Cyberinfrastructure

10th Team Meeting
Berkeley, CA
January, 2011

John Helly

CIWG Objectives

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

CIWG Agenda

- Current Status
  - NSF Architecture Planning Proposal (pending)
  - Teragrid Computing Allocations
    - Hugh Morrison's Experiments and Plans*
    - Community account for MMF runs
    - Exa-scale Data Transpose
  - Teragrid Allocation Calendar
- CMMAP CI Architecture (Roadmap)
  - Subversion MMF Development Repository (Mark B.)
  - iRODS web-browser and parallel data transfer service
  - MMF (SP-CAM) Community Account Portal
- Data Transportation
  - Parallel File I/O (Jeff Daily, Karen Schuchardt)
  - Data transpose project on Dash and Triton (prelim performance)
- Discussion

Current Status

Current Allocations

Home | **My TeraGrid** | Resources | User Forums | Documentation | Training | Consulting | Allocations

Allocations/Usage Accounts Profile Tickets Registered DNs Change Portal Password Add/Remove User Community Account SSH Terminal Citation Info

Allocations/Usage

Projects — Show Inactive Projects | Show Expired Allocations

Modeling Global Climate Variability with the Multi-scale Modeling Framework New parameterizations of Cloud Micro-physics and Developing Community Community Accounts Portal for Running the MMF

Show Project Details

Allocations*

Start Date	End Date	Resource	SUs Remaining	SUs Awarded	My Usage (SU)	% Remaining	Alloc. Type	State
2010-07-01	2011-06-30	abe-queenbee-steele-lonestar.teragrid	2,307,000	2,307,000	0.0	100 %	new	active
						Show Users on abe-queenbee-steele-lonestar.teragrid		
2010-07-01	2011-06-30	NCSA Tape	5	5	0.0	100 %	new	active
						Show Users on NCSA Tape		

Data Transposition Development for Exa-scale Data in Memory

Show Project Details

Allocations*

Start Date	End Date	Resource	SUs Remaining	SUs Awarded	My Usage (SU)	% Remaining	Alloc. Type	State
2010-05-04	2010-11-07	dash.sdsc.teragrid	29,970	30,000	30.0	100 %	supplement	active
						Show Users on dash.sdsc.teragrid		
2009-11-07	2010-11-07	Spur	30,000	30,000	0.0	100 %	new	active
						Show Users on Spur		

Regionalization of Anthropogenic Climate Change Simulations

Show Project Details

Allocations*

Start Date	End Date	Resource	SUs Remaining	SUs Awarded	My Usage (SU)	% Remaining	Alloc. Type	State
2009-04-01	2010-09-30	Ranger	43,186	3,950,000	0.0	1 %	new	active
2009-04-01	2010-09-30	Spur	397	500	103.0	79 %	new	active

Modeling Global Climate Variability with the Multi-scale Modeling Framework: The Boundary-layer Cloud Problem

Show Project Details

Allocations*

Start Date	End Date	Resource	SUs Remaining	SUs Awarded	My Usage (SU)	% Remaining	Alloc. Type	State
2009-04-01	2010-09-30	Steele	103,900	950,000	0.0	11 %	new	active
						Show Users on Steele		

Leveraging National & Partner Resources

	Organization	Resource	2007	2008	2009	2010	2011
Data Allocations	San Diego Supercomputer Center (SDSC)	Disk	15 Terabytes	15 Terabytes	30 Terabytes	45 Terabytes	45 Terabytes
		BlueGene			30,000 SUs*		
		Triton				30,000 SUs	30,000 SUs
Computing Allocations	Teragrid (multi-institution)	SDSC DataStar (IBM SP4)	600,000 SUs	1,200,000 SUs			
		Grid Roaming			600,000 SUs	2,703,000 SUs	
		LSU Steele			900,000 SUs		2,307,000 SU
		SDSC (Dash)					60,000 SU
	Lawrence Berkeley National Laboratory (LBNL)	National Energy Research Scientific Computing Center (NERSC)			700,000 SUs		
	Oak Ridge National Laboratory (ORNL)	Cray XT			2,000,000 hrs	3,000,000 hrs	
	National Center for Atmospheric research (NCAR)	BlueIce IBM Power5			500,000 SUs		

TeraGrid Allocation Calendar

TeraGrid Home > User Support > Access >> Allocations & Accounts > Allocation Calendar

	Startup/Educational Allocation		Research Allocation (TRAC)	
Units Requested *Service Units (SUs) on compute resources (TeraGrid Resource Catalog)* *Terabytes (TBs) on Data Resources*	**System size**	**Maximum compute request**	30,000 – Unlimited	
	< 100 TFLOPS	Up to 30,000 SUs		
	>= 100 TFLOPS	Up to 200,000 SUs		
	▪ Aggregate request for multiple compute resources cannot exceed 200,000 SUs			
	▪ Storage on disk: 5 TB ▪ Storage on tape: 25 TB			
Deadlines	N/A		**Open Submissions** Dec. 15 Mar. 15 [1] Jun. 15 Sept. 15	**Close Submissions** **Jan. 15** [1] **Apr. 15** **Jul. 15** **Oct. 15**
Allocations Begin	Two weeks after submission		April 1 July 1 October 1 January 1	
Review Cycle	Within one week		Quarterly [2]	
Typical Use	Classroom or training accounts and startup accounts requiring small amounts of time		Experienced users with research projects	

CMMAP Architecture Roadmap

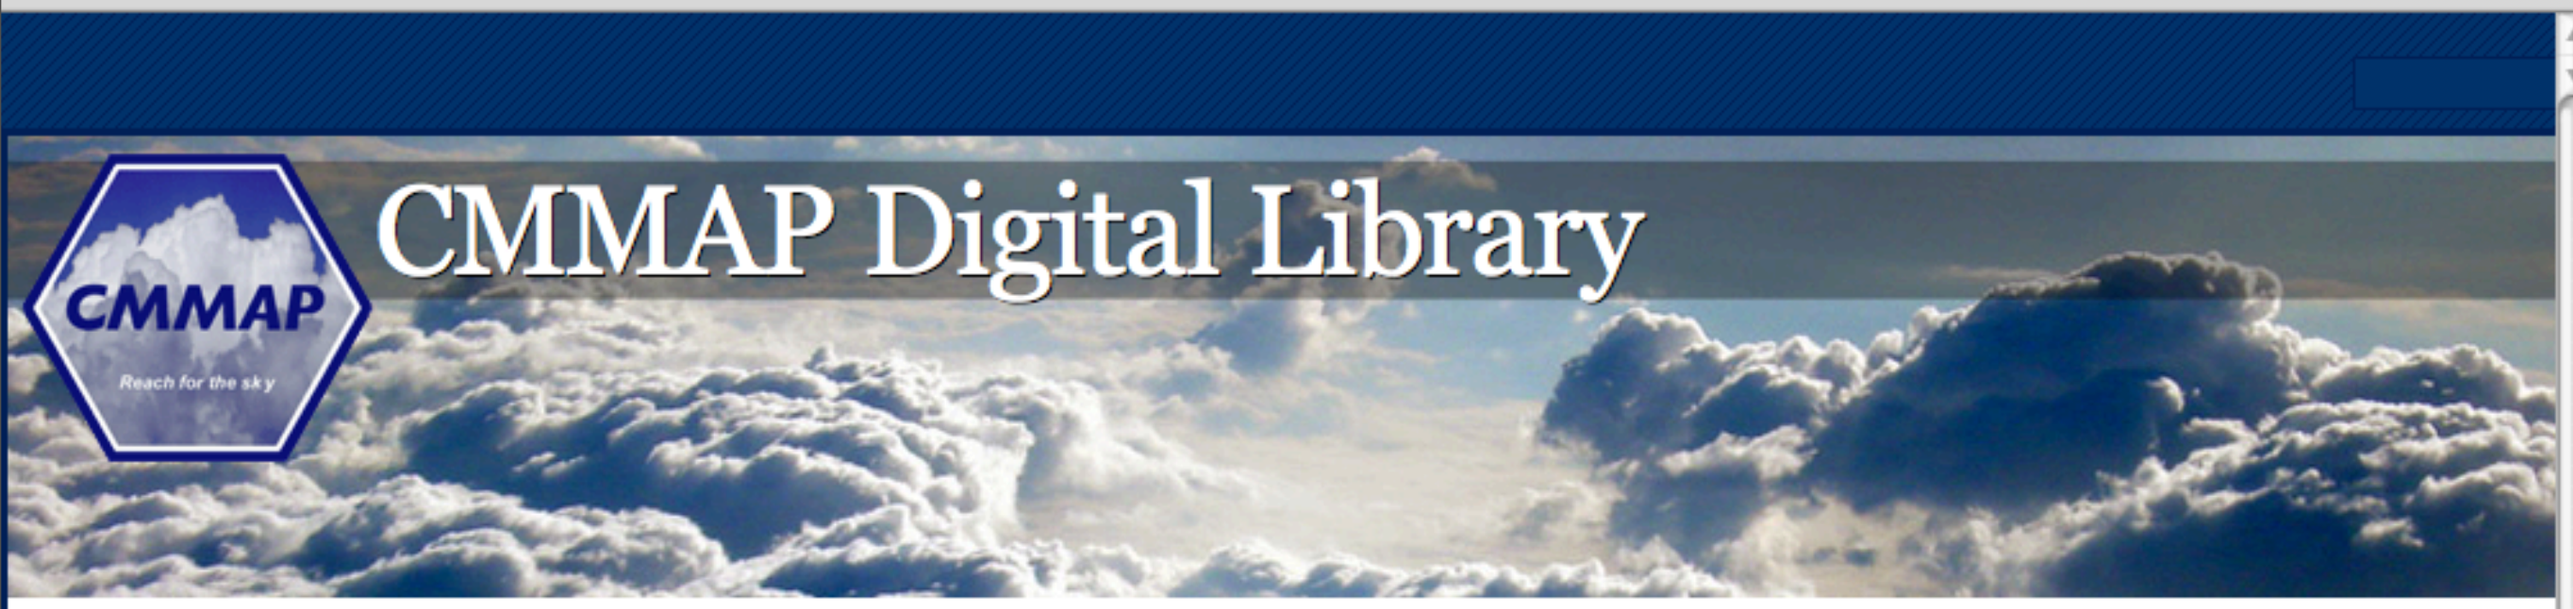

DATA RESOURCES

- Data Collections Browser
- Search Metadata Catalogue

COMPUTING RESOURCES

- NSF Teragrid
- DOE INCITE

SOFTWARE

- Bulk Data Transfer Client
- Client software for accessing CMMAP data holdings.
- Model Development Team
- Multi-scale Modeling

Data Collections Browser now available from CMMAP Digital Libary

This browser makes it possible to conveniently browse the data holdings of the CMMAP Digital Library. An account is r
so contact John Helly (hellyj@ucsd.edu) or Mark Branson (mark@atmos.colostate.edu) for access.

hellyj's blog Add new comment

Subversion repository account creation or password resetting.

Repository URL:

https://svn.sdsc.edu/repo/cmmap

Obtaining or Re-setting a password.

1. To generate a new password from any unix host, please run the following:
htpasswd -mn

2. The output should look something like this: jd:$apr1$L7wBD/..$l.koeYBEZ3TfM.qOW6fXr0

3. Copy and paste that output into an email to jd@sdsc.edu with the subject:

Please add or replace this user in the CMMAP subversion repository.

INCITE Resources & Allocations

Subversion Repository

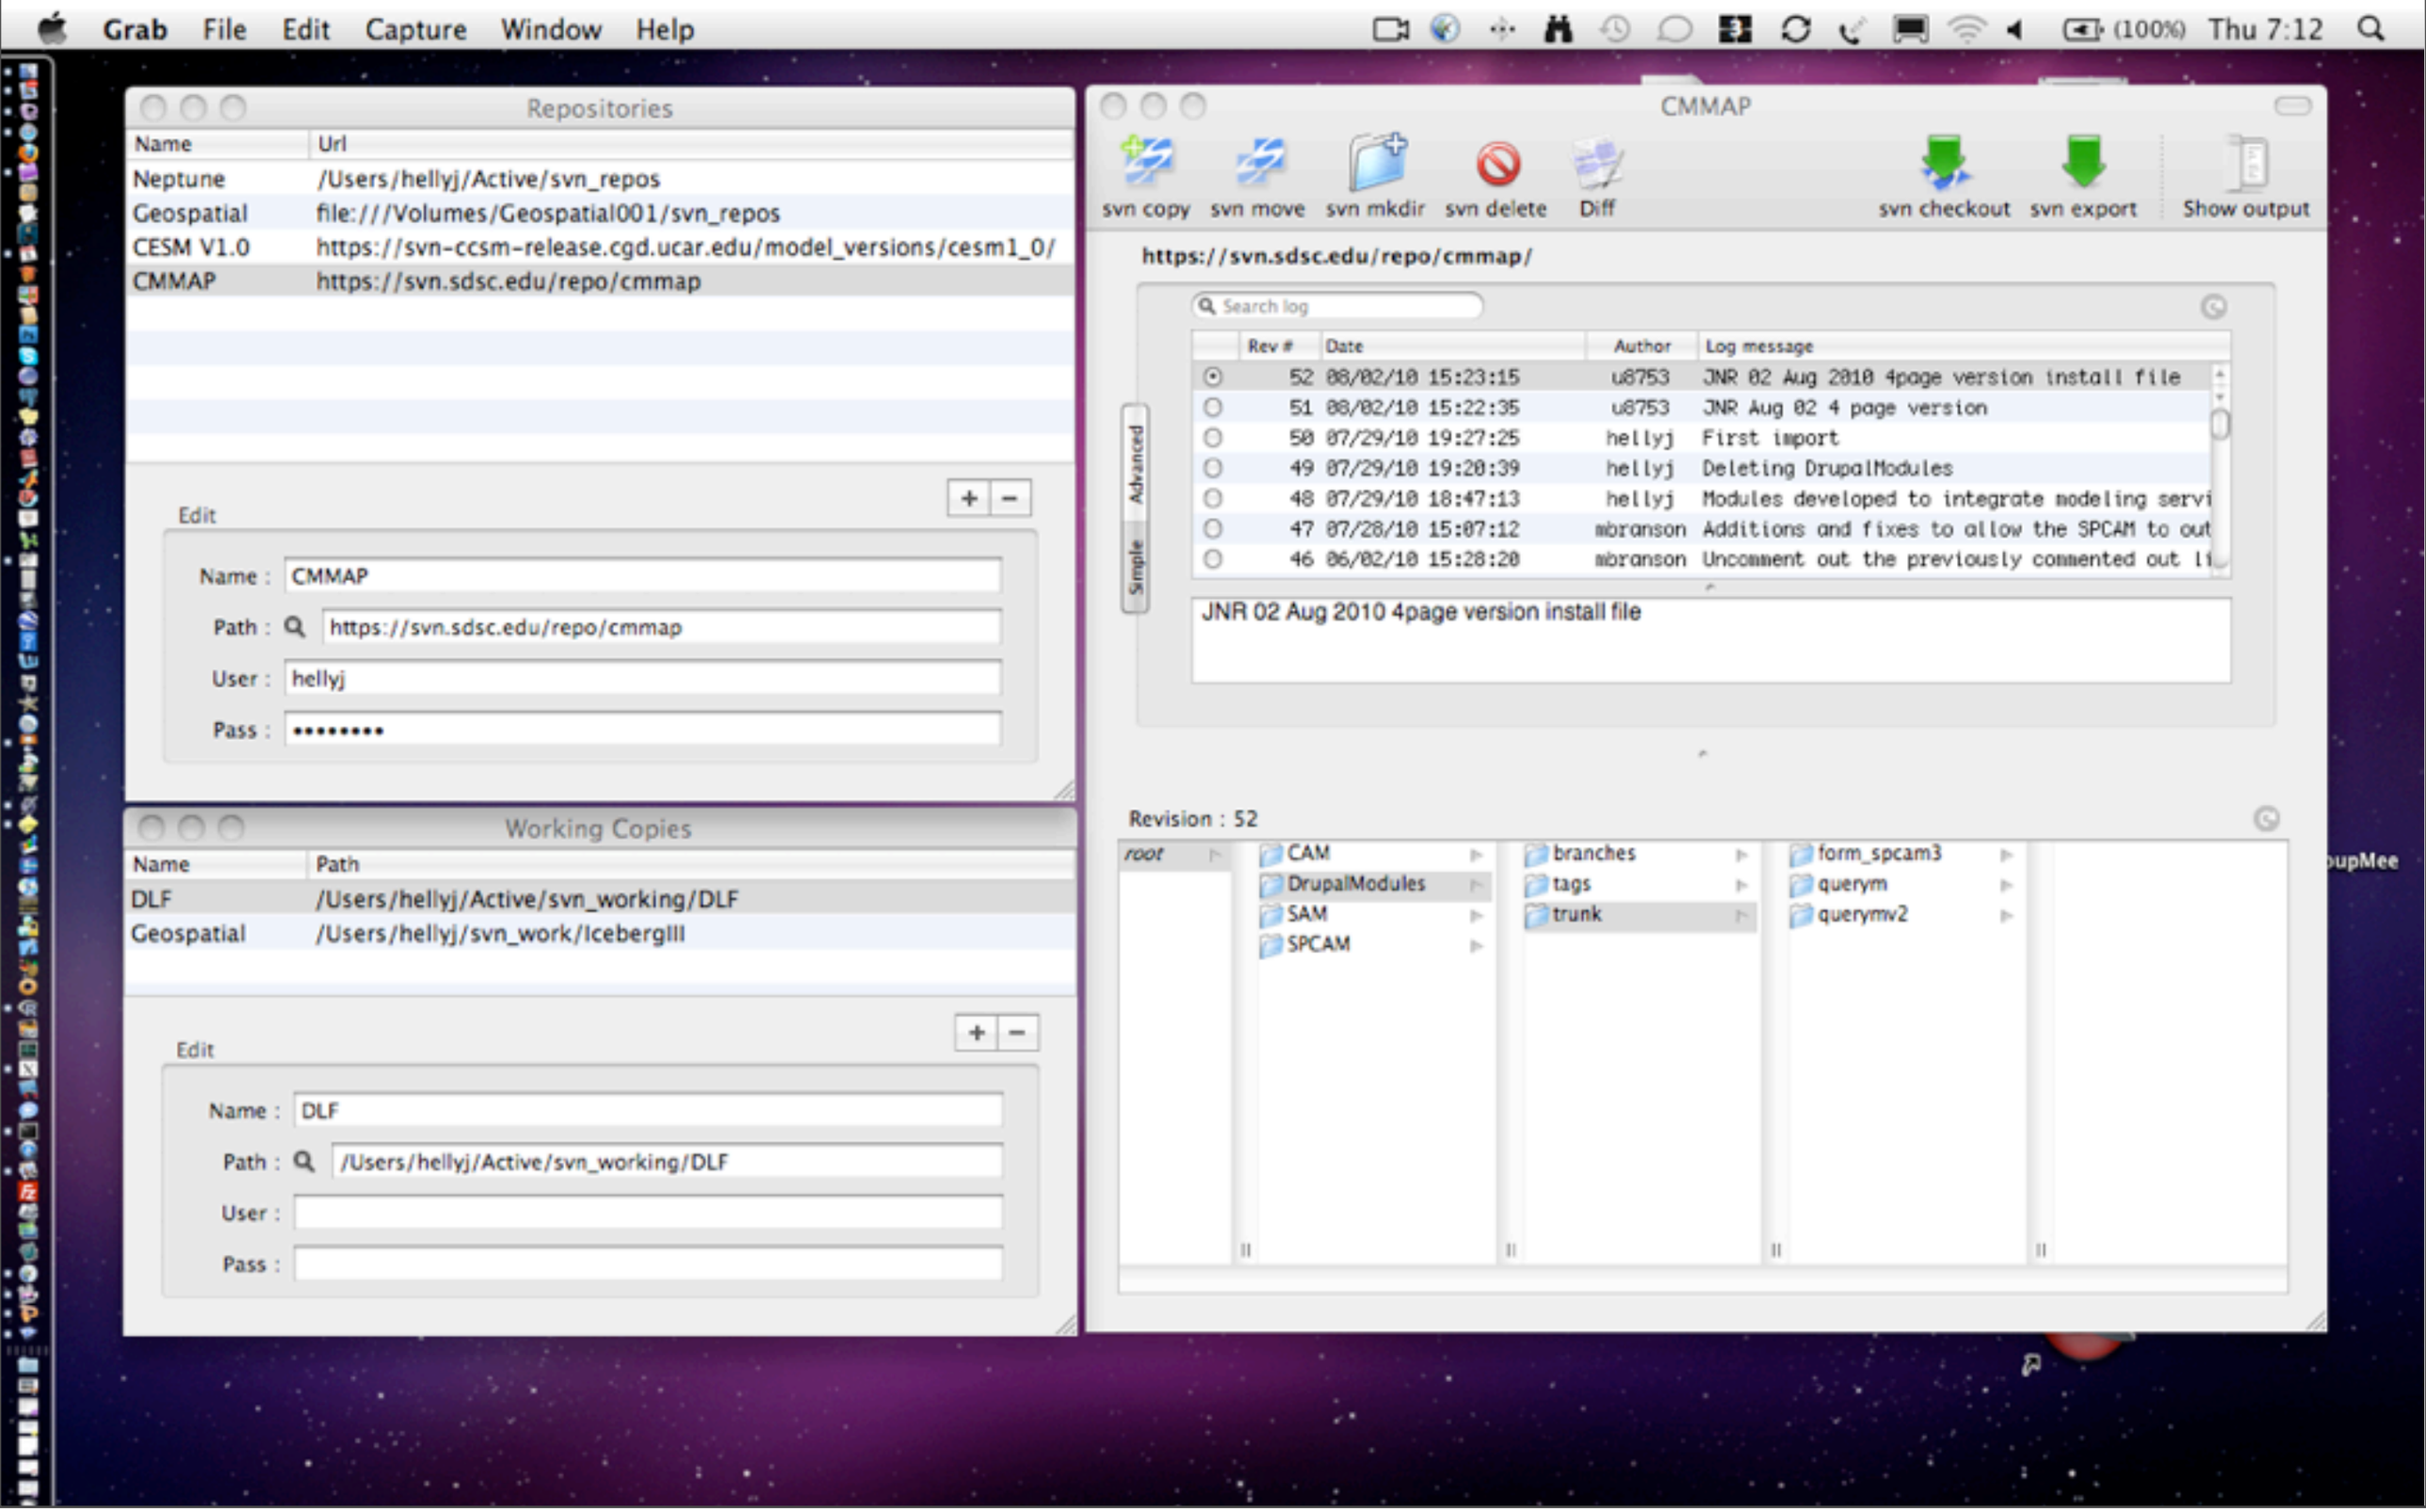

iRODS Command-line Client: Scriptable

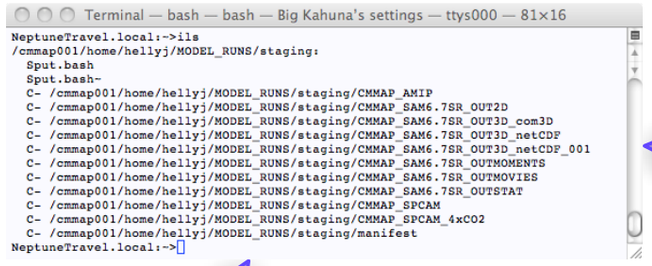

iRODS Federation Mechanism

CMMAP Phase II CI @CSU

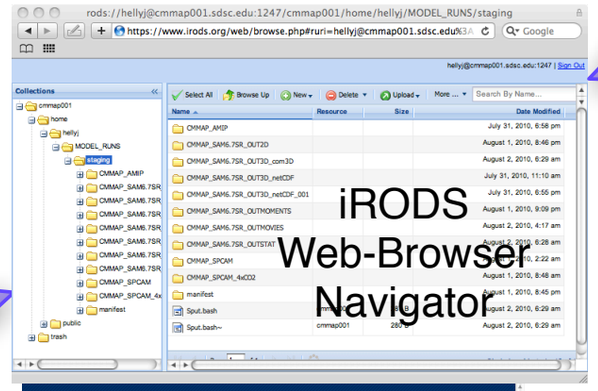

iRODS Web-Browser Navigator

MMF Community Portal

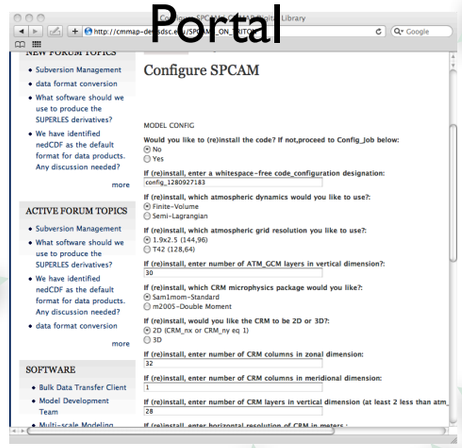

CMMAP Disk Allocation @SDSC

Drupal: programmability

- iquery
- MMF Workbench

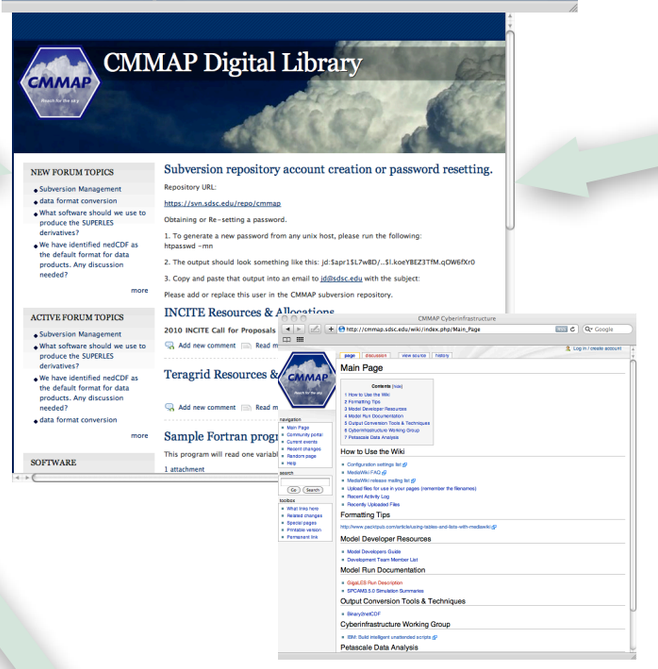

MediaWik: Documentationi

FRE
Workflow Control
Collaboration w/GFDL

others

Community Account

Steele@Teragrid

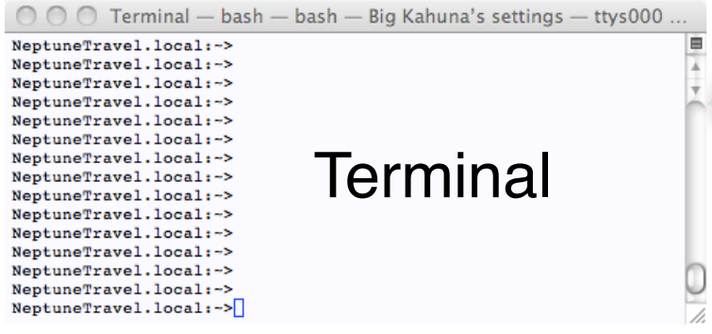

Terminal

Computing

Data

iRODS Command-line Client: Scriptable

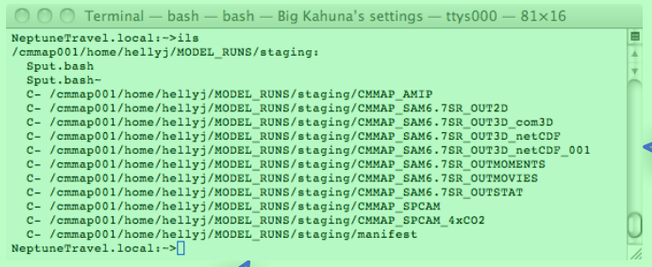

iRODS Federation Mechanism

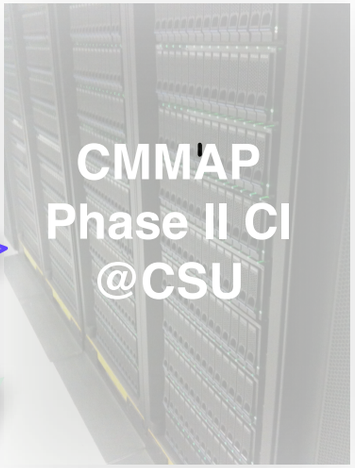

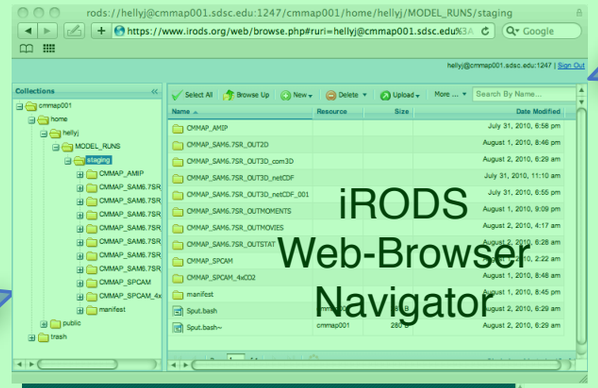

iRODS Web-Browser Navigator

Drupal: programmability

- iquery
- MMF Workbench

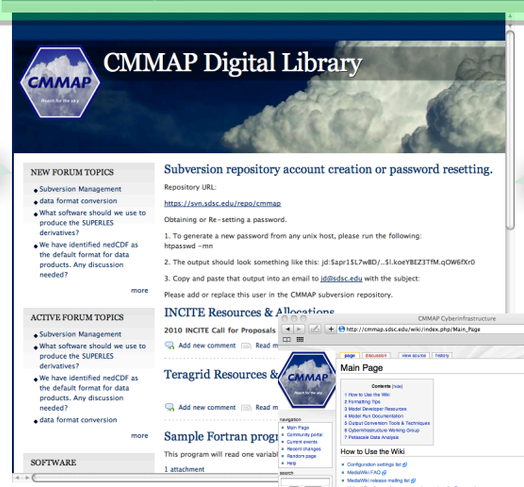

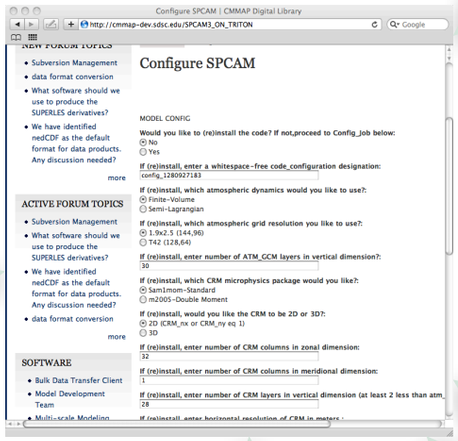

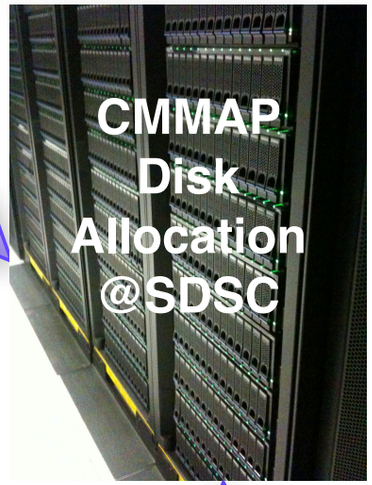

MediaWik: Documentationi

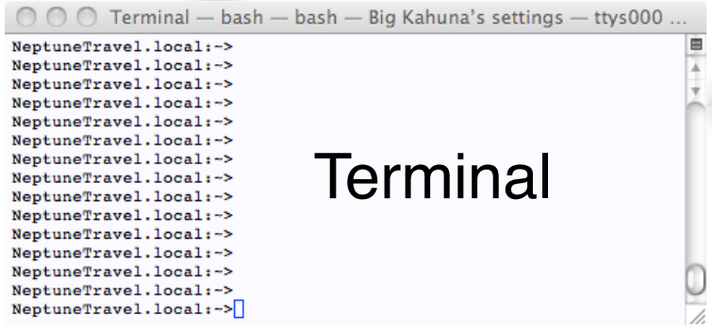

Terminal

FRE
Workflow Control
Collaboration w/GFDL

others

Community Account

Steele@Teragrid

Computing

Data

iRODS Web-broswer

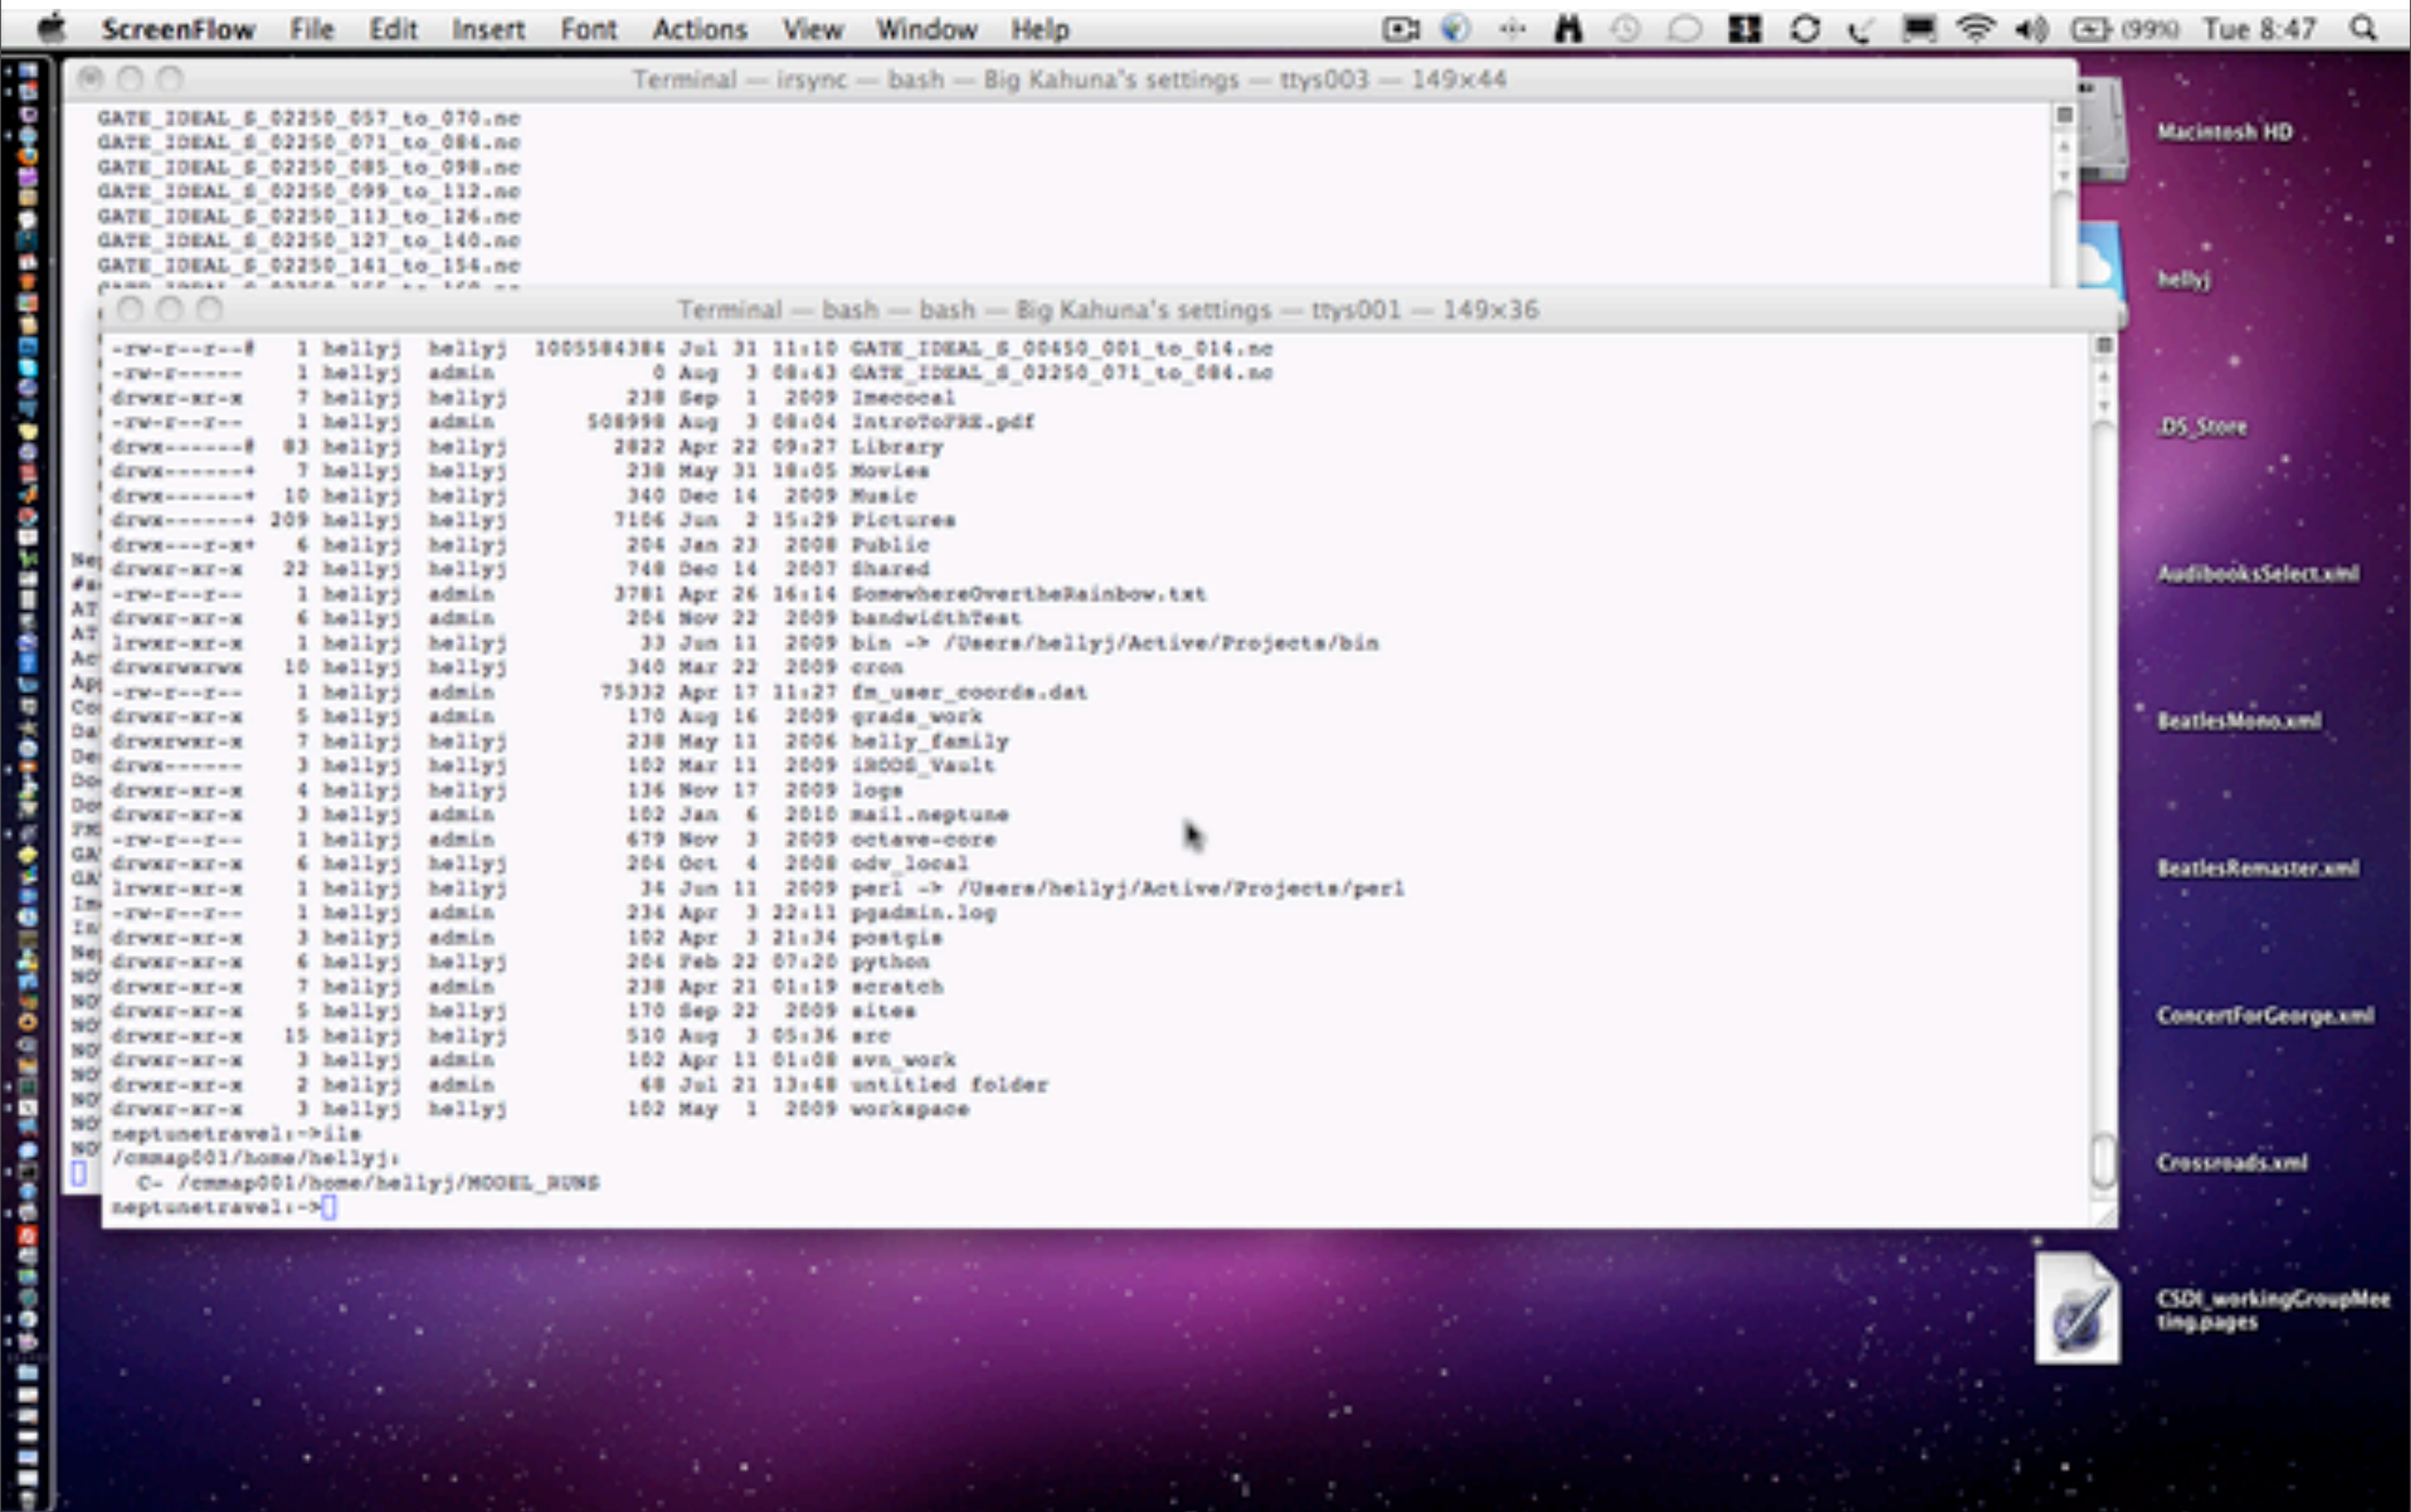

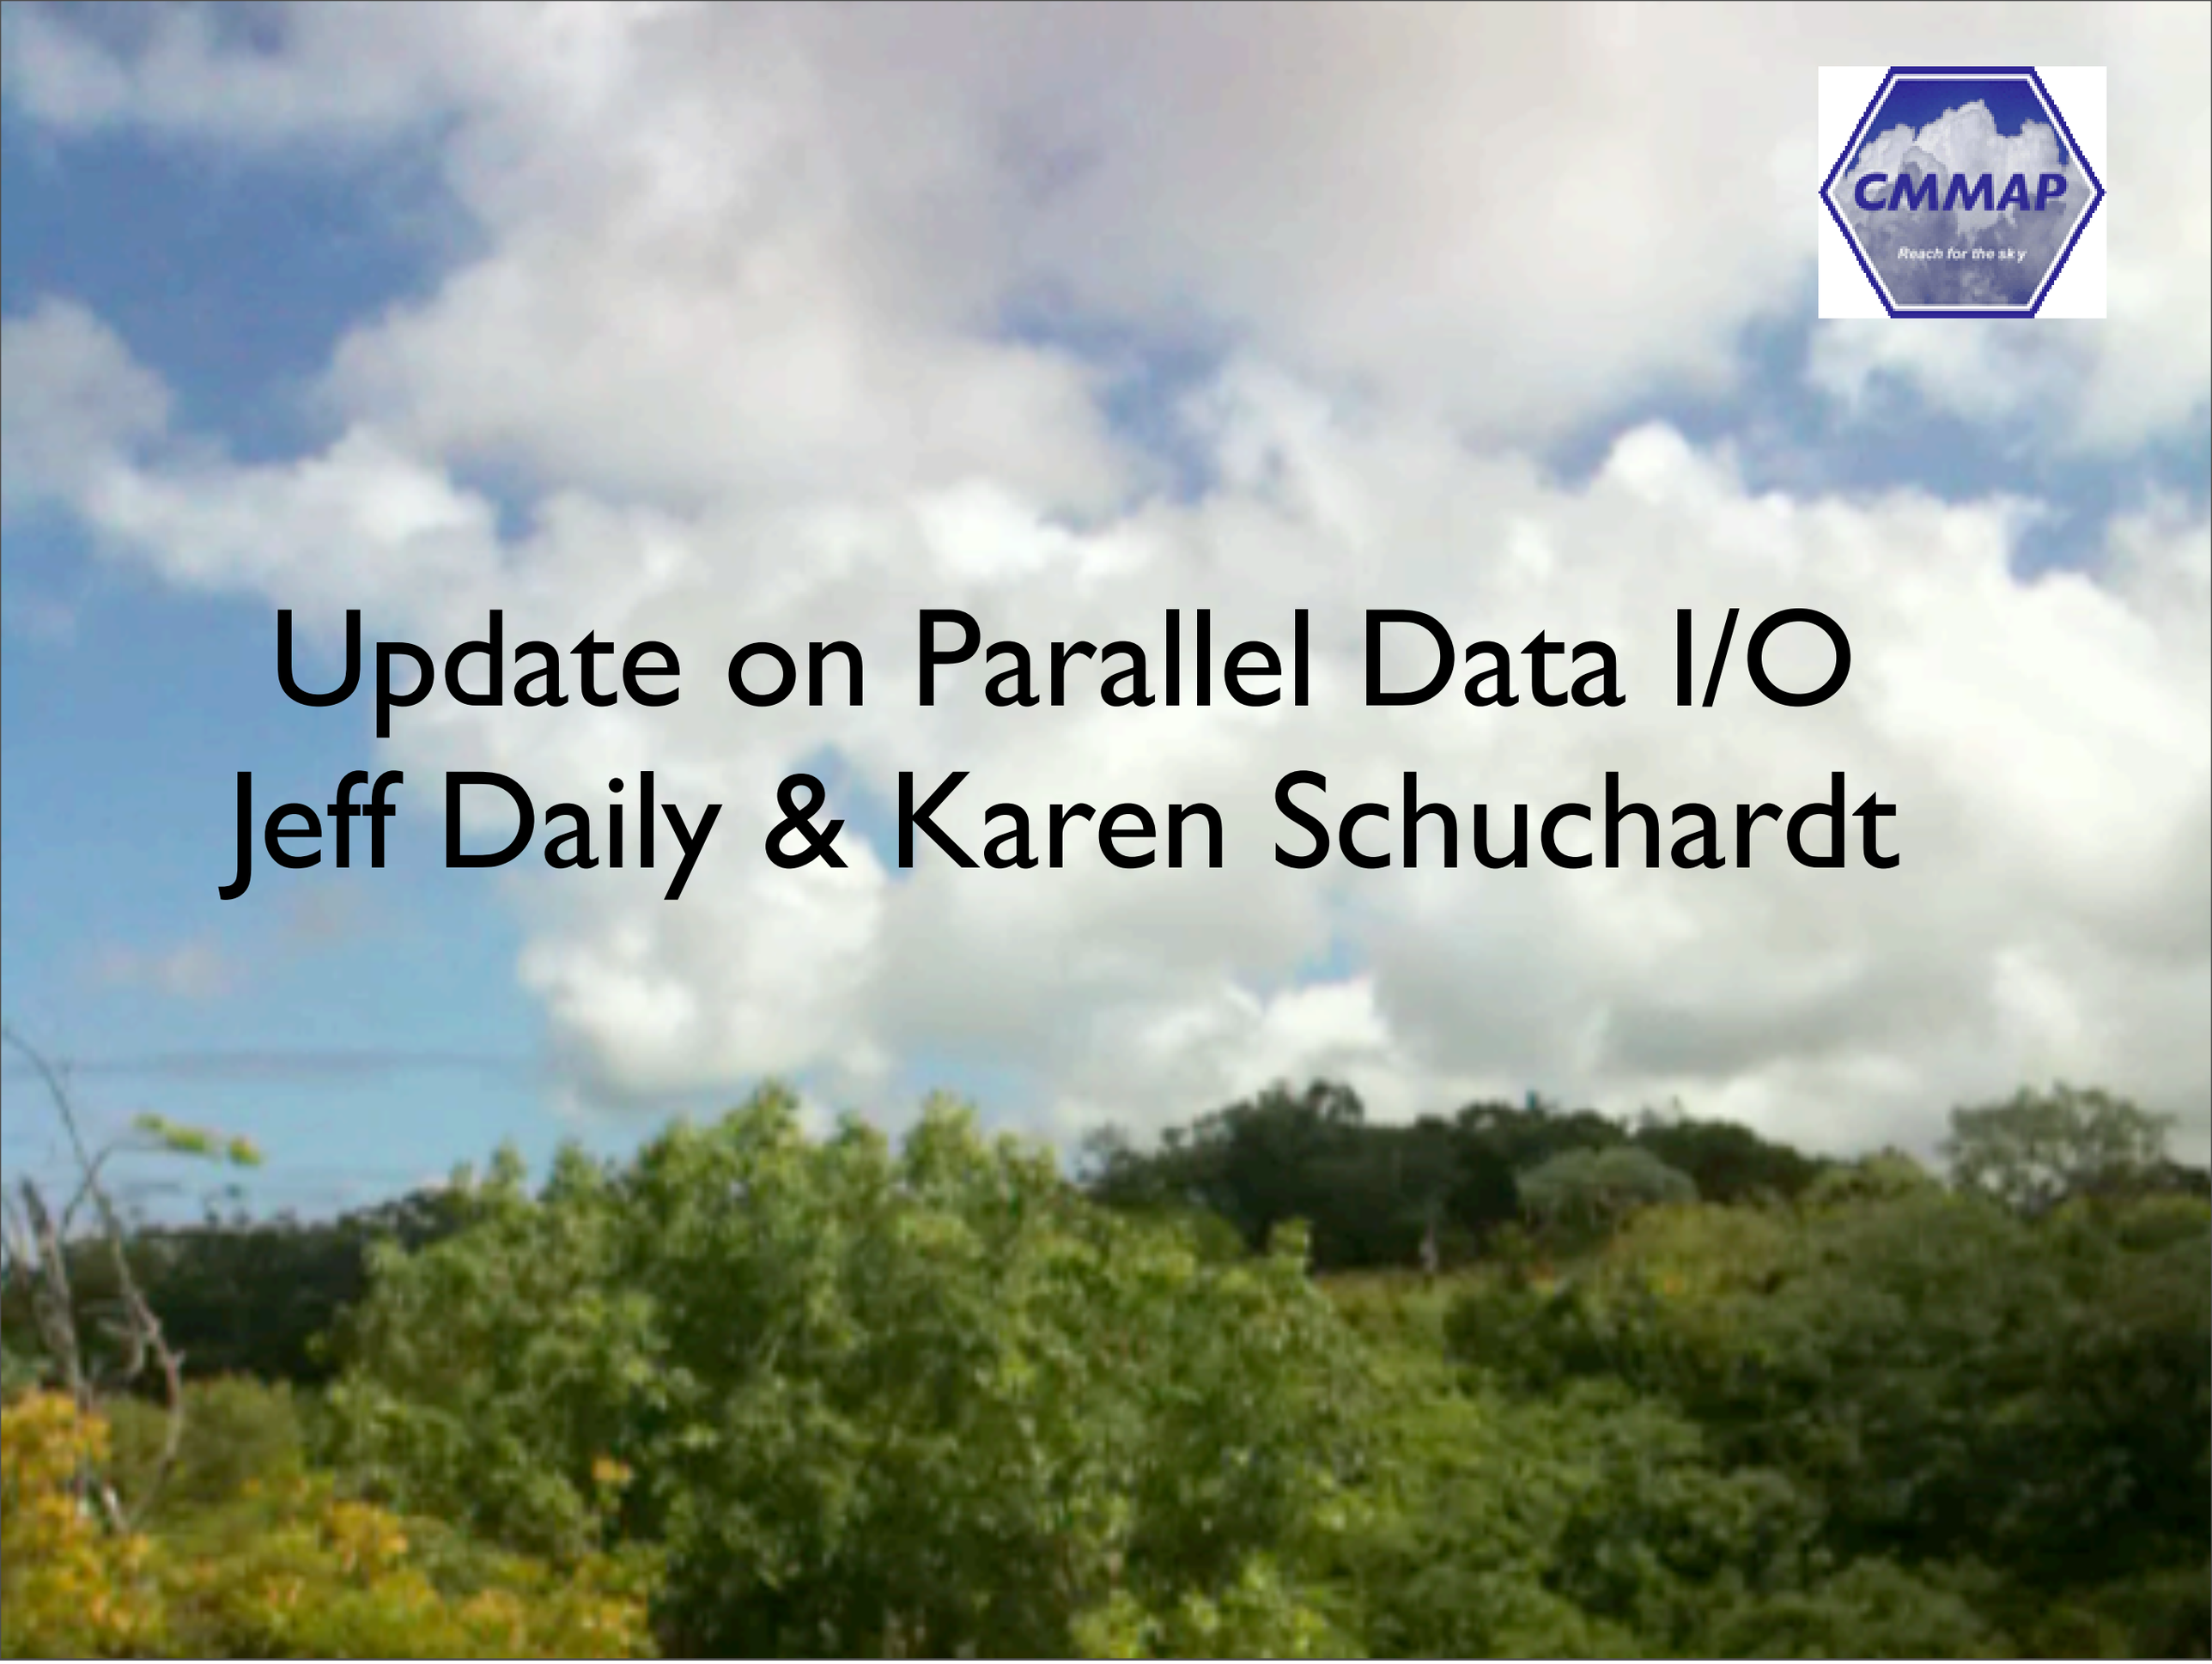

Update on Parallel Data I/O

Jeff Daily & Karen Schuchardt

GCRM Data Problem

The GCRMs and Giga-LES models are conceptually complex, but in addition they pose problems that are technical, practical, and fiscal, rather than conceptual in nature. This is where the need for new infrastructure arises. Our proposed infrastructure project relates to data management, analysis, and visualization:

- GCRMs produce terabytes to petabytes of model output. The data is created at supercomputer centers. It must be archived, curated, and made available to users at remote sites.
  - Many difficult choices must be made; for example, choosing which fields to output, and what subsetted spatial and temporal resolutions to save, are complex.
  - Routinely saving global model output with high temporal resolution is not practical.
  - A possible strategy is to save regional model output (for one or more selected regions) with high temporal resolution, and full spatial resolution, and global model output with lower temporal resolution and perhaps even reduced spatial resolution.
- Extraction of useful information from GCRM output is complicated by the sheer volume of data produced, the wide range of scales represented, and the diverse phenomena included. New methods are needed for comparison of model output with a variety of observations, including satellite data.
- New methods are needed for the efficient and effective visualization of GCRM results. The range of scales is so large that "zooming" capabilities are essential. New approaches are needed to visualize and analyze the time evolution of complex three-dimensional structures (such as large rotating convective clouds) that are associated with multiple interacting fields, including vector fields.

In short, *the very large models used in cloud-climate studies must be supported by a suitably designed infrastructure for data management, analysis, and visualization*. These needs are community wide and should be addressed in a coordinated fashion that serves the community as a whole.

Internet Data Transfer Capacity

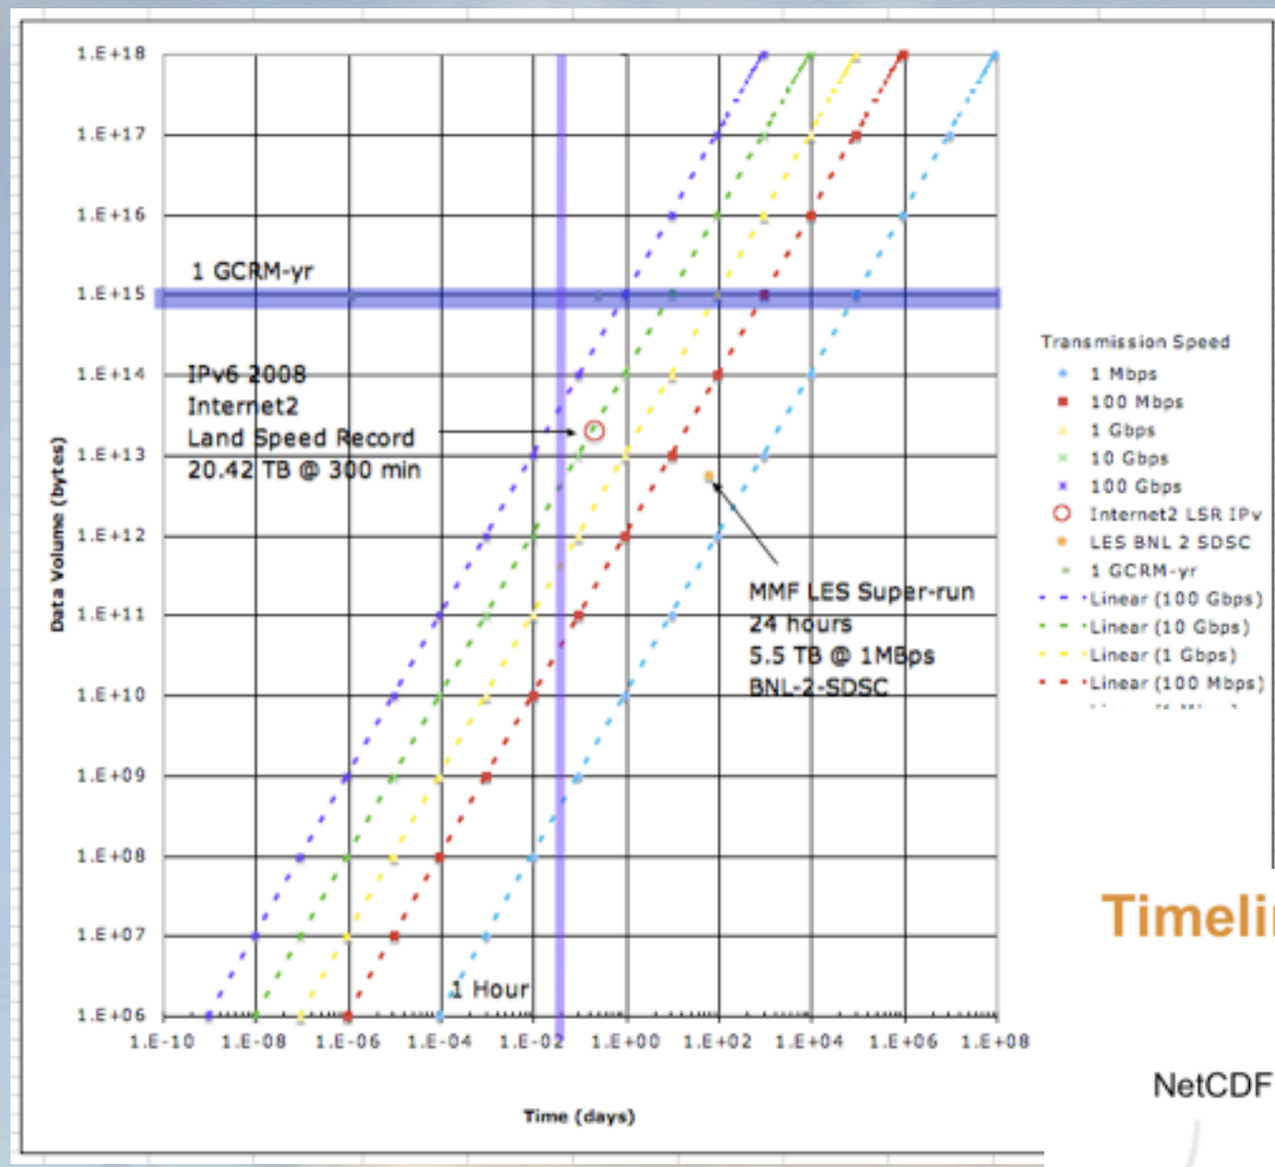

Parallel Input/ Output Technology Progress

Timeline

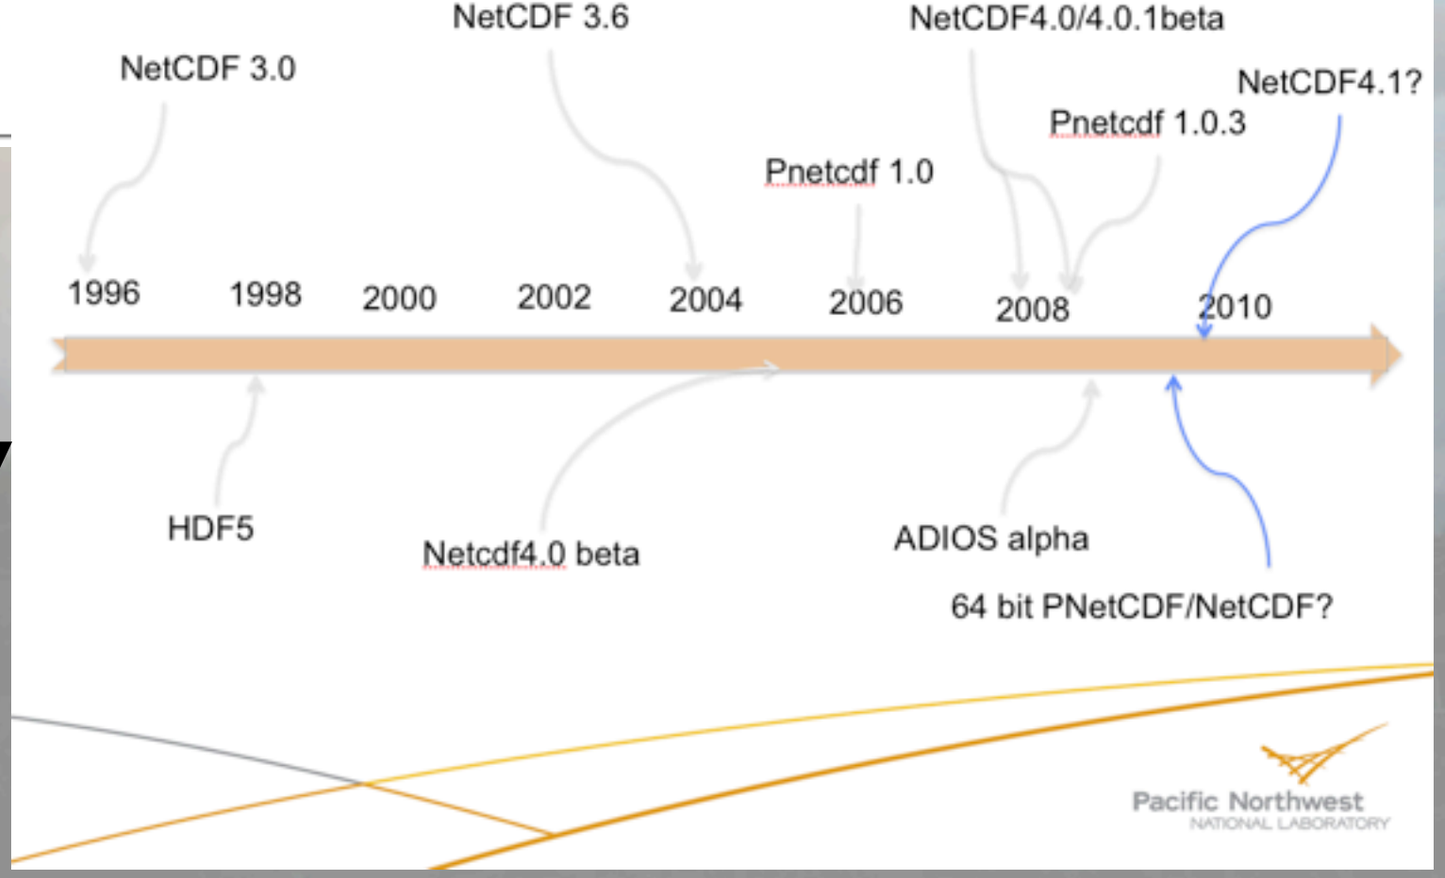

Pacific Northwest
NATIONAL LABORATORY

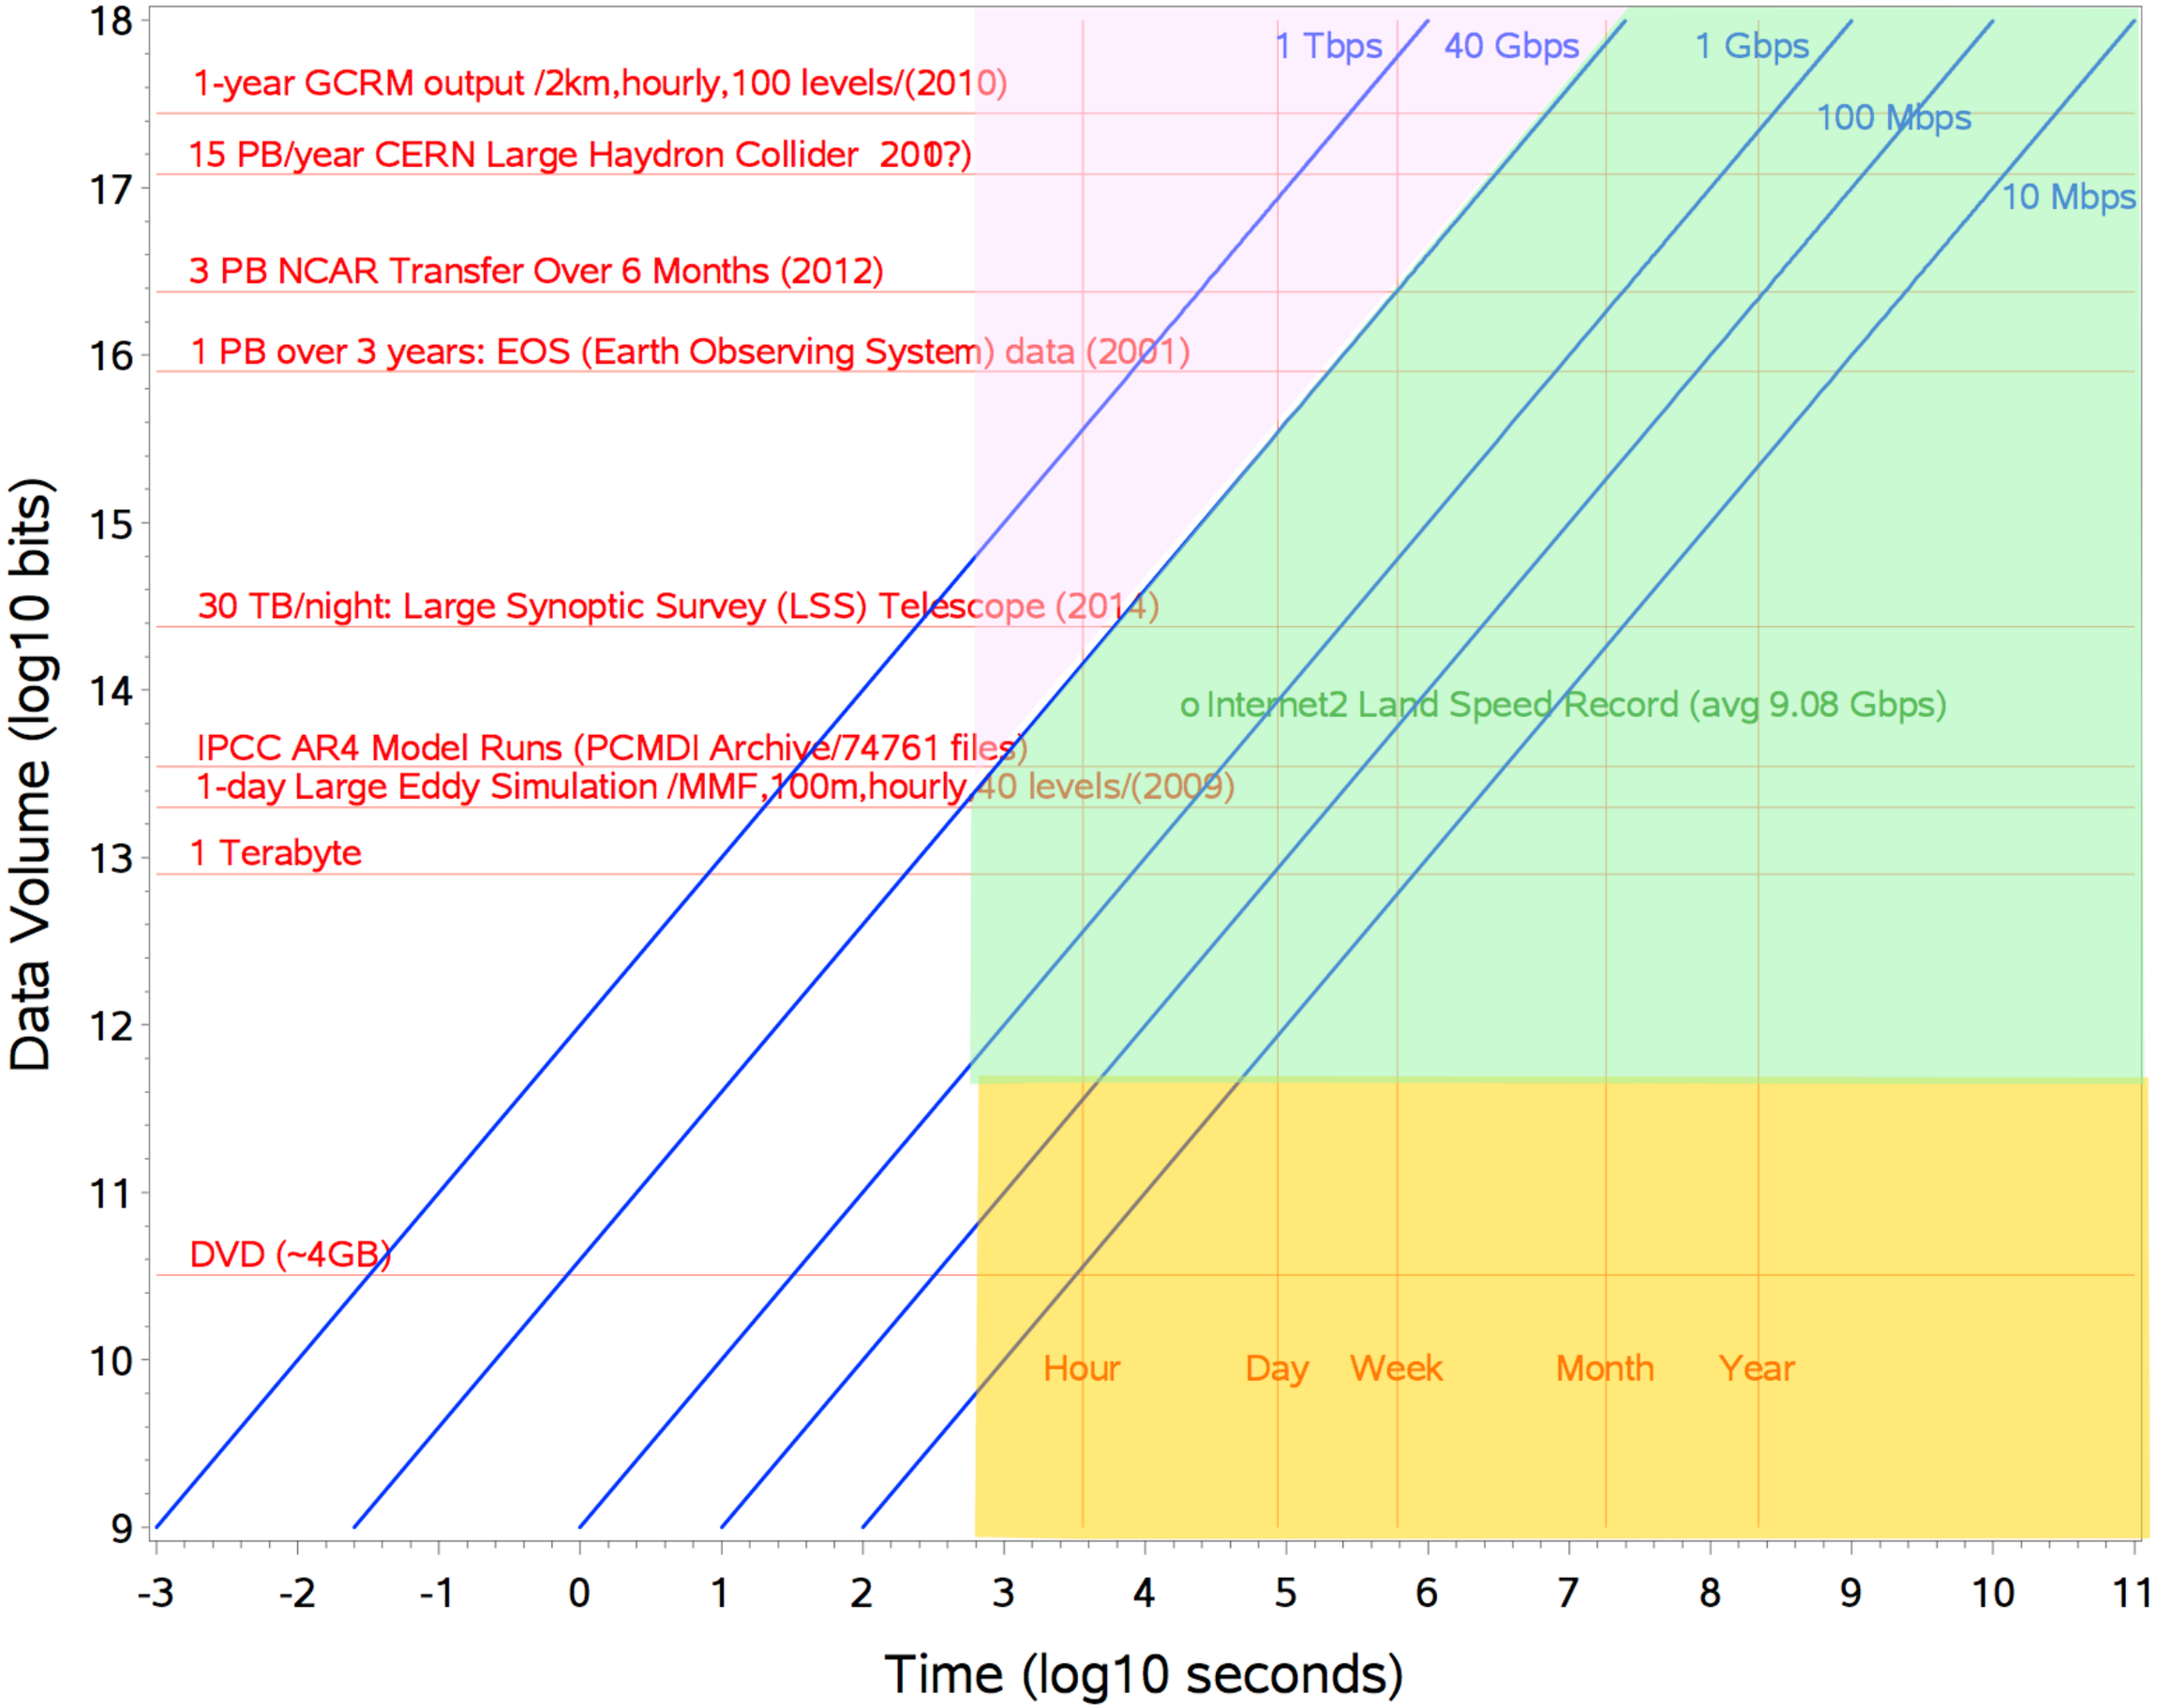
1 Tbps
40 Gbps
1 Gbps
100 Mbps
10 Mbps
1-year GCRM output /2km,hourly,100 levels/(2010)
15 PB/year CERN Large Haydron Collider 200?)
3 PB NCAR Transfer Over 6 Months (2012)
1 PB over 3 years: EOS (Earth Observing System) data (2001)
30 TB/night: Large Synoptic Survey (LSS) Telescope (2014)
o Internet2 Land Speed Record (avg 9.08 Gbps)
IPCC AR4 Model Runs (PCMDI Archive/74761 files)
1-day Large Eddy Simulation /MMF,100m,hourly,40 levels/(2009)
1 Terabyte
DVD (~4GB)
Hour
Day
Week
Month
Year
Data Volume (log10 bits)
Time (log10 seconds)
18
17
16
15
14
13
12
11
10
9
-3
-2
-1
0
1
2
3
4
5
6
7
8
9
10
11

CMMAP Data Services News

Karen Schuchardt
UC Berkeley
January 2011

Pacific Northwest
NATIONAL LABORATORY

PNetCDF

- defines cdf5 format based on cdf3 with fixes for large variables
- New psuedo non-blocking IO that improves performance
- Support for setting the header padding size (important for basic metadata editing)
- Integers are 4 bytes; no longs
- In theory, can be specified as the mechanism to use via the NetCDF4 interface
- Up to 10GB/s franklin (of possible 16GB)

- Pretty robust and good performance (version 1.2)
- Data can only be processed with pnetcdf-based analysis tools
  - Future support model not clear

Pacific Northwest
NATIONAL LABORATORY

NetCDF4

- Based on widely used HDF5 data model
  - HDF5 format is a rich data format with filesystem-like constructs
  - Lots of tunable features like compression, chunking
  - Can be restricted to netcdf data model constructs

- Fortran interface (still) has 32bit restriction
- Not stress tested
- Performance currently lags PNetCDF

- New DOE Exascale project to optimize HDF5
  - LBNL (Prabhat – PI), HDF5 Group, PNNL
  - Being done in context of real applications
    - GCRM
    - Pore scale simulators (groundwater, physics)

Pacific Northwest
NATIONAL LABORATORY

Runs

- Z anelastic model (run to date)
  - 15km, 26 interfaces,18 days, 3 hourly
  - Jablonowski test case, 10000 processors on franklin
  - About 8 variables ~= 44GB
- Z anelastic model (by spring/summer)
  - 4 or 7 km (or both), 36? interfaces,18 days, 3? hourly
  - Initial conditions TBD, ~10000/40000 processors on hopp2
  - Will have physics added (Don Dazlich)
    - 8? variables ~= **TBD**

NOTE: time free only until April

Pacific Northwest
NATIONAL LABORATORY

Analysis - ParCal

- New DOE Project for Parallel Data Analysis Tools
- ANL (Rob Jacob PI), NCAR, PNNL, Sandia
- Major outcomes
  - Version of NCL which transparently turns current NCL data arrays and operations into parallel operations
  - Climate specific compression to be tested in PNetCDF

Pacific Northwest
NATIONAL LABORATORY

Cyber-infrastructure Working Group: Pagoda

Jeff Daily and Karen Schuchardt, PNNL

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Parallel Analysis of Geodesic Data

- a.k.a Parallel Analysis of Geoscience Data
- C++ API for developing custom analysis
  - Most similar to Java NetCDF API
- Data-parallel command-line tools
  - Mimics the NetCDF Operators (NCO)

NCO	pagoda
ncks	pgsub
ncra	pgra
ncea	pgea
ncbo	pgbo
ncflint	<soon>
ncwa	<soon>
ncrcat	
ncrename	
ncatted	
ncpdq	

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Pagoda Design

- Focus on parallel IO and large variables
  - Do what NCO does but when your data is too large for your workstation
- Handles regular and geodesic grids
  - Geodesic grids are described using an explicit topology
  - Explicit topology needed for analysis/visualization e.g. VisIt
- Reads and writes classic NetCDF via Parallel NetCDF
- Reads and writes NetCDF4
- Runs on workstations, clusters, HPC systems e.g. hopper

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

pgsub Strong Scaling

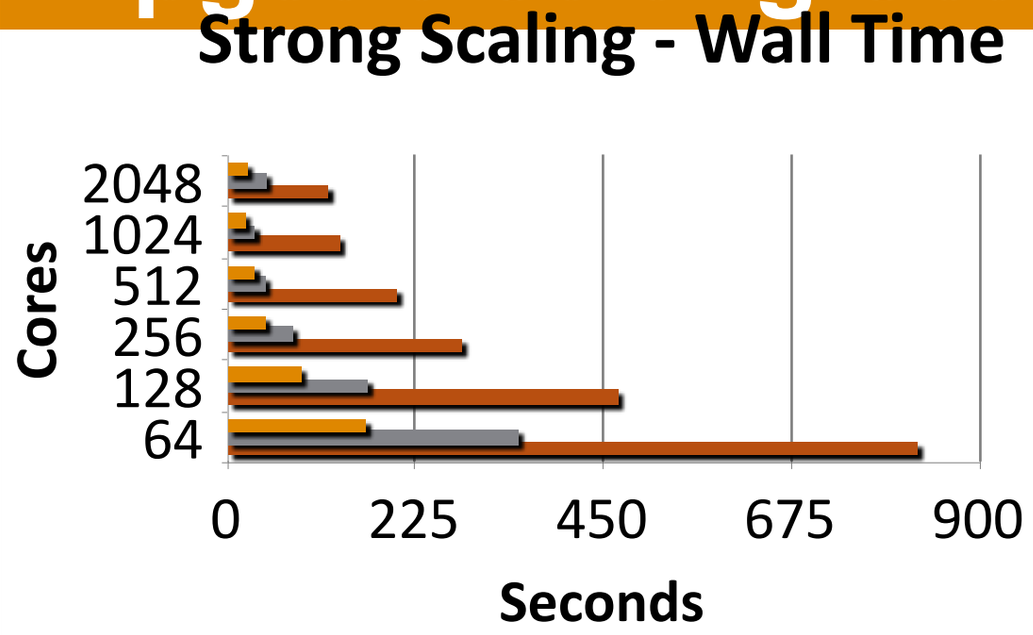

- Shown to scale up to 2K cores
- Shows that IO is a major bottleneck
- Write bandwidth nearly 5GB/s on franklin
- Our first optimization shows importance of efficient use of IO

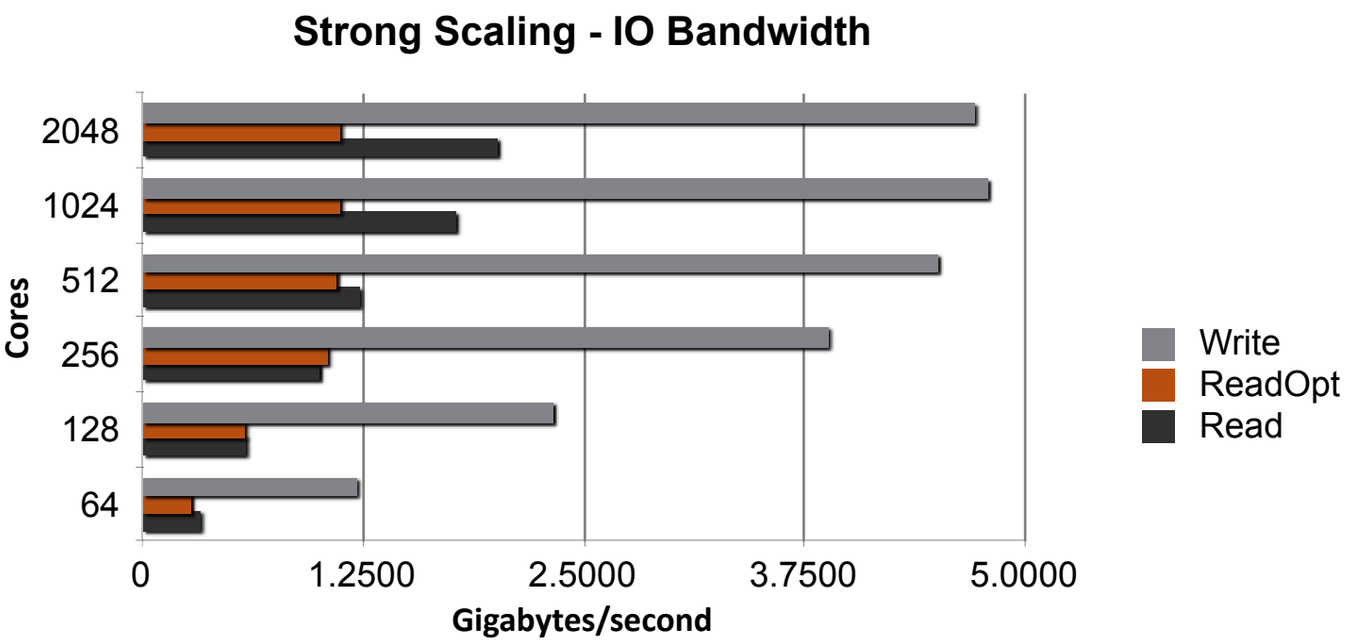

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Future Directions

- "make it easy" – A higher level API
- New language bindings? Python? Fortran?
- Handle additional conventions e.g. missing_value
- Finish pgflint (ncflint), pgwa (ncwa)
- Grid interpolation
- Other operators?
  - What if header isn't big enough and data is too large?
  - What if pnetcdf's "CDF5" format is used?
- We need more users and user input on what's needed
  - Already in use/testing by CSU, ANL, NCAR

Proudly Operated by Battelle *Since 1965*

Thanks

- http://svn.pnl.gov/gcrm/wiki/pagoda
- pagoda-dev@googlegroups.com

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Model Run Management App

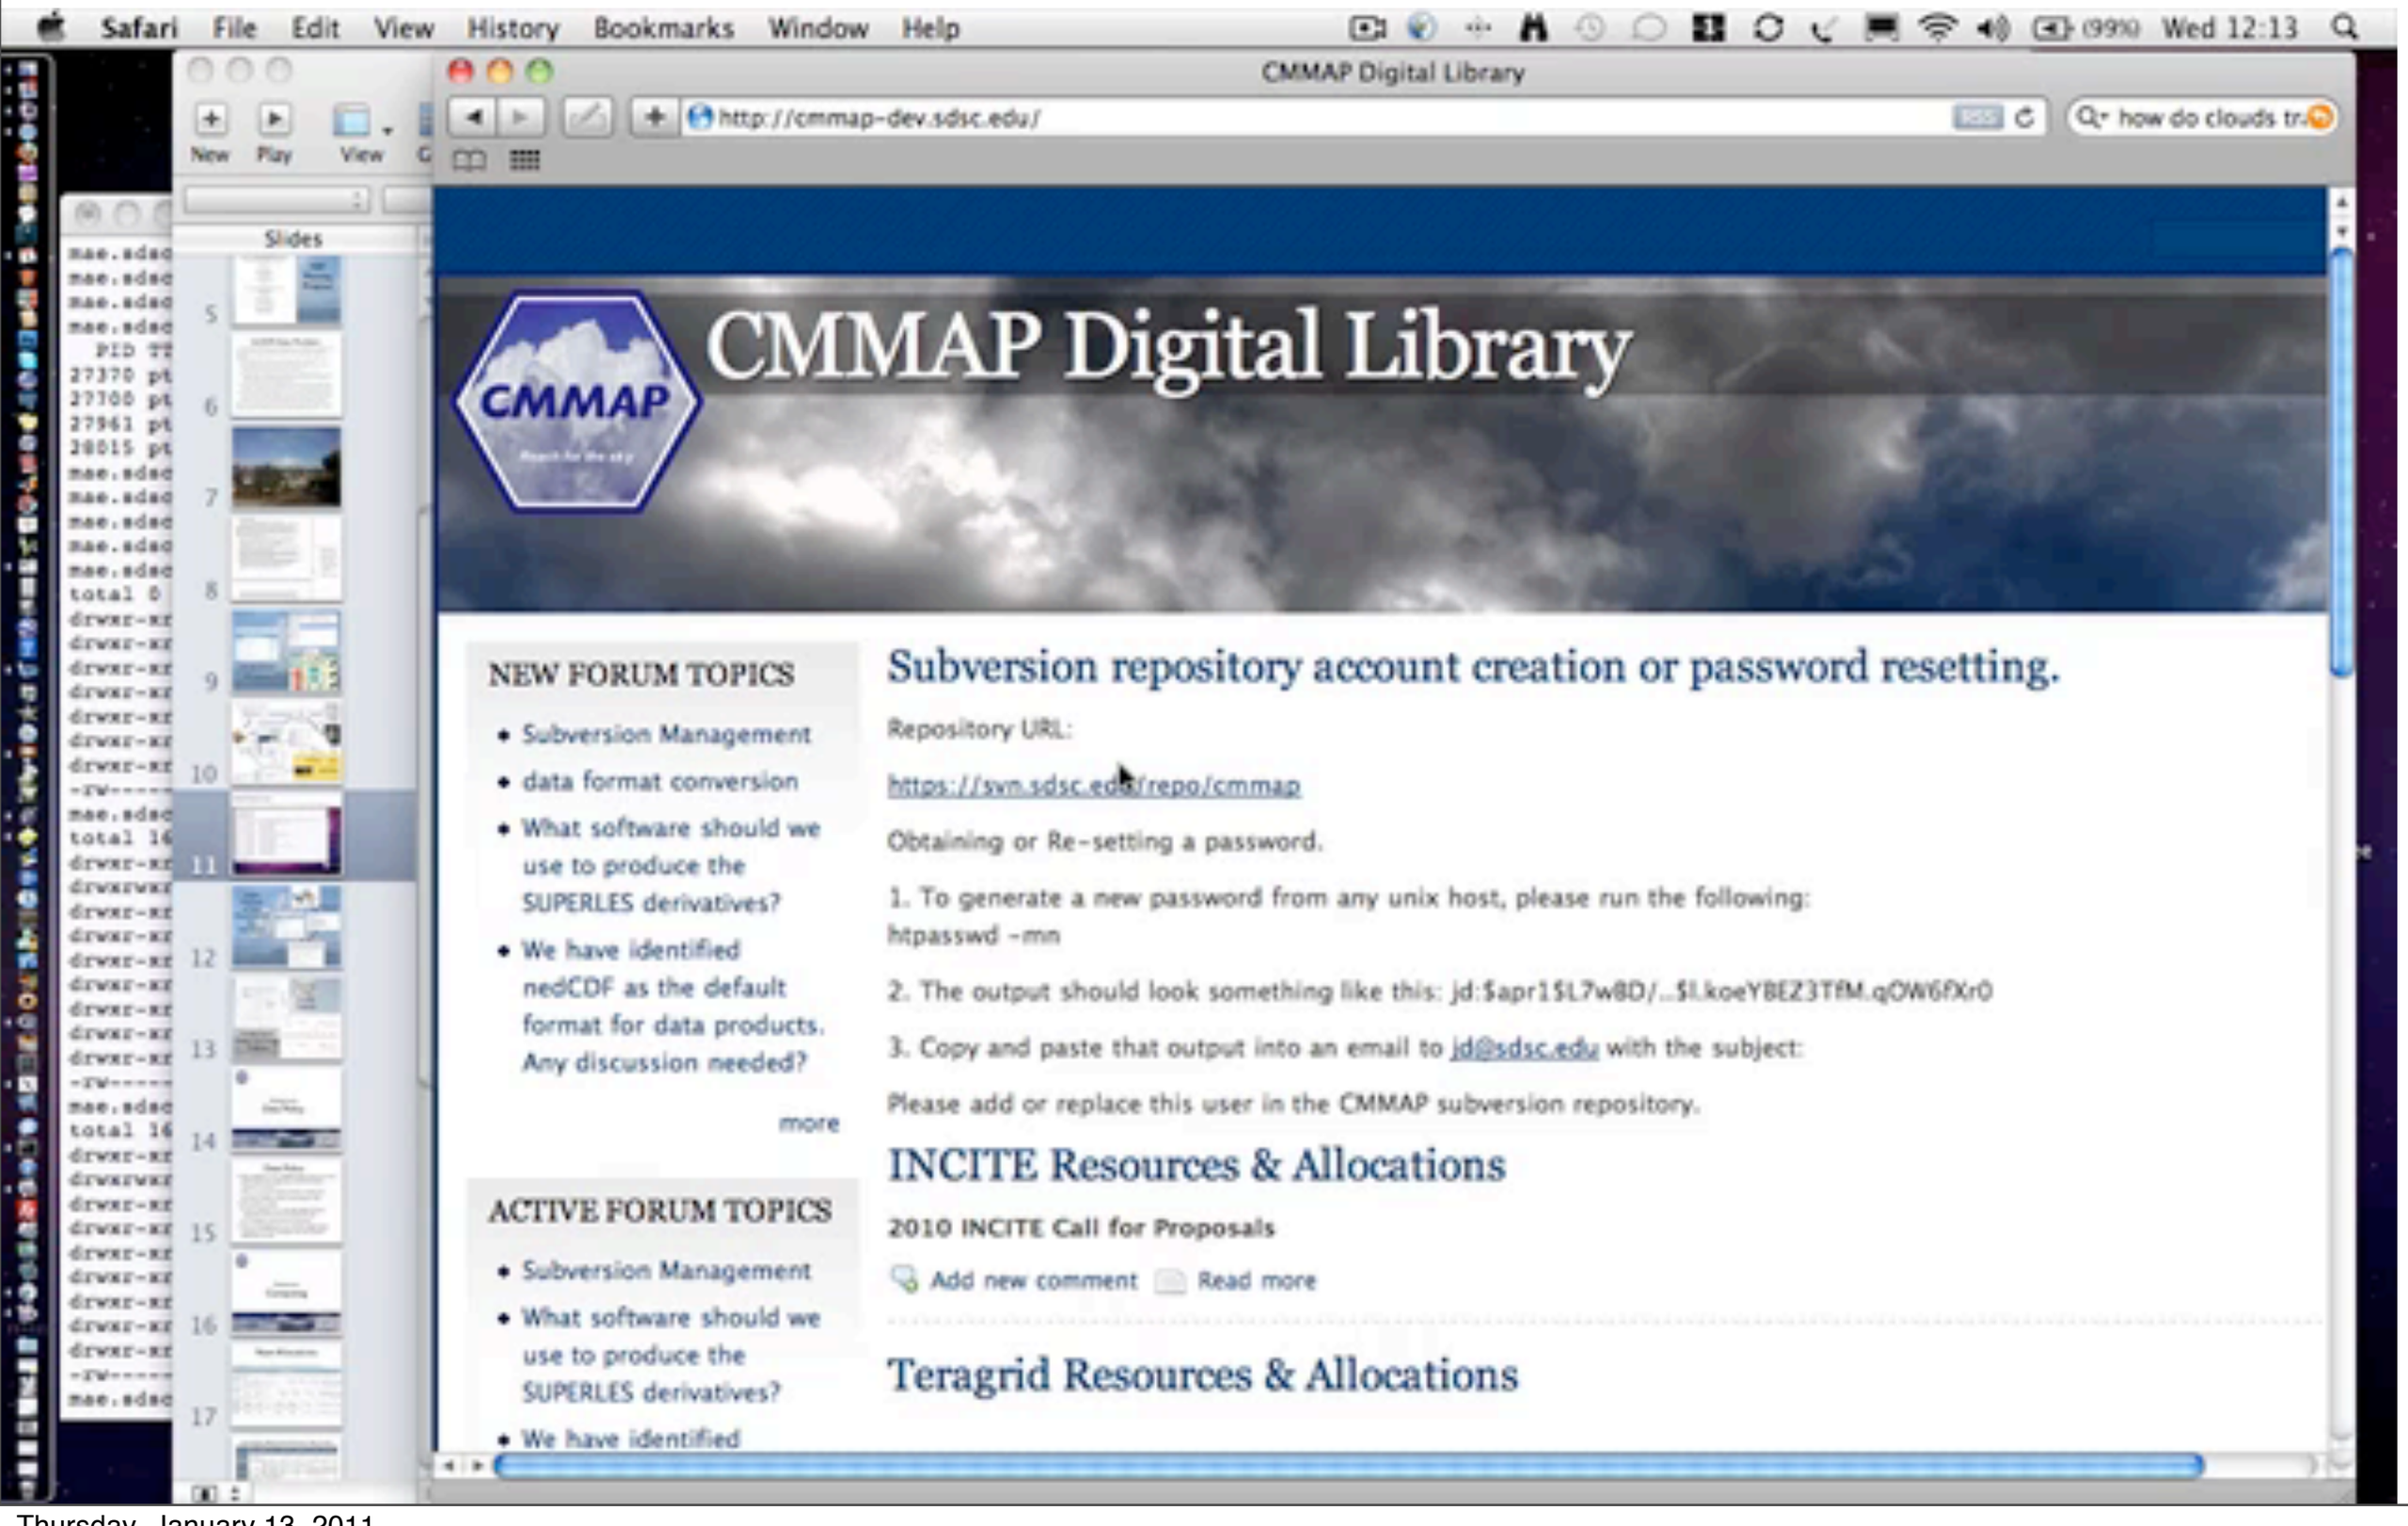

iRODS Command-line Client: Scriptable

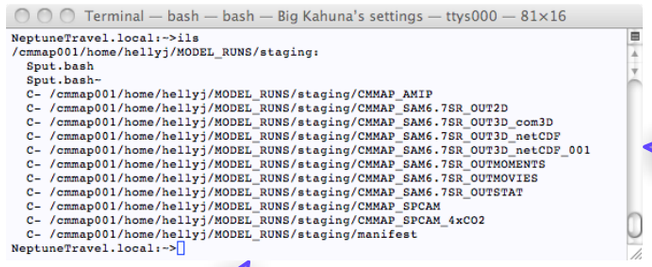

iRODS Federation Mechanism

CMMAP Phase II CI @CSU

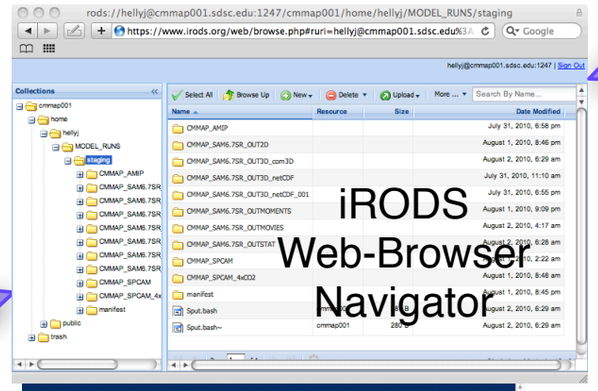

Drupal: programmability

- iquery
- MMF Workbench

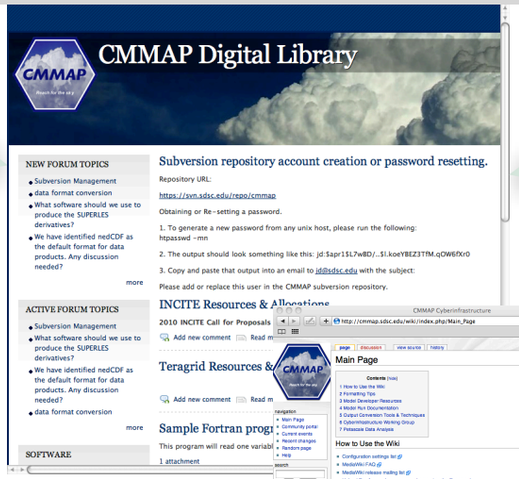

MediaWik: Documentationi

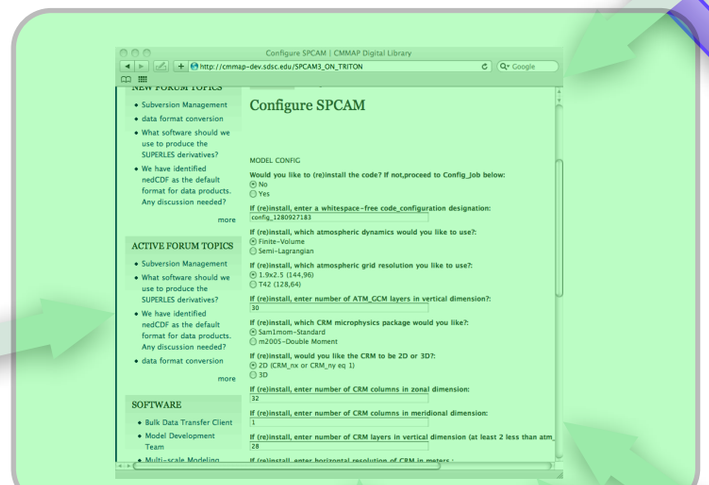

CMMAP Disk Allocation @SDSC

FRE
Workflow Control
Collaboration w/GFDL

others

Community Account

Steele@Teragrid

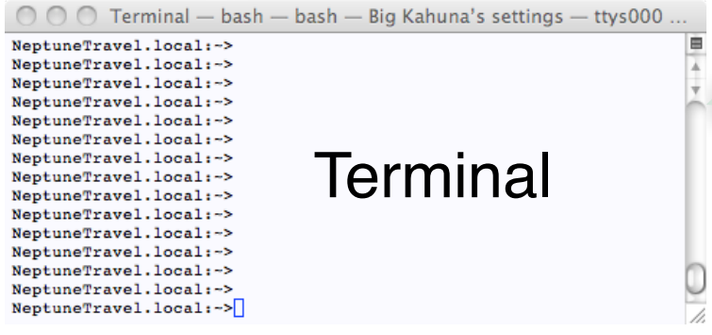

Computing

Data

Model Run Management App

iRODS Command-line Client: Scriptable

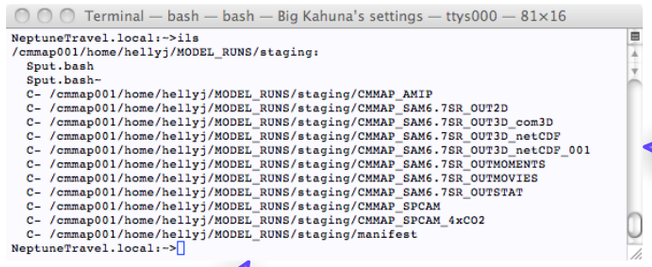

iRODS Federation Mechanism

CMMAP Phase II CI @CSU

Drupal: programmability

- iquery
- MMF Workbench

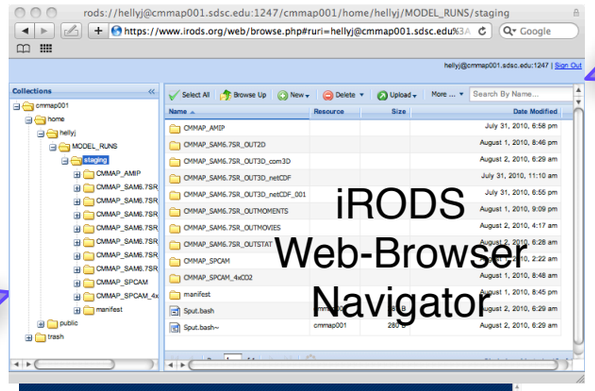

iRODS Web-Browser Navigator

MediaWik: Documentationi

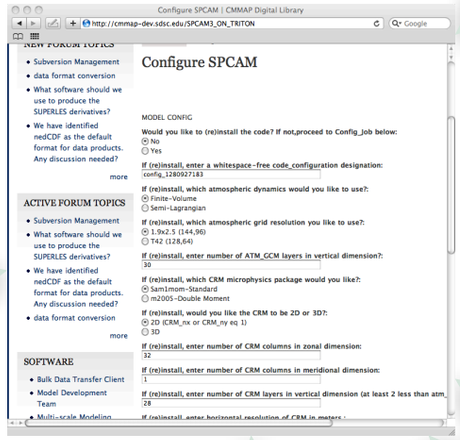

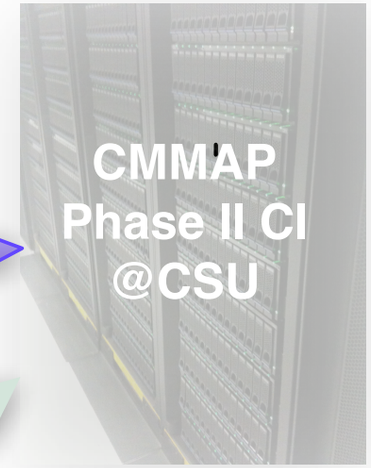

FRE Workflow Control Collaboration w/GFDL

others

Community Account

Steele@Teragrid

Terminal

Computing

Data

Introduction to FRE: The Flexible Modeling System Runtime Environment

Developed at GFDL, 2002-2010, by:

Amy Langenhorst *Amy.Langenhorst@noaa.gov*

Aleksey Yakovlev *Aleksey.Yakovlev@noaa.gov*

V. Balaji *V.Balaji@noaa.gov*

Introduction to FRE

- The FMS Runtime Environment (FRE) is a toolset for managing experiments from start to finish
  - acquire source code, compile (*fremake*)
  - launch jobs to run models (*frerun*)
  - postprocessing the output (*frepp*)

The Mechanics of FRE

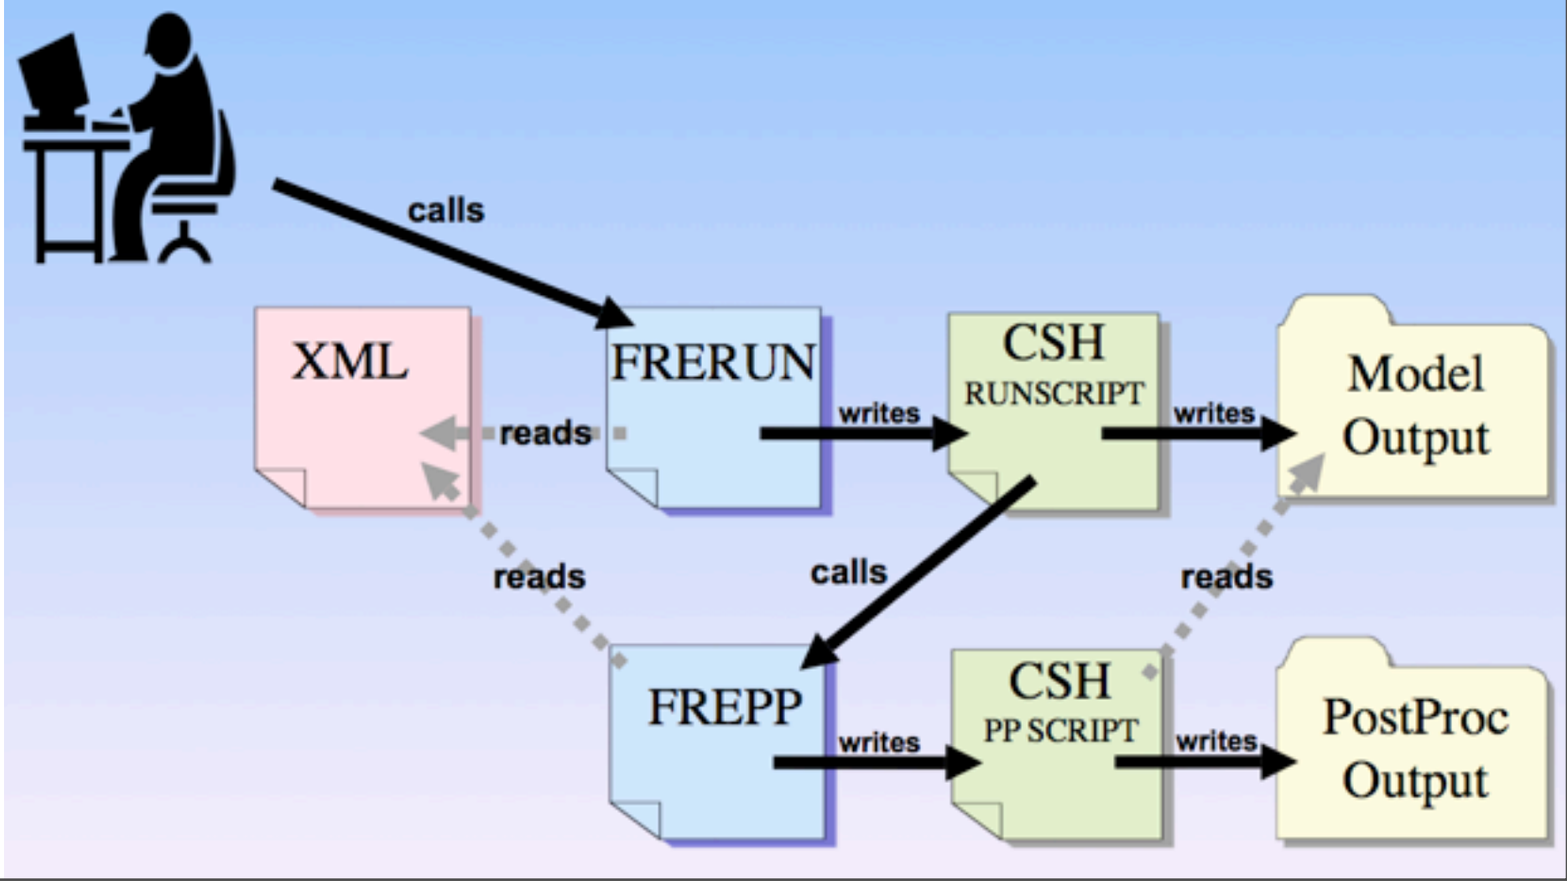

Hugh Morrison

Proposed tests of microphysics in MMF using Teragrid

- Recent request for large Teragrid allocation was reviewed and accepted (allocation through July 2011)

- Broadly, goal is to address two separate but related issues: 1) sensitivity to key parameters in scheme, 2) tuning

• anticipate ~ 50 sensitivity runs (1 week spinup + 4 additional weeks) for initial tests, ~ 10 tuning runs (2 year runs + 1 month spin up) based on configurations identified from sensitivity tests, and a ~25-year AMIP- type run – total of 633 months simulation time, coupled run?

• anticipate tests using 1.9 x 2.5 degree fv core, 64 CRM columns per large grid

- Based on timing tests from Mike Pritchard on Purdue Steele cluster, 3.5 million SU's were requested

• ~ 2.5 x more expensive using 2-moment microphysics compared to standard SpCAM, estimated cost is 5600 SU per month simulation.

- Tests to focus on:

1) Reducing costs of microphysics
- larger microphysics timestep
- reduced # of prognostic (advected) variables

2) Sensitivity to new microphysics developments and parameter settings
- Morrison and Grabowski (2008) ice microphysics (has 6 prognostic variables versus 9 in M2005 → improved cost)
- Parameter settings identified as being important in CRM tests (e.g., graupel density and fallspeed)

3) Tuning

Next steps:

Tests in stand-alone SAM (for significant code changes, such as reduction in # of prognostic variables)

Working with GUI interface to compile/run SpCAM
- approaches for modifying/testing code (e.g., svn)

Analysis of output
- do we want to develop a "standard" diagnostics package, a la the CAM diagnostics?

CMMAP Data Transposition Code

Overview of Code

- LES (SP-CAM) Data files in NETCDF format.
- OpenMP code setup to read one file per core simultaneously.
- Each file corresponds to a spatial partition at a given time. The code reads in all the files into memory to do the transposition.
- Tested on
  - Triton regular [8 core, 24GB] and large memory nodes [32 cores, 512GB] w/ data oasis [lustre].
  - Dash regular [8 core, 48GB] and vSMP node [128 cores, 650GB] w/ GPFS-WAN.

Initial Recovery of Time-series

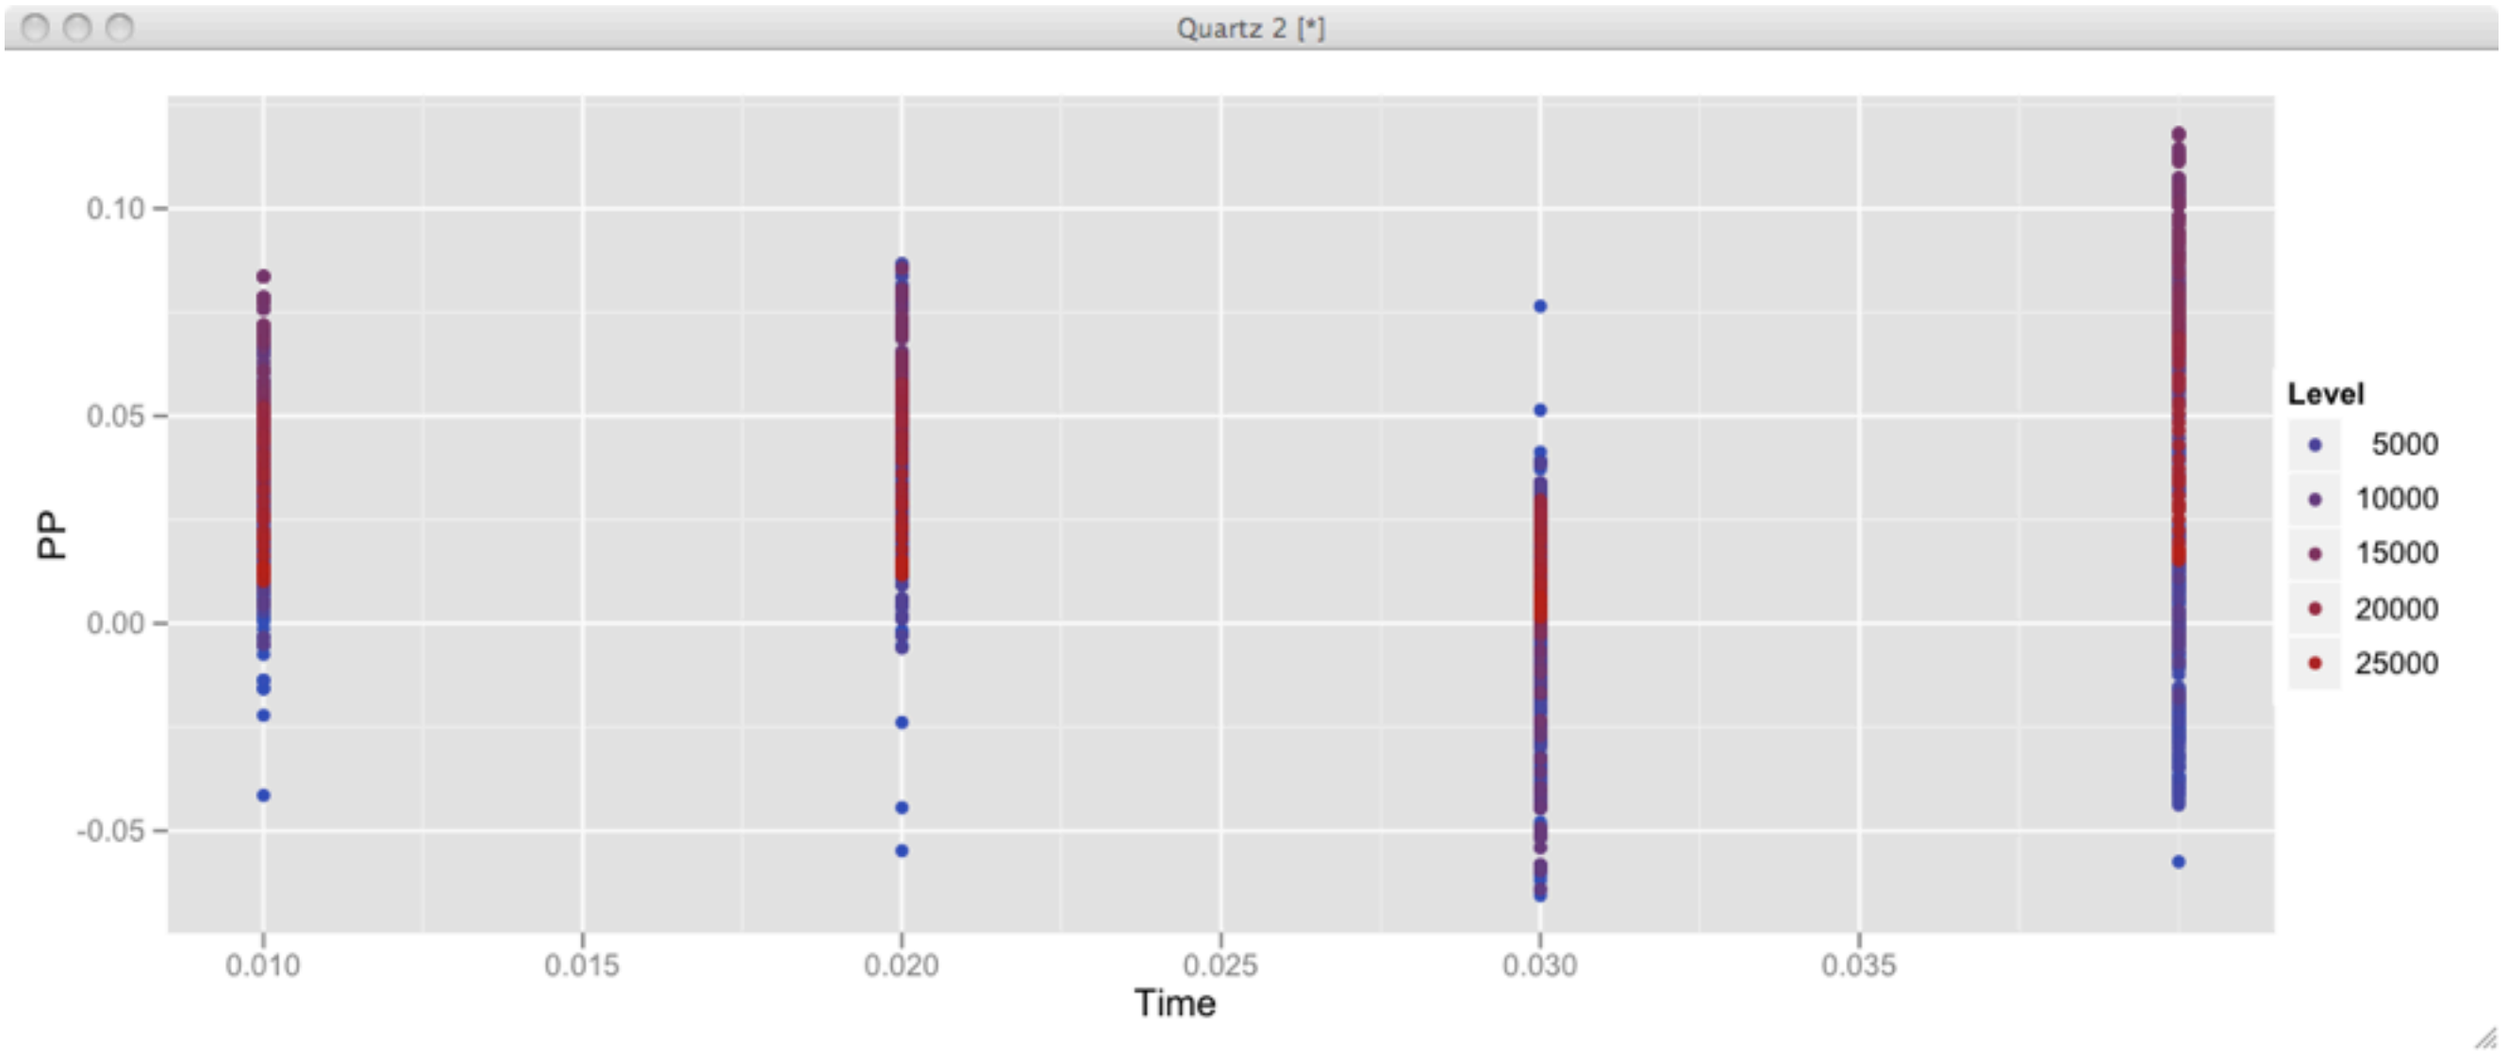

Triton Results

- Code tested with 76 files. Total size of data read : 128GB.
- Run times are dominated by I/O performance.

No. of Cores	8-core node [2 Quad Nehalems]	32-core node [8 Quad Shanghais]
1	347s	505s
2	197s	222s
4	122s	118s
8	105s	154s

Dash Results

- Tested on regular compute nodes w/ GPFS-WAN. vSMP node testing in progress.

No. of Cores	8-core compute node (w/GPFS-WAN)
1	838s
2	496s
4	301s
8	245s

Summary and Future Work

- OpenMP code tested and results verified on Triton and Dash.
- Current performance limited by I/O performance of filesystem on given node.
- Achieved ~1.2GB/s w/ lustre on Triton node. The maximum achievable is 1.25GB/s [Myrinet card peak]. GPFS-WAN performance on Dash node is lower due to network setup. Lustre testing on Dash is in progress.
- Peak performance achieved using 4 cores on 32-way node => it might be useful to limit number of threads reading. Can still use more threads for analysis part.
- Developing hybrid (MPI + OpenMP) code to make use of more nodes and get better I/O performance [Lustre on Triton can do over 7GB/s on reads].

Discussion

- Upcoming allocation proposals (new LES run?)
- Additional model runs for digital library
- High-volume data visualization
- Time-series data recovery from model runs for validation support

Backup

CI-P-1_asSubmitted.pages (Word Processing)

Proposal to the National Science Foundation for

Community Infrastructure Planning:

Cyber-Infrastructure for the Cloud-Climate Community

Prepared in response to
CISE Computing Research Infrastructure (CRI) Program Solicitation 08-570

Principal Investigator:

David A. Randall

Department of Atmospheric Science
Colorado State University
Fort Collins, Colorado 80523

tel: (970) 491-8474
email: randall@atmos.colostate.edu

Co-Investigators:

John Helly

San Diego Supercomputer Center and Scripps Institution of Oceanography
University of California, San Diego
10100 Hopkins Drive
La Jolla, CA 92093-0505

Kelley Wittmeyer

Department of Atmospheric Science
Colorado State University
Fort Collins, Colorado 80523

Michelle Strout and Sangmi Pallickara

Department of Computer Science
Colorado State University
Fort Collins, Colorado 80523

92% 8,942 Words Page 1 of 26

NSF Planning Proposal

Thursday, January 13, 2011

4. Planning process

Our goal is to submit the CI proposal in the summer of 2011.

Planning will begin at CMMAP's Team Meeting in August 2010, the same week that this planning proposal will be submitted to NSF. We will also take advantage of CMMAP's Team Meeting in January 2011, in Berkeley, California, where we systematically collect ideas from the CMMAP team, and also engage the local talent at NERSC, including the developers of ViSIT. The planning activity will be a major component of the January 2011 meeting.

The tasks to be completed before the submission of the CI proposal include:

- Scoping the hardware and software systems, including the determination of expected life-cycle cost to operate over its useful lifetime.
- Developing a plan for maintaining and enhancing the proposed infrastructure up to and beyond the sunset of the planned CI grant.
- Outreach to solicit input from the CMMAP community and the national HPCC centers on the design of the hardware and software components and external and internal interfaces of the proposed infrastructure.
- Soliciting input from the CMMAP community on the design of the education and outreach activities that will be associated with the proposed infrastructure.
- Developing a site plan for the hardware, in cooperation with the CSU administration, including the university's Facilities office.

- Developing an operations plan, which must take into account the education and outreach activities.Initiation of a planning process to determine hardware and software configuration and a likely data loading model to size the system against

5. Management Plan

- Many difficult choices must be made; for example, choosing which fields to output, and what subsetted spatial and temporal resolutions to save, are complex.
- CMMAP standard?

The Program Planning Prism

- Management
  - CIWG (Cyberinfrastructure Working Group)
  - Data Policy
  - Software Policy
- Resource Planning
  - Computing
  - Data
  - People
- Community Support
  - Data interoperability
  - Model code portability

R: Analysis and Plotting

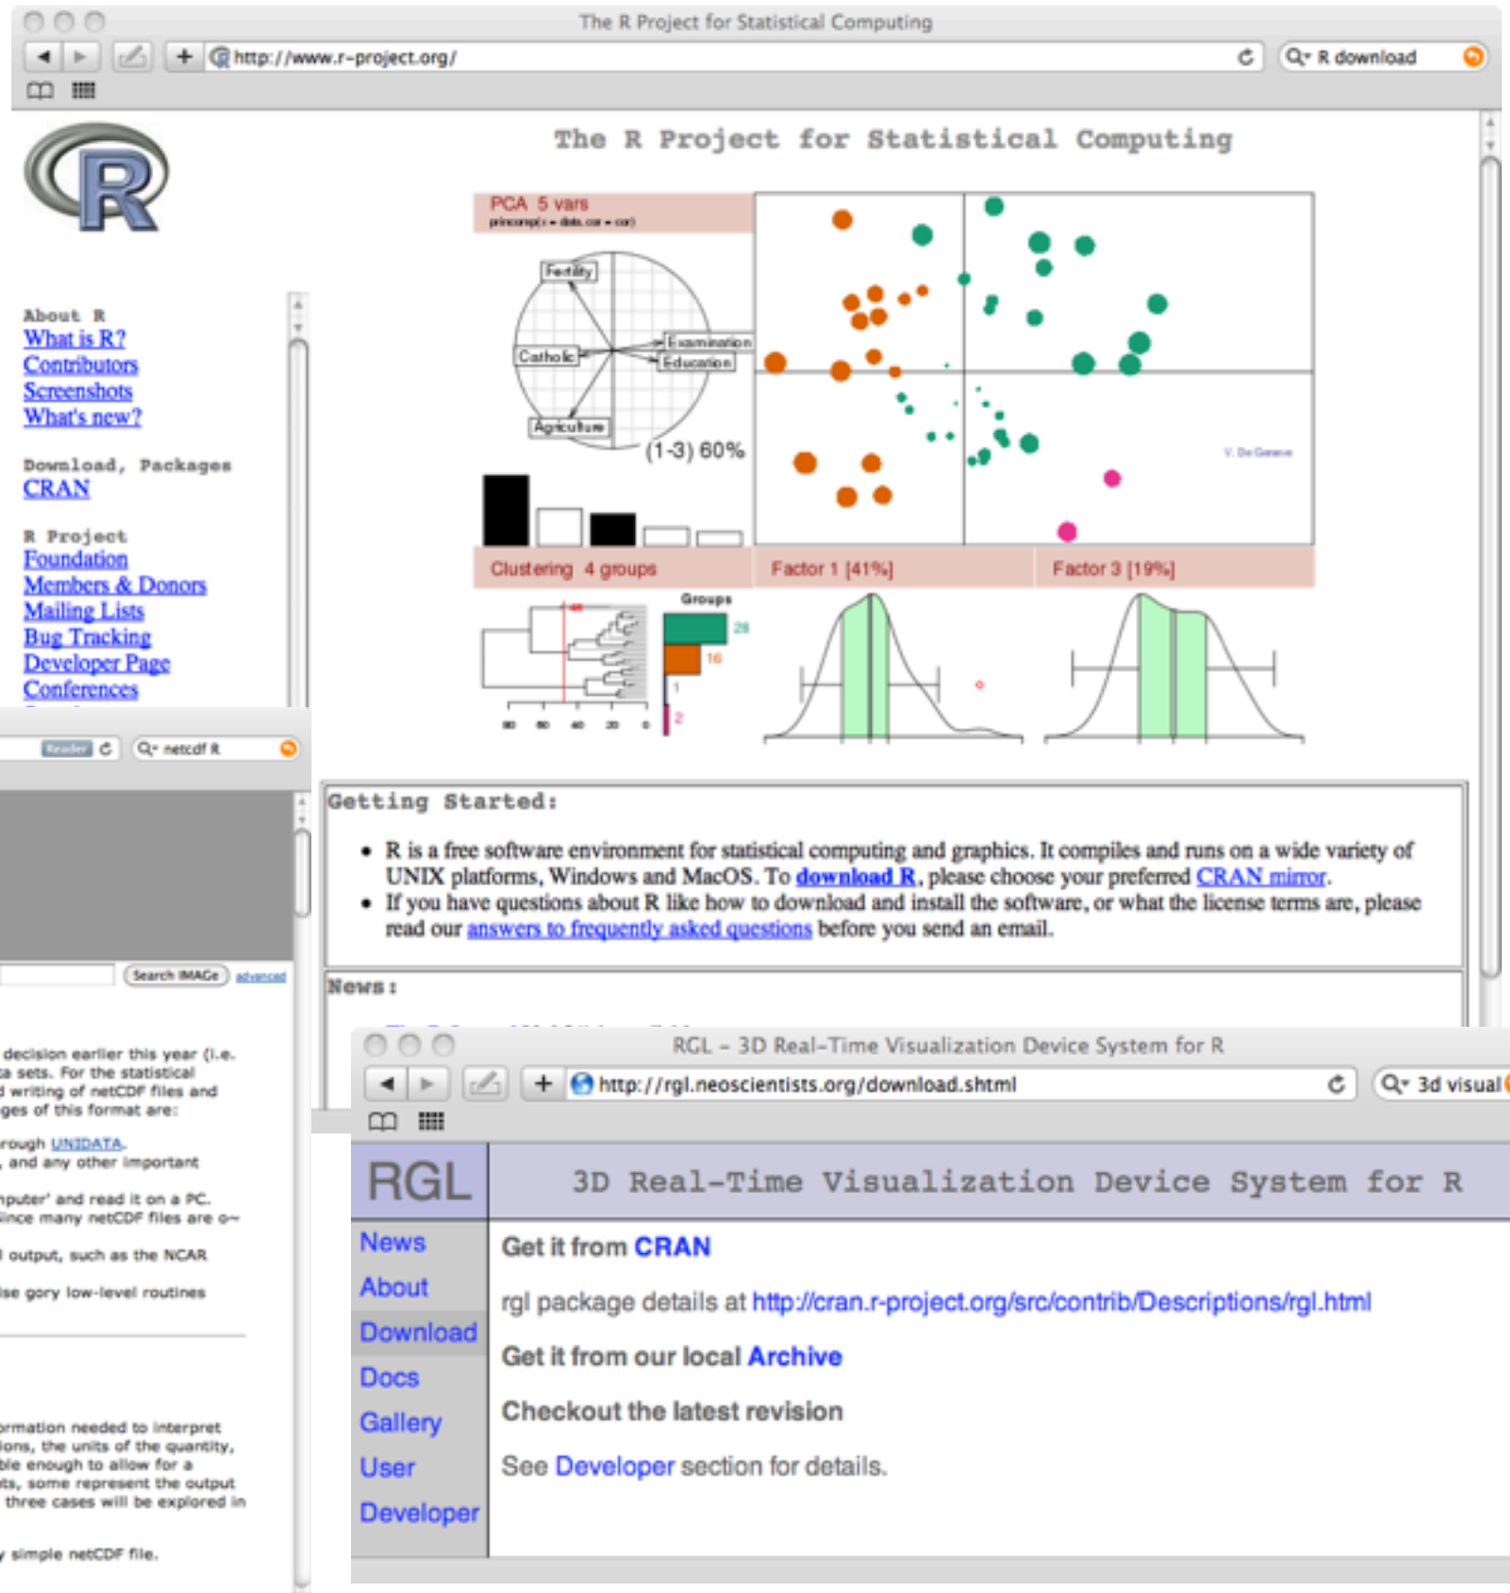

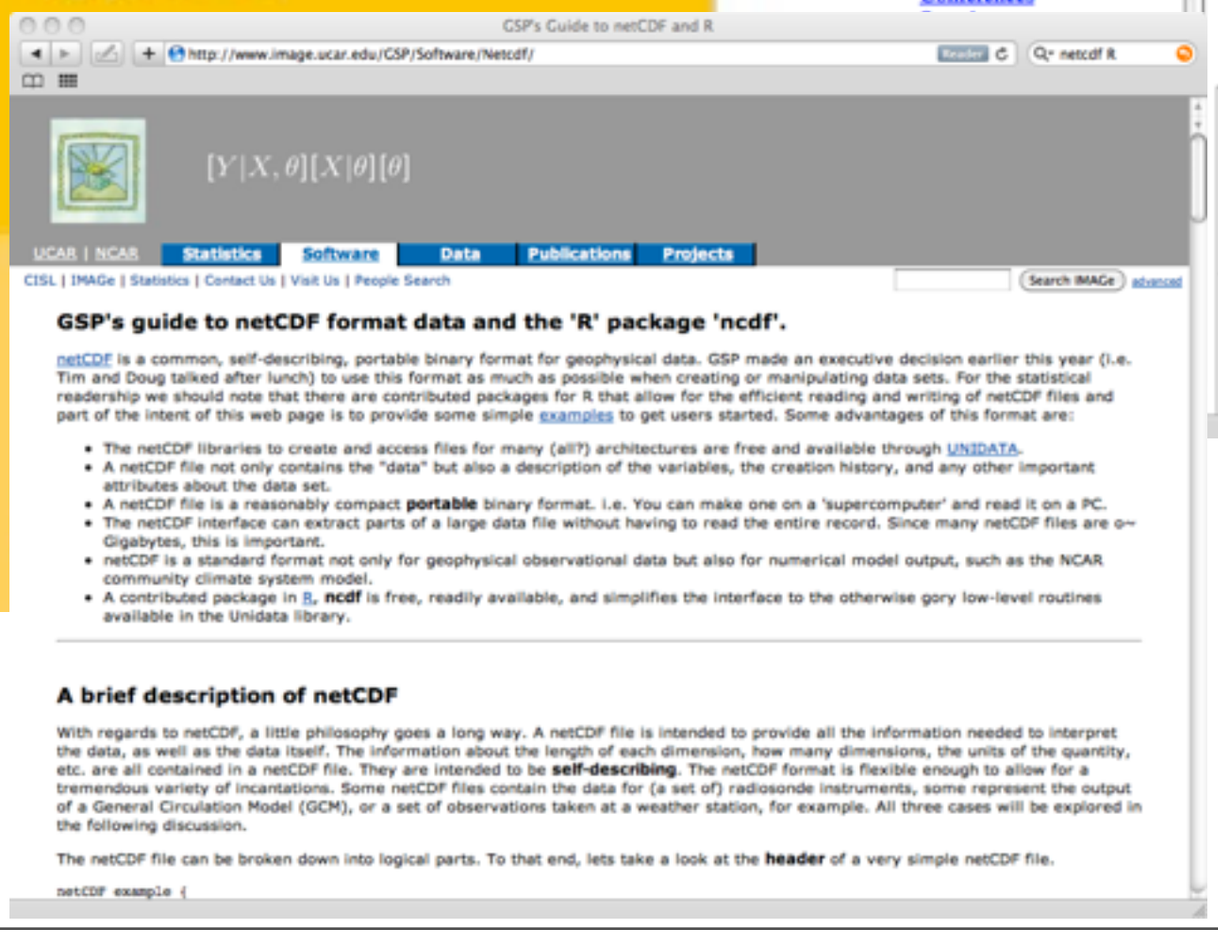

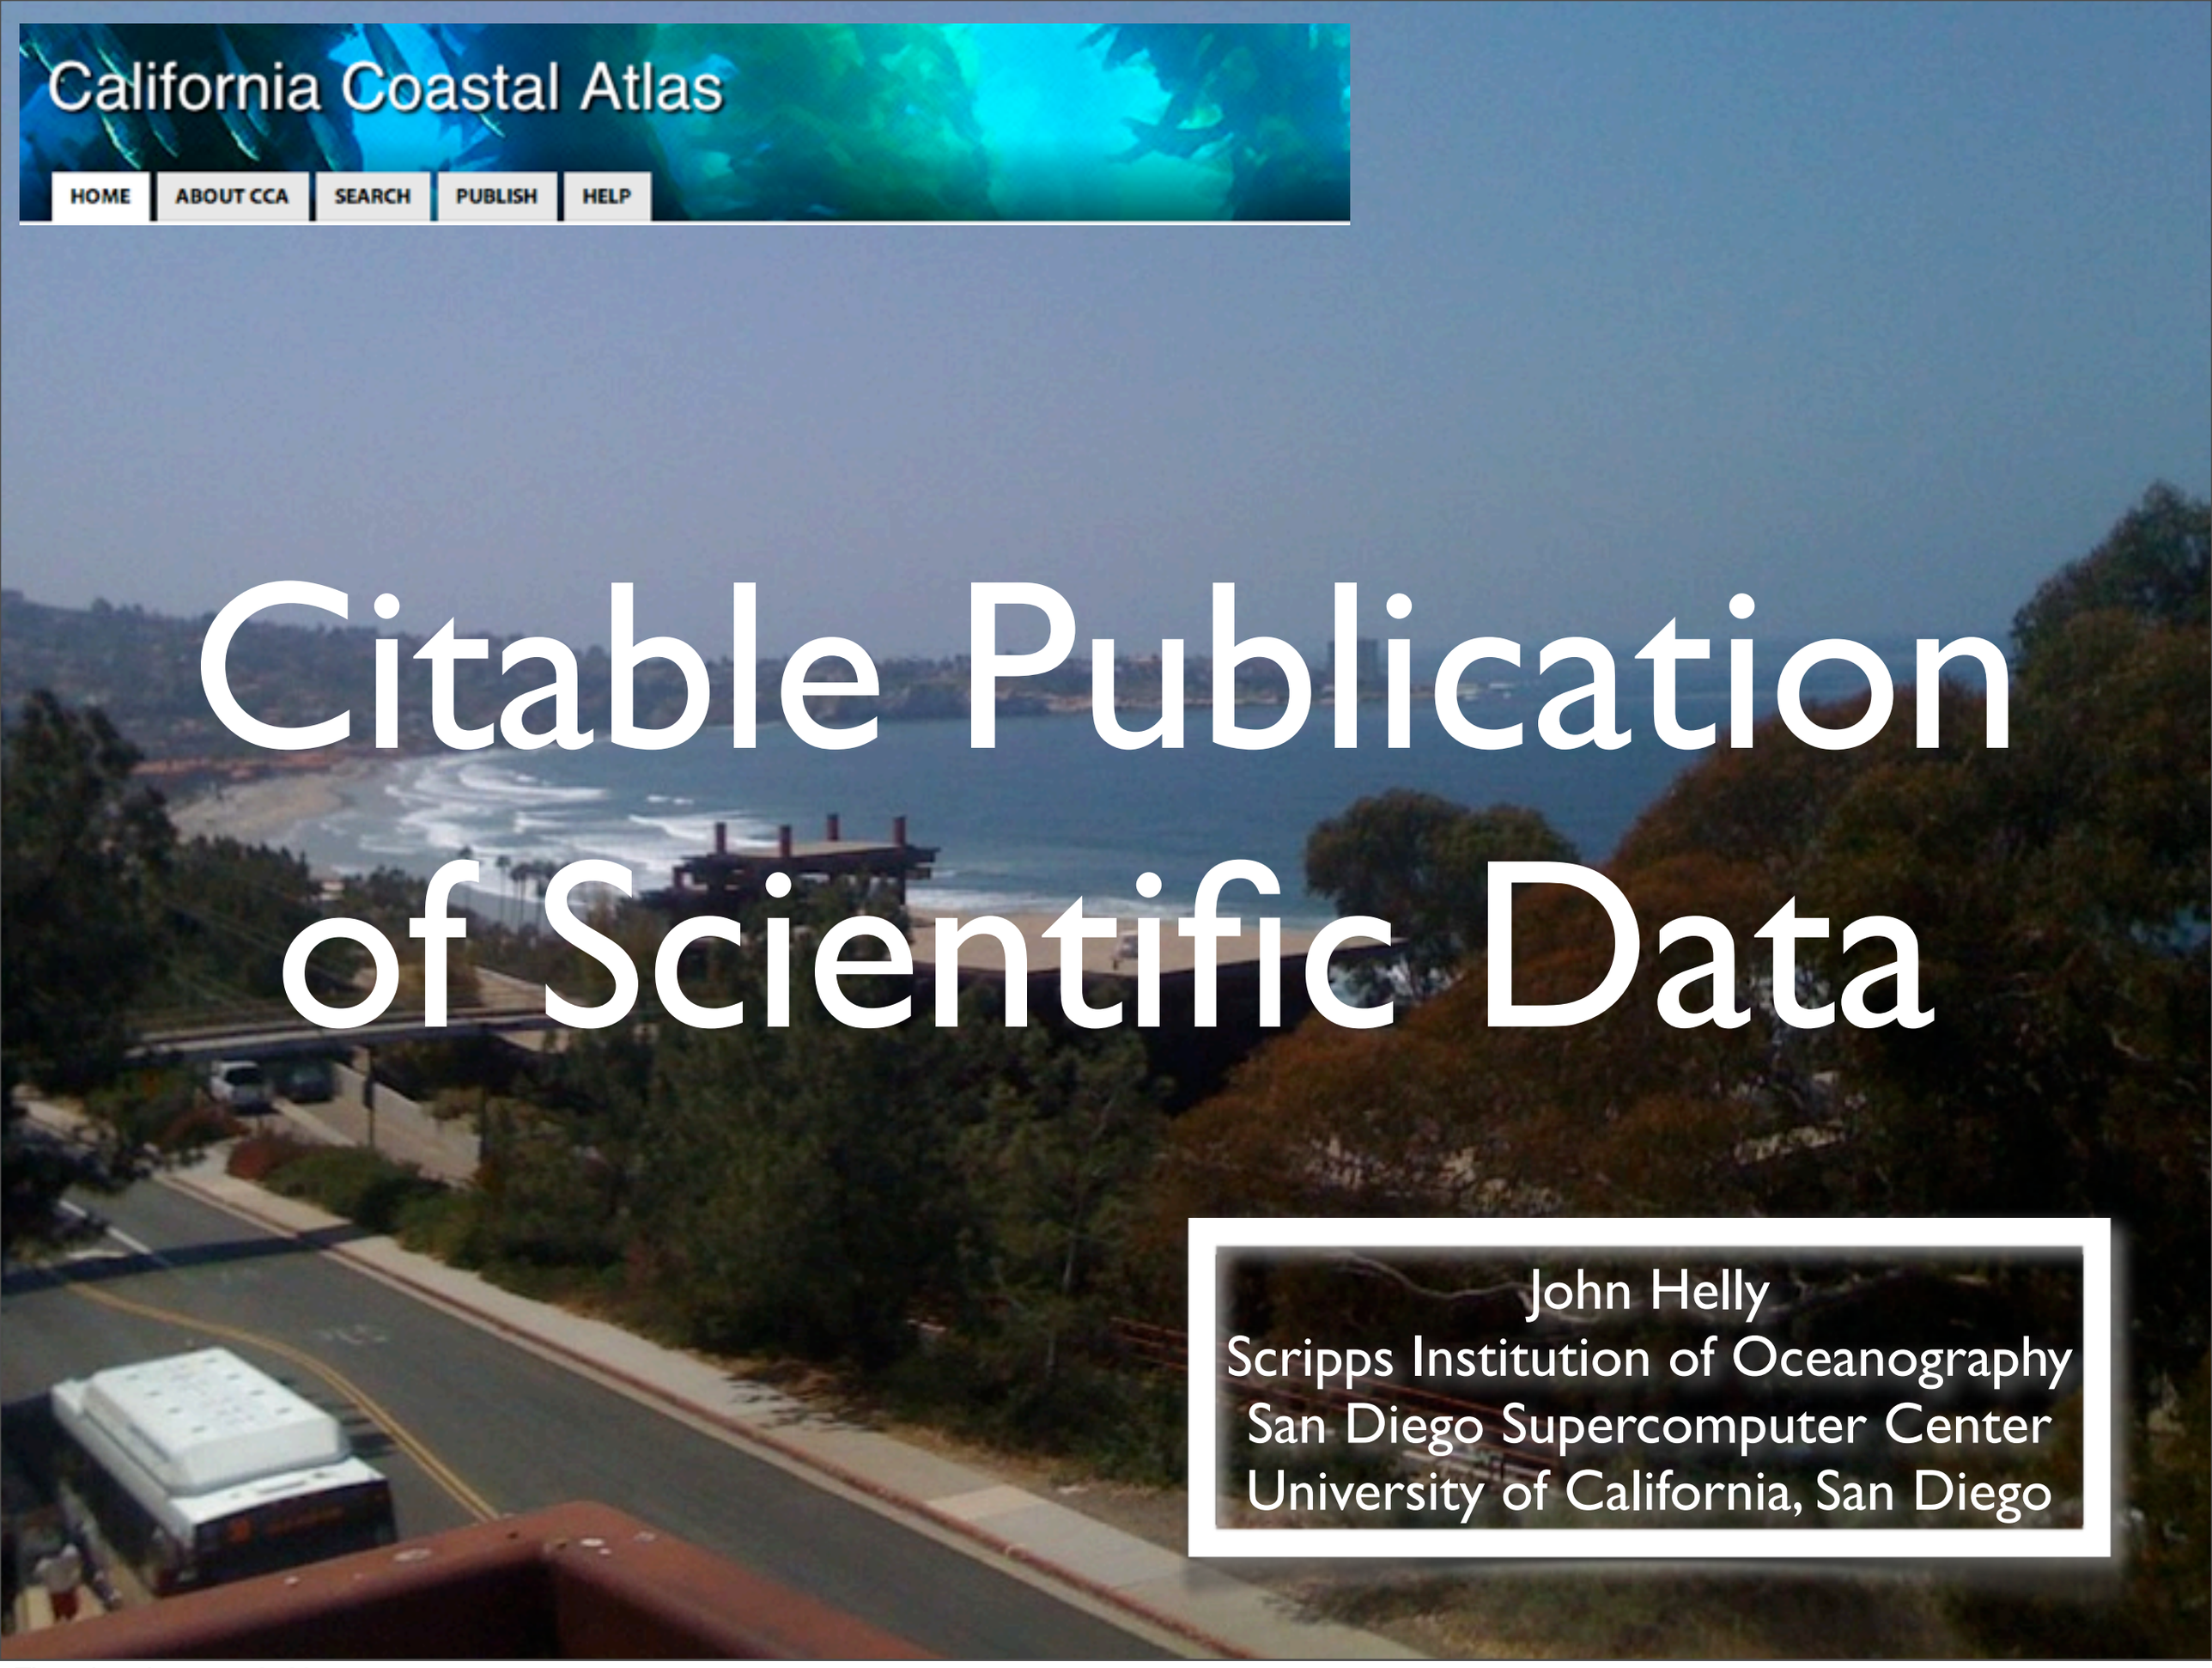

Citable Publication of Scientific Data

John Helly
Scripps Institution of Oceanography
San Diego Supercomputer Center
University of California, San Diego


Optiputer@SDSC

RSS

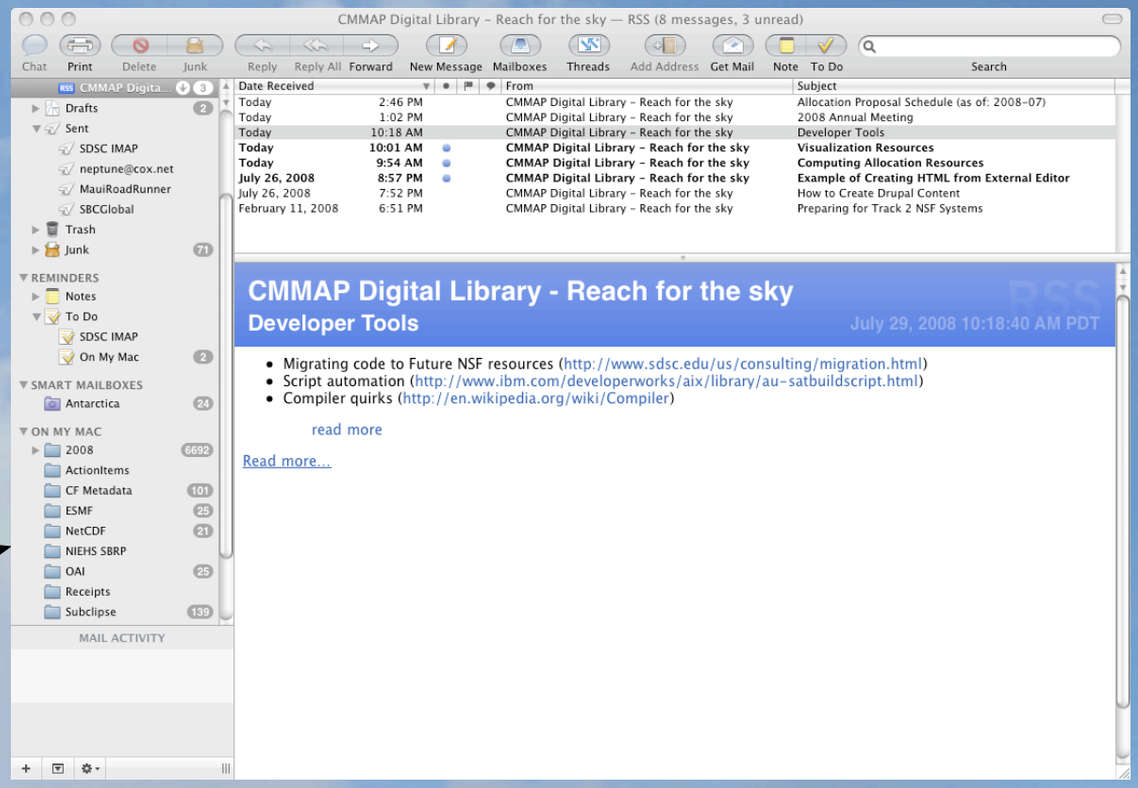

URL data retrieval

OAI metadata

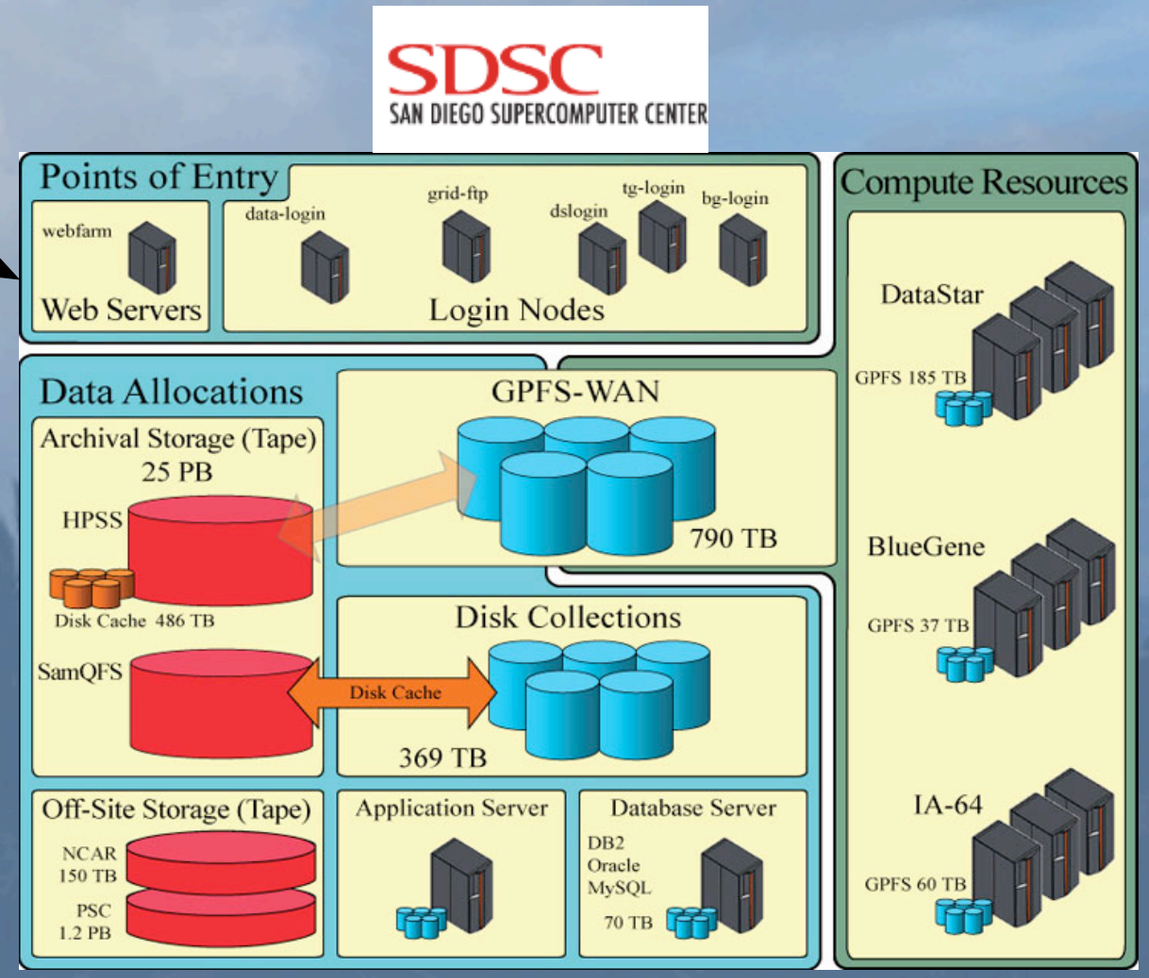

Management

Data Policy

Thursday, January 13, 2011

Data Policy

- Data published on the CMMAP Digital Library is in the public domain but registration and authorization required to access it
  - this is to prevent hacking and bot-crawling and
  - provide tracking of who is accessing the data
- All metadata is public
  - will be published via a new OAI (Open Archive Initiative) service to be instituted this year
  - CF metadata conventions are followed
- We are investigating the use of DOIs (digital object identifiers) for data consistent with the scholarly publication process

Community Support

Data Interoperability

Resources

Visualization of Very Large Datasets

Kitware ParaView 3.2.0

File Edit View Sources Filters Animation Tools Help

Time: 0

Result Surface

Pipeline Browser

cs://mpc1074:11111
comet_p4.pvtu
Calculator2
Contour2

Object Inspector

Properties Display Information

Apply Reset Delete

Isosurfaces

Value Range: [6.79195e-11, 7.55638]

6.79e-11
0.329
0.657
0.986
1.31
1.64
1.97
2.3
2.63

Delete
Delete All
New Value
New Range
Scientific

B
7.56
5.67
3.78
1.89
6.79e-11

Z Y X

ParaView

Thursday, January 13, 2011

3D visualization of geodesic data

VisIT

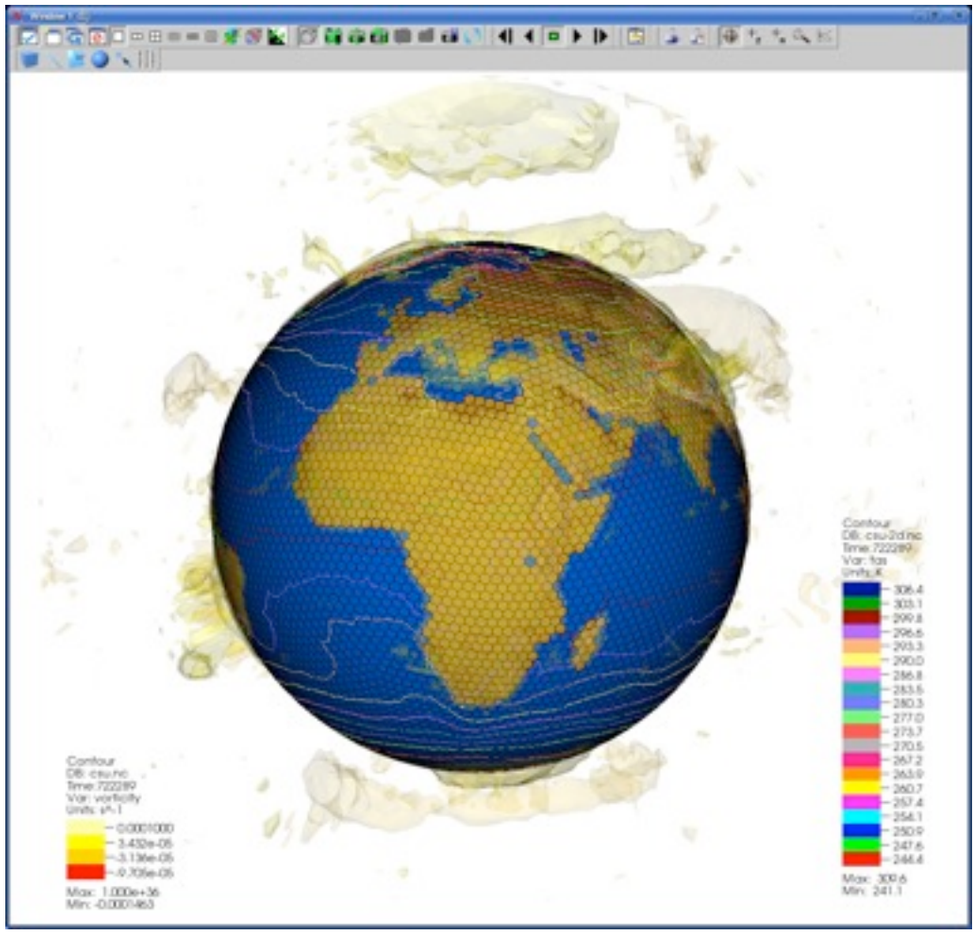

3D isocontours of vorticity.

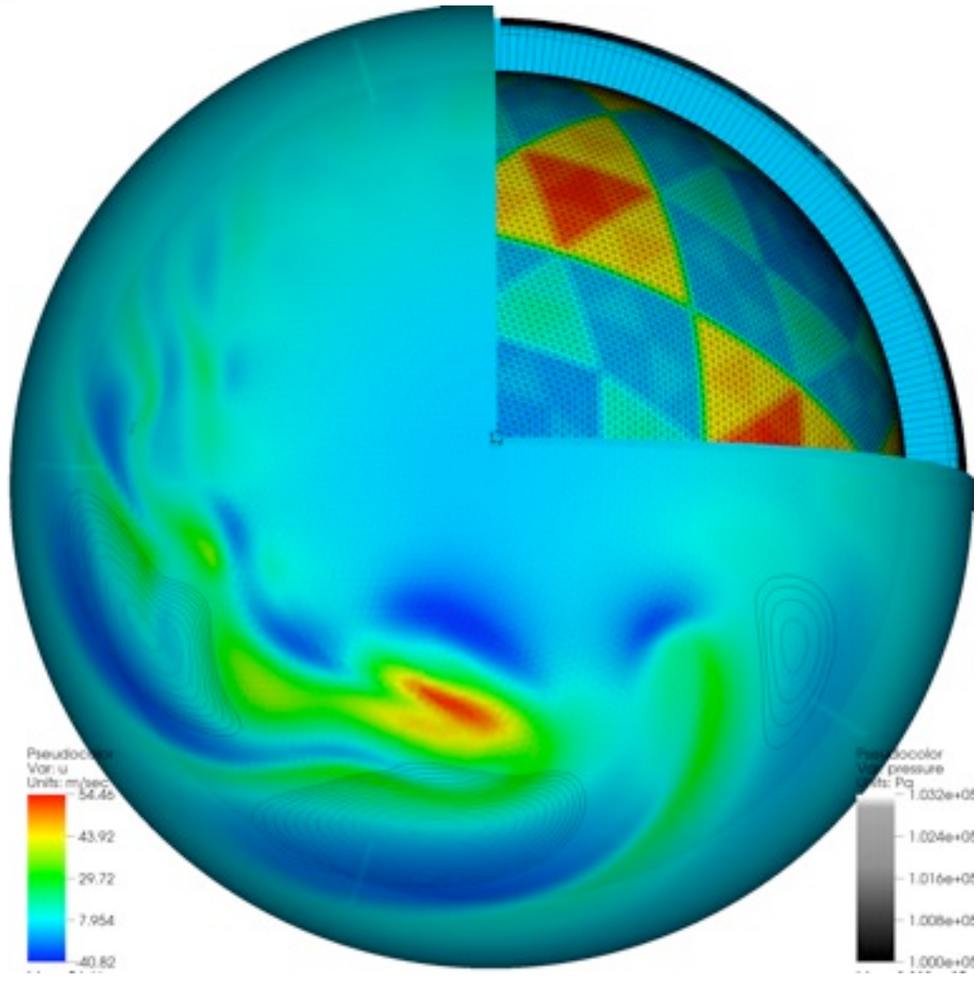

Composite plot of multiple mesh types and variables in the geodesic grid. Cell area (2D cell-centered data) and wind velocity (3D corner-centered on layers) data is shown by pseudocolor plots. Pressure (3D cell-centered on layers) is shown by contour lines.

Plots and movies courtesy of Prabhat (lbnl)

Community Support

Model Code Portability

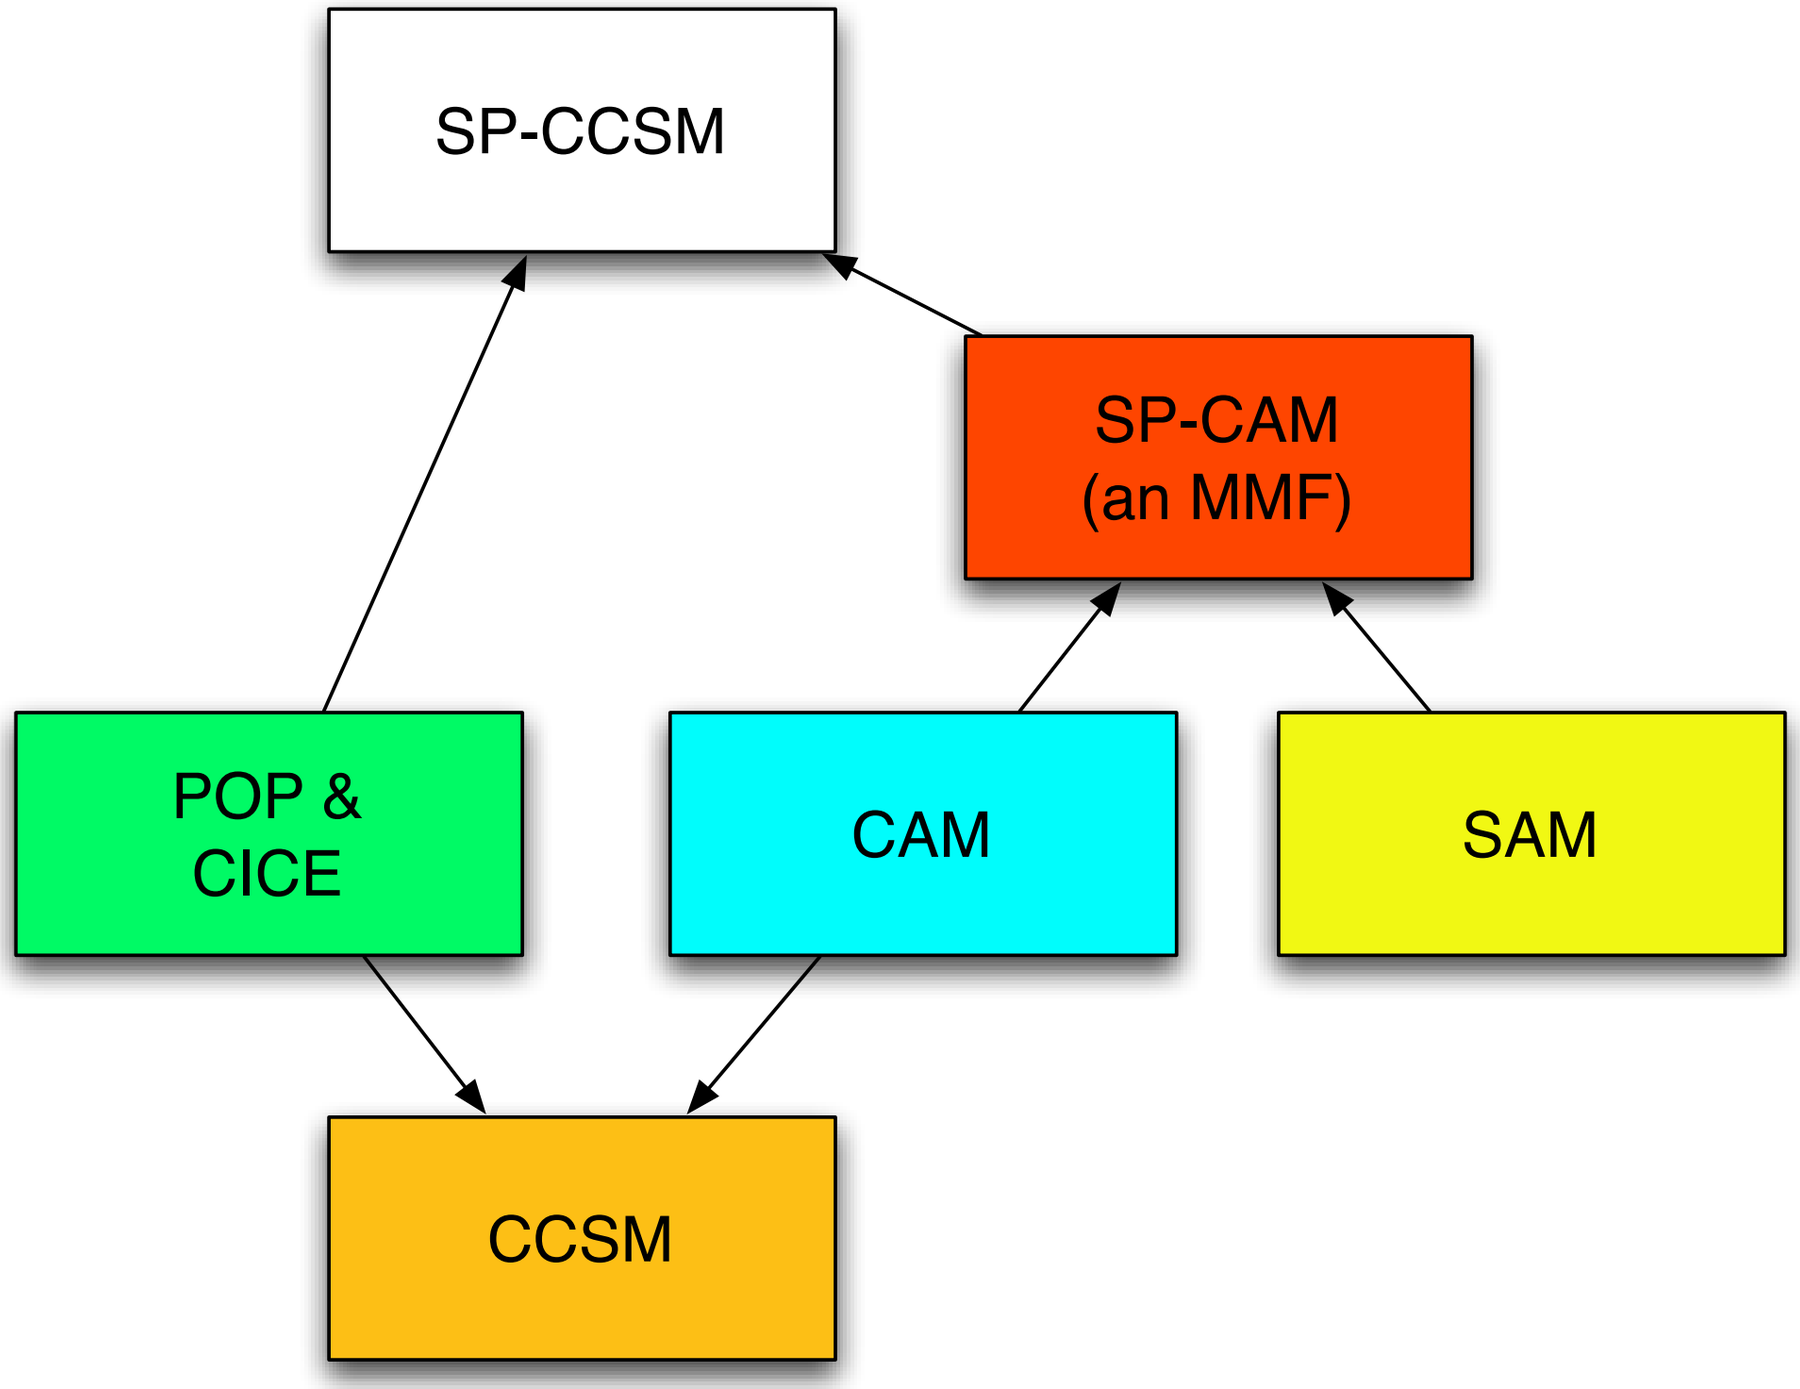
SP-CCSM
SP-CAM
(an MMF)
POP &
CICE
CAM
SAM
CCSM

Future Testing

- Remote viz on Teragrid
  - –TACC (UT Austin) / SPUR
  - –Super-LES data
  - –ParaView (client-server, remote X-session over ssh)
- UCSD Cave
- SDSC high-capacity network connections
- CSU network connections
- Other interested parties?