
# Open Science Grid Overview

D. Olson, LBNL
5 April 2007
DOSAR Workshop

# Contents

- The Consortium
- The Project
- The Grid
- The Plans (The Process)
- The Conclusion

- Some slides for your later perusal.

# The Consortium

- Chartered December 2004
  - The vision of the Open Science Grid Consortium (OSG) is one of a persistent national grid infrastructure for large scale US science: the Open Science Grid.
- Founding members
  - ANL, BNL, CCR SUNY-Buffalo, FNAL, JLab, LBNL, SLAC, US Atlas, BaBar, BTeV, CDF, US CMS, D0, GRASE, LIGO, SDSS, STAR, US Atlas s&c, US CMS s&c, Condor, Globus, PPDG-Common, SRM, VDT, Griphyn, iVDGL, PPDG

# The Project

6 mo. into 5 year funding from NSF & DOE at ~ $5M/yr

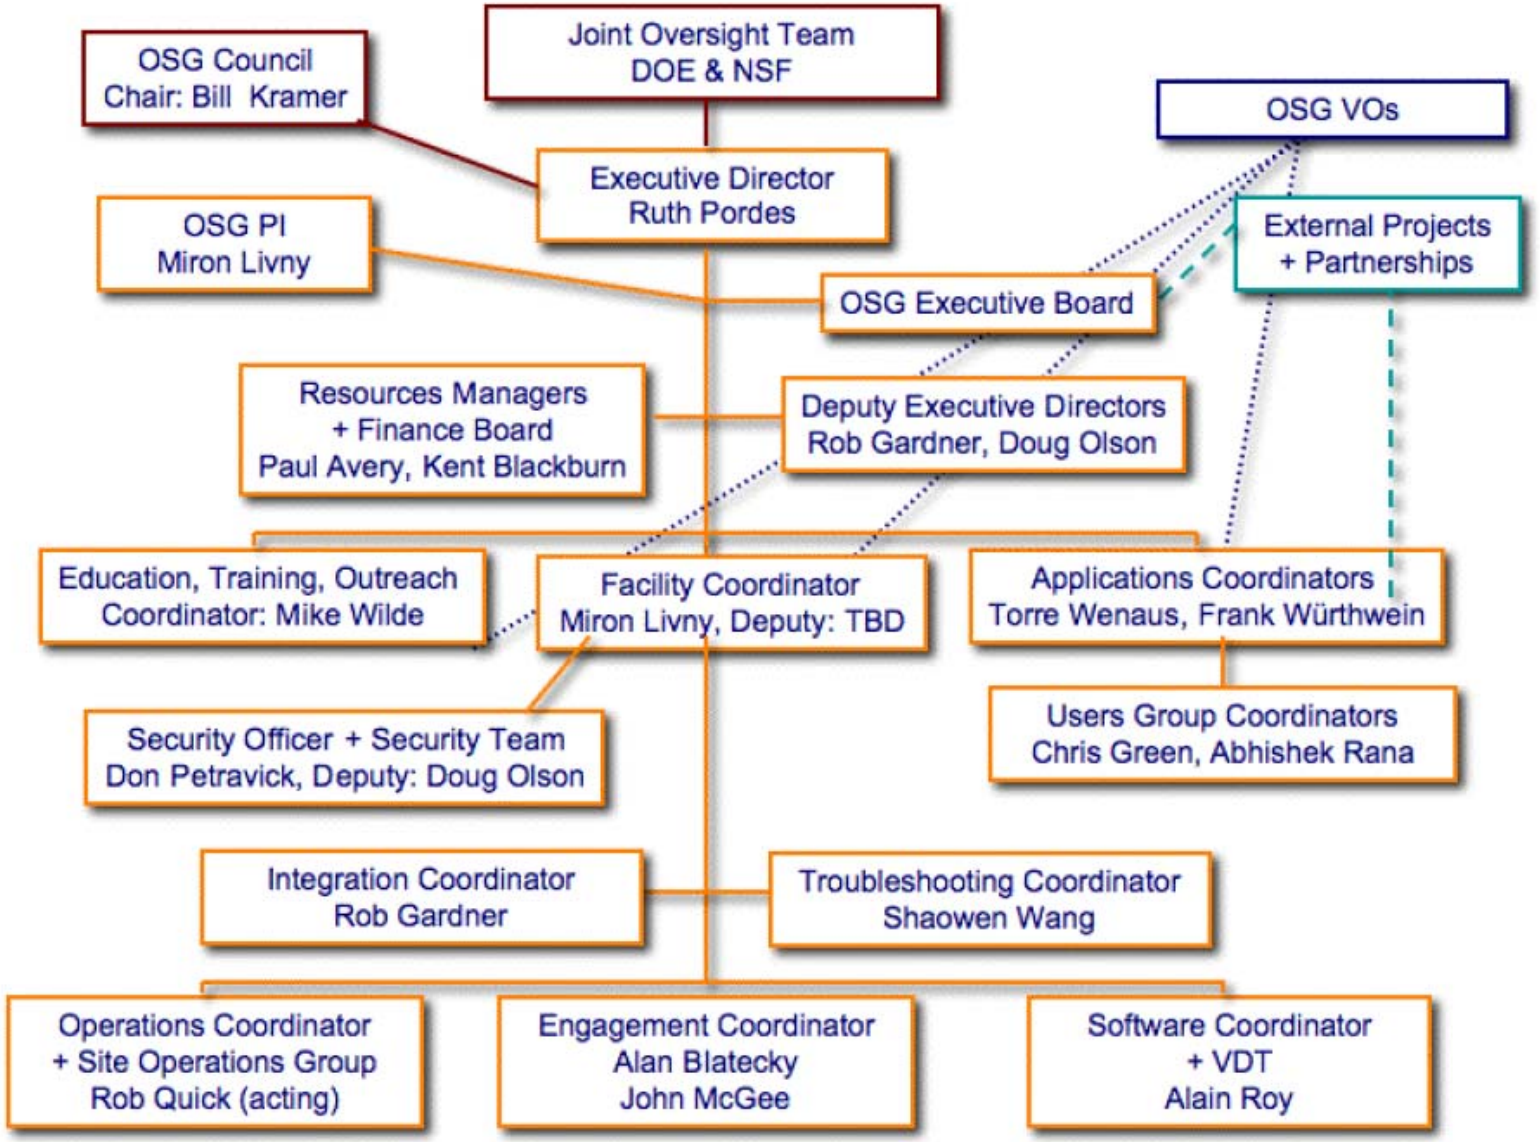

# External Projects Contributing to OSG, represented on the Executive Board

| | |
|---|---|
| Center for Enabling Distributed Petascale Science (CEDPS) | Jenny Schopf |
| Community Driven Improvement of Globus Software (CDIGS)/Globus | Dan Fraser |
| Condor | Todd Tannenbaum |
| dCache | Patrick Fuhrmann |
| Data Intensive Science University Network(DISUN) | Frank Würthwein |
| Disk Resource Manager (DRM) | Alex Sim |
| ESNET | *under discussion* |
| Internet2 | *under discussion* |
| LIGO Physics and the Information Frontier | Patrick Brady |
| OSG Accounting | Philippe Canal |
| OSG Privilege/Authorization | Gabriele Garzoglio |
| UltraLight/advanced networks | Frank Lingen |
| U.S. LHC + US representation to the WLCG | Michael Ernst, Ian Fisk, |

# Partners

The OSG works actively with partners -- including grid and network organizations as well as international, national, regional and campus grids -- to create a grid infrastructure that spans the globe.

- Data Intensive Science University Network (DISUN)
- Enabling Grids for E-SciencE (EGEE)
- Grid Laboratory of Wisconsin (GLOW)
- Grid Operations Center at Indiana University
- Grid Research and Education Group at Iowa (GROW)
- Nordic Data Grid Facility (NorduGrid)
- Northwest Indiana Computational Grid (NWICG)
- TeraGrid
- Texas Internet Grid for Research and Education (TIGRE)
- TWGrid (from Academica Sinica Grid Computing)
- Worldwide LHC Computing Grid Collaboration (WLCG)

- Discussion at recent all hands meeting with DOSAR, NYSGrid, DES, GIN/Pragma

# The Grid

- The Users
  - VOs
- The Facility
  - Operations
    - interface to sites and support centers
  - Software
  - Integration
  - Troubleshooting
  - Integrated Security

# It takes VOs to make OSG work!

| | |
|---|---|
| cdf | Collider Detector at Fermilab |
| cms | Compact Muon Solenoid |
| compbiogrid | CompBioGrid |
| des | Dark Energy Survey |
| dosar | Distributed Organization for Scientific and Academic Research |
| dzero | D0 Experiment at Fermilab |
| engage | Engagement |
| fermilab | Fermi National Accelerator Center |
| fmri | Functional Magnetic Resonance Imaging |
| gadu | Genome Analysis and Database Update |
| glow | Grid Laboratory of Wisconsin |
| gpn | Great Plains Network |
| grase | Group Researching Advances in Software Engineering |
| gridex | Grid Exerciser (GEx) |
| grow | Grid Research and Education Group at Iowa |
| gugrid | Georgetown University Grid |
| i2u2 | Interactions in Understanding the Universe Initiative |
| ligo | Laser Interferometer Gravitational-Wave Observatory |
| mariachi | Mixed Apparatus for Radar Investigation . . . |
| nanohub | nanoHUB Network for Computational Nanotechnology (NCN) |
| nwicg | Northwest Indiana Computational Grid |
| osg | Open Science Grid |
| osgedu | OSG Education Activity |
| sbgrid | Structural Biology Grid |
| sdss | Sloan Digital Sky Survey |
| star | Solenoidal Tracker at RHIC |
| usatlas | United States ATLAS Collaboration |

# Distributed Operations

- Grid Operations Center at IU
- Support Centers for all VOs, sites, services
- Ticket tracking and routing
- Metrics
- Security processes

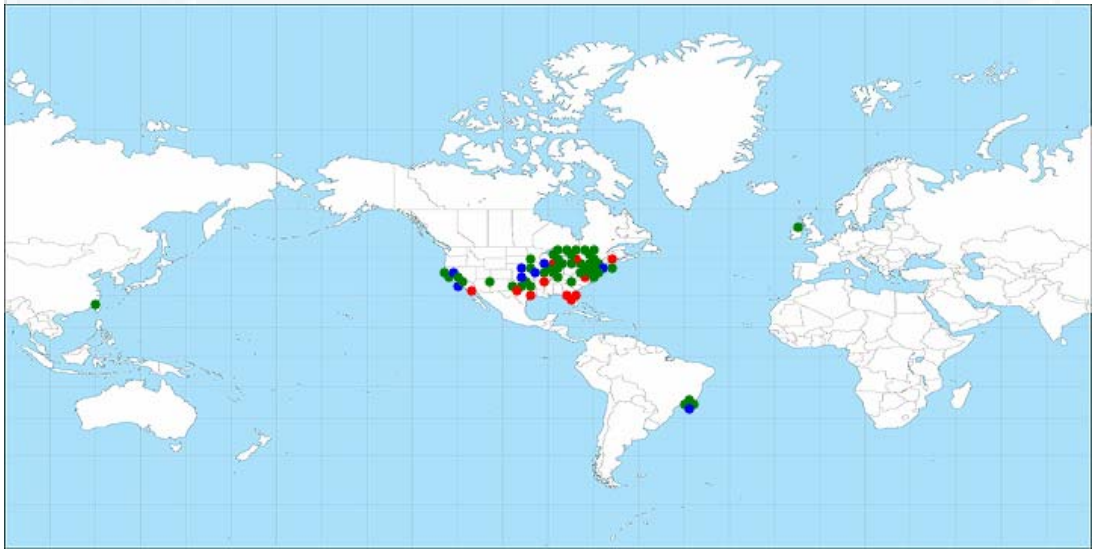

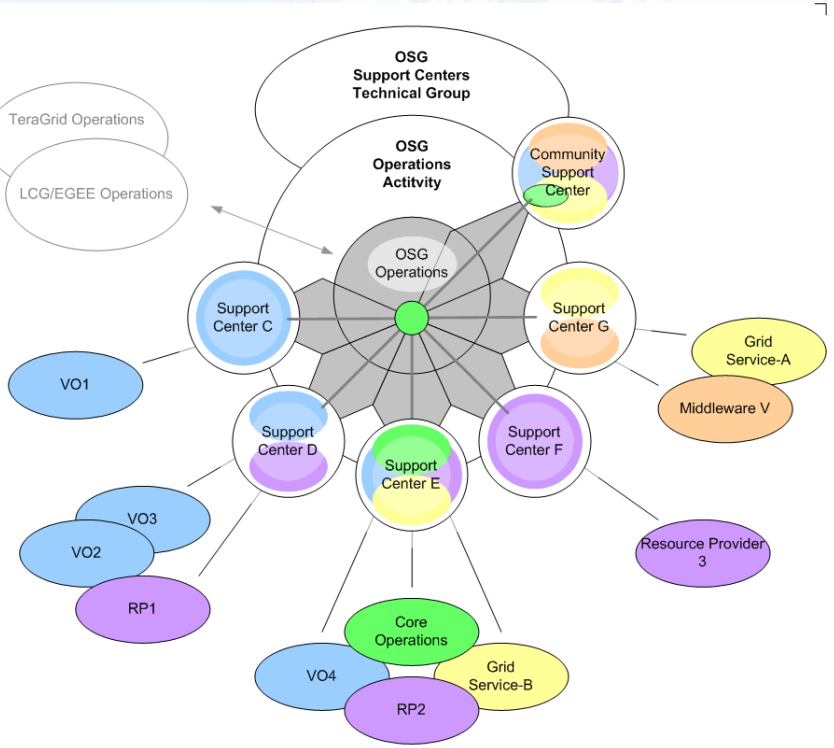

# OSG Middleware

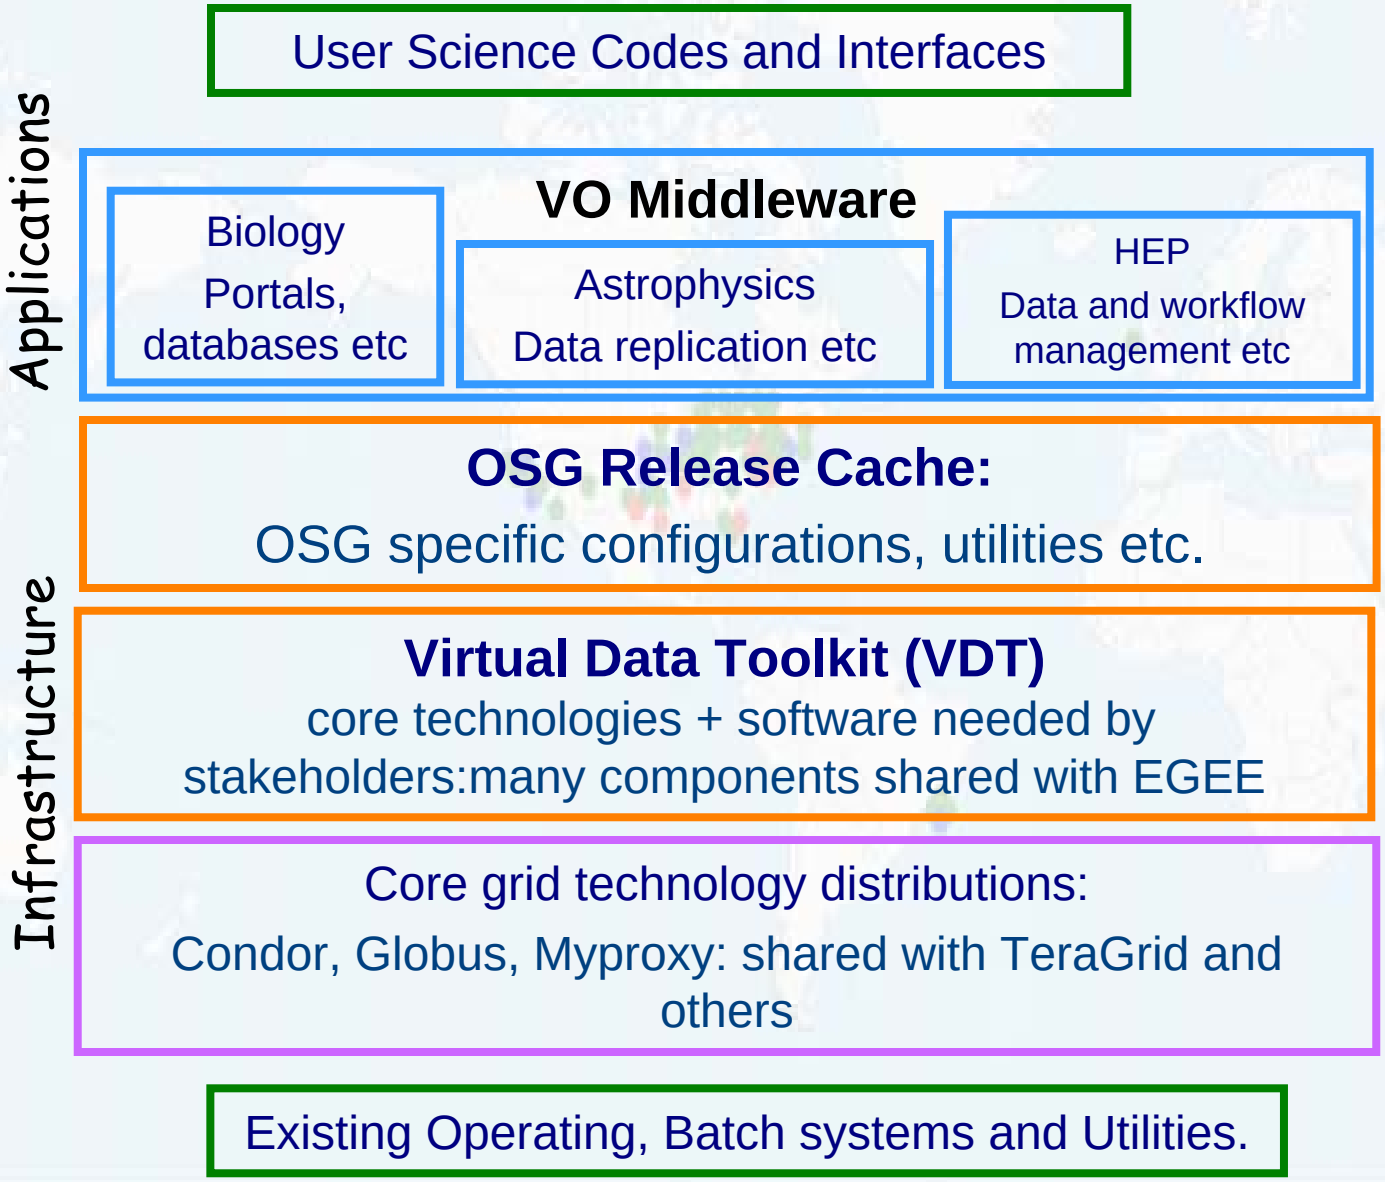

# Integration

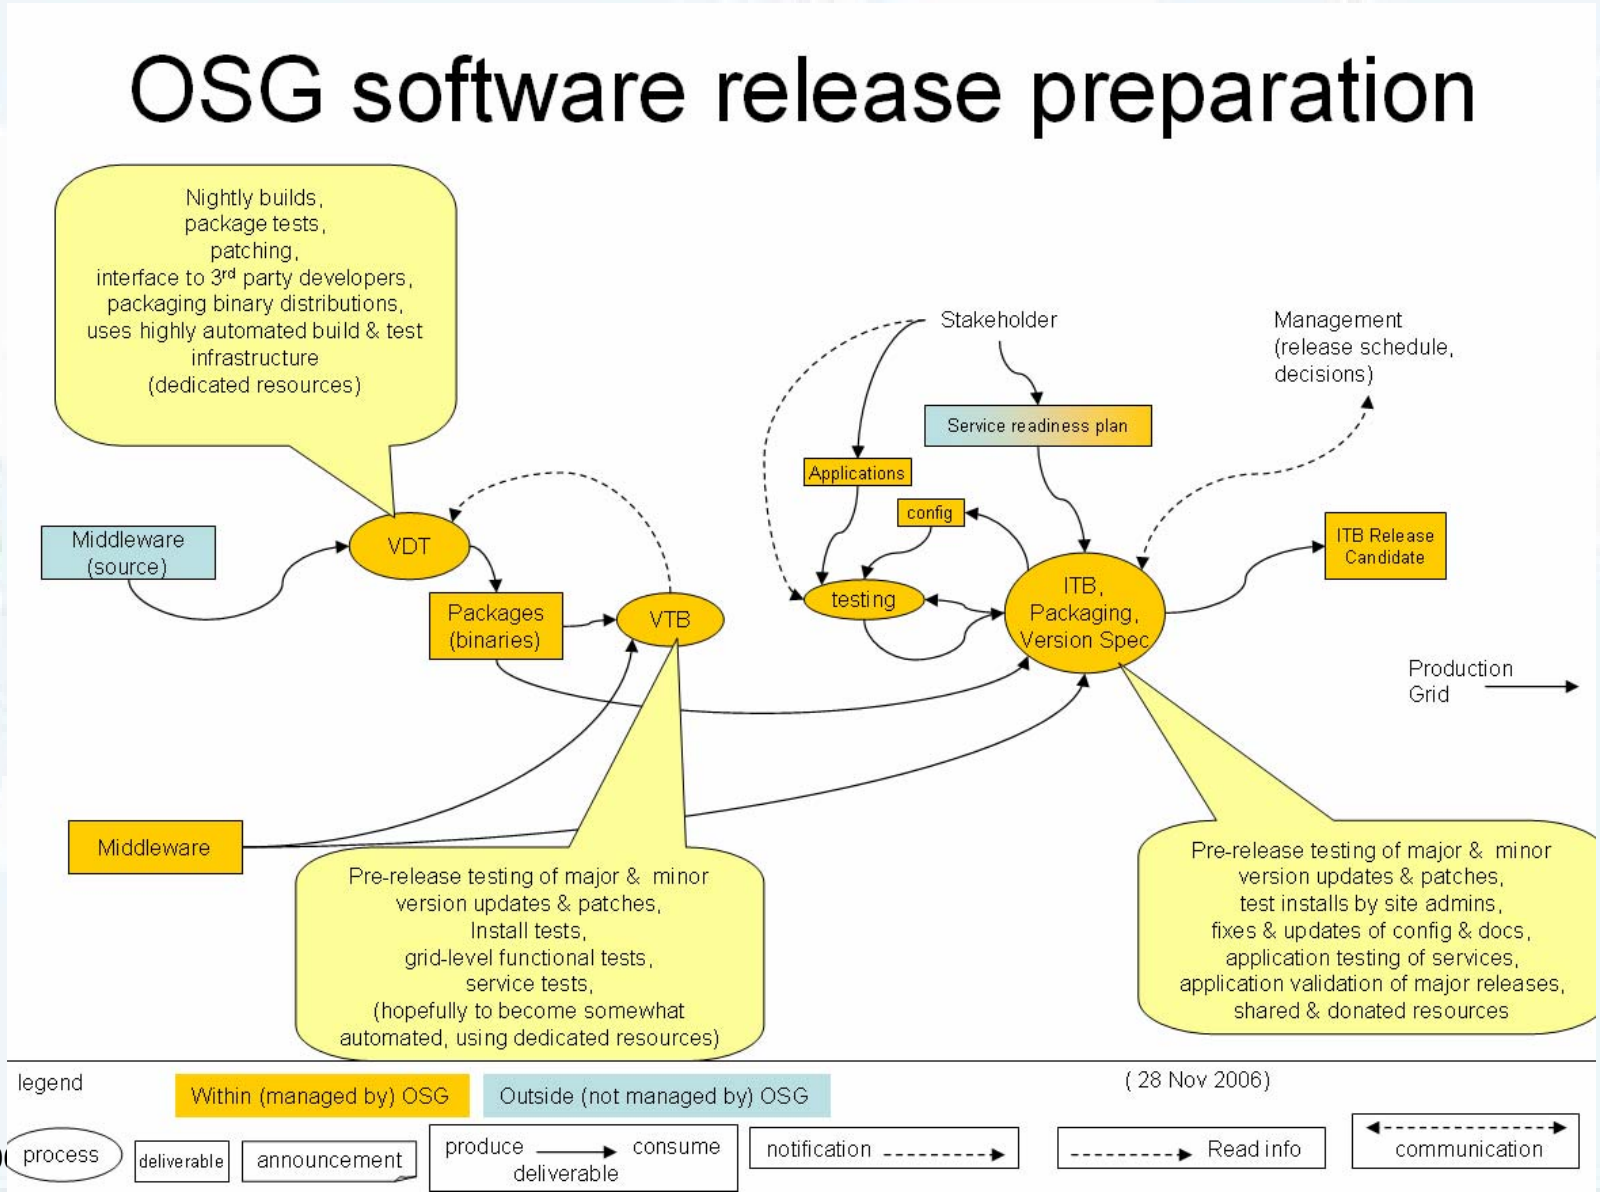

# Troubleshooting

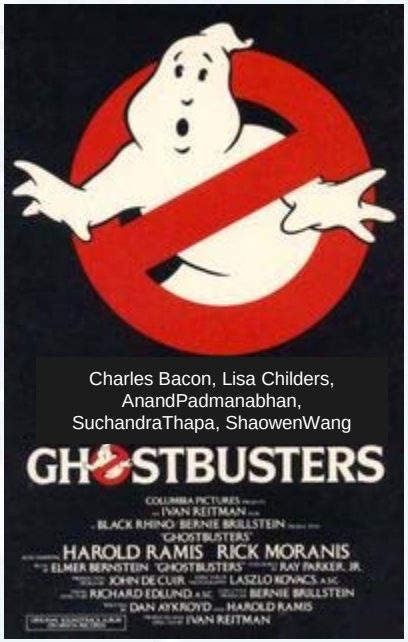


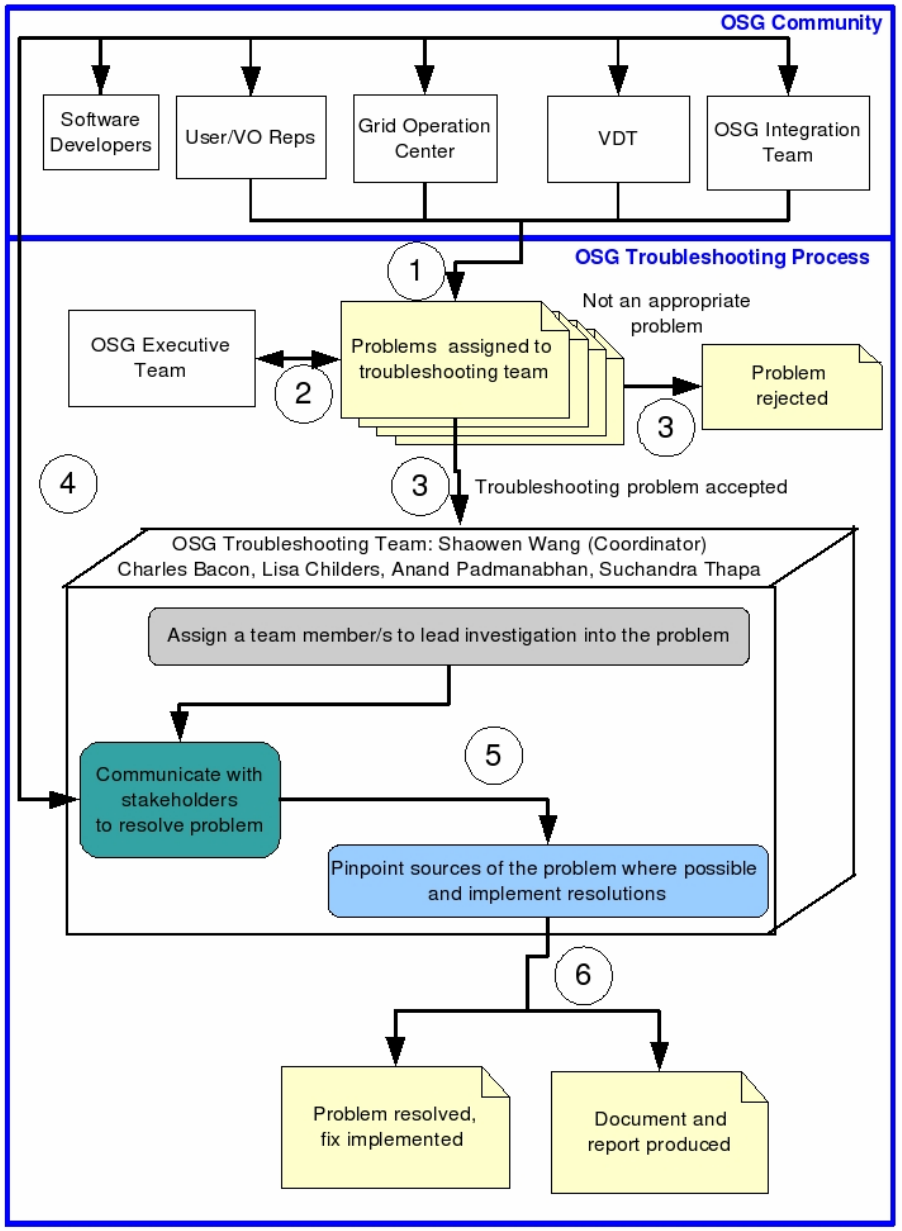

# What's New in OSG 0.6

- Gratia – accounting system
- CEMon – aids resource selection
- Increasing emphasis on storage elements
  - but still no SEInstallGuide
- Squid (web proxy caching service, optional)
- See full documentation at
  https://twiki.grid.iu.edu/twiki/bin/view/ReleaseDocumentation/DocumentationTable

Open Science Grid

MonAlisa has been used to get an accounting-like view.

Gratia will provide the accurate resource usage accounting.
Still some issues to work out.

http://gratia-osg.fnal.gov:8880/gratia-reporting/

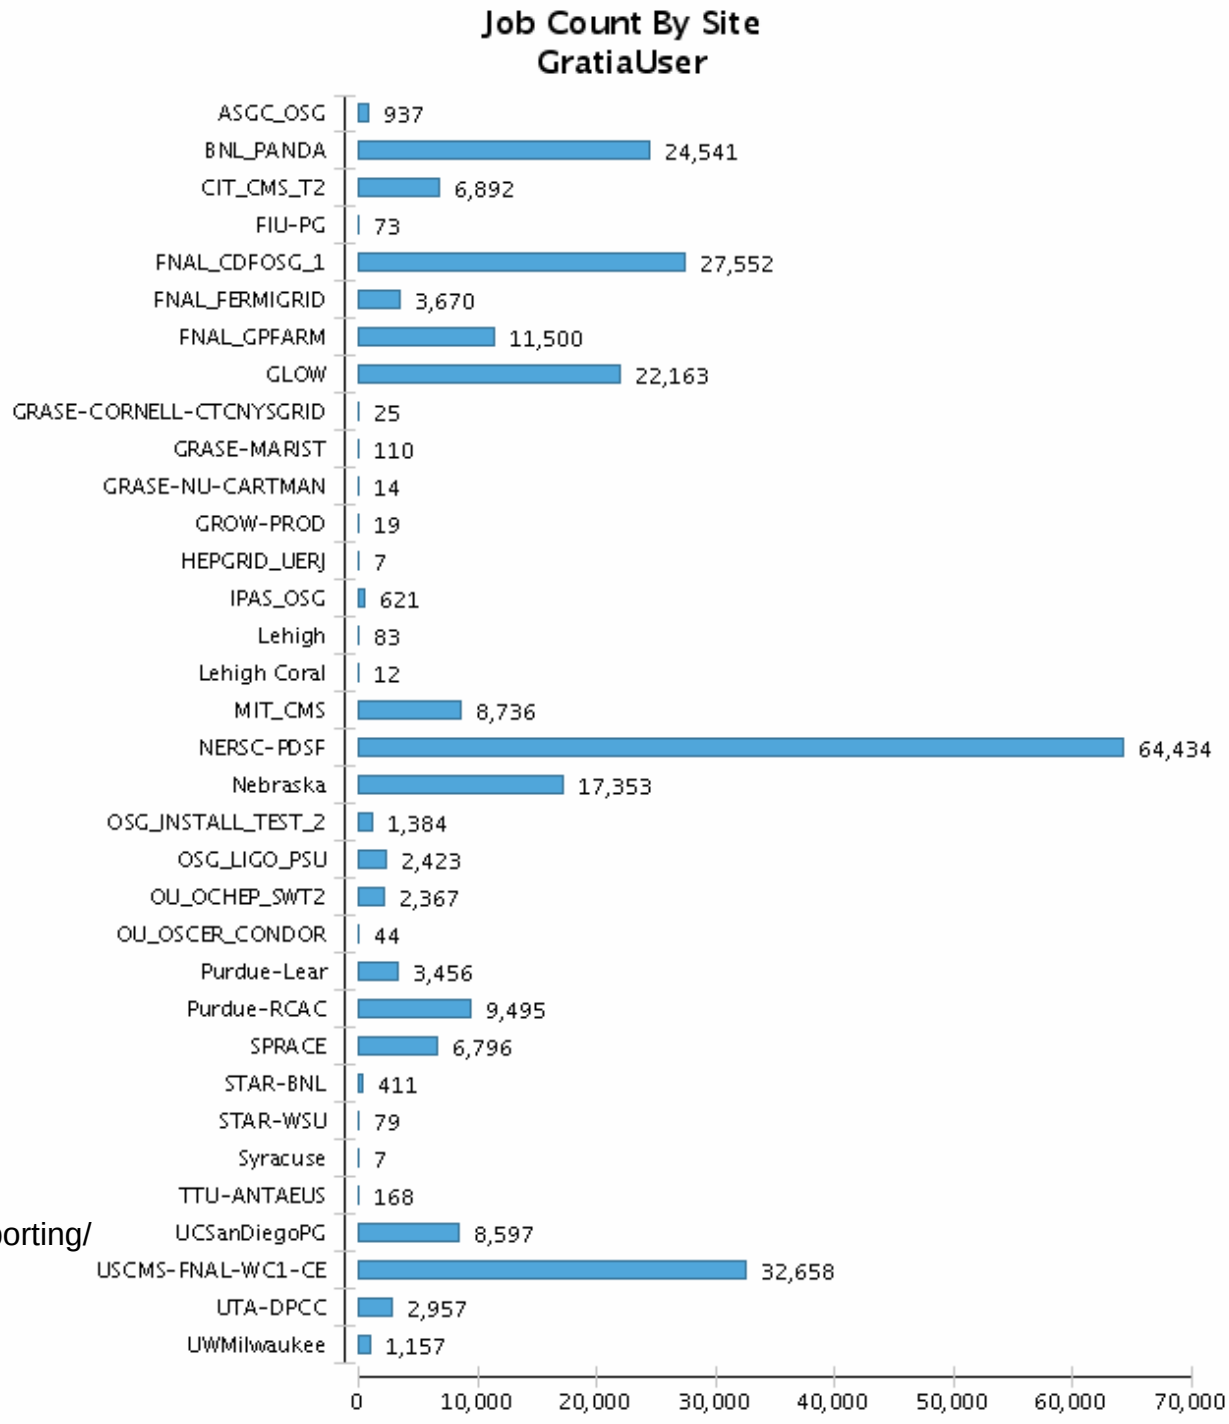

# CEMon/ReSS

https://twiki.grid.iu.edu/twiki/bin/view/ResourceSelection/

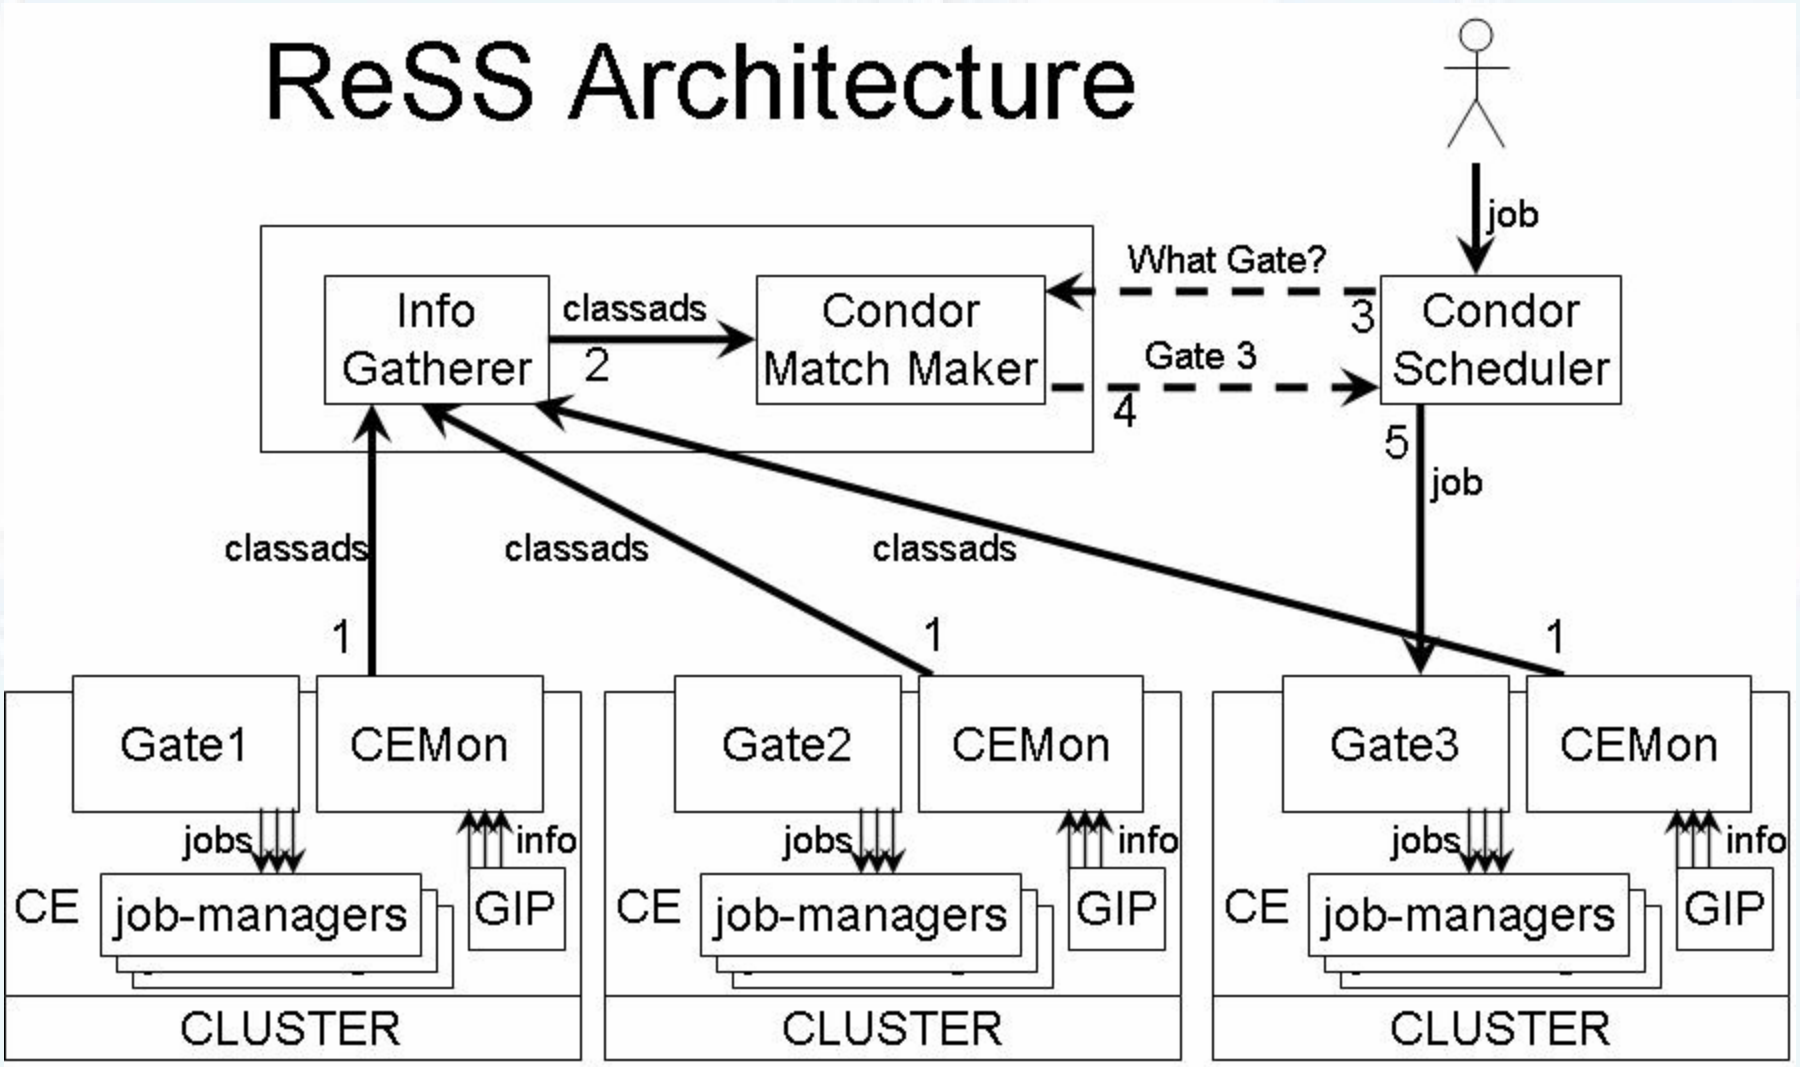

# The Plans (The Process)

- Next release (OSG 0.8) expected Sept. 2007
  - VOMRS
    - Registration service, can be used with VOMS
  - Condor NFS Lite Jobmanager
    - No shared FS between CE and WN
  - dCache
  - New DRM (Java version, BeStMan)
  - Various VDT improvements
  - Final contents depends on user requirements and effort

# Longer Term – Extensions & Contributing Projects

- Panda for general OSG users
- Work with CEDPS
  - Center for Enabling Distributed Petascale Science
  - integrate logging/troubleshooting tools
    - syslog-ng entering ITB now
  - Storage management/data placement
  - Scalable Services (Globus Workspaces)


# Panda Dashboard

Configuration

6 min old Update

Panda monitor

Quick guide, twiki

**Jobs** - search
Recent running, activated, waiting, assigned, defined, finished, failed jobs
Select analysis, production, test jobs

**Quick search**
Job
Dataset
Task
File

**Summaries**
Blocks: days
Errors: days
Nodes: days
Daily usage

**Tasks** - search
Generic Task Req
EvGen Task Req
CTBsim Task Req
Task list
Task browser

**Datasets** - search
Dataset browser
New datasets
Panda subscriptions
All subscriptions

**Datasets Distribution**
AODs
RDOs

**Dashboards:** Production DDM AutoPilot Sites & Grids Analysis Physics data Usage & Quotas Plots ArdaDash

Not logged in. List users

## Panda Based Distributed Analysis Dashboard

Information and tools for distributed analysis with Panda

pathena supports user submission to Panda from the (p)athena command line

**Analysis jobs:** New-style listing of analysis jobs. Old-style listing is here if you prefer it (tell me why!). To look up a particular Panda job by ID use the quick search at left or click a PandaID in the job listing.

**Analysis users:** User list (also linked at top right, or above if you've logged in) shows analysis usage, ordered by most recent. From there you can go to your page (you're on the list if you've run a Panda job); if you 'log in' you'll get easier access to your page from a new menu at the top of the page.

**Groups:** Groups are supported to organize users by role, physics working groups etc. and support collaborative work, accounting rights etc. (Not much used yet.)

**Data access:** See the physics data page linked above

**Analysis job summary, last 12 hours** (Node details)

| Site | Nodes | Jobs | Latest | defined | assigned | waiting | activated | running | holding | transferring | finished | failed | tot | trf | other |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| All | 115 | 323 | 04-04 16:38 | 9 | 0 | 0 | 0 | 9 | 160 | 0 | 135 | 10 | 7% | 2% | 5% |
| ANALY_BNL_ATLAS_1 | 20 | 67 | 04-04 16:38 | 9 | 0 | 0 | 0 | 3 | 4 | 0 | 45 | 6 | 12% | 6% | 6% |
| ANALY_BNL_ATLAS_2 | 0 | 0 | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | | | |
| ANALY_LONG_BNL_ATLAS | 42 | 256 | 04-04 16:28 | 0 | 0 | 0 | 0 | 6 | 156 | 0 | 90 | 4 | 4% | 0% | 4% |
| ANALY_UTA-DPCC | 0 | 0 | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | | | |
| Unassigned | 0 | 0 | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | | | |

**Analysis job error report, last 12 hours** Job wall time: 105 hrs Error losses: trans: 0 (0.2%) panda: 0 (0.0%) ddm: 0 (0.0%) other: 7 (6.4%)

| Error type (type count) | Count | CPU-hrs | Latest | Code: | Description |
|---|---|---|---|---|---|
| **All** | defined:9 | assigned:0 | waiting:0 | activated:0 | running:9 holding:160 transferring:0 finished:135 failed:10 (6.9%) |
| exeErrorCode (3) | 3 | 2.8 | 04-03 16:21 | 1154: | *Unknown error code* |
| pilotErrorCode (4) | 1 | 4.0 | 04-04 06:35 | 1156: | Pilot could not recover job |
| pilotErrorCode (4) | 3 | 2.8 | 04-03 16:21 | 1158: | Reached maximum number of recovery attempts |
| taskBufferErrorCode (3) | 3 | 0.0 | 04-04 15:17 | 100: | Job expired and killed six days after submission (or killed by user) |
| transExitCode (6) | 3 | 0.2 | 04-04 15:06 | 20: | *Unknown error code* |

# CEDPS Logging

## A Case for Unified Logging

- Unified logging is needed for the Grid
  - Auditing
  - Forensics
  - Troubleshooting
- Example Use Cases:
  - My Grid Job failed. Why? Need to look at error logs on several hosts and several sites.
  - What is the list of hosts/resources that a given compromised user credential accessed in the past week

## syslog-ng

- Features:
  - Can filter logs based on level and *content*
  - Arbitrary number of sources and destinations
  - Provides remote logging
    - Can act as a proxy, tunnel thru firewalls
  - Execute programs
    - Send email, load database, etc.
  - Built-in Log rotation
  - Timezone support
  - Fully qualified host names
  - Secure via stunnel (http://www.stunnel.org)
    - allows you to encrypt arbitrary TCP connections inside SSL

## Sample Site Deployment

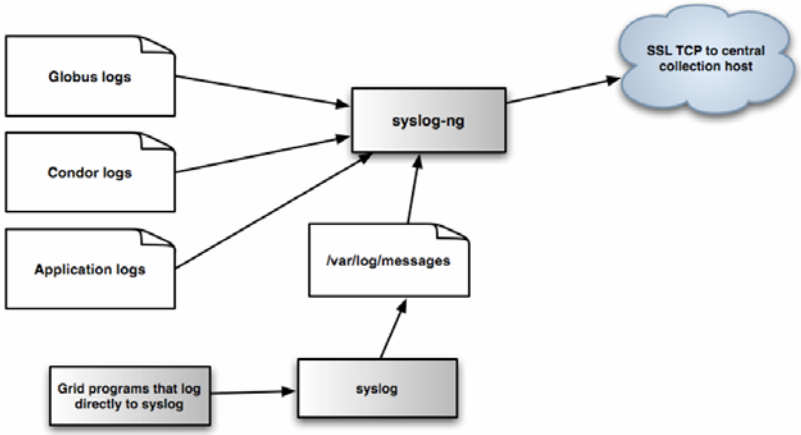


## Sample Grid Deployment

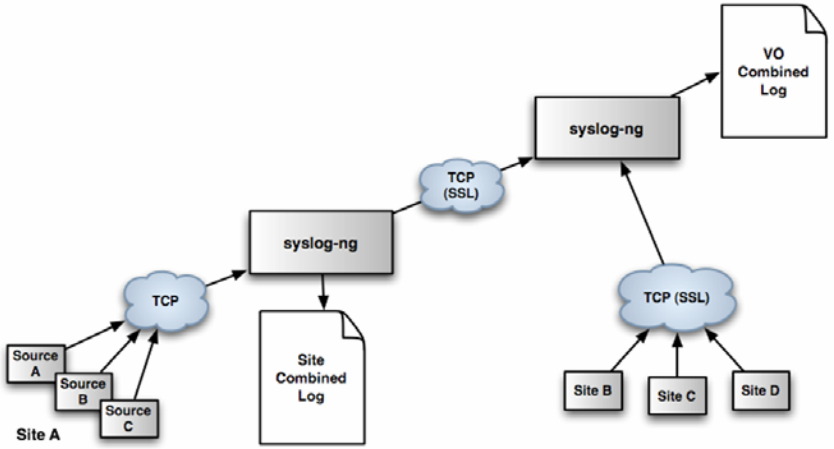


G for DO

# The Conclusion

- OSG is
  - open to collaboration, participation, cooperation (as you know).
  - driven by community stakeholder needs.
  - managed as a project with schedules, milestones, reporting
  - dedicated to achieving production national grid infrastructure for science.