CMS and gLite

Nicolò Magini
IT-GS-EIS
CERN 'White Areas' Lectures
Friday March 27th, 2009

www.eu-egee.org

e-infrastructure

Outline

- **Introduction on the CMS computing model**
- **Data Management**
  - Data bookkeeping – DBS, DLS
  - Data transfers – PhEDEx, **FTS**
  - Storage – **SRMv2**
  - Data transfer commissioning
- **Workload management**
  - Monte Carlo production - PRODAGENT
  - Analysis – CRAB and CRABSERVER
  - Using the **gLiteWMS**
- **Monitoring and commissioning**
  - **Dashboard** monitoring
  - Site commissioning – **SAM**, JobRobot, **SiteStatusBoard**
- **Conclusions**

CMS Computing Model

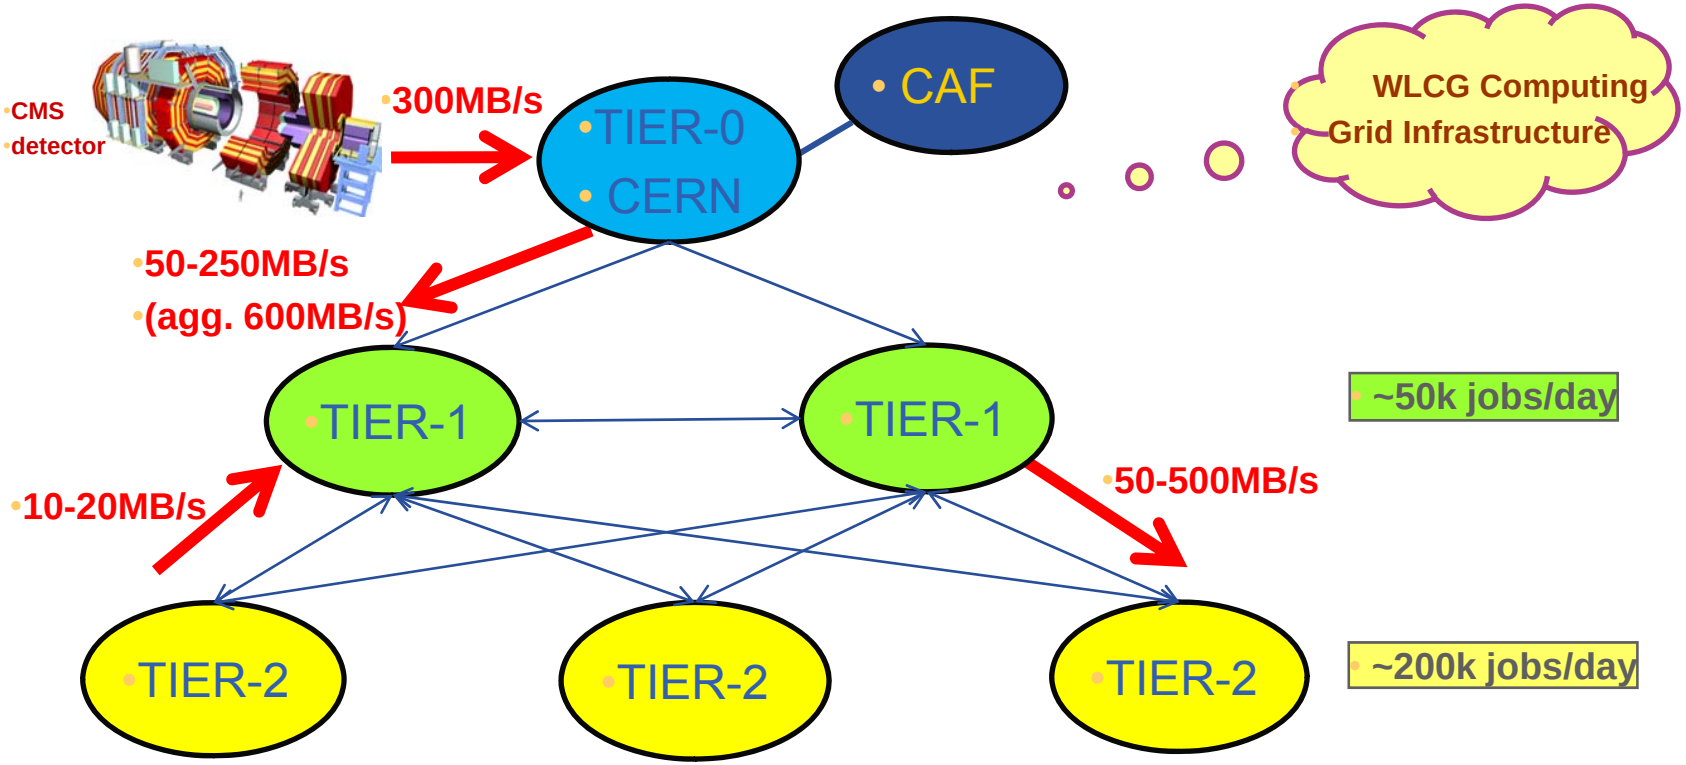

- **Tier-0**
  - *(the accelerator centre)*
  - Data acquisition & initial processing
  - Long-term mass data storage
  - CMS CERN Analysis Facility
  - (latency critical data processing, high priority analysis)
  - Distribution of data → Tier-1 centres
- **7 Tier-1s**
  - *("online" to the DAQ)*
  - High availability centres
  - Custodial mass storage of share of data
  - Data reconstruction and reprocessing
  - Data skimming & selection
  - Distribute analysis data → Tier-2s
- **~50 Tier-2s**
  - *in ~20 countries*
  - End-user physics analyses
  - Detector Studies
  - Monte Carlo Simulation → Tier-1

Data Organization

- **CMS expects to produce large amounts of data (events)**
  - O(PB)/year
- **Event data in files**
  - average file size reasonably large (≥ GB)
    - output files merged when too small
    - avoid scaling issues with storage systems and catalogues when dealing with too many small files
  - O(10^6) files/year
- **Files are grouped in fileblocks**
  - group files in blocks (1-10 TB) for bulk data management reasons
    - exist as a result of either MC production or data movement
  - 10^3 Fileblocks/year
- **Fileblocks are grouped in datasets**
  - Datasets are large (100 TB) or small (0.1 TB)
    - Dataset definition is mostly physics-driven

Data Tier	Description (size estimate in 2009)
GEN	Event kinematics
SIM	Simulation with detector material
DIGI	Simulation of detector readout
RAW	Raw detector output (~1.5 MB/evt)
RECO	Reconstructed event (~500 kB/evt)
AOD	High level analysis objects (~100 kB/evt)
USER	Final user files

Data Management components

- **Data Management System to discover, access and transfer event data in a distributed computing environment**
  - Fileblock-based data replication and tracking
  - Local (trivial) file catalogue
  - Data location based processing
    - Data is placed at a site through explicit CMS policy, not moved around in response to job submission

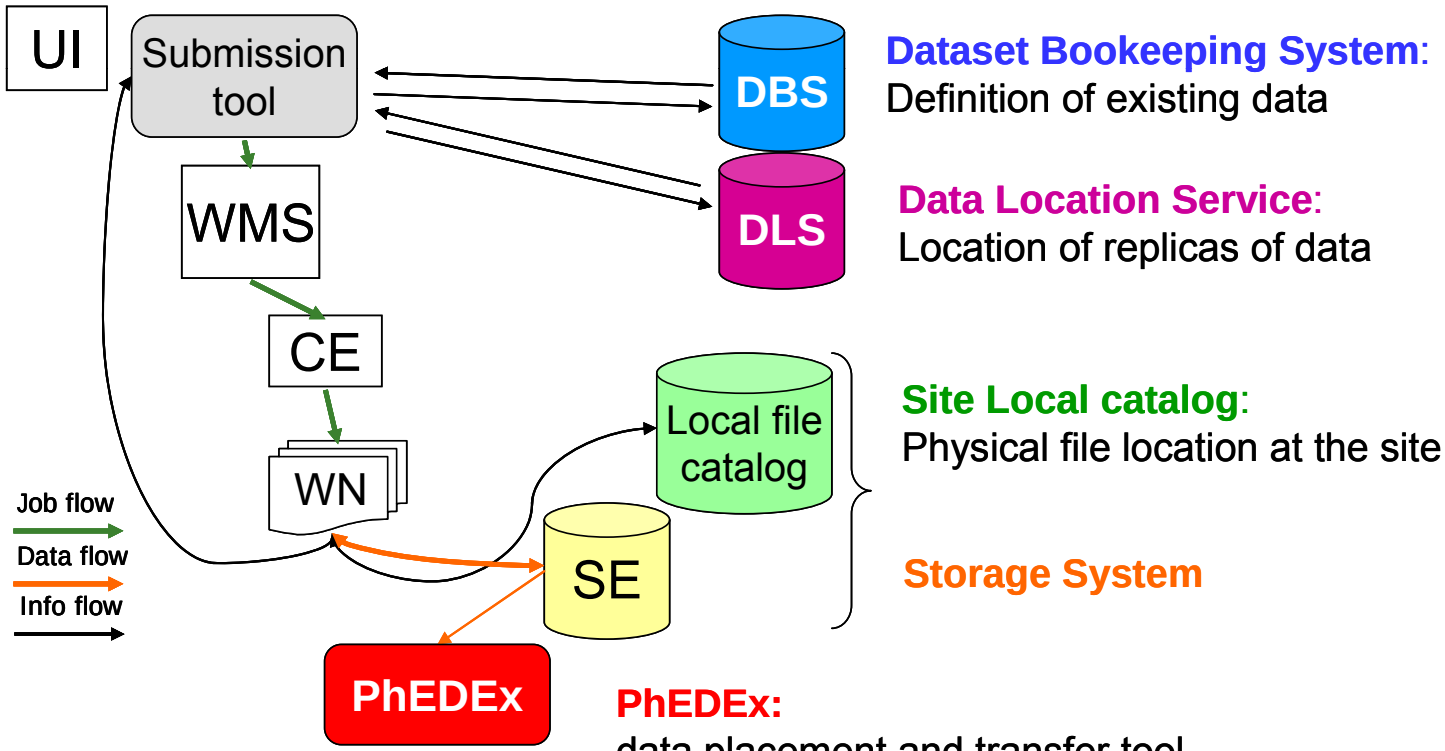

Dataset Bookkeeping System (DBS)

- **Data definition**
  - specifies what the data is and how it was produced
- **Data discovery**
  - knows what data exists and how it is organised in fileblocks
- **Deployed at**
  - Tier-0 (central DBS)
  - User space (private DBS)
- **Implemented as**
  - Oracle DB at T0 (central DBS)
  - SQLite DB (private DBS)
- **Accessed through**
  - Webpage
  - Python API - CLI

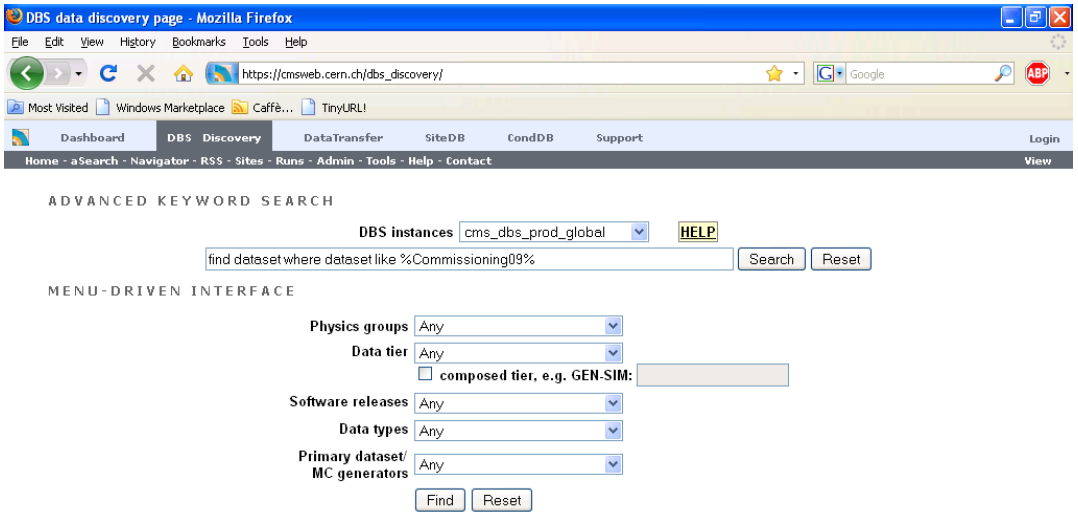

Data Location Service (DLS)

- **Allows to find out at what sites the data is available**
  - maps fileblocks to SEs
  - does not contain the physical location of files
- **No longer an independent component**
- **Current implementations use other DM components**
  - Tier-0 (global DLS)
    - Now implemented as a read-only view of PhEDEx replica tables
  - Local DBS instances
    - Now incorporated in DBS

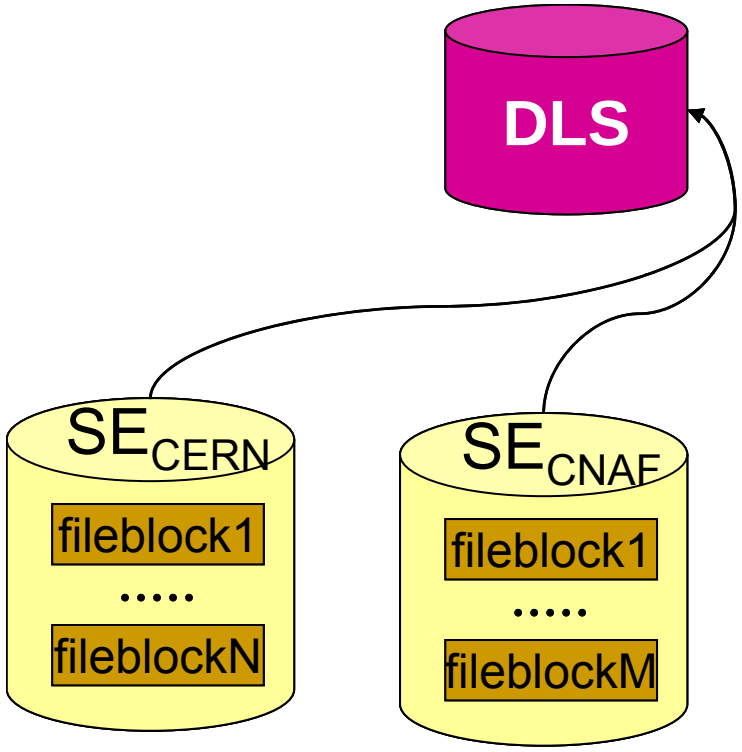

Local ("trivial") File Catalogue

- **“site-local information remains site-local”**
- **Contains the physical location of files**
  - maps LFNs to PFNs
- **Implementation**
  - “trivial” LFN→PFN mapping
  - file preinstalled on experiment area at each site

```
<!-- Below is the default rule -->
<lfn-to-pfn protocol="direct"
  path-match="/+(.*)"
  result="/castor/cern.ch/cms/$1"/>
<lfn-to-pfn protocol="stageout" chain="direct"
  path-match="(.*)"
  result="$1"/>
<lfn-to-pfn protocol="rfio" chain="direct"
  path-match="(.*)"
  result="rfio:$1"/>
<lfn-to-pfn protocol="srmv2" chain="direct" destination-match="T1_CH_CERN_.*" is-custodial="y"
  path-match="(.*)"
  result="srm://srm-cms.cern.ch:8443/srm/managerv2?SFN=$1"
  space-token="CMS_DEFAULT"/>
<lfn-to-pfn protocol="srmv2" chain="direct" destination-match="T1_CH_CERN_.*" is-custodial="n"
  path-match="(.*)"
  result="srm://srm-cms.cern.ch:8443/srm/managerv2?SFN=$1"
  space-token="CMS_DEFAULT"/>
<lfn-to-pfn protocol="srmv2" chain="direct" destination-match="T2_CH_CAF"
  path-match="(.*)"
  result="srm://srm-cms.cern.ch:8443/srm/managerv2?SFN=$1"
  space-token="CMS_CAF"/>
<lfn-to-pfn protocol="srmv2" chain="direct"
  path-match="(.*)"
  result="srm://srm-cms.cern.ch:8443/srm/managerv2?SFN=$1"/>

<pfn-to-lfn protocol="direct"
  path-match="/+castor/cern\.ch/cms/(.*)"
  result="/$1"/>
<pfn-to-lfn protocol="rfio" chain="direct"
  path-match="^rfio:/+(.*)"
  result="/$1"/>
<pfn-to-lfn protocol="srmv2" chain="direct"
  path-match=".*\?SFN=(.*)"
  result="$1"/>
```

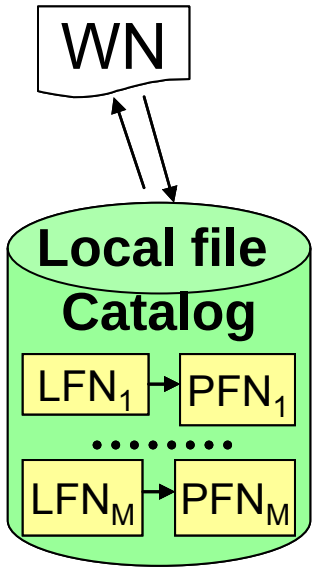

- *Extract from the CERN TrivialFileCatalogue*

CMS transfer management system

- **PhEDEx - Physics Experiment Data Export**
  - reliable, scalable dataset replication system

- **Each CMS site runs a set of software agents**
  - Loosely coupled, highly-specialized components designed to fulfill a specific "simple" task in a reliable way
    - Central agents: routing, task assignment, …
      - *Run at CERN*
    - Site-specific agents: download, export, mass storage staging and migration …
      - *Run at every CMS site on a gLite VOBOX (to simplify proxy management) or UI*
- **Agents inter-communicate with a central blackboard**
  - Transfer Management DB (TMDB), Oracle RAC backend
    - Block replica location & file mapping
    - Block subscription and allocation
    - File metadata information (filesize, cksum, etc)
    - Transfer state (at node; in transfer; wanted; available)
- **Transfer status monitored from PhEDEx webpage**
- **DataService to retrieve information from the database through simple http access**

- **1 T0, 7 T1's, 49 T2's + T3's in current PhEDEx transfer topology**

PhEDEx Design Overview

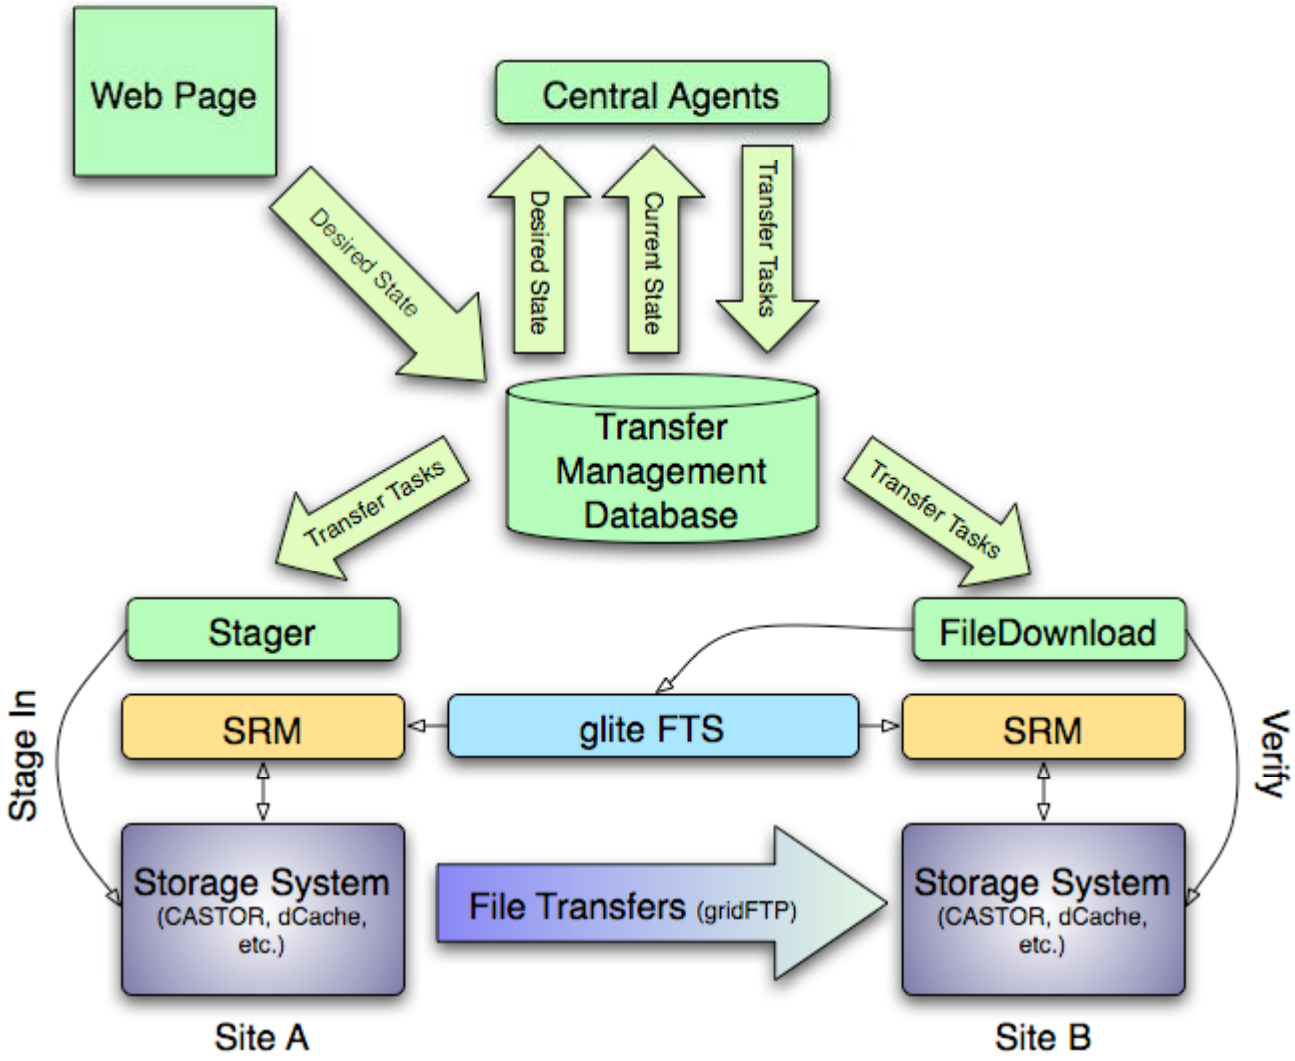

FTS backend in PhEDEx

- Transfers in PhEDEx are started by the DownloadAgent running at the destination site
  - All sites are pulling data in CMS transfers
- The DownloadAgent is interfaced with gLite FTS since ~ 3 years
  - Can also use other transfer tools as backend
    - Example: srmcp between OSG sites
- New FTS backend deployed in March 2008
  - Main features
    - Possiblity to configure limits on pending/active files on a per-link basis
    - By default, FTS channels are kept full by limiting new submissions on pending transfers rather than active
    - Improved transfer status monitoring with constant rate polling to FTS server

FTS usage policy in CMS

- T0 → T1 transfers
  - T1 site submits to FTS-T0-EXPORT service at CERN
- T1 → T1 transfers
  - Destination T1 site submits to its own FTS server
- T1 → T2 transfers
  - T2 site submits to source T1 FTS server
- T2 → T1 transfers
  - T1 site submits to its own FTS server
- Other transfers
  - Destination site submits to its associated T1's FTS server

Storage Resource Manager usage

- **SRMv2.2 deployed at CMS sites in CCRC'08 Phase 1**
  - Variety of different implementations deployed
    - CASTOR, dCache, DPM, StoRM, BeStMan, even a private implementation!
  - 'SRMv1-like' deployment: no Space Tokens required, but sites are free to define them when needed

- **Remote access to SRM**
  - Mostly through FTS (EGEE) or srmcp (OSG) in PhEDEx transfers
    - PhEDEx can discover the destination SpaceToken to use in transfers if it's published in the destination site's TrivialFileCatalogue
    - The site can publish in TFC different SpaceTokens to be used by PhEDEx in different contexts. Typical example at a T1:
      - TAPE SpaceToken for Custodial data import
      - DISK SpaceToken for non-Custodial data import
  - High-level tools (lcg-cp/srmcp) for remote stageout of user job output
- **Local access to storage**
  - Normally local access protocol defined in site local configuration
    - rfio, dcap, direct posix access, etc.
  - Can also use SRM clients (srmcp/lcg-cp), but unaware of Space Tokens

Data transfer commissioning

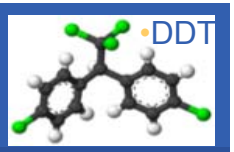

- **Debugging Data Transfers**
  - A program to maintain a high-quality transfer network for CMS
  - Uses PhEDEx LoadTest generator
  - Metric to determine if a link is usable for data operations
    - Link achieves COMMISSIONED status when:
      - *For links from T0 and T1s:*
        - Transfer burst of 1.65 TB in 24h (20 MB/s average)
      - *For links from T2s to T1s:*
        - Transfer burst of 0.42 TB in 24h (5 MB/s average)
    - Interact with sites to solve issues encountered in commissioning
      - *FTS channel configuration issues, FTS timeouts*
      - *SRM "errors" of some kind, gridftp timeouts, storage system issues*
      - *myproxy problems, or simply proxy expired*
      - *Network failures*
      - *File exists or error in path at destination*
      - *Filesystem authentication problems, delegation issues*
      - *PhEDEx agents down un-noticed, or misconfigured*
      - *LoadTest samples not available*
      - *etc…*
    - Provide contacts to experts and support teams and document known issues

Improvement in transfer rate

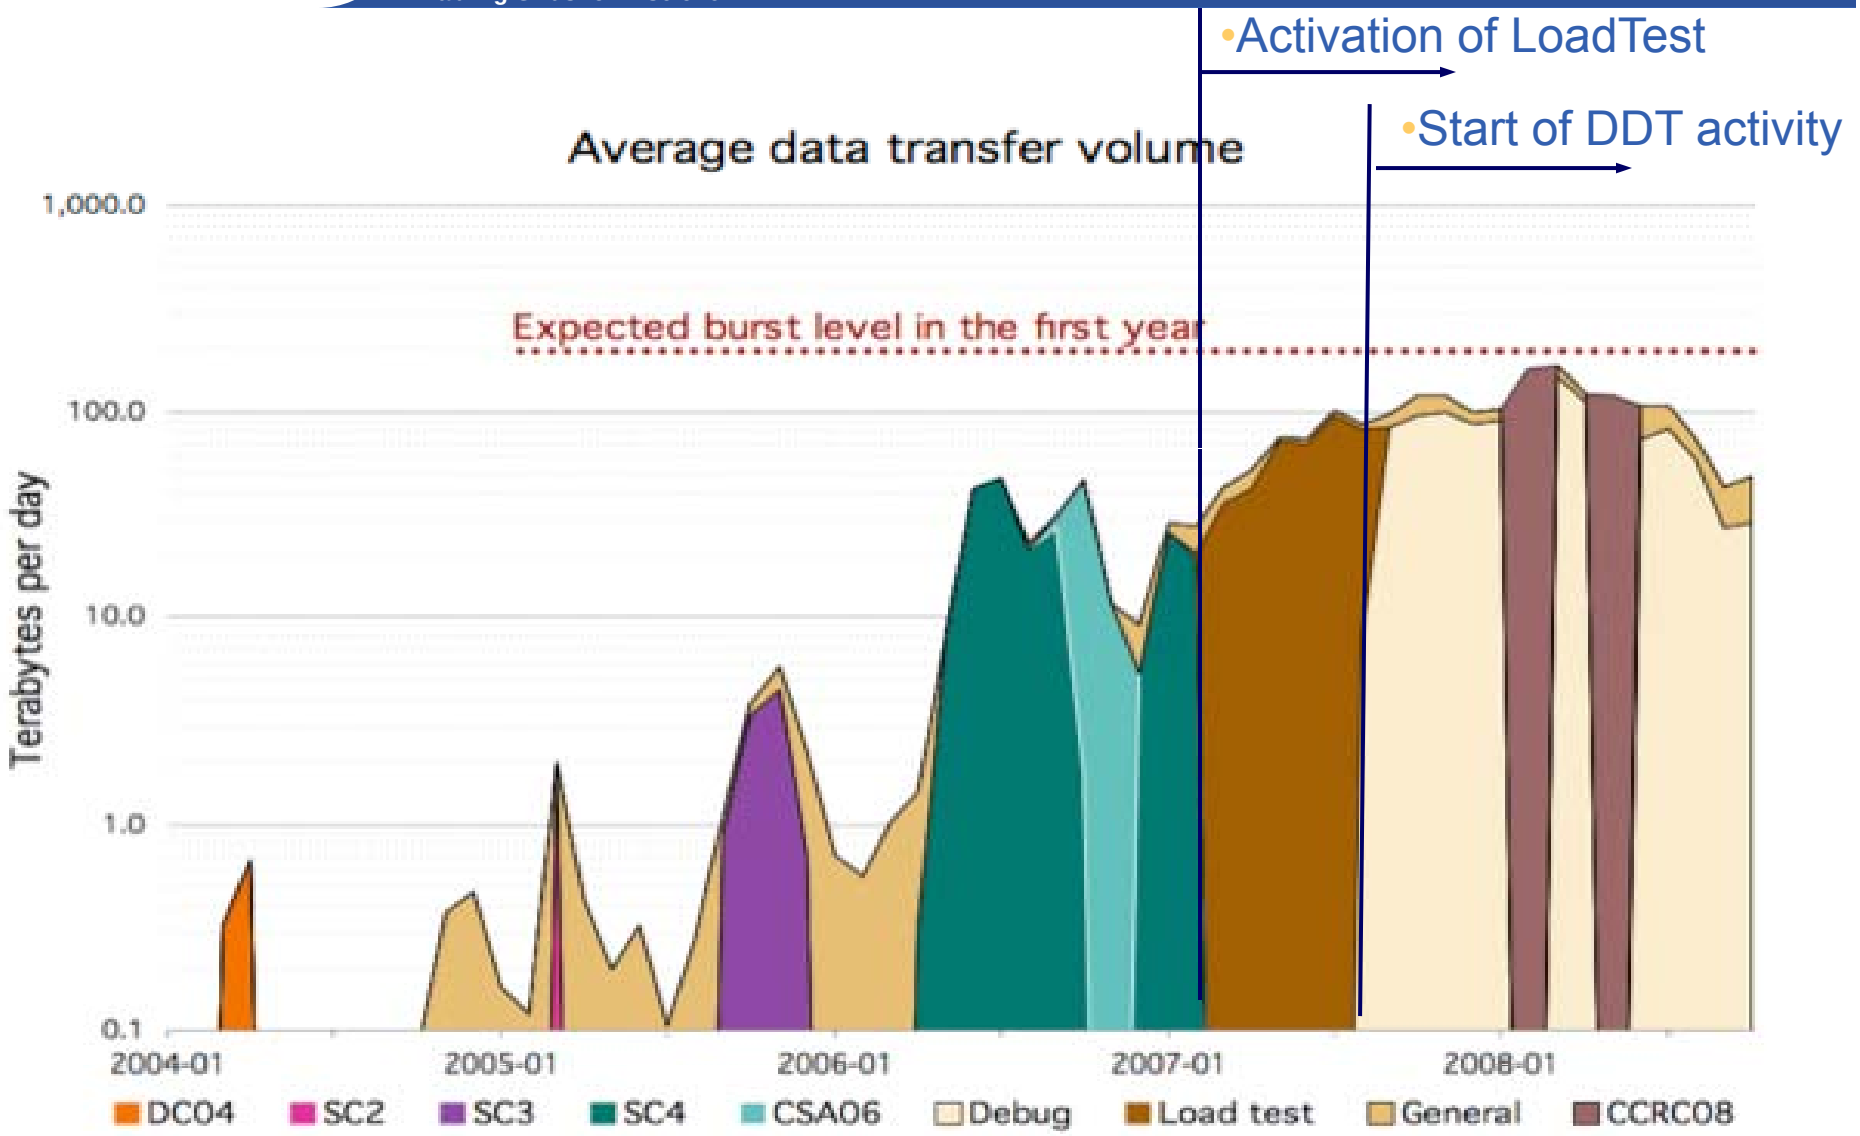

Link quality monitoring

- ***Transfer quality monitoring enabled on all links***
- ***To guarantee a minimum number of attempts per day, a 'heartbeat' style transfer is enabled on all links by setting a low rate LoadTest injection***
  - *~1 file transferred every 3 hours.*
- ***Current target: quality >50% on at least half of the links***

- ***Transfer quality map for T1 → T2 links***

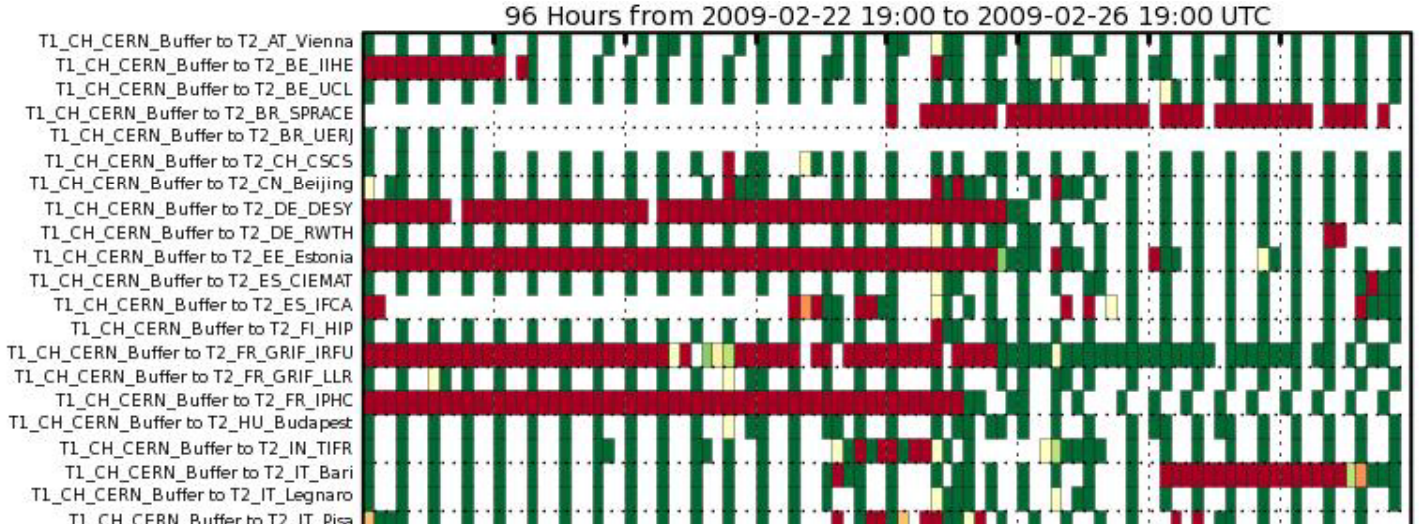

DDT current status

Status in February 2009: 523 COMMISSIONED links

- **all 56** T0 → T1 and T1-T1 cross-links COMMISSIONED
- 300/352 (85%) T1 → T2 downlinks COMMISSIONED
  - 43/49 (88%) T2s have at least one downlink
- 140/352 (40%) T2→T1 uplinks
  - 41/49 (83%) T2s have at least one uplink
- Plus 27 T2-T2 cross-links (not in Computing Model)

Data Management Summary

- **FTS massively used in PhEDEx transfers**
- **SRMv2.2 deployed**
  - But exploitation of functionality new to SRMv2.2 is restricted to a limited usage of SpaceTokens
  - CMS likes to ‘keep it local’ – SRM bypassed through local protocol access when possible
- **LFC not used at all**
  - Using DLS_PHEDEX + local TrivialFileCatalogue instead
- **Use of high-level data management clients for single file stagein/stageout**
  - lcg_util
  - dCache srm client
  - No direct usage of GFAL APIs

Workload Management: CMS use-cases

- **Monte Carlo Production:**
  - Automatic, parallelized system for simulation/reconstruction of huge data samples
    - Goal 1.5 x 10^9 Simulated Events per year at 40+ CMS T2s
- **Official data processing (RAW→RECO→AOD)**
  - Reprocessing, skimming

- **Basic analysis tasks (single user):**
  - Transparent usage of the Grid infrastructure as well as local batch system, integrated with the CMS workload management system
- **Regime Analysis and intensive analysis tasks**
  - Centralized system dealing with huge tasks, automating the analysis workflow, optimizing Grid usage
  - High concurrency system for multiuser environment
    - O(2000) users on O(50) sites

Production system

- **Used for Monte Carlo production and official data processing**

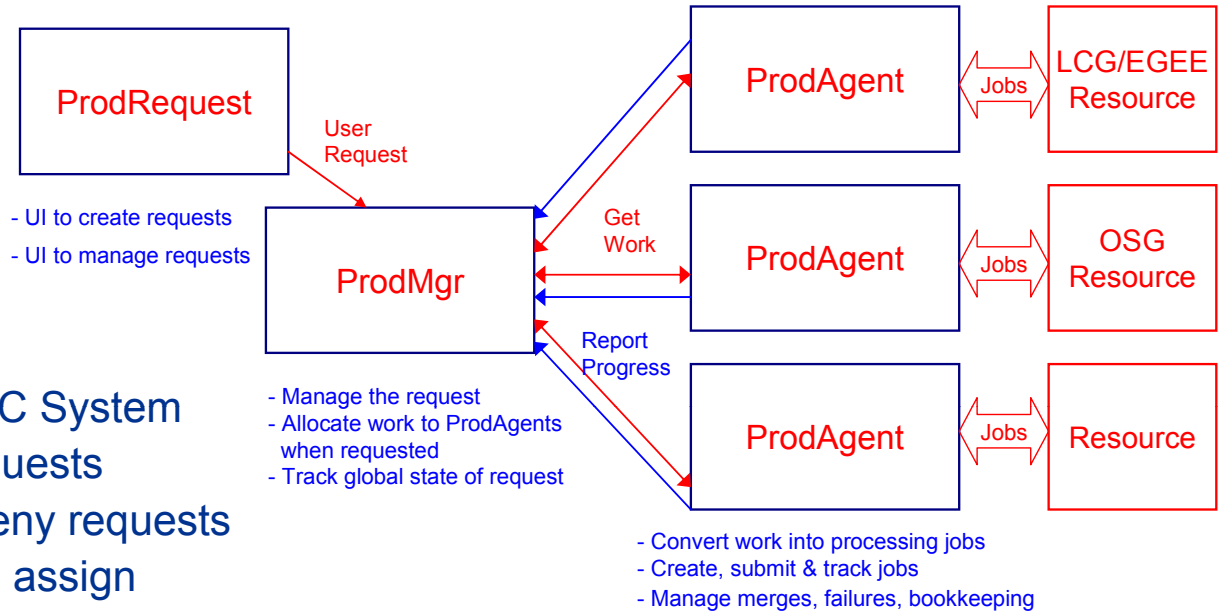

- **ProdRequest**
  - Global Point of entry to MC System
  - Users can make/track requests
  - Managers can approve/deny requests
  - Production Managers can assign requests to ProdAgents
  - ProdMgr retrieves workflow specs and feed work to the ProdAgents
- **ProdMgr**
  - manages MC processing requests
  - assigns work to the Production Agents
  - evaluates the status of requests (finished, unfinished, failed, ...)

- Production Agent
  - Acts as a common front end to different resources (grids, farms)
  - Splits work in atomic components
  - Submits and tracks jobs
  - Merges output files
  - Publishes output data into DBS and DLS
  - Injects produced data in PhEDEx

ProdAgent workflow

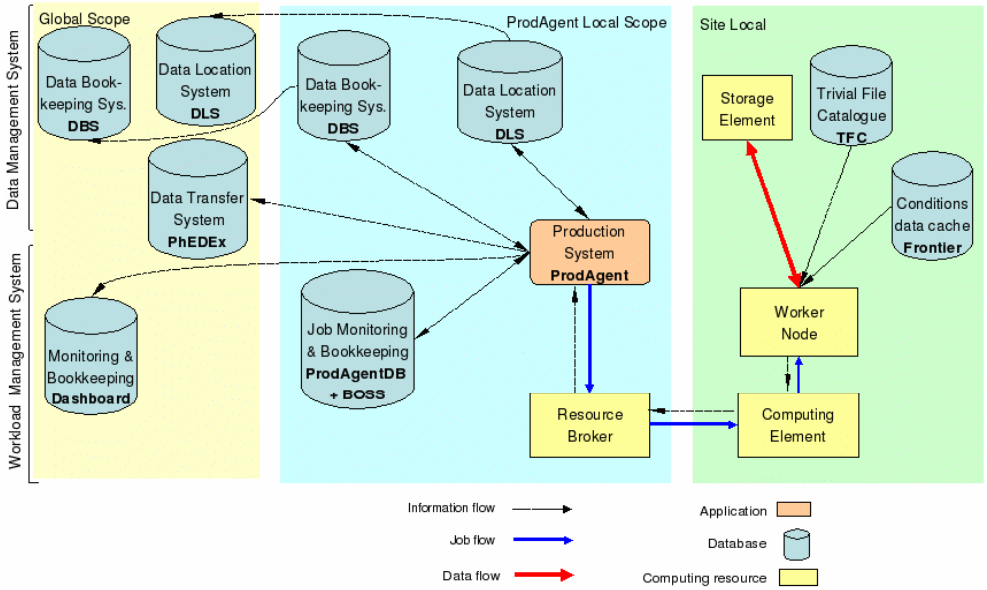

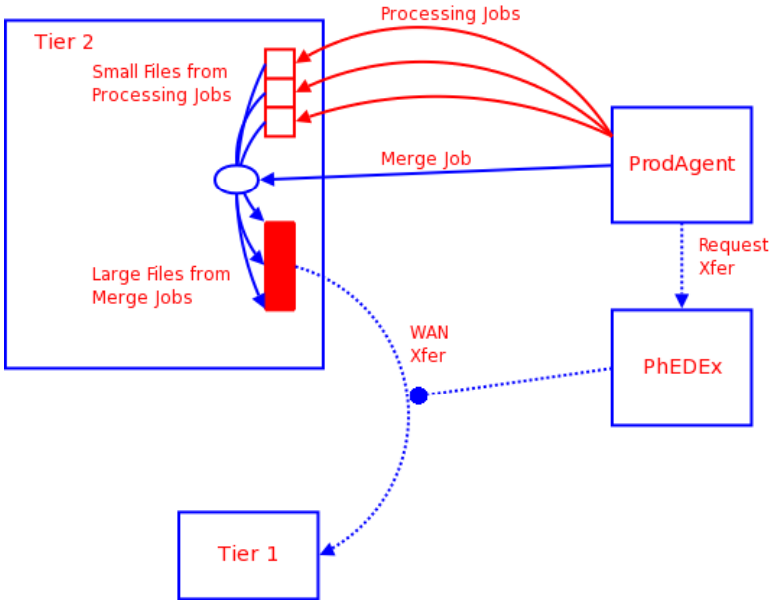

- Work split in atomic loosely coupled components
- Data processing, bookkeeping, tracking and monitoring occurs in local-scope
- Data bookkeeping and location information promoted to global-scope databases and data transfer system after successful processing
- Scaling achieved by running any number of concurrent ProdAgent instances

Typical analysis workflow

- **The user develops his analysis code and builds the executable and the libraries**
  - done on a UI
- **The user defines which data he wants to analyse**
  - data registered in DBS
- **The user submits an analysis task to the Grid (or a local farm)**
  - transparent access to different Grids
- **The task is divided into many parallelized jobs**
  - jobs allocated to run near the data
- **The user checks the status of his task**
- **When the task is finished, the user collects the results**
  - stored on user space or a SE

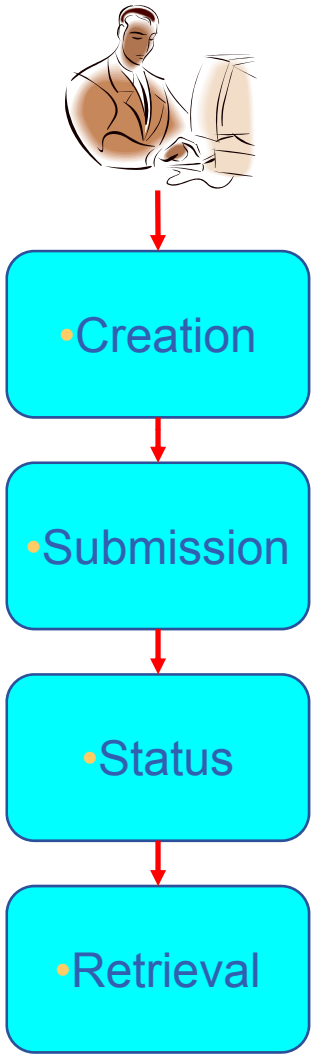

CMS Remote Analysis Builder

- Used for small analysis tasks

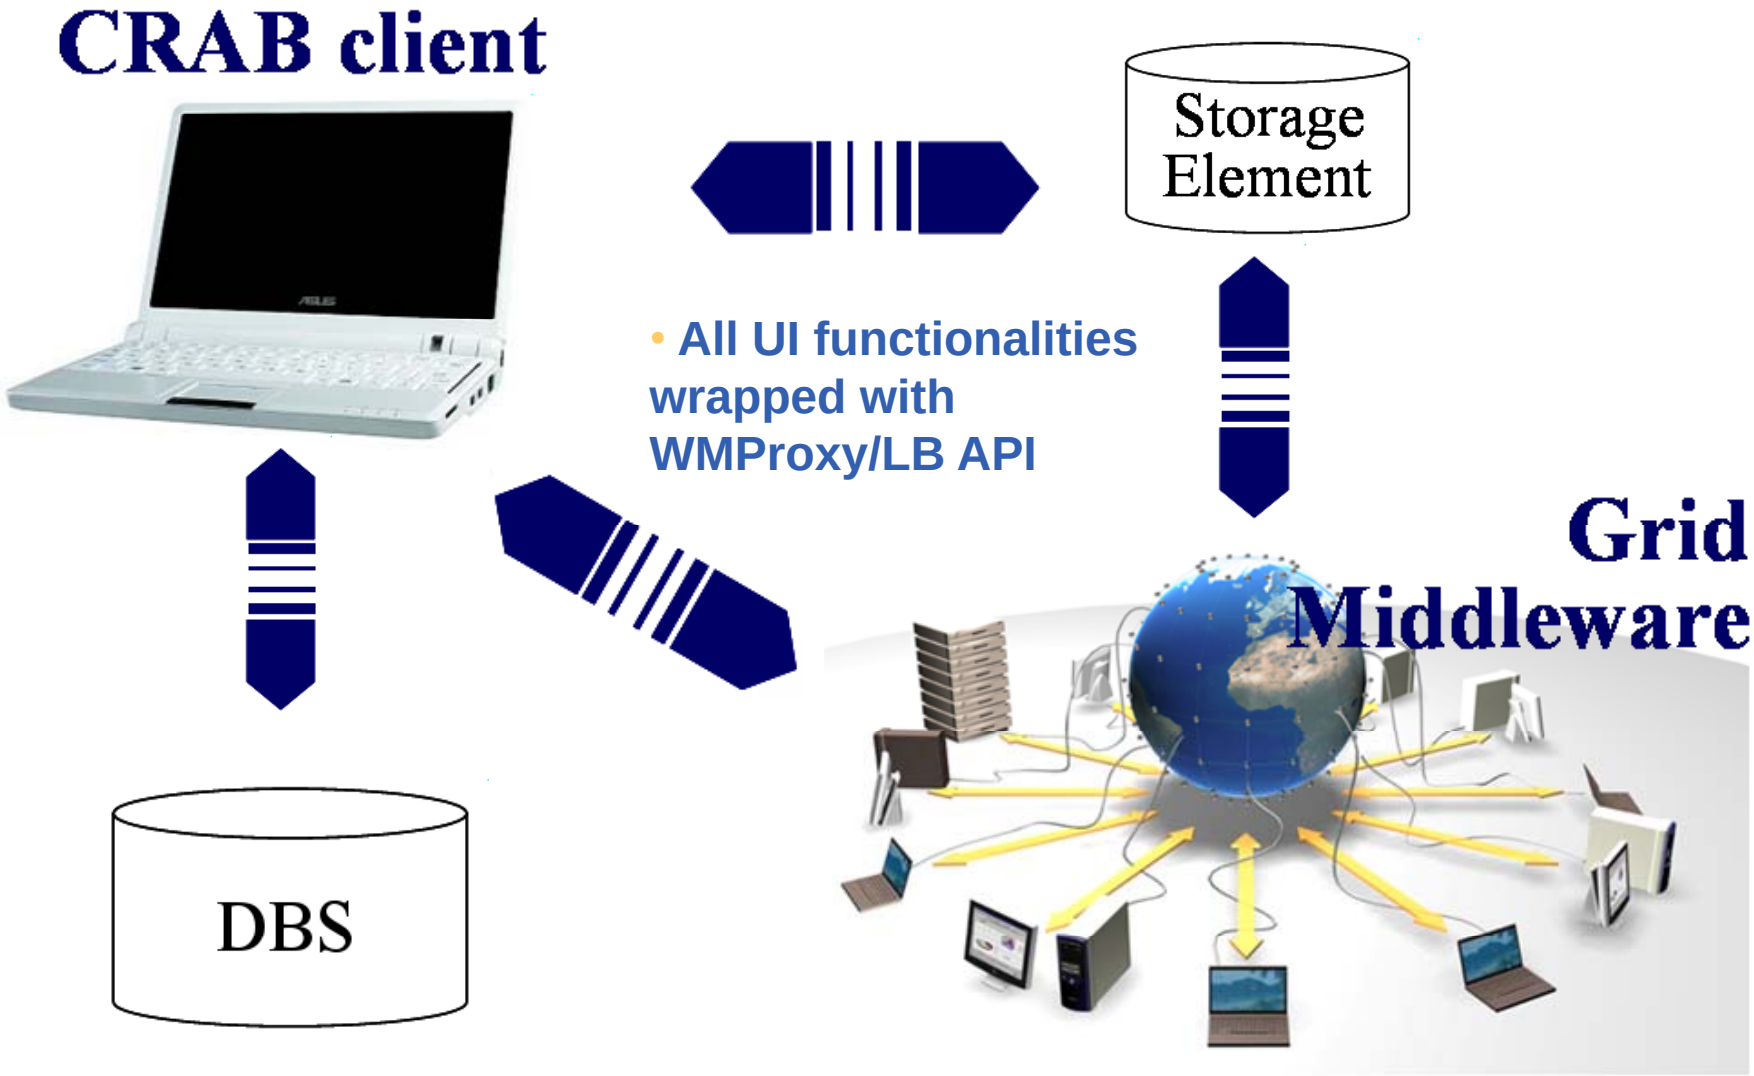

CRAB Analysis Server

- Used for larger scale analysis tasks

CRAB client

- Multi Threaded Output Handling

Storage Element

- Multi Threaded Job Submission (collections, many users concurrently)
- Direct ISB/OSB Shipping from WN

DBS

CrabServer

Grid UI

Delegation

- Multi threaded Status Query

Grid Middleware

Interaction with WMS systems

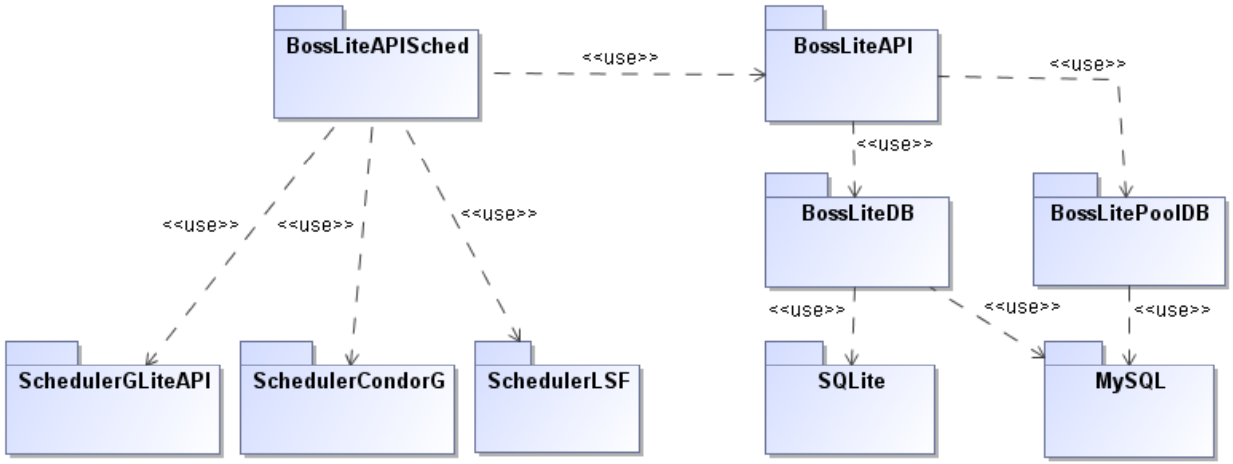

- **BossLite: basic common CRAB and PRODAGENT component to interface with Grid and batch systems**

- **Plugins for transparent interaction with Grid** [WLCG, OSG] **and local batch systems** [LSF, PBS…]

- **DataBase backend for logging and bookkeeping**

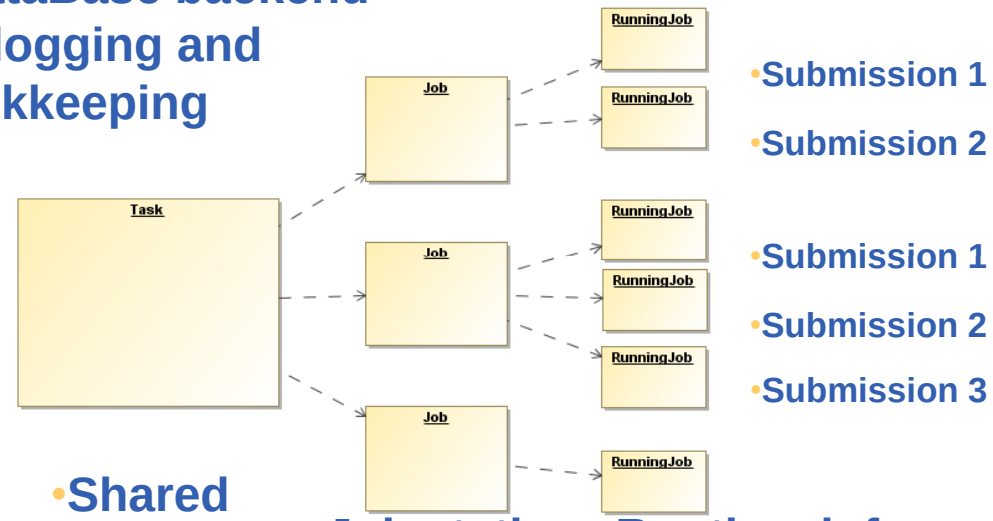

- **User Task Description:** identical jobs accessing different part of a dataset or producing a part of a MC sample

BossLite plugins

- **glideinWMS**
  - glidein = Condor pilot job
  - Used exclusively for organized reprocessing with PRODAGENT at T1s
  - Tested also for analysis with CRAB
- **gLiteWMS**
  - Accesses gLite WMS through WMProxy python API
  - Accesses LB through API for easy extended monitoring
  - Supports bulk submission, bulk match-making, bulk status query
  - Used for Monte Carlo production with PRODAGENT at T2s
  - Used for analysis with CRAB
- **LSF**
  - Used for analysis at CAF
- **...**

Single gliteWMS usage

- MC production and analysis jobs are balanced over many WMS: currently 7 for the analysis, 4 for the production

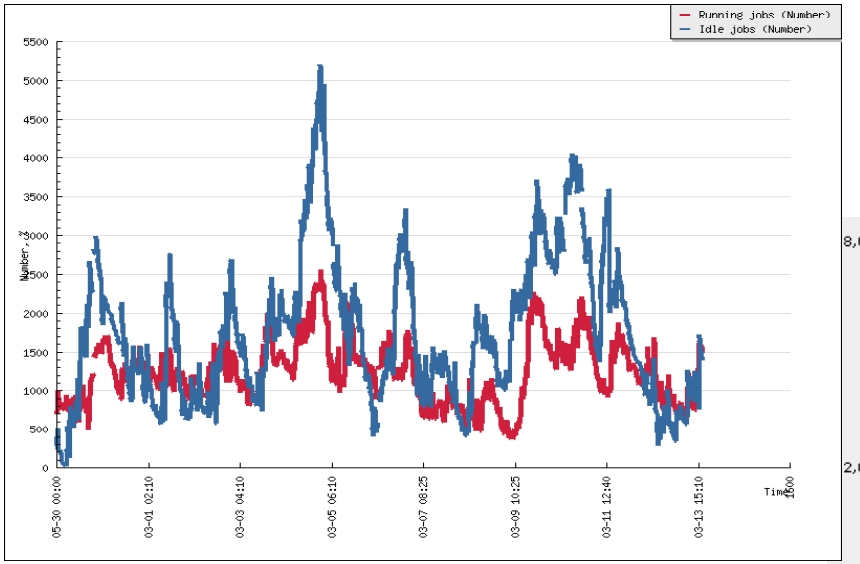

- Daily job rate, including ended jobs
- Typical job load for a single WMS may already reach a steady rate of 15k jobs per day
- Stress tests reached 30 kJobs per day without breaking point signal for a single WMS!

- Typical instantaneous load of a single WMS (jobs running/idle)
  - Up to 5 kJobs simultaneously handled in every day analysis per WMS

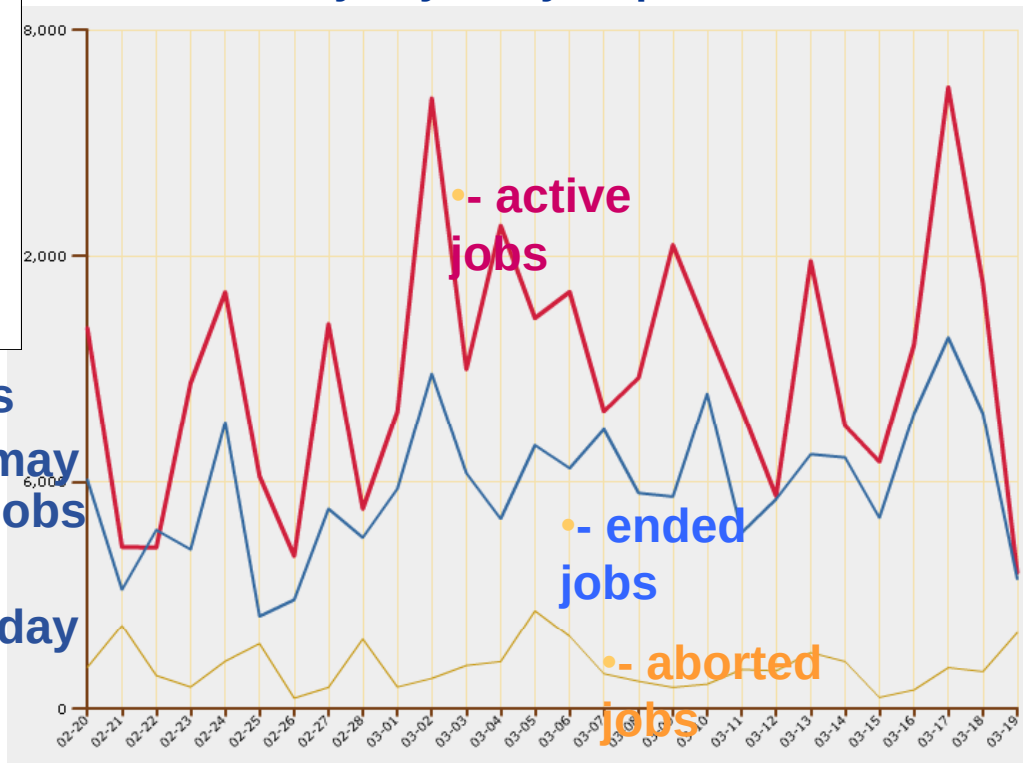

Overall WMS Performance

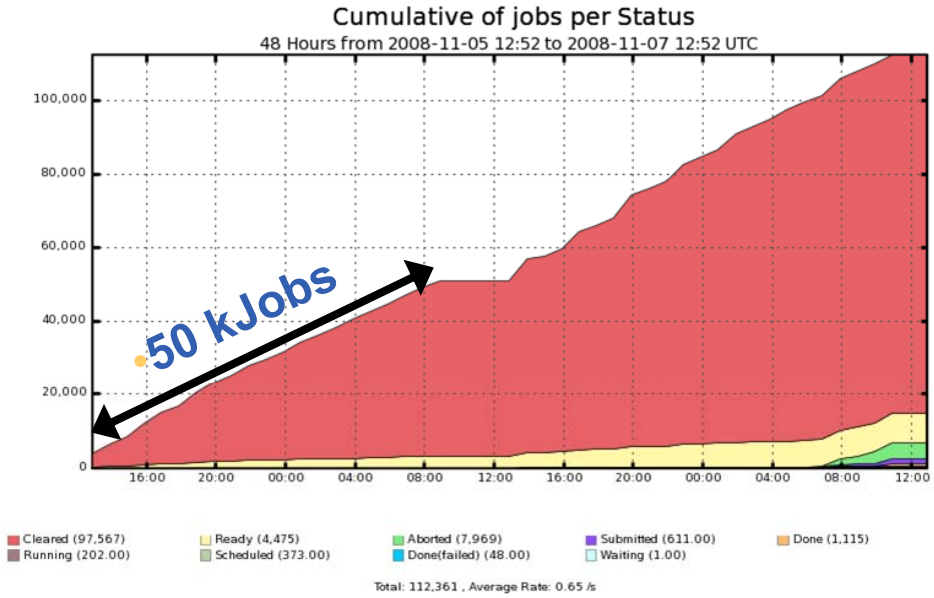

- **Single CRAB Server instance in multi use mode reached 50KJobs per day using 2 WMS**
- **A single ProdAgent instance reached around 30kJobs per day**
  - Lower performance in the output copy from the WMS: plan to reduce the size and number of the files to be retrieved

- **Reached limits are mainly due to tracking and output retrieval/handling**
- **Some optimizations already in place, other small tweaks possible**
- **The WMS architecture is such that the system scales linearly with the number of WMSs**
  - Add as many WMSs to a CMS service as needed
- **The CMS architecture is similar:**
  - Deploy as many instances of PA and CRAB Server as needed:
- **No scale problems foreseen at the expected rates**
  - 50/100 kJobs/day for production and 100/200 kJobs/day for analysis

Job distribution per activity

From May 2008 to March 2009 : 23M total jobs submitted

- **about 78% of the total analysis jobs with gLite WMS (the rest mainly CondorG)**
- **~ 600 distinct users over the last 3 months**

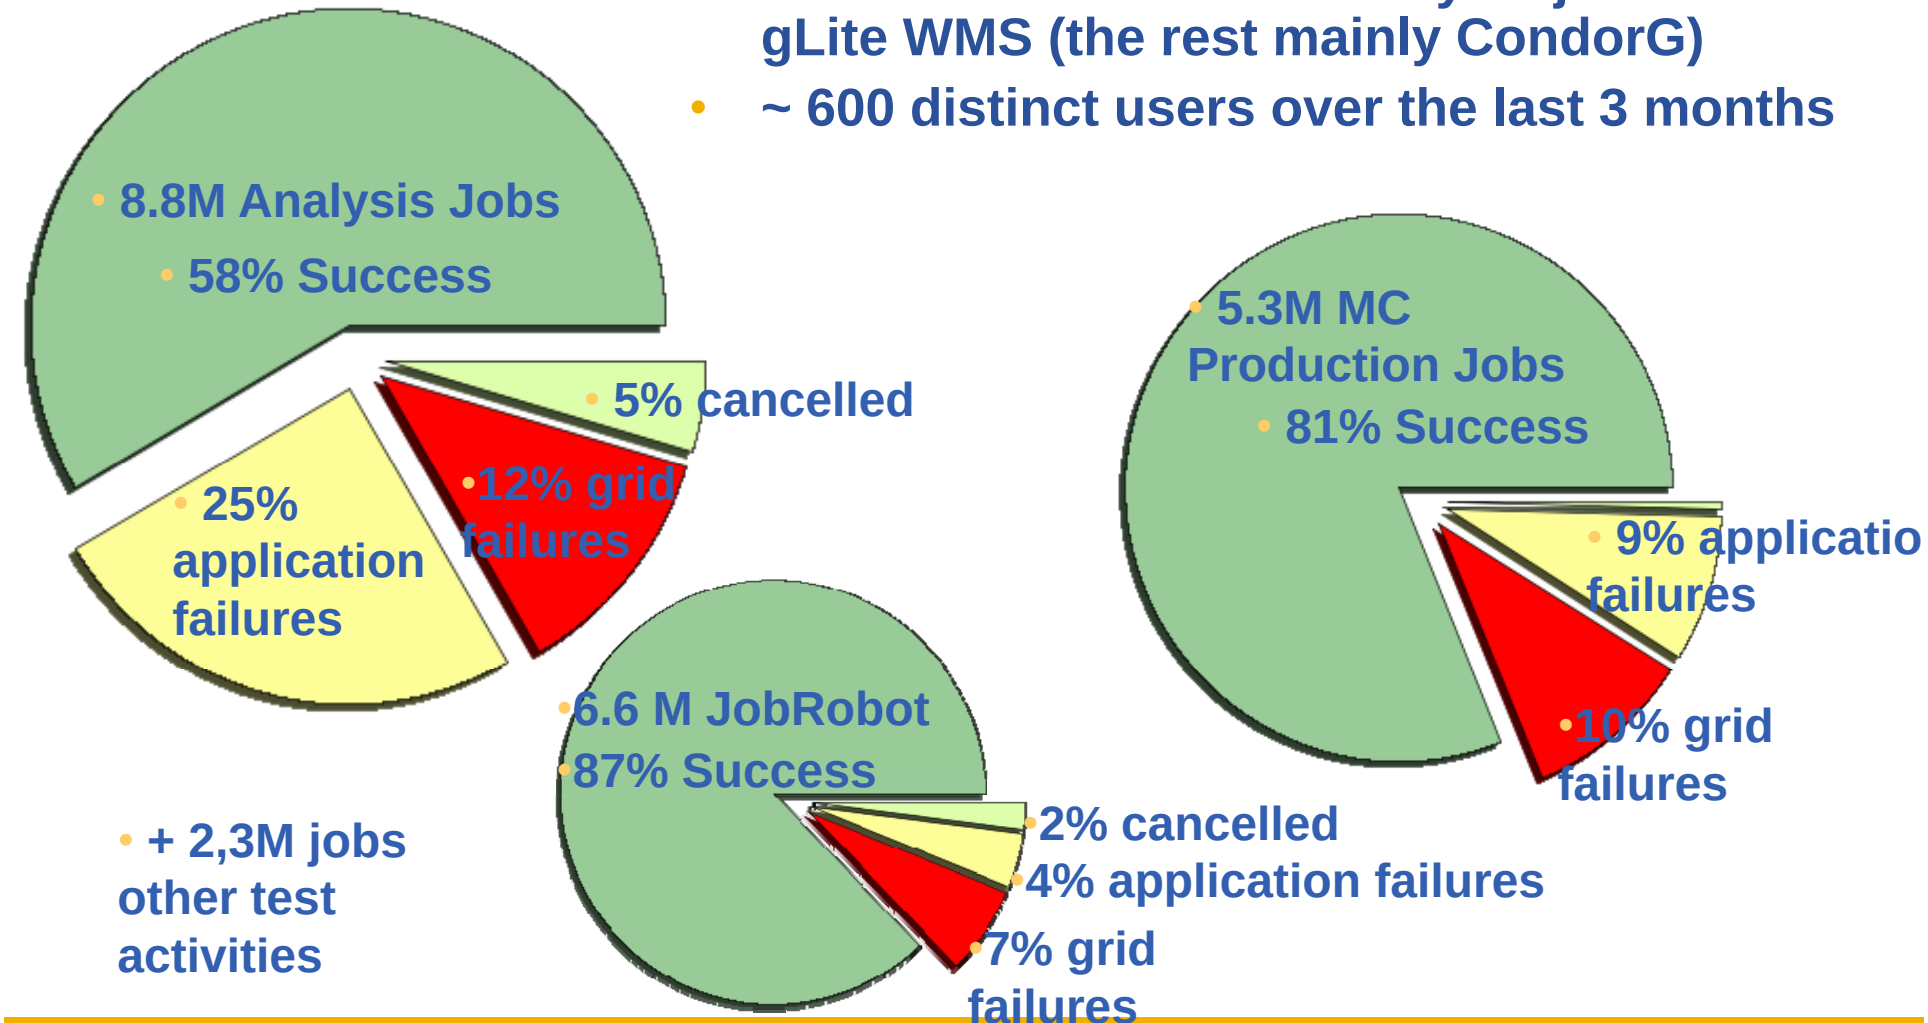

Analysis support task force

- **Improve reliability of analysis jobs**
- **Analysis of failures, user support/feedback, improve monitoring, etc**

dashboard QAOES SEND A COMMENT GET MORE INFO

The page will be refreshed every hour.

Timerange: last 6 hours — Description: jobs in the selected time range are the data set input for the data mining process.
Min number of jobs: 100 — Description: minimum number of jobs, to list a created rule. Low number results in long calculation time!
Confidence: 80 — Description: confidence% of jobs in the input data set that contain the attributes in the antecedent also contain the consequent attribute.

Show Rules

Link to jobSummary	Antecedent	Consequent	Support in % /number of jobs	Confidence in %	Lift
click!	dataset=**/QCDpt30/W...**	ERROR=60303	0.639/179	100.000	31.890
click!	ce=**cmsgrid02.hep.wisc.edu** queue=**cmsgrid02....** dataset=**/QCDpt300/...**	ERROR=60303	0.800/224	98.678	31.468
click!	user=**RachelAWilken** ce=**cmsgrid02.hep.wisc.edu** queue=**cmsgrid02....**	ERROR=60303	2.229/623	98.578	31.436
click!	dataset=**/QCDpt470/...**	ERROR=60303	0.729/203	97.608	31.127
click!	user=**RachelAWilken** site=**T2_US_Wisconsin**	ERROR=60303	2.693/754	96.667	30.827
click!	user=**RachelAWilken**	ERROR=60303	2.707/758	95.466	30.444
click!	dataset=**/QCDpt300/...**	ERROR=60303	1.207/337	94.944	30.277
click!	user=**RachelAWilken** ce=**cmsgrid01.hep.wisc.edu** queue=**cmsgrid01....**	ERROR=60303	0.464/130	88.435	28.202
click!	ce=**cmsgrid01.hep.wisc.edu** queue=**cmsgrid01....** dataset=**/QCDpt300/...**	ERROR=60303	0.407/113	88.372	28.181
click!	user=**Asciaba** site=**T2_RU_SINP** dataset=**/QCD_pt_0_...**	ERROR=8001	0.839/235	81.597	10.711
click!	site=**T2_RU_SINP**	ERROR=8001	0.839/235	81.034	10.637

Workload Management summary

- **CMS successfully uses the gliteWMS in push model for job submission in Monte Carlo Production and Analysis tasks**
  - More than 30k jobs with a single gliteWMS
  - Up to 50k jobs from a single CRAB server using parallel gliteWMS servers
- **Positive feedback cycle with gliteWMS developers**
- **CMS also uses glideinWMS in pull model**

Calibration and alignment

- **Definition**
  - It is the process of applying, during event reconstruction, corrections to the measured quantities to eliminate any bias
  - Some correction factors are determined online, others offline from analysis of some amounts of data
    - typically, a reprocessing is needed to account for all corrections
  - Calibration data is “non-event” data and is needed for event reconstruction and MC generation

Calibration Use Cases

- **Online calibration data is transferred from online database (at experiment site) to offline database (at CERN IT)**
  - based on Oracle Streaming
- **Offline calibration data is produced at the CMS CERN Analysis Facility (CAF)**
- **Calibration data is distributed to all Tier-1 and Tier-2**
- **Calibration data is read-only**

Frontier

- **A system to cache requests for calibration data at remote sites using HTTP as transport protocol and a standard proxy server caching mechanism**
- **Multi-tiered system**
  - Tier 0
    - Central offline Oracle database
    - Frontier server (based on Tomcat)
    - ≥ 1 Squid servers (HTTP proxy servers)
  - Tier 1
    - Squid servers cache data from T0
    - Two redundant servers for high availability
  - Tier 2
    - Squid servers cache data from T1
    - One server sufficient

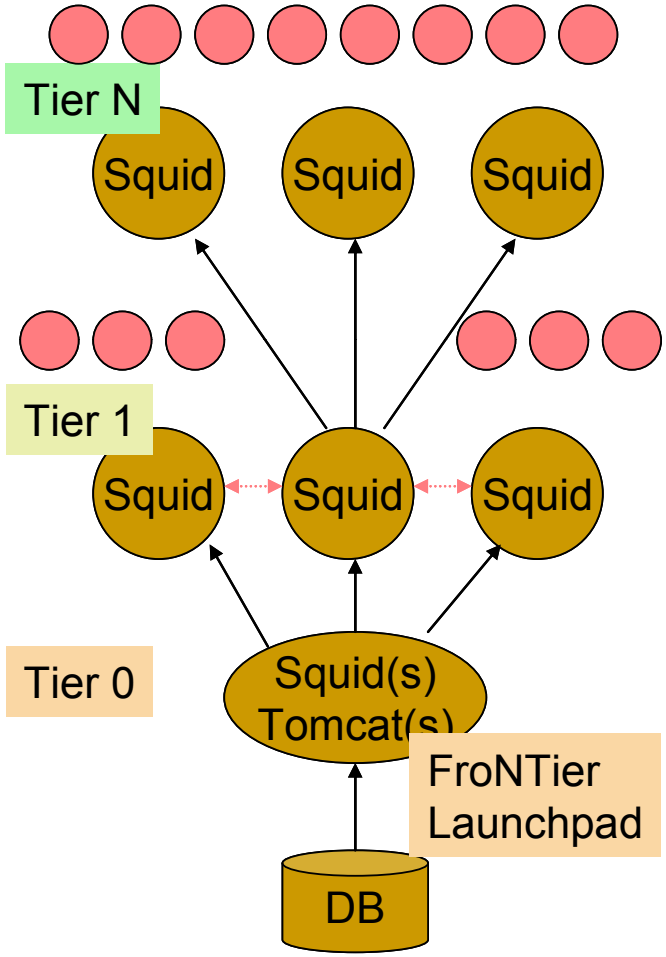

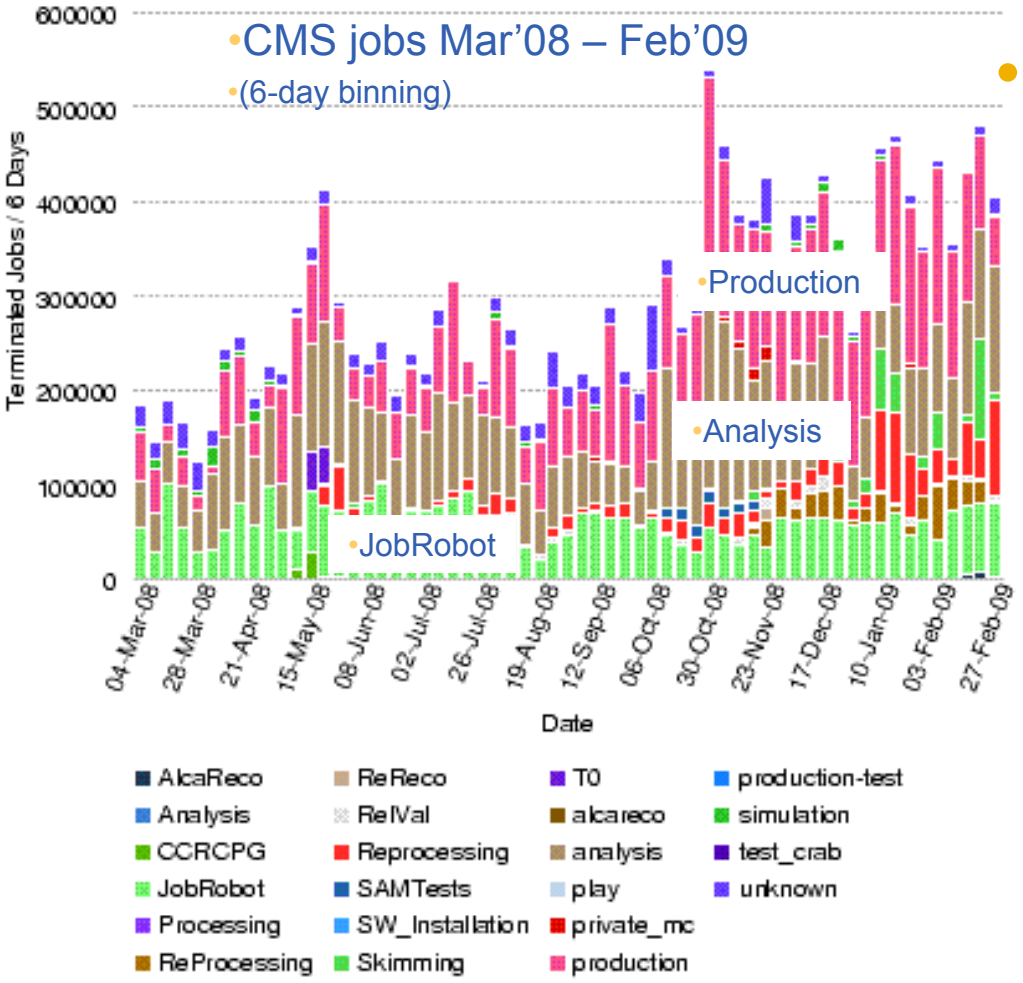

- ***Dashboard job monitoring***

- **Dashboard**
  - Main entry point to a complete view of the CMS Grid activities
    - Job monitoring
      - *Now also real-time with WMCore API*
    - Site monitoring
      - *SAM*
      - *SiteStatusBoard*
    - …
  - Used in daily shifts

Site Commissioning: SAM tests

- **Site Availability Monitoring – CMS SAM tests**
  - High priority jobs submitted every hour
  - Test CE, SRMv2, experiment software, conditions cache, data read, stage out, etc.

Sitename	Service Type	Service Name	mc	js	swinst	squid	analysis	basic	frontier	jsprod	lcg-cp	get-pfn-from-tfc
T1_DE_FZK	CE	ce-1-fzk.gridka.de	ok	ok	ok	ok	ok	ok	ok	ok		
		ce-2-fzk.gridka.de	ok	ok	ok	ok	ok	ok	ok	ok		
		ce-3-fzk.gridka.de	ok	ok	ok	ok	ok	ok	ok	ok		
		ce-4-fzk.gridka.de	ok	ok	ok	ok	ok	ok	ok	ok		
		ce-5-fzk.gridka.de	ok	ok	ok	ok	ok	ok	ok	ok		
	SRMv2	gridka-dCache.fzk.de									ok	ok
T1_ES_PIC	CE	ce-test.pic.es		ok	info	ok	ok	ok	ok	error		
		ce05.pic.es	ok	ok	ok	ok	ok	ok	ok	ok		
		ce06.pic.es	ok	ok	ok	ok	ok	ok	ok	ok		
		ce07.pic.es	ok	ok	ok	ok	ok	ok	ok	ok		
	SRMv2	srmcms.pic.es									ok	ok
T1_FR_CCIN2P3	CE	cclcgceli03.in2p3.fr	ok	ok	ok	ok	ok	ok	ok	ok		
		cclcgceli04.in2p3.fr	ok	ok	ok	ok	ok	ok	ok	ok		
	SRMv2	ccsrm.in2p3.fr									ok	ok
T1_IT_CNAF	CE	ce04-lcg.cr.cnaf.infn.it	ok	ok	ok	ok	ok	ok	ok	ok		
		ce05-lcg.cr.cnaf.infn.it	ok	ok	ok	ok	ok	ok	ok	ok		
		ce06-lcg.cr.cnaf.infn.it	ok	ok	ok	ok	ok	ok	ok	ok		
	SRMv2	srm-v2-cms.cr.cnaf.infn.it									ok	ok
		storm-fe-cms.cr.cnaf.infn.it									ok	ok
T1_TW_ASGC	CE	lcg00125.grid.sinica.edu.tw	warn	ok	ok	ok	ok	ok	ok	ok		
		w-ce01.grid.sinica.edu.tw	warn	ok	ok	ok	ok	ok	ok	ok		
		w-ce02.grid.sinica.edu.tw	warn	ok	ok	ok	ok	ok	ok	ok		
		w-ce03.grid.sinica.edu.tw	warn	ok	warn	ok	error	ok	ok	ok		
	SRMv2	srm2.grid.sinica.edu.tw									error	ok
T1_UK_RAL	CE	lcgce03.gridpp.rl.ac.uk	ok	ok	ok	ok	ok	ok	ok	ok		
		lcgce04.gridpp.rl.ac.uk	ok	ok	ok	ok	ok	ok	ok	ok		
	SRMv2	srm-cms.gridpp.rl.ac.uk									ok	ok
T1_US_FNAL	CE	cmsosgce.fnal.gov	ok	ok	ok	ok	ok	ok	ok	ok		
		cmsosgce2.fnal.gov	ok	ok	ok	ok	ok	note	ok	ok		
		cmsosgce4.fnal.gov	ok	ok	ok	ok	ok	note	ok	ok		
	SRMv2	cmssrm.fnal.gov									ok	ok

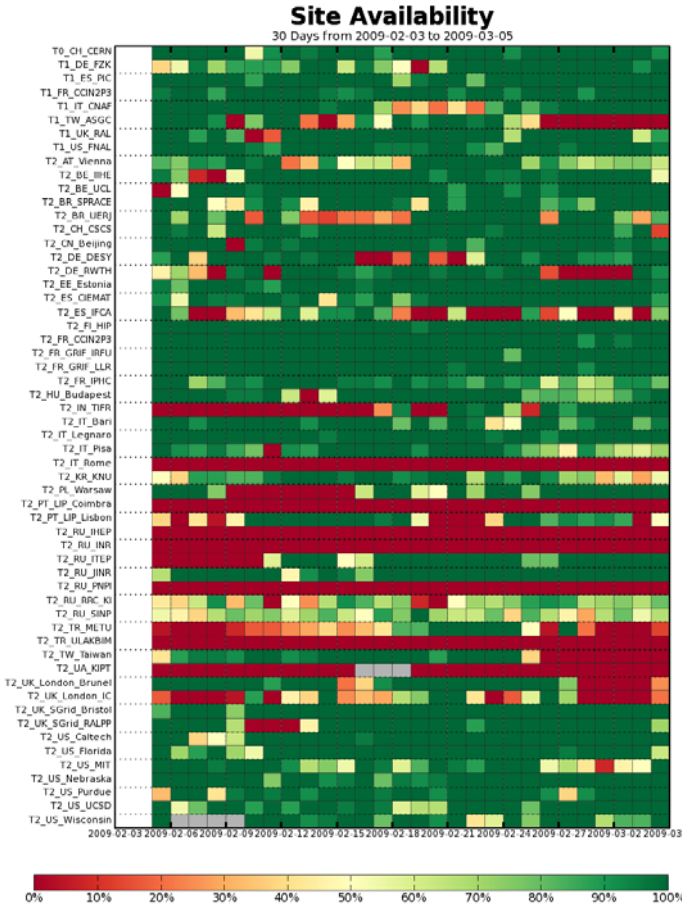

Site Commissioning: JobRobot

- **Job Robot load generator**
  - Tool for automatic job preparation, submission, collection, evaluation
  - Few hundred jobs/site/day to more than 50 sites (~15k jobs/day)
  - Simple jobs reading data
  - Two operations modes
    - Monitoring: constant low rate jobs submission
    - Stress: filling sites with jobs

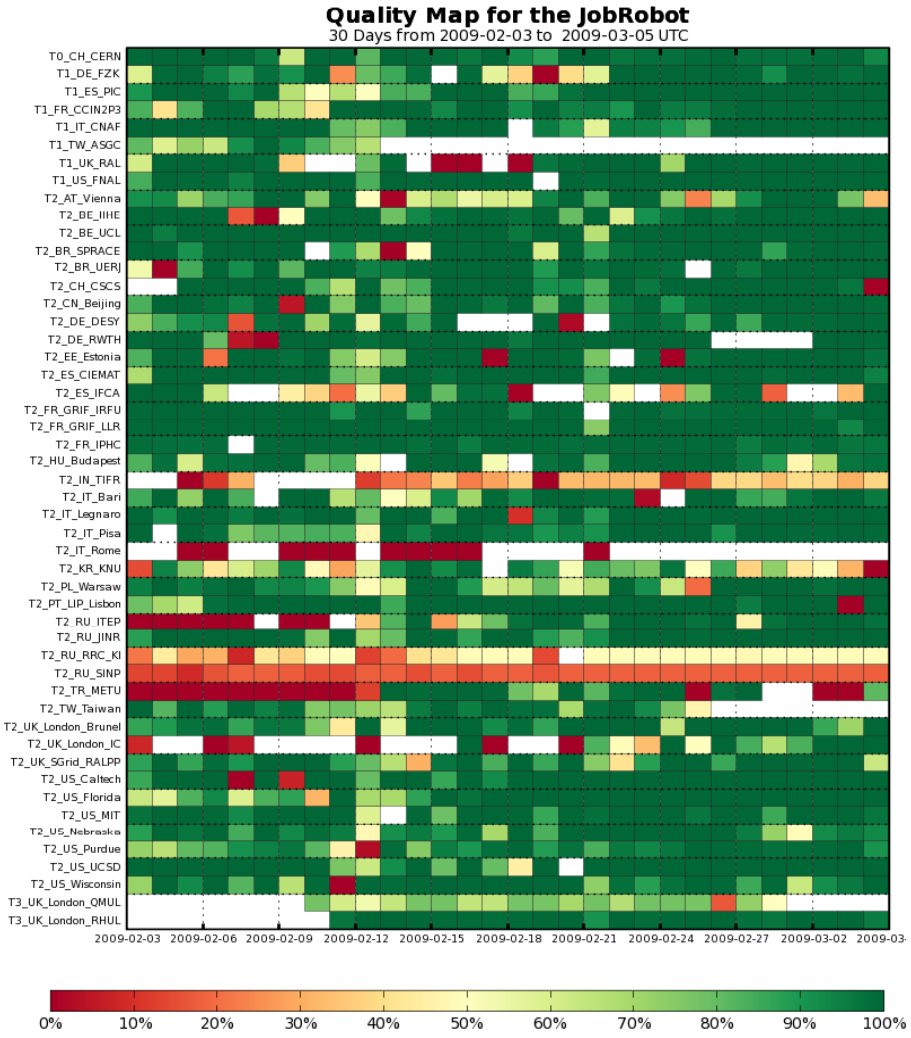

Site Commissioning: aggregation

Enabling Grids for E-sciencE

- **Collect and display all site commissioning information in Dashboard Site Status Board**
- **Combine all metrics into single daily 'site readiness status'**
  - Ready, warning, not-Ready, Scheduled Downtime
- **Improvement seen over the last 6 months**

CMS dashboard — **Site Status for the CMS sites** — Found a bug? HELP

Index	**Expanded Table**	Gridmap	Alternative views

Put the mouse over any column header to get the description of the column
Clicking on a column header will display the evolution of that column over the last 24 hours
: information is more than 24h old

Site Name	SiteComm JR	Commissioned Links (expand this column)	Site availability	SiteReadiness Status	Maintenance (expand this column)	Good links
T0_CH_CERN	98%(699)	n/a	100%	n/a	GOCDB	n/a
T1_CH_CERN	n/a	combined	100%	n/a	n/a	yes
T1_DE_FZK	91%(594)	combined	100%	R	GOCDB	yes
T1_ES_PIC	100%(600)	combined	100%	R	GOCDB	yes
T1_FR_CCIN2P3	99%(600)	combined	80%	W	GOCDB	no
T1_IT_CNAF	100%(600)	combined	100%	R	GOCDB	yes
T1_TW_ASGC	n/a	combined	0%	NR	GOCDB	n/a
T1_UK_RAL	100%(600)	combined	100%	R	GOCDB	yes
T1_US_FNAL	100%(600)	combined	100%	R	OIM	yes
T2_AT_Vienna	99%(600)	combined	64%	NR	GOCDB	yes
T2_BE_IIHE	100%(600)	combined	100%	R	GOCDB	no
T2_BE_UCL	100%(600)	combined	100%	R	GOCDB	yes
T2_BR_SPRACE	67%(392)	combined	76%	R	OIM	no
T2_BR_UERJ	100%(475)	combined	100%	R	OIM	no
T2_CH_CAF	n/a	combined	n/a	R	n/a	n/a

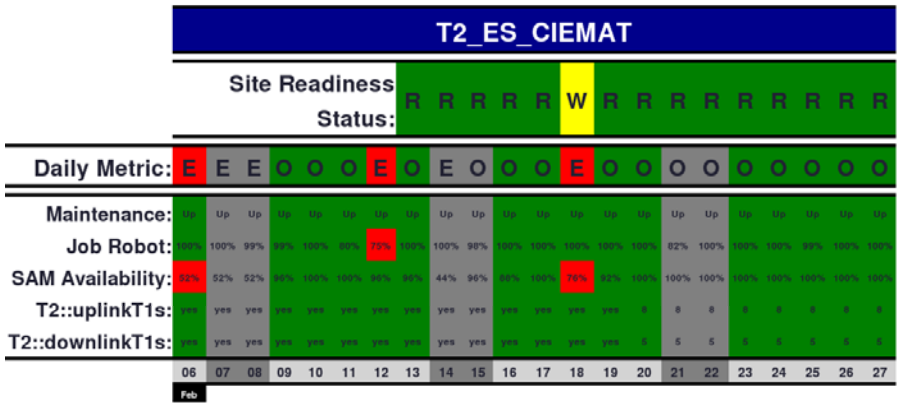

Conclusions

- **Data Management**
  - Using CMS-specific systems for bookkeeping
  - PhEDEx integrated with FTS and SRMv2
  - Transfer debugging and operations need dedicated effort
- **Workload Management**
  - Using CMS-specific systems for task management
  - Both CRAB and PRODAGENT integrated with gliteWMS
  - Also integrated with glideinWMS for pilot jobs
  - Need to improve job efficiency especially for analysis
- **Monitoring and commissioning**
  - Dashboard and SAM heavily used for daily monitoring and for site commissioning
  - Commissioning activity results in constant improvement in site reliability

Acknowledgements

- **Material provided by:**
  - Andrea Sciaba’
  - Daniele Bonacorsi
  - Vincenzo Miccio
  - Giuseppe Codispoti
  - Jose Hernandez Calama
  - James Letts