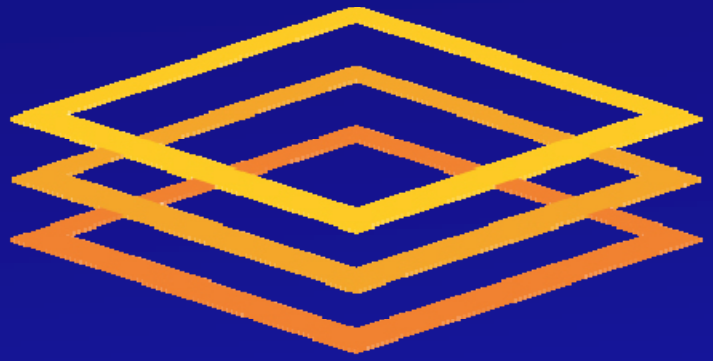

Open Science Grids

Analysis of the Use, Value and Upcoming Challenges for the Open Science Grid

Ruth Pordes, OSG Executive Director, Fermilab

Project supported by the Department of Energy Office of Science SciDAC-2 program from the High Energy Physics, Nuclear Physics and Advanced Software and Computing Research programs, and the National Science Foundation Math and Physical Sciences, Office of CyberInfrastructure and Office of International Science and Engineering Directorates.

Analysis of the Use

93 registered Compute resources;

25 Storage registered storage resources

>30% of US ATLAS, US CMS cycles in 2008 ran on resources accessible through OSG.

>2500 people used resources through the OSG in past year.

CDF & D0 getting good usage for monte carlo.

Non-physics use ~10% very cyclic.

VO throughput over past year
>40 facilities contributing

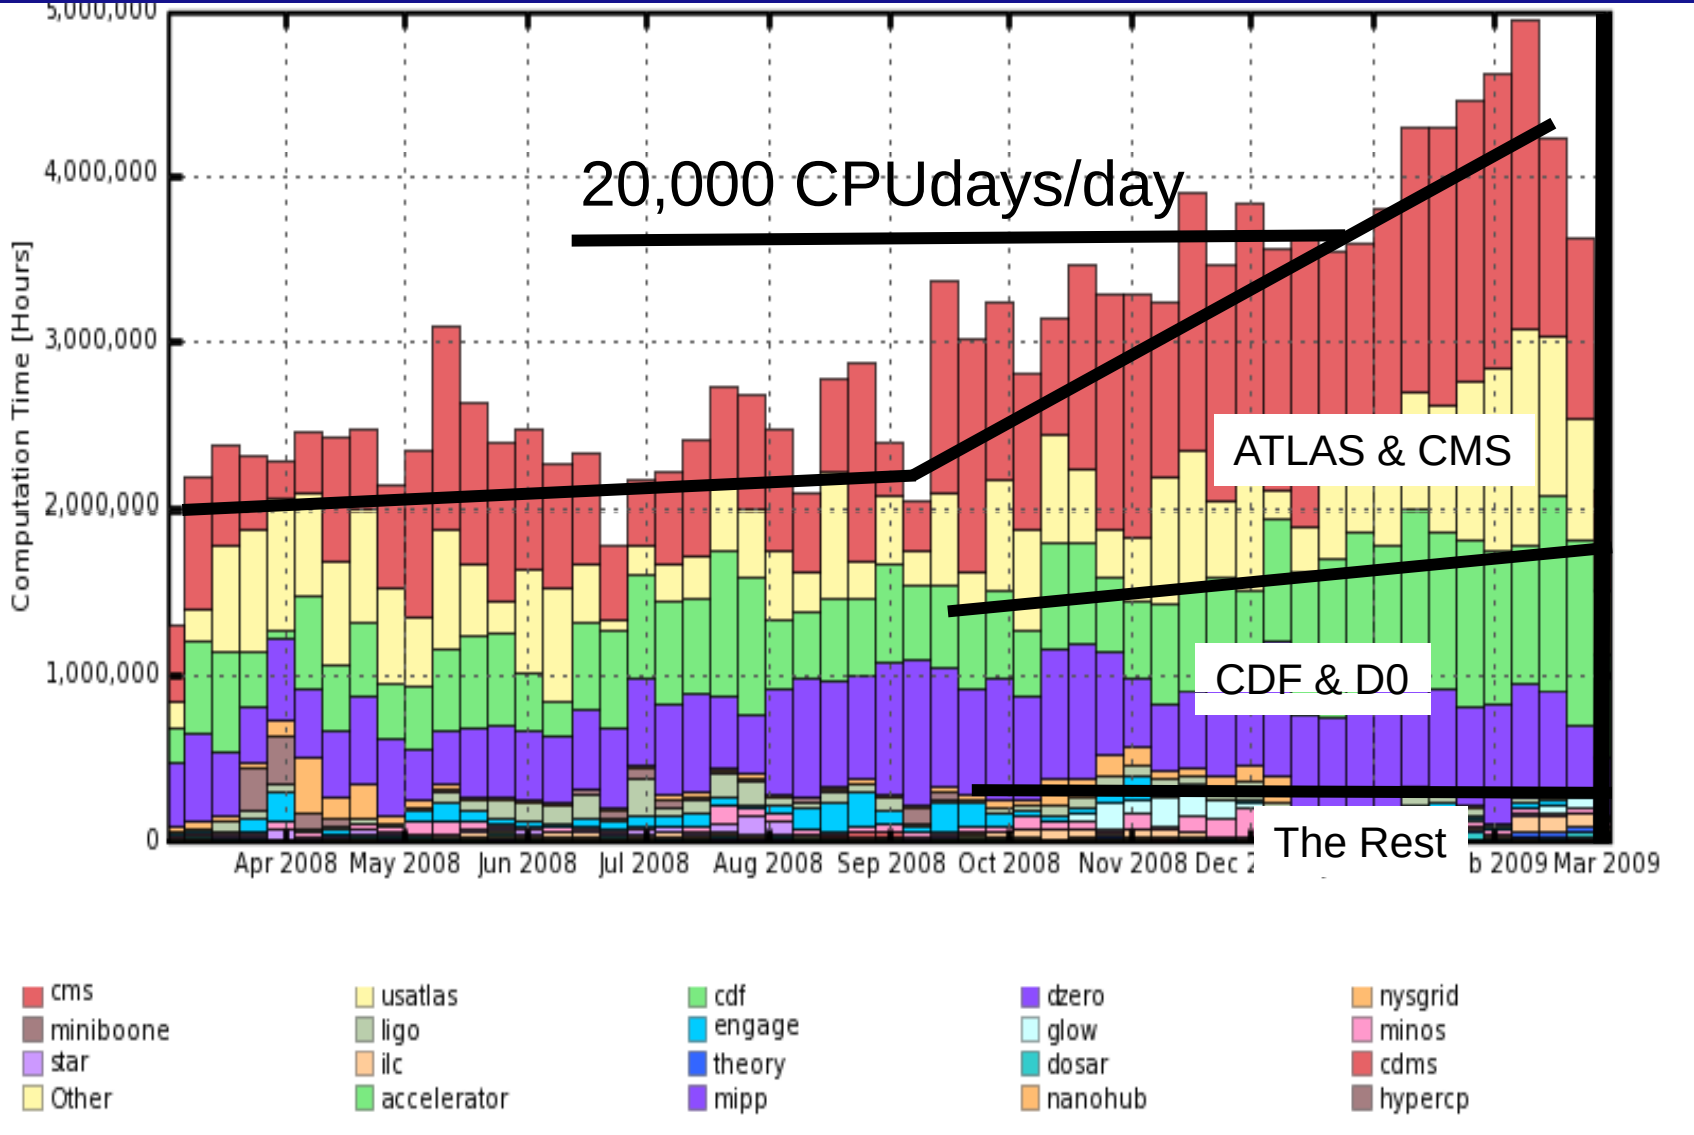

Current Map of OSG Resources (in the US)

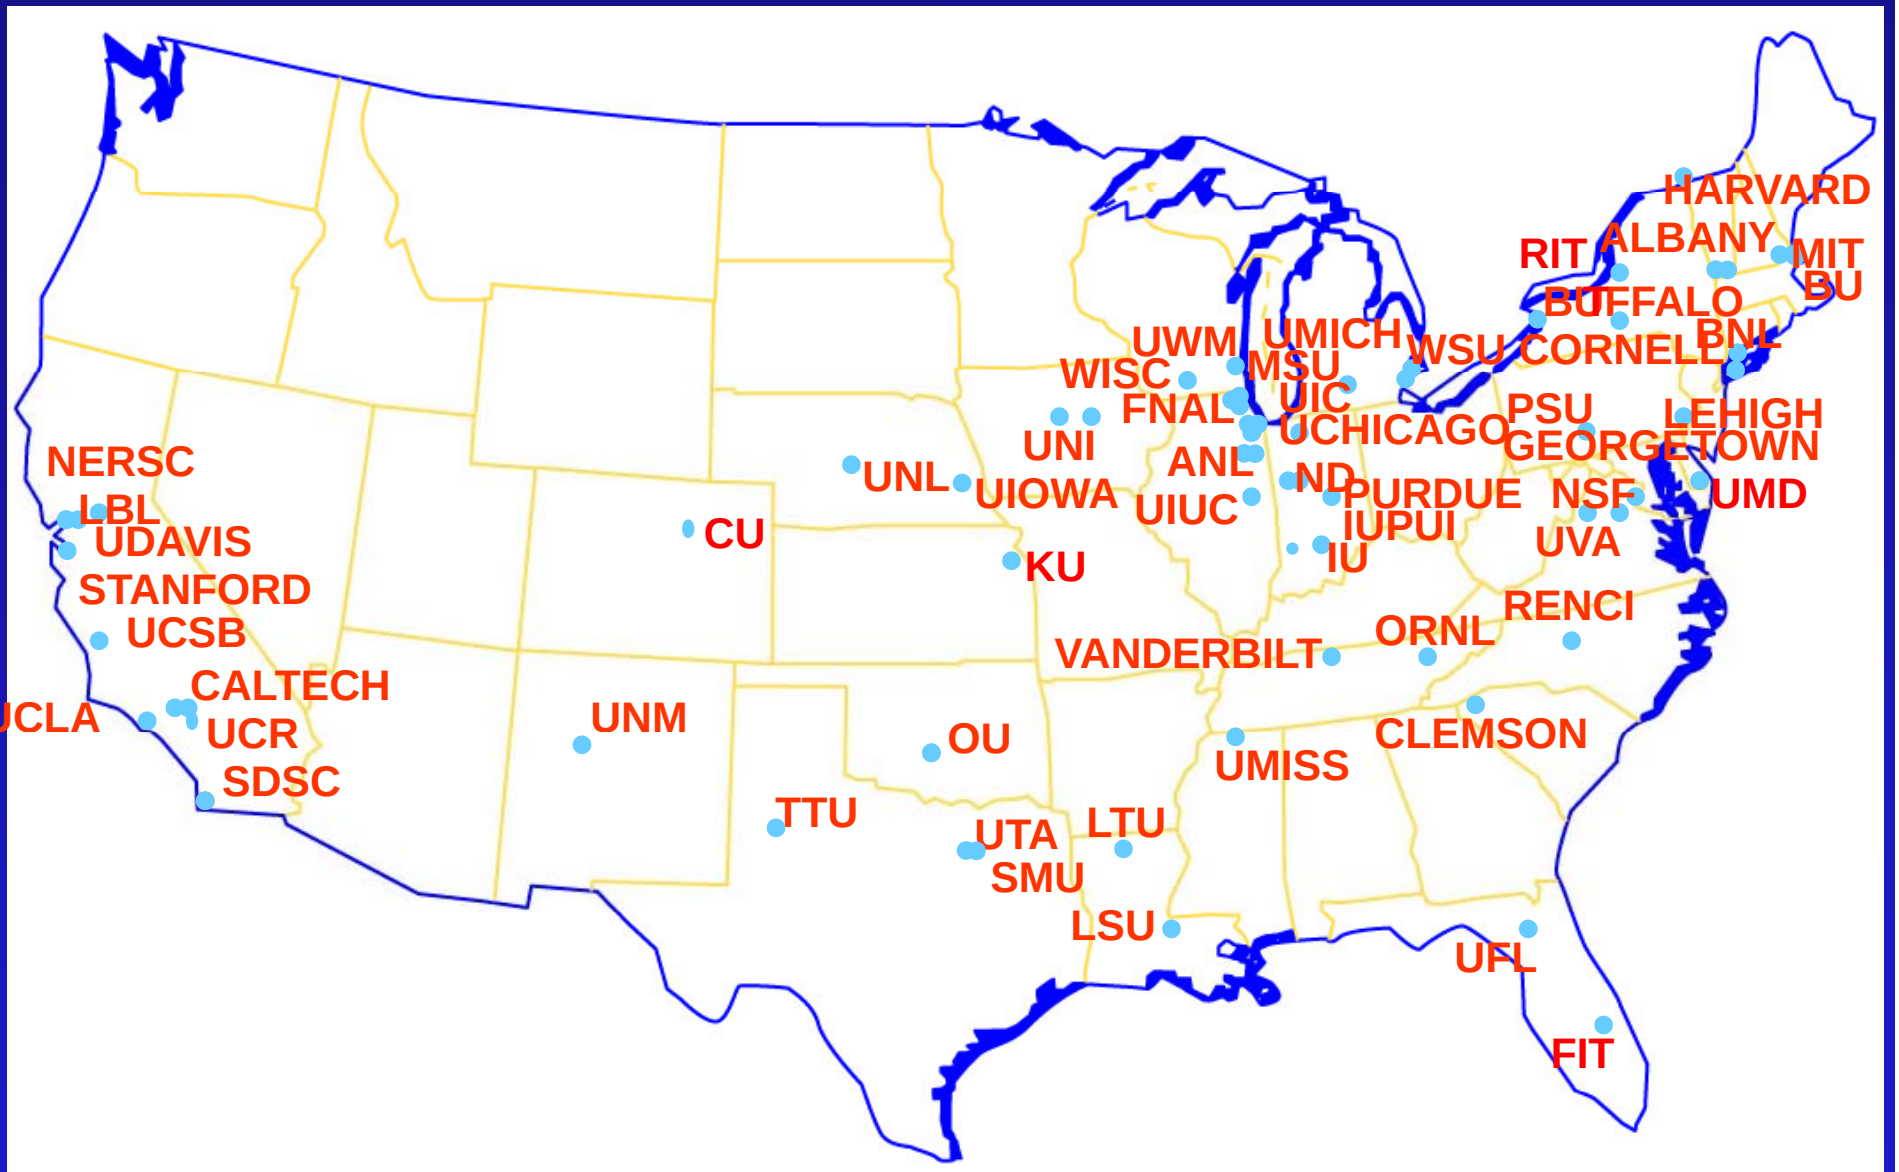

Improvements in Technologies in 2008 Integrated into the Software Stack the Virtual Data Toolkit (now at ~70 components)

Initial use of opportunistic storage for science output (D0)

Early adoption in physics and generalization of “overlay” job scheduling or “pilot” technologies

Resource service validation framework and (security, functionality) probes for monitoring site configurations and services

Site Resource selection and OSG matchmaking services in more general use

Non-Physics VOs & Usage – 2008

Applications smoothly come in, use resources and disappear (e.g. Glow, protein mapping)

At max <~4,000 cpudays/day queue times increased.

Total shows significant “cycles of use”.

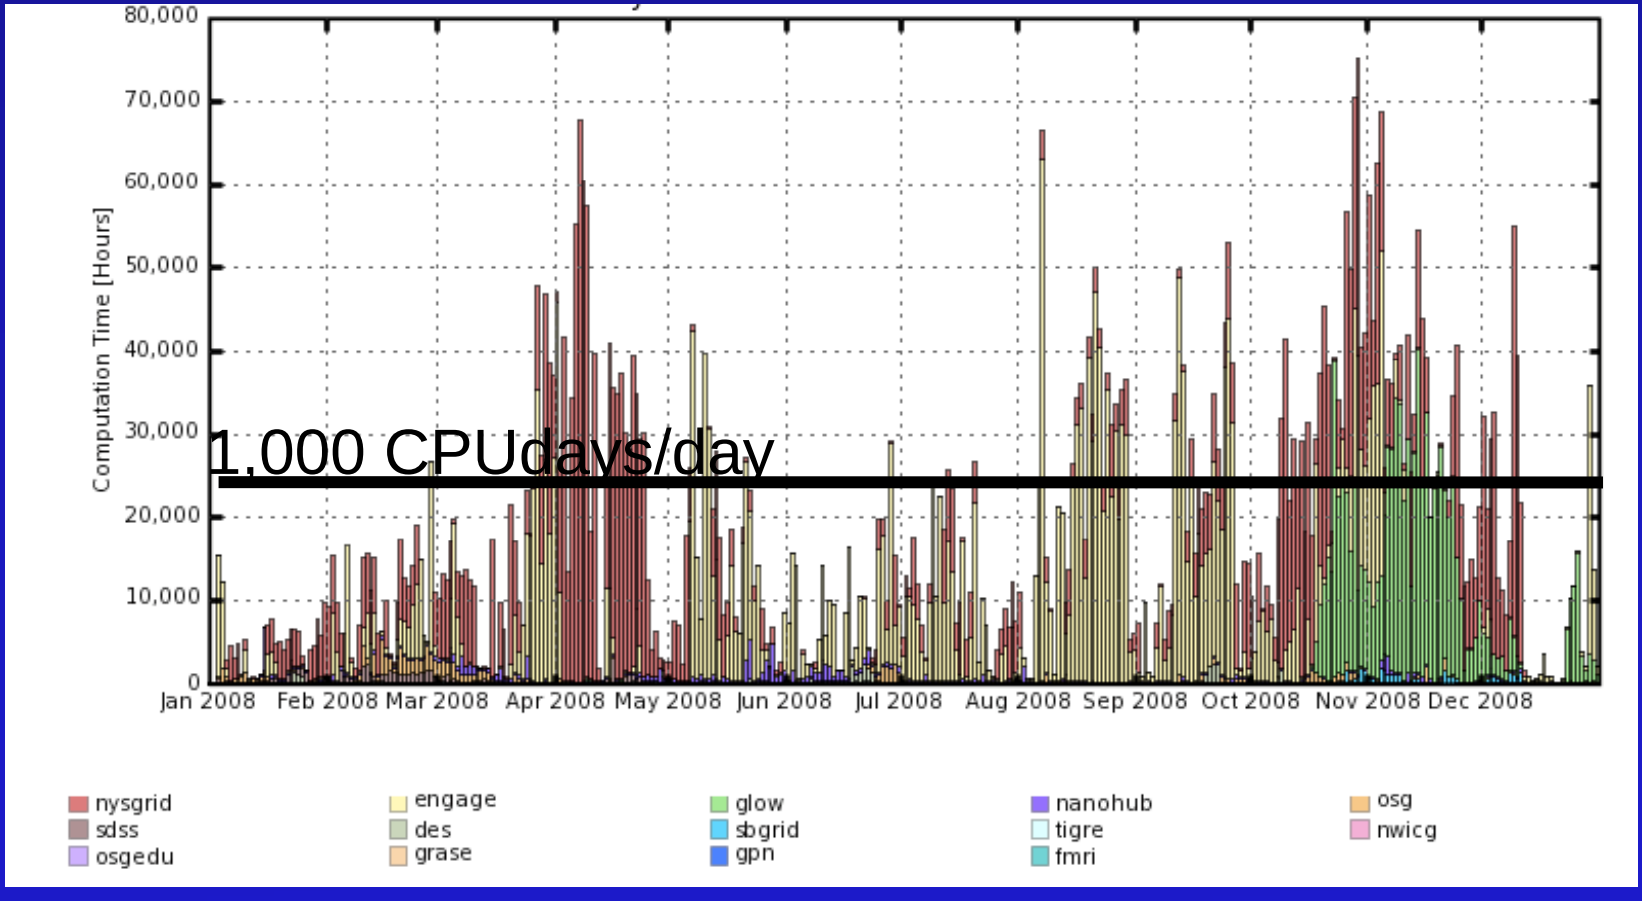

Science Output

~5 non-physics publications.

+ “raptor” protein structure prediction ranked 2 in the world.
  (publication in review)

CDF + D0 >100 physics publications in 2008;

STAR >10 publications in 2008.

Several LIGO publications.

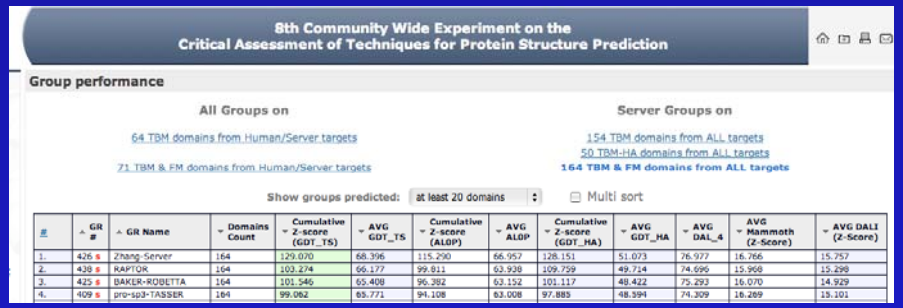

8th Community Wide Experiment on the
Critical Assessment of Techniques for Protein Structure Prediction

Group performance

All Groups on

64 TBM domains from Human/Server targets

71 TBM & FM domains from Human/Server targets

Server Groups on

154 TBM domains from ALL targets

50 TBM-HA domains from ALL targets

164 TBM & FM domains from ALL targets

Show groups predicted: at least 20 domains Multi sort

#	GR #	GR Name	Domains Count	Cumulative Z-score (GDT_TS)	AVG GDT_TS	Cumulative Z-score (ALOP)	AVG ALOP	Cumulative Z-score (GDT_HA)	AVG GDT_HA	AVG DAL_4	AVG Mammoth (Z-Score)	AVG DALI (Z-Score)
1.	426 s	Zhang-Server	164	129.070	68.396	115.290	66.957	128.151	51.073	76.977	16.766	15.757
2.	438 s	RAPTOR	164	103.274	66.177	99.811	63.930	109.759	49.714	74.696	15.968	15.298
3.	425 s	BAKER-ROBETTA	164	101.546	65.408	96.382	63.152	101.117	48.422	75.293	16.070	14.929
4.	409 s	pro-sp3-TASSER	164	99.062	65.771	94.108	63.008	97.885	48.394	74.309	16.269	15.101

Protein Structure Prediction Center
Sponsored by the US National Library of Medicine (NIH/NLM)
Please address any questions or queries to: casp@predictioncenter.org
© 2008, University of California, Davis

Analysis of the Value

Goal

Develop and estimate the benefit and cost effectiveness, and thus provide a basis for discussion of the value of, the Open Science Grid (OSG).

Initial strawman to discuss to of the economic, intellectual, and scientific value that OSG provides.

- Benefit: the gain (usually tangible, but can be intangible) that accrues to the customer from the product or service.

- Value: numerical quantity (measured, assigned, computed); quality that renders something desirable.

OSG Doc 813.

5 Areas of Benefit/Value

Supports Collaborative Research from Small to Large Scales.

Provides a Sustained US Cyber Infrastructure for Scientists.

Contributes to Computer Science and Software Body of Knowledge.

Sustains and Enhances US Expertise.

Creates an Environment for Opportunistic Computing.

e.g. Common S/W Integration, testing, release

	# Major releases a year
Condor and the Globus core grid middleware	~2 Condor, ~1 Globus
Three storage service implementations	~2 each (storage is an important focus of WLCG in particular)
Information and job management services	~2 for each of 4 components
Security infrastructure	~2 for each of 4 components
Underlying common toolkits	~10 in total`
Total Effort expended	35 – 140 FTE weeks (0.67-2.7 FTE)

Table 4: Rate of and Effort on New Software Releases

OSG has sustained and experienced effort to bring to this work. If we take each of the 3 large stakeholders and the 8 at-large VOs and the users of the Engagement VO (12 organizations), then the central software packaging, distribution and support group provides an effort saving of some large fraction of 12* the direct effort used.

Cost savings from centralized software packaging, distribution, support	8-24 FTE

e..g Operating the Core Services

There is some overhead of interacting with a central operations group rather than community specific ones. We take a strawman that each community saves 33% of the 3 FTEs that each community would otherwise need.

	Effort Saved
Round the clock operations for 3 main stakeholders, 5 at–large VOs (3 of the VOs have combined operations at Fermilab), and engaged users (represented by the Engage VO and counted as 1) for a total of 7 communities	7 FTEs

Reduction in Security Risk

	Effort Saved
Security incident response	0.4-1.3 FTE
Policy Development	1.2-2.4 FTE
Direct benefit in tools developed for site administrators	0.5 FTE
Provides sustained expertise in security realm, and thus reduce risk, for all OSG members	

Contributes to Computer Science Knowledge

Expand computer science (CS) knowledge through use of CS technologies in a real, multi-disciplinary environment.

Provide an at-scale laboratory environment and existing collaboration for the testing and validation of CS methods and research directions.

Synthesize the evolving needs of the stakeholders and the problems encountered, thus bringing real world needs to CS research and development.

Value proposition from Opportunistic Use

	Fermi	Nebraska	Oct 1, 2007 - Oct 31, 2008
Cost Item	**$K/Year**	**$K/Year**	**Notes/Definition**
Facility $/ CPU	$92	$117	Used as a sanity check of Facility and Power Costs
System $/ CPU	$128	$150	Used as a sanity check of System Costs
Staff $/ CPU	$191	$840	This is expected by economy of scale. (Nebraska doubled their capacity this year with no staff increase. With that increased capacity for next year, the staff cost/CPU will be $388.)
Total Cost/ CPU Hour	**$0.047**	**$0.126**	For FY 2008
2009 $ / CPU Hour	**$0.040**	**$0.071**	Estimate for FY 2009 (Assuming full year costs with current capacity)

Whats Next?

The WLCG

OSG, working with the LHC experiments in the US, is committed to be an effective contributor to the WLCG for at least the next 10-15 years.

Prepare for WLCG Data Taking

Tier-1s and Tier-2s have demonstrated data taking throughput and robustness.

Tier-3s:

- Specifically supported by OSG.
- Currently ~20 and this is expected to grow to >70 in the next year.
- Not uniform; challenge for OSG to be flexible to the range of needs.

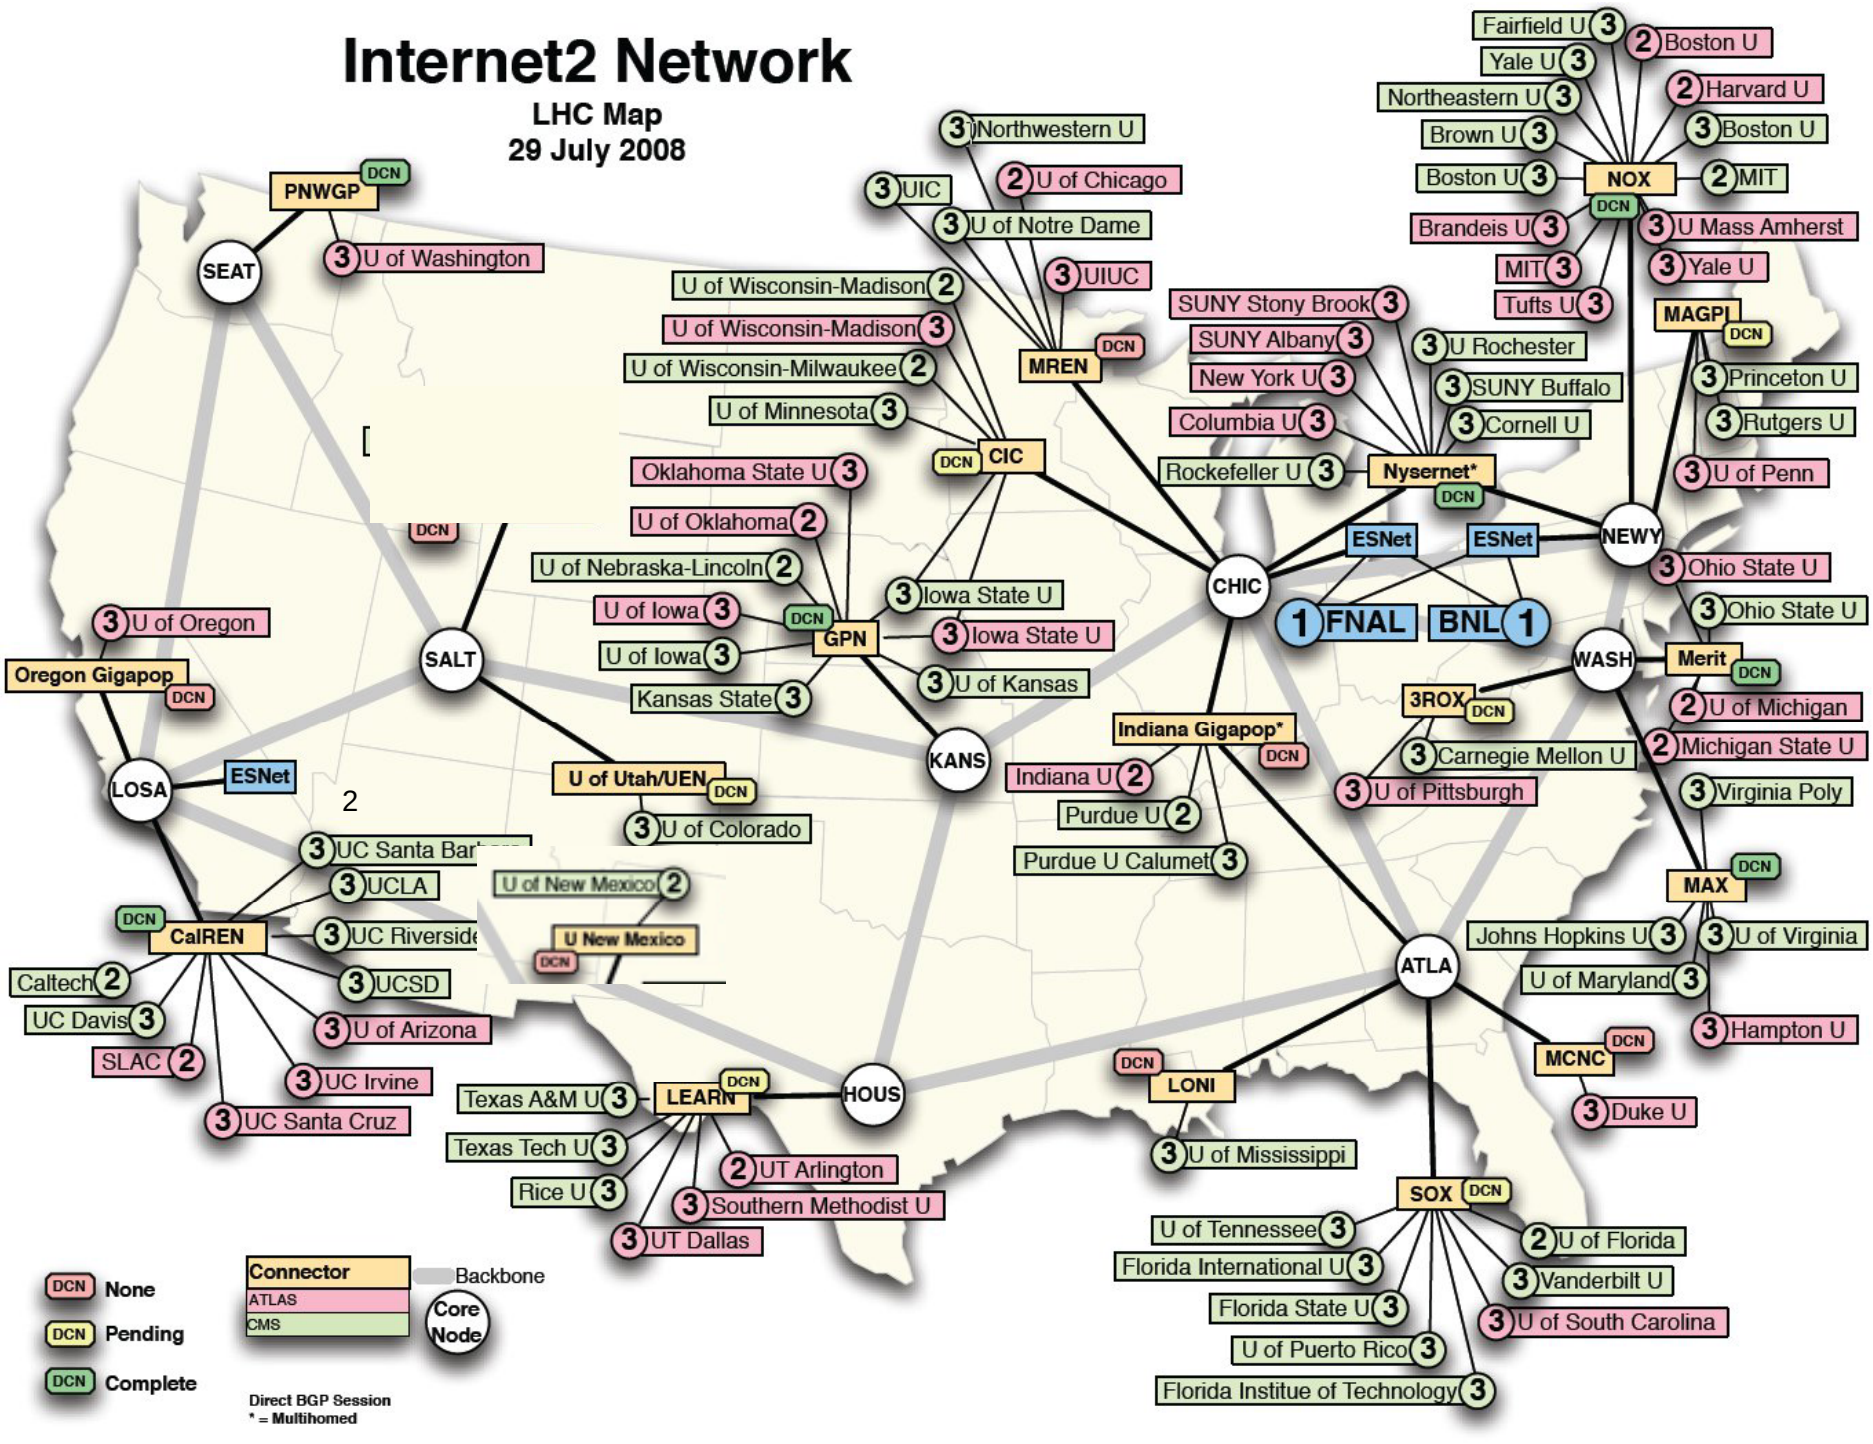
Internet2 Network
LHC Map
29 July 2008
PNWGP
DCN
SEAT
U of Washington
U of Oregon
Oregon Gigapop
DCN
LOSA
ESNet
SALT
DCN
U of Utah/UEN
DCN
U of Colorado
UC Santa Bar
UCLA
U of New Mexico
CalREN
DCN
UC Riverside
U New Mexico
DCN
Caltech
UC Davis
SLAC
UCSD
U of Arizona
UC Irvine
UC Santa Cruz
Northwestern U
UIC
U of Chicago
U of Notre Dame
UIUC
MREN
DCN
U of Wisconsin-Madison
U of Wisconsin-Madison
U of Wisconsin-Milwaukee
U of Minnesota
CIC
DCN
Oklahoma State U
U of Oklahoma
U of Nebraska-Lincoln
U of Iowa
U of Iowa
Kansas State
GPN
DCN
Iowa State U
Iowa State U
U of Kansas
KANS
Indiana Gigapop*
DCN
Indiana U
Purdue U
Purdue U Calumet
CHIC
1 FNAL
BNL 1
ESNet
ESNet
Fairfield U
Boston U
Yale U
Harvard U
Northeastern U
Boston U
Brown U
Boston U
NOX
DCN
MIT
Brandeis U
U Mass Amherst
MIT
Yale U
Tufts U
SUNY Stony Brook
MAGPI
DCN
SUNY Albany
U Rochester
Princeton U
New York U
SUNY Buffalo
Rutgers U
Columbia U
Cornell U
Rockefeller U
Nysernet*
DCN
U of Penn
NEWY
Ohio State U
Ohio State U
WASH
Merit
DCN
3ROX
DCN
U of Michigan
Carnegie Mellon U
Michigan State U
U of Pittsburgh
Virginia Poly
MAX
DCN
Johns Hopkins U
U of Virginia
U of Maryland
Hampton U
ATLA
MCNC
DCN
Duke U
LONI
DCN
U of Mississippi
SOX
DCN
U of Tennessee
U of Florida
Florida International U
Vanderbilt U
Florida State U
U of South Carolina
U of Puerto Rico
Florida Institue of Technology
HOUS
LEARN
DCN
Texas A&M U
Texas Tech U
Rice U
UT Arlington
Southern Methodist U
UT Dallas
DCN None
DCN Pending
DCN Complete
Connector
ATLAS
CMS
Backbone
Core Node
Direct BGP Session
* = Multihomed

Ad-hoc, short-lived VOs

New communities discussing with OSG need ad-hoc, short lived VOs or collaborating groups: SNS/nanotechnology; bio-energy.

Our infrastructure does not support this now; must look for development project to contribute the what is needed.

VOSS: Cross Organizational Analysis

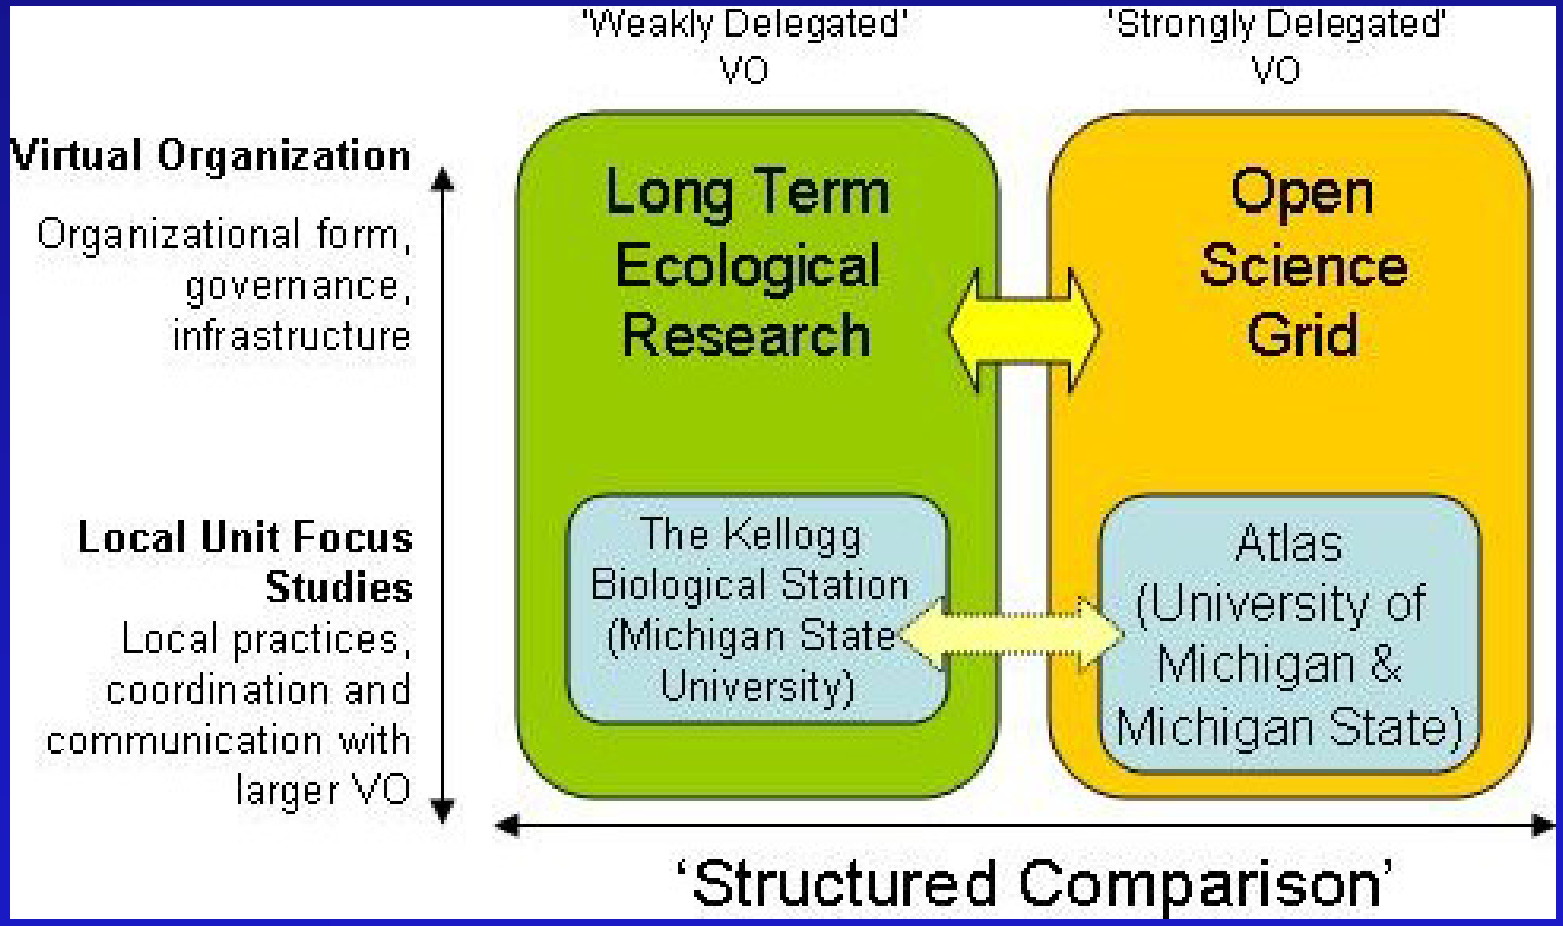

from David Ribes and Thomas Finholt

Improve and Sustain the Software Stack

Through the (new) software tools group we will be able to provide **more focus and effort on interactions with software providers, enabling new software** capabilities, and transition new techologies to the production infrastructure.

Based on **our unique environment for testing and use of technologies (at scale, under fire, and by multiple communities**) we will continue to look for new software providers who can best advance the common technologies our stakeholders need.

Sustaining OSG for the stakeholders

DOE SciDAC/NSF Review Jan 2009 gave Encouragement to Proceed:

"Continue to engage in the larger strategic questions of national cyberinfrastructure(s) among relevant parties. Continue to iterate on the best rules and framework for federation – evolutionary and strategic for all large cyberinfrastructures – be focused and need-driven (i.e. archival storage). Create conditions at the right levels for discussion of frontier issues of cyberinfrastructures."

Future Needs..

Expanding the communities and partnerships especially with Campus and Teragrid

Managing change and increased heterogeneity

Open Science Grid**S**

Security and trust (attention, defense, protection)

Really (really) provide for ad-hoc group collaborations

Extended computing paradigms – Workspaces, Clouds, VMs

OSG and the National Cyberinfrastructure ?

How can OSG's experience on the Campuses be of most value?

What is OSG's role and place in partnering with TeraGrid?

How can OSG contribute to the usability of the Leadership Class Facilities?

How does OSG contribute to software sustainability?

Is OSG an exemplar from which Europe in its next phase of National Grids can benefit?

From the Director of the Office of Cyberinfrastructure (NSF)

OSG as Model "Campus Bridge"

- NSF very interested in creating "bridges" from campus to national CI
- OSG is a *national* CI, *locally* deployed...perfect model of this
- We are very interested in...
  - Exploring ways to integrate campuses better with national centers, instruments
  - TeraGrid-OSG cooperation
    - Driven by applications!
  - Understanding example science communities that can benefit from, drive this
  - Related international cooperation: EGEE/EGI, etc

National Science Foundation
Where Discoveries Begin

Edward Seidel
hseidel@nsf.gov

Office of Cyberinfrastructure

Clouds & Grids
Views of Distributed Computing

OSG accommodates Clouds in the model as (another) interface to delivery of well-specified storage, processing, data, information and work management services.

OSG automated management of job execution resource selection and dispatch prepares us for engagement with implementations of Clouds.

At a practical level demonstrations exist

- Integrating Condor pools with Amazon elastic clouds
- Running STAR applications on the cloud through virtualization.

in discussion following VM workshop at All Hands Meeting

Users of of the OSG are actively starting to use Condor-Glide-in, Cloud computing both commercial and installed on OSG resources, & Virtual Machine technologies, and as a way of improving throughput & usability & masking the heterogeneity of the underlying resources.

If all can be supported on the common infrastructure then it is OSG's position that the choice of technology is up to the individual community (Community Grid or VO).

At present the technologies are such that in general sites must be configured through "root" access by system administrators to support these techniques. For example, in many cases "glexec" is installed to support the security associated with the use of Condor-Glideins; for XEN the OS must be configured in a particular fashion;

OSG will not deploy technologies or policies that prevent the use of virtualization, glide-ins or cloud technologies or mechanisms. We will work with VOs requesting use of these technologies to assess and support their usability & security on OSG accessible resources and the integration of Commerical Clouds as potential resource providers.

How to Manage Policy and Effectively "fill the facility"

OSG Roadmap 2009-2010 - WLCG

Operate an effective infrastructure and

- Continue to extend the site monitoring, validation, accounting, information services for sites & VOs towards complete coverage of needed capabilities & fault checking; Continue work with interoperation/joint approaches with WLCG monitoring group.
- Complete SLAs and Procedures for core services (security, policy etc).

Work closely with the experiments and

- Grid-wide deployment of pilot/pull-mode job routing and execution technologies (glidein-wms, panda). Meet the security requirements of the "experiment frameworks" WLCG group for interoperability between EGEE and OSG for glexec/SCAS enabled pilots.
- Provide software/VDT updates released in a timely fashion, can be installed incrementally, and can be rolled-back when/as needed.
- Develop policy/priority mechanisms (inter-VO and intra-VO) for when the resources become fully or over subscribed.
- Complete evaluation of use of OSG by ALICE and support their production needs as requested. Ensure good integration of any additional needs from Heavy Ion parts of the experiment.

OSG Roadmap 2009-2010 – Tier-3s

Ramp up support for US LHC Tier-3s. Expect ~70 within a year.

- Understand different types of Tier-3 and adapt OSG technologies, and processes to accommodate their needs. US ATLAS has identified 4 types of Tier-3: with Grid Services; Grid End-point; Workstation; and Analysis. We will work with them to understand what this means.
- Identify Tier-3 liaison within the OSG staff, and work with the US ATLAS and US CMS Tier-3 support organizations.

Develop support for experiment analysis across OSG sites.

As needed, support opportunistic use of other OSG resources for US LHC physics.

OSG Roadmap 2009-2010 - broadening.

- Integrate other Identity systems as starting-points in the end-to-end security infrastructure (Shibboleth, OpenID, Bridge-CA)
- Solidify and extend the use of storage – currrently evaluting Bestman/Hadoop as an SE for Tier-2 and Tier-3s.
- Ensure continued interoperation with EGEE including support for CREAM CE.
- Improve the usability of the infrastructure for new small communities; as new communities come to the table add tasks to meet their needs (Structural Biology Grid, other Biology groups)
- Provide initial support for ad-hoc, dynamic VOs (initial stakeholder evaluation by SNS)

OSG Roadmap 2009-2010 - sustaining

Work with US agencies, US LHC, LIGO and other stakeholders to understand needs and plan for sustaining the OSG after the end of the current project in 2011.

- Work to a model of “satellite” projects which can feed technologies, applications and new services into the core of the OSG.
- Understand interfaces to TeraGrid.

Work with WLCG, EGEE in the EGI-NGI era, bringing experience and methods of OSG to the table to inform and learn.

Lots of interesting things to do!