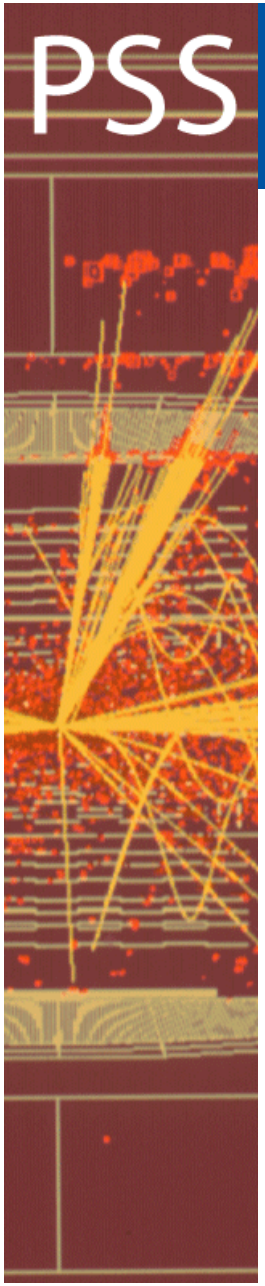

PSS Physics Services Support

CERN IT Department

Grids, middleware, tools and experiment activities at CHEP07

Andrea Sciabà

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

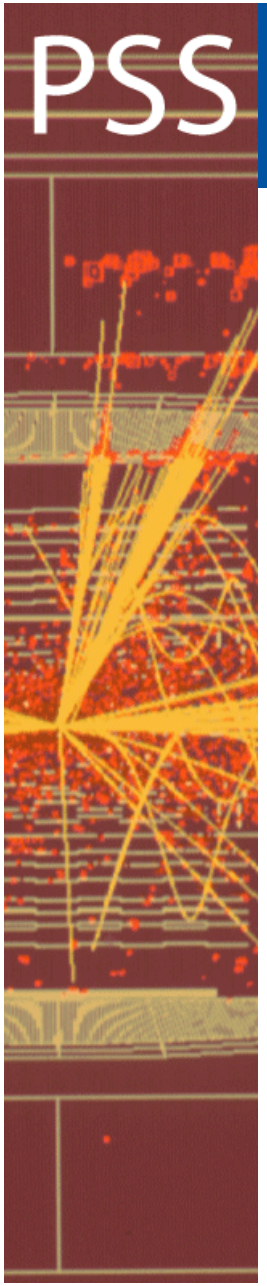

Outline

- Introduction
- Production Grids
- Workload Management Systems
- Data Management
- Monitoring
- Experiment Tools
- Experiment Challenges and Activities
- Conclusions

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Introduction

- Several sessions at CHEP07 were centered on Grid computing and its applications
  - Computer Facilities, Production Grids and Networking
  - Grid Middleware and Tools
  - Distributed Data Analysis
- This presentation tries to summarize a personal selection of the reports on some of the areas covered
  - The main focus will be though on the usage of the Grid by the LHC experiments
  - Many topics have been left out…

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

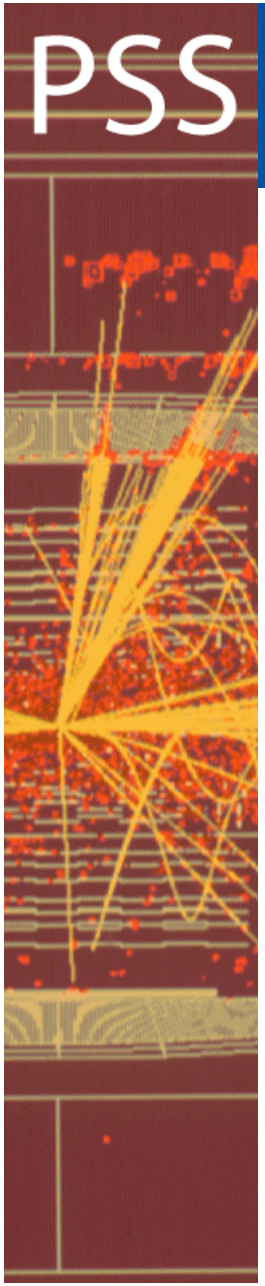

Production Grids

- OSG (R. Pordes)
- GridPP (J. Coles)
- Nordic Data Grid Facility (NDGF) (L. Fischer)

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

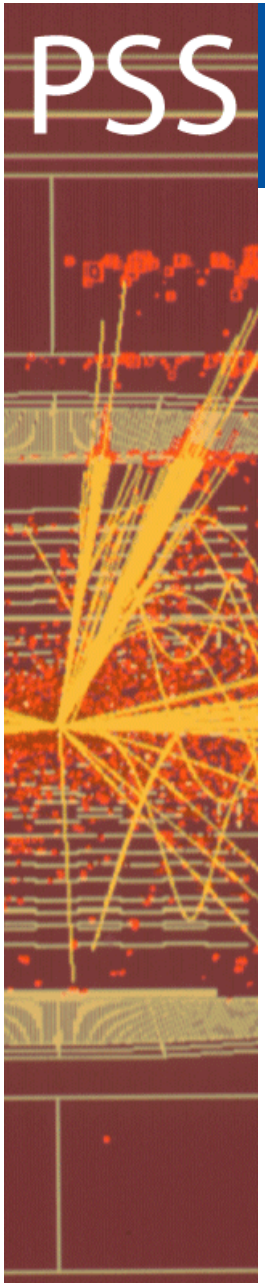

OSG status and plans

- Supports ATLAS and CMS
- But also CDF, D0, neutrino physics, STAR, LIGO, etc.
- Operates the infrastructure and supports site administrators and users
- 75 processing resources in US, 4 in Brazil, 1 in Mexico, UK, Taiwan and South Korea
  - ~30,000 CPU cores
  - 75,000 jobs/day
  - 240,000 CPU wallclock hours/day
  - CPU time/wallclock time ~0.75
  - 20%-30% usage from non-owners!
- 15 storage resources
  - ~4 PB of disk
  - 100 TB of data moved per day
  - 1 PB/day next year
- ~2x increase in accessible resources by end of 2009

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

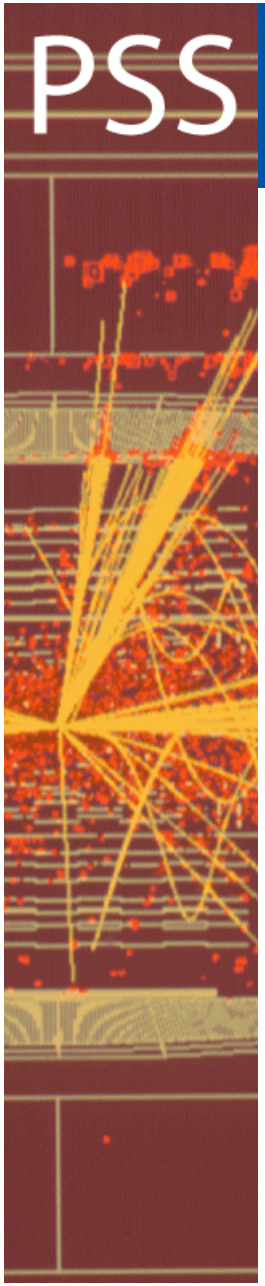

OSG middleware

- The middleware stack is the Virtual Data Toolkit (VDT)
- Interfaces
  - Job execution
    - Condor-G
      - Submitting to WS, pre-WS GRAM, Condor glide-in
  - Storage
    - SRM
      - Full support to SRM 2.2 from next February
  - Data transfer
    - GridFTP
  - VO management
    - VOMRS, VOMS shared with EGEE
- Achieving interoperability with WLCG
  - Information system publishes to the WLCG BDII
  - OSG CEs can accept jobs from gLite WMS

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Nordic Data Grid Facility

- The Tier-1 is distributed over the 7 largest

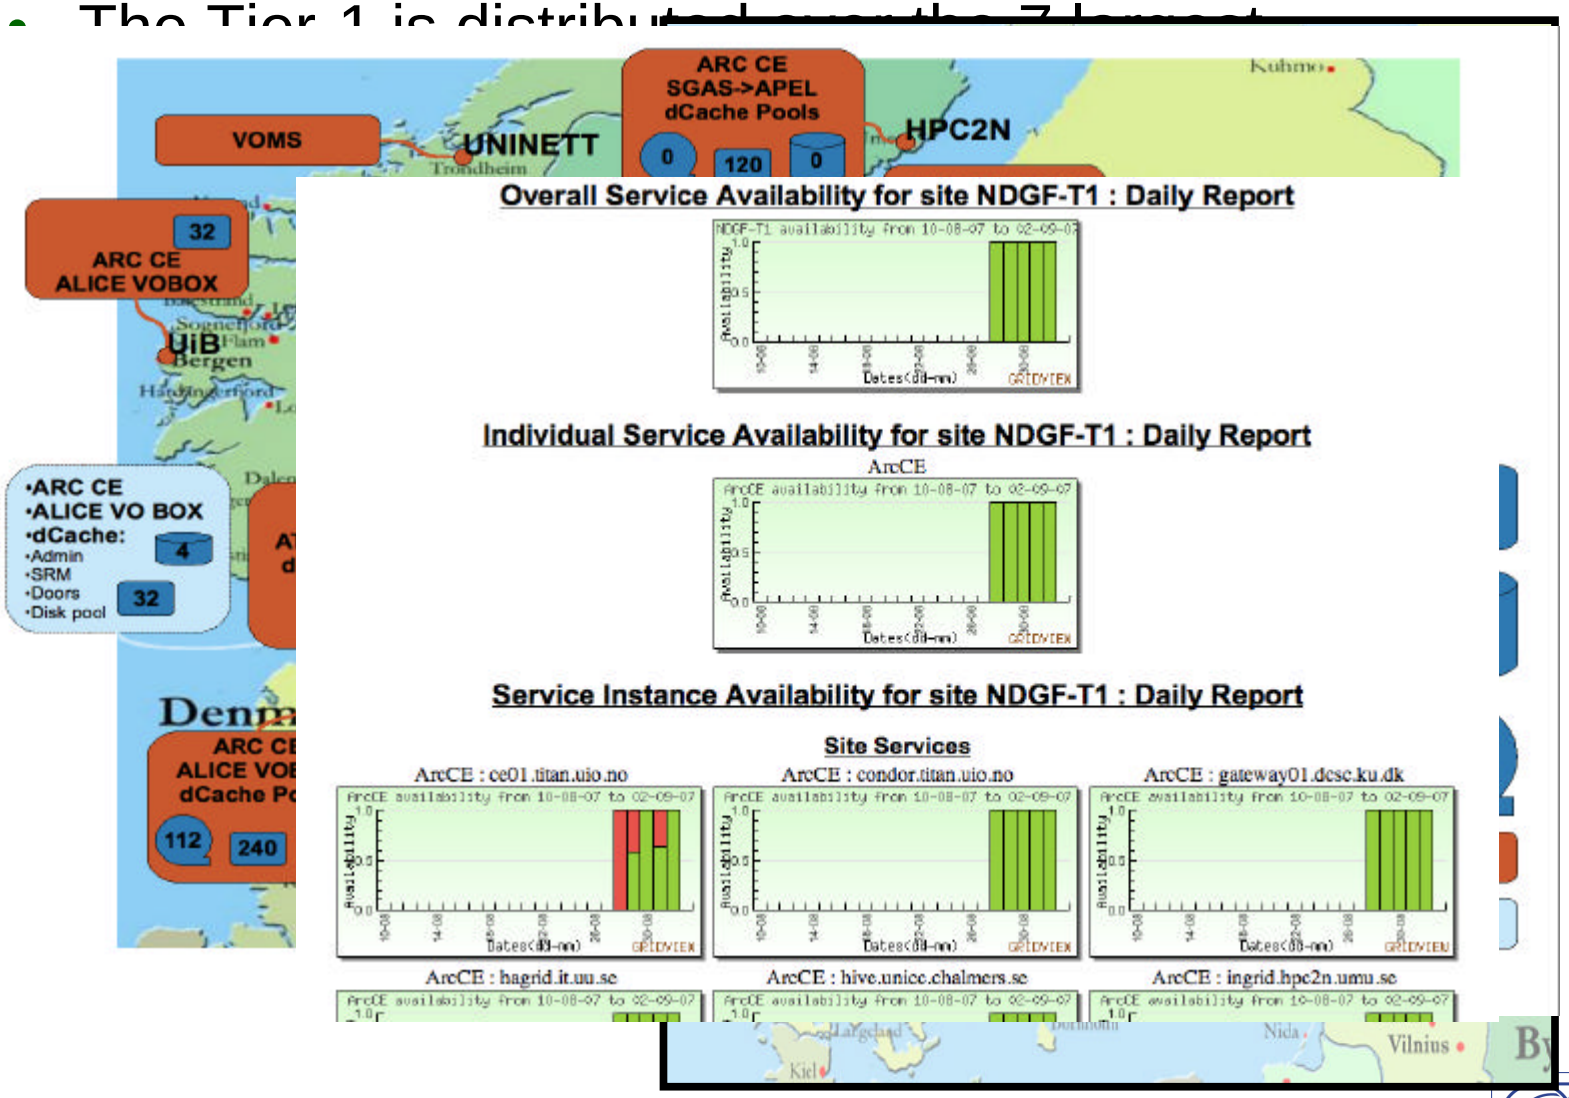

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

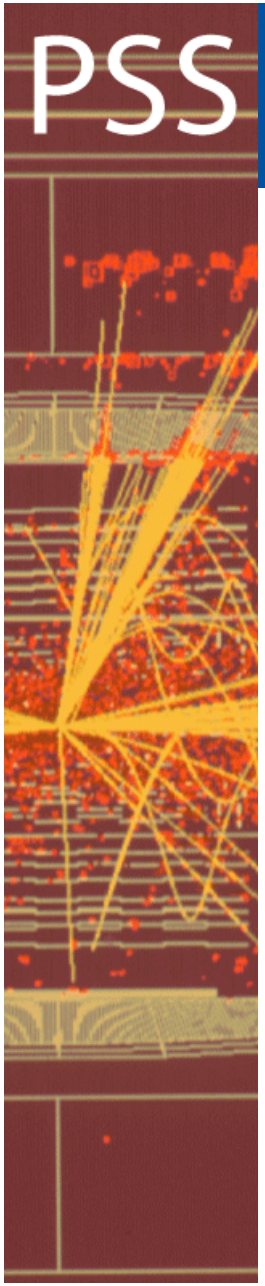

Workload Management

- A plethora of WM systems exist and are used
  - "Push“ models
    - **gLite WMS** (M. Cecchi, S. Campana)
    - Condor-G
  - “Pull”, or “Pilot” models
    - A job sits on a job slot and pulls "real" jobs from a central task queue
    - "standard" pilots
      - DIRAC (S. Paterson)
      - **PaNDA** (Y. Smirnov, P. Nilsson, T. Maeno)
      - AliEN (P. Saiz)
    - Condor Glide-in
      - **glideinWMS** (I. Sfiligoi)
      - **CRONUS** (S. Padhi)
  - Other Grids
    - Xgrid from Apple (A. Kocoloski)
    - Sun Grid (M. Potekhin)
- Several talks spread over different sessions

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

gLite WMS

- Main functions
  - Match user jobs to available resources
  - Submit the job to the resource through the appropriate interface
    - Condor-G for the LCG CE
    - ICE for CREAM
  - "Manage" the job on behalf of the user
    - Log events to the Logging & Bookkeeping service
    - Automatically resubmit failed jobs
- Several functionalities make it better than the old LCG Resource Broker
  - **Bulk submission and bulk matchmaking**
  - Shared input sandbox among similar jobs
  - Local cache of the Information System ("Information Supermarket")
  - Task queue to hold submission requests for which resources are not available
  - Load limiter
  - A new Web Service interface ("WMProxy")

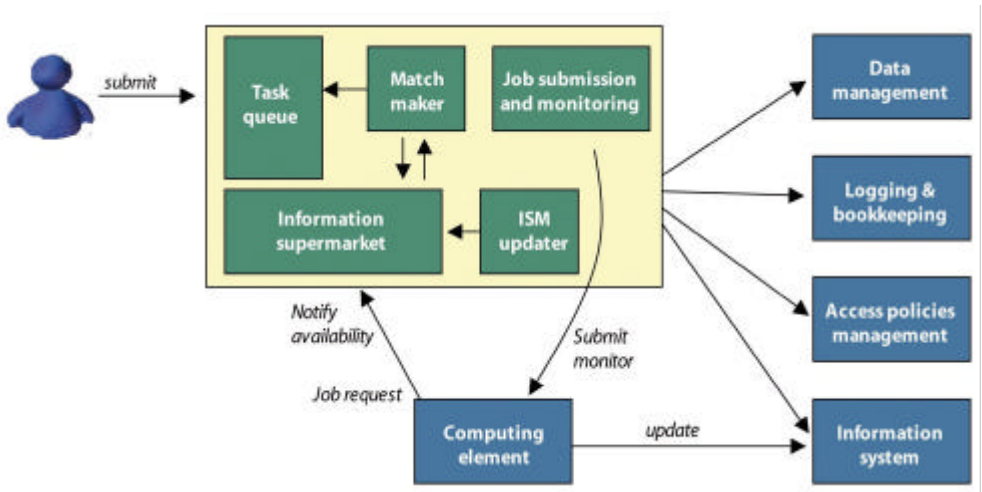

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

gLite WMS performance

- After a very intense testing and bug fixing, the WMS demonstrated performance levels compatible with the experiment needs
  - ~16K jobs/day over one week
  - Stable behaviour and performance
  - Negligible fraction (0.3%) of jobs with problems due to the WMS
- Used in some way or another by all the LHC experiments
  - ATLAS, CMS to submit real jobs
  - ALICE, LHCb to submit pilot jobs

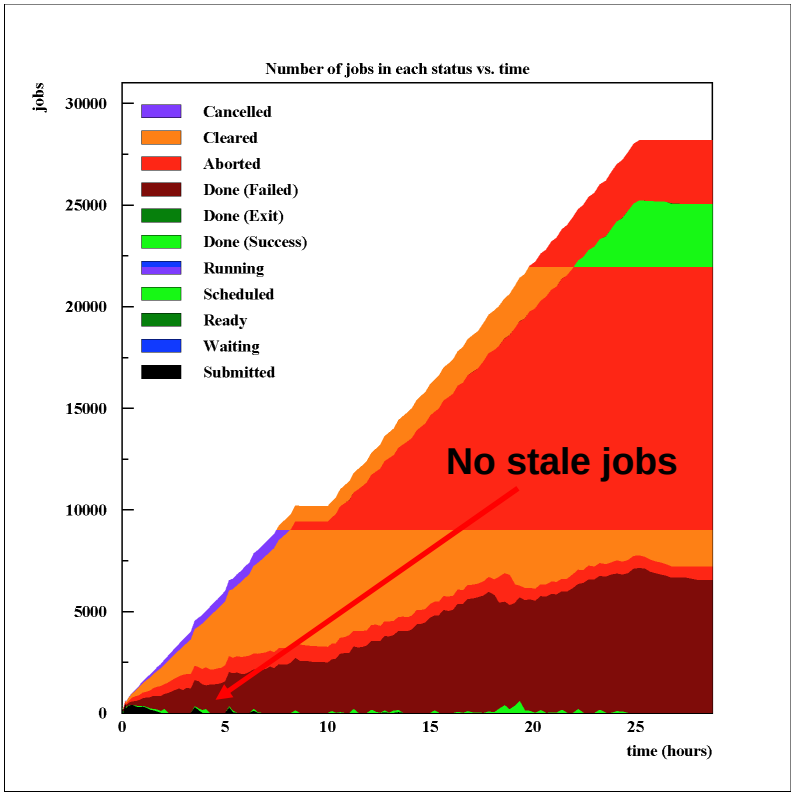

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

The gLite WMS in the ATLAS MC production

- The ATLAS production system submits part of the work to EGEE using the gLite WMS
  - Extremely painful at the beginning
- Good performances after the recent WMS improvements
  - Up to 20K jobs/day
  - Job efficiency ~60%
    - 13% due to the “WMS” (including problems with the CEs)
  - **Negligible** wallclock time inefficiency
    - Most jobs fail at the beginning and are resubmitted

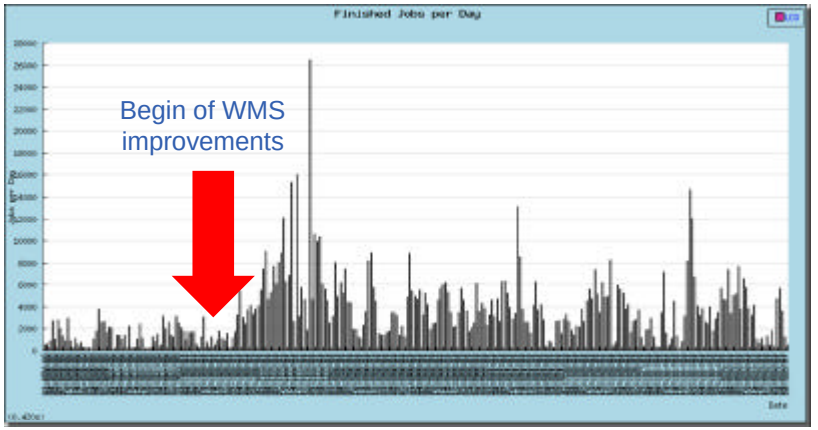

Job errors (7/8 2007)

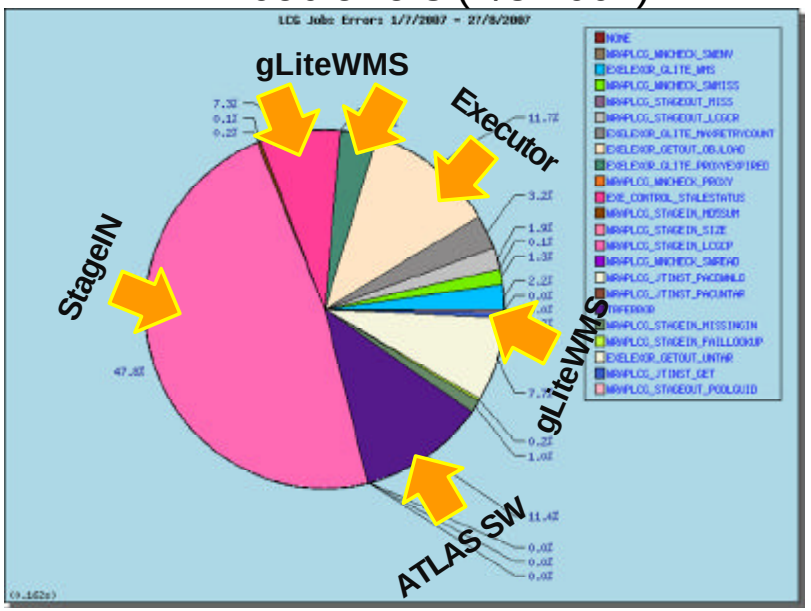

gLite WMS: ~13% Data Management: 47% ATLAS SW: 11%

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

The WMS in CMS data analysis

- CMS supports submission of analysis jobs via WMS
  - For CSA07 the goal is to submit at least 50K jobs/day via WMS
  - The Job Robot (a load generator simulating analysis jobs) is successfully submitting up to 20K jobs/day

Submission rate

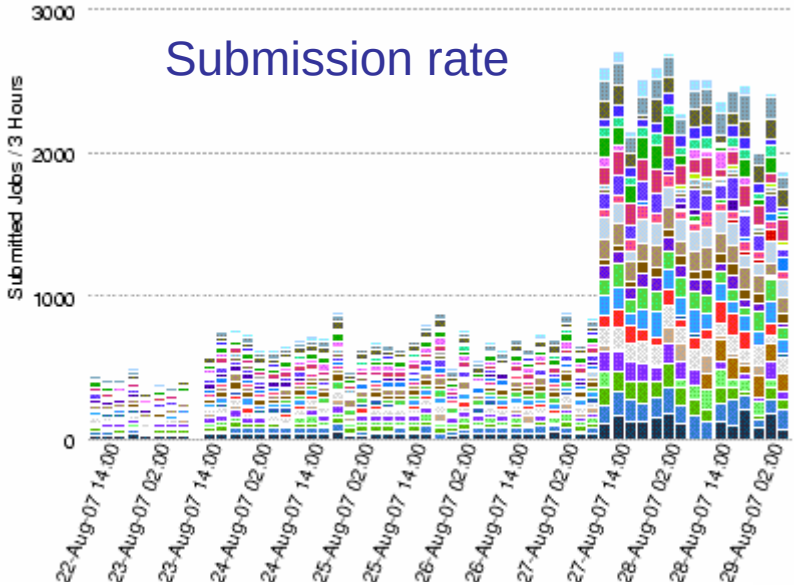

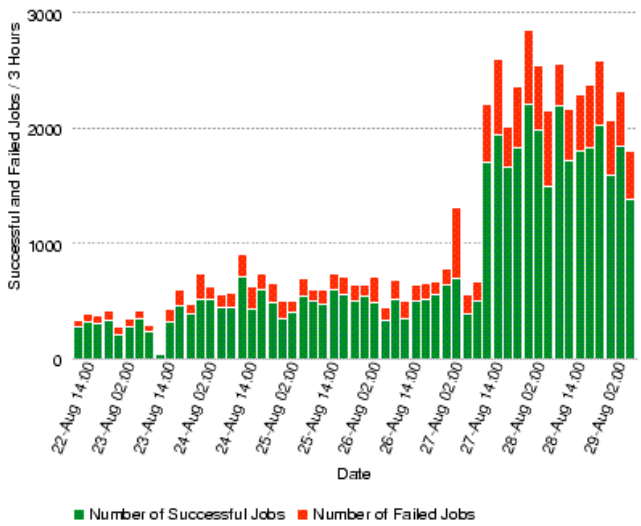

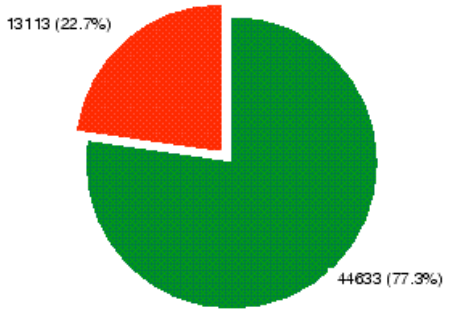

Success rate

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PANDA: Production ANd Distributed Analysis system

- Used by ATLAS in OSG for production and analysis
  - Submits to Condor-G, PBS and Condor
- Central server Apache-based
  - Priority control
  - Resource allocation
  - Job scheduling
- For production jobs it takes care to send data transfer requests to DQ2 when needed
  - Data transfers managed asynchronously
- Performances
  - ~30M simulated events, >30% of total
  - Currently ~15K jobs/day, scales up to at least 40K jobs/day
  - Panda errors < 3% of total
  - 85% job efficiency
  - 91% WCT efficiency
- Plans
  - Extend support to EGEE
  - Add glideins (see next slide)

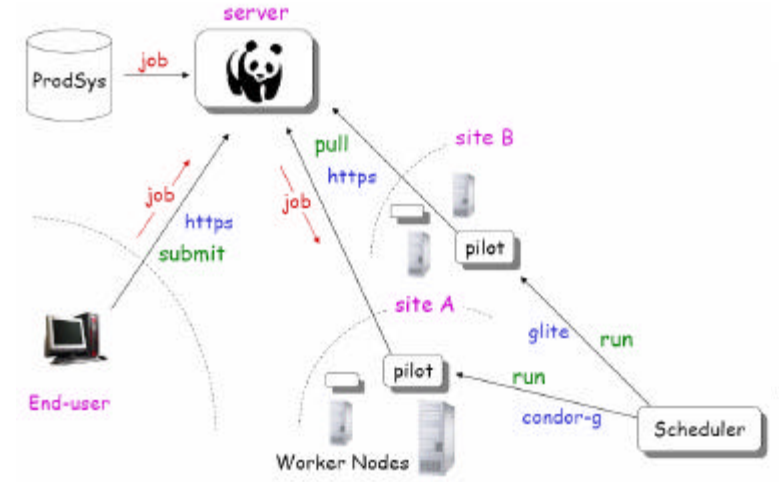

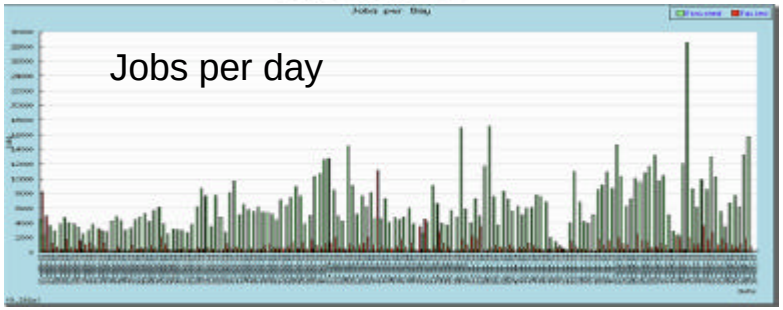

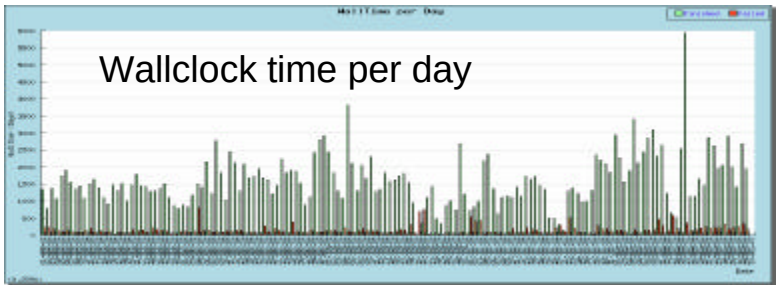

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Condor Glideins

- Based on the Condor architecture
- A glidein is a starter submitted as a Grid job that runs on a worker node
  - A virtual private Condor pool running on WNs!
- Glideins can also take care of
  - Sanity checks
  - Discovery and publish batch slot characteristics
  - Preparation of the job environment

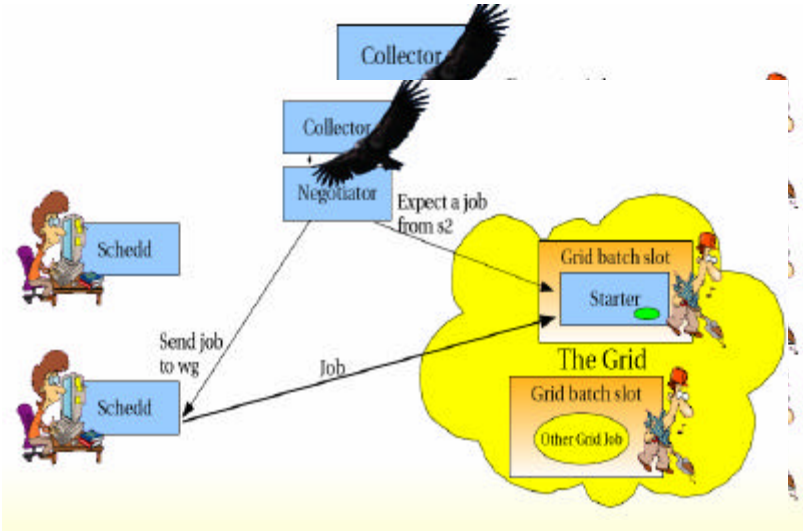

- Advantages
  - Easy to use for Condor users
  - Hides the Grid intricacies
  - Protects from obvious errors
- Disadvantages
  - Starters need an open connection to a central Condor proxy server (the GCB)
  - The user and the starter run under the same local account
  - Resource hungry (1.5 MB/running process on submit node)
  - Some components still immature (the GCB)
  - Debugging can be challenging!

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

glideinWMS

- A WMS built on top of glideins by CMS
  - A VO frontend matches jobs in the queue to sites
  - A glidein factory submits the needed glideins on the matched site
  - The rest is the same and the Condor negotiator has the final word
- Advantages
  - Keeps a steady number of idle glideins in site as long as there are waiting jobs
- Do glideins scale?
  - Not very well
  - Up to 1K jobs in parallel for CMS
- And glideinWMS?
  - No problem there

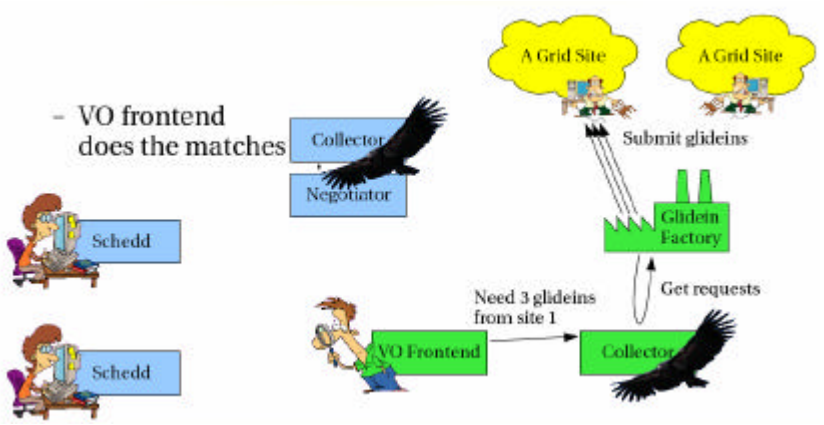

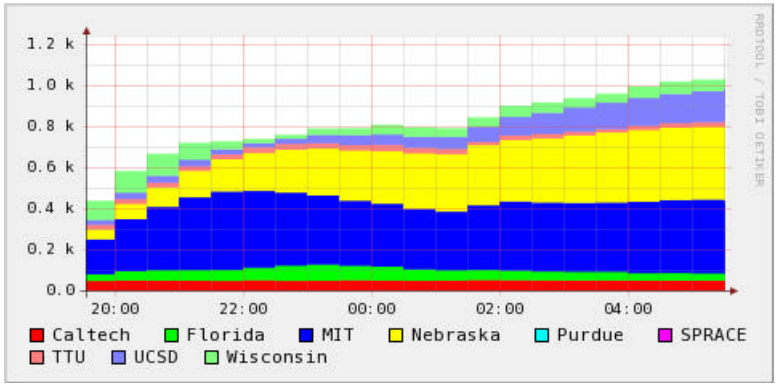

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CRONUS: ATLAS MC production with glideins

- Pilots submitted to OSG (via Condor-G) and to EGEE (via gLite WMS)
  - NDGF recently added
- The resulting virtual cluster is uniform across Grids
- Performances
  - ~645K jobs in 2007
  - WCT efficiency ~87%
  - Most errors related to data I/O and ATLAS software
  - Scaled up to ~5200 concurrent jobs
- Major problems
  - Loss of connection between starter and submit host

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

gLExec

- Problem: pilot jobs run their payload (= user jobs) with the pilot's Grid credentials
  - Impossible to trace things back to the real job owner
  - Impossible to ban single users from sites
  - Users can steal or corrupt each other's data
- Solution: gLExec allows to switch UID for the running job given the real user's credentials
  - It will also keep a log of all the authorization requests
- Advantages
  - Solves the problem!
- Disadvantages
  - The VO must do the accounting
  - Policies are on local files

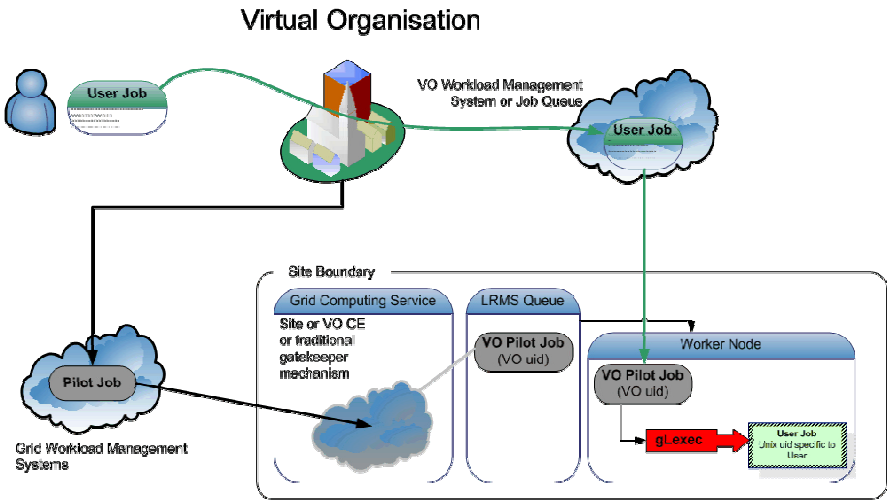

- Deployed and tested at FNAL with good results so far

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

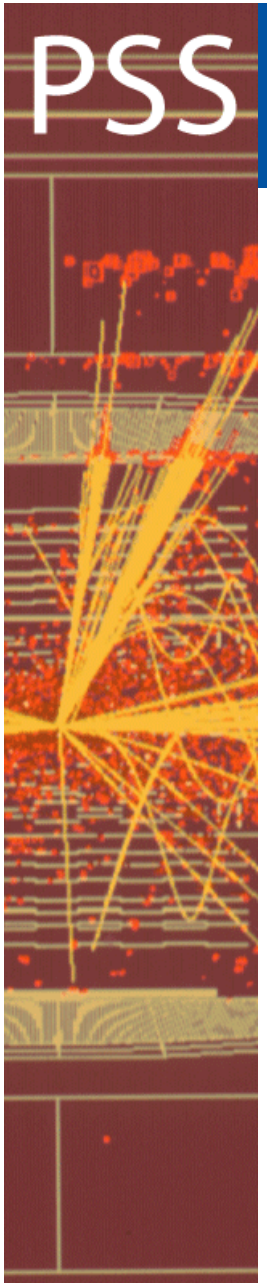

Other pilot job systems

- DIRAC
  - Used by LHCb in EGEE for production and analysis
  - Pilots are submitted via LCG RB or gLite WMS
  - Up to 10K concurrent jobs
  - Uses standard batch system components (e.g. Maui) for prioritization of jobs in the Task Queue
- AliEn
  - full middleware stack developed in ALICE

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Data Management: the middleware

- Interfaces
  - SRM 2.2 (F. Donno)
- Storage solutions
  - dCache (T. Perelmutov, P. Fuhrmann, G. Berhmann)
  - DPM (M. Schulz)
  - CASTOR2 (G. Lo Presti, J. Van Eldik)
- File catalogues
  - LFC (M. Schulz)
- Data transfer tools
  - FTS (M. Schulz)
- User clients
  - GFAL, lcg_util (M. Schulz)

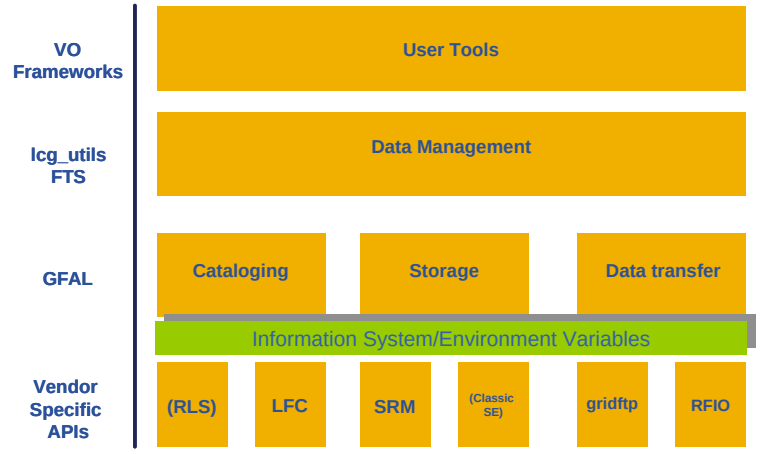

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

SRM 2.2 testing

- SRM is a standard interface for space allocation and file management on the Grid
  - CASTOR2, DPM, dCache, StoRM, BeStMan
- Tests done with the S2 suite developed at RAL, to ensure
  - Validation of use cases
  - Consistency
  - Compliance with the specification
  - Performance and reliability
  - 52 tests run 4 times/day on 21 endpoints with different configurations

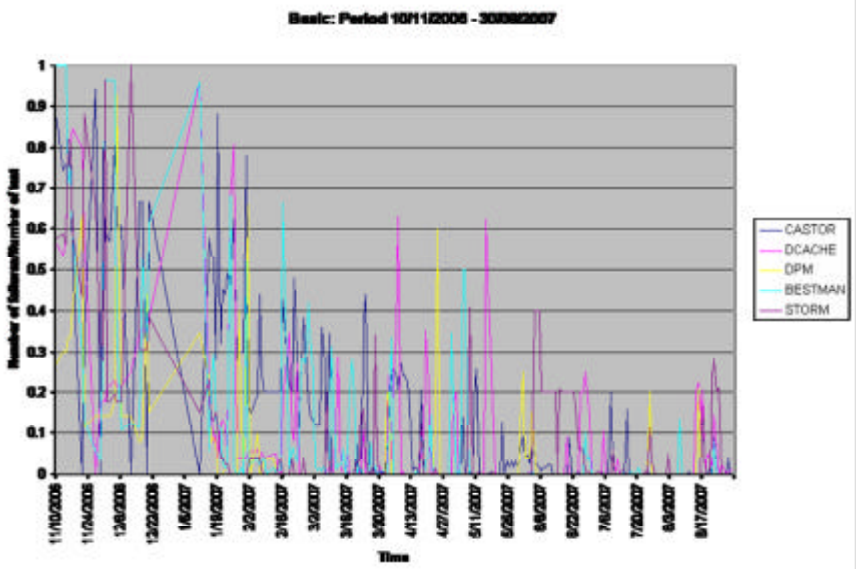

- Other groups testing also
  - Higher level tools
    - FTS, lcg_util, etc.
  - Transparent access (v1 vs. v2)
  - handling of VOMS groups
  - VO data access patterns
  - …

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

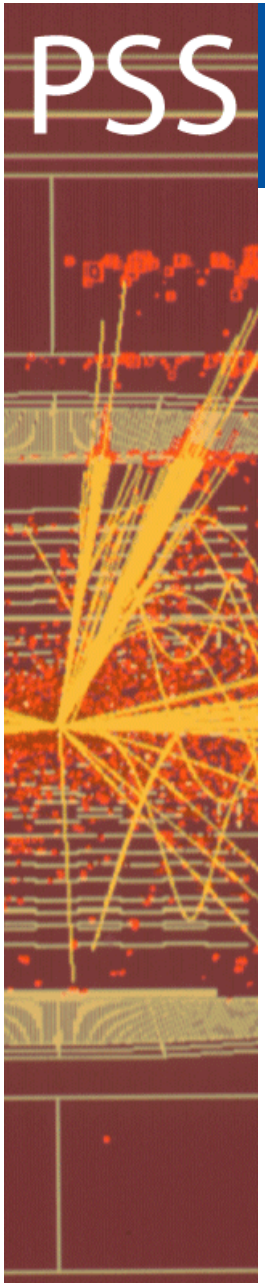

Status of the Storage Elements

CERN IT Department

- dCache
  - Used at 7 Tier-1's and several Tier-2's
  - Supports GridFTP, (gsi)dCap, xRoot, SRM, nfs2/3
  - Working on Chimera, a new namespace system to replace PNFS
    - >2 GB files, ACLs, user defined queries, fast
  - Working also on NFS 4.1
    - Faster, GSS authentication, ACLs
    - Clients come with all major OS vendors
- DPM
  - Disk-only SE for Tier-2's
    - 96 in production
  - Supports GridFTP, secure RFIO, xRoot (limited), http(s)
  - Full support of VOMS for authorization (file permissions, ACLs)
  - Working on file encryption for medical applications

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN

LFC, lcg_util

- LCG File Catalog
  - Stores user file names, location, owner and permissions
    - Hierarchical namespace
  - Supports GSI security and VOMS based ACL
    - LFC is used as internal catalogue by DPM
  - The effort is mainly on optimization
    - Bulk methods can reduce round trips ⇒ up to 20x faster
- lcg_util
  - C API + CLI to cover most common use cases
    - Replication, catalogue interaction, etc.
  - Work ongoing to make them work independently from LFC
    - The goal is to have a unique and official "SRM client" for users
- GFAL
  - For POSIX-like file access
  - Supports now SRM 2.2

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

File Transfer System

- The tool to transfer
endpoints
  - Handles retries in
    - Failures of the
  - Channels are "qu
    - Inter-VO and in
  - Different roles
    - User, VO prod
- Deployed at Tier-0
  - Over 9 petabytes
- FTS 2.2 improveme
  - Better security mo
  - Improved monitor
  - Beta SRM 2.2 sup
  - Database optimiz

FTS Report

Disclaimer
This page contains a report generated from information stored in the FTS Database and is intended for reporting purposes only. Since the format will probably change in the future, it's therefore recommended not to use parsing robots on it.

Statistics concerning all the transfers performed yesterday
Between 2006-10-12 08:00:00 +02:00 and 2006-10-13 08:00:00 +02:00

CERN-* Filter Show VO details

Channel Name	VO Name	Total	% Failures	# Succ.	# Fail.	1st Failure Reason	% 1st Failure Reason	2nd Failure Reason	% 2nd Failure Reason	Avg. Size (GB)	Avg. Duration (sec)	Avg. Tx Rate (MB/sec)	Eff. Tx Bytes (GB)	Tx Bytes (GB)
CERN-PIC	[All]	12282	73.97	3192	9070	Dest SRM	58.22	Other	37.53	0.63	283.03	1.62	1700.41	1700.41
	atlas	8932	99.92	7	8925	Dest SRM	57.13	Other	38.09	0	220	0	0	0
	cms	208	0	208	0					2.7	767.55	3.64	561.26	561.26
	dteam	974	0.51	969	5	Other	80	Source SRM	20	0.95	356.31	2.88	923.83	923.83
	lhcb	2115	6.53	2005	140	Source SRM	99.29	Other	0.71	0.11	165.85	0.81	215.32	215.32
	ops	3	0	3	0					0	202.67	0	0	0
CERN-RAL	[All]	8688	58.26	3544	5195	Other	84.91	Source SRM	14.88	0.85	478.22	2.59	3026.81	3027.57
	alice	1155	82.6	201	954	Other	99.58	Dest SRM	0.31	1.86	1805.05	1.11	372.95	372.95
	atlas	4512	88.52	518	3994	Other	84.85	Source SRM	15.15	1.79	1428.04	1.49	926.26	926.57
	cms	227	3.08	220	7	Dest SRM	85.71	Source SRM	14.29	2.53	348.65	10.08	555.61	555.61
	dteam	1077	3.99	1034	43	Other	86.05	Source SRM	9.3	0.95	276.64	4.01	980.47	980.91
	lhcb	1725	9.1	1568	157	Source SRM	99.36	Other	0.64	0.12	146.03	1.16	191.52	191.52
	ops	3	0	3	0					0	27	0.01	0	0
CERN-BARA	[All]	9792	42.55	5051	3741	Dest SRM	93.77	Source SRM	12.22	1.34	108.02	15.4	6777.95	6784.92
	alice	3134	15.12	2660	474	Source SRM	57.17	Dest SRM	41.14	1.66	169.53	18.43	4426.44	4430.29
	atlas	2018	53.32	942	1076	Dest SRM	72.4	Source SRM	16.54	1.15	144.44	9.42	1085.07	1087.6
	dteam	3488	61.32	1349	2139	Dest SRM	98.74	Other	0.98	0.93	91.91	14.66	1260.74	1261.32
	lhcb	148	35.14	96	52	Dest SRM	92.31	Other	3.85	0.06	76.1	0.93	5.7	5.7
	ops	4	0	4	0					0	97.25	0.02	0	0
CERN-INFN	[All]	11492	42.31	6630	4862	Dest SRM	43.95	Other	37.7	1.13	395.77	3.21	7514.29	7514.94
CERN-CERN	[All]	1536	39.71	926	610	Source SRM	58.36	Dest SRM	15.9	0.07	287.71	0.38	67.89	68.08
CERN-ASCC	[All]	6851	23.54	5236	1613	Source SRM	50.84	Other	28.89	1.14	1086.6	1.08	5855.81	6080.58
CERN-GRIDKA	[All]	12755	21.36	10028	2727	Source SRM	64.36	Other	32.53	0.67	371.97	3.19	6762.02	6767.53
CERN-TRIUMF	[All]	2244	20.63	1781	463	Other	61.77	Source SRM	31.1	1.04	395.15	3.63	1847.25	1917.13
CERN-BNL	[All]	13975	19.42	11261	2714	Source SRM	69.97	Other	24.24	0.44	190.38	3.41	4851.58	4860.34
CERN-IN2P3	[All]	11697	13.76	10087	1610	Source SRM	48.57	Other	47.45	1.22	286.21	5.33	12328.63	12328.63
CERN-FNAL	[All]	917	4.58	875	42	Transfer	97.62	Other	2.38	0	379.88	0	0	0

Click on the Channel Name to show the VO details

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Data management: the experiments

- PhEDEx (L. Tuura)
  - The data movement system of CMS, very robust and performant by design
    - Manages not single files but blocks (5-10 TB)
    - Protects the system from failures
    - All communications among agents through a central DB, very strongly optimized
    - Thorough validation procedure
    - Routinely transfers ~100 TB/day, 50K files/day
    - Focus is now on just improving the interface

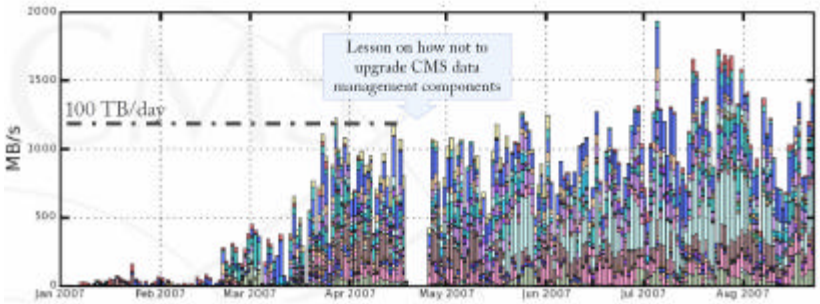

PhEDEx aggregate data transfers

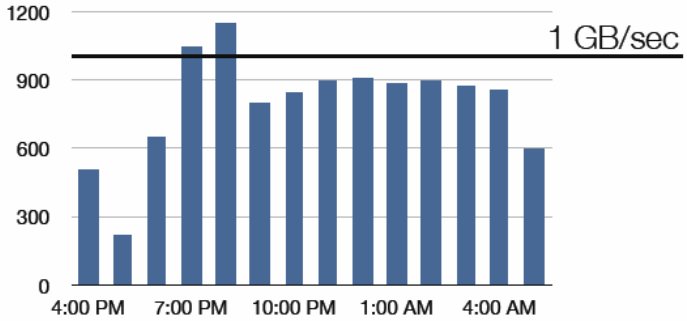

ATLAS Tier-0 export

- ATLAS DDM (M. Lassnig)
  - Takes care of data movement and data bookkeeping
  - Based on a central catalogue and some site services
  - Some recent improvements
    - Improved transfer reliability (failover strategies, etc.)
    - Improvement of fairshare
    - New database schema
    - More efficient SQL queries
  - **Almost ready for data taking**

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Monitoring

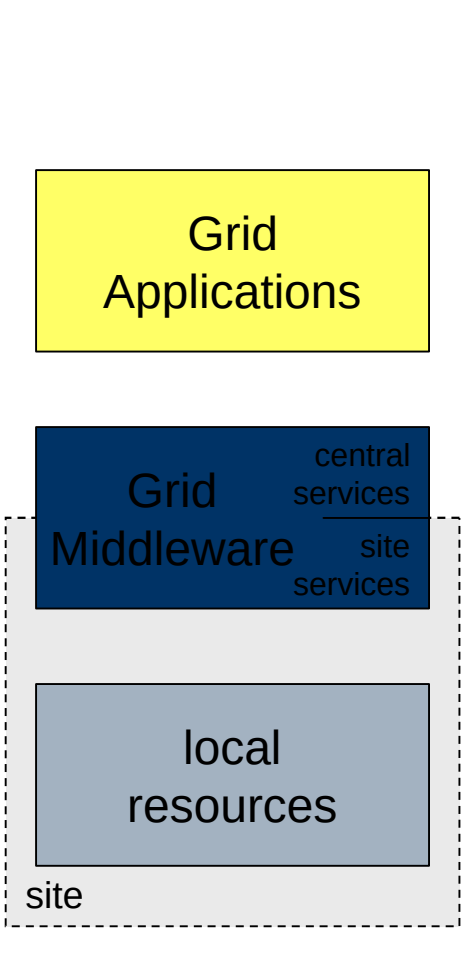

Domain	Monitoring Tools in use
Application monitoring	Experiment Dashboards ...
Grid Services monitoring	Gstat SAM/GridView GridICE GridPP Real Time Monitor ...
Local monitoring	Lemon/SLS Nagios Ganglia...

3 WLCG Monitoring Working Groups

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Grid service monitoring
(J. Casey)

- Provide monitoring views suited for
  - Site admins
  - Grid operators
  - VO management
  - Grid project management
- Allow the site to see information collected also by remote probes

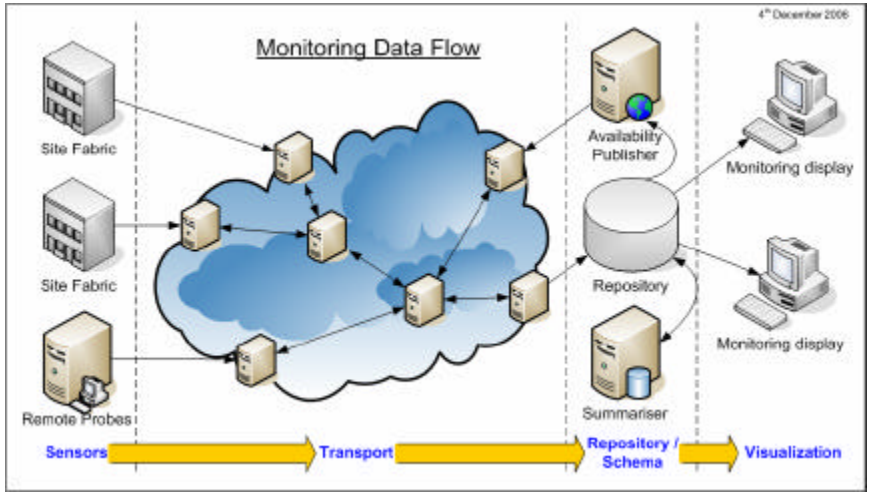

- Standardize the monitoring of Grid services
  - Define "standard" probes
  - Define a standard Data Exchange format for publishing and consuming
- Prototype based on Nagios
  - Can be already be tried out!

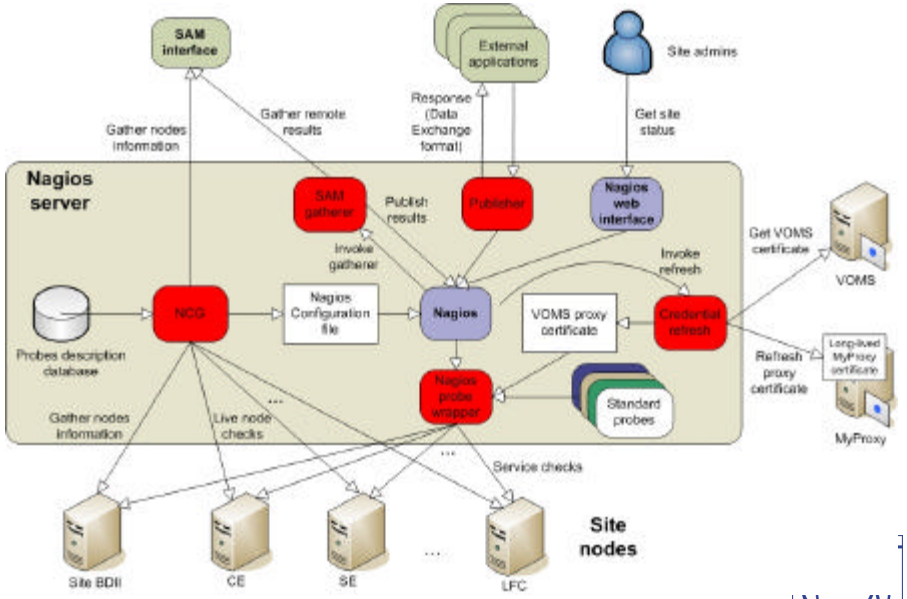

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Application monitoring
(J. Andreeva)

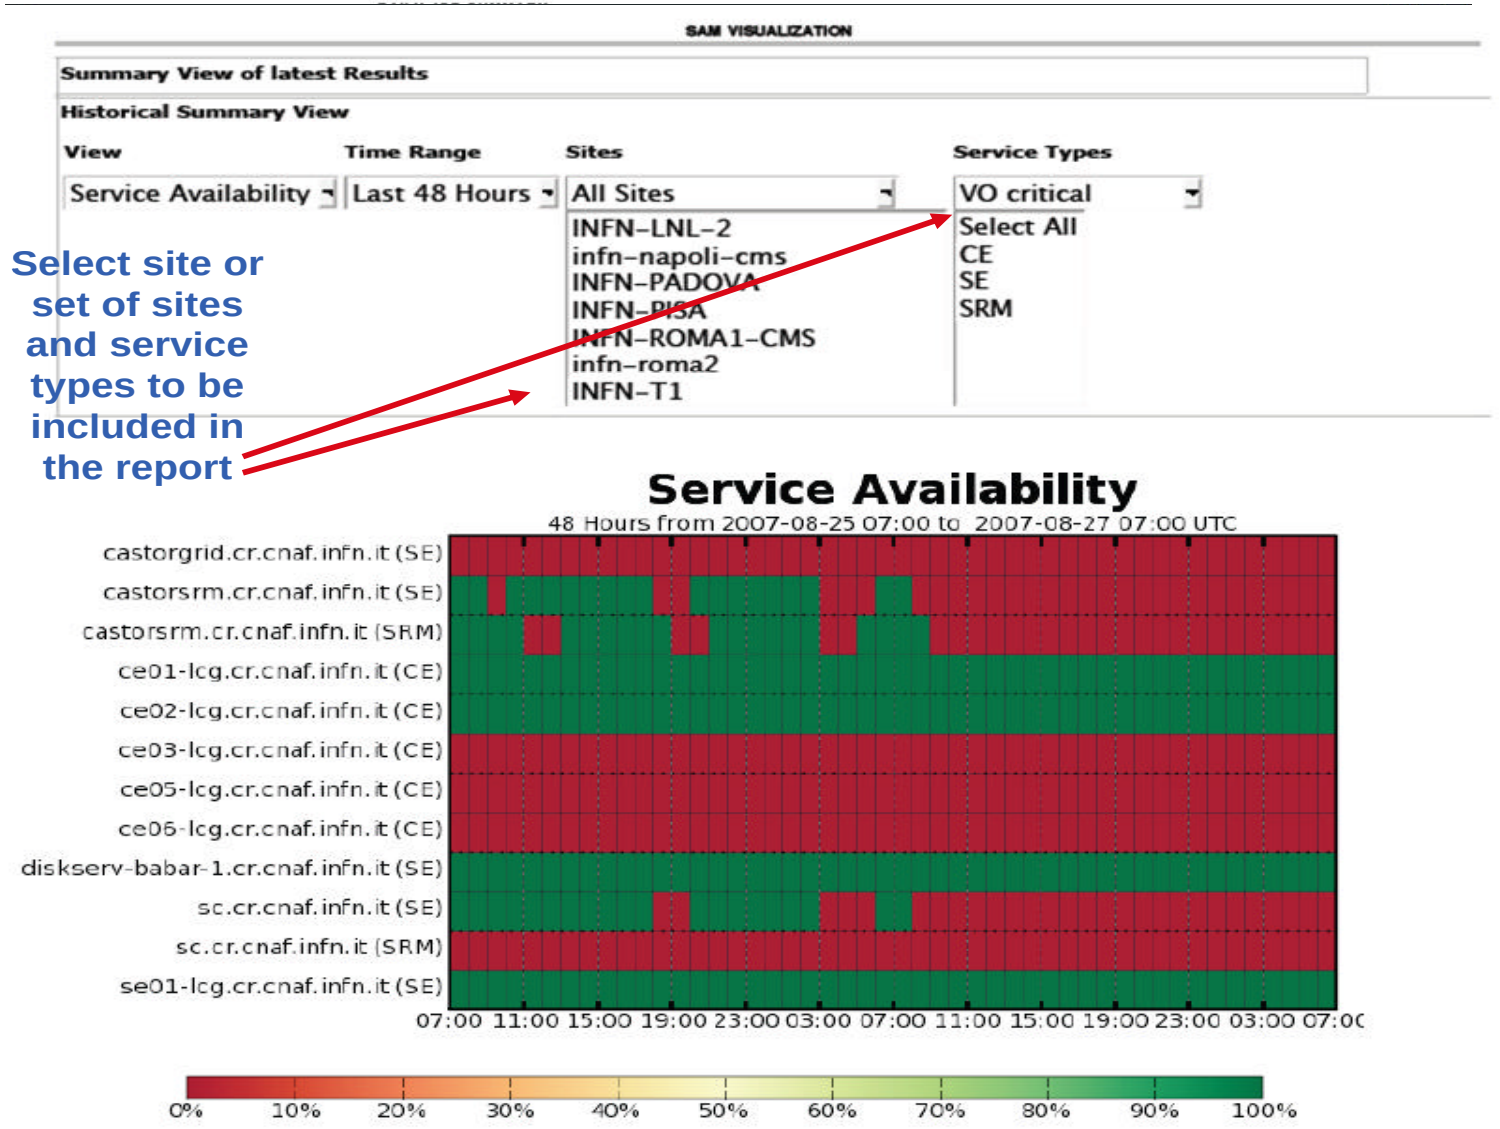

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Grid and application monitoring: SAM (A. Retico)

- A framework developed to monitor remote Grid services for the LCG operations, and now used also by
  - Site Availability calculations
    - GridView
  - Middleware certification
  - Experiment specific tests

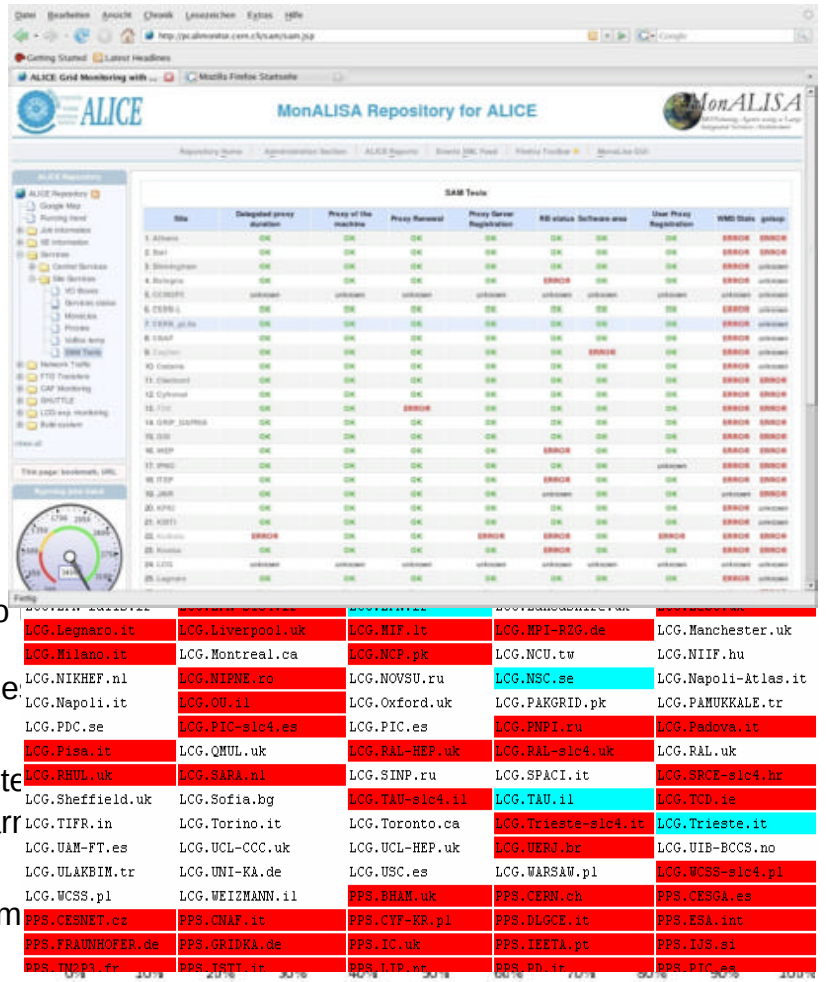

- CMS
  - "plugs" custom tests for the CE and SRM to the worker node to the local storage, etc.
  - Calculates a "CMS availability" based on the
- ALICE
  - Uses SAM to publish the results of custom te
  - Results integrated in ALICE monitoring, alarr
- LHCb
  - Critical tests for LHCb applications are subm SAM
  - Test results are used to build a "production ma
- ATLAS
  - Now only "standard" tests run with an ATLAS proxy
  - Developing custom tests for the SRM (approach similar to CMS)

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Analysis tools: Ganga
(A. Maier, J. Helmsheuser)

- An application to configure – prepare – submit and monitor jobs
  - Submitted to Grid or to a batch system
- Plugins for several applications
  - DaVinci, Athena, etc.
- Use case: LHCb
  - Test analysis code on local machine
  - Analyze partial sample on batch system
  - Analyze full sample on the Grid via DIRAC backend
- ~ 840 users (~10% of LCG users)
  - Mainly from ATLAS and LHCb but 20% non-HEP

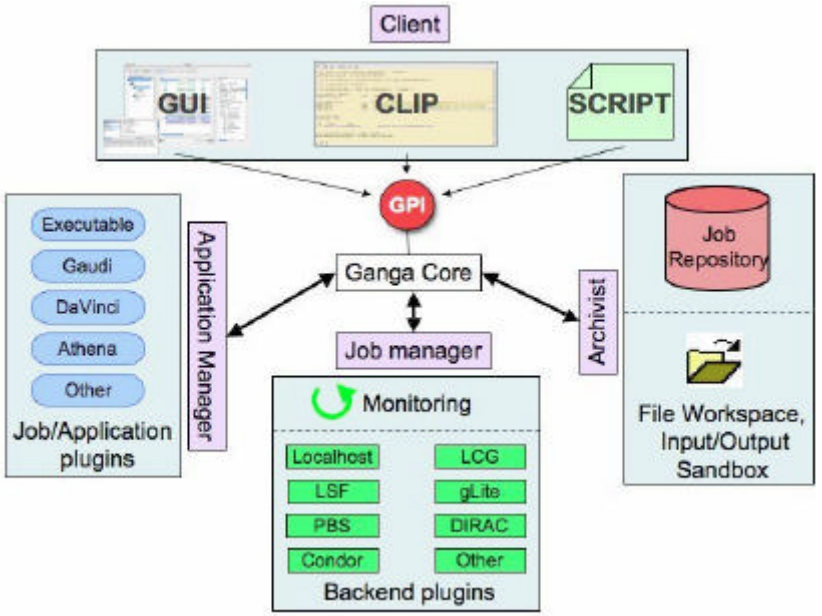

- New developments
  - Ganga robot: to perform complex tasks like
    - Submitting jobs, extracting data about them, reporting stats
  - Windows port
    - Windows UI
    - Windows backend

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

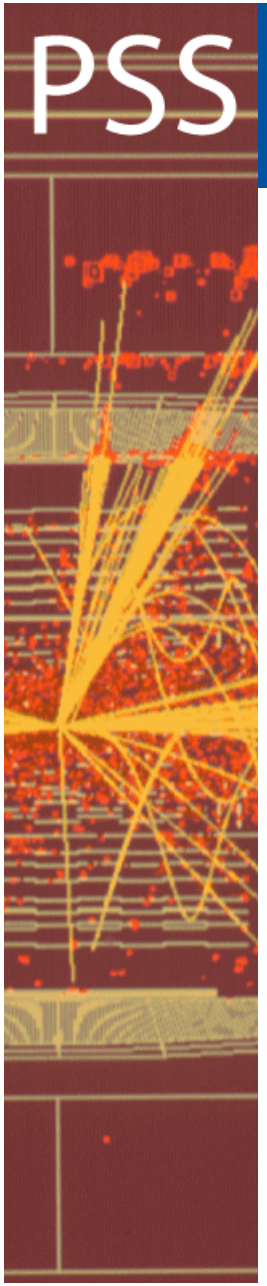

Experiment challenges and activities

CERN IT Department

- CMS Monte Carlo production
  - J. Hernández, D. Evans
- ATLAS Monte Carlo production
  - X. Espinal, S. Campana, Y. Smirnov

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN

CMS Monte Carlo Production

- 2007 goal: 50M events/month
- 2008 goal: $1.5 \cdot 10^9$ simulated events/year
- ProdAgent, a new production system
  - Rewritten from scratch
  - Scalable, avoids single points of failure, supports many grids, better automation
  - Used also for data processing (reconstruction, skimming, etc.)

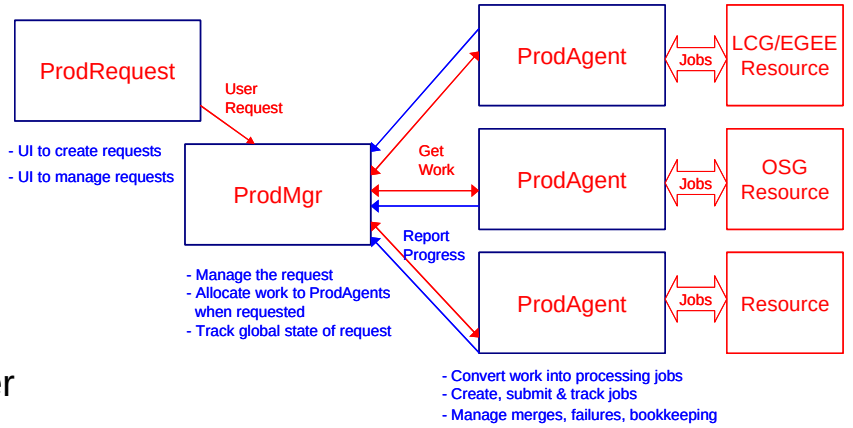

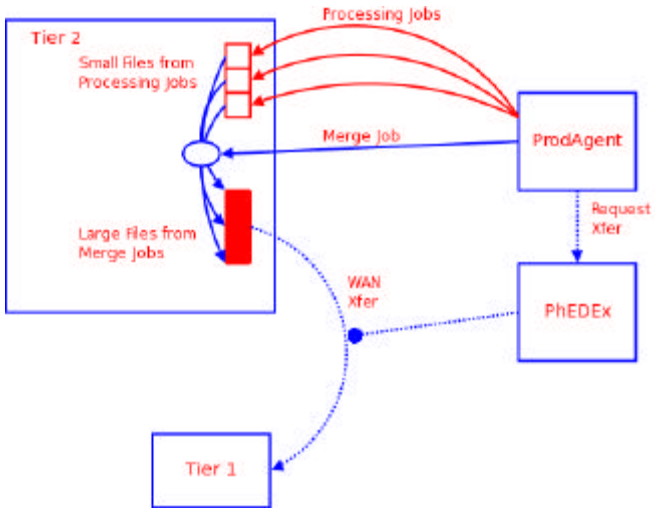

Production Period	Number of MC events	Production Rate	Production Teams / PAs
Winter07	90 M	35M/month	6 / 10
Spring07	67 M	46M/month	6 / 9
Summer07	145 M	64M/month	5 / 7

Overall MC production @ CMS in 2007

~5K slots available at T0/T1

~6K slots available at T2

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CMS MC production performances

- ~175 M events in last 3 months
- ~ 2/3 done at Tier-2 sites
- ~ 50% done at OSG sites

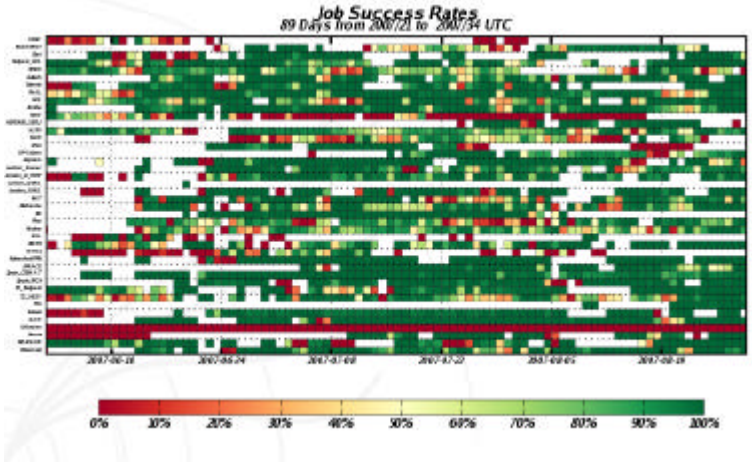

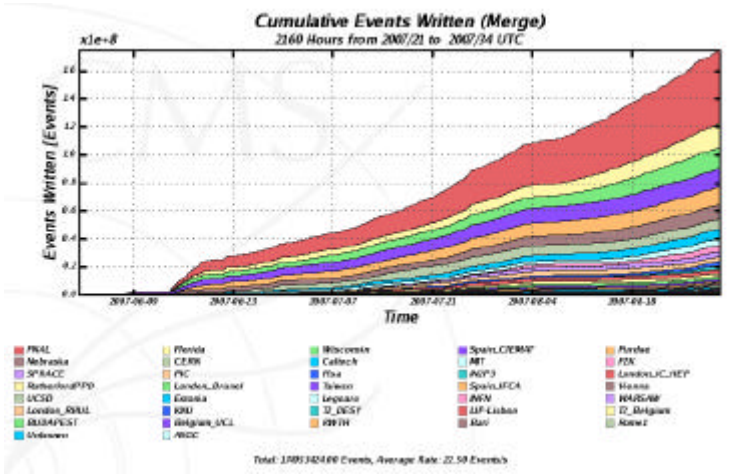

- 40+ sites used
- > 20K/day production jobs
- Average job efficiency ~ 75%, including application and Grid

- Ramp up from June (~ 5000 slots used in average)
- In average ~50% resource occupation. Production inefficiencies, no automatic resource management (manual job release by operators), workflow management (many sites, many small production requests)

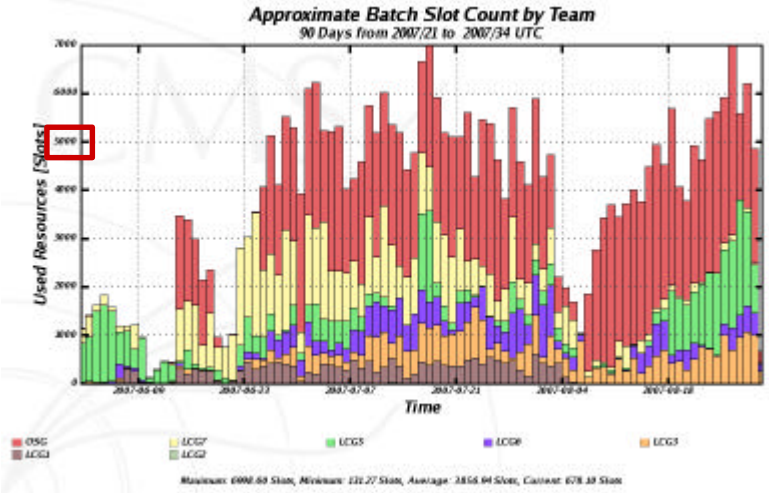

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

ATLAS Monte Carlo Production in EGEE

- Prodsys Architecture
  - ProdDB: common database
  - Eowyn: common supervisor
  - Executors: Lexor, CondorG, CRONUS, Dulcinea, Panda
  - Integrated with the Distributed Data Management
- Operations
  - Production coordinators: assign tasks to clouds, monitor activities
  - Shifters
    - WM: monitor jobs
    - DM: control data flow
  - Need for more automation!
    - Error spotting and reporting
    - Monitoring
    - Automatic task assignment
- Performances
  - 60M events produced in 1st quarter of 2007
  - Up to 55K finished jobs/day
  - Job efficiency ~60%
  - WCT efficiency improving with time and reaching 90%
    - 65% of the inefficiency due to data I/O

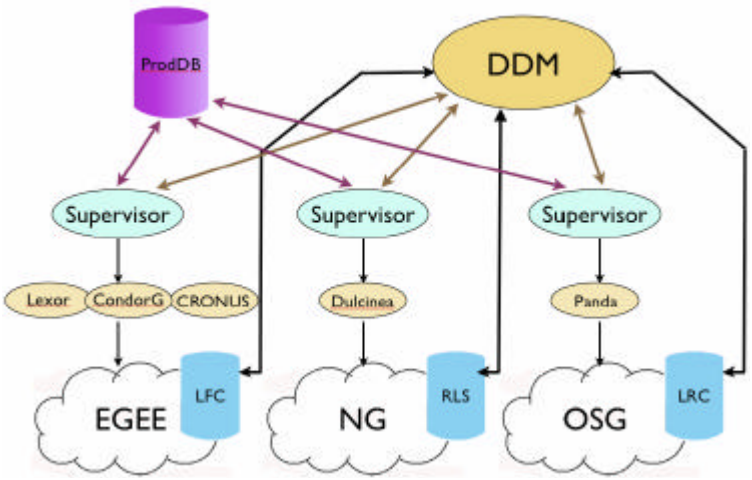

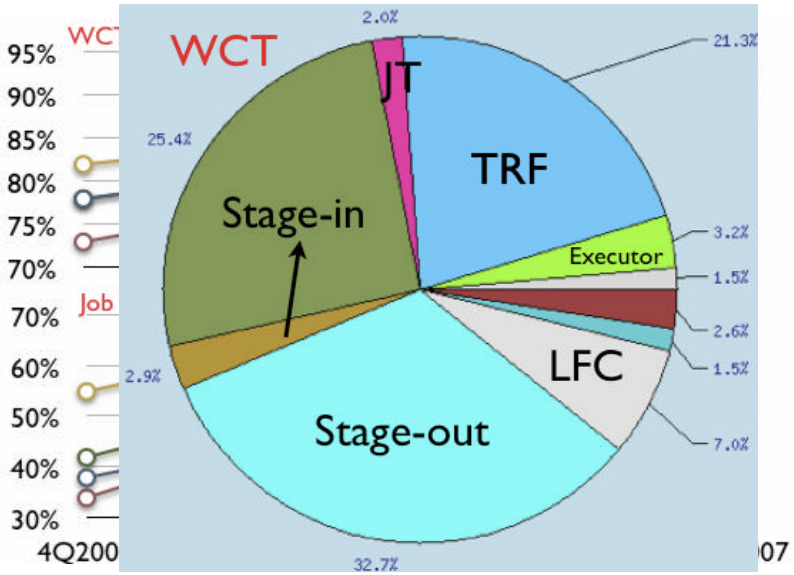

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

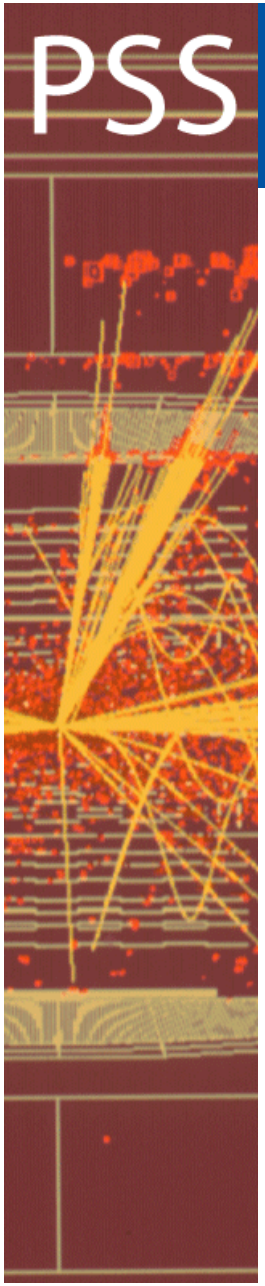

Conclusions

- The parallel sessions provided a very comprehensive coverage of the HEP Grid computing
  - Some topics where particularly represented, even on more sessions
    - WMS systems, tools for production/analysis, storage systems, etc.
- Workload management
  - There are many systems developed by the experiments, but very few which are part of a middleware stack
  - Systems based on pilot jobs are becoming more popular, mainly because they succeed in "hiding" the Grid from the users
  - Still, "push-mode" systems are as reliable and scalable as needed

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

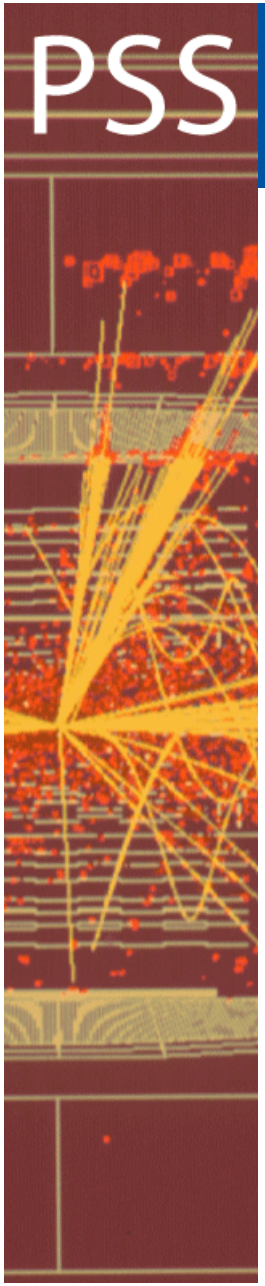

Conclusions

- Data Management
  - Here the top priority is on releasing production-quality SRM 2.2-compliant versions of the SEs
  - The experiments seem to be quite satisfied with their DM systems
- Monitoring
  - Still a lot of work to do to have standardized monitoring systems, but also a lot of new developments
    - SAM and the Dashboard rapidly expanding in functionality and effectiveness
- Experiment activities
  - Taking as example the MC production systems, they are able to reach the targets, but at the expense of a lot of manpower
  - Most of the development goes into improving automation and error recovery

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it