

# NCCS User Forum

June 15, 2010



# Agenda



**Welcome & Introduction**  
**Phil Webster, CISTO Chief**

Analysis Updates & 3D Demo  
Tom Maxwell, Analysis Lead

Current System Status  
Fred Reitz, HPC Operations

Questions and Comments  
Lynn Parnell/Phil Webster

NCCS Compute Capabilities  
Dan Duffy, Lead Architect

Visualization Wall Demos &  
NCCS Facility Tours

User Services Updates  
Tyler Simon, User Services

Informal Breakout Sessions



# NASA Center for Climate Simulation



- New name for a Goddard Space Flight Center organization that has provided supercomputing resources to NASA scientists and engineers for more than 25 years.
- With climate and weather modeling representing the bulk of NCCS computing, the new name reflects emphasis on supporting NASA Earth science.
- Data-centric science: create datasets and make them available for science researchers around the world.
- The NCCS is continuing to expand its services to meet NASA's growing climate data needs in creation, utilization, and exploitation of Earth science model data.



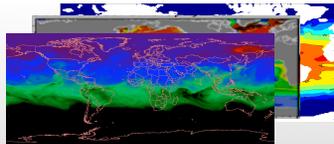


# NCCS Data Centric Climate Simulation Environment



## 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

## Analysis & Visualization\*

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



## DATA Storage & Management

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

## 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

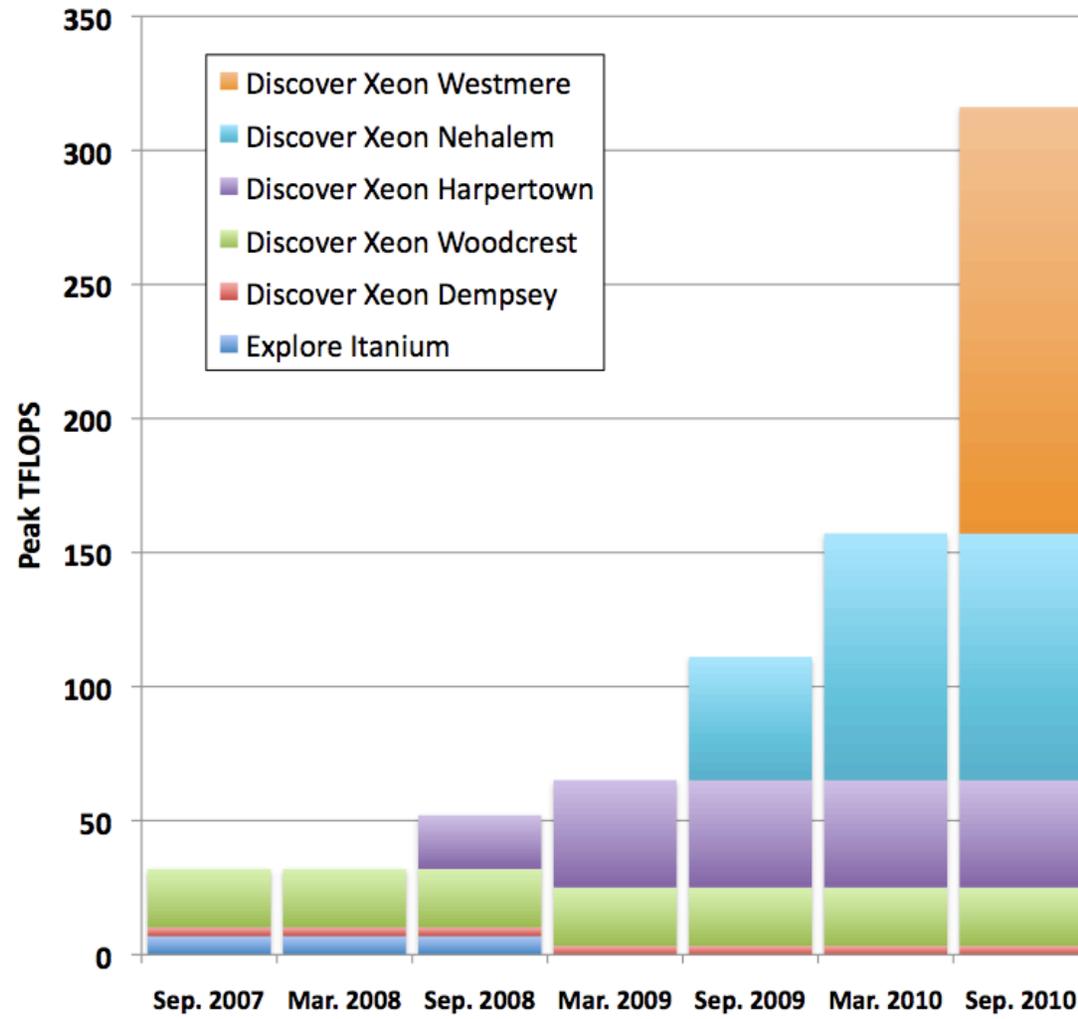
## Data Archival and Stewardship

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





# NCCS Compute Capacity Evolution 2007-2010





# Highlights of the New NCCS



- Discover SCU7 hardware, Fall 2010
  - Approximately doubles Discover capacity, to 300+ peak TFLOPS
- Visualization Wall in the Data Exploration Theater
  - Present highly detailed depictions of scientific results
- Data Management via Integrated Rule-Oriented Data System (iRODS)
  - MODIS and IPCC pilot projects
- Earth System Grid and other Data Portal data sharing and publishing services
- Advances in Data Analysis software in close collaboration with scientists





# NCCS Staff Additions



Maria Ealey, Computing Allocations



# Meeting Logistics Questions



- How many for Vis Wall Demo and NCCS Facility Tours?
- How many for the following Informal Breakouts:
  - Archive data flow
  - Discover workflow evaluation
  - 3D visualization
  - Discover nobackup plans discussion
  - Sysadmin confab



# Agenda



Welcome & Introduction  
Phil Webster, CISTO Chief

Analysis Updates & 3D Demo  
Tom Maxwell, Analysis Lead

**Current System Status**  
**Fred Reitz, HPC Operations**

Questions and Comments  
Lynn Parnell/Phil Webster

NCCS Compute Capabilities  
Dan Duffy, Lead Architect

Visualization Wall Demos &  
NCCS Facility Tours

User Services Updates  
Tyler Simon, User Services

Informal Breakout Sessions



# Key Accomplishments



## General

- GLOPAC mission support
- Data Exploration Theater

## Discover/Analysis Environment

- Returned SCU5 to normal queue structure
- Additional storage

## Dataportal

- Upgraded GPFS
- Upgraded OS to SLES 11
- Established two new FTP sites – g5\_cld, oceanwinds

## DMF

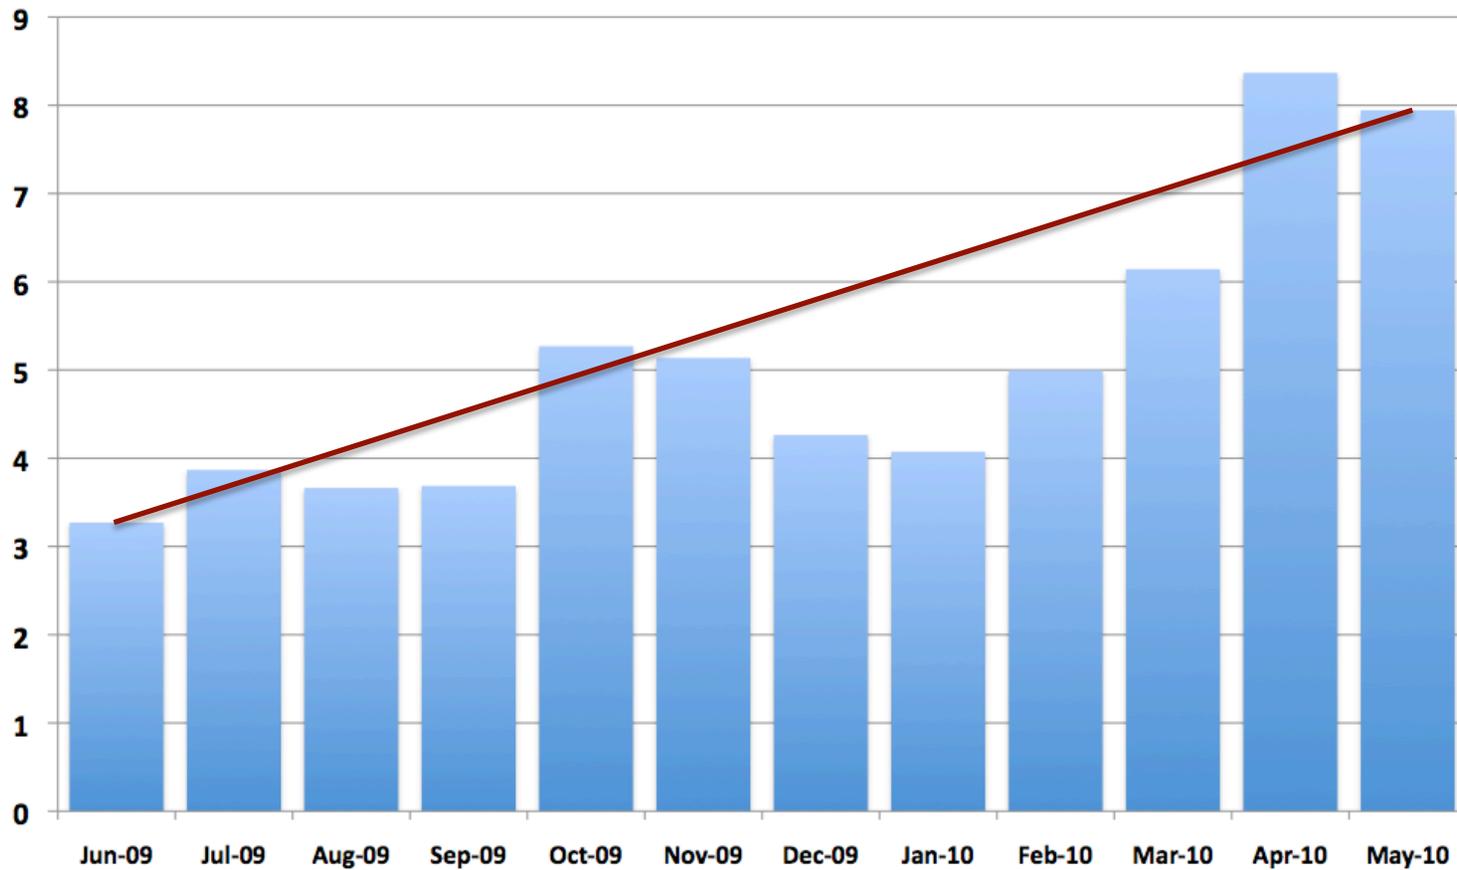
- Upgraded OS to SLES 10 SP3
- Removed 64 processors from palm/dirac



# Discover Total CPU Consumption Past 12 Months (CPU Hours)

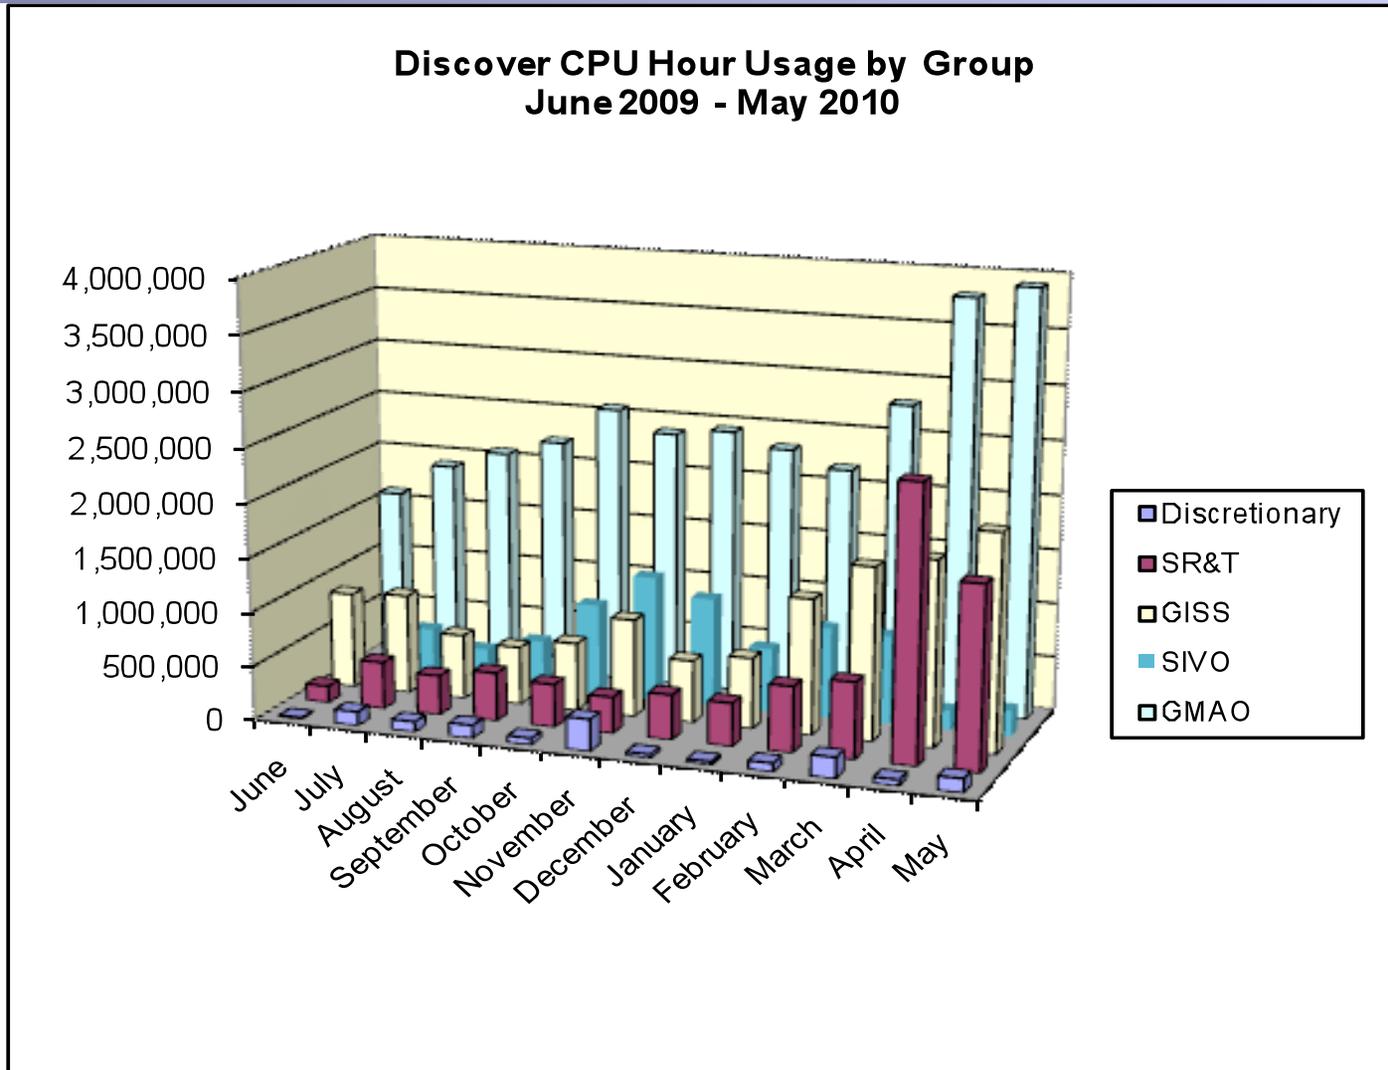


**Discover Monthly Millions of CPU Hours Consumed  
June 2009 - May 2010**





# Discover CPU Consumption by Group Past 12 Months

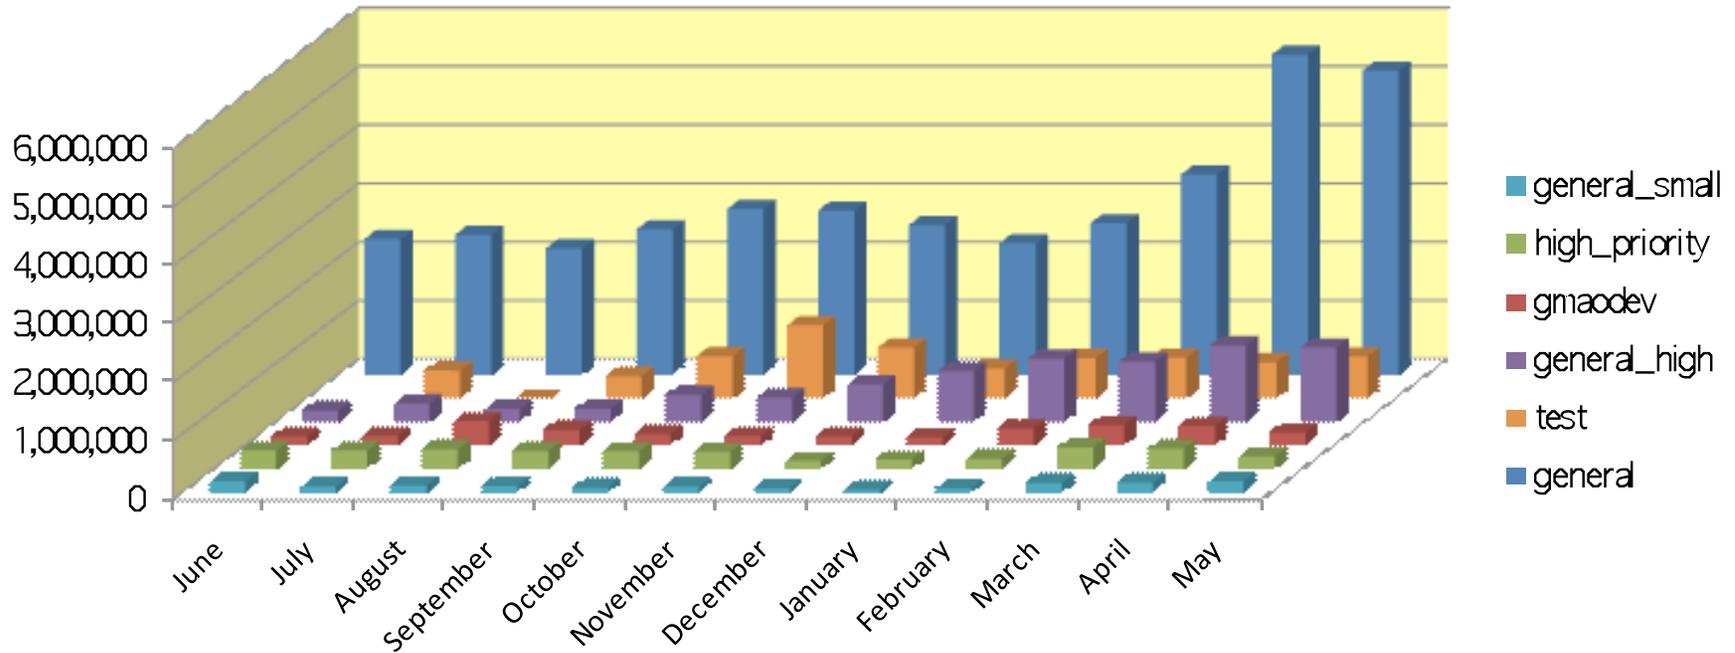


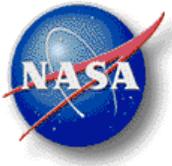


# Discover CPU Consumption by Queue Past 12 Months

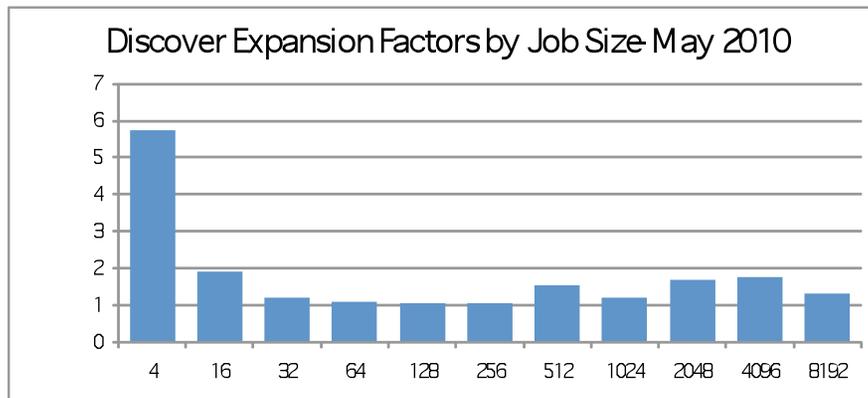
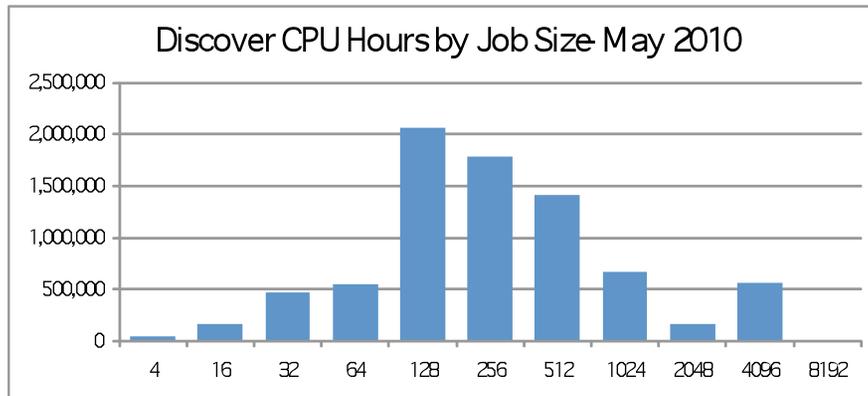
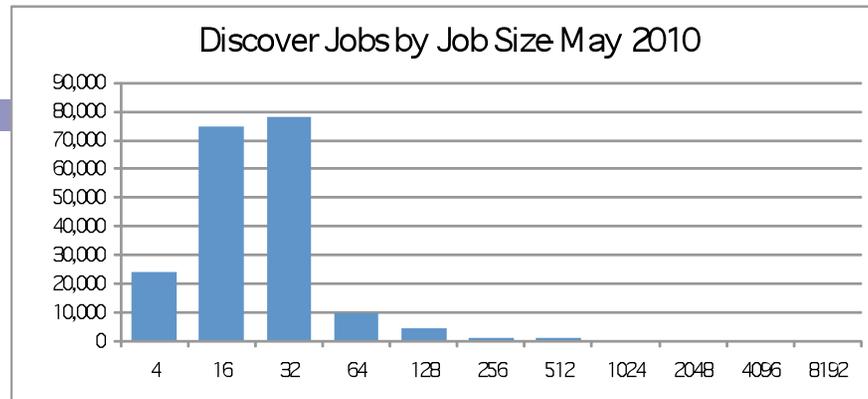


### Discover CPU Hours by Top Queues June 2009- May 2010



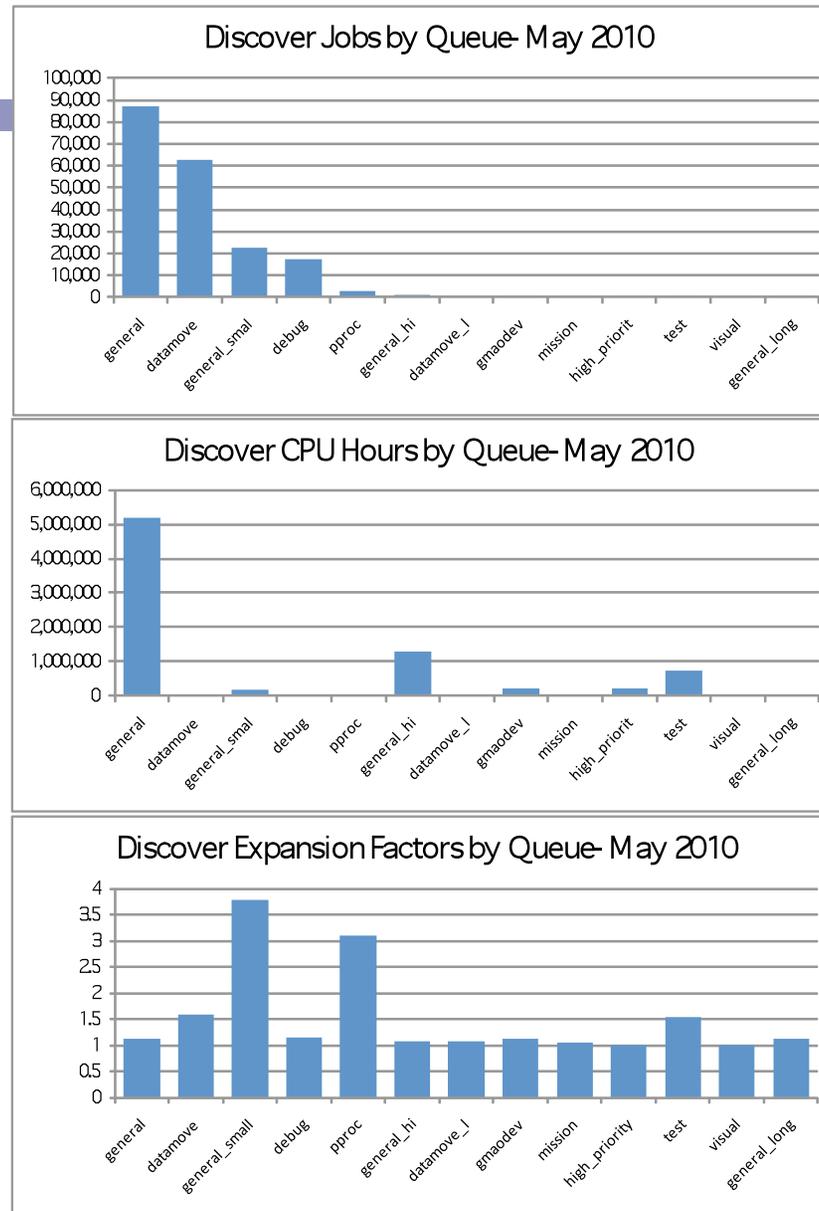


# Discover Job Analysis – May 2010



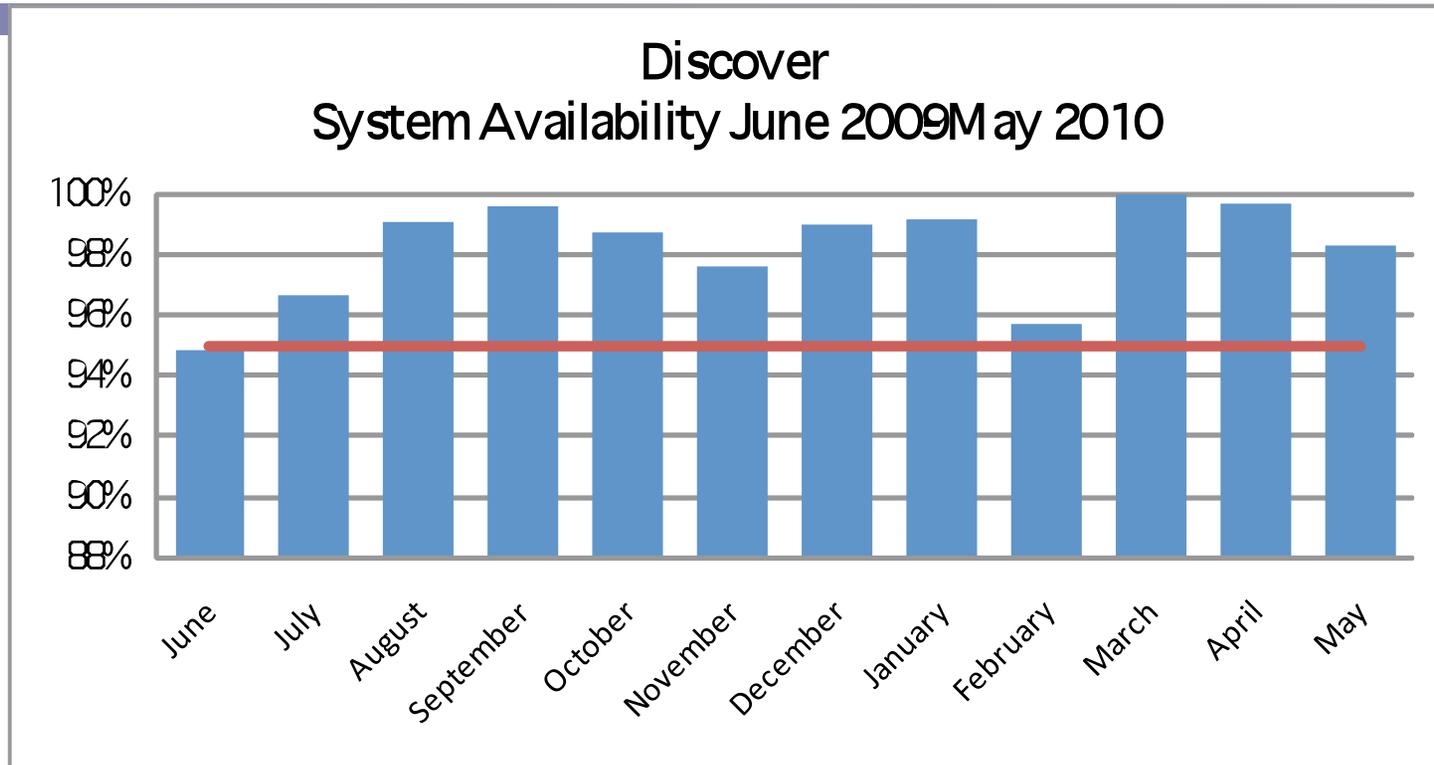


# Discover Job Analysis – May 2010





# Discover Availability – Past 12 Months



## Scheduled Maintenance: March-May

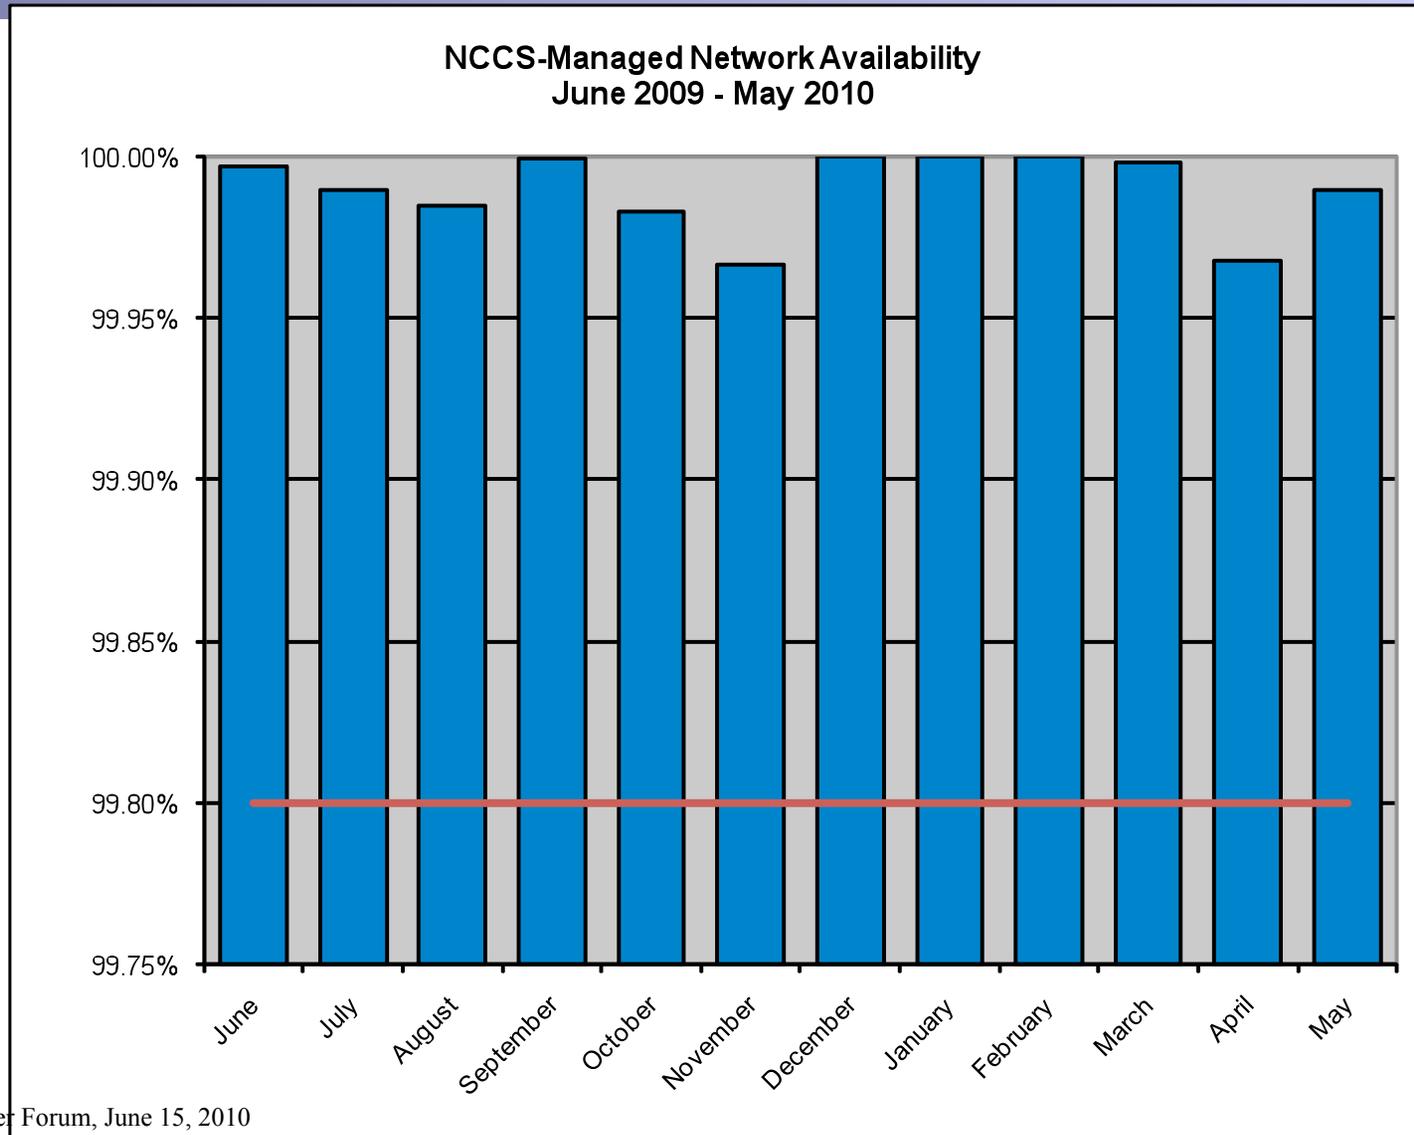
- 22 April – 1 hour  
Network switch maintenance
- 12 May – 12.5 hours  
InfiniBand, disk controller firmware

## Unscheduled Outages: March-May

- 21 April – 1 hour 15 minutes  
GPFS hang



# NCCS Network Availability Past 12 Months





## Future Enhancements



- Data Exploration Theater performance, functionality
- Discover Cluster
  - PBS v10
  - SLES11
  - GPUs
- Data Portal
  - iRODS
  - Additional storage
  - New database servers
- Mass Storage
  - Additional storage
  - Platform upgrade



# Agenda



Welcome & Introduction  
Phil Webster, CISTO Chief

Analysis Updates & 3D Demo  
Tom Maxwell, Analysis Lead

Current System Status  
Fred Reitz, HPC Operations

Questions and Comments  
Lynn Parnell/Phil Webster

**NCCS Compute Capabilities**  
**Dan Duffy, Lead Architect**

Visualization Wall Demos &  
NCCS Facility Tours

User Services Updates  
Tyler Simon, User Services

Informal Breakout Sessions



# SCU7 Discover Cluster Augmentation



- SCU7 will augment Discover compute nodes with Intel Xeon Westmere processors within the existing SCU5 and SCU6 InfiniBand communication fabric.
  - Dual-socket, hex-core processors.
  - 12 MB L3 cache, 50% larger than Nehalem's.
  - 24 GB or RAM.
  - Quad-Data Rate (QDR) Infiniband
- Initial performance of the GEOS-5 Cubed-Sphere benchmark showed a slight improvement on Westmere processors compared to Nehalem.
  - Test compared use of 4 cores per socket on both Nehalem and Westmere.
- Expecting equipment delivery July 2010



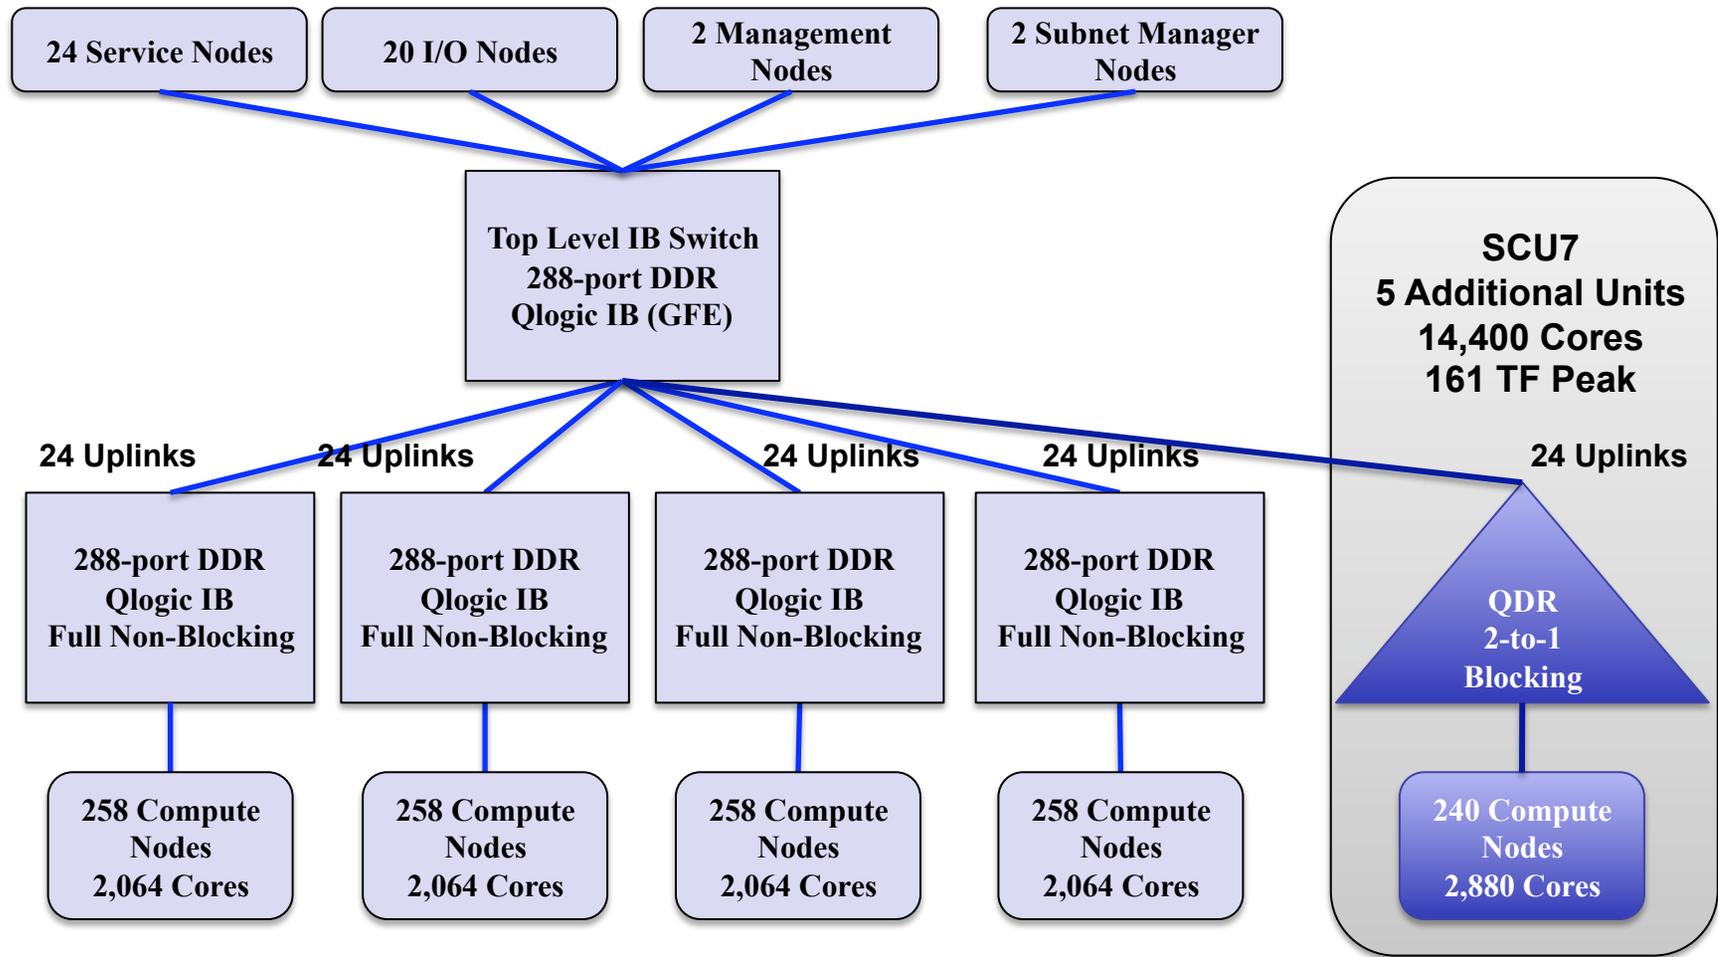
*Intel Xeon 5600 processor series: "Westmere"*

Xeon Processor	Speed	Cores/Socket	L3 Cache Size
Nehalem	2.8 GHz	4	8 MB
Westmere	2.8 GHz	6	12 MB

*Discover's SCU7 will augment the compute capacity available for high resolution GEOS-5 Cubed Sphere runs and for GISS and GMAO simulations in support of the IPCC 5<sup>th</sup> Reassessment.*

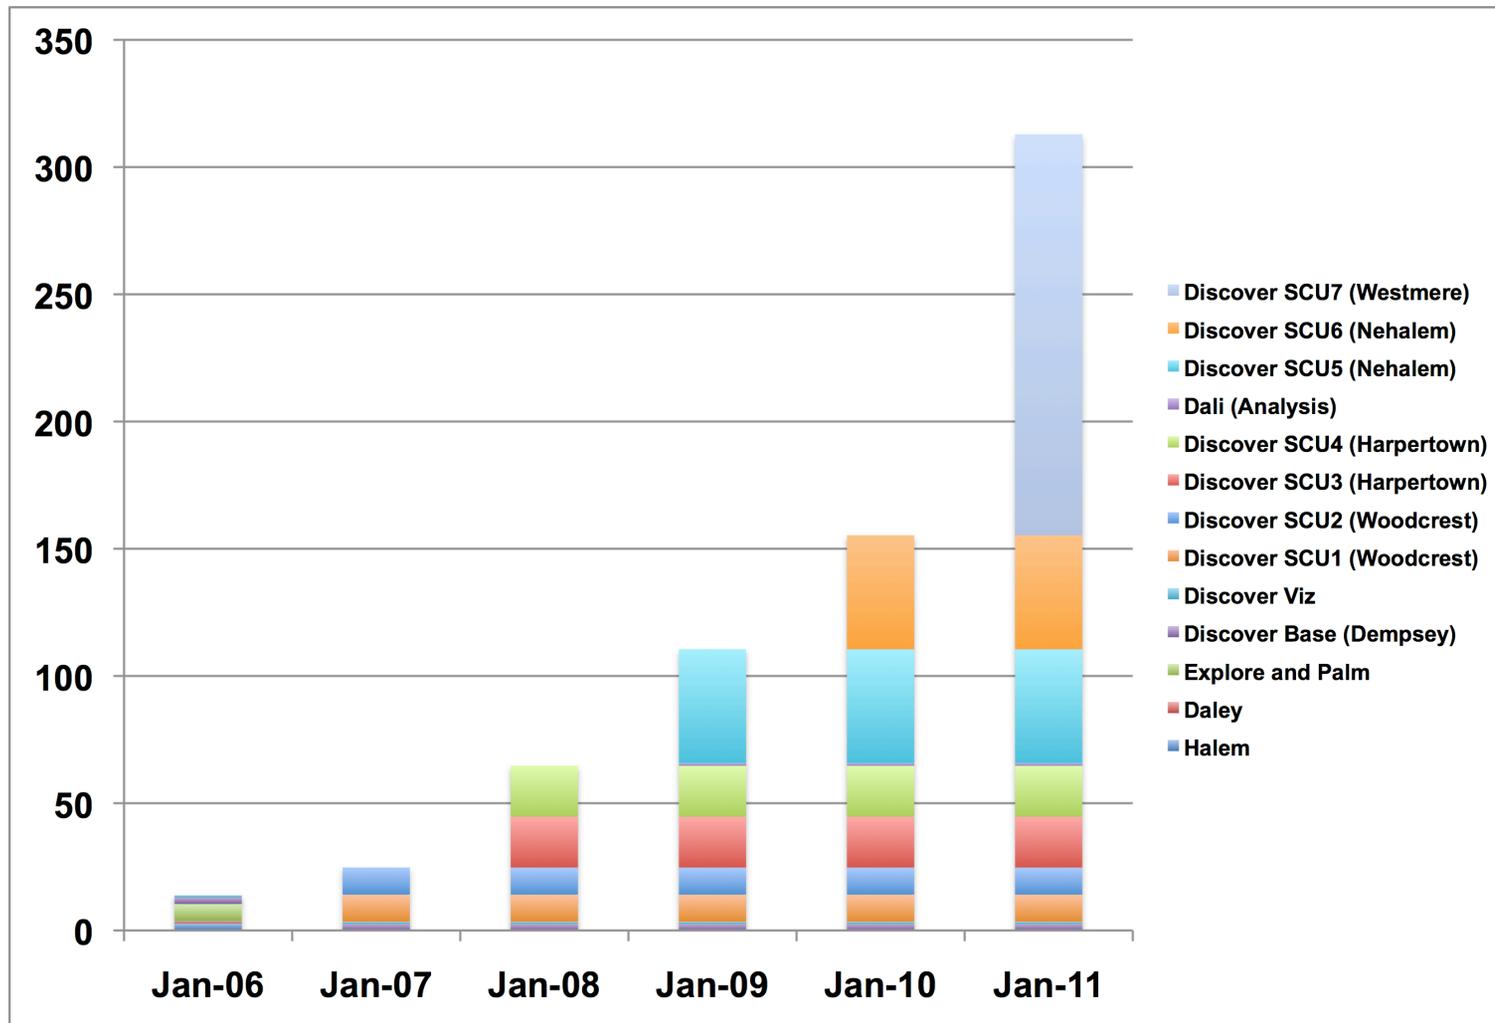


# Extending the IB Fabric to Include SCU7





# NCCS Peak Computing (TF) Over Time





# NCCS Visualization Capability



- The NCCS High Performance Visualization Capability was installed and made operational a very short time following the early March equipment delivery, and includes:
  - High Performance 3x5 Hyperwall (6' x 17')
    - Dell servers with Nvidia FX1700 GPUs.
    - Samsung 46-inch, very small bezel televisions, 720p.
  - Two fixed and three mobile Scientific Analysis Workstations
    - Dell workstations with Nvidia FX4800 GPUs.
    - High-definition 46-inch, 1080p Samsung monitor attached to the Nvidia GPU, will be 3D capable.
- NCCS and SIVO's Scientific Visualization Studio staff are providing tools and assistance to help scientists exploit the Hyperwall.
- Mobile workstations can be deployed for periodic use in scientists' offices or meeting spaces.



*The NCCS High Performance Visualization Capability will expedite highly detailed depictions of scientific results to facilitate the development of new insights and discoveries.*



## Upgrades to the Archive (Dirac)



- Dirac is currently on a system that is 5 years old
- Upgrades
  - New DMF server to replace the Bx2
  - Parallel Data Movers for Tape
  - Two Login Nodes for interactive access to the archive data
  - Two NFS Nodes for improved NFS access to the archive data from Discover
  - All systems are commodity Xeon based
  - Upgrade to the primary disk cache to 600 TB RAW
  - Upgrade to the disk cache for the DMF databases for improved serviceability and access to data
- What does this mean to the users?
  - A much higher performing archive
  - Outages for moving file systems and databases between systems (please be patient)
- Time Line
  - Upgrades to be complete by the end of August



# Agenda



Welcome & Introduction  
Phil Webster, CISTO Chief

Analysis Updates & 3D Demo  
Tom Maxwell, Analysis Lead

Current System Status  
Fred Reitz, HPC Operations

Questions and Comments  
Lynn Parnell/Phil Webster

NCCS Compute Capabilities  
Dan Duffy, Lead Architect

Visualization Wall Demos &  
NCCS Facility Tours

**User Services Updates**  
**Tyler Simon, User Services**

Informal Breakout Sessions



# Ticketing System Upgrade



- NCCS Users will be able to login (via the web) and create new tickets, add a new description or add an attachment to an existing open ticket.
- Users will be able to search the built-in knowledge base for solutions to common problems.
- Follow current system issues in real time by subscribing to global tickets.



# Discover Hardware/Software Migrations

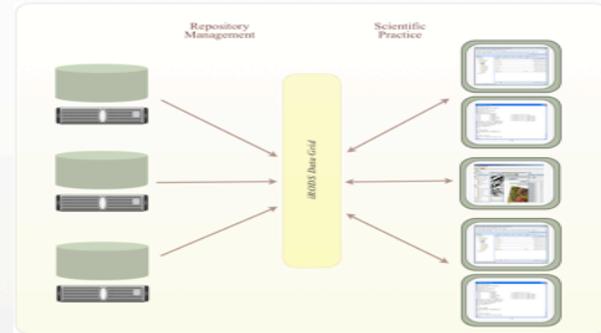
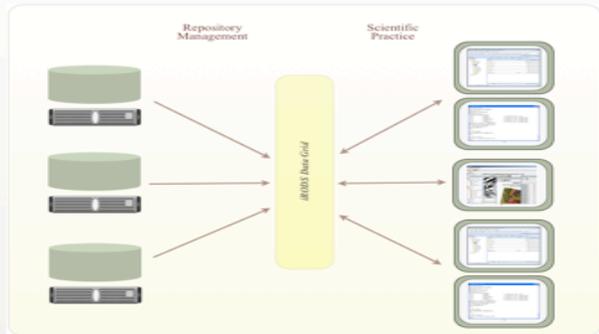


- Deadline to migrate off of ScaliMPI nodes is July 6.
- We will also be moving away from the Visual nodes and Queue at a later date TBD.
- Testing new versions of Intel MPI 4.0 with debugging tools.

[support@nccs.nasa.gov](mailto:support@nccs.nasa.gov)



# iRODS



NCCS  
Integrated iRODS  
Testbed





## iRODS



- Analyze the potential for iRODS to federate climate simulation data and observational data
- Provide one-stop shopping to remote and local data while providing data management services).
- Metadata is catalogued by iRODS and that existing user interfaces let users manage the data and metadata (e.g., add, search, extract, annotate).



# iRODS



Name	Resource	Size	Date Modified
MYD05_OC			May 27, 2010, 1:36 pm
MOD03			May 27, 2010, 1:36 pm
MOD07_L2			May 27, 2010, 1:36 pm
MOD07_OC			May 27, 2010, 1:36 pm
MOD35_L2			May 27, 2010, 1:36 pm
MOD35_OC			May 27, 2010, 1:36 pm
MODCSR_G			May 27, 2010, 1:36 pm
MOD021KM			May 27, 2010, 1:36 pm
MOD02HKM			May 27, 2010, 1:36 pm
MOD02QKM			May 27, 2010, 1:36 pm
MYD05_L2			May 27, 2010, 1:37 pm
MYD04_L2			May 27, 2010, 1:37 pm
MYBA0D_D3			May 27, 2010, 1:37 pm
MYBARS_D3			May 27, 2010, 1:37 pm
MYBCER_D3			May 27, 2010, 1:37 pm
MYBCFR_D3			May 27, 2010, 1:37 pm
MYBCIR_D3			May 27, 2010, 1:38 pm
MYBCOT_D3			May 27, 2010, 1:38 pm
MYBCTR_D3			May 27, 2010, 1:38 pm
MYBCTT_D3			May 27, 2010, 1:38 pm
MYBWIR_D3			May 27, 2010, 1:38 pm
MYBWSW_D3			May 27, 2010, 1:38 pm
MYD08_D3			May 27, 2010, 1:38 pm
MYDATML2			May 27, 2010, 1:38 pm
MODCSR_D			May 27, 2010, 1:38 pm
MODCSR_8			May 27, 2010, 1:38 pm
MYBCFR			May 27, 2010, 1:39 pm
MYBCTP			May 27, 2010, 1:39 pm
MYBCTT			May 27, 2010, 1:39 pm
MYBRGB			May 27, 2010, 1:39 pm



# Agenda



Welcome & Introduction  
Phil Webster, CISTO Chief

**Analysis Updates & 3D Demo**  
**Tom Maxwell, Analysis Lead**

Current System Status  
Fred Reitz, HPC Operations

Questions and Comments  
Lynn Parnell/Phil Webster

NCCS Compute Capabilities  
Dan Duffy, Lead Architect

Visualization Wall Demos &  
NCCS Facility Tours

User Services Updates  
Tyler Simon, User Services

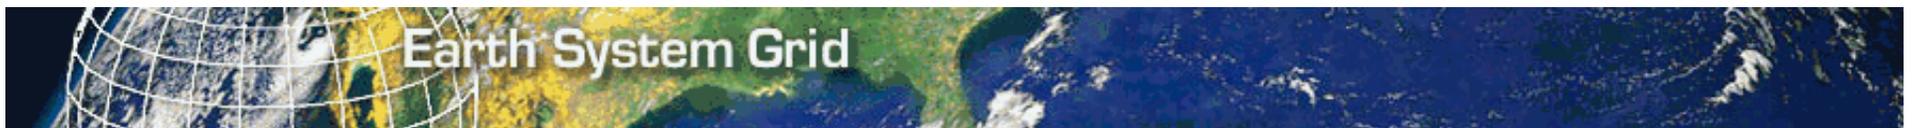
Informal Breakout Sessions



# Earth System Grid Data Node Update



- ESG Data Node software (2.4) received from PCMDI.
  - Installed and tested on the NCCS Dataportal.
  - Successfully published to PCMDI ESG Gateway.
- Developing plans for publishing MERRA data .
  - Publishing virtual files not currently feasible.
  - Will copy and publish monthly means.
- Will publish GISS & GMAO IPCC AR5 contributions.
  - Scheduled date for publication of initial data: August 2010

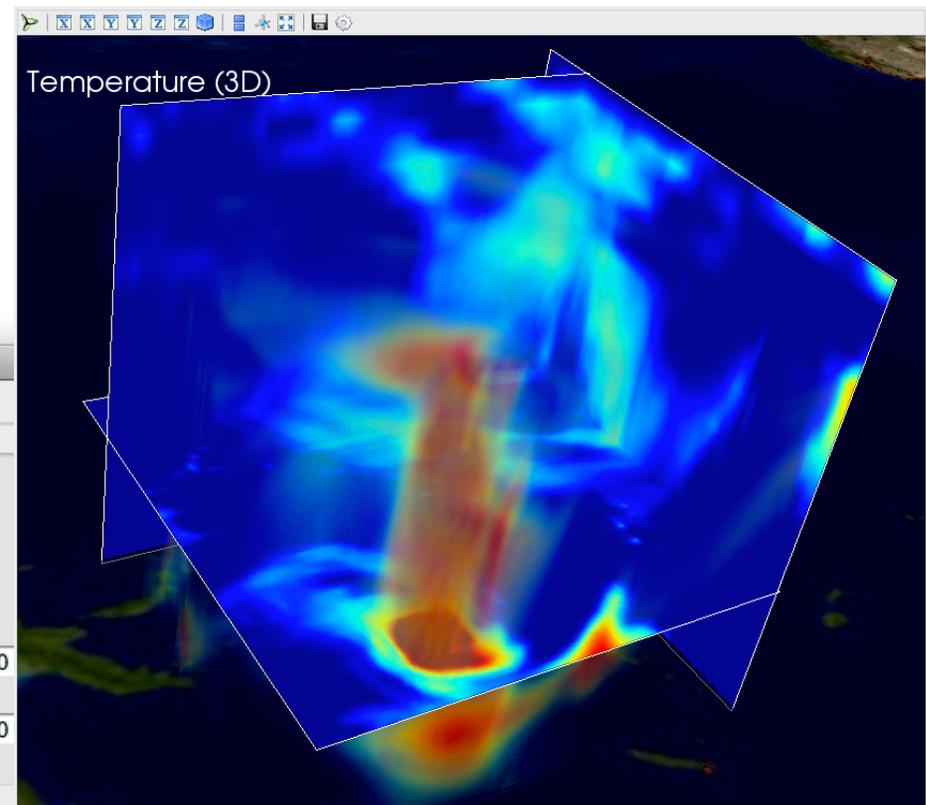
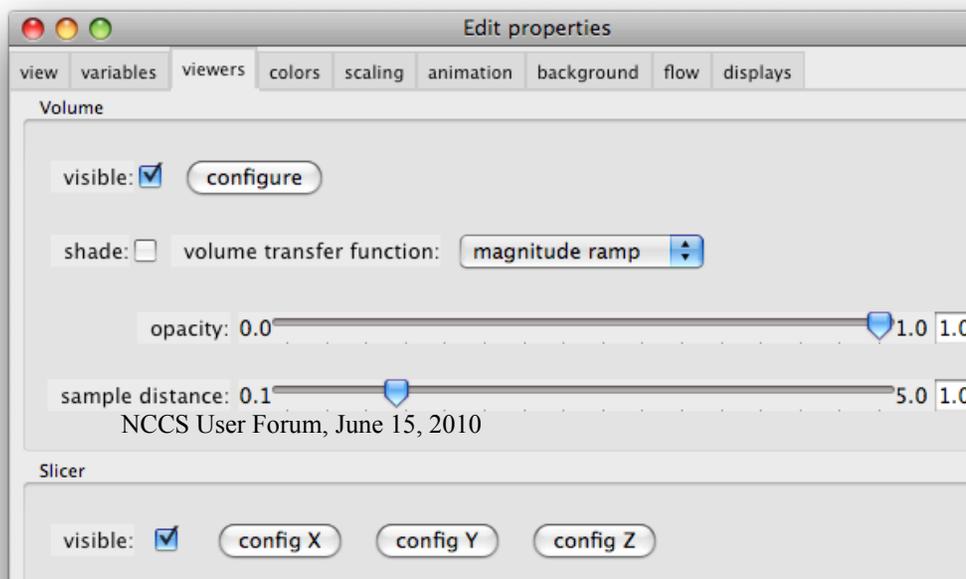




# DV3D



- Interactive 3D visualization of simulation data
  - Tailored for climate scientists
  - Interface to NCCS/ESG Analysis Products Server
- Python app built on MayaVi / VTK
  - Runs inside Grads
  - Access all Grads variables
- Developed at NCCS
  - Hyper wall demo to follow





## Analysis Products Server (Under Development)



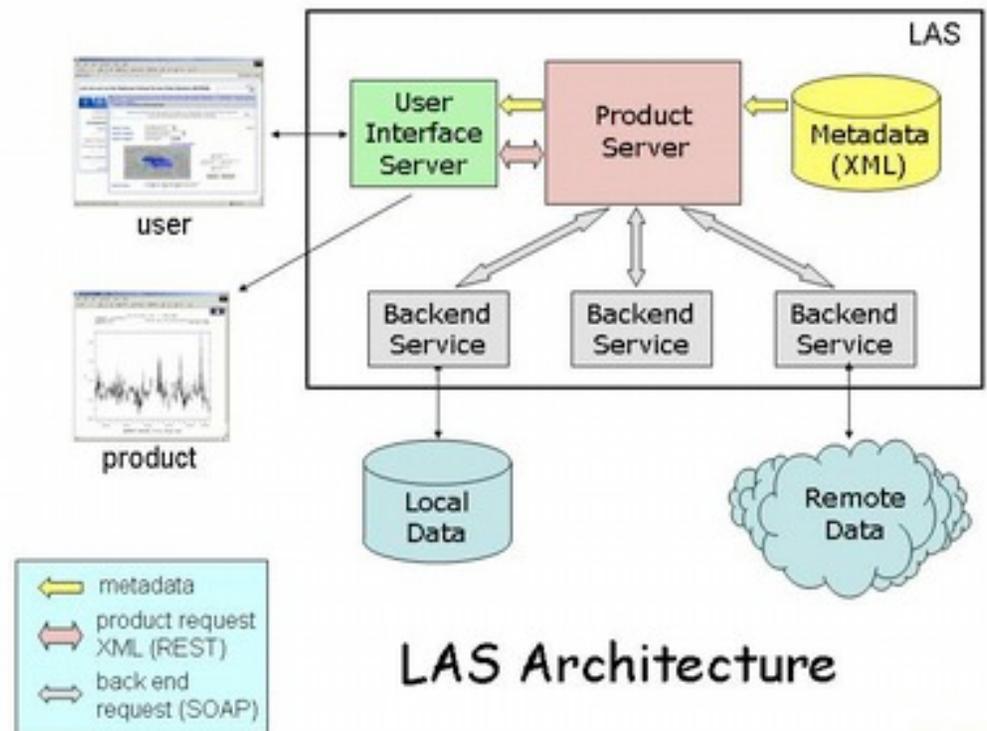
- Server-side Analysis Services executed on Dali
  - Python client for GrADS/DV3D
  - DET: “Gateway to Dali”
    - Phase 1: Access saved simulation data
    - Phase 2: Access running simulation.
  - Service APIs:
    - Live Access Server (LAS) request protocol
    - IRODS
  - Planned Services
    - Subsetting, regridding/rescaling, reformatting, aggregation, etc.
    - Differencing, averaging (e.g. zonal means)
    - Data transforms ( wavelet, FFT, etc. )



# Live Access Server



- ESG Analysis module
  - ESG LAS not yet available
- Three Tiered architecture
  - Interface Server
    - Installed on dp1
  - Product server
    - Installed on dp1
  - Back End Server (BES)
    - Installed on dp7 (future: dali)
- Current BES uses Ferret
  - Potential BE Services:
    - Grads, CDAT, NCO, IRODS



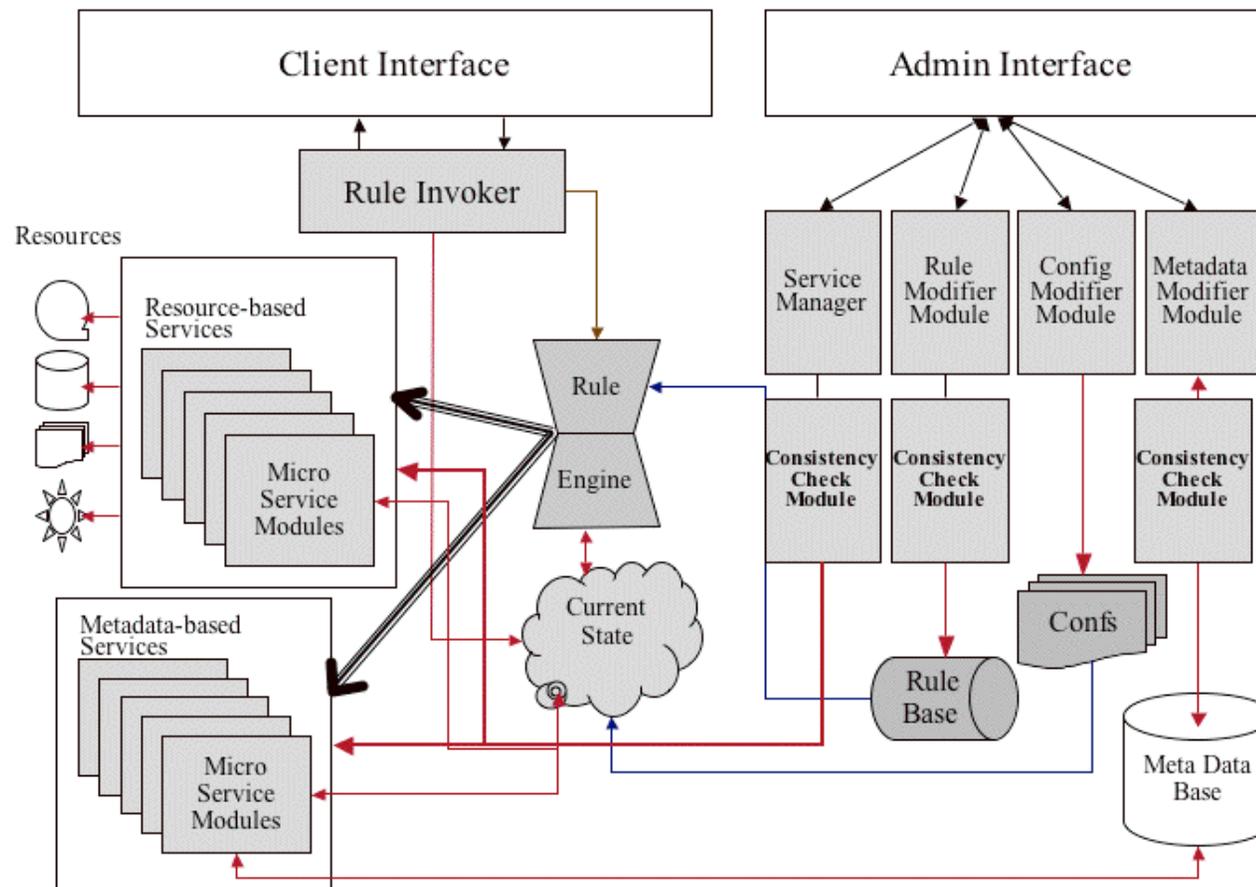
Steve Harkin



# IRODS



- Integrated Rule Oriented Data System
- Rule Engine applies user-defined Rules and Services
- Nodes to be installed on Dali and Data Portal





# Agenda



Welcome & Introduction  
Phil Webster, CISTO Chief

Analysis Updates & 3D Demo  
Tom Maxwell, Analysis Lead

Current System Status  
Fred Reitz, HPC Operations

**Questions and Comments**  
**Lynn Parnell/Phil Webster**

NCCS Compute Capabilities  
Dan Duffy, Lead Architect

Visualization Wall Demos &  
NCCS Facility Tours

User Services Updates  
Tyler Simon, User Services

Informal Breakout Sessions



# Important Contact Information

---



NCCS Support:

[support@nccs.nasa.gov](mailto:support@nccs.nasa.gov)

(301) 286-9120



# Agenda



Welcome & Introduction  
Phil Webster, CISTO Chief

Analysis Updates & 3D Demo  
Tom Maxwell, Analysis Lead

Current System Status  
Fred Reitz, HPC Operations

Questions and Comments  
Lynn Parnell/Phil Webster

NCCS Compute Capabilities  
Dan Duffy, Lead Architect

**Visualization Wall Demos &  
NCCS Facility Tours**

User Services Updates  
Tyler Simon, User Services

Informal Breakout Sessions



# Vis Wall Demo and NCCS Facility Tours



- *No Food or Drink, please*
- Vis Wall in S214/Data Exploration Theater
  - Introduction (Phil Webster)
  - 3D Demonstrations (Tom Maxwell)
- NCCS Facility Tours
  - Begin at the bottom of front stairs
- Informal Breakouts to follow in E210, S216, and staff offices



# Agenda



Welcome & Introduction  
Phil Webster, CISTO Chief

Analysis Updates & 3D Demo  
Tom Maxwell, Analysis Lead

Current System Status  
Fred Reitz, HPC Operations

Questions and Comments  
Lynn Parnell/Phil Webster

NCCS Compute Capabilities  
Dan Duffy, Lead Architect

Visualization Wall Demos &  
NCCS Facility Tours

User Services Updates  
Tyler Simon, User Services

**Informal Breakout Sessions**