



Open Science Grid

Open Science Grid School Operations Support and Tools

or

“What to Do When You Have Problems and Other Interesting Stuff”

**10 December - 2010
Sao Paulo, Brazil**

**Rob Quick
OSG Operations Coordinator
Indiana University**



Gedankenexperiment

- You have geographically distributed resource consumers (users) working under different administrative domains who want to run “grid enabled” jobs.
- You have geographically distributed compute resources operating under different administrative domains that allow “grid enabled” jobs to be submitted.
- Given these two positions are true... What next?



Economies of Scale

There are factors that cause a producer's average cost per unit to fall as the scale of output is increased. "Economies of scale" is a long run concept and refers to reductions in unit cost as the size of a facility and the usage levels of other inputs increase.

-Wikipedia

“The Open Science Grid (OSG) advances science through open distributed computing. The OSG is a multi-disciplinary partnership to federate local, regional, community and national Cyberinfrastructures to meet the needs of research and academic communities at all scales.”

OSG delivered across 81 sites

In the last 24 Hours

514,000 Jobs

1,018,000 CPU Hours

136,000 Transfers

933 TB Transferred

In the last 30 Days

13,568,000 Jobs

35,405,000 CPU Hours

41,911,000 Transfers

27,434 TB Transferred

In the last Year

158,110,000 Jobs

345,571,000 CPU Hours

515,600,000 Transfers


217,660 TB Transferred

What is Operations?

- Business Operations Definition

“Operations transform resource or data inputs into desired goods, services, or results, and create and deliver value to the customers.”

- businessdictionary.com



"NOTICE HOW HE LIKES HIS PERSONAL COMPUTER ALMOST AS WELL AS HIS SECURITY BLANKET..."

What is OSG Operations?

- Community
 - Communication
 - Support
 - Infrastructure



Who Uses OSG Operations?

- Users
- Resource Administrators
- VOs
 - VO Administrators
 - Application Developers
- Other Interests
 - OSG Management
 - Funding Agencies
- Anyone with an OSG or “Grid” Issue or Problem

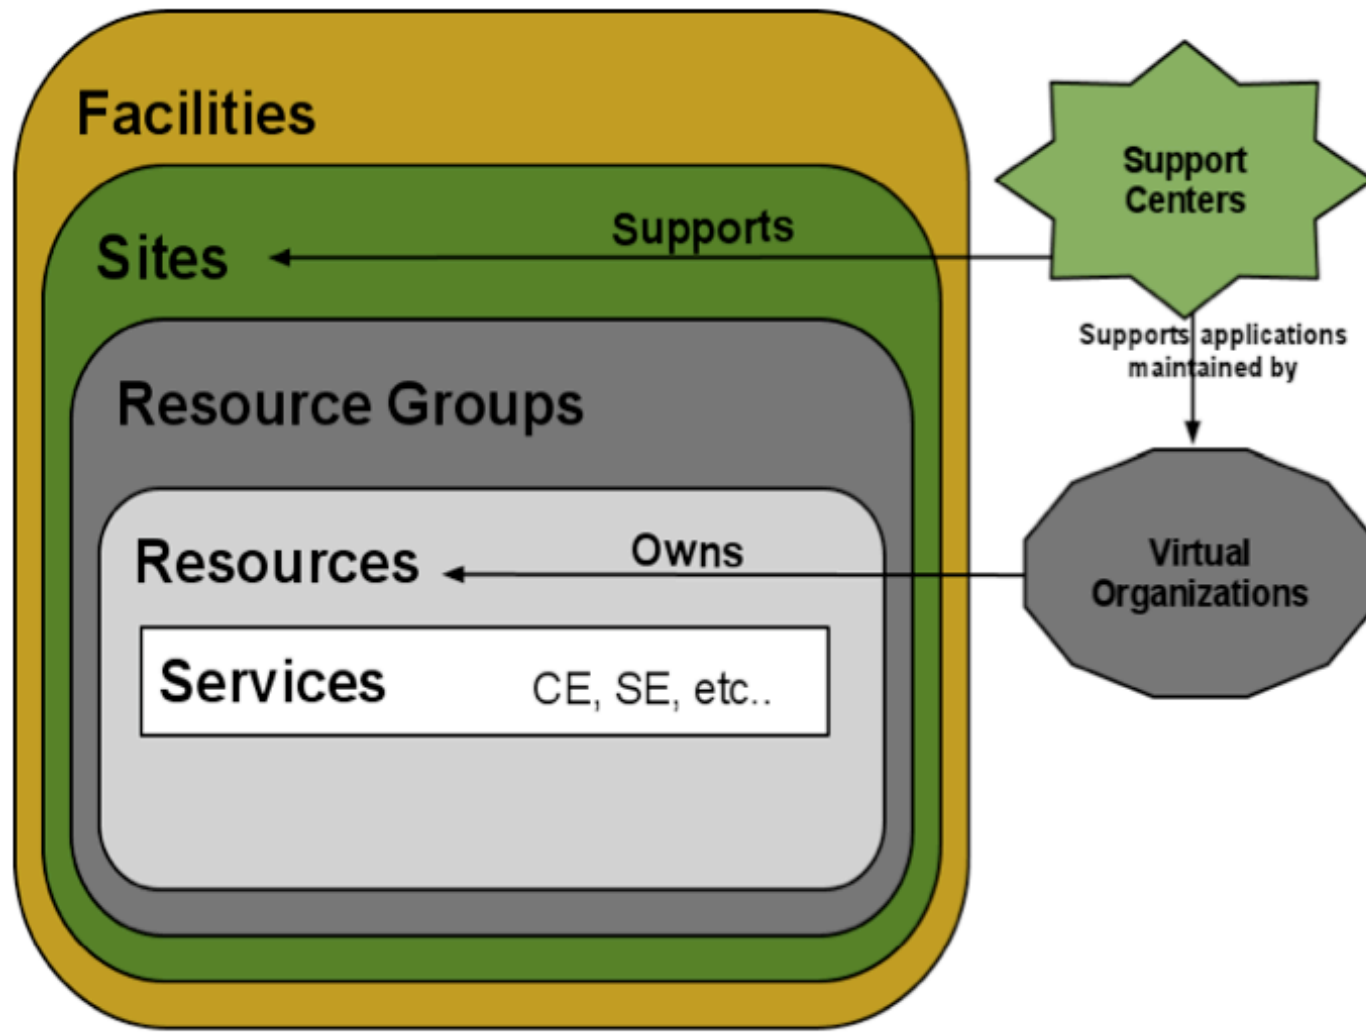
A User Walkthrough

- Issuing Credentials
- Registration with OSG and/or VO
- Documentation for getting started
 - Extracting and Using Credentials
 - Client Installation
- Software Cache
- Resource Discovery
 - BDII, ReSS, OSG Matchmaker
- Problem Reporting/Troubleshooting
- Accounting

A Resource Administrator Walkthrough

- Issuing Credentials
- Installation Documentation
- Software Cache
- Resource Registration
 - Administrative
 - Security
- Publishing
- Monitoring
- Reporting Issues/Troubleshooting
- Accounting

Some Operational Terminology



OIM Registration Database

- Describes OSG Topology
- Record of Administrative Data
 - Administrative Contacts
 - Emergency Contacts
 - Other Descriptive Information
- Semi-static Technical Information
 - Important endpoints (Resource and VO)
 - Information Publishing
 - Part of WLCG
- Downtime Information
- <https://oim.grid.iu.edu>

TWiki Collaborative Documentation

- Collaborative Documentation Environment
 - Installation Documentation
 - OSGRA Documentation (How to Install a Cert)
 - Details on all OSG Work Areas
 - Operations
 - Engagements
 - VOs
 - Many Others
 - Technical Documentation
- Weekly Meeting Agendas and Minutes
- Weak Search Engine
- <https://twiki.grid.iu.edu>

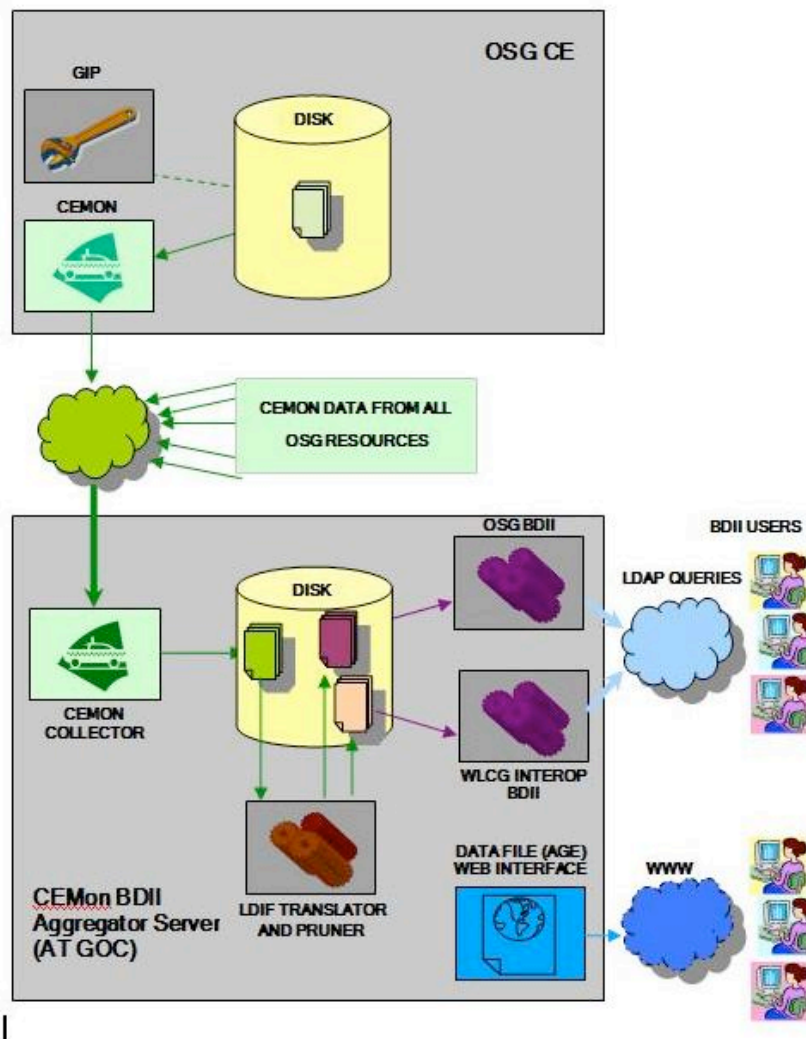
OSG Middleware Cache

- <http://software.grid.iu.edu/>
- Pointers to VDT Software Packages
- Certificate Authority Distribution (CADist)
 - <http://software.grid.iu.edu/pacman/cadist/>
- VO Package
- Certificate Requests
 - <https://software.grid.iu.edu/cert/certreg.php>

Information Systems

- BDII
 - GLUE Schema LDIF
 - <http://is.grid.iu.edu>
 - How to write an LDAP Search (Coming Up)
- ReSS
 - GLUE Schema Class Ad
 - <https://osg-ress-1.fnal.gov:8443/ReSS/>
- OSG Matchmaker
 - Mats Already Covered This
- MyOSG
 - Raw GLUE Data
 - Along with other important information

BDII Architecture



Schematic of CEMon BDII Information Service on the OSG

Hands On BDII



Constructing a BDI Query

- `ldapsearch -h is.grid.iu.edu -p 2170 -x -b mds-vo-name=local,o=grid`
- `ldapsearch -h is.grid.iu.edu -p 2170 -x -b mds-vo-name=GridUNESP_CENTRAL,mds-vo-name=local,o=grid`
- `ldapsearch -h is.grid.iu.edu -p 2170 -x -b mds-vo-name=local,o=grid "(&(objectClass=GlueCE)(GlueCEAccessControlBaseRule=*gridunesp))"`
- `ldapsearch -h is.grid.iu.edu -p 2170 -x -b mds-vo-name=local,o=grid "(&(objectClass=GlueCE)(GlueCEAccessControlBaseRule=*gridunesp))"`
`GlueCEInfoLRMSType GlueCEStateFreeCPUs`
`GlueCEInfoApplicationDir GlueCEInfoDataDir`

Monitoring

- Resource and Service Validation (RSV)
 - Locally run probes on baseline functionality
 - Results published to the RSV Collector
 - Some results pushed forward to WLCG

Current RSV Status

GridUNESP_CENTRAL OSG Production Resource Group

GridUNESP_CENTRAL ce.grid.unesp.br

✓ No issues found for this resource.

CE Service Status

✓ No issues found for this service.

Critical Metrics

✓ CACert Expiry ⓘ	Reported 2 hours ago
Checks if the resource's CA certificate is valid.	
Show Detail	
✓ OSG Directory Permissions ⓘ	Reported 59 minutes ago
Checks if permissions are set correctly on important user-accessible OSG directories.	
Show Detail	
✓ OSG Version ⓘ	Reported 16 minutes ago
Determines the installed OSG version.	
Show Detail	
✓ Ping ⓘ	Reported 10 minutes ago
Check if resource responds to pings.	
Show Detail	
✓ GRAM Authentication ⓘ	Reported 5 minutes ago
Checks if available proxy for the resource is valid and not about to expire.	
Show Detail	

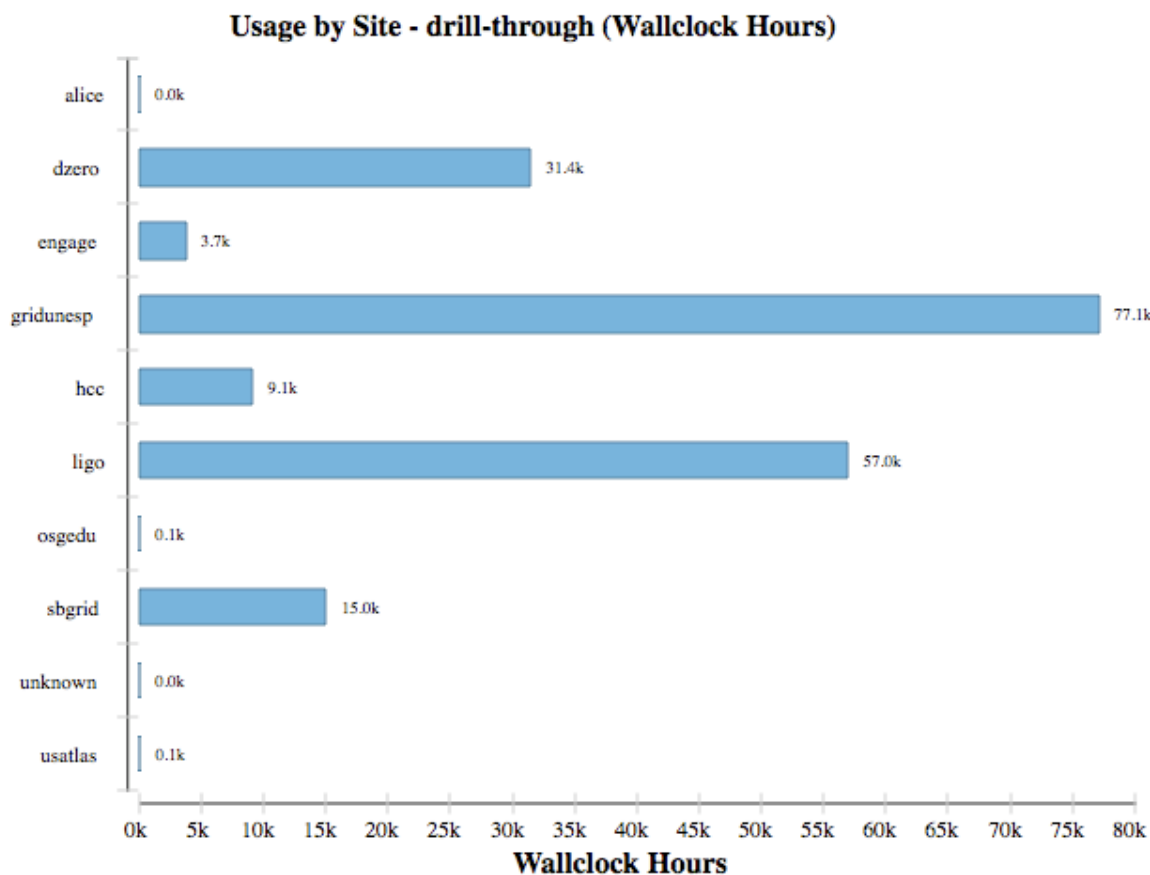
Uh oh! Trouble!!

- Something goes wrong somewhere in the chain...
 - <https://ticket.grid.iu.edu>
 - ~8-10 tickets per day
 - Same interface to submit and track tickets
- If you ever get stuck, get an error, can't accomplish something submit a ticket!

Accounting

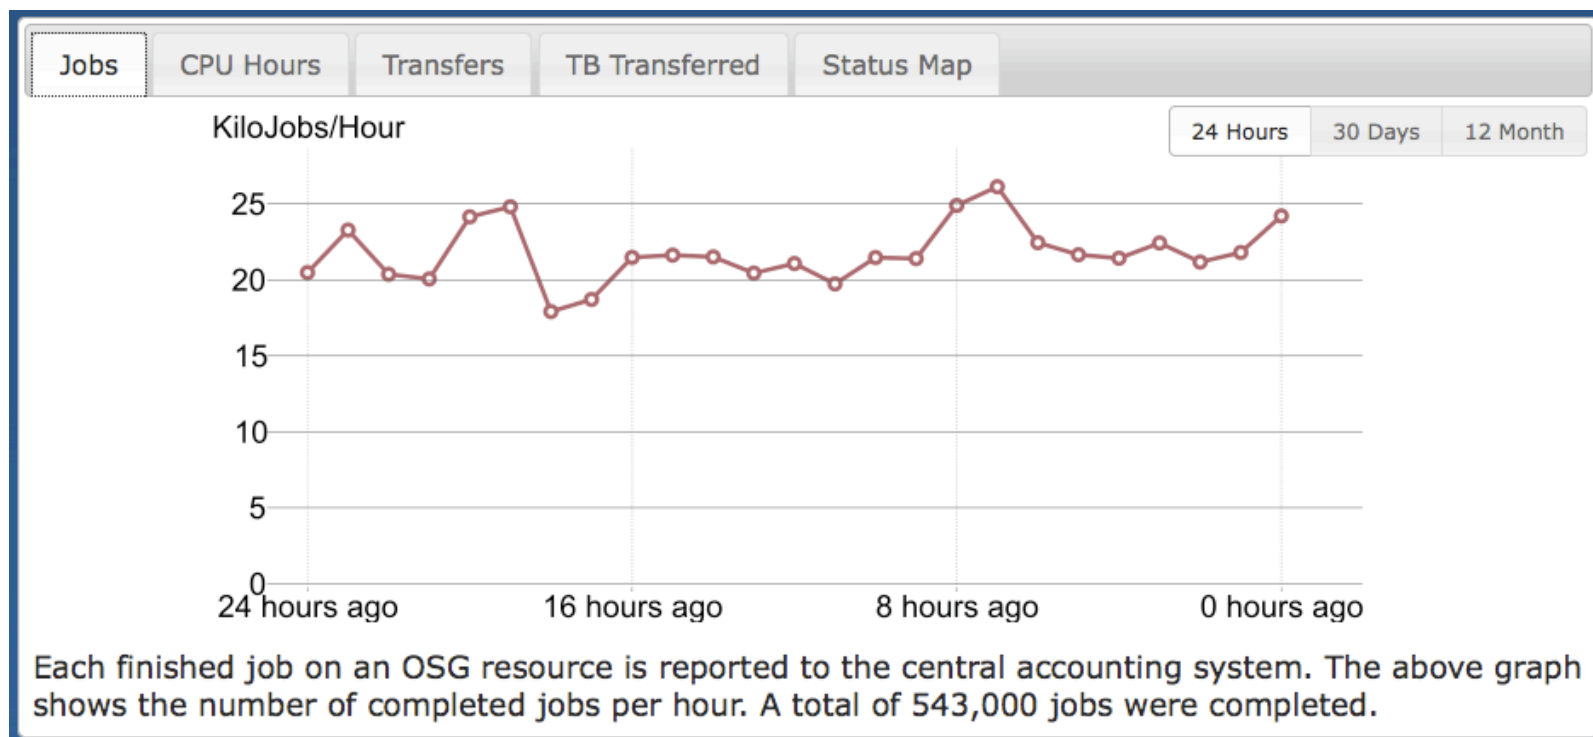
- <http://gratia-osg-prod-reports.opensciencegrid.org/gratia-reporting/>

Date range: 2010-11-30 - 2010-12-07
Site = GridUNESP_CENTRAL



OSG Display

- <http://display.grid.iu.edu/>
- Accounting and Status Monitoring Overview
 - Created specifically for the DOE



More Communication

- Community Notifications
 - <http://osggoc.blogspot.com/>
- OSG Web Pages
 - <http://www.opensciencegrid.org/>
- Document Database
 - <http://osg-docdb.opensciencegrid.org/>
- Trouble Ticket Exchange
 - GGUS, USATLAS, VDT
- VOMS for MIS and OSGEDU

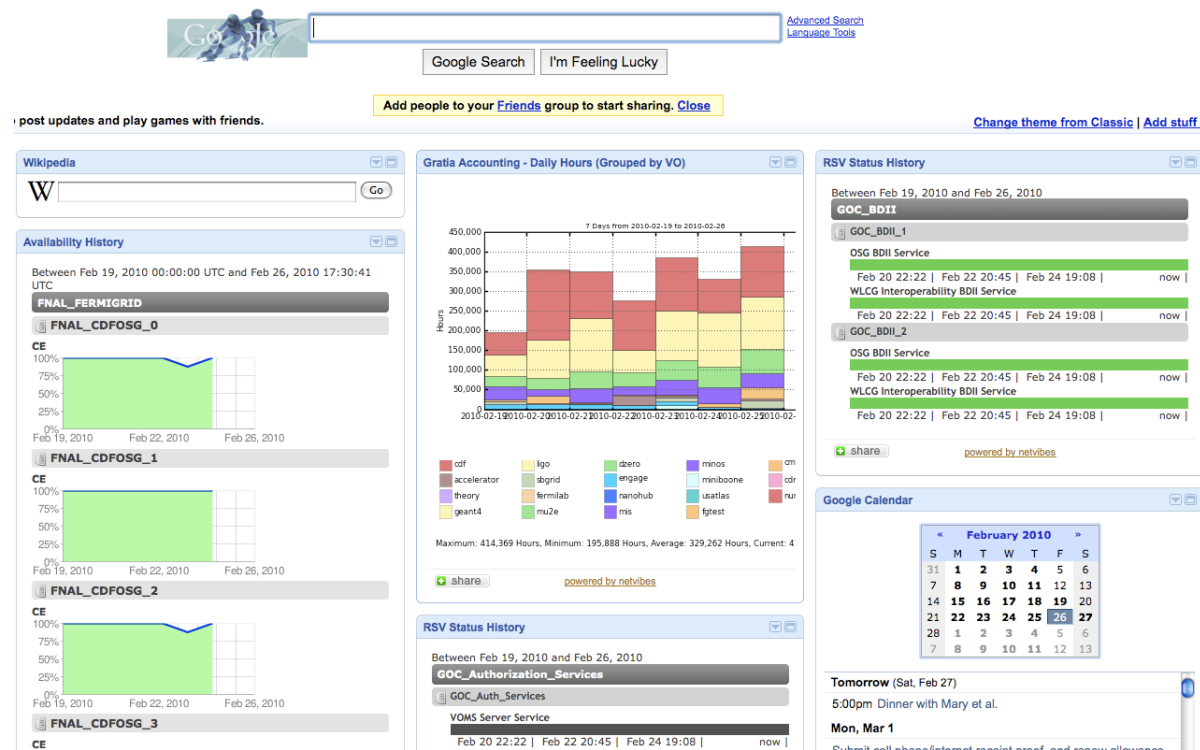


The ITB and Operations Services

- Most services mentioned have a mirror Integration Testbed (ITB) Service
 - New features and bug fixes
 - Testing
- Change Management Schedule
 - 1st, 3rd Tuesdays ITB Release
 - 2nd, 4th Tuesdays Production Release

MyOSG Presentation

- Ties most services together at a single address
 - <https://myosg.grid.iu.edu>



The screenshot displays the MyOSG web interface. At the top, there is a search bar with a Google logo and a search button. Below the search bar, there are several panels:

- Availability History:** A section showing the availability of various services. It includes a line graph for 'FNAL_FERMIGRID' and three bar charts for 'FNAL_CDFOSG_0', 'FNAL_CDFOSG_1', and 'FNAL_CDFOSG_2'. The y-axis represents availability percentage (0% to 100%) and the x-axis shows dates from Feb 19, 2010, to Feb 26, 2010.
- Gratia Accounting - Daily Hours (Grouped by VO):** A stacked bar chart showing daily hours for various Virtual Organizations (VOs) from Feb 19, 2010, to Feb 26, 2010. The y-axis represents hours (0 to 450,000). A legend identifies VOs such as cdf, ligo, ibero, minos, cm, accelerator, sbgrid, engage, miniboone, cdf, theory, fermilab, nanohub, usatlas, rui, geant4, mu2e, and ms. Statistics provided: Maximum: 414,369 Hours, Minimum: 195,888 Hours, Average: 329,262 Hours, Current: 4.
- RSV Status History:** A section showing the status of various services. It includes a table for 'GOC_BDII_1' and 'GOC_BDII_2' with columns for service name, dates, and status (now). A 'powered by netvibes' logo is visible.
- Google Calendar:** A calendar view for February 2010. The date Feb 27 is highlighted. Below the calendar, there is a reminder for 'Tomorrow (Sat, Feb 27) 5:00pm Dinner with Mary et al.' and 'Mon, Mar 1'.

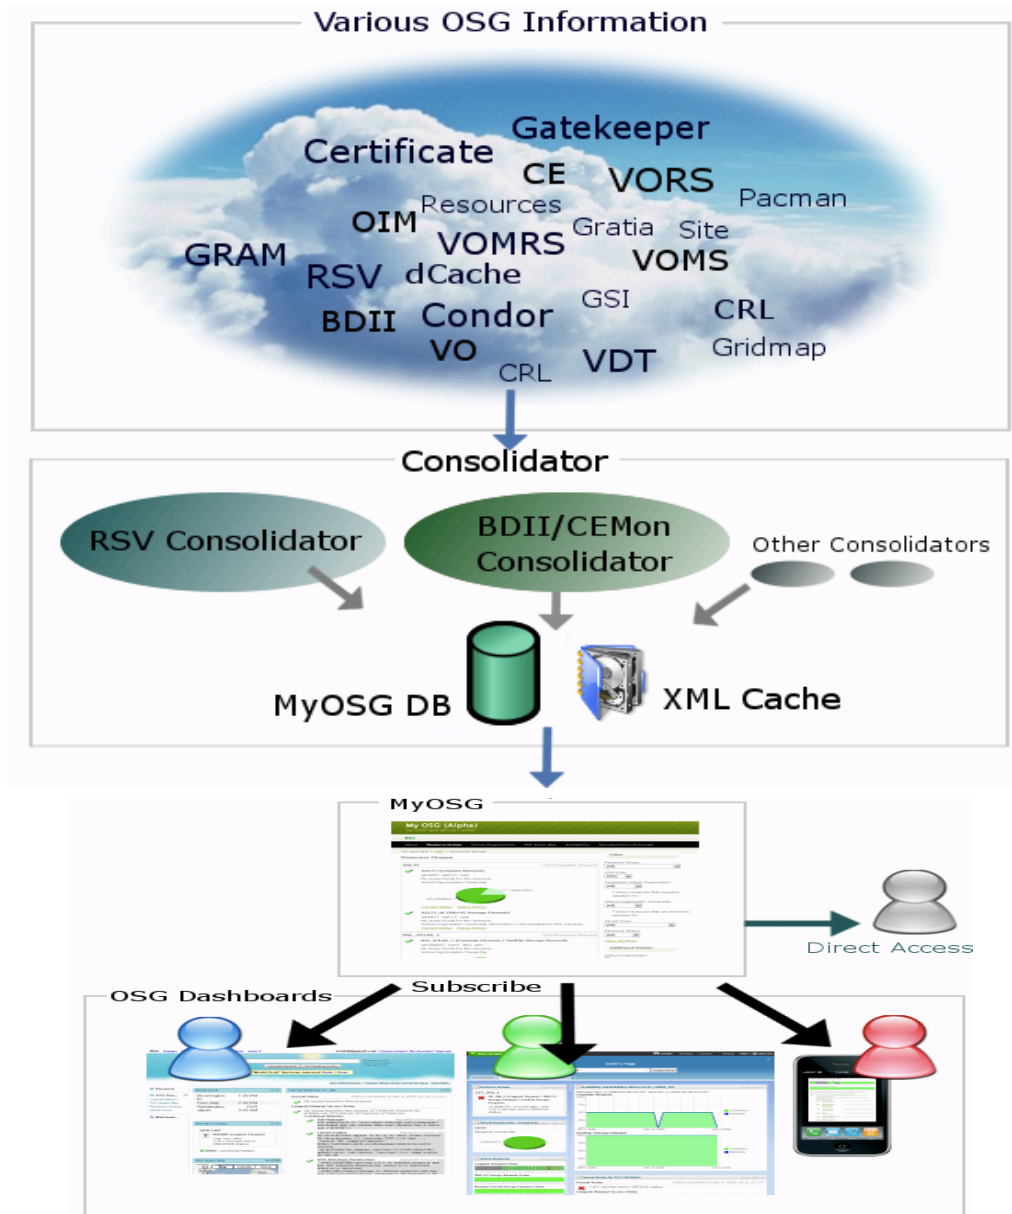
MyOSG

- The **Problem** we were trying to solve
 - There are a plethora of **Operational Tools**
 - Administrative, Monitoring, Accounting, Information Systems
 - Several different groups are **Developing** these tools independently
 - Users find these tools **Useful** for various information
 - These tools are **Distributed** out among a variety of sources
 - **New collaborators** didn't know each of these tools and sometimes are unaware of the information published about there resource
 - **Experienced collaborators** had to manage several bookmarks to each source of information
 - Different sources could sometime be publishing **contradictory** information

Other Considerations

- Each consumer of Operational Data needs a **Unique** view
 - The site administrator, the application submitter, the Operations Coordinator, etc.
- Data should be available in the form the user can adopt into their **own daily workflow**
- New tools services should be **easily added** as they are developed and become useful to the community
- All the information should be available from a **single** location

Original Conceptual Design



Experiences

- **Widgets** (UWA format) didn't catch on like we expected
 - Though it is still available
- **Mobile content, XML** or using the MyOSG **web pages** directly are most popular forms of access
- New consolidators are **rarely** necessary
 - Though easy enough to add
- **Merlin Mann**

“...open your heart to the remote possibility that any mythology of personal failure that involves messiahs periodically arriving to make everything “easy” for you might not *really* be helping your work *or* your mental health *or* your long-standing addiction to using tools solely to ship new excuses.”

Hands On MyOSG

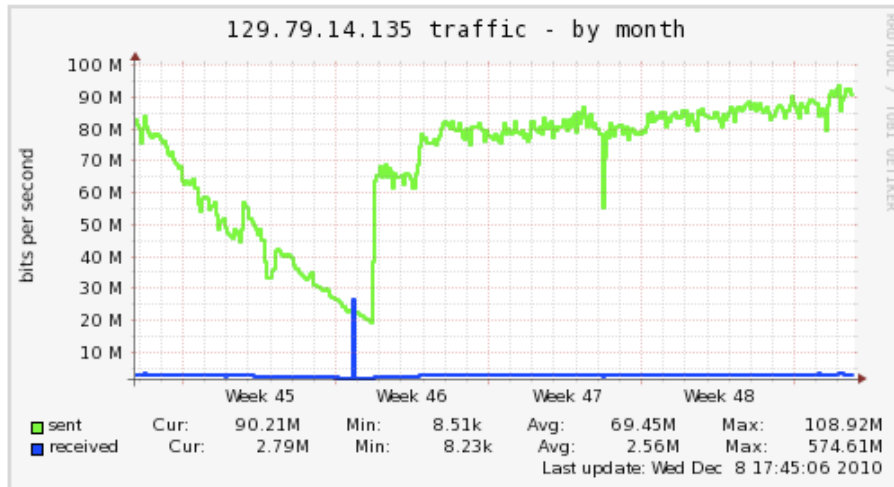


Some Things to Try

- Filter on an Interesting Visualization
 - Health, Accounting, Administrative, Information
- Filter on an Interesting Set of Resources
- Export in an interesting Format
 - Universal Widget API (UWA), XML, Mobile

Internal Monitoring

- Internal Monitoring and Documentation



Webs

- Documentation
- BDI
- Misc
- MyOSG
- OIM
- RSV
- ServersAndServices
- SoftwareCache
- TWiki
- Ticketing
- Training
- VOMS
- Main
- Meetings
- Support
- Sandbox
- TWiki

What Other Services are Needed?



Human Support Services

- Answering Phone/Email/Tickets
 - 24x7x365 Coverage
- Troubleshooting and Routing of Issues
- Emergency Response
 - Certificate Revocation
- Organizing Meetings
- Building Relationships
- Engaging the OSG Community