

CMS Grid operations




Andrea Sciabà

CERN




Grid Operations Workshop
13-15 June, 2007
Stockholm




Outline

- Service operations and support
 - CMS services
 - Support channels
 - Support for Grid problems
- Activity operations
 - Data challenge planning
 - An example: Monte Carlo production operations
- User support
 - Current system
 - Future evolution
- Monitoring and other tools
- Conclusions




Services for the CMS computing

- The CMS computing depends on lots of different types of services
 - CMS-specific services
 - **Central**: calibration/alignment database (FroNtier), data catalogues (DBS), production tools, PhEDEx transfer database
 - **Sites**: PhEDEx agents, Squid server, site local configuration files
 - Grid services
 - **Central**: VOMS/VOMRS, FTS, BDII, LCG RB, LCG WMS, Myproxy
 - **Sites**: CEs, SRM servers, site BDIIs
 - Other auxiliary services
 - CMS Central web server, build machines, HyperNews, AFS, Savannah, etc.




Support channels

- CMS service support
 - Mostly done via Hypernews ("community support")
 - Satisfies most needs
 - Savannah is used for problem tracking and bugs
- CMS support at remote sites
 - CMS is active only at sites directly involved in it, which routinely participate to discussions about computing issues
 - CMS sites have a "CMS site contact"
 - Acts as interface between the experiment and the site
 - Manages local CMS services
- Grid middleware and services
 - "Grid experts" in CMS or in EIS team
 - CMS site contacts
 - via GGUS, or via OSG support



Dealing with Grid problems

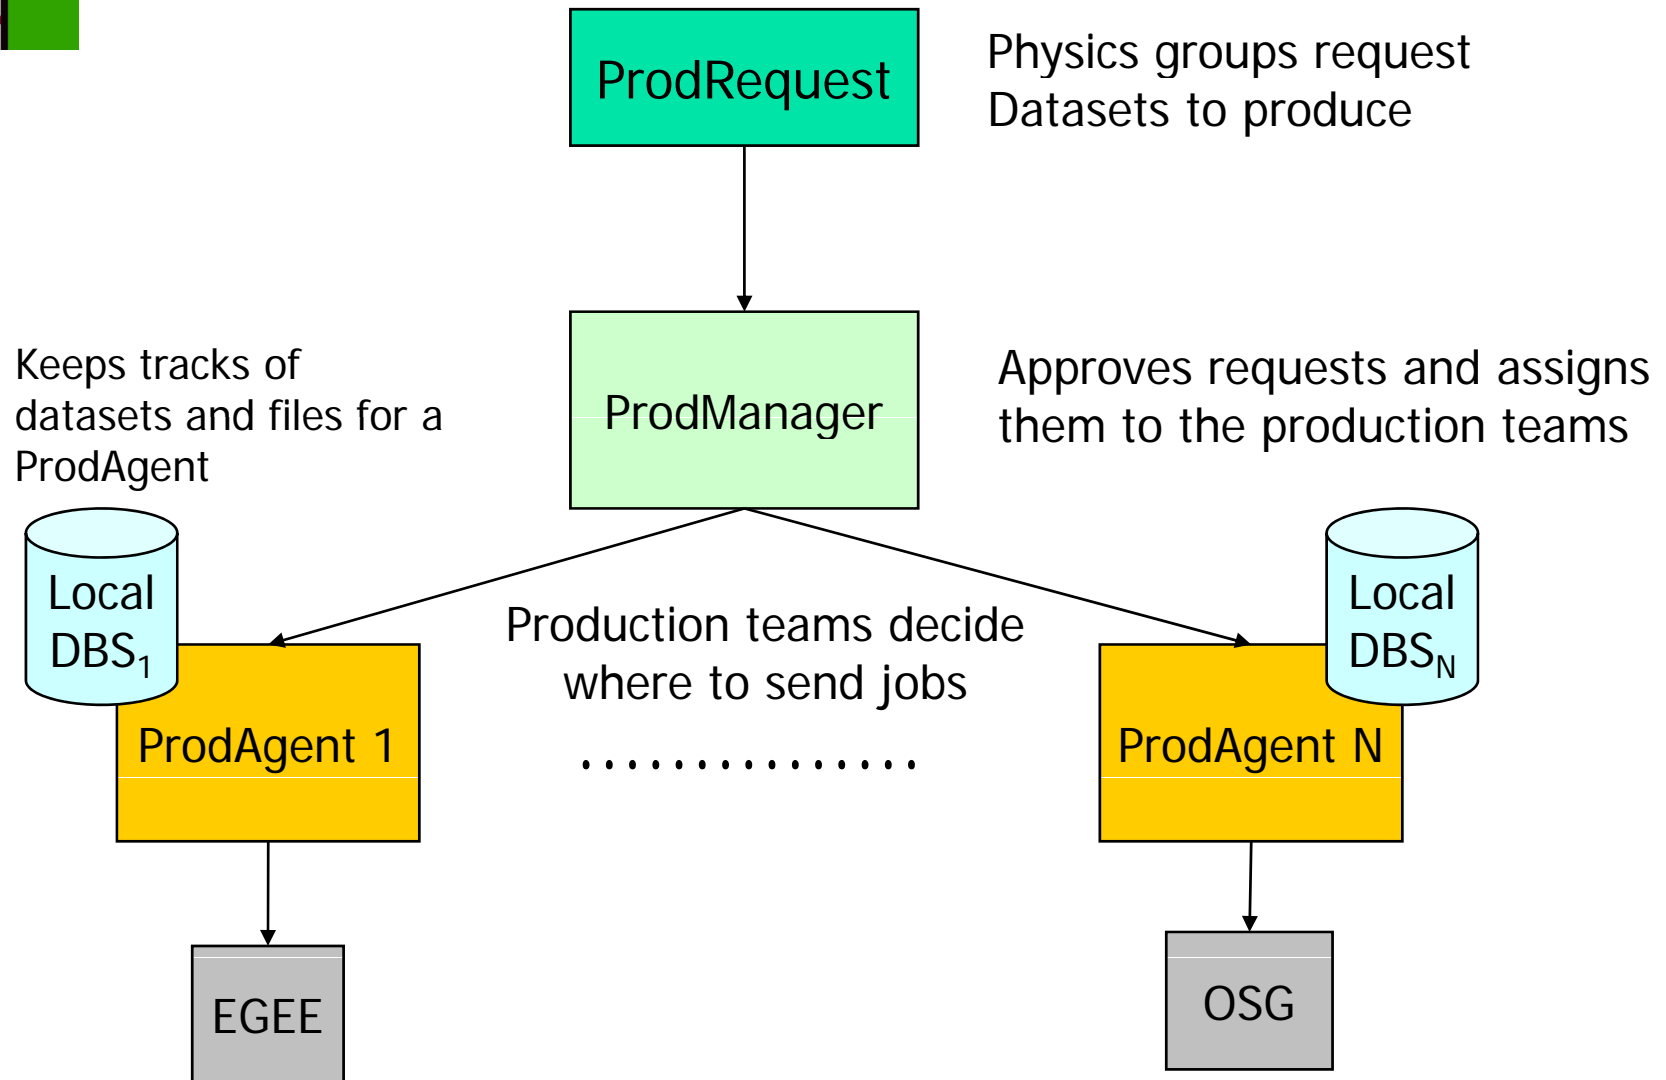

- Workload management issues
 - Mostly encountered by “end users” who submit analysis jobs
 - Normally go through CMS “grid experts”
 - Harder to debug because
 - more components are involved
 - these are more mature, hence problems are subtle
- Data Management issues
 - Most problems are obvious site/storage issues
 - Normally go through via CMS site contacts, which follow up with site managers



WLCG coordination

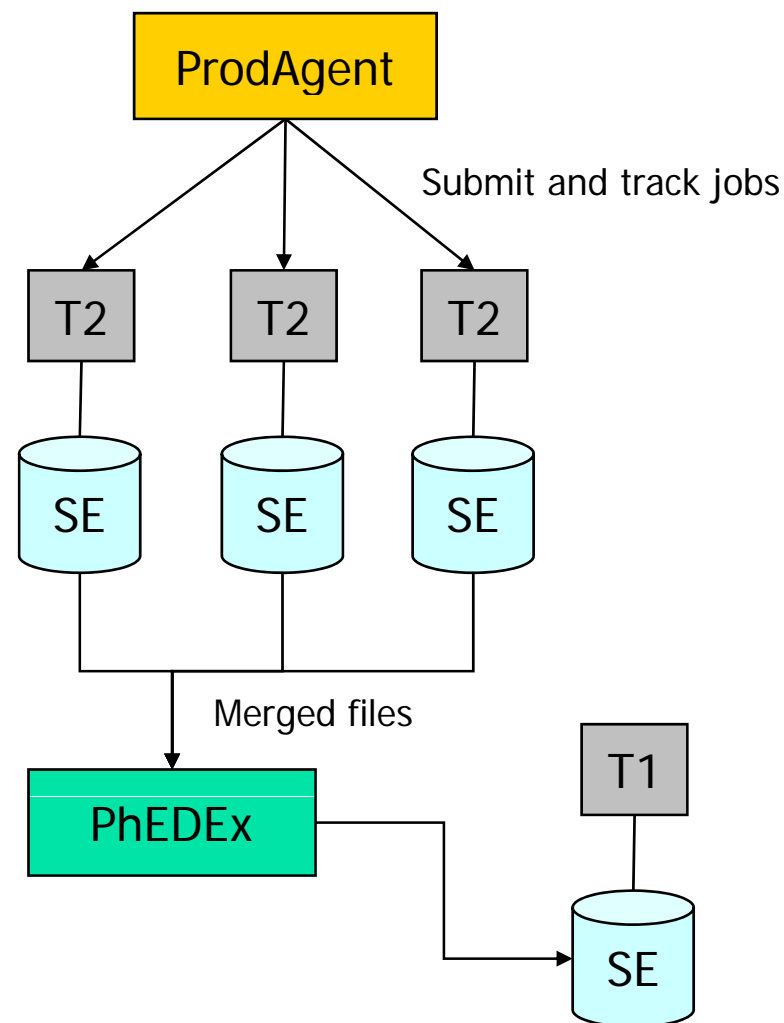
- Communication channels
 - Sites for resource negotiation/allocation
 - CMS-WLCG task force meetings to discuss technical details
 - E.g. SLC4 WNs, production shares, FTS deployment, etc.
 - Experiment Coordination Meeting (ECM) for more general discussions about planning with WLCG and the other experiments


An example: Monte Carlo production operations workflow



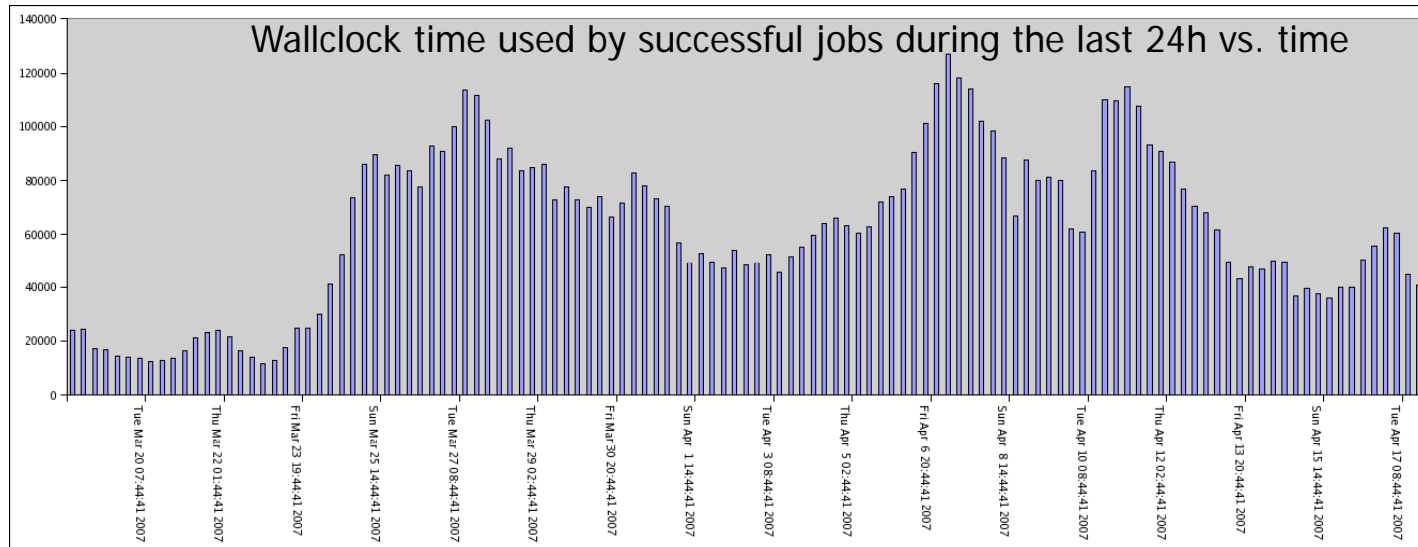
ProdAgent workflow

- Production team \Leftrightarrow a ProdAgent instance
- Job splitting based on job duration (~ 24 h)
- Output files are merged based on desired file size (~ 1 GB)
- MC files are then injected in PhEDEx (the CMS data movement system) and replicated to Tier-1






MC production monitoring




- Several monitoring pages allow to have a global view of how the MC production is going on
 - For example:
 - No. of job slots used in average $\cong 3000$
 - Maximum achieved $\cong 7000$



MC production operational issues


- Problems are taken care of by the production teams
- Interaction with sites goes through local CMS contact persons
- The CMS contact person chooses his preferred communication channel with the site
 - Typically direct interaction with site managers
 - Might also be GGUS
- Current issues
 - CMS problems
 - Bugs in CMSSW, limitations in ProdAgent or DBS, ...
 - Site problems
 - SE, CMS software deployment, data transfers, ...
 - "Grid" problems
 - workload management problems, as usual, are harder to debug
 - Often jobs are just resubmitted...

User support: how it worked so far




- Almost all requests for help are addressed to a Hypernews forum
- Advantages
 - It often provides the quickest response
- Disadvantages
 - Nobody is really committed to provide a solution
 - The user might not know to which forum address his question
 - People may lose track of whether a problem was eventually solved and how
 - It is inconvenient to browse over the forum archives to find the solution for a similar problem

User support: how it should be done from now




- A “physical” CMS helpdesk exists at CERN and FNAL
- The single access point for user support is the CMSSW Savannah page
 - Advantages
 - Perfect if the user does not know whom to contact
 - Problems are followed until solved
 - A solution is guaranteed
 - Disadvantages
 - It needs finding people committed to solve user problems!
 - Status
 - Already functional, but still building expertise
- Hypernews will still be the preferred choice for
 - Discuss development
 - Exchange ideas




Grid computing support (I)

- The idea is that eventually the end users will not be directly exposed to Grid problems
 - Communications with Grid only through “experts”
 - Currently goes through the CRAB forum
 - CRAB is the user interface to submit and manage CMS analysis jobs
 - It works remarkably well thanks to the CMS developers’ support
 - It should eventually move to the general user support



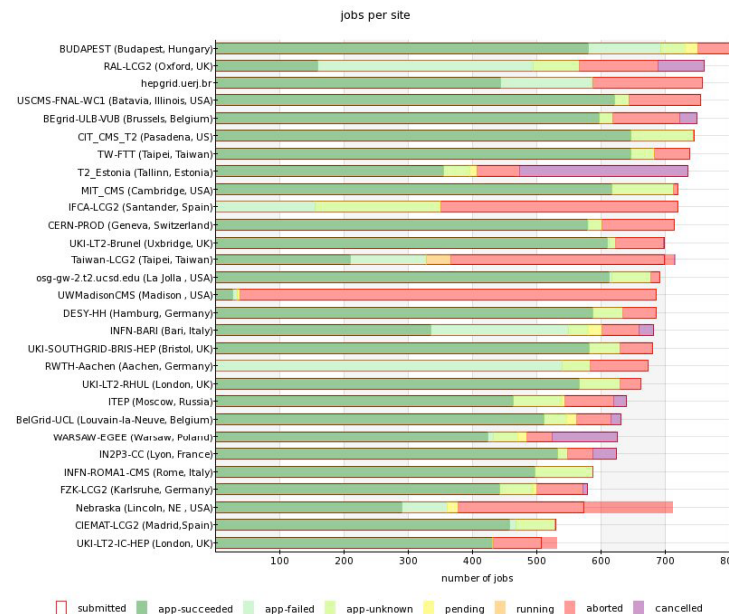
Grid computing support (II)

- CMS support people will fill GGUS tickets on behalf of the user (who will stay in the loop) whenever needed
- Open issues
 - Still not clear how CMS site contact persons fit into the picture of user support
 - Should one be able to assign a ticket to a CMS site contact person through GGUS?
 - How are the sites supposed to contact the central CMS support? Are EGEE broadcasts enough?
 - Avoid to end up having parallel channels for user support and operation support
 - In fact, discussions are going on in CMS!





Monitoring tools

- ARDA dashboard
 - Global monitor of job successes/failures
 - Already covered by Julia's talk
- SAM tests
 - See next talk
- Job Robot
 - Submits "fake" analysis jobs to sites using CRAB
- PhEDex monitoring
 - Monitors everything related to data movement



Job Robot results in June





Other tools: SiteDB

RAL :: Site Details

RAL is a Tier 1 site based in Didcot, Oxon, UK. It runs the LCG grid middleware.

To view the pledged resources for RAL click [here](#).



CMS Contacts

[Chris Brew](#)

"Site Admin"

tel 1: 01234, tel 2: 56789

[Neil Geddes](#)

"Site Executive"

tel 1: None, tel 2: None

[Sebastien Greder](#)

"Data Manager"

tel 1: None, tel 2: None

[Simon Metson](#)

"Site Admin"

tel 1: +441179288714, tel 2: 71641

[Dave Newbold](#)

"Site Executive"

tel 1: 01234, tel 2: 56789

Software on RAL

Site has development releases installed.

Site installs software manually.

lcgce01.gridpp.rl.ac.uk

- CMSSW_1_4_2
- CMSSW_1_3_1
- CMSSW_1_3_0

Site Configuration

GOC information

[GOCDB Link](#)

Compute Element(s)

lcgce01.gridpp.rl.ac.uk

Storage Element(s)

Click to see DBS entries

[raismra.rl.ac.uk](#)

PhEDEx configuration

Click to subscribe data

[TI_RAL_Buffer](#) [TI_RAL_MSS](#)


Associated Sites

Child Sites


- [Estonia](#)
- [London](#)
- [Rutherford PPD](#)
- [Bristol](#)

Site monitoring


PhEDEx transfer rate to RAL
- click plot for larger version



RAL PhEDEx transfer quality
- click plot for larger version



Some plots from Dashboard, number of file block registered at the site etc



Conclusions

- Operational support in CMS goes through several channels
 - But procedures are often not clear
 - The communications flow needs to evolve towards a more mature model
- User support is already making a significant progress from a simple “community support” model
 - GGUS is seen as a fundamental component of the system
- Several monitoring systems
 - Increasingly effective in helping to bring sites in good shape