Open Science Grid Update for DOSAR

Keith Chadwick,
Fermilab

Outline

Open Science Grid

- Cover status/happenings on OSG since last talk.
- A representative Campus Grid - FermiGrid
- DOSAR collaboration with OSG?

Current OSG deployment

Open Science Grid

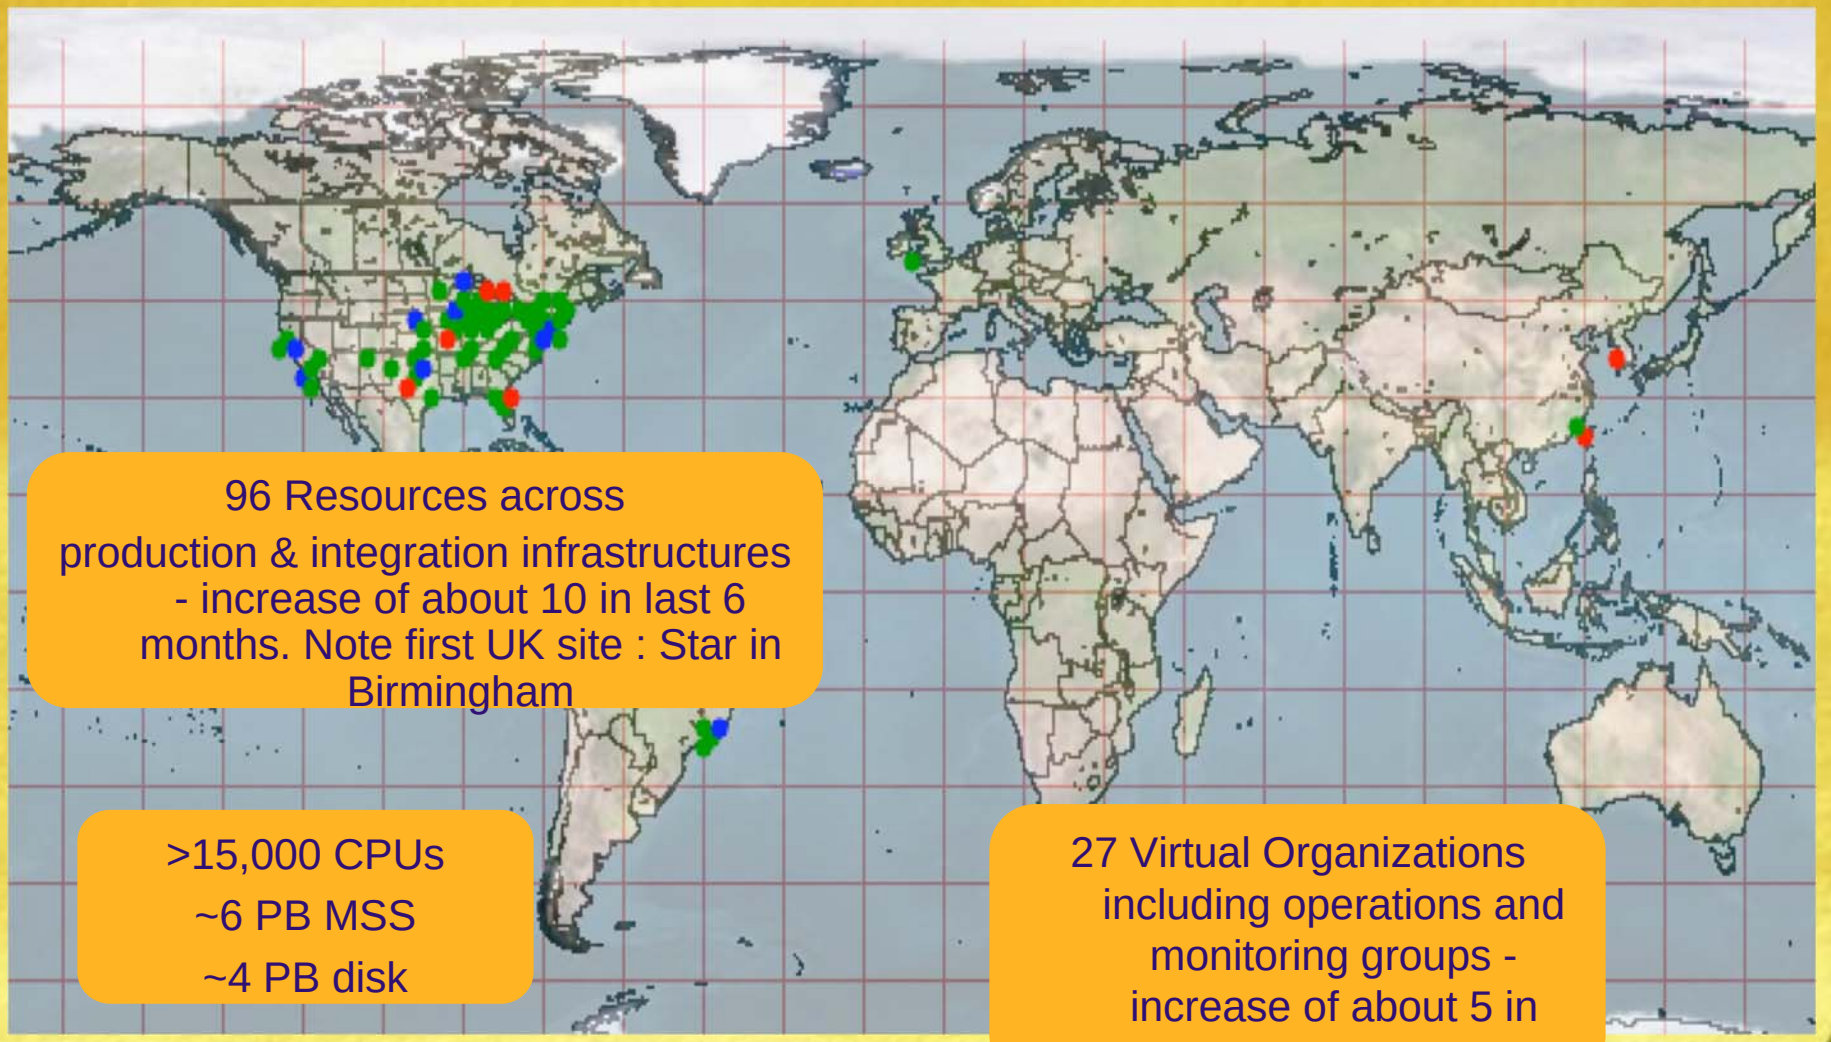

96 Resources across production & integration infrastructures - increase of about 10 in last 6 months. Note first UK site : Star in Birmingham

>15,000 CPUs
~6 PB MSS
~4 PB disk

27 Virtual Organizations including operations and monitoring groups - increase of about 5 in last 6 months

OSG Funding

Open Science Grid

- OSG will be co-funded by DOE and NSF for 5 years at $6M/year starting in Sept ‘06
- Includes major deliverables for US LHC, LIGO and STAR; well used by CDF and D0; and potentially other experiment major deliverables in the future.
- Commitment to engage communities beyond Physics.
- Commitment to collaborate with EGEE and TeraGrid.
- Reminder:
  - Consortium policies are to be open to participation by all researchers.
  - Project responsibilities are to operate, protect, extend and support the Distributed Facility for the Consortium.

Project Management

Open Science Grid

- CS/technology providers part of management teams.
- LHC well represented in the management structure.
- External project technical leads (Globus, CDIGS, SOS, LIGO-PIF) part of the Executive Board.
- Partners (WLCG, TeraGrid, EGEE etc.) represented through liaisons on the Council.

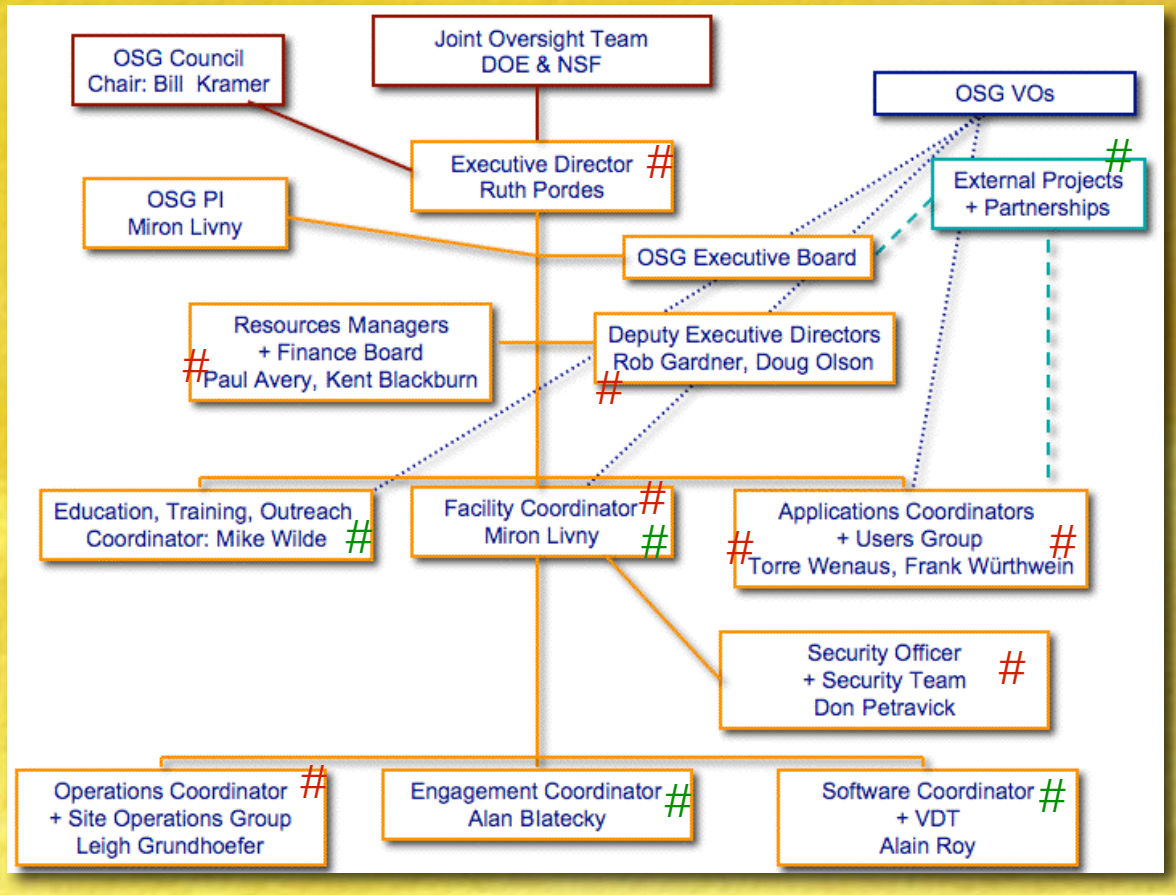

Effort Distribution in Year 1

Open Science Grid

- Will review distribution each year

	FTEs
Facility operations	5.8
Security and troubleshooting	4.5
Software release and support	6.5
Engagement (Year 1 only)	2.5
Education, outreach & training	2.5
Extensions in capability and scale.	8.5
Management & Staff	4.0
Total FTEs	**34.3**

OSG support for non-physics communities.. update

Open Science Grid

- Seeing DOSAR VO jobs at OU submitted through OSG services and accounted on MonaLisa.
- Alan Blatecky's group at RENCI is porting the "award winning **Bioportal**" to OSG.
- >100 **Nanotechnology** jobs -- that run from 10-40 days -- are being executed on LIGO, ATLAS and CMS sites.
- OSG is discussing partnership with the Northwest Indiana Computing Grid (NWICG) -- which riases the issue of **Gaussian.** When we start talking to Computational Chemistry we quickly run into licencing issues. Yes, we also say it is the responsibility of the project/VO .. But there are 50 sites on OSG.
- The P-Grade portal has been interfaced to a version of **CHARMM** molecular dynamics simulation package. Some versions of this also have licencing issues.
- Work on **Campus Grids** enabling Crimson Grid, NWICG, New York State Grid (NYSG) , GPN (Nebraska Education/Training grid) partnerships. (Note: Partners do not have to contribute resources; collaboration can equally be in software, procedures, education, training, security etc.)

OSG Middleware

Open Science Grid

- OSG Middleware is deployed on existing farms and storage systems.
- OSG Middleware interfaces to the existing installations of OS, utilities and batch systems.
- VOs have VO scoped environments in which they deploy applications (and other files), execute code and store data.
- VOs are responsible for and have control over their end-to-end distributed system using the OSG infrastructure.

Applications

Infrastructure

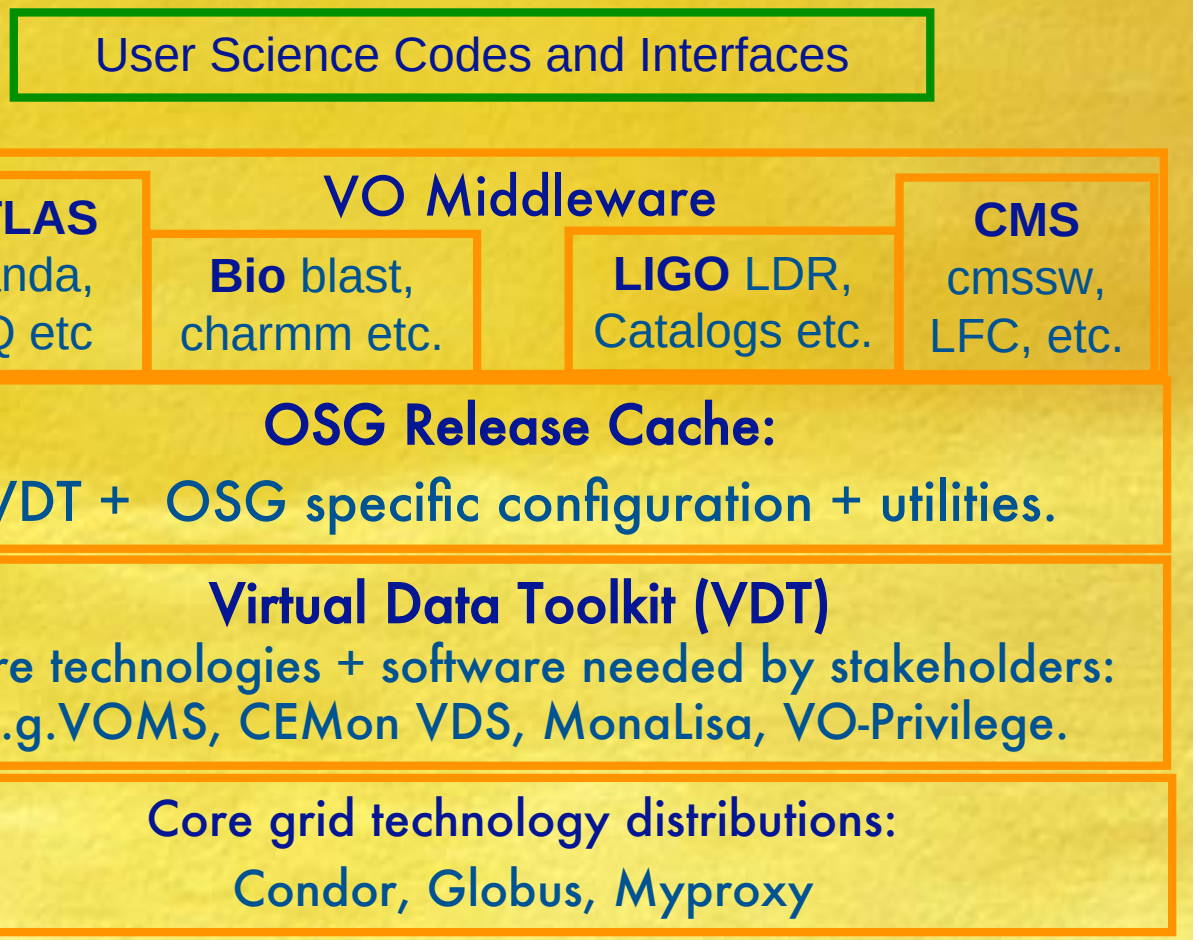

OSG Software Releases

Open Science Grid

- OSG 0.4.1 released in May.
- Security patch for Globus released in September.
- Still supporting some sites at 0.2.1
- Future releases:
  - OSG Release 0.6.0 Schedule and function priorities driven by needs of US LHC.
  - Balanced by need to prepare for medium-longer term sustainability and evolution.
  - Goal of 2 major releases a year with easier minor updates inbetween.

Upcoming VDT Work

Open Science Grid

- Provide smooth updates to a running production service:
  - In-place vs. on-the-side
  - Preserve old configuration while making big changes.
  - Make updates an hour not a days work.
- Support for more platforms?
  - More linux variants and needs of EGEE.
  - AIX needed by NERSC.
  - Mac OS X being tested by STAR, TIGRE, …
  - Solaris?
- Add dCache/SRM packaged in RPMs and tested on existing storage/data management testbed.

OSG 0.6.0

- Expect to rely on VDT 1.3.12 or 1.4.0
- Release early 2007.
- Some components being tested for earlier installation:
  - Gratia accounting
  - CEmon for Resource Selection Service.
- OU early tester of Gratia -- we appreciate this a lot.

Task	Priority
Gratia	1
Improve ease of upgrade of one VDT version to another	1
BDII	2
About 50 small bug fixes and enhancements	2
New GIPs	2
gLite client tools for interoperability purposes	2
New VDS	2
dCache/SRM as "RPM"	3
tclGlobus	3
Update platforms we build on	3
New Globus - WSGRAM and GT4	3
CEMON	4
SQUID	4
Globus Virtual Workspaces,for Edge Services Framework	5
Job managers/GIPs for Sun Grid Engine	5
Switch from MySQL to MySQL-Max	5
Upgrade Monalisa to the latest version	5
MDS 4	5
VOMRS	5
glexec	6
Grid Exerciser	6
Production release of the VDT (1.4.0)	6

Risk based view of the world

Open Science Grid

- Organizations implement controls over their activities so as to obtain acceptable residual risk. Organizations: Sites, VOs and Grids.
  - Each has a security process lifecycle.
  - Satisfaction jointly and severally.
- Each organization is captain of its own ship.
  - However, constrained to interoperate.
- Standards (e.g. OSG User AUP, Service AUP, VO AUP) aid interoperation.

Site-VO Interoperation

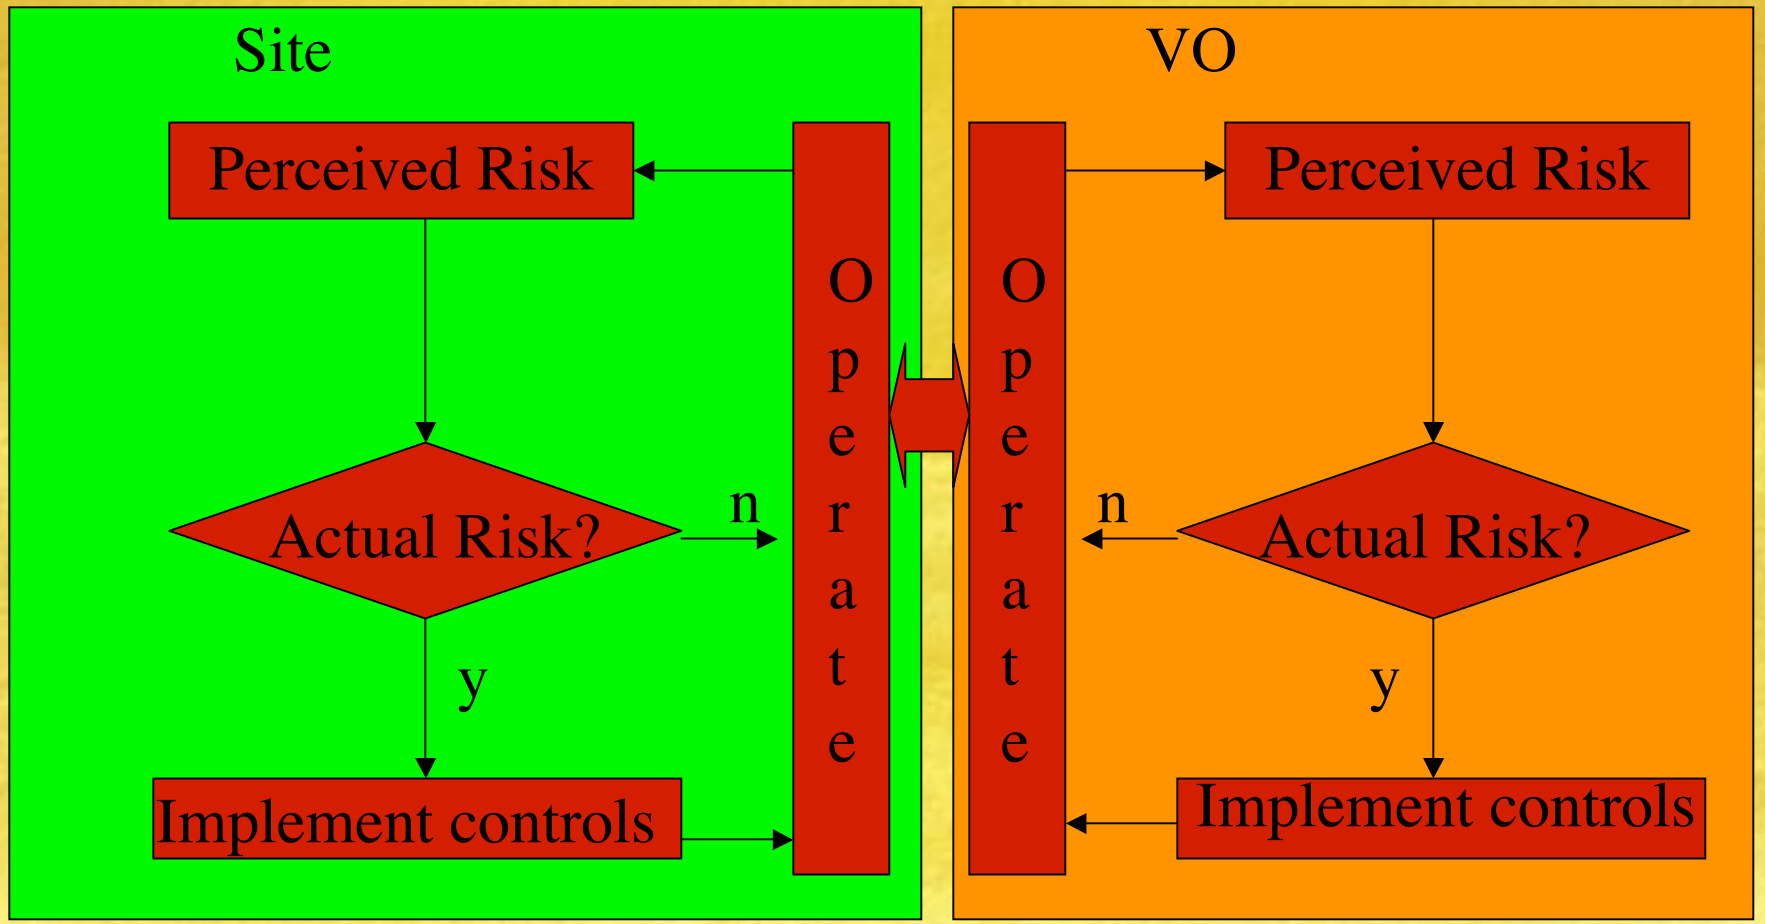

Scaling :-(

Open Science Grid

- Every organization has to reduce its risk to acceptable residual risk, and the security has to interoperate.
  - Identified trust.
    - Something specific is relied on
    - Is not checked each time.
  - Apropos Managerial, Operational and Technical Controls, for trust items.
- Scaling is a problem, work needs to be done
  - Common standards -- I.e OSG AUP.
  - Fewer entities
    - Aggregate small organizations into larger organizations
    - Larger organizations have greater resources

Liaison work

Open Science Grid

- TAGPMA, DOEgridPMA, EUgridPMA
  - A PMA is Policy Management Authority
- MWSWG -- Meeting at SLAC with emphasis on interoperability.
- GGF -- Global Grid Forum
- JSPG -- Joint Security Policy Mgemt Group
- DOE Security CET
- DOE Grid Briefing
- User AUP, Service AUP, VO AUP
- Grid Joint Security Policy working group.

Science Grid Communications

Broad set of activities

- (Katie Yurkewicz)
- News releases, PR, etc.
- Science Grid This Week, 1.5 years (>1000 subscribers)

www.interactions.org/sgtw

http://www.interactions.org/sgtw/

SCIENCE GRID THIS WEEK

MARCH 8, 2006

ABOUT SGTW | SUBSCRIBE | ARCHIVE | CONTACT SGTW

Calendar/Meetings

March

7-8, Second CLEANER All-Hands Meeting, Arlington, Virginia

8-10, GridChem Workshop: Distributed Computational Chemistry (on the Grid), Austin, Texas

13-15, ISSSE 06: International Symposium on Secure Software Engineering, Washington, D.C.

26-28, PRAGMA 10: Pacific Rim Applications and Grid Middleware Assembly Tenth Workshop, Townsville, Queensland, Australia

Full Calendar

Image of the Week

Indian President A.P.J. Abdul Kalam using VRVS at CHEP06. (Click on image for larger version.)
Image Courtesy Phillipe Galvez

On Feburary 17, Indian President A.P.J. Abdul Kalam visited the Computing in High Energy and Nuclear Physics (CHEP06) conference in Mumbai, India. The President's speech to the conference highlighted grid computing in India and around

Feature Story

Simulating Supersymmetry

The ATLAS detector under construction.
Image © CERN

One of the discoveries eagerly anticipated by particle physicists working on the world's next particle collider is that of supersymmetry, a theoretical lost symmetry of nature. Supersymmetry, often called SUSY, predicts the existence of a superpartner particle for every known particle.

Why the big hunt for SUSY's "sparticles"? Recent experiments have suggested that most of the matter in our universe is not made of familiar atoms, but of some new sort of "dark matter." Discovering a hidden world of sparticles will shed light on the nature of this dark matter, connecting observations performed at earth-based accelerators with those performed by astrophysicists and cosmologists.

Physicist Sanjay Padhi, a Chancellor Fellow at the University of Wisconsin-Madison, searches for SUSY using the ATLAS detector at the Large Hadron Collider. Although the LHC and ATLAS won't start collecting experimental data until 2007, he and his colleagues are already hard at work generating the

BBC Project Takes on Climate Change

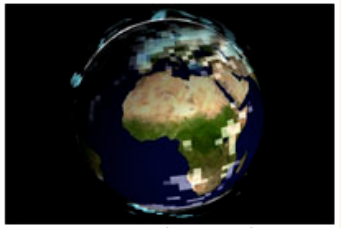

Image Courtesy Climateprediction.net

It seems like new distributed computing projects are popping up every day, and it seems like the world is catching on to the possibilities made possible by harnessing the power of thousands of connected PCs. Case in point: Last month, the British Broadcasting Corp. teamed up with ClimatePrediction.net to launch a distributed computing project that is running, initially at least, concurrent with the channel's "Climate Chaos" season of programming.

GRID today

Dubbed the "BBC Climate Change Experiment," this project, according to ClimatePrediction.net chief software architect Carl Christensen, is a little different than other similar projects. Unlike other distributed computing experiments where user computers will perform a "workunit" before moving on to another task, computers on the BBC experiment run the entire climate model—from start to finish. The experiment takes three months "on the fastest PCs out there today," he said, whereas tasks on other projects can be completed in a matter of hours.

There is a twofold reason for this

OSG Newsletter

Monthly newsletter

- (Katie Yurkewicz)
- 7 issues now

www.opensciencegrid.org/osgnews

Open Science Grid News

FEBRUARY 2006

OSG HOME | SUBMIT NEWS | SUBSCRIBE | ARCHIVE | ABOUT OSG NEWS

Meetings and Events

Condor Week 2006
April 24-27, 2006

GGF17
May 9-12, 2006

HPDC 2006
June 19-23, 2006

View Full Calendar
Add New Event

Operations Report

The Operations team provides a bi-weekly report to the Council to keep them up to date and to bring issues to the table for discussion. Here are a few items from recent postings.

With CMS and ATLAS increasing their number of submitted jobs to a site, several scaling problems have arisen. These issues are being addressed by extensions to Condor and Condor-G, and with the Globus software which is run on the head node for the job-manager.

More than 45% of sites have upgraded to OSG 0.4.0, and 26% are reporting to the MonaLisa (ML) accounting. The daily usage reports are based on ML, so while it remains an optional component, if you want your site to be included in the accounting you will need to install and configure it. The operations team will be happy to help with this.

The education project MARIACHI and

From the Executive Director

Dear OSG Consortium and Friends,

I am very pleased to announce that Bill Kramer has been selected by the Council as its new chairperson. As head of the NERSC computing center, Bill brings a wealth of experience and understanding to our program, and we are already keeping him busy. As one of the new applications coordinators, Frank Würthwein continues to be part of the OSG's core team and we will continue to benefit from his contributions and insights. I look forward to working with each and every member of the Executive Board.

At the beginning of this month we submitted the OSG program of work as an unsolicited proposal to the NSF's Mathematical and Physical Sciences Directorate, and we are in the process of submitting the same proposal to the DOE SciDAC-2 program. The proposal focuses on three key areas: the OSG facility; education, outreach and training; and science-driven extensions.

The Consortium meeting saw the presentation and discussion of many aspects of the use and provisioning of the facility, including the contents and schedule of the next two OSG releases and VDT 1.3.10. The local organizers—Paul Avery, DeeDee Carver and Jorge Rodriguez—did a superb job. CMS is ramping up OSG activity once again and DZero is validating one site at a time to run SAMGrid-based reprocessing jobs. Mike Wilde is working with Soma Mukherjee and UTB on the logistics and schedule for this year's summer school; please contact him if you are interested in contributing.

Sincerely,
Ruth Pordes, OSG Executive Director

Applications - Reprocessing D0 Data

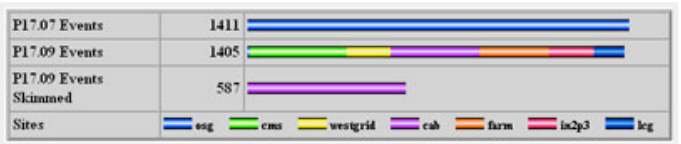

Reprocessing status as of February 20.
(Click on image for larger version.)

D0's latest reprocessing of its Run IIa data used several OSG sites, which together processed more than 10 million events. D0 has used resources from collaborating institutions for several years for reprocessing, since their Fermilab resources are busy processing newly collected data and Monte Carlo simulations are always ongoing. The addition of OSG and LCG resources

OSG Summary

Open Science Grid

- The Open Science Grid infrastructure is a part of the WLCG Collaboration.
- OSG supports and evolves the VDT core middleware in support of the EGEE and the WLCG as well as other users.
- The OSG software stack is deployed to interoperate with the EGEE and provide effective systems for the US LHC experiments among others.

What is FermiGrid?

Open Science Grid

FermiGrid is:

- An example of a campus Grid.
- A set of common services, including:
  - The Fermilab site globus gateway, VOMS, VOMRS, GUMS, SAZ, MyProxy, Gratia Accounting, etc.
- A forum for promoting stakeholder interoperability and resource sharing within Fermilab.
- The portal from the Open Science Grid to Fermilab Compute and Storage Services:
  - Production: fermigrid1, fngp-osg, fcdfosg1, fcdfosg2, docabosg2, sdss-tam, FNAL_FERMIGRID_SE (public dcache), stken, etc…
  - Integration: fgtest1, fnpcg, etc…

People

Open Science Grid

FermiGrid Operations Team:

- Keith Chadwick (CD/CCF/FTP) – Project Leader - Monitoring and Metrics
- Steve Timm (CD/CSS/FCS) – Linux OS Support, Condor, Globus Gatekeeper
- Dan Yocum (CD/CCF/FTP) – Middleware Support - VOMS, VOMRS, GUMS
- Neha Sharma (CD/CCF/FTP) - Middleware Support - Storage, Accounting

FermiGrid Stakeholder Representatives:

- Keith Chadwick – FermiGrid Common Services Representative
- Doug Benjamin – CDF Stakeholder Representative
- Ian Fisk – U.S.CMS Representative
- Amber Boehnlein & Alan Jonckheere – D0 Stakeholder Representatives
- Nickolai Kouropatkine – DES & SDSS Stakeholder Representative
- Steve Timm – Fermilab General Purpose Farms Representative
- Ruth Pordes – OSG Stakeholder Representative
- Eileen Berman & Rob Kennedy - Storage (enstore & dcache) Representatives

FermiGrid Web Site & Additional Documentation:

- http://fermigrid.fnal.gov/

FermiGrid Resources

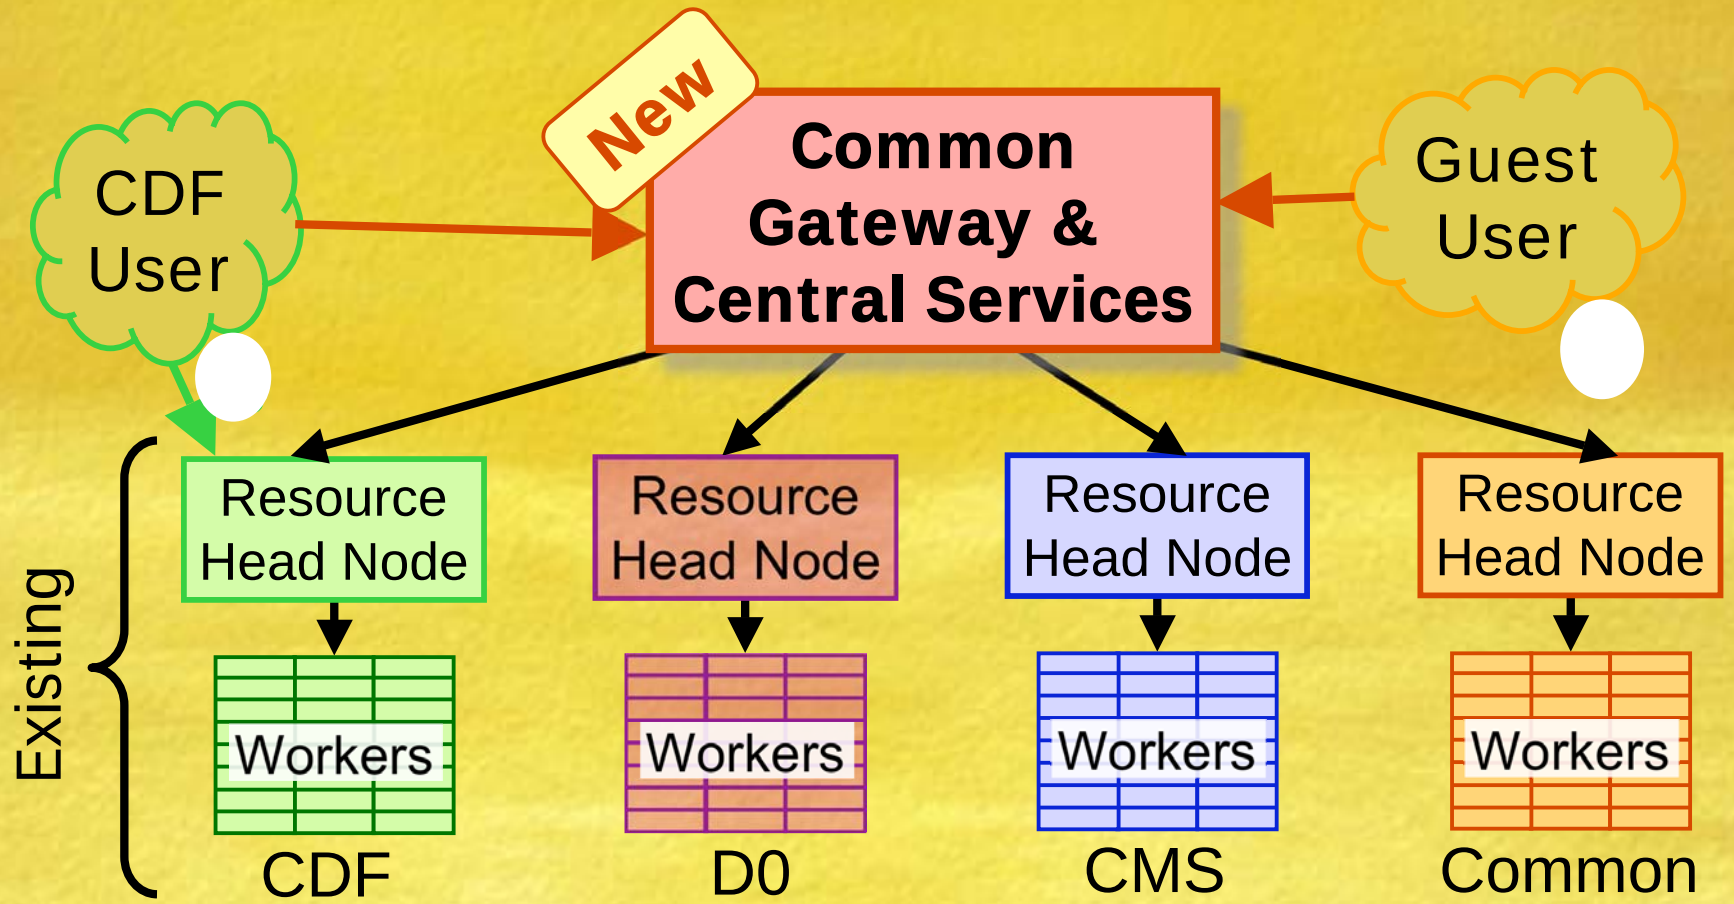

New Resources to be added to “Common”

Virtual Organizations

Open Science Grid

FermiGrid currently hosts the following Virtual Organizations:

- auger http://www.auger.org/
- des http://decam.fnal.gov/
- dzero http://www-d0.fnal.gov/
- fermilab http://www.fnal.gov/
- gadu http://www-wit.mcs.anl.gov/Alex/GADU/Index.cgi
- nanohub http://www.nanohub.org/
- sdss http://www.sdss.org/
- ilc http://ilc.fnal.gov/
- lqcd http://lqcd.fnal.gov/
- I2u2 http://www-ed.fnal.gov/uueo/i2u2.html

SAZ - Site AuthoriZation Module

Open Science Grid

- We are in the process of implementing our SAZ (Site AuthoriZation) module for the Fermilab “Open Science Enclave”.
- SAZ will enable Fermilab to make Grid job authorization decisions for the Fermilab site based on the users DN, VO, Role and CA.
- Our current plan is to operate in a default accept mode for user credentials that are associated with known and “trusted” VOs and CAs.

Moving to Resource Sharing @ FNAL

Open Science Grid

User Resources

as of 9/06

User Access

Stakeholder Resources

Stakeholder	CDF	USCMS	D0	GP Farms	OSG	Fermilab Storage
CDF	✓	✓	✗	✓	✓	✓
USCMS	✓	✓	i	✓	✓	✓
D0	✗	✓	✓	✓	i	✓
GP Farms	✓	✓	i	✓	✓	✓
OSG	✓	✓	i	✓	✓	✓

Key:

Icon	Description
✓	Task Completed & Interoperability Verified
i	Task In Process
✗	Task Not Completed or Interoperability Not Present

OSG Interfaces for Fermilab

Open Science Grid

- FermiGrid offers opportunistic use of FNAL computing resources through Open Science Grid (OSG) interfaces.
- Most of this work to accomplish this happened in the context of the installation and configuration of the Fermilab Common Grid Services and deployment and integration on the GP Farm, CDF OSG1&2, DO OSG CAB2 clusters.
- The Fermilab “job-forwarding” gateway has caused some problems for users that assume that a job submitted to the fork jobmanager will have access to the same filesystem that a job submitted to the condor-g jobmanager has.
  - We are in the process of deploying a multi-cluster filesystem based on the BlueArc NFS server appliance to provide a unified filesystem view from all clusters participating in FermiGrid job forwarding.

Metrics and Service Monitors

Open Science Grid

In addition to the normal operation effort of running the various FermiGrid services, significant effort has been spent over the last year to collect and publish operational metrics and instrument service monitors:

- Globus gatekeeper calls by jobmanager per day
- Globus gatekeeper IP connections per day

- VOMS calls per day
- VOMS IP connections per day
- VOMS service monitor

- GUMS calls per day
- GUMS service monitor

- Resource selection and acceptance of the fermilab VO at OSG sites

Metrics typically run once a day.

Service monitors typically run multiple times per day and are equipped to detect problems with the service that they are monitoring, notify administrators and automatically restart services as necessary to ensure continuous operations:

http://fermigrid.fnal.gov/fermigrid-metrics.html

Metrics - Globus Gatekeeper - fermigrid1

Open Science Grid

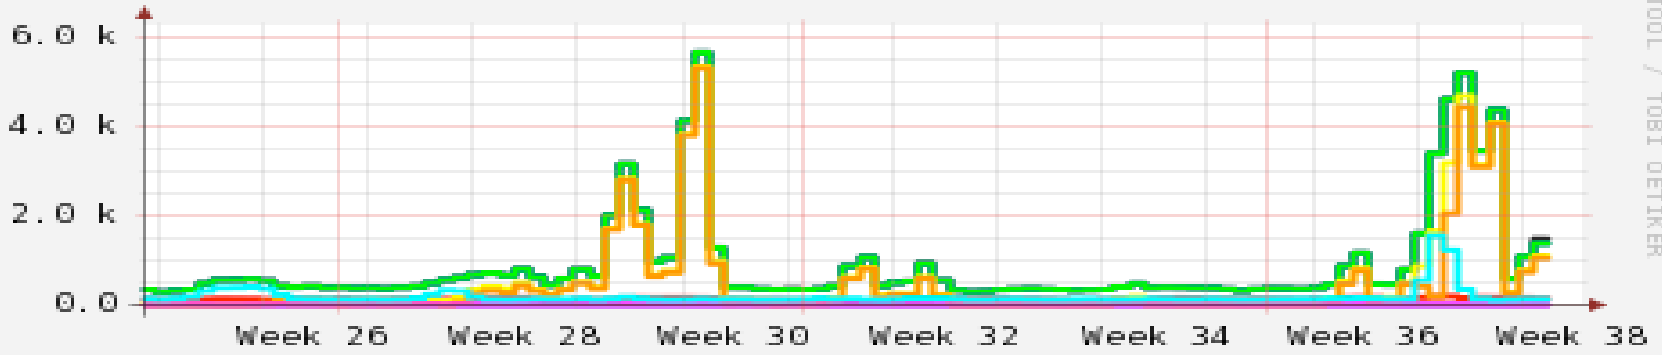

	Maximum	Average	Minimum	LastVal
authenticated	5662	885	295	1487
authorized	5662	882	295	1379
jobmanager	5662	882	295	1379
forks	190	139	70	146
managedforks	0	0	0	0
condor	5351	544	26	1090
condor-g	5317	451	0	1049
mis	14	11	7	9
default	1555	189	116	134

Data for fermigrid1 between 18-Jun-2006 and 21-Sep-2006
Plot generated on 21-Sep-2006

Metrics - Gatekeeper IP Connections

Open Science Grid

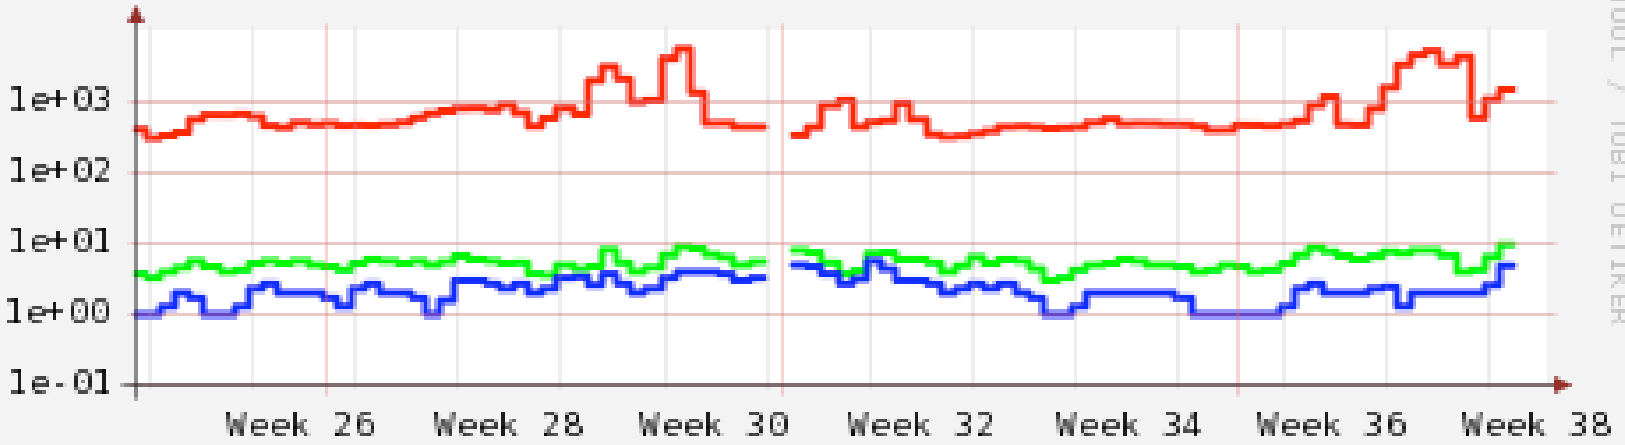

	Maximum	Average	Minimum	LastVal
total	5692	955	311	1501
unique	10	6	3	10
fermilab	6	2	1	5

Data for fermigrid1 between 18-Jun-2006 and 21-Sep-2006
Plot generated on 21-Sep-2006

Metrics – VOMS - fermigrid2

Open Science Grid

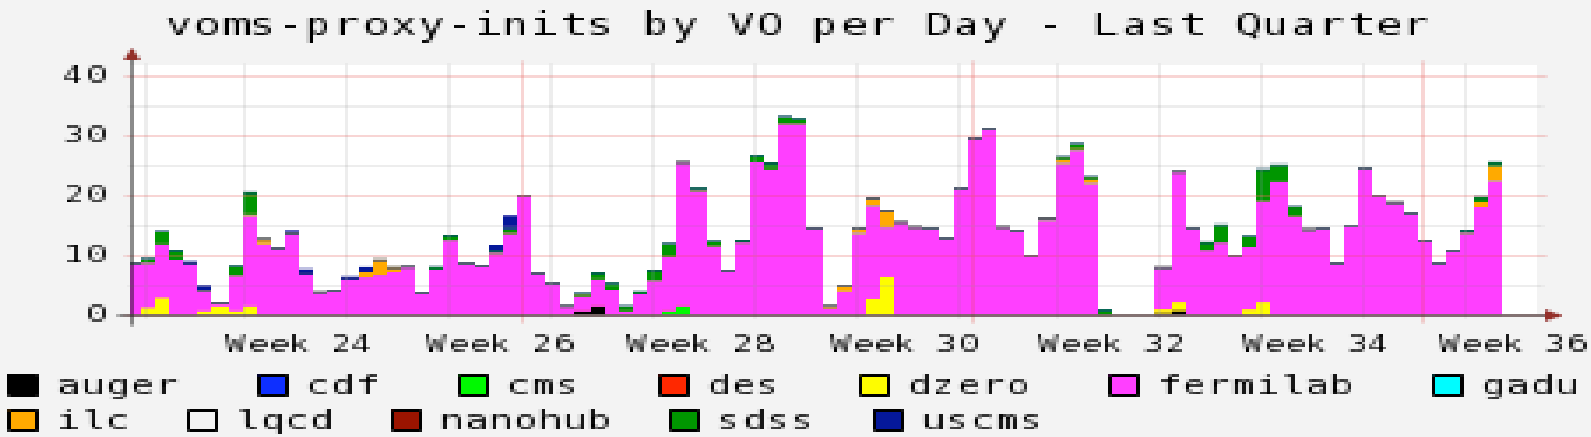

	Maximum	Average	Minimum
auger	1	0	0
cdf	0	0	0
cms	1	0	0
des	0	0	0
dzero	6	0	0
fermilab	32	12	0
gadu	0	0	0
ilc	3	0	0
lqcd	0	0	0
nanohub	0	0	0
sdss	5	0	0
uscms	2	0	0

Data between 04-Jun-2006 and 08-Sep-2006
Plot generated on 08-Sep-2006

Metrics - VOMS - fermilab VO groups

Open Science Grid

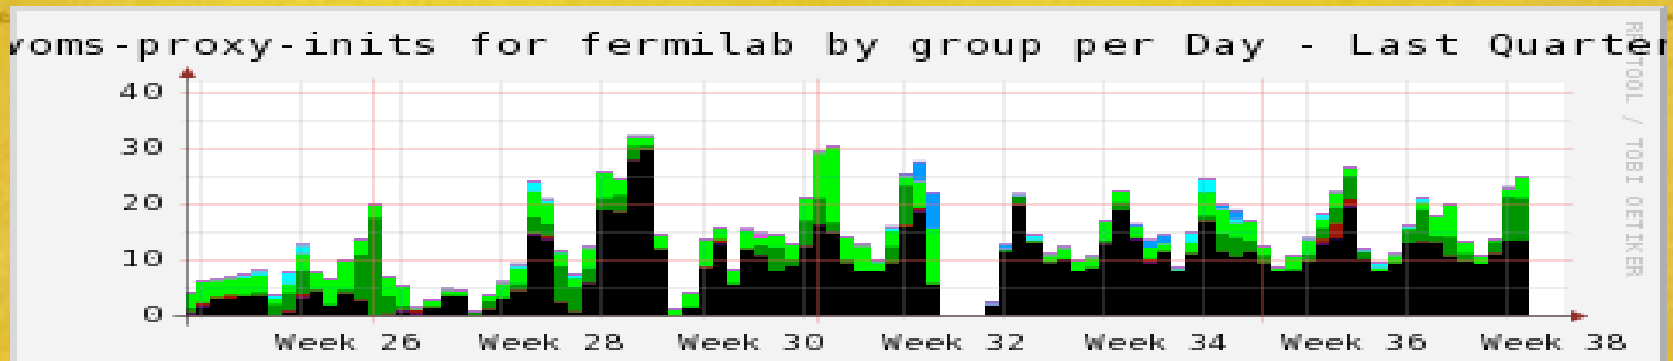

	Maximum	Average	Minimum	LastVal
null	30	9	0	13
accelerator	0	0	0	0
astro	0	0	0	0
cdms	3	0	0	0
hypercp	0	0	0	0
ktev	17	2	0	8
miniboone	14	2	0	4
minos	1	0	0	0
mipp	2	0	0	0
nova	0	0	0	0
numi	0	0	0	0
patriot	0	0	0	0
test	6	0	0	0
theory	0	0	0	0

Data between 18-Jun-2006 and 21-Sep-2006
Plot generated on 21-Sep-2006

Metrics - VOMS IP Connections

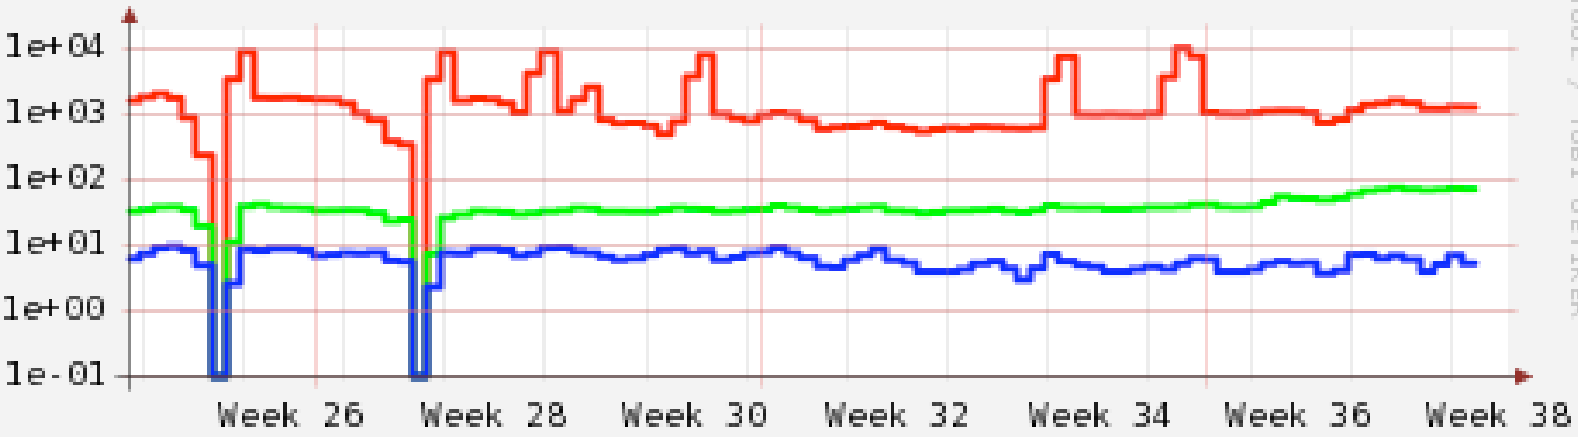

	Maximum	Average	Minimum	LastVal
total	10341	1778	0	1286
unique	76	38	0	73
fermilab	10	6	0	5

Data for fermigrid2 between 18-Jun-2006 and 21-Sep-2006
Plot generated on 21-Sep-2006

Monitoring - VOMS - fermigrid2

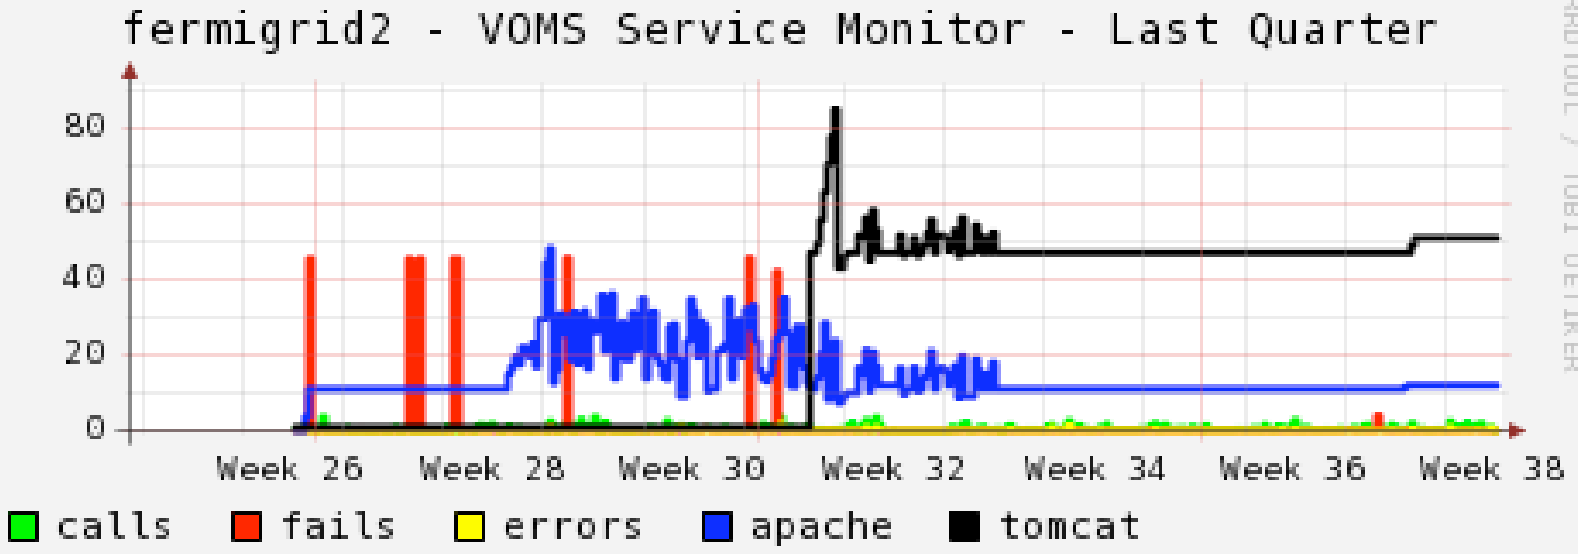

	Maximum	Average	Minimum	LastVal
calls	4	1	0	0
fails	46	1	0	0
errors	2	0	0	0
apache	48	14	0	12
tomcat	85	28	1	51

Data for fermigrid2 between 18-Jun-2006 and 21-Sep-2006
Plot generated on 21-Sep-2006

Metrics – GUMS - fermigrid3

Open Science Grid

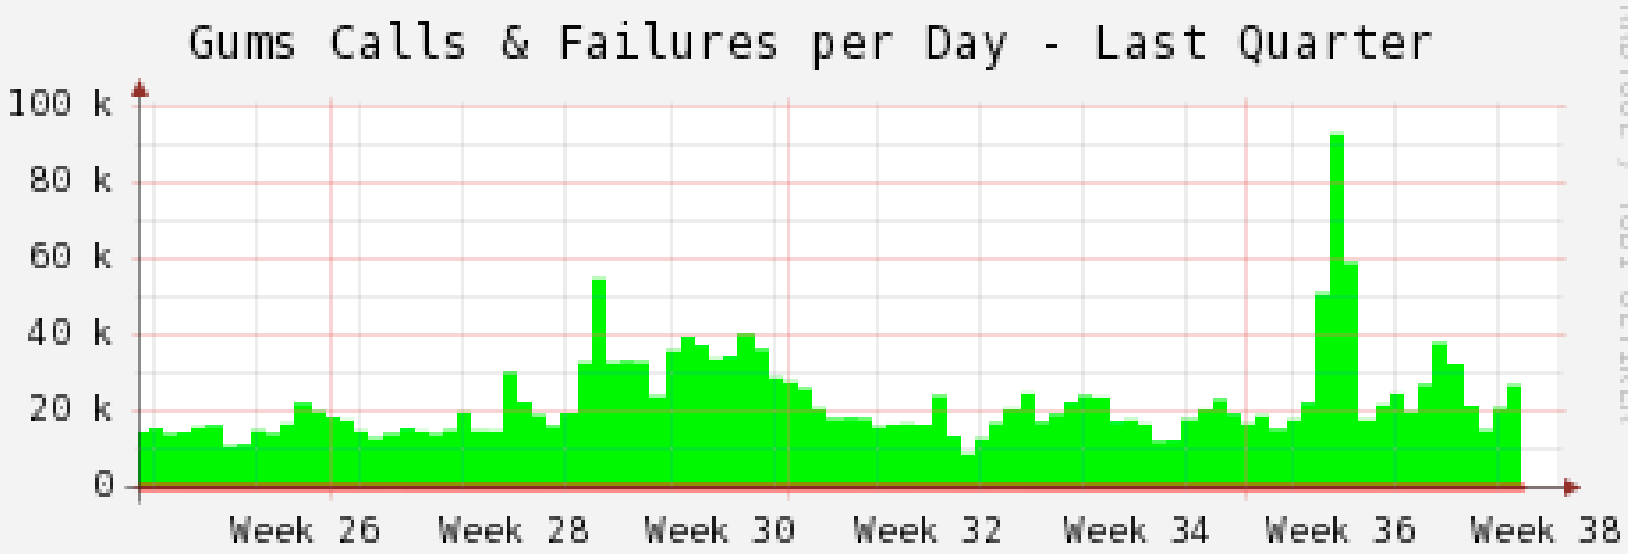

	Maximum	Average	Minimum	LastVal
Calls	92099	21965	8069	26249
Failures	11	0	0	1

Data between 18-Jun-2006 and 21-Sep-2006
Plot generated on 21-Sep-2006

Monitoring - GUMS - fermigrid3

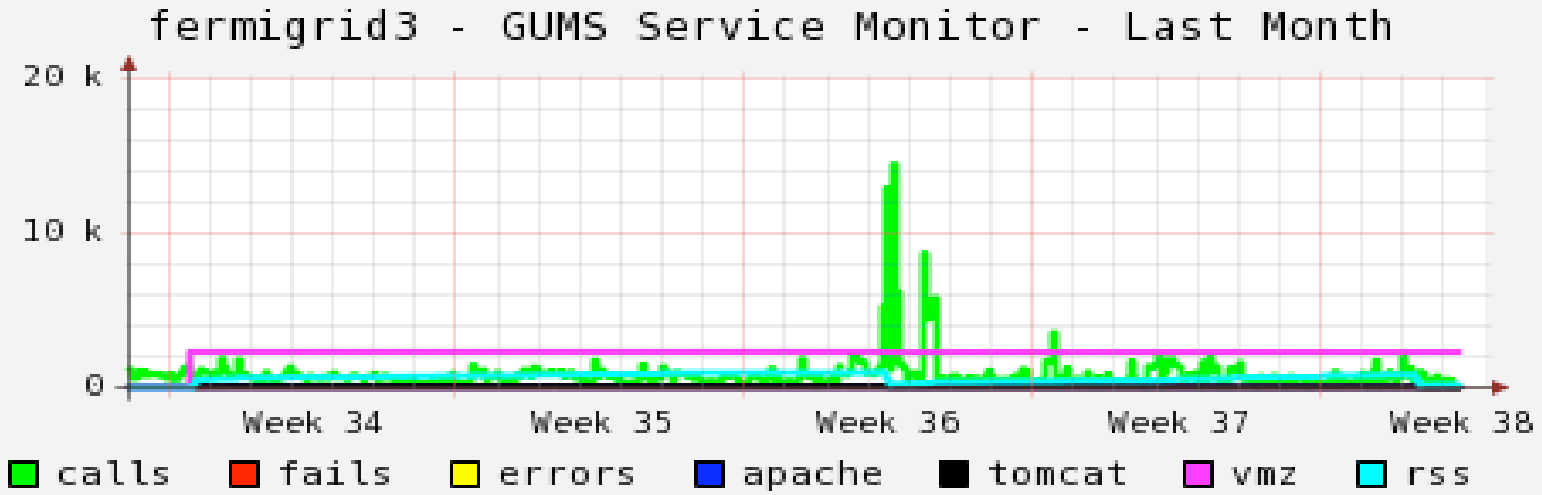

	Maximum	Average	Minimum	LastVal
calls	14467	866	0	0
fails	2	0	0	0
errors	0	0	0	0
apache	12	11	10	12
tomcat	78	59	1	1
vmz	2318180	2195604	0	2291244
rss	1002724	648203	0	179222

Data for fermigrid3 between 20-Aug-2006 and 21-Sep-2006
Plot generated on 21-Sep-2006

fermilab VO acceptance within the OSG.

Open Science Grid

- We have recently started development of a fermilab VO acceptance probe.
- Goals:
  - See where members of the fermilab VO can execute grid jobs in the OSG.
  - Cross check VO Resource Selector report.
  - Record detailed results on success / failure.
  - Provide a historical record.
  - Serve as a tool for Fermilab management.

Actions performed by the fermilab VO Probe

Open Science Grid

•Fetches the current Gridcat list of sites (production or integration).

•For each Compute Service entry in the Gridcat list of sites, the fermilab VO probe fetches the Gridcat information corresponding to the "Site Name" and parses out the:

- The Site Name
- The Gatekeeper Host Name
- The Site Application Directory ($APP)
- The Site Data Directory ($DATA)
- The Site Temporary Directory ($TMP)
- The Site Worker Node Temporary Directory ($WNTMP)

•The VO probe then runs the following series of tests against the Site:

- ping -c 1 gatekeeper_host_name or globusrun -a -r (if the ping failes).
- globus-run-job gatekeeper_host_name /usr/bin/printenv
- globus-url-copy -v -cd local_file gsiftp:gatekeeper_host_name:2811/temporary_file_name
- globus-url-copy -v -cd gsiftp:gatekeeper_host_name:2811/temporary_file_name /dev/null
- globus-run-job gatekeeper_host_name /bin/rm temporary _file_name

•The start and end time of the tests is recorded, as are the results of the individual probes.

•If an individual test in the sequence fails, then all remaining tests are skipped.

•Finally a detail report, summary report and trend plots are generated from the probe results.

fermilab VO probe - Production

Open Science Grid

http://fermigrid.fnal.gov/monitor/fermigrid0-fermilab-vo-production-monitor.html

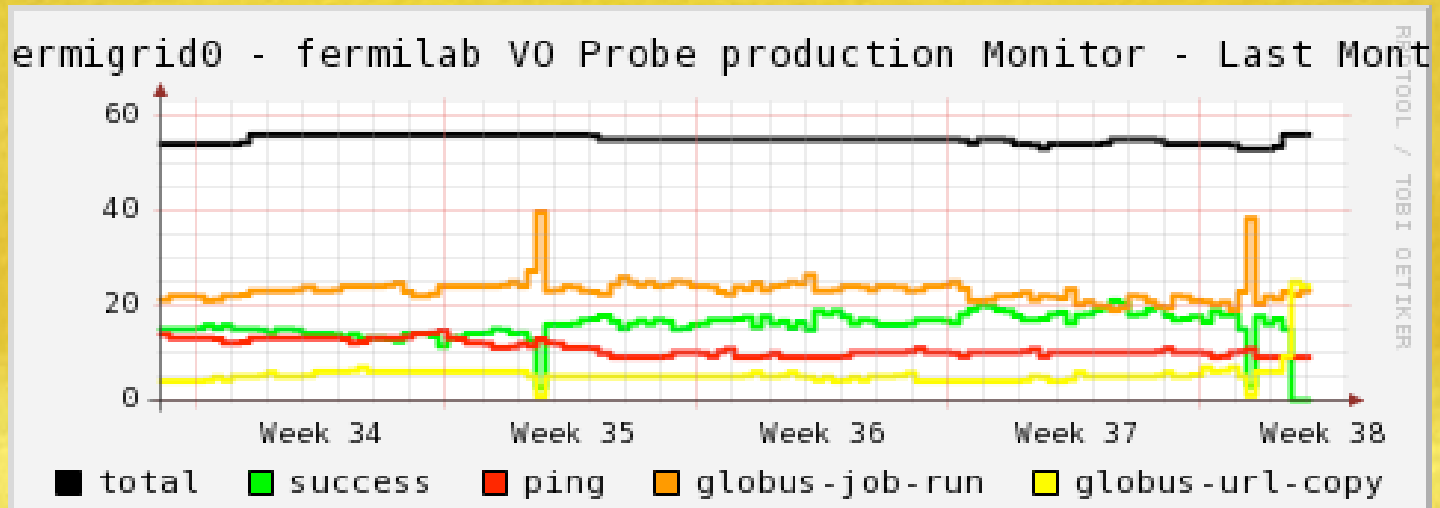

	Maximum	Average	Minimum	LastVal
total	56	55	53	56
success	21	16	0	0
ping	15	11	9	9
globus-job-run	40	23	19	23
globus-url-copy	25	5	1	24

Data for fermigrid0 between 20-Aug-2006 and 21-Sep-2006
Plot generated on 21-Sep-2006

fermilab VO probe - ITB

Open Science Grid

http://fermigrid.fnal.gov/monitor/fermigrid0-fermilab-vo-integration-monitor.html

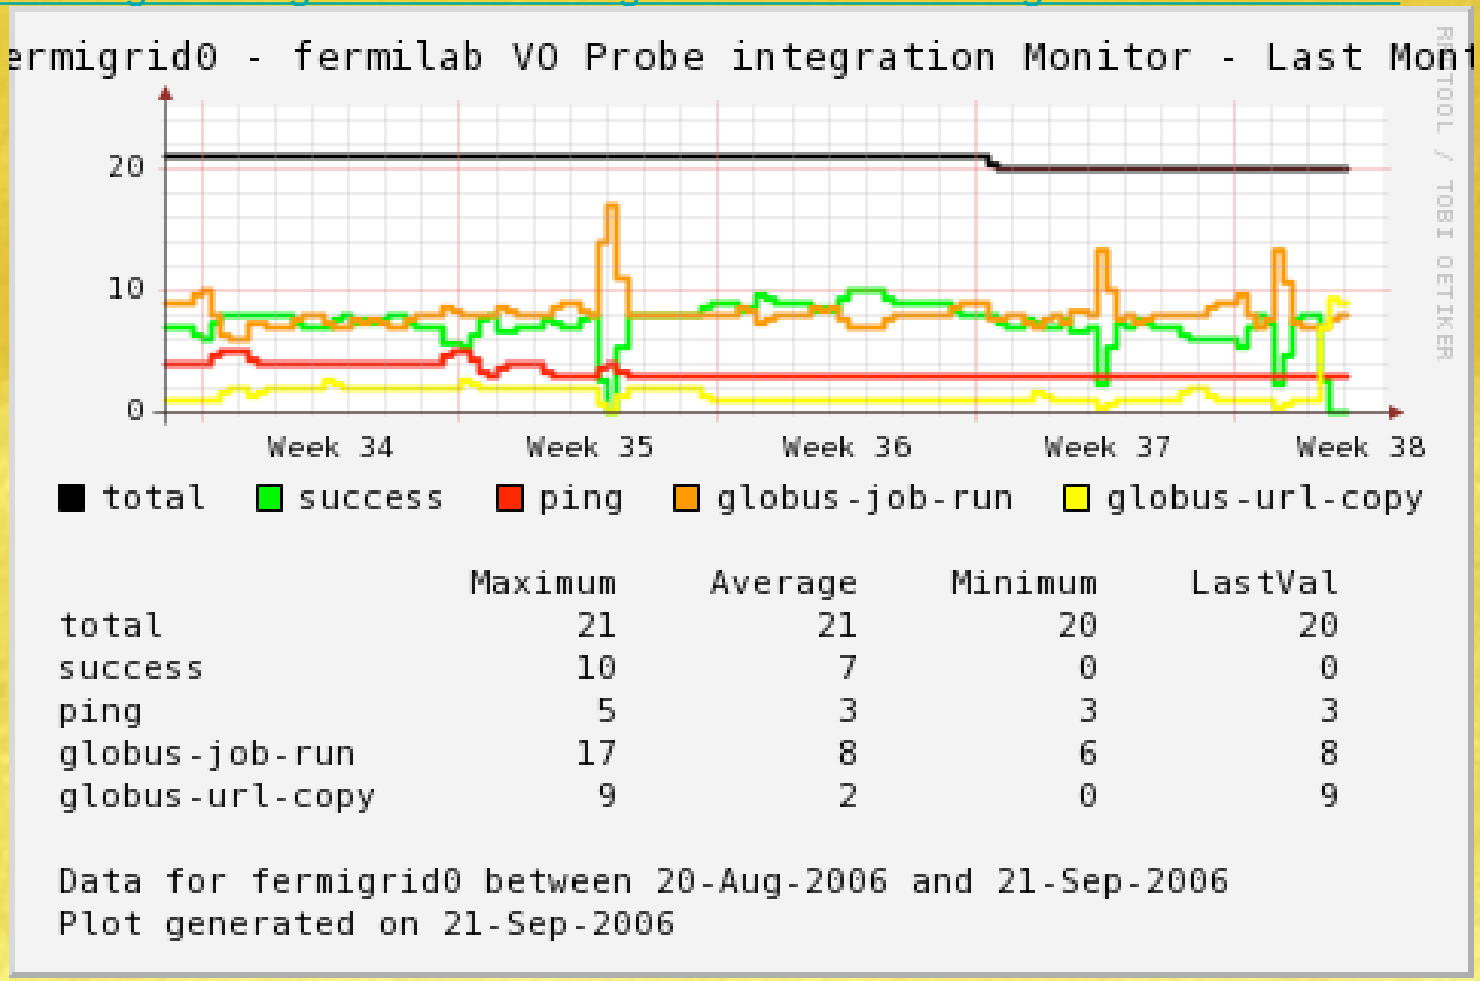

	Maximum	Average	Minimum	LastVal
total	21	21	20	20
success	10	7	0	0
ping	5	3	3	3
globus-job-run	17	8	6	8
globus-url-copy	9	2	0	9

Data for fermigrid0 between 20-Aug-2006 and 21-Sep-2006
Plot generated on 21-Sep-2006

FermiGrid - Current Work

Open Science Grid

- Policy, Security and Authorization Infrastructure:
  - Complete development of Site AuthoriZation (SAZ) Service and deployment (target is Oct 1, 2006).
  - gLexec development and deployment (target is Oct 1, 2006).
  - Fermilab Open Science Enclave enumeration, trust relationships and policy.
- Public storage and storage element:
  - FNAL_FERMIGRID_SE - public dcache storage element - 7 TBytes.
  - Migration to BlueArc NFS appliance storage.
- Further Metrics and Service Monitor Development:
  - Generate appropriate tickets to the OSG GOC for sites which are listed in VORS but fail fermilab VO probe.
  - Enhance and extend Metrics and Service Monitors.
- Service Failover:
  - At the present time the services are running in non-redundant mode.
  - We are thinking about the best ways to implement service failover.
  - We have a test fault tolerant GUMS server in operation.
  - Linux-HA clustering, BlueArc, XEN, and other technologies are being looked at.
- Research & Development & Deployment of future ITBs and OSG releases:
  - Ongoing work…
  - testbed systems are essential so that we can perform the research, development & integration without impacting our production services.
  - Looking at XEN to facilitate this work also.