Cyberinfrastructure

11th Team Meeting
Fort Collins, CO
August, 2011

John Helly

CMMAP
Reach for the sky

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
- CMMAP CI Architecture (Roadmap)
- Data Transportation
- Discussion

Current Status

Blue Waters computing system

Blue Waters Project
Petascale Science and Engineering
Computing System
National Petascale Computing Facility
Education, Outreach and Training
Industry Partnerships
Events
Newsletters
Project Office
Blue Waters Contributes to Campus

Blue Waters Update—NCSA/IBM Joint Statement

Effective August 6, 2011, IBM terminated its contract with the University of Illinois to provide the supercomputer for the National Center for Supercomputing Applications' Blue Waters project.

NCSA is confident that its goal of building a sustained-petascale supercomputer remains achievable in a timely manner. NCSA is coordinating with the National Science Foundation to ensure project continuity and that the goals of the project are achieved.

The University of Illinois and NCSA selected IBM in 2007 as the supercomputer vendor for the Blue Waters project based on projections of future technology development. The innovative technology that IBM ultimately developed was more complex and required significantly increased financial and technical support by IBM beyond its original expectations. NCSA and IBM worked closely on various proposals to retain IBM's participation in the project but could not come to a mutually agreed-on plan concerning the path forward.

IBM will return money received to date and NCSA will return equipment delivered by IBM per terms of the contract.

The University of Illinois at Urbana-Champaign is home to a modern, energy-efficient data center that provides high-bandwidth connectivity to national and international networks and a massive archival storage system. NCSA also has deep staff expertise in computer and computational science that will ensure the science and engineering community can take full advantage of any new supercomputer.

IBM, the University of Illinois, and NCSA will explore other opportunities to continue the strong working relationship established during the Blue Waters project.

For more information, contact Bill Bell, 217.265.5102, or Trish Barker, 217.265.8013.

The Blue Waters project is supported by the National Science Foundation and the University of Illinois.

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		

The Teragrid is Dead. Long Live the Teragrid

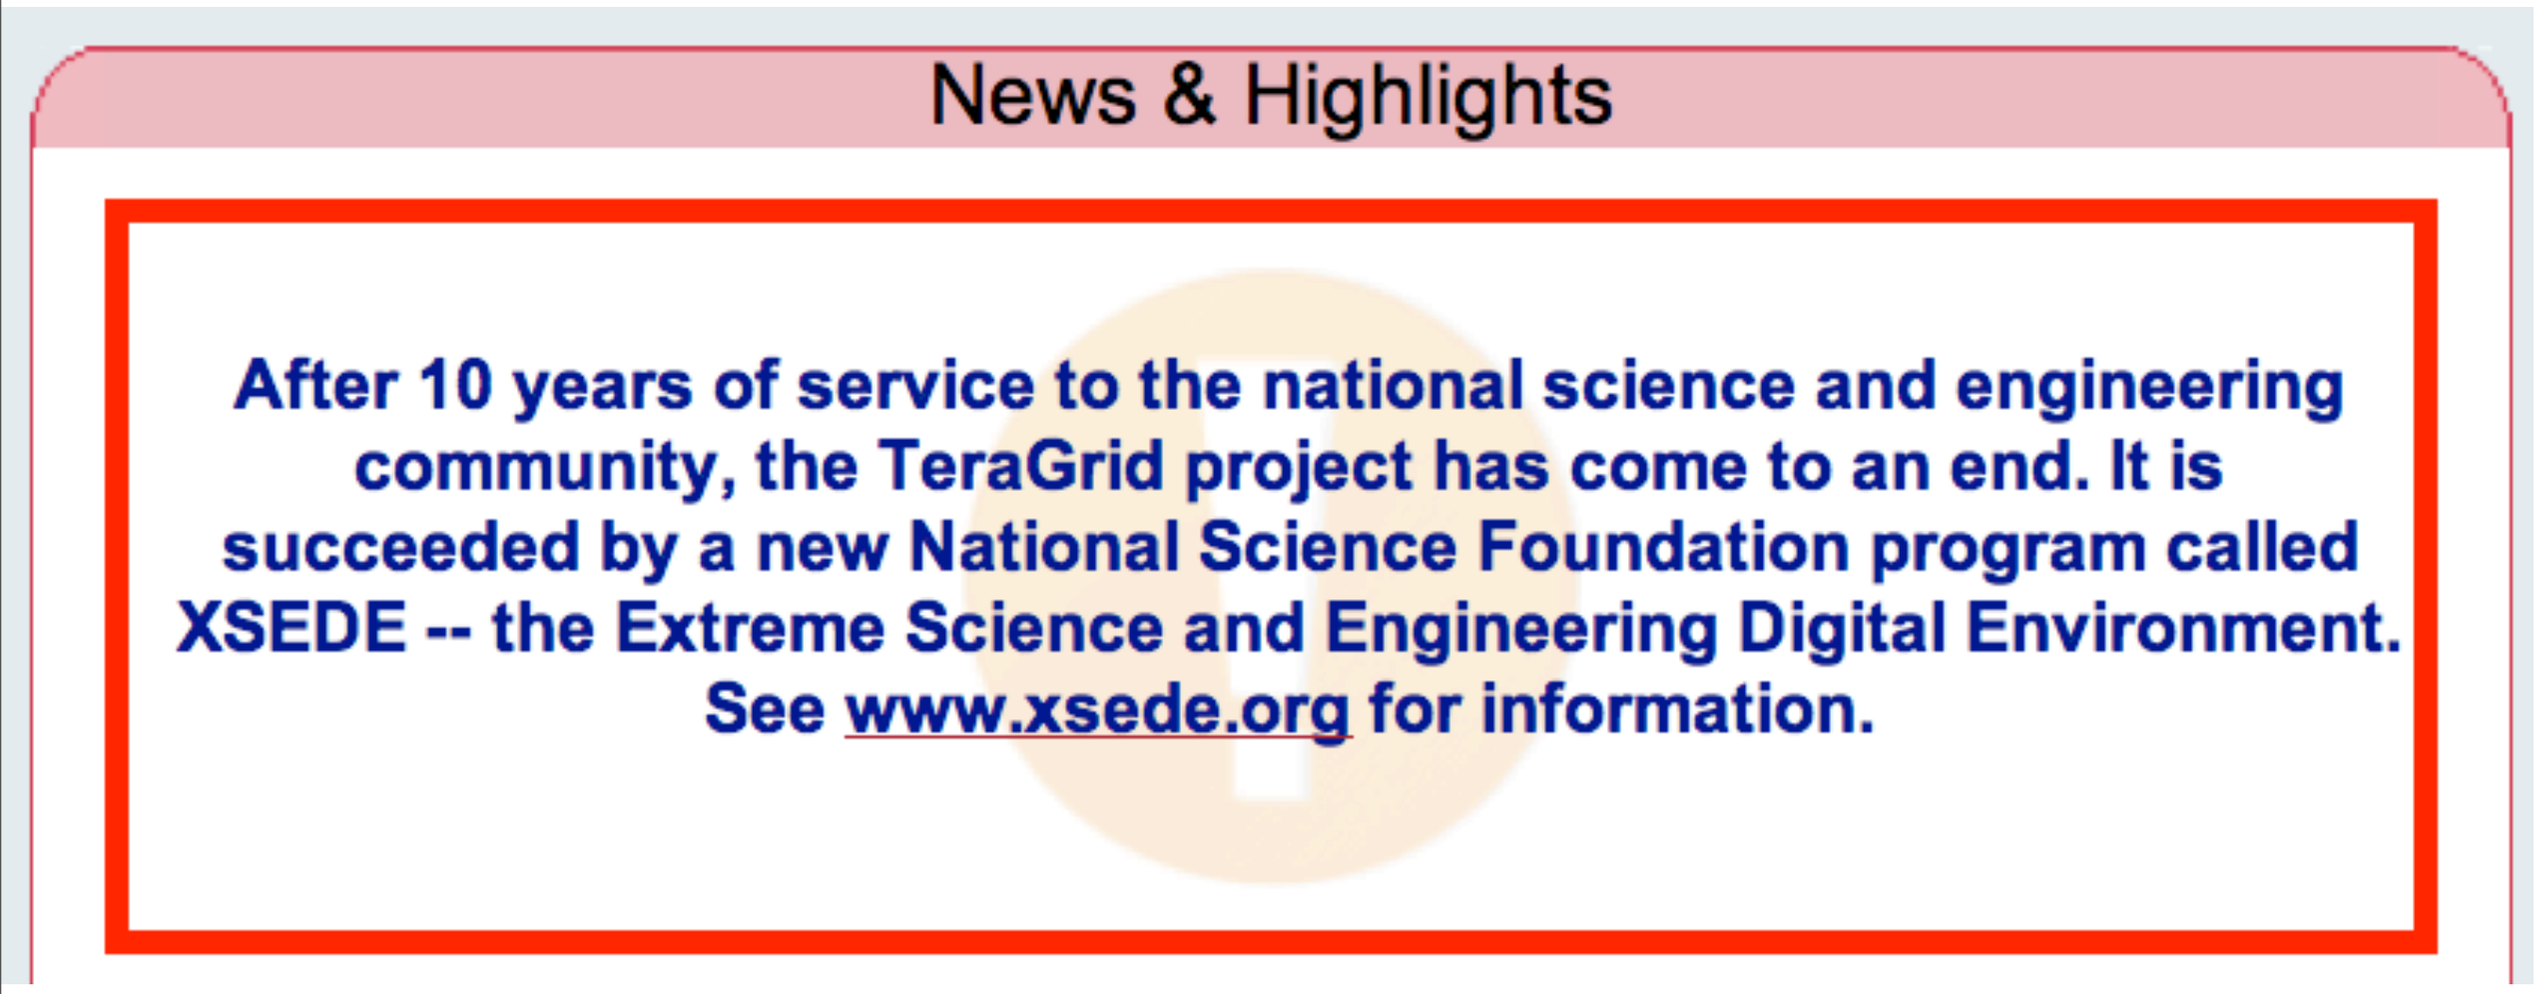

Projects

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

Project PI: Helly, John

Charge No.: TG-ATM100027

RESOURCE	SUS AWARDED	SUS REMAINING	% REMAINING	MY USAGE (SU)	START DATE	END DATE	ALLOC. TYPE	STATE
asta.teragrid	7	7	100%	0.0	2010-12-24	2011-12-30	transfer	active
abe-queenbee-steele.teragrid	2,307,000	2,277,261	99%	0.0	2010-07-01	2011-12-30	new	active
NCSA Tape	5	5	100%	0.0	2010-07-01	2011-12-30	new	active

Regionalization of Anthropogenic Climate Change Simulations. - Atmosphere ocean coupled downscaling

Project PI: Kanamitsu, Masao

Charge No.: TG-ATM090032

RESOURCE	SUS AWARDED	SUS REMAINING	% REMAINING	MY USAGE (SU)	START DATE	END DATE	ALLOC. TYPE	STATE
Ranger	2,808,000	1,158,222	41%	0.0	2010-10-01	2011-09-30	renewal	active

HPC SYSTEMS | ADVANCED VIS SYSTEMS | STORAGE SYSTEMS | SPECIAL PURPOSE SYSTEMS

NAME	INSTITUTION	SYSTEM	PEAK TFLOPS	MEMORY TBYTES	STATUS	LOAD	RUNNING JOBS	QUEUED JOBS	OTHER JOBS
Kraken	NICS	Cray XT5	1174.00	147.00	Up		503	312	142
Ranger	TACC	Sun Constellation Cluster	579.40	123.00	Up		411	509	170
Lonestar	TACC	Dell Linux Cluster	302.00	45.00	Up		0	309	78
Trestles	SDSC	Appro AMD Magny-Cours Cluster	100.00	20.25	Up		141	0	180
Steele	Purdue	Dell Intel 64 Linux Cluster	60.00	12.40	Up*				
Lincoln	NCSA	Dell/Intel PowerEdge 1950	47.50	3.00	Up		30	27	1
Blacklight	PSC	SGI UV	37.20	32.00	Up		94	164	1
Pople	PSC	SGI Altix 4700	5.00	1.54	Up*		40	0	21
Dash	SDSC	Appro Intel Nehalem Cluster	4.90	3.00	Up		2	3	0
		Total:	**2310.0**	**387.19**			**1221**	**1324**	**593**

**Indicates failure of one or more status test.*

Hover mouse pointer over Resource Name, Resource Status, and headings to see additional information.

Gordon User Guide
- Gordon Home & Technical Summary
- Logging In
- Jobs
- Storage
- Software

Gordon User Guide: Technical Summary

Gordon will be a powerful 1024-node supercomputer, unique in its extensive use of flash memory and virtual shared-memory "supernodes". Gordon's architecture will allow it to reduce solution times of data-intensive problems. It is scheduled to be installed in mid-2011. Gordon is the follow-up to Dash, the first supercomputer to use flash memory.

Technical Summary

- 1024 compute nodes
- 32 supernodes, each with:
  - 32 compute nodes at 240 gigaflops per node
  - 2 I/O nodes with 4 TB of flash memory per node
  - 10 TB of memory (2 TB of DRAM and 8 TB of flash memory)
  - Peak performance of 7.7 teraflops
- Peak system performance of 245 teraflops
- 64 TB of DRAM
- 256 TB of flash memory and 4 PB of disk storage

Contact Us
consult@sdsc.edu
1-866-336-2357
(Toll-Free)

Report a security incident

Important Links
Quick Comparison: Storage Resources

Did You Get What You

XSEDE | USER PORTAL
Extreme Science and Engineering Discovery Environment

Search XSEDE...

HOME MY XSEDE RESOURCES DOCUMENTATION ALLOCATIONS TRAINING USER FORUMS HELP

Overview Course Calendar Online Training

Several online training opportunities are available for both XSEDE and the various technology areas within XSEDE. Please view the list of available training classes below:

PROGRAMMING
Introduction to Linux (Cornell Virtual Workshop)
An Introduction to C Programming (Cornell Virtual Workshop)
An Introduction to Fortran Programming (Cornell Virtual Workshop)
MATLAB Programming (Cornell Virtual Workshop)

SCRIPTING TOPICS
Python (Cornell Virtual Workshop)
Balancing Scripts and Compiled Code in Scientific Applications (Cornell Virtual Workshop)

PARALLEL PROGRAMMING CONCEPTS
Parallel Computing Explained (CI-Tutor)
Parallel Programming Concepts and High-Performance Computing (Cornell Virtual Workshop)
Multilevel Parallel Programming (CI-Tutor)

MESSAGE PASSING INTERFACE (MPI)
Message Passing Interface (MPI) (Cornell Virtual Workshop)
Introduction to MPI (CI-Tutor)
Intermediate MPI (CI-Tutor)

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**	**Close Submissions**
			Dec. 15 Mar. 15 [1] Jun. 15 Sept. 15	**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	

allocation.

- **Research** allocation requests are reviewed at quarterly XRAC meetings.
  Research allocations may be requested for any computation resource and require a formal request document and CVs (for PI/Co-PIs). Research allocations are typically appropriate as follow-ons to Startups; but a PI need not request a Startup prior to submitting a Research request. Requests may be submitted the quarter before they are needed during designated submission windows. The XSEDE Resource Allocations Committee (XRAC) meets quarterly to review requests based on merit and award available SUs.

OPEN SUBMISSIONS	CLOSE SUBMISSIONS	ALLOCATIONS BEGIN
Dec 15	Jan 15	April 1
Mar 15	Apr 15	Jul 1
Jun 15	Jul 15	Oct 1
Sep 15	Oct 15	Jan 1

Writing Your Research Allocation Request

Well written requests contain all the information the review panel requires to assess your project qualifications. Details for writing a request to use any XSEDE resource are available in the XSEDE Resource Allocation Policies.

Well written research allocation requests have had the following noted characteristics:

- The research was summarized in context of the current state of the art; it outlined the computational algorithms to be used; and it related those algorithms to subsections of the request.
- The reviewers were provided sufficient information, but not overwhelmed by details.
- The justification for the request was clear, and closely coupled to computational experiments and needs, so that if the committee needed to reduce the original request, it could be done rationally with minimum disruption to the investigator.
- Results from relevant previous allocations, including manuscripts published, accepted, submitted, or in preparation, were summarized and related to the current request.

Sample Research allocation requests are available as examples of how requests should be presented. Please view the How to write a winning proposal page for advice on Science Gateway proposals.

Submitting Your Research Allocation Request

Requests for computing allocations must be submitted electronically via the XSEDE User Portal.

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

iRODS Command-line Client: Scriptable

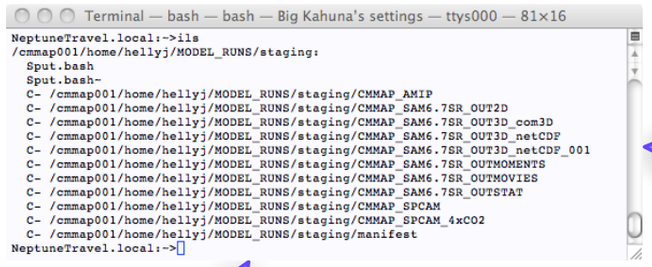

iRODS Federation Mechanism

CMMAP Phase II CI @CSU

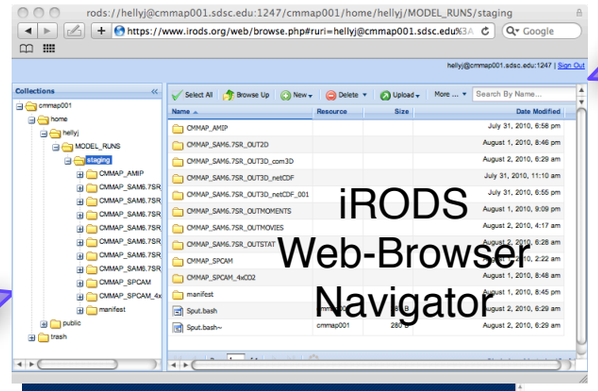

iRODS Web-Browser Navigator

MMF Community Portal

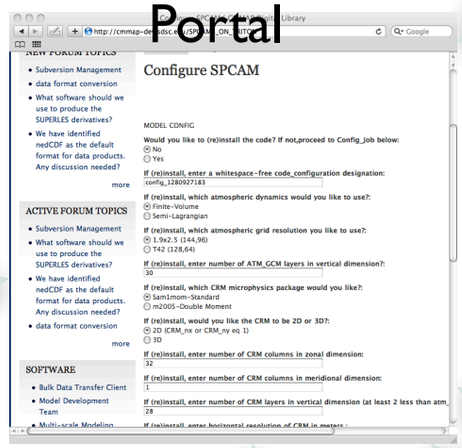

CMMAP Disk Allocation @SDSC

Drupal: programmability

- iquery
- MMF Workbench

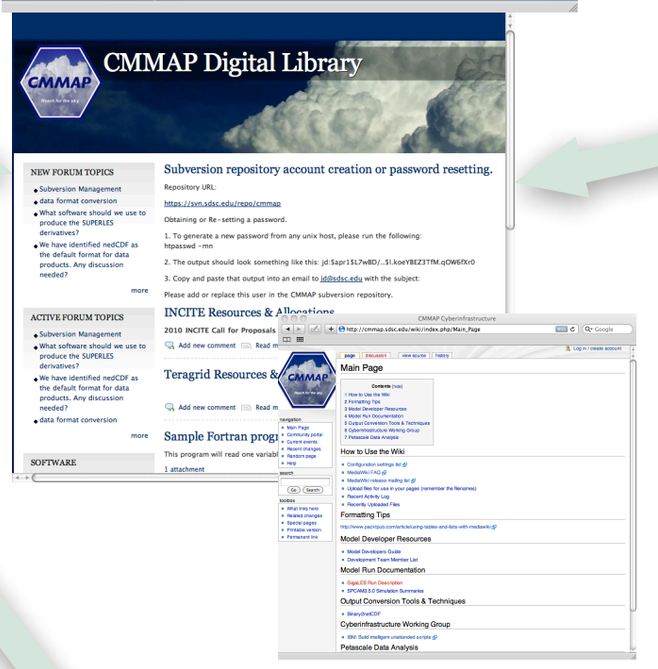

MediaWik: Documentationi

FRE
Workflow Control
Collaboration w/GFDL

others

Community Account

Steele@Teragrid

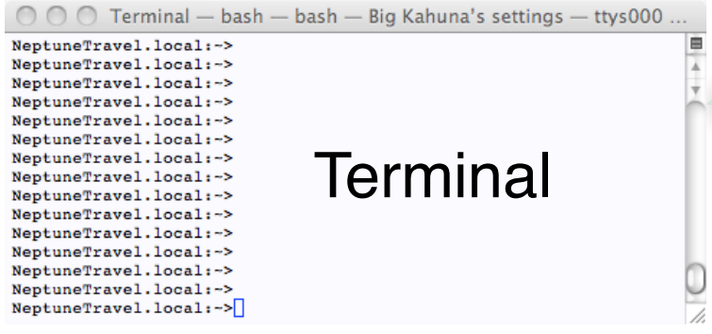

Terminal

Computing

Data

Subversion Repository

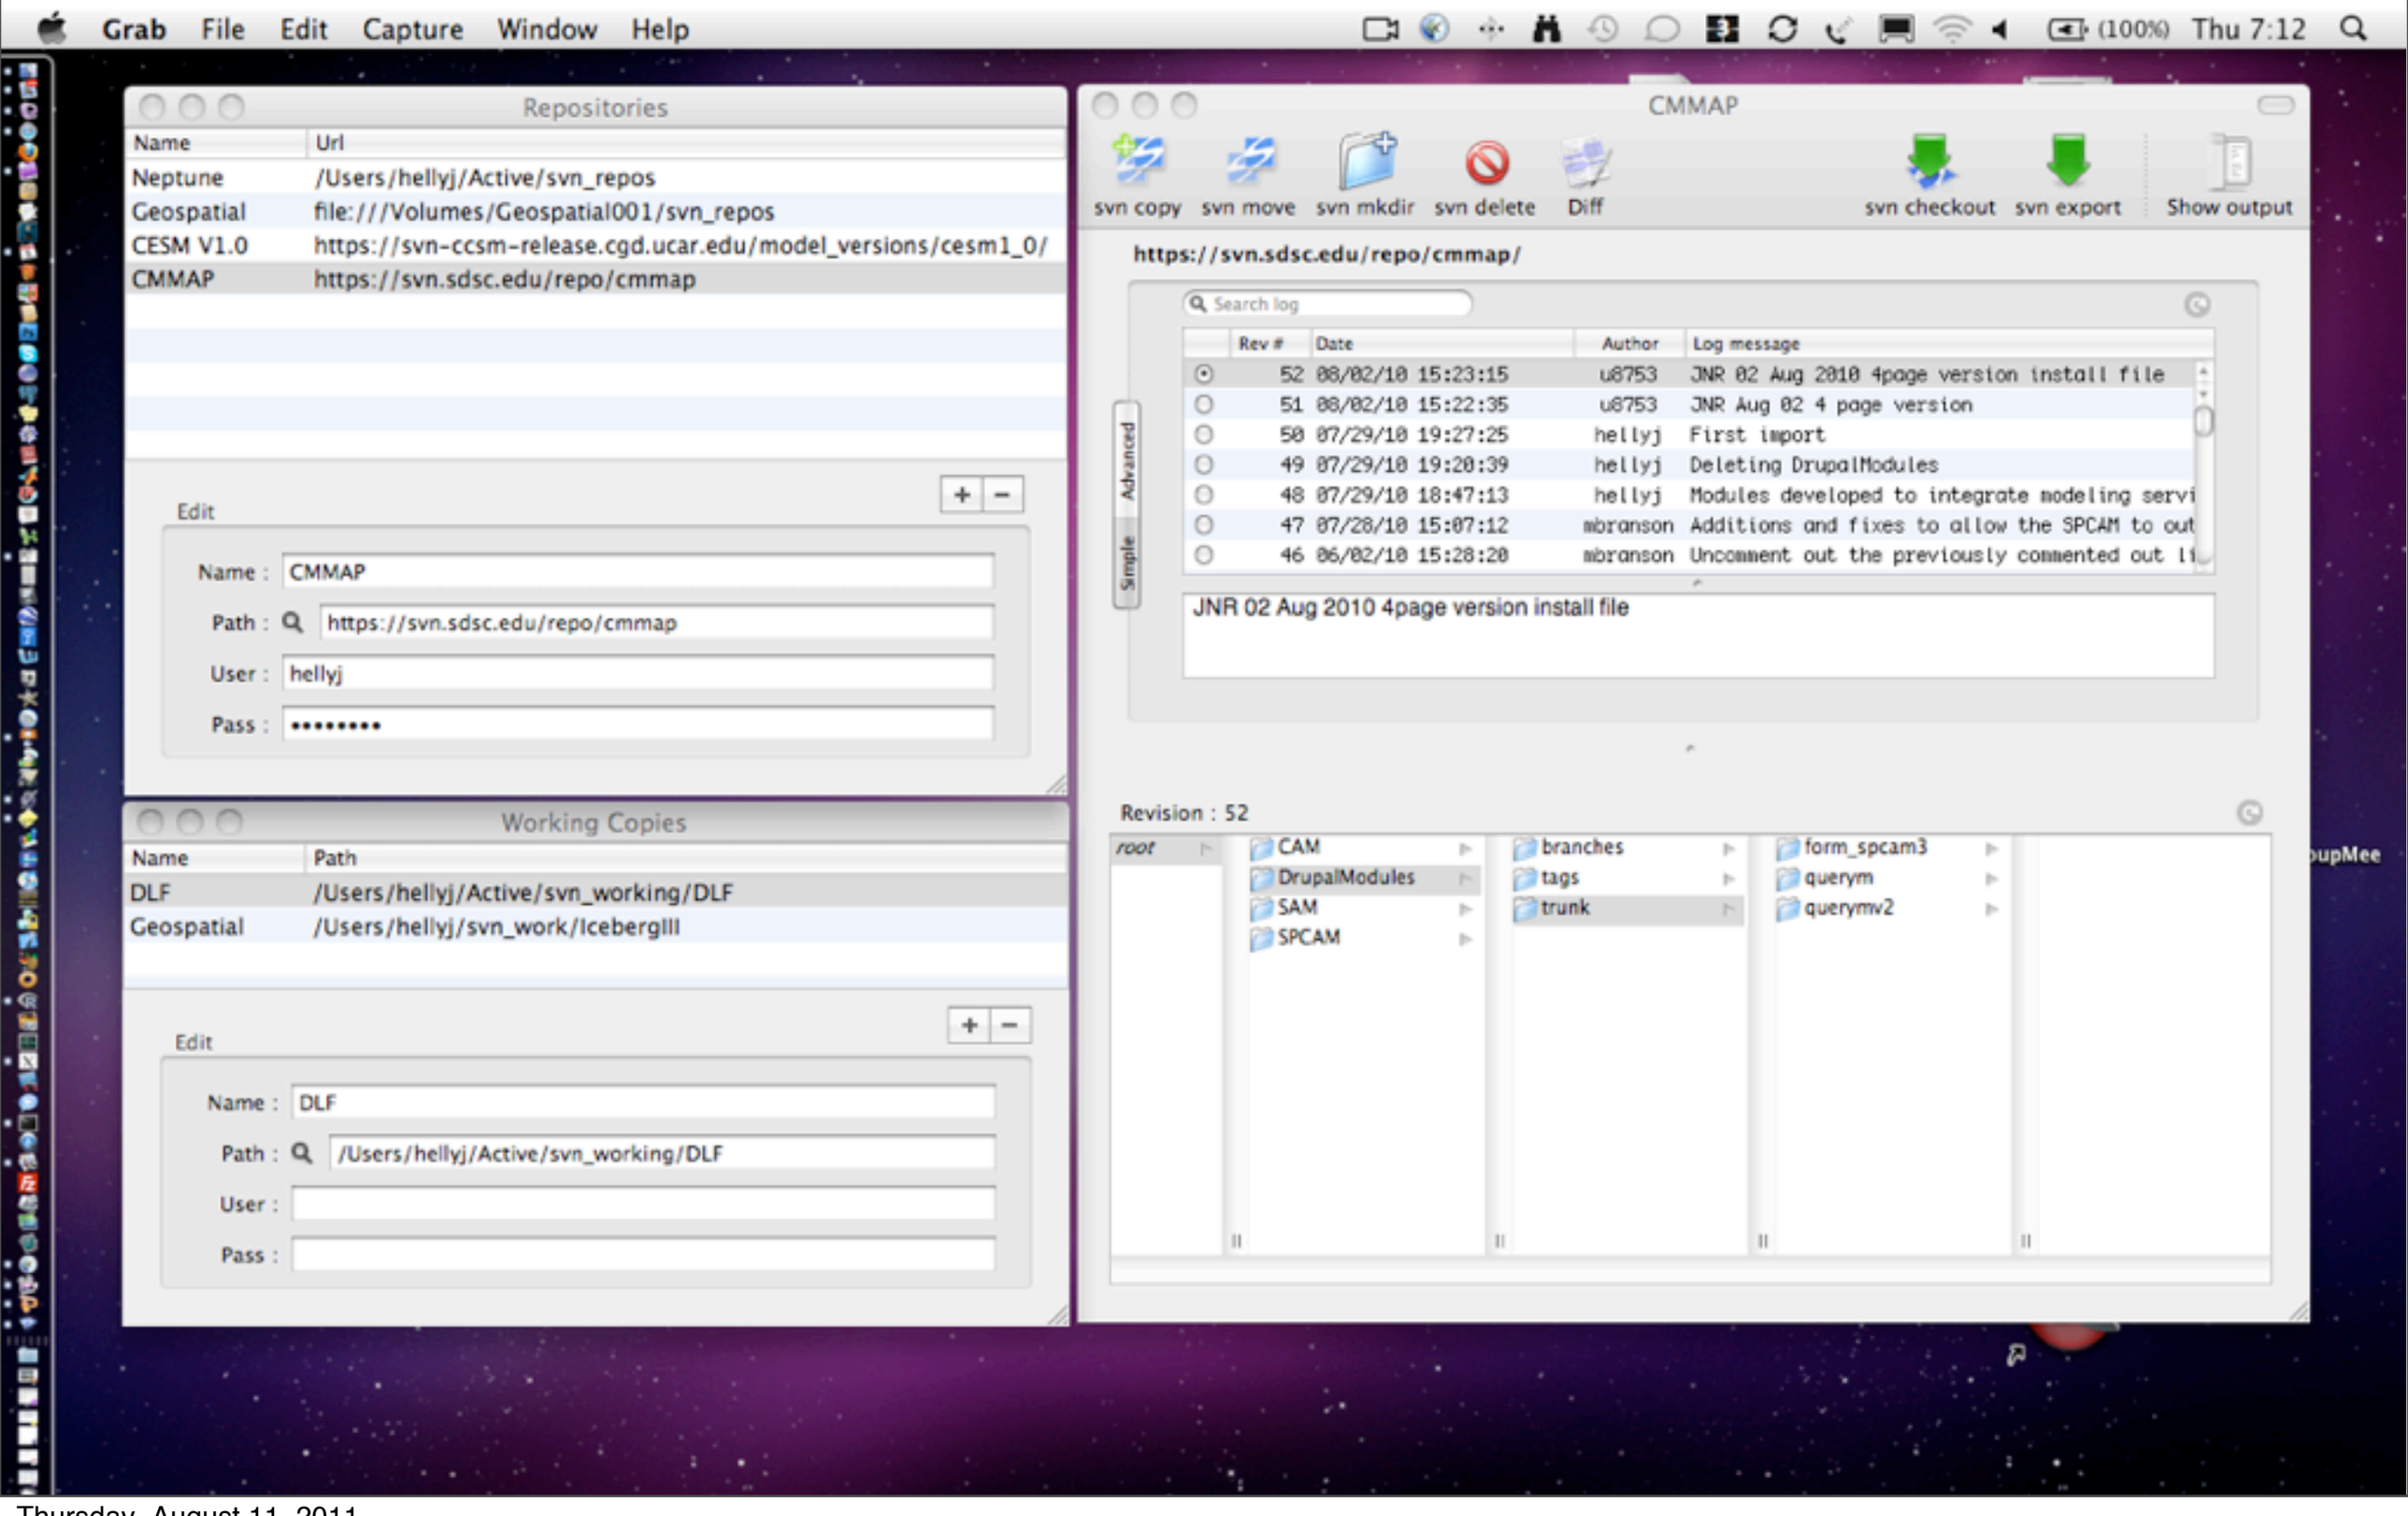

iRODS Command-line Client: Scriptable

```
Terminal — bash — bash — Big Kahuna's settings — ttys000 — 81×16
NeptuneTravel.local:~>ils
/cmmap001/home/hellyj/MODEL_RUNS/staging:
  Sput.bash
  Sput.bash~
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_AMIP
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_SAM6.7SR_OUT2D
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_SAM6.7SR_OUT3D_com3D
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_SAM6.7SR_OUT3D_netCDF
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_SAM6.7SR_OUT3D_netCDF_001
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_SAM6.7SR_OUTMOMENTS
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_SAM6.7SR_OUTMOVIES
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_SAM6.7SR_OUTSTAT
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_SPCAM
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/CMMAP_SPCAM_4xCO2
  C- /cmmap001/home/hellyj/MODEL_RUNS/staging/manifest
NeptuneTravel.local:~>
```

iRODS Federation Mechanism

CMMAP Phase II CI @CSU

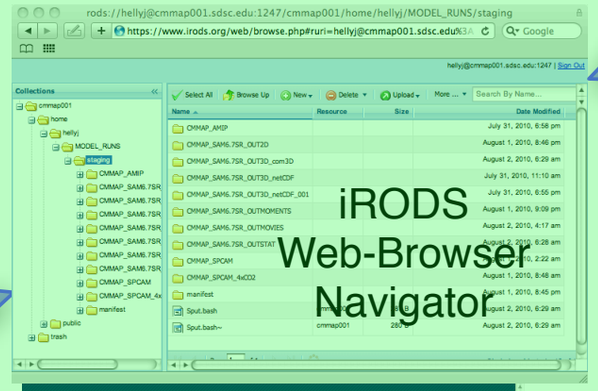

iRODS Web-Browser Navigator

Drupal: programmability

- iquery
- MMF Workbench

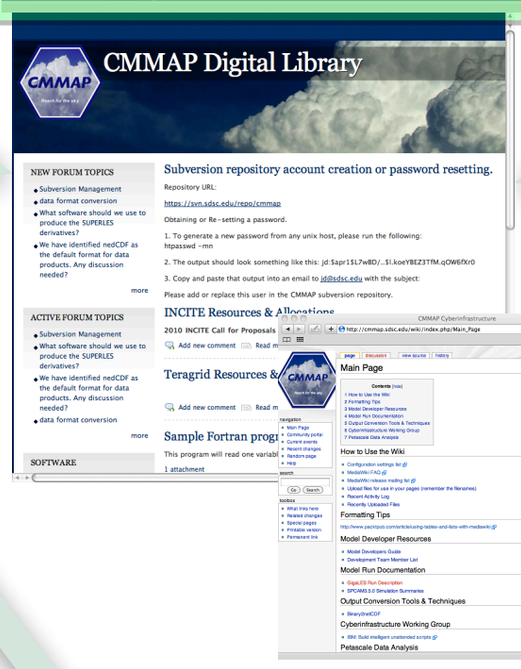

MediaWik: Documentationi

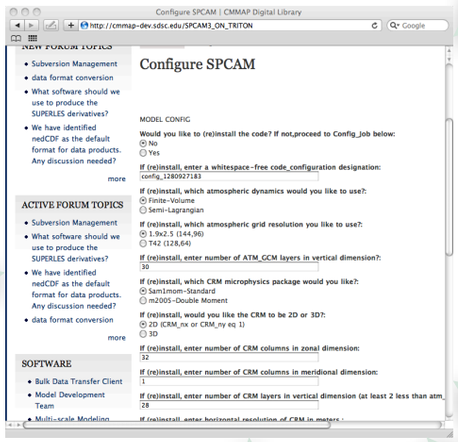

CMMAP Disk Allocation @SDSC

FRE
Workflow Control
Collaboration w/GFDL

others

Community Account

Steele@Teragrid


```
Terminal — bash — bash — Big Kahuna's settings — ttys000 ...
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
NeptuneTravel.local:~>
```

Terminal

Computing

Data

iRODS Web-broswer

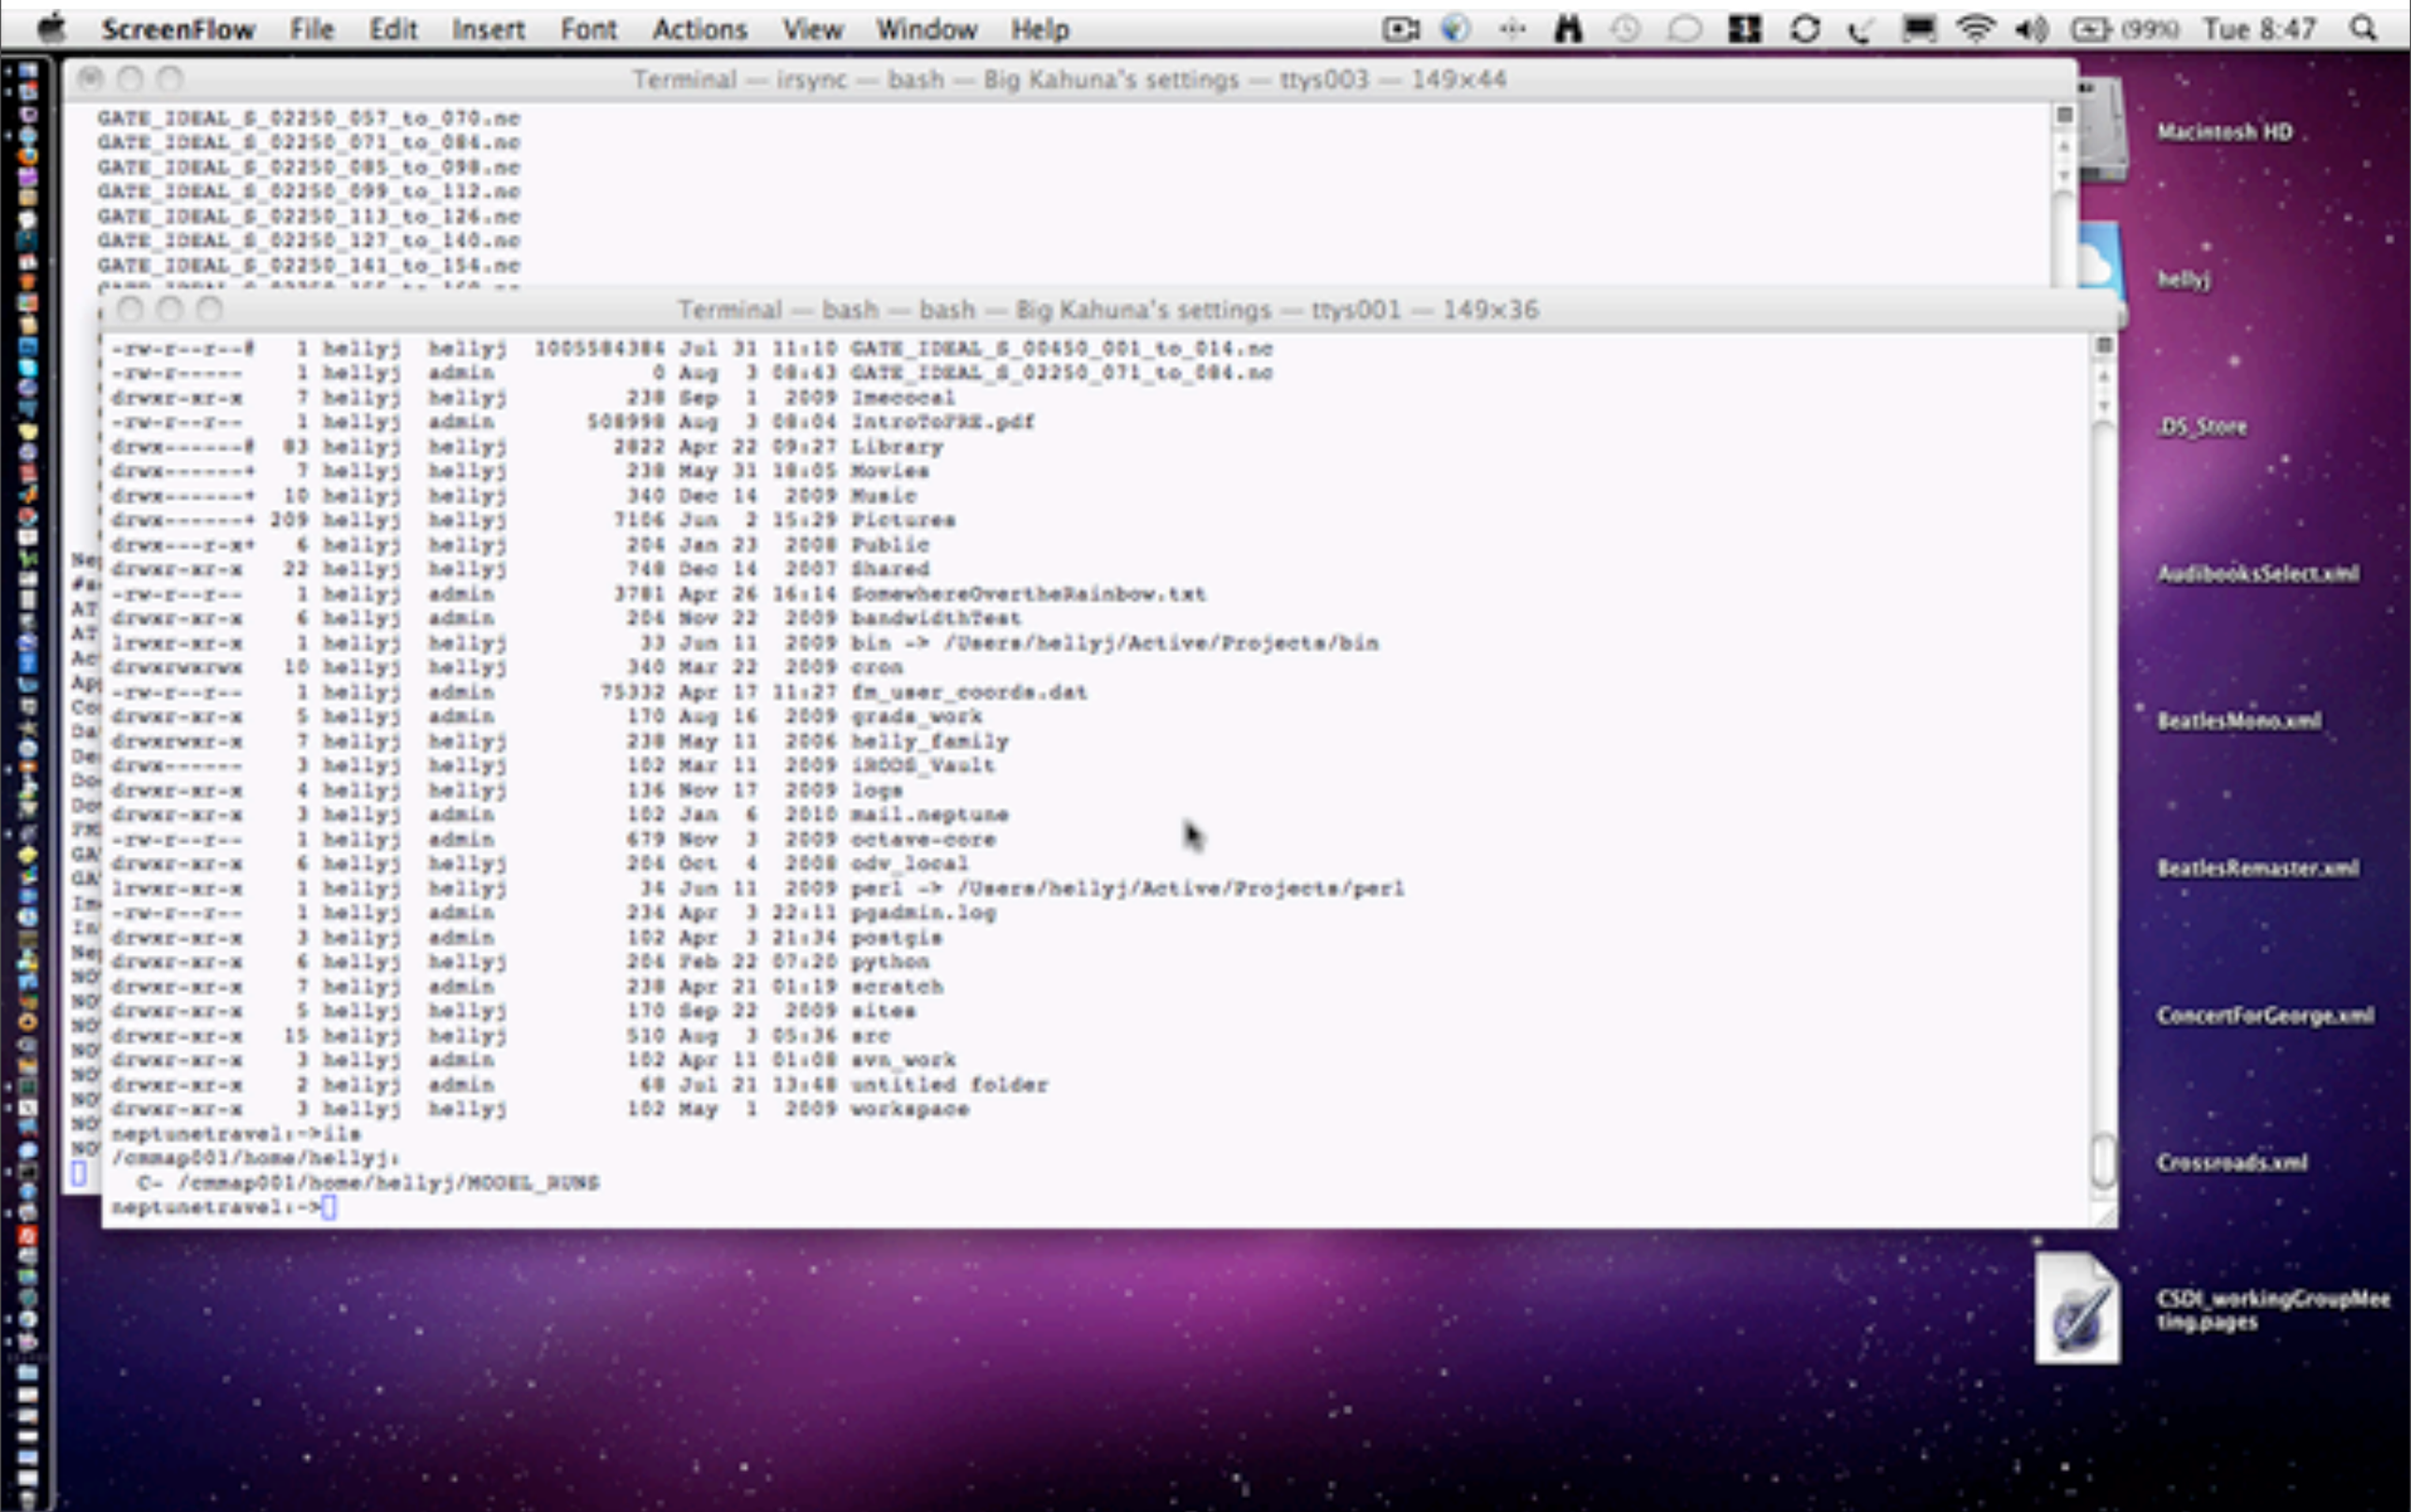

Model Run Management App

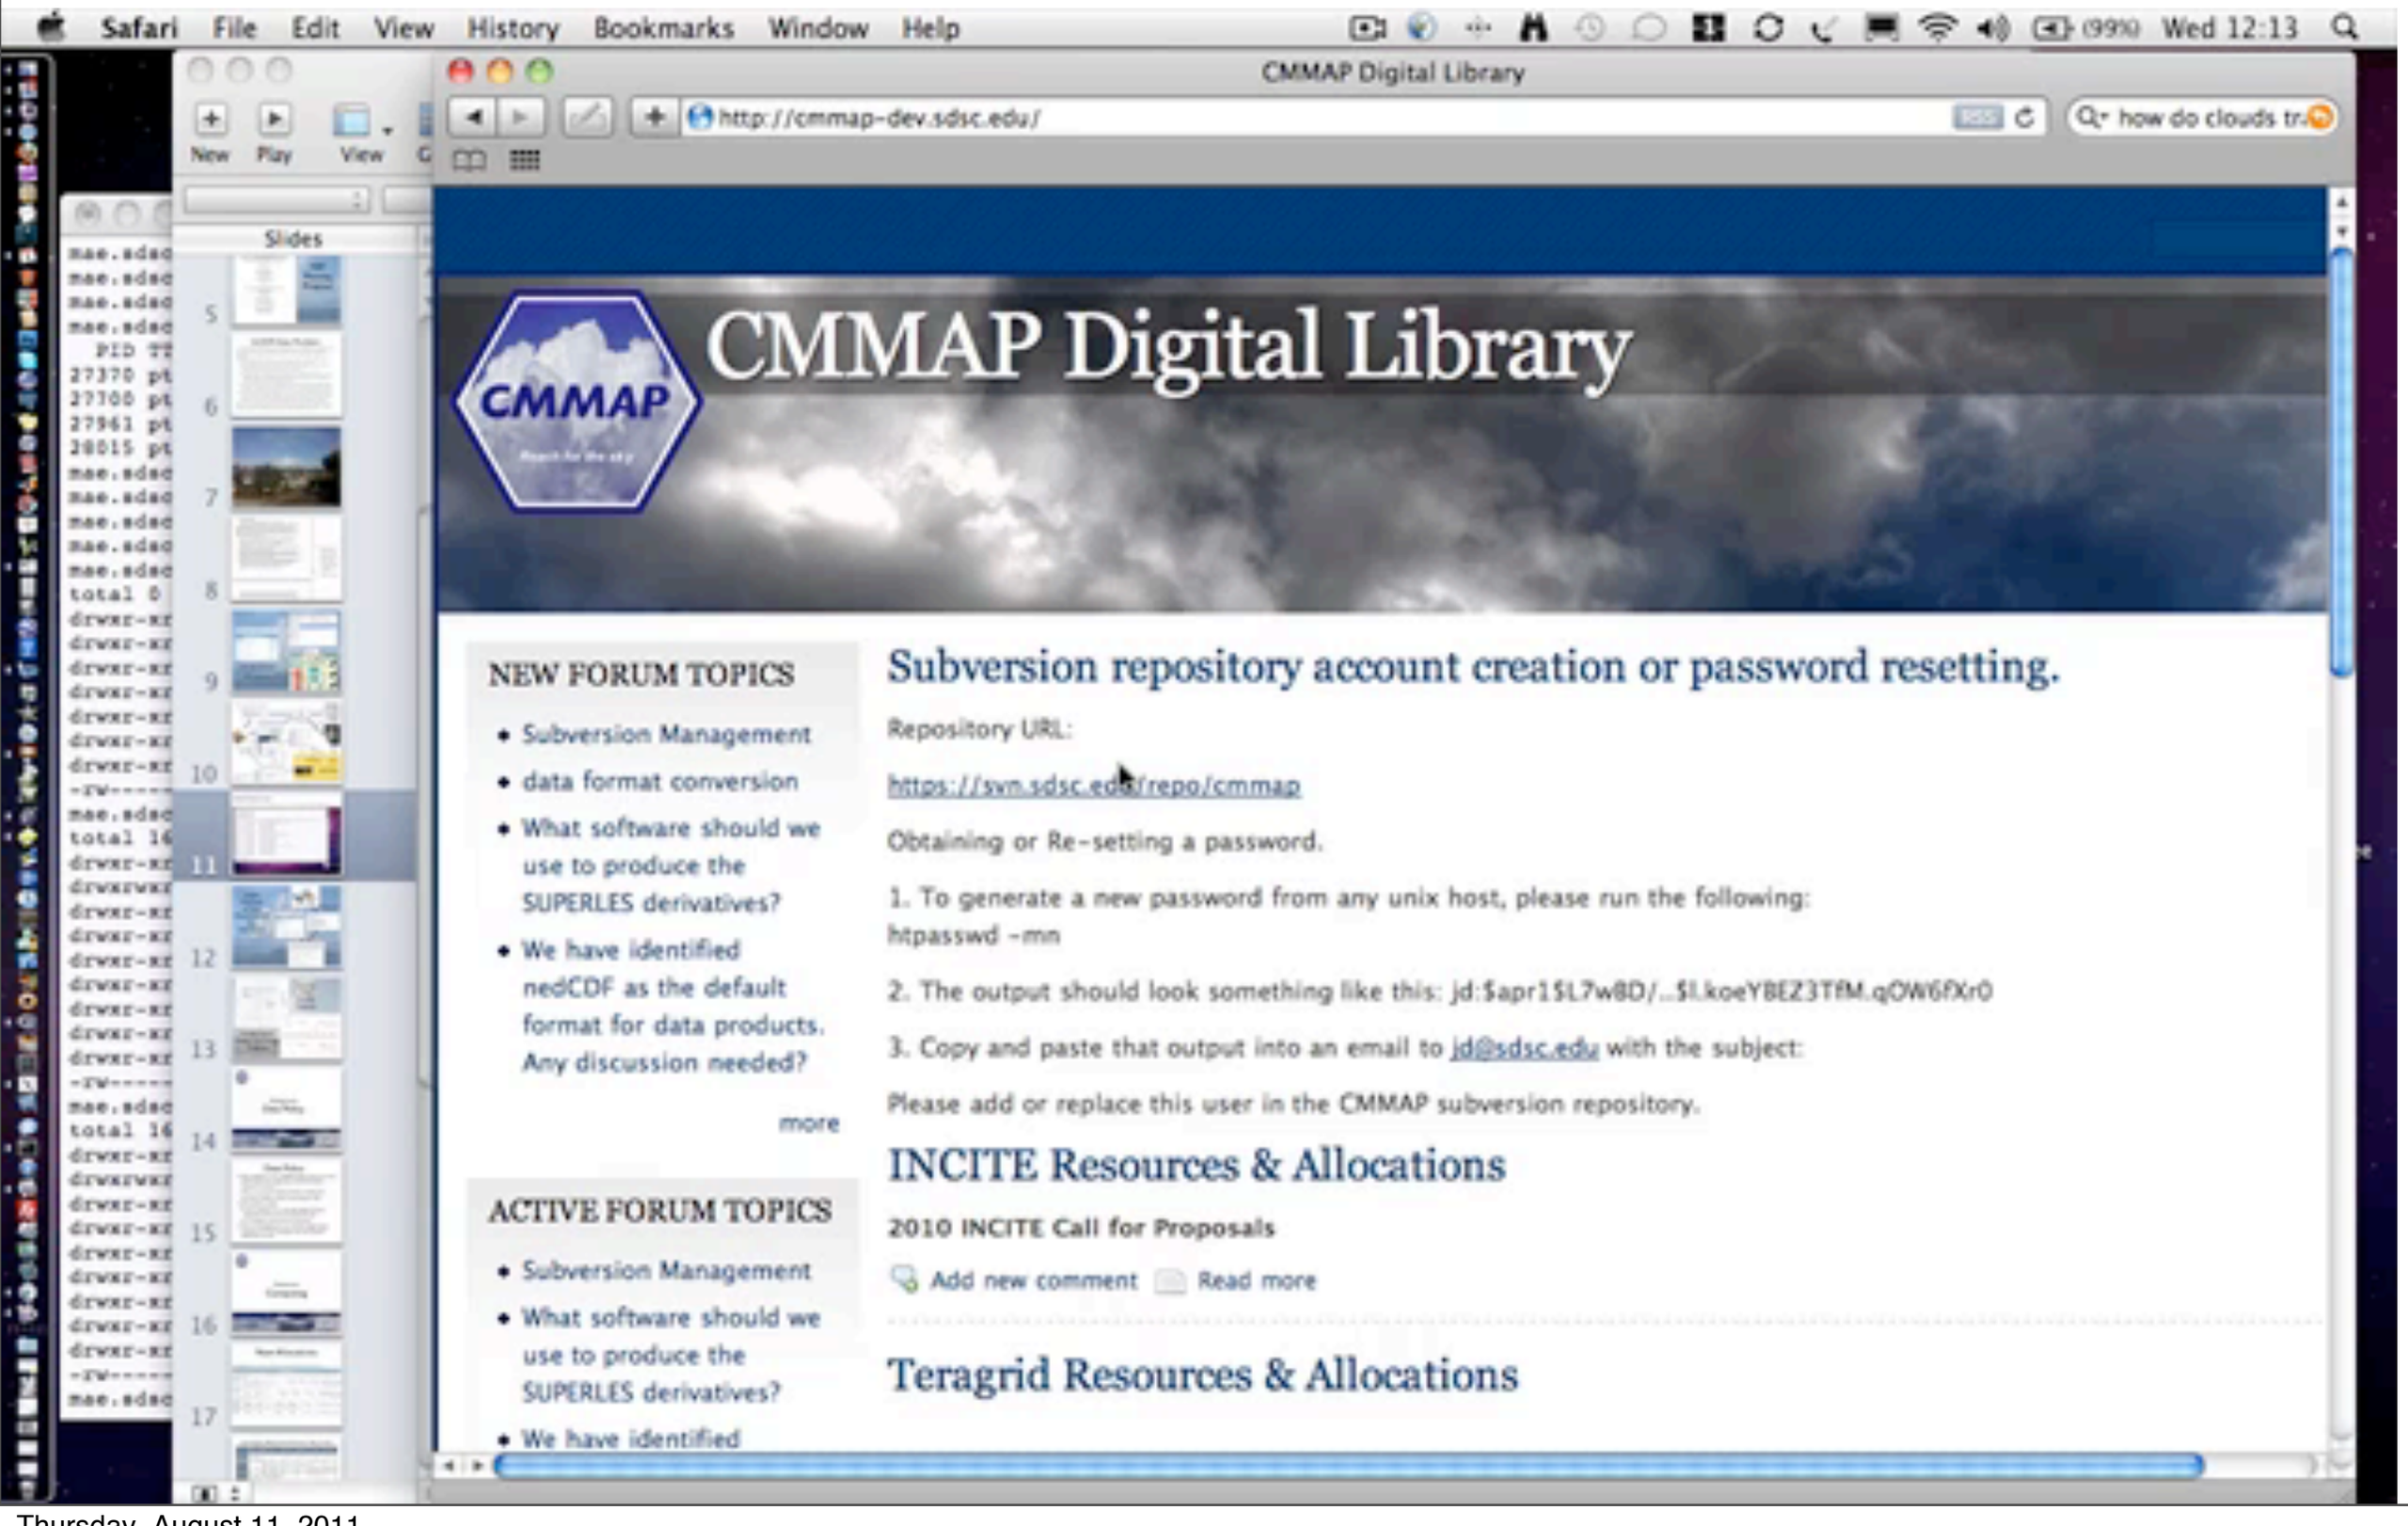

iRODS Command-line Client: Scriptable

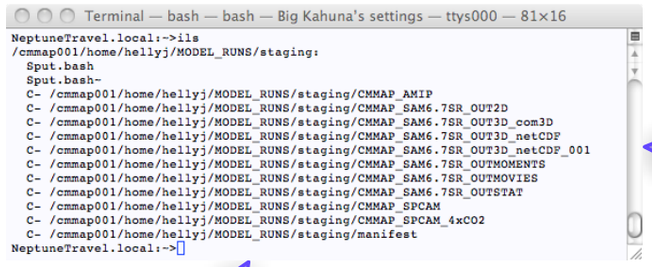

Output returned to SDSC for post-processing analysis and publication

San Diego Supercomputer Center

MMF CMS module

CMMAP Digital Library

Subversion Code Repository

Job Submission and Monitoring

Subversion Working Copy

Remote HPCC Computer (e.g., Steele, Kraken)

Web Browser

Local terminal for microphysics code development

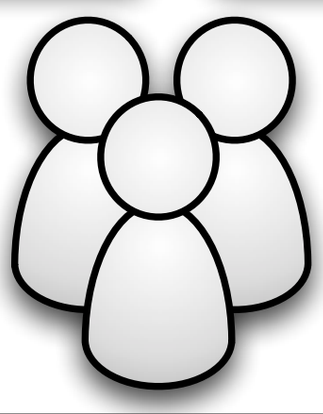

Extension to Data Analysis and Visualization on other HPCC resources such as Trestles and Gordon (SDSC) using new high capacity storage systems (Data Oasis) and virtual point-to-point network connections (DYNES)

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

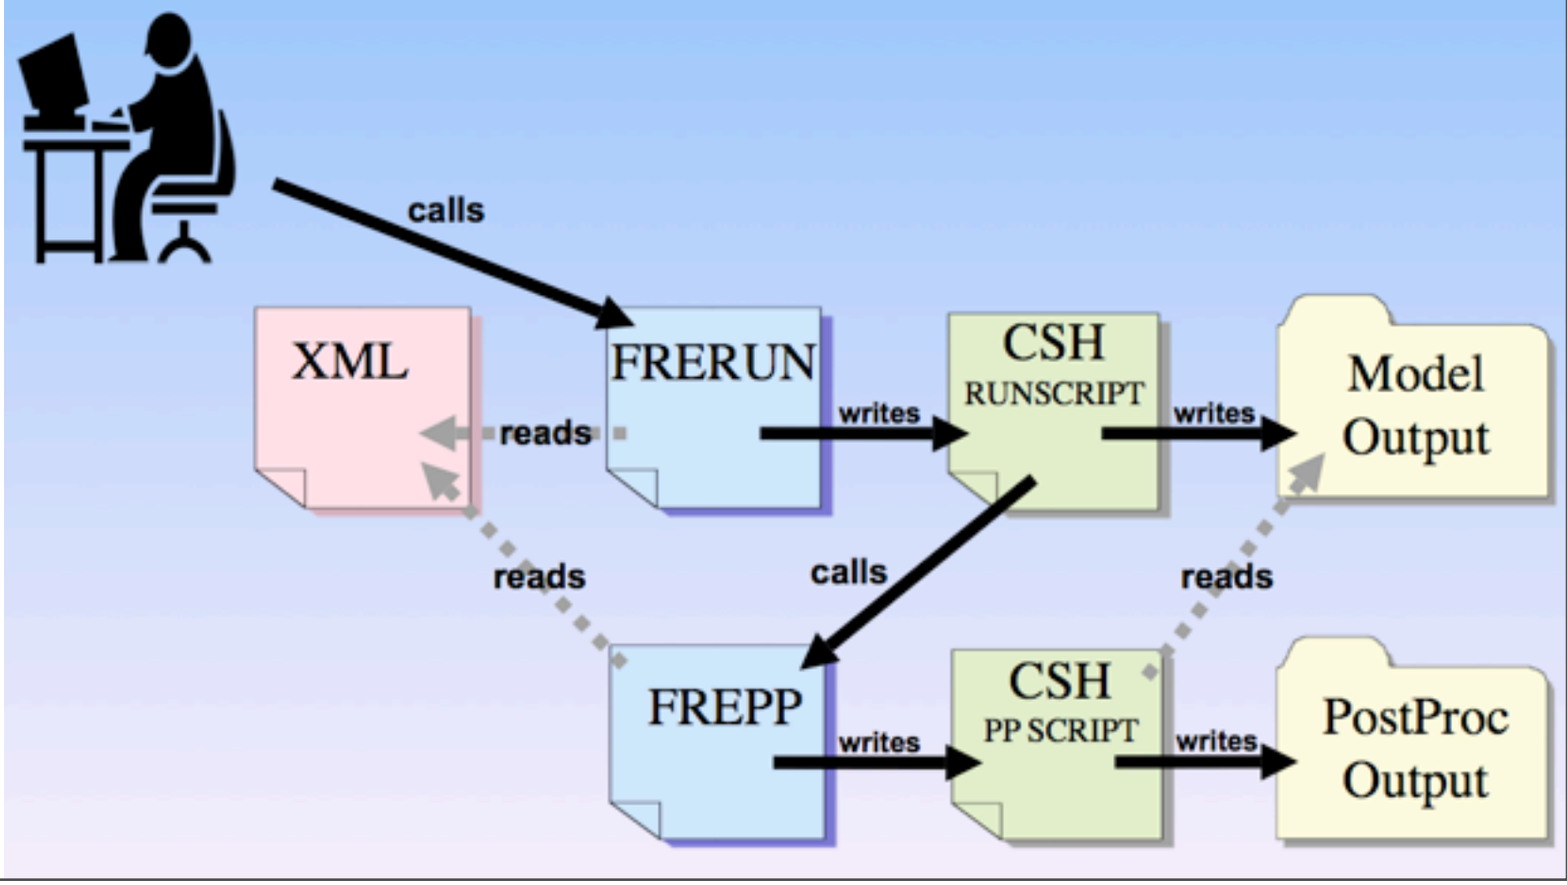

Community Support

Model Code Portability and Future Architectures

Multiple Levels of Parallelism

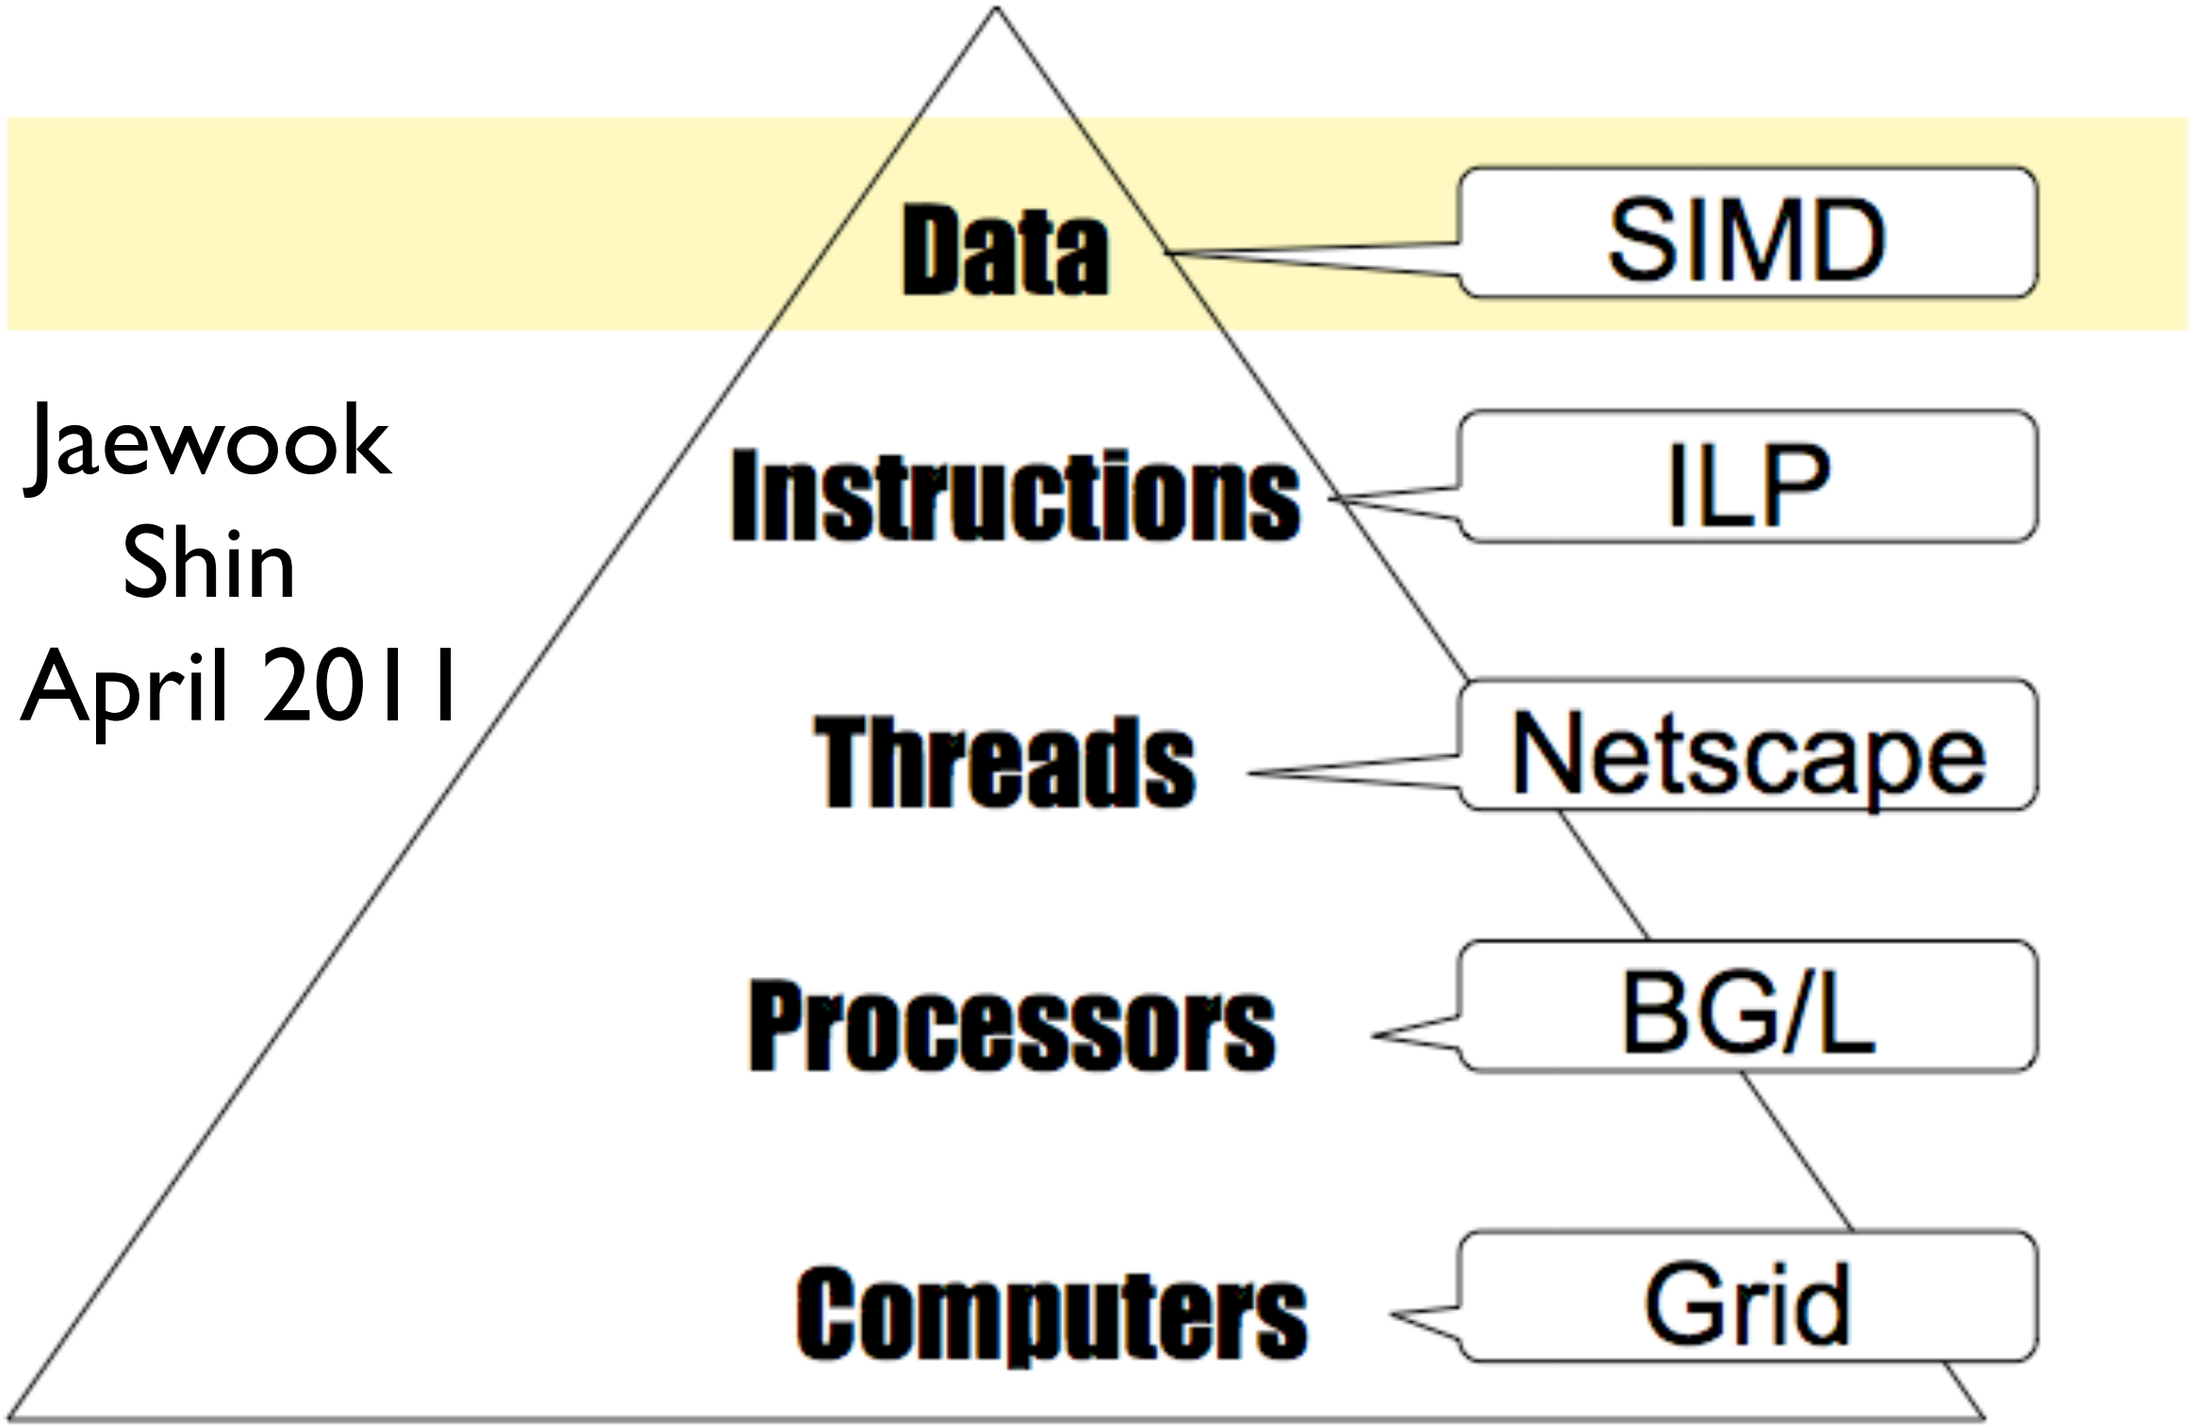

Jaewook Shin

1 April 2011

CMMAP GCRM Development Status?

- After January meeting we convened a telecon to review the parallel implementation strategy for the 'GCRM'
- Main issues relate to the future architectures evolved strong SIMD characteristics and language dependencies

Community Support

Data Interoperability

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.

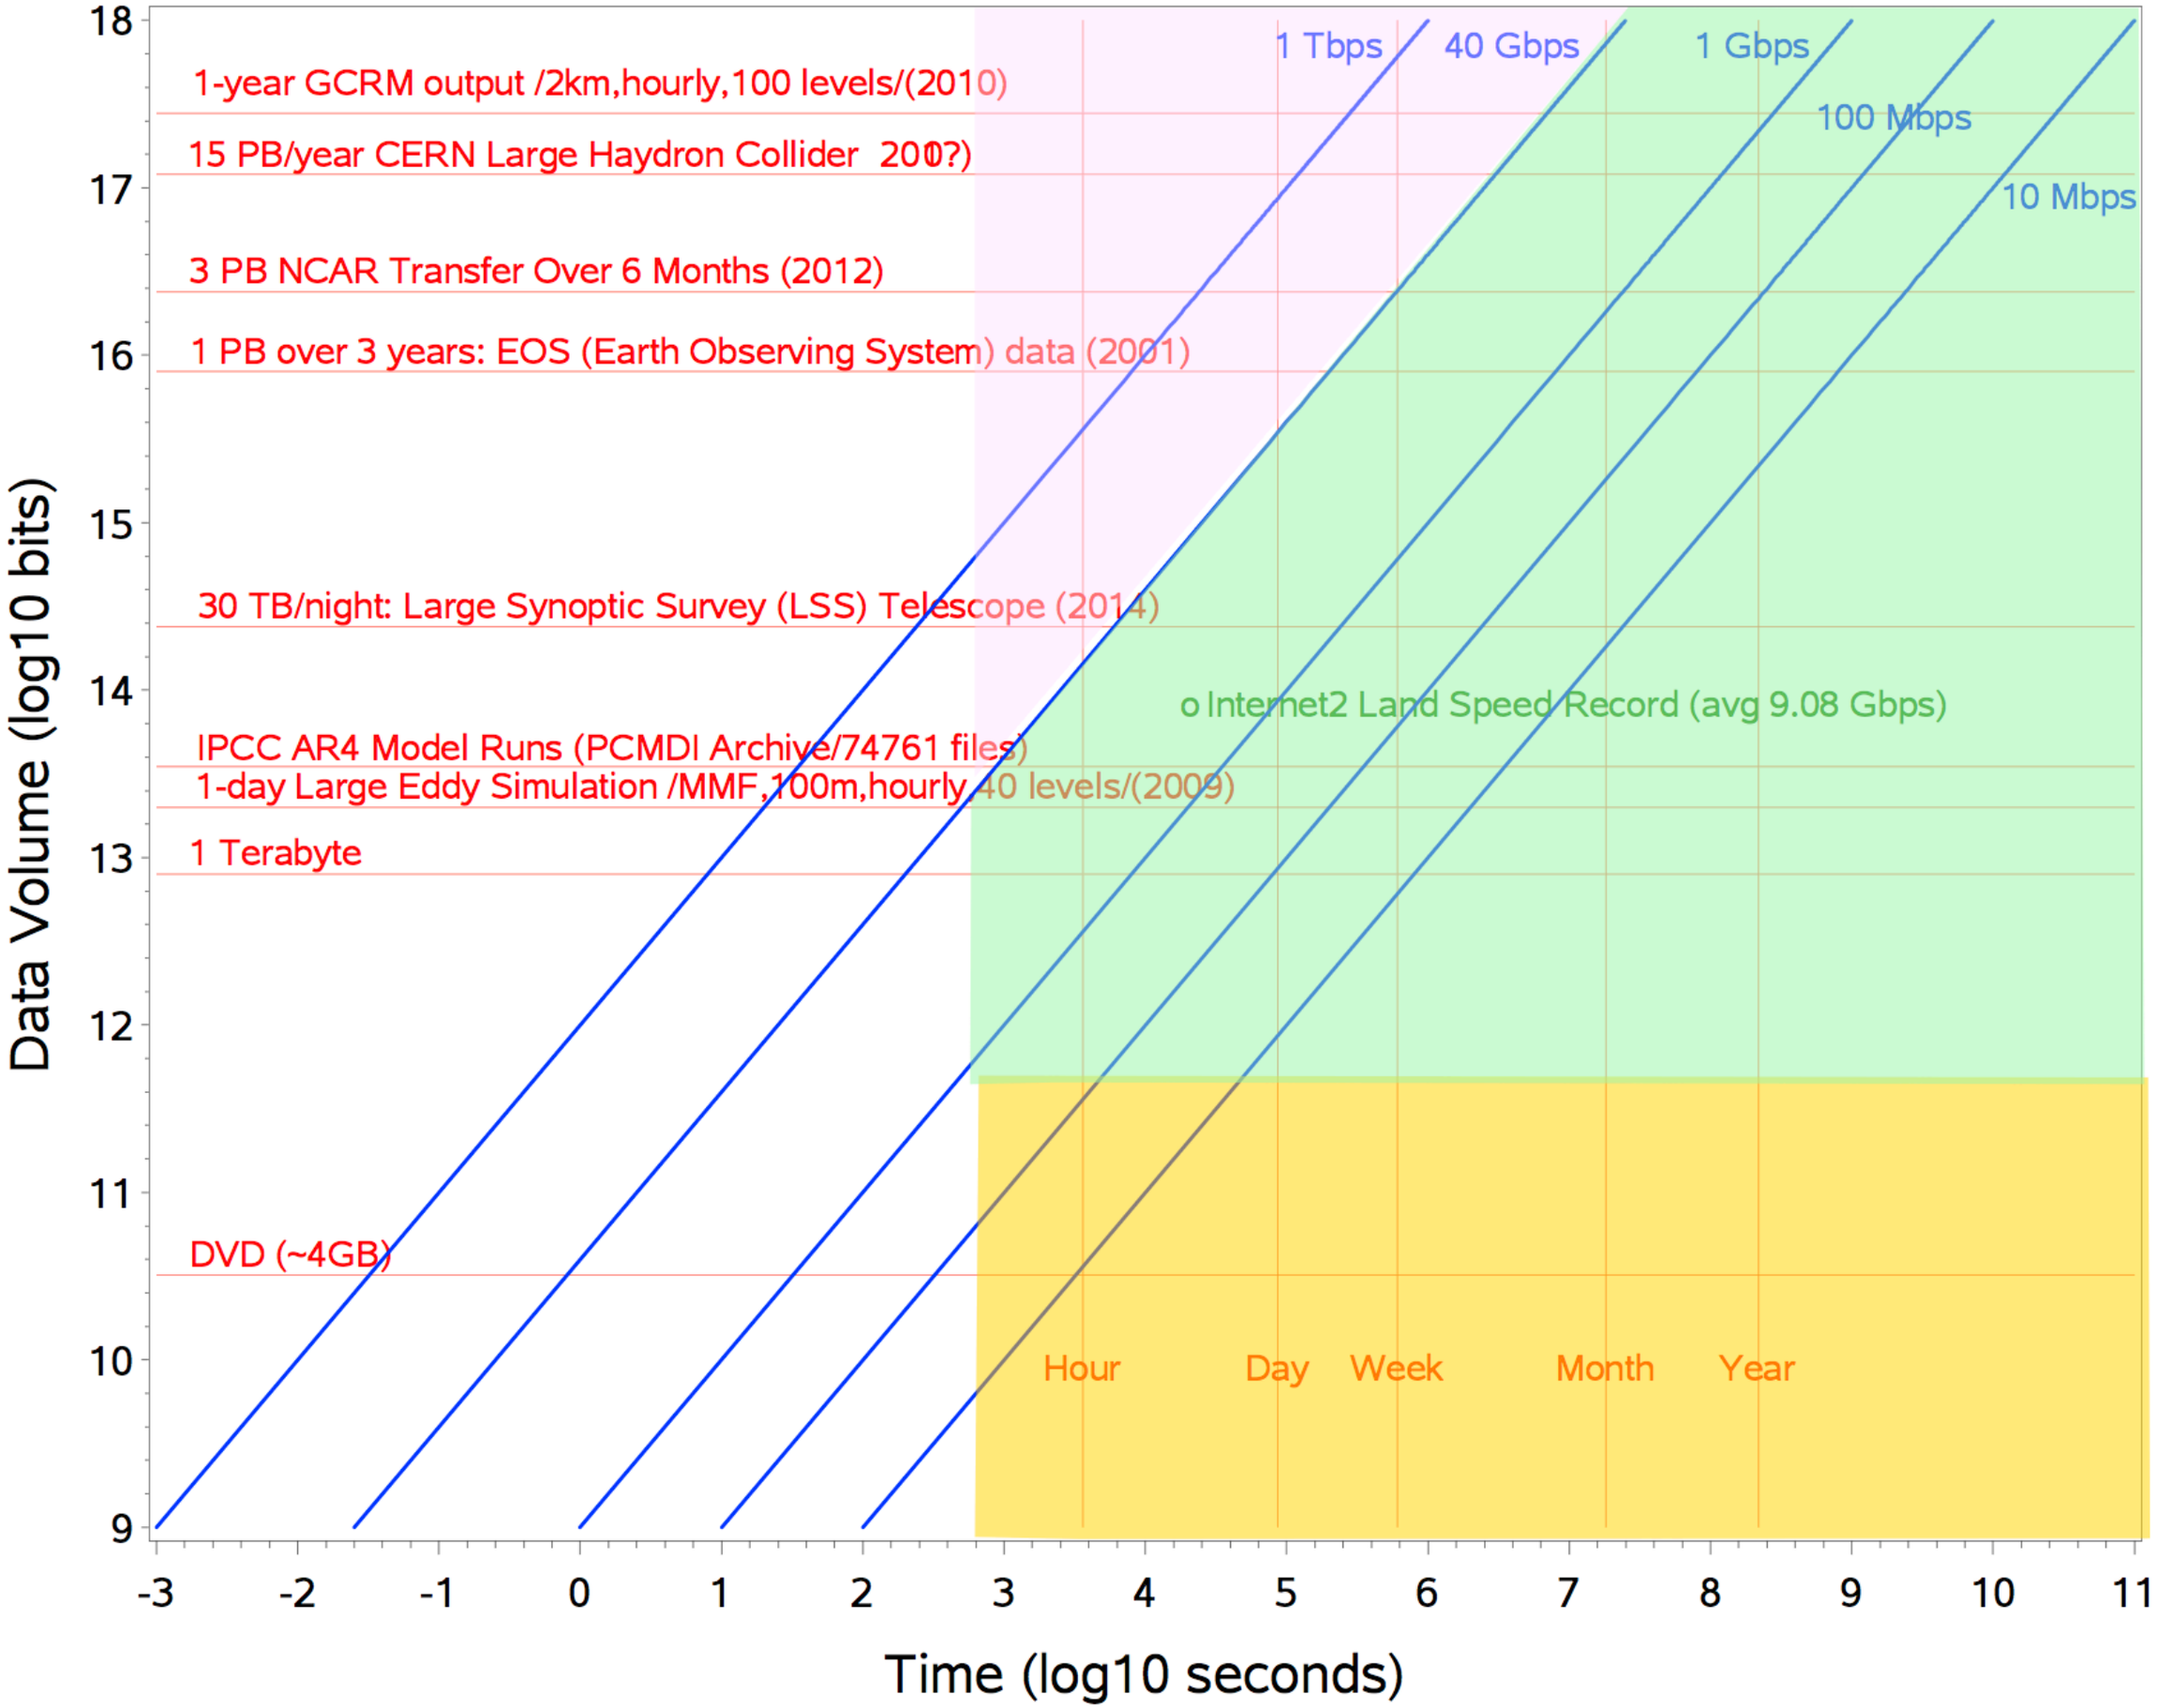
1 Tbps
40 Gbps
1 Gbps
100 Mbps
10 Mbps
1-year GCRM output /2km,hourly,100 levels/(2010)
15 PB/year CERN Large Haydron Collider 200?)
3 PB NCAR Transfer Over 6 Months (2012)
1 PB over 3 years: EOS (Earth Observing System) data (2001)
30 TB/night: Large Synoptic Survey (LSS) Telescope (2014)
o Internet2 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

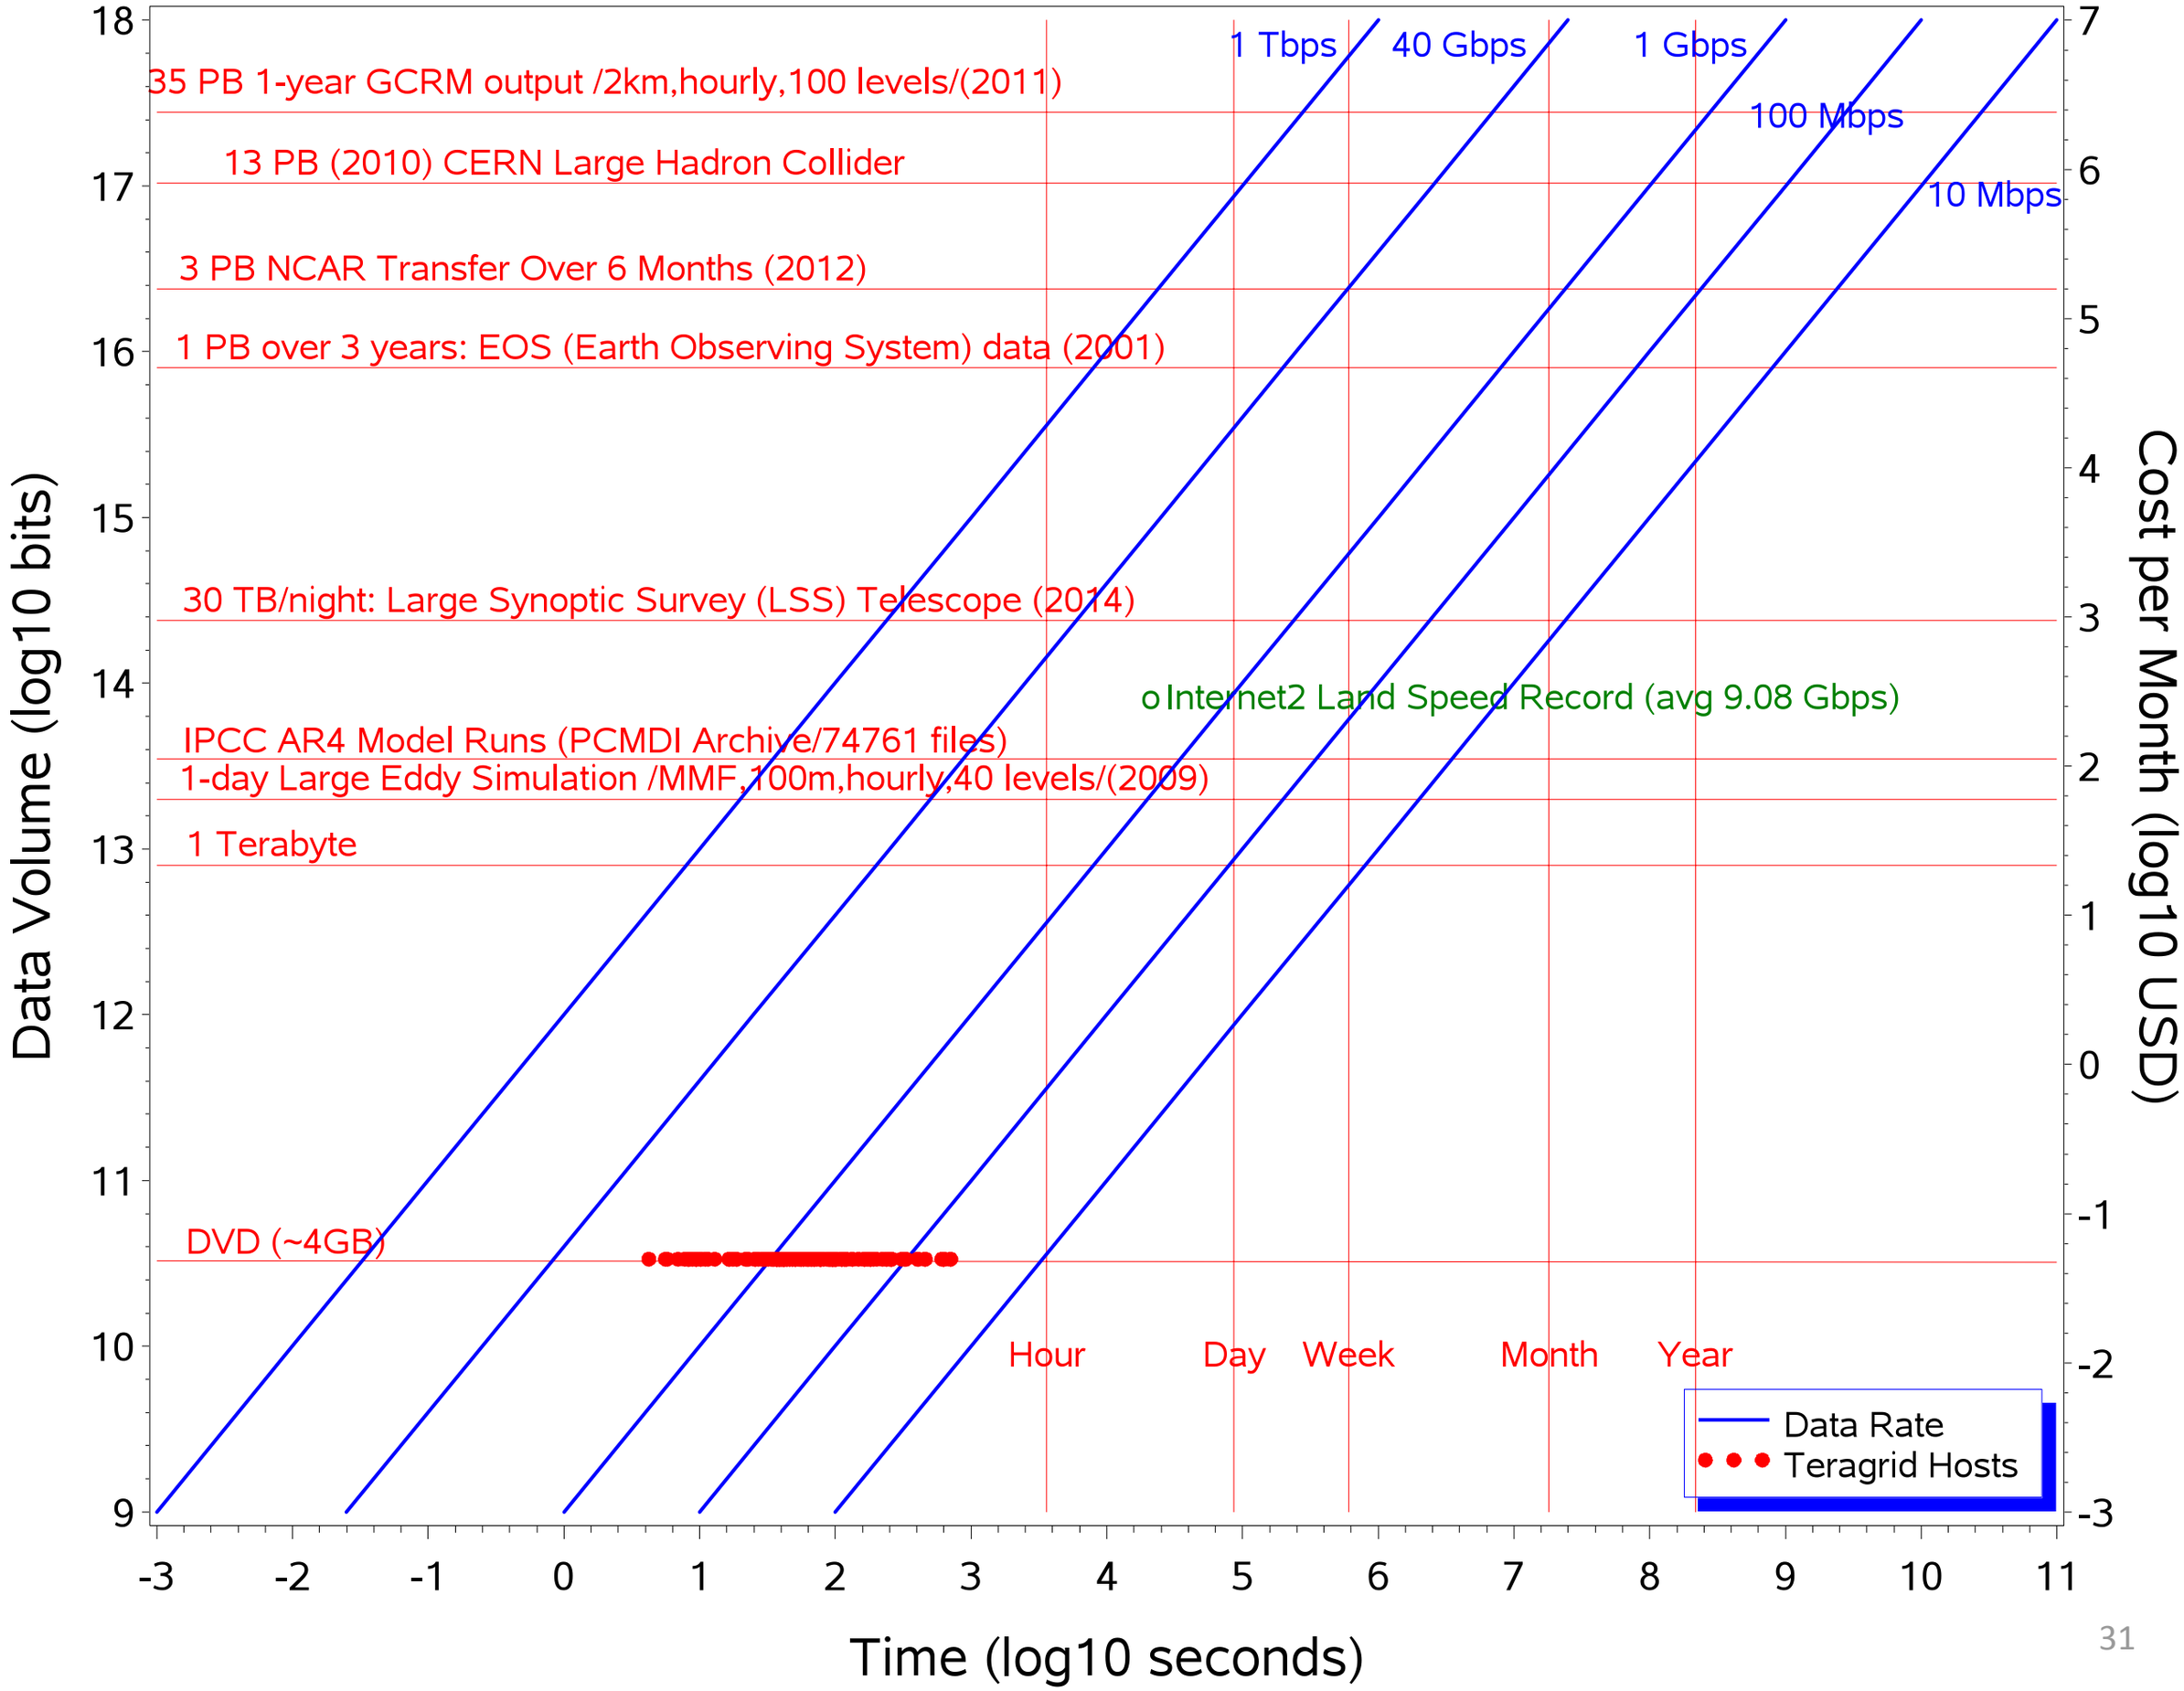
35 PB 1-year GCRM output /2km,hourly,100 levels/(2011)
13 PB (2010) CERN Large Hadron Collider
3 PB NCAR Transfer Over 6 Months (2012)
1 PB over 3 years: EOS (Earth Observing System) data (2001)
30 TB/night: Large Synoptic Survey (LSS) Telescope (2014)
IPCC AR4 Model Runs (PCMDI Archive/74761 files)
1-day Large Eddy Simulation /MMF,100m,hourly,40 levels/(2009)
1 Terabyte
DVD (~4GB)
1 Tbps
40 Gbps
1 Gbps
100 Mbps
10 Mbps
Internet2 Land Speed Record (avg 9.08 Gbps)
Hour
Day
Week
Month
Year
Data Rate
Teragrid Hosts
Data Volume (log10 bits)
Cost per Month (log10 USD)
Time (log10 seconds)
18
17
16
15
14
13
12
11
10
9
7
6
5
4
3
2
1
0
-1
-2
-3
-3
-2
-1
0
1
2
3
4
5
6
7
8
9
10
11

INTERNET2

July 11th 2011 – Summer Joint Techs

Eric Boyd, Internet2 Deputy CTO

Jason Zurawski, Internet2 Research Liaison

NSF MRI-R2: DYnamic NEtwork System (DYNES, NSF #0958998)

DYNES Projected Topology (July 2011)

- Based on applications accepted
- Showing peerings to other Dynamic Circuit Networks (DCN)

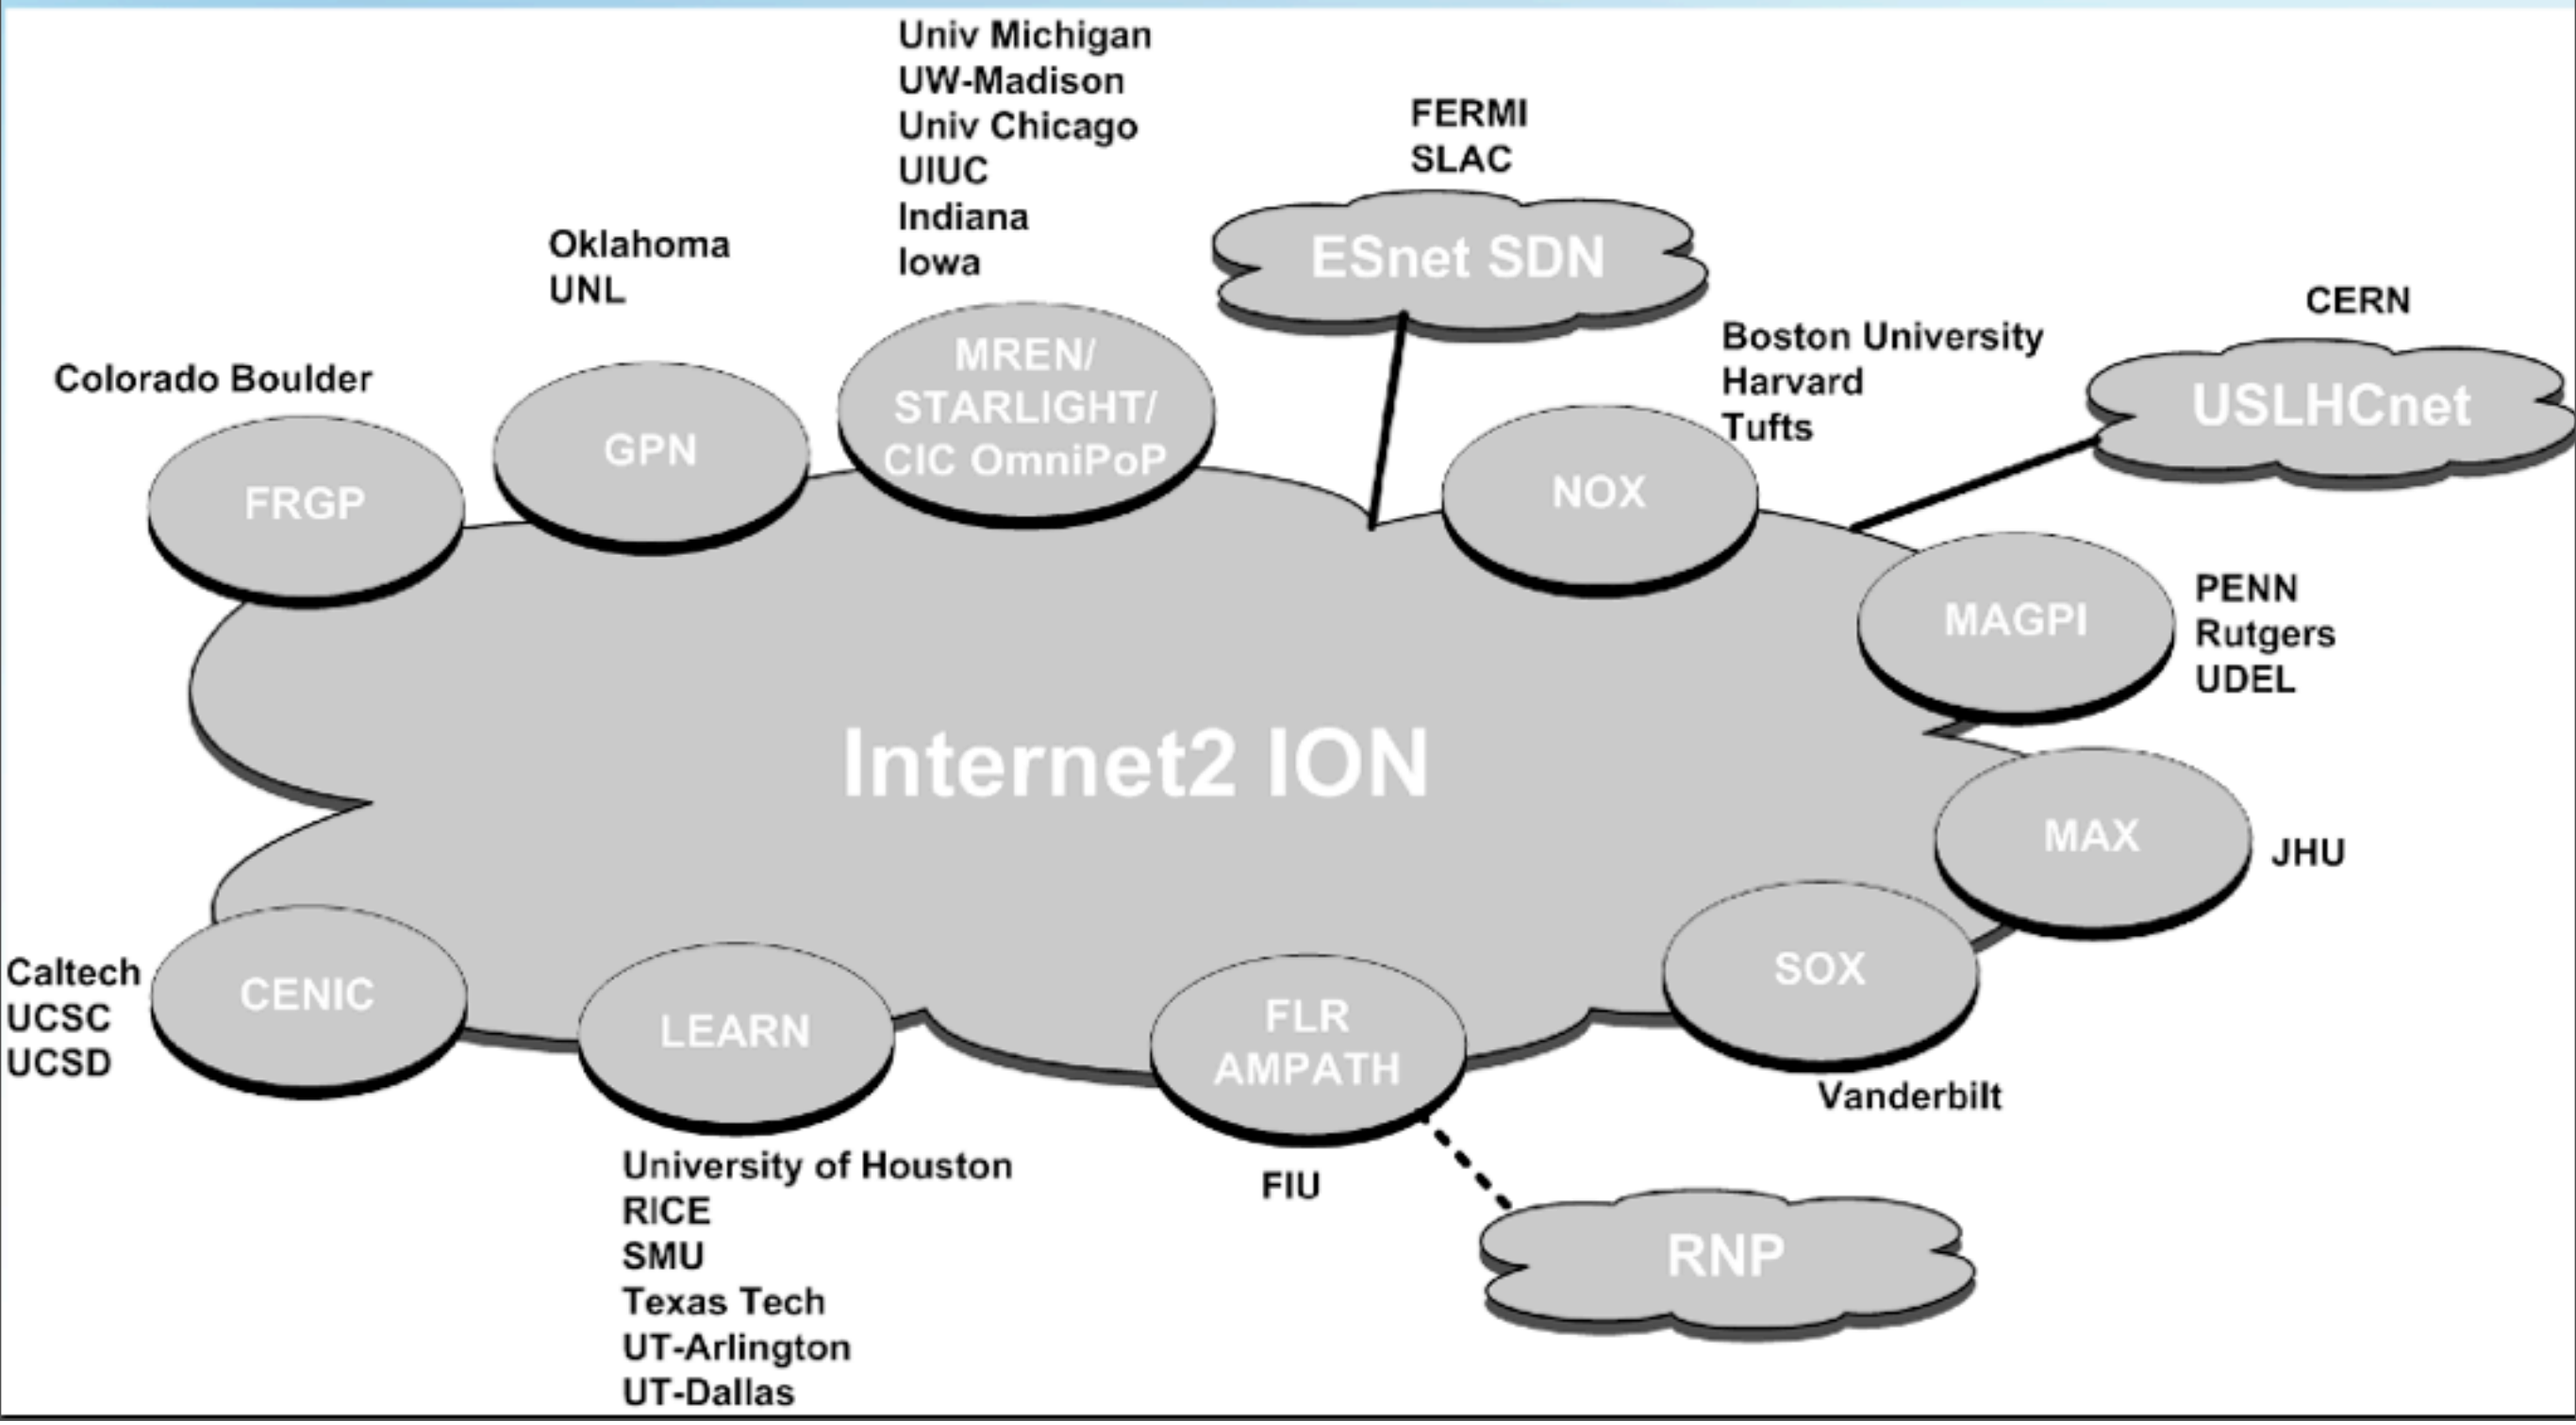

Greetings;

We are writing to offer a brief update on the DYNES process, to build on some of the topics and discussion that resulted from the recent meeting at Joint Techs in Fairbanks AK. The DYNES BoF slides are available in this location for future reference:

http://events.internet2.edu/2011/jt-uaf/agenda.cfm?go=session&id=10001881&event=1151

As a reminder, we have partitioned the DYNES deployment into 3 groups:

- Group A - Completing end of July 2011
- Group B - Starting end of July 2011, completing end of August 2011
- Group C - Starting beginning of Sept 2011, completing October 2011

The groups are listed at this location: http://www.internet2.edu/ion/docs/20110412-DYNES-Groups.pdf, the University of California, San Diego is in Group C.

The process for "Group Deployment" involves a warm up phone call with the end sites and their regional partner to describe the ordering and installation process, along with additional calls (as needed) to gather specific requirements in each site regarding equipment and configuration. We are asking that if you are planning to make any significant changes to your network in expectation of the DYNES equipment (e.g. making special purchases, or thinking of purchasing DYNES equipment on your own), please let us know as soon as possible.

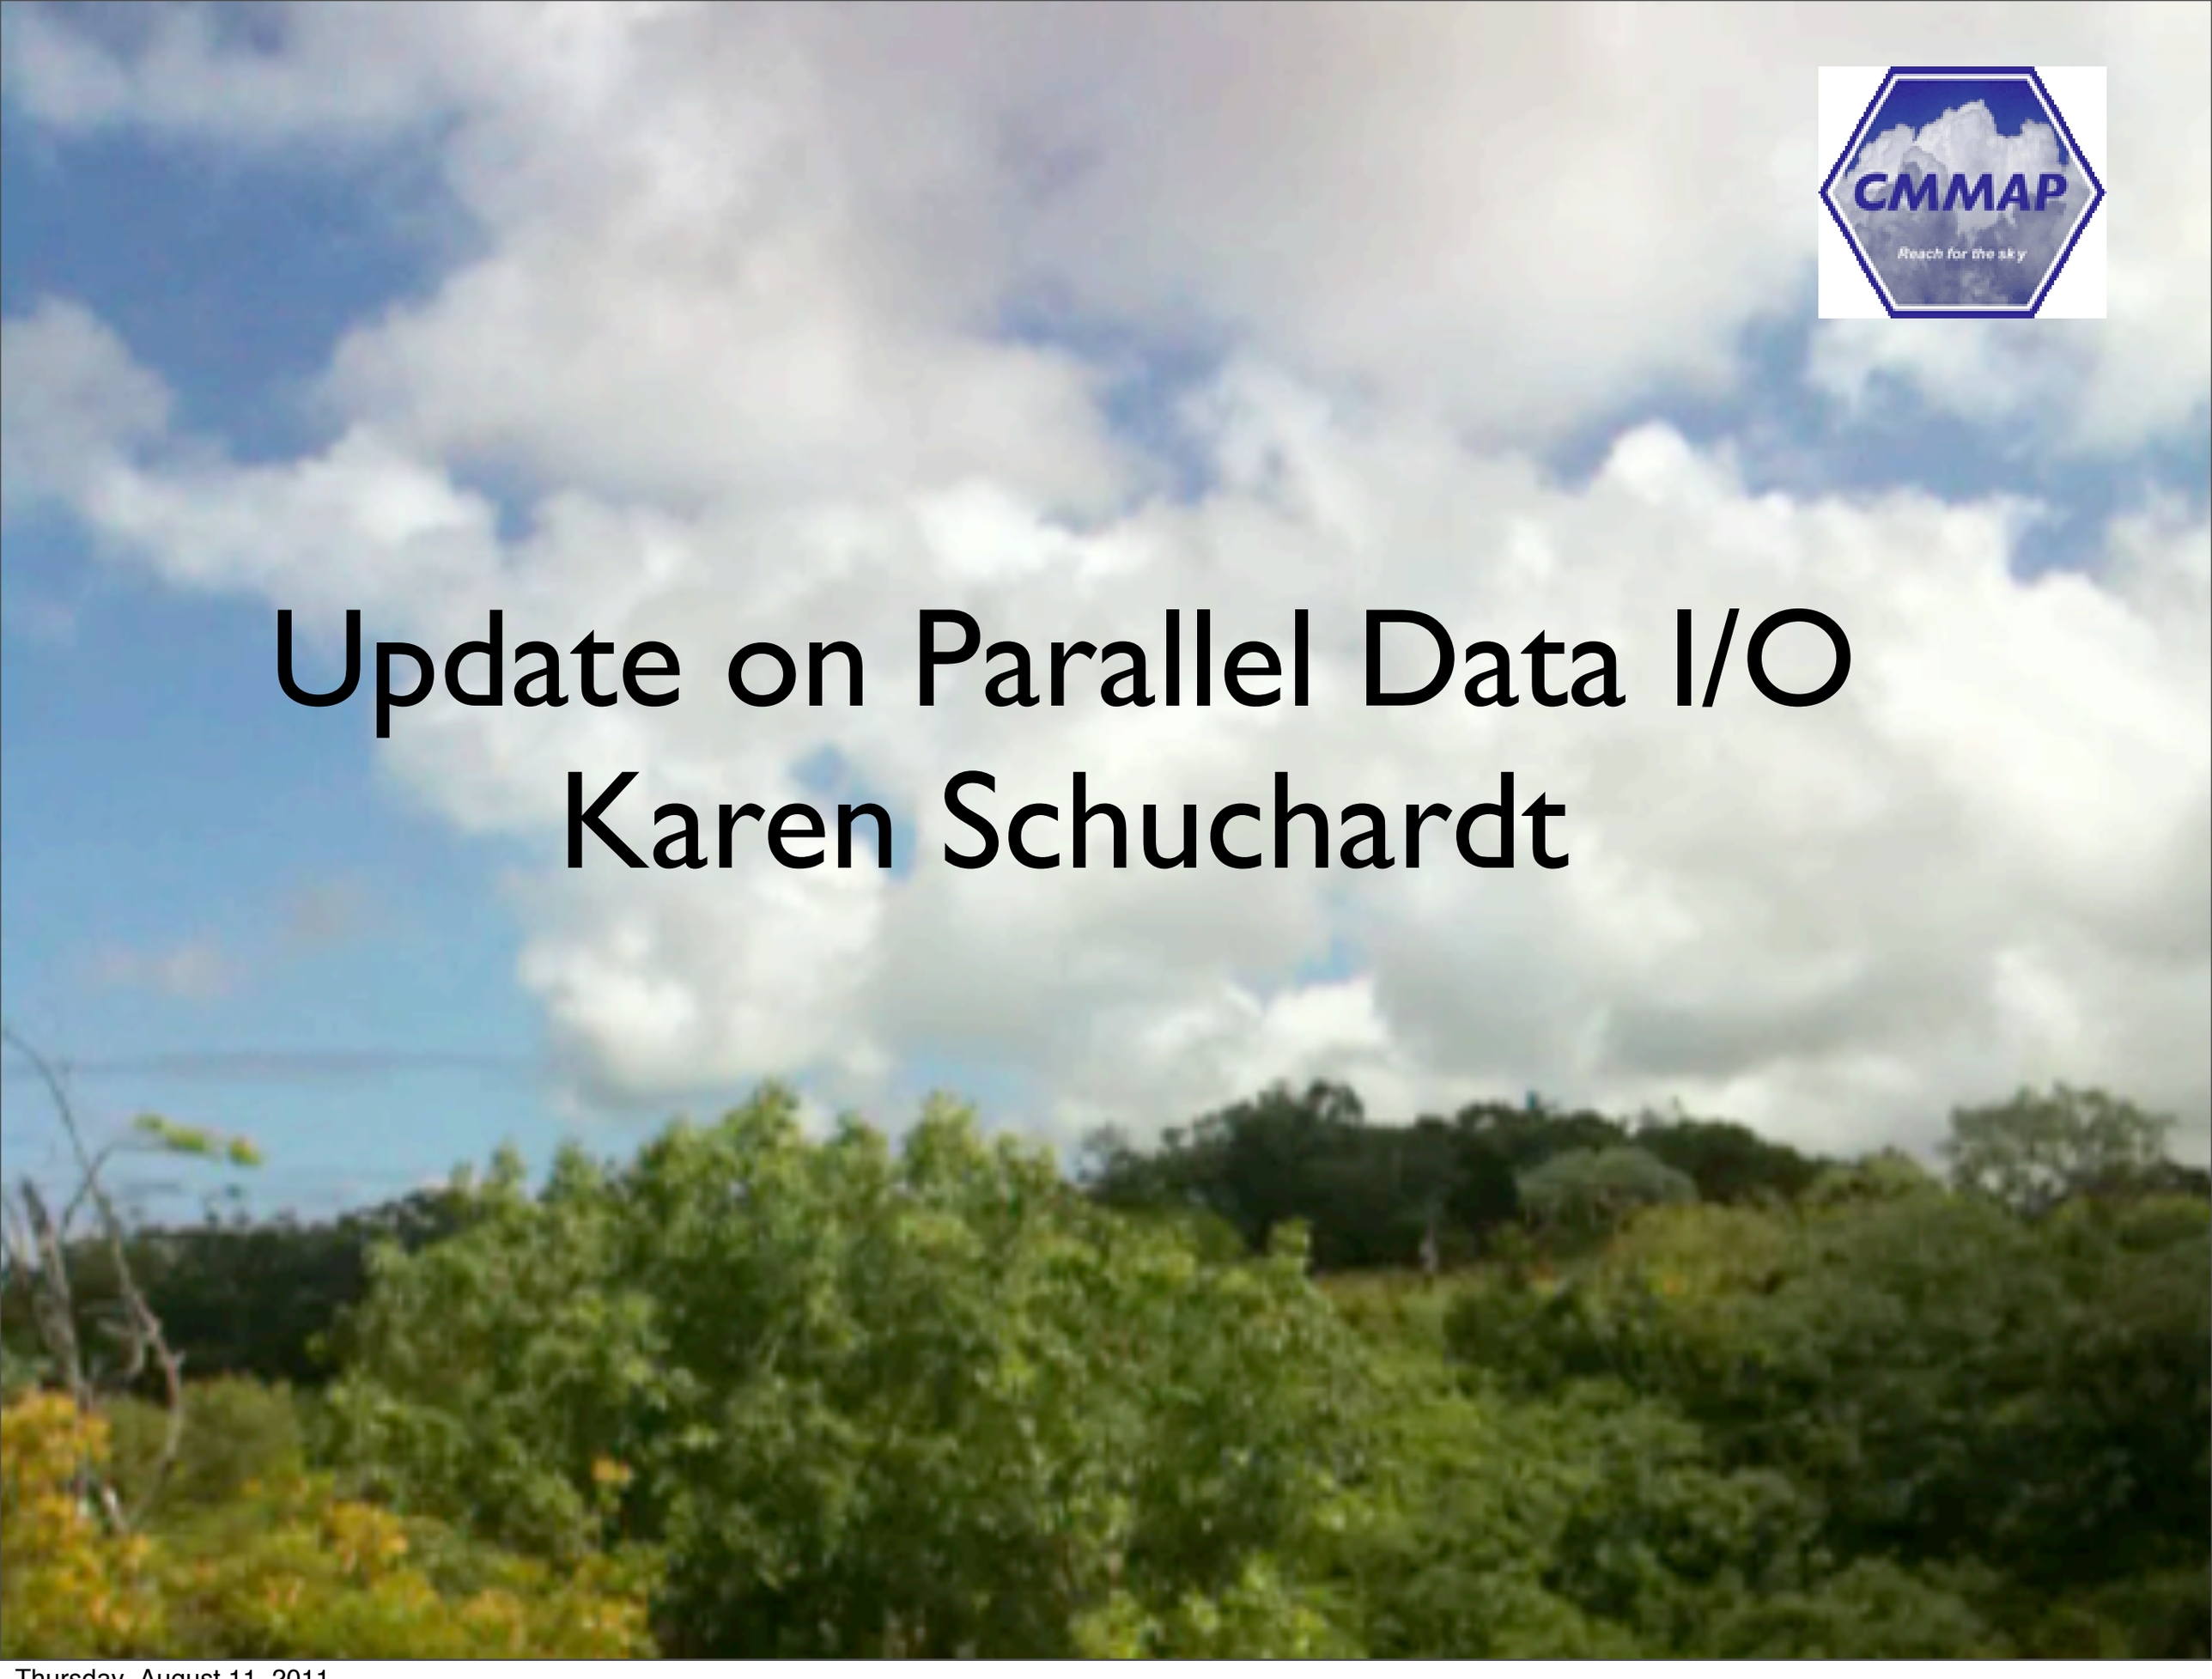

Update on Parallel Data I/O

Karen Schuchardt

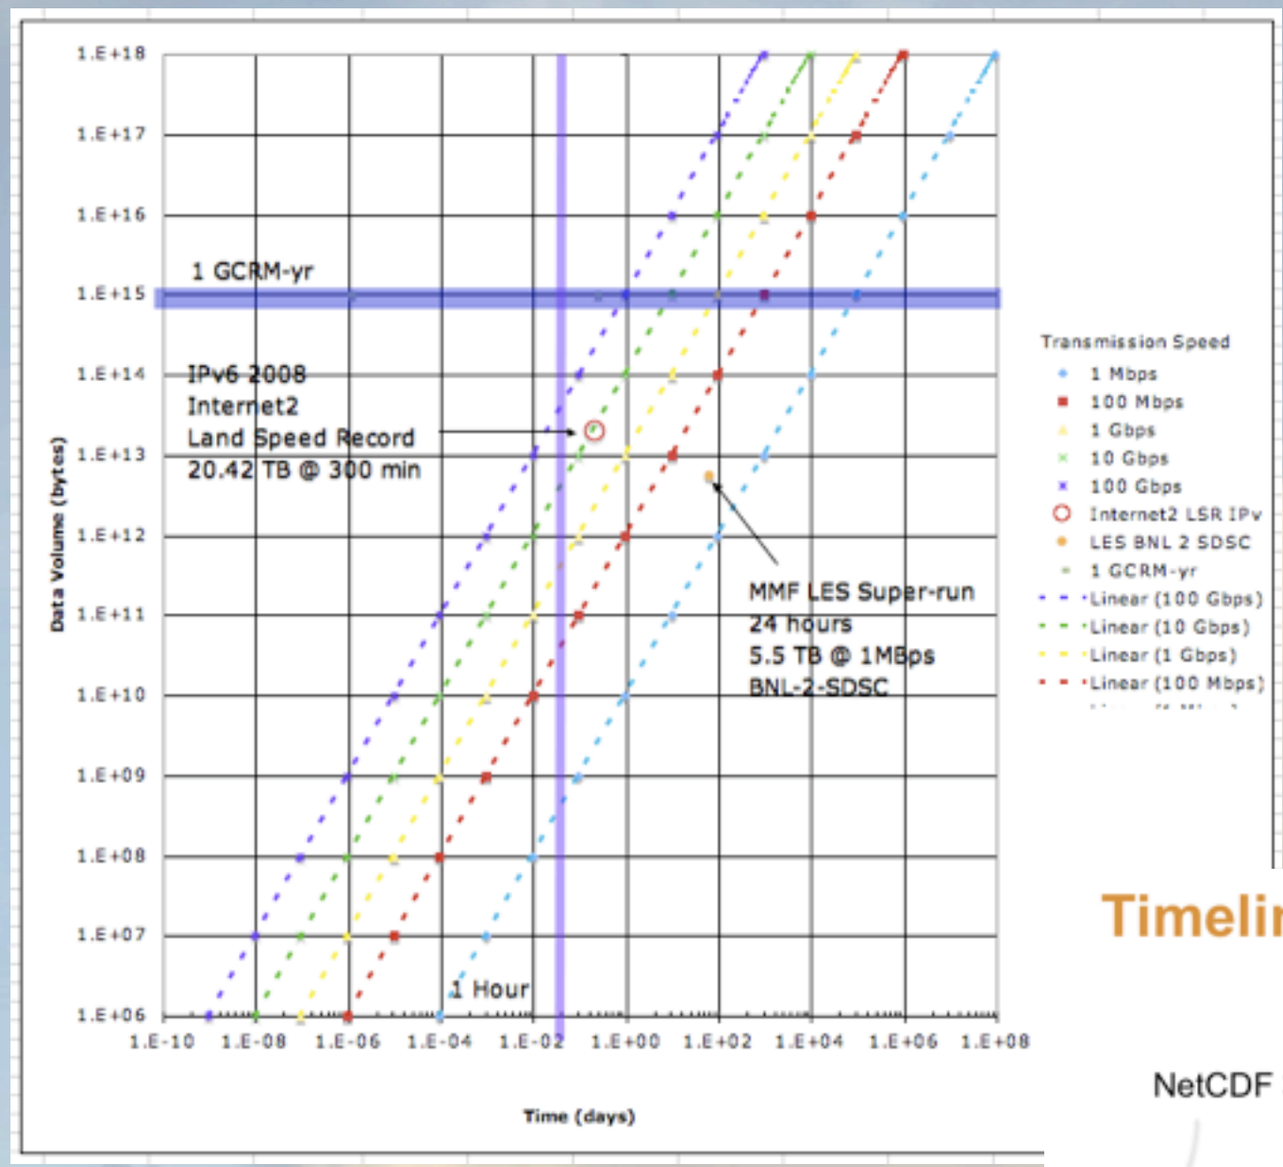

Internet Data Transfer Capacity

Parallel Input/ Output Technology Progress

Timeline

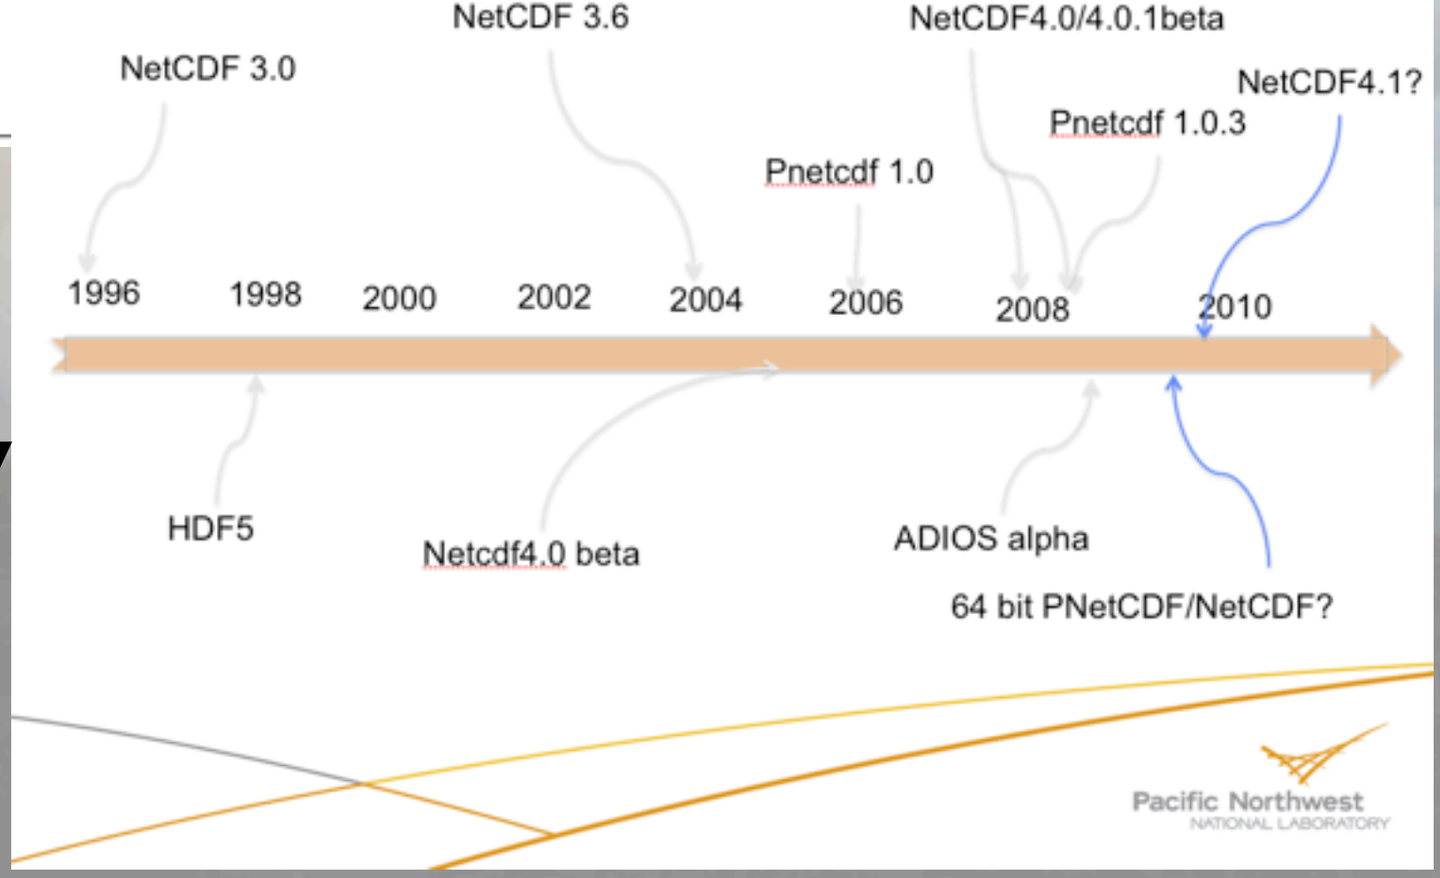

Pacific Northwest
NATIONAL LABORATORY

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

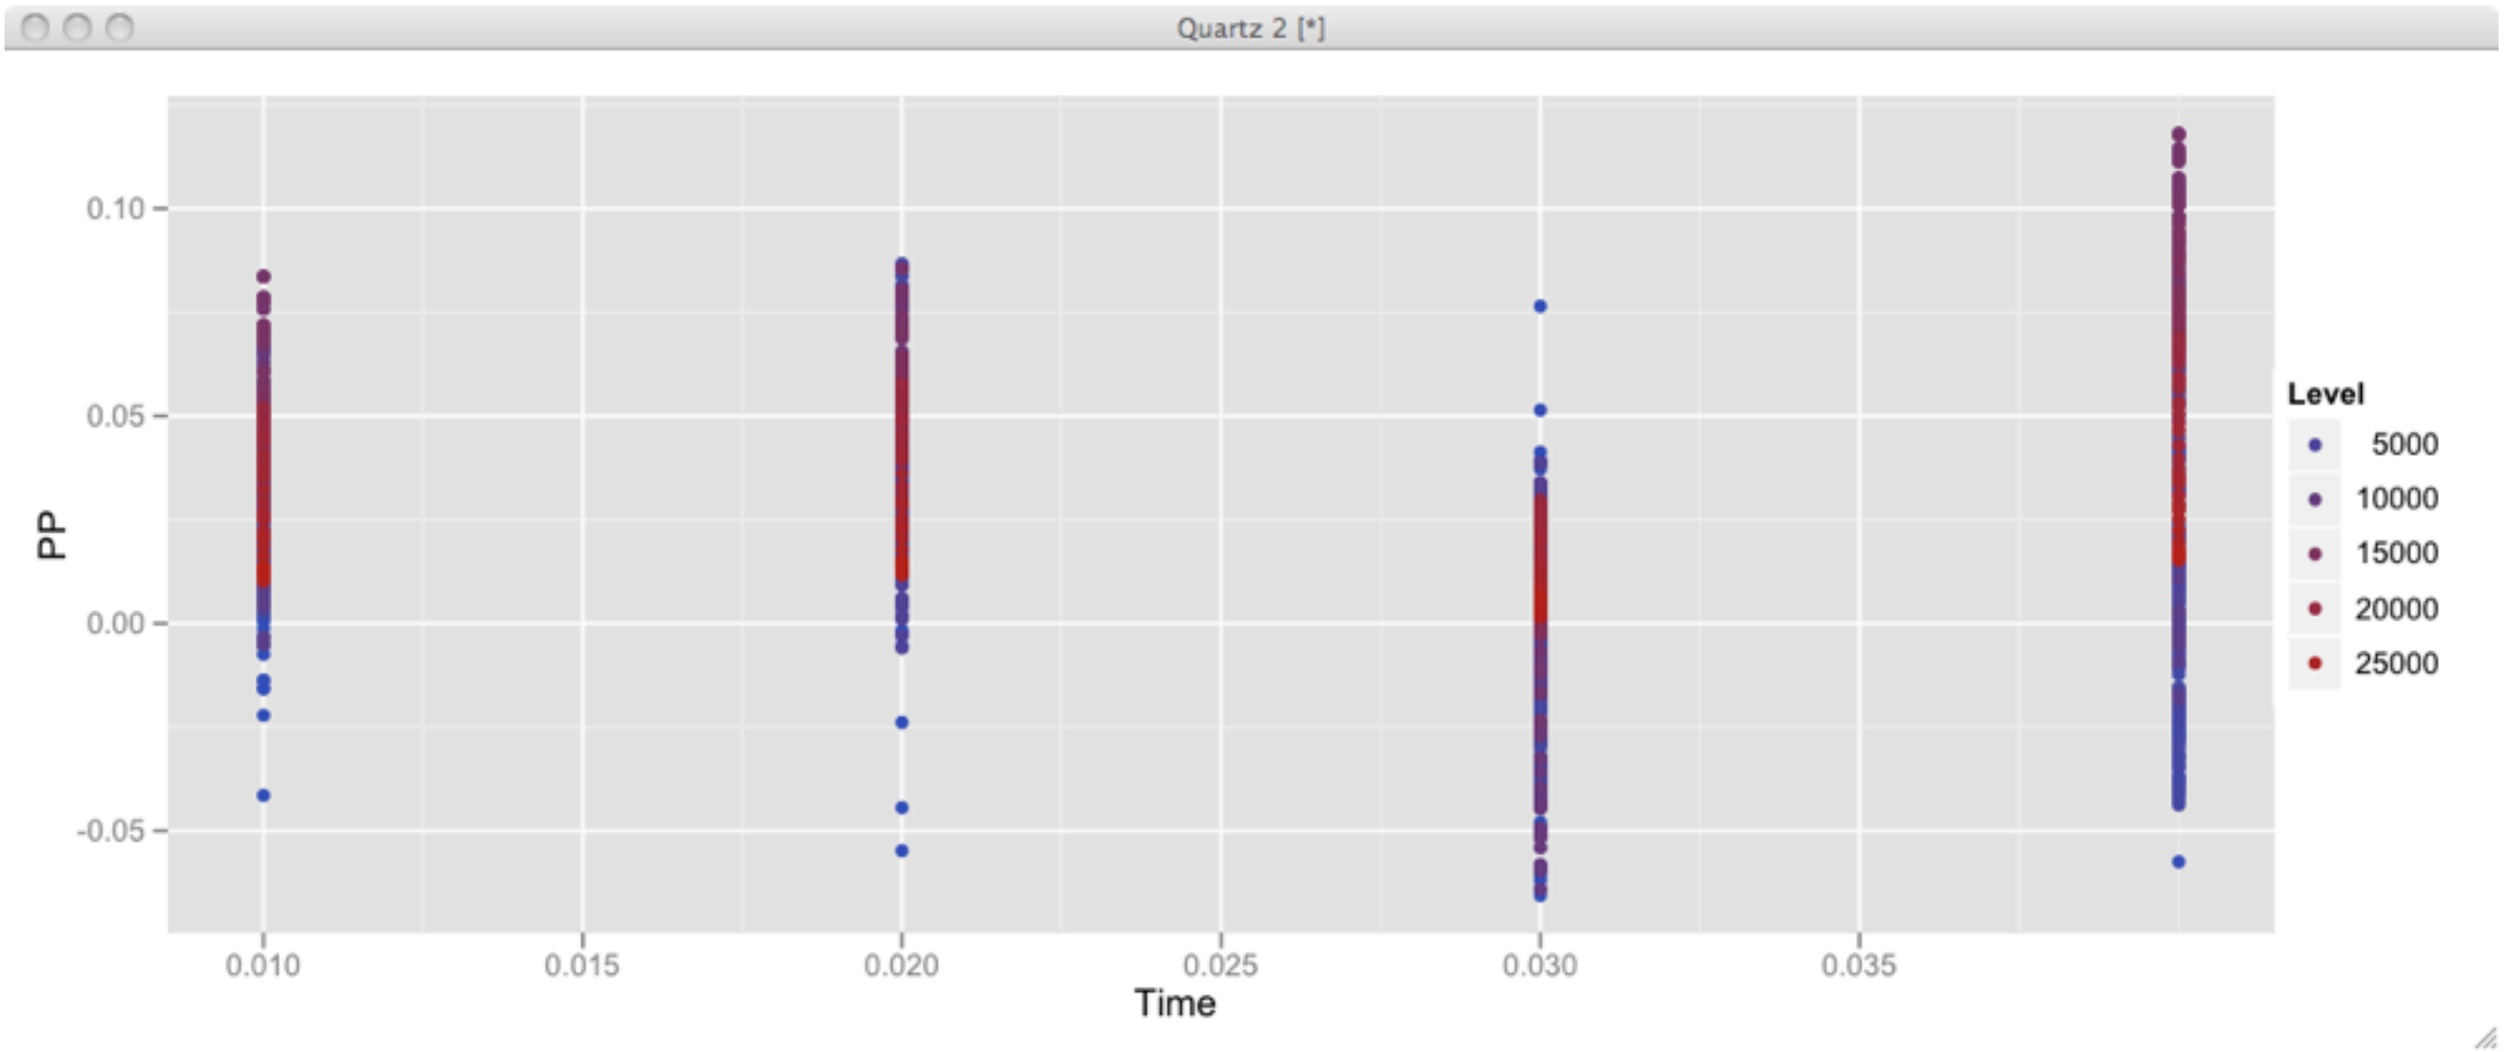

Triton Results

- 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].

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, August 11, 2011

3D visualization of geodesic data

VisIT

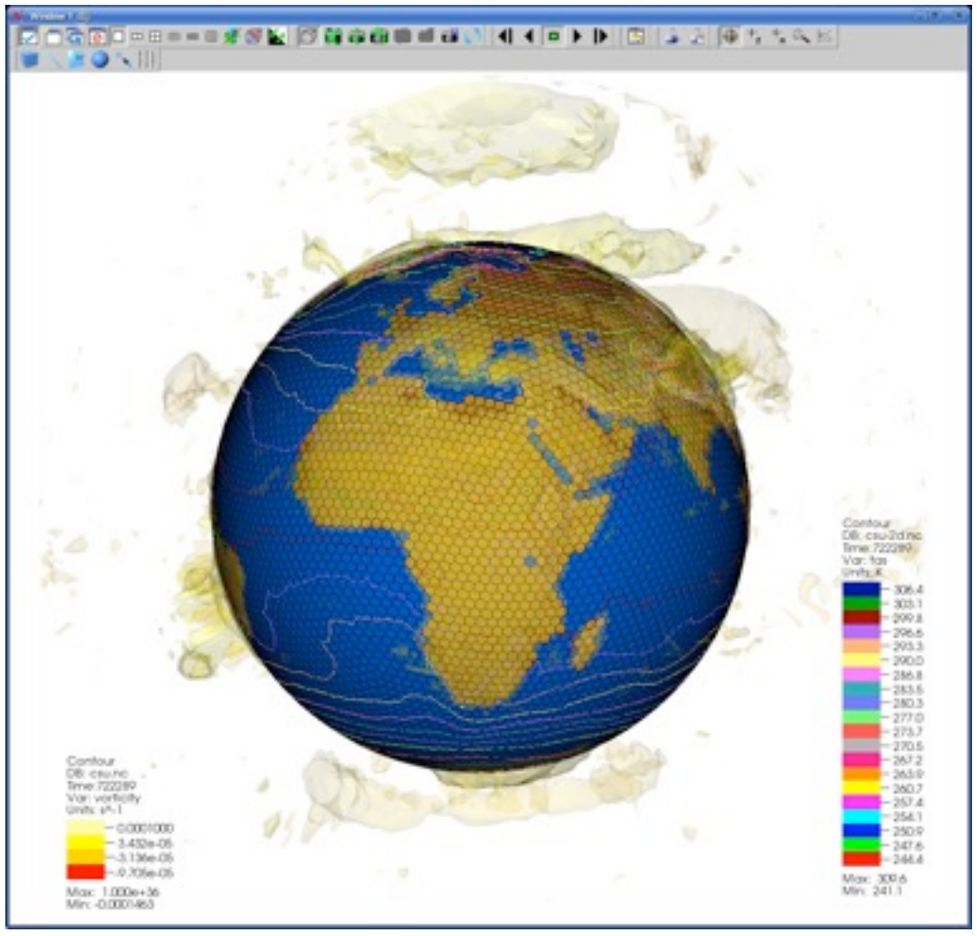

3D isocontours of vorticity.

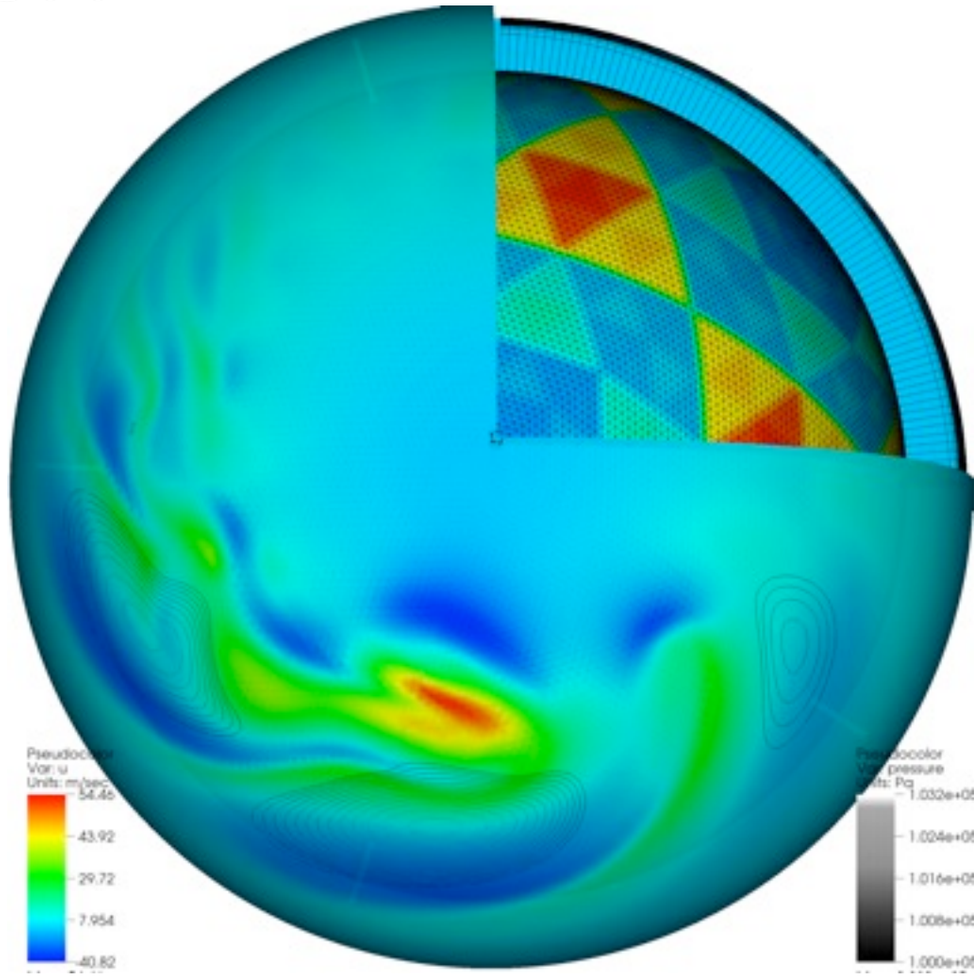

Composite plot of multiple mesh types and variables in the geodesic grid. Cell area (2D cell-centered data) and wind velocity (3D corner-centered on layers) data is shown by pseudocolor plots. Pressure (3D cell-centered on layers) is shown by contour lines.

Plots and movies courtesy of Prabhat (lbnl)