



NASA Center for Computational Sciences

NCCS User Forum

13 September 2007



Agenda



NASA Center for Computational Sciences

- Introduction — Phil Webster
- Systems Status—Mike Rouch
- Discover SCU2 – Mike Rouch
- Visualization Services – Carrie Spear
- Data Sharing Services—Ellen Salmon
- User Services — Sadie Duffy
- Questions or Comments



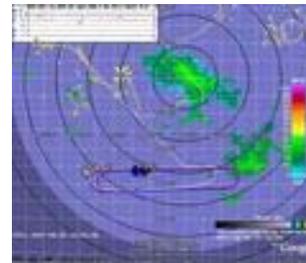
NCCS Supports NASA TC4 Mission (Tropical Composition, Cloud and Climate Coupling)



NASA Center for Computational Sciences

TC4 campaign from July 16 – August 12, 2007

- Study the tropical tropopause transition layer (TTL) to understand chemical, dynamical, and physical processes associated with climate change and atmospheric ozone depletion.
- Complement NASA A-train satellite data with project-specific observational data.
- TC4 deployed 25 DC-8, ER-2 and WB-57 flights; 292 weather balloons; and 93 dropsondes.
- Over 200 scientists, engineers, and mission support personnel were based in Costa Rica and Panama. This large international experiment united researchers from 8 NASA centers, over 14 universities, and more than 20 U.S. and international agencies.



NCCS support to TC4

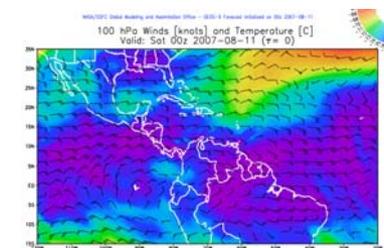
Computation Services - NCCS hosted:

- Real-time GEOS5 analyses and forecasts,
- Meteorological and forecasts,
- Real-time estimates/forecasts of aerosols, CO and CO2 tracers,
- Special high resolution forecasts to aid flight planning.

Data Services – NCCS provided datasets via:

- Web-based Weather Map tools,
- Comprehensive Web Map Services (SVS WMS),
- OPeNDAP server providing subsetting to field scientists, and
- Anonymous FTP access to all the data.

These capabilities were an integral part of TC4 mission operations in Costa Rica, and play an important role for mission analysis and interpretation of results.



9/13/2007

NCCS User Forum



Halem Status



NASA Center for Computational Sciences

- Halem (the man) retired in 2002
 - Emeritus position as Chief Information Research Scientist
 - Halem Emeritus I
- *Halem* (the machine) retired **May 1**
 - 40 million CPU hours for Earth Science Research
 - Four years of service
 - Self maintained for over 1 year
 - Halem Emeritus II
- Replaced by *Discover*
 - Factor of 5 capacity increase
 - All users successfully migrated





Discover SCU2



NASA Center for Computational Sciences

- 23 July 2007 - NCCS took delivery of additional nodes for *Discover* from Linux Networx.
- Increased capacity includes:
 - 256 dual processor, dual core Intel Woodcrest nodes
 - along with additional specialty login, management and data migration nodes and
 - an additional 70TB of user storage.
- System integration followed pre-defined acceptance test plan.
 - All components run as standalone system to address initial hardware failure due to shipping
 - System connected in early August to current test and validation system to configure nodes with production software stack
 - Nodes moved to production environment in mid-August for further testing
 - 12 September 2007 – Commence 30-day acceptance test



NCCS increased overall capacity of commodity linux cluster by ~11TF

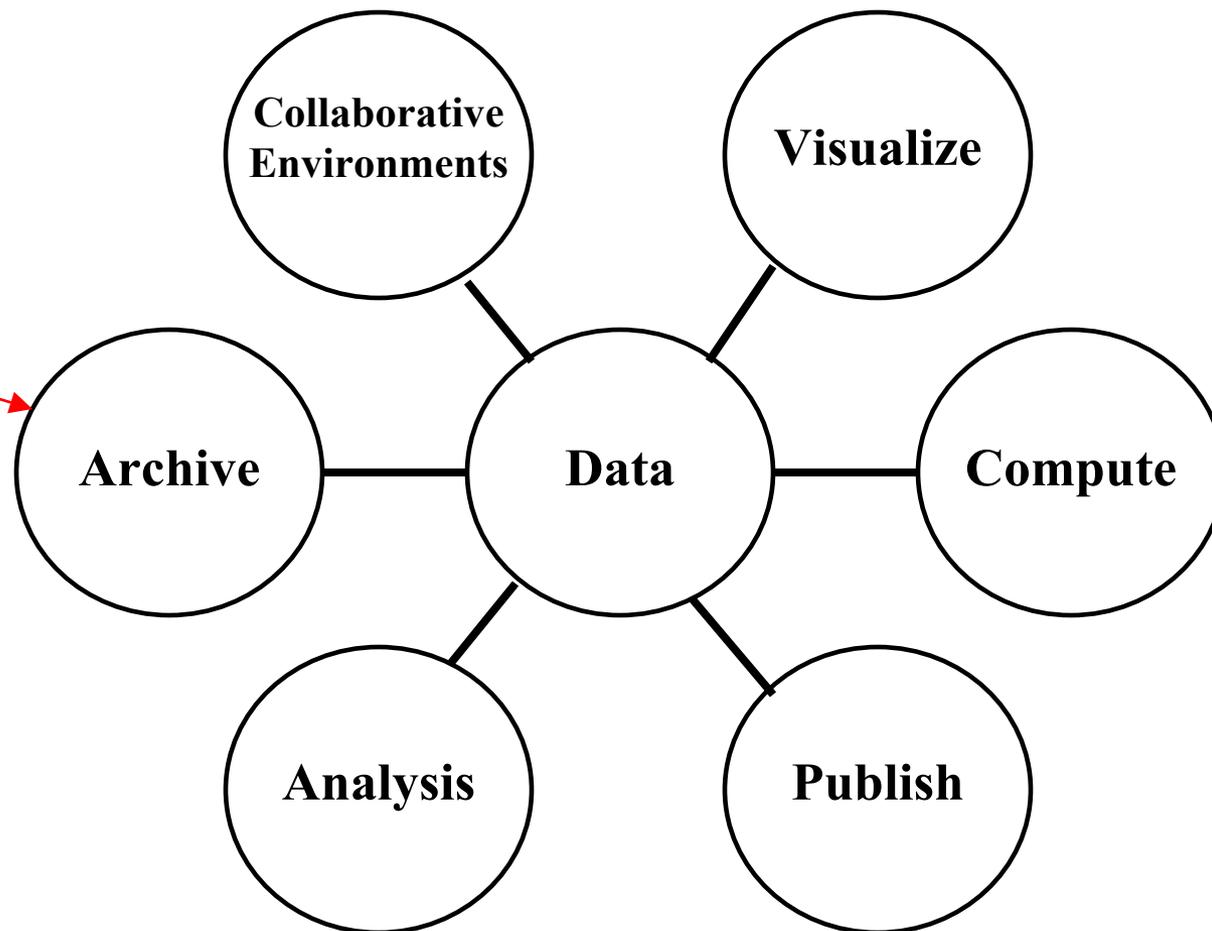
Discover system now ~25TF



Conceptual Architecture



NASA Center for Computational Sciences

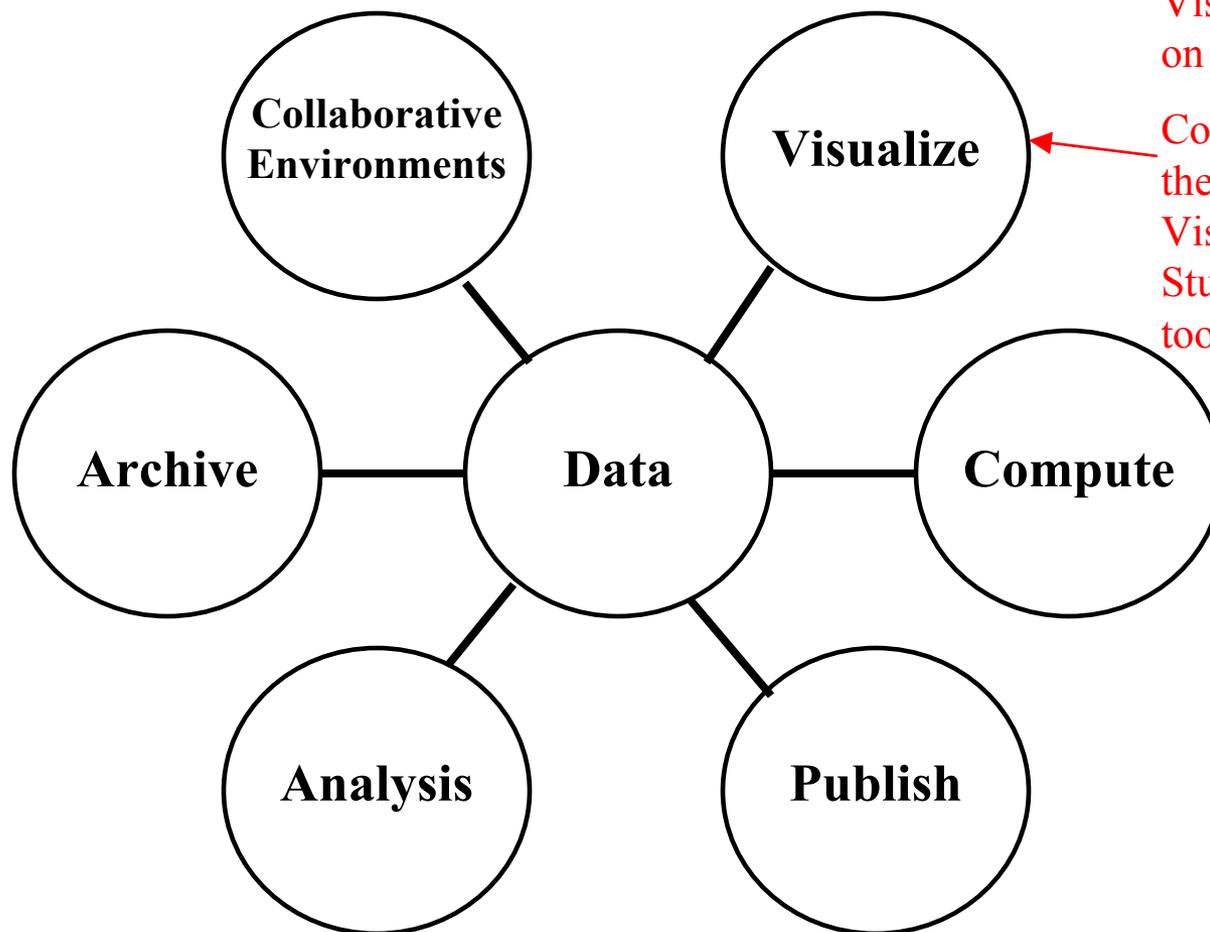




Conceptual Architecture



NASA Center for Computational Sciences



Visualization nodes on *discover*

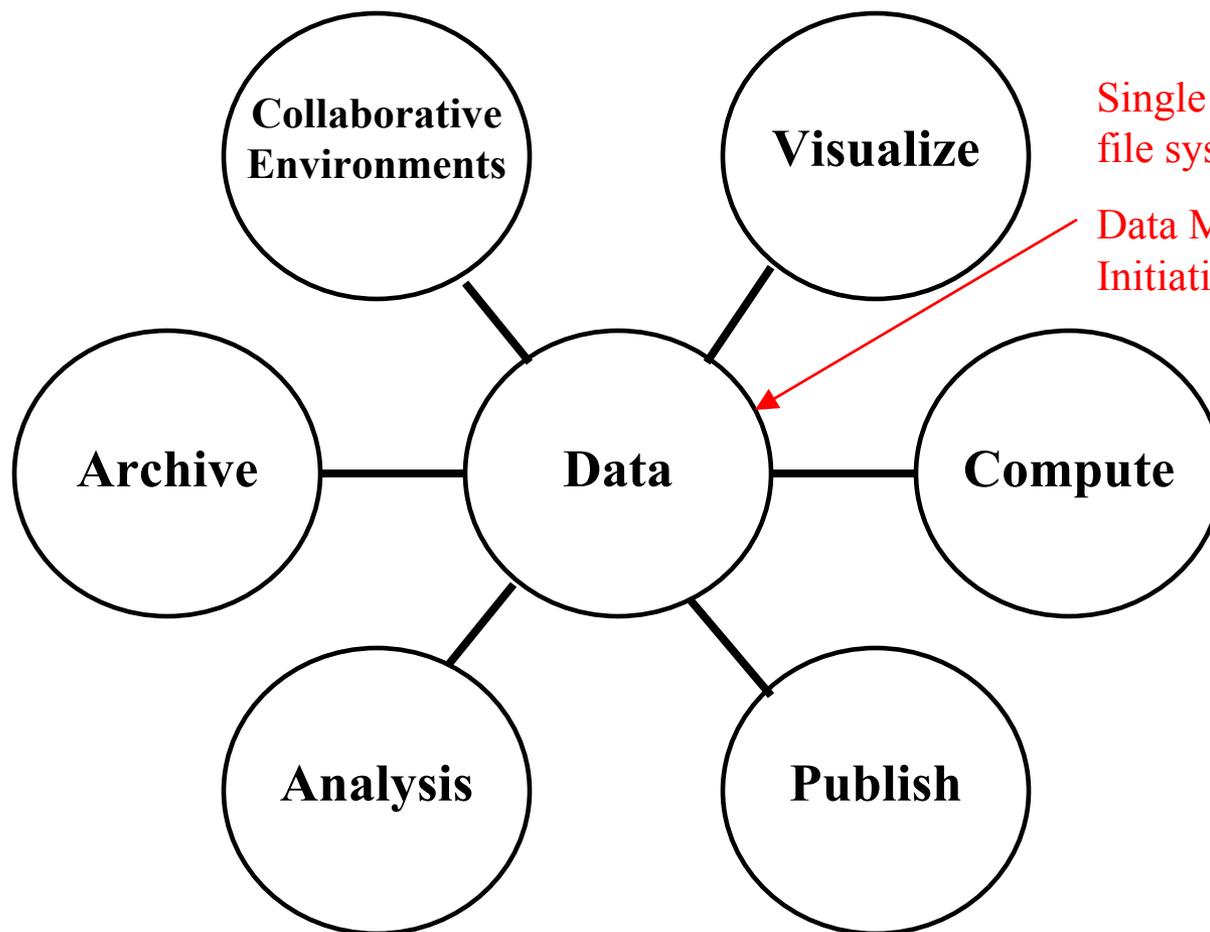
Collaboration with the Scientific Visualization Studio to provide tools



Conceptual Architecture



NASA Center for Computational Sciences



Single NCCS wide file system in FY09

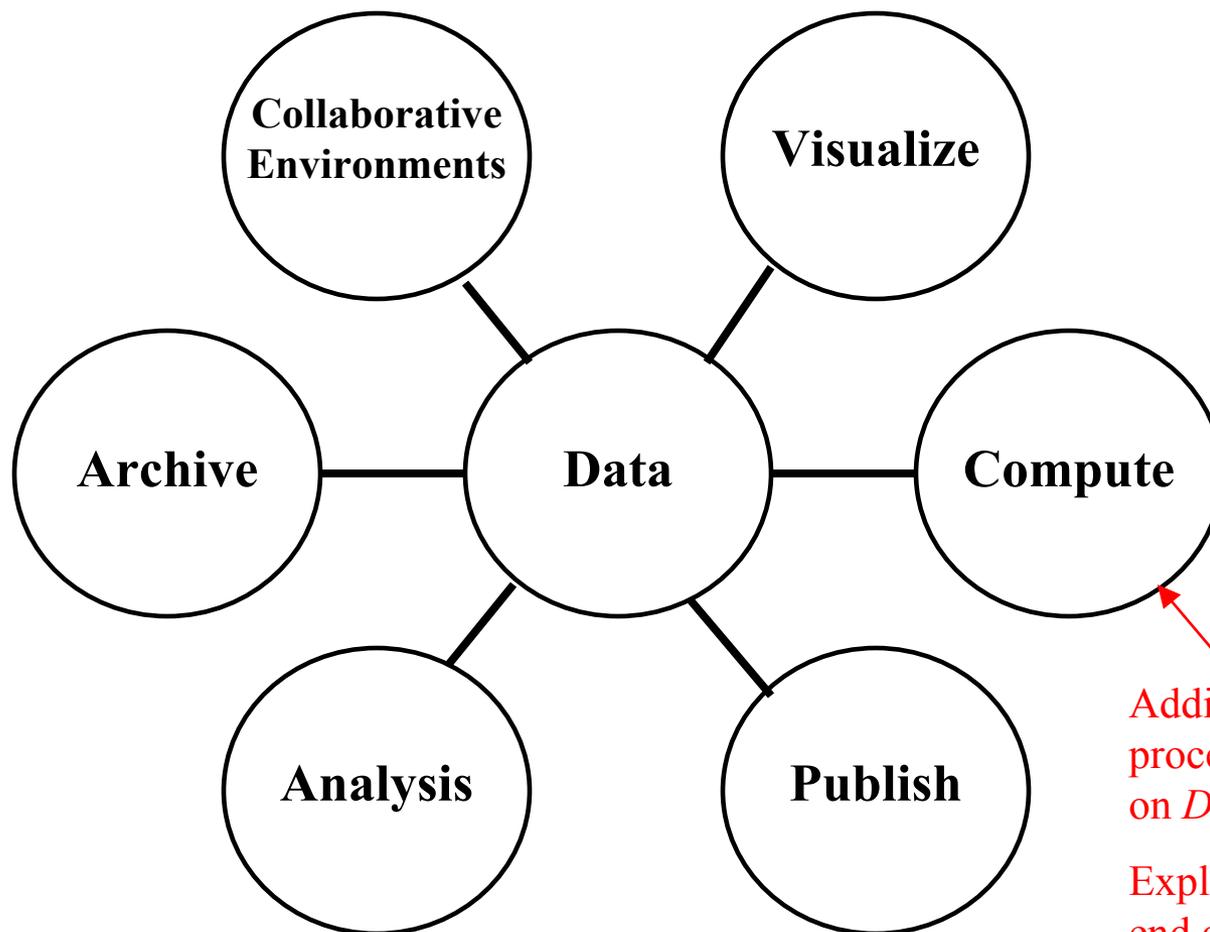
Data Management Initiative in FY08



Conceptual Architecture



NASA Center for Computational Sciences



Additional 1024 processing elements on *Discover*

Explore will retire end of FY08

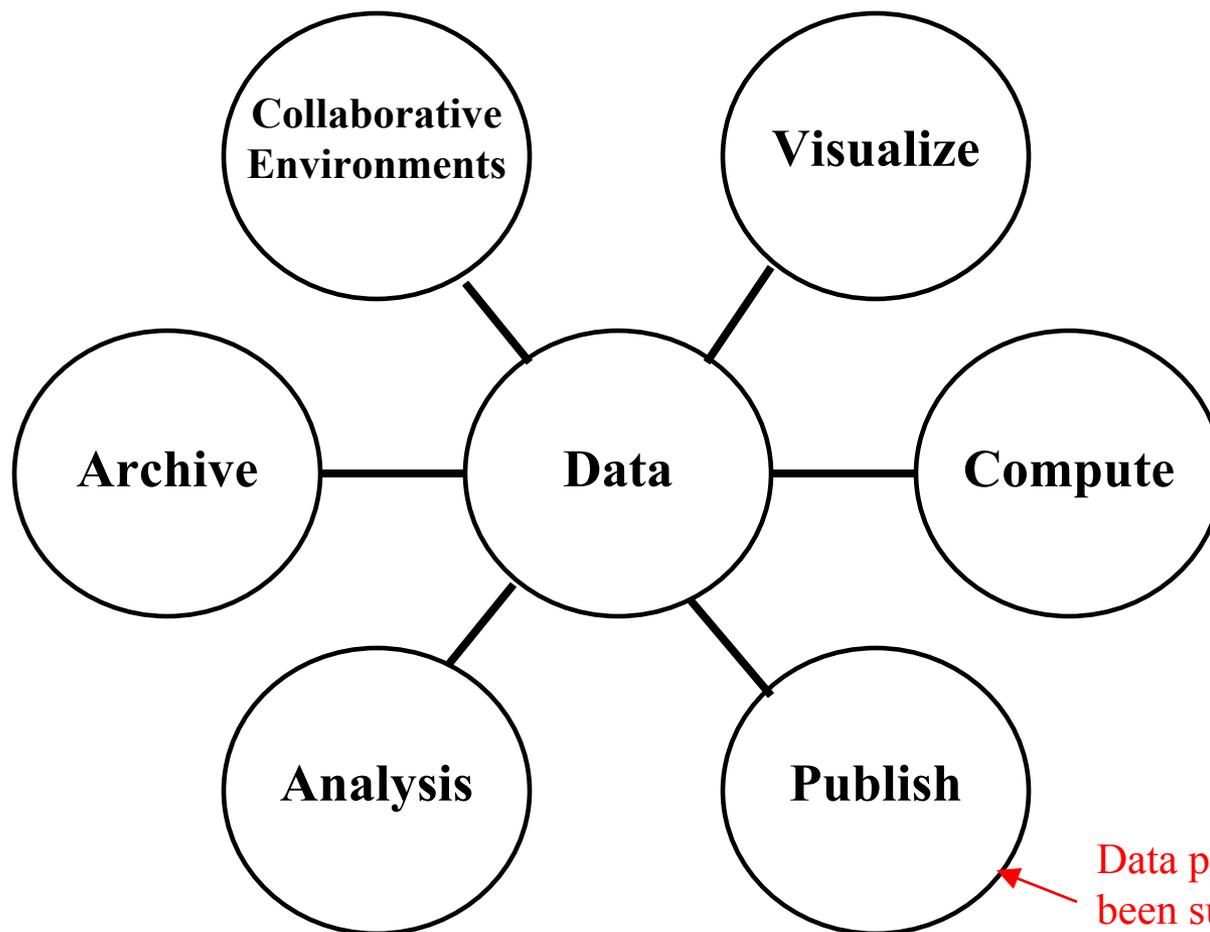
Halem retired



Conceptual Architecture



NASA Center for Computational Sciences



Data portal prototype has been successful

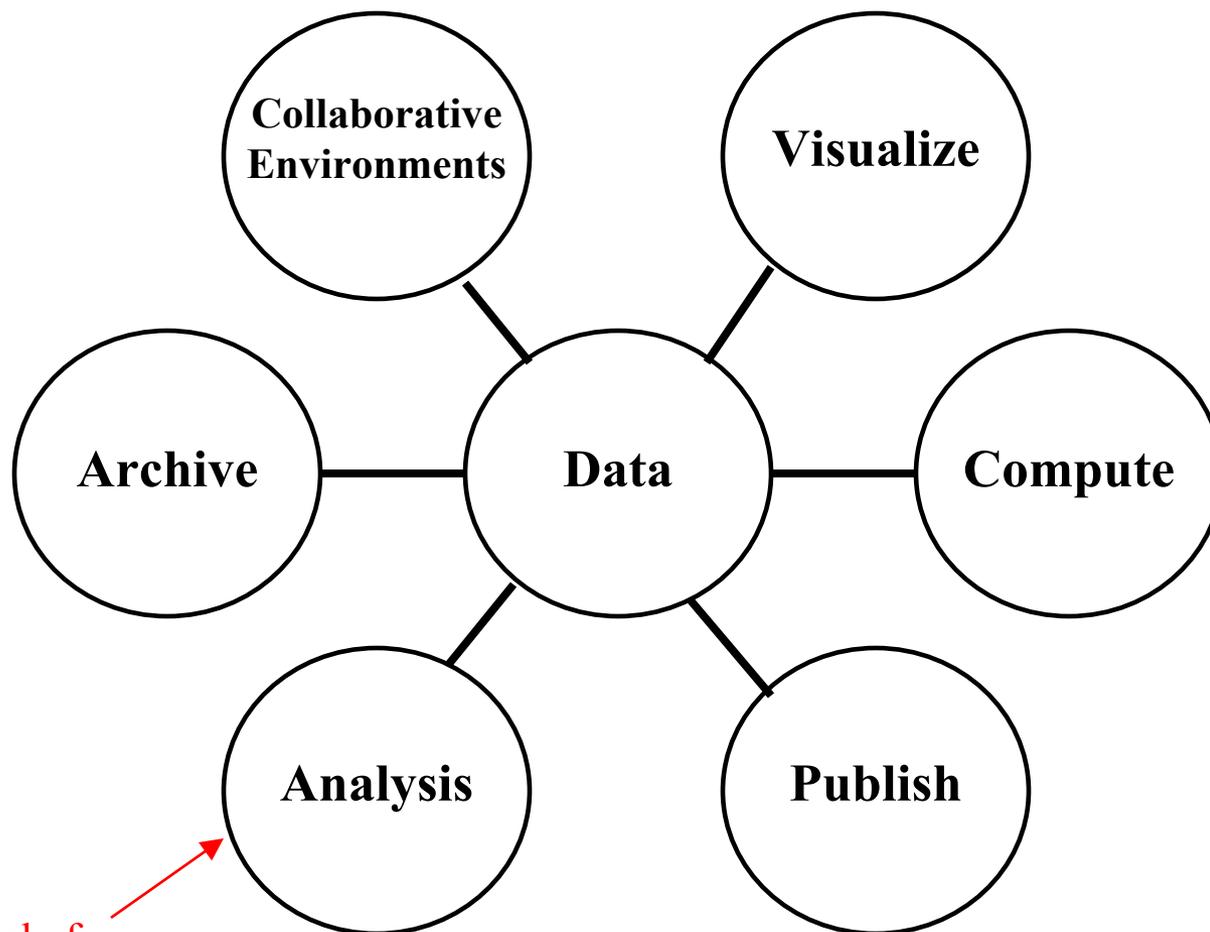
Developing requirements for follow on system



Conceptual Architecture



NASA Center for Computational Sciences



Conceptual framework for an Analysis Environment



Agenda



NASA Center for Computational Sciences

- Introduction — Phil Webster
- Systems Status—Mike Rouch 
- Discover SCU2 – Mike Rouch
- Visualization Services – Carrie Spear
- Data Sharing Services—Ellen Salmon
- User Services — Sadie Duffy
- Questions or Comments



Systems Status



NASA Center for Computational Sciences

- Courant Status
- Explore
 - Utilization
 - System Availability
 - Usage
 - Issues/Resolutions
- Discover
 - Utilization
 - System Availability
 - Usage
 - Issues/Resolutions
 - What's New



Courant Status



NASA Center for Computational Sciences

- System will be decommissioned - *Jan 31, 2008*

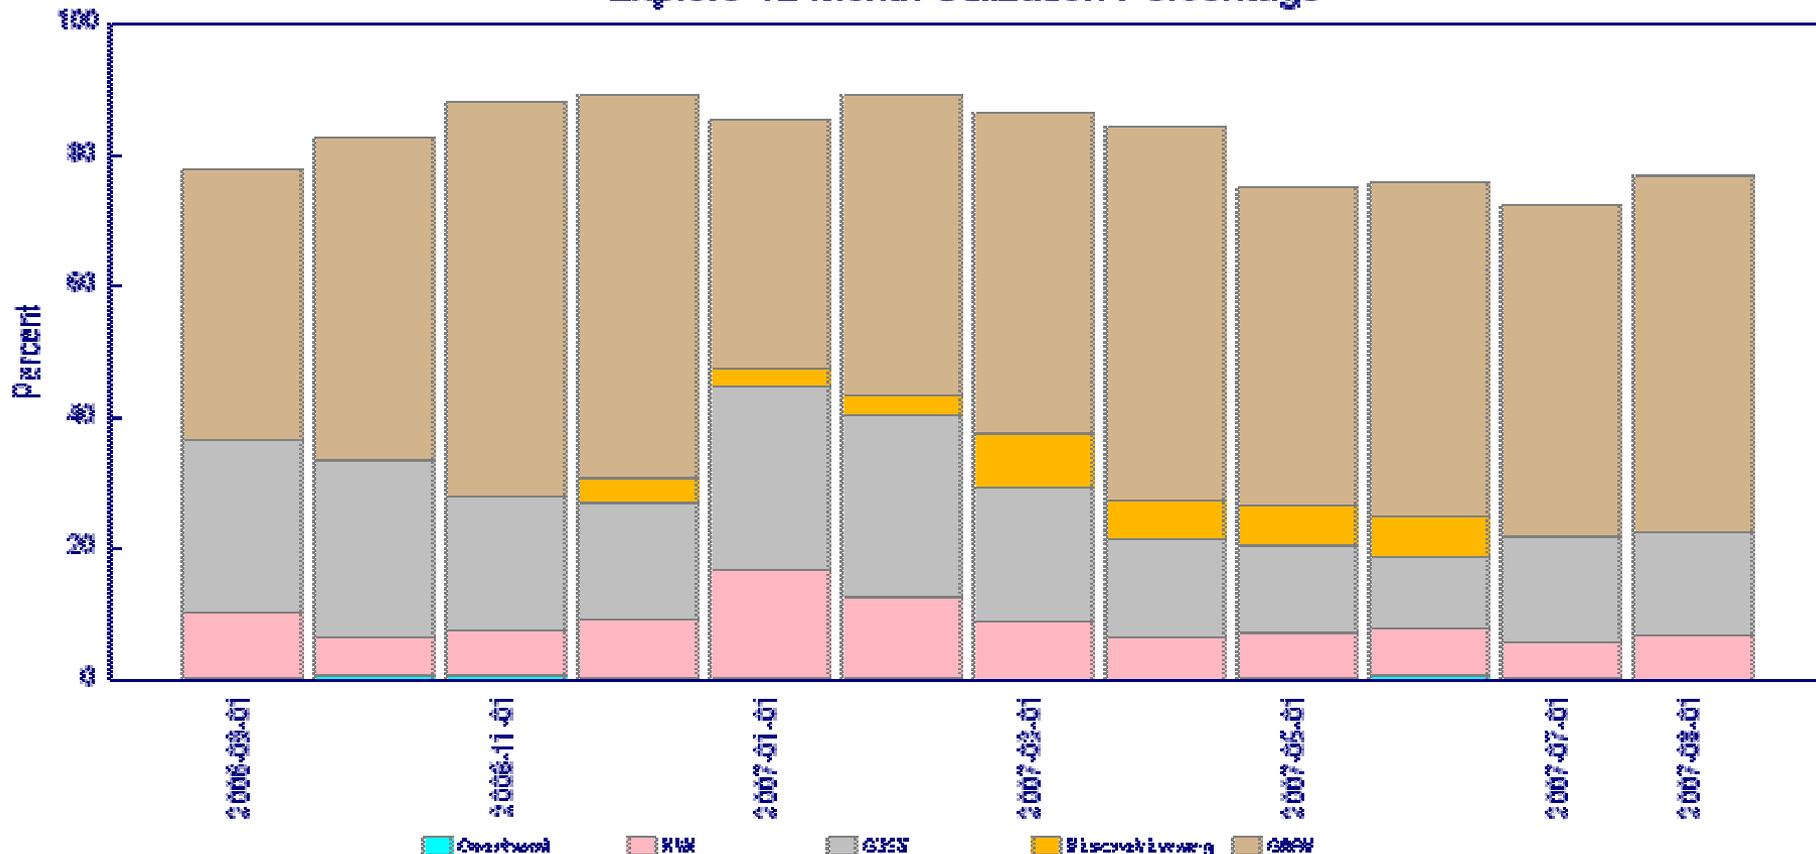


Explore Utilization Past 12 Months



NASA Center for Computational Sciences

Explore 12 Month Utilization Percentage

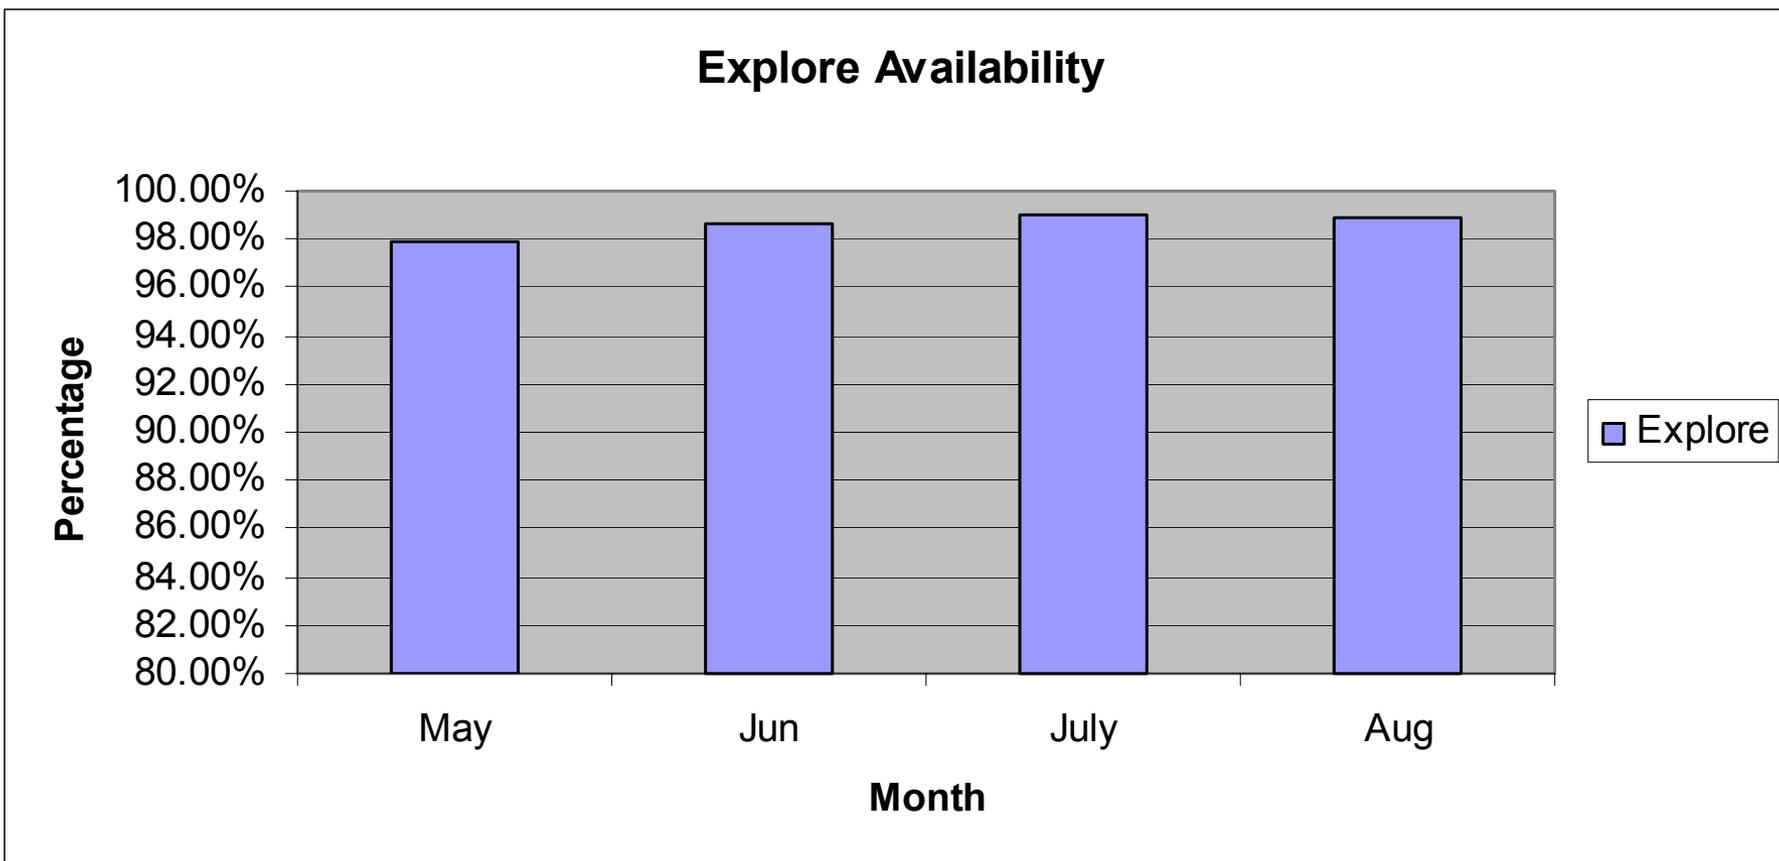




SGI Explore Availability



NASA Center for Computational Sciences





Explore Queue Expansion Factor

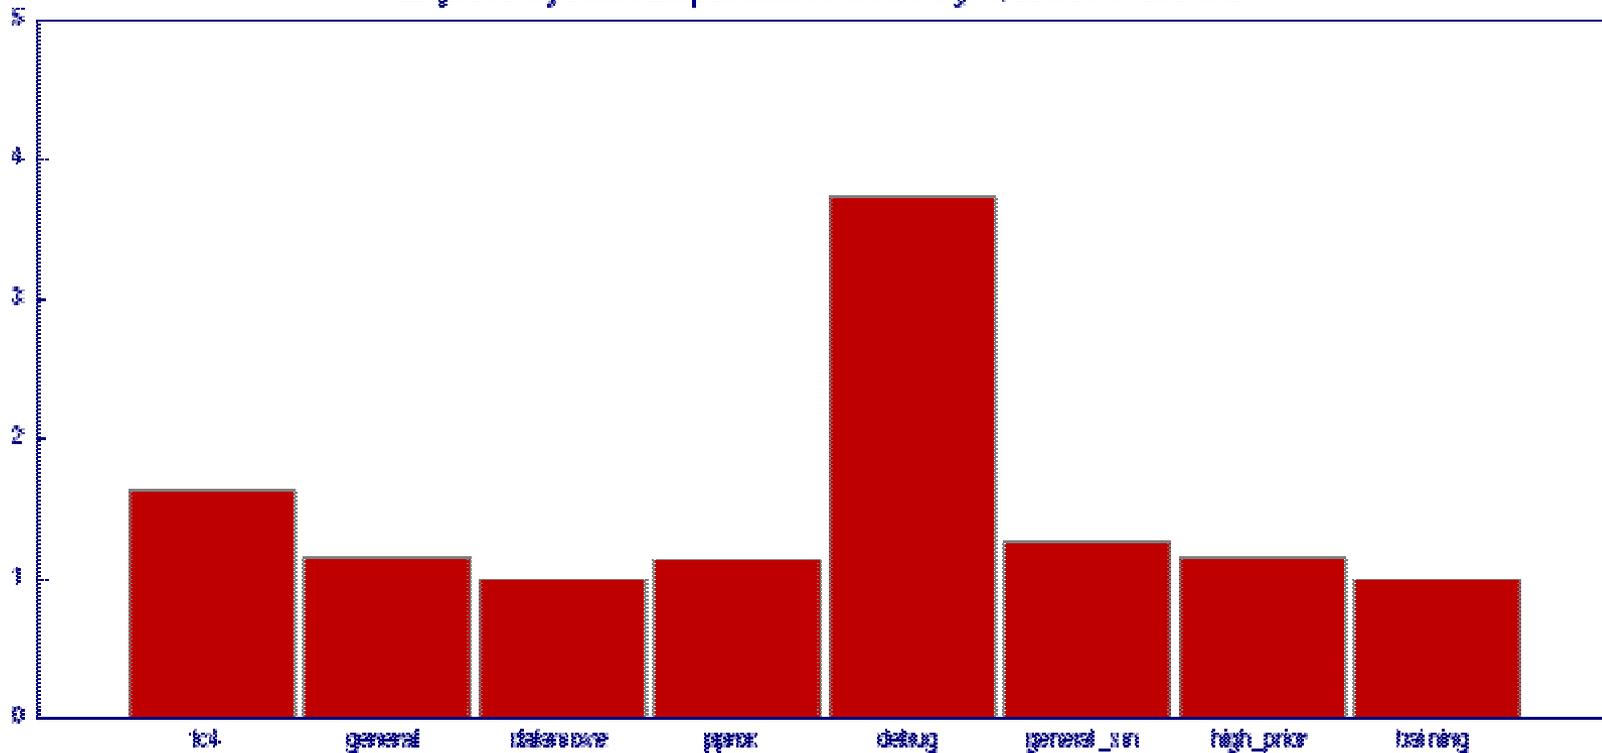


NASA Center for Computational Sciences

$$\frac{\text{Queue Wait Time} + \text{Run Time}}{\text{Run Time}}$$

Weighted over all queue's for all jobs

Explore system Expansion Factor by Queue: 6 Weeks





Explore Issues



NASA Center for Computational Sciences

Eliminate Data Corruption – SGI Systems

Issue: Files being written at the time of an SGI system crash **MAY** be corrupted. However, files appear to be normal.

Interim Steps: Careful Monitoring

- ✓ Install UPS – **COMPLETED 4/11/2007**
- ✓ Continue Monitoring
- ✓ Daily Sys Admins scan files for corruption and directly after a crash
- All affected users are notified

Fix: SGI will provide XFS file system patch

- Awaiting fix – Progress being made by SGI
- Will schedule installation after successful testing



Recent Explore Improvements



NASA Center for Computational Sciences

- **Improving File Data Access – Completed July 2007**
 - ✓ Increase File System Data Residency from Days to Months
 - ✓ Analysis completed; New File System being created
 - ✓ Scheduling with users to move data into new file systems
- **Enhancing Systems – Completed May 2007**
 - Software OS & CxFS upgrades to Irix
 - ✓ Irix 6.5.29
 - ✓ CxFS 4.04 Server
 - Software OS & CxFS upgrades to Altix
 - ✓ Latest SLES .282 Kernel and Patches
 - ✓ CxFS 4.0.4 Client

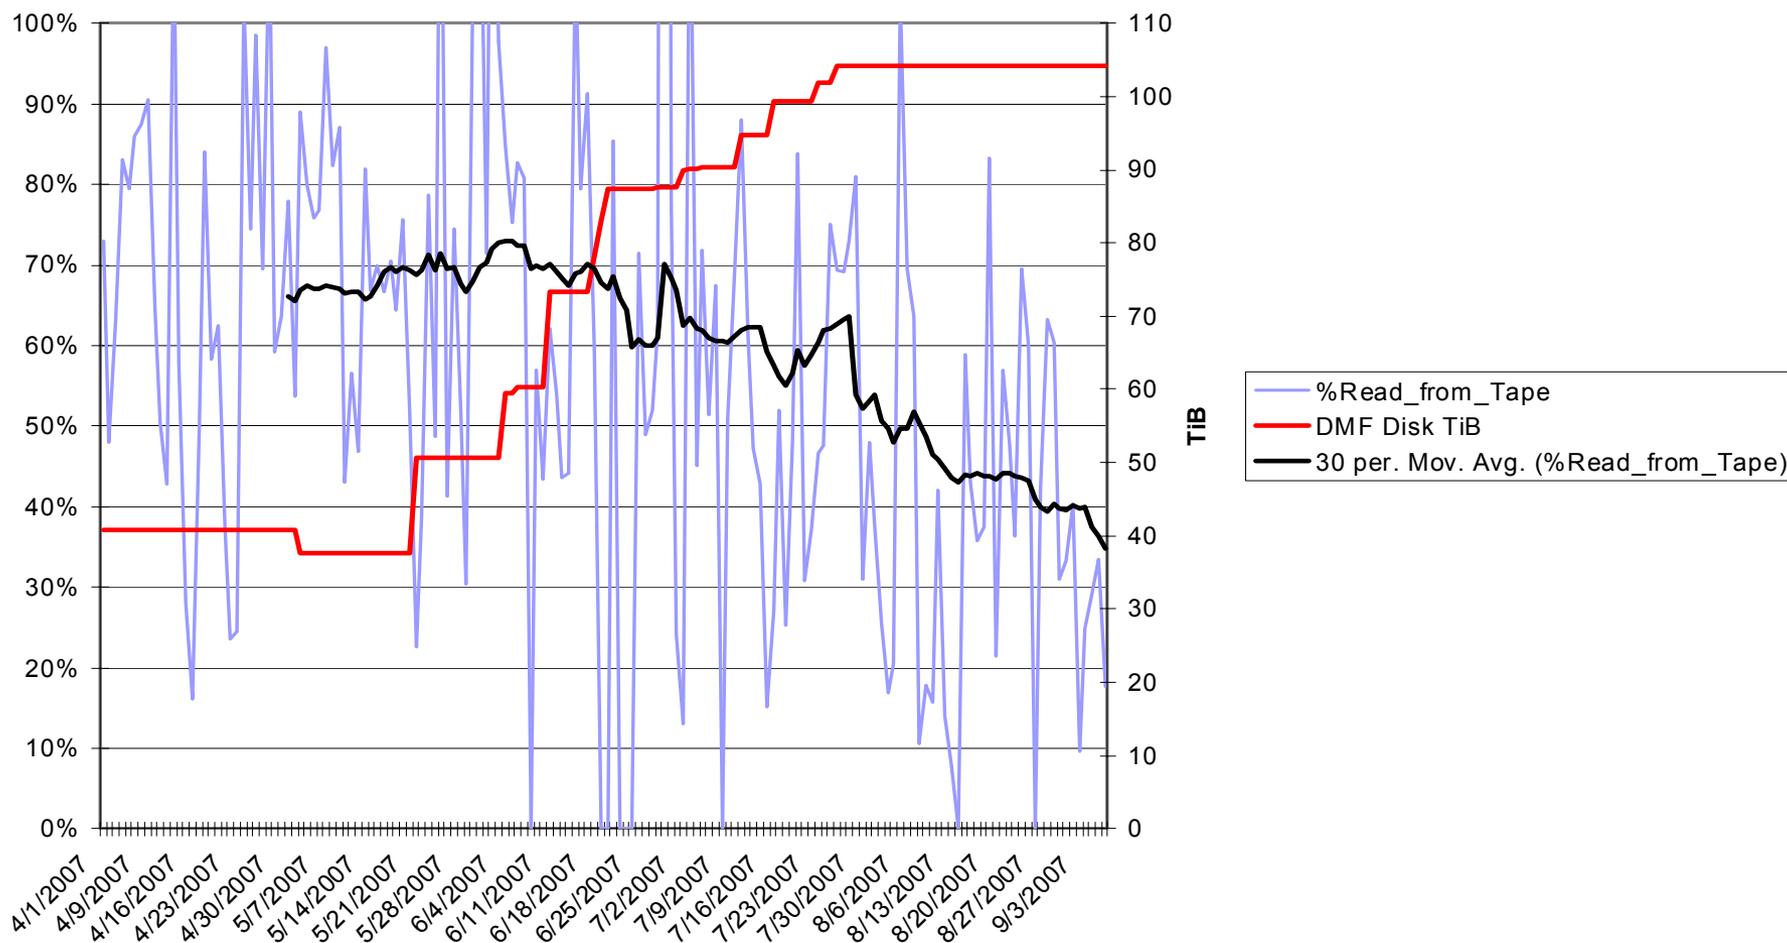


Improved Archive File Data Access



NASA Center for Computational Sciences

**NCCS Tape Retrieve Trends:
Percentage of Data Read from Tape Decreases as DMF Disk Cache Grows**





Recent Explore Improvements



NASA Center for Computational Sciences

– Explore

- LDAP – Completed *Aug 2007*
- Upgraded PBS to 8.0 - Completed *May 2007*



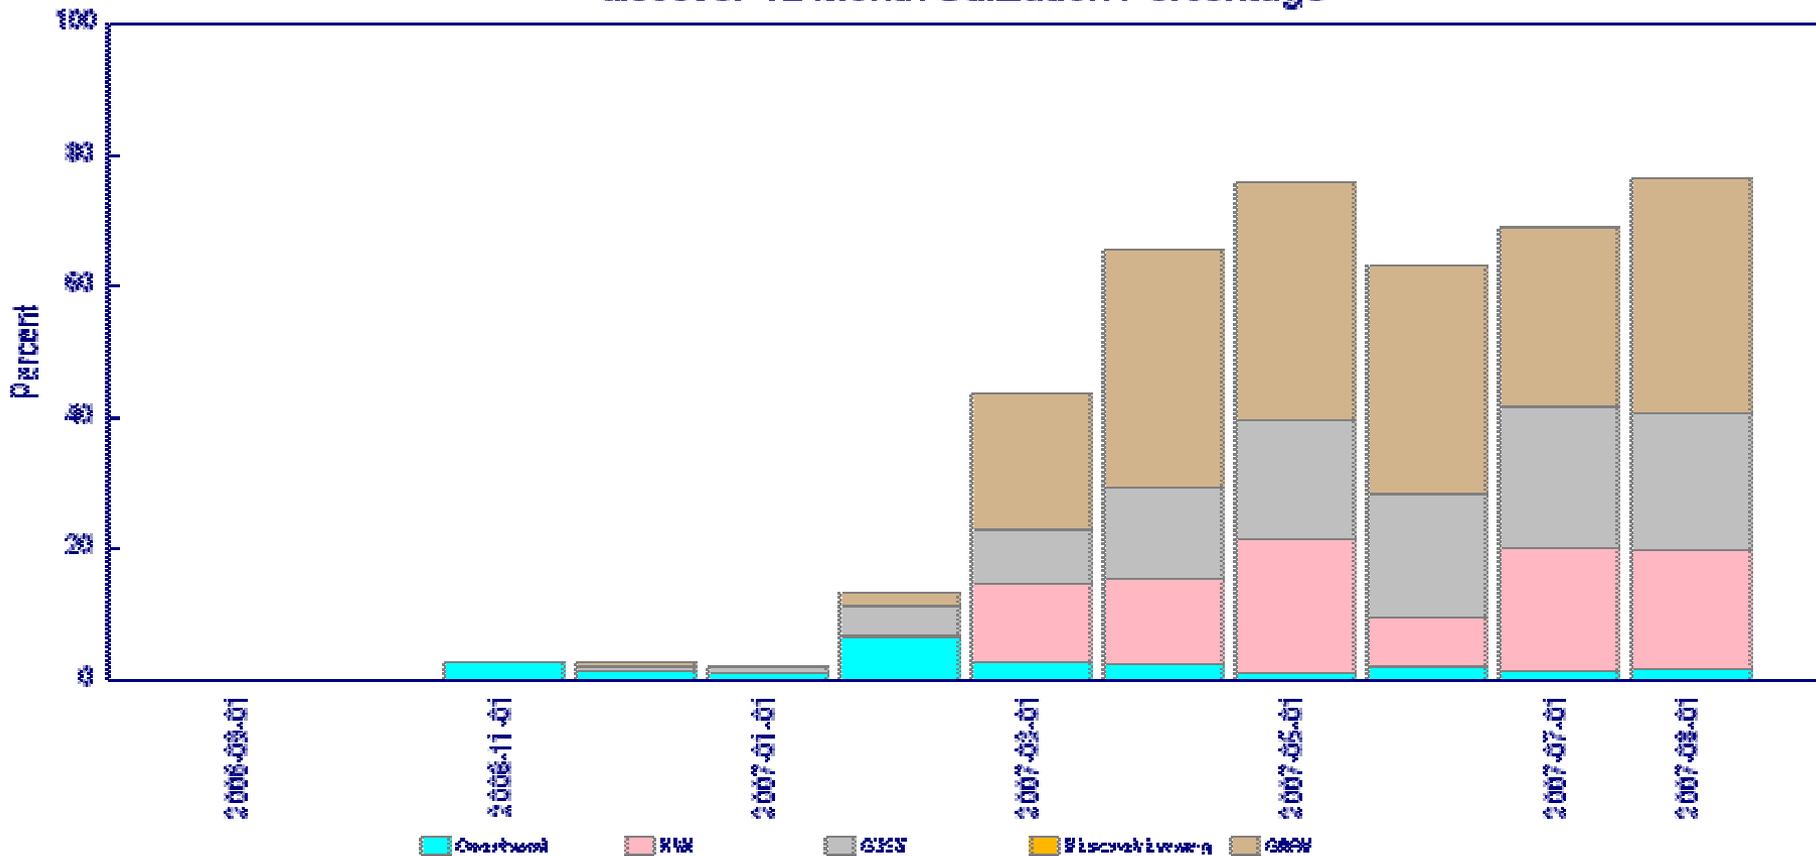
Discover Utilization

Jan – Aug 2007



NASA Center for Computational Sciences

discover 12 Month Utilization Percentage

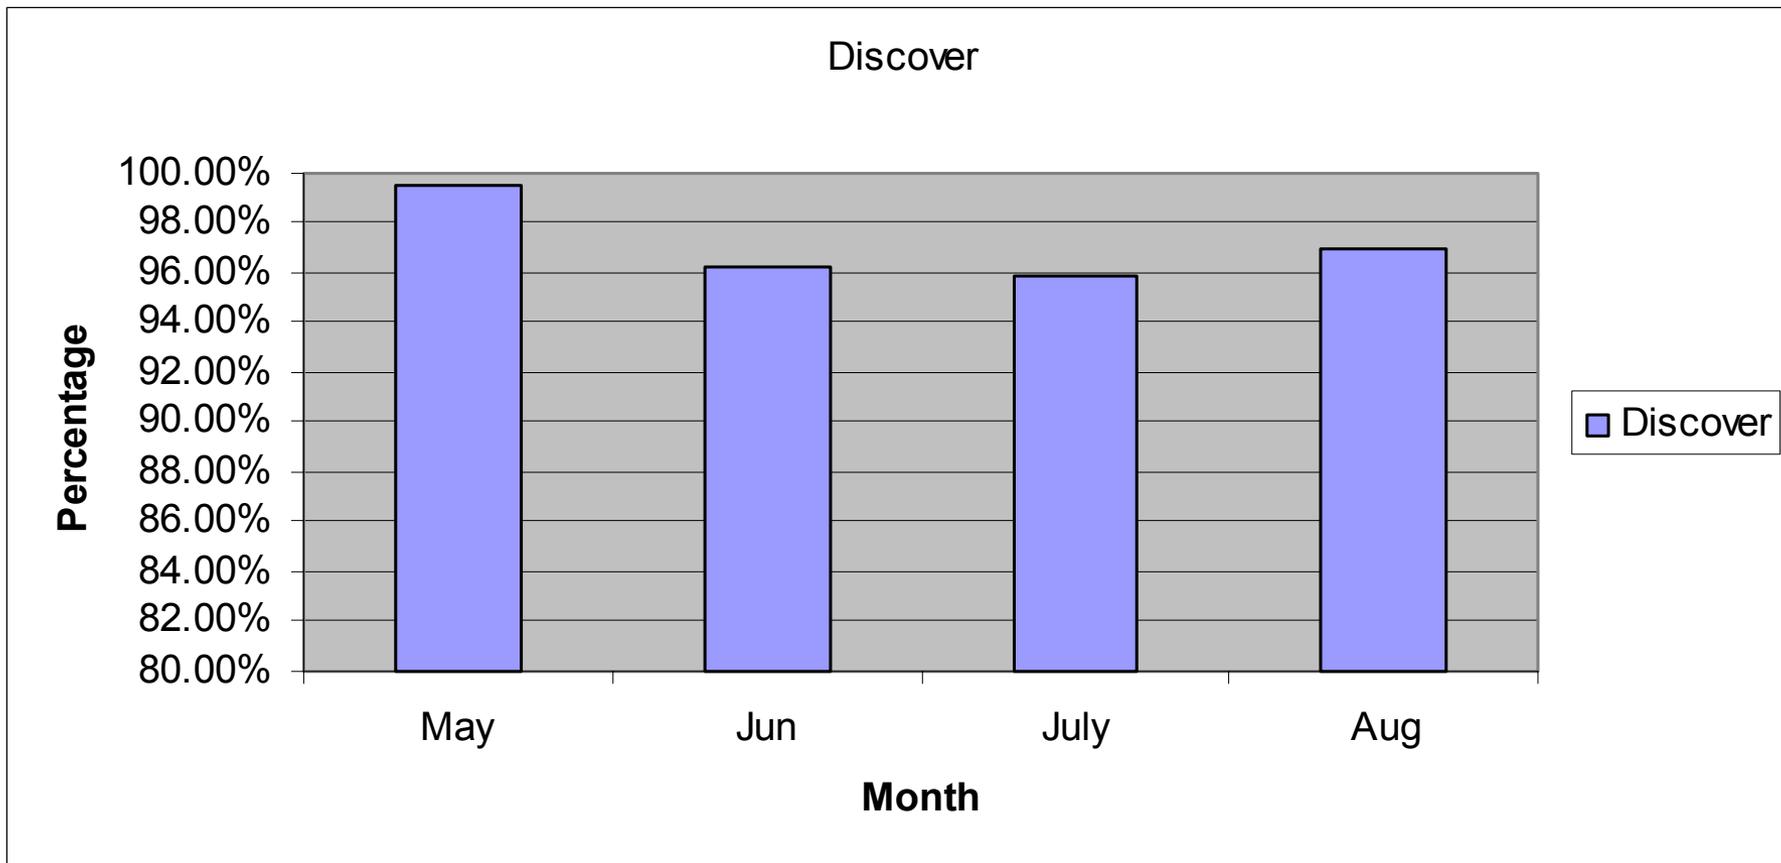




Discover Cluster Availability



NASA Center for Computational Sciences





Discover Queue Expansion Factor

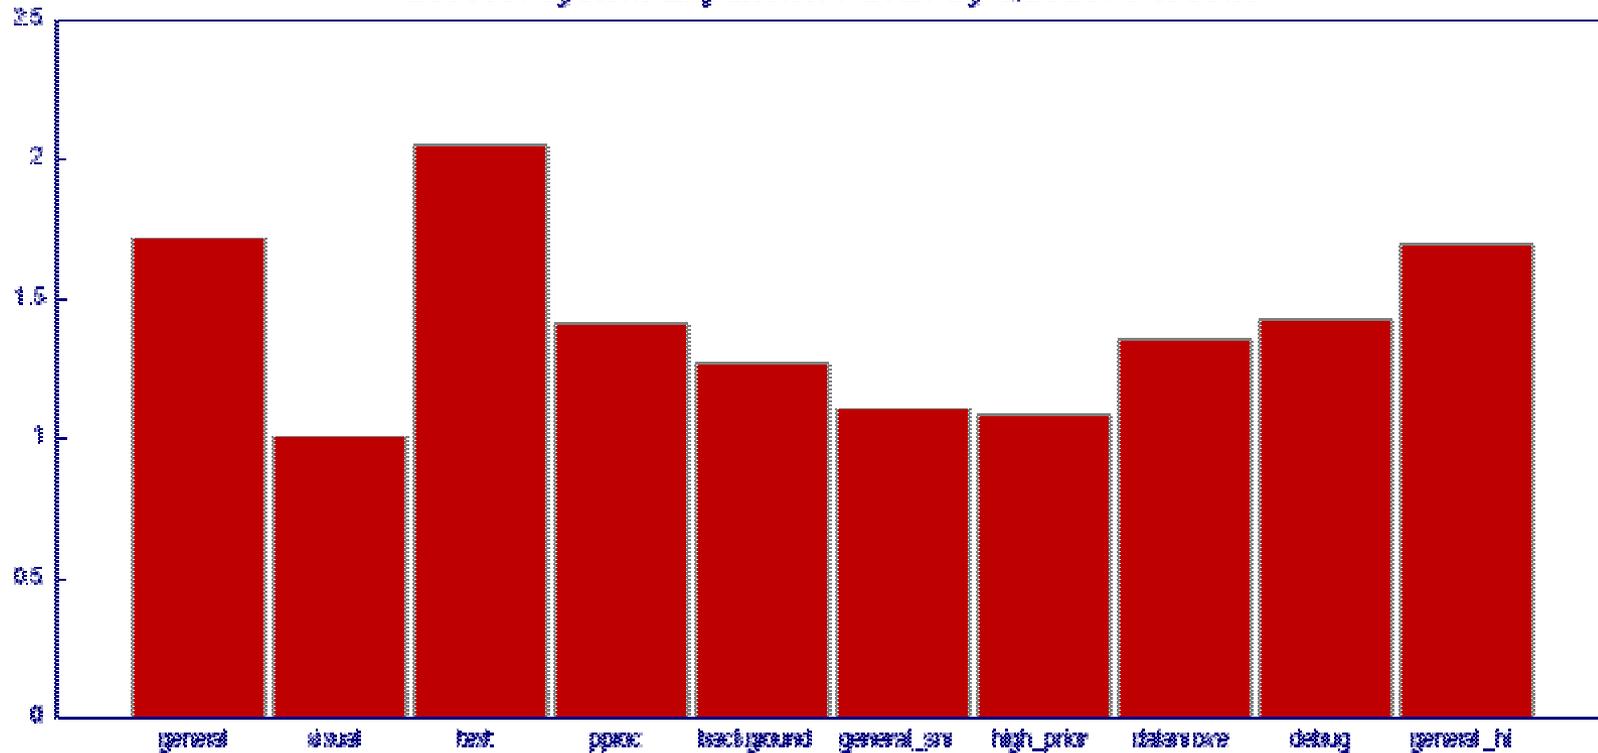


NASA Center for Computational Sciences

$$\frac{\text{Queue Wait Time} + \text{Run Time}}{\text{Run Time}}$$

Weighted over all queue's for all jobs

discover system Expansion Factor by Queue: 6 Weeks





Discover SCU2



NASA Center for Computational Sciences

- 23 July 2007 - NCCS took delivery of additional nodes for *Discover* from Linux Networx.
- Increased capacity includes:
 - 256 dual processor, dual core Intel Woodcrest nodes
 - along with additional specialty login, management and data migration nodes and
 - an additional 70TB of user storage.
- System integration followed pre-defined acceptance test plan.
 - All components run as standalone system to address initial hardware failure due to shipping
 - System connected in early August to current test and validation system to configure nodes with production software stack
 - Nodes moved to production environment in mid-August for further testing
 - 12 September 2007 – Commence 30-day acceptance test



NCCS increased overall capacity of commodity linux cluster by ~11TF

Discover system now ~25TF



Discover Status



NASA Center for Computational Sciences

- SCU2 unit in 30-day acceptance testing
- Open for general use
 - No changes required to user code
 - PBS queues up and running jobs!
 - 1536 cpus when you went home, 2560 cpus when you came in
- We are here to help if you need it

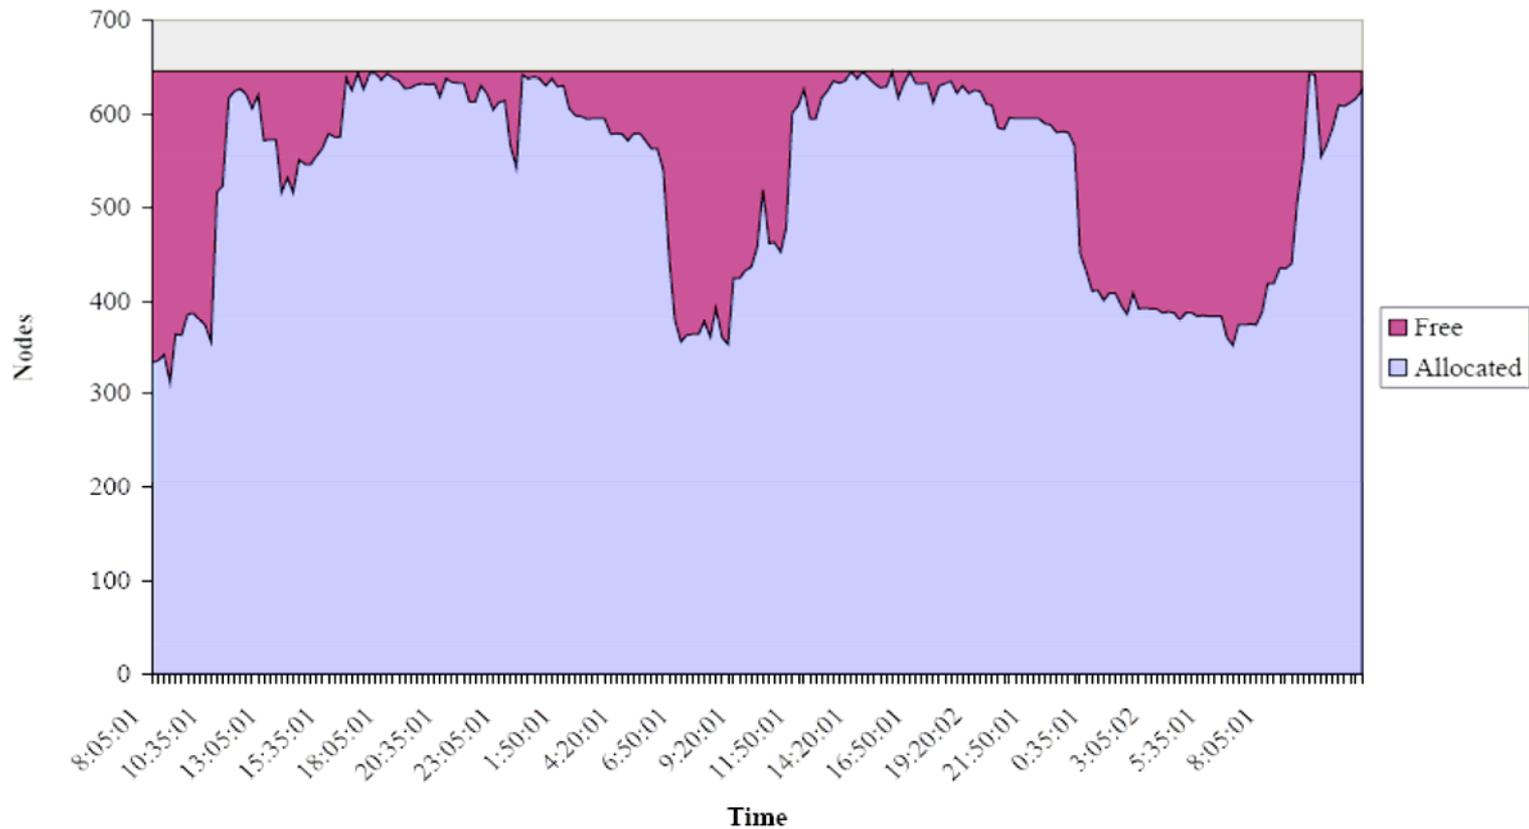


Discover utilization after SCU2



NASA Center for Computational Sciences

Discover Node Usage Since The Addition Of SCU-2





Current Issues

Discover



NASA Center for Computational Sciences

- **Job goes into Swap**

- **Symptom:** When a job is running, one or more nodes goes into a swap condition
- **Outcome:** The processes on those nodes runs very slow causing the total job to run slower.
- **Progress:** Monitoring is in place to trap this condition. The monitoring is working for majority instances. As long as the nodes do not run out of swap, the job should terminate normally.



Current Issues

Discover



NASA Center for Computational Sciences

- **Job Runs Out of Swap**
 - **Symptom:** When a job is running, one or more nodes run out of swap
 - **Outcome:** The nodes become hopelessly hung, requires a reboot and the job dies.
 - **Progress:** Monitoring in place to catch this condition, kill the job before it runs out of swap, notify the user and examine the job. The monitoring is working for majority instances. Also, scripts are in place to cleanup after this condition and it is also working for majority instances.
 - **NOTE:** If your job fails abnormally please call User Services so we can determine why the monitoring scripts did not catch the failure and we can improve error checking.



Future Enhancements



NASA Center for Computational Sciences

- **Enhancing Systems**
 - **Discover Cluster**
 - Software OS
 - SLES 9 SP3 .283 Kernel *Nov 2007*
 - SLES 10 *Jan 2008*
 - **Dirac**
 - LDAP *in the near future*



Agenda



NASA Center for Computational Sciences

- Introduction — Phil Webster
- Systems Status—Mike Rouch
- Discover SCU2 – Mike Rouch
- Visualization Services – Carrie Spear ←
- Data Sharing Services—Ellen Salmon
- User Services — Sadie Duffy
- Questions or Comments



Visualization Services - Discover



NASA Center for Computational Sciences

- **Hardware**
 - 16 Nodes (Currently 8 available through PBS)
 - AMD Processor, 8 GIG memory
 - Graphics hardware acceleration is not available except through a physically connected monitor.
 - Rendering GPU available for applications that leverage this capability.
- **Access**
 - Currently only accessible through PBS on the “visual” queue
 - Has access to all the same GPFS file systems as the rest of discover
 - Would you like them to be externally accessible?
- **Software**
 - IDL (hardware acceleration not available), Ferret
 - What software would you like to see made available?

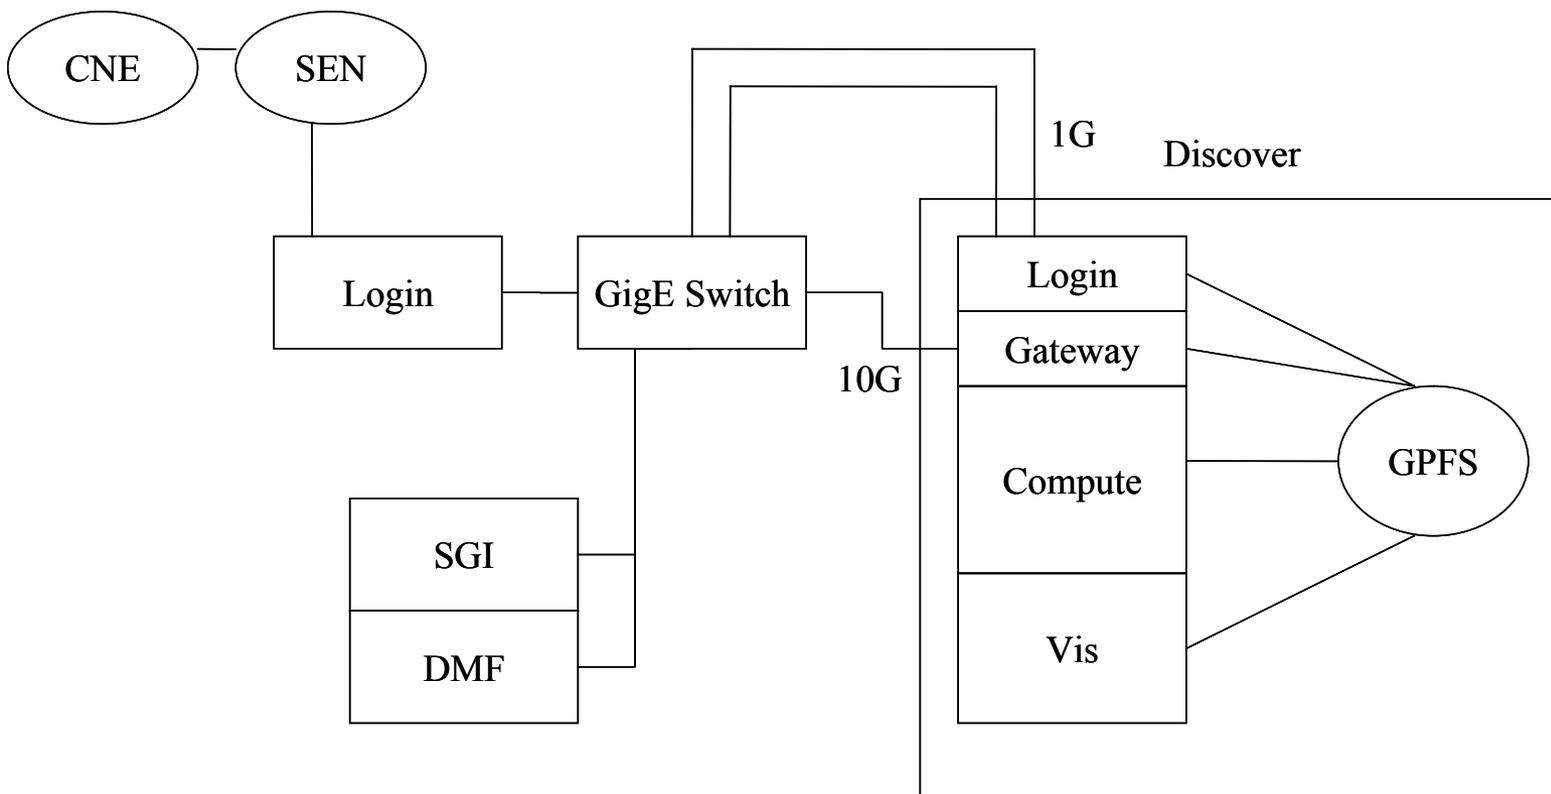
You can contact Carrie through the user services group at support@nccs.nasa.gov



Conceptual Diagram: Discover



NASA Center for Computational Sciences





Visualization Features



NASA Center for Computational Sciences

- User access to viz nodes via login host
- Connect to viz node via PBS (either batch or interactive)
- Direct access to system-wide GPFS file system
- Insight to model output during job execution
- Monitoring capability through analysis/visualization function
- Hyperwall capabilities planned
- Remote display back to user desktop
- Viz output archival to DMF



Agenda



NASA Center for Computational Sciences

- Introduction — Phil Webster
- Systems Status—Mike Rouch
- Discover SCU2 – Mike Rouch
- Visualization Services – Carrie Spear
- Data Sharing Services—Ellen Salmon
- User Services — Sadie Duffy
- Questions or Comments





Data Sharing Services



NASA Center for Computational Sciences

- Data Sharing services
 - Share results with collaborators without requiring NCCS accounts
 - Capabilities include: web access to preliminary data sets with limited viewing and data download
- General Characteristics
 - Data created by NCCS users
 - Support to active SMD projects with finite data sharing requirements
 - Not an on-line archive (future access to NCCS archived data)
- Approach
 - Evolve capabilities for specific projects and generalize for public use
 - Data portal resources managed by the NCCS
 - NASA security/privacy/web/data requirements managed by the NCCS
 - Web access, display, and download features supported by NCCS



Data Sharing Services - Status



NASA Center for Computational Sciences

- Services
 - Web registration (under revision per NPG 1382.1)
 - Directory listings
 - Data download (http, ftp, bbftp)
 - Limited data viewing/display (GrADs, IDL)
- Projects under development
 - TC4
 - OSSE
 - Cloud Library
 - MAP WMS
 - GEOS5 validation
 - Coupled Chemistry
 - GMI



Data Sharing Service Request



NASA Center for Computational Sciences

- **Project:** *SMD Project Name*
- **Sponsor:** *Sponsor Requesting Data Sharing Service*
- **Date:** *Date of Request*

- **Overview:** *Description of the specific SMD project producing data that are needed by collaborators outside of NCCS.*
- **Data:** *Information about data types, owners, and expected access methods to support data stewardship & protection planning. Export Control documentation required.*
- **Access:** *Define collaborators eligible to access data.*
- **Resources:** *Estimate required data volumes and CPU resources.*
- **Duration:** *Define project lifecycle and associated NCCS support.*
- **Capability:** *Description of incremental service development.*
Example:
 - » Web interface to display directory listings & download data
 - » Evaluate usage & data demands
 - » Add thumbnail displays to better identify data files
 - » Implement data subsetting capabilities to reduce download demands on remote users
 - » Reach back into NCCS archive for additional data holdings



Discussion



NASA Center for Computational Sciences

Contact us if you want to explore data sharing opportunities.

Ellen.Salmon@nasa.gov

301-286-7705

Harper.Pryor@GSFC.nasa.gov

301-286-9297



Agenda



NASA Center for Computational Sciences

- Introduction — Phil Webster
- Systems Status—Mike Rouch
- Discover SCU2 – Mike Rouch
- Visualization Services – Carrie Spear
- Data Sharing Services—Ellen Salmon
- User Services — Sadie Duffy ←
- Questions or Comments



User Services



NASA Center for Computational Sciences

- Allocations
 - FY08Q1 allocations due by September 26th, 2007 online at <https://ebooks.reisys.com/gsfc/nccs/submission/index.jsp?soId=27>
- LDAP passwords
 - LDAP in use on discover and explore, if you need your LDAP password please contact us at 301-286-9120 or email us at support@nccs.nasa.gov
- Downtime emails by subscription
 - Every user added by default
 - You can unsubscribe if you do not wish to get these notifications



Login Time outs



NASA Center for Computational Sciences

- As of the 19th of September all inactive login sessions will have a expire after 60 minutes.
- Due to NIST Special Publication 800-53 “Recommended Security Controls for Federal Information Systems”
- Idle is defined as no data being sent to your screen or data being input from your keyboard.
- Messages will be sent prior to session termination



- Questions?
- Comments?