

CISL Update

Operations and Services

CISL HPC Advisory Panel Meeting

29 October 2009

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





Overview

- **Staff Comings and Goings in OSD**
- **Computing, Storage, Data and Software Updates**
- **University Workshops**
- **Preparations for NWSC already underway**
 - *Security Review*
 - *User Surveys*
 - *Accounting/Allocation Systems*
 - *HPSS Migration*
 - *RFP Process for NWSC Procurement*

Staff Comings and Goings...

• Changes

- *Director of Operations and Services*
 - Anke Kamrath, Joined NCAR July, 2009
- *Consulting Services Group (CSG)*
 - Si Liu, consultant replaced Wei Huang (moved to VETS)
- *Computer Production Group (CPG)*
 - Ronald Reed, Operator replaced Ken Albertson (moved to RAL)
- *Data Analysis Systems Group (DASG)*
 - Kendall Southwick, Student Assistant on Vapor
- *Supercomputer Systems Group (SSG)*
 - Irfan Elahi, promoted to group head
 - Russell Gonsalves
 - Nate Rini



• Openings

- CSG Consulting Position. SE III (Mike Page Replacement)
- DASG Vapor Position. SE I opening.
- Co-location Manager Position



Computing Update

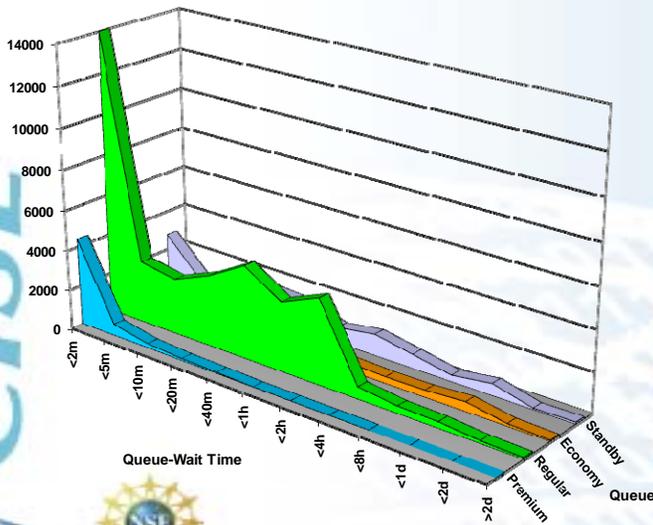
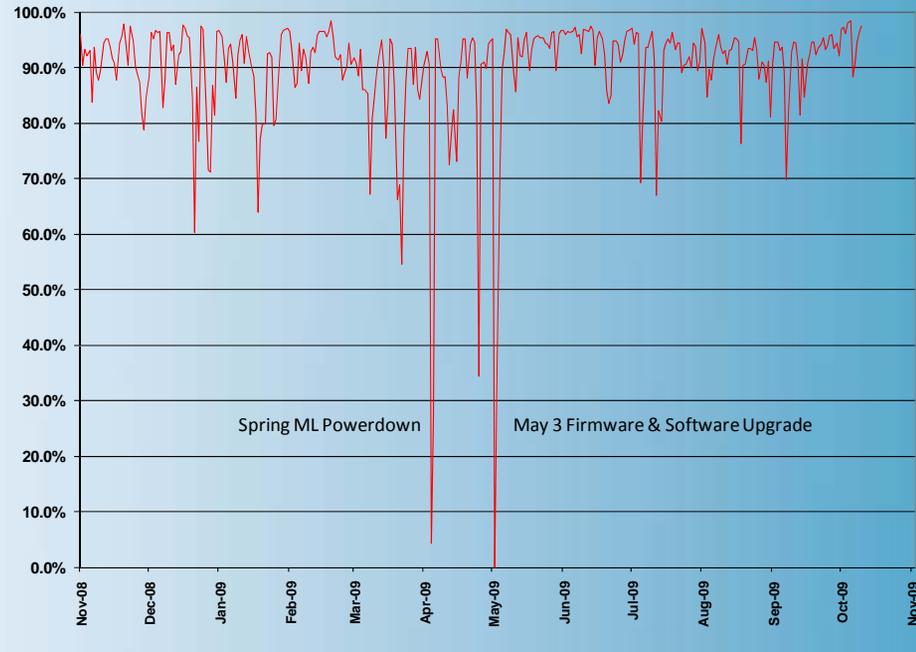
- **Strong Demand for Bluefire**
 - *2 Large memory bluefire nodes redeployed into batch pool. Bluefire now has 120 32-processor batch nodes*
 - *IPCC AR5 Ramping up (Sept 09-Dec 10)*
 - 12M GAUs needed to complete
 - Will Generate 1.3PB data
 - *High Utilization, Low Queue Wait Times*
- **Decommissioning Lightning and Pegasus**
- **3 Racks Added to Frost (BlueGene System)**
 - *TeraGrid resource (>20 Tflops)*



High Utilization, Low Queue Wait Times

- Bluefire system utilization is routinely >90%
- Average queue-wait time < 1hr

Bluefire System Utilization (daily average)



Average Queue Wait Times for User Jobs at the NCAR/CISL Computing Facility

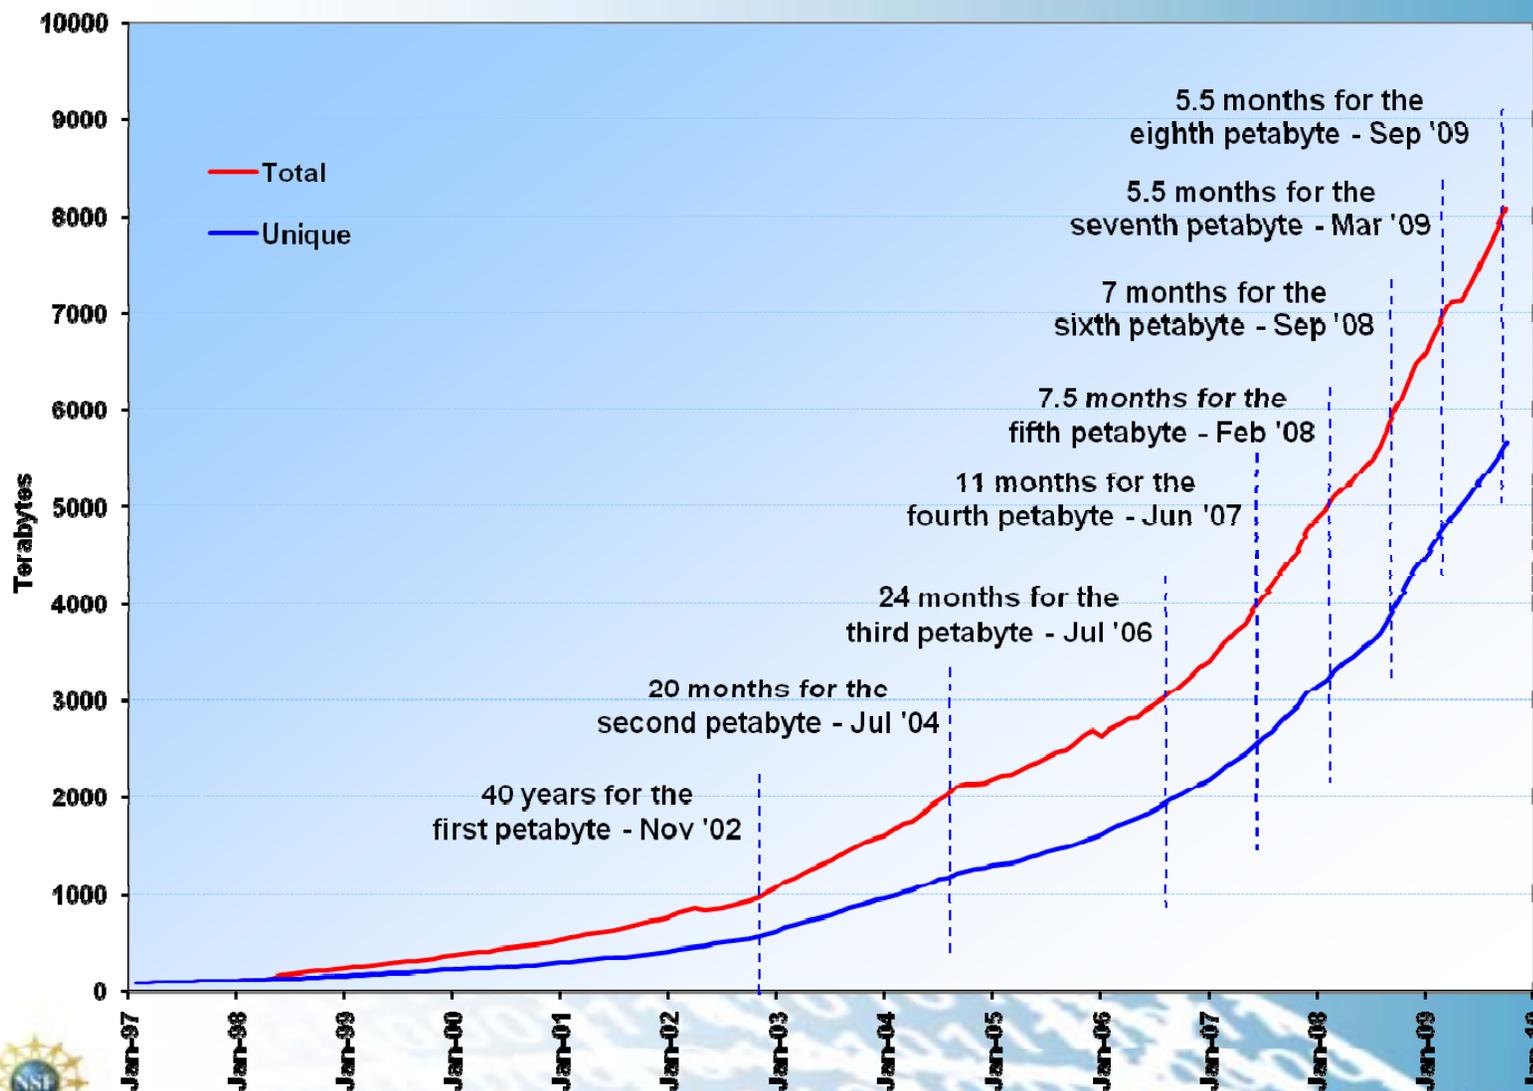
	bluefire since Jun'08	lightning since Dec'04	blueice Jan'07-Jun'08	bluevista Jan'06-Sep'08
Peak TFLOPs >	77.0 TFLOPs	1.1 TFLOPs	12.2 TFLOPs	4.4 TFLOPs
	all jobs	all jobs	all jobs	all jobs
Premium	00:05	00:11	00:11	00:12
Regular	00:36	00:16	00:37	01:36
Economy	01:52	01:10	03:37	03:51
Stand-by	01:18	00:55	02:41	02:14

Archive Update

- **NCAR MSS exceeds 8 PBs total data stored – Sep 2009**
- **MSS Transition to HPSS**
 - *Target date Jan 2011*
 - *Erich to present more on this.*
- **AMSTAR Data Ooze (migrating data to new tape technology)**
 - *Optimized ooze software deployed July 2009. Streams tape to tape, creates two copies in parallel.*
 - *Oozing on average 20+ TBs per day, 6 PBs of data to ooze of which 4 PBs are unique*
 - *Estimated completion date Mar-Apr 2010 well ahead of the Dec 2010 EOL on Powderhorn tape libraries*

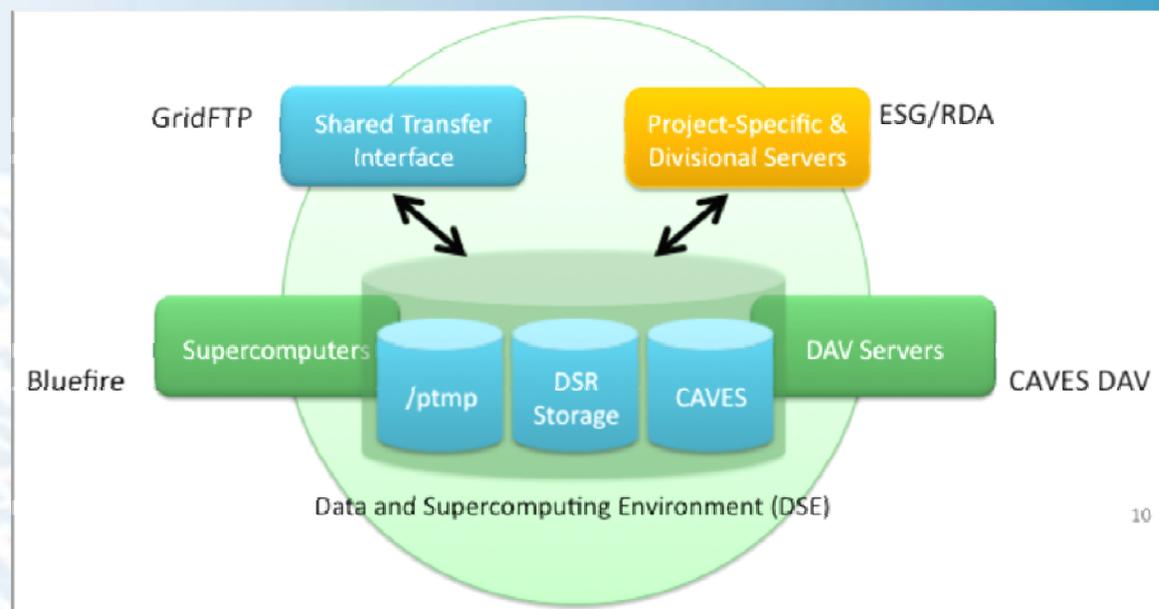


NCAR MSS - Total Data in Archive



Storage Update: *Data Services Redesign*

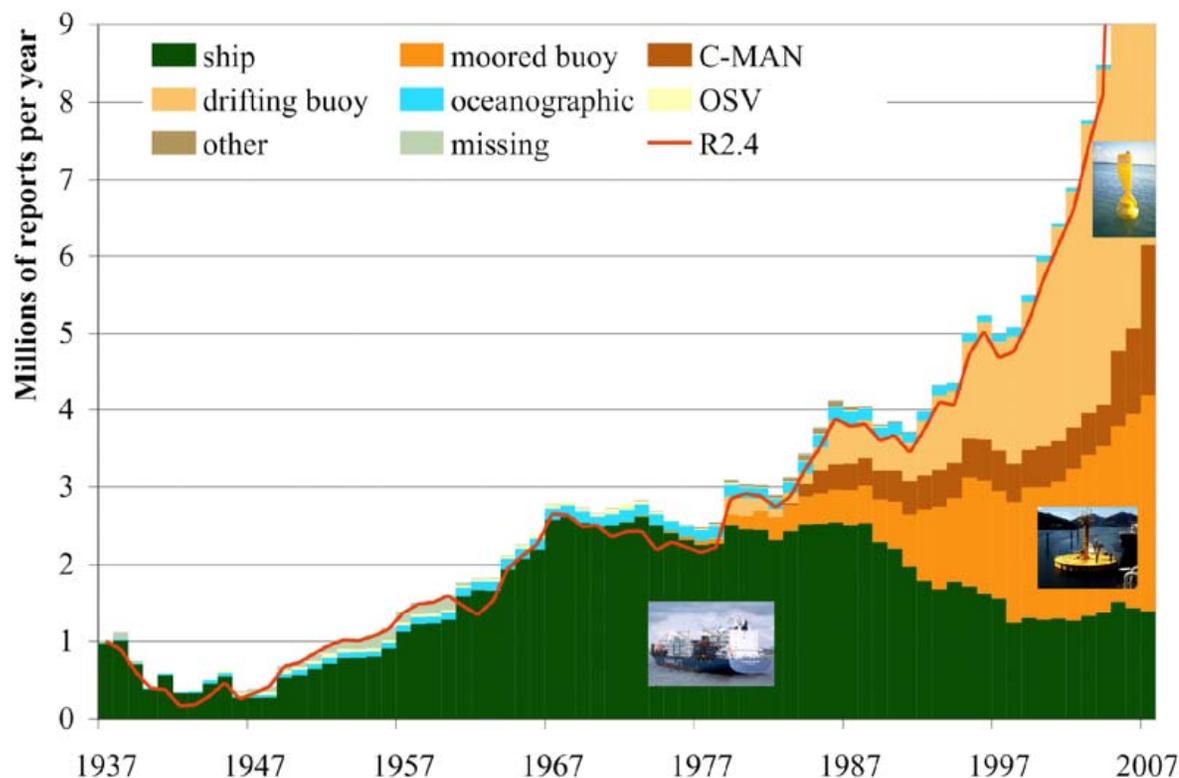
- Pending management review
- **Near-term Goals:**
 - *Creation of unified and consistent data environment for NCAR HPC*
 - *High-performance availability of central filesystem from many projects/systems (RDA, ESG, CAVS, Bluefire, TeraGrid)*
- **Longer-term Goals:**
 - *Filesystem cross-mounting*
 - *Global WAN filesystems (GPFS, and/or Lustre)*
- **Tentative Schedule:**
 - *Phase 1: March 2010*



Data Collections Update:

Major update of International Comprehensive Ocean-Atmosphere Data Set (ICOADS) in the RDA

- Release 2.5 covers 1662-2007
- Many data sources added over the full period of record
- **NEW** monthly automated update procedure adds preliminary data, now for 2008 - Sept. 2009
- ICOADS is a 25-year collaborative project with NOAA



Annual distribution (1937-2007) of major platform types in Release 2.5 shown as millions of reports per year.

Software Update

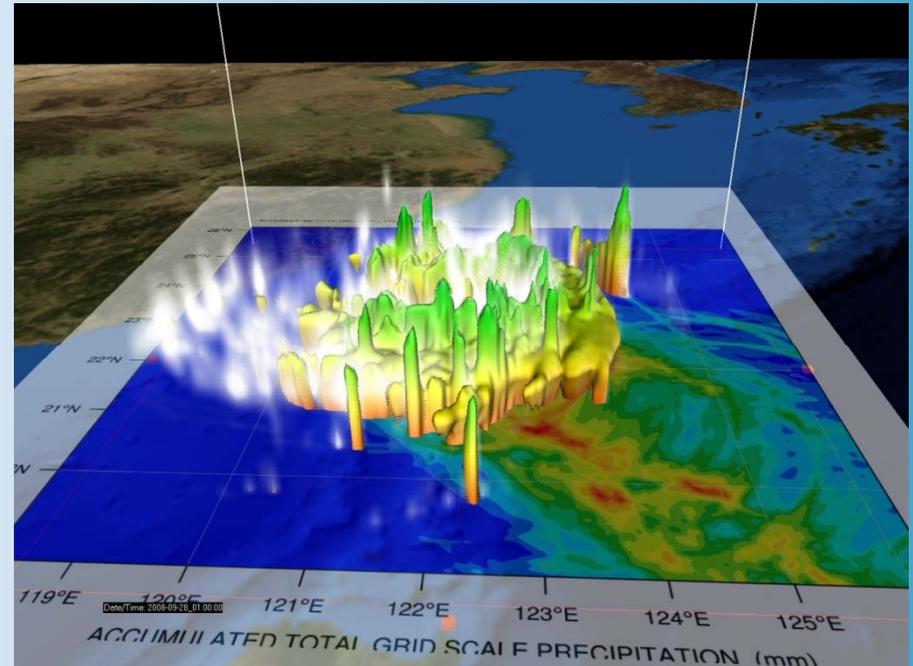


Version 1.5.0 released

- Numerous new features aimed at WRF community
 - Geo-referenced imagery
 - Integration with NCL
 - Moving domains
- Block structured AMR Grids
- Ease of use improvements

New funding streams

- NSF XD Vis Award
 - NCAR, TACC, Utah, Purdue
 - 2009-2012, \$386k
 - Focus: petascale DAV
- TeraGrid GIG Award
 - 2009-2011, \$112k
 - Focus: progressive access data models



This typhoon Jangmi image from MMM's Bill Kuo and Wei Wang illustrates several new capabilities. A moving, nested domain is tracked by VAPOR. Satellite imagery and plots generated by NCL are correctly positioned in the 4D scene with geo-referencing.

Support for Summer '09 Workshops

- **Climate Modeling Primer, July 27-31**
 - *Organized by CGD (Gettelman lead) and held at NCAR*
 - *41 participants, mostly graduate students*
 - *Ran low resolution versions of CAM; some work with CCSM*
 - *Provided 10-15 dedicated nodes and overview of using bluefire to participants*



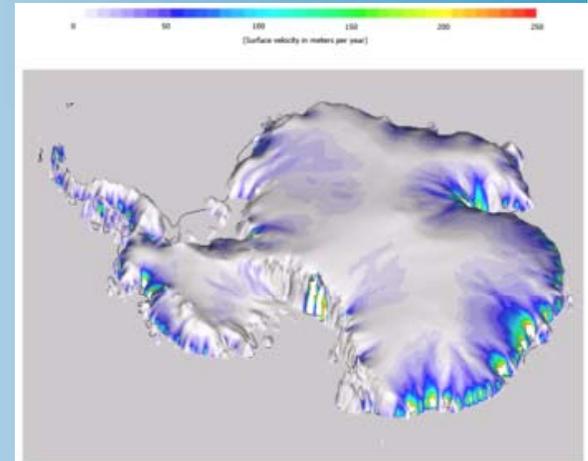
Support for Summer '09 Workshops

- **Scaling to Petascale, August 3-7**
 - *Organized by Great Lakes Consortium for Petascale Computation*
 - *Supported by the NSF through the Blue Waters Project (NCSA and others)*
 - *Participants joined in from four locations using high-definition streaming video*
 - *Ran a variety of applications*
 - *Provided 16 dedicated nodes during hands-on sessions*



Support for Summer '09 Workshops (continued)

- **Ice Sheet Modeling Workshop, August 3-14**
 - *Organized by Christina Hulbe (Portland State University), Jesse Johnson (U of Montana), and Cornelis van der Veen (U of Kansas); supported by NSF*
 - *Brought current and future ice-sheet scientists together to develop better models for the projection of future sea-level rise*
 - *Provided accounts for each participant and dedicated share queue on bluefire since using single processor codes*
 - *Some participants will continue to use bluefire based on Johnson's NSF award for this workshop*



http://websrv.cs.umt.edu/isis/index.php/Summer_Modeling_School

Preparations for NWSC Already Underway

Storage Environment

- HPSS Migration
- Data Services Redesign (DSR)
- Data Workflow Requirements Analysis

Security Review

- Balancing "Science and Security"
- Phase 1: Underway for DSR
- Phase 2: Global Filesystems

Accounting/Allocation Systems

- Current System built in 70s.
- Reviewing requirements
- Leveraging other products (e.g., Gold, AmieGold)

RFP Process for NWSC Procurement

- Tom Engel to Discuss Further

User Survey



User Survey Plans

Q1 2010 survey of active users

- *Use to baseline current environment for NWSC*
- *Understand where we're doing well and where we can improve*
 - Investigate any area with >10% dissatisfaction
- *Evaluate future needs*
- *Topics*
 - Ease of access within current security environment
 - Ease of use (initially and on-going)
 - Consulting (include methodology of service delivery)
 - Documentation for users
 - Training desired and method of delivery



Follow-up survey ~4 months after NWSC computer available to users

On the road to a PetaFlop Computer...

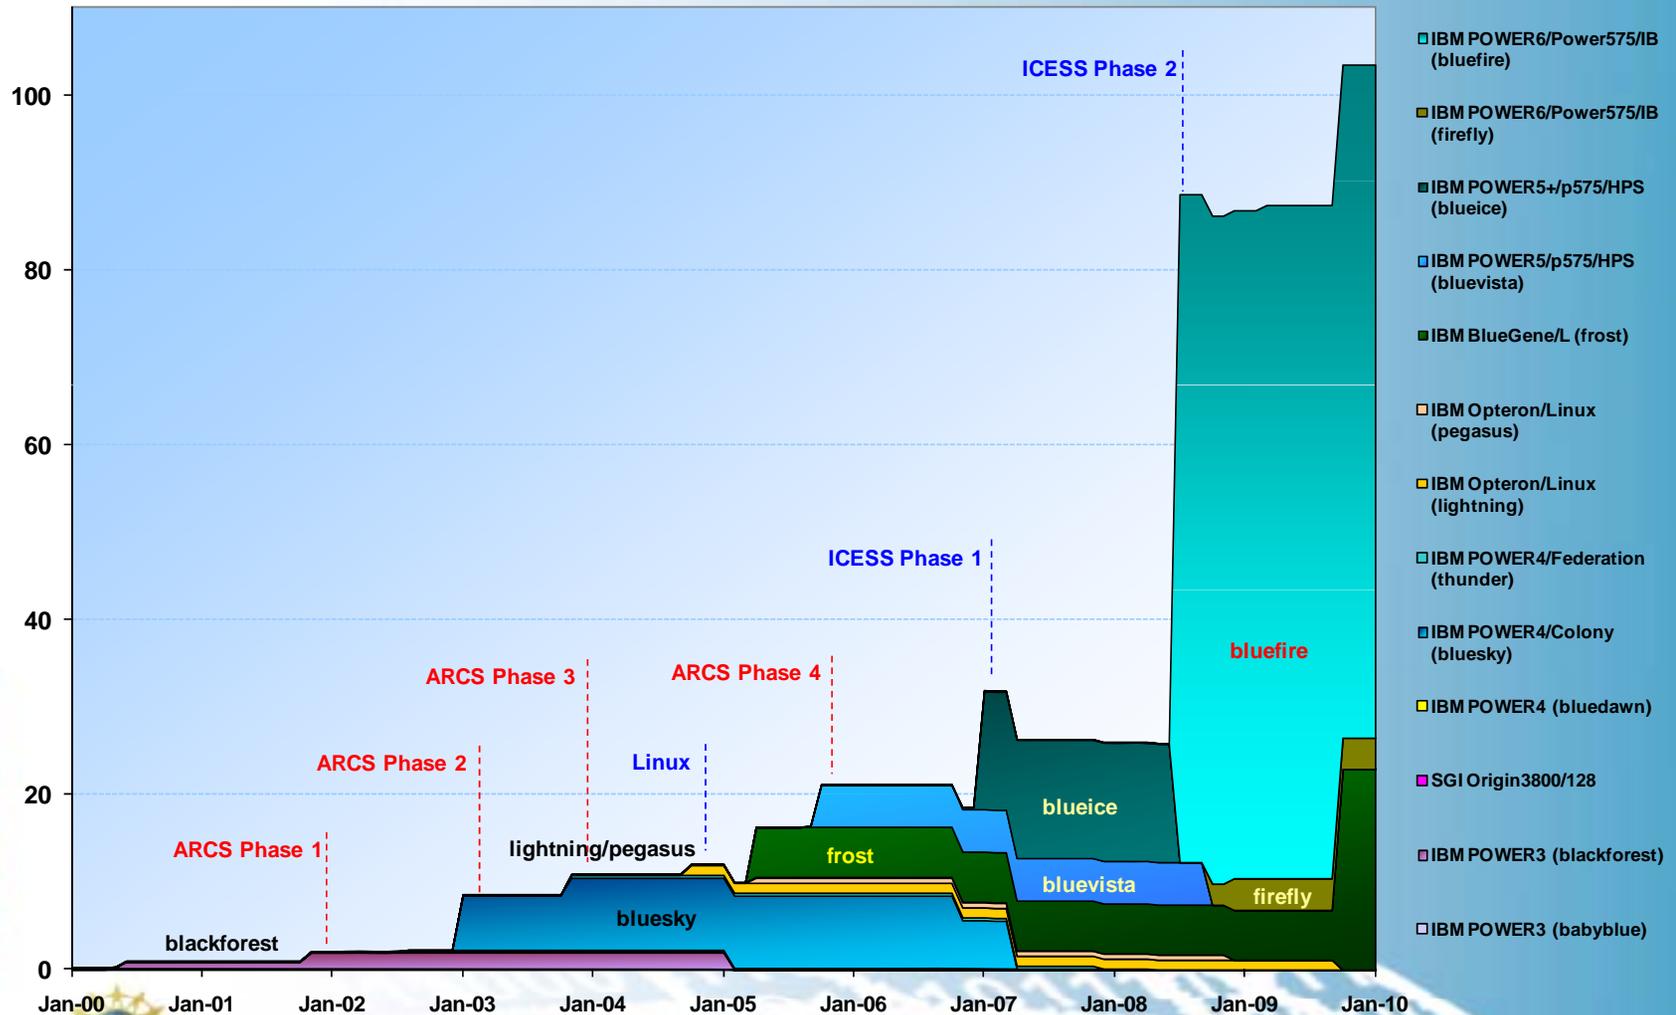
- And 6-15 Petabyte Filesystem
- And +100 Petabyte Archive



CURRENT NCAR COMPUTING > 100TFLOPS

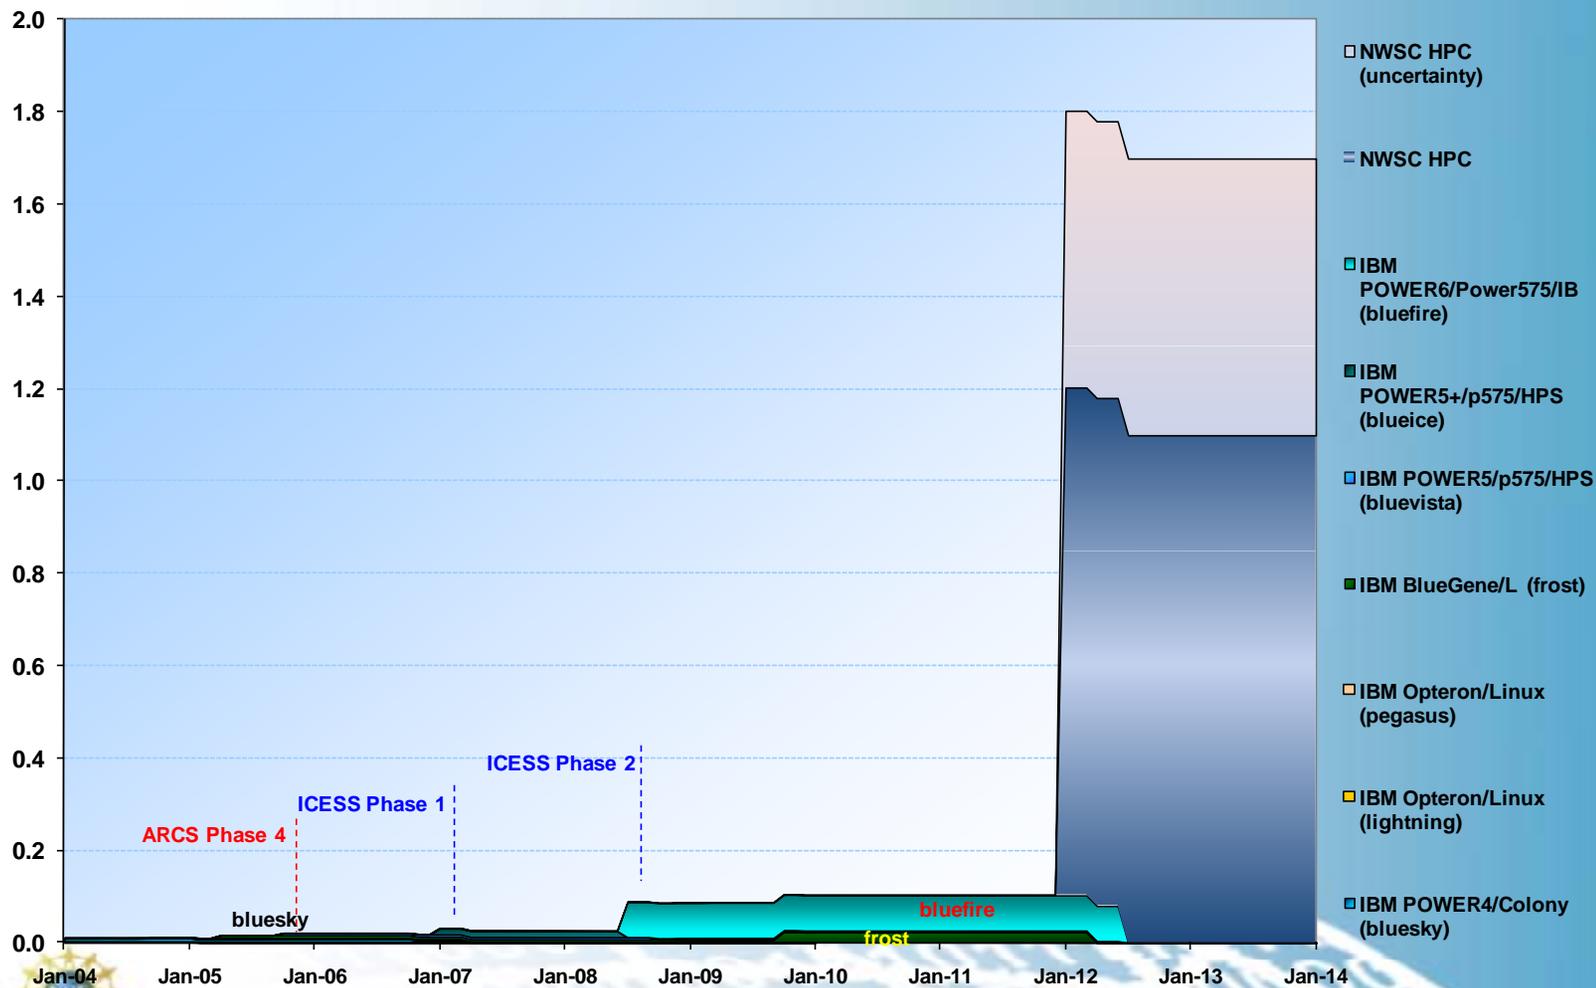
Peak TFLOPs at NCAR (All Systems)

Computational & Information Systems Laboratory



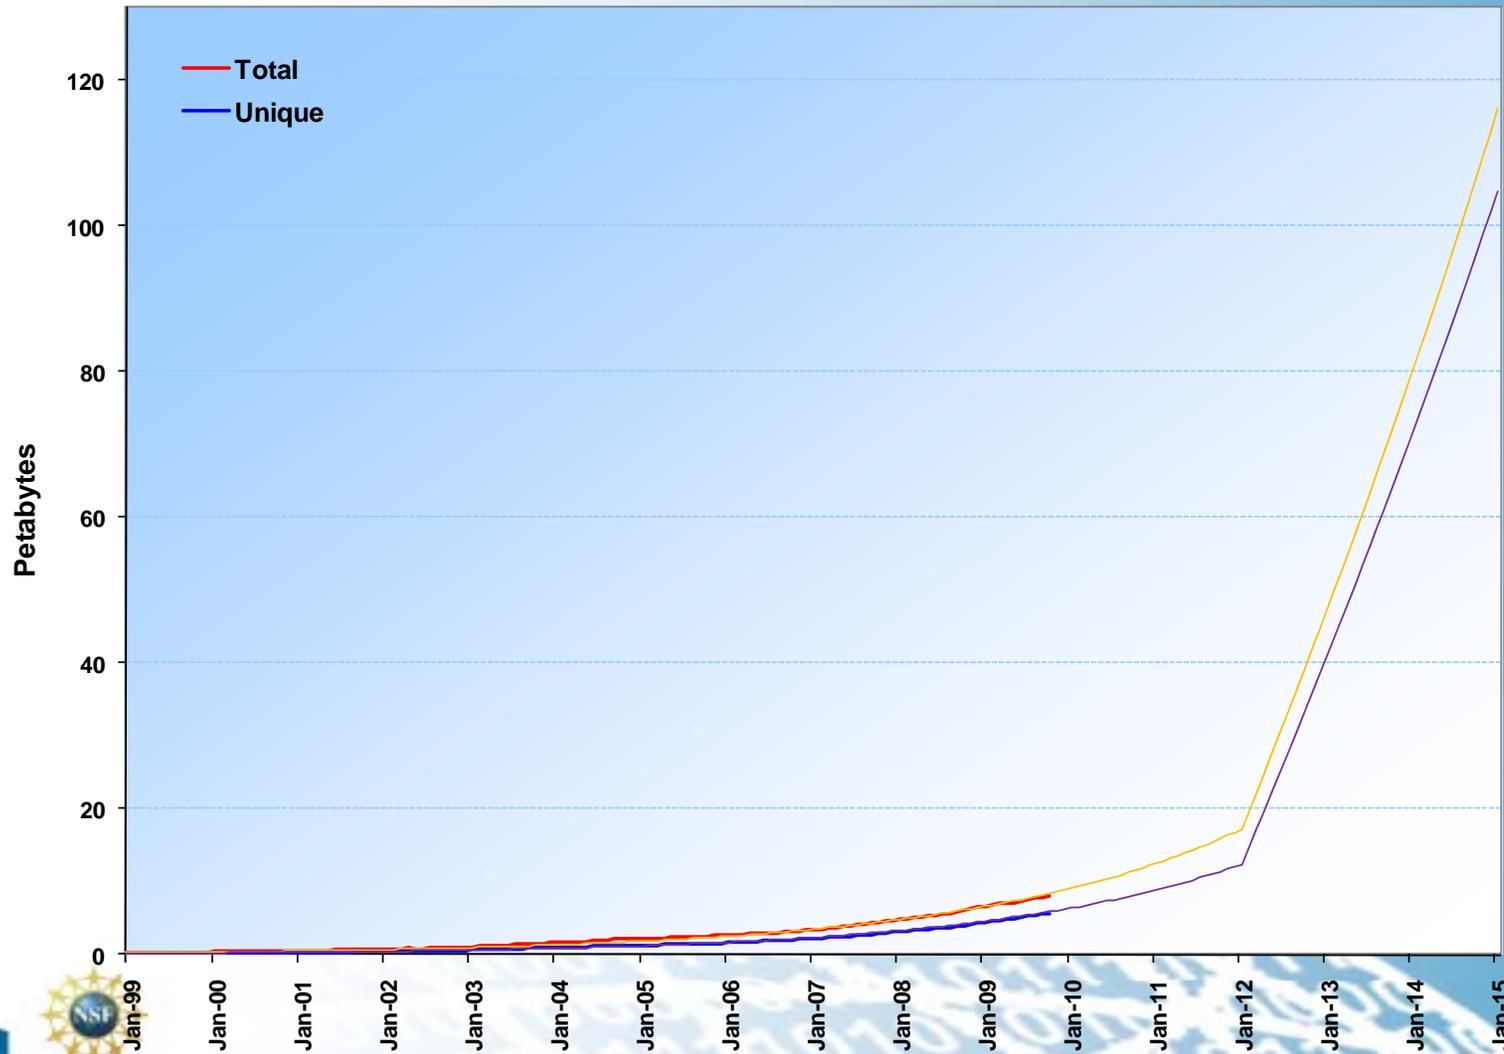
NWSC HPC Projection

Peak PFLOPs at NCAR



NCAR Data Archive Projection

Total Data in the NCAR Archive (Actual and Projected)





Questions and Discussion