Cyberinfrastructure

John Helly
San Diego Supercomputer Center
University of California, San Diego

Outline

- Management
- Resources
- Research

Management

CIWG Objectives

- Make efficient use of computing and data resources
  - acquire resources
  - coordinate resource utilization
  - collaborate to leverage joint efforts
- Validate goals and provide advice anto Executive Committee

CIWG Management Resources

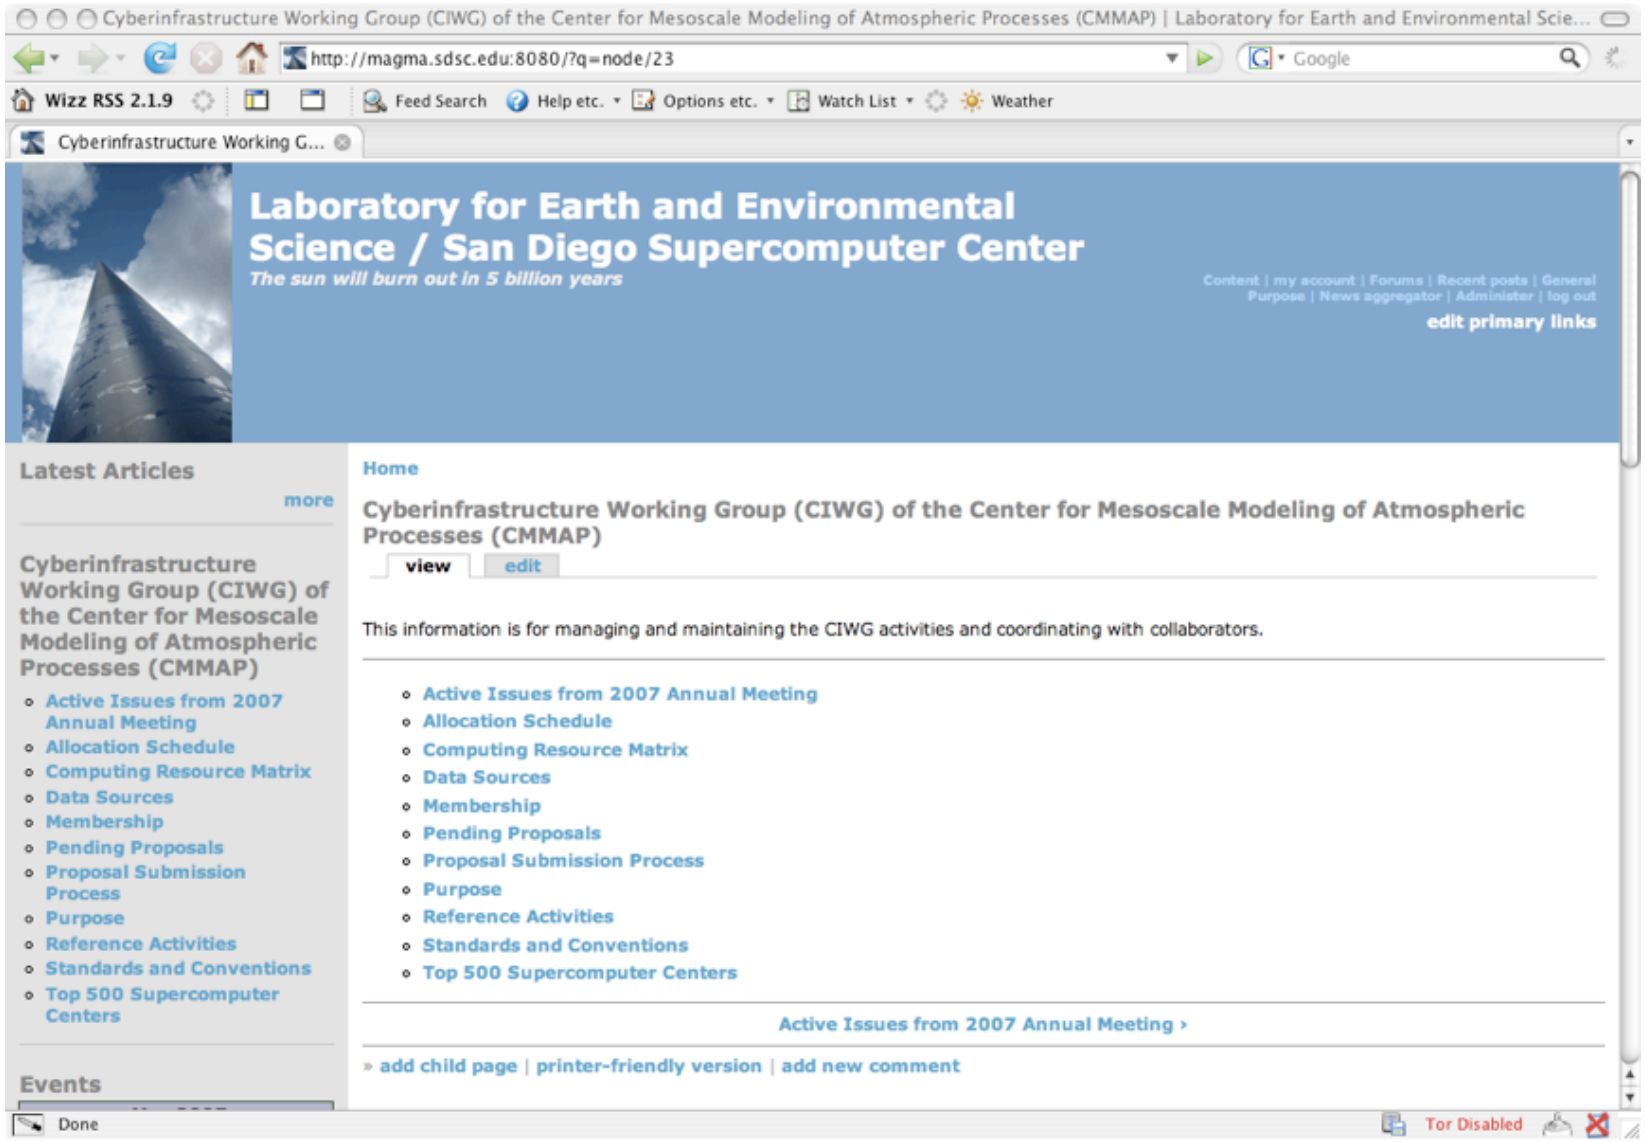

Cyberinfrastructure Organization

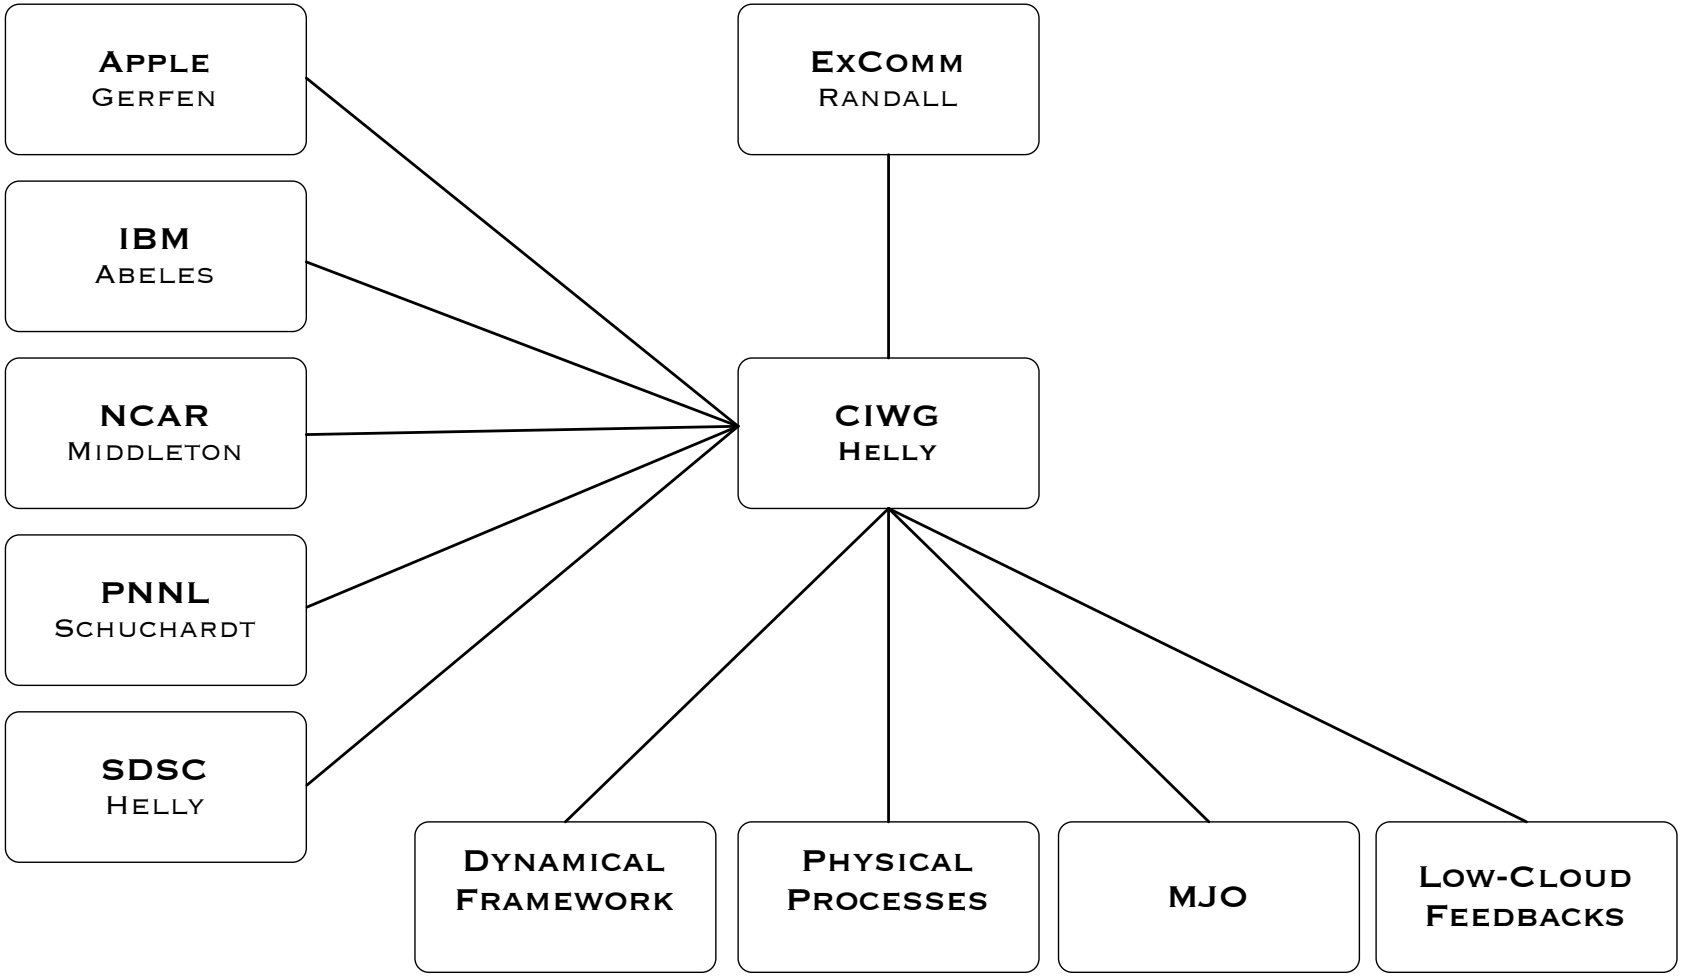

Architectural Goals

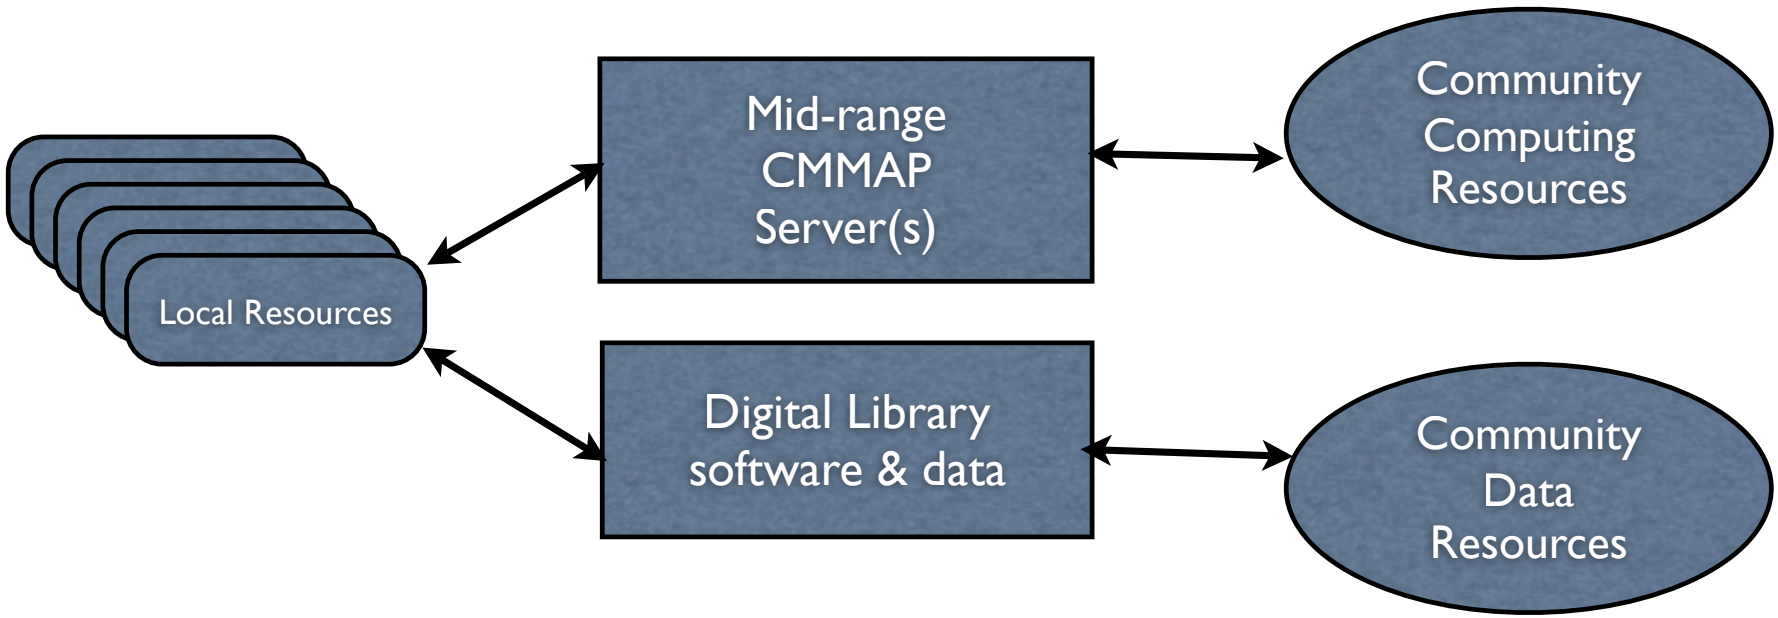

Home Institution

CMMAP Machine Room

External Resources

Proposal Efforts

- 2006 MRI Not-selected by CSU
- 2007 NSF SDCI Modeler's Toolbox (not selected; reviewers said Teragrid is doing this)
- 2008 NSF MRI (could not match)
- 2008 NSF MSI Community Computing and Data Resource (not selected; do not have reviews yet)

Resources

Background

- CMMAP cyberinfrastructure is shared, not owned
  - Collaboration and leveraging are key to acquisition and efficient use of resources
  - CMMAP-owned computing and data resources will come from separate proposals
- Training the next generation of scientists in a state-of-the-art computing environment is essential
  - Want graduate student involvement in every aspect of computing activities

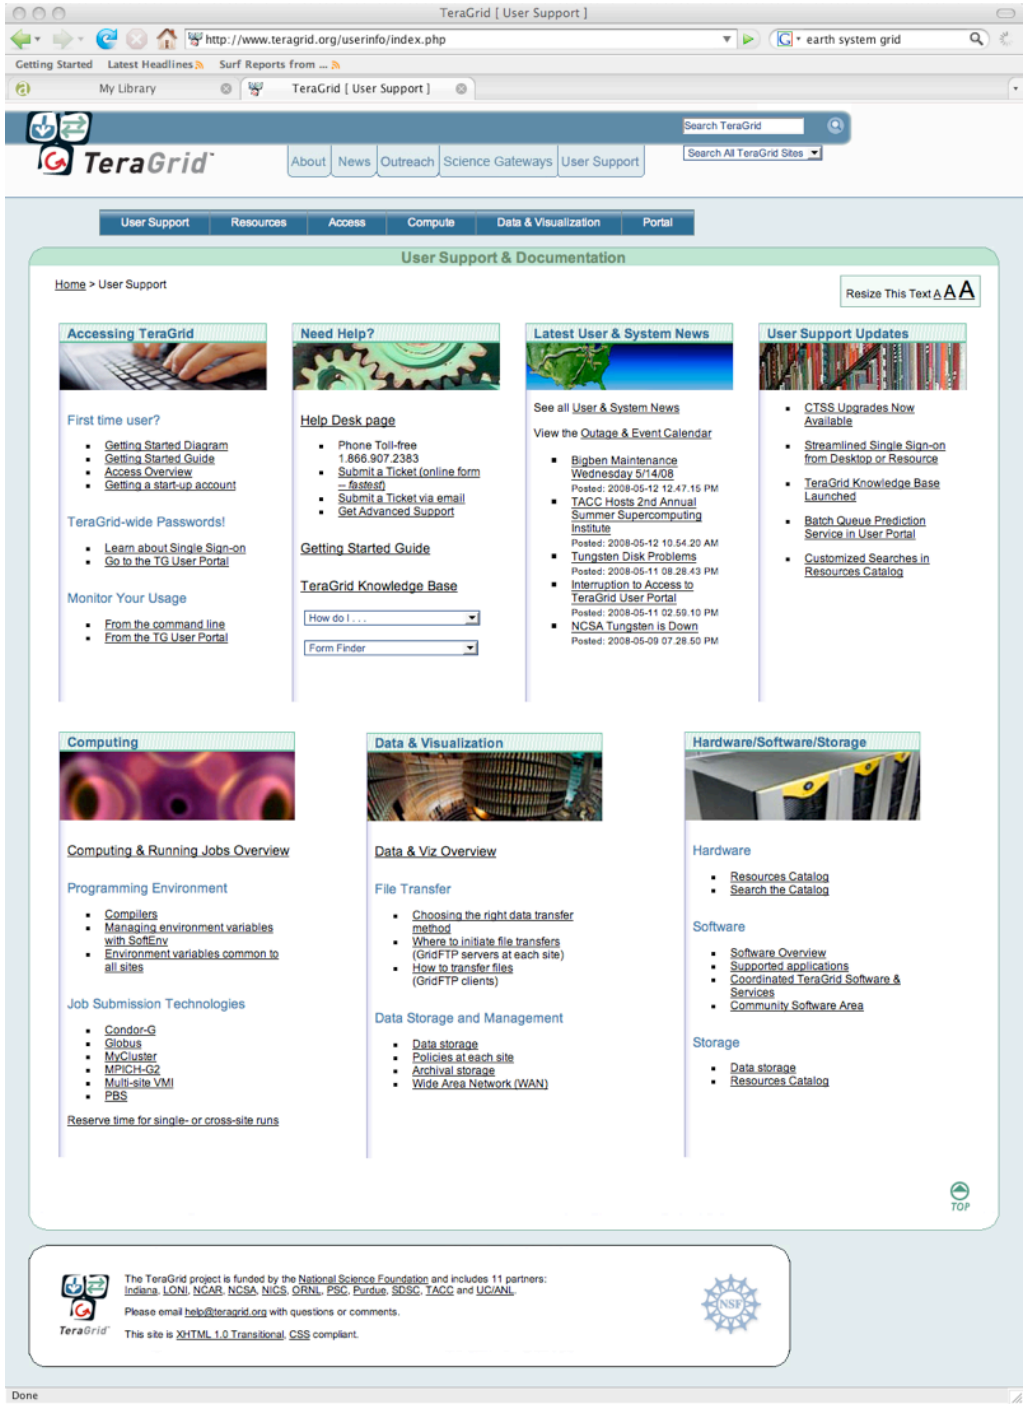

NSF Teragrid Resources

High Performance Computing Systems

Name	Institution	System	CPUs	Peak TFlops	Memory TBytes	Disk TBytes	Load	Jobs* R	Jobs* Q	Jobs* O
Ranger	TACC	Sun Constellation	62976	504.00	123.00	1730.00		1	291	67
Abe	NCSA	Dell Intel 64 Linux Cluster	9600	89.47	9.38	100.00		281	239	82
Lonestar	TACC	Dell PowerEdge Linux Cluster	5840	62.16	11.60	106.50		106	222	1
Queen Bee	LONI	Dell Intel 64 Linux Cluster	5440	50.70	5.31	100.00		52	1	0
Steele	Purdue	Dell Intel 64 Linux Cluster	6496	40.00	12.40	170.00		146	12	38
Big Red	IU	IBM e1350	3072	30.60	6.00	266.00		200	0	265
BigBen	PSC	Cray XT3	4136	21.50	4.04	100.00		0	0	0
Blue Gene	SDSC	IBM Blue Gene	6144	17.10	1.50	19.50		3	0	6
Tungsten	NCSA	Dell Xeon IA-32 Linux Cluster	2560	16.38	3.75	109.00		139	246	67
DataStar p655	SDSC	IBM Power4+ p655	2176	14.30	5.75	115.00		19	83	28
TeraGrid Cluster	NCSA	IBM Itanium2 Cluster	1744	10.23	4.47	60.00		100	34	0
Cobalt	NCSA	SGI Altix	1024	6.55	3.00	100.00		60	351	0
Frost	NCAR	IBM BlueGene/L	2048	5.73	0.51	6.00		11	5	0
TeraGrid Cluster	SDSC	IBM Itanium2 Cluster	524	3.10	1.02	48.80		12	22	2
DataStar p690	SDSC	IBM Power4+ p690	192	1.30	0.88	115.00		4	3	13
TeraGrid Cluster	UC/ANL	IBM Itanium2 Cluster	128	0.61	0.24	4.00		0	0	0
NSTG	ORNL	IBM IA-32 Cluster	56	0.34	0.07	2.14		0	0	0
Rachel	PSC	HP Alpha SMP	128	0.31	0.50	6.00		31	46	0
		Total:	**114284**	**874.38**	**193.42**	**3157.94**		**1165**	**1555**	**569**

High Throughput Computing Systems

Name	Institution	Active/Available Nodes	Active/Available CPUs	Peak TFlops	Memory GBytes	Disk GBytes	Resource Details	Jobs
Condor Pool	Purdue	**5212** / 6687	**13318** / 17327	14	11588	230073		
	Total:	**5212 / 6687**	**13318 / 17327**	**14**	**11588**	**230073**		

Advanced Visualization Systems

Name	Institution	System	CPUs	Peak TFlops	Memory TBytes	Disk TBytes	Graphics HW
TeraGrid Cluster	UC/ANL	Intel Xeon Cluster	192	0.61	0.38	4.00	nVIDIA GeFORCE 6600GT AGP graphics cards
Maverick	TACC	Sun E25K	128	0.27	0.50	0.56	16 nVIDIA QuadroFX 3000G graphics cards
		Total:	**320**	**0.88**	**0.88**	**4.56**	

**Jobs Key*: R - Number of Jobs Running, Q - Number of Jobs Queued, O - Number of Jobs in an Other State

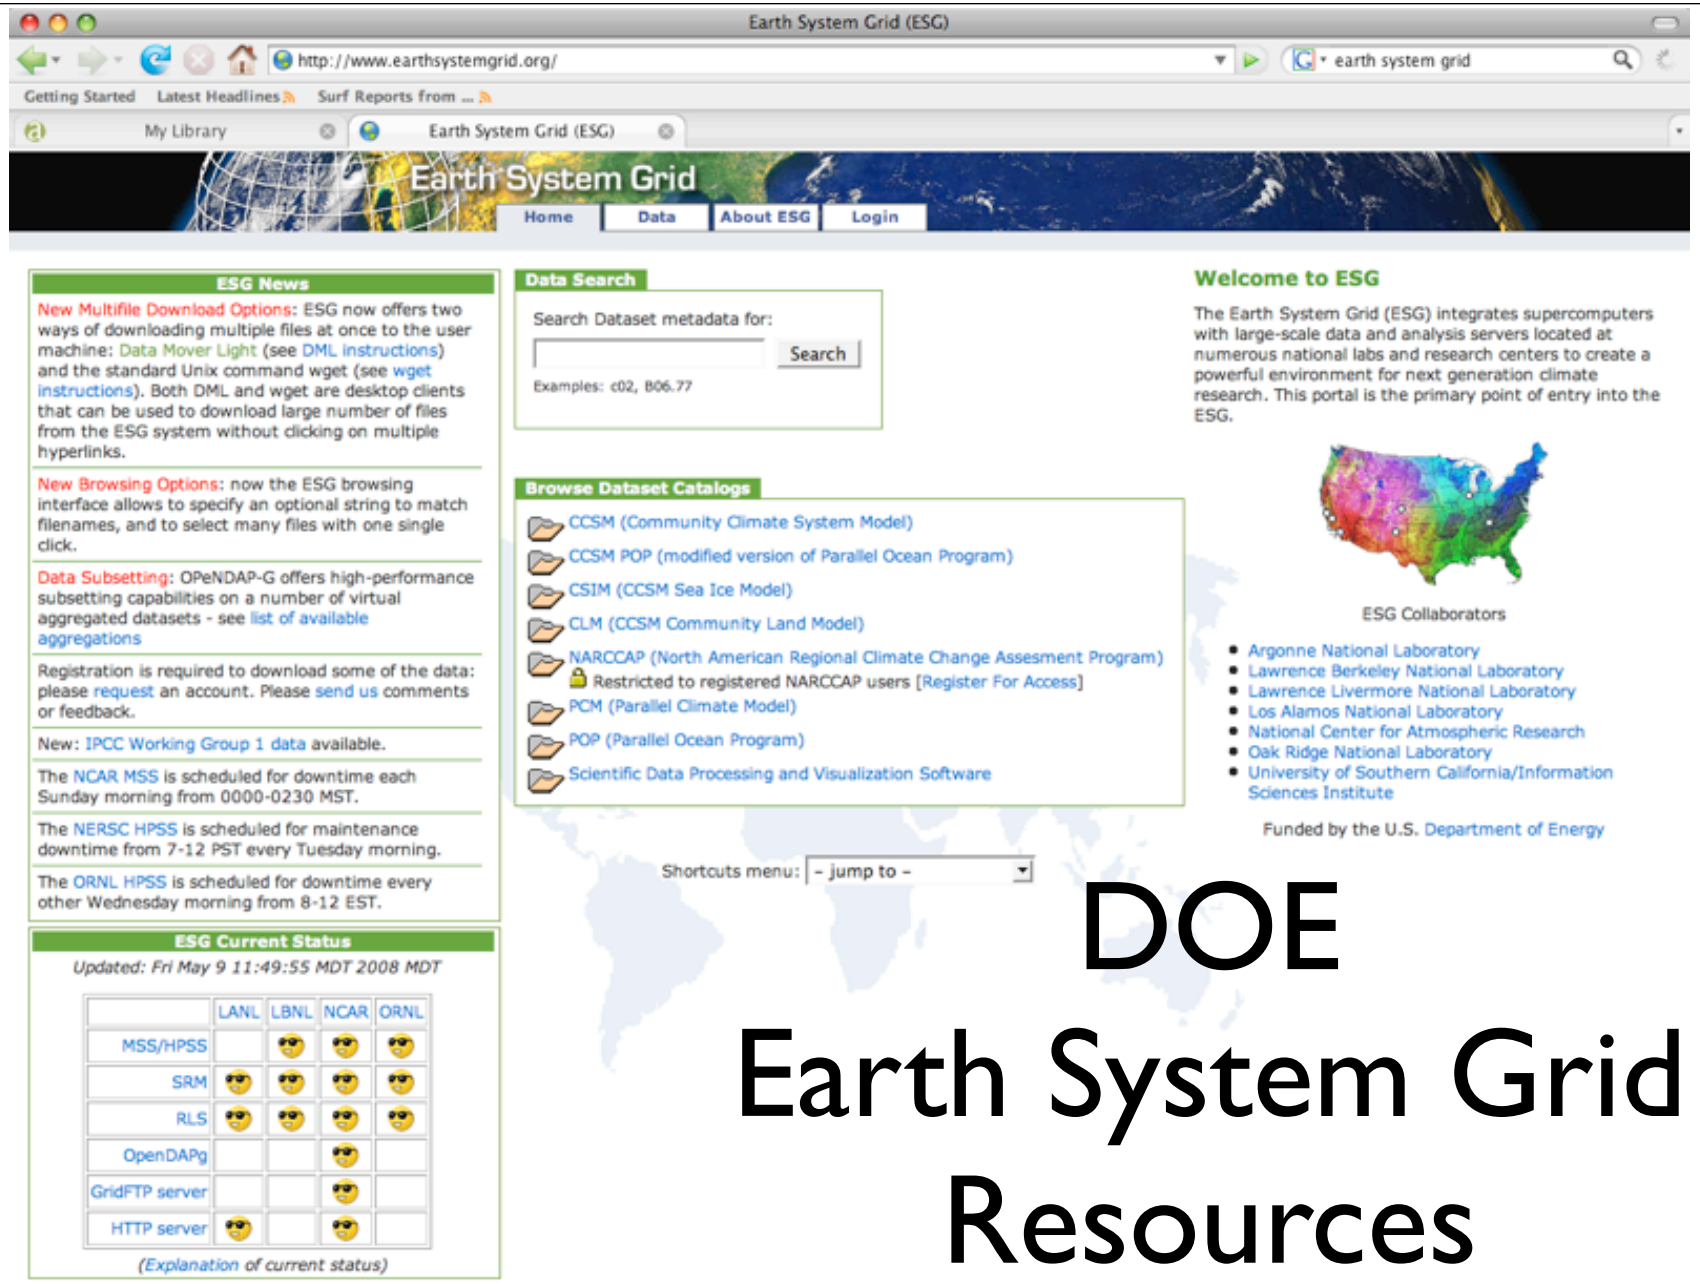

DOE Earth System Grid Resources

Leveraging National & Partner Resources

	Organization	Resource	Amount
Data Allocations	San Diego Supercomputer Center (SDSC)	Disk	15 Terabytes
		BlueGene	30,000 SUs*
Computing Allocations	Teragrid (multi-institution)	SDSC DataStar (IBM SP4)	600,000 SUs
		Grid Roaming	600,000 SUs
	Lawrence Berkeley National Laboratory (LBNL)	National Energy Research Scientific Computing Center (NERSC)	700,000 SUs
	National Center for Atmospheric research (NCAR)	BlueIce IBM Power5	500,000 SUs
	IBM Watson Research Center	BGW - eServer Blue Gene Solution	TBD
	Stonybrook		TBD

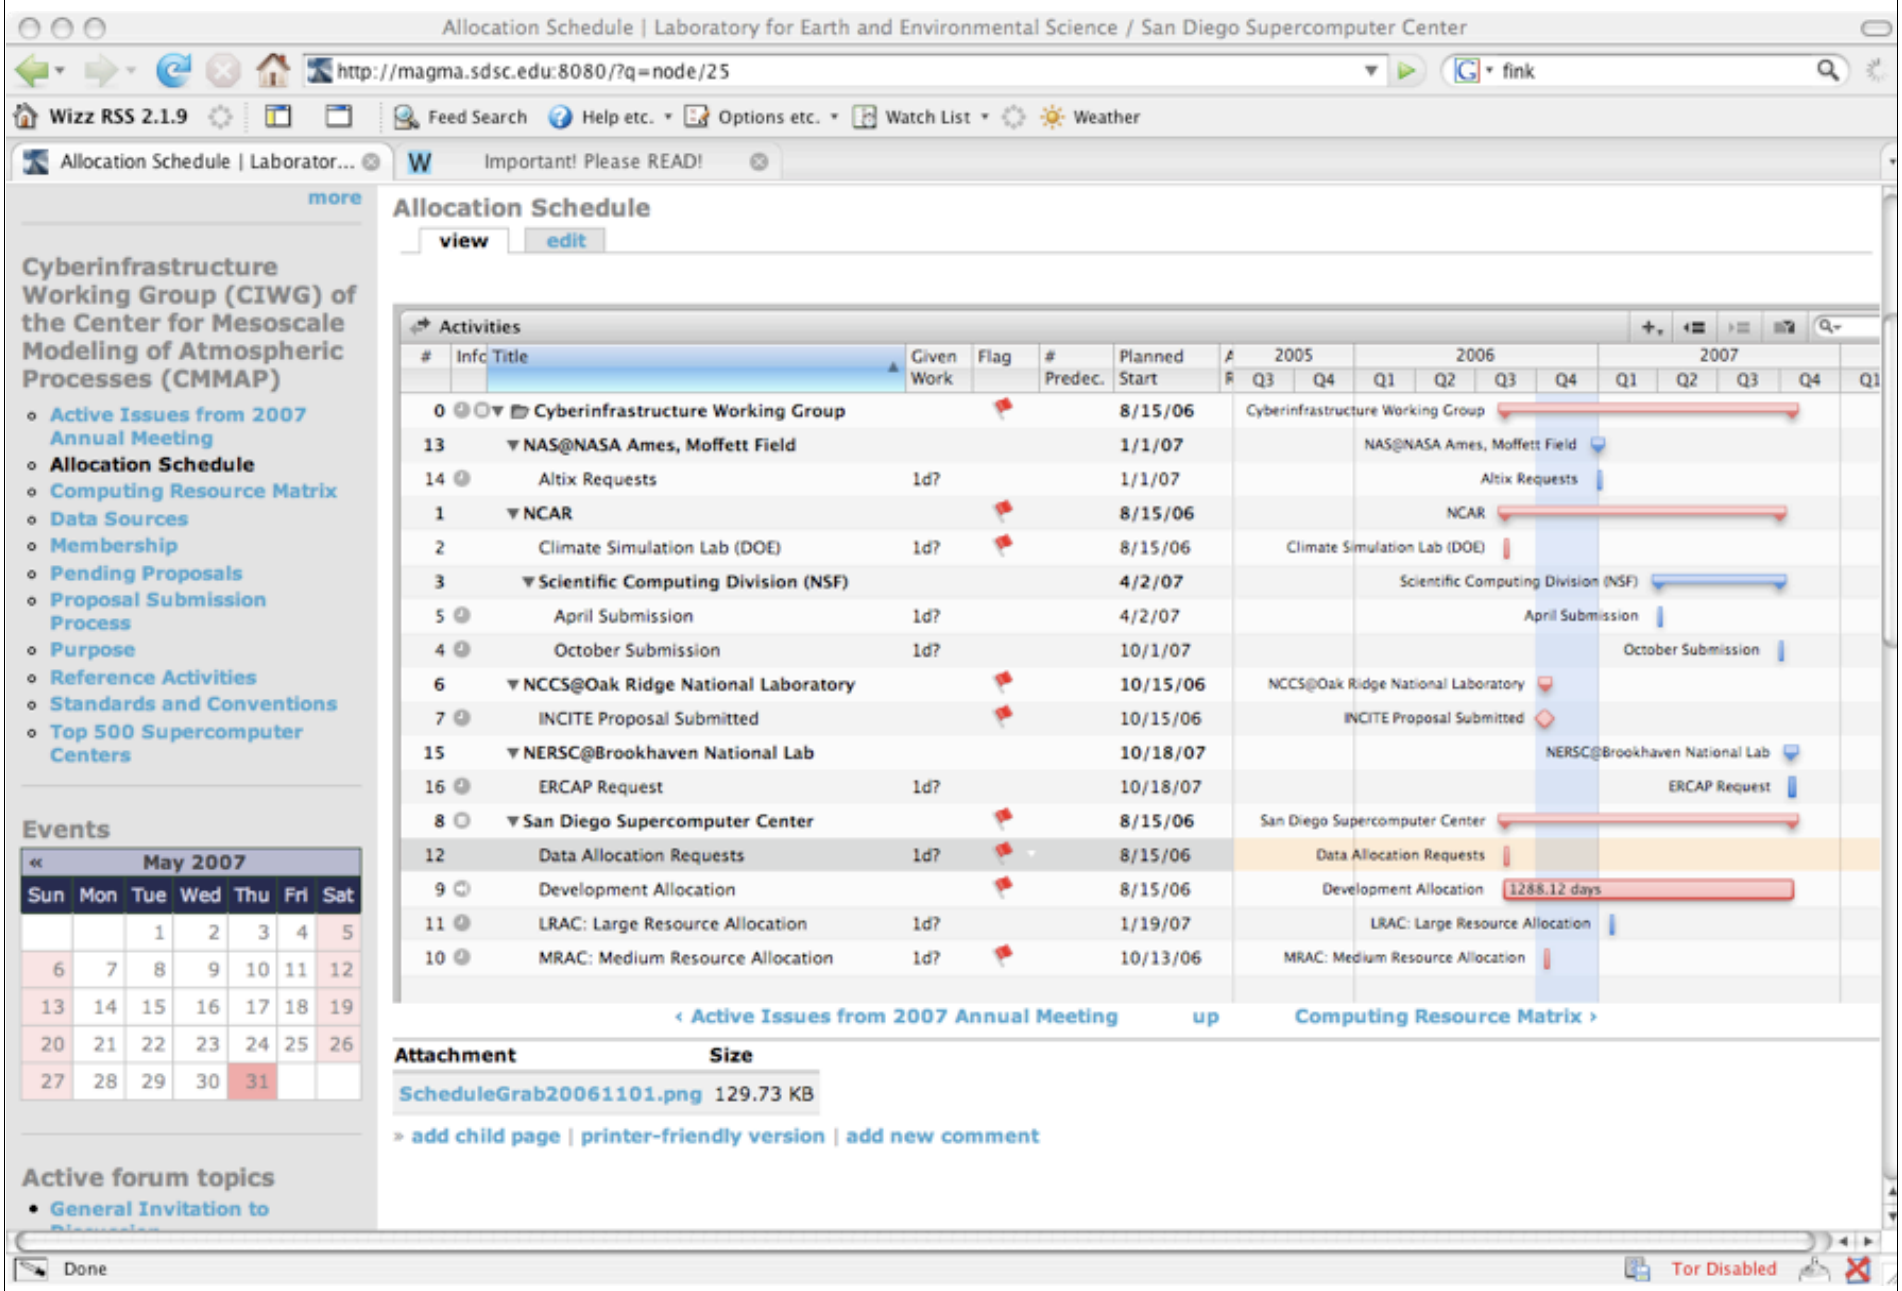
Allocation Schedule | Laboratory for Earth and Environmental Science / San Diego Supercomputer Center
http://magma.sdsc.edu:8080/?q=node/25
fink
Wizz RSS 2.1.9
Feed Search
Help etc.
Options etc.
Watch List
Weather
Allocation Schedule | Laborator...
Important! Please READ!
more
Cyberinfrastructure Working Group (CIWG) of the Center for Mesoscale Modeling of Atmospheric Processes (CMMAP)
Active Issues from 2007 Annual Meeting
Allocation Schedule
Computing Resource Matrix
Data Sources
Membership
Pending Proposals
Proposal Submission Process
Purpose
Reference Activities
Standards and Conventions
Top 500 Supercomputer Centers
Events
May 2007
Sun Mon Tue Wed Thu Fri Sat
Active forum topics
General Invitation to
Allocation Schedule
view
edit
Activities
Info Title Given Work Flag # Predec. Planned Start
2005 2006 2007
Q3 Q4 Q1 Q2 Q3 Q4 Q1 Q2 Q3 Q4 Q1
0 Cyberinfrastructure Working Group 8/15/06
13 NAS@NASA Ames, Moffett Field 1/1/07
14 Altix Requests 1d? 1/1/07
1 NCAR 8/15/06
2 Climate Simulation Lab (DOE) 1d? 8/15/06
3 Scientific Computing Division (NSF) 4/2/07
5 April Submission 1d? 4/2/07
4 October Submission 1d? 10/1/07
6 NCCS@Oak Ridge National Laboratory 10/15/06
7 INCITE Proposal Submitted 10/15/06
15 NERSC@Brookhaven National Lab 10/18/07
16 ERCAP Request 1d? 10/18/07
8 San Diego Supercomputer Center 8/15/06
12 Data Allocation Requests 1d? 8/15/06
9 Development Allocation 8/15/06
11 LRAC: Large Resource Allocation 1d? 1/19/07
10 MRAC: Medium Resource Allocation 1d? 10/13/06
1288.12 days
‹ Active Issues from 2007 Annual Meeting
up
Computing Resource Matrix ›
Attachment Size
ScheduleGrab20061101.png 129.73 KB
» add child page | printer-friendly version | add new comment
Done
Tor Disabled

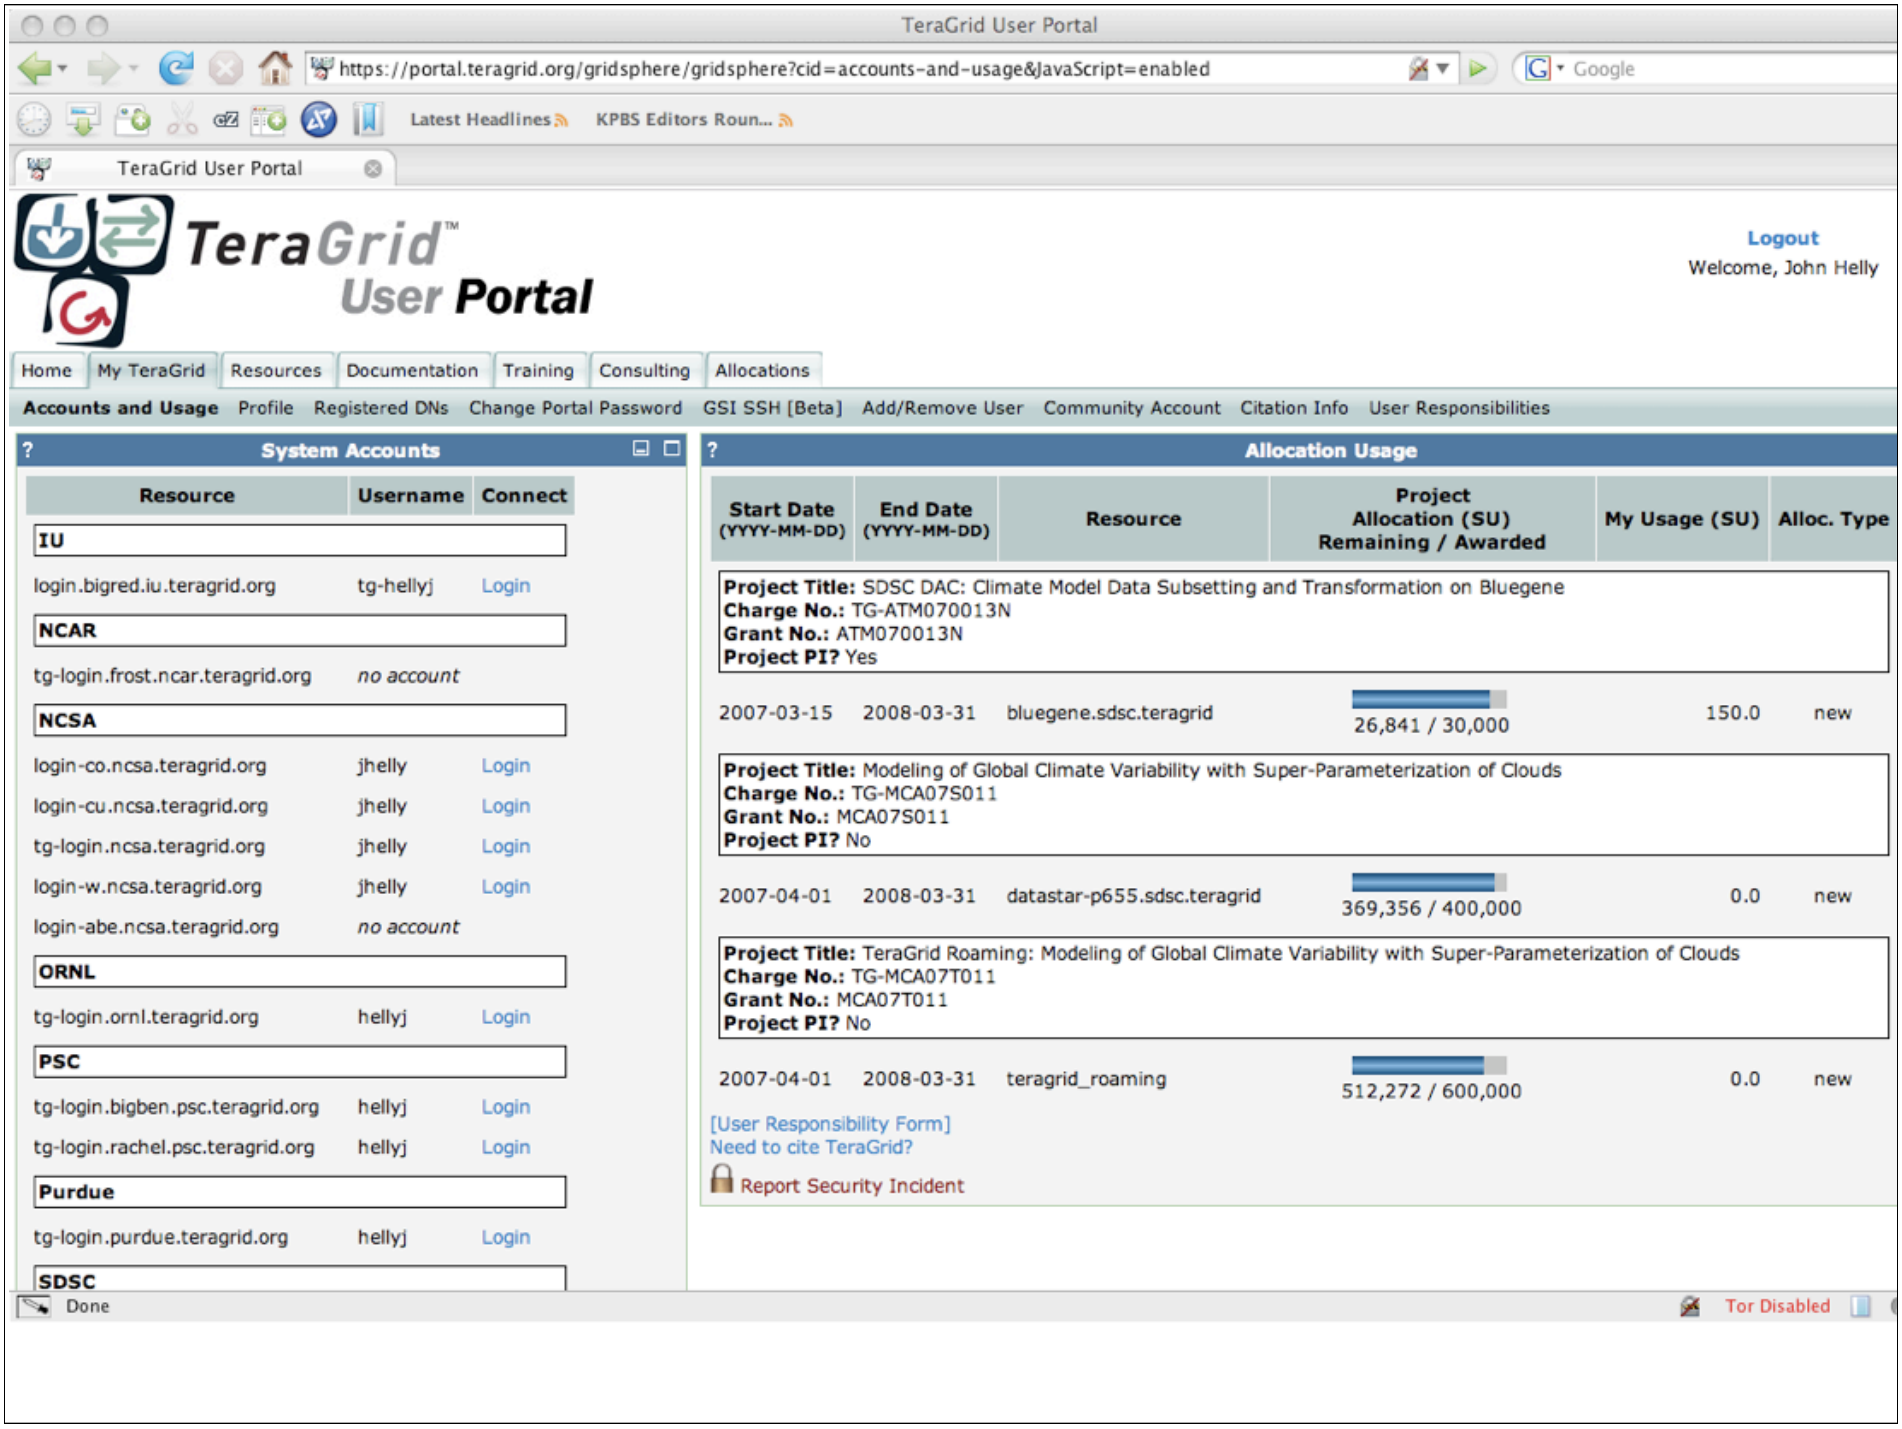
TeraGrid User Portal
https://portal.teragrid.org/gridsphere/gridsphere?cid=accounts-and-usage&JavaScript=enabled
Google
Latest Headlines
KPBS Editors Roun...
TeraGrid User Portal
TeraGrid User Portal
Logout
Welcome, John Helly
Home
My TeraGrid
Resources
Documentation
Training
Consulting
Allocations
Accounts and Usage
Profile
Registered DNs
Change Portal Password
GSI SSH [Beta]
Add/Remove User
Community Account
Citation Info
User Responsibilities
System Accounts
Resource
Username
Connect
IU
login.bigred.iu.teragrid.org
tg-hellyj
Login
NCAR
tg-login.frost.ncar.teragrid.org
no account
NCSA
login-co.ncsa.teragrid.org
jhelly
Login
login-cu.ncsa.teragrid.org
jhelly
Login
tg-login.ncsa.teragrid.org
jhelly
Login
login-w.ncsa.teragrid.org
jhelly
Login
login-abe.ncsa.teragrid.org
no account
ORNL
tg-login.ornl.teragrid.org
hellyj
Login
PSC
tg-login.bigben.psc.teragrid.org
hellyj
Login
tg-login.rachel.psc.teragrid.org
hellyj
Login
Purdue
tg-login.purdue.teragrid.org
hellyj
Login
SDSC
Allocation Usage
Start Date (YYYY-MM-DD)
End Date (YYYY-MM-DD)
Resource
Project Allocation (SU) Remaining / Awarded
My Usage (SU)
Alloc. Type
Project Title: SDSC DAC: Climate Model Data Subsetting and Transformation on Bluegene
Charge No.: TG-ATM070013N
Grant No.: ATM070013N
Project PI? Yes
2007-03-15
2008-03-31
bluegene.sdsc.teragrid
26,841 / 30,000
150.0
new
Project Title: Modeling of Global Climate Variability with Super-Parameterization of Clouds
Charge No.: TG-MCA07S011
Grant No.: MCA07S011
Project PI? No
2007-04-01
2008-03-31
datastar-p655.sdsc.teragrid
369,356 / 400,000
0.0
new
Project Title: TeraGrid Roaming: Modeling of Global Climate Variability with Super-Parameterization of Clouds
Charge No.: TG-MCA07T011
Grant No.: MCA07T011
Project PI? No
2007-04-01
2008-03-31
teragrid_roaming
512,272 / 600,000
0.0
new
[User Responsibility Form]
Need to cite TeraGrid?
Report Security Incident
Done
Tor Disabled

Petascale Computing

Emerging Track 1,2 Systems

August 8, 2007

- NSF's Track 1 petascale system is expected to be able to make arithmetic calculations at a sustained rate in excess of a sizzling 1,000-trillion operations per second (a petaflop per second)
- **Track 1: A Leadership-Class System**
  - **University of Illinois at Urbana-Champaign (UIUC) will receive $208 million over 4.5 years to acquire and make available a petascale computer "Blue Waters"**
  - **500 times more powerful than today's typical supercomputers.**
  - **The system is expected to go online in 2011.**
- **Track 2: Mid-Range High-Performance Computing Towards the Petascale**
  - University of Tennessee at Knoxville Joint Institute for Computational Science (JICS).
  - $65 million, 5-year project will include partners at Oak Ridge National Laboratory, the Texas Advanced Computing Center, and the National Center for Atmospheric Research.
  - **peak performance of just under one petaflop that is almost four times the capacity of the current NSF-supported Teragrid**.
  - 2006 a 5-year, $59 million Track 2 award to the Texas Advanced Computing Center at The University of Texas at Austin and Arizona State University and Cornell University (Ranger).
  - **Track 2 systems will be integrated into the TeraGrid.**

Track 1 Applications Managed Through Proposals

Leadership-Class System Acquisition - Creating a Petascale Computing Environment for Science and Engineering

PROGRAM SOLICITATION
NSF 06-573

National Science Foundation

Office of Cyberinfrastructure

Preliminary Proposal Due Date(s) ***(required):***

September 08, 2006

Full Proposal Deadline(s) (due by 5 p.m. proposer's local time):

February 02, 2007

SUMMARY OF PROGRAM REQUIREMENTS

General Information

Program Title:

Leadership-Class System Acquisition - Creating a Petascale Computing Environment for Science and Engineering

Synopsis of Program:

NSF's goal for high performance computing (HPC) in the period 2006-2011 is to enable petascale science and engineering through the deployment and support of a world-class HPC environment comprising the most capable combination of HPC assets available to the academic community. The petascale HPC environment will enable investigations of computationally challenging problems that require computing systems capable of delivering sustained performance approaching 10^{15} floating point operations per second (petaflops) on real applications, that consume large amounts of memory, and/or that work with very large data sets. Among other things, researchers will be able to perform simulations that are intrinsically multi-scale or that involve the simultaneous interaction of multiple processes.

Research

Collaborations

Participants	MMF Bluegene	Data Extraction & Subsetting	Data & Metadata Interoperability
IBM			
NCAR			
PNNL			
SDSC			
Stonybrook			

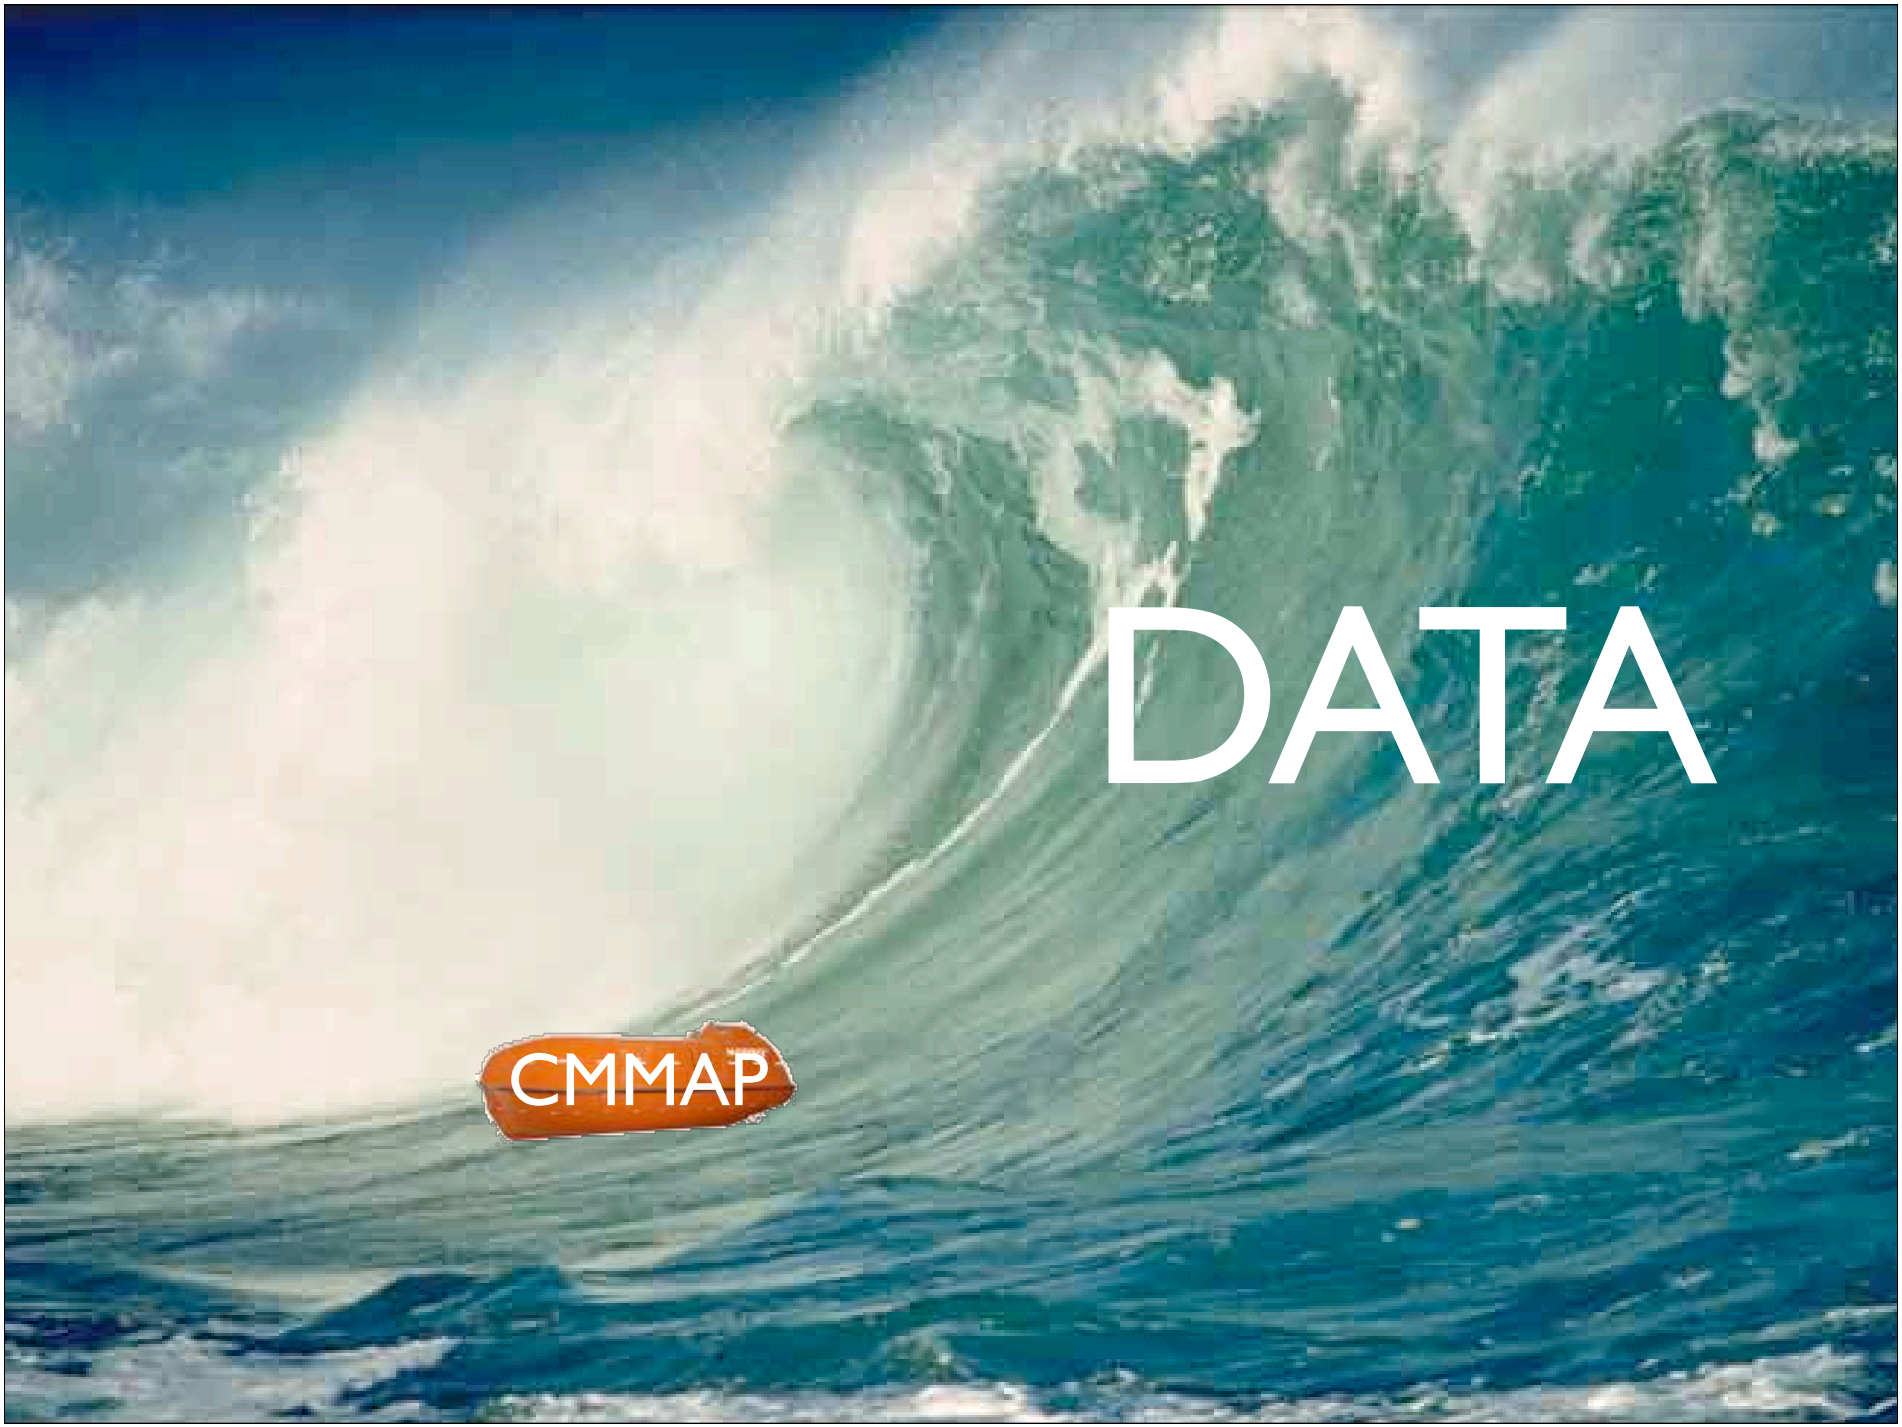
DATA
CMMAP

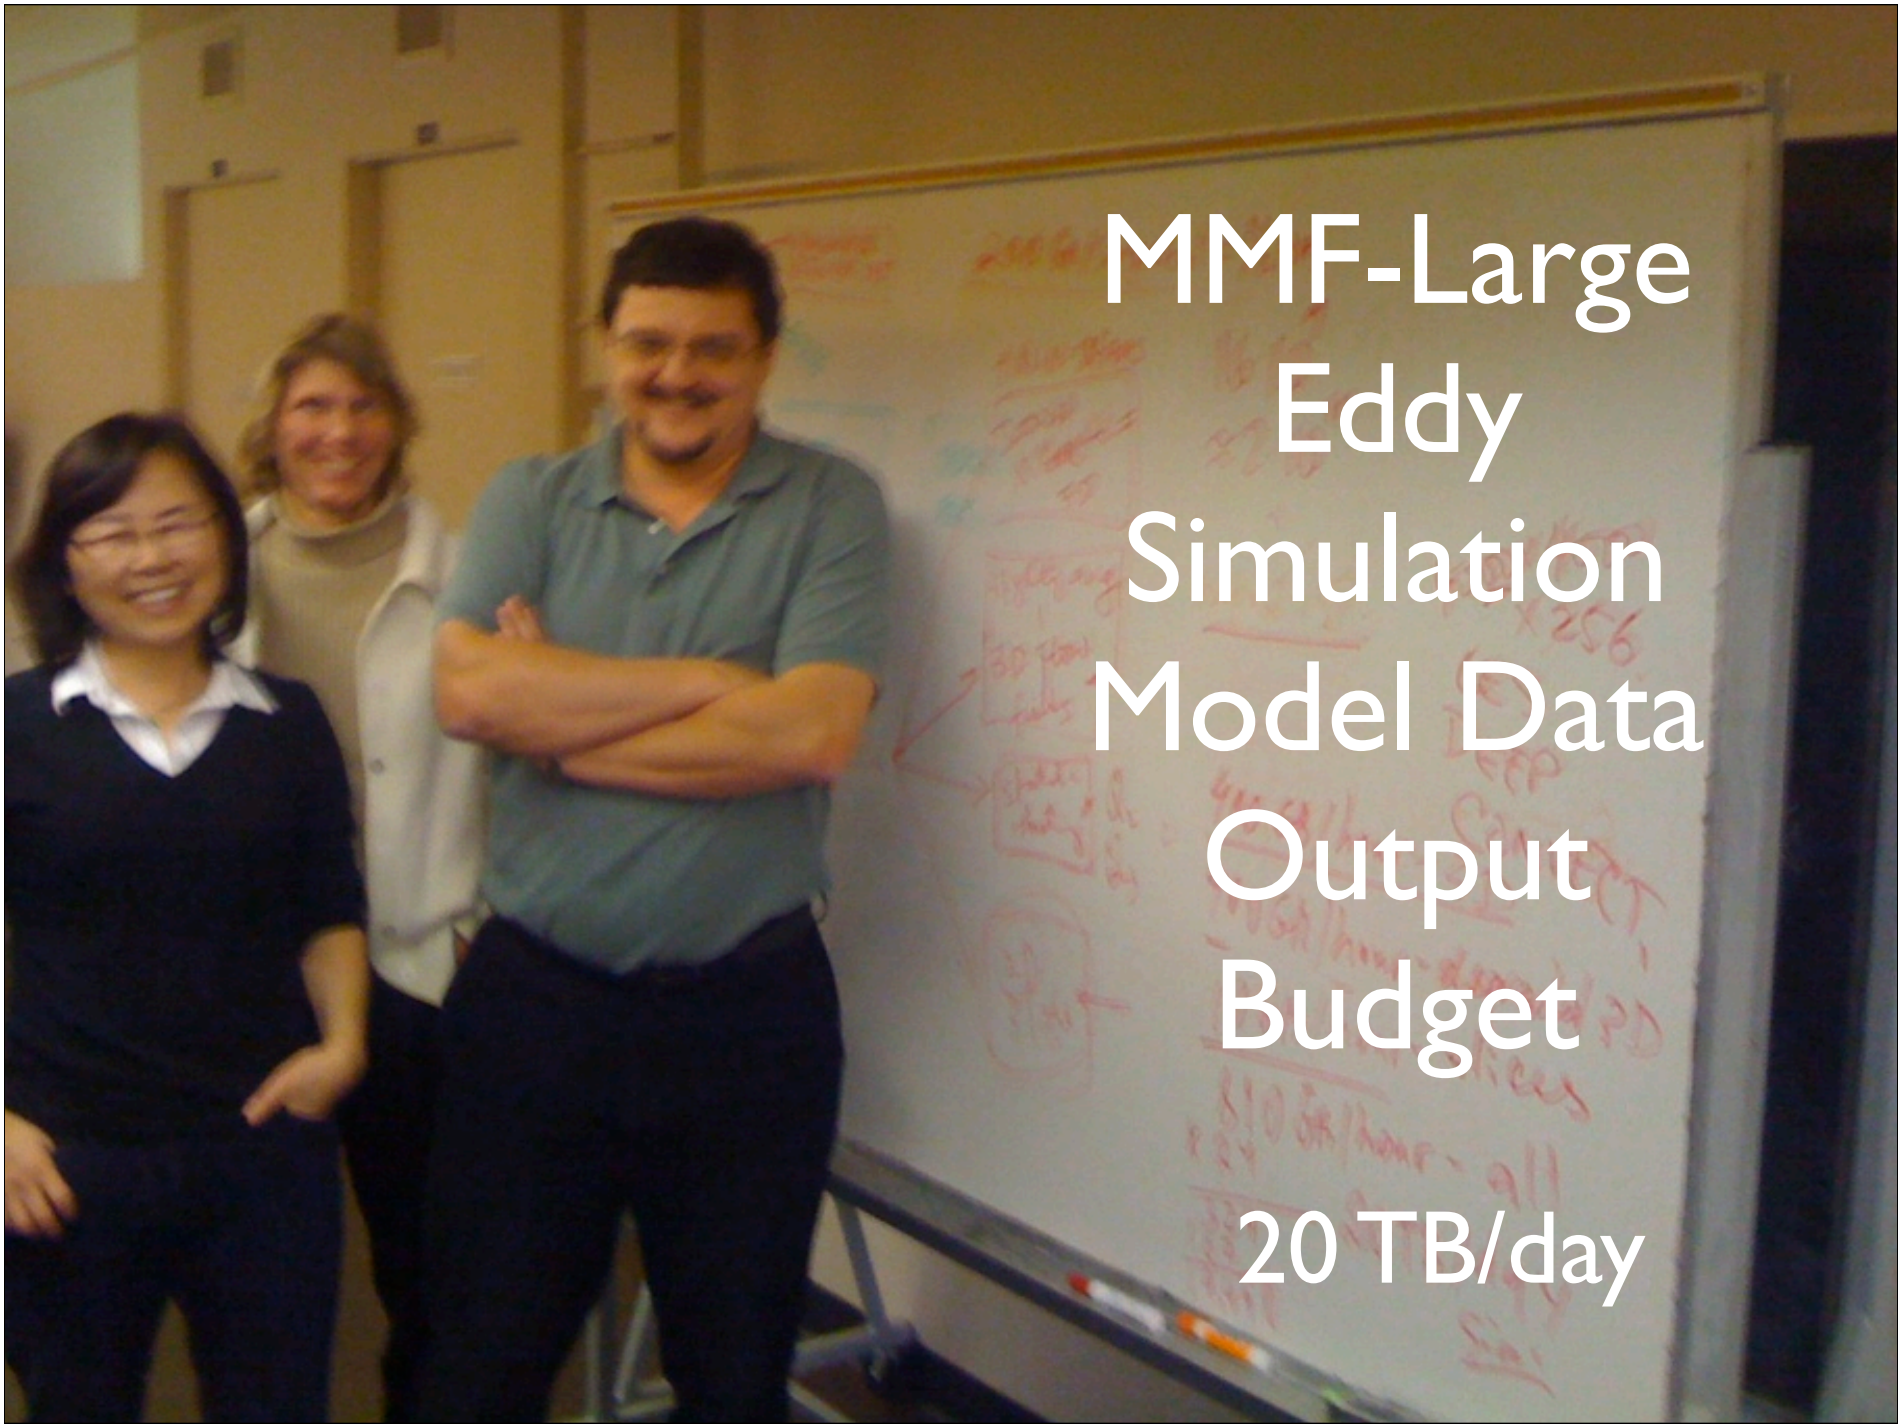
MMF-Large
Eddy
Simulation
Model Data
Output
Budget
20 TB/day

The GCRM Tsunami

2 km, 100 levels, hourly data

- ~4 TB / simulated hour
- ~100 TB / simulated day
- ~35 PB / simulated year

4 km, 100 levels, hourly data

- ~1 TB / simulated hour
- ~24 TB / simulated day
- ~9 PB / simulated year

Other Data Tsunami

- 30 TB/night: Large Synoptic Survey (LSS) Telescope (2014)
- 15 PB/year: CERN's Large Haydron Collider (May 2008)
- 1 PB over 3 years: EOS (Earth Observing System) data (2001)

(from K. Schuchart presentation January, 2008)

Data Challenges of GCRM

- Extremely high volumes of data
  - 10 GB/ variable / step
  - 1-10 petabytes / simulated year
  - Can’t just move data to local systems
  - Data will have to be on-offline
  - 4 byte offsets exceeded
  - Huge number of files per simulation
  - Model for running analysis on the entire data set needed

Data Challenges of the GCRM (cont)

- Geodesic Grid
  - Preliminary (but not sufficient) support in some analysis tools
  - Standards for complete description not defined
  - Hyperslab-ing on coordinate values not supported and very costly
  - Grid itself is large (~ 2GB)

Data Challenges of the GCRM (cont)

- Current scalar analysis tools break down
  - Insufficient memory
  - Assumptions made for smaller data sets no longer valid
  - Screen resolution exceeded

Strategic Planning

- CIWG is developing a strategy for handling model output with existing and projected resources
  - how much to write out?
  - how to subset and extract the right chunks of data for analysis?
  - how much can be stored for how long?
  - how much can be transported to what locations?

Dissemination and Interoperability


CMMAP
Digital Library

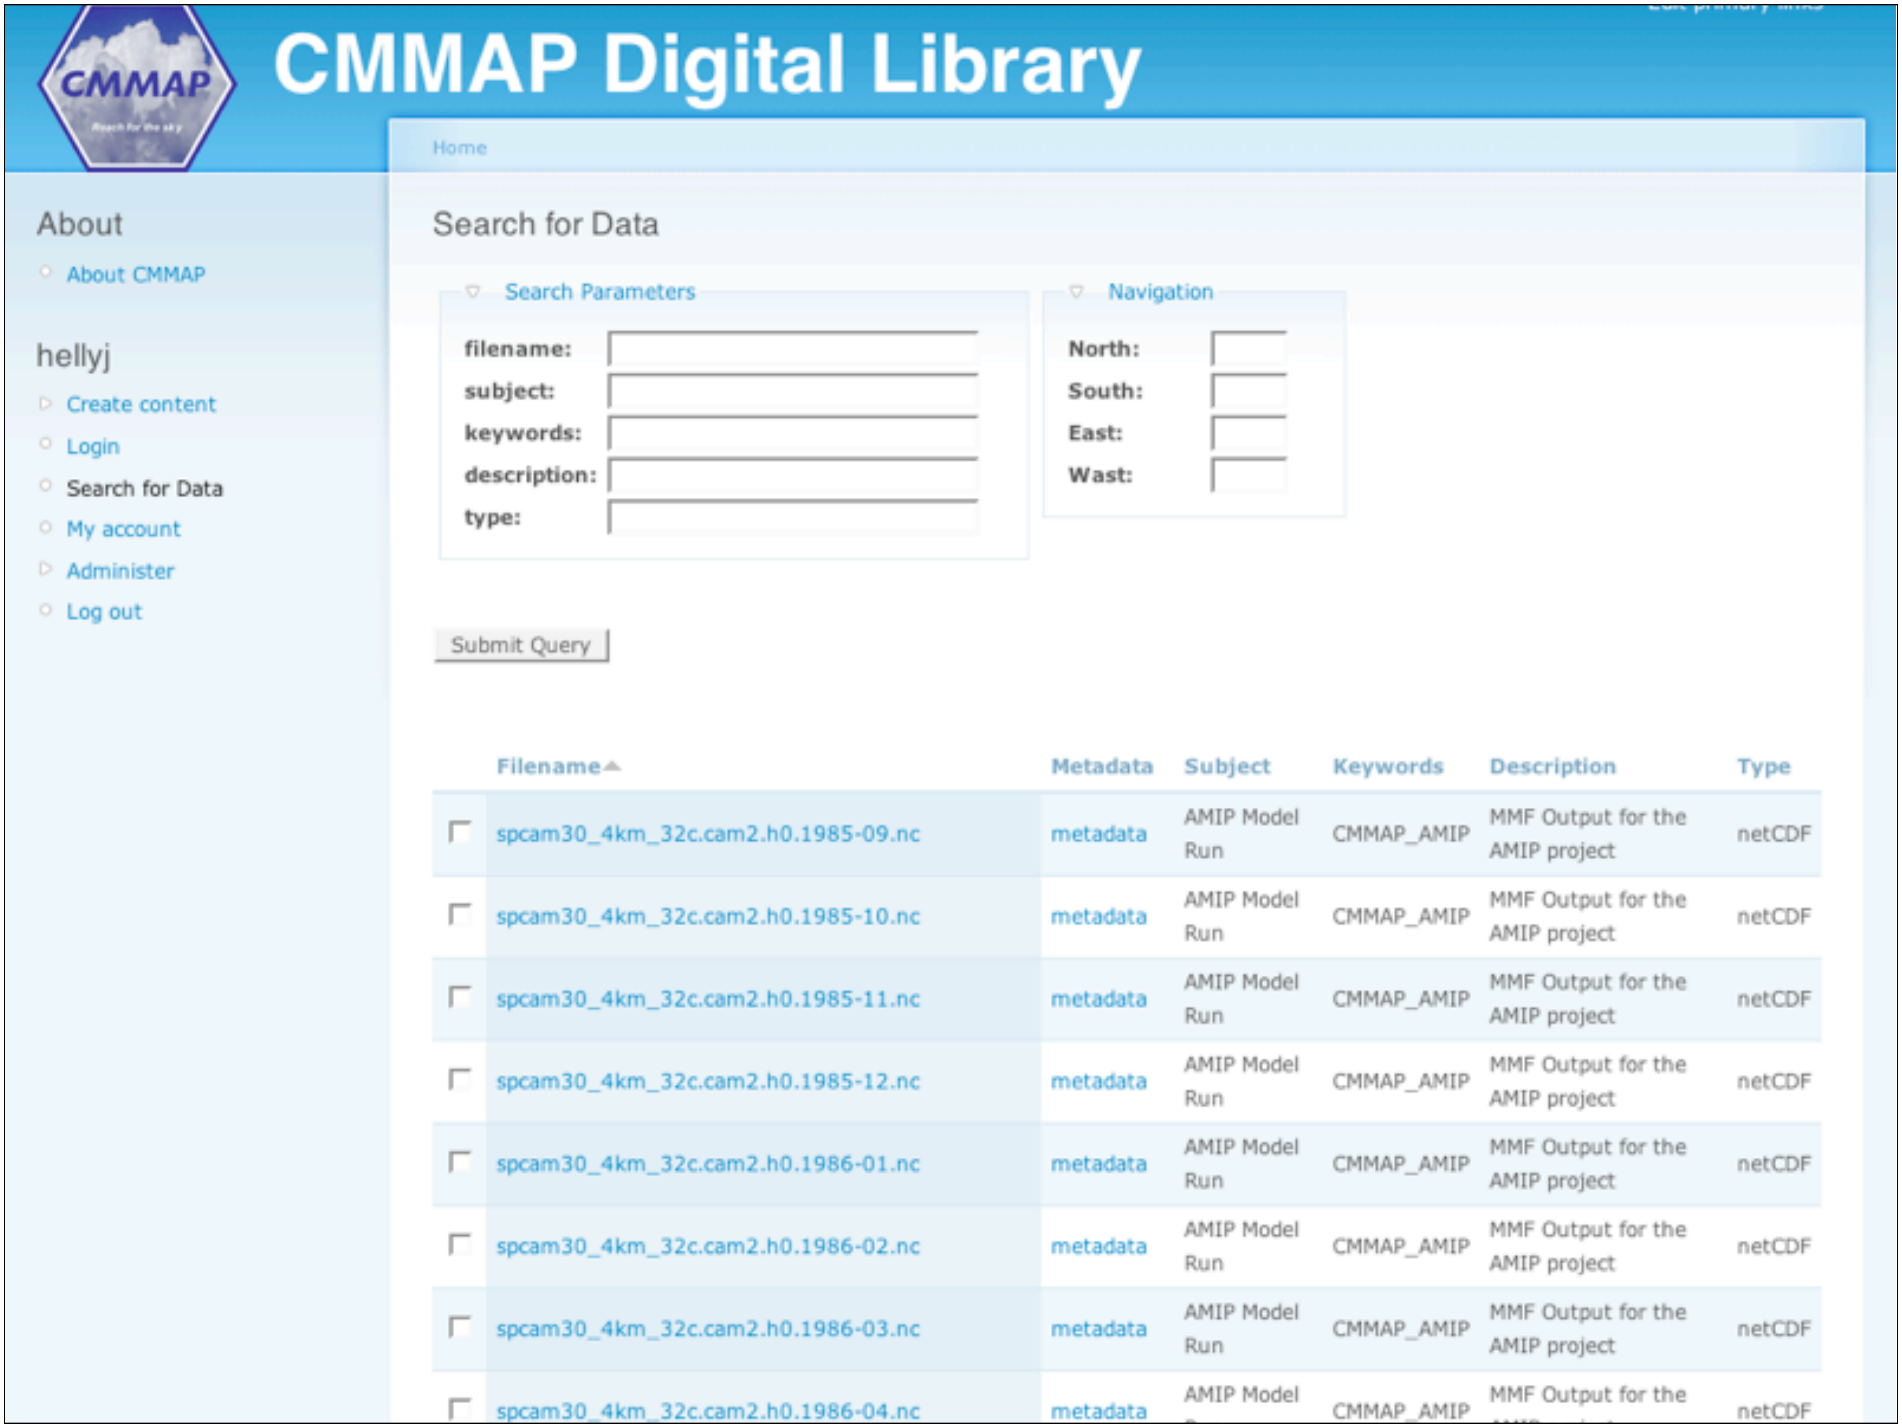
CMMAP Digital Library
Home
About
About CMMAP
hellyj
Create content
Login
Search for Data
My account
Administer
Log out
Search for Data
Search Parameters
filename:
subject:
keywords:
description:
type:
Navigation
North:
South:
East:
Wast:
Submit Query
Filename
Metadata
Subject
Keywords
Description
Type
spcam30_4km_32c.cam2.h0.1985-09.nc metadata AMIP Model Run CMMAP_AMIP MMF Output for the AMIP project netCDF
spcam30_4km_32c.cam2.h0.1985-10.nc metadata AMIP Model Run CMMAP_AMIP MMF Output for the AMIP project netCDF
spcam30_4km_32c.cam2.h0.1985-11.nc metadata AMIP Model Run CMMAP_AMIP MMF Output for the AMIP project netCDF
spcam30_4km_32c.cam2.h0.1985-12.nc metadata AMIP Model Run CMMAP_AMIP MMF Output for the AMIP project netCDF
spcam30_4km_32c.cam2.h0.1986-01.nc metadata AMIP Model Run CMMAP_AMIP MMF Output for the AMIP project netCDF
spcam30_4km_32c.cam2.h0.1986-02.nc metadata AMIP Model Run CMMAP_AMIP MMF Output for the AMIP project netCDF
spcam30_4km_32c.cam2.h0.1986-03.nc metadata AMIP Model Run CMMAP_AMIP MMF Output for the AMIP project netCDF
spcam30_4km_32c.cam2.h0.1986-04.nc metadata AMIP Model CMMAP_AMIP MMF Output for the netCDF

Standards and Conventions

Table 2. Tabulation of resources of interest identified by workshop participants.

Category	Common Name	Resource Description or Reference
Standards and Conventions	Climate and Forecast (CF) Standard Names	http://www.cgd.ucar.edu/cms/eaton/cf-metadata/CF-current.html#sname
	COARDs	ftp://ftp.unidata.ucar.edu/pub/netcdf/Conventions/COARDS
	Digital Object Identifier System	http://www.doi.org/
	Dublincore	http://dublincore.org/index.shtml
	EML	http://knb.ecoinformatics.org/software/eml/eml-2.0.1/index.html
	Fedora	http://www.fedora.info
	FGDC	http://www.fgdc.gov/
	Grib	http://www.wmo.ch/web/www/WDM/Guides/Guide-binary-2.html
	Geographic Markup Language	http://www.opengis.net/gml/
	Geoscience Markup Language	http://www.opengis.net/GeoSciML
	HDF	http://hdf.ncsa.uiuc.edu/
	ISO 19115	Defines the schema required for describing geographic information and services.
	ISO 19139	Geographic information -- Metadata -- XML schema implementation
	MMI	http://marinemetadata.org
	NetCDF	http://www.unidata.ucar.edu/software/netcdf/
	OAI	http://www.openarchives.org/
	OBIS	http://www.iobis.org/
	OpenGIS	http://www.opengeospatial.org/
	SensorML	http://www.opengeospatial.org/projects/groups/sensorweb
	STD-DOI	http://www.std-doi.de
Transport Protocol	GridFTP	http://www.globus.org/grid_software/data/gridftp.php
	HTTP	http://www.w3.org/Protocols/
	OPeNDAP	http://www.opendap.org/
	REST	http://en.wikipedia.org/wiki/Representational_State_Transfer#References
	SOAP	http://www.w3.org/TR/soap/

Adopting Community-based Standards and Conventions

Visualization

University of Melbourne has constructed a massive 96 million pixel "OptIPortal" visualization wall - known affectionately as the 'OzIPortal' - constructed from twenty-four 30-inch LCD screens.