

# Gratia Status

Penelope Constanta, Chris Green,  
Jeff Mack, Ken Schumacher (FNAL/CD)  
Greg Richardson (Piocon Technologies)

Philippe Canal (FNAL/CD)

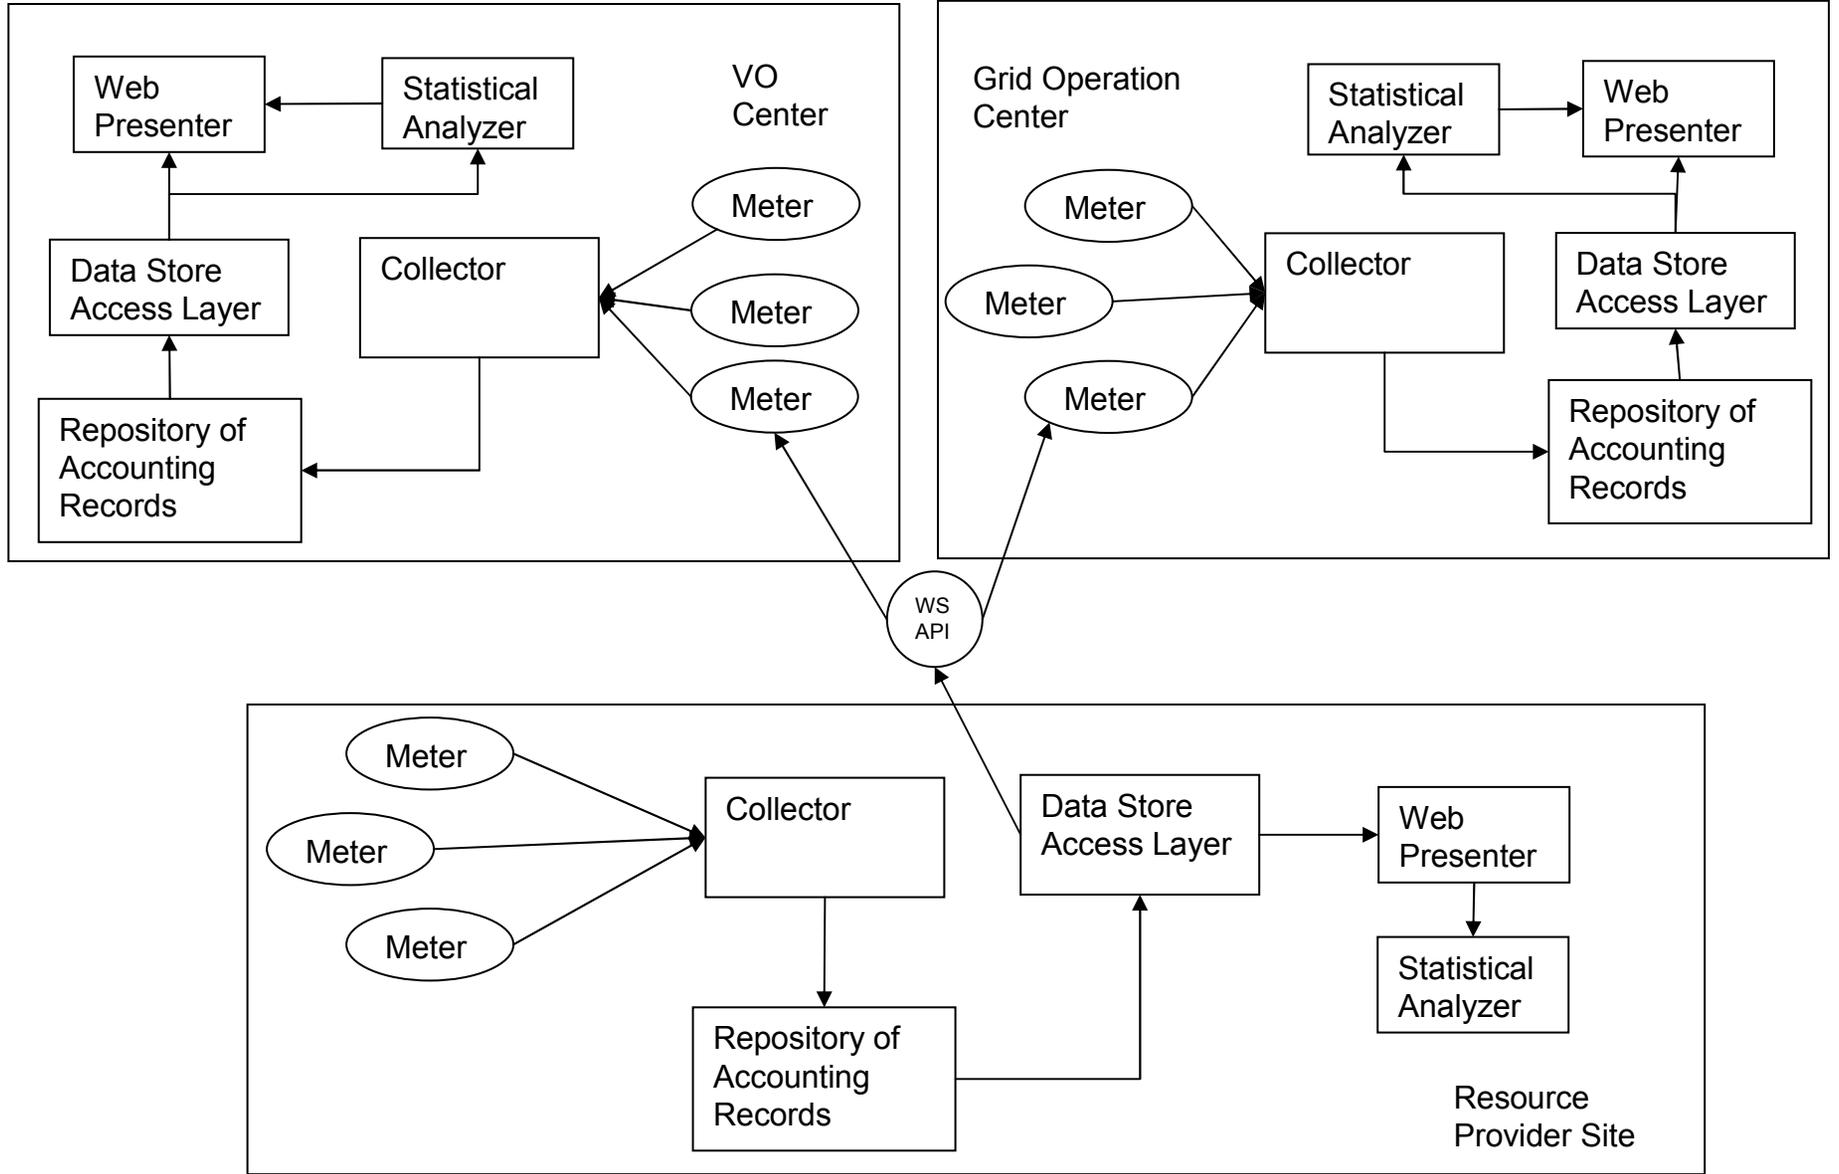
November 21<sup>st</sup>, 2006

# Gratia Status

- Gratia Architecture Overview
- New Fermilab Process Accounting
- OSG Grid Accounting
- Grid Accounting Inter-operability
- Gratia Project Plan
- OSG Grid Metrics

# Contacts

- More Information
  - [Project Definition](#), Requirements and Design Documents
  - [OSG Accounting TWiki](#) and [mailing list](#) (osg-accounting@openscience.org)



# Collectors

- Main Functionalities
  - Keep track of which Probes are reporting
  - Receive information from the Probes
  - Correlate or complete the information
    - For example looking up which VO a UID is mapped to.
  - Apply filters and policies
  - Store the information in the 'local accounting data store'

# Probes

- Ideally they would be the data producers themselves
- Fall back solution is the typical log parsing with all its drawbacks
  - When is it run? Where is the log file? How long is it kept? Does it really contain all the information needed? etc.
- Another alternative is active query
  - condor\_q for example
  - but there are performance issues in some cases
- In either cases, the ‘Meter’ should only depend
  - on the ‘tool’ (condor, pbs, etc.)
  - and the simple, lightweight Accounting Interface Library
  - (and should be eventually distributed with the tool)

# Probe – Collector Connection

- A Probe uses a secure channel to send usage information to the Collector, the usage data header contains:
  - creation time of the record
  - time interval covered by the accounting data
  - Resource/Service ID
  - UID or DN
  - Optional Global Request ID
  - Structure of the accounting data that follow and size
- Main information entered (mostly) via a name/value pair
  - The accounting connection library will format this into an XML files
  - This XML file follows the OGF UR XML Format

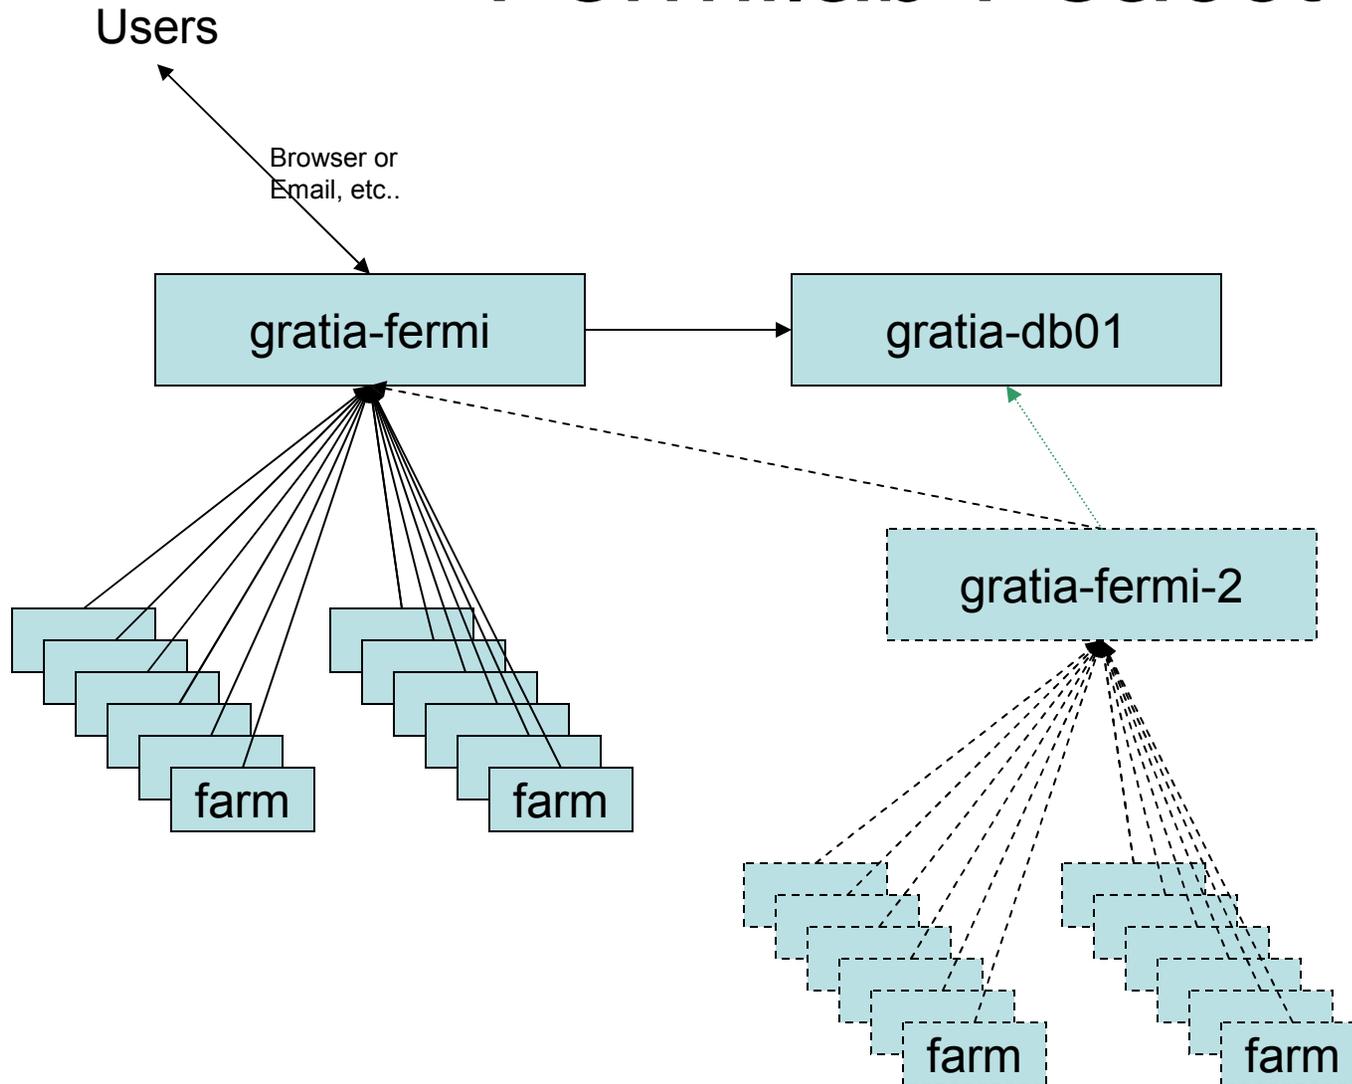
# Avoiding Data-loss

- One of our goals is to reduce as much as possible the amount of data loss due to transient problems.
- Another goal is to make the Accounting Interface as simple to use as possible
- In consequence we will provide a library which will establish and manage the network connections
- In case where the connection between two elements is broken the information that would have been sent is record in a local data file.
- As soon as the connection is restored the local data file is send to the next element.
  - Example of elements pairs:
    - Probe-Collector
    - Collector – Local Accounting store

# Fermilab Psacct

- Probe ready for deployment
- Has been running on USCMS-FNAL-WC1-CE since May 2006.
- Will be installed next on the general purpose farm.

# Fermilab Psacct



November 21st, 2006

Gratia Status

10

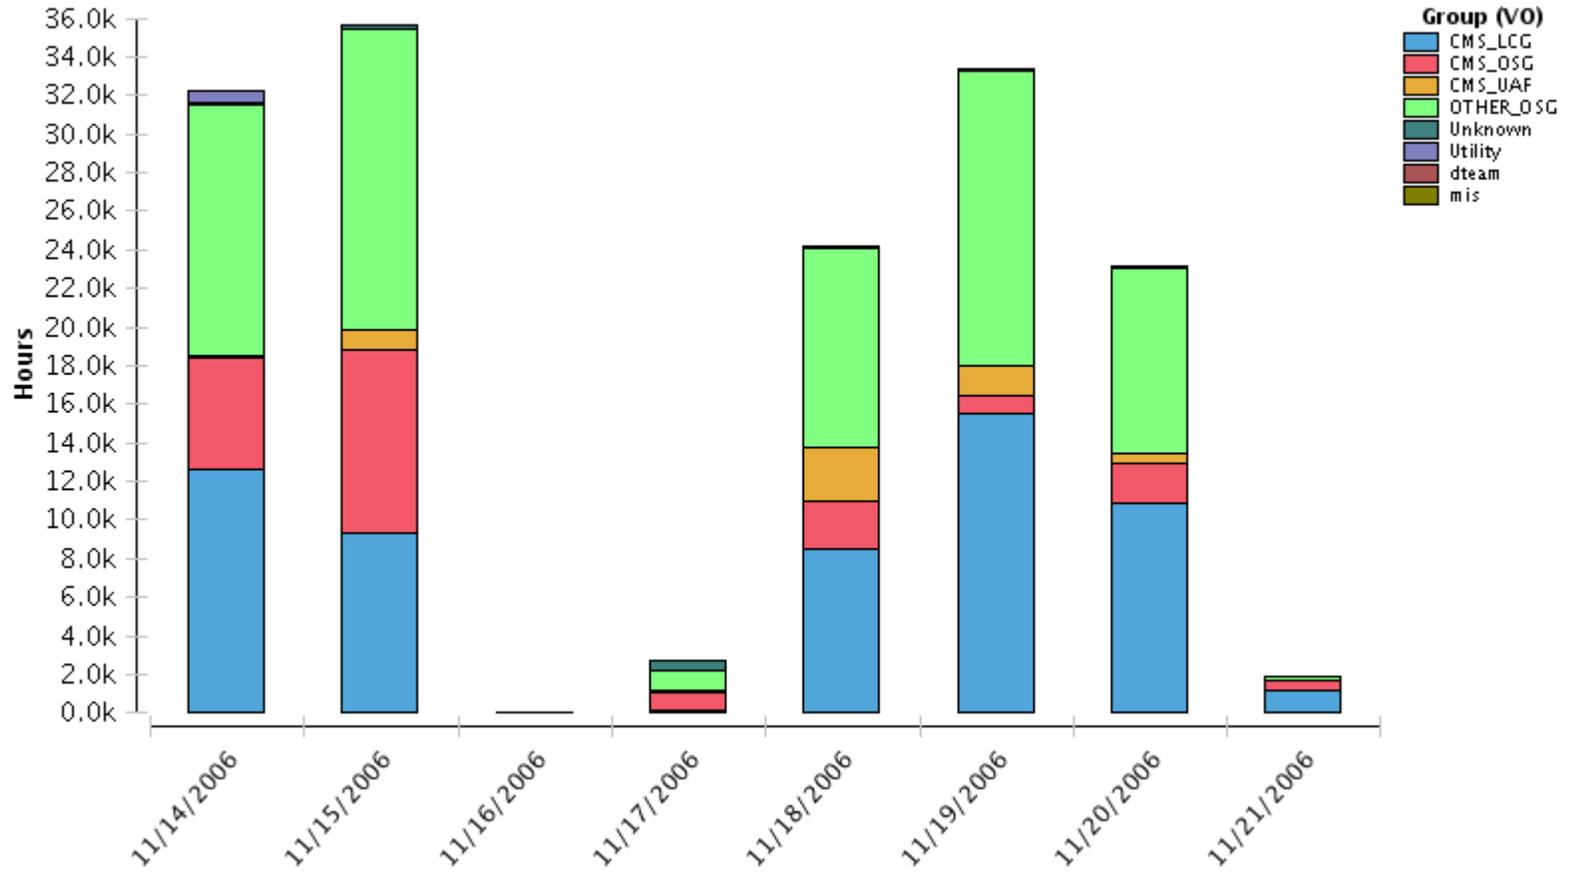
# Fermilab Psacct

- Installation:
  - Instructions for installation and configuration can be found in the [OSG Twiki](#)
  - 2 simple rpm (Gratia Common and Gratia psacct)
  - Relies on standard linux psacct
  - Supersedes 'facct' (which needs to be disabled)
    - Can 'catch-up' the information for the current calendar month
  - On the worker nodes the Gratia probe:
    - Processes the accounting file once day and contact to server
    - If communication with server fails, worker node keeps the message until the next (successful) data upload.
    - Keeps backup copy of raw accounting file for 30 days (configurable)
  - 'Adding' a node to an existing farm only requires the installation of the Gratia psacct probe.

# Fermilab Psacct

- Administration:
  - PC Performance Ratings.
    - Probe send information about the worker node
    - Server keeps a list of performance rating for each PC description (not hostname).
  - Groups:
    - Link 'userid' and 'group' stored only in the central server on a 'per farm' basis.
    - The fact that a worker node belong to a 'farm' is entered in the probe configuration file on the worker node
  - *Possible enhancement*
    - *If there was a way to retrieve the list of node that are intended to belong to a farm in an automatic way (misnet, equipdb), we could add another helpful piece of information:*
      - *list of node that never reported to Gratia (usually would indicate installation problem)*

### CPU (User + System) - Process Accounting



```
bash-3.1$ ./cmsweekly.py 2006/11/01 2006/12/01
```

```
Weekly
```

VO	Wall Hours	Norm Wall	CPU Hours	Norm CPU	Wall Load	Norm Wall	CPU Load	Norm CPU
CMS_LCG	0	0	66946	91247	0.0%	0.0%	4.3%	4.4%
CMS_OSG	0	0	53632	70835	0.0%	0.0%	3.4%	3.4%
CMS_UAF	0	0	17441	23079	0.0%	0.0%	1.1%	1.1%
dteam	0	0	0	0	0.0%	0.0%	0.0%	0.0%
mis	0	0	0	0	0.0%	0.0%	0.0%	0.0%
OTHER_OSG	0	0	118984	160929	0.0%	0.0%	7.6%	7.7%
Unknown	0	0	6965	8933	0.0%	0.0%	0.4%	0.4%
Utility	0	0	1442	1848	0.0%	0.0%	0.1%	0.1%
User Total	0	0	263967	355024	0.0%	0.0%	16.8%	17.1%
Total	0	0	265408	356872	0.0%	0.0%	16.9%	17.2%

```
Other :
```

```
# of CPUS : 2182
```

```
Date : 11/2006 (30 days)
```

```
bash-3.1$ ./condor.py 2006/11/01 2006/12/01
```

```
From Condor
```

VO	Wall Hours	Norm Wall	CPU Hours	Norm CPU	Wall Load	Norm Wall	CPU Load	Norm CPU
cern	13	0	0	0	0.0%	0.0%	0.0%	0.0%
cms	94751	0	77584	0	12.0%	0.0%	9.8%	0.0%
dteam	5	0	0	0	0.0%	0.0%	0.0%	0.0%
escience	1	0	0	0	0.0%	0.0%	0.0%	0.0%
grid	1	0	0	0	0.0%	0.0%	0.0%	0.0%
User Total	94772	0	77584	0	12.0%	0.0%	9.8%	0.0%
Total	94772	0	77584	0	12.0%	0.0%	9.8%	0.0%

```
Other :
```

```
# of CPUS : n/a
```

```
Date : 11/2006 (30 days)
```

# Fermilab Psacct

– Next:

- Better reporting
- Web interface to centrally manage 'groups'.

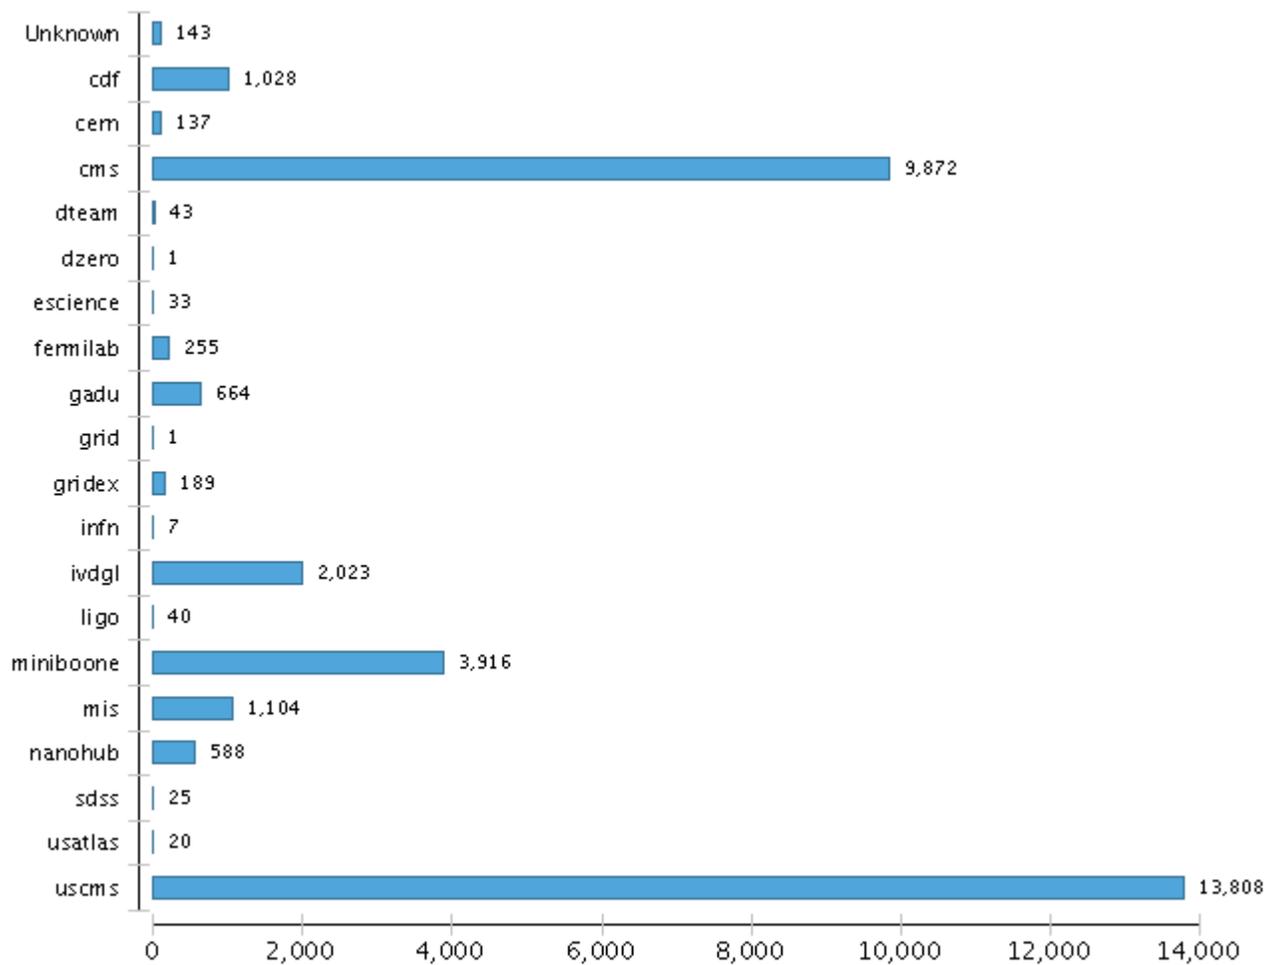
– Request:

- ***Official permission to 'operate' this service.***
- Missing info:
  - Automated list of node per farms.

# OSG Accounting

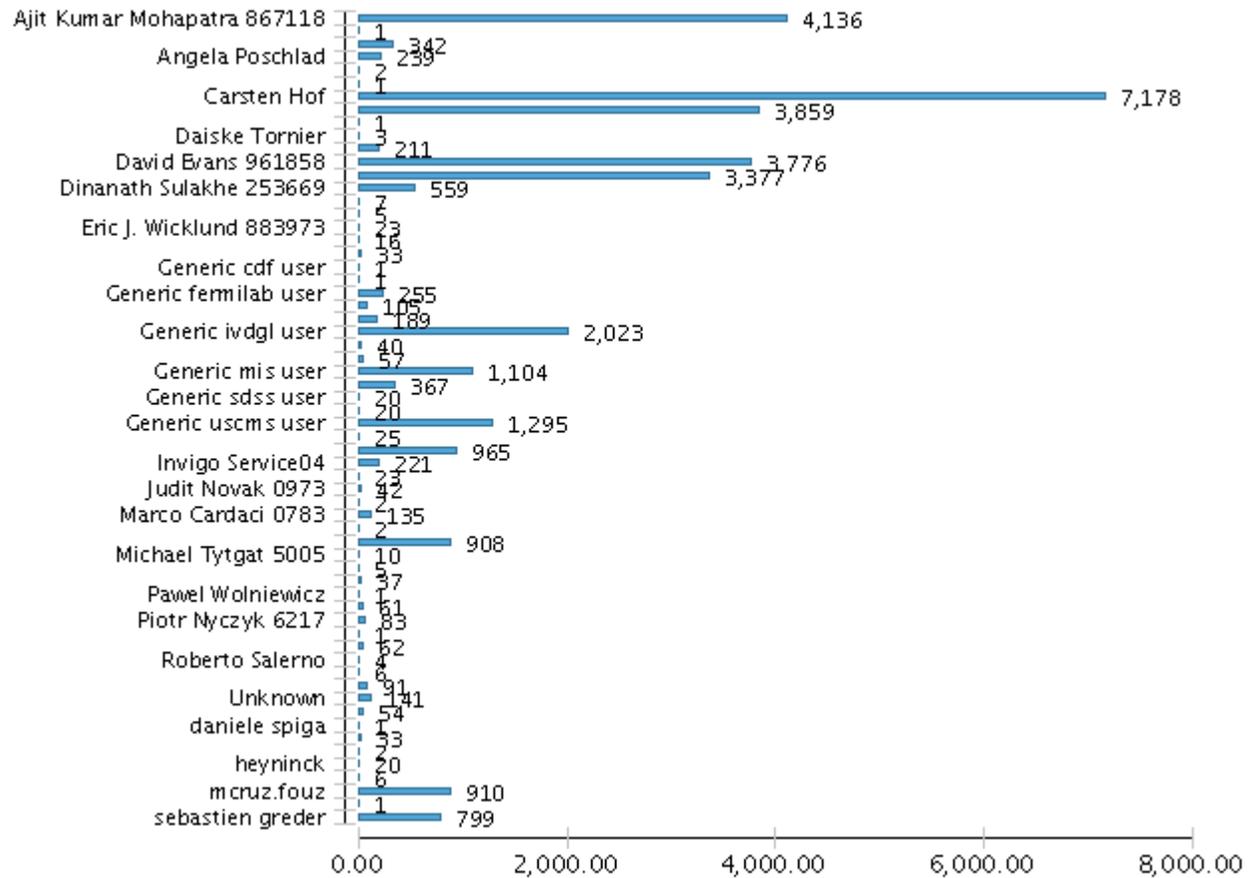
- Probe ready for deployment
  - Require patching of VDT 1.3.10 (available)
  - Part of VDT since 1.5.0
  - Being tested on the OSG ITB
  - Started push for install on cms tier 2.
  - Development for Condor 6.9.x done by the Condor team itself.
  - Still might need update of the probe for VDT 1.3.9 and WS-GRAM.
- Collector
  - Have done stress testing
  - Ready to enable authentication on both the upload side and the report side.
    - However no roles yet due to technical difficulties with voms-proxy
  - Still need to be connected to the OSG Alarm system (footprint)
    - stalled/phone tag
  - Good working relationship with VDT and Condor team
- Next:
  - Automatic discovery of local user id/VO relation ship (either via GUMS or via grid map file)
  - Flesh out the role and policy design and implementation on the report side.
  - Better/Extended user interface.
  - Deployment of Collectors (requires more testing, authentication features to be more mature)

### Job Count by VO



Date range: 2006-11-14 00:00:00 GMT - 2006-11-21 23:59:59 GMT

### Job Count by User



For all jobs finished on 2006-11-12 (Central Time)

	Site	# of Jobs	Wall Duration	Delta jobs	Delta duration
1	Cmslsgce	75	8	-6	-11
2	Cmsosgce	1,769	9,605	-2,716	-4,781
3	FNAL_FERMIGRID	352	47	20	1
4	Purdue-Lear	5	3	-216	-71
5	Purdue-ITaP	11	15	-130	-7
	All sites	2,212	9,678	-3,048	-4,869

For all jobs finished on 2006-11-12 (Central Time)

	VO	# of Jobs	Wall Duration	Delta jobs	Delta duration
1	cdf	388	1,105	-183	-75
2	cern	14	1	-5	0
3	cms	53	7	0	-11
4	dteam	4	0	-1	0
5	escience	4	0	0	0
6	fermilab	43	24	1	1
7	gadu	6	0	6	0
8	gridex	2	15	-25	-6
9	ivdgl	4	0	-250	-1
10	miniboone	908	4,066	113	-2,696
11	mis	166	0	14	0
12	nanohub	147	23	2	0
13	osg	1	0	1	0
14	uscms	472	4,437	-2,721	-2,081
	All VOs	2,212	9,678	-3,048	-4,869

For all jobs finished on 2006-11-12 (Central Time)

	Site	VO	# of Jobs	Wall Duration	Delta jobs	Delta duration
1	Cmslsgce	cern	14	1	-5	0
2	Cmslsgce	dteam	4	0	-1	0
3	Cmslsgce	cms	53	7	0	-11
4	Cmslsgce	escience	4	0	0	0
5	Cmsosgce	uscms	469	4,434	-2,650	-2,010
6	Cmsosgce	gadu	3	0	3	0
7	Cmsosgce	miniboone	908	4,066	113	-2,696
8	Cmsosgce	osg	1	0	1	0
9	Cmsosgce	cdf	388	1,105	-183	-75
10	FNAL_FERMIGRID	gadu	3	0	3	0
11	FNAL_FERMIGRID	mis	166	0	14	0
12	FNAL_FERMIGRID	fermilab	36	24	1	1
13	FNAL_FERMIGRID	nanohub	147	23	2	0
14	Purdue-ITaP	ivdgl	2	0	-105	-1
15	Purdue-ITaP	fermilab	7	0	0	0
16	Purdue-ITaP	gridex	2	15	-25	-6
17	Purdue-Lear	uscms	3	3	-71	-71
18	Purdue-Lear	ivdgl	2	0	-145	0
	All sites	All VOs	2,212	9,678	-3,048	-4,869

From the results sent directly to the Summary database

For all jobs finished on 2006-11-12 (Central Time)

	Site	VO	# of Jobs	Wall Duration	Delta jobs	Delta duration
1	OUHEP	DOSAR	1	83	0	46
2	OUHEP	DZero	1	876	0	-47
3	OU_OCHEP_SWT2	ATLAS	1	1,968	0	25
4	OU_OSCER_CONDOR	DOSAR	1	422	0	-79
5	OU_OSCER_CONDOR	DZero	1	53	0	-6
	All sites	All VOs	5	3,402	-1	-66

From the results sent to the detailed Gratia database

For all jobs finished on 2006-11-12 (Central Time)

	Site	VO	# of Jobs	Wall Duration	Delta jobs	Delta duration
1	Cmslsgce	cern	14	1	-5	0
2	Cmslsgce	cms	53	7	0	-11
3	Cmslsgce	dteam	4	0	-1	0
4	Cmslsgce	escience	4	0	0	0
5	Cmsosgce	cdf	388	1,105	-183	-75
6	Cmsosgce	gadu	3	0	3	0
7	Cmsosgce	miniboone	908	4,066	113	-2,696
8	Cmsosgce	osg	1	0	1	0
9	Cmsosgce	uscms	469	4,434	-2,650	-2,010
10	FNAL_FERMIGRID	fermilab	36	24	1	1
11	FNAL_FERMIGRID	gadu	3	0	3	0
12	FNAL_FERMIGRID	mis	166	0	14	0
13	FNAL_FERMIGRID	nanohub	147	23	2	0
14	Purdue-ITaP	fermilab	7	0	0	0
15	Purdue-ITaP	gridex	2	15	-25	-6
16	Purdue-ITaP	ivdgl	2	0	-105	-1
17	Purdue-Lear	ivdgl	2	0	-145	0
18	Purdue-Lear	uscms	3	3	-71	-71
	All sites	All VOs	2,212	9,678	-3,048	-4,869

November 21st, 2006

Gratia Status

# Gratia Performance

Note: The tests were run (for the most part) with default settings.

1. Gratia can handle somewhere between **100K -> 120K transactions per hour** with a local database (a database on the same machine as gratia). This is the gratia-> database rate.
2. This rate will drop roughly 20% if the database is on a separate machine (network latency?)
3. The overall cpu utilization during testing was about 60% with the timing being split between gratia and mysql. Note that this was mysql 5.x with the fancy triggers in place.
4. On the input side (probe -> gratia), gratia appears to be able to handle about **300K+ transactions per hour**.
5. There is a significant memory increase required as the number of simultaneous, high volume inputs increases. In a word, if 500 CE's suddenly decide to replicate 100K records to gratia at the same time, we will have problems.
6. The transaction rate (1) might be increased by increasing the number of database update threads. I'll test this when psg5 comes online.

As a result of (5), I've added a "governor" to gratia. If any queue size exceeds some magic number (currently 10K), I will shut down input processing - e.g. - the probes won't be able to send additional information to gratia until the queue size has been cut in half. Note that by default gratia comes up with 3 queues.

In our current production, the queue size should usually be close to 0. In general, the production gratia service is being hit with about **5K transactions per day**.

Finally, mysql appears to behave in a very strange fashion when one runs out of disk space (I did). It hangs. Stops. Dead on arrival. Threads lock. I'm still investigating this.

# Gratia Performance

Note: 2<sup>nd</sup> series of tests we 3 separate servers: PSG1 (data supplier) -> PSG3 (gratia) -> PSG5 (database)

Test 1: Single gratia instance running 10 database update threads:

PSG3 cpu utilization = 30%  
PSG5 cpu utilization = 10%  
Transaction Rate = 237K per hour

Test 2: Single gratia instance running 20 database update threads

PSG3 cpu utilization = 40%  
PSG5 cpu utilization = 10 -> 15%  
Transaction Rate = 247K per hour

Test 3: 2 gratia instances running 10 update threads each

PSG3 cpu utilization = 50 -> 70%  
PSG5 cpu utilization = 10 -> 15%  
Transaction Rate = 340K per hour

Upon further investigation there appears to be a fair amount of thread blocking taking place within tomcat as well as within hibernate. Tomcat is generally considered a high volume server. Hibernate does all sorts of synchronization things within itself. As a result, if one expects to handle very high volumes, it's better to have multiple gratia's running with 5 -> 10 update threads each rather than a single gratia with a great many update threads.

We weren't able to detect any network bottlenecks or mysql bottlenecks during this process. We were using the "huge" model of mysql which we suspect can handle much higher update loads than presented in these tests.

# Gratia WBS

- Amongst the progress made:
  - stress testing
  - Progress: Implementation of data filtering
  - Worked on: reproducing data quality issues
  - Optimized report timing
    - Automatic generation of 'summary' tables
  - Upgrade disaster recovery mechanism
  - Finalized PBS probe.
  - Finalized rpm and patch mechanism
- Items that were done and were not explicitly in the previous WBS:
  - Need for good, easy to use release mechanism outside of VDT (to speed up deployment before OSG 0.6.0) [This is effort is now done]
  - Need to fix problems and finish PBS probe [This is effort is now done]
  - Special Deployment on QCD cluster [This is effort is now done]
- Things that were added or changed since the last WBS:
  - Revised estimate for Storage probe. (6.1)
  - Revised estimate for Implementation of Role based authentication (2.1.1)
  - Revised estimate for analyzing the 'apparent' data discrepancies (1.1)
  - Revised estimate for role implementation 2.1.3
  - Added section on OSG reports (4.6)
  - Added section on interoperability (7)
  - Added section on Grid Metrics (9) ...
  - Reflected progress made since the last wbs.
- Total FTE counted is on 2.5, from which must be deducted .25 (see 1.2)
- Total FTE described below 79 FTE weeks, not including the work on Metrics.
- At the current staffing level this is 35 weeks.

# WBS: Short Term for the Gratia Project.

## Date: November 17.

### 1. Improve Robustness and data quality in Gratia.

- 1.1 Discover a way to reproduce the case where condor under report the job being run.  
5 weeks FTE - started but difficulties and more oddities found.
- 1.2 On going support, new install, new VO, new features, meeting, etc.  
Infinite (and to be subtracted from the available FT.)  
.25 week FTE (every week)
- 1.3 Follow through on PBS/LSF probe installation. Make sure Rosario incorporates our patches (1 FTE week) - Done.
- 1.4 Work with the Condor team to improve/finalize the condor probe.  
2 week FTE (on our side) - started
- 1.5 Work with BNL to deploy SunGrid probe  
1 week FTE
- 1.6 Update existing condor probe to support WS-Gram (and maybe VDT 1.3.9)  
2 week FTE

### 2. Authentication, Authorization, Privacy and Role checking.

- 2.1 Authorization:
  - 2.1.1 Implement (possibly using the code from SASZ) the layer extracting information from the VOMS certificated.  
5 weeks FTE - Difficulties in getting the voms-proxy from the user to the server.
  - 2.1.2 Design the set of roles and policies needed for proper protection of the reports and services.  
2 weeks FTE
  - 2.1.3 Implement the set of roles and policies needed for proper protection of the reports and services.  
3 weeks FTE - started
- 2.2 Encryption: Implement encryption of the DN during transfer.  
2 weeks FTE

### 3 Server Development

- 3.1 Polishing and Documentation 'interfaces' to manage the probes and compute elements.  
2 weeks FTE
- 3.2 Check and Strengthen the 'probe black out detection' mechanism.  
4 weeks FTE
- 3.3 Continue to test reliability and scalability.  
1 weeks FTE
- 3.4 Implement proper connection to the gums. Currently this is emulated by storing on the server side the reversed grid map file (local user id to VO) per compute element (file name UserVoMap?.hostname were hostname is the fully qualified name of the head node).  
2 weeks FTE
- 3.5 Implement a series of test to check that an install is fully functional.  
4 weeks FTE
- 3.6 Connect the Gratia servers to the OSG Alarm systems.  
2 weeks FTE - stalled
- 3.7 Implement DN lookup in the case of jobs that were submitted to condor by means OTHER than the Globus gram. Possibly relying on the upcoming (standard) log file containing the need information. Length depend on the availability of this file
- 3.8 Write an administration guide for the collector and the reporting tools  
1 weeks FTE
- 3.9 Work with the VDT team to properly package and update the Probe and Collector for the next production release  
2 weeks FTE

# WBS: Short Term for the Gratia Project.

## Date: November 17.

### 4. Reports

- 4.1 Continue development of the PSacct reports required by Fermilab Computing Division.  
7 weeks FTE
- 4.2 Update the reports to take in consideration the information gathered from the 'probe black out detection'.  
4 weeks FTE
- 4.3 Write an administration guide fo the reporting tools  
1 weeks FTE
- 4.4 Implement policies (check) in the reporting tools.  
2 weeks FTE
- 4.5 Monthly reports to LCG in behalf of CMS  
1 day FTE per month (includes verifying and checking data).
- 4.6 Polish OSG reports  
3 weeks FTE

### 5. Deployment

- 5.1 Deploy PSACCT probe at Fermilab.  
1 week FTE (on our side)
- 5.2 Write an implementation/administration plan for PSacct at fermi  
1 weeks FTE
- 5.3 Write a deployment/integration/operation plan for OSG  
1 weeks FTE
- 5.4 Deploy the Condor and PBS Probe across the OSG.  
5 weeks FTE

### 6. Medium term project.

- 6.1 Write/test Storage probe  
10 weeks FTE

### 7. Interoperability

- 7.1 Monthly coordination meeting with TeraGrid, EGEE and OSG accounting effort.  
.125 week FTE per month
- 7.2 Standardization effort  
Develop with TeraGrid, EGEE and OMII a common proposal for an extension of the standard for atomic job usage rercord and for a new aggregate usage record xml format.  
2 weeks FTE

### 8. Long Term project

- 8.1 Write/test Network probe  
TBA

### 9. Grid Metrics

- 9.1 Review and understand the general OSG requirement  
Includes figuring out which requirements can be satisfied and/or approximated with the existing tools.  
6 weeks FTE
- 9.2 Review planning and long term prospects.  
(Most likely will extend the information provided below)  
.5 weeks FTE
- 9.3 Implement the requirements that relies on Gratia  
3 weeks FTE
- 9.4 Push/Prode/Help other information providers to extend (if necessary) their reports  
10 weeks FTE

# WBS: Conclusion

- Still quite a bit to do
- Increase in number of reporting sites (more than 10 fold) **could** uncover new difficulties.
- Need to preserve current staffing for 6 to 8 months
  - Addition of Grid Metrics will shuffle some of the timeline.
- Greg is focusing on server development and documentation
  - This would end up being transitioned to Philippe and Chris.

# EGEE/TeraGrid/OSG

- Maintain contact between the 3 efforts
  - Some clear overlap (UR XML format, role based access to the data, data collection).
  - Probe development effort is the main area of collaboration.
    - EGEE wrote the PBS/LSF probes
    - Fermilab wrote the Condor probe, OSG will write the SunGrid probe.
  - Meet once a month via VRVS.
- UR Record Work
  - Current GGF UR XML is a very good foundation
  - Only support 'atomic' record about a single job
  - Need to slightly extend it to completely fit our need:
    - VOName, 'Global' Resource Identifier, etc.
  - Does not yet address 'Aggregate' which we need for summarized inter-grid exchange.
  - Our input not really addressed by the standard committee.
  - Will send in a join (EGEE, TeraGrid, OSG, OMII) request/proposal to the committee. Hope to send it in time for presentation at the January GGF meeting.

# What's next

- Automate UserId, VOName link retrieval
- ***Continue to 'verify' data.***
- Deployment, deployment and deployment
  
- Improve Server user interface
- Extend and automate testing.
- Extend documentation.
- Storage probe (and network?).
- Report, reports and more reports
  - extend the report to fully replace the old facct reports
  - extend the report to include an error estimate (based on detected black out periods).
  
- [OSG] Grid Metrics

# Grid Metrics

- Next 2 slides describes the current status of the OSG regarding their need for Metrics
  - Courtesy of Ruth Pordes

# Deliverables of OSG Metrics?

- *Availability of Resource and Services - as measured by the OSG site functional tests and validation infrastructure.*
- *Resource use of each VO – wall clock time, data moved, storage used – on a monthly bases - measured by the Accounting system.*
- *Resources provided by each Site – throughput in terms of processing and data movement, and storage used - measured by the Accounting system.*
- *Shared resource usage – what storage and processing is used by VOs on resources that they don't own. Measured by the Accounting system.*
- *Throughput of OSG – number of submitted and executed jobs per day, number of files transferred per day - measured by the Accounting system.*
- *Latency of the OSG – the time a job spends in the wait queue, the wait time for a request for transfer from storage or between resources. Measured by the VOs?*
- *Success rate of jobs per VO and per site. Success is to be measured by the VO in terms of end-to-end success rate and at the entrance and exit of the interface to grid middleware and services. Included in this metric is the “retry” rate of automated retrial of actions that either time out or have other errors. Measured by the VOs?*

# Roadmap for OSG metrics?

- *Agency reportable milestones to specify the metrics, with numeric goals, in the yearly project plan and to meet the metrics by the end of each year of the project.*
- *First year project plan to Agencies by 1 Dec (probably 20th Dec in practice).*
- *First Joint Oversight Team meeting expected in February -- agencies still organizing themselves.*
- *Finance Board meetings will review and feedback on the metrics on an ongoing basis (probably co-scheduled with Executive Board meetings every 6 weeks).*
- *So--- first metrics report is naturally aligned with Consortium all hands week of 5th March 2007.*

# Remarks on OSG Metrics

- 4 out of 7 of the metrics comes naturally from Gratia (Number of jobs per CE, number of work nodes used, number of VOs served on each CE, etc..) but calls for a storage probe (not yet available)
- Overlap in concept and people involved and/or to be contacted with the Grid Accounting projects.
- In this context it makes some sense to extend the charge of the Fermilab Grid Accounting group to look into OSG metrics
- Dateline (March 2007) is extremely tight given the current staffing.

# Concluding Remarks

- Sound Architecture
- Many “*details*” to be filled out
- Good Collaboration with most partners
- Grid Metrics?
- Question/Comment/Grips?