



CISL Update Operations and Services

CISL HPC Advisory Panel Meeting 21 October 2010

**Anke Kamrath
anke@ucar.edu
Operations and Services Division
Computational and Information Systems Laboratory**

Overview

- **Staff Comings and Goings in OSD**
- **Updates:**
 - *NWSC-1 RFP Update*
 - *GLADE Deployment Status*
 - *Lynx Status*
 - *Science and Supercomputers*
 - *U Wyoming Update*
 - *Export Controls*
 - *Allocation Timeline*
 - *Rolling out Storage Allocations*



OSD Staff Comings and Goings...

- **Changes**

- *Departures*

- *Juli Rew (retired May)*
 - *Sean McCreary (DSG) (left June)*
 - *Lynda McGinley (DSG) (retiring October 2010)*
 - *Ginger Caldwell (retired August)*
 - *Russell Gonsalvez (SSG) (left September)*
 - *Dan Lagreca (DASG/VAPOR) (left September)*

- **New Staff**

- *User Services Section Head – David Hart (August)*
 - *New DB Person – Rob Perry (July)*
 - *New DSG Group Leader - Mark Bradford (July)*
 - *New Lustre Expert - Junseong Heo (June)*
 - *New DSS/TIGGE Staffer – Thomas Cram (August)*
 - *New Student Assistant - Ben Golden (Sept)*

- **Openings**

- *UCAR Security Position*
 - *CISL/NWSC Security Position*
 - *3 User Services/Help-desk Positions*
 - *SSG Positions (SEIII, SEII)*
 - *VAPOR Program (SEI)*



NWSC-1 RFP Update

- **SAP (Science Advisory Panel)**
 - *Recommended Benchmark Suite and Workload Distribution*
 - *Reviewed by CISL Council*
 - *SAP on hiatus until proposals come back in 2011*
- **TET (Technology Evaluation Team)**
 - *Packaging Benchmarks*
 - *Refining RFP*
 - 3 Pieces
 - *Compute*
 - *DAV++*
 - *Central Filesystem*
 - *Archive piece split out. Considering staggering from Compute procurement (e.g., do in 2014 and extend AMSTAR for now).*
 - *Vendor NDA's this month and meetings at SC next month*
 - *RFP release in Dec 2010*
- **BET (Business Evaluation Team)**
 - *Strategizing on how to structure 3 elements of proposal and to allow bifurcation if necessary*
 - *Reviewing AMSTAR options versus new Archive RFP*



Capacity and Capability Benchmarks

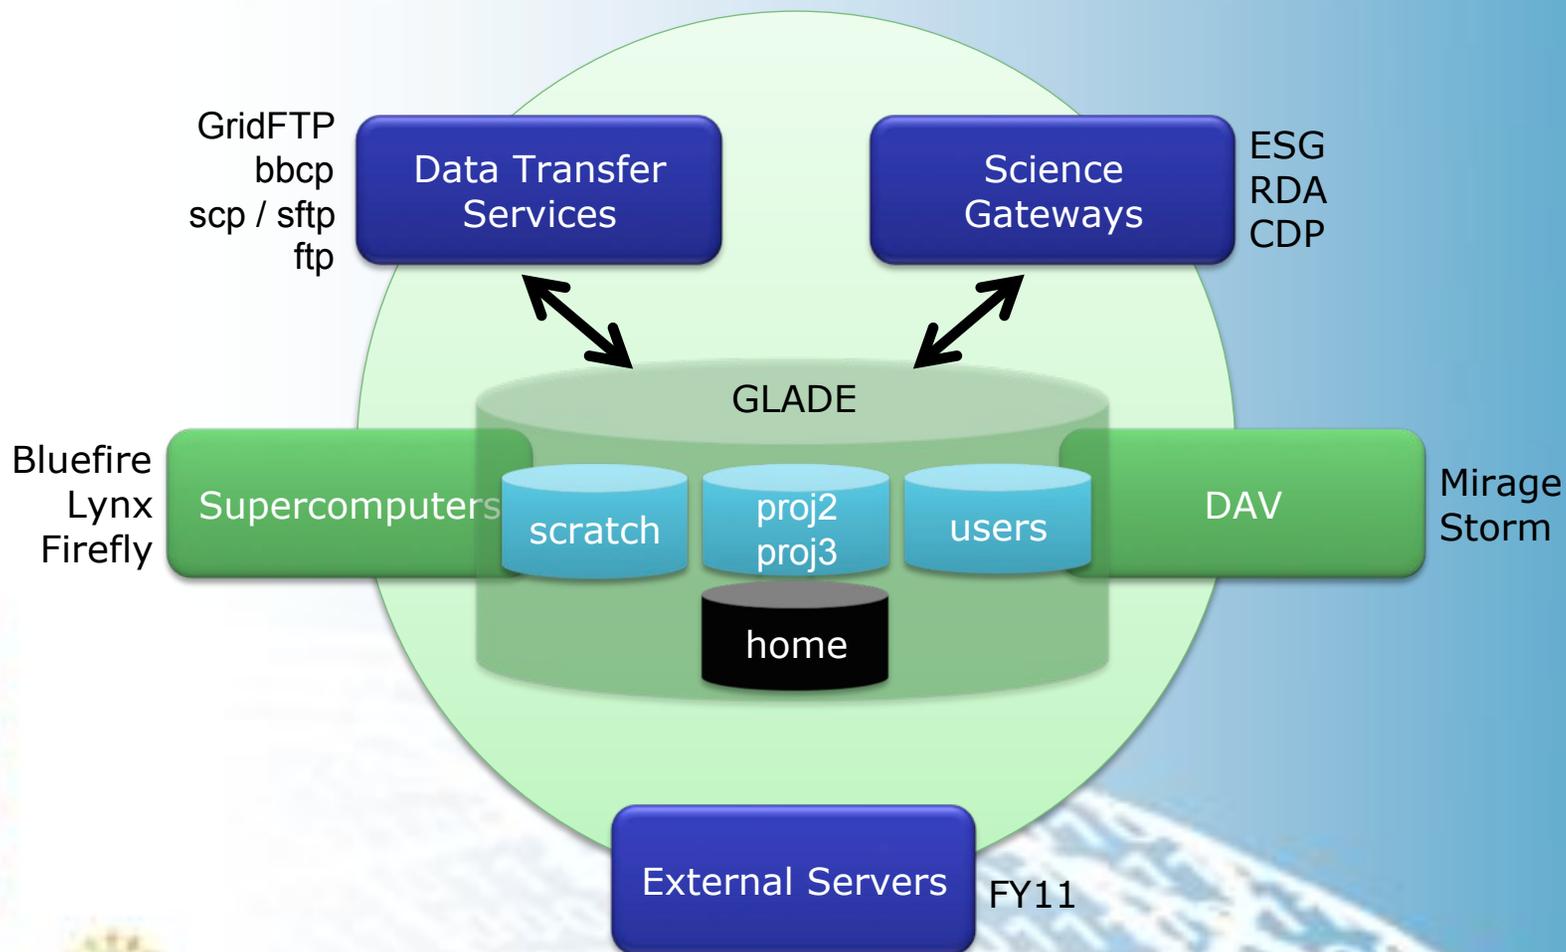
Capacity benchmarks and weighting factors:

- WRF 3.2 (20%) - 2.5 km CONUS
- DART (10%) - Configuration TBD
- CESM1 (40%)
 - FV core 0.47 x 0.63
 - WACCM4 1.9 x 2.5 x 90L
- LES (10%) - 1024³
- MuRAM (10%) - 2x10⁹ points
- PMCL3D (10%)

Capability benchmarks

- NRCM - 1.33 km
- CESM1 - HOMME dycore + CAM5 physics (res TBD)
- LES - 4096³

GLADE Data Architecture Deployed



GLADE User Data Spaces

Computational Laboratory

scratch
204TB

users
70TB

project
554TB

All users get access	By request during allocation cycle	By request during allocation cycle
Included in compute/dav GAUs	Charged for actual usage	Charged for allocation per month
Data older than 14-30 days purged	Data older than 3 months scrubbed (optional)	No CISL data management
No CISL backups	No CISL backups	No CISL backups
No quota	300GB, 500GB, 1 TB quotas	> 1TB
SHORT DURATION	MEDIUM DURATION	LONG DURATION

home
10TB

All users get an allocation
Included in compute/dav GAUs
No scrubbing
CISL provides backups
10GB, 20GB quotas
DURATION OF ACCOUNT

CISL
Computational Laboratory



GLADE Project Update

- **Completed**
 - *All storage hardware has been brought online*
 - *Installation of initial data transfer gateways*
 - *GLADE access to all HPC resources*
- **To be completed – [November 1]**
 - *Configuration of ESG & RDA science gateways*
 - *Expansion of data transfer services to two nodes*
 - *Implementation of new /glade/home backup services*
 - *Consolidation of DAV, lynx home directories into /glade/home*
- **To Be Completed – [1Q CY2011]**
 - *Consolidation of bluefire home directories into /glade/home*

GLADE FY11 Futures

- **Develop charging algorithms and code**
- **Develop allocation process**
- **Develop usage metrics and tools**
- **Develop scheduling algorithms and resource management tools**
- **Expand into WAN based services for divisional servers**



Lynx – Cray XT5m

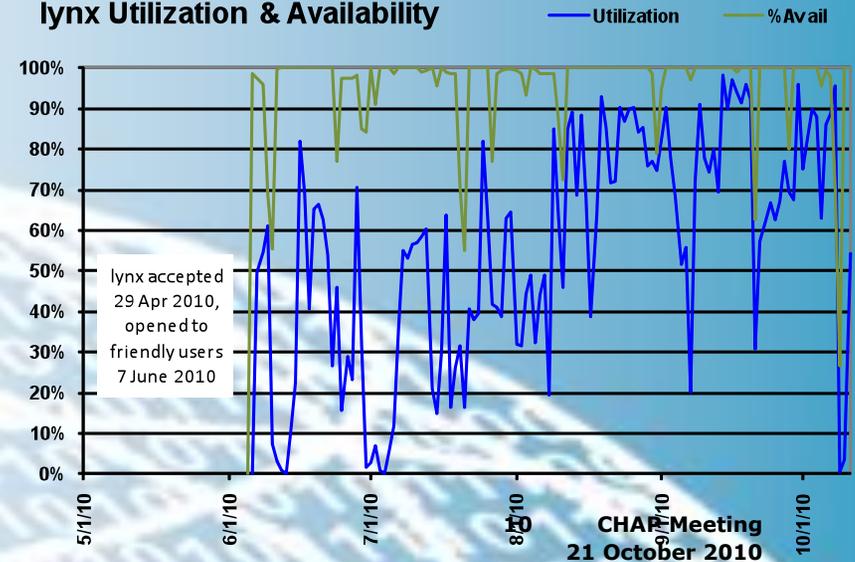
- **76 compute nodes / 912 2.2 GHz AMD cores / 8.0 teraflops peak**
- **“Test” system, since August: 70% utilization, 97% availability**
- **Friendly users (are, overall, happy – most often complaint: other users are hogging the machine)**

- 1 ACD
- 13 CGD
- 3 HAO
- 26 IMAGE
- 13 MMM
- 2 RAL
- 6 UNIV

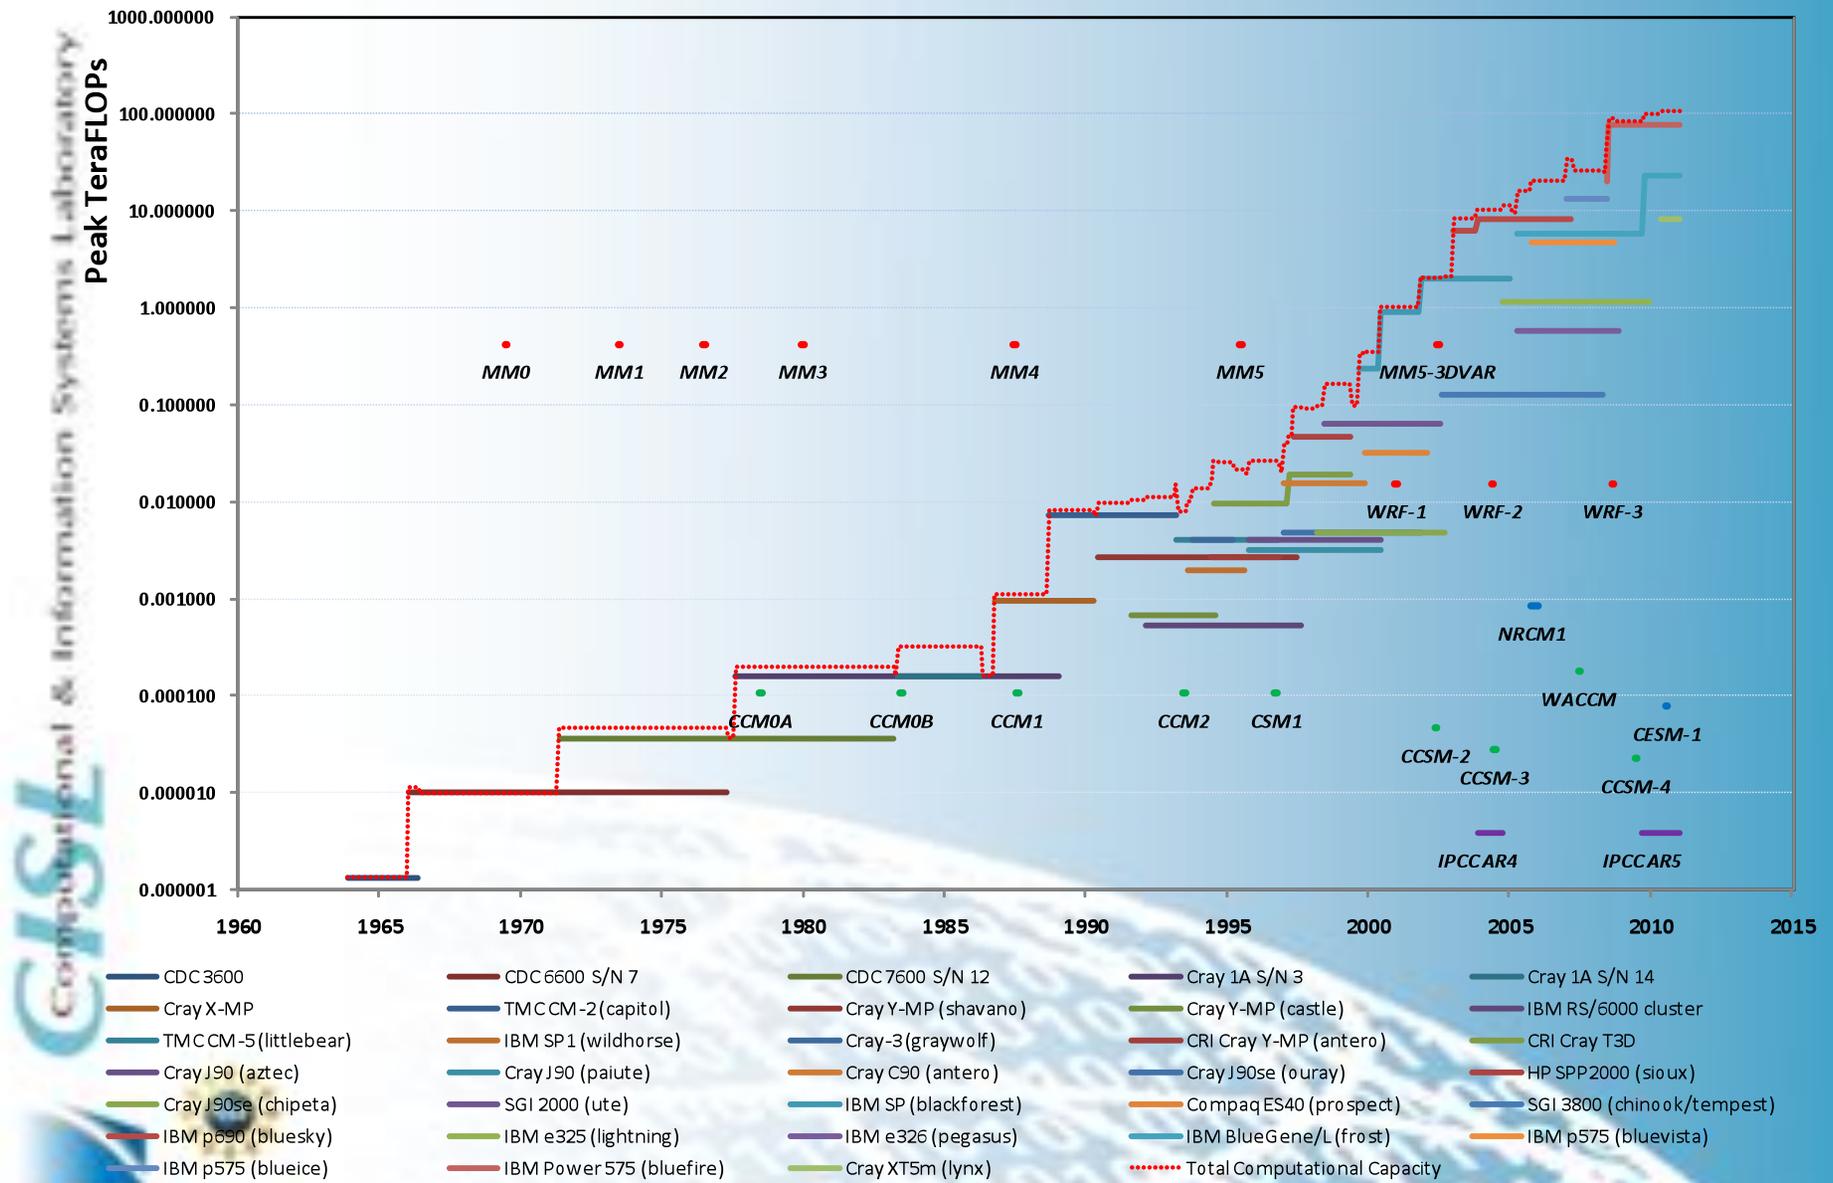


- **/glade filesystems via Cray’s Data Virtualization Service (DVS)**
- **Moab/Torque batch subsystem provides fair-share & backfill scheduling**

lynx Utilization & Availability



Supers & Science at NCAR



Interactions with U Wyoming Underway

- **U Wyoming User POC identified – Bryan Shader**
- **Working with 4 early users/apps**
 - 1. *Mohamed Piri (carbon sequestration)*
 - 2. *Po Chen (seismic wave propagation)*
 - 3. *Stefan Heinz (CFD/combustion simulation)*
(visiting NCAR soon)
 - 4. *Dimitri Mavriplis (CFD) (currently visiting NCAR)*
- **Using some CISL director's reserve GAUs to help early U Wyoming users**
- **Planning IT field trip visit from NCAR to U Wyoming December 2-3.**
 - *Discussion of NWSC System Architecture and possible integration opportunities (e.g., central filesystem, security) for UWy*
 - *Computational Science Applications/Users*
 - *User Support and Allocation Models*
 - *Recruiting at NWSC (e.g., from Community College?)*



Export Controls Update

- **Recall April Information on new requirements**
 - *No "services to"*
 - Embargoed Nations (Cuba, N. Korea, Iran, Sudan, Syria)
 - Prohibited end-user (six lists)
 - *Services include compute, storage, phone, email, ..*
- **UCAR hired a consultant to advise NCAR. Current understanding:**
 - *If a user (from these countries) is in US on a student visa and is using resources/services new requirements don't apply*
 - *For users from outside US – self identifying their country of origin (e.g., via registration like RDA does) is fine. No further verification is needed.*
 - *Waiting for further "Formal" guidance from UCAR/consultant but at this point it appears that CISL has little exposure.*
 - *Some software distribution sites (e.g., VAPOR, NCAL) may need to create registration page.*



Allocation Timeline

- **Due to step function in resources in 2012 and massive over-request (based on Bluefire GAUs) – only annual allocations being made.**
- **CHAP Oct 10 Allocations:**
 - *For 1 year (Jan 1, 2011-Dec 31, 2011)*
 - *Bluefire only*
- **CHAP Apr 11 Allocations:**
 - *For 1 year (Jul 1, 2011-June 20, 2012)*
 - *Bluefire only*
- **CHAP Oct 11 Allocations**
 - *For multi-year (but with annual computing plans)*
 - *Bluefire decommissioning June 2012.*
 - *Allocations for Bluefire (6 mos) and new NWSC-1 System*



Rolling out Storage Allocations

- **Disk and Archive systems represent 1/3 of the NWSC investment— *nearly \$13 million.***
- **Three goals**
 - *Ensure that use of scarce and costly resources are directed to the most meritorious projects*
 - *Minimize user hurdles and reviewer burden*
 - *Build on familiar process for requesting HPC allocations*
- **Thus, storage allocations and policies will focus on the small number of largest-scale consumers**
 - *All projects will get "default" archive allocation, on the order of 20 TB*
 - *May only review/award project disk space above a minimum threshold.*



Storage Allocations Timeline

- **NCAR projects will request disk allocations as part of compute requests in early 2011.**
- **CSL and CHAP processes will ramp up to formal allocations for NWSC in 2012.**
 - *CHAP, Oct. 2010 — Submitters required to identify archive needs above 10 TB.*
 - *CSL, Jan. 2011 — Submitters identify archive needs >10 TB, asked to describe data sharing plan.*
 - *CHAP, Apr. 2011 — Submitters will be able to request and justify project space on GLADE. More substantial justification of archive needs over 10 TB.*
 - *CHAP, Oct. 2011 — Continued refinement of request instructions and review criteria for storage needs.*
 - *CSL and CHAP, spring 2011 — First official disk and storage allocations for NWSC resources.*

Storage Allocation Action Items

- **Work with CHAP and CSLAP to define request instructions and review criteria**
 - *Efficiently uses the (storage) resources*
 - *Has a well-defined data management plan*
 - *Appropriately supports the science objectives*
- **Revise storage charging for greater transparency and alignment with costs**
- **Complete work on new Accounting Management System to support storage allocations and accounting model**



Questions and Discussion