


The Open Science Grid

Ruth Pordes

OSG Executive Director,
Fermilab

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 Cyber-Infrastructure and Office of International Science and Engineering Directorates.

The OSG Vision

Transform compute and data intensive science through a national cyber- infrastructure that includes from the smallest to the largest organizations.

Bridge local, national, & international cyber-infrastructures.

Engage communities, facilities and software providers working together towards the common goals.

The OSG Reality

Broad community building through a Consortium. Support of end-to-end distributed systems for research communities through a funded Project collaborating with the stakeholders.

Deliverables, schedule, priorities driven by the stakeholder members of the governing Council and "project managed" through the Executive Board.

Commitments to the LHC experiments, & to the World Wide LHC Computing Grid Collaboration, subject to the agreements with US ATLAS and US CMS.

OSG is a funded project

Starting in September 2006 OSG is funded for five years by DOE and NSF.

Funding covers 34 FTEs at 18 institutions.

Joint Oversight Team from 6 funding agency program offices.

Mix of physics and computer science in management roles.

Science Advisory Committee of 4 senior scientists.

OSG Today

75 Processing Resources spanning the US, 4 in Brazil, 1 site each in Mexico, UK, Taiwan, Korea, providing accessibility to ~30,000 cores used subject to the local owners policies.

15 Storage Resources providing accessibility to ~4 Petabytes disk with static allocation by local owners.

75,000 jobs/day

240,000 CPUWallclockhours/day;10,000 batch slots dc

CPUUsage / Wallclockhours ~0.75

100 Terabytes/day data movement

Effectiveness Measures

Failures measured at “Grid interface” vary from 20-50%.

Each new VO has to debug each site they use.

Automated resource selection for only a few.

Sharing of storage resources embryonic.

Information publishing correct at ~50% level.

Published accounting: 50% for processing; 0% for storage.

OSG Tomorrow

- LHC 100,000 jobs/day, data transfers of >1 Petabyte/day (2008).
- Support STAR, LIGO “x10” increase in processing by 2009.
- Continue support for CDF, D0,Neutrino physics.
- Two new (non-physics) communities to production usage per year.
- Policies to manage resource oversubscription.
- x2 in accessible resources by end of 2009.

The OSG Context: a Distributed facility interacting with peers

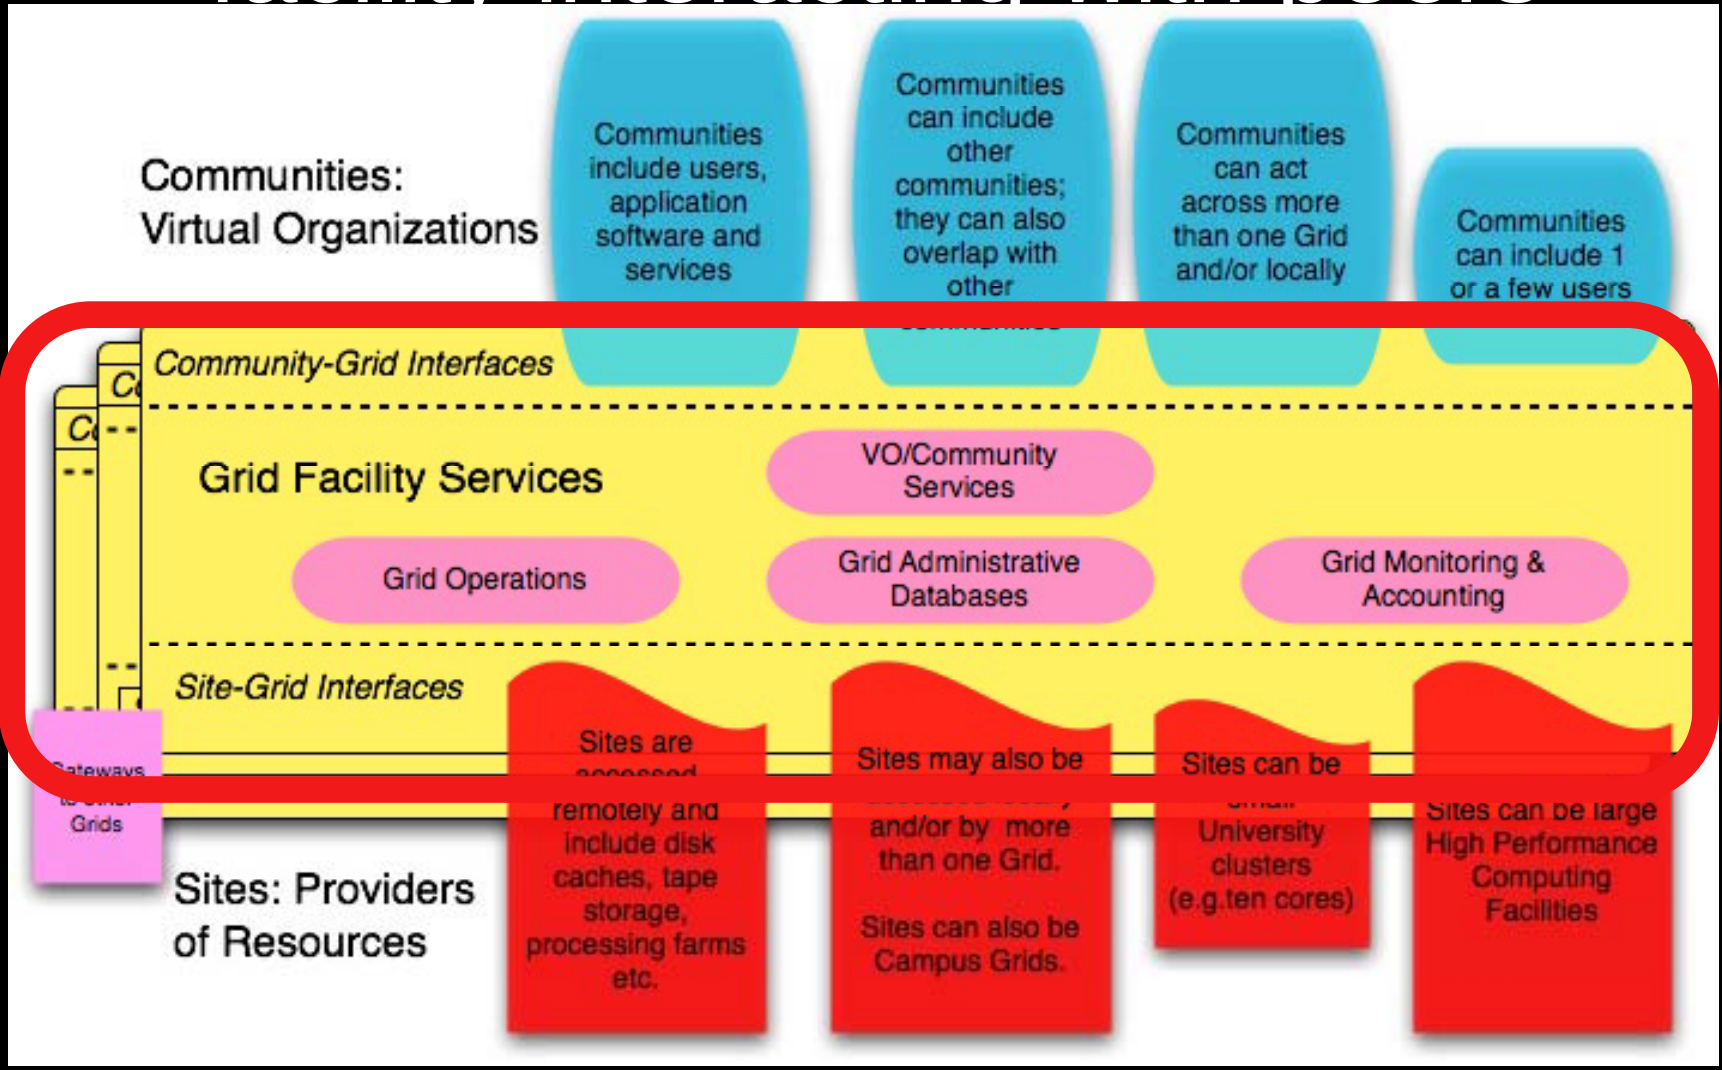

The OSG Context: Community systems across shared resources

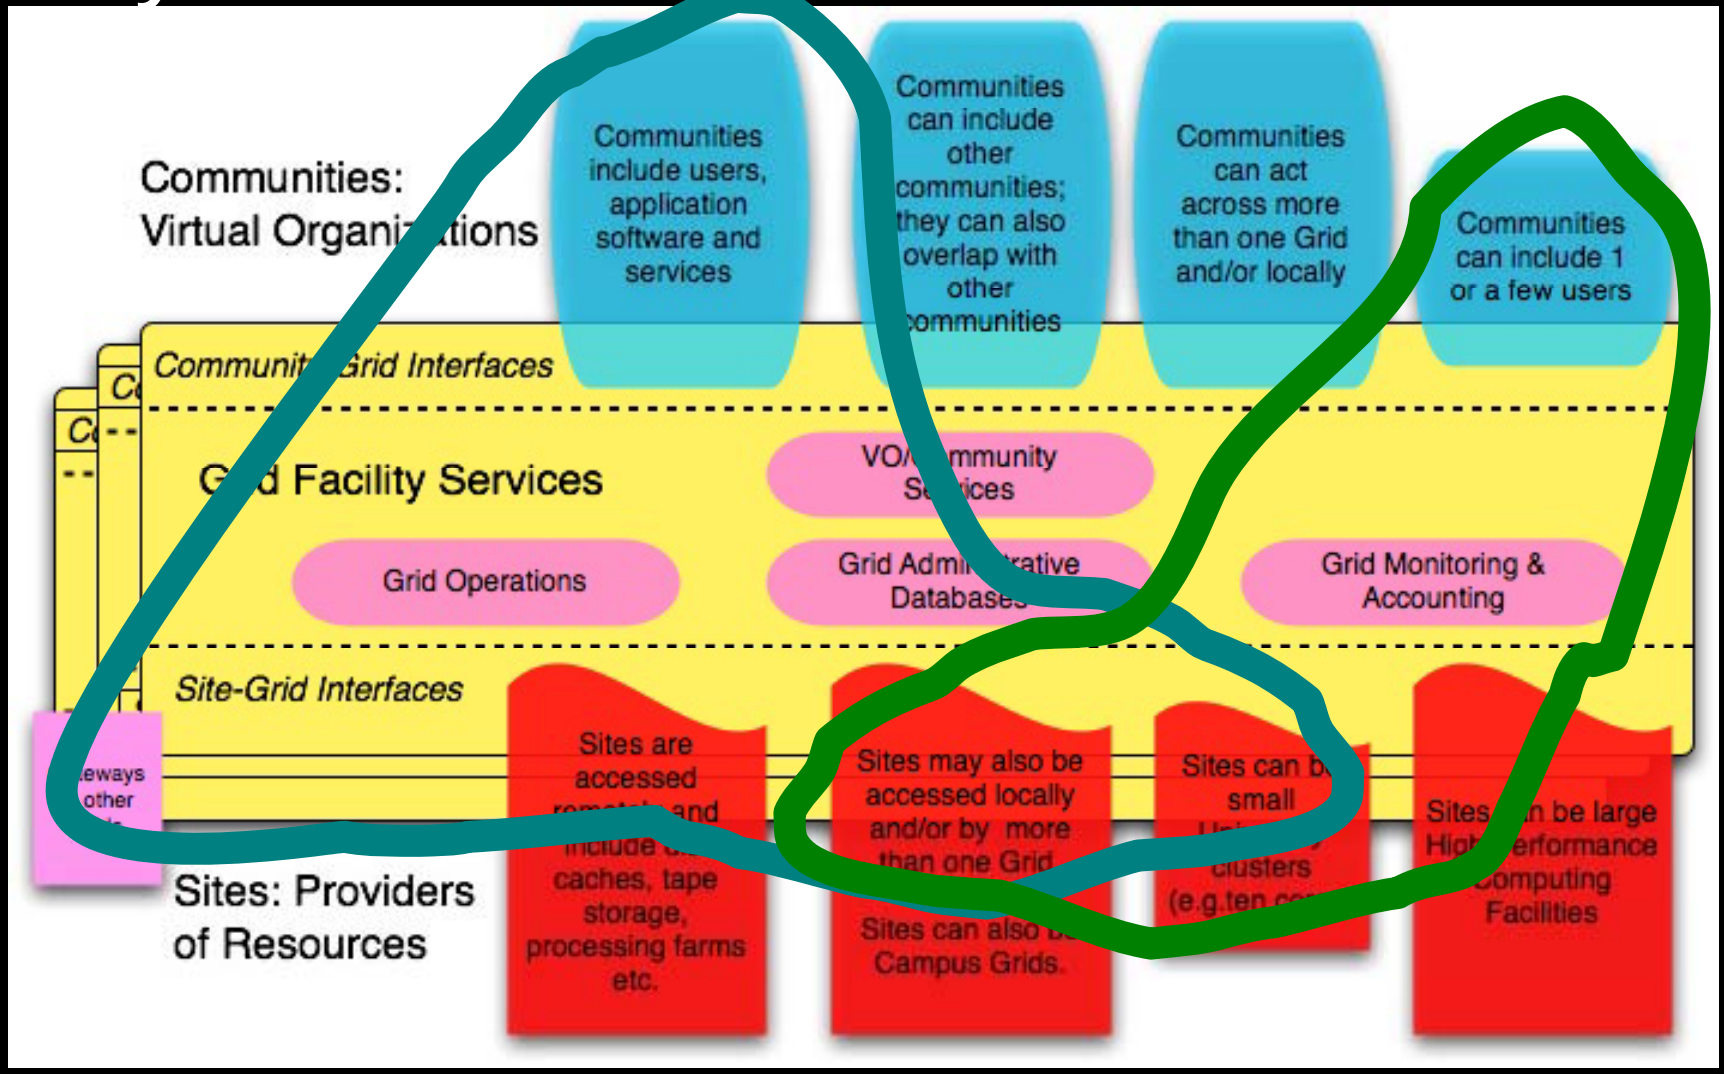

OSG Does Not

Own the Resources: The farms and storage are contributed by the Consortium members. Use commodity (and research) networks.

Develop the Software: middleware and applications are developed by contributors and external projects.

Make a “one size fits all”: We define interfaces which people can interface to and provide a reference software stack which people may use.

OSG Activities Include

Operating the distributed infrastructure & grid-wide services.

Support for site administrators and users.

Release, deployment and support of the Virtual Data Toolkit.

Provision of operational security and policies.

Troubleshooting “swat-team” for user and system problems.

Integration and testing of new software on the integration grid.

Engagement of new communities, education of students and the workforce.

Extensions in the capability and scale of the infrastructure.

OSG Implementation Architecture: Functional Interfaces

Storage Management: SRM.

- 2 implementations of common community standard

Data Transfer: GridFTP

- 3 implementations of OGF standard

Local Data Access:

-- Environment variables on the site let VOs know where to put and leave files.

Job Execution: single client (Condor-G) multiplexing to multiple servers

- web service and non-web service GRAM, Condor-glide-ins.

VO Management:

- 1 implementation based on EGEE VOMS interface.

Accounting:

OSG Implementation Architecture: Software Stack

OSG Virtual Data Toolkit (VDT) packaging and distribution of ~40 components, built for ~15 Linux versions.

Collections for Compute Element, Processing Worker Node, Storage Element, VO management, Client.

OSG & WLCG

OSG Resource Information to WLCG Information Collectors (BDII) filtered through central collector.

Service and Resource Validation (SRV) provided to WLCG in next OSG release.

Accounting information automatically uploaded based on configuration files.

Transparent data movement across global facility.

Transparent job scheduling across global facility.

Horizontal and Vertical Integrations

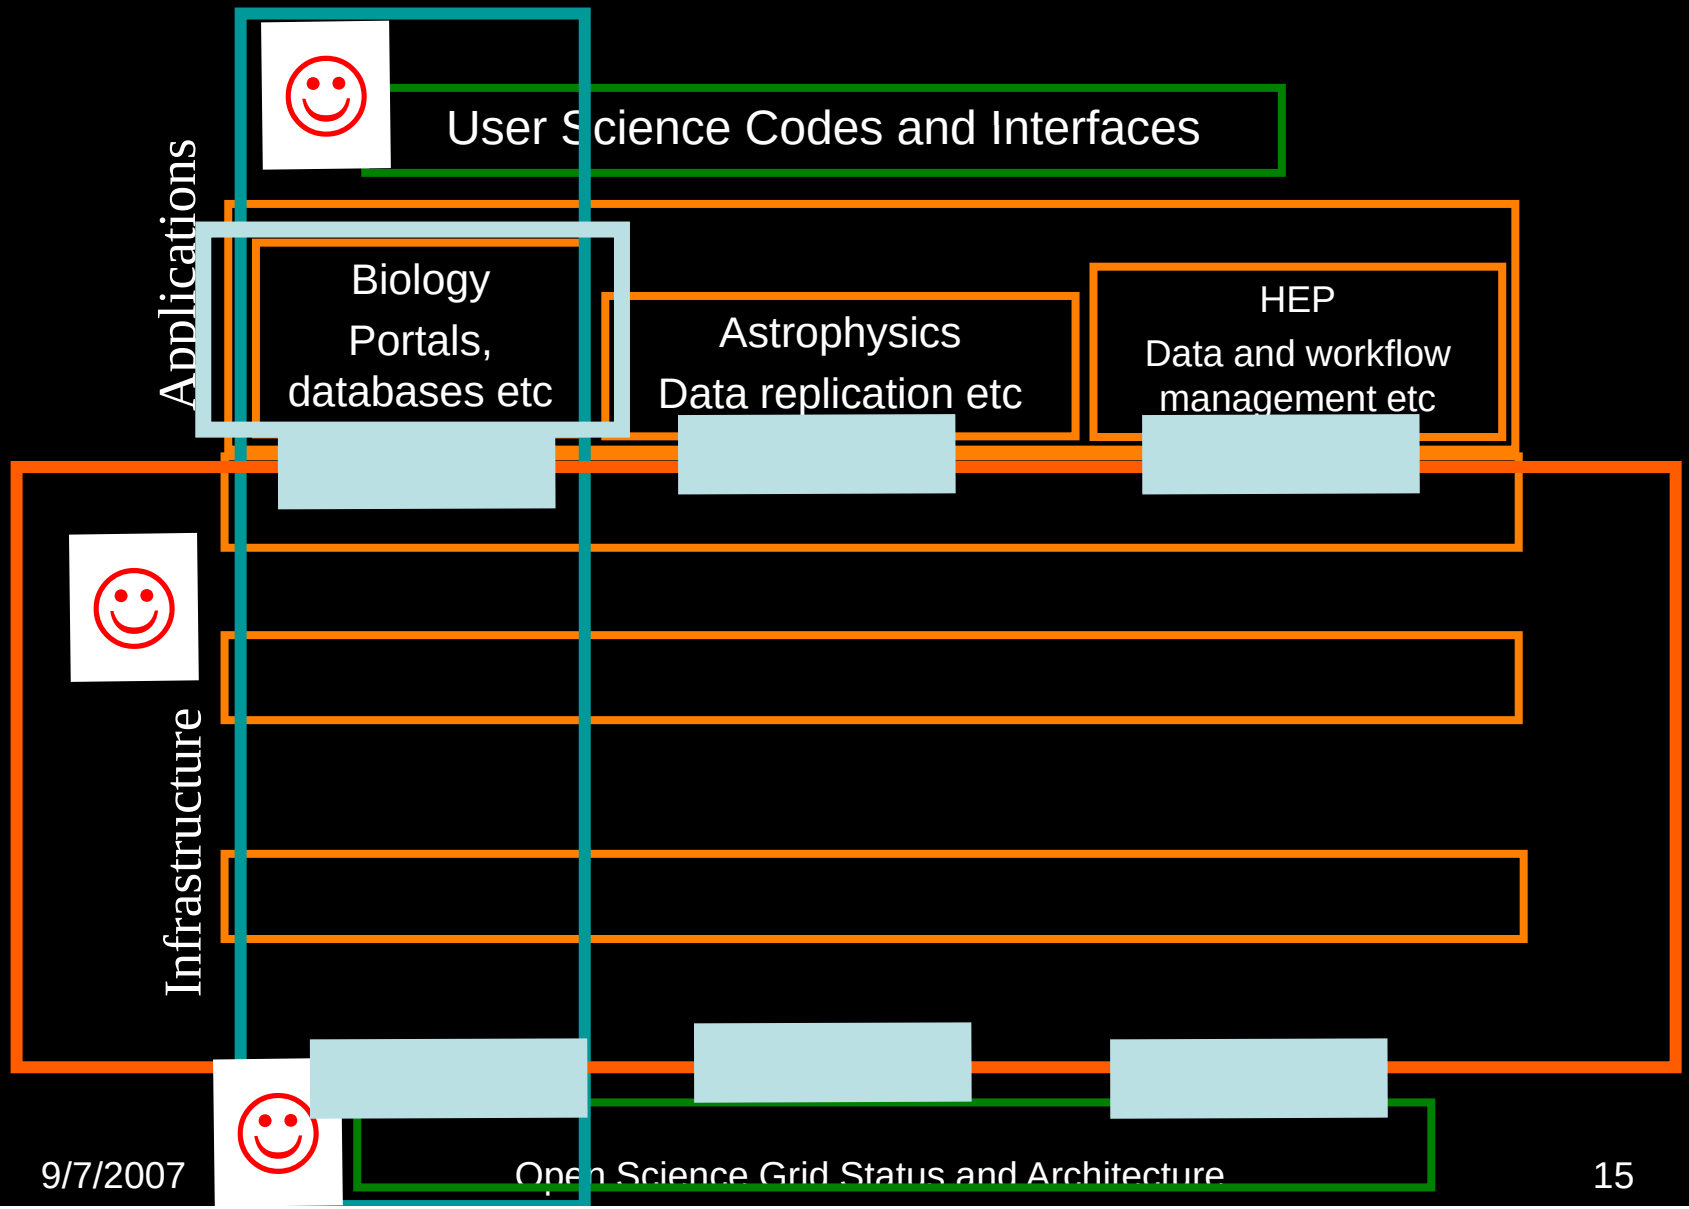

Software Release Timetable

Major release every ~six months.

OSG 0.8.0 in provisioning testing now.

- Scheduled release in September
- First release with full support from VDT for storage services.
- Most storage services at SRM V1.1

OSG 1.0 for 2/2008

- Release for LHC data taking.
- Full support for SRM V2.2
- Expect minor version releases after this.

Rely on External Software Providers

Model is open source software. Groups -- communities, projects -- contribute to the software for value gained for them.

Software questions for the next year

VDT Clients for AIX, MacOSX.

Windows sites now on OSG
-- what s/w support should OSG provide?

New communities demand MPI
-- what s/w support should OSG provide?

New communities need licensed software.
-- how will we provide this?

OSG Architecture: Opportunistic Use

Resources support opportunistic use of otherwise unused storage and processing for multiple OSG communities.

Processing handled through user VO/VO group/VO role mapping to Unix Accounts + standard batch queues set up by the site administrators.

Storage will be handled through support of SRM V2.2 features. (OSG driving additional functionality than that specified for WLCG).

Resource Management

Many resources are owned or statically allocated to one user community.

- The institutions which own resources typically have ongoing relationships with (a few) particular user communities (VOs)

Remainder of an organization's available resources can be "used by everyone or anyone else".

- organizations can decide against supporting particular VOs.
- OSG staff are responsible for monitoring and, if needed, managing this usage.

Our challenge is to maximize good - successful - output from the whole system.

Past Year has seen successes in Opportunistic Use

Accounting measures use by “non-owners”: currently around 20-30% of usage.

Biggest beneficiary was D0. (see the talk for the gain and the pain)

We are committed to improve the effectiveness and increase this service through LHC commissioning, data taking and analysis.

Rosetta from Kuhlman Lab

"peaks-and-valleys" usage pattern is typical though of course time constant varies.

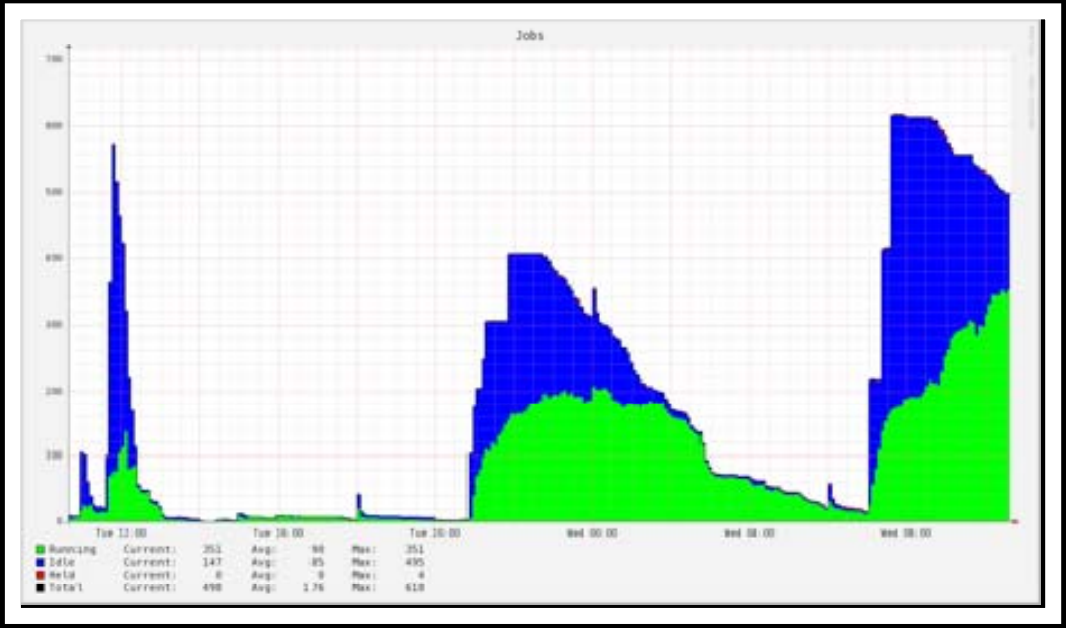

NanoHub

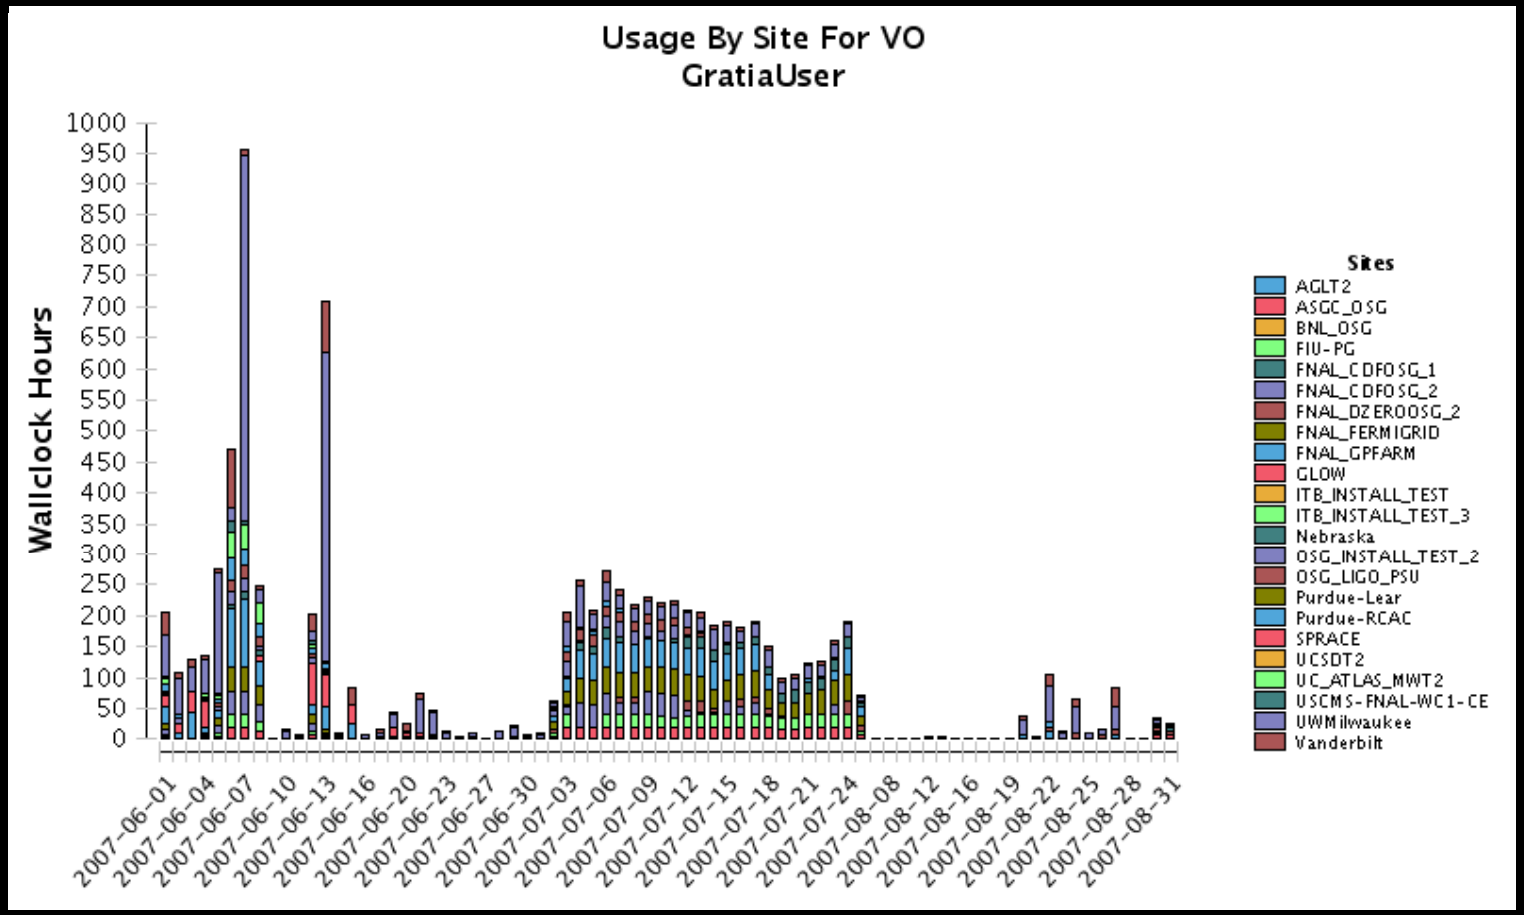
Usage By Site For VO
GratiaUser
Wallclock Hours
1000
950
900
850
800
750
700
650
600
550
500
450
400
350
300
250
200
150
100
50
0
2007-06-01
2007-06-04
2007-06-07
2007-06-10
2007-06-13
2007-06-16
2007-06-20
2007-06-23
2007-06-27
2007-06-30
2007-07-03
2007-07-06
2007-07-09
2007-07-12
2007-07-15
2007-07-18
2007-07-21
2007-07-24
2007-08-08
2007-08-12
2007-08-16
2007-08-19
2007-08-22
2007-08-25
2007-08-28
2007-08-31
Sites
AGLT2
ASGC_OSG
BNL_OSG
FIU-PG
FNAL_CDFOSG_1
FNAL_CDFOSG_2
FNAL_DZEROOSG_2
FNAL_FERMIGRID
FNAL_GPFARM
GLOW
ITB_INSTALL_TEST
ITB_INSTALL_TEST_3
Nebraska
OSG_INSTALL_TEST_2
OSG_LIGO_PSU
Purdue-Lear
Purdue-RCAC
SPRACE
UCSDT2
UC_ATLAS_MWT2
USCMS-FNAL-WC1-CE
UWMilwaukee
Vanderbilt

OSG Architecture: Campus Infrastructure: gateway to national grid

The OSG software stack aims to support:

Sharing of resources across multi-organization University Campus. (In the US this also involves a cultural change and model of trust, which OSG can help facilitate).

Single bi-directional points of entry between campus and national infrastructure allowing centralized local management, monitoring and support.

Gateways that enable protocol translation while maintaining performant end-to-end transport of data and jobs between infrastructures.

OSG is on the Road!

Significant % of ATLAS and CMS data and job throughput supported by the infrastructure.

US ATLAS and US CMS make significant contributions in terms of resources and effort.

STAR, CDF, D0, LIGO,MiniBoone all acknowledge contributions and themselves contribute.

Many talks and posters refer to work on and contributions to the Consortium.