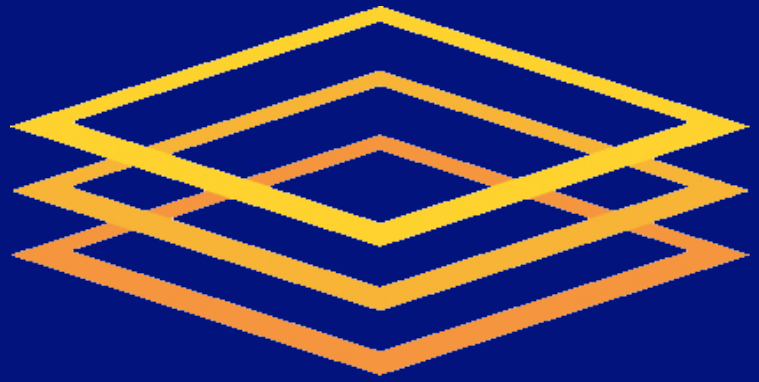

Open Science Grid

Collaborative Science on a

High-Throughput Distributed Facility

Ruth Pordes, OSG Executive Director, Fermilab

NLIT May 13th 2008

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.

Since 1980s…Fermilab Science Mission supports Science with Large Data and Collaboration Needs

Data stored in 2007 for the 2 Tevatron Experiments (CDF, DZero) if on DVDs 4 GB, 1.2 mm high

Early adopters of networks for data distribution, web information access, & distributed computing.

DZero Collaboration

In 2000s… Fermilab also hosts data, computing, physics for CERN LHC CMS Experiment

Data export from Fermilab to CMS Universities over a month

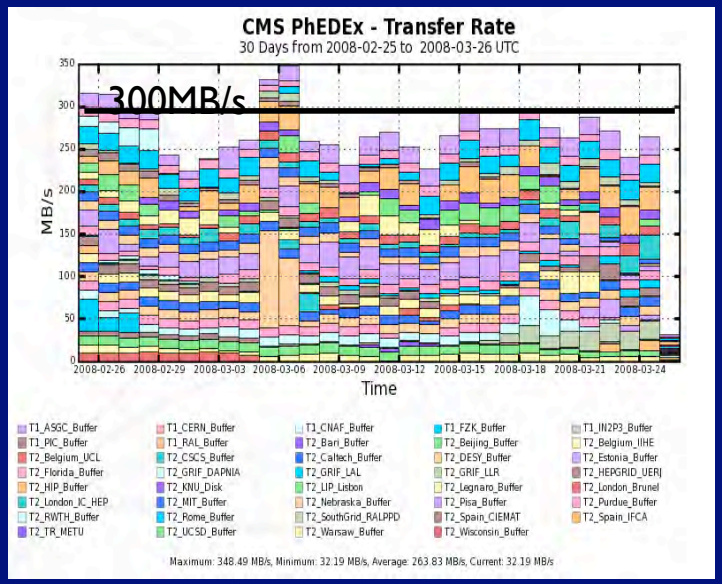

CMS: 2310 Scientific Authors
38 Countries
175 Institutions
1 of 4 LHC experiments

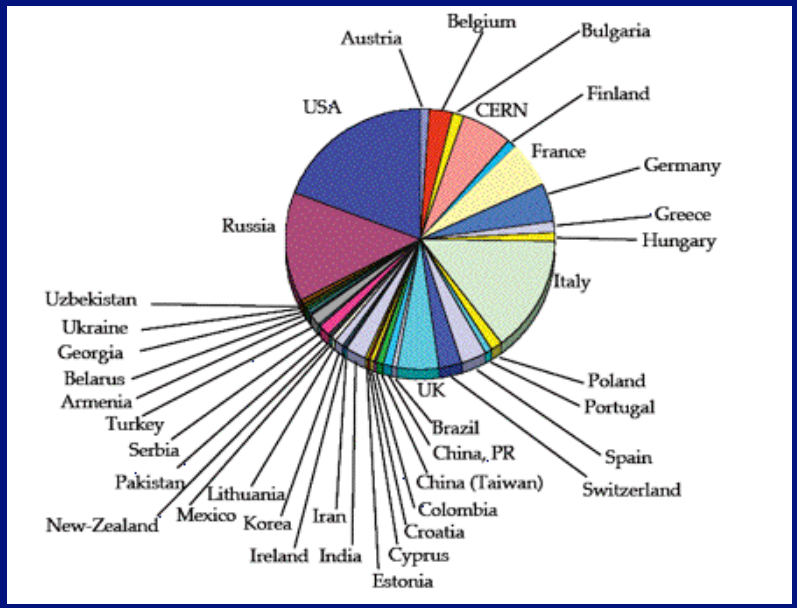

LHC@FNAL, 24x7
Remote Operations Center

Emergence of Grid Projects (for LHC)

LHC adopted a world wide computing Grid model for data distribution, storage, processing and analysis.

Computer Science development groups promoted Grid concept and software as general solution (silver bullet).

Funders sponsored Collaborative projects - computer science, physics, other sciences - to leverage and promulgate general and widespread solutions.

- EU : European Data Grid (2000-2004), EGEE-I (2004-2006), EGEE-II (2006-2008), EGEE-III (2008-1010), EGI?
- US: *DOE SciDAC-1:* Particle Physics Data Grid, DOE Science Grid (~1999-2006); *NSF ITR:* Grid Phyiscs Networks, International Virtual Data Grid Laboratory
- And more….

Open Science Grid: emerged from ad-hoc cooperation of Science Collaborations, Grid Projects, & IT Facilities

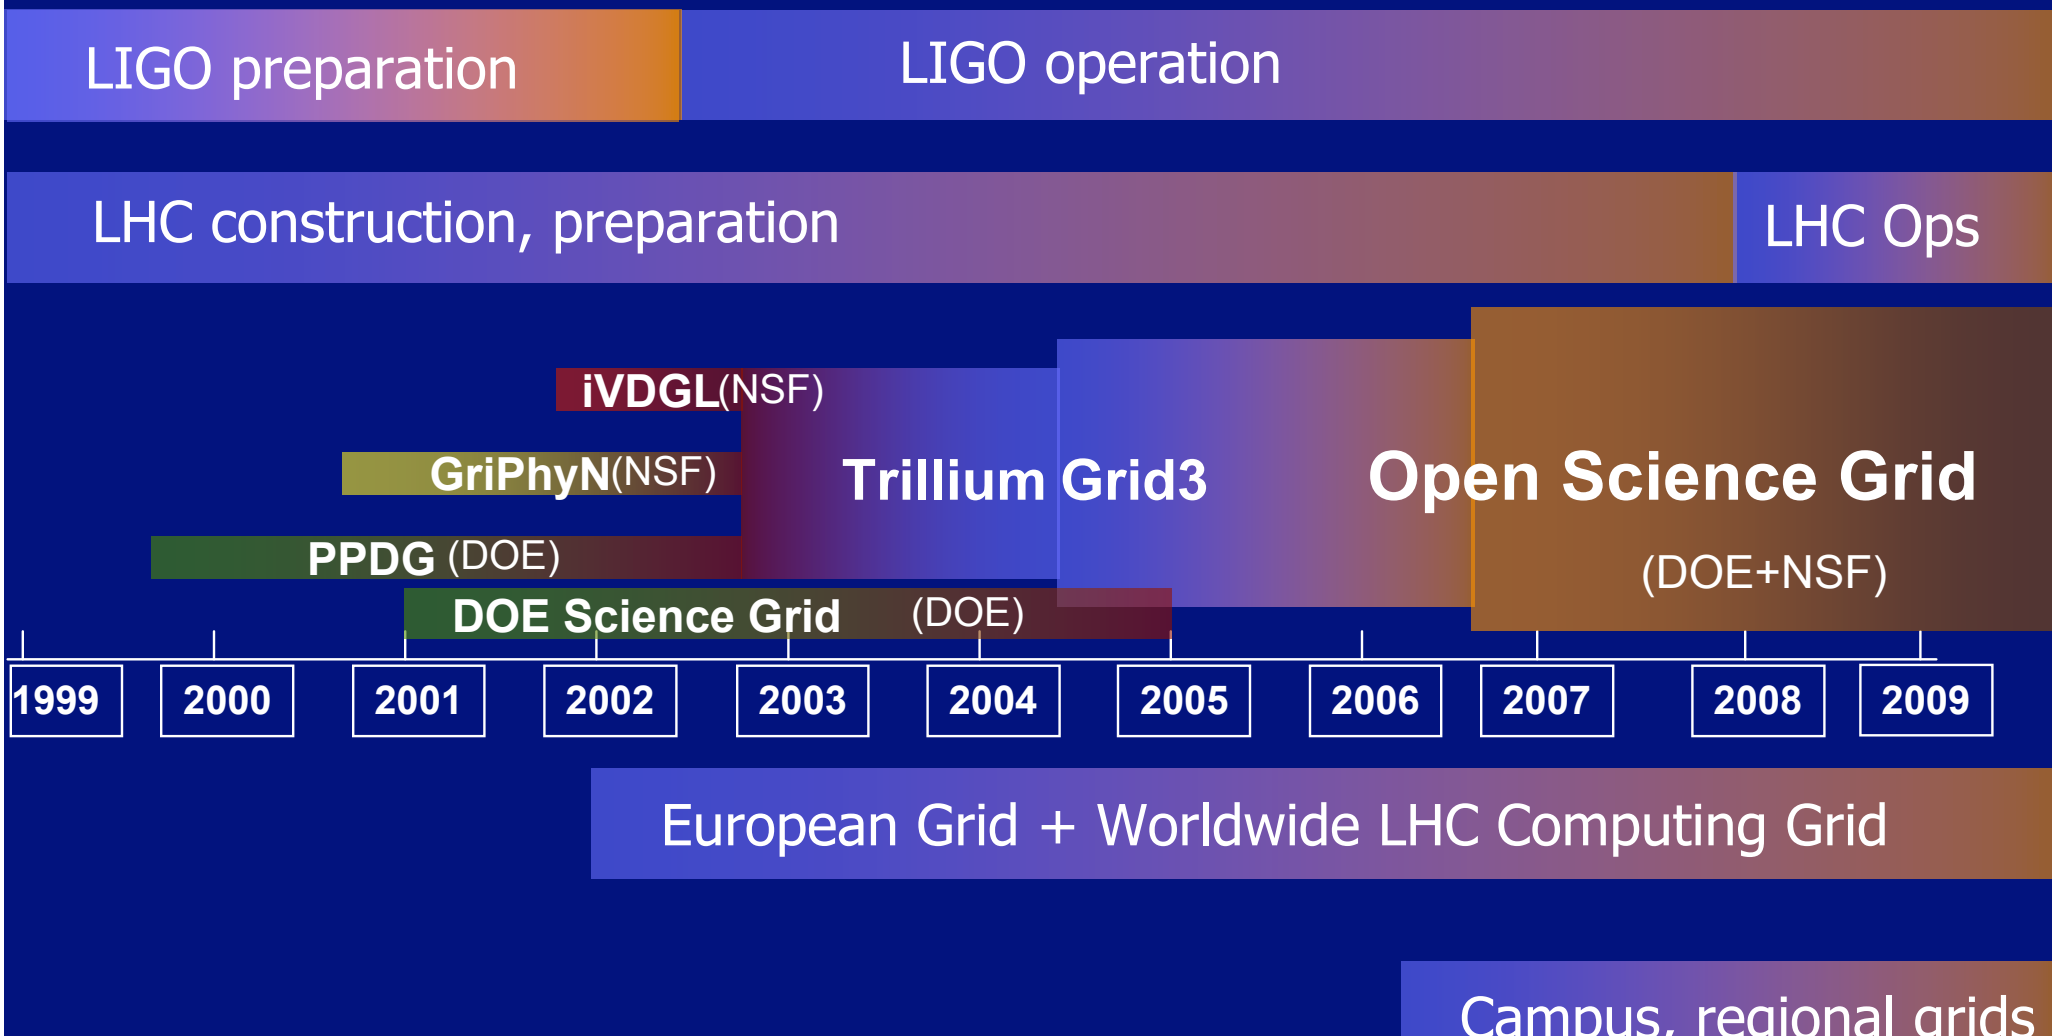

Open Science Grid characteristics

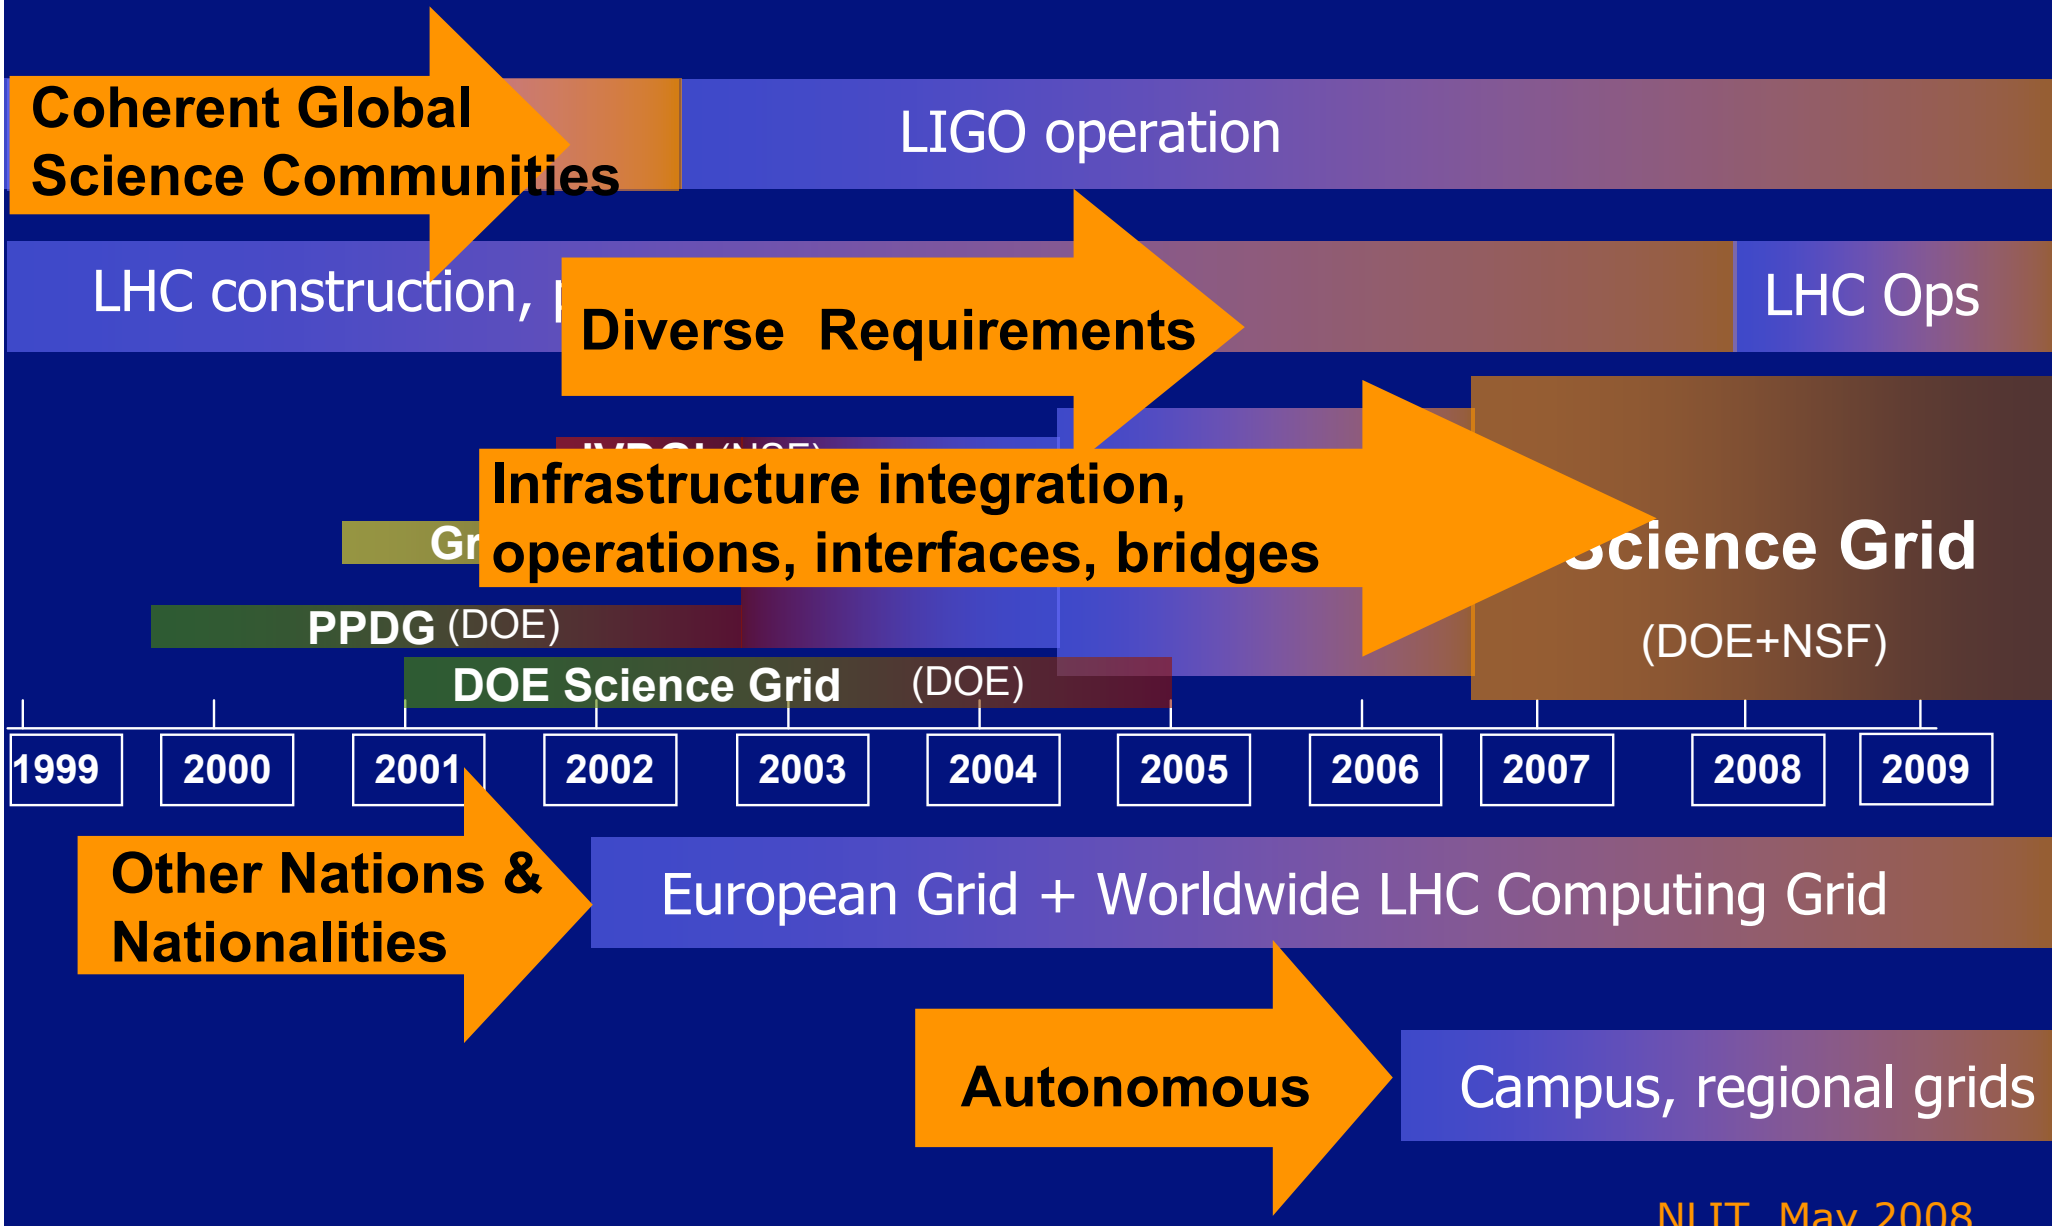

Open Science Grid's Goal

To provide a cross-domain self-managed national **distributed high-throughput computing facility**…

that brings together **campus and community infrastructures at all scales**…

to provide services, support and facilitate the needs **of scientific and research communities at all scales…**

It's a non-goal to:

- own processing, storage, networking hardware.
- develop software.

OSG and the WorldWide LHC Computing Grid

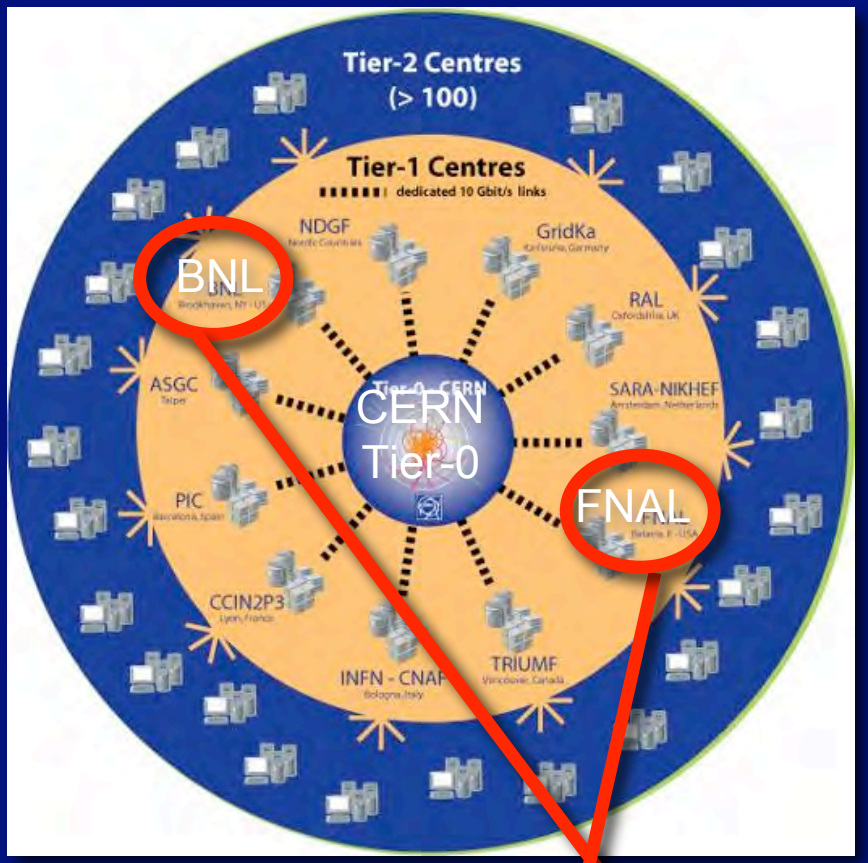

US ATLAS,
US CMS
Tier-1s

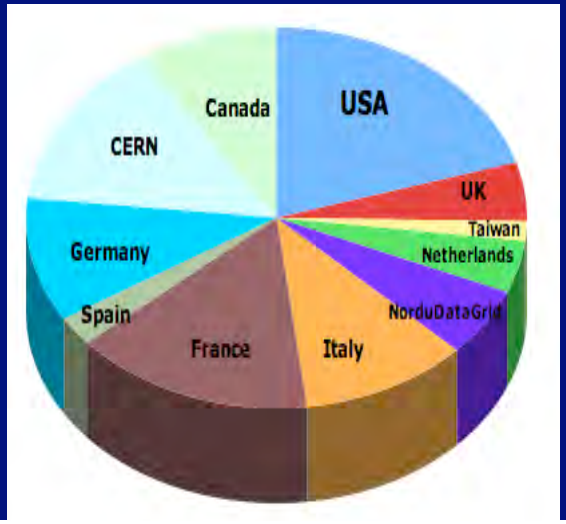

ATLAS and CMS processing usage for the past year (WLCG Apel reports).

Open Science Grid

bridges

DOE & NSF

Integration of high-throughput and high-performance computing and storage from each (e.g. NERSC, NCSA)

Shared resources between OSG and TeraGrid (e.g. Purdue, ORNL NCSA coming online)

Trust (single signon) between heterogeneous security domains.

Dynamic sharing of available resources or ad-hoc allocation agreements across agencies.

Meeting point for different software toolkits.

Integration of small to large scale facilities

Snapshot of the OSG

Research Participation

- Majority from physics : Tevatron, LHC, STAR, LIGO.
- Use by 10 other (small) research groups.
- >10 software development groups contribute to commonly supported software stack (Virtual Data Toolkit).

Core Project

- 5 years funding (if we pass our reviews) at total of $30M.
- 35 FTE effort across 16 institutions (try to make staff contributions >50% level each).
- Full time Project Manager.
- Leverage ~equivalent effort in contributions.

Computational resources accessible at

- 5 DOE Labs - BNL, Fermilab, NERSC, ORNL, SLAC.
- 65 Universities.
- 5 partner campus/regional grids.

Snapshot of the OSG cont

Accessible resources of

- >43,000 cores,
- 6 Petabyes disk cache,
- 10 Petabyes tape stores

Usage

- 15,000 CPU WallClock days/day
- 1 petabyte data distributed/month.
- 100,000 application jobs/day.
- 20% cycles through resource sharing, opportunistic use.

Map of the Sites (5/4/08)

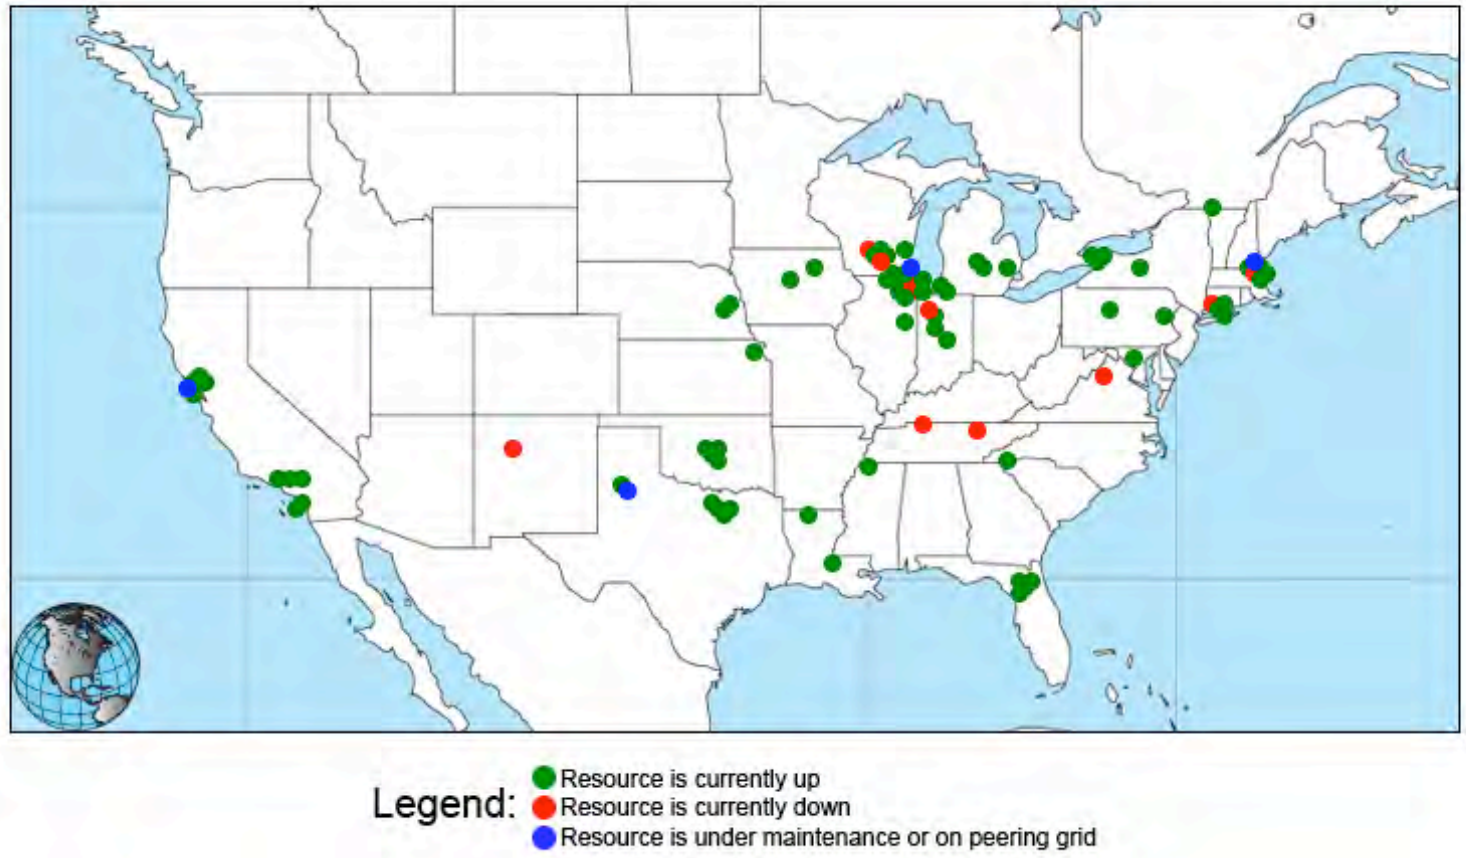

(+3 in Brazil; 2 in Mexico; 2 in Tawain; 1 in the UK)
Grows by 10-20 per year.

Use by Community

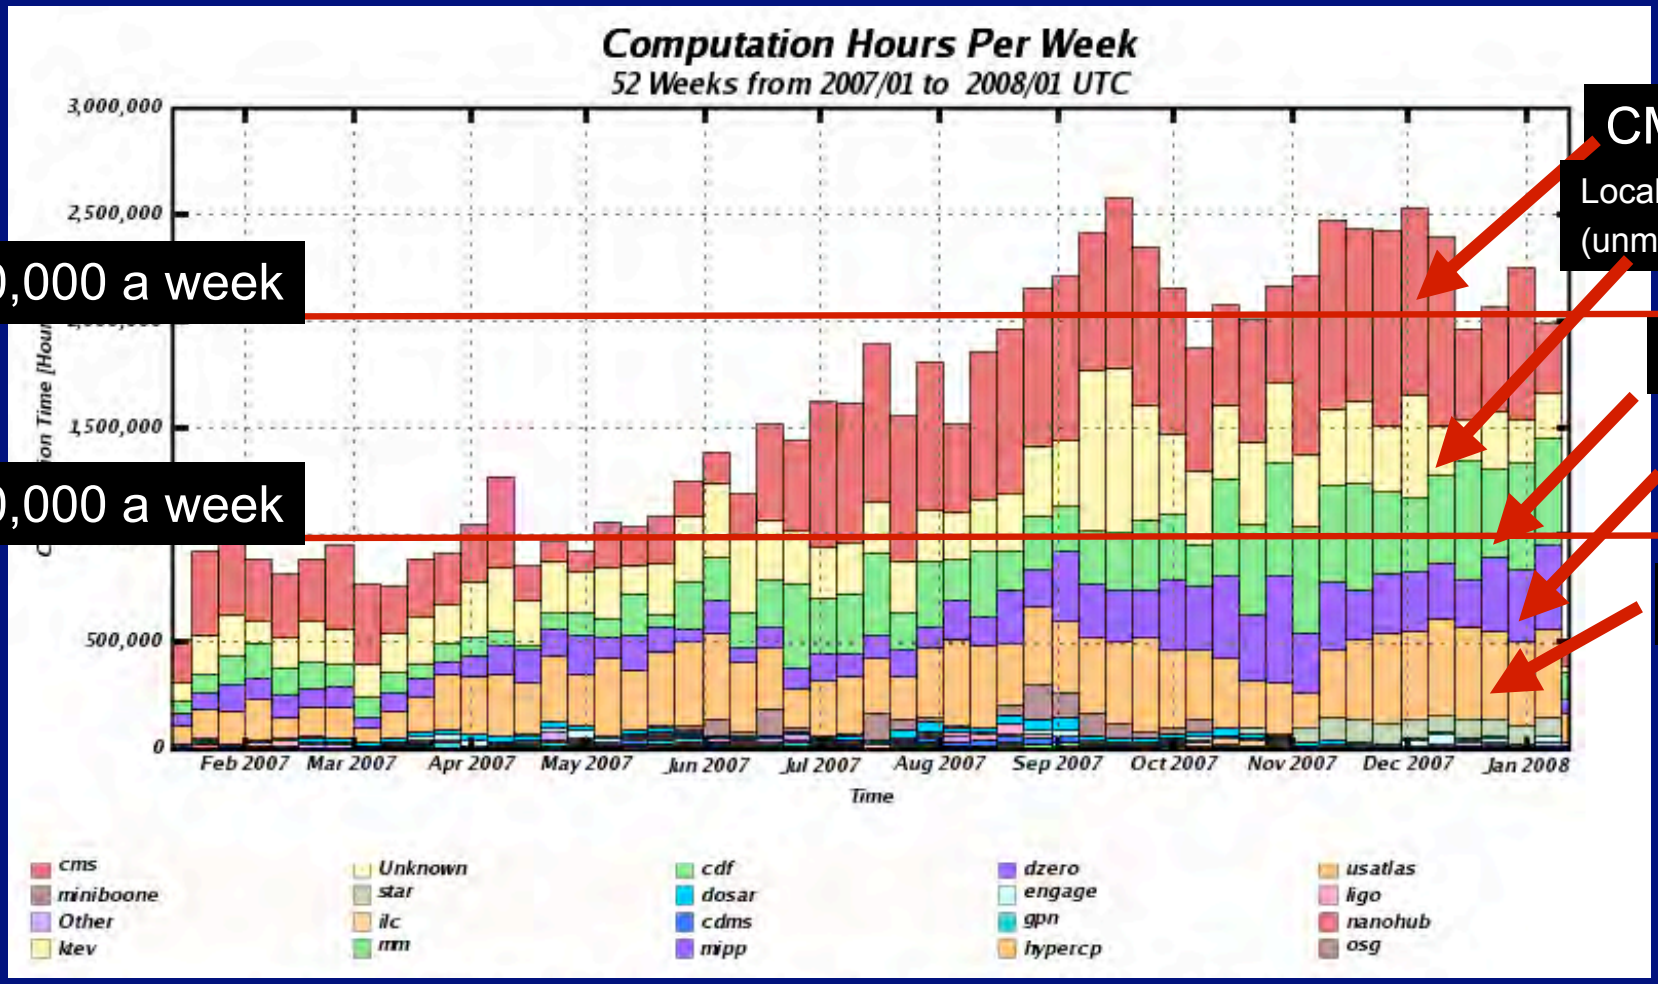

2,000,000 a week

1,000,000 a week

CMS

Local Usage & bugs (unmapped to VO)

CDF

D0

ATLAS

with constant expansion…
Laser Inteferometer Gravitational Observatory (LIGO) BOINC based Einstein@home

LIGO science topic to search for deviations in pulsar signals. Runs across Boinc home computers, German and US grids.

Testing use of a new service on OSG & getting science output.

Using 3,000 nodes steadily before accounting system adapted to report it!

User stats
Users 1 - 100
Sorted on Total credit
Total number of users: 197947

Einstein@Home

▸▸ 100 ▸▸ 500 ▸▸ 1000 ▸▸ 10000

last month	last week	last day	Pos.	User Name	Total credit	Credit /day	Credit /week	Credit /month	Average credit	Over take	Options Compare
→0	→0	→0	1	Steffen Grunewald, for Merlin/Morgane	206,921,717	606,039	4,571,176	24,452,071	726,257	-	
→0	→0	→0	2	Bruce Allen	158,541,694	0	0	0	1	-	
→0	→0	→0	3	AEI eScience group, for the German Grid (D-Grid) and the Open Science Grid (OSG)	139,527,258	2,801,237	21,153,060	55,293,054	2,425,140	8	
→0	→0	→0	4	UW-Madison CAE	68,334,005	30,108	259,335	1,396,695	38,629	-	
→0	→0	→0	5	L&S IT Office	62,610,351	131,061	1,005,620	4,097,124	138,993	59	
→0	→0	→0	6	Erik A. Espinoza	48,553,082	23,000	282,851	2,231,881	50,335	-	
→0	→0	→0	7	LIGO Livingston LDAS	45,309,169	948	8,538	225,044	7,080	-	
→0	→0	→0	8	UTSC Computing Labs	42,474,911	3,790	19,898	81,018	2,741	-	
→0	→0	→0	9	UITS Computer Labs	42,065,464	33,886	250,956	1,093,532	36,025	13	

Use by Site

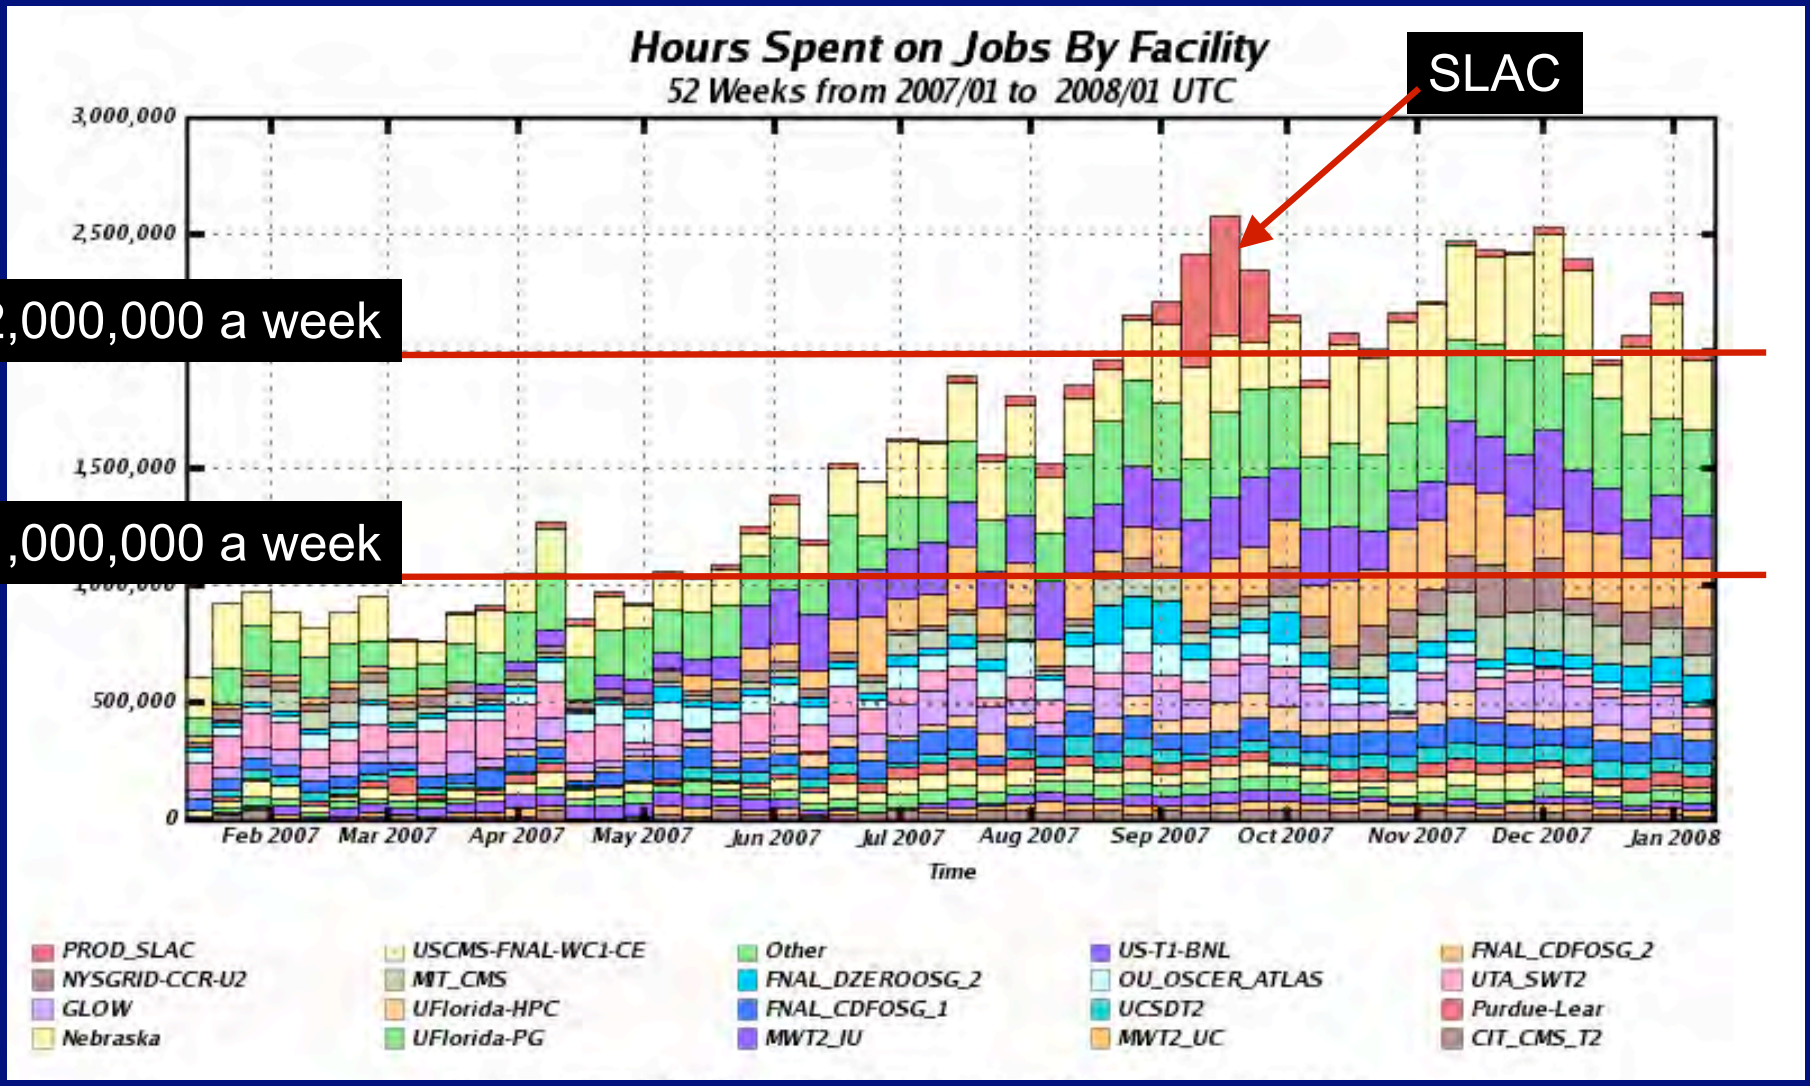
Hours Spent on Jobs By Facility
52 Weeks from 2007/01 to 2008/01 UTC
SLAC
2,000,000 a week
1,000,000 a week
3,000,000
2,500,000
1,500,000
500,000
0
Feb 2007
Mar 2007
Apr 2007
May 2007
Jun 2007
Jul 2007
Aug 2007
Sep 2007
Oct 2007
Nov 2007
Dec 2007
Jan 2008
Time
PROD_SLAC
NYSGRID-CCR-U2
GLOW
Nebraska
USCMS-FNAL-WC1-CE
MIT_CMS
UFlorida-HPC
UFlorida-PG
Other
FNAL_DZEROOSG_2
FNAL_CDFOSG_1
MWT2_IU
US-T1-BNL
OU_OSCER_ATLAS
UCSDT2
MWT2_UC
FNAL_CDFOSG_2
UTA_SWT2
Purdue-Lear
CIT_CMS_T2

The 3 Strata of OSG

Virtual Organizations

Common Services

Resources and Sites

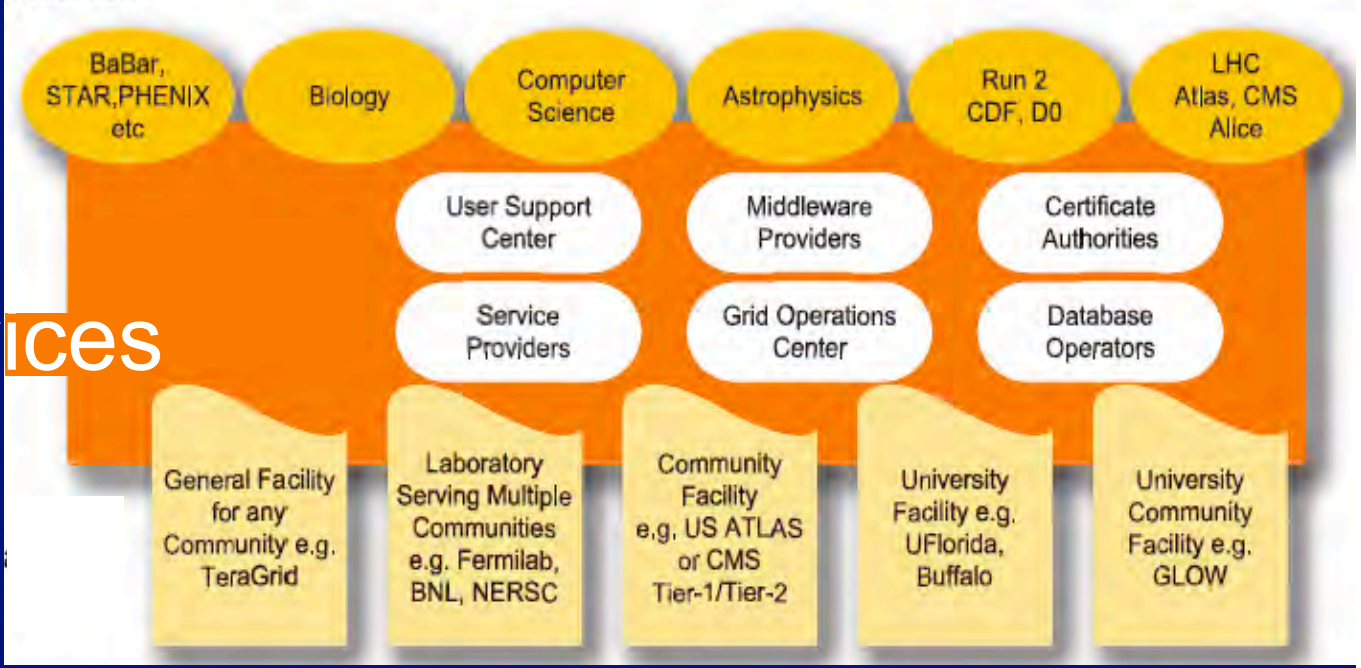

Providing End-to-end Integration

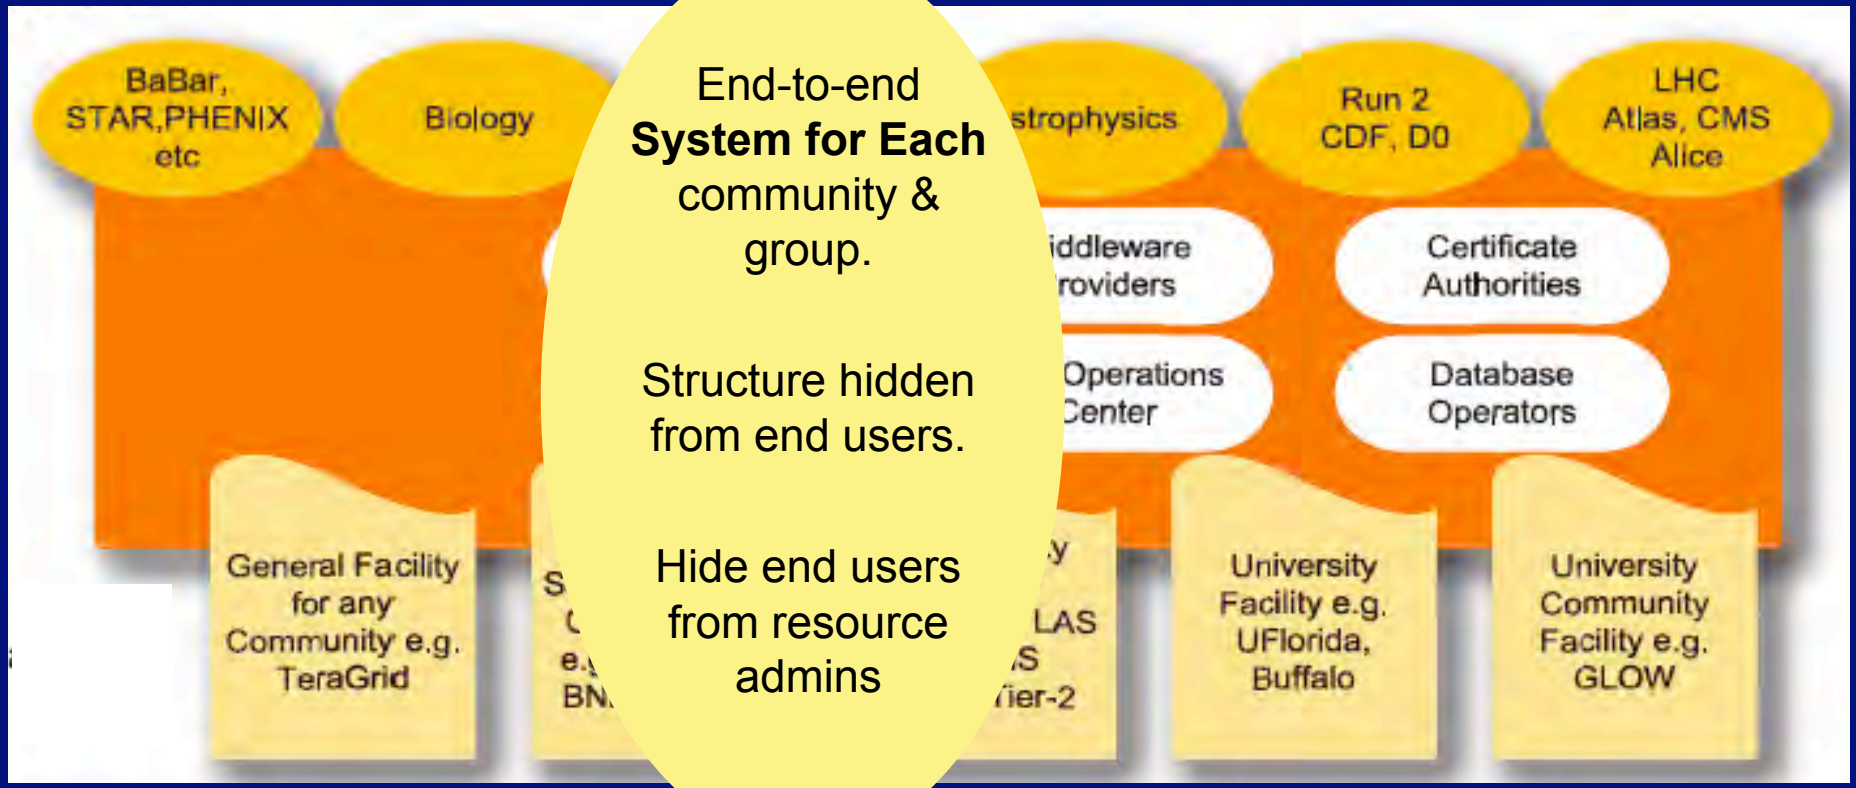

And Bridging to Other Grids

OSG Acts as an Agent

Experiments, Collaborations, Researchers,
Faculty, Educators, Students,

Development and research projects sponsored by DOE,
NSF, ASCR, CISE, universities, labs, industry etc.

Sites & Resources
Ownership, Hardware, Software

Site owners retain local management of their resources, including responsibility for security and policies of use.

- A site configured to be used through the OSG retains support for local use.
- Site administrators contribute to the OSG monitoring, maintenance of effective operation of the distributed facility.

Processing hardware is (different flavors of) Linux farms.

- Exceptions are the Clemson Windows cluster and Nersc AIX machines.

Tape stores & disk caches are unique to each provider:

- Fermilab Enstore; BNL, NERSC, SLAC HPSS.
- Variety of disk caches with largest at the labs & LHC university sites.

OSG software installs (for compute, storage, user management etc.) provide for remote access, administration and participation in the distributed facility.

OSG and DOE HPC Facilities

NERSC resources accessible to OSG.

- OSG has an allocation and manages the policies across the OSG Consortium.

OSG provides ad-hoc support for MPI.

- User needs driving pragramatic steps to support;
- Stimulates design/architectural development of full support;
- Software and distributed infrastructure support currently underway.

OSG provides the “glue” to move data and jobs across diverse resources at different stages of the end-to-end workflow.

- to/from unique HPC facilities and/or local desktops and/or local-area infrastructures.

Networks

We rely on today's state of the art networks for science, advanced experimental research networks, and end-to-end “last mile” connectivity.

OSG works with ESNET, Internet2, Campus network groups to help ensure that the functionality and performance needed by the user communities are met.

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.

A Virtual Organization

Is an autonomous and self-organized collaborative community of people and things they share.

Delegates physical identity management to the participating “Real” Organizations:

- Verification of the actual identity of the people, and the purchase and maintenance of physical hardware.

Ranges in scale and complexity:

- from “heavy-weight” long-lived, with strong governance and policies e.g. the LHC experiments,
- to “ad-hoc” self-organizing, groups of individuals who dynamically build on a few common ideas in much the same way as internet-based social networks e.g. students in a Grid school class.

OSG acts as an Agency for Virtual Organizations

A Virtual Organization participates in the OSG through agreement.

OSG brokers and supports relationships and expectations between and among Virtual Organizations, resource, facility, service and software providers.

A (relevant) example: multi-layered, multi-national; Complex

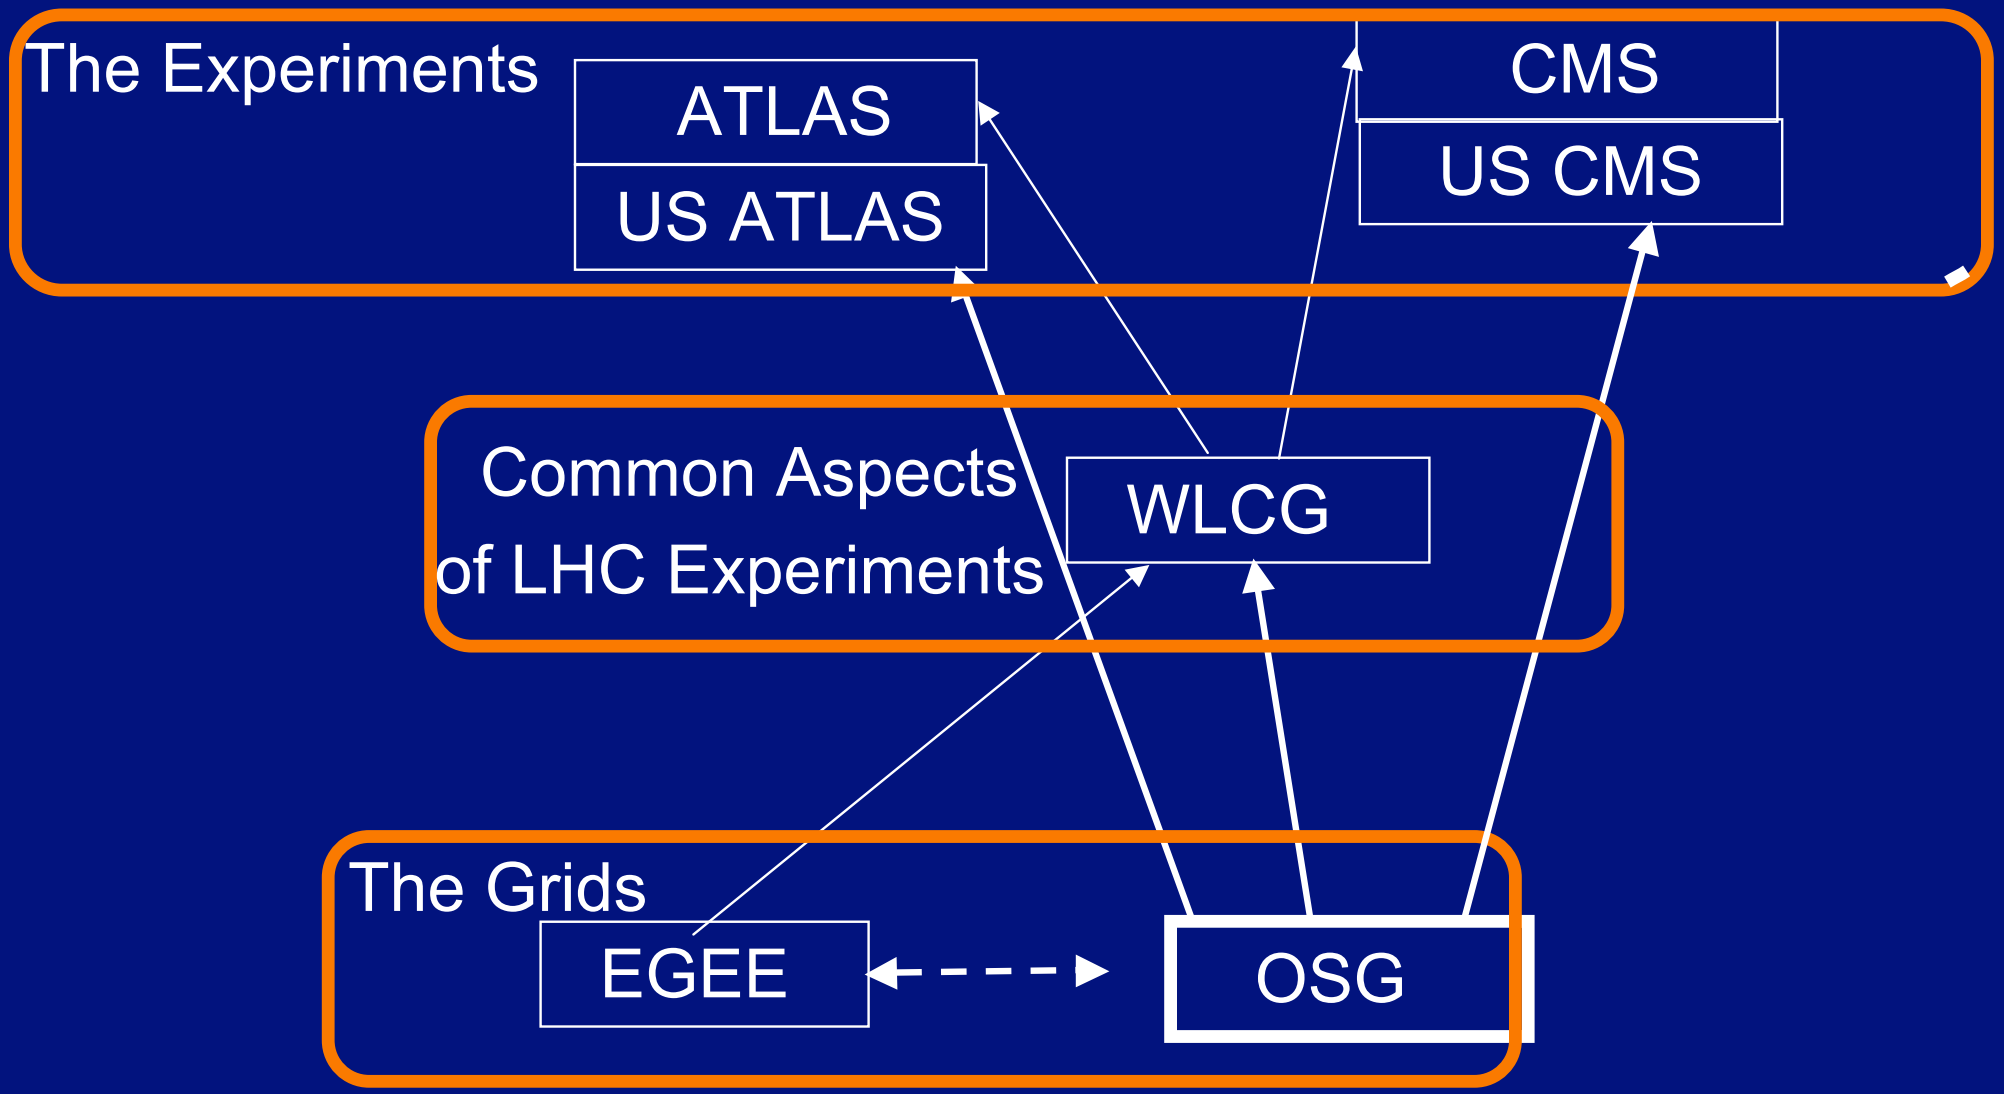

Why is regarding OSG as a VO useful?

Open Science Grid is a collection of people and the common services/software and assets that it owns:

- Software repositories.
- Databases of information.
- Services for use by the organization.

OSG includes other VOs in the same vein:

- **Engage, Education, OSG, MIS, OPS**.

Governance, procedures, security policies developed for OSG itself act as templates for the other VOs.

Common Services

Operations, Grid-wide services

Troubleshooting & Diagnosis

Software Toolkits

Integration & Testing

Security

Engagement (way of helping new/novice recruits)

Operations and Grid-Wide Services

OSG provides typical services of a Data Center: resource monitoring and alarms; accounting; information management; helpdesk/ticket system..

.. while allowing local site and organizational control.

OSG also provides coordination, expertise and support to make contributions effective, low effort (& secure).

Software Toolkit
The Open Science Grid Virtual Data Toolkit (VDT)

Integrated set of software easy to install, configure, operate, update.

OSG provides frontline support for all components.

Grid agnostic - components configured & supported for use by TeraGrid, EGEE, UK Escience.

Includes:

- *Core Grid middleware*: Condor and the Globus Toolkit.
- *Information & Monitoring*: Information providers, LDAP repositories, Gratia accounting, RSV resource validation scripts.
- *Storage:* BeStMan(LBNL), dCache(FNAL/DESY), XRootd(SLAC).
- *Security:* X509 certificate management, VOMS VO management, GUMS/Prima Certificate to Unix ID mapping.
- *Support Software:* Apache, Tomcat, Berkeley DB, MySQL, OpenLDAP, PHP, Metronome build and test, Squid web caching, etc. & tools to help administrators and users.

Major External Software Providers

BeStMan	LBNL
Community Driven Improvement of Globus Software (CDIGS)	Argonne National Laboratory
Center for Enabling Distributed Petascale Science (CEDPS)	Argonne National Laboratory
Condor	University of Wisconsin, Madison
dCache	DESY/Fermilab
Data Intensive Science University Network (DISUN)	UCSD, Caltech, Wisconsin, Florida
Enabling Grids for EScience (EGEE-II/III)	Europe
Energy Sciences Network (ESnet)	ESnet/LBNL, ESnet
Internet2	Internet 2
LIGO Physics and the Information Frontier	University of Wisconsin, Milwaukee
OSG Accounting	Fermilab
OSG Privilege/Authorization	Fermilab/Brookhaven

Software release process is well developed & has been exercised multiple times

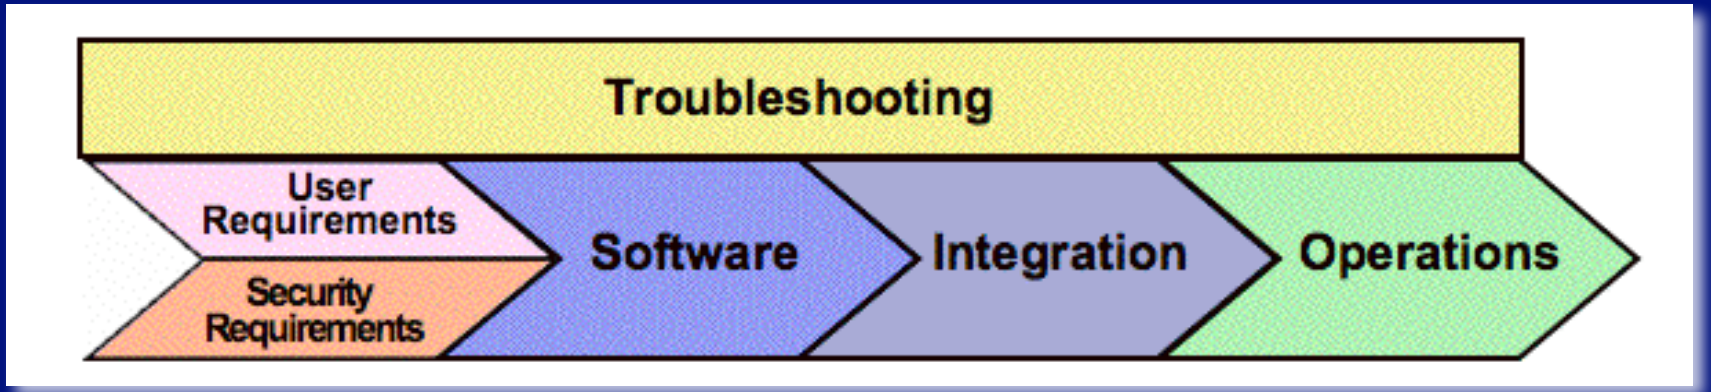

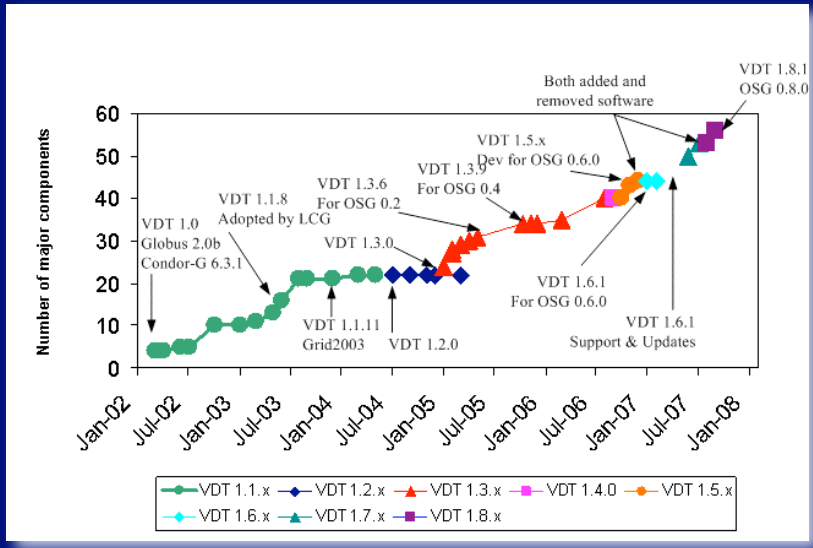

VDT components over time: built for 15 Linux Versions

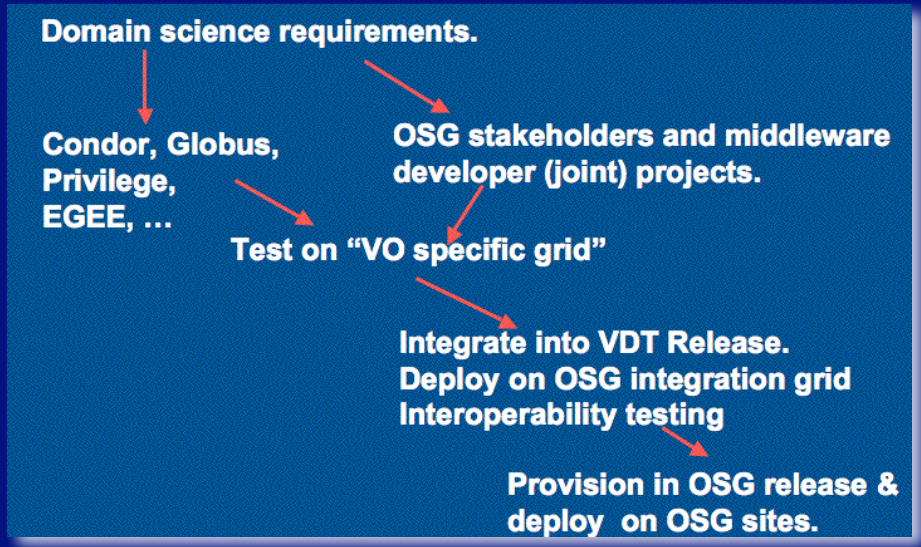

Development & testing

Integrated End-to-End Security

Virtual Organizations, Sites, Open Science Grid each responsible for their own assets.
Local DOE/NSF security teams helping with policies & forensics.

Check vulnerabilities

Develop/announce patches

MOUs with providers

Education

Incident response & fire drills

Monitoring & auditing

Policies & agreements.

Participate in standards bodies to maintain inter-grid compatibility.

Cooperate on response & operation teams.

Collaborate on policies & software development.

Engagement - our word for immersive, embedded help and training

Bring the power of the OSG to scientists, educators, and campus beyond the physics domain…

and to use these experiences with new users and domains to drive new requirements to evolve the OSG.

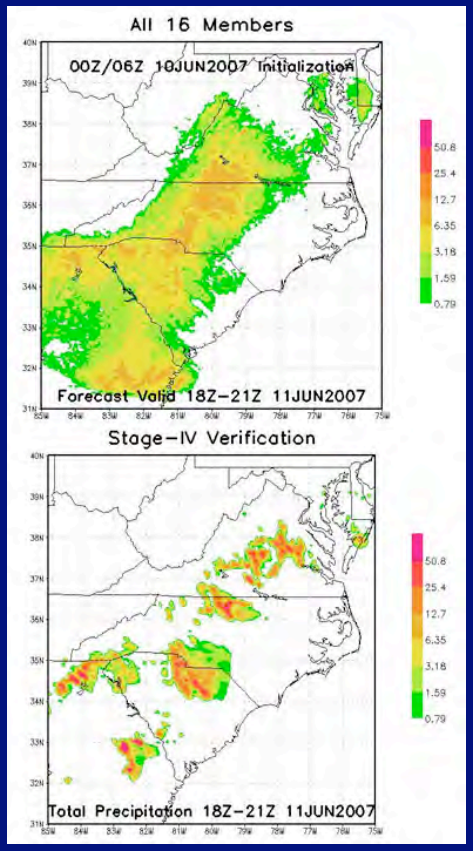

16 member ensemble run for fine grained (4km) forecast using Weather Research Forecast application. Resulted in Publication.

John McGee - Engagement Coordinator- "Best are applications that are pleasingly parallel & easily chunked into various sizes for distribution to the remote resources. We have most easily found researchers ready for big runs with these characteristics in: biochemistry, genetics, information & library science."

Kuhlman: Rosetta protein modelling

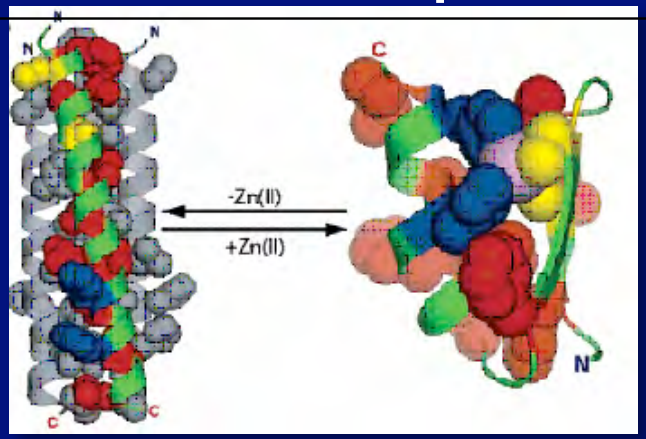

Luettich: Coastal Modeling

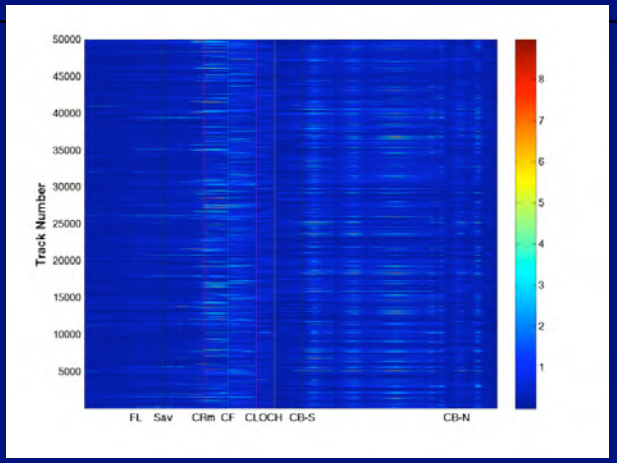

Nov-1-2007 through Jan-15-2008

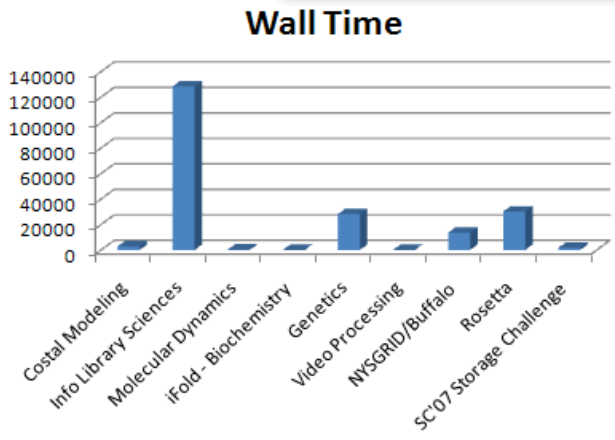

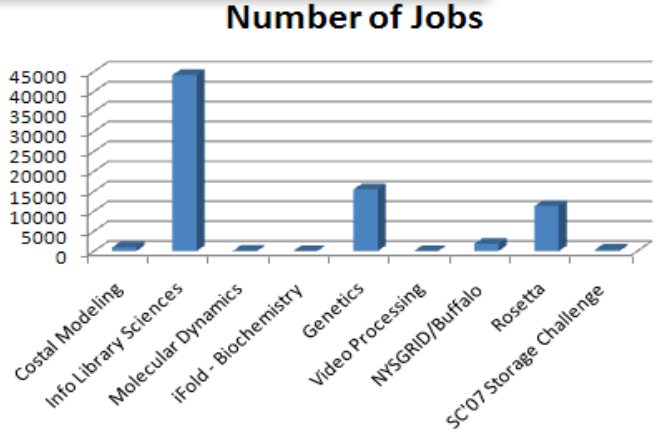

Integrating Campus, Local & the National Grids - some examples

Windows cluster across Clemson campus interfaces to OSG through single Linux gateway node.

- Machines in 27 different locations on Campus ~1,700 job slots; >1.8M hours served in 6 months ; users from Industrial and Chemical engineering, and Economics;

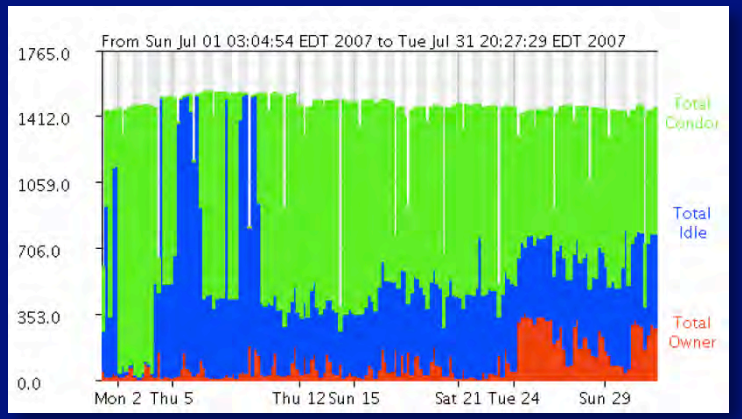

Department clusters at Madison (Grid Laboratory Of Wisconsin) interfaced to provided local sharing with automated “condor upload” capability to OSG.

Campus grid of Fermilab clusters (FermiGrid) provides internal sharing using OSG s/w & protocols as well as use of/by OSG.

New York State Grid of local universities internally uses OSG s/w, with sites optionally registered with OSG.

A Challenge: Communication & Governance

It took us a year to organize ourselves with by-laws, management plan, project structure;

Meetings, twiki, mail-lists, documents, chat rooms - & still people “don’t know whats going on”;

Policies & information collection need to balance open-ness with privacy & quality.

Broader communication: International Science Grid This Week (iSGTW) e-newsletter

The Value of OSG

Improving efficiencies of resource use through dynamic sharing and load-balancing.

Increasing return on US investments in physics and computer science.

Encouraging multi-disciplinary approaches through collaborative work .

Succeeding in US competitiveness with EU (EGEE) peers.

Broadening the use & expertise in cyber-infrastructure for science and research.