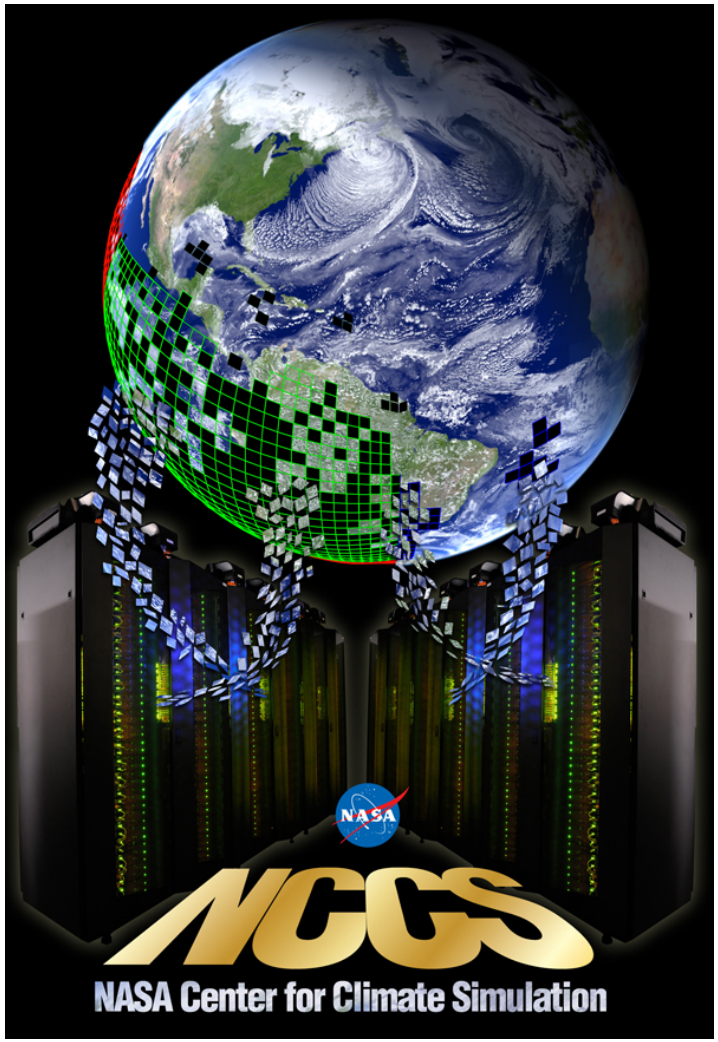

NCCS User Forum

September 24, 2013

Agenda

- Introduction
- Special Topics
- NCCS Updates
- NCCS User Survey
- NCCS Operations and User Services Updates
- Questions and Answers
- Breakout Session: SLURM

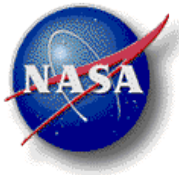

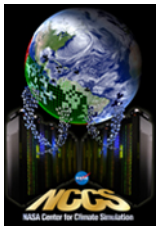

Introduction

Dan Duffy,
HPC Lead and NCCS Lead Architect

Requests for Computer Time

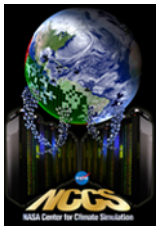

- For NASA SMD-supported research
- Computer time requests were due Sep. 20th
- One-year allocations beginning November 1, 2013
- Principal Investigators: see https://hec.reisys.com/hec/computing/index.do

- For more information on the request process, see:
  - Eligibility for Getting Computing Time
    http://hpc.nasa.gov/request/request.html
  - NASA High-End Computing (HEC) Program Systems & Services
    http://www.hec.nasa.gov/user/systems.html

Staff Additions

Welcome to New Members of the NCCS Team:

Eric Winter

Dennis Lazar

Greg Fitzgerald

Recent Accomplishments

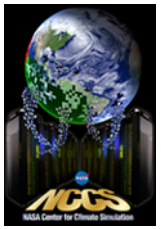

- SCU8 is in production
- SCU9 Intel Xeon SandyBridge (480 nodes) integrated
- Ramping up support for the GMAO Nature Run
- Continued development of ODAS Data Management Services
- Production prototype for remote visualization based on UV-CDAT

NCCS Compute Capacity Evolution
September 2007- September 2013

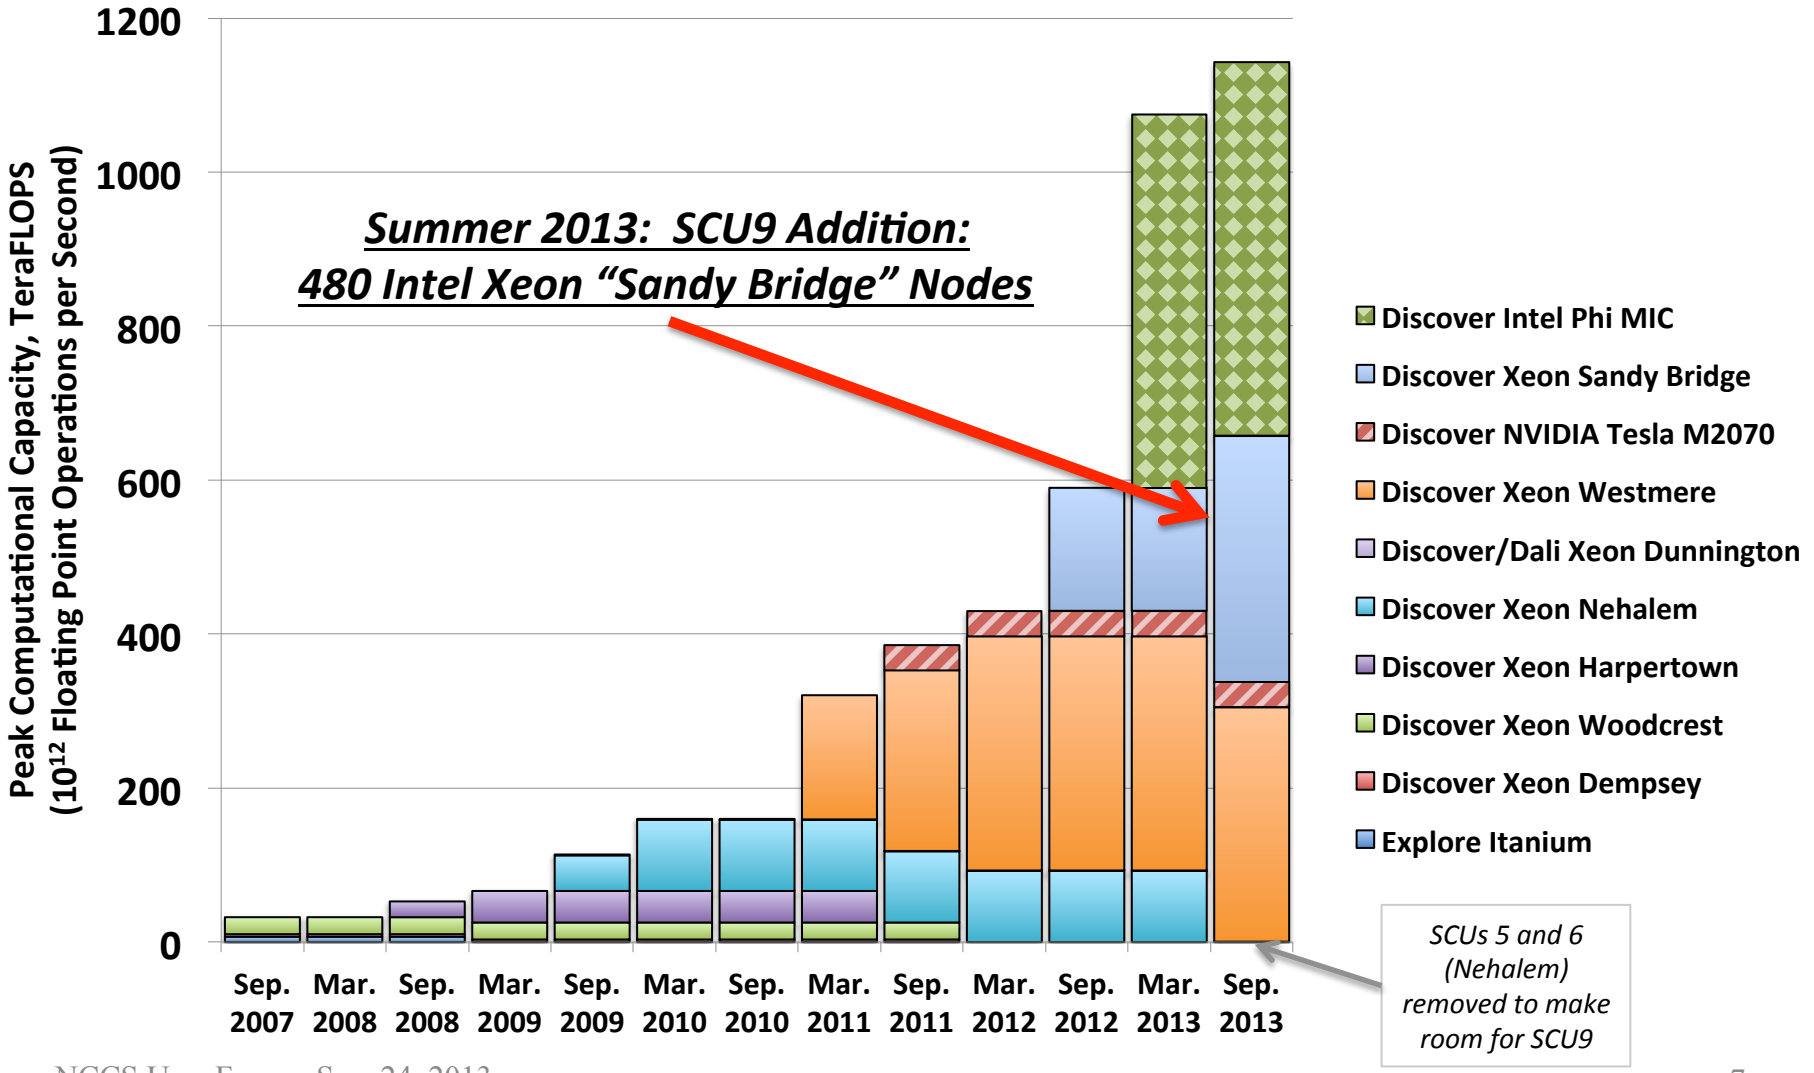

Support for Special Projects

- GEOS-5 Nature Run
  - 2-year Nature Run at 7.5 km resolution
  - 3-month Nature Run at 3.5 km resolution
  - Will generate about 4 PB of data
- Various Campaigns
  - HS3
  - Discover-AQ
  - SEAC4RS
  - NOMADSS

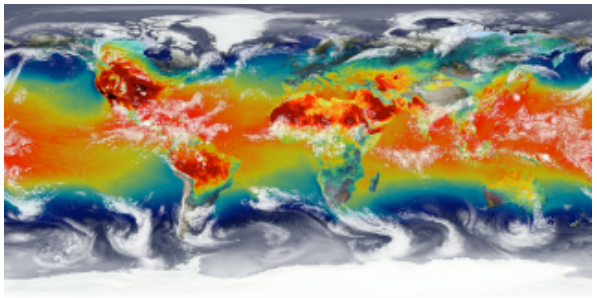

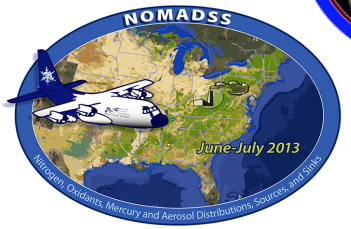

Data Services for GMAO's Decadal Ocean Data Assimilation Data System (ODAS)

- GEOS-5 Ocean Data Assimilation System
  - Monthly data from January 1961 through August 1996 (so far)
- Active NCCS Data Portal Services for ODAS
  - Anonymous FTP
  - Thematic Real-time Environmental Distributed Data Services (THREDDS)
  - Live Access Server (LAS)
- In development
  - GRaDS Data Server (GDS)
  - Integrated Rule-Oriented Data Systems (iRODS)
- GMAO Ocean Data Portal points to these services
  - http://gmao.gsfc.nasa.gov/research/ocean/DATA_PORTAL/index.php

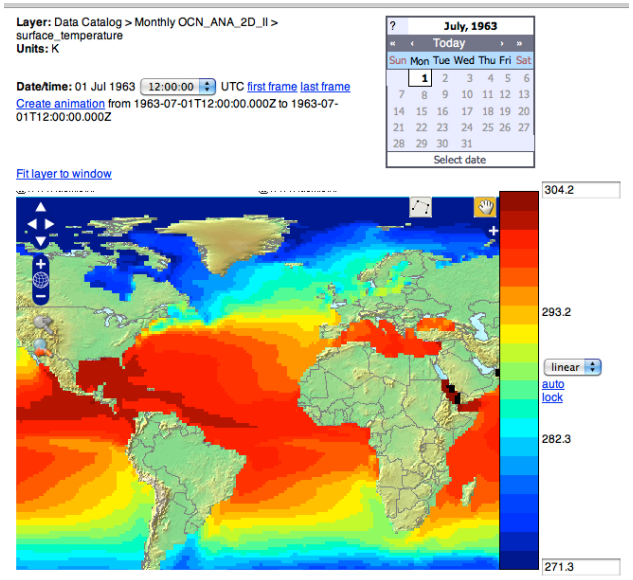

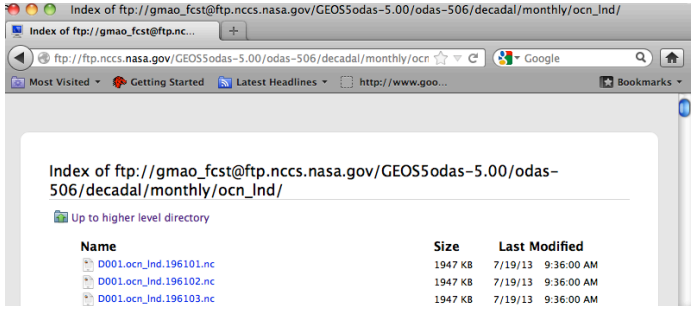

UV-CDAT Remote Visualization Prototype

- Goal – Provide a high performing interface to UV-CDAT running on NCCS systems
  - Ultrascale Visualization Climate Data Analysis Tools (UV-CDAT) http://www.nccs.nasa.gov/images/UVCDAT-Brown-Bag.ppt
    http://uvcdat.llnl.gov/
- Production prototype in place
  - Actually evaluating two competing technologies (one open source and one commercial)
  - Would like users to access the prototypes and provide feedback
  - ***Submit a ticket if you are interested***
- IEEE Computer Journal Article, Sept. 2013, highlights UV-CDAT
  - "Ultrascale Visualization of Climate Data"

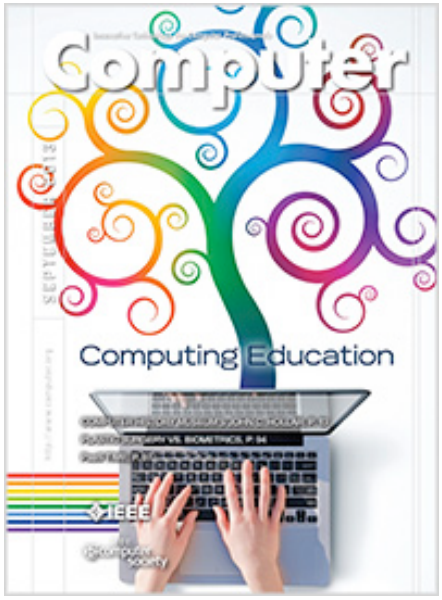

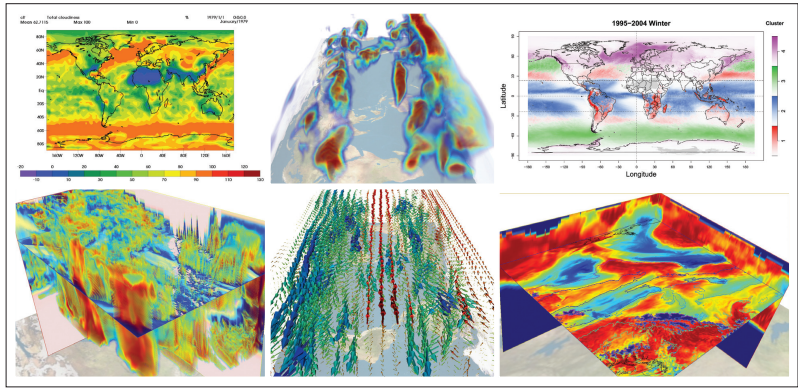

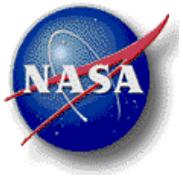

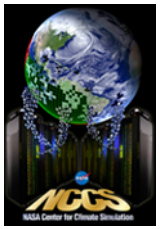

Special Topics

Dan Duffy,
HPC Lead and NCCS Lead Architect

Special Topics

- NCCS Strategic Plan
  - Create a 5-year strategic plan for the NCCS
  - Would like involvement from the user community
  - "Design thinking" approach
  - Draft by February 2014

- Decommission of SCU5/6 Components
  - Approximately half of the SCU5/6 system has been claimed from NASA excess by UMBC (Milt Halem's Group)
  - Other components will be reused internally within the NCCS, providing batch compute nodes for data services, upgrading the SVS render nodes, and potentially the Goddard animation group.

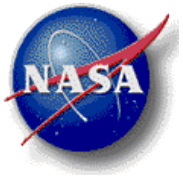

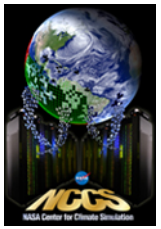

NCCS Updates

Dan Duffy,
HPC Lead and NCCS Lead Architect

Discover Resource Manager

- The NCCS is transitioning away from PBS to SLURM (Simple Linux Utility for Resource Management)
  - Best value to the NCCS based on current and future technical requirements and cost
- Complete transition to SLURM by **October 31, 2013**
  - Striving to make this as transparent as possible
  - Most PBS scripts and q-commands will be able to be used without any changes
  - Pioneer users are already testing SLURM
  - Breakout session on SLURM at end of the User Forum
- Coming soon: NCCS web pages with SLURM info
- References
  - Website
    - http://slurm.schedmd.com/
  - The Rosetta Stone document to compare different resource managers
    - http://www.schedmd.com/slurmdocs/rosetta.pdf
  - An advanced usage tutorial
    - http://slurm.schedmd.com/slurm_ug_2011/Advanced_Usage_Tutorial.pdf

SCU8 Status

- SCU8 – in general use since early/mid-summer
  - 2.6 GHz Intel Xeon SandyBridge processors (16 cores per node)
  - 32 GB of RAM (2 GB per core)
  - One Intel Phi accelerator per node
  - Quad Data Rate (QDR) InfiniBand
- Intel Phi Coprocessors
  - Offload use available on all other SCU8 nodes.
  - Direct use of Phi is now available via special "native" queue, by request
- Training:
  - A number of NCCS Brown Bags so far
    - Content available on NCCS web site
    - Training will be repeated upon request
  - Contact support@nccs.nasa.gov

Discover SCU8

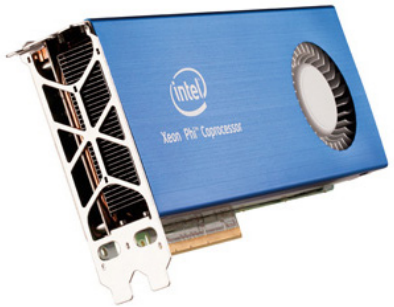

Intel Phi Coprocessor

SCU9 Status

- SCU9
  - Decommissioned SCU5/6 and integrated SCU9 earlier this year
  - 480 compute nodes
    - 2.6 GHz Intel Xeon SandyBridge processors (16 cores per node)
    - 64 GB of RAM (4 GB per core)
    - Fourteen Data Rate (FDR) Infiniband
    - Each capable of holding up to two Intel Phi units or NVIDIA GPUs – None installed at this point
  - Upgraded all I/O nodes to SandyBridge processors with this integration effort as well
- Dedicated to the GMAO Nature Run for now
  - Hope to have more general access available by early 2014

Discover SCU9

GPUs and Phis

- NVIDIA Graphical Processing Units (GPUs)
  - warp queue
    - ~30 nodes, each with two NVIDIA M2070 units
  - warp-test queue – newest GPUs
    - 2 nodes available with NVIDIA K10 units
    - 1 node available with NVIDIA K20
- Intel Phi Co-processors
  - native queue
    - ~20 nodes available to run offload, native, or symmetric mode
    - All other nodes within SCU8 can run in offload mode
- Want access or want help using these advanced computing systems?
  - Contact support@nccs.nasa.gov

NVidia GPU

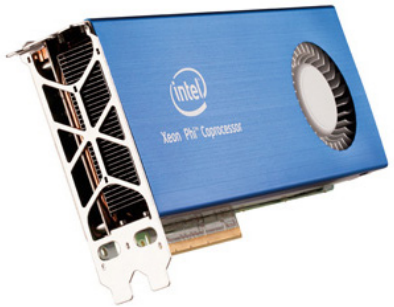

Intel Phi Coprocessor

Nature Run Disk

- GEOS-5 Nature Run
  - 2-year nature run at 7.5-km resolution
  - 3-month nature run at 3.5-km resolution
  - Will generate about 4 PB of data
- Disks will arrive by the end of the month
  - 5,760 TB of usable capacity
  - 1,800 by 4-TB disk drives in 3 racks
  - NetApp solution similar to what is currently in operations
- In production (at least for the Nature Run) within 1 month after arrival
  - Takes 2 weeks just to format all the disks!

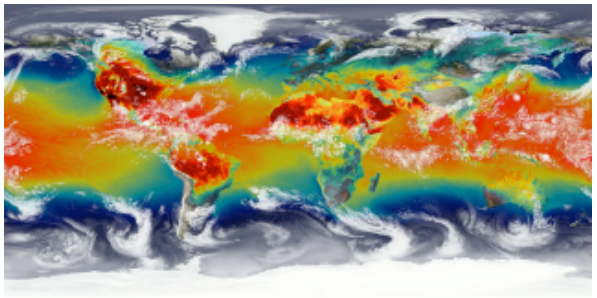

NetApp 5400

GPFS Metadata Disk

- Discover GPFS Metadata Storage Upgrade
  - Goal is to dramatically increase the aggregate metadata performance for GPFS
  - Current aggregate IOPs is ~8K (on a good day)
  - Cannot speed up a single metadata operation, but can speed up the combination of all operations
- Purchased a solution from Whiptail
  - Solid state disks
  - Targeting a minimum of 40K IOPs and hoping to do much better!
- Delivery by the end of the month
  - NCCS will coordinate closely with the user community throughout the integration of this new capability
  - Fully integrated before the holidays

Whiptail Accela

Data Portal Capacity and Throughput

- Disk Capacity
  - Increased by ~100 TB of usable disk
- Additional Servers
  - Adding 2 servers with 10 GbE capability
  - In support of higher-speed access between Discover and the Data Portal
- Discover File System Servers
  - Upgrading the discover file system servers to 10 GbE
- Fully integrated before the holidays

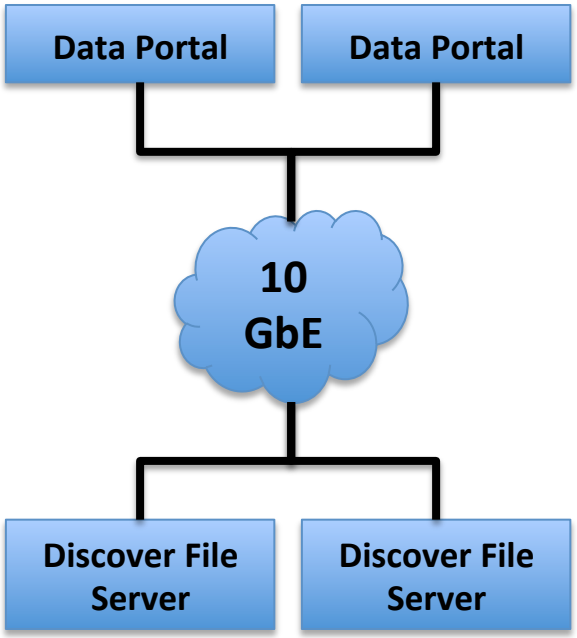

Archive Software Stack Changes

- Upgrade to the Dirac Archive Operating System, DMF software, and CXFS-related software packages in November
  - NCCS will coordinate with the user community and schedule an appropriate time for the upgrade

Oracle SL8500 Tape Library

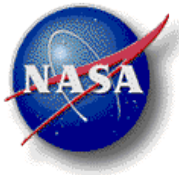

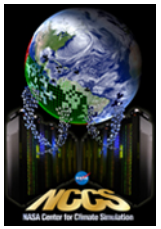

User Survey

Dan Duffy,
HPC Lead and NCCS Lead Architect

Update on NCCS Response to User Survey

1. External data transfer performance and convenience
   - Upgrade of the Data Portal *(by end of CY 2013)* and Discover File System Servers *(completed)* to 10 GbE capabilities
   - Analysis of the GISS to NCCS network and recommendations for upgrades *(completed)*
   - Upgrade of the NCCS's SEN to CNE link to 10 GbE *(completed)*
2. More timely notifications of problems or unplanned outages
   - Web dashboard for system status is still under development
3. "Architecting for More Resiliency," especially the Discover storage file systems
   - Beginning of a strategic planning effort; asking for support from the user community
   - Continued evaluation of many alternative computing platforms, including cloud and even on-demand HPC

2013 User Survey

- Annual survey will be announced by the middle of October
  - Survey will be open for about 3 weeks
  - Expect it to take between 15 to 20 minutes to complete your responses
- Expect some changes to the questions based on our lessons learned
  - Will be asking you to rank the importance of the question as well
- Would like more contact information
  - Not required, but helps the NCCS to follow through on specific comments

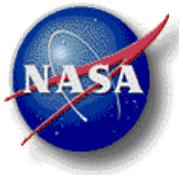

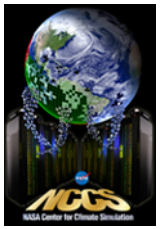

NCCS Operations & User Services Update

Ellen Salmon

- Intel Phi Code Porting
- Upcoming & Status
- Ongoing Investigations
- NCCS Brown-Bag and SSSO Seminars

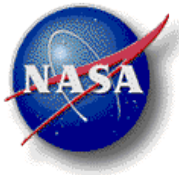

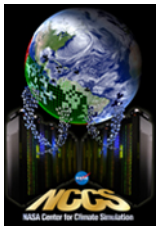

Intel Phi Code Porting

NCCS Code-Porting Efforts for Discover Intel Phi Many Integrated Core (MIC)

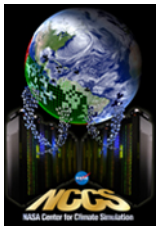

- NCCS staff, SSSO, vendor, and external community members continue to work on the following codes:
  - GEOS-5 components (GMAO, NOAA/GFDL, et al.)
  - WRF (with NOAA/ESL, NCAR, et al.)
  - Hall3D Magnetohydrodynamics code (Dorelli, Goldstein, et al.)
  - Ice Melt code (Kwo-Sen Kuo et al.)
  - GRAIL (high degree/high order solutions of lunar gravity field)
- Contact support@nccs.nasa.gov .

Intel Phi: Preliminary Results for GEOS5 Dynamical Core

Run Parameters	
Grid Resolution	C96 (roughly 1°)
Vertical levels	71
Test Case	Held-Suarez
Model Time	6 hrs
Precision	64 bit

Hard-ware	6x SNB Nodes	6x KNC 5110P
MPI per Node	16	1
OMP per MPI	0	240 (Max)

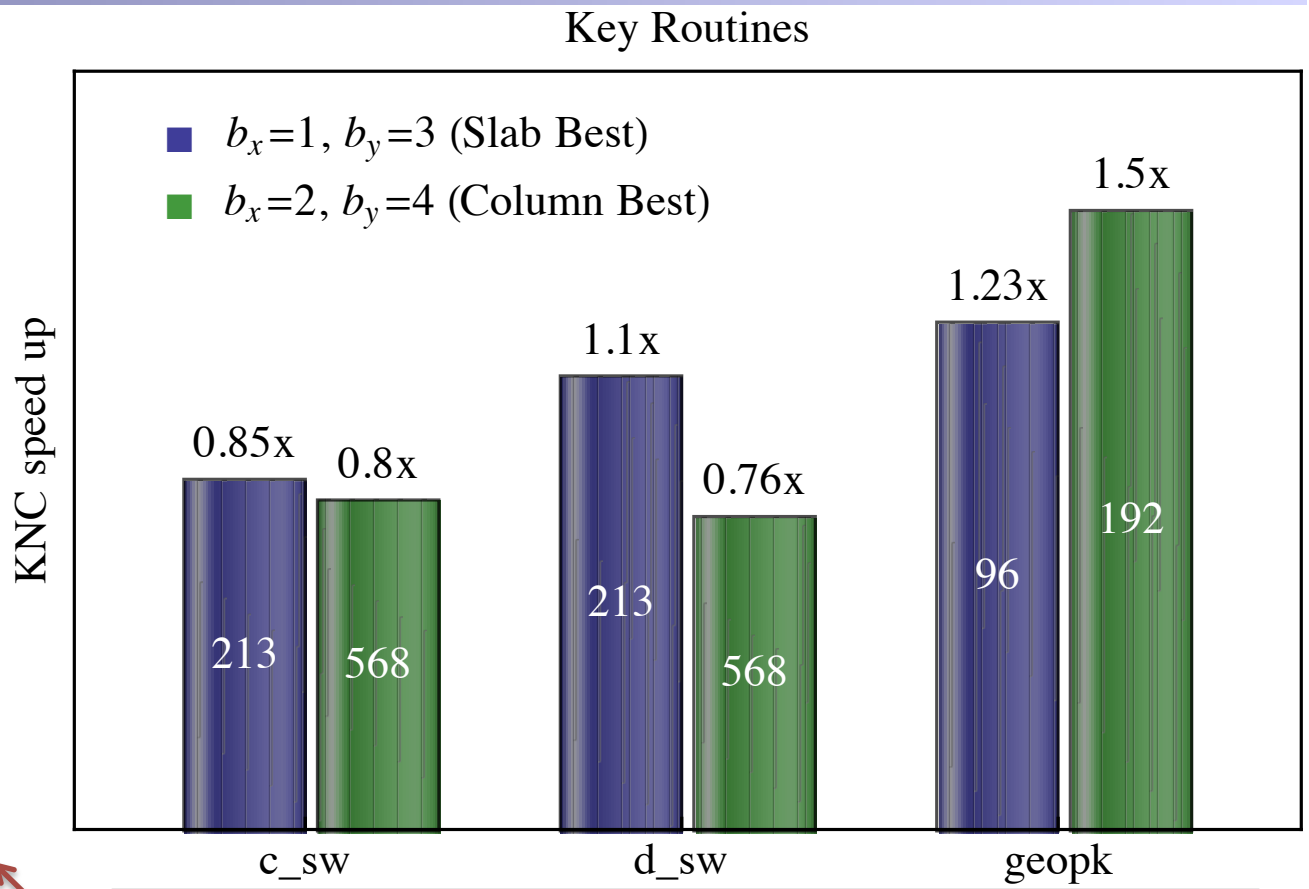

SNB == Intel Xeon Sandy Bridge (16 cores per node)
KNC == Intel Xeon Phi "Knight's Corner" 5110P (60 cores each)
OMP == OpenMP

GEOS-5 preliminary results provided by SSSO's Craig Pelissier and Kareem Sorathia.

Intel Phi: Preliminary Performance Comparison: WRF 3.5

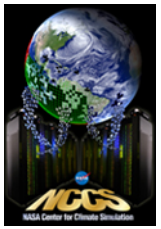

CONUS 12-km Test Case	Configurations (nodes x MPI ranks x OMP threads)	Compute Time (Seconds)
Xeon SNB (Sandy Bridge) Only (16 cores per node)	1 x 4 x 4	106.1
Xeon Phi (KNC) Only (60 cores per node)	1 x 1 x 180	122.2
Symmetric (Sandy Bridge SNB and Xeon Phi KNC)	nodes x (SNB + KNC) = 1 x (8 x 2 + 7 x 34)	**92.0**
CONUS 2.5-km Test Case	**Configurations (nodes x MPI ranks x OMP threads)**	**Compute Time (Seconds)**
Xeon SNB (Sandy Bridge) Only (16 cores per node)	8 x 4 x 4	931.2
Symmetric (Sandy Bridge SNB and Xeon Phi KNC)	nodes x (SNB + KNC) = 8 x (8 x 2 + 6 x 40)	**759.6**

OMP == OpenMP CONUS = US Lower 48 States

Intel Phi and GPUs: Preliminary Results for Hall3D Magnetohydrodynamics (MHD) Code

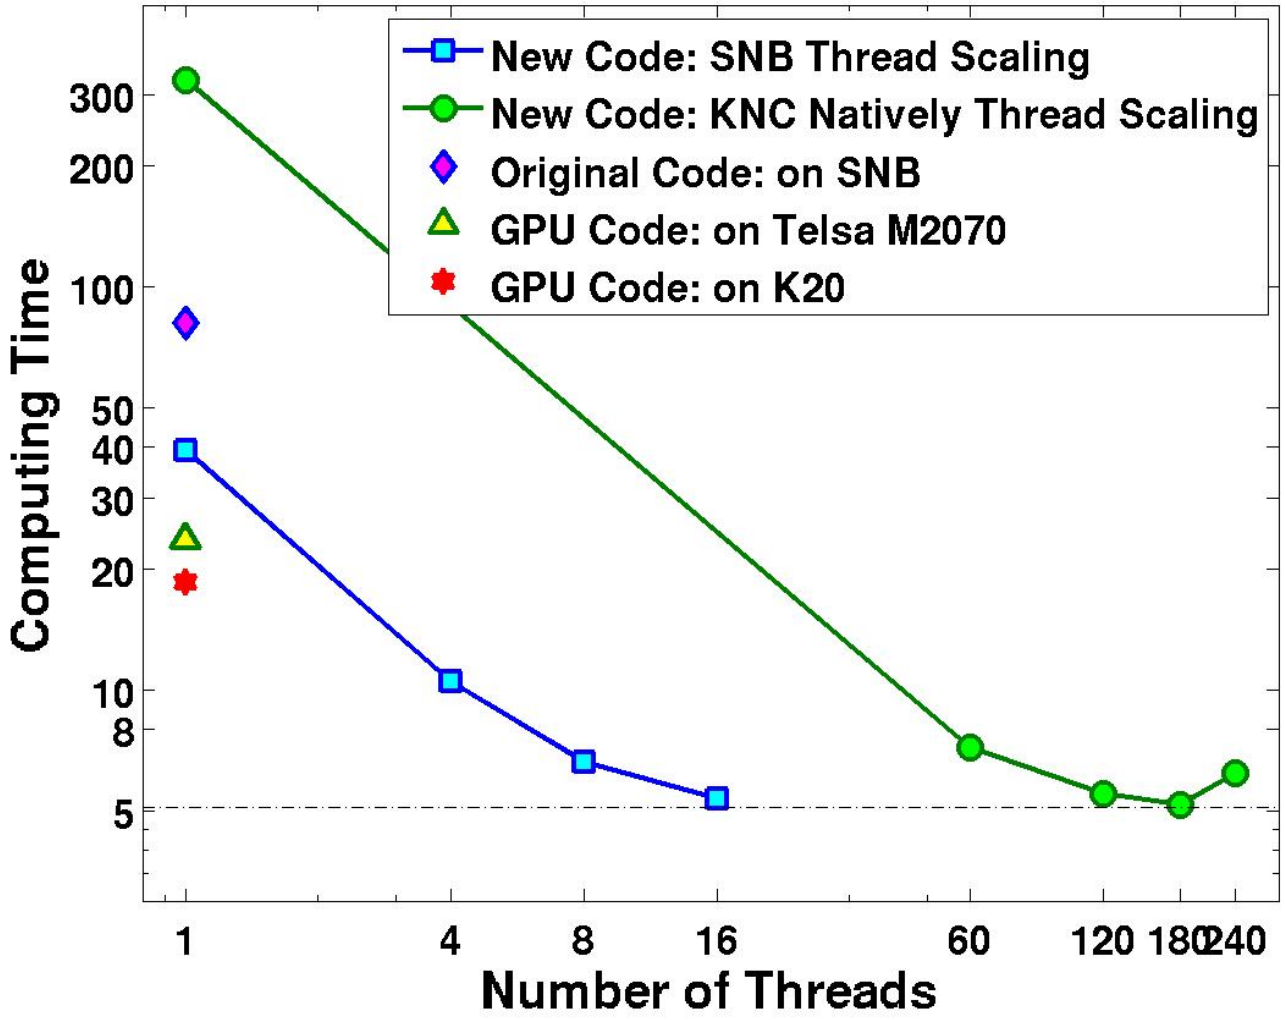

Intel Xeon:
SNB: Xeon Sandy Bridge (16 cores per node)

Intel Many Integrated Core (MIC) Coprocessor:
KNC: Xeon Phi "Knight's Crossing" 5110P (60 cores each)

NVIDIA GPUs:

Tesla M2070:
(~448 CUDA cores each)

K20:
(~2500 CUDA cores each)

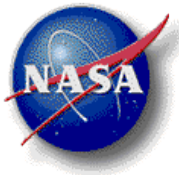

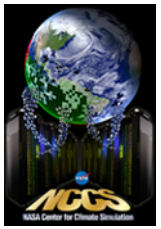

Upcoming & Status

Upcoming Maintenance: Outages and Updates

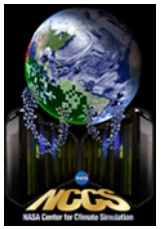

- **Discover resources**
  - SLURM resource (batch) manager transition
    - When: "rolling" queue/partition migrations through October 31
  - DDN controller replacement (disk arrays)
    - When: late October; possible day-long Discover-wide outage
  - GPFS software
    - When: TBD; hoping for minimal-impact "rolling" updates
  - GPFS metadata Solid State Disk
    - When: TBD "rolling" deployment through end of CY2013, no Discover-wide outage expected
- **Dirac Archive**
  - Software updates: Linux operating system, DMF, CXFS, other related packages
    - When: November; day-long Dirac Archive outage

SLURM and Older Intel MPI, MVAPICH, and OpenMPI Versions

If you use the versions below, please contact support@nccs.nasa.gov ASAP for help in migrating to SLURM-compatible versions.

- Intel MPI
  - mpi/impi-3.1.*
  - mpi/impi-3.2.011
  - mpi/impi-3.2.1.009
  - mpi/impi-4.*-beta
- MVAPICH
  - other/mpi/mvapich2-1.6rc2/ intel-11.1.072
  - other/mpi/mvapich2-1.7*
  - other/mpi/mvapich2-1.8/*
  - other/mpi/mvapich2-1.8a2/*
- OpenMPI
  - mpi/openmpi-1.2.5/intel-9
  - mpi/openmpi-1.2.5/intel-10
  - other/mpi/openmpi/1.4.*
  - other/mpi/openmpi/1.6-gcc_4.8-20120401_nag-5.3-854
  - other/mpi/openmpi/1.6.0*
  - other/mpi/openmpi/1.6.3*
  - other/mpi/openmpi/1.6.4*

*These old (and beta) versions are **incompatible with SLURM**.*

OFED (software stack)

- Discover InfiniBand OFED (software stack) changes:
  - 2/3 of Discover is already on the new OFED release
    - All computational nodes on InfiniBand Fabric 2: SCU7 Westmere nodes, and all SCU8 and SCU9 Sandy Bridge nodes
  - Rolling, announced, gradual changeovers of other parts of Discover (e.g., via PBS queues or properties).
    - SCUs 1 through 4, handful of remaining interactive and Dali nodes
  - ***Recompile is recommended***.
    - Some codes work fine without a recompile.
    - Other codes require a recompile to take advantage of some advanced features.

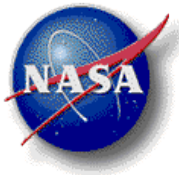

Discover Compiler / Library Recommendations

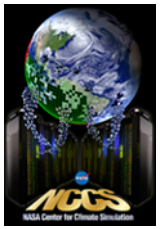

- Use **current libraries and compilers** to get many benefits:
  - Executables created with older versions can experience problems running on Discover's newest hardware.
  - Often, simply rebuilding with current compilers and libraries ***(especially Intel MPI 4.x and later)*** fixes these issues.
  - Current versions can enable use of new features like Sandy Bridge nodes' advanced vector extensions (AVX) for improved performance.
  - Use of current versions greatly increases NCCS staff's abilities to track down other problems…
    - Especially when seeking vendor support to fix problems.

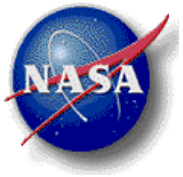

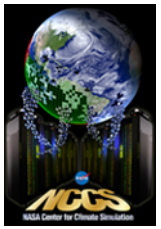

Ongoing Investigations

Ongoing Discover Investigations

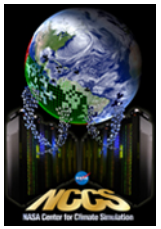

- GPFS slowness due to heavy parallel I/O activity. *Mitigations:*
  - New Sandy Bridge I/O nodes (completed Summer 2013):
    - More cores per I/O node—16 cores, rather than 8—improved concurrency.
    - More total memory channels—4, rather than 3, per "socket"—better for data moving.
    - More total I/O "lanes" per I/O node.
  - Many-month (perpetual) effort: background "rebalancing" of data among filesystems to better accommodate workloads.
- Heavy GPFS metadata workloads. *Mitigations:*
  - Solid State Disk acquisition: improve responsiveness in handling many concurrent small, random I/O actions (e.g., for directories, filenames, etc.). Installation/integration begins soon.
- PBS "Ghost Jobs". *Mitigations:*
  - PBS 12 upgrade (completed) and migration to SLURM (through October 31, 2013).
- Slow InfiniBand Memory Leak. *Mitigations:*
  - Enhanced automated monitoring and pro-active node reboots between jobs.
- Discover Computational Allocation Issue. *Corrected.*

Ongoing Recovery: Archive Tape Media Issue

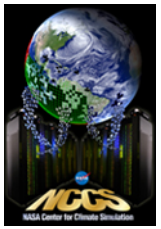

September 2013: Oracle continues to recover files from damaged Archive tapes.

- Crinkled/damaged archive tapes caused a number of "Please examine/replace these archive files" tickets in the last several months.
- ***Damage is no longer occurring.***
- Oracle identified **faulty tape motor on a single tape drive** as the cause, and:
  - Replaced that tape drive and 11 others to proactively remediate the problem, prior to pinpointing the cause.
  - Providing data recovery services to extract usable data from damaged tapes.
  - Replacing all tape media affected by the problem.
- ~12 tapes sent, so far, to Oracle Tape Recovery Services.
  - So far 7 [September 2013 update] have been returned.
    - Recovered all data on 2 of the tapes and much of the data on the other 3 tapes.
- Larger list of tapes was damaged, but NCCS staff was able to recover files from those because second copies of files still existed on separate (unaffected) tapes.
- Reminder: **dmtag –t 2 <filename>** to get two tape copies of archive files, where needed.

June 2013

Brown Bag Seminars

NCCS Brown Bag Seminars

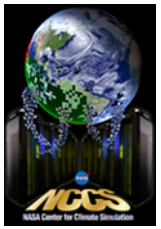

- ~Twice monthly in GSFC Building 33 (as available).
- Content is available on the NCCS web site following seminar:

  https://www.nccs.nasa.gov/list_brown_bags.html

- Current emphases:

  Migration from PBS to SLURM

  Using Intel Phi (MIC) Coprocessors

- Current/potential *Intel Phi* Brown Bag topics:

  - ✓ Intro to Intel Phi (MIC) Programming Models
  - ✓ Programming on the Intel MIC Part 2 – How to run MPI applications
  - ✓ Running WRF on Intel Phi MIC
  - Advanced Offload Techniques for Intel Phi
  - Maximum Vectorization
  - Performance Analysis via the VTune™ Amplifier
  - Performance Tuning for the Intel Phi

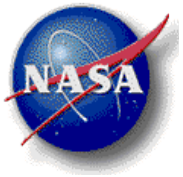

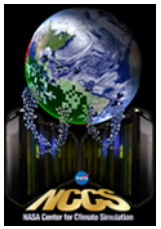

Questions & Answers

NCCS User Services:

support@nccs.nasa.gov

301-286-9120

https://www.nccs.nasa.gov

Contact Information

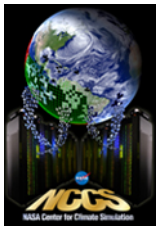

NCCS User Services:

support@nccs.nasa.gov

301-286-9120

https://www.nccs.nasa.gov

http://twitter.com/NASA_NCCS

Thank you

Supporting Slides

NASA Center for Climate Simulation Supercomputing Environment

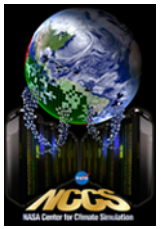

Supported by HQ's Science Mission Directorate

① ***Discover*** Linux Supercomputer, June 2013:

- Intel Xeon nodes
  - ~3,200 nodes
  - ~42,100 cores
  - Peak ~624 TFLOPS general purpose
  - 97 TB memory (2 or 4 GB per core)
- Coprocessors:
  - Intel Phi MIC
    - 480 units
    - ~485 TFLOPS
  - NVIDIA GPUs
    - 64 units
    - ~33 TFLOPS
- Shared disk: **7.2** PB

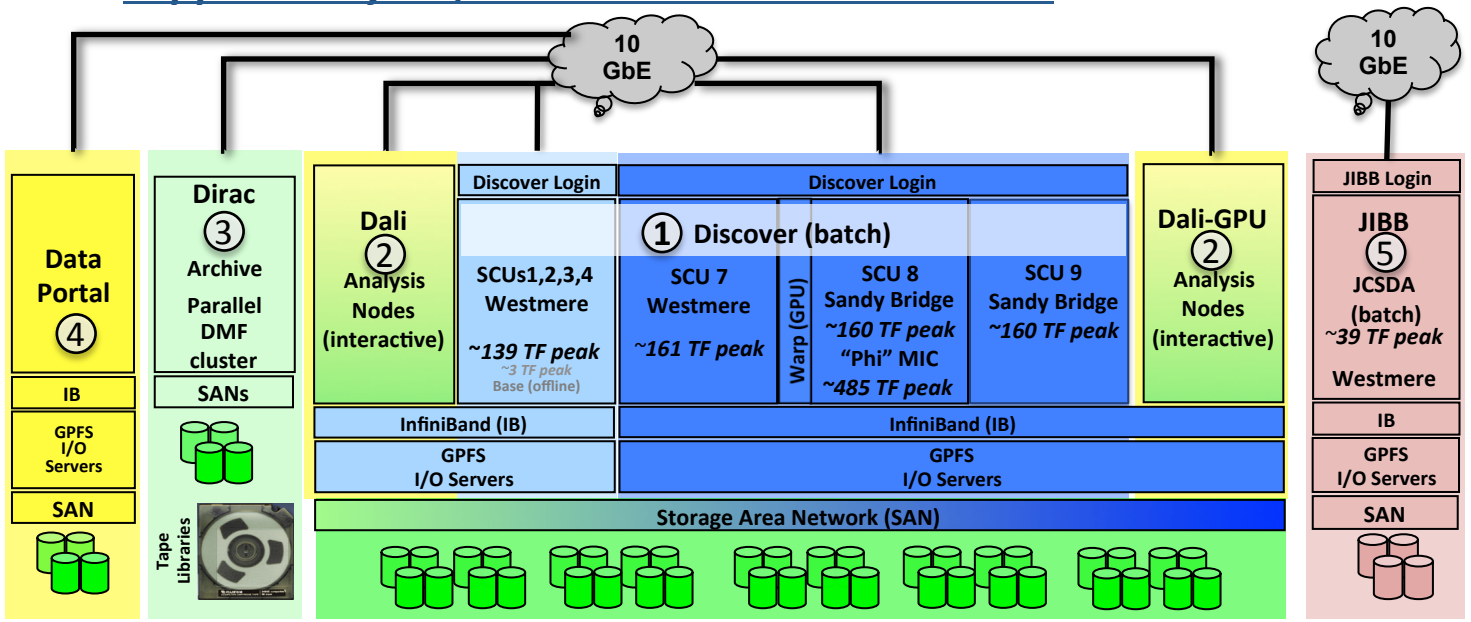

② ***Dali*** and ***Dali-GPU*** Analysis

- 12- and 16-core nodes
- 16 GB memory per core
- Dali-GPU has NVIDIA GPUs

③ ***Dirac*** Archive

- 0.9 PB disk
- ~70 PB robotic tape library
- Data Management Facility (DMF) space management

④ ***Data Portal*** Data Sharing Services

- Earth System Grid
- OPeNDAP
- Data download: http, https, ftp
- Web Mapping Services (WMS) server

⑤ ***JIBB***

- Linux cluster for Joint Center for Satellite Data Assimilation community

March 1, 2013

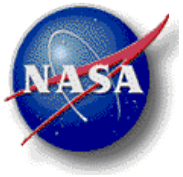

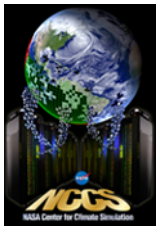

NCCS Metrics Slides
(Through August 31, 2013)

NCCS Discover Linux Cluster Utilization Normalized to 30-Day Month

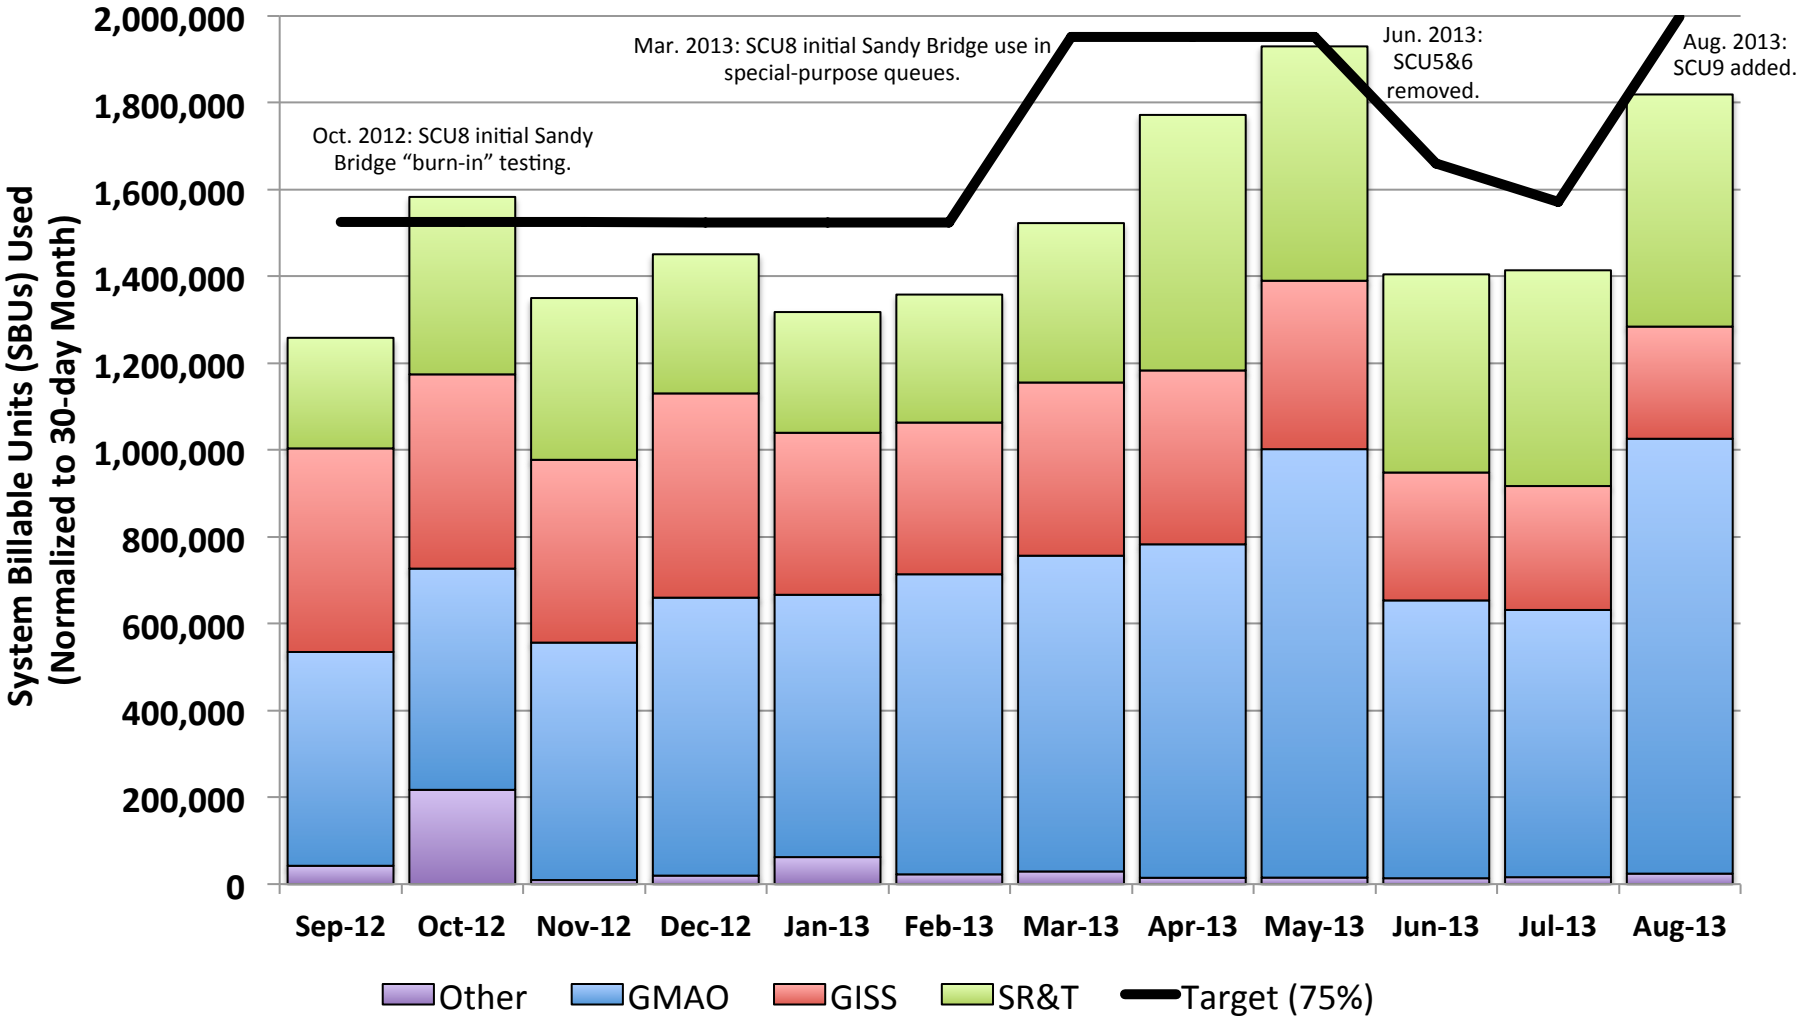

Discover Linux Cluster Expansion Factor

Expansion Factor = (Queue Wait + Runtime) / Runtime

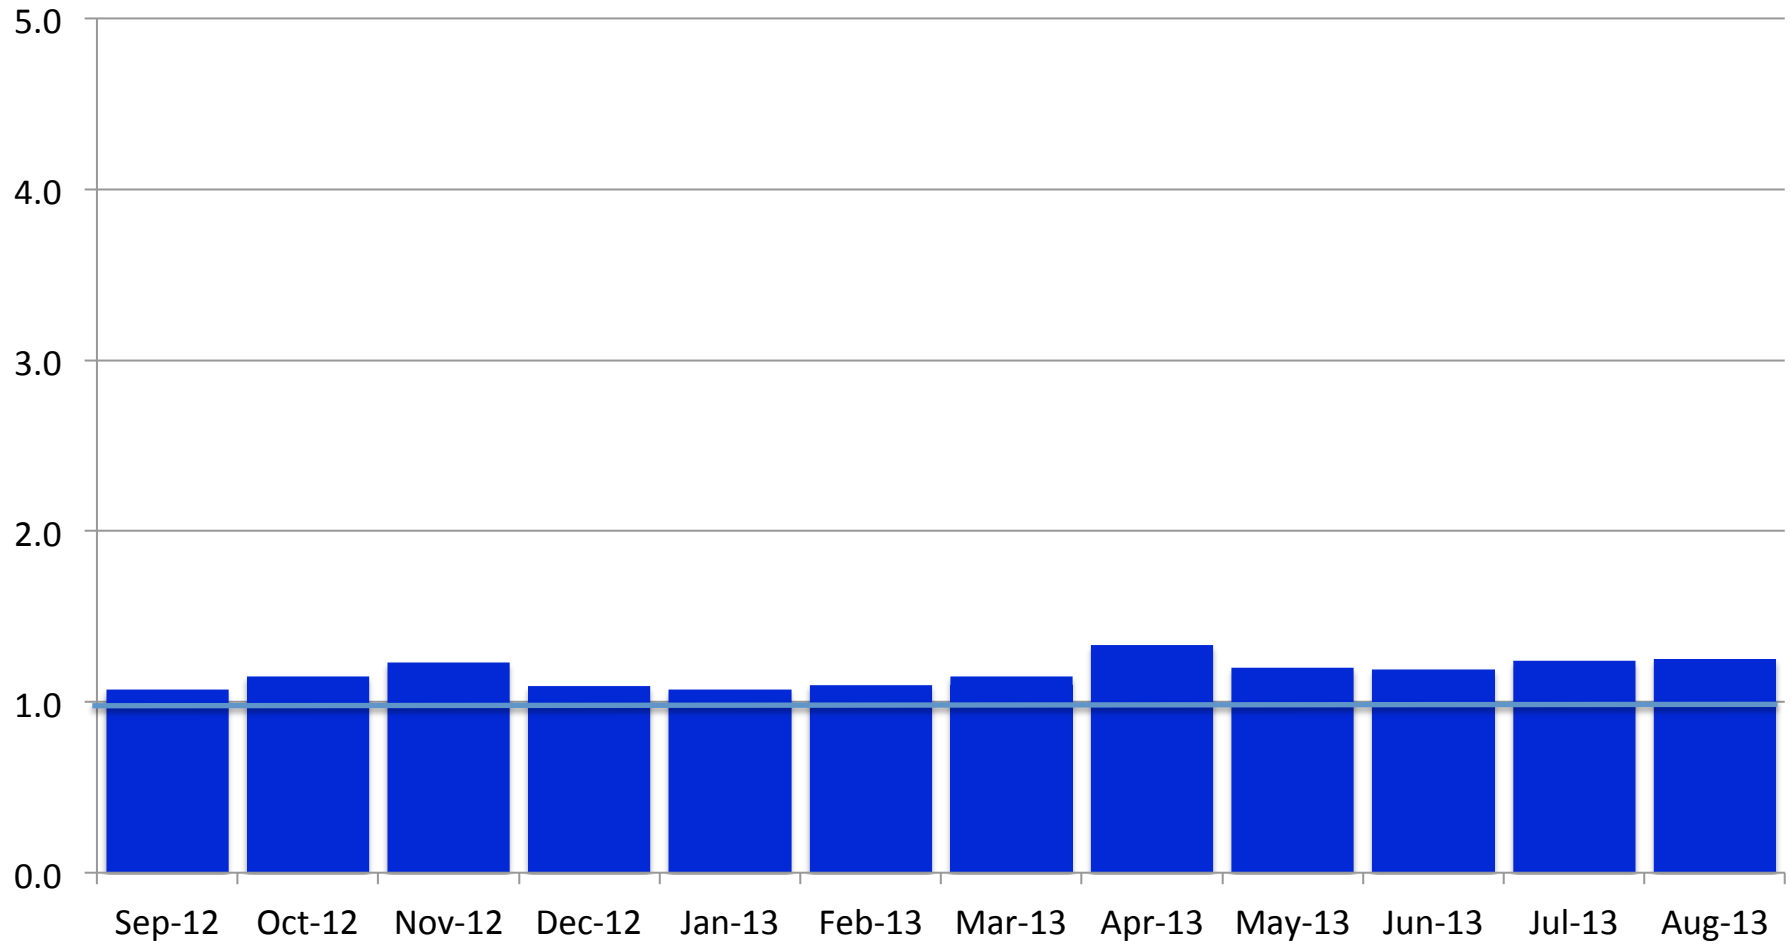

Discover Linux Cluster Downtime

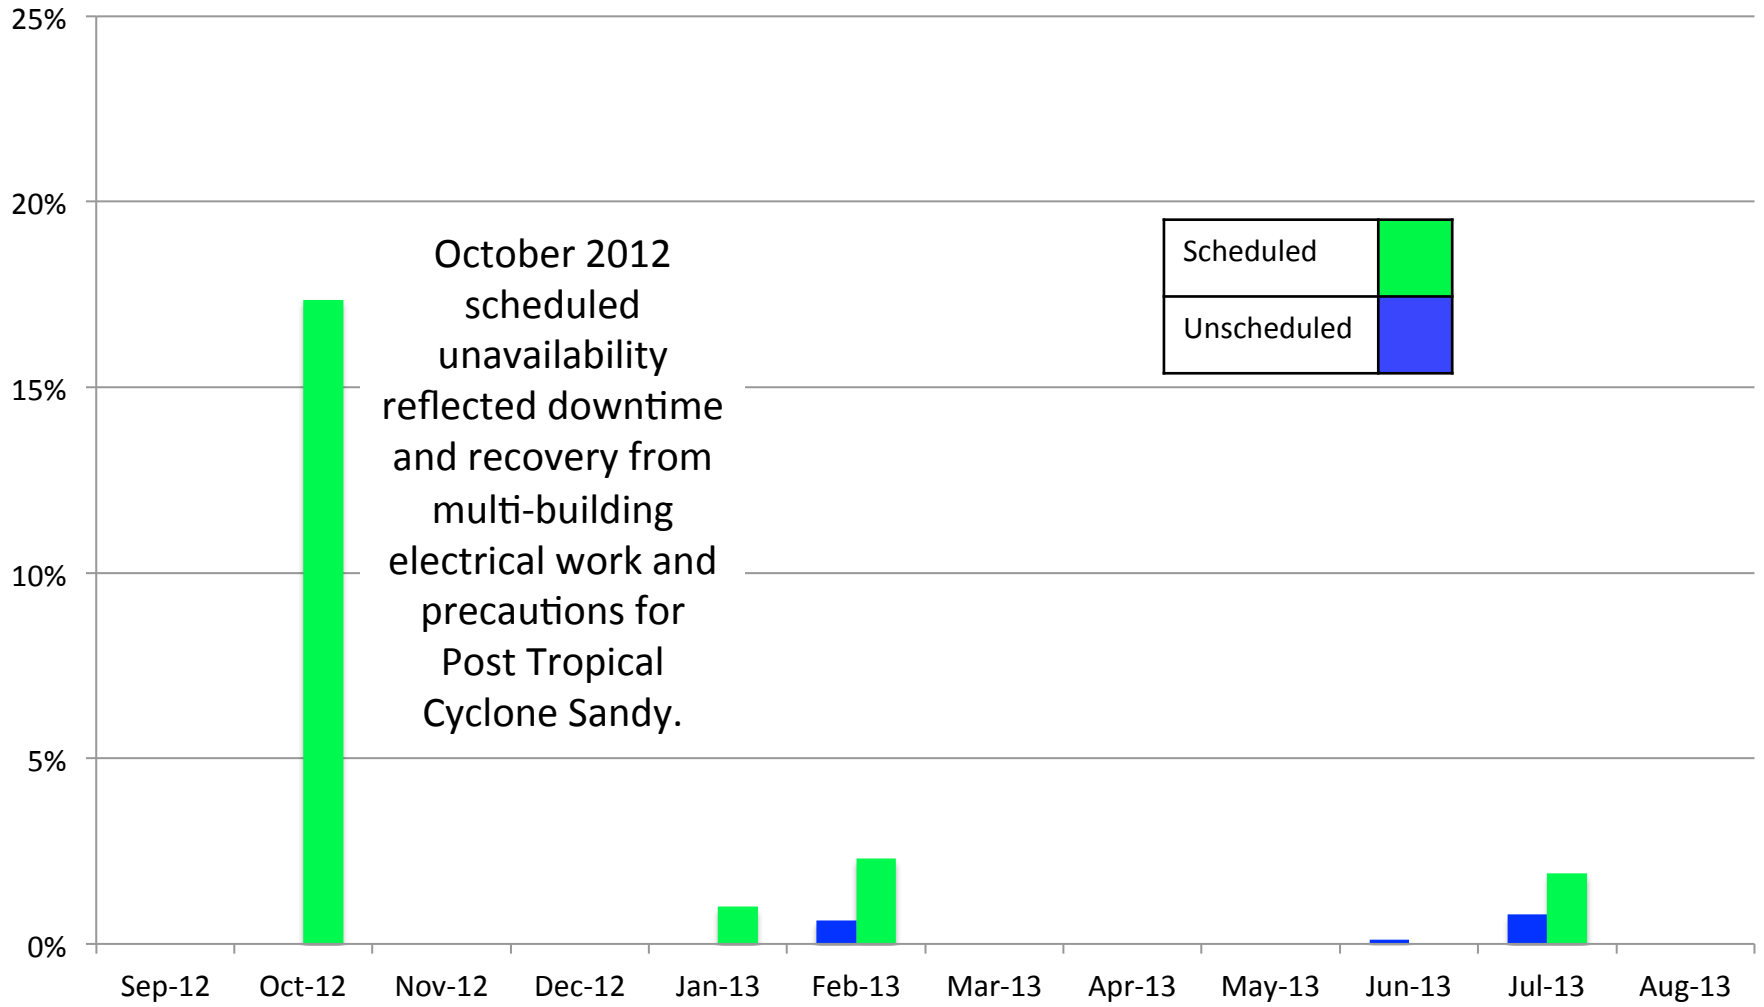

NCCS Mass Storage

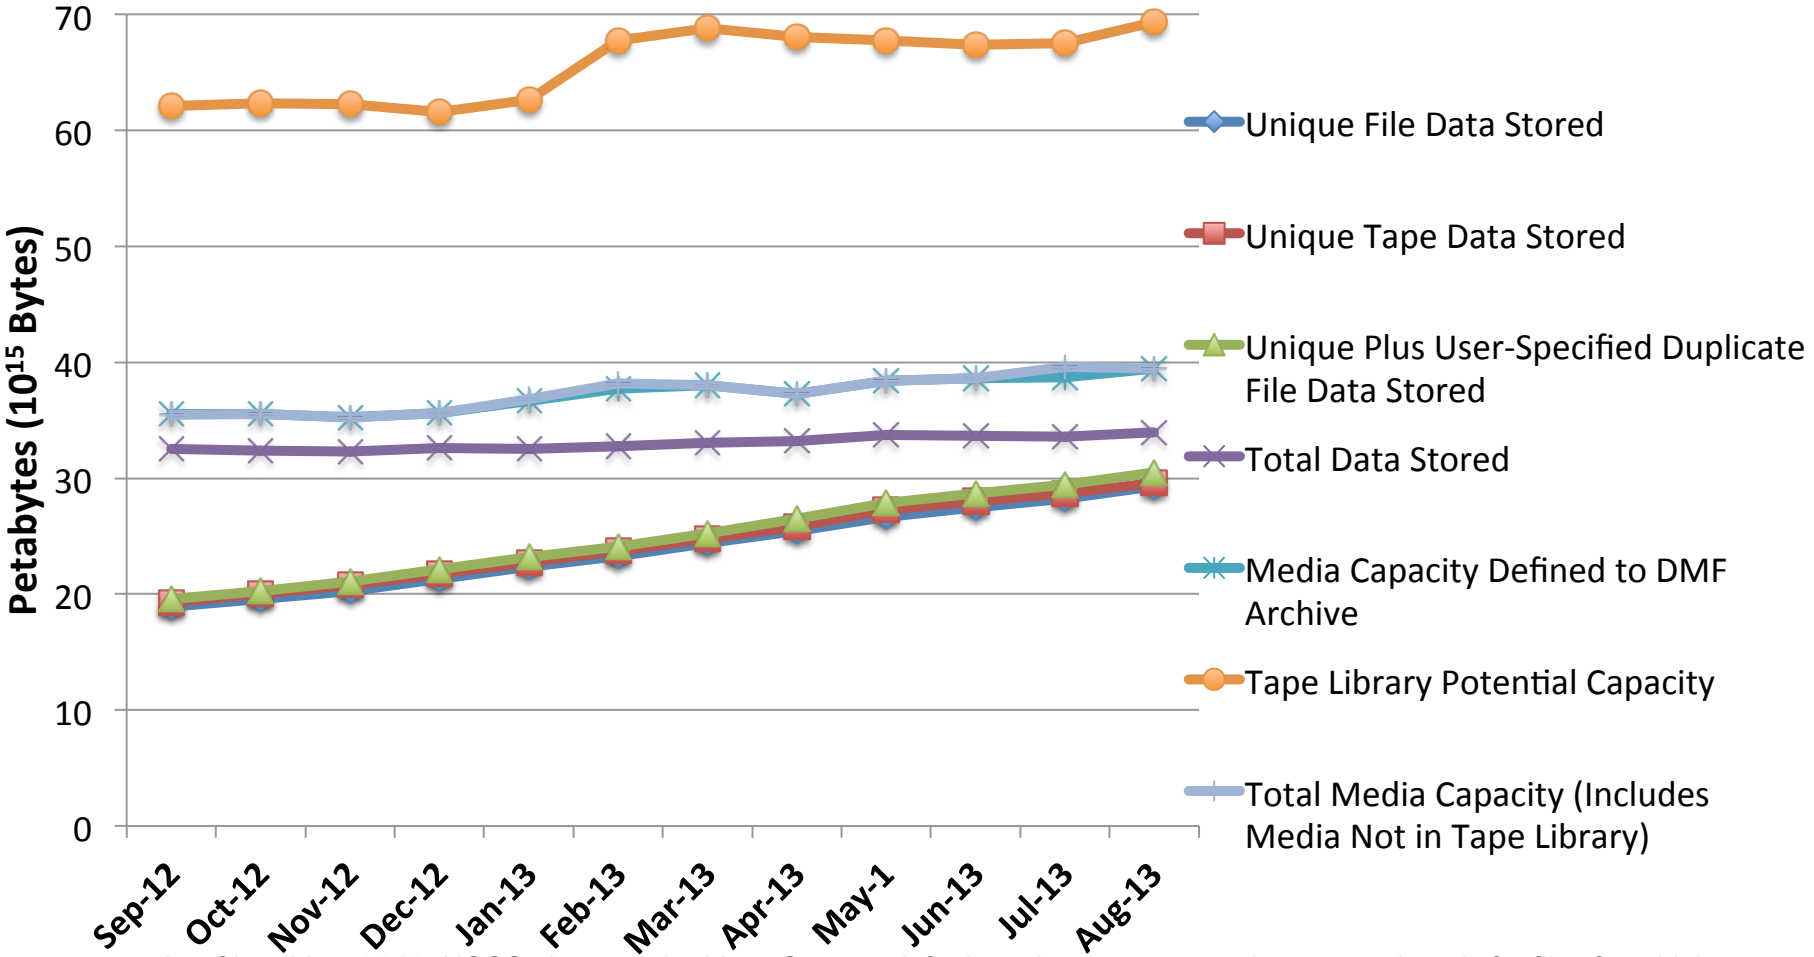

As of late May, 2012, NCCS changed the Mass Storage default so that two tape copies are made only for files for which two copies have been explicitly requested. NCCS is gradually reclaiming second-copy tape space from legacy files for which two copies have not been requested.

NCCS Earth System Grid Federation Services for NASA's and Peer Organizations' Climate Simulations, Selected NASA Observations, and Selected Analyses

- **GISS and GMAO** researchers are using the NCCS Discover cluster for simulations in support of the fifth phase of the Coupled Model Intercomparison Project (**CMIP5**), which supports the **Intergovernmental Panel on Climate Change's Fifth Climate Assessment (IPCC AR5)** and related research.
- The research community accesses data via the NCCS's **Earth System Grid Federation (ESGF)** node http://esgf.nccs.nasa.gov/ .

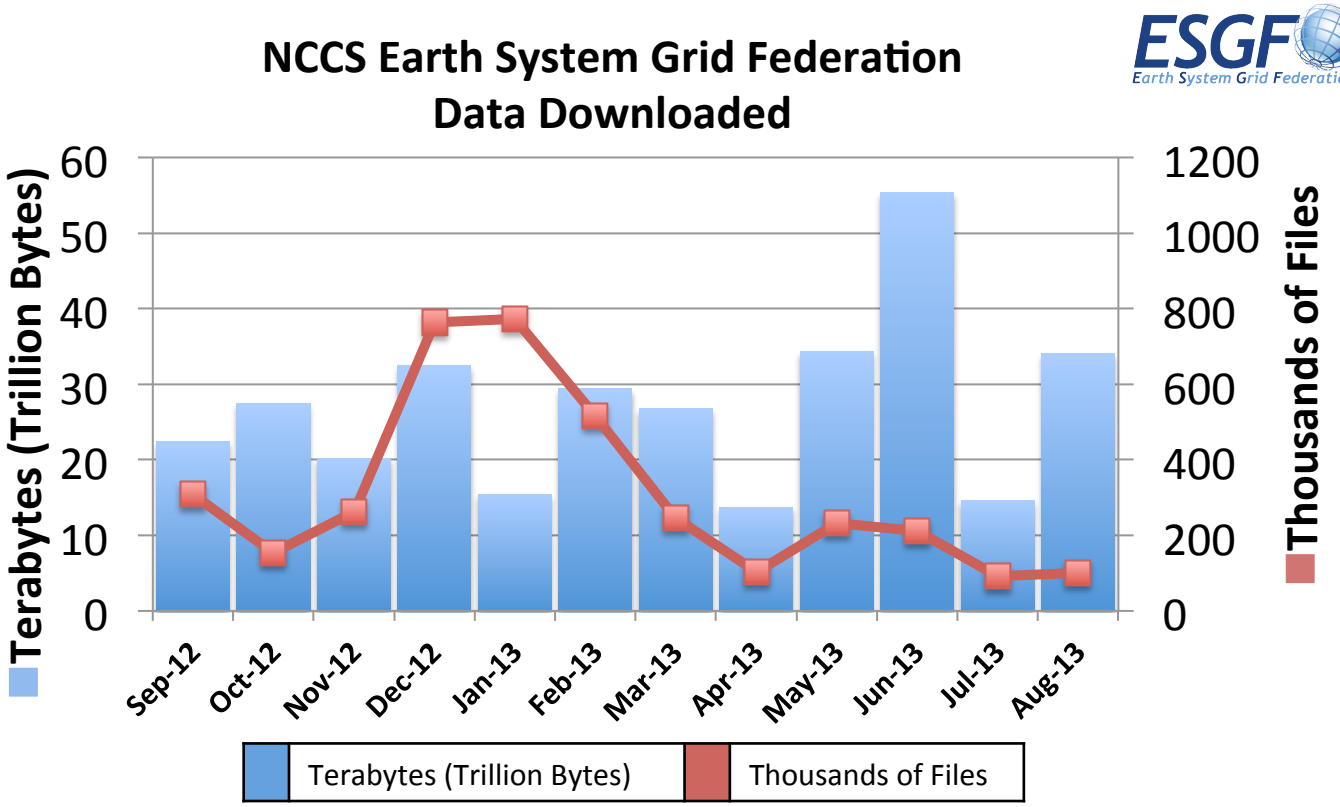

- The **NCCS Data Portal** serves data in CF-compliant format to support these Earth System Grid Federation Projects:
  - **CMIP5**: Long-term NASA GISS simulations, and decadal simulations from NASA's **GMAO**; **NOAA NCEP**; and **COLA** (Center for Ocean-Land-Atmosphere Studies).
  - **Obs4MIPs**: selected satellite observations from NASA's **GPCP**, **TRMM**, **CERES-EBAF, Terra MODIS** and **MISR aerosol**.
  - **Ana4MIPs**: analyses from NASA/GMAO's Modern Era Retrospective-Analysis for Research and Applications (**MERRA**).
  - **NEX-DCP30**: bias-corrected, statistically downscaled (0.8-km) CMIP5 climate scenarios for the conterminous United States.

Dataportal Utilization – File Downloads

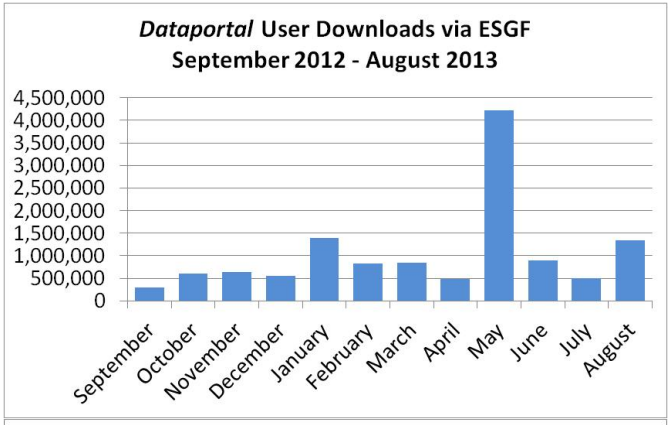

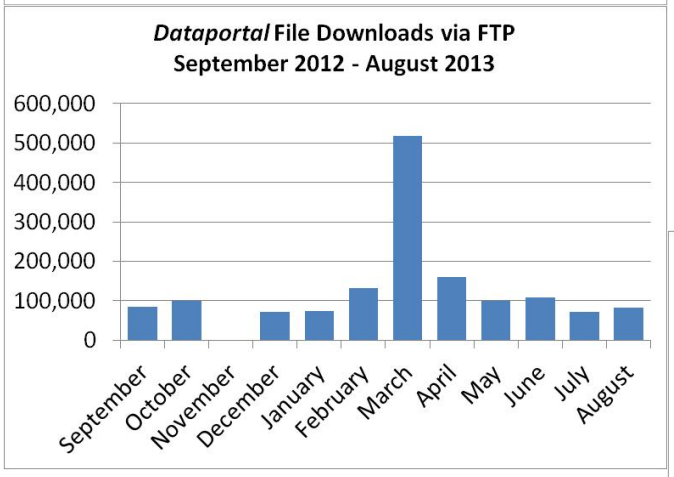

Download Mechanism	Downloaded Data Files			Available Data (TB)
	June	July	August	
Web	5,126,271	7,818,789	3,440,985	91
ESGF	893,415	495,401	1,339,053	154
FTP	108,780	71,487	82,640	144
GDS	2,879,037	2,813,855	3,590,158	76

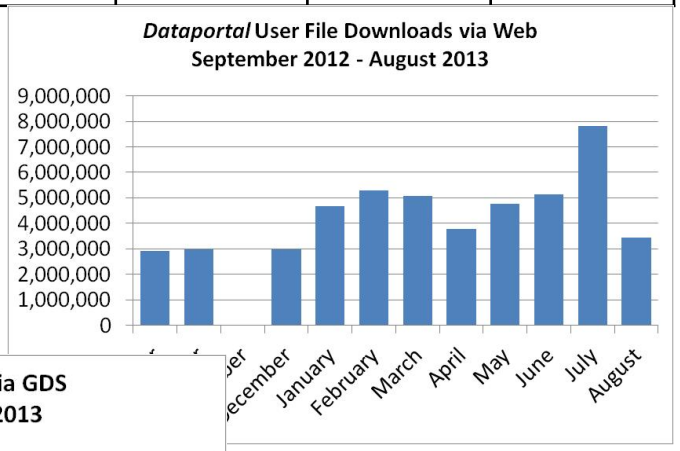

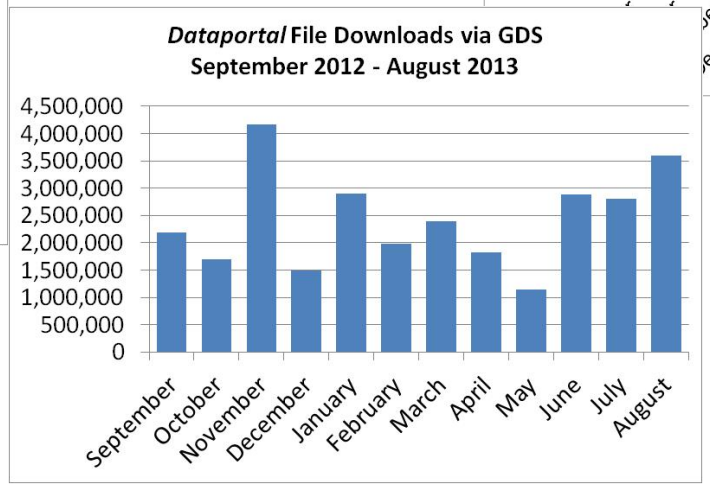

Dataportal Utilization – Data Accessed

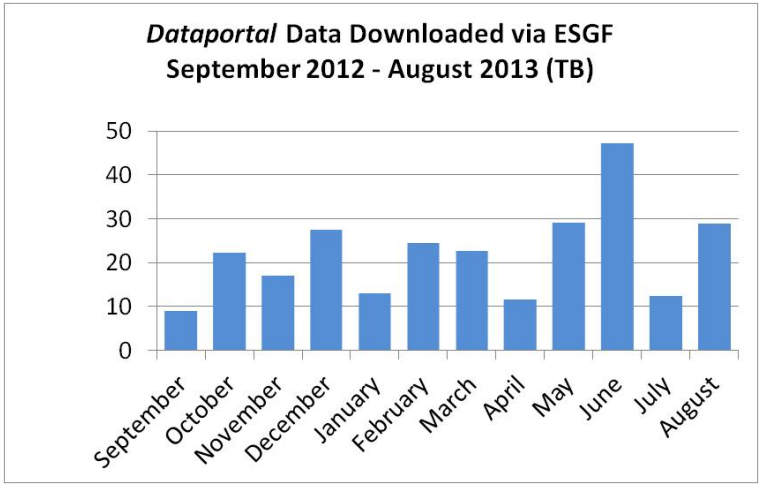

Download Mechanism	Data Accessed (TB)			Available Data (TB)
	June	July	August	
Web	2.7	2.1	21.4	91
ESGF	47.2	12.5	29.0	154
FTP	8.6	5.5	4.9	144

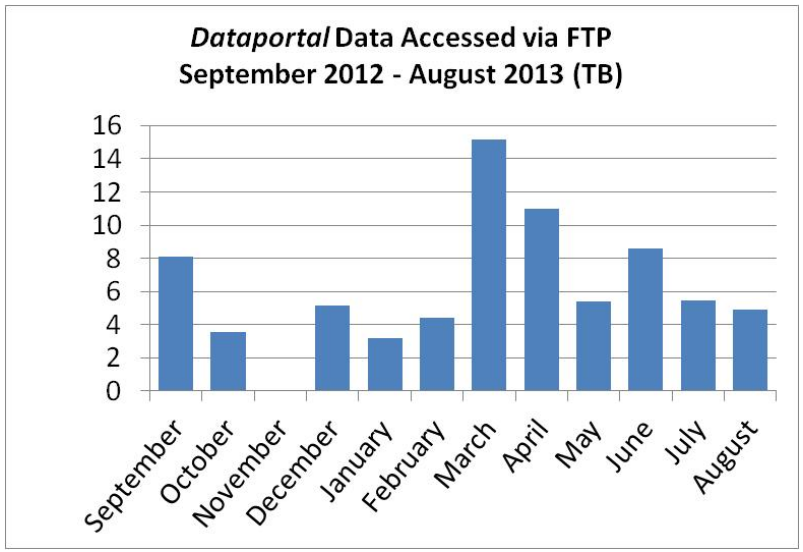

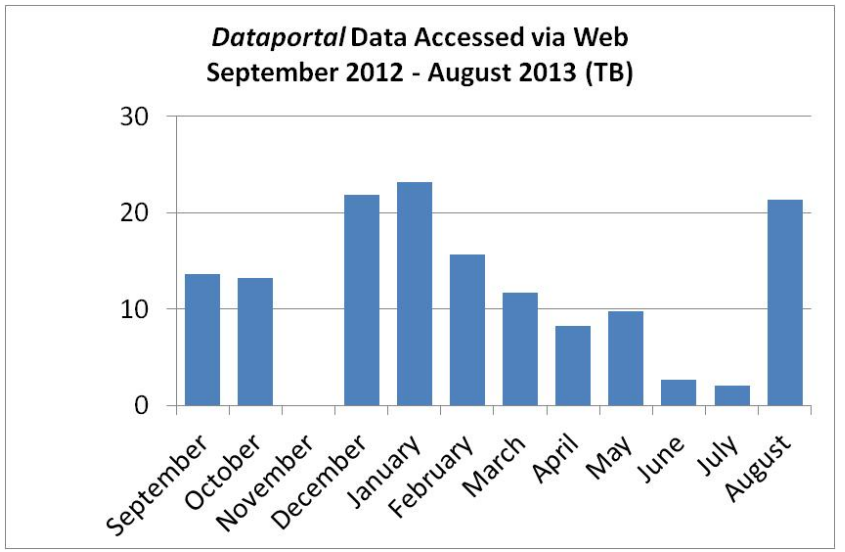

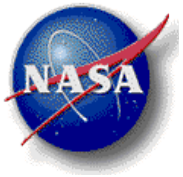

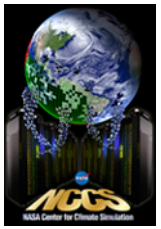

Some Discover Updates Slides (SCU8's Intel Sandy Bridge) from September 25, 2012 NCCS User Forum

Discover SCU8 Sandy Bridge: AVX

- The Sandy Bridge processor family features:

  Intel Advanced Vector eXtensions

- Intel AVX is a wider, new 256-bit instruction set extension to Intel SSE (Streaming 128-bit SIMD Extensions), hence higher peak FLOPS with good power efficiency.
- Designed for applications that are floating point intensive.

Discover SCU8 Sandy Bridge: User Changes

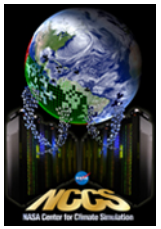

- Compiler flags to take advantage of Intel AVX (for Intel compilers 11.1 and up)

  -xavx:

  - Generate an optimized executable that runs on the Sandy Bridge processors ONLY

  -axavx –xsse4.2:

  - Generate an executable that runs on any SSE4.2 compatible processors but with additional specialized code path optimized for AVX compatible processors (i.e., run on all Discover processors)
  - Application performance is affected slightly compared to with "-xavx" due to the run-time checks needed to determine which code path to use

Sandy Bridge *vs.* Westmere: Application Performance Comparison – Preliminary

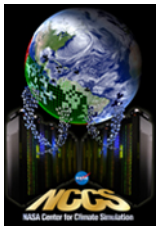

Sandy Bridge Execution Speedup Compared to Westmere				
WRF NMM 4km	**Same executable**		**Different executable (compiled with –xavx on Sandy Bridge)**	
	Core to Core	**Node to Node**	**Core to Core**	**Node to Node**
	1.15	**1.50**	**1.35**	**1.80**
GEOS5 GCM half degree	**Same executable**		**Different executable (compiled with –xavx on Sandy Bridge)**	
	Core to Core	**Node to Node**	**Core to Core**	**Node to Node**
	1.23	**1.64**	**1.26**	**1.68**

Discover SCU8 – Sandy Bridge Nodes

- 480 IBM iDataPlex Nodes, each configured with
  - Dual Intel SandyBridge 2.6 GHz processors (E5-2670) 20 MB Cache
  - 16 cores per node (8 cores per socket)
  - 32 GB of RAM (maintain ratio of 2 GB/ core)
  - 8 floating point operations per clock cycle
  - Quad Data Rate Infiniband
  - SLES11 SP1
- Advanced Vector Extensions (AVX)
  - New instruction set (http://software.intel.com/en-us/avx/)
  - Just have to recompile

Sandy Bridge Memory Bandwidth Performance

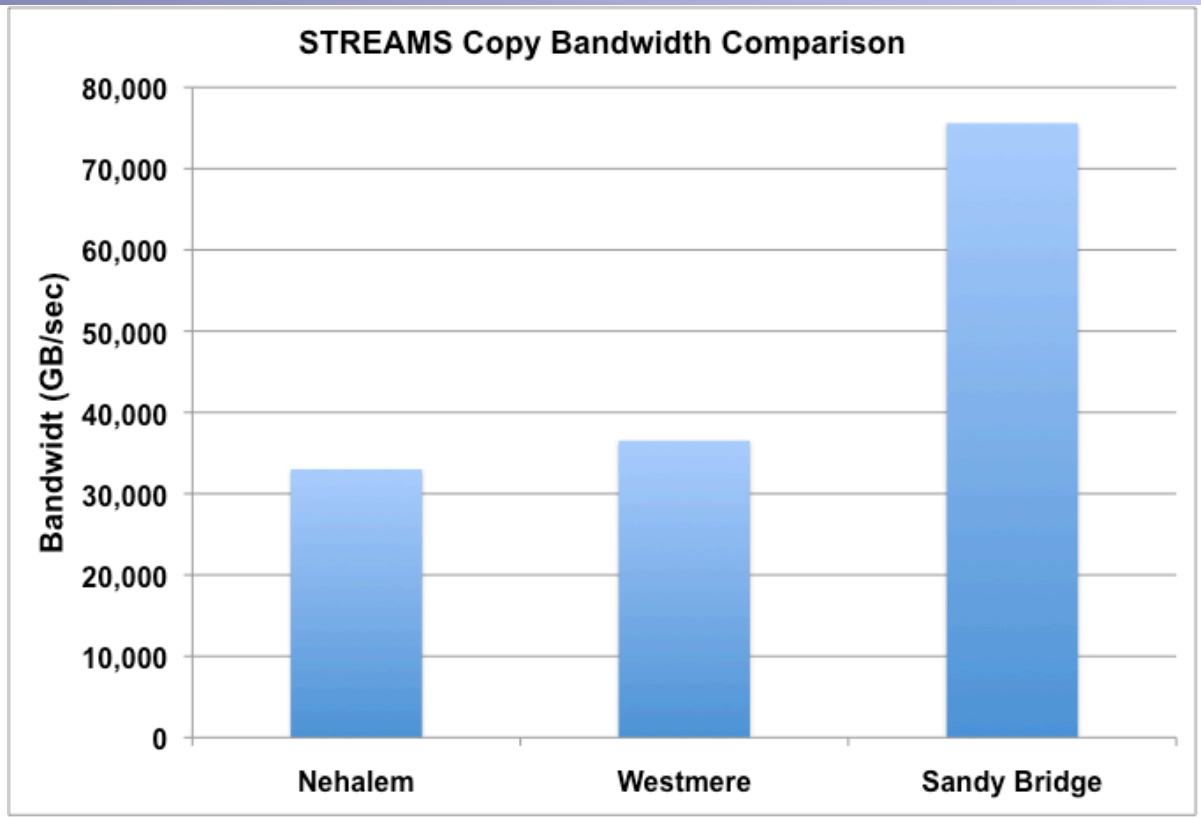

- STREAMS Copy Benchmark comparison of the last three processors
  - Nehalem (8 cores/node)
  - Westmere (12 cores/node)
  - SandyBridge (16 cores/node)

SCU8 Sandy Bridge
Finite Volume Cubed-Sphere Performance

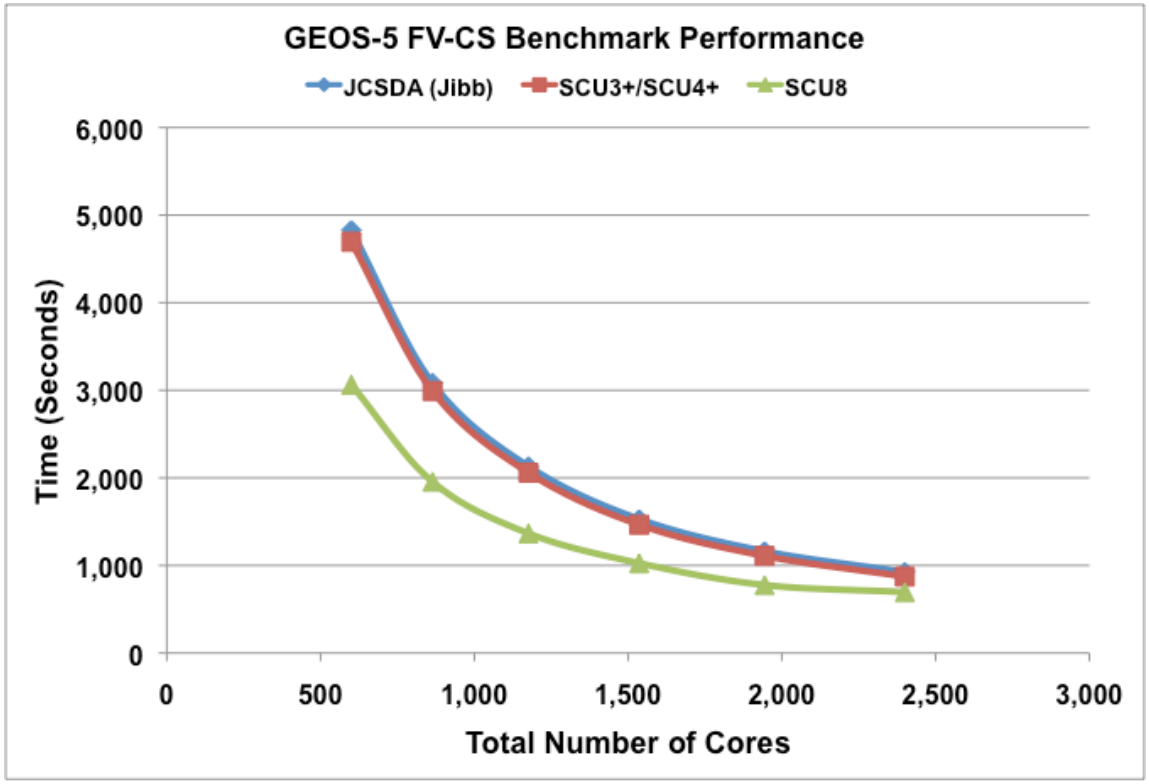

JCSDA (Jibb): Westmere

Discover SCU3+/SCU4+: Westmere

Discover SCU8: Sandy Bridge

- Comparison of the performance of the GEOS-5 FV-CS Benchmark 4 shows an improvement of 1.3x to 1.5x over the previous systems' processors.