



NCCS User Forum

25 September 2008

GSEFC
NCCS



Agenda

Welcome & Introduction
Phil Webster, CISTO Chief
Scott Wallace, CSC PM

Current System Status
Fred Reitz, Operations Manager

Compute Capabilities at the NCCS
Dan Duffy, Lead Architect

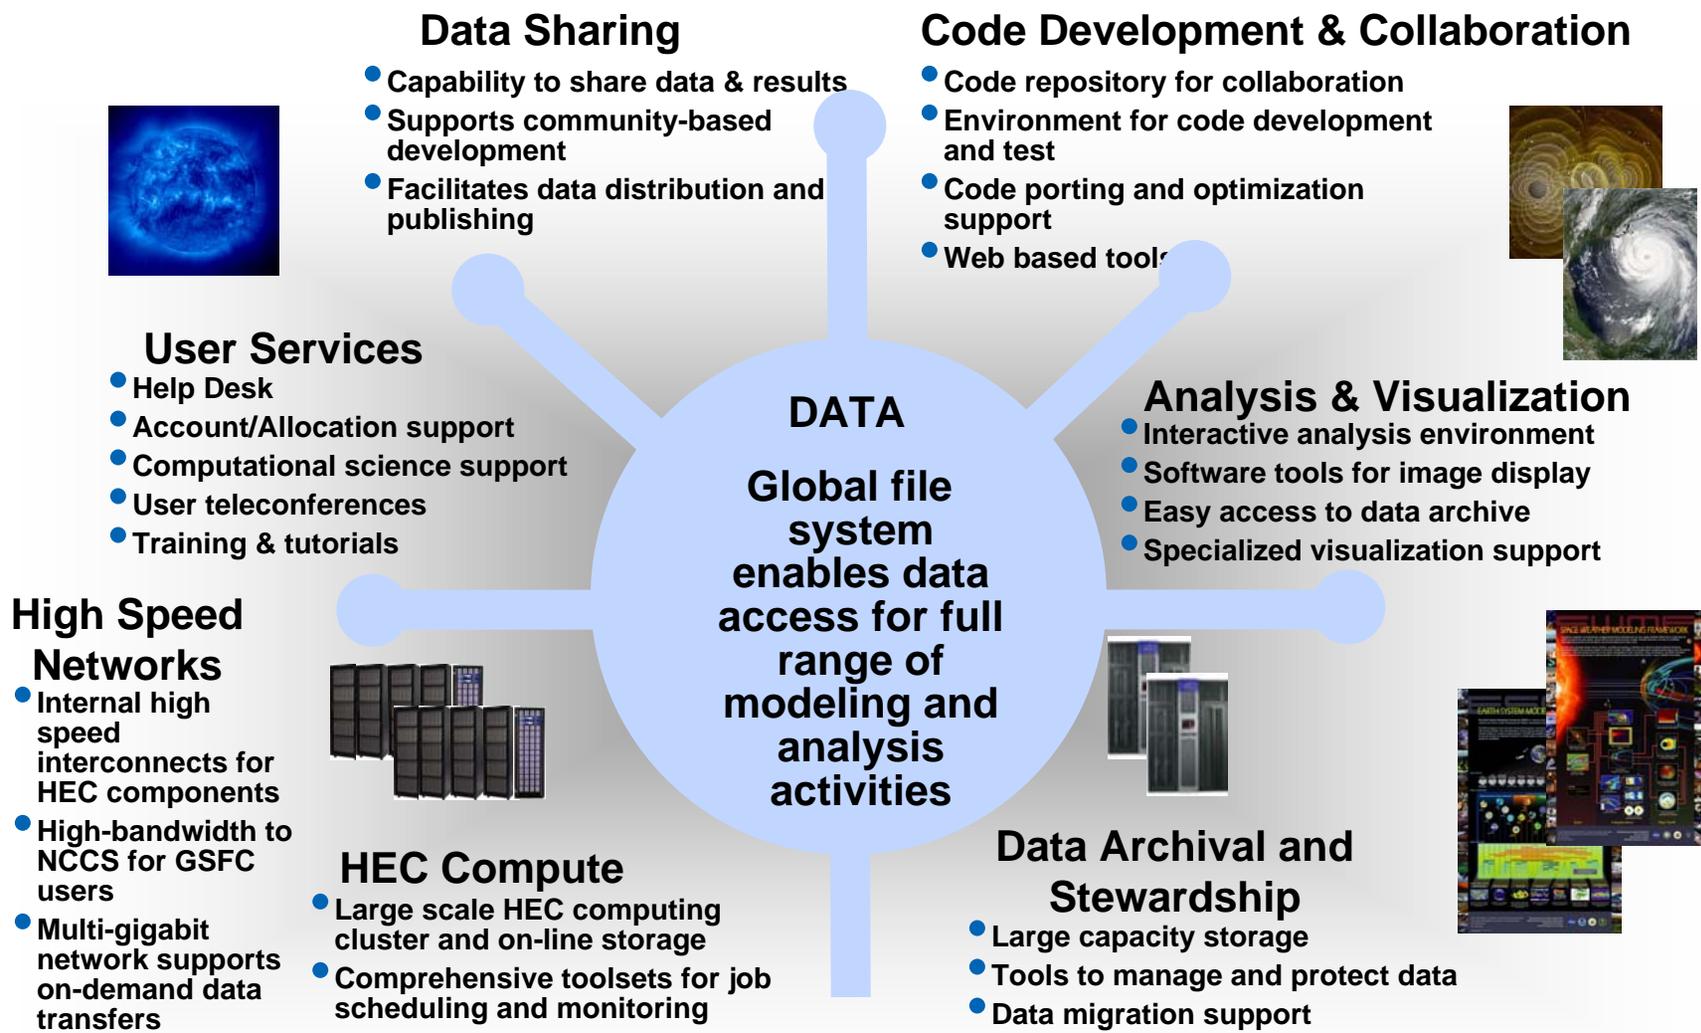
SMD Allocations
Sally Stemwedel, HEC Allocation Specialist

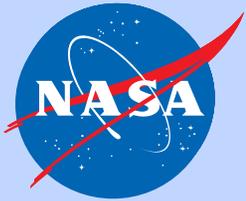
User Services Updates
Bill Ward, User Services Lead

Questions and Comments
Phil Webster, CISTO Chief

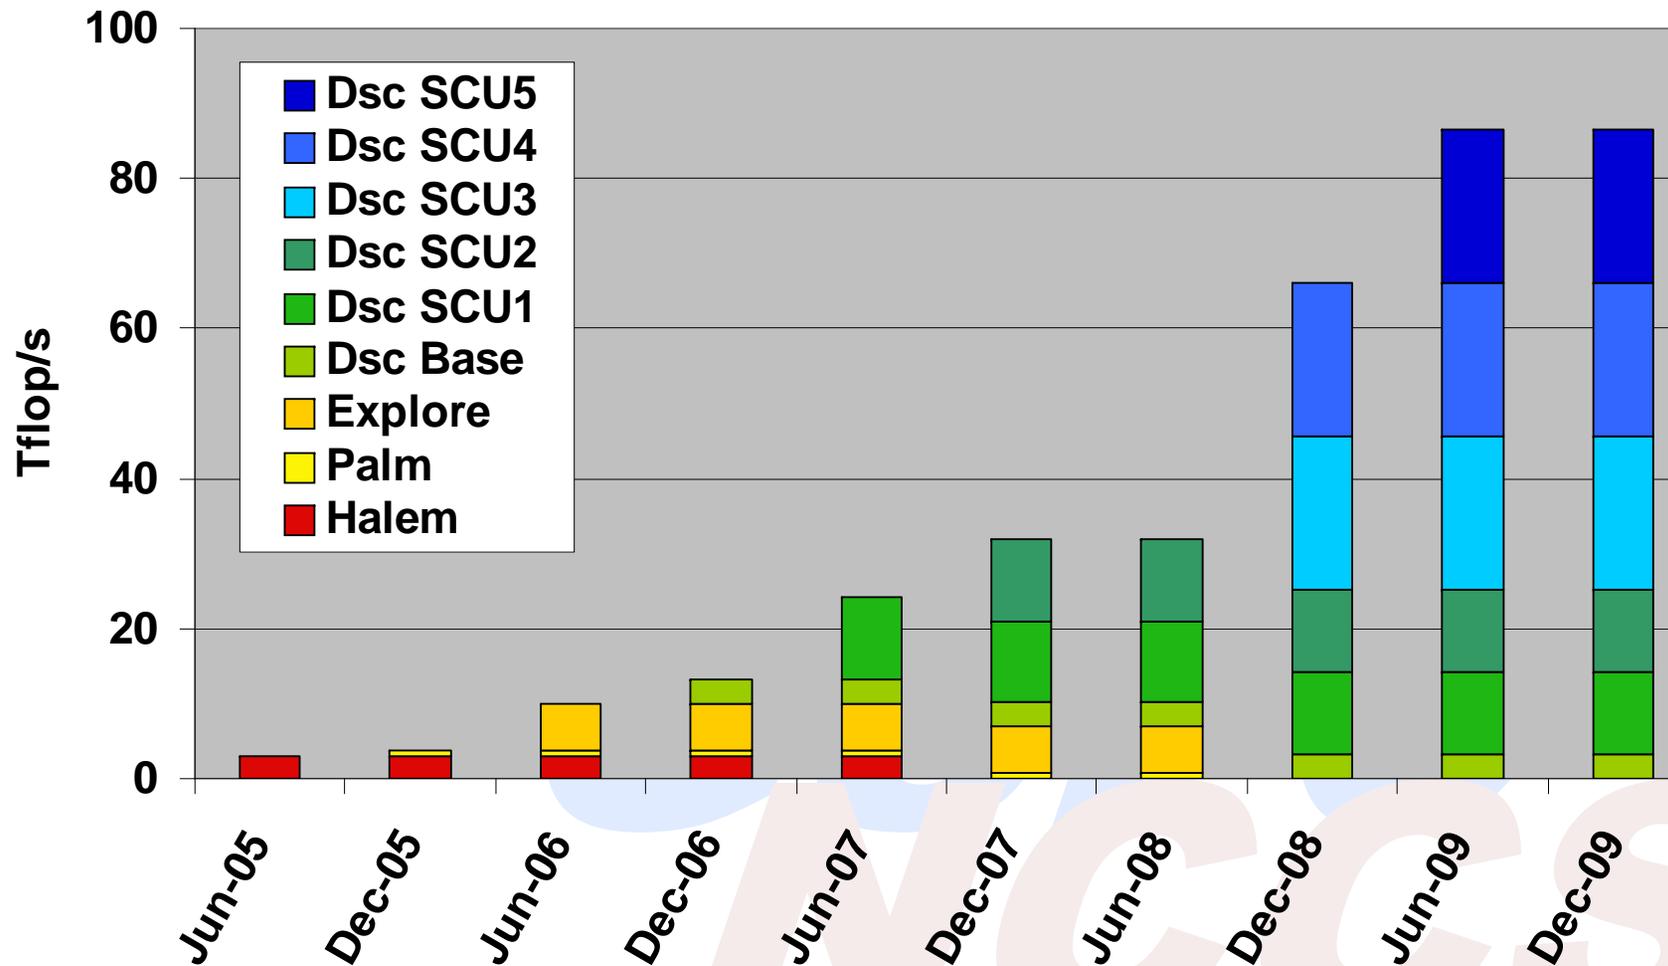


NCCS Support Services





Resource Growth at the NCCS





NCCS Staff Transitions

- New Govt. Lead Architect: Dan Duffy
(301) 286-8830, Daniel.Q.Duffy@nasa.gov
- New CSC Lead Architect: Jim McElvaney
(301) 286-2485, James.D.McElvaney@nasa.gov
- New User Services Lead: Bill Ward
(301) 286-2954, William.A.Ward@nasa.gov
- New HPC System Administrator: Bill Woodford
(301) 286-5852, William.E.Woodford@nasa.gov



Key Accomplishments

- SLES10 upgrade
- GPFS 3.2 upgrade
- Integrated SCU3
- Data Portal storage migration
- Transition from Explore to Discover

GSEFC
NCCS



Integration of Discover SCU4

- SCU4 to be connected to Discover 10/1 (date firm)
- Staff testing and scalability 10/1–10/5 (dates approximate)
- Opportunity for interested users to run large jobs 10/6–10/13 (dates approximate)

GSEFC
NCCS



Agenda

Welcome & Introduction
Phil Webster, CISTO Chief
Scott Wallace, CSC PM

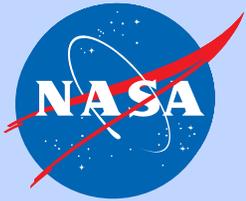
Current System Status
Fred Reitz, Operations Manager

Compute Capabilities at the NCCS
Dan Duffy, Lead Architect

SMD Allocations
Sally Stemwedel, HEC Allocation Specialist

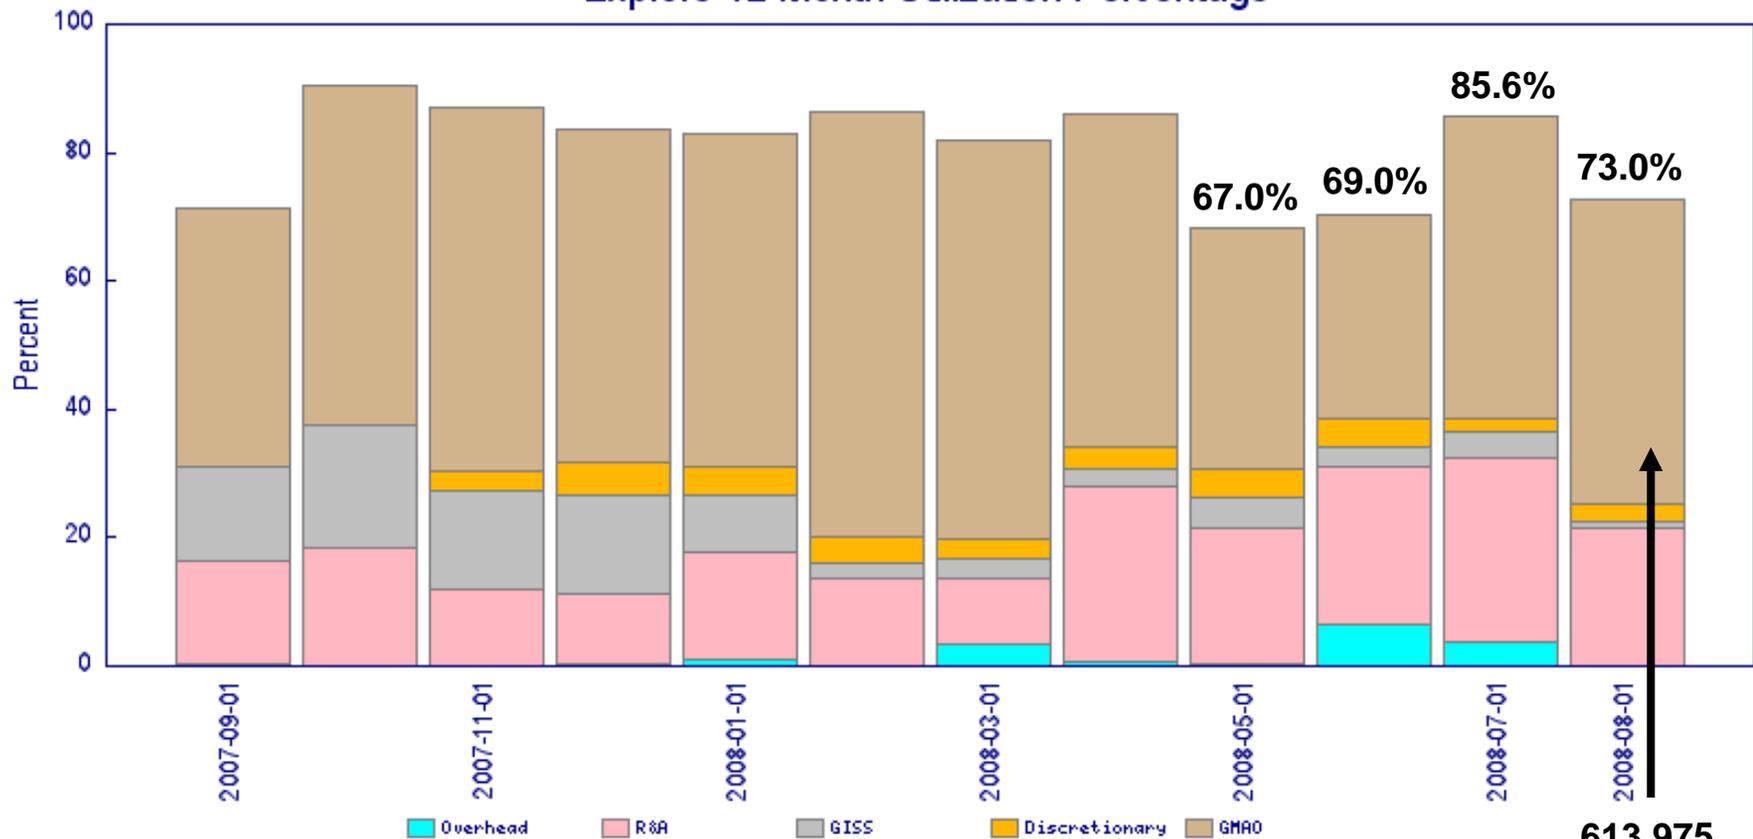
User Services Updates
Bill Ward, User Services Lead

Questions and Comments
Phil Webster, CISTO Chief

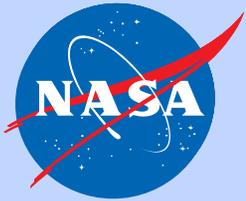


Explore Utilization Past Year

Explore 12 Month Utilization Percentage



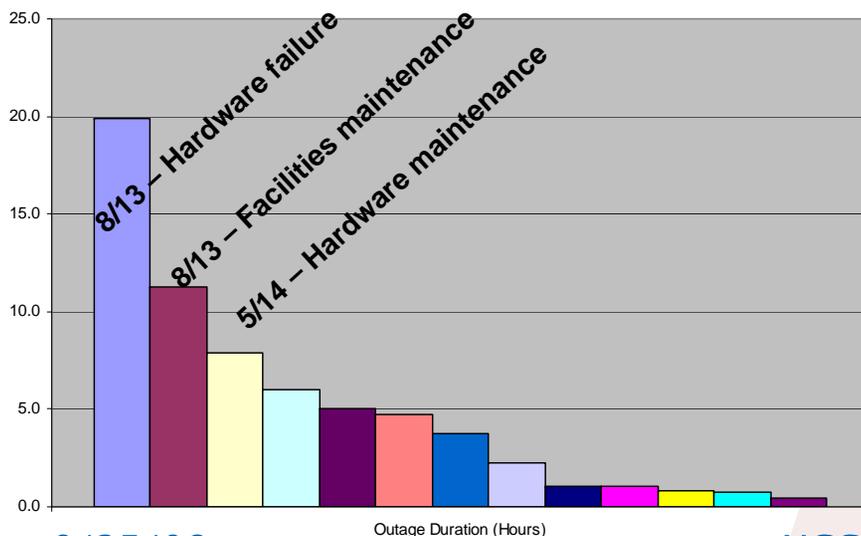
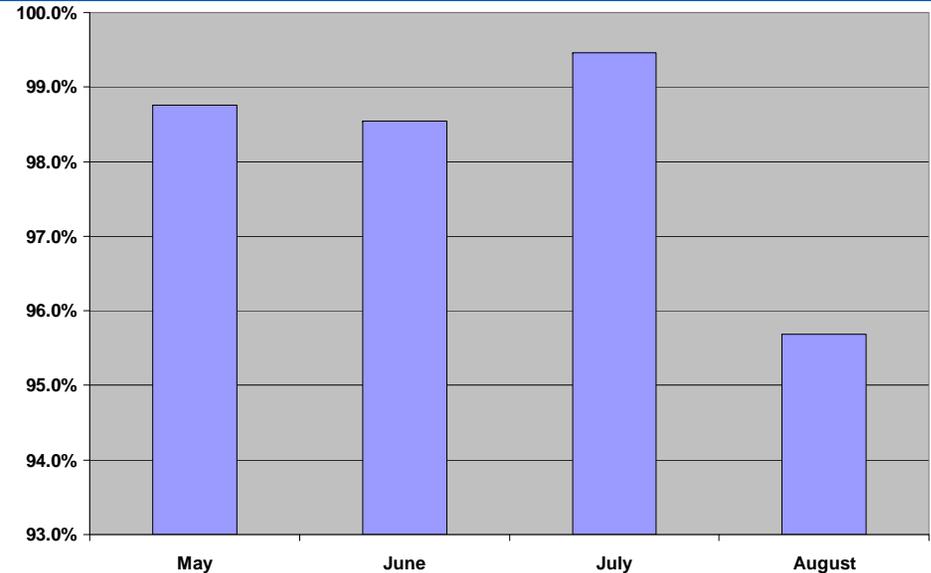
613,975
CPU hours



Explore Availability

May through August availability

- 13 outages
 - ▶ 8 unscheduled
 - ◆ 7 hardware failures
 - ◆ 1 human error
 - ▶ 5 scheduled
- 65 hours total downtime
 - ▶ 32.8 unscheduled
 - ▶ 32.2 scheduled



Longest outages

- 8/13 – Hardware failure, 19.9 hrs
 - ▶ E2 hardware issues after scheduled power down for facilities maintenance
 - ▶ System left down until normal business hours for vendor repair
 - ▶ Replaced/unseated PCI bus, IB connection
- 8/13 – Hardware maintenance, 11.25 hrs
 - ▶ Scheduled outage
 - ▶ Electrical work – facility upgrade

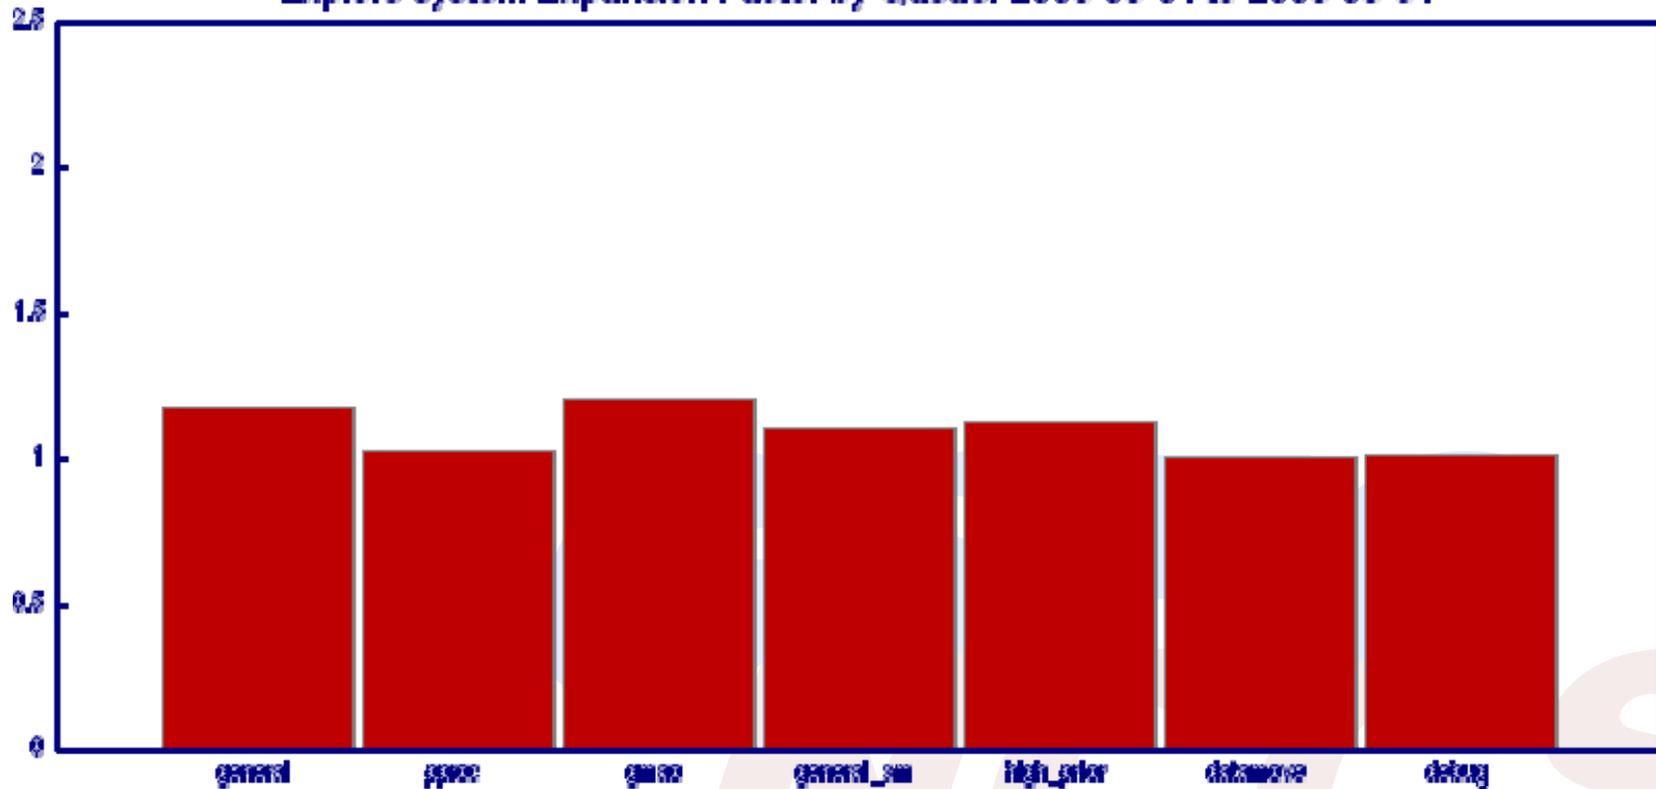


Explore Queue Expansion Factor

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

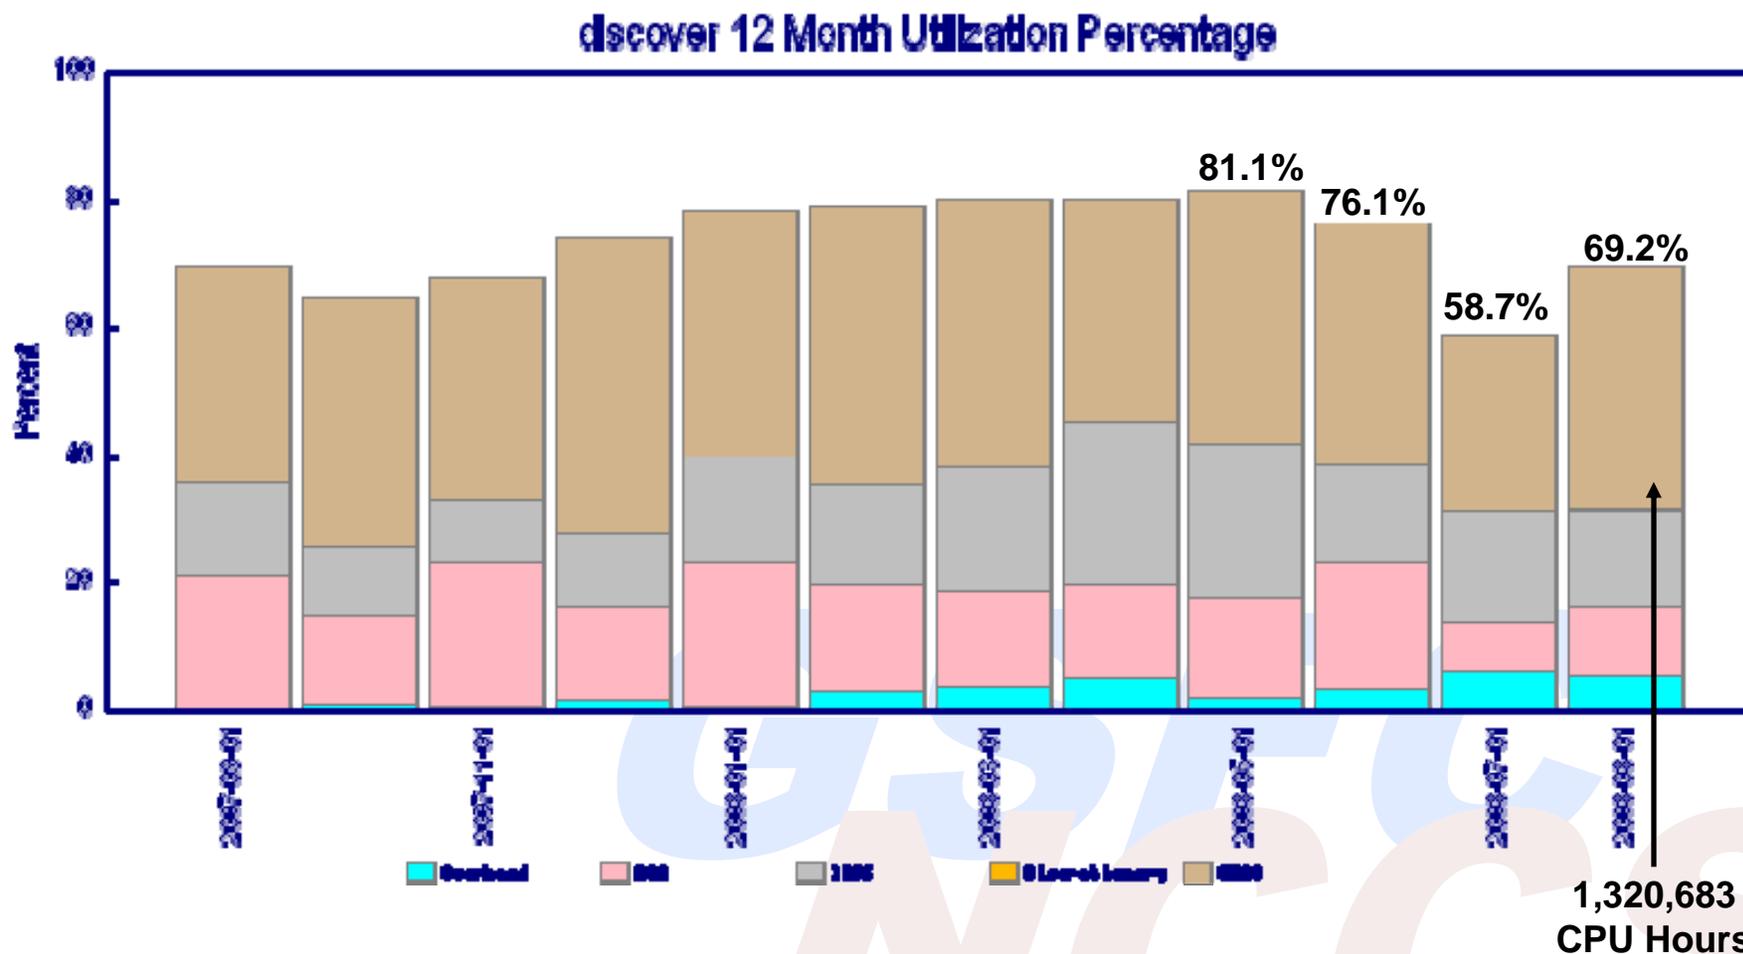
Weighted over all queues for all jobs
(Background and Test queues excluded)

Explore system Expansion Factor by Queue: 2008-08-01 to 2008-08-31



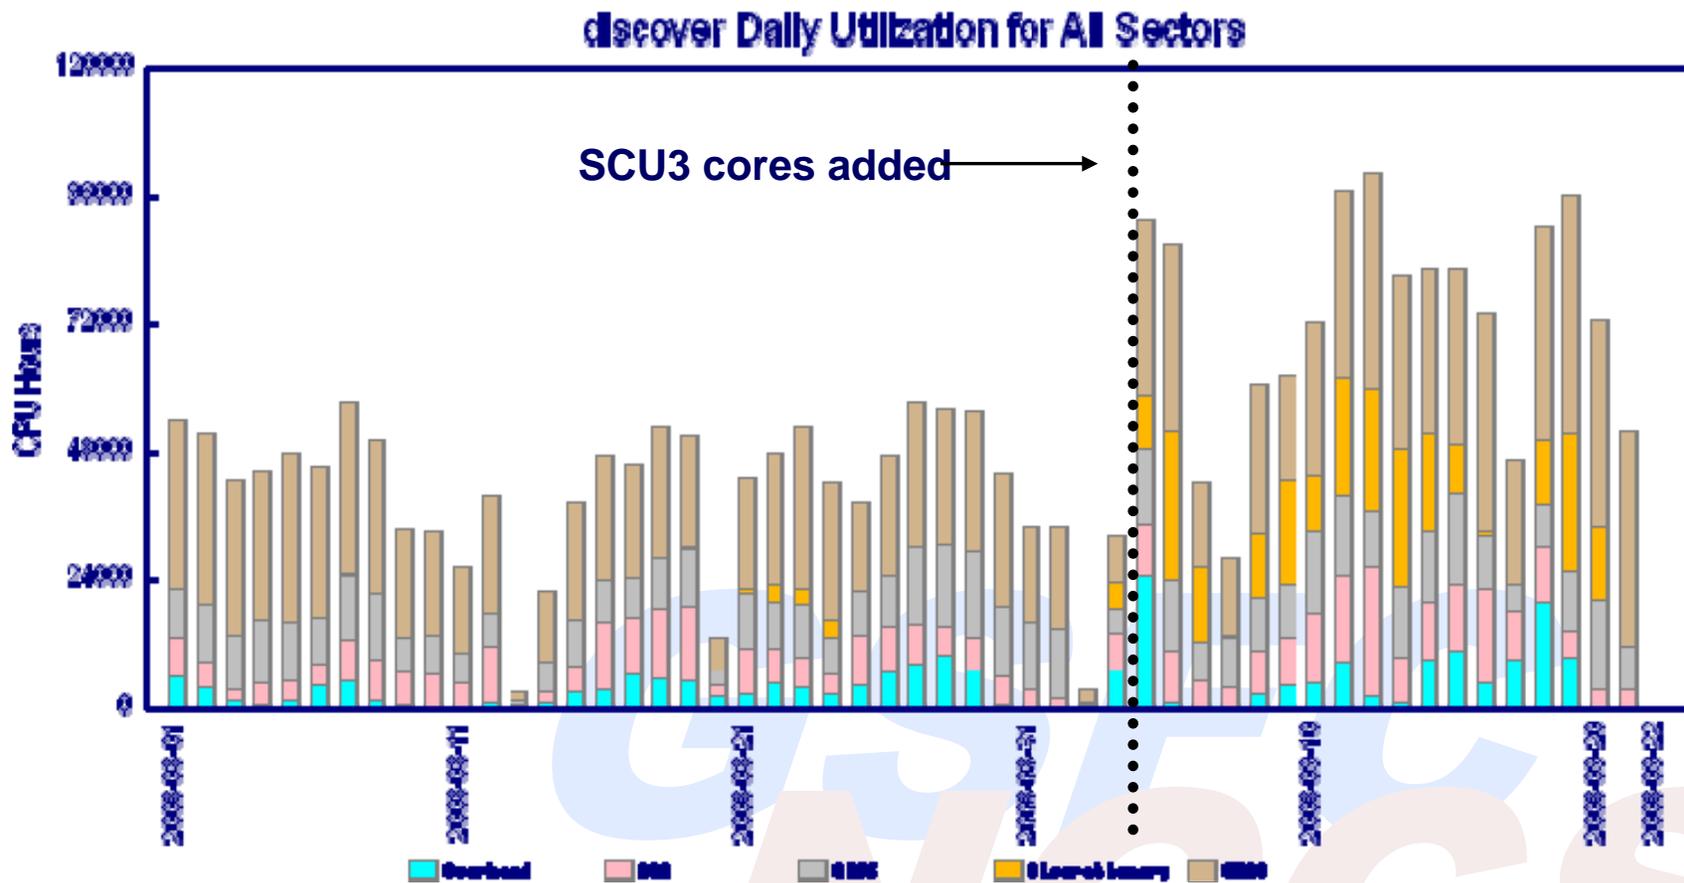


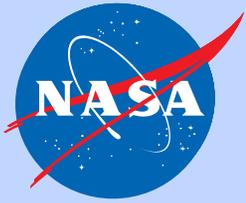
Discover Utilization Past Year





Discover Utilization

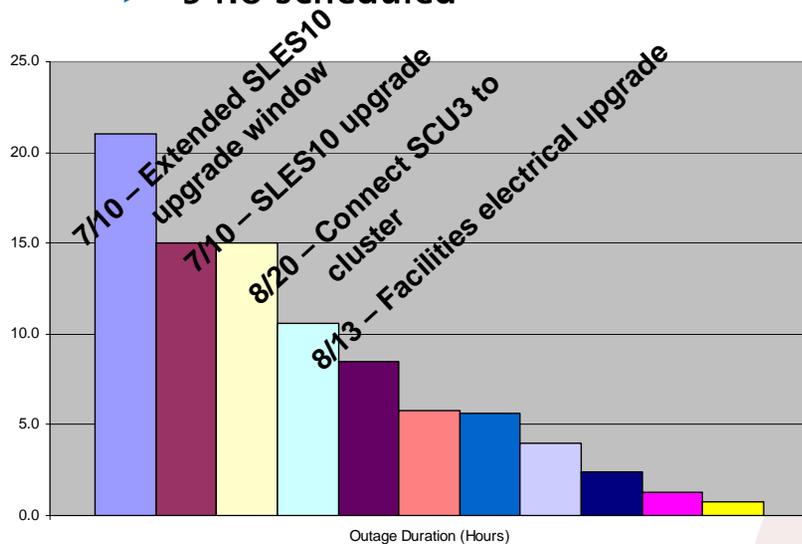
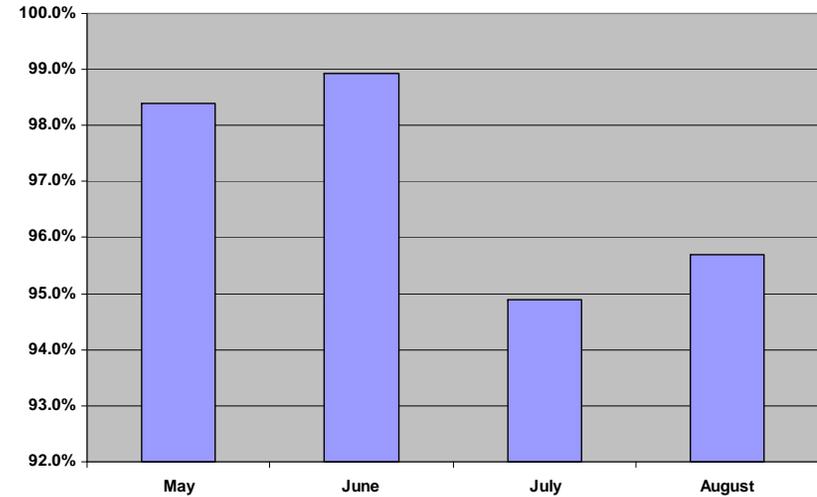




Discover Availability

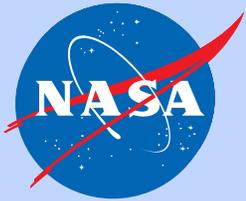
May through August availability

- 11 outages
 - ▶ 6 unscheduled
 - ◆ 2 hardware failures
 - ◆ 2 software failures
 - ◆ 3 extended maintenance windows
 - ▶ 5 scheduled
- 89.9 hours total downtime
 - ▶ 35.1 unscheduled
 - ▶ 54.8 scheduled



Longest outages

- 7/10 - SLES 10 upgrade, 36 hrs
 - ▶ 15 hours planned
 - ▶ 21 hour extension
- 8/20 - Connect SCU3, 15 hrs
 - ▶ Scheduled outage
- 8/13 - Facilities maintenance, 10.6 hrs
 - ▶ Electrical work - facility upgrade

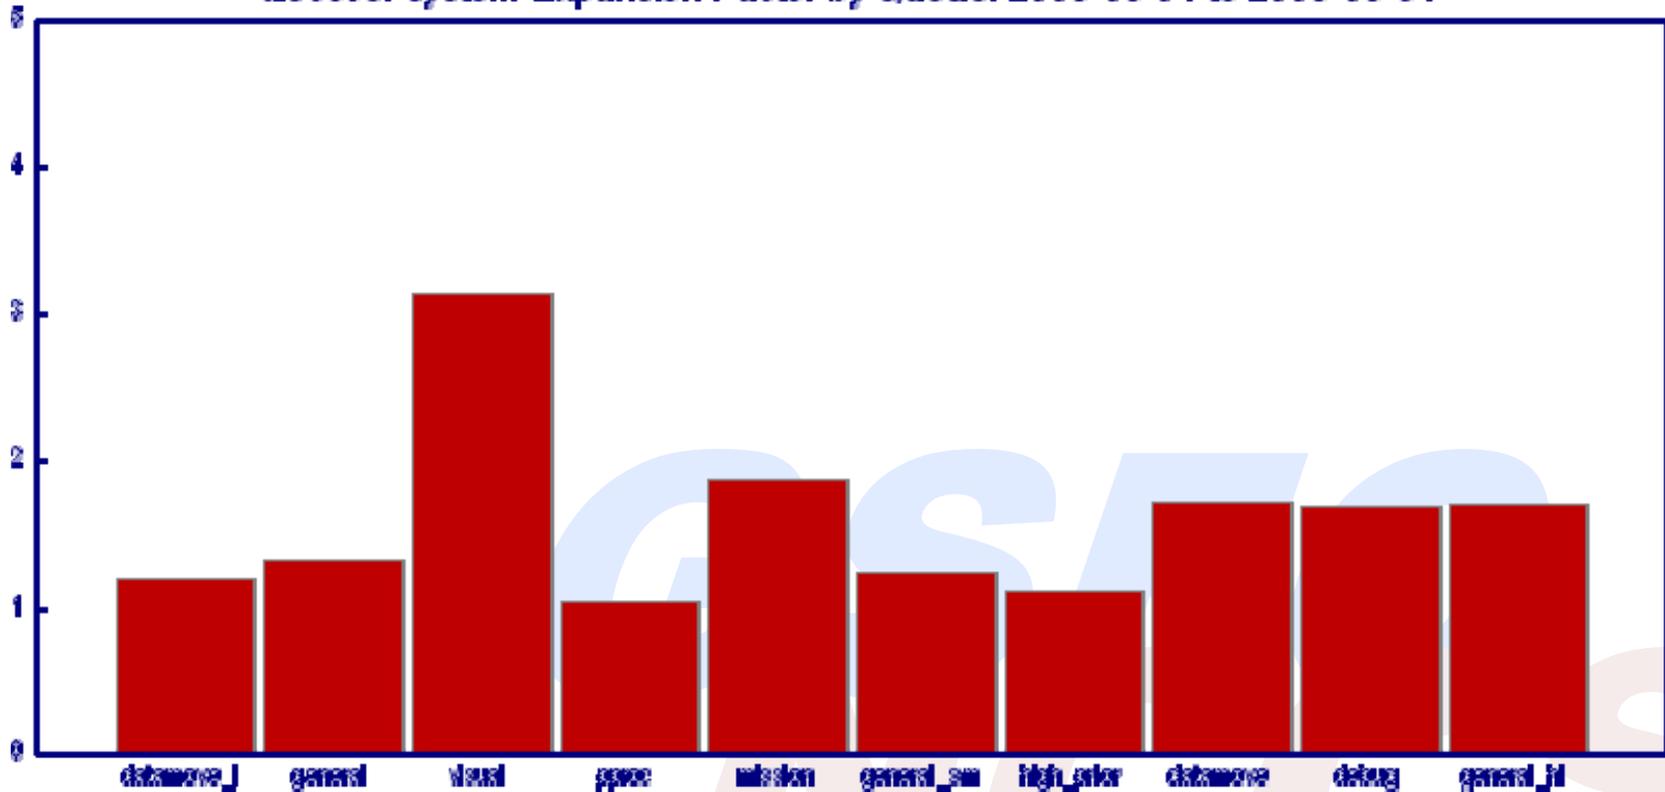


Discover Queue Expansion Factor

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

Weighted over all queues for all jobs
(Background and Test queues excluded)

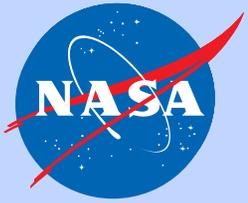
discover system Expansion Factor by Queue: 2008-08-01 to 2008-08-31





Current Issues on Discover: Infiniband Subnet Manager

- **Symptom:** Working nodes erroneously removed from GPFS following Infiniband Subnet problems with other nodes.
- **Outcome:** Job failures due to node removal
- **Status:** Modified several subnet manager configuration parameters on 9/17 based on IBM recommendations. Problem has not recurred; admins monitoring.



Current Issues on Discover: PBS Hangs

- **Symptom:** PBS server experiencing 3-minute hangs several times per day
- **Outcome:** PBS-related commands (qsub, qstat, etc.) hang
- **Status:** Identified periodic use of two communication ports also used for hardware management functions. Implemented work-around on 9/17 to prevent conflicting use of these ports. No further occurrences.



Current Issues on Discover: Problems with PBS -V Option

- **Symptom:** Jobs with large environments not starting
- **Outcome:** Jobs placed on hold by PBS
- **Status:** Investigating with Altair (vendor). In the interim, requested users not pass full environment via -V, instead use -v or define necessary variables within job scripts.



Current Issues on Discover: Problem with PBS and LDAP

- **Symptom:** Intermittent PBS failures while communicating with LDAP server.
- **Outcome:** Jobs rejected with bad UID error due to failed lookup
- **Status:** LDAP configuration changes to improve information caching and reduce queries to LDAP server; significantly reduced problem frequency. Still investigating with Altair.



Future Enhancements

- Discover Cluster
 - ▶ Hardware platform – SCU4 10/1/2008
 - ▶ Additional storage
- Discover PBS Select Changes
 - ▶ Syntax changes to streamline job resource requests
- Data Portal
 - ▶ Hardware platform
- DMF
 - ▶ Hardware platform

GSEFC
NCCS



Agenda

Welcome & Introduction
Phil Webster, CISTO Chief
Scott Wallace, CSC PM

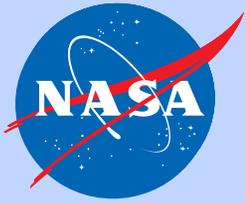
Current System Status
Fred Reitz, Operations Manager

Compute Capabilities at the NCCS
Dan Duffy, Lead Architect

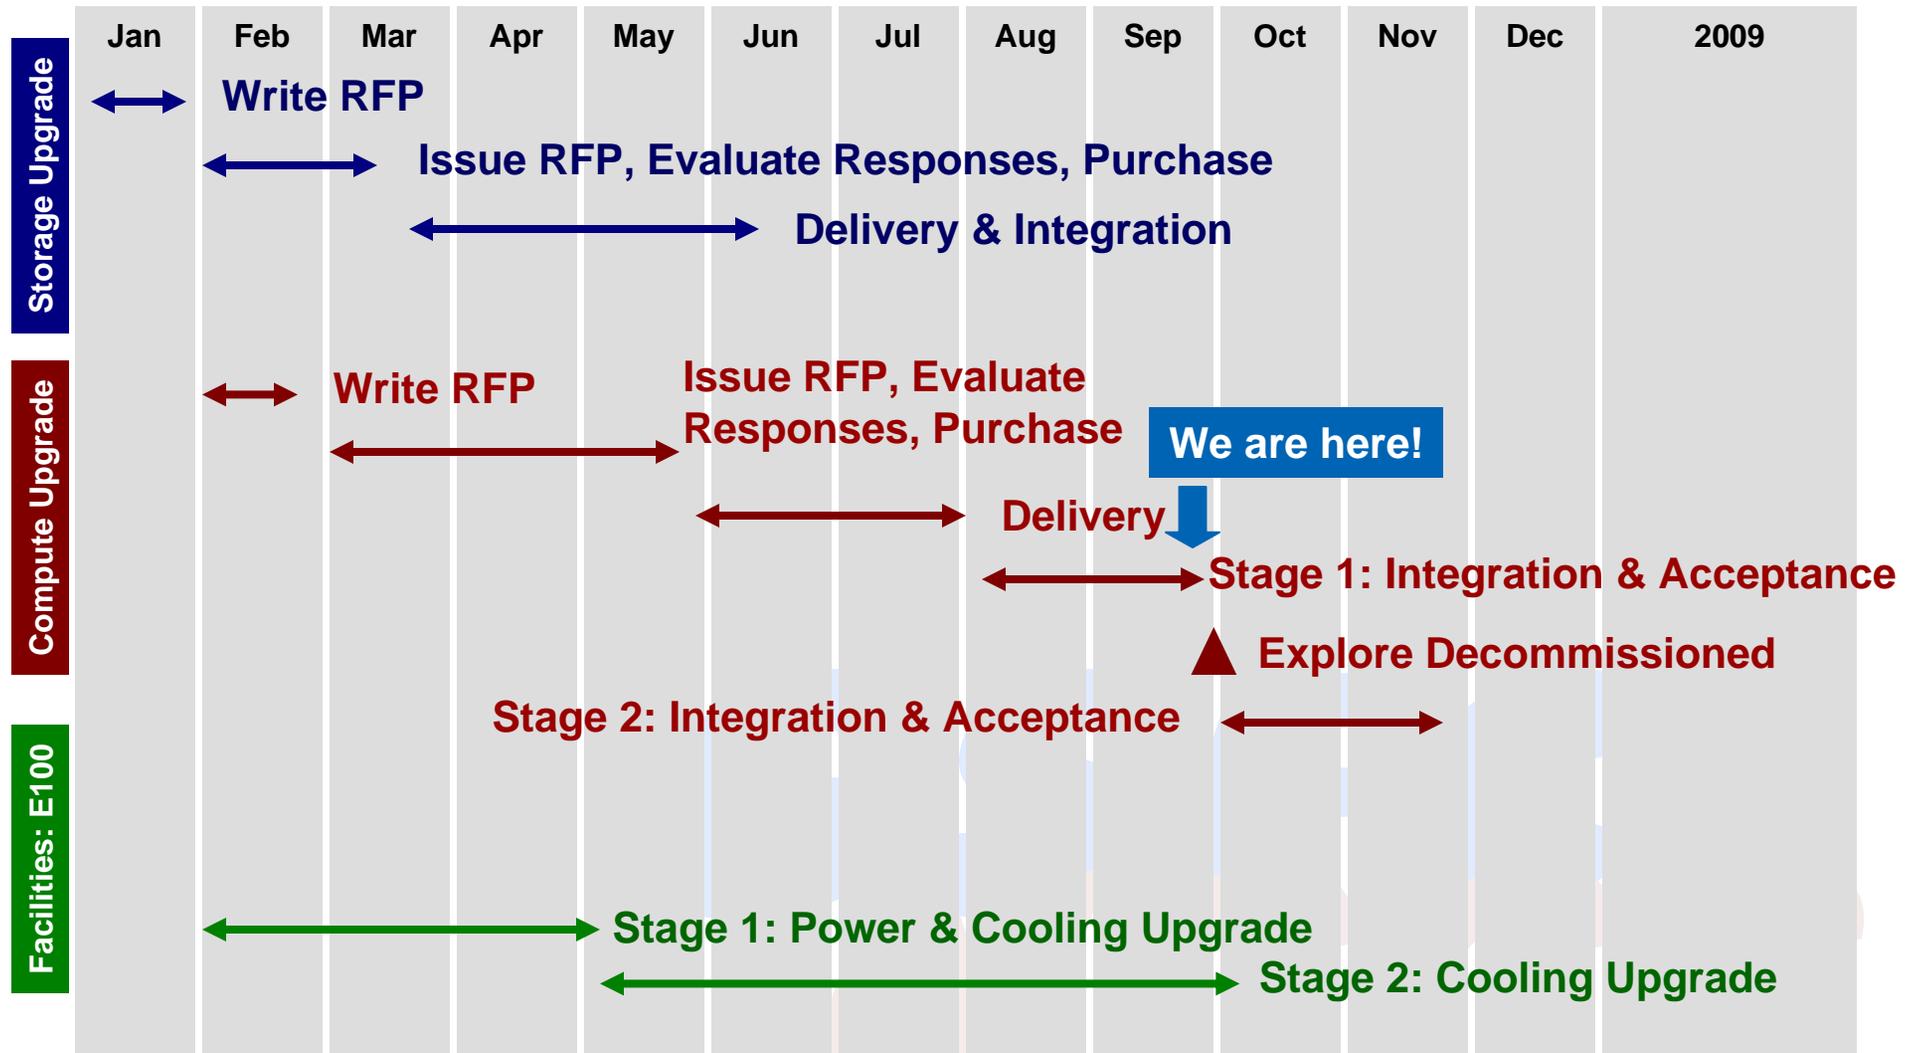
SMD Allocations
Sally Stemwedel, HEC Allocation Specialist

User Services Updates
Bill Ward, User Services Lead

Questions and Comments
Phil Webster, CISTO Chief

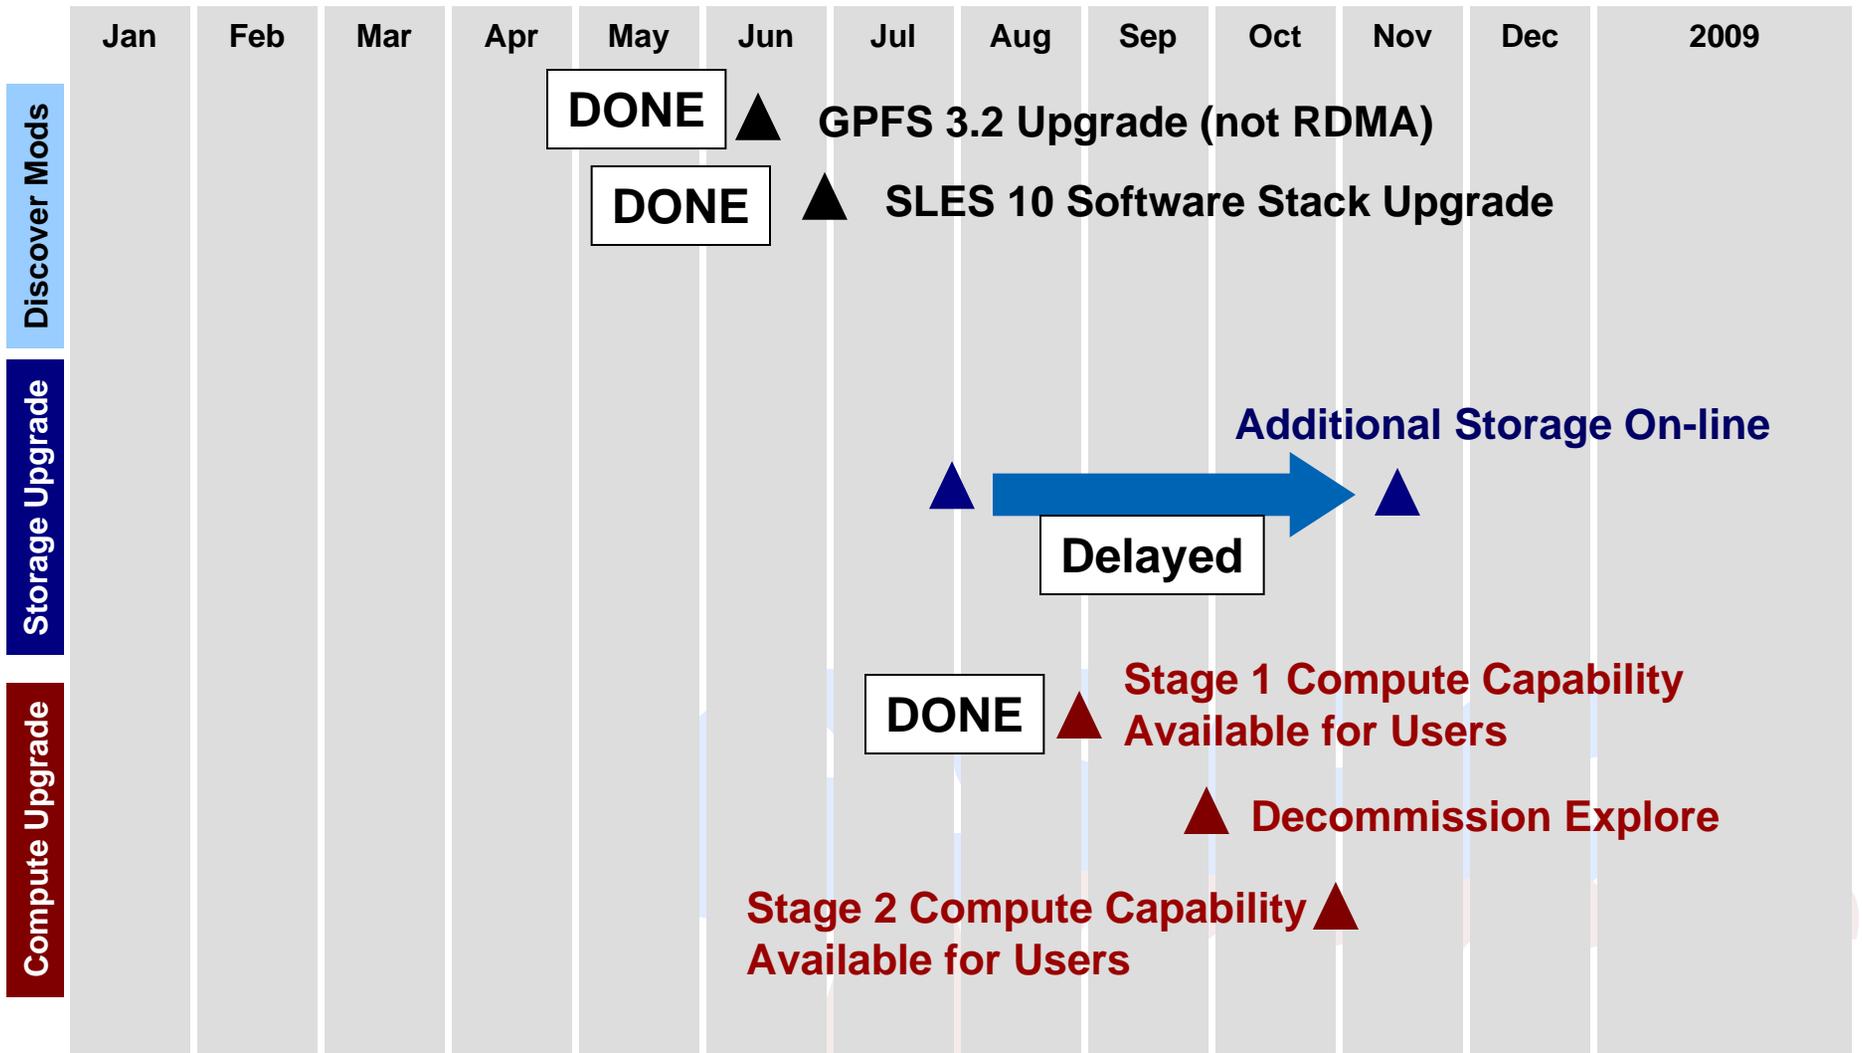


Overall Acquisition Planning Schedule 2008 - 2009





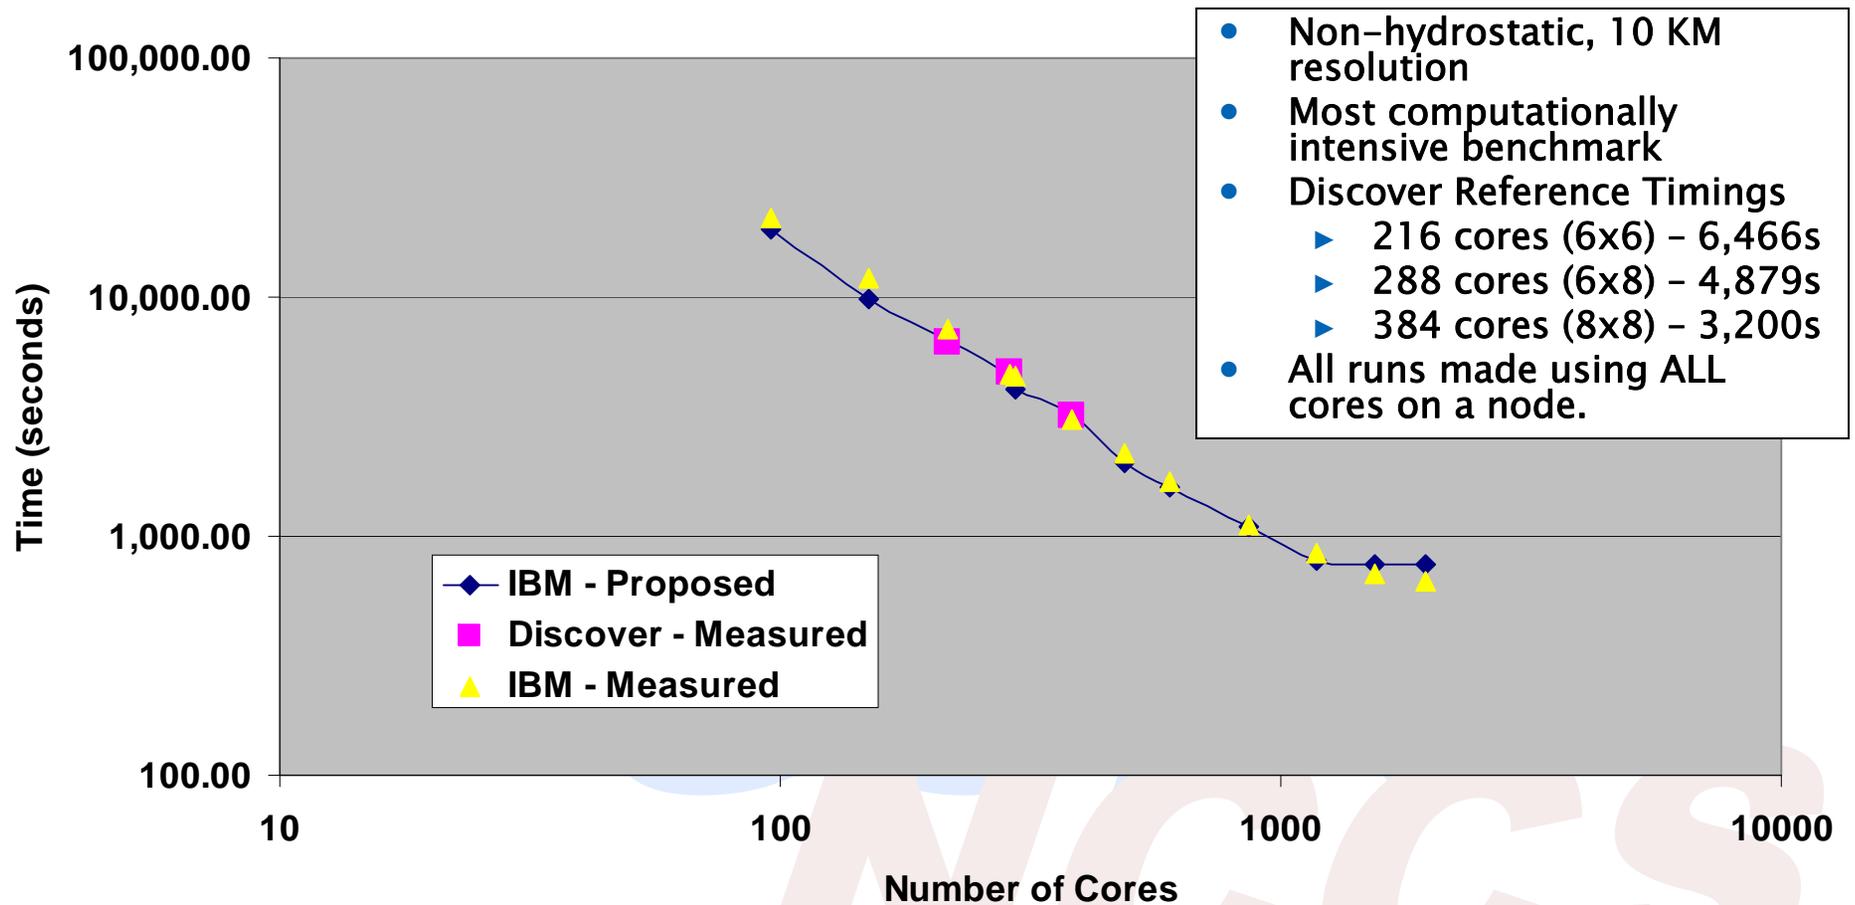
What does this schedule mean to you? Expect some outages – Please be patient





Cubed Sphere Finite Volume Dynamic Core Benchmark

Cubed Sphere - Benchmark 3





Near Future

- Additional storage to be added to the cluster
 - ▶ 240 TB RAW
 - ▶ By the end of the calendar year
 - ▶ RDMA pushed into next year
- Potentially one additional scalable unit
 - ▶ Same as the new IBM units
 - ▶ By the end of the calendar year
- Small IBM Cell application development testing environment
 - ▶ 2 to 3 months



Agenda

Welcome & Introduction
Phil Webster, CISTO Chief
Scott Wallace, CSC PM

Current System Status
Fred Reitz, Operations Manager

Compute Capabilities at the NCCS
Dan Duffy, Lead Architect

SMD Allocations
Sally Stemwedel, HEC Allocation Specialist

User Services Updates
Bill Ward, User Services Lead

Questions and Comments
Phil Webster, CISTO Chief



SMD Allocation Policy Revisions

8/1/08

- 1-year allocations only during regular spring and fall cycles
 - ▶ Fall e-Books deadline 9/20 for November 1 awards
 - ▶ Spring e-Books deadline 3/20 for May 1 awards
- Projects started off-cycle
 - ▶ Must have support of HQ Program Manager to start off-cycle
 - ▶ Will get limited allocation expiring at next regular cycle award date
- Increases over 10% of award or 100K processor-hours during award period need support of funding manager; email support@HEC.nasa.gov to request increase.
- Projects using up allocation faster than anticipated are encouraged to submit for next regular cycle.



Questions about Allocations?

- Allocation POC
Sally Stemwedel
HEC Allocation Specialist
sally.stemwedel@nasa.gov
(301) 286-5049
- SMD allocation procedure and
e-Books submission link
www.HEC.nasa.gov



Agenda

Welcome & Introduction
Phil Webster, CISTO Chief
Scott Wallace, CSC PM

Current System Status
Fred Reitz, Operations Manager

Compute Capabilities at the NCCS
Dan Duffy, Lead Architect

SMD Allocations
Sally Stemwedel, HEC Allocation Specialist

User Services Updates
Bill Ward, User Services Lead

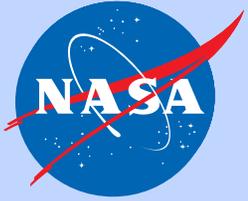
Questions and Comments
Phil Webster, CISTO Chief



Explore Will Be Decommissioned

- It is a leased system
- e1, e2, e3 must be returned to vendor
- Palm will remain
- Palm will be repurposed
- Users must move to Discover

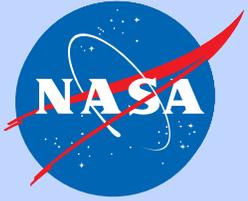
GSEFC
NCCCS



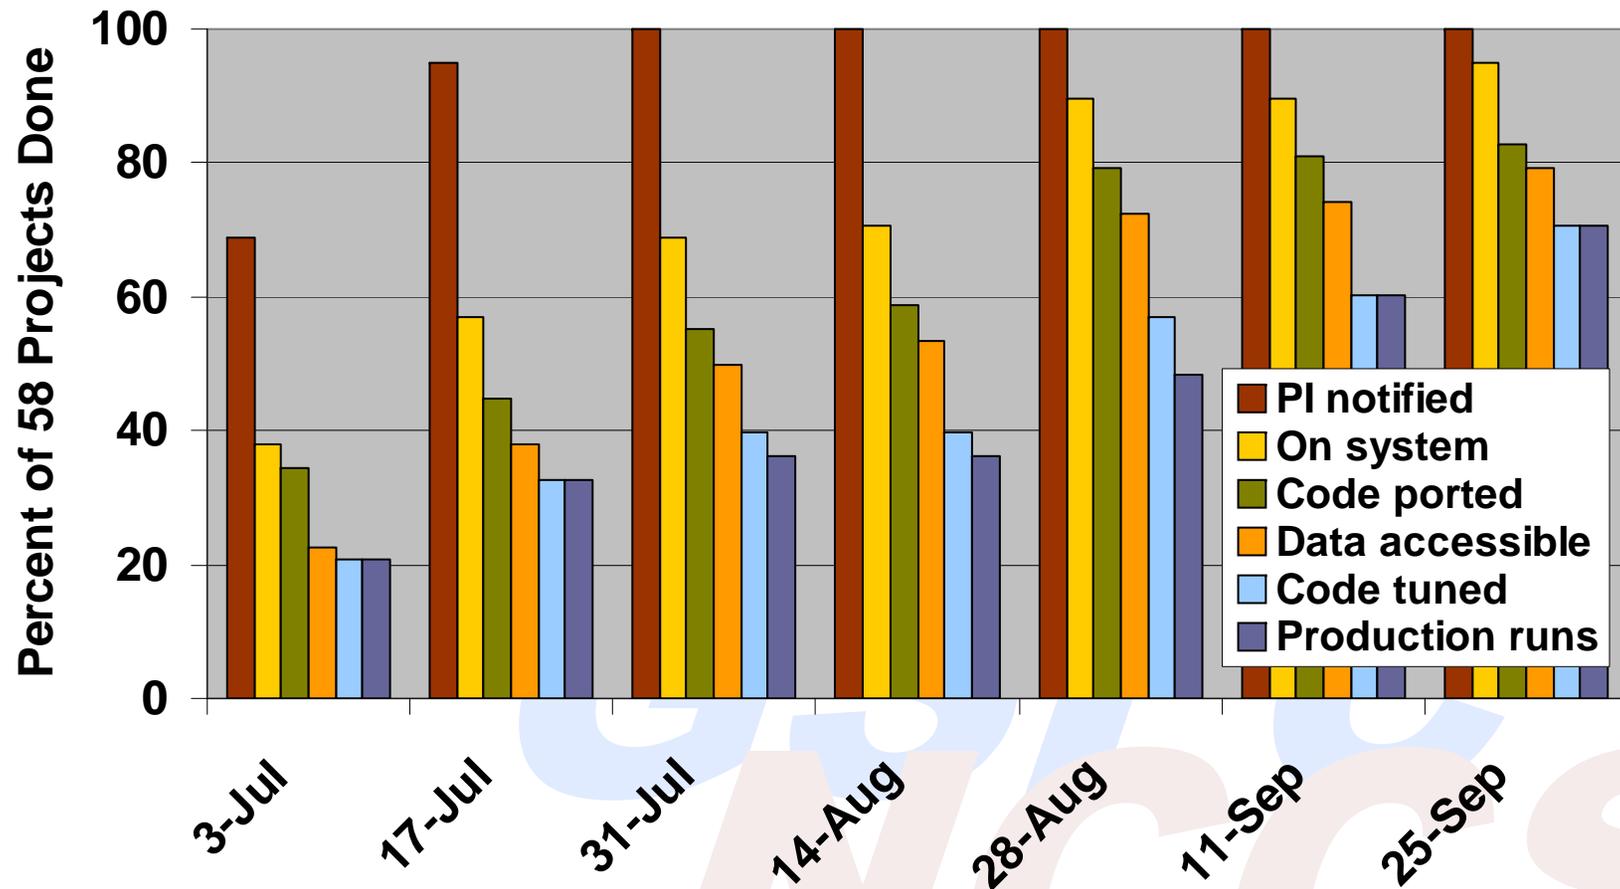
Transition to Discover – Phases

- PI notified
- Users on Discover
- Code migrated
- Data accessible
- Code acceptably tuned
- Performing production runs

GSFC
NCCS

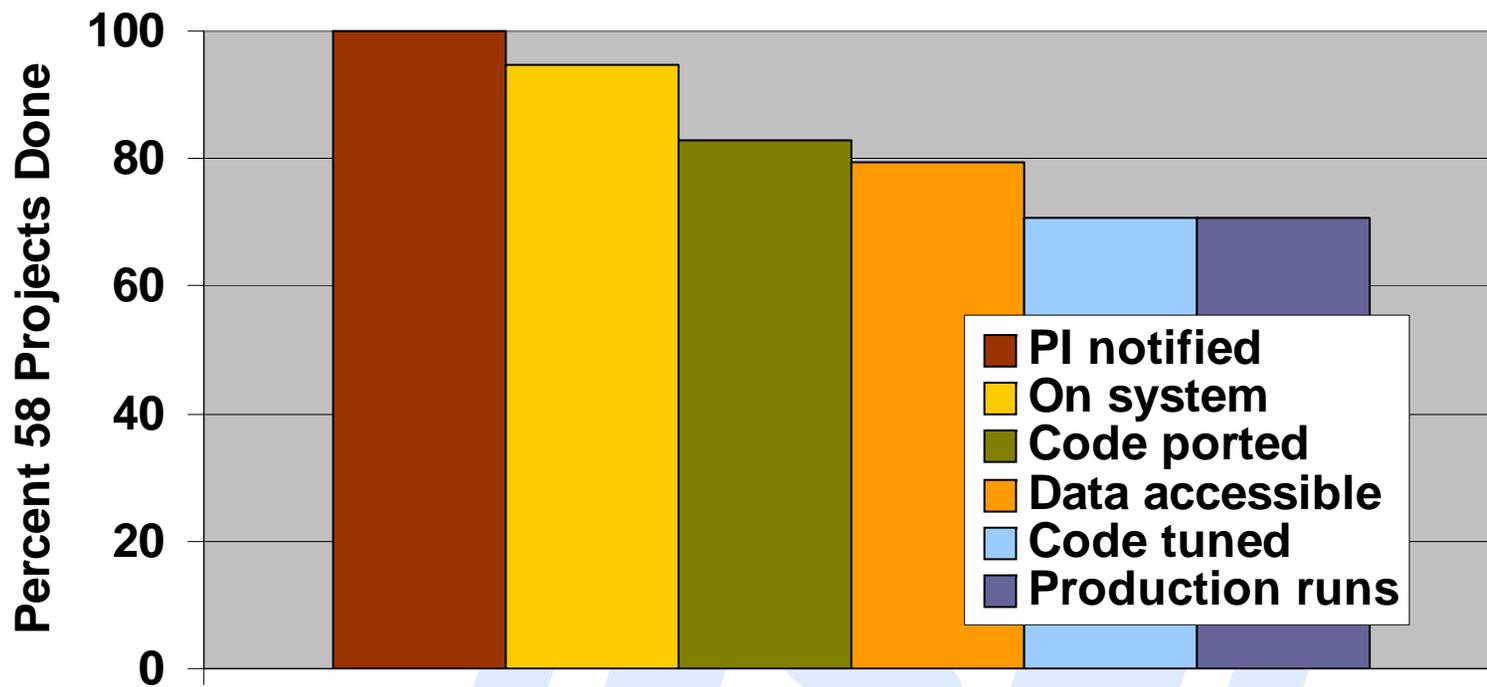


Transition to Discover – Status





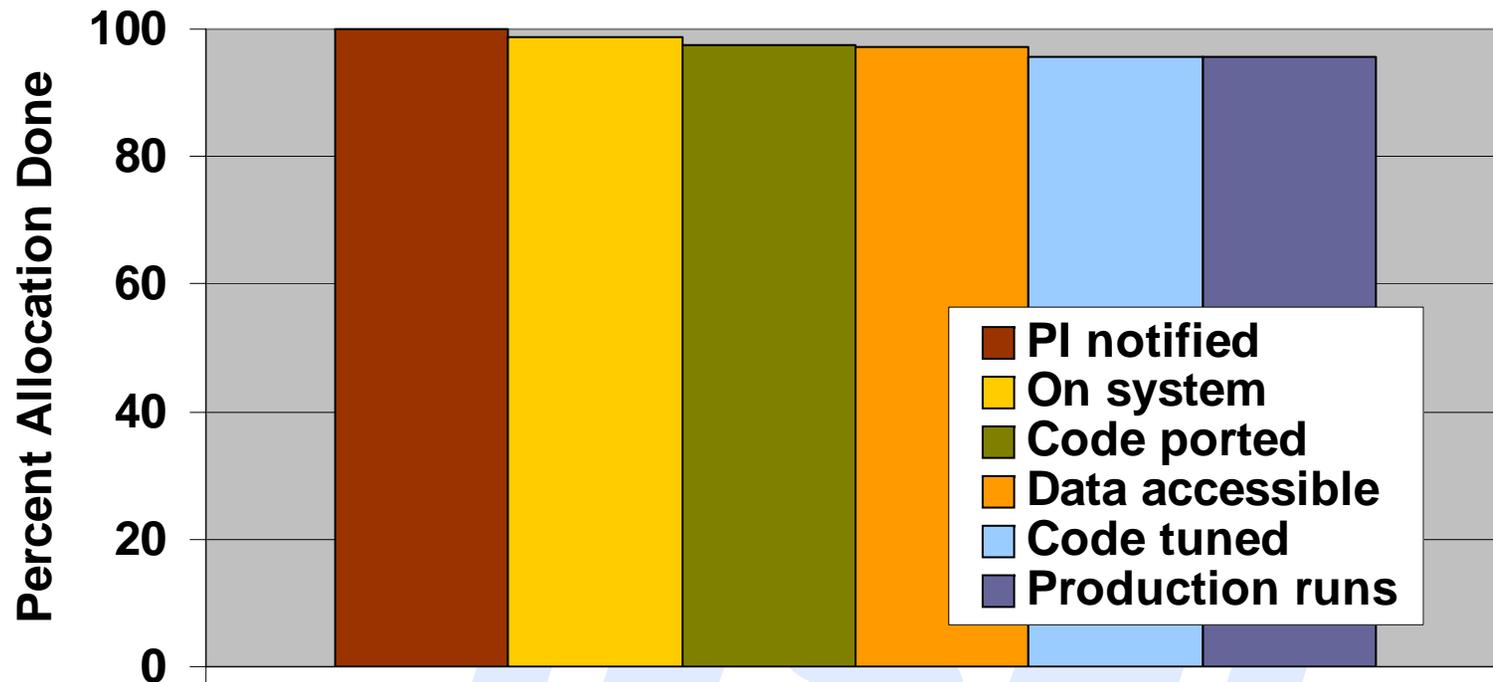
Transition to Discover – Status



25-Sep



Transition to Discover – Status



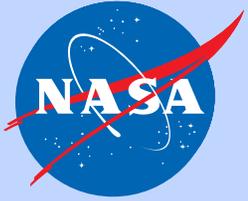
25-Sep



Accessing Discover Nodes

- We are in the process of making the PBS select statement more simple and streamlined
- Keep doing what you are doing until we publish something better
- For most folks, changes will not break what you are using now

GSFC
NCCS



Discover Compilers

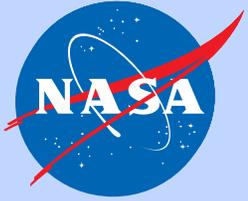
- `comp/intel-10.1.017`
preferred
- `comp/intel-9.1.052`
if intel-10 doesn't work
- `comp/intel-9.1.042`
only if absolutely necessary

GSFC
NCCS



MPI on Discover

- **mpi/scali-5**
 - ▶ Not supported on new nodes
 - ▶ `-l select=<n>:scali=true` to get it
- **mpi/impi-3.1.038**
 - ▶ Slower startup than OpenMPI
 - ▶ Catches up later (anecdotal)
 - ▶ Self-tuning feature still under evaluation
- **mpi/openmpi-1.2.5/intel-10**
 - ▶ Does not pass user environment
 - ▶ Faster startup due to built-in PBS support



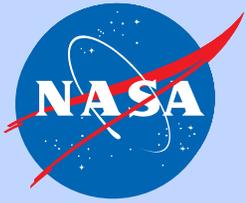
Data Migration Facility Transition

- DMF hosts Dirac/Newmintz (SGI Origin 3800s) to be replaced by parts of Palm (SGI Altix)
- Actual cutover Q1 CY09
- Impacts to Dirac users:
 - ▶ Source code must be recompiled
 - ▶ Some COTS must be relicensed
 - ▶ Other COTS must be rehosted



Accounts for Foreign Nationals

- Codes 240, 600, and 700 have established a well-defined process for creating NCCS accounts for foreign nationals
- Several candidate users have been navigated through the process
- Prospective users from designated countries must go to NASA HQ
- Process will be posted on the web very soon
<http://securitydivision.gsfc.nasa.gov/index.cfm?topic=visitors.foreign>



Feedback

- Now – Voice your ...
 - ▶ Praises?
 - ▶ Complaints?
- Later – NCCS Support
 - ▶ support@nccs.nasa.gov
 - ▶ (301) 286-9120
- Later – USG Lead (me!)
 - ▶ William.A.Ward@nasa.gov
 - ▶ (301) 286-2954



Agenda

Welcome & Introduction
Phil Webster, CISTO Chief
Scott Wallace, CSC PM

Current System Status
Fred Reitz, Operations Manager

Compute Capabilities at the NCCS
Dan Duffy, Lead Architect

SMD Allocations
Sally Stemwedel, HEC Allocation Specialist

User Services Updates
Bill Ward, User Services Lead

Questions and Comments
Phil Webster, CISTO Chief



Open Discussion Questions Comments

GSEFC
NCSS