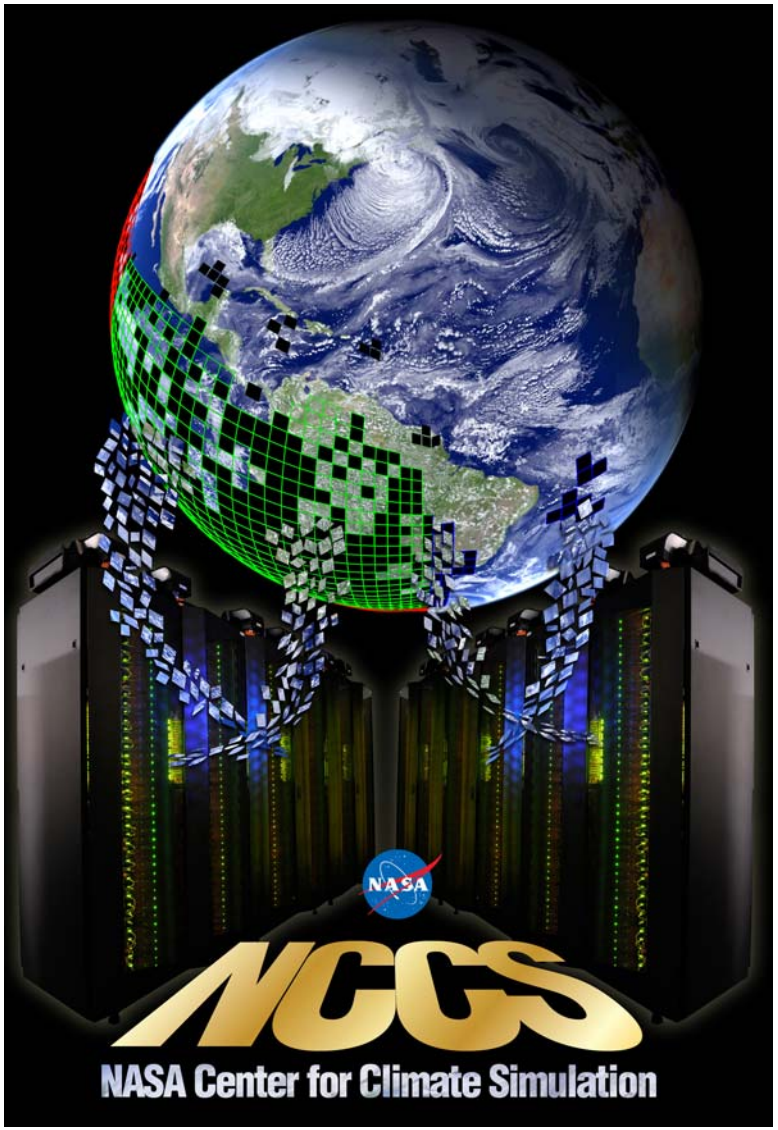

Data Serving Climate Simulation Science

at the

NASA Center for Climate Simulation (NCCS)

MSST2011, May 24-25, 2011

Ellen Salmon (Ellen.Salmon@nasa.gov)
High Performance Computing, Code 606.2
Computational & Information Science and Technology Office (CISTO)
NASA Goddard Space Flight Center
Greenbelt, MD 20771
http://nccs.nasa.gov
http://cisto.gsfc.nasa.gov

Standard Disclaimers and Legalese Eye Chart

- All Trademarks, logos, or otherwise registered identification markers are owned by their respective parties.
- Disclaimer of Liability: With respect to this presentation, neither the United States Government nor any of its employees, makes any warranty, express or implied, including the warranties of merchantability and fitness for a particular purpose, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.
- Disclaimer of Endorsement: Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government. In addition, NASA does not endorse or sponsor any commercial product, service, or activity.
- The views and opinions of author(s) expressed herein do not necessarily state or reflect those of the United States Government and shall not be used for advertising or product endorsement purposes.
- All errors in this presentation are inadvertent and are the responsibility of the primary author.

Presentation Roadmap

- Climate Simulation Science at NASA
- About the NCCS
- NCCS Evolution
- NCCS Archive
- NCCS Data Sharing / Data Portal
- Future NASA Climate Simulation Science

Climate Simulation Science at NASA

NASA Science Mission Directorate

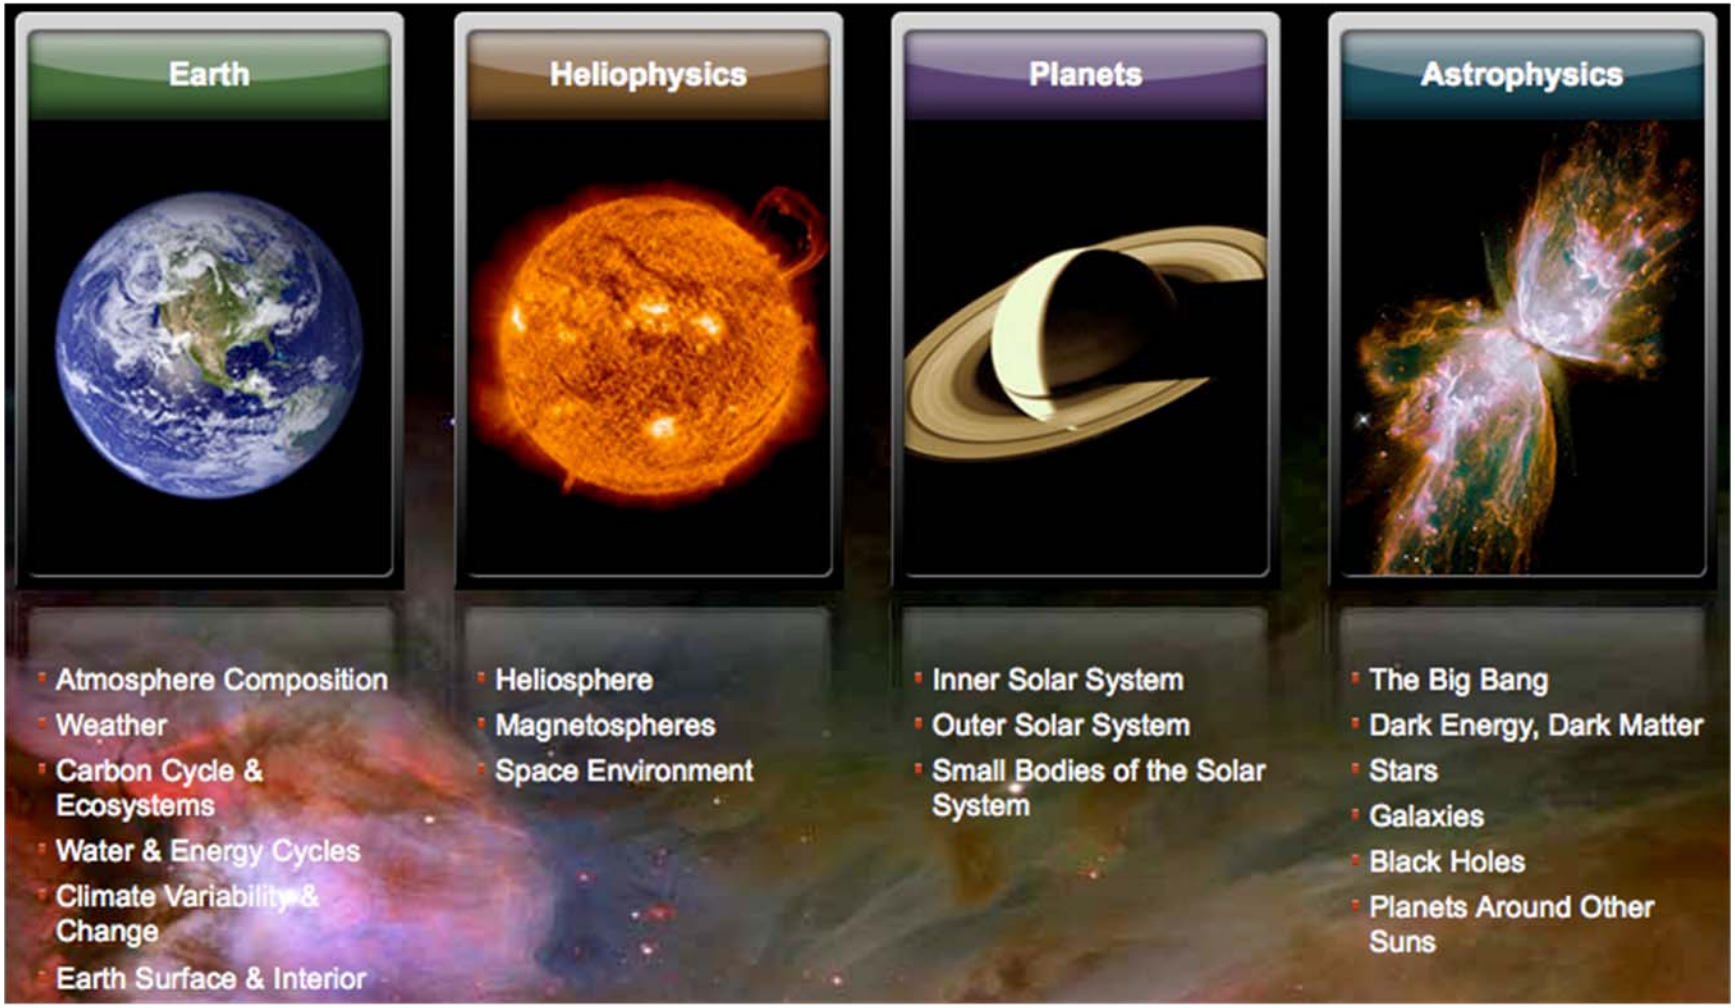

NASA Climate Science Computing

"Data Intensive". . . . . .

- Spans timescales from weather to short-term climate prediction to long-term climate change.
- Brings models and observations together through data assimilation and simulation.
- Creates products to support NASA instrument teams and atmospheric chemistry community.
- Reanalysis results in vast data sets (100s of terabytes) for the scientific community.
- Climate models produce large data sets (100s of terabytes) for the scientific community as well as decision makers.

. . . . . And Requires "Data Centric" Computing

- Designed for effective manipulation of large data sets.
- Global file system makes data available to all services. Effective data management tools are required.
- Efficient data analysis needs to have "supercomputing" capability with data sets online.
- Data sets must be made easily accessible to "external users" with analysis and visualization capability.

Climate Simulation Data Communities

NASA Scientific Community
- Simulation data consumers
- Advance scientific knowledge
- Direct access to systems
- *Supercomputer* capability required for effective analysis

NASA Modeling Community
- Model development, testing, validation, and execution
- Data creation
- Largest HPC usage
- Requires observational data as input

Simulation Data

External Applications Community
- Huge opportunity for impact from climate change data
- Simulation data consumers
- Limited ES data expertise
- Web-based access to systems

External Scientific Community
- Simulation data consumers
- Advance scientific knowledge
- Web-based access to data

Each community has different capabilities and data usage requirements.

Data Centric Architecture

Analysis and Visualization

Terascale *environment with tools to support interactive analytical activities*

High Performance Computing

Building toward ***Petascale*** *computational resources to support advanced modeling applications*

Data Storage and Management

Petabyte *online storage plus technology-independent software interfaces to provide data access to all NCCS services*

Data Archiving and Stewardship

Petabyte *mass storage facility to support project data storage, access, and distribution, access to data sets in other locations*

Data Sharing and Publication

Web-based environments to support collaboration, public access, and visualization

About the NCCS

NASA Center for Climate Simulation (NCCS)

- Operated by NASA Goddard's High Performance Computing (HPC) Group.
- Provides HPC and data services designed for climate, ocean, and weather simulation and other NASA science research.
- Serves a research community based at NASA centers and laboratories and universities across the country and internationally.
- Maintains advanced data capabilities and facilities that allow researchers to access the enormous volume of data generated by weather and climate models running on the facility's supercomputers.
  - mass-storage system (archive)
  - data distribution technologies
  - high-speed networks
- Makes available data analysis and visualization tools needed to interpret modeling data.

NCCS Data Centric Climate Simulation Environment

Data Sharing and Publication

- Capability to share data & results
- Supports community-based development
- Data distribution and publishing

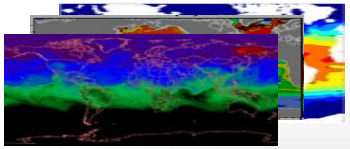

Code Development*

- Code repository for collaboration
- Environment for code development and test
- Code porting and optimization support
- Web based tools

User Services*

- Help Desk
- Account/Allocation support
- Computational science support
- User teleconferences
- Training & tutorials

DATA Storage & Management

Global file system enables data access for full range of modeling and analysis activities

Analysis & Visualization*

- Interactive analysis environment
- Software tools for image display
- Easy access to data archive
- Specialized visualization support

Data Transfer

- Internal high speed interconnects for HPC components
- High-bandwidth to NCCS for GSFC users
- Multi-gigabit network supports on-demand data transfers

HPC Computing

- Large scale HPC computing
- Comprehensive toolsets for job scheduling and monitoring

Data Archival and Stewardship

- Large capacity storage
- Tools to manage and protect data
- Data migration support

** Joint effort with SIVO*

Some Recent NCCS Highlights

Discover – First production IB commodity cluster in NASA HEC

SL8500 tape libraries and T10000A tape drives

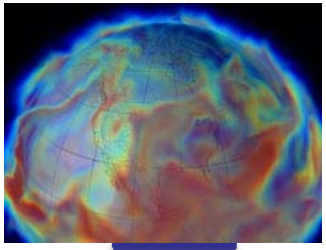

Introduction of Dali analysis environment, interactive access to large data sets

Climate Data Theater New tool for scientific analysis and visualization

Dirac moves to Parallel DMF Intel Xeon Cluster, increases disk to 400 TB

2007 **2009** **2010**

Introduction of the Data Portal System and Services
WMS, MAP, Cloud Library

Dirac DMF moves to SGI Linux Itanium, increases disk to 130 TB

Discover upgrades support 3 KM global resolution model – highest resolution ever run to date

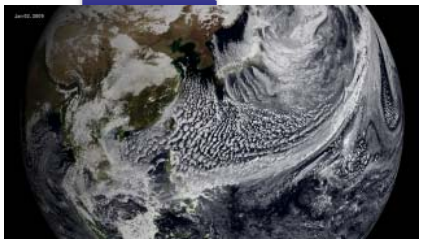

T10000B tape drives

Advancement of Data Services with the Data Management System (DMS) and the Earth Systems Grid (ESG)

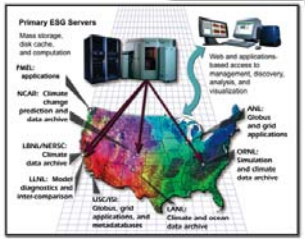

NCCS Archive

Notional NCCS Architecture

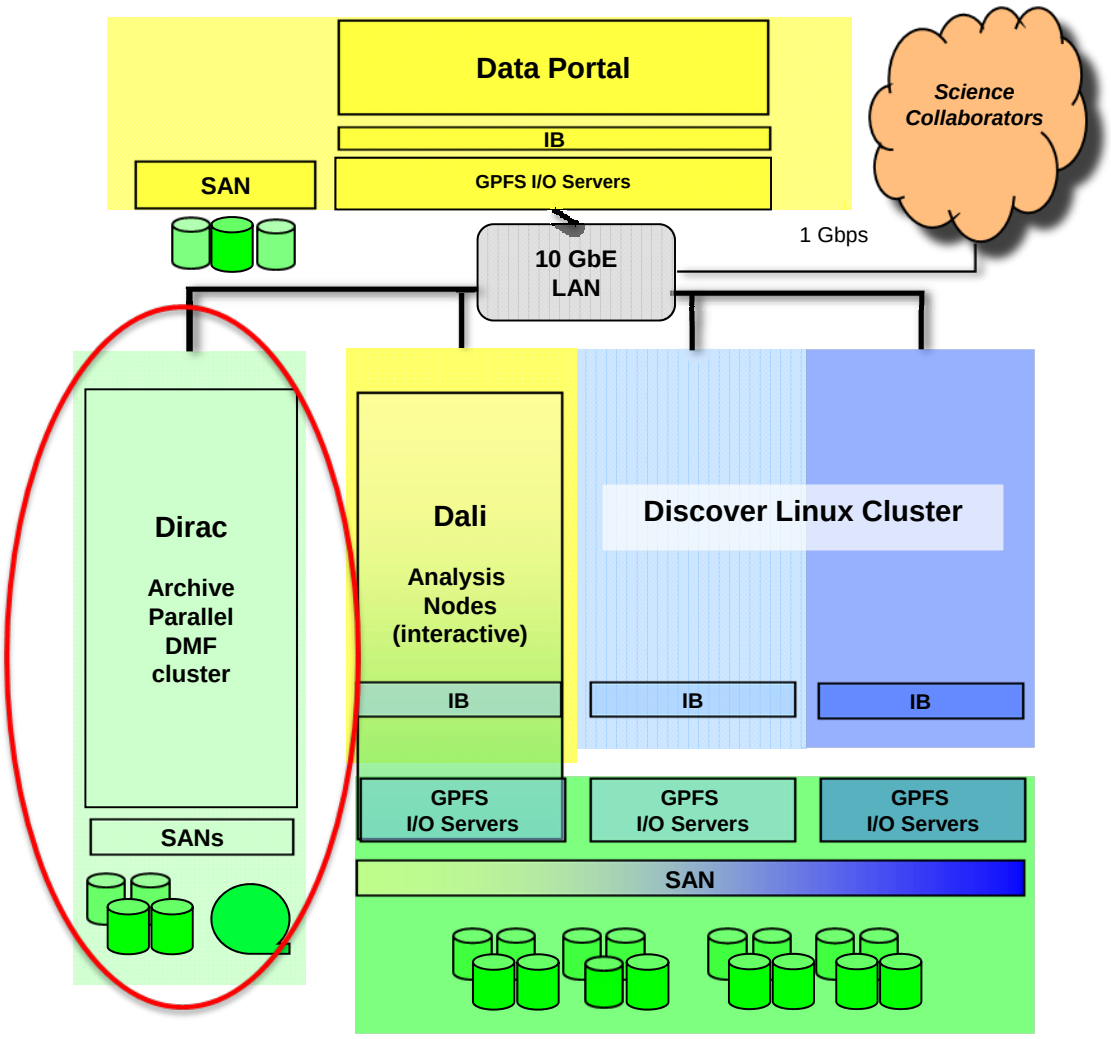

Today's Dirac: Parallel DMF Cluster

- SGI DMF 5.3, ISSP 2.1, SLES11
- Parallel DMF Cluster
  - SGI Altix XE (Intel Xeon Westmere)
    - 2 @ High-Availability DMF Metadata Servers (Active/Passive)
  - SGI C1103-TY12 (Intel Xeon Nehalem)
    - 2 @ NFS Edge Servers
    - 3 @ Parallel Data Movers
    - 2 @ User Login Nodes
- Disk Arrays
  - DMF Disk Cache on SGI IS15000 (600 TB raw)
  - DMF Databases, user /home directories on SGI IS220 (13.5 TB raw)
- Tape Libraries
  - Primary copies
    - 2 Oracle / Sun / StorageTek SL8500 libraries
      - 16 @ 9940B drives (media 200 GB uncompressed)
      - 16 @ T10000B drives (media 1 TB uncompressed)
    - 6 StorageTek 9310 libraries (Powderhorns)
      - 11 @ 9940B tape (media GB uncompressed)
  - Risk Mitigation Duplicate copies
    - 3 StorageTek 9310 libraries
      - 6 @ 9940B drives (media 200 GB uncompressed)
      - 12 @ T10000A drives (media 500 GB uncompressed)
- Near-term upgrades and plans
  - Additional 600 TB (raw) disk
  - Additional SL8500 library for duplicate copies
  - T10000C tape drives for primary and duplicate copies (media 5 TB uncompressed)
  - Continue migrating from old media, libraries

NCCS Archive: Millions of Files and TiB (2^{40} Bytes) Stored, June 2006 – May 2011

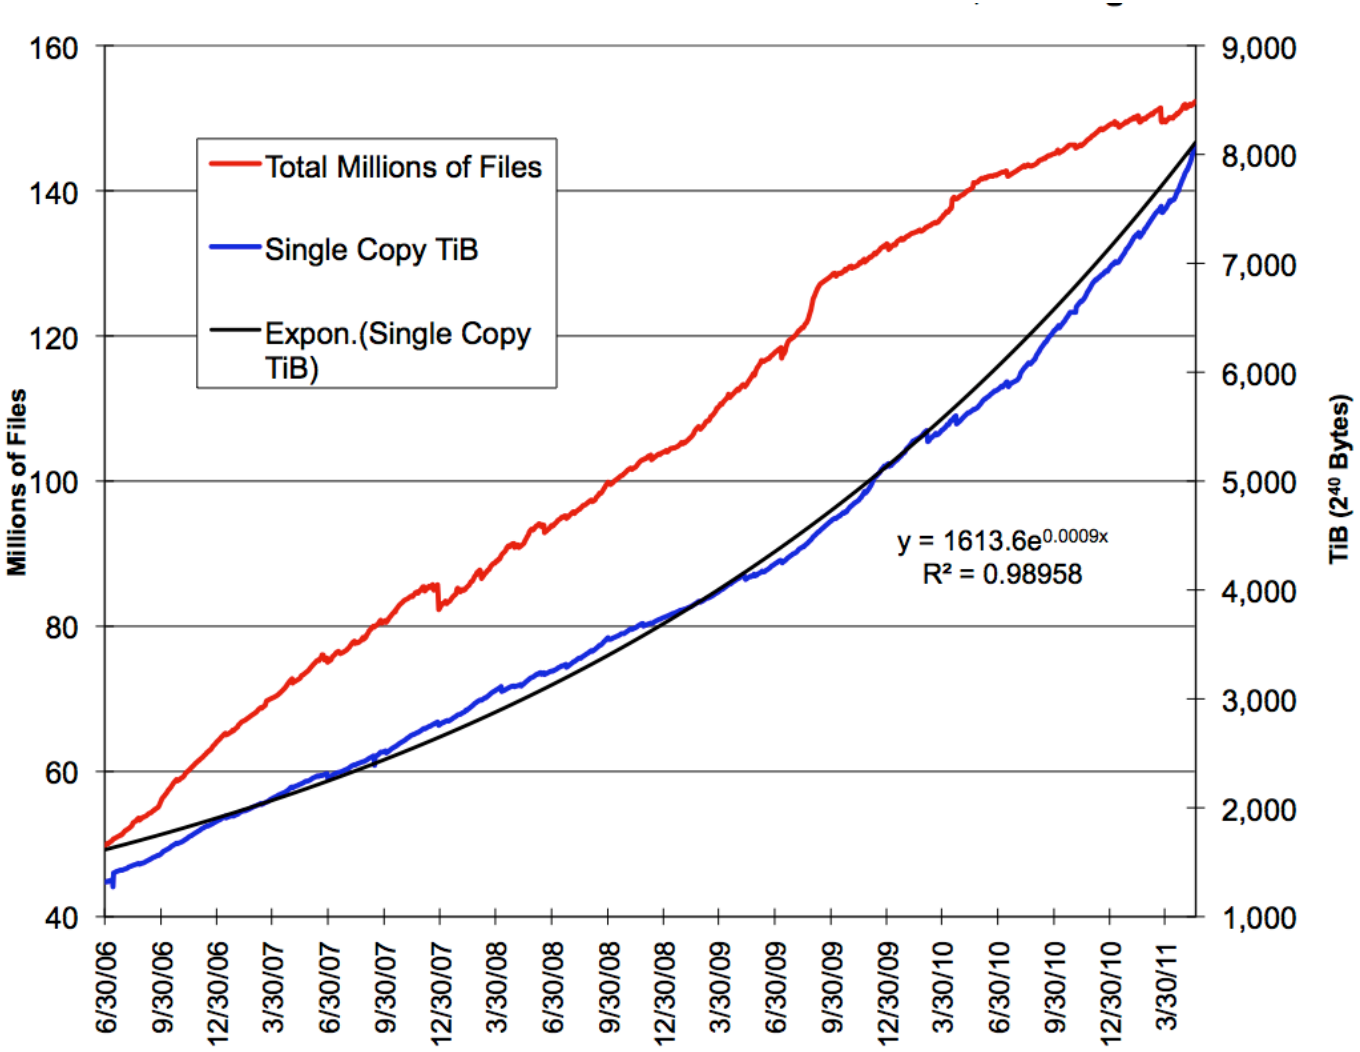

NCCS Archive: Age of Files

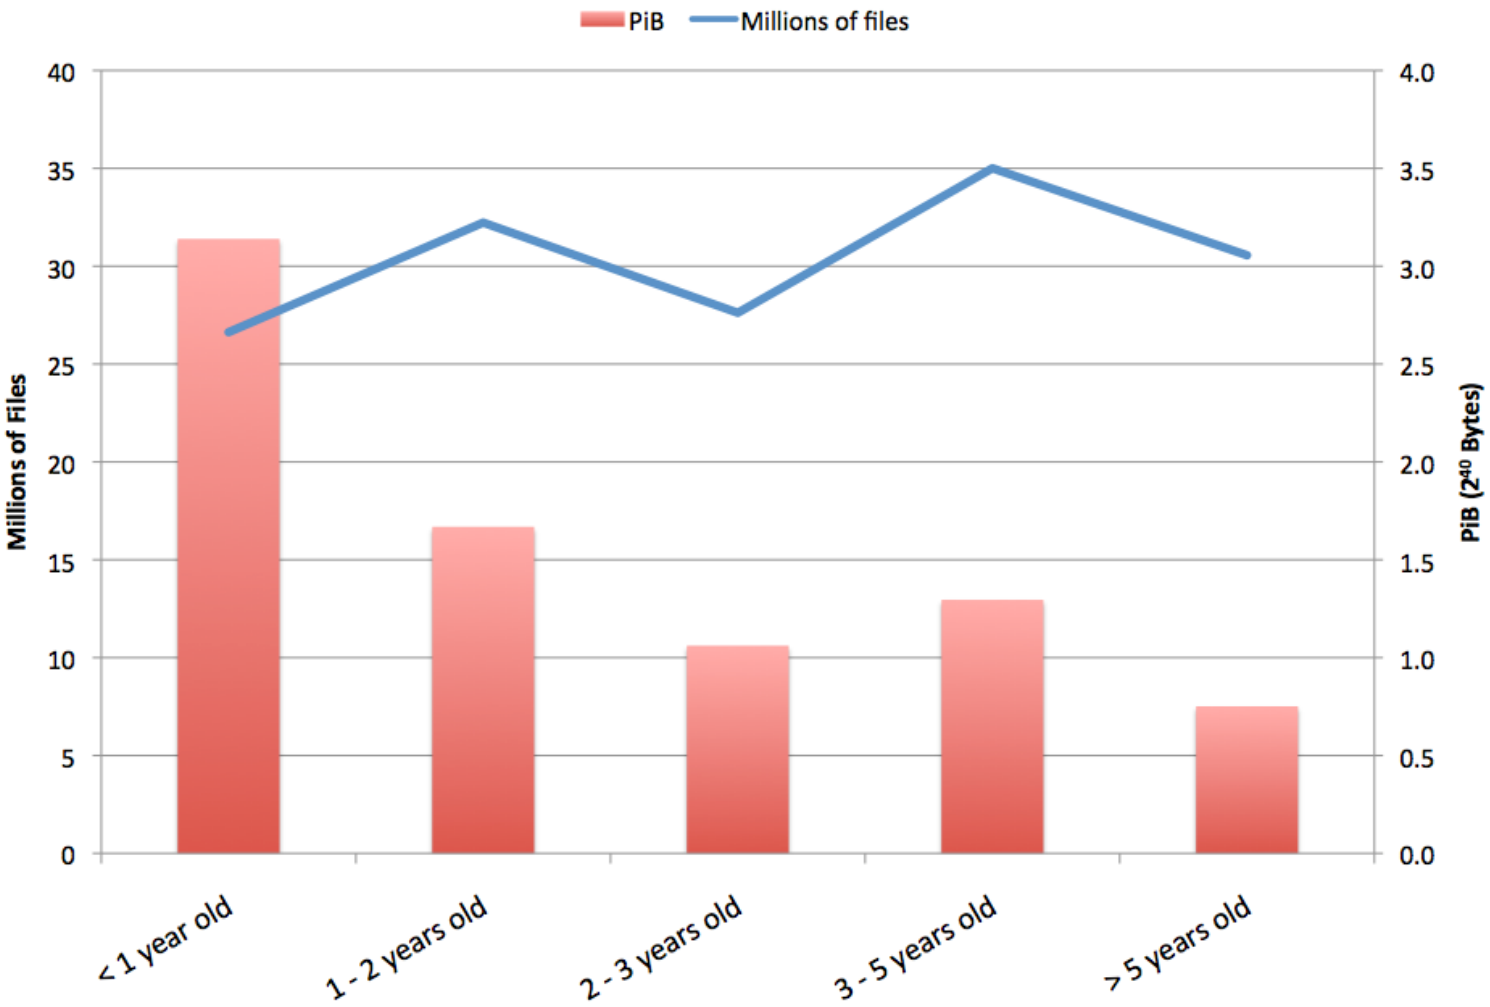

NCCS Archive: Age of Retrieved Files, November 2010 – May 2011

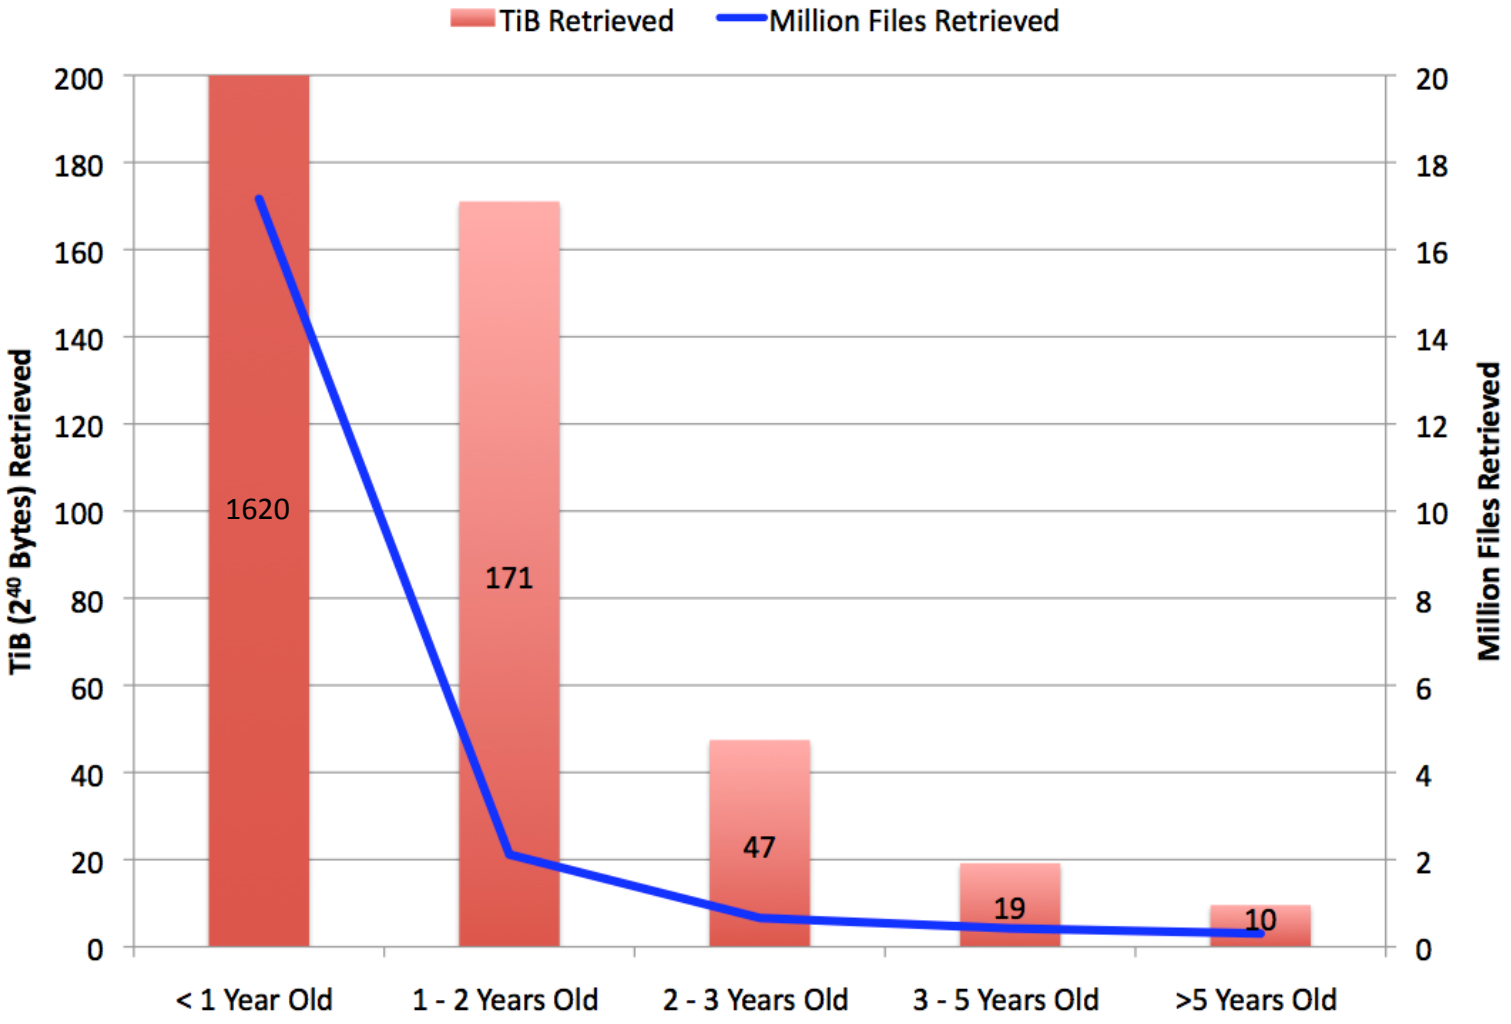

NCCS: Archive Monthly Net Growth and Average High Performance Computer TFLOPS, 2006 – 2011

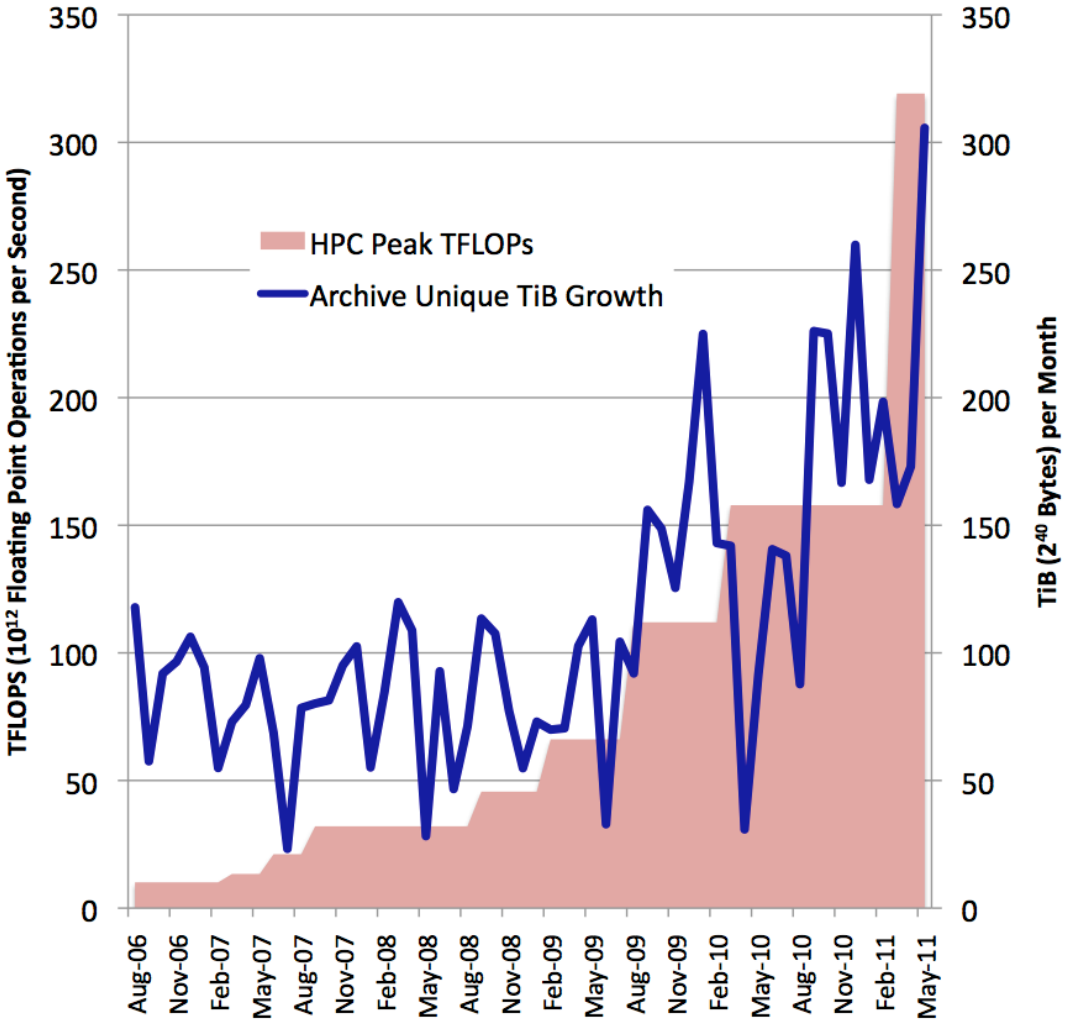

NCCS Archive – Monthly User Read / Write Activity and High Performance Computer Core-Hours, 2008 – 2011

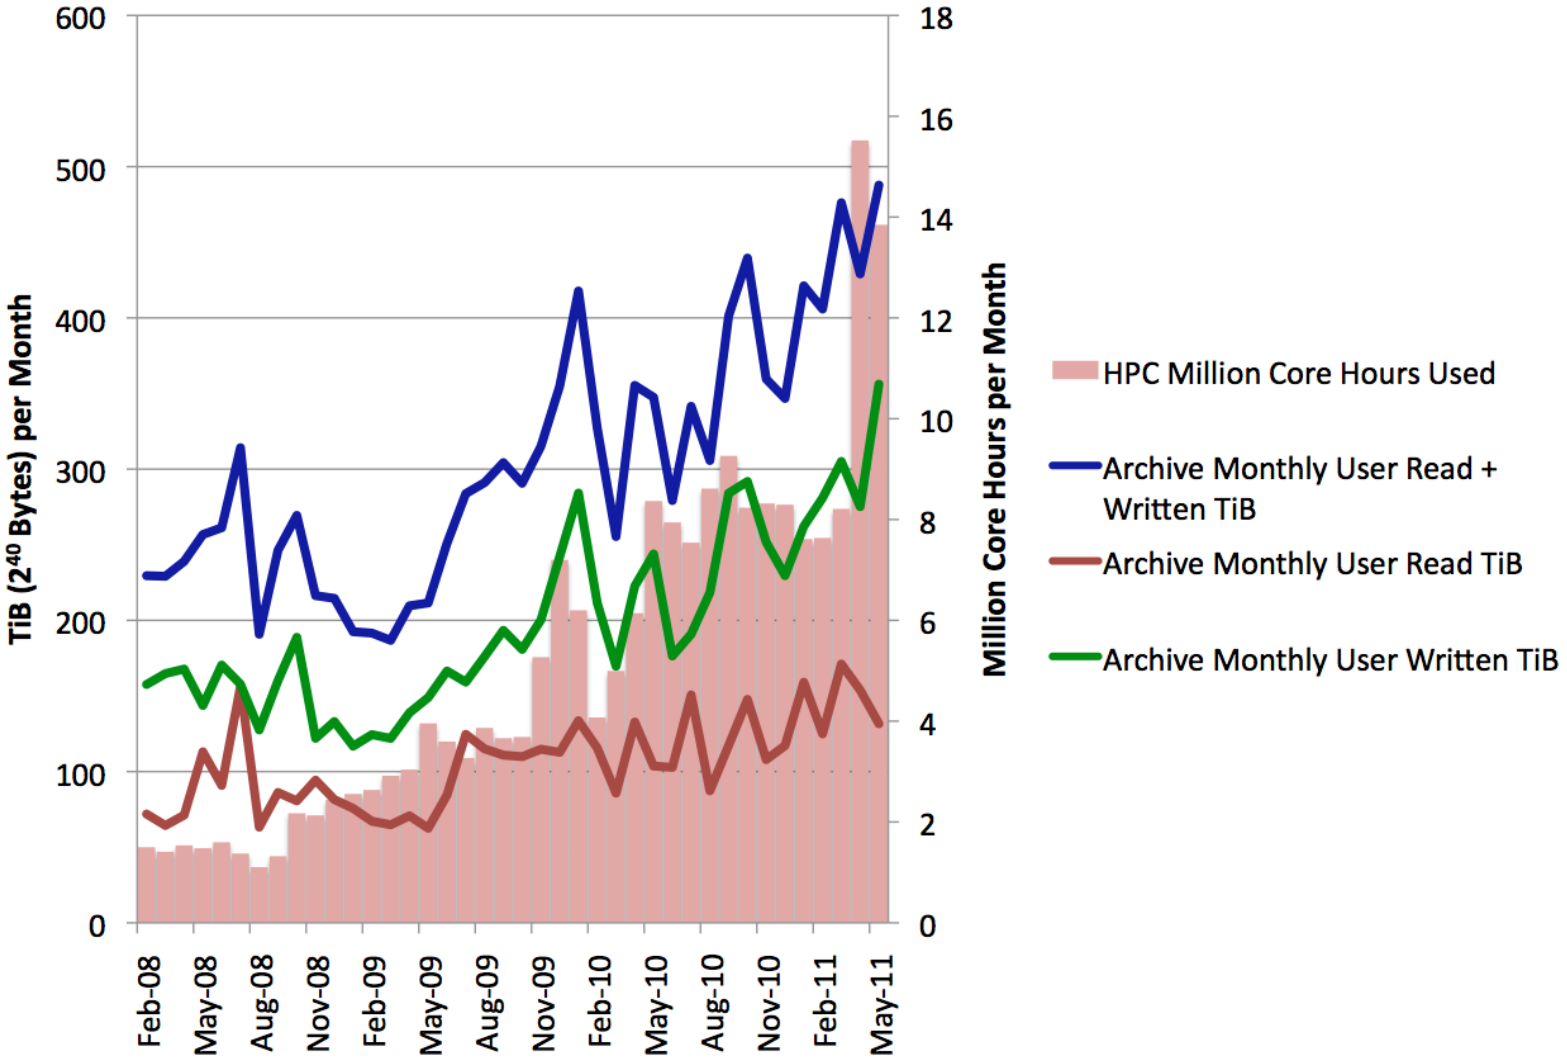

NCCS Archive – Tape Media Migration October 2009 – April 2010

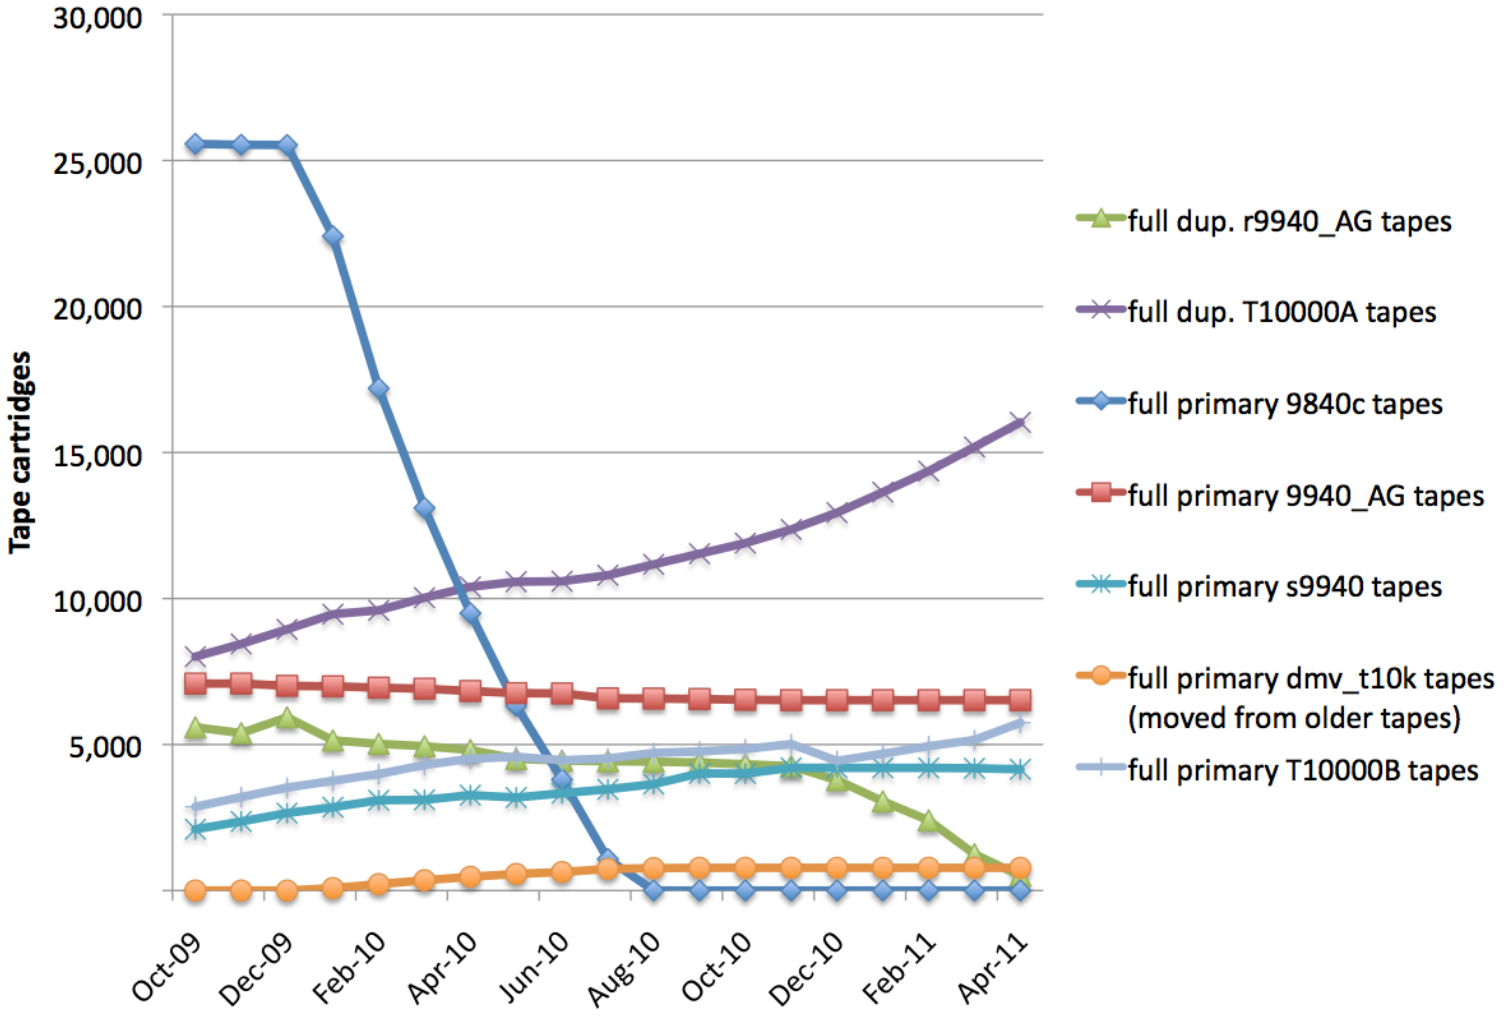

Large-Scale Climate Simulations on Discover's Intel Xeon Westmere SCU7

Discover's recent SCU7 upgrade is being used for climate research experiments and simulations supporting the Fifth Reassessment (AR5) for Intergovernmental Panel on Climate Change (IPCC).

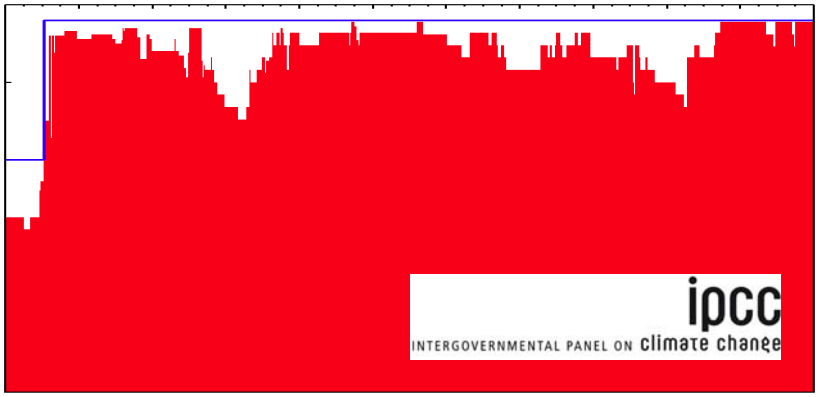

Utilization (red) of Discover SCU7 nodes designated (blue line) to enable GISS simulations supporting AR5, March 21-31, 2011.

GISS, GMAO IPCC AR5 support:

- SCU7's additional resources enabled GISS researchers to approach completion on one set of ModelE2-R historical experiments for IPCC AR5.
- GMAO researchers continued to refine their specialized 10- and 30-year cases that will use coupled atmosphere-ocean-land initialization.

From April to June 2011, the NCCS will advance SCU7's use for SMD's large-scale science investigators by extending access to researchers who intend to scale codes to 1000 or more cores, with no charge against their allocations.

3.5-km GEOS-5

2009 Atlantic Hurricane Bill

GEOS-5 fvCubed Sphere:

- Completed ~9 months of planned 24-month 10-km GEOS-5 Nature Run by end of March.
- Tuning high-resolution physics to better resolve tropical deep convection and hurricane structure.
- Using Goddard Chemistry Aerosol Radiation and Transport (GOCART) "interactive" chemistry at 3.5- to 10-km resolutions globally for the first time.
- 10-km global 5-day forecasts complete within a 3-hour operational window,3.5-km global 3-day forecasts complete within 24 hours.
- Scales to 13,824 cores of SCU7's 14,400 total.

NCCS Data Sharing / Data Portal

NCCS Data Sharing Services: Data Portal

- General Characteristics
  - Web-based environments to support collaboration, public access, and visualization.
  - Supports active NCCS User projects, data is created by NCCS users.
  - Allows viewing preliminary results without transferring data to local workstations.
  - Allows sharing results with collaborators without requiring NCCS users accounts.
- Current Services
  - Simple anonymous ftp and http download
  - Cataloging, subsetting, *limited* data viewing/display
    - THREDDS Data Server (TDS)
    - GrADS Data Server (GDS)
    - Live Access Server (LAS)
    - Web Mapping Services (WMS) viewers
    - Specialized "wxmaps" viewers
- Typical New Service Development Approach
  - Capabilities developed for specific projects in "user space."
  - Services promoted for production use after successful web audit and readiness reviews.
  - Offerings generalized for other projects as warranted.

Earth System Grid

ipcc INTERGOVERNMENTAL PANEL ON climate change

- *New:* Earth System Grid (ESG) Data Node to serve NASA's climate simulation contributions to the 5th Reassessment for the Intergovernmental Panel on Climate Change (IPCC AR5).

NCCS Data Portal Architecture

- Server and disk hardware
  - HP C7000 BladeSystem: 16 nodes Xeon Harpertown (128 cores)
  - 2 additional dedicated database servers
  - ~200 TB (usable) RAID disk (SGI/LSI IS220 and Promise RAID)
- GPFS filesystems shared among Data Portal nodes.
- Direct connection to select NCCS HPC directories for prompt publication; minimizes data movement and multiple copies of data.
- Sufficient compute capability for initial data analysis.

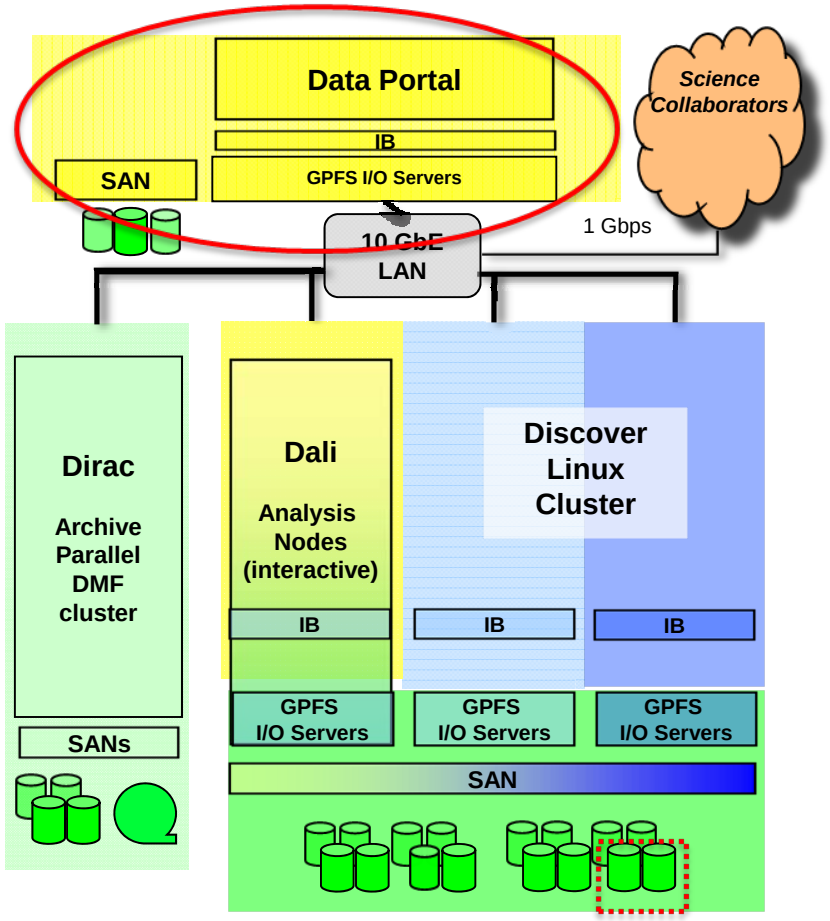

Future NASA Climate Simulation Science

Tomorrow's Climate Science

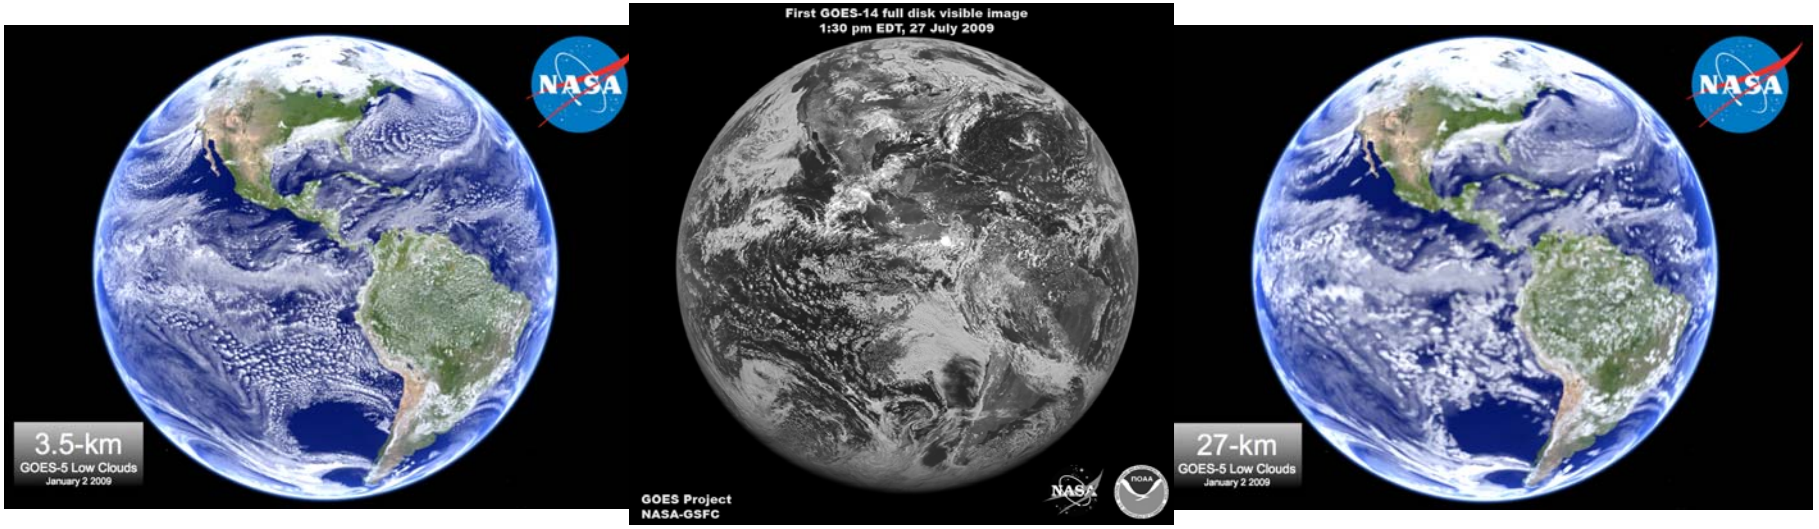

- **New missions** – increased sensing of the earth's climate system as recommended by decadal studies; more data and more types of data
- **Advanced assimilation** – use more data to produce better model initialization
- **Higher resolution** – better representation of atmospheric processes to improve prediction
- **Greater complexity** - understanding and predicting/projecting future climate
- **Coupled ocean-atmosphere-land models** - including full carbon cycle
- **Increased collaboration** – of models, model output, simulation & observational data sets

Climate Simulation Drivers

- Current Snapshot:
  - **One-day 3-km** global climate simulation takes 1 wall-clock day
  - Computational requirement: **4,000 cores**
  - **10 TB** data generated

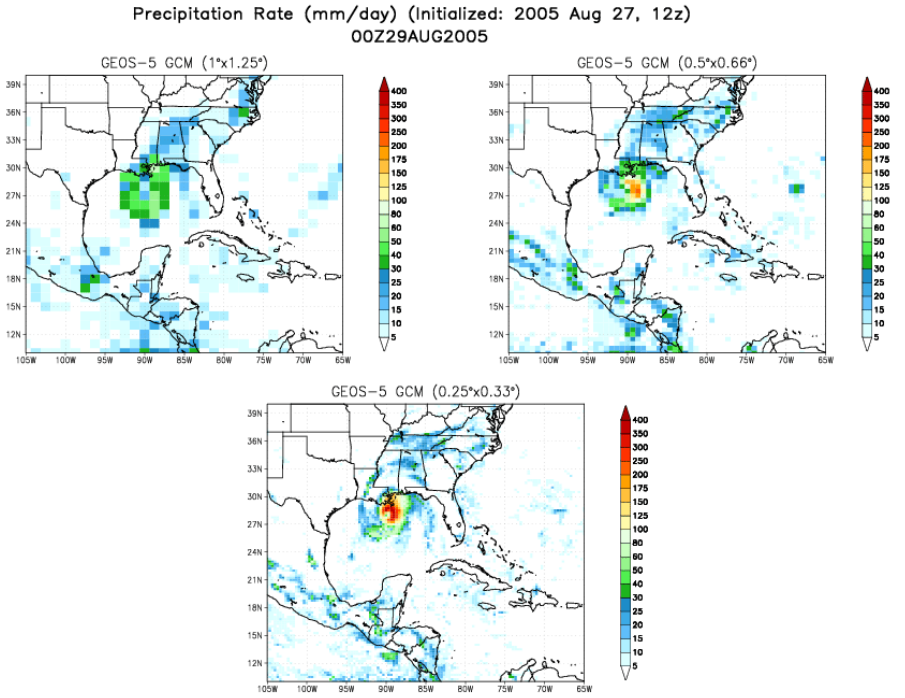

- Goal:
  - **One-year 1.5-km** global climate simulation in 1 wall-clock day
  - Computational requirement: **~6 million cores**
  - **15 PB** data generated

Thank You

Supporting Slides

Acronyms

- AR5 – Fifth Climate Reassessment
- CLIVAR – World Climate Research Programme project that addresses Climate Variability and Predictability
- CMIP5 – Coupled Model Intercomparison Project, phase 5
- CUDA – Compute Unified Device Architecture, a parallel programming framework by NVIDIA
- DMF – Data Migration Facility
- GEOS-5 – Goddard Earth Observing System Model, Version 5
- GISS – Goddard Institute for Space Studies
- GMAO – Global Modeling and Assimilation Office
- GOCART – Goddard Chemistry Aerosol Radiation and Transport
- GPGPU – General Purpose Graphical Processing Unit
- HPC – High Performance Computing
- IPCC – Intergovernmental Panel on Climate Change
- JCSDA – Joint Center for Satellite Data Assimilation
- LTO – Linear Tape Open format
- PCMDI – Program for Climate Model Diagnosis and Intercomparison
- PGI – The Portland Group
- SCU – Scalable Computational Unit
- SIVO – Software Integration and Visualization Office
- SR&T – Supporting Research & Technology
- TB – Trillion Bytes
- TFLOPS – Trillion Floating Point Operations per Second
- WCRP – World Climate Research Programme
- WGCM – Working Group on Coupled Modeling

HPC for Earth Science – the Science Drivers

- **Atmospheric Data Assimilation Systems**
  - **specialized products for NASA instrument teams •**
  - **satellite data impact on weather, air quality, and climate prediction •**
  - **observing system science, including future mission planning •**
  - **climate data record of essential climate variables •**
  - **emerging systems include aerosols, carbon species and reactive gases.**
- **Ocean, ice, and land data assimilation systems**
  - **role of ocean, cryosphere, and land processes in climate •**
  - **initialize the slow components of the climate memory •**
  - **transport of carbon species, nutrients and biota •**
  - **climate data record of essential climate variables.**
- **Coupled ocean-atmosphere-land-sea-ice models•**
  - **climate simulation and prediction at subseasonal-to decadal time scales, •**
  - **climate-weather interactions.**
- **Coupled Chemistry/Biogeochemistry Climate Models•**
  - **simulation and prediction of ozone hole recovery •**
  - **chemistry-climate feedbacks • carbon cycle feedbacks • **
  - **provide information on climate forcings and climate change projections.**
- **Mesoscale regional models•**
  - **high-resolution weather prediction and pollution transport, understand hurricane formation**
- **Prototypes of next-generation systems needed for resolving processes on fine scales.**

HPC for Earth Science – the Resource Drivers

- **Resolution•**
  - **better representation of physical processes •**
  - **better comparison with and use of satellite data •**
  - **regional impacts of climate variability and change •**
  - **better input to future mission design •**
  - **better forecasts**
- **Complexity - *modeling the Earth as a System* •**
  - **climate feedbacks•**
  - **applications to decision support - health, energy •**
  - **air quality, water availability and quality •**
  - **better forecasts**
- **Data Volume•**
  - **sharing PB of data between climate centers [AR5: est. up to 15 PB archive, > 1TB/day data transfer]•**
  - **supporting scientific community - providing access to existing simulations •**
  - **visualization and analysis of model output[2013: est. 1PB/day short-lived data; 37PB/year archive data]**

NCCS Data Management System

- The NCCS Data Management System is built upon NSF- and NARA-funded open source iRODS (Integrated Rule-Oriented Data System).
- iRODS microservices and Rule Engine provide the means to implement policies and workflows that are based on (meta)data attributes, offering great flexibility.
- Initial NCCS implementation uses iRODS to improve access to shared climate data by federating simulation and observational data via a data grid, effectively intermediating between the operational data and scientific communities.
- Operational challenges include
  - Defining metadata and the required mappings.
  - Automating the capture and publishing of metadata.
  - Creating a catalog of NCCS policies to be mapped into iRODS rules.
  - Creating an architecture for workflows to be mapped into iRODS microservices.

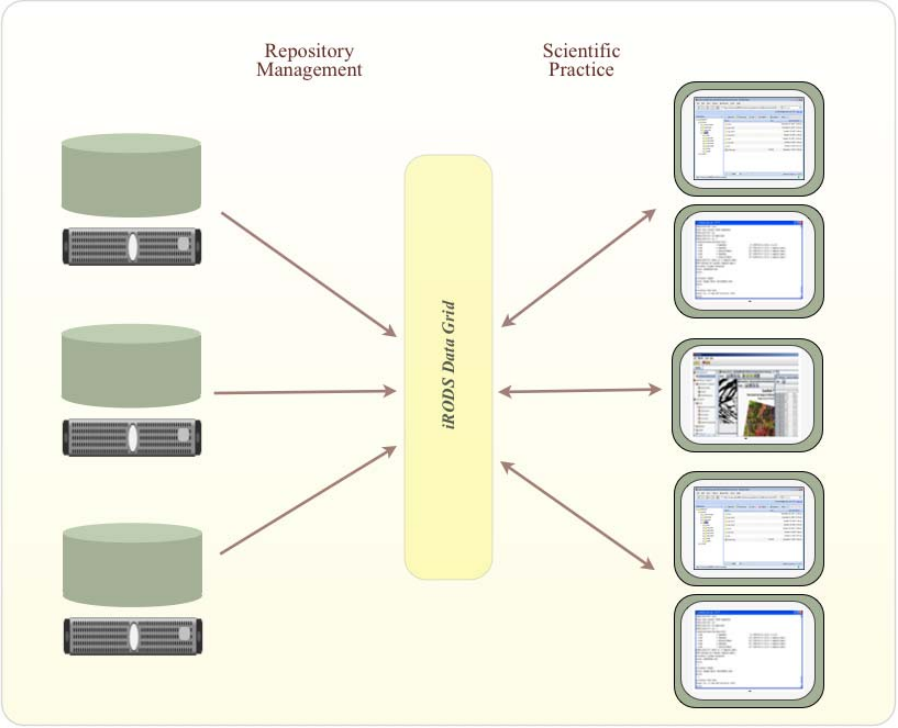

Goal: Improving access to shared observational and simulation data through the creation of a data grid