



CISL Update

Operations and Services

CISL HPC Advisory Panel Meeting

29 April 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*
 - *Data Services Vision for NWSC*
 - *Preparing for the Changing Data Workflow*
 - Archival Migration: MSS to HPSS
 - Cray XT5 Rack – Lynx
 - GLADE Deployment (Presentation from Pam)
 - Storage Accounting (Discussion with CHAP)
 - *Export Controls*
 - *Users and Allocations*
 - University Publications in 2009
 - NWSC Allocation Plans – WNA (Wyoming NCAR Allocations)
 - Accessing Science Merit of WNA proposals
 - *Discussion with CHAP*
 - *New TIGGE Award*

Staff Comings and Goings...

- **Changes**

- *Retirements*

- *Ed Arnold, DSG (April)*
 - *Juli Rew, CSG (end of May)*
 - *Julie Chapin, WEG (June)*
 - *Leif Magden, WEG (April)*



- *ESS - Jasen Boyington, Co-location Manager*
 - *HSS/CSG - Rory Kelly moved from TDD to CSG (replacing Michael Page)*
 - *HSS/DASG - Dan Lagreca on Vapor project*
 - *NETS - Blake Caldwell, Network Engineer 1*



- **Openings**

- *Elevating User Service role to Section Level*
 - *Job Opening Posted. Interviews underway.*
 - *DBA II Open*
 - *DSG Group Leader (Lynda McGinley returning to technical role)*



NWSC-1 RFP Update

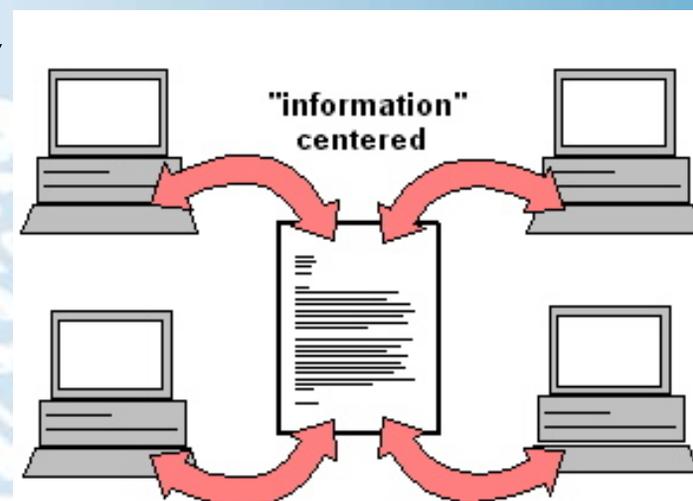
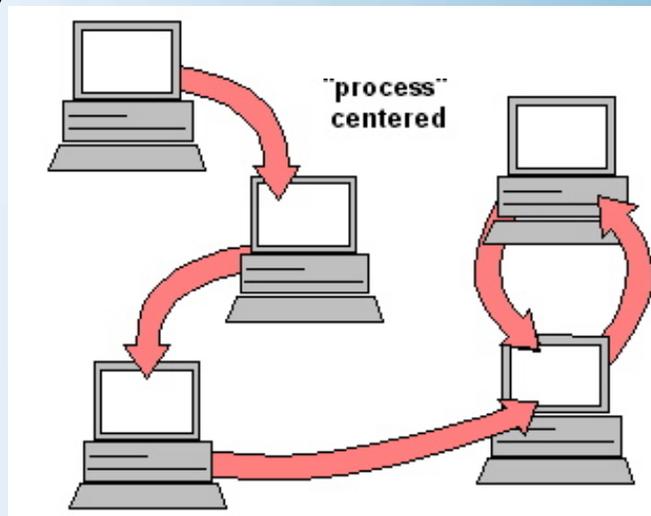
- **SAP (Science Advisory Panel)**
 - *Met mid-April 2010*
 - *Refining Science Requirements*
 - Via questionnaire (see handout)
 - Meeting in May
 - CHAP: Bert is representing CHAP; will share questionnaire
 - *Analyze results and provide to TET*
 - *Identify key Benchmark applications*
- **TET (Technology Evaluation Team) and BET (Business Evaluation Team) kicking off their efforts soon.**

A Changing Scientific Data Workflow

"Process Centric"

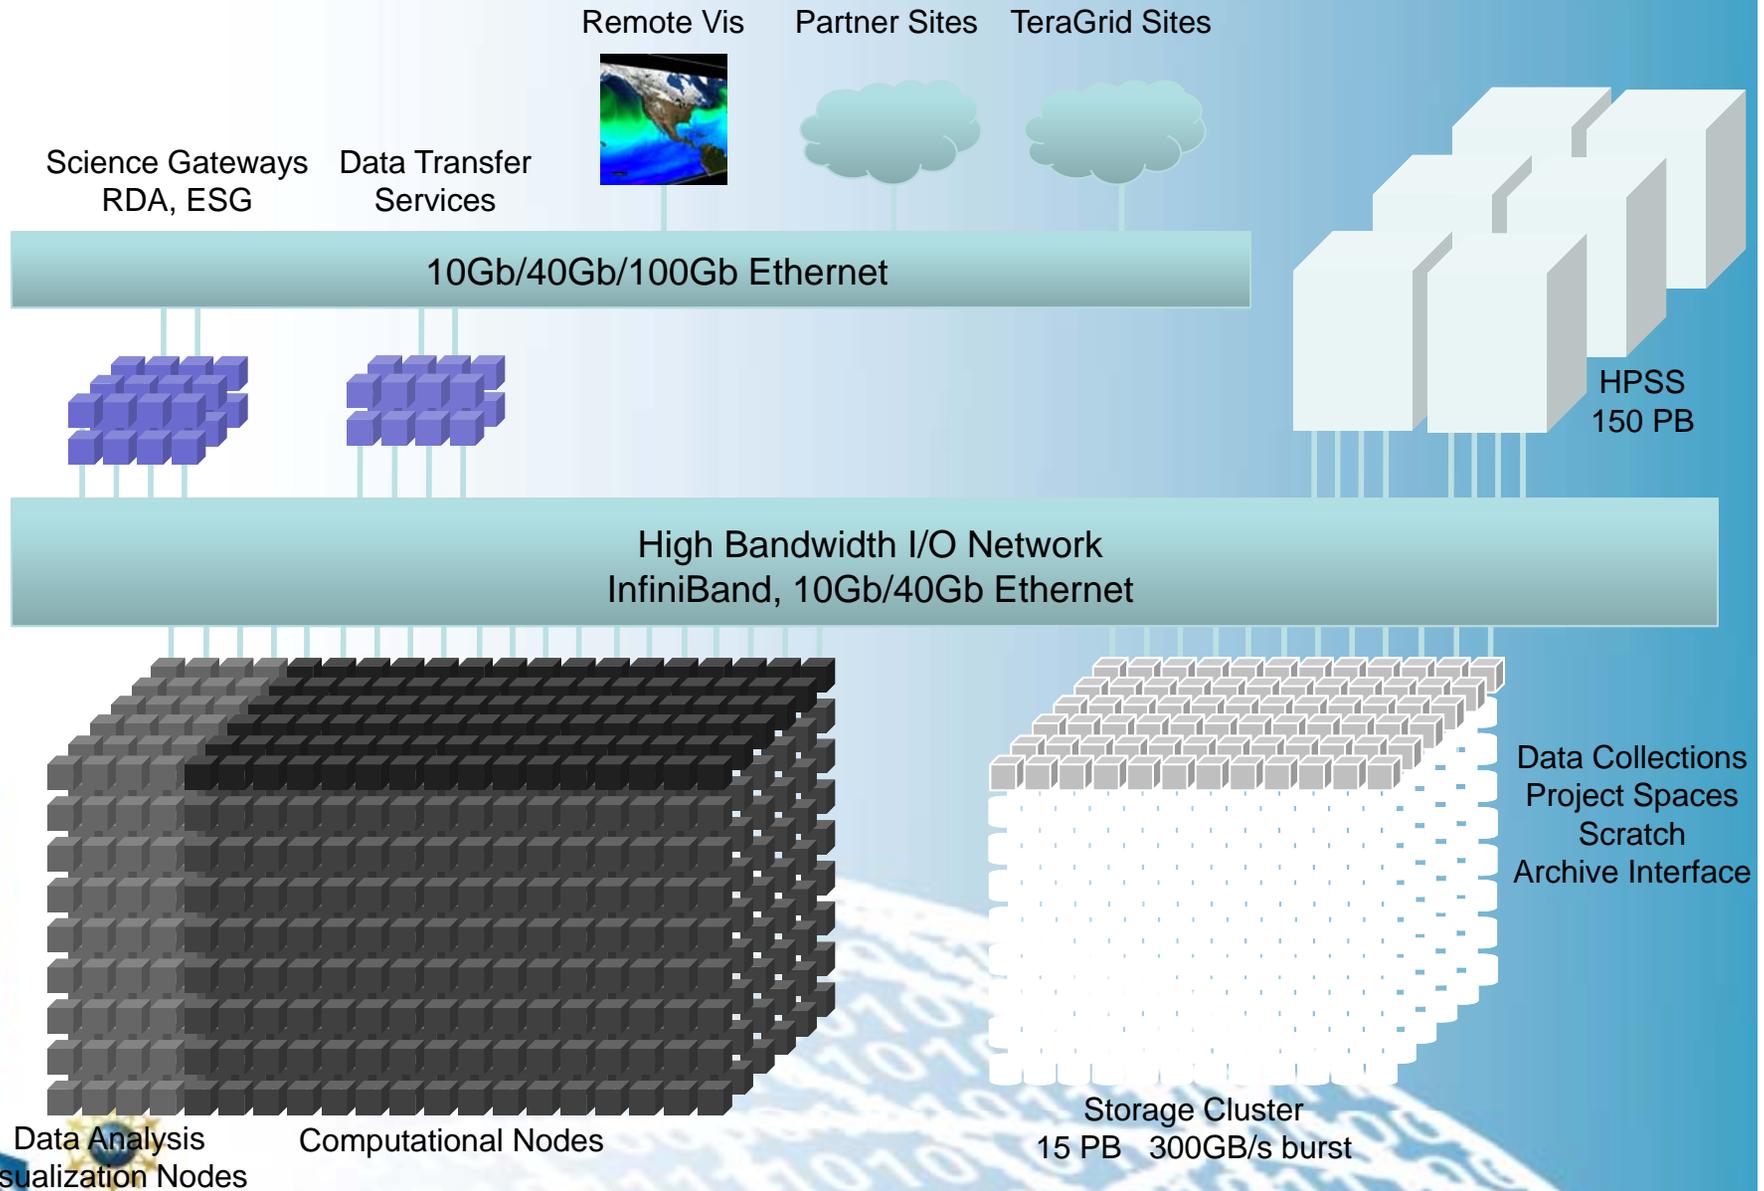
to

"Information Centric"



NWSC Conceptual Data Architecture

Computational & Information Systems Laboratory



Challenges and Trade-offs for Procurement

- How to split up \$\$.
- \$30M Compute/Disk; \$6.4M Archive
- One approach is to set I/O BW and Filesystem Size. See what FLOPS come out.

Config	I/O Sus	I/O Max	Size	STORAGE (\$M)	PFLOPS (max \$/PF)	PFLOPS (avg \$/PF)	PFLOPS (min \$/PF)	COMPUTE (\$M)
1	225 GB/s	300 GB/s	15 TB	\$10.12	0.43	0.83	1.39	\$19.88
2	225 GB/s	300 GB/s	6 TB	\$9.03	0.45	0.88	1.46	\$20.97
3	112 GB/s	150 GB/s	15 TB	\$7.51	0.49	0.94	1.57	\$22.49
4	112 GB/s	150 GB/s	6 TB	\$6.67	0.50	0.98	1.63	\$23.33



Preparing for the Changing Data Workflow

- **MSS-HPSS Migration**
- **GLADE (Globally Accessible Data Environment)**
 - *Pam to present*
- **Storage Allocations**
 - *Discussion during Pam's presentation*
- **System and Filesystem Evaluations**
 - *GPFS*
 - *Lustre*
 - *WAN filesystems*
 - *CRAY XT5 (Lynx)*

Migrating from MSS to HPSS (Update)

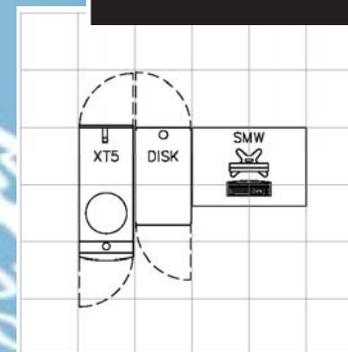
- Expansion of HPSS library from 1PB to 5 PB (capacity) in Jan 2010.
- Acquisition of 150 TB of disk for HPSS data movers (to be purchased Q2)
- Currently have ~50 users.
- Monthly meetings to consult with users on the NCAR MSS to HPSS transition.
- Migration to be complete Jan 2011
- Projects:
 - *Legacy tape project (accessing MSS data from HPSS) - on track, additional testing, code development done*
 - *Gatekeeper/load balancing - currently in evaluation phase (this is for distributing load across host clients like bluefire vs. data analysis machines to ensure that each gets a "fair share" of the HPSS traffic)*



Cray XT5 Comes to NCAR

Lynx (a small Jaguar)

- Arrived Monday, April 26.
- Test & evaluate Cray products, administration, Lustre, GLADE (GPFS/DVS), alternate batch subsystems & resource management
- Local resource comparable to ORNL "Jaguar", NICS "Kraken" & NERSC "Franklin" – participating in IPCC AR5
- **Single cabinet Cray XT5m** (8.03 TFLOPs)
 - 76 compute nodes (912 total processors)
 - 2 hex-core 2.2 GHz AMD Opteron (Istanbul) chips (12 cpus/node)
 - 16 GBytes memory per node (1.33 GB/core)
 - 2 Login nodes (4 total processors)
 - 1 dual-core 2.6 GHz AMD Opteron
 - 8 GBytes memory (4 GB/core)
 - 8 I/O nodes
 - 4 reserved for Cray system functions & LSI/Lustre filesystems
 - 2 for GPFS & Cray DVS testing (each optical 10-GigE)
 - 2 for Lustre & Cray DVS testing (each optical 10-GigE)
- **Single cabinet LSI disk subsystem** (32 TBytes)
- Cray Linux Environment, Cray, PGI & EKOPATH compilers & libraries, MPI, SHMEM & OpenMP, optimized math & scientific libraries, Lustre, MOAB/Torque software





Export Controls

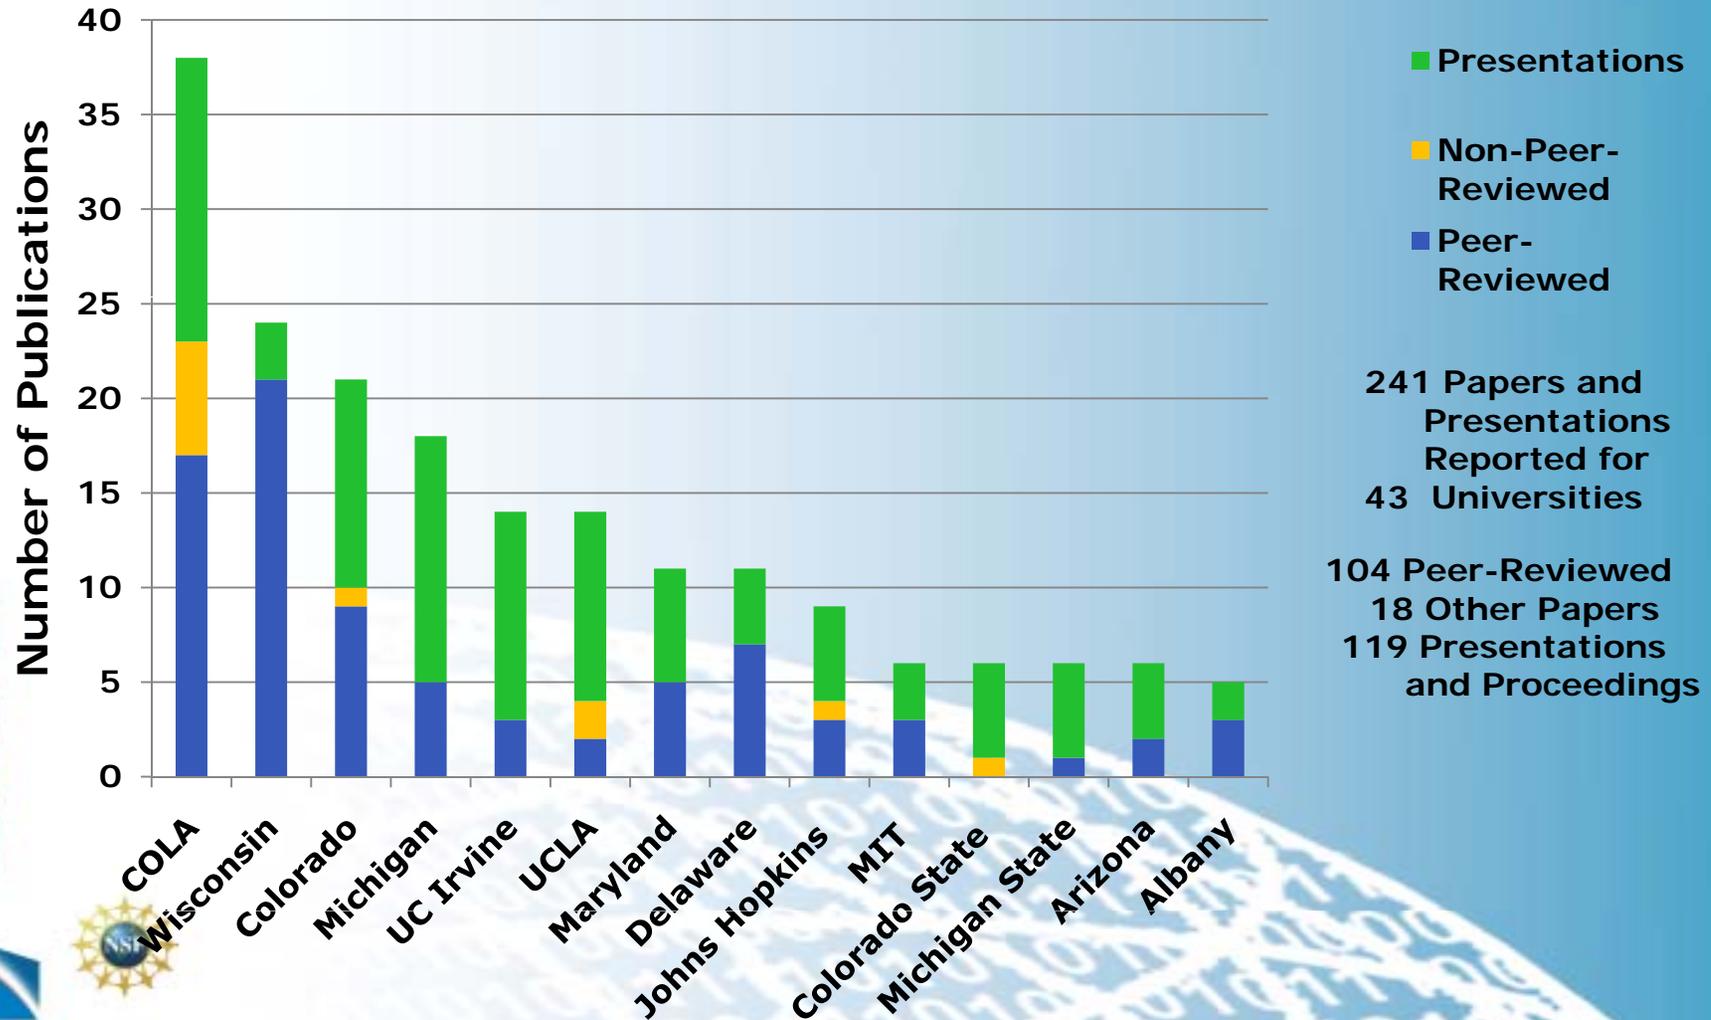
- **New requirements**
 - *No “services to”*
 - Embargoed Nations (Cuba, N. Korea, Iran, Sudan, Syria)
 - Prohibited end-user (six lists)
 - *Services include compute, storage, phone, email, ..*
- **Many universities are struggling with requirements (no consensus within TeraGrid)**
- **CISL Developing Phased Approach**
 1. *HPC / Storage Users*
 - First new users
 - Then existing users
 2. *Data Collection Users*
 3. *Software Users*



Users and Allocations

- **FY09 University Publications/Presentations**
- **Strong University Usage**
- **NWSC Allocations**
 - *New Category - Wyoming-NCAR-Allocations (WNA)*
- **Assessing Science Merit of Wyoming Allocations**

FY2009 University Publications and Presentations which used CISL Resources



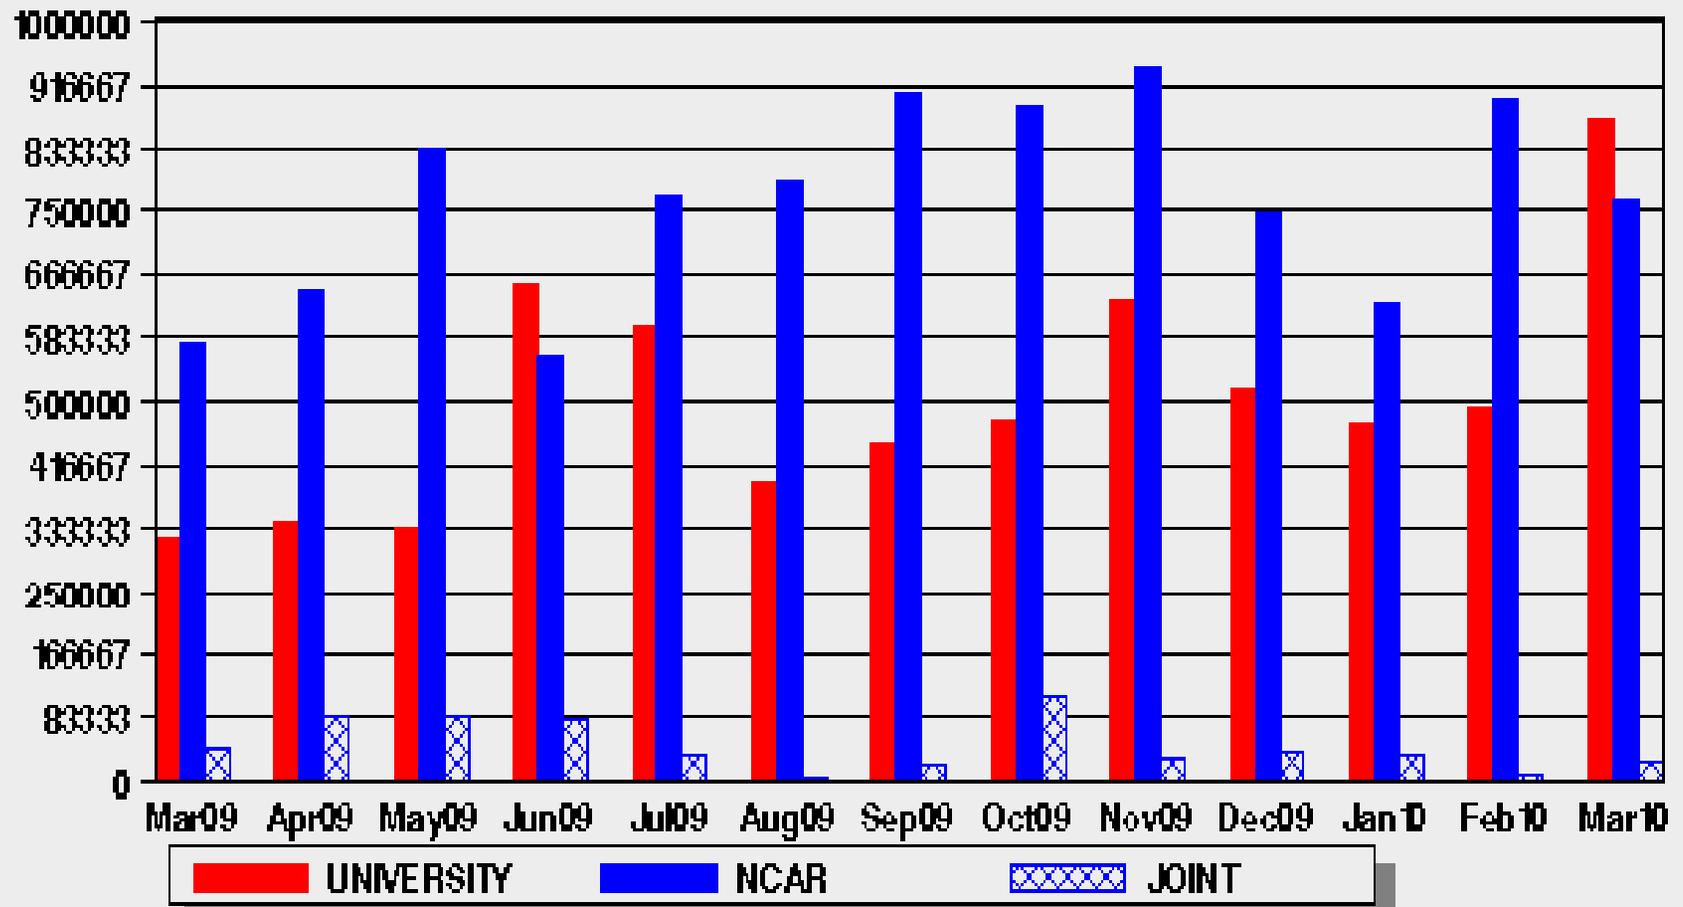
241 Papers and Presentations Reported for 43 Universities

104 Peer-Reviewed
 18 Other Papers
 119 Presentations and Proceedings

Compute Gaus Used by General Community

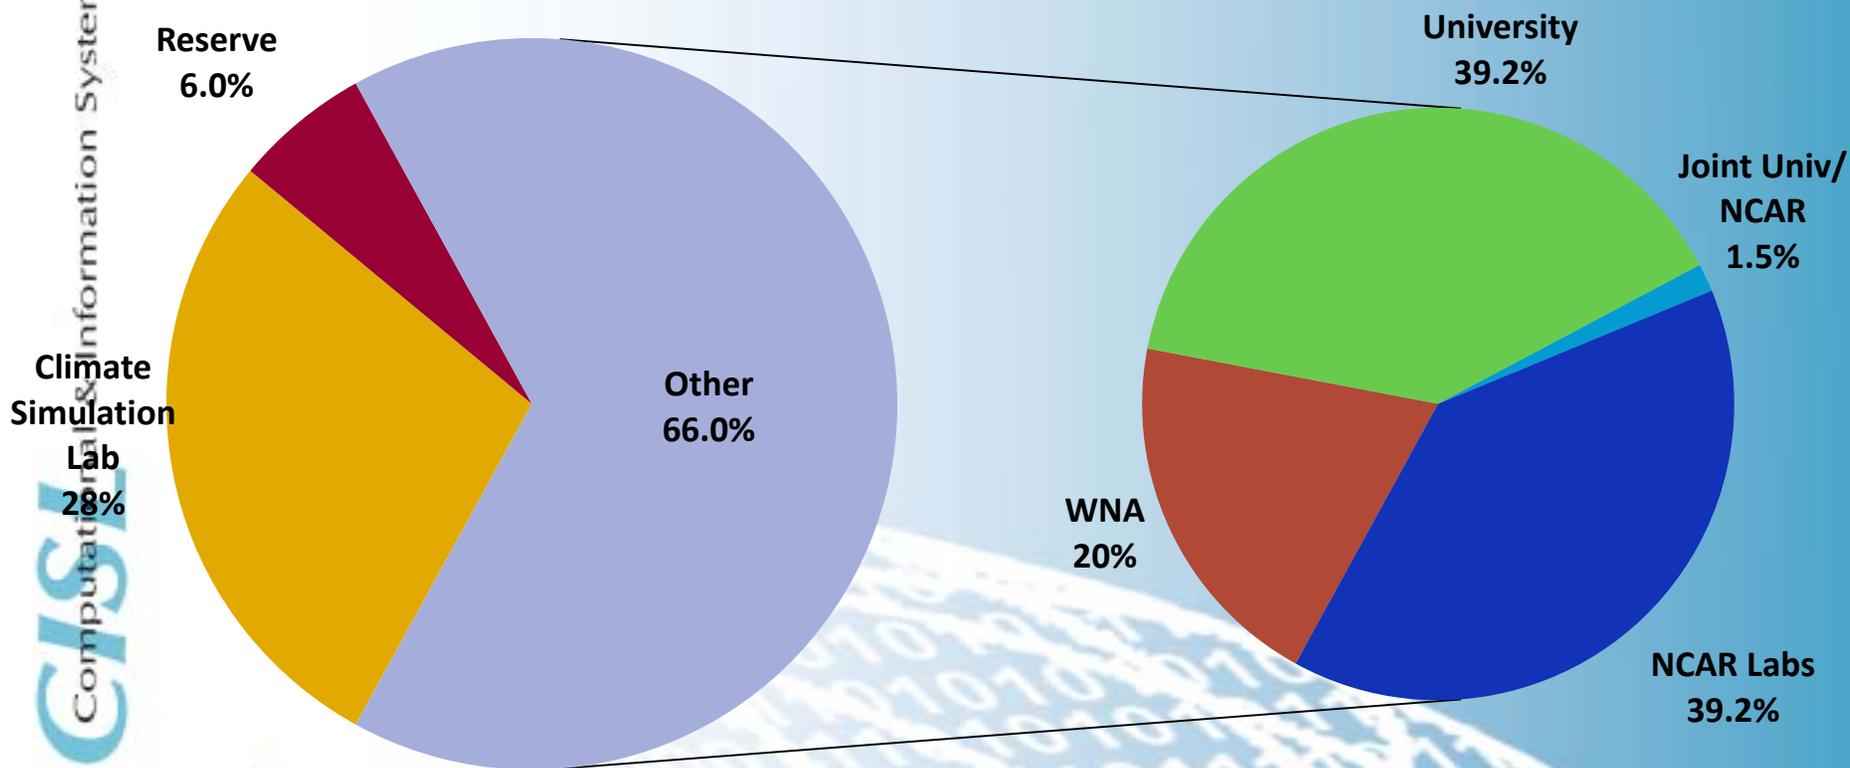
University, Joint and NCAR Projects Excluding Director's Reserve
ASD, UCC, and NCC Included

Compute — Gaus



NWSC Allocations

Wyoming-NCAR Allocations (WNA) – 20%



Computational & Information Systems Laboratory



NCAR



Assessing Scientific Merit of WNA Allocations

- Need CHAP's help as we built out WNA Process.
- ...criteria to be considered by the WNA Allocation Committee shall include the following:
 - Overall scientific merit (*If a proposal has been peer reviewed and received an NSF or other Federal agency grant, it will be deemed to have scientific merit; however, other mechanisms may be used by the WNA Allocation Committee to determine scientific merit. **Proposals not funded by NSF or other Federal agency grant, will initially be referred to the CISL High performance computing Advisory Panel (CHAP) for review and approval.***)
 - Are for research in Earth System science subject matter areas that are of substantial interest to the WNA (in identifying such areas, the WNA shall consider areas that are of significant interest to the State of Wyoming, as determined by UW);
 - Have substantial involvement of both UW and NCAR researchers;
 - Include UW researchers as the principal or co-principal investigator;
 - Strengthen UW's research capacity directly or through collaboration with other entities;
 - Directly or through collaboration, strengthens university computational science capacity in EPSCOR states; and
 - Are in new or emerging Earth System science research areas .



RDA : TIGGE Archive Access Improvements and Validation Data Portal (Ensemble weather forecast data from 10 NWP centers worldwide)

- **Two-year funding from NSF to support a SE II**
- **CISL will contribute GLADE storage and computing**
- **Service improvements**
 - ***Faster access to the 300+ TB archive***
 - For CISL HPC and web services using GLADE and HPSS
 - ***Faster turn around on user requests***
 - For subsetting and re-gridding using a Linux cluster or Lynx
 - ***Customized access to RDA weather observations***
 - For model forecast validation.





Questions and Discussion