Service, Operations and Support Infrastructures in HEP

Processing the Data from the World's Largest Scientific Machine

Patricia Méndez Lorenzo (IT-GS/EIS), CERN

ISGC 2008, Academia Sinica Taipei

Worldwide LHC Computing Grid

Distributed Production Environment for Physics data Processing

Outlook

GRID Environment

Users

How to support, attract... and maintain them

Services

Maintenance, operations

Application to WLCG

Experiences we gained within that Grid infrastructure

- User support and high reliable services are key aspects of any Grid project
- In this talk we will address the status of the application support, the operations and the services applied to WLCG and the lessons we learnt with the gridification of the LHC experiments

The LHC Grid Computing

- LHC will be completed this year and run for the next 10-15 years
- Experiments will produce about **15** **M**
(a
- Ll
co
ar
p
- R
ce
2(

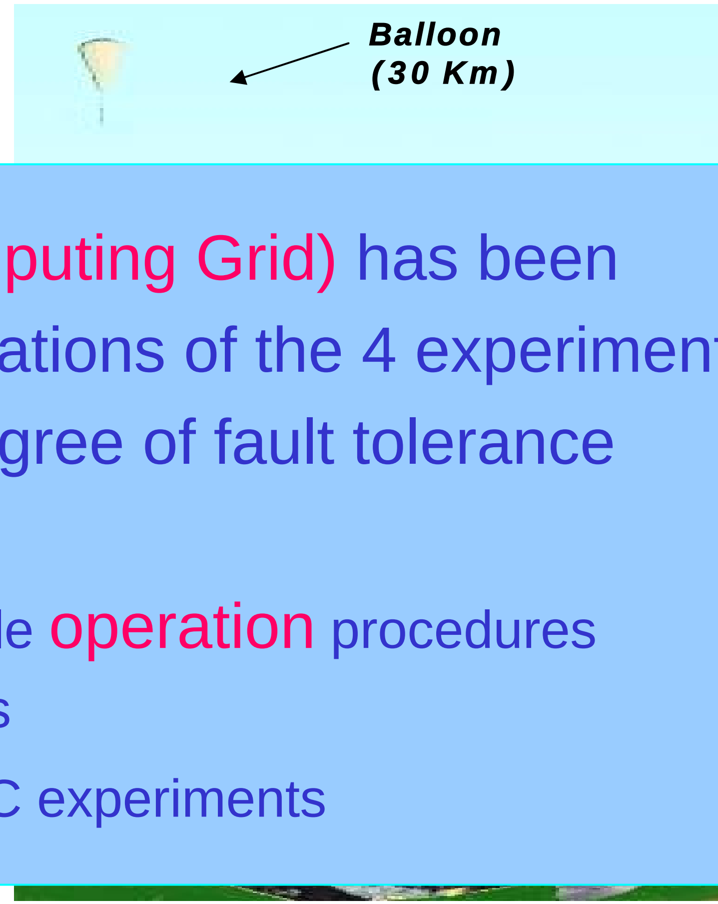

- The WLCG (Worldwide LHC Computing Grid) has been built to cover the computing expectations of the 4 experiments
- WLCG needs to ensure a high degree of fault tolerance and reliability of the Grid services
  - Through a set of effective and reliable operation procedures
    - Applicable to T0, T1 and T2 sites
  - Individualized support for the 4 LHC experiments

The Grid solution

- The WLCG technical report lists:
  - Significant costs of maintaining and upgrading the necessary resources … more easily handled in a distributed environment, where individual institutes and organizations can fund local resources whilst contributing to the global goal
    - Clear from the funding point of view
  - … no single points of failure. Multiple copies of the data, automatic reassigning of tasks to resources … facilities access to data for all scientists independent of location. … round the clock monitoring and support.
    - Similar to the 3rd point of the "Grid Checklist" following I. Foster`s book
    - … "*To deliver nontrivial qualities of service*"… The utility of the combined system is significantly greater than the sum of its parts
      - Clear example with FTS: Able to stop transfers by its own in case one SE is in maintenance

Projects we use for LHC Grid

- WLCG depends on 2 major science grid infrastructures provided by:
  - EGEE – Enabling Grid for E-SciencE (160 communities)
  - OSG – US Open Science Grid (also supporting other communities beyond HEP)
  - Both infrastructures are federations of independent GRIDs
  - WLCG uses both OSG resources and many (but not all) from EGEE

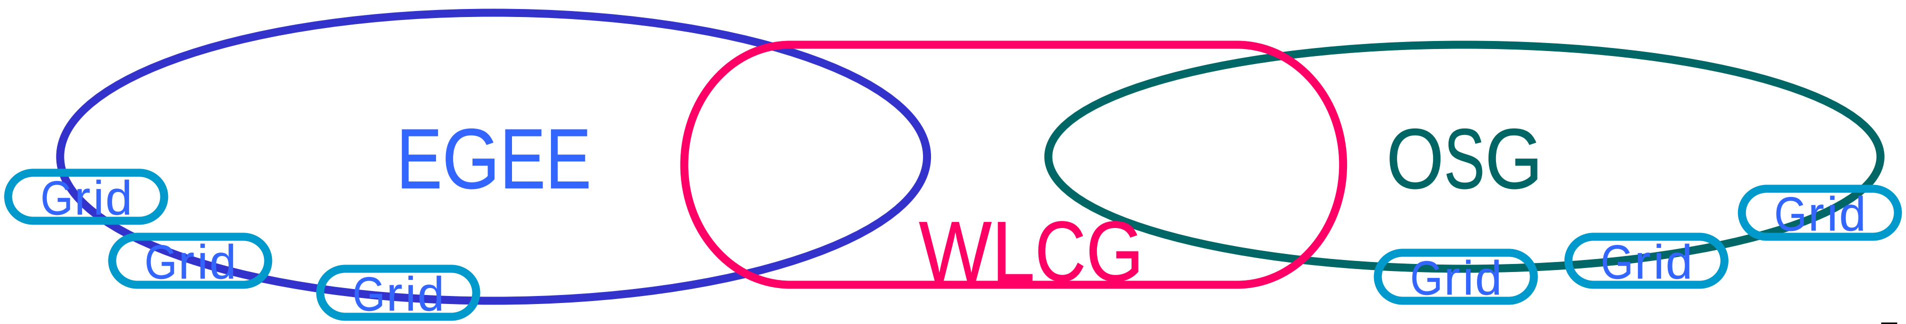

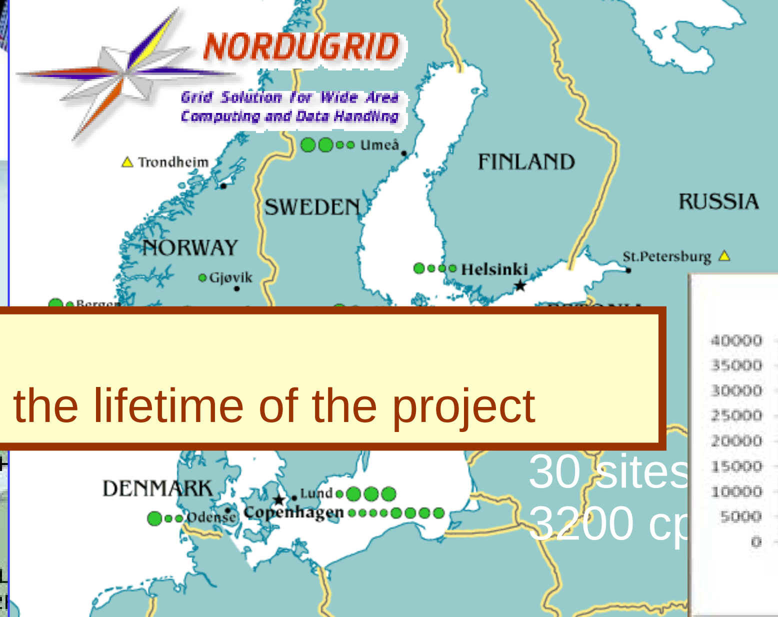

EGEE:
Steady growth over the lifetime of the project

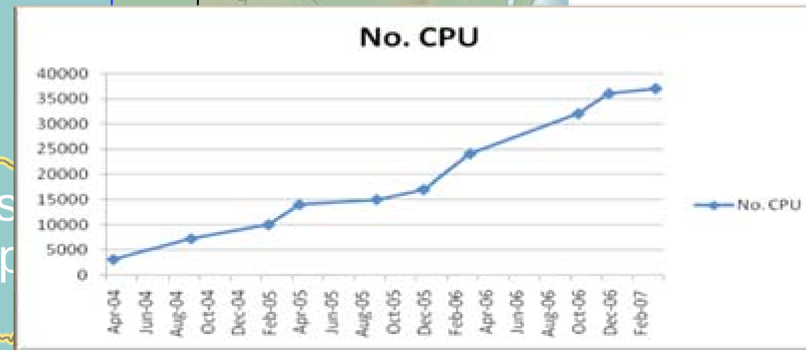

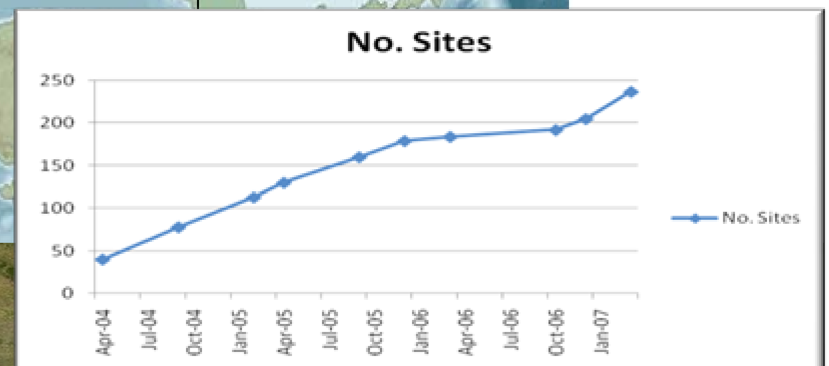

EGEE:

- > 237 sites, 45 countries
- > 35,000 processors,
- ~ 10 PB storage

The WLCG Distribution of Resources

Tier-0 – the accelerator centre

- Data acquisition and initial Processing of raw data
- Distribution of data to the different Tier's

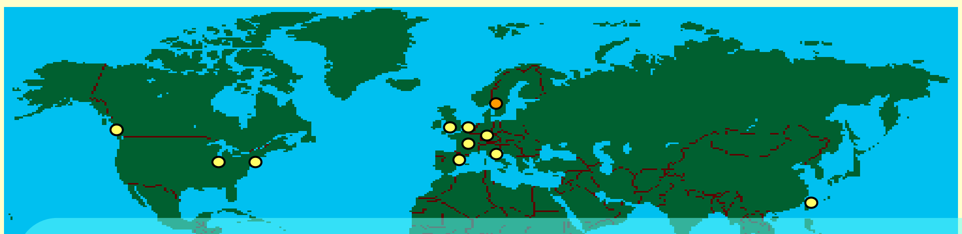

Canada – Triumf (Vancouver)
France – IN2P3 (Lyon)
Germany – Forschunszentrum Karlsruhe
Italy – CNAF (Bologna)
Netherlands – NIKHEF/SARA (Amsterdam)
Nordic countries – distributed Tier-1

Spain – PIC (Barcelona)
Taiwan – Academia SInica (Taipei)
UK – CLRC (Oxford)
US – FermiLab (Illinois)
– Brookhaven (NY)

Tier-1 (11 centers) – "online" to the data acquisition process → high availability

- Managed Mass Storage – → grid-enabled data service
- Data-heavy analysis
- National, regional support

Tier-2 – ~200 centres in ~40 countries

- Simulation
- End-user analysis – batch and interactive

Towards a production level service

- Level of the Grid services provided by WLCG
  - Signed by all sites in the MoU
    - Contains the services provided by each site per experiment, the time schedule for recovery and the level of service
  - Ranges from 99% for key services at T0 up to 95% for several services at T2 sites
    - Defined for high level functionalities more than individual services
- Specifications of the experiments: Critical services
  - Including degradation of the experiment in case of failures
- A general procedure has been defined to ensure a good level of service
  - Checklist for new services, recommendations for middleware and DB developers, operational techniques and procedures

Building a Robust Service

- Essential elements towards a robust service
  - Careful planning complemented with a set of techniques
  - Deployment and operational procedures included
- Keep in mind always the following ideas
  - **Think Service:** service is much more than a piece of middleware
  - **Think Grid:** A change to a service at a given site might affect remote users and services
- A robust service means minimum human intervention
  - Mandatory for scalability and to handle huge amount of data

The Check-list for a new service

- The following aspects must be ensured by any new service (Grid or non Grid):
  - Software implementation
    - Short glitches have to be coped with
    - Load-balanced setup among different servers is a good solution also for maintenance procedures
    - In case of a DB backend, the re-establish of lost connections is mandatory
  - Hardware setup
    - Avoid single points of failure through:
      - Power supplies
      - Network connections
  - Operational procedures and basic tests are needed
    - Standard diagnosis procedures should be ensured
    - Site manager will also have the possibility to define their own tests

Example: The FTS at the T0

- Hardware:
  - Split into several components
    - Web-service
      - DNS load-balanced over 3 nodes
      - Checked by monitoring every minute and drop problematic nodes from the load balance
    - Data-transfers
      - Channel agent daemons
        - » Balanced over 3 nodes with no redundancy
        - » Good partition: problems in one channel do not affect the rest of channels
      - VO agent daemons
        - » All available in one node with no redundancy
    - Proxy renewal
      - Load balanced in several servers
    - Monitoring
      - Placed on a single node
      - Not critical for the service operation
- Software:
  - FTS middleware can transparently handle loss of one or multiple servers
  - Service components well de-coupled from each other
    - Each component will keep running although other one is down

Operations Meetings

- The experience gained at CERN for many years shows the value of running daily operations meetings
  - No longer than 10min
  - The alarms of the day before are checked and discussed
  - Assignment of open issues to the responsible
  - If any weakness in the system, this will be exposed
  - Experiments are also welcome to join!
- Weekly operations meeting with the ROC responsible
  - Escalation of open problems mostly opened by the experiments
  - Each experiment is represented during the meeting
- Besides this procedure any operation in the system must be performed as quick and transparent as possible
  - Otherwise the degradation of the system can cause serious problems to the users and other services
- Also once the intervention is finished a post-mortem and documentation procedure is required for future interventions

Middleware deployment techniques

- Robustness and resilience must be aspects to observe during middleware design
  - High-availability setup onto an already deployed software is much more difficult
- Basic characteristics to observe
  - Modular design
    - Allows an easier debug
    - The different operational features of the system are ensure in case of individual failures
    - Example: FTS. Each daemon agent can be taken apart of individual maintenance ensuring the performance and resilience of the rest of elements
  - System architecture should be ready for glitches
    - Possible errors must be known and documented
  - Maintenance state must be minimized as much as possible
  - Apply the rule: “think service”
    - Can the system be cleanly and transparently stopped for the user?
    - Can the node be drained with no service effects?
  - Apply the same rules to a Grid environment

DB deployment techniques

- Critical part of any Grid environment
- Rules to follow:
  - Connection retries
    - The software should be ready to retry the connection if this becomes unusable
    - Multi-threaded applications. Make sure the application is written to make use of the DB High Availability features
  - Use the DB to enforce all known integrity constraints
  - Testing the new versions previously to the deployment
    - This procedure will ensure smooth service operations
  - Testing the new versions in a realistic framework
    - Experiments should participate
    - Example DB at CERN: A pre-production setup is available and experiments can test their own applications during one month before final deployment
  - For the users: always get in touch with your DBA
    - This procedure will also help you to develop more advanced applications

Critical Services

- Experiments provide their own services to implement Grid services need by their computing models
  - Failures on these services can also compromise the experiment performance
  - The same rules applied to the Grid services will have also to be applied to experiment services
- We already know that during the first months of the real data taking we will face several instabilities
  - For critical services a 24x7 service will have to be provided
  - On-call procedures will be defined for those services
  - For the rest a contact person will be defined (expert call out by console operators)

Grid vs. Experiment software

- Even in LHC the experiment requirements in terms of services are quite different
- Two different middleware categories to be considered:
  - Generic middleware
    - Common to all applications
  - Application-oriented middleware
    - Requires a deep knowledge of the experiment models and their Grid use
- The experiment support has to be defined in the same terms
  - Each support member will need to know the computational model of the supported experiment
  - This requires a big effort by the Grid support team and a hand-by-hand work and collaboration
  - In the case of HEP applications and also beyond dedicated support members are assigned to each application

Support strategies

- Grid technologies has been adopted by many research communities as their computational model
  - It covers computational and storage requirements for worldwide distributed communities
  - Ensure a good level of service at all sites
  - Provides the necessary manpower to ensure the required level of service
  - **Strong Points**
    - Excellent way of federating resources in quasi-homogeneous infrastructures
    - Matchmaking computational requirements in many applications
    - Decentralization of resources; small portion available on the lab
    - Worldwide access to data and resources
  - **Weak Points**
    - Reliability and scalability are open issues
    - Difficult to use, lack of user-friendly applications worldwide available

Support Infrastructure

- Experiment Integration Support (EIS) team
  - Born in 2002 with the LHC Computing Grid project: ~10 members, all high energy physicists
- EIS main goals
  - Helping the four LHC experiments (ALICE, ATLAS, CMS, LHCb) in Grid-related issues
  - Help the VO with running the "production": Manage Grid and VO-specific services
  - Creation of documentation (gLite User Guide)
  - Contribute to end user support if required
  - Provide expertise for understanding and solving Grid-related issues
    - Site and middleware level
    - Deep knowledge of both Grid infrastructure and VO computing model
  - Integration and testing of middleware
    - Development of missing components
    - Evaluation of middleware functionality and scalability testing
  - Acting as interface between the experiments and WLCG
  - Not only for HEP: also Biomed and other VOs
- ARDA team
  - Devoted to development for the experiments (and not only!)
    - Monitoring dashboard, GANGA, Diane, AMGA, ...

VO Monitoring

- VOs need to have the full picture of the Grid status…
  - Grid services
  - VO-specific services
- … and to know how they are doing on the Grid
  - Job status, success/failure statistics
  - Status of data transfers
- Exactly at the boundary between the application and the Grid domain
- Examples from the LHC computing:
  - Usage of the Service Availability Monitoring (SAM) for VO-specific service monitoring
  - ARDA dashboard to monitor the experiment workflows

SAM Infrastructure

- One of the major projects of the EIS team
- Adapting to the applications the SAM framework, used by EGEE Grid operations to monitor the site status
  - Define specific tests for different aspects
    - VO software installation, VO services and processes, etc.
  - Produce an adequate GUI, or integrate in an existing one
- What can SAM provide?
  - Calculation of service availability/reliability
  - Alarms
  - Overall status of the Grid from the VO point of view

displays

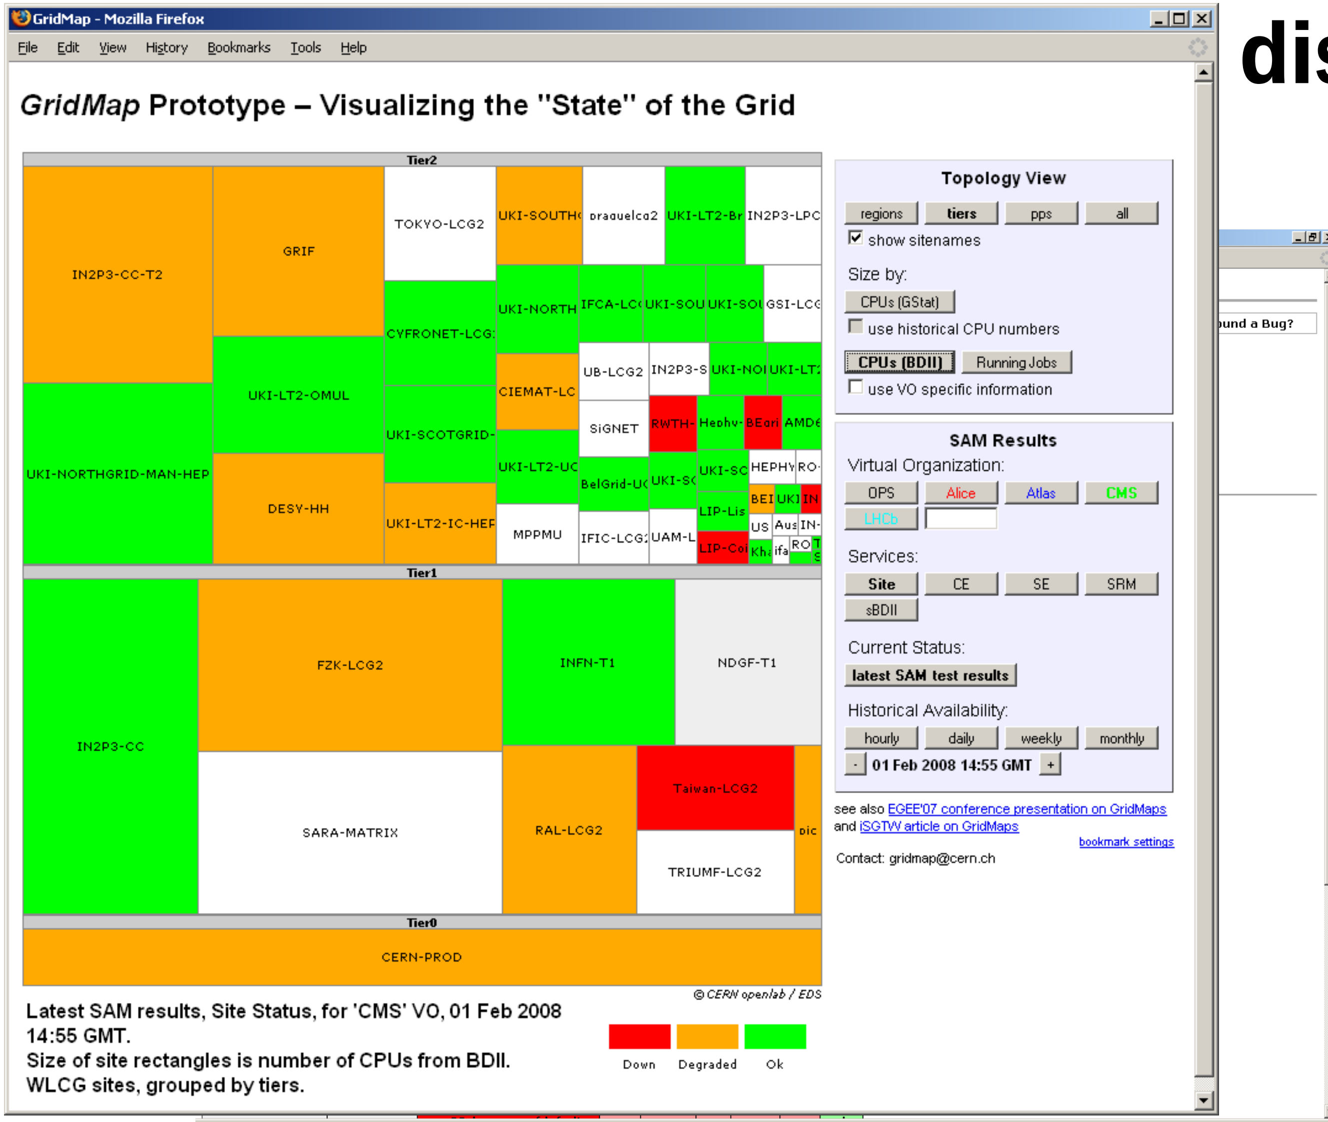
GridMap - Mozilla Firefox
File Edit View History Bookmarks Tools Help
GridMap Prototype – Visualizing the "State" of the Grid
Tier2
IN2P3-CC-T2
GRIF
TOKYO-LCG2
UKI-NORTHGRID-MAN-HEP
UKI-LT2-QMUL
DESY-HH
CYFRONET-LCG2
UKI-SCOTGRID-
UKI-LT2-IC-HEP
CIEMAT-LC
SiGNET
MPPMU
UB-LCG2
IFIC-LCG2
Tier1
IN2P3-CC
FZK-LCG2
INFN-T1
NDGF-T1
SARA-MATRIX
RAL-LCG2
Taiwan-LCG2
TRIUMF-LCG2
Tier0
CERN-PROD
© CERN openlab / EDS
Latest SAM results, Site Status, for 'CMS' VO, 01 Feb 2008 14:55 GMT.
Size of site rectangles is number of CPUs from BDII.
WLCG sites, grouped by tiers.
Down
Degraded
Ok
Topology View
regions
tiers
pps
all
show sitenames
Size by:
CPUs (GStat)
use historical CPU numbers
CPUs (BDII)
Running Jobs
use VO specific information
SAM Results
Virtual Organization:
OPS
Alice
Atlas
CMS
LHCb
Services:
Site
CE
SE
SRM
sBDII
Current Status:
latest SAM test results
Historical Availability:
hourly
daily
weekly
monthly
01 Feb 2008 14:55 GMT
see also EGEE'07 conference presentation on GridMaps
and iSGTW article on GridMaps
bookmark settings
Contact: gridmap@cern.ch

Dashboard

JOB SUMMARY

You found a bug? You have a suggestion?

jobsummary | waitingtime | runningtime

any user
any site
any ce
any submissiontool
any dataset
any application
any rb
any activity
any grid
☐ unk ☐ pend ☐ run ☐ term
☐ done ☐ canc ☐ abort ☐ g-unk
☐ succ ☐ fail ☐ a-unk
☐ donesuccess
2008-01-01 17:23:50 to
2008-02-01 17:23:50
☐ all jobs regardless submission time
sort by site
bars in the plot
submit

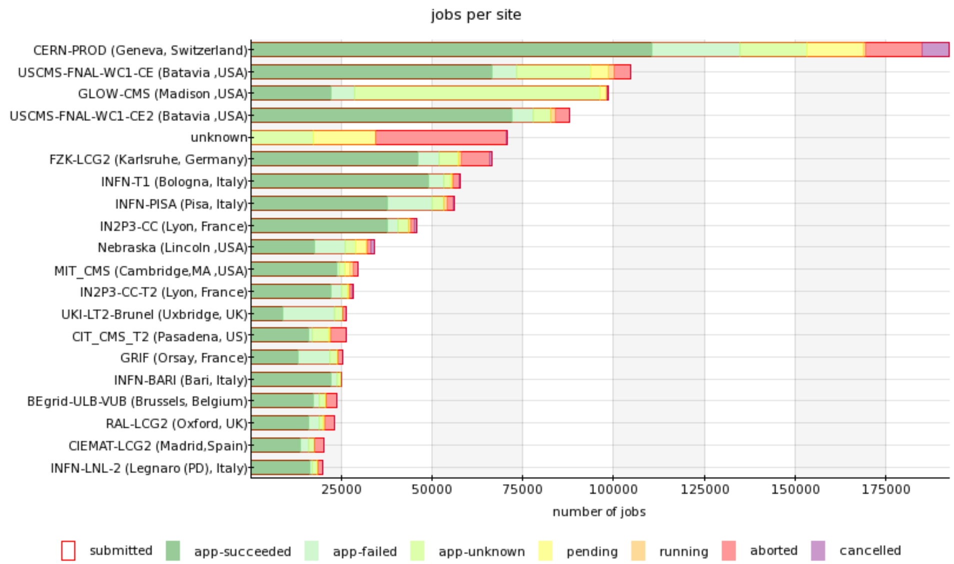

	dashboard					grid					application				overall	
site↓	Sub	Unk	Pend	Run	Term	Done	Canc	Abort	Unk	Grid%	Succ	Fail	Unk	App%	D/S	Overall%
AMD64.PSNC.PL (Poznan, Poland)	1	0	0	0	1	1	0	0	0	100	0	0	1	0	0	0
BEIJING-LCG2 (Beijing, China)	7934	0	153	28	7753	5132	1	2231	389	69.7	4268	963	2522	81.59	4127	53.23
BEgrid-ULB-VUB (Brussels, Belgium)	23756	0	256	14	23486	17736	70	2875	2805	86.05	17404	1658	4424	91.3	15467	65.86
BG-INRNE (Sofia, Bulgaria)	1	0	0	0	1	1	0	0	0	100	0	0	1	0	0	0
BG01-IPP (Sofia, Bulgaria)	1	0	0	0	1	0	0	1	0	0	0	0	1	0	0	0
BUDAPEST (Budapest, Hungary)	7278	0	4	19	7255	5972	1	794	488	88.26	6017	321	917	94.94	5643	77.78
BelGrid-UCL (Louvain-la-Neuve, Belgium)	12618	0	0	14	12604	9719	0	2035	850	82.69	9666	618	2320	93.99	8910	70.69
CERN-PROD (Geneva, Switzerland)	192649	0	15493	602	176554	118155	7730	15539	35130	88.38	110494	24520	41540	81.84	90655	51.35
CESGA-EGEE (Santiago de Compostela, Spain)	4	0	0	0	4	4	0	0	0	100	0	1	3	0	0	0
CGG-LCG2 (Massy, France)	1	0	0	0	1	1	0	0	0	100	0	1	0	0	0	0
CIEMAT-LCG2 (Madrid,Spain)	20371	0	49	48	20274	15065	108	2567	2534	85.44	13812	2246	4216	86.01	12778	63.03
CIT_CMS_T2 (Pasadena, US)	26281	0	180	376	25725	13268	7	4169	8281	76.09	16162	778	8785	95.41	11679	45.4
CSCS-LCG2 (Manno, Switzerland)	10368	0	0	0	10368	7911	48	1539	870	83.71	4881	167	5320	96.69	4353	41.98
CYFRONET-IA64 (Cracow, Poland)	1	0	0	0	1	1	0	0	0	100	0	0	1	0	0	0
total	*13752978*		*49671*	*11185*	*1314433*	*807726*	*13169*	*114845*	*378693*	*87.55*	*823700*	*162523*	*328210*	*83.52*	*622768*	*47.38*

- A co… …s for t…
- Diffe…
  - Gl…
  - St…
  - St…
  - …

Data management monitoring for ATLAS

Accounting information from Apel and Gratia for ATLAS (prototype)

…ecific

Summary

- WLCG has been born to cover the high computer and storage demand of the experiments from 2008
- We count today more than 200 sites in 34 countries
- The whole project is the results of a big collaboration between experiments, developers and sites
- It is a reality, GRID is being used in production
- Covering a wide range of research fields, from HEP applications until humanitarian actions

Patricia Méndez Lorenzo patricia.mendez@cern.ch