs/>
- <http://www.arm.gov/docs/data>



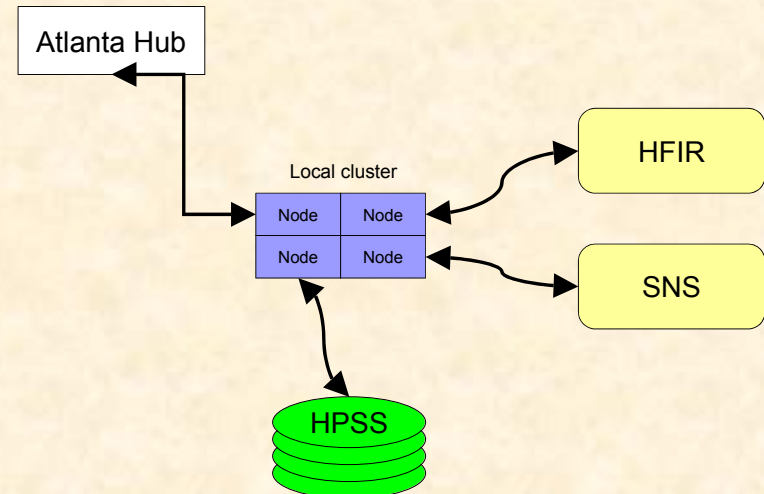
The ARM Program's award-winning Millimeter Wave Cloud Radar (MMCR) is an active sensor that can measure cloud boundaries and microphysical properties with high space and time resolution. This data sample shows a thunderstorm passing above the Central Facility at ARM's Southern Great Plains site

**ORNL intends to bring neutron science to the TeraGrid through the Southeastern TeraGrid Extension for Neutron Science (SETENS).**

- **Provide access to active experimental tools**
  - High Flux Isotope Reactor (HFIR)
  - Spallation Neutron Source (SNS)
- **Access to live and archived neutron scattering data**
  - Real-time access to neutron science data
  - Archived data stores through HPSS

**For more information see:**

- <http://www.csm.ornl.gov/>
- <http://www.sns.gov/>
- <http://www.teragrid.org>





- **Cray X1 (aka SV2) (phoenix)**
  - 256 processors (12.8 GF per processor)
  - 1 TB memory
  - 32 TB disk
  - 3.2 TeraOps Peak



- **IBM pSeries 690 (cheetah)**
  - 864 processors (1.3GHz Power4)
  - 1.1 TB memory
  - 40 TB disk
  - 4.5 TeraOps Peak
  - Federation switch (moving from Colony switch)



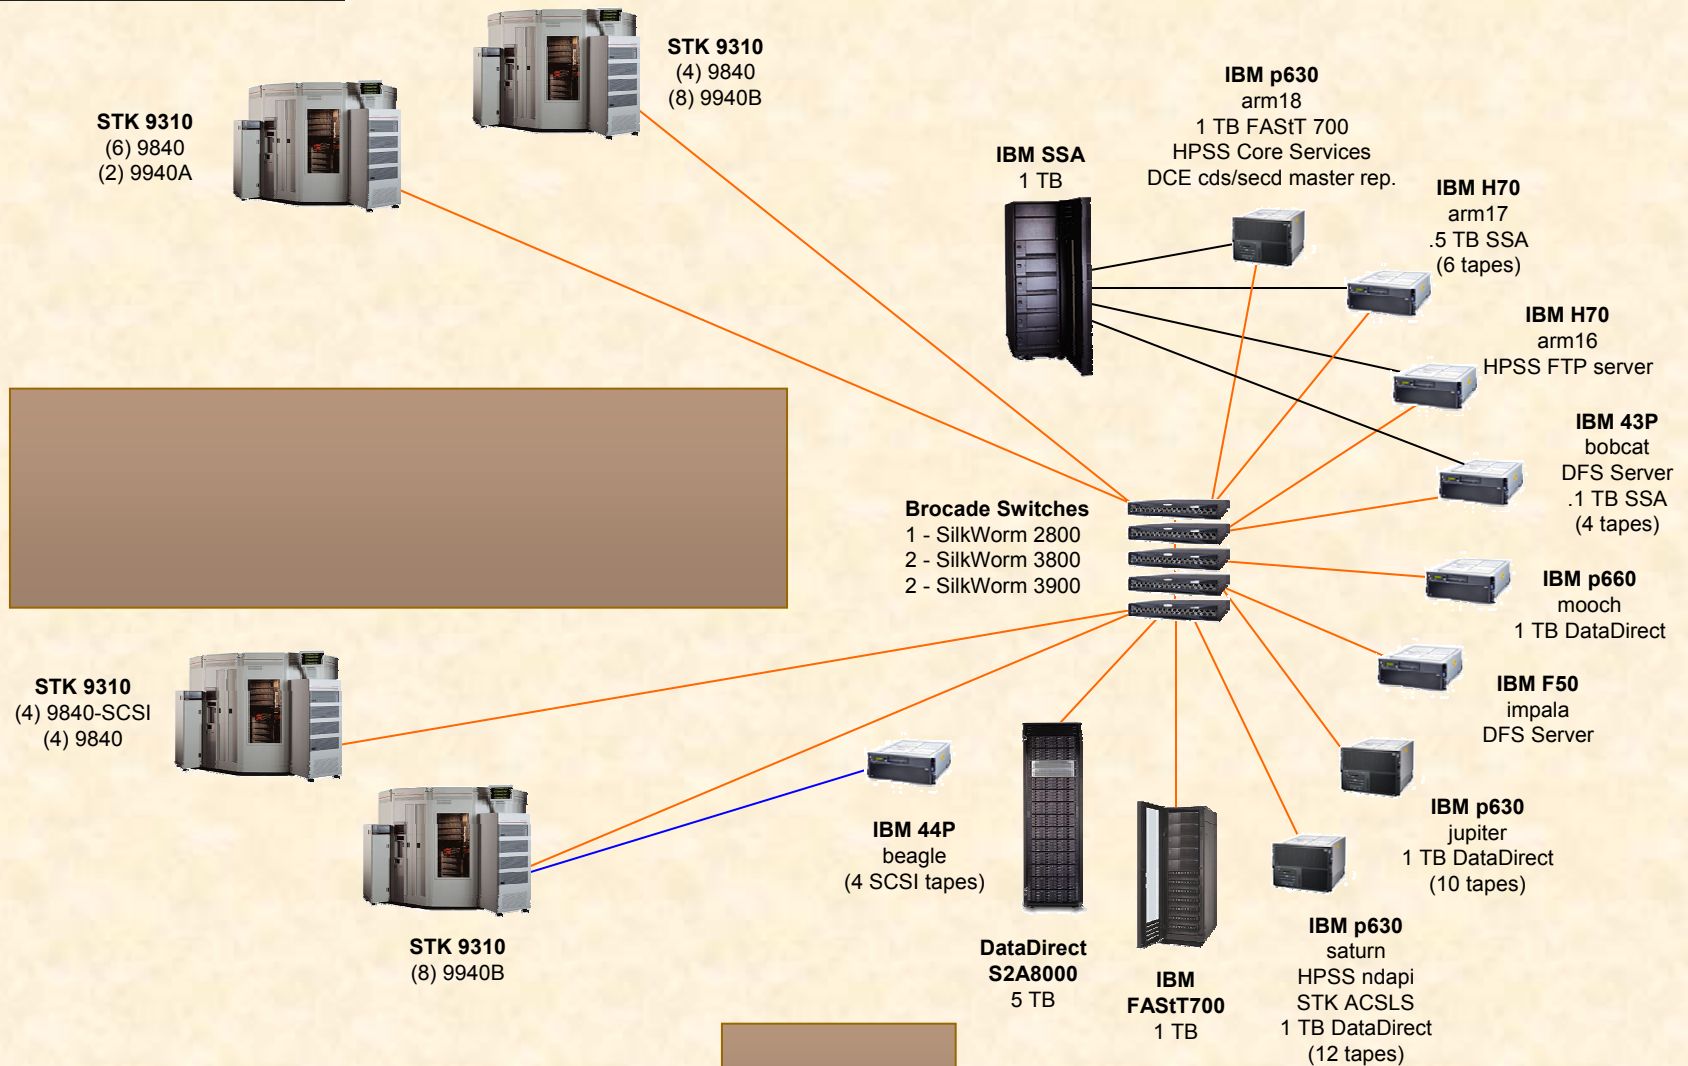
- **SGI Altix (RAM)**
  - 256 processors (1.5GHz Itanium2)
  - 2 TB memory (Single system image)
  - 12 TB disk
  - 1.53 TeraOps Peak



- **IBM SP (eagle)**
  - 704 compute processors (375 MHz Power3-II Winterhawk)
  - 372 GB memory
  - 9.2 TB disk
  - 1.08 TeraOps Peak

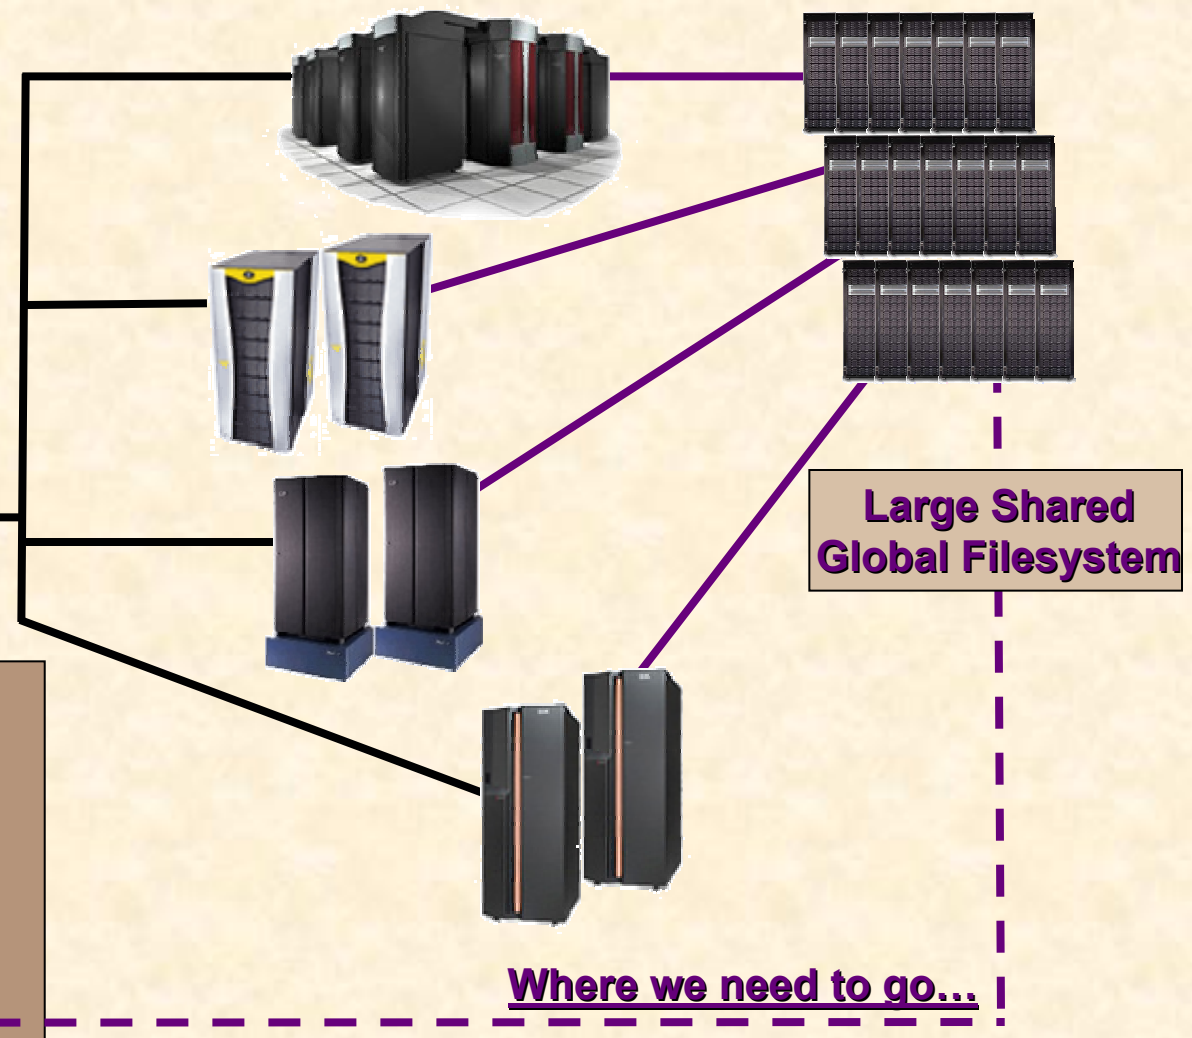


# ORNL HPSS Storage Environment



## Testing Shared File Systems

Likely to test  
ADIC's  
StorNext File System  
& Lustre File System  
3Q - 2004



Where we are today

Large Shared Global Filesystem

Where we need to go...

**Questions?**

---

**THE END**