

User Forum

NASA Center for Climate Simulation
High Performance Science

July 22, 2014



Agenda



- Introduction
- Hardware Updates & Procurements
- User Survey
- Archive
- Operations and User Services Updates
- Questions and Answers



Staff Additions



Welcome to New Members of the NCCS Team:

Jordan Robertson

George Britzolakis

Dan'l Pierce

Steve Ambrose

Welcome to Summer Interns:

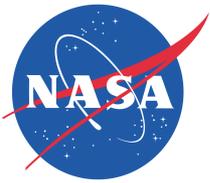
Mira Holford

Winston Zhou

Caitlin Ross

Joseph Clamp

Posters presented on Thursday, July 31st, B28 Atrium



Recent Accomplishments Systems and Operations



- Hosted full-day Alinea workshop (MAP, DDT) (Mar 2014)
- Integration Efforts
 - Nature Run Storage on Discover: 7,200 TB RAW disk (Nov 2013)
 - JIBB Upgrades: ~40 TF SandyBridge and ~400 TB RAW disk (Feb-Apr 2014)
 - ESGF data node on new SandyBridge node with 10 Gbps (Feb 2014)
 - “Authorization to Operate” (ATO) completed and signed for 3 more years (Apr/May 2014)
 - Migration out of 9 legacy Tape Libraries (June 2014)
- Discover Cluster Efforts
 - SLURM migration (October 2013)
 - IB Fabric congestion reduction – cable replacements and configuration changes
- Archive Growth and Policy Recommendations Study (June 2014)
- Pre-ABoVE on proof-of-concept NCCS High Performance Science Cloud (ongoing)



Recent Accomplishments Campaigns and Special Support



- Field Campaigns

- DISCOVER-AQ Fall 2013
- HS3 Summer/Fall 2013
- ATTREX (Guam) Winter 2014
- IPHEX 2014 (Smokey Mountains) May/June 2014
- DISCOVER-AQ FRAPPE (Colorado) Ongoing 2014



View from the NASA ER-2 during an IPHEX 2014 flight, May 24, 2014 (image credit: NASA)

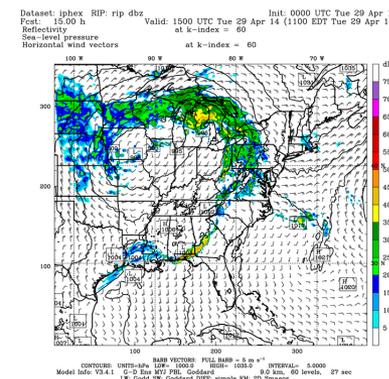


- Upcoming Field Campaigns

- ARISE and HS3 2014

- Other Special Support:

- SMAP Level 4 Root Zone and Carbon product generation support
- DSCOVR EPIC processing (ongoing)
- GEOS-5 two-year, 7-km Nature Run
- MERRA2
- ABoVE



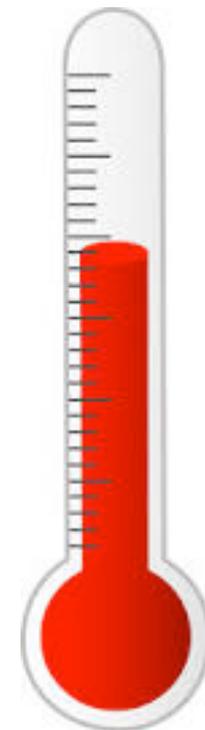
NU-WRF's outer (9-km) domain forecast for 1100 EDT April 29, 2014, depicting simulated radar reflectivity and sea level pressure and wind vectors. When compared with operational models for this forecast, NU-WRF better simulated diminished precipitation over the IPHEX 2014 study region.



GSFC-Wide Chilled Water Outage (Cooling for NCCS Hardware) July 2014



- Center-wide chilled water outage July 8 (began 19:41) due to lightning strike in Building 24 that affected the West Campus pumps
 - NCCS Facilities team arrived on site shortly after to assess the situation
 - Upon realization that the chilled water would be out for an indefinite amount of time, the operations team began bringing down all HPC systems
 - Users were notified as quickly as possible
 - Room temperatures rose rapidly and exceeded 120 F within a short time period before the systems were shut down
- FMD addressed power issues and started pumps
 - The pumps were started back up several hours after the event
 - Took several hours for the water to reach normal operating temperatures
 - Took several hours for the rooms to reach normal operating temperatures
- Operations team began restoring service early July 9th
 - Discover available July 9th at 17:10 (without SCU8)
 - Archive available July 11th at 19:00 (after significant disk rebuilds)
- NCCS lessons learned held on July 15th





Hardware Updates and Procurements

Dan Duffy,
HPC Lead and NCCS Lead Architect

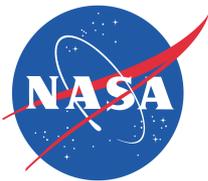


FY14-FY15 Cluster Upgrade



- Combined funding from FY14 and FY15
 - Taking advantage of new Intel processors – double the floating point operations over SandyBridge
 - Decommission SCU7 (Westmeres)
- Scalable Unit 10
 - Target to effectively double the NCCS compute capability
 - 128 GB of RAM per node with FDR IB (56 Gbps) or greater
 - Benchmarks used in procurement include GEOS5 and WRF
- Target delivery date ~Oct 2014





FY14 NCCS Wide File System



- Augment storage along with the cluster upgrade
 - Targeting about 10 PB or more (depends on cost)
- Creation of an NCCS wide file system
 - Separate from GPFS
 - Available even when there are issues with GPFS
 - Possible NFS solution (exploring options)
 - Many applications will benefit from client side caching
 - Move home directories and other file systems into this storage solution
 - Accessible by all Discover nodes (including compute) and Archive
 - Will provide data to portal services (just like GPFS)
- Procurement
 - To be released early August
 - Target installation late fall 2014



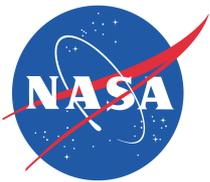


Archive Upgrades



- Increased DMF License Capacity (45 PB)
- Tape Storage Area Network (SAN)
 - Upgraded switch capacity and speeds (16 Gbps)
- 20 New Tape Drives
 - Capable of 8 TB per tape
 - To be installed in August 2014
- Migration of Tapes to new Drives (constant)
- Archive capacity planning study – more on this later





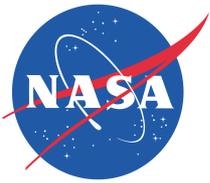
Nature Run Storage – Installed Late Fall 2013



- Integrated 7,200 TB RAW disk capacity for the GMAO Nature Run
- 2-year Nature Run at 7.5 KM resolution
 - Completed
- 3-month Nature Run at 3.5 KM resolution
 - Just starting
- Will generate about 4 PB of data (compressed)
- All data to be publically accessible
- <ftp://G5NR@dataportal.nccs.nasa.gov/>



COURTESY: NSTRPP

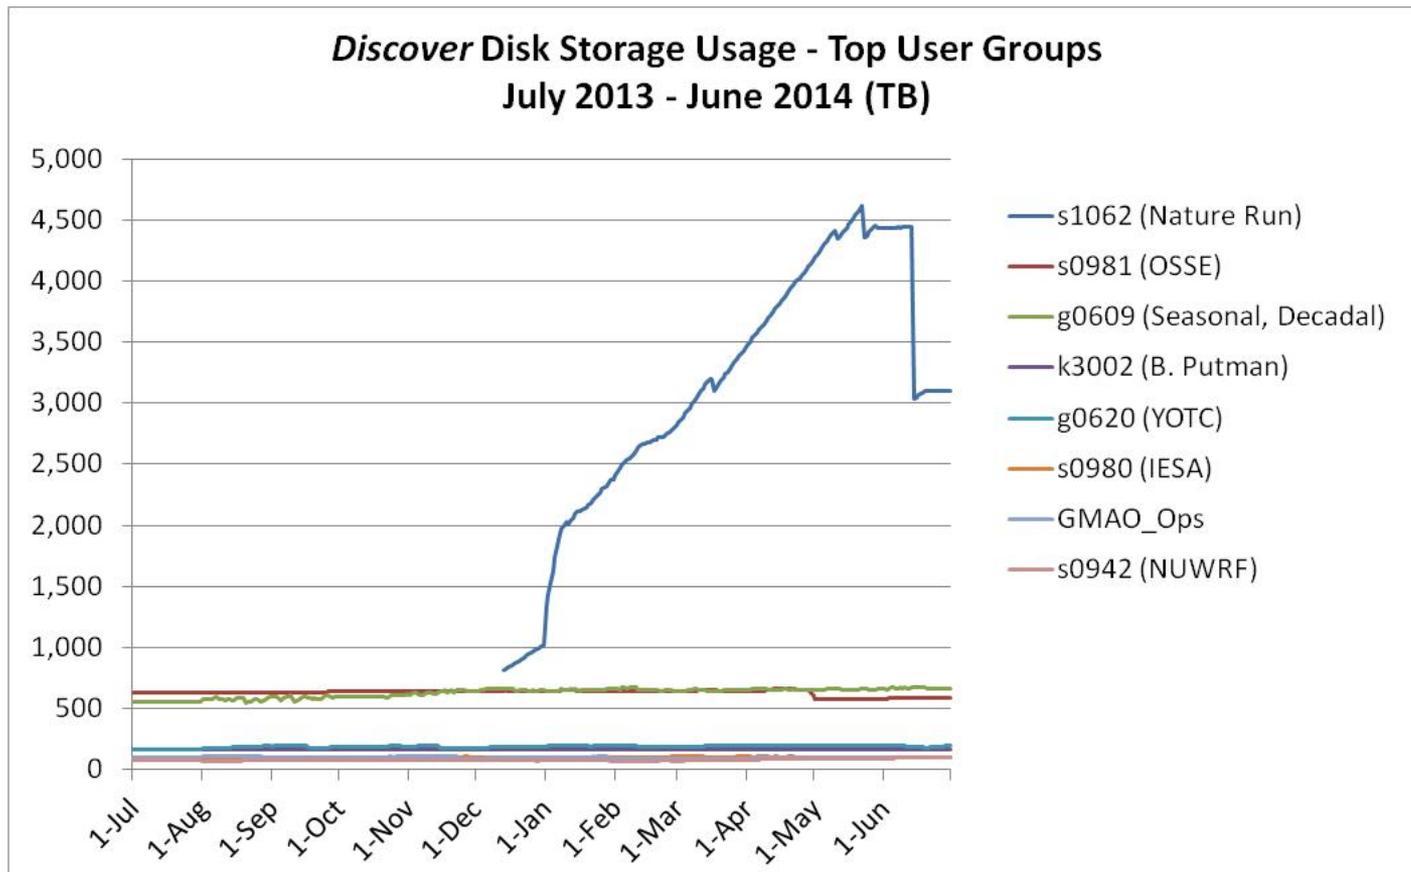


Discover Storage Breakdown

June Operational Analysis



- Nature Run storage (s1062) on new filesystem (dnb03)
 - Was rapidly growing, leveled off, then clean-up after completion of run (another to start soon)





Hyperwall Monitors Installed June 2014



- Upgraded 4-year old monitors
 - 15 high resolution monitors
 - New mounting mechanism
- Next Steps
 - Content being updated for HD
 - Servers to be update in 2015
- Please feel free to request scheduling of the wall for:
 - Presentations
 - Tours
 - Family
 - School groups



Lori Perkins (Science Visualization Studio) describes a visualization of aerosols simulated by GEOS-5 and displayed on the new Visualization Wall in the NCCS's Data Exploration Theater. (Photo credit: Jarrett Cohen, CISTO/GST. Aerosol image provided by Bill Putman, Global Modeling and Assimilation Office, GSFC Code 610.1)

To schedule the wall, contact:

Heidi Dewan
heidi.dewan@nasa.gov
301-286-9426

NCCS User Services:
support@nccs.nasa.gov
301-286-9120



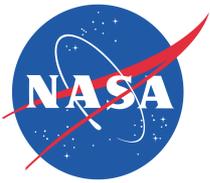
JIBB Upgrade – Early 2014



- Doubled the Compute Capacity to ~77 TF Peak
 - Additional 120 Compute Nodes
 - 1,920 cores; 39 TF
 - 2.6 GHz Intel Sandybridge with 64 GB of RAM
 - Fourteen Data Rate Infiniband Network (56 Gpbs) in a 2-to-1 blocking fabric
- Doubled Storage Capacity to ~800 TB
- 2 New Login Nodes
- Nature Run mounted on login nodes
 - Exploring options to extend the nature run to the compute nodes



Upgraded to approximately double the computational and storage capacity. Received funding through NOAA from the Hurricane Sandy Relief bill.



NCCS User Survey Results & Responses

Al Settell,
CSC Program Manager, CISTO-SCTS



Comparison to 2012 Survey

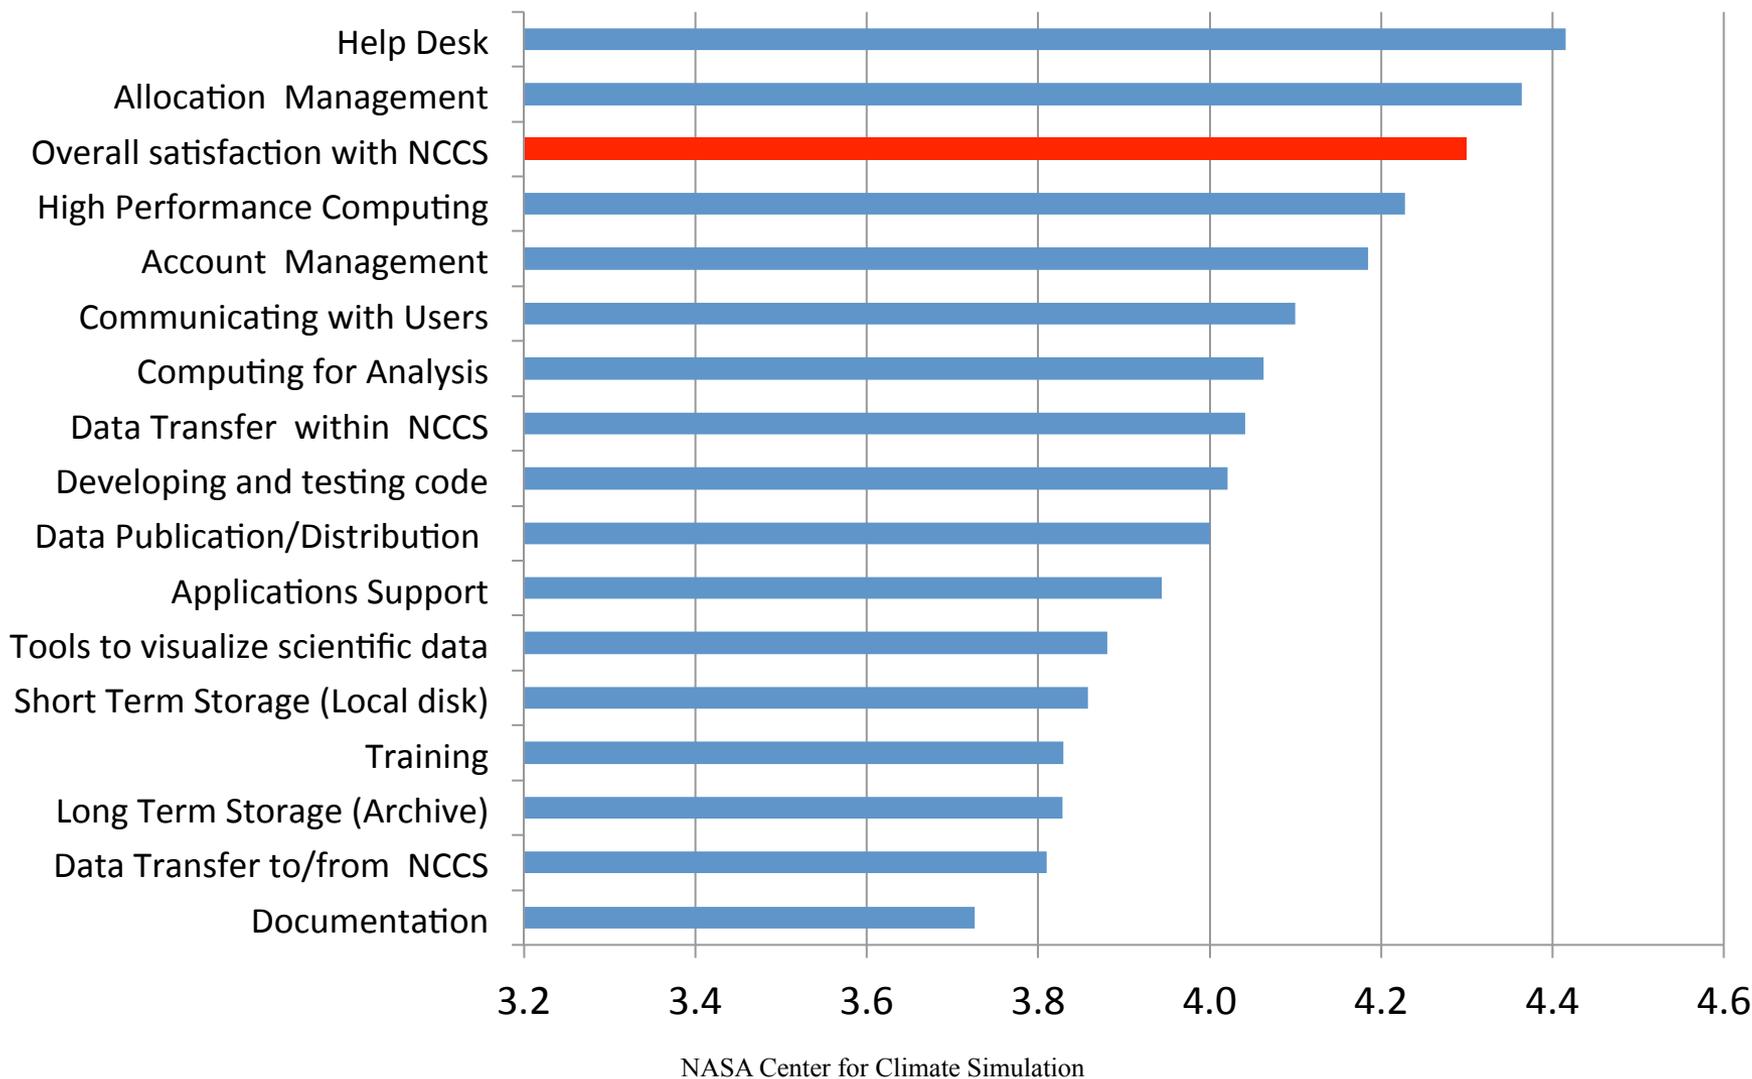


Area	2013	2012
Overall satisfaction with NCCS	4.30	3.93
High Performance Computing	4.23	4.07
Computing for Analysis	4.06	3.8
Long Term Storage (Archive)	3.83	3.63
Short Term Storage (Local disk)	3.86	3.58
Data Transfer to/from NCCS	3.81	3.32
Data Transfer within NCCS	4.04	
Data Publication/Distribution	4.00	
Help Desk	4.42	4.22
Account Management	4.18	3.93
Allocation Management	4.36	4.00
Applications Support	3.94	3.79
Documentation	3.73	3.38
Training	3.83	3.41
Communicating with Users	4.10	3.92
Tools to visualize scientific data	3.88	3.52
Developing and testing code	4.02	3.72

NCCS User Survey 2013

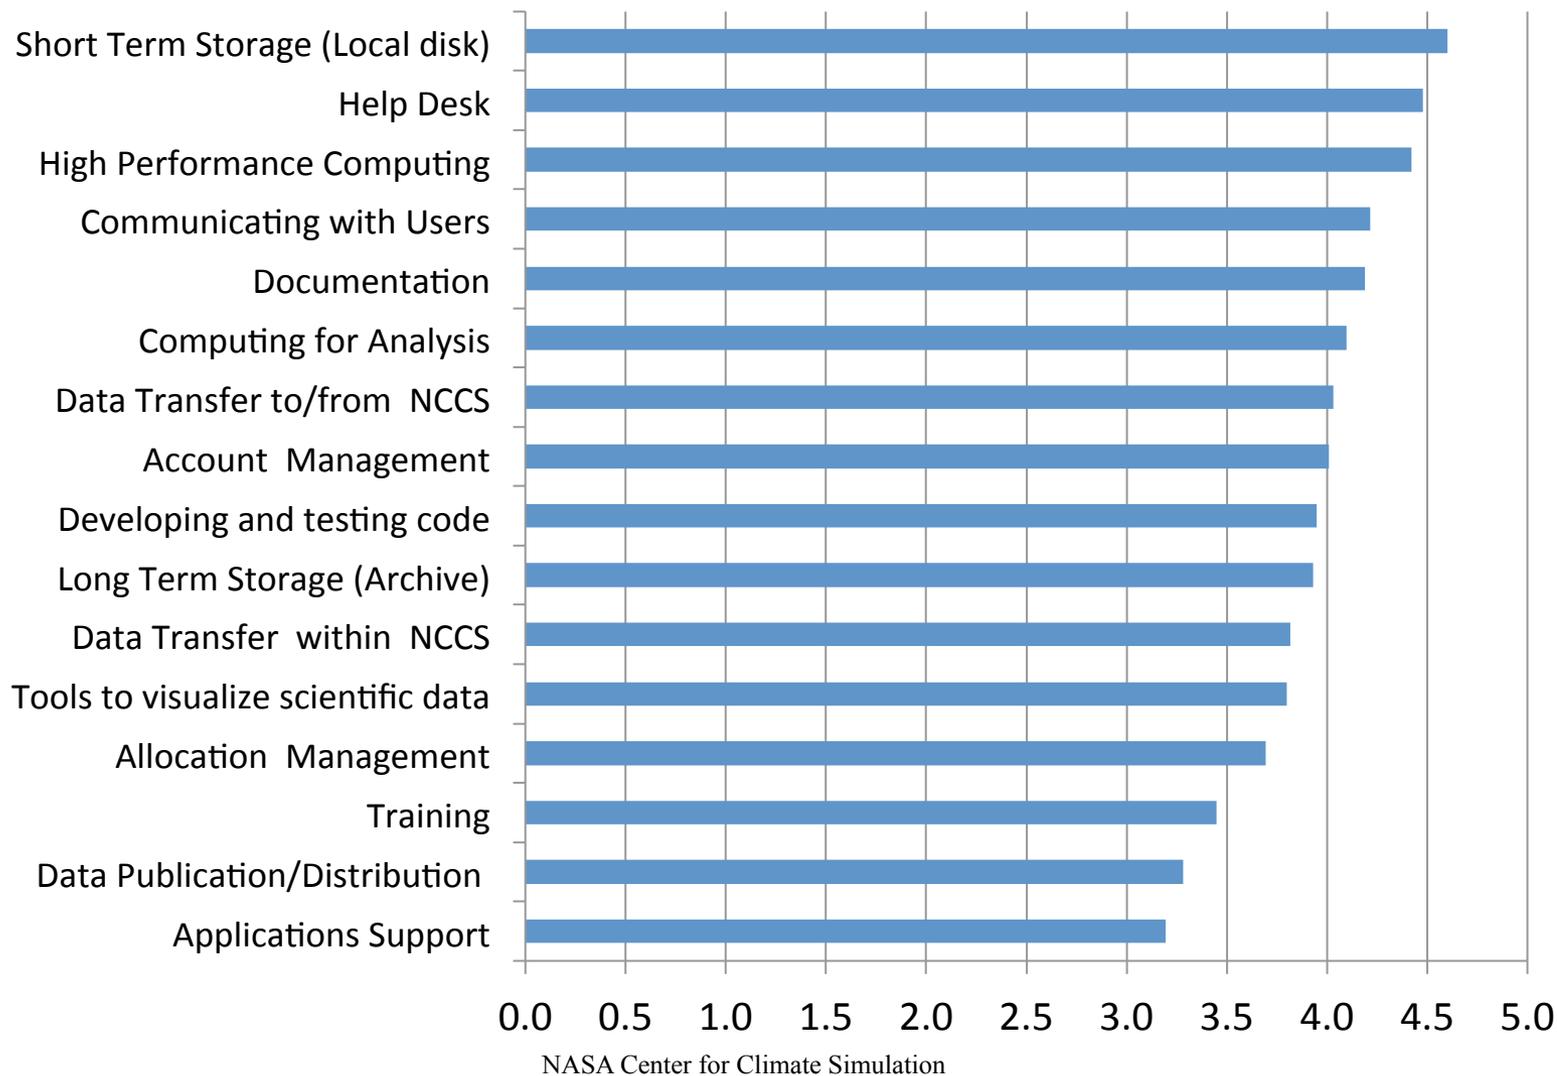


Results by Service Area - Performance





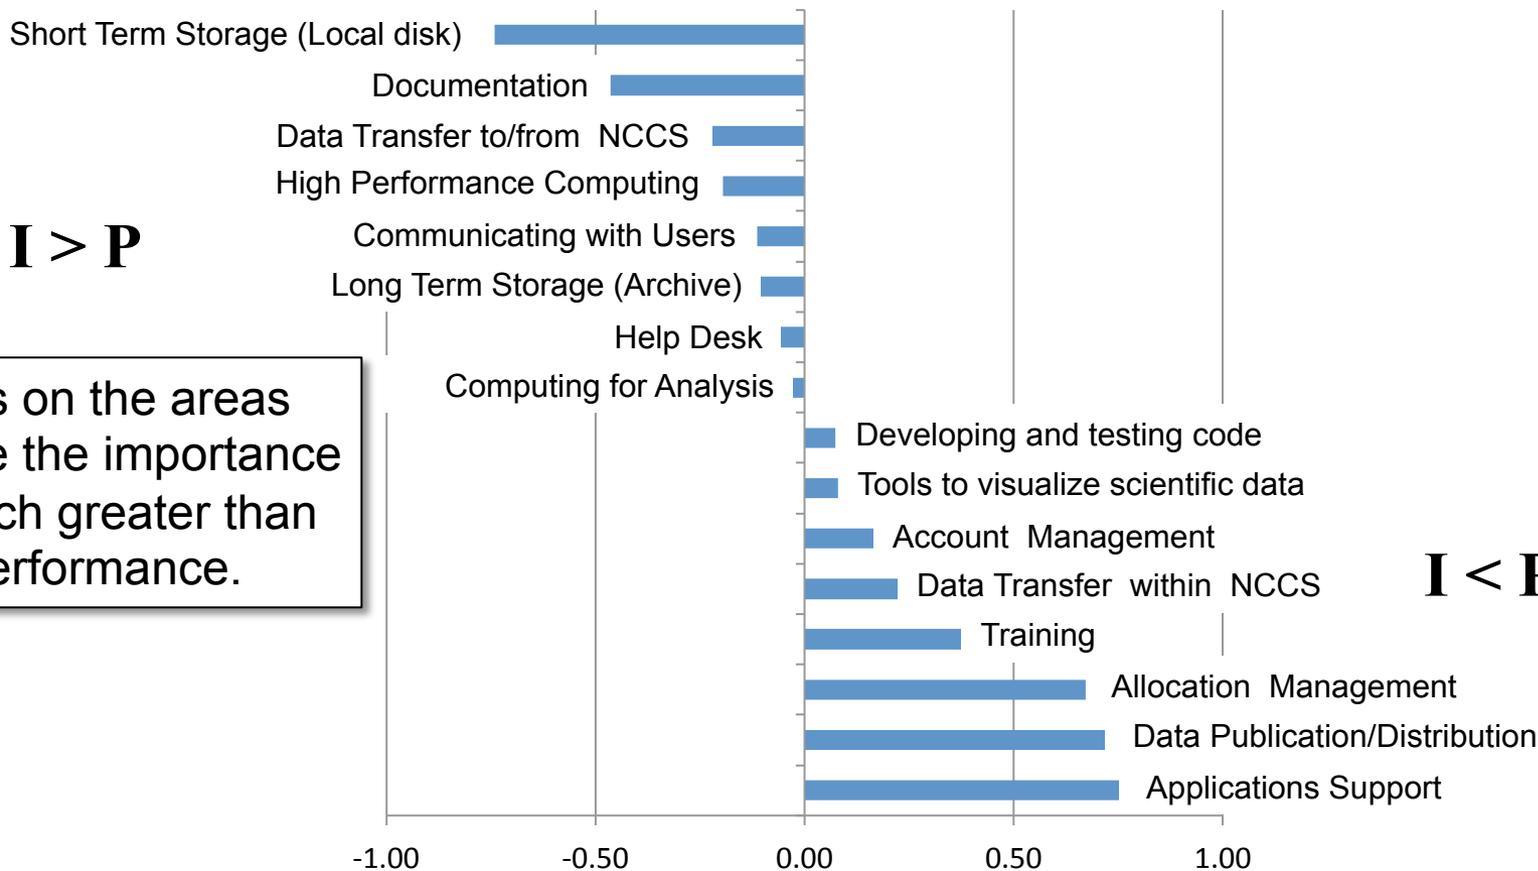
Results by Service Area - Importance



NCCS User Survey 2013



Results by Service Area - Performance (P) Minus Importance (I)



I > P

Focus on the areas where the importance is much greater than the performance.

I < P

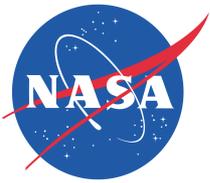
NCCS User Survey 2013



Themes – Based on Scores and User Comments



- Communications
 - Improved documentation/support, e.g., more examples in primer
 - User Notification improvement (more timely and consistent notifications)
 - Ticketing system improvements
- Discover
 - Longer running jobs
 - More scratch space
 - Process improvements, e.g., quicker response to requests for increased disk
- Archive
 - Improved reliability and data restore timeliness
 - Performance
- New Services
 - Remote visualization
 - Remote GUI-interactive improvements
 - Expanded licensing



Action Plan

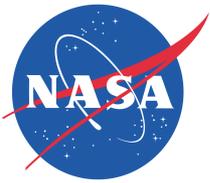


- **Communications**
 - Created Communications and Marketing Plan
 - Website and virtual presence improvements
 - Business process improvements for notifications
- **Discover**
 - Longer running jobs via SLURM Quality of Service (QoS)
 - NCCS center-wide file system
 - Business process improvements for disk requests
- **Archive**
 - Archive Study and Planning Improvements (ongoing)
 - Storage Area Network (SAN) and Tape Drive Upgrades
 - More is coming
- **New Services**
 - Remote visualization servers and software being delivered in near future
 - Explore remote desktop capabilities to improve GUI interactive response on Discover
 - Tracking license usage and “denials” of license for better capacity planning



NCCS Archive

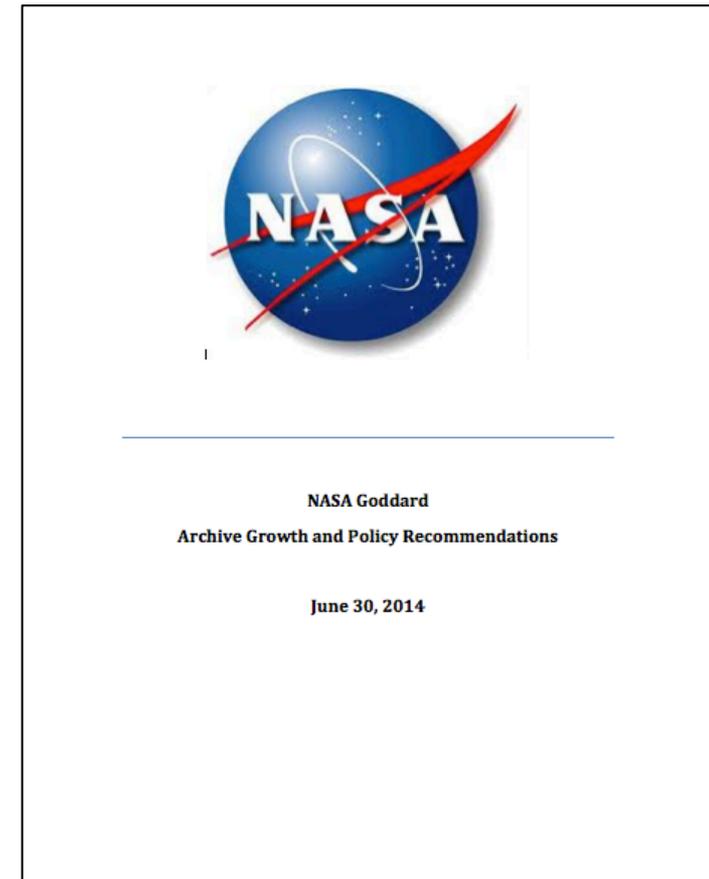
Tom Schardt



Archive Capacity Planning Study

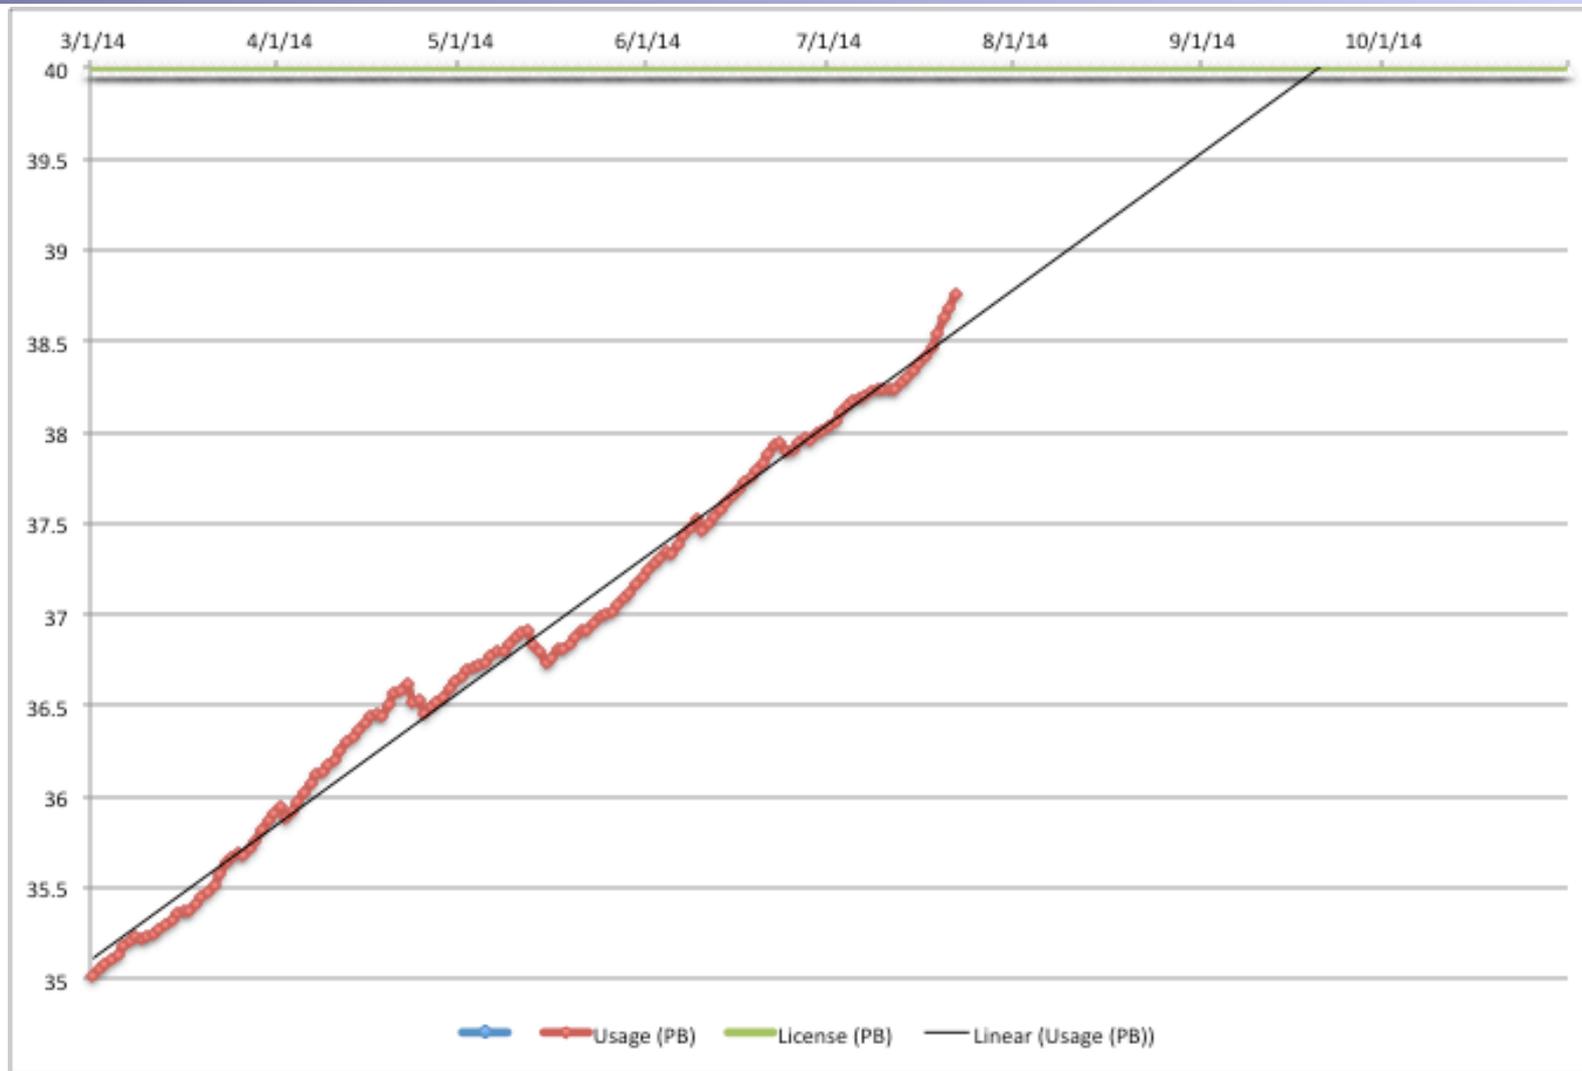


- Archive capacity planning study was completed in June 2014
 - Person from outside the NCCS was commissioned for the study
- The study took into account
 - Current architecture
 - Growth projections
 - Options for performance improvements
 - Specific and general suggestions
 - Projected growth and budget forecasts





Projected Growth





Noted Areas of Concern



- Thrashing of archive file systems (using archive as scratch)
- Data does not remain resident very long on the disk cache
- The large number of small files in DMF cause problems
- Large amounts of files/data are stored and never recalled
- Constant migration of data to newer tape media puts a load on the system above and beyond the users
- Tape libraries are almost full; new libraries are very expensive and take up large amounts of space in the computer room
- Overall cost to maintain growth



Areas Under Consideration Based on the Study



- Perform a full analysis of the archive solution, including the following
 - Policies
 - Architecture
 - Budget
 - Performance Improvements
 - Hardware
 - Operations
 - Functionality
 - User Advisory Group
- Identify improvements, prioritize, and implement
 - Not a lengthy process



Capping the Growth of the Archive



- Quotas are needed to control the growth of the archive and therefore maintain budgetary constraints
- Additional policies under consideration include
 - Data expiration
 - Other (TBD)
- These are under preliminary evaluation
 - Communication and coordination with the users is critical to the successful implementation of any policy
- User Advisory Group
 - The NCCS is looking for users who would like to take part in an advisory group on archive changes
 - This group would be a start on an overall NCCS Advisory Group for all services
 - If you are interested, please let us know



NCCS Operations & User Services Update

Ellen Salmon



Building 33 “Offices Hours” for NCCS Technical User Services Staff



- An NCCS representative typically holds Wednesday office hours in building 33 room C116
- Purpose is to provide face-to-face technical user support to assist in
 - Troubleshooting (e.g., steps to minimize swapping)
 - Optimizing code
 - Optimizing use of NCCS resources
 - Facilitating NCCS responses to user requests
- Schedule for next four weeks:
 - 7/23: George Britzolakis
 - 7/30: Hamid Oloso
 - 8/6: Denis Nadeau
 - 8/13: Eric Winter
- Feel free to stop by with questions and problems



New Batch Job Capabilities via “Native” SLURM



- New capabilities coming, via “native” SLURM
- For example, Quality of Service (qos), which can enable many features, e.g.:
 - Longer job wall-time
 - *Users must request to be enabled (email support)*
 - More to come
- These advanced features are available via “native” SLURM and will not be “back-ported” to the PBS wrapper.
- Watch for the Brown Bag Seminar on July 31 about converting PBS scripts to “native” SLURM.



Upcoming Brown Bag Seminar



How to Convert Your Discover Job Scripts from PBS to SLURM

- July 31, 2014, 12 noon, Bldg. 33, H118
- Review of issues & techniques involved when migrating PBS job scripts to “native” SLURM scripts.
- “Native” SLURM scripts allow use of advanced features like Quality of Service (qos).



Miscellaneous



- Discover to Data Portal 10 GbE Connections upgrade
 - Formerly 4 by 1 GbE and now 2 by 10 GbE (5x improvement)
- Dali
 - Monitoring Page
 - Idea to create a similar page as the NCCS job monitor for Dali nodes
 - Will assess the feasibility and potential implementation
 - Load Balance/Round Robin
 - No load balance currently; round robin login across the different Dali nodes
 - Will assess the feasibility of load balancing
- Tour of the NCCS (individuals, groups, school groups, family)
 - Please schedule through Heidi Dewan and/or send an email to support@nccs.nasa.gov



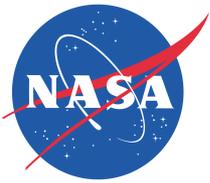
Questions & Answers

NCCS User Services:

support@nccs.nasa.gov

301-286-9120

<https://www.nccs.nasa.gov>



Contact Information



NCCS User Services:

support@nccs.nasa.gov

301-286-9120

<https://www.nccs.nasa.gov>

http://twitter.com/NASA_NCCS

Thank you