ATLAS Cloud Operations

Outline (1/3)

- Main ADC activities. Short introduction for:
  - Simulation and reprocessing (ProdSys) [Xavi]
  - Distributed Analysis [Xavi]
  - Distributed Data Management [Ikuo]
- Central operations and clouds relationship [Ikuo]
  - Shifts teams: ADCoS, DAST, Point-1
    - Communication channels: mailing lists, regular meetings, VCR
  - Cloud squads (ATLAS cloud support)
    - Interventions in case of cloud-related issues
      - i.e. Drain and hold ProDsys and DDM
  - Site interventions:
    - Announcing and preparing downtimes and resuming operations
    - Interventions for simulation:
      - drain queues in Panda and prevent brokering
    - Interventions for data transfers:
      - FTS channel playing for Storage/Batch downtimes
  - Site communications
    - When problem found, how to communicate with central operations team and shifts teams

Outline (2/3)

- Monitoring your site
  - DDM [Ikuo]
    - Basic monitoring of the activities
    - Disk used and space expectations
    - Data transfers situation and efficiencies
      - Simulation, Functional Tests, DAQ runs
    - Data integrity:
      - Searching for dark data, lost data and LFC integrity
    - Monitoring ATLAS deletion service
  - Simulation and data processing [Xavi]
    - Basic monitoring of the activities
    - Disentangling simul from repro
    - Single tasks progress and spotting blocked tasks/jobs
  - Centralized tests and critical services [Ikuo]
    - wLCG/ATLAS "Big brothers"
      - wLCG SAM tests
      - ATLAS-VO specific SAM tests
  - Local "critical" ATLAS services: [Ikuo]
    - FTS, LFC, SiteServices
- Pilot factory [Eric]

Outline (3/3)

- Monitoring Distributed Analysis [Xavi]
  - Gangarobot and HammerCloud interfaces
    - test jobs and hammercloud status
  - Guaranteeing User Analysis health:
    - Boosting and tuning your site for distributed analysis in HC interface
      - [WMS, Panda]x[streaming, local_copy]
    - Scheduling and configuring HC tests
  - Panda interface for UA
    - Status and progress of panda-based UA jobs
  - ATLAS SW releases: status and management
  - CondDB access [Xavi]
    - Short intro to new Frontier/squids schema
- Criticality of ATLAS "services" [Xavi]
  - "Refreshing" the wLCG commitments/SLAs for T1 and T2
- Site exclusion policy [Xavi]
  - Metrics
  - (near future) Dashboard SSB

ATLAS production system

ATLAS simulated events production

- Due to LHC's high SNR (10^{-9}) ATLAS needs a huge amount of simulated data
  - Currently running about 40k concurrent jobs
    - Ending about 150k jobs per day involving about 50 different sites and a huge number of different tasks

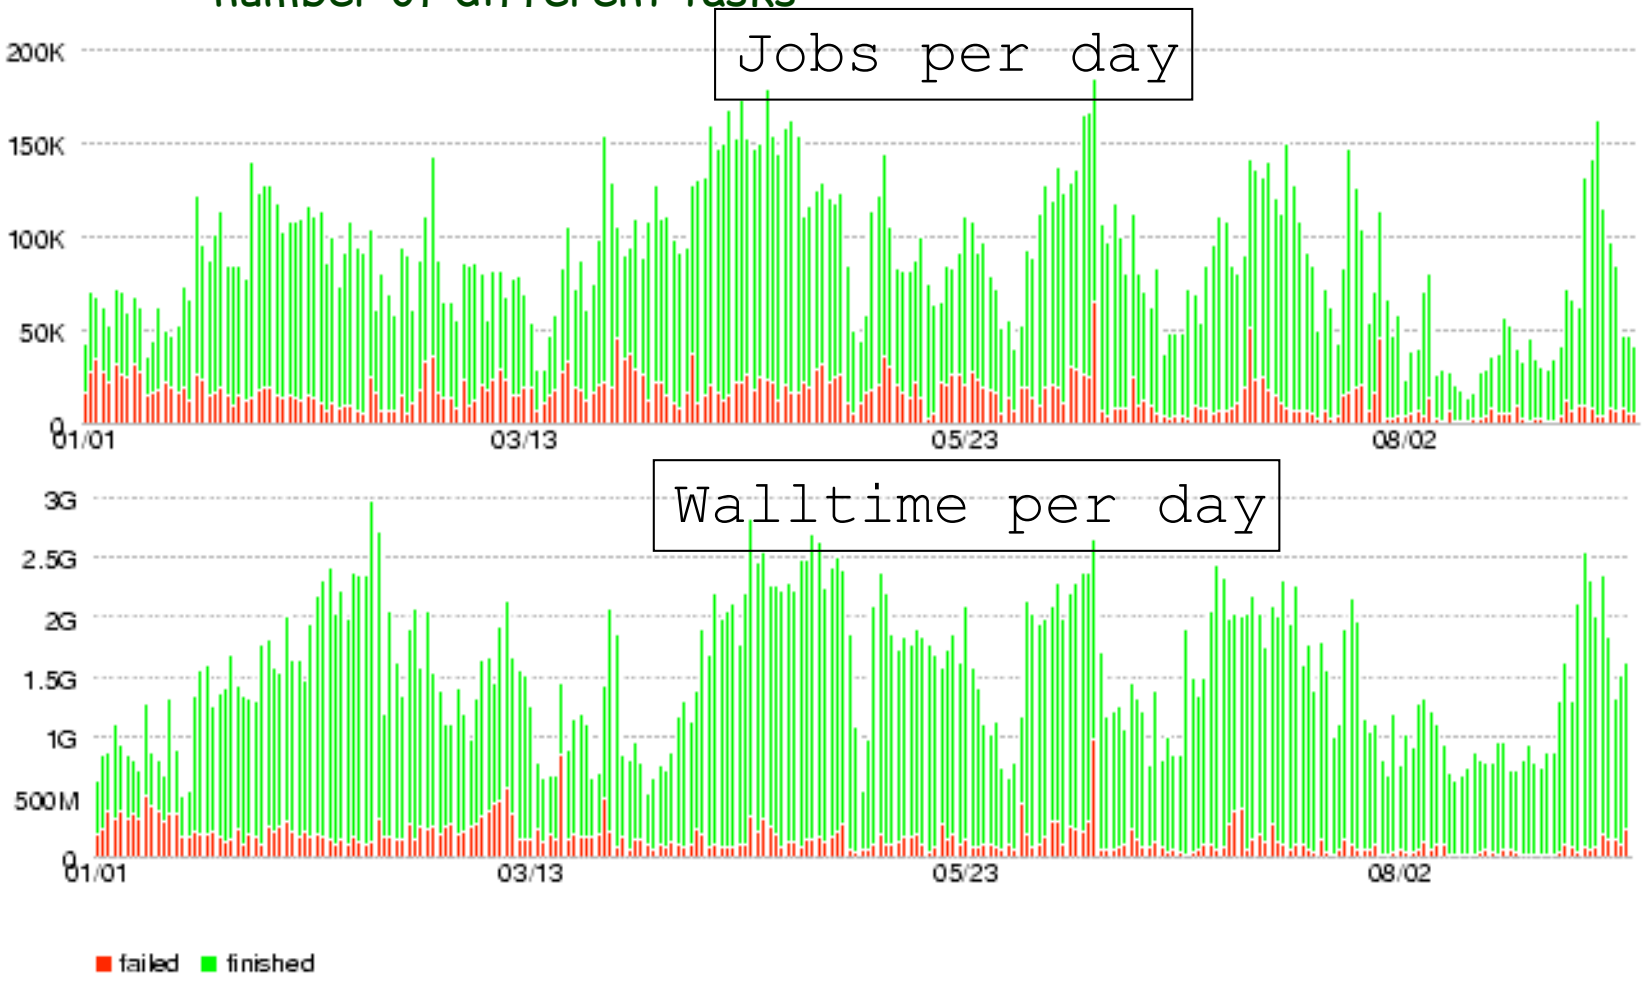

ATLAS ProdSys: workflow

One single Production Database (ProdDB)
and PanDA server at CERN

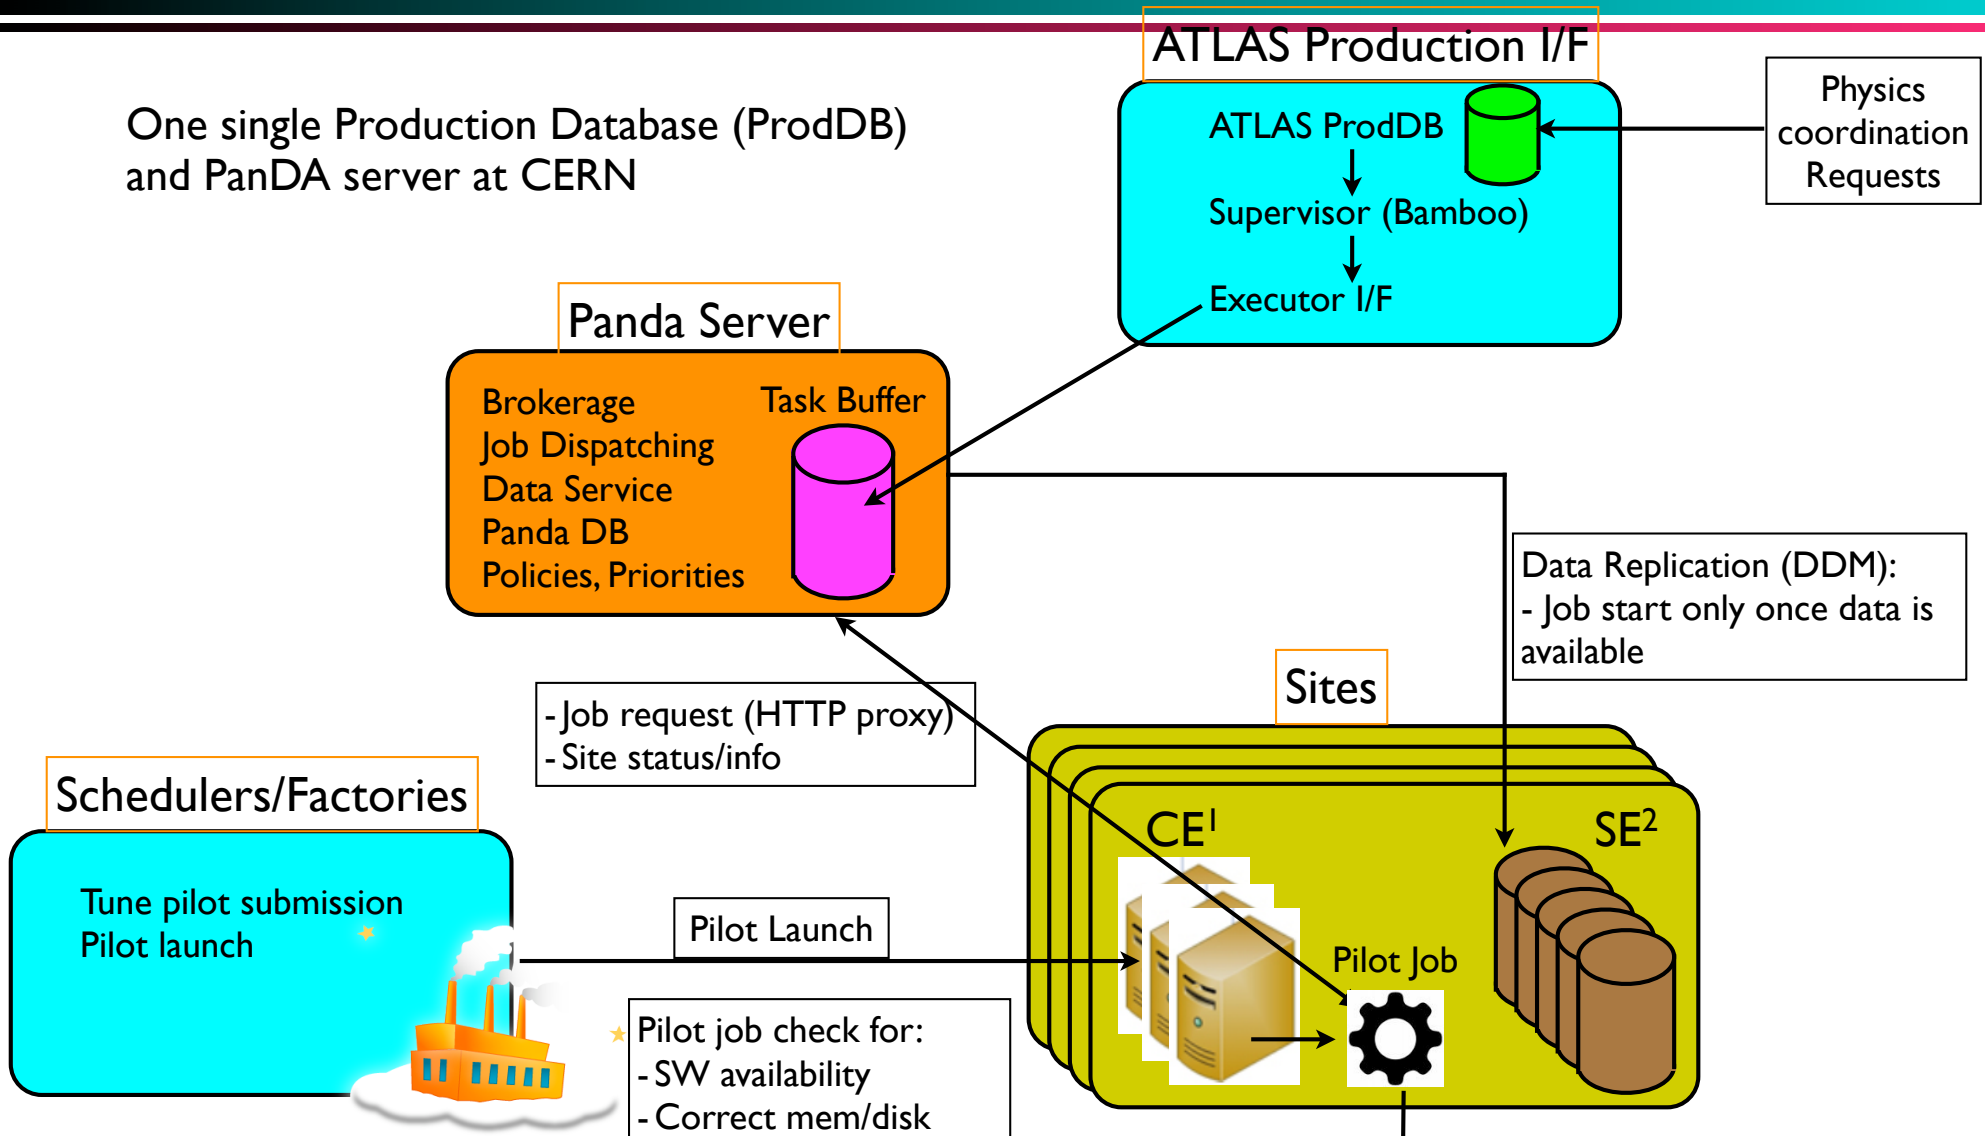

[1]CE: Computing Element
[2]SE: Storage Element

ATLAS ProdSys: Bamboo

- Is the Brain of the Production System
  - Unique process directly interacting with ProdDB
- Bamboo does:
  - Select free jobs from the ProdDB
  - Pass the jobs to PanDA (who take care of submission)
  - Ask about the job status
  - Validate finished jobs
    - Job finishing OK: job done. Otherwise: released again for next attempts
  - Update job entries in the ProdDB

ATLAS ProdSys: Brokerage

- Brokerage:
  - Assign jobs to sites:
    - T1s: defined -> activated (data already present), T2s: defined -> assigned (data subscribed)
  - Calculate weights for brokering using:
    - The number of WNs running jobs for last 3 hours
    - The number of WNs which requested jobs for last 3 hours
    - The number of jobs per WN
    - The number of assigned/activated jobs
    - The number of input files already available at the site
    - Available ATLAS releases
  - Jobs tend to go to the site which has large weight (RW):
    - RW=(#active WNs) x (# jobs per WN) / (# waiting jobs)
- Dispatch/destination datasets
  - Temporary DS created at the time of job submission (20 jobs or 20 files)
    - Dispatch DS: _disXYZ, dispatch input files to T2s, frozen when created and DQ2 transfer files and send callbacks to activate the jobs
    - Destination DS: _subXYZ. Transfer output files to the T1s (holding)

ATLAS ProdSys: Dispatch and Destination Datasets

- Temporary datasets created when jobs are submitted
  - Typically one dispatch/destination dataset per 20 jobs or 20 files
- Dispatch datasets
  - _disXYZ
  - Dispatch input files to T2
  - Get frozen when they are created
  - DQ2 or PandaMover transfers files and then sends callbacks to activate jobs
- Destination datasets
  - _subXYZ
  - transfer output files to T1
  - Empty at beginning. Files are added when jobs are finished
    - For T1: holding → finished/failed
    - For T2: holding → transferring
- Subscription is made when the first file is added
- Get frozen when all jobs contributing to the dataset are finished/failed
- DQ2 transfers files and sends callbacks
  - transferring → finished/failed
  - Panda scans LFs for transferring jobs every day and change job status if all output files are available at T1

ATLAS ProdSys: Pilots and autopilots (1/2)

- Autopilot is a scheduler to submit pilots to sites via condor-g/glidein
  - Pilot → Gatekeeper
  - Job → Panda server
- Pilots are scheduled to the site batch system and pull jobs as soon as CPUs become available
  - Panda server → Job → Pilot
- Pilot submission and Job submission are different
  - Job = payload for pilot
- How pilot works
  - Sends the several parameters to Panda server for job matching (HTTP request)
    - CPU speed
    - Available memory size on the WN
    - List of available ATLAS releases at the site
- Retrieves an `activated' job (HTTP response of the above request)
  - activated → running
- Runs the job immediately because all input files should be already available at the site
- Sends heartbeat every 30min
  - Each heartbeat is a single HTTPS session
  - There isn't a permanent connection between pilot and Panda server
  - If pilot dies silently, panda will set the job status to 'holding' 6 hours later

ATLAS ProdSys: Pilots and autopilots (2/2)

- Sends jobStatus='finished'/'failed' at the end of the job
  - Copy output files to SE
  - Register files to LFC
  - running → holding
- Pilot itself doesn't access DQ2
  - Panda server adds output files to DQ2 datasets
  - holding → transferring/finished/failed
- Sends jobStatus = 'holding' if the pilot cannot copy output files to SE or cannot register files to LFC
  - running → holding
  - Then the pilot terminates immediately
  - Another pilot will try to recover the job if the site supports

Distributed Analysis

Distributed Analysis: data replication and distribution

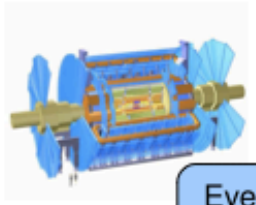

Event Filter

RAW

Tier 0

RAW/ AOD/ESD

Reprocessing of full data, Managed Tape
Access: RAW, ESD, AOD
Disk Access: AOD, parts ESD

Production of MC events
User Analysis
Disk Storage: AOD

CERN Analysis Facility

1st pass calibrations, Reconstruction, 24h Data export

Tier 1

10 Sites

AOD

Tier 2

~40 sites

+ many Tier3

Limited Access:
ESD, RAW, calibration
Analysis

Johannes Elmsheuser (LMU München) Distributed Analysis in ATLAS using GANGA 24/03/2009 2 / 17

Distributed Analysis: event data model

Refining the data by: Add higher level info, Skin, Thin, Slim

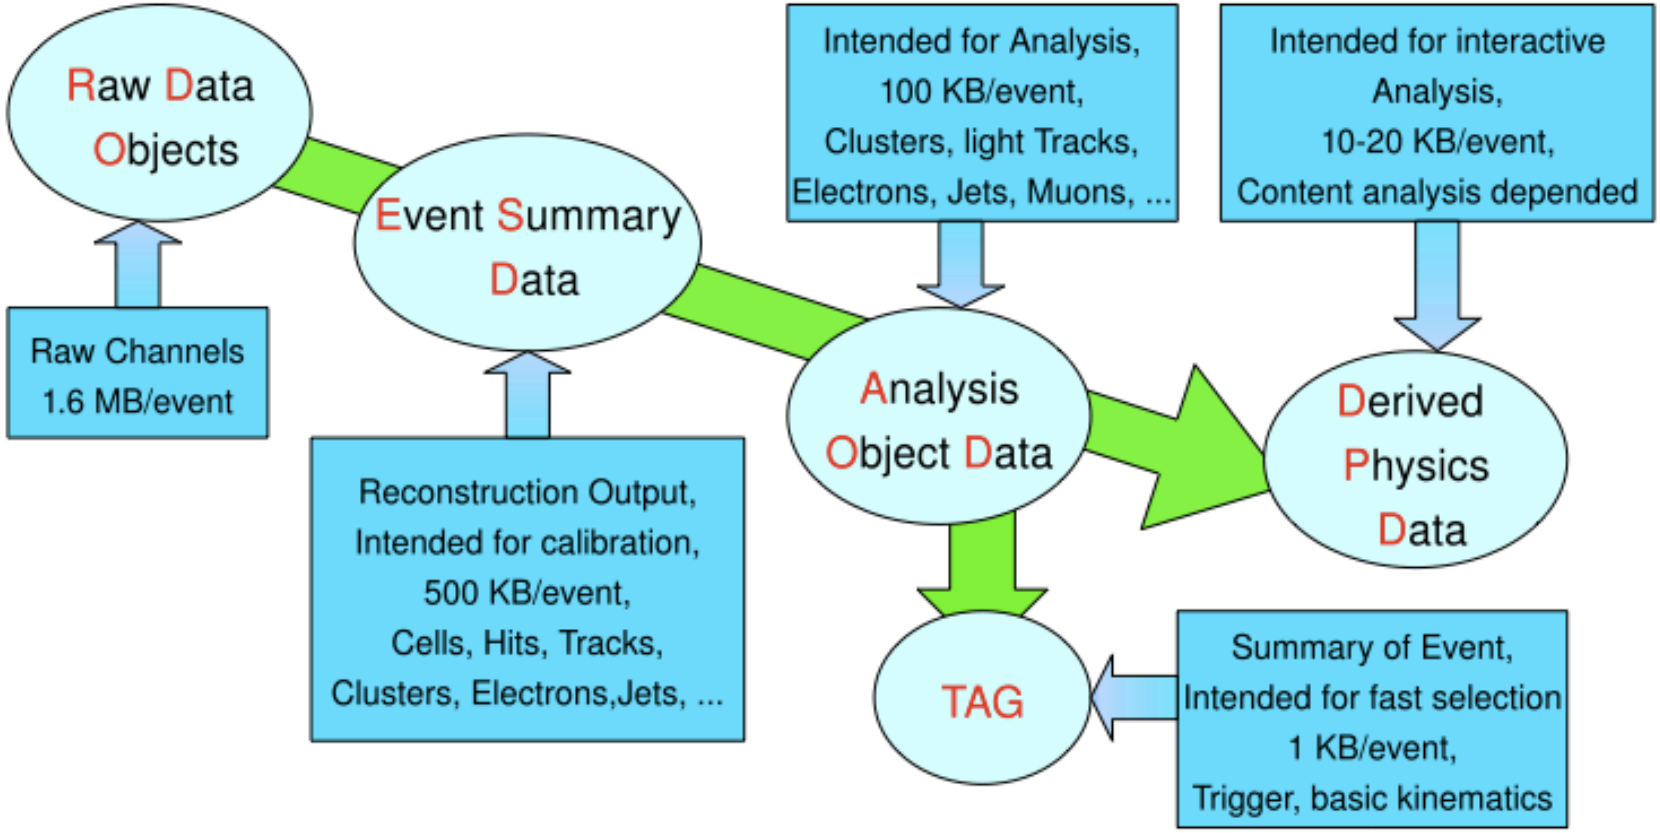

Distributed Analysis model

The distributed analysis model is based the ATLAS computing model

- Data is distributed to Tier1 and Tier2 facilities by default by the ATLAS Data Distribution system DQ2
  - available 24/7
  - Automated file management, distribution and archiving throughout the whole grid using a Central Catalogue, FTS, LFCs
  - Random access needs a pre-filtering of data of interest, e.g. Trigger or ID streams or TAGs (event-level meta data)
- user jobs are sent to the data
  large input data-sets (several TBs)
- Results must be made available to the user
  potentially already during processing
- Data is added with meta-data and bookkeeping in catalogues

Distributed Analysis workflows

The distributed analysis model is based the ATLAS computing model

- Data is distributed to Tier1 and Tier2 facilities by default by the ATLAS Data Distribution system DQ2
  - available 24/7
  - Automated file management, distribution and archiving throughout the whole grid using a Central Catalogue, FTS, LFCs
  - Random access needs a pre-filtering of data of interest, e.g. Trigger or ID streams or TAGs (event-level meta data)
- user jobs are sent to the data
  large input data-sets (several TBs)
- Results must be made available to the user
  potentially already during processing
- Data is added with meta-data and bookkeeping in catalogues

Distributed Analysis

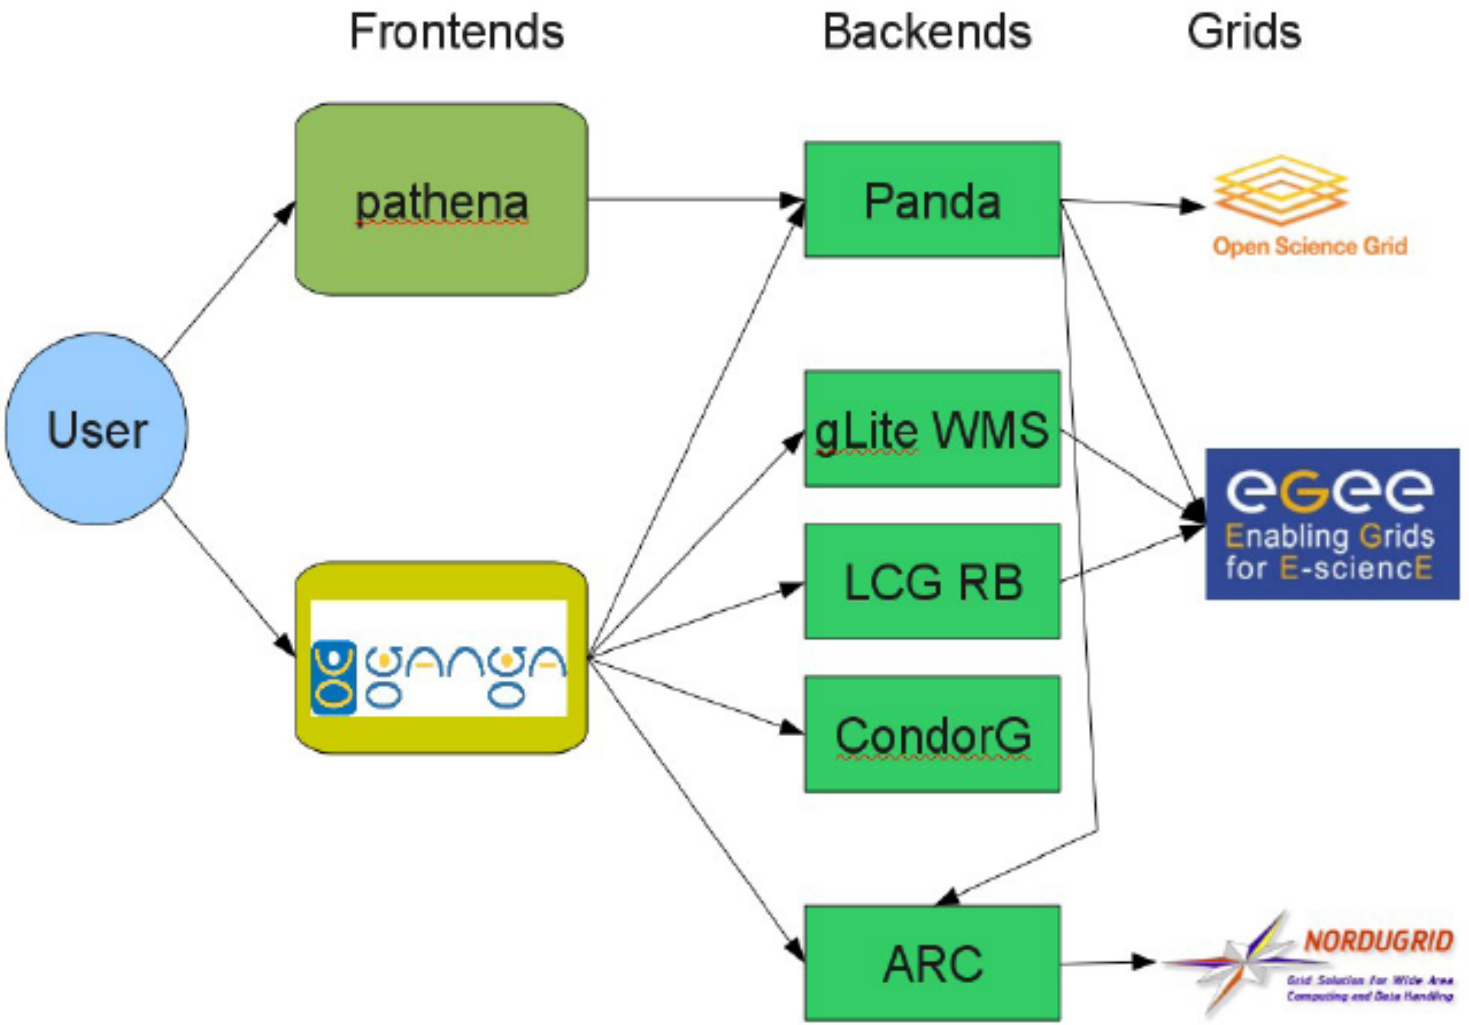

3-layers:

User interface

Gateways to resources

Execution infrastructure

Distributed Data Management (DDM)

I.Ueda

DDM Concepts

- Dataset
  - Contains Files (GUID, LFN)
  - Has Replica Locations (the sites where replicas reside)
- File
  - Entry in a Dataset (GUID, LFN)
  - Physical entity at sites (SURL, TURL)
- Dataset Container
  - Contains Datasets
  - No Replica/Location information in itself
- Subscription
  - Request for dataset placement

DDM Components

- DQ2
  - Central Catalogs (dataset and replica location)
  - Site Services - to process subscriptions
- Subscription Tools
  - to tell DQ2 which datasets should be replicated to which sites
- FTS - for file transfers
  - One per cloud (at T1) for to-cloud and within-cloud transfers
  - Two at CERN
    - T0-T1 exports and transfers to CERN
    - T0-calibration_sites exports
- LFC - as the file catalog (GUID,LFN,SURL)
  - One per cloud (at T1) for files in the cloud
  - One at CERN (for files at CERN)

DDM Workflow

- Subscriptions are created by various tools
  - T0 export, T1-T1, T1-T2, production input/output, user request
- A Site Service finds a subscription to a site which it serves
  - There are several site services, basically per cloud, but some clouds are treated together
- Resolves the dataset contents and dataset replica locations
  - central catalogs
- Finds which files are not in the destinatio site
  - the cloud LFC
- Choose a source site from the available replica locations
- Submit a transfer job to the proper FTS
- Poll the transfer status, and verify when it is done
- Register the file at the destination to the LFC

Central operations and clouds

I.Ueda

Overview

- ADC Central Operations Team
  - A group of experts of various components of ADC
- ADC Expert On-Call (from the operations team)
  - Main responsible for the intervention
  - Interface between shifters and the experts
- ADC Shifts
  - Watch the monitors (and do some instructed intervention)
  - report to the expert
  - notify the sites (ggus) and the cloud squad (email, savannah)
- Cloud Squad
  - Treat cloud issues and ATLAS-specific issues at the sites
  - Interface between the sites and the central operations
- Sites
  - Treat site issues, may consult cloud squad for ATLAS-specific issues

ADC Shifts Teams

- ADC@Point-1: Shifts on Atlas Distributed Computing at Point-1
  - Data export from T0
  - Real data distribution (T1-T1, T1-T2) -- some overlap with ADCoS
  - Central Services (DDM) -- some overlap with ADCoS
  - atlddmsh@mail.cern.ch
- ADCoS: Shifts on Atlas Distributed Computing
  - Official production (MC, Reprocessing)
  - Real and MC data distribution (T1-T1, T1-T2) -- some overlap with ADC@Point-1
  - Central Services (DDM, ProdSys) -- some overlap with ADC@Point-1
  - atlas-project-adc-operations-shifts@cern.ch
- DAST: Shifts on Distributed Analysis
  - User Analysis on GRID
  - User Data Access, User Data Replication
  - atlas-project-adc-operations-analysis-shifts@cern.ch

Cloud Squads

- What is a squad?
  - A cloud squad is a support team for the corresponding cloud in the Atlas Distributed Computing (ADC) : there is one squad team per cloud.
  - a.k.a. ATLAS cloud support
  - email: atlas-support-cloud-<cloud>@cern.ch
    - eg. for FR cloud: atlas-support-cloud-fr@cern.ch
  - savannah 'squad' for ddm-ops
    - eg. for FR cloud: https://savannah.cern.ch/users/dq2-ddm-ops-cloud_fr

Cloud Squads

- What are they supposed to do?
  - Treat the issues within the cloud
    - Cloud-wide issues
    - ATLAS-specific site issues for the sites in the cloud
      - ➡ (ATLAS files, ATLAS jobs, ...)
  - Interface between the sites and the ATLAS central operations
    - Translate ATLAS language into Site/WLCG language
    - Supplement the lack of knowledge in the central operations about the sites and the cloud
  - SquadHowTo twiki to describe all about it
    - under construction...
    - being filled by the FR squad people. Merci!

Communications to Cloud/Site

- ggus ticket to the site with cloud support in CC
  - ADC@Point-1 in every ticket
  - ADCoS in case of missing files, no space on storage
  - DAST?
- email
  - Issues in pilot factory, Cloud brokeroff in panda
- savannah
  - DDM operations (https://savannah.cern.ch/bugs/?group=dq2-ddm-ops)
  - assigned to the cloud squad, not to the site (can be put in CC)
- ELOG (https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/)
  - Not a communication tool, but for the record of actions taken
    - but clouds/sites can have a look
  - Maybe referred from the tickets

Communications from Cloud/Site

- email
  - ADC expert (atlas-adc-expert@cern.ch)
    - in general
    - 14 members, the ADC expert on-call is one of them
  - ADCoS shifters (atlas-project-adc-operations-shifts@cern.ch)
    - for the panda queues treatment
  - ADC operations ML (atlas-project-adc-operations@cern.ch)
    - for wider announcements
    - 155 members (as of 13 Jan 2010, not only the operations team)

- savannah
  - DDM operations (https://savannah.cern.ch/bugs/?group=dq2-ddm-ops)
  - No need for login (but better with it)

- ELOG
  - for the record of actions taken, solutions made (not for conversations)
  - Login needed to write (so probably not for the sites, but for cloud squads)

Reporting errors/problems

- ADC eLog:
  - Site interventions/problems noticed etc.
    - https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/
- ATLAS software:
  - Validation
    - https://savannah.cern.ch/bugs/?group=validation
  - Panda:
    - https://savannah.cern.ch/bugs/?group=panda
- ADCoS mailing list
  - atlas-project-adc-operations-shifts(at)cern.ch
- ADC VCR:
  - Skype chat

Cloud Issues

- LFC down
  - No file registration, No SURL retrieval
    - Basically no jobs can run in the cloud
      - The cloud should be set "brokeroff"
        - done by the ADC-expert
    - No file transfer can happen
      - All the sites in the cloud should be excluded from the DDM
        - done by the ADC-expert (or DDM-operations)
- FTS down
  - No transfers to/within the cloud (except from T0)
    - All the sites should be excluded from the DDM
      - done by the ADC-expert (or DDM-operations)
- In either cases,
  - Contact the ADC-expert so that they can do the work with confidence
  - For a scheduled downtime, to be done beforehand
    - the job and/or transfer queues are to be drained

Cloud Issues

- eg.

  To: <atlas-project-adc-operations@cern.ch>

  Subject: set DE offline

  Please can someone set DE offline(for the downtime including LFC).

  To: <atlas-adc-expert@cern.ch>

  Subject: [Fwd: set DE offline]

  Can you do it please + remove all DE sites from Site Services.

  To: <atlas-adc-expert@cern.ch>

  Subject: DE cloud still offline in PandaMon

  DE cloud should be back from downtime, could you please set it online for production?

  PandaMon shows it still offline

Cloud Issues

- eg.

  To: <atlas-project-adc-operations@cern.ch>

  Cc: <atlas-project-adc-operations-shifts@cern.ch>,

  <atlas-project-adc-operations-analysis-shifts@cern.ch>

  Subject: LFC and FTS database migration - SARA

  tomorrow SARA will have a downtime for moving backend database of LFC and FTS to the oracle rac. The downtime information: https://goc.gridops.org/downtime/list?id=48755447

  Shall we set the whole cloud offline during the downtime?

  Cc: <atlas-project-adc-operations@cern.ch>, <atlas-project-adc-operations-shifts@cern.ch>, <atlas-project-adc-operations-analysis-shifts@cern.ch>

  Subject: Re: LFC and FTS database migration - SARA

  the FTS and LFC services are back.

Site issues

- Production/Analysis queues
  - Sites can treat their queues (drain, offline, online) to prevent job submission
  - ADCoS shifters can set the panda queues "offline", when they see errors
    - then set "test" and send test jobs, when the issue is gone.
    - and then set "online" if the test jobs complete successfully
  - Cloud squad can do the same before the shifters see the errors?

- DDM data transfers
  - Sites cannot prevent the access to the storage (by itself)
  - ADC expert excludes the site from the DDM when they see errors
    - then put it back when the issue is gone
  - Cloud squad or the site can set the FTS channel inactive?
    - so that the transfers do not happen without excluding the site from DDM

Site issues

- Missing files
  - Files registered on LFC, but the physical files do not exist at the site
    - First check the files are really missing at the site
    - If yes, need to clean the LFC catalog entries, and subscribe for the dataset
    - If not, investigate why not accessible
      - temporarily unaccessible (disk server offline, etc.)?
      - accessible with lcg-cp, but not to the job?
    - report to the ADC expert
      - may need further actions depending on how serious the problem is
- Corrupted files
  - Files with checksum different from the value registered in DQ2 and LFC
    - First, check the file at the site is really corrupted or not.
    - Then, check if the file is corrupted at the "original T1"
      - the files produced in a cloud are aggregated to its T1, and then replicated to the other clouds/ sites
    - If corrupted, to be assigned to that cloud, to be investigated there, then removed from the dataset
    - If not, delete the corrupted file with the LFC entry, and then subscribe for the dataset

Example Workflows

- Shifter finds a problem, sends a ticket with cloud-support in CC
  - Site reacts, consulting cloud-support if necessary
  - cloud-support translates the ATLAS language into WLCG/Site language if necessary

- Shifter finds a disk full, submits a savannah ticket
  - the DDM operations assign the ticket to the cloud squad
    - the site will be set offline in DDM
  - the cloud squad asks the site if the space can be increased
    - if yes, the site does the work, and cloud squad respond to the ticket
    - if not, the cloud squad consult the ADC operations to make a space

Example Workflows

- Shifter finds a problem, send a ggus ticket without cloud-support in CC
  - The shifter sets the site offline in the production system
  - Site reacts, solving the problem, closes the ticket (status "solved")
  - Shifters don't notice and the site stays offline
    - they follow "open" tickets, but "solved" ones are often overlooked.
    - the one on shift when the issue is solved is likely different from the one who submitted the ticket.
  - The site sees no jobs (or, finds it offline). what to do?
    - Consult the cloud squad
  - The cloud squad contacts the ADCoS (elog or email), and the shifter sets the site in "test" and sends test jobs
  - Seeing success of test jobs, the shifter set the site "online"

Example Workflows

- A site schedules a downtime, and notifies the cloud squad about it
  - The cloud squad does what are necessary
  - For SE
    - set the FTS channel inactive (if possible)
    - or ask the ADC expert to exclude the site from DDM (if the downtime is longer than ~1 day)
  - For CE
    - make sure that the site starts draining the queue well in advance
    - ask the ADCoS to set the site offline in the prodsys (panda)
      - ➡ or, should/can the cloud squad do it?

What to do to prepare an SD?

- Potential unavailability of services related to simulation production should be announced in advance.
  - The main dependencies are:
    - T1: LFC, Storage
    - T2: batch system, storage
- Cloud squad or ADCoS shifters can set the T2s sites to:
  - brokeroff : stop new jobs from being assigned (draining, used for SD)
    - At least 24h in advance
    - curl -k --cert /tmp/x509up_u`id -u` 'https://panda.cern.ch:25943/server/controller/query?tpmes=setbrokeroff&queue=queuename&moduser=your_name&comment=Site%20Intervention'
  - offline: stop pilots and stop any in the queue from picking up jobs (used when storage is failing, aka uSD)
- ADC expert take care of T1:
  - Critical as LFC holder, if LFC is down, *cloud* should be set offline immediately.
  - Storage: Critical to feed T2s with input data, T1 site offline and cloud brokeroff, T2s can still compute the running jobs and store outputs in PRODDISK buffer.

What to do to prepare a downtime ?

- Coming back after an SD (brokeroff to brokeron):
  - Activate brokering
    - curl -k --cert /tmp/x509up_u`id -u` 'https://panda.cern.ch:25943/server/controller/query?tpmes=setbrokeron&queue=queuename&moduser=your_name&comment='
- Coming back after an uSD (offline to online):
  - Site need to test it's queues in Panda. Set site into test mode:
    - curl -k --cert /tmp/x509up_u`id -u` 'https://panda.cern.ch:25943/server/controller/query?tpmes=settest&queue=queuename&moduser=your_name&comment='
  - Schedule test jobs at your site:
    - https://twiki.cern.ch/twiki/bin/view/Atlas/ADCoS#How_to_send_test_jobs_procedure
  - Check for results in Panda:
    - http://panda.cern.ch:25980/server/pandamon/query?job=*&type=test&hours=3
  - If jobs succeeded, set site online:
    - curl -k --cert /tmp/x509up_u`id -u` 'https://panda.cern.ch:25943/server/controller/query?tpmes=setonline&queue=queuename&moduser=your_name&comment='

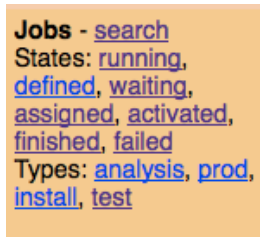

Monitoring tools

for your cloud/site

Monitoring tools

- Why monitor by yourself?

- A shifter finds a problem, and inform you (usually with GGUS)
  - The cloud squad or the site may need to look at the monitoring tools and to understand what the problem is

- A site sees no jobs running, or no data coming in
  - The site or the cloud squad should look at the monitoring tools to understand the situation

Site Status Board

- To summarize all the major monitoring results for sites

- Under Construction…
  - So, please be patient and go through these many monitoring pages to see the site/ cloud status
  - Xavi will show you the first version in the context of the "Site exclusion policy"

Monitoring: Sites Exclusion

- Panda (for production jobs and pathena analysis jobs) : http://panda.cern.ch:25980/server/pandamon/query?dash=clouds
  - Clouds treated only by the ADC experts
  - Sites treated by the ADCoS (and possibly ADC experts, cloud squads)
- Ganga (for ganga analysis jobs) : http://gangarobot.cern.ch/blacklist.html
  - Done automatically
- DDM (for data distribution) :
  - Treated only by the ADC experts
  - Site Services: http://cern.ch/project-atlas-ddm-sls/SS_sites-shares_monitor.txt
  - Functional tests (T1-T1, T1-T2): http://atladcops.cern.ch:8000/drmon/ftmon_TiersInfo.html
  - Real data distribution: http://atladcops.cern.ch:8000/drmon/crmon_TiersInfo.html
  - MC data distribution: http://panda.cern.ch/?mode=listAODReplications&summary=Clouds
- Check the ELOG to see the reason
- Site should be informed with a GGUS ticket for panda queue
- Cloud Squad should be informed with a savannah ticket for DDM

Monitoring: Sites Exclusion

- PANDA http://panda.cern.ch:25980/server/pandamon/query?dash=clouds

ELOG number

IN2P3-LAPP	ANALY_LAPP	ANALY_LAPP	6			online	auto
	IN2P3-LAPP-lapp-ce01-atlas-pbs	LAPP	6			online	auto
IN2P3-LPC	ANALY_LPC	ANALY_LPC	2			online	auto
	IN2P3-LPC-clrlcgce01-atlas-lcgpbs	LPC *https://savannah.cern.ch/bugs/*	2			online	auto
	IN2P3-LPC-clrlcgce02-atlas-lcgpbs	LPC *6495*	2			offline	manual
	IN2P3-LPC-clrlcgce03-atlas-lcgpbs	LPC *https://savannah.cern.ch/bugs/*	2			online	auto
IN2P3-LPSC	ANALY_LPSC	ANALY_LPSC	2			online	auto
	IN2P3-LPSC	IN2P3-LPSC	2			online	auto
	IN2P3-LPSC-lpsc-ce-atlas-pbs	IN2P3-LPSC	2			online	auto
LYON_REPRO	LYON_REPRO	LYON_REPRO	1			online	auto
NIPNE	ANALY_ROMANIA02	ANALY_ROMANIA02	6			online	auto
	ANALY_ROMANIA07	ANALY_ROMANIA07	6			online	auto
	RO-02-NIPNE-tbat01-atlas-lcgpbs	ROMANIA02 *downtime*	6			online	auto
	RO-07-NIPNE-tbit01-atlas-lcgpbs	ROMANIA07	6			online	auto
TOKYO-LCG2	ANALY_TOKYO	ANALY_TOKYO	2			online	auto
	ANALY_TOKYO_RFIO	ANALY_TOKYO_RFIO *Do.Not.Set.ONLINE.Please Eric.Lancon*	2			offline	manual
	TOKYO-LCG2-lcg-ce01-atlas-lcgpbs	TOKYO	2			online	auto

Or, Some comments

Monitoring: Sites Exclusion

- GANGA (for ganga analysis jobs) : http://gangarobot.cern.ch/blacklist.html

Sitest currently blacklisted

```
GRIF-SACLAY_MCDISK
GRIF-LAL_PHYS-SUSY
GRIF-LPNHE_HOTDISK
GRIF-LPNHE_MCDISK
GRIF-LAL_PRODDISK
GRIF-SACLAY_LOCALGROUPDISK
GRIF-LAL_HOTDISK
GRIF-LPNHE_SCRATCHDISK
GRIF-SACLAY_HOTDISK
GRIF-SACLAY_SCRATCHDISK
GRIF-LAL_MCDISK
GRIF-SACLAY_DATADISK
GRIF-LPNHE_LOCALGROUPDISK
GRIF-LPNHE_DATADISK
GRIF-LAL_SCRATCHDISK
GRIF-LPNHE_PRODDISK
GRIF-SACLAY_PHYS-TOP
GRIF-SACLAY_PRODDISK
GRIF-LAL_LOCALGROUPDISK
GRIF-LAL_DATADISK
GRIF-LPNHE_PHYS-SM
RO-02-NIPNE_HOTDISK
RO-02-NIPNE_DATADISK
RO-02-NIPNE_PRODDISK
RO-02-NIPNE_SOFT-TEST
RO-02-NIPNE_MCDISK
RO-02-NIPNE_SCRATCHDISK
ITEP_SCRATCHDISK
ITEP_PRODDISK
ITEP_HOTDISK
```

- The DDM locations rather than the CE names
  - "the jobs go to the data"
  - So, GANGA choose the CE according to on which DDM locations the input data resides

Monitoring: Sites Exclusion

- DDM Site Services: http://cern.ch/project-atlas-ddm-sls/SS_sites-shares_monitor.txt

```
atlddm28: sites= GRIF-SACLAY_MCDISK,GRIF-SACLAY_DATADISK,GRIF-SACLAY_PRODDISK,GRIF-SACLAY_HOTDI
atlddm28: shares= default,input,production,tier1
atlddm28: blacklistedDestinations=
atlddm28: blacklistedSources=

atlddm29: sites= NIKHEF-ELPROD_DATADISK,SARA-MATRIX_DATADISK,SARA-MATRIX_DATATAPE,SARA-MATRIX_M
atlddm29: shares= tier0
atlddm29: blacklistedDestinations= SARA-MATRIX_DATADISK,SARA-MATRIX_DATATAPE,SARA-MATRIX_MCDISK
atlddm29: blacklistedSources= SARA-MATRIX_DATADISK,SARA-MATRIX_DATATAPE,SARA-MATRIX_MCDISK,SARA

atlddm35: sites= PIC_DATADISK,PIC_DATATAPE,PIC_MCDISK,PIC_MCTAPE,PIC_HOTDISK,PIC_SCRATCHDISK,PI
atlddm35: shares= default,input,production,tier1
atlddm35: blacklistedDestinations=
atlddm35: blacklistedSources=

voatlas42: sites=
voatlas42: shares=
voatlas42: blacklistedDestinations=
voatlas42: blacklistedSources=

lxvm0349: sites= NIKHEF-ELPROD_DATADISK,RAL-LCG2_DATADISK,RAL-LCG2_DATATAPE,SARA-MATRIX_DATADIS
lxvm0349: shares= test
lxvm0349: blacklistedDestinations= SARA-MATRIX_DATADISK,SARA-MATRIX_DATATAPE
lxvm0349: blacklistedSources= SARA-MATRIX_DATADISK,SARA-MATRIX_DATATAPE,NIKHEF-ELPROD_DATADISK
```

Site Service for FR cloud

No exclusion

SARA excluded (for downtime)

Monitoring: Sites Exclusion

- DDM Functional tests (T1-T1, T1-T2): http://atladcops.cern.ch:8000/drmon/ftmon_TiersInfo.html

T1	tier	%	stream	tiers status
CANADA	AUSTRALIA-ATLAS_DATADISK	5%		registered: 6 ready: 5
	CA-ALBERTA-WESTGRID-T2_DATADISK	15%		
	CA-SCINET-T2_DATADISK	45%		
	SFU-LCG2_DATADISK	20%		
	VICTORIA-LCG2_DATADISK	2%		
	ALBERTA-LCG2_DATADISK			
ES	IFAE_DATADISK	21%		registered: 6 ready: 6
	IFIC-LCG2_DATADISK	41%		
	LIP-COIMBRA_DATADISK	6%		
	LIP-LISBON_DATADISK	6%		
	NCG-INGRID-PT_DATADISK	5%		
	UAM-LCG2_DATADISK	21%		
FRANCE	BEIJING-LCG2_DATADISK	10%		registered: 11 ready: 10
	GRIF-LAL_DATADISK	24%		
	GRIF-LPNHE_DATADISK	50%		
	GRIF-SACLAY_DATADISK	26%		
	IN2P3-CPPM_DATADISK	5%		
	IN2P3-LAPP_DATADISK	15%		
	IN2P3-LPC_DATADISK	15%		
	RO-02-NIPNE_DATADISK	10%		
	RO-07-NIPNE_DATADISK	10%		
	TOKYO-LCG2_DATADISK	100%		
	IN2P3-LPSC_DATADISK			

Legend:

- ** ready status
- ** waiting status
- ** not validated status
- ** site in downtime

Associated tiers groups are highlighted with different colors.

Total Tier2s in 10 clouds: 75 (ready: 60)

Excluded from FT

Monitoring: Sites Exclusion

- Real data distribution: http://atladcops.cern.ch:8000/drmon/crmon_TiersInfo.html

	SFU-LCG2_DATADISK	20%		
DE	CSCS-LCG2_DATADISK	15%		registered: 10 ready: 10
	CYFRONET-LCG2_DATADISK	15%		
	DESY-HH_DATADISK	50%		
	DESY-ZN_DATADISK	50%		
	GOEGRID_DATADISK	14%		
	LRZ-LMU_DATADISK	14%		
	MPPMU_DATADISK	14%		
	PRAGUELCG2_DATADISK	15%		
	UNI-FREIBURG_DATADISK	14%		
	WUPPERTALPROD_DATADISK	14%		
ES	NCG-INGRID-PT_DATADISK			registered: 6 ready: 3
	IFAE_DATADISK	21%		
	IFIC-LCG2_DATADISK	41%		
	UAM-LCG2_DATADISK	21%		
	LIP-COIMBRA_DATADISK			
	LIP-LISBON_DATADISK			
FR	BEIJING-LCG2_DATADISK	10%		registered: 8 ready: 8
	GRIF-LAL_DATADISK	45%		
	GRIF-LPNHE_DATADISK	25%		
	GRIF-SACLAY_DATADISK	30%		
	IN2P3-LAPP_DATADISK	15%		
	IN2P3-LPC_DATADISK	15%		
	RO-07-NIPNE_DATADISK	10%		
	TOKYO-LCG2_DATADISK	100%		

Total Tier2s in 9 clouds: 60 (ready: 45)

Excluded from the real data distribution although included in the FT

Monitoring: Sites Exclusion

- MC data distribution: http://panda.cern.ch/?mode=listAODReplications&summary=Clouds

- *Green* - site has a complete dataset replicas (data transfer is done)
- *Orange* - if site has an incomplete dataset replicas. It also means that subscription is processed
- *Red* - site has 0 files
- *Magenta* - site not subscribed or subscription is not processed

Excluded from the MC data distribution

ATLAS AODs on sites : ASGC CERN BNL CNAF FZK LYON NG NIKHEF PIC RAL SARA TRIUMF

T2 subscription requests (AOD) within clouds : Last Data Subscription to Tier-2s : Thu Jan 14 19:11:41 2010

CA	ALBERTA-LCG2_MC	AUSTRALIA-ATLAS_MC	SFU-LCG2_MC	TORONTO-LCG2_MC	VICTORIA-LCG2_MC			
replicas/subscriptions(req. %)	8/0(15)	0/16(20)	9/8(20)	61/0(10)	126/0(10)			
DE	CSCS-LCG2_MC	CYFRONET-LCG2_MC	DESY-HH_MC	DESY-ZN_MC	GOEGRID_MC	LRZ-LMU_MC	MPPMU_MC	PRAGUELCG2_MC
replicas/subscriptions(req. %)	13/1(14)	10/4(14)	32/16(50)	8/32(50)	5/10(15)	20/0(15)	8/6(14)	12/14(15)
ES	IFAE_MC	IFIC-LCG2_MC	LIP-COIMBRA_MC	LIP-LISBON_MC	NCG-INGRID-PT_MC	UAM-LCG2_MC		
replicas/subscriptions(req. %)	19/0(21)	14/0(41)	371/35(6)	1/0(6)	2/0(5)	3/0(21)		
FR	BEIJING-LCG2_MC	GRIF-LAL_MC	GRIF-LPNHE_MC	GRIF-SACLAY_MC	IN2P3-LAPP_MC	IN2P3-LPC_MC	RO-07-NIPNE_MC	TOKYO-LCG2_MC
replicas/subscriptions(req. %)	0/20(10)	7/79(45)	3/45(25)	10/48(30)	0/29(15)	0/29(15)	0/20(10)	19/171(100)

Monitoring: DDM Dashboard

- http://dashb-atlas-data.cern.ch/dashboard/request.py/site

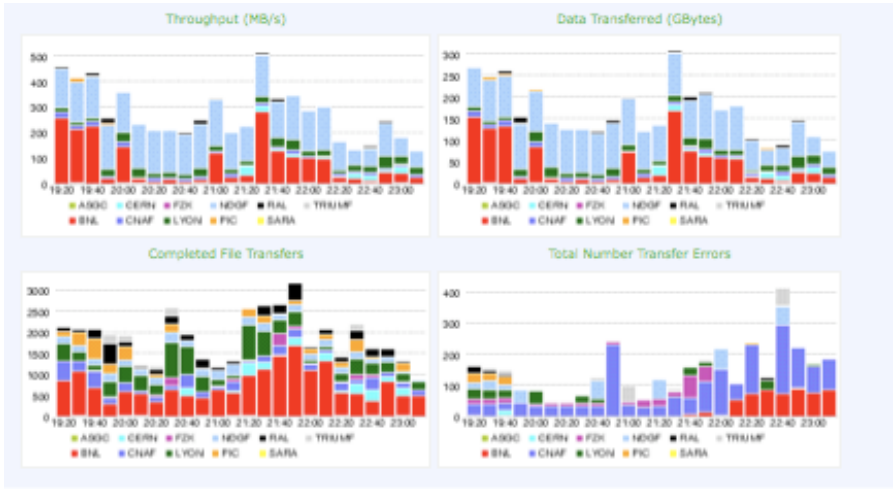

Activity Summary ('2010-01-13 19:20' to '2010-01-13 23:20' UTC)

Click on the cloud name to view list of sites

	Transfers			Registrations		Errors			Services
Cloud	Efficiency	Throughput	Successes	Datasets	Files	Transfer	Registration	Services	Grid
ASGC	100%	9 kB/s	42	18	42	0	0	0	
BNL	97%	91 MB/s	18600	906	18619	548	0	0	
CERN	98%	7 MB/s	2055	207	2061	35	0	0	
CNAF	71%	13 MB/s	4759	351	4751	1981	0	0	
FZK	76%	1 MB/s	1030	606	1031	322	0	0	
LYON	97%	25 MB/s	8048	865	8042	271	0	0	
NDGF	90%	148 MB/s	4015	212	4018	437	0	0	
PIC	96%	3 MB/s	3149	469	3130	127	0	0	
RAL	98%	6 MB/s	3868	677	3873	92	0	0	
SARA	0%	0 kB/s	0	0	0	0	0	0	
TRIUMF	83%	866 kB/s	966	274	968	193	0	0	

- To watch the transfer statistics for clouds/sites
- To see the errors
- To see the datasets/files in transfers
- To see the history of transfers for a file

Monitoring: Subscriptions and Transfers

- http://atladcops.cern.ch:8000/drmon/ftmon.html (for Functional Tests)
- http://atladcops.cern.ch:8000/drmon/crmon.html (for Real Data Distribution)

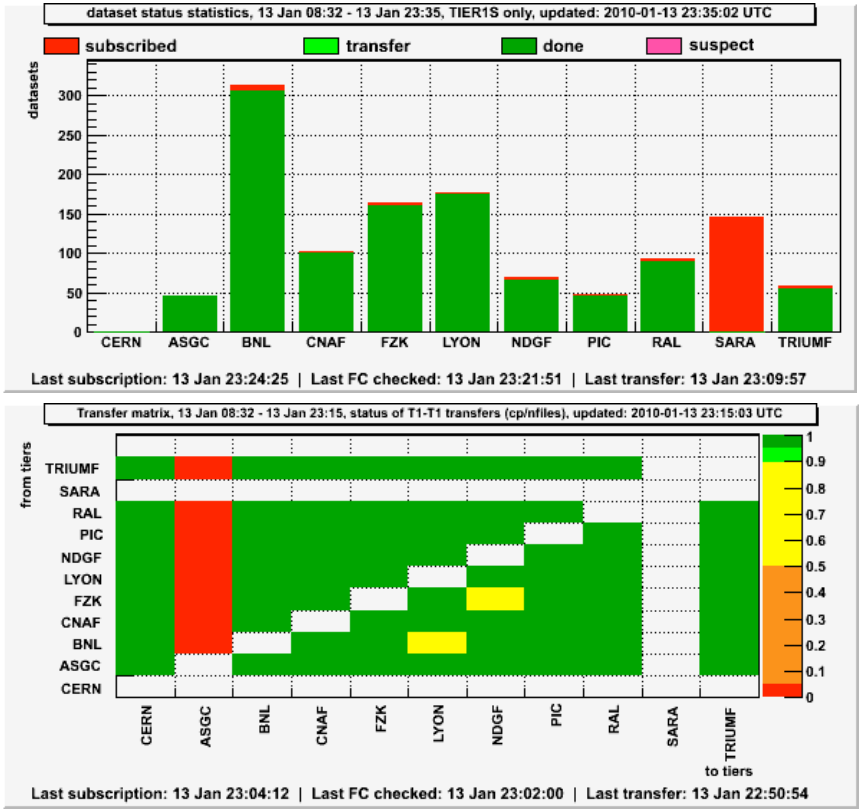

- To see datasets subscribed but transfers not yet completed
  - SARA is in downtime
  - ASGC has something wrong in FTS
- To see 'slowness' in a specific channel
  - BNL - LYON is slow...

Monitoring: FTS

- FTS
  - Useful when we don't see any transfers on the DDM dashboard
    - Really no transfers?
    - or Transfer jobs taking too long?
  - Necessary when we don't see anything wrong on the DDM dashboard
    - Transfers successful
    - but eg.
      - somehow overall transfer rate is low...
        - maybe there are transfers taking too long
      - or, Functional Test transfers do not proceed
        - maybe there are many other transfers in the queue

Monitoring: FTS

- FTS (https://cctools2.in2p3.fr/stockage/fts/monitoring/ftsmonitor.php)

[Back HOME] FTS Monitor v1.2.0 - last updated: 15/01/2010 03:29:21 [updated every 600 seconds] [About]

channel: BNL-IN2P3 for VO atlas Submit Query

Channel configuration

Between **BNL-ATLAS** and **IN2P3-CC**
20 concurrent files on **6** streams each
The channel is **Active**
Last used at **15/01/10 02:29:20,000000 +00:00**
Agent running on **cclcgftsli07**
Go to history of channel management

Check here if there is no transfers

CHANNEL SHARE	CHANNEL LIMIT	CHANNEL VOSTATE
100 %	0	ON

CHANNEL MANAGERS
/DC=org/DC=doegrids/OU=People/CN=Hironori Ito 564424
/O=GRID-FR/C=FR/O=CNRS/OU=CC-LYON/CN=Catherine Biscarat
/O=GRID-FR/C=FR/O=CNRS/OU=CC-LYON/CN=Ghita Rahal
/O=GRID-FR/C=FR/O=CNRS/OU=LAPP/CN=Stephane Jezequel

Jobs statistics (last 14 days)

Ready	Active	Finished	FinishedDirty	Failed	Canceled
1275				326	

Files statistics (number transferred per hour during last 24h)

01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	00	Total
406	370	691	1163	1103	389	412	863	486	1026	899	1164	379	1276	672	894	1281	269	289	372	277	267	935	826	16709

Job details (last 24 hours)

Active job is not too old?

JOB ID	SUBMIT TIME	JOB STATE	SOURCE SE	DEST SE	STORAGE CLASS
4af8238d-017d-11df-b697-9f56045f39a2	15/01/10 02:25:48	Ready	dcsrm.usatlas.bnl.gov	ccsrm.in2p3.fr	ATLASMCDISK
8f2f298c-0178-11df-b697-9f56045f39a2	15/01/10 01:51:55	Ready	dcsrm.usatlas.bnl.gov	ccsrm.in2p3.fr	ATLASDATADISK
889f51c5-0178-11df-8df4-c72f3951a8f6	15/01/10 01:51:44	Ready	dcsrm.usatlas.bnl.gov	ccsrm.in2p3.fr	ATLASDATADISK
7599a124-0178-11df-8df4-c72f3951a8f6	15/01/10 01:51:12	Active	dcsrm.usatlas.bnl.gov	ccsrm.in2p3.fr	ATLASMCDISK
5712d7bc-0178-11df-b987-fb395cd16068	15/01/10 01:50:21	Finished	dcsrm.usatlas.bnl.gov	ccsrm.in2p3.fr	ATLASDATADISK

Monitoring: FTS

- FTS (https://cctools2.in2p3.fr/stockage/fts/monitoring/ftsjob.php?jobid=42b...)
  - inside the job -- per file information

[Back HOME] FTS Monitor v1.2.0 - last updated: 15/01/2010 03:48:09 [updated every 600 seconds] [About]

jobid: 42b72d95-0170-11df-b697-9f56045f39a2 Submit Query

Job information

SUBMIT TIME	JOB STATE	CHANNEL NAME	SPACE TOKEN	STORAGE CLASS	REASON
15/01/10 00:52:31	FinishedDirty	BNL-IN2P3		ATLASMCDISK	One or more files failed. Please have a look at the details for more information

File transfers

FILE ID	SOURCE FILE	FILE STATE	SIZE IN MB	MB PER SECOND	DURATION IN SEC	ERROR SCOPE	ERROR PHASE	REASON
41690281	/TAG.098205._000007.pool.root.1	Finished	1	.07	22			
41690282	/RDO.105458._000060.pool.root.1	Finished	35	.84	42			
41690283	/TAG.098205._000015.pool.root.1	Failed	1	.00	923	TRANSFER	TRANSFER	globus_ftp_client: the server responde with an error 426 Transfer aborted (Transfer was killed)
41690284	/TAG.098989._000047.pool.root.1	Finished	1	.07	21			
41690285	/RDO.105458._000376.pool.root.1	Finished	32	.77	42			
41690286	/RDO.105458._000026.pool.root.1	Finished	33	.73	46			
41690287	/TAG.099022._000006.pool.root.1	Finished	1	.08	22			
41690288	/RDO.105458._000226.pool.root.1	Finished	34	.97	36			

Monitoring: Disk Usage

- http://atlddm02.cern.ch/dq2/accounting/cloud_view/FRANCESITES/30/

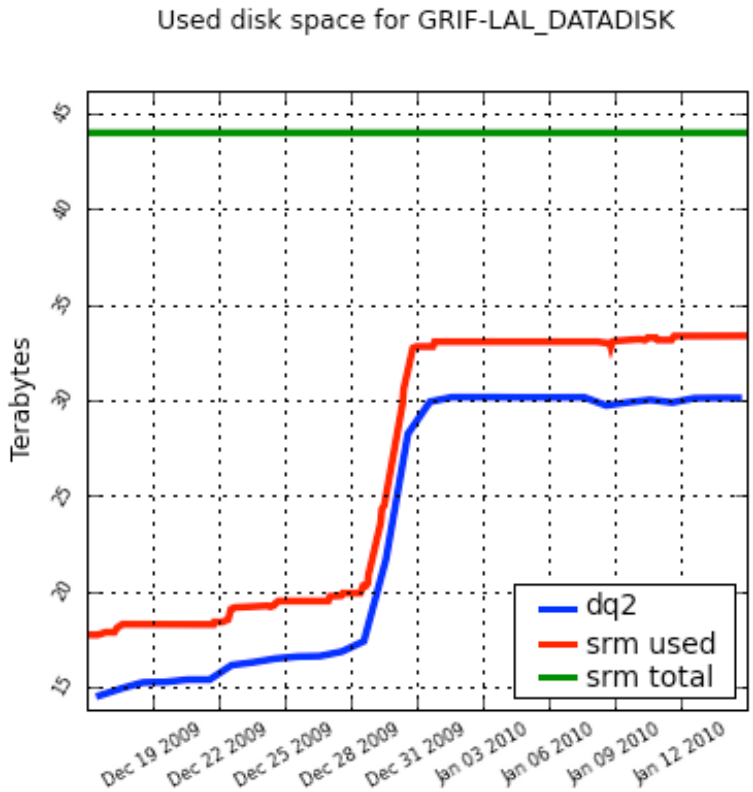

- "srm used" constantly larger than "dq2" indicating "dark data".

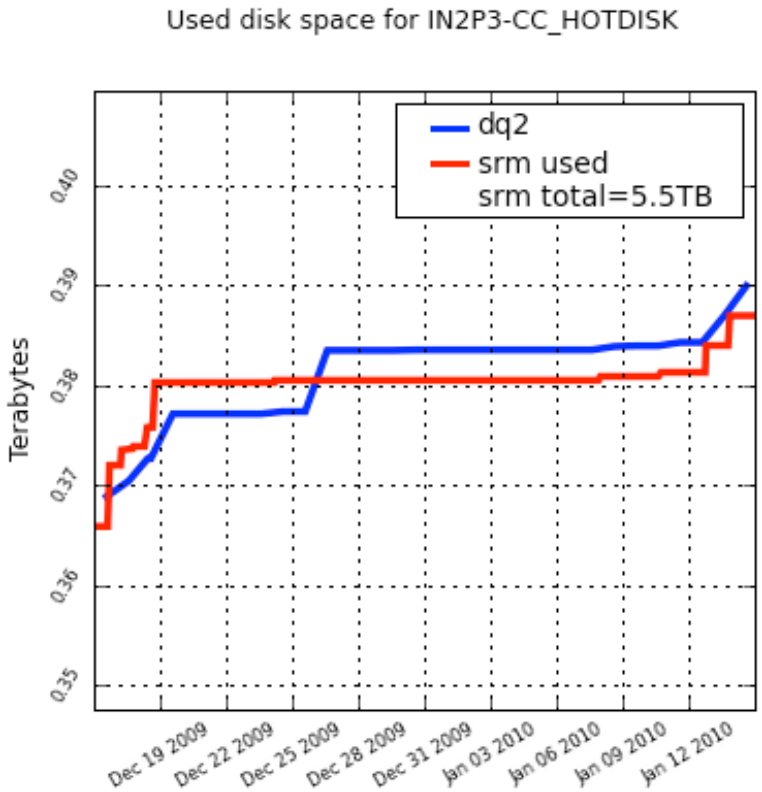

- "srm used" constantly smaller than "dq2" suggesting "incomplete datasets".

Monitoring: Disk Usage

- Dark data
  - Not registered in DQ2, but on the SE
    - remaining files from failed transfers
    - failed upload
    - dq2 registration deleted by accident
    - deletion service removed the catalog entries but failed to delete the physical files
    - ...
  - Need to clean up
    - Registered in LFC, but not in DQ2 :
      - https://twiki.cern.ch/twiki/bin/view/Atlas/DDMOperationsScripts#Detection_of_LFC_files_not_regis
    - LFC - SE inconsistency check :
      - https://twiki.cern.ch/twiki/bin/view/Atlas/DDMOperationsScripts#Dump_of_SE
- Incomplete datasets
  - Need to resubscribe

Monitoring: DDM Deletion Service

- http://atlddm02.cern.ch/dq2/deletion/

DDM Deletion Activity Overview

Clouds	Submitted datasets	Waiting datasets	Deleted files for last hour	Errors for last hour
+ CANADASITES	0	362	1620	0
+ CERN	1	10237	0	22458
- FRANCESITES	5	0	0	0
Sites				
IN2P3-CC_DATADISK	0	0	0	0
IN2P3-CC_DATATAPE	0	0	0	0
IN2P3-CC_HOTDISK	0	0	0	0
IN2P3-CC_LOCALGROUPDISK	0	0	0	0
IN2P3-CC_MCDISK	0	0	0	0
IN2P3-CC_MCT				
IN2P3-CC_PERI				
IN2P3-CC_PERI				
IN2P3-CC_PERI				
IN2P3-CC_PHY				
IN2P3-CC_PHY				
IN2P3-CC_SCR				
BEIJING-LCG2_				
BEIJING-LCG2_				
BEIJING-LCG2_				
BEIJING-LCG2_				

Result page **1** - **1** of about **5** results for * **at GRIF-LAL_LOCALGROUPDISK**. (**0.0270771980286** seconds)

Result Page: : 1

Site	Dataset	State	Storage Cleanup	Catalog Cleanup	Creation date	Last Update
GRIF-LAL_LOCALGROUPDISK	user09.LaurentVacavant.mc09-00900GeV-105012-j3.evgen.e466.v0_der1260976202	Submitted	N	N	2010-01-08 14:36:05	2010-01-08 14:36:05
GRIF-LAL_LOCALGROUPDISK	data08_cosmag.00091639.physics_L1Calo.recon.ESD.o4_r602_tid032002	Submitted	N	N	2009-12-29 16:40:50	2009-12-29 16:40:50
GRIF-LAL_LOCALGROUPDISK	data08_cosmag.00091639.physics_L1Calo.recon.ESD.o4_r602_tid031194	Submitted	N	N	2009-12-29 16:33:32	2009-12-29 16:33:32
GRIF-LAL_LOCALGROUPDISK	data08_cosmag.00091639.physics_RNDM.recon.ESD.o4_r602_tid032006	Submitted	N	N	2009-12-29 16:40:51	2009-12-29 16:40:51
GRIF-LAL_LOCALGROUPDISK	data08_cosmag.00091639.physics_RNDM.recon.ESD.o4_r602_tid031195	Submitted	N	N	2009-12-29 16:33:34	2009-12-29 16:33:34

Monitoring simulation and data processing

- ATLAS production Dashboard:
- Monitoring your site
- Monitoring tasks at your site
  - Disentangle site/task errors
  - Top bottom architecture: task/site -> job
- Monitoring site/cloud status in Panda
- Simulation Functional Tests
- Why my site is not running jobs ?
- Preparing a downtime, interventions for an uSD
- Reporting problems/errors

Monitoring production: monitoring your cloud

http://dashb-atlas-prodsys-test.cern.ch/dashboard/request.py/overview

ATLAS dashboard

Data: All Activities | Jobs: Production | Jobs: Analysis | Panda: Production | SLS: Central Services

Tasks | Grid jobs | Summaries | Shifters | Functional tests | Admin | User Guide | Feedback

find

view
by grid
by cloud
by dest_cloud
by executortype
by executor
by site
by cluster
by tasktype
by task

select cloud
BNL
None
LYON
RAL
FZK
PIC
SARA
TRIUMF
CNAF
NDGF
ASGC
CERN

Activity in ...

2010-01-12 21:00:00 — 2010-01-13 09:59:59

jobs

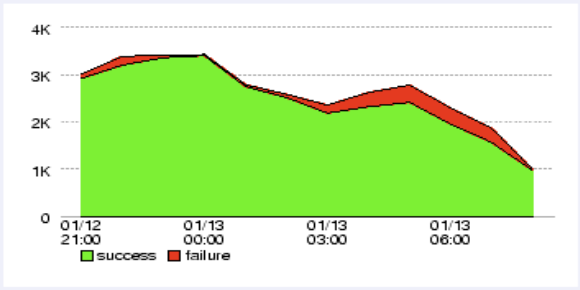

queued jobs

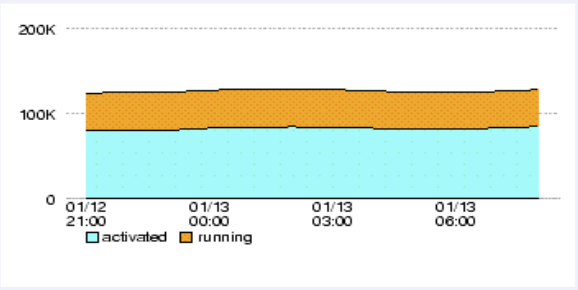

jobs

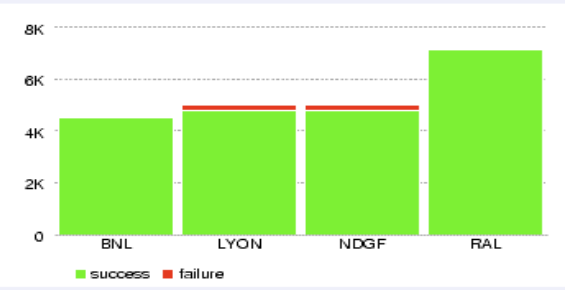

errors (jobs)

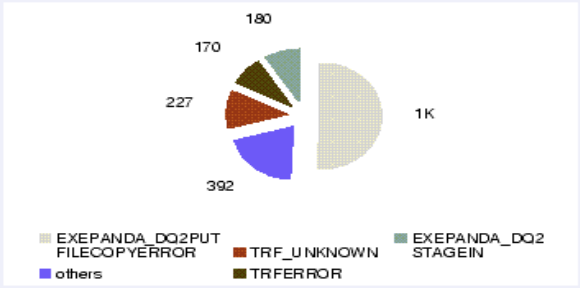

walltime (seconds)

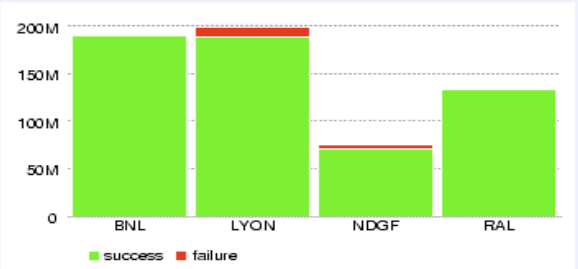

cloud	defined	assigned	waiting	activated	running	holding	transferring	success	failure	efficiency
RAL	0	25	0	6531	5366	402	1983	7057	46	99.4%
NDGF	0	0	0	0	4493	119	486	4694	240	95.1%
LYON	0	316	0	16018	10226	684	3101	4677	213	95.6%
BNL	0	76	0	17087	9370	1522	3842	4438	16	99.6%
TRIUMF	0	95	0	2540	1230	462	2137	3795	154	96.1%
PIC	0	0	0	2858	3158	353	3458	2260	15	99.3%
ASGC	0	0	0	1358	1107	416	73	1678	8	99.5%
CNAF	0	48	0	4122	2060	53	2318	429	1217	26.1%
None	2	684	69	30200	2354	367	2146	428	55	88.6%
CERN	0	10	0	51	22	1	0	40	0	100%
FZK	0	1	0	2676	2784	212	9300	11	2	84.6%
SARA	0	4537	0	211	1332	627	4843	0	0	-
total	*2*	*5792*	*69*	*83652*	*43502*	*5218*	*33687*	*29507*	*1966*	*93.8%*

CRITICAL | WARNING | NORMAL | GOOD | NO_ACTIVITY

Monitoring production: monitoring your cloud

http://dashb-atlas-prodsys-test.cern.ch/dashboard/request.py/overview

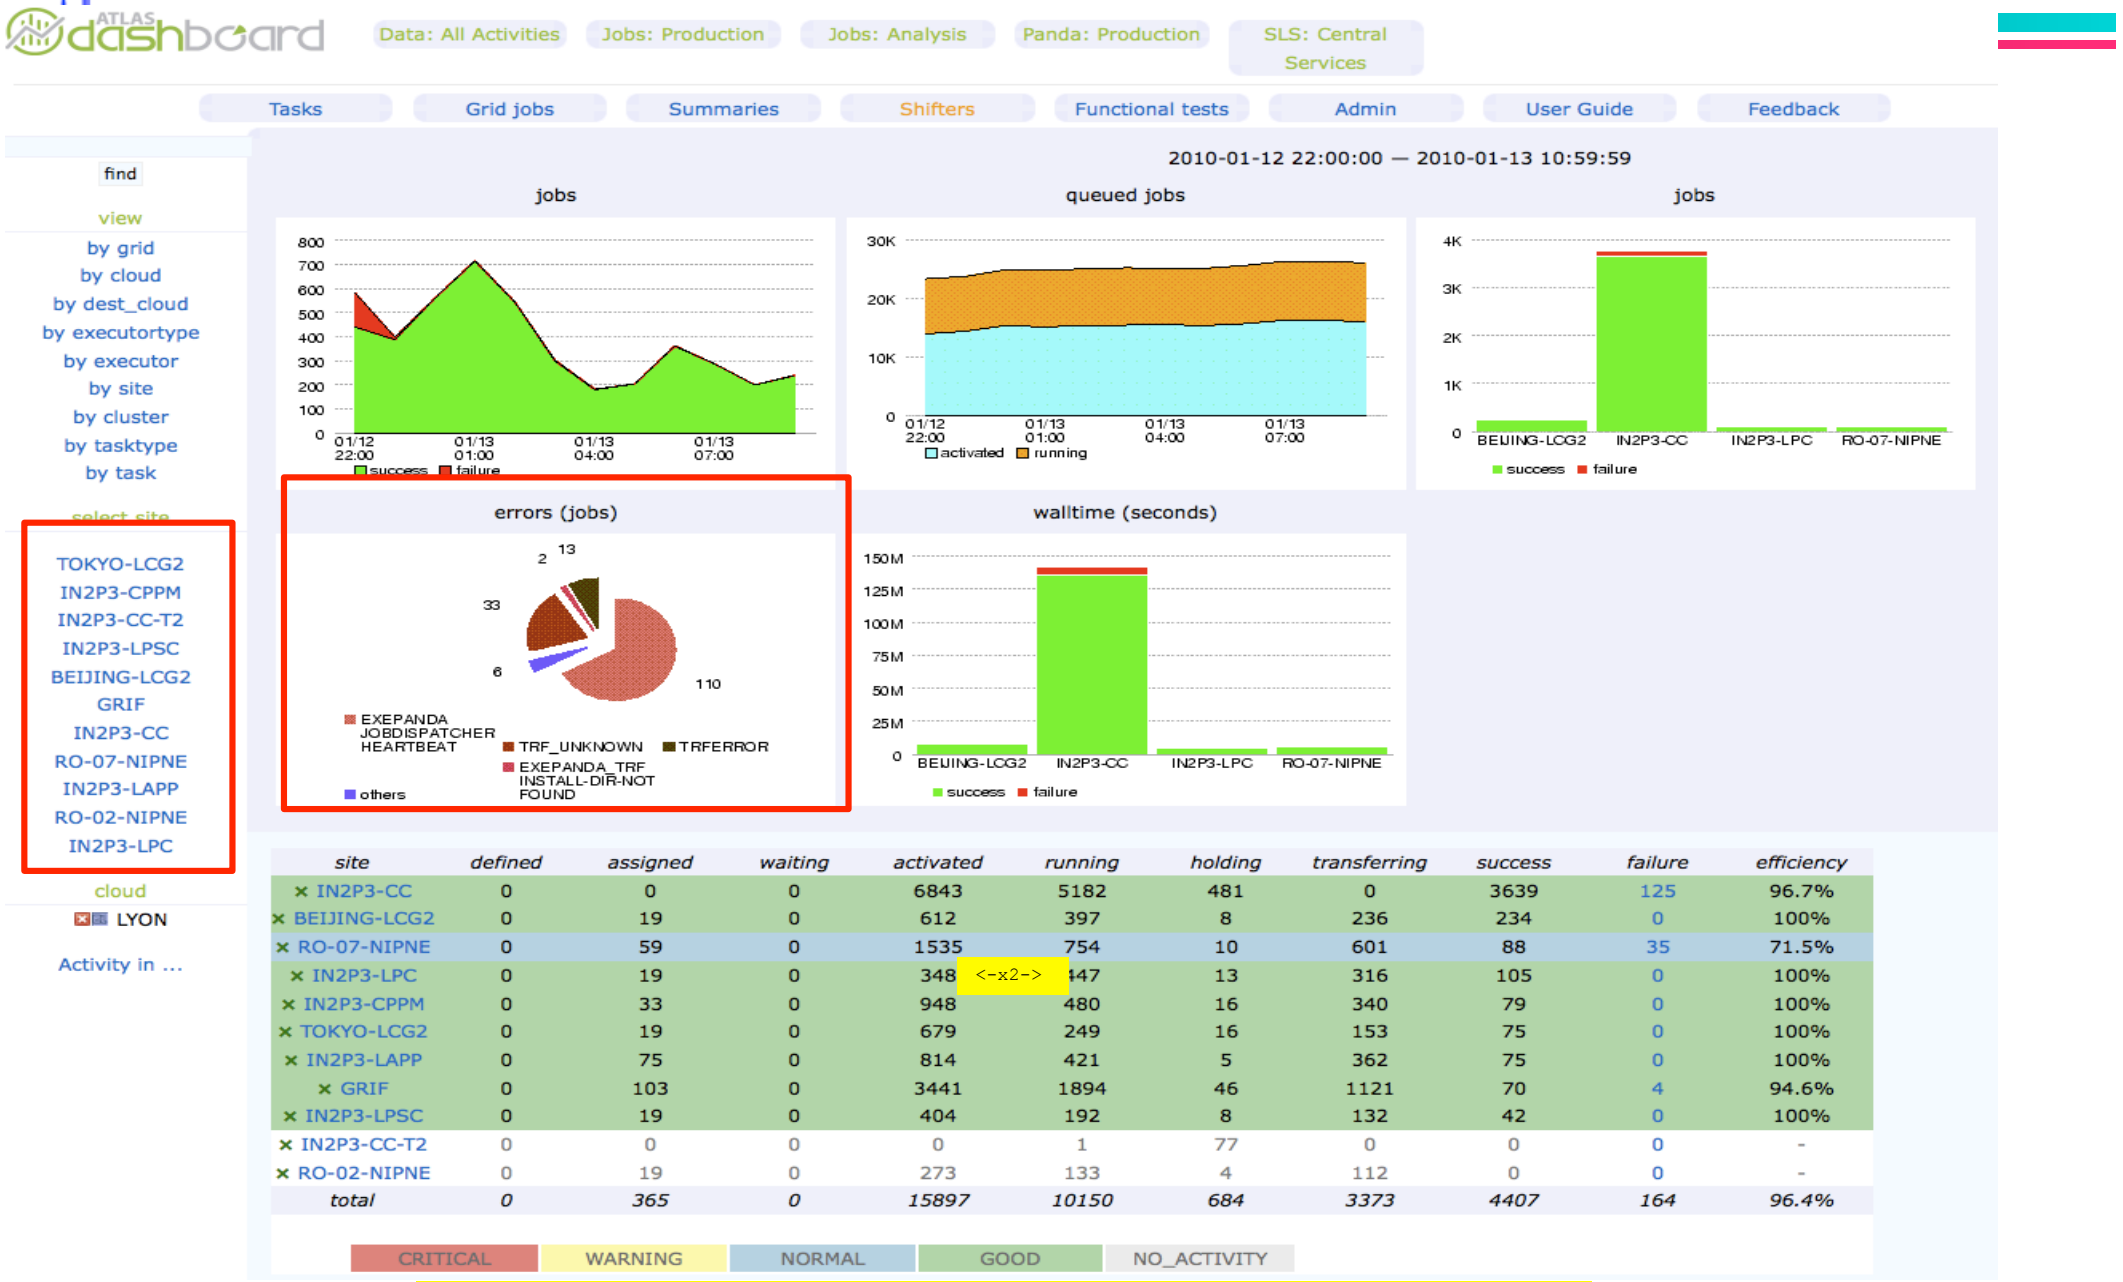

site	defined	assigned	waiting	activated	running	holding	transferring	success	failure	efficiency
× IN2P3-CC	0	0	0	6843	5182	481	0	3639	125	96.7%
× BEIJING-LCG2	0	19	0	612	397	8	236	234	0	100%
× RO-07-NIPNE	0	59	0	1535	754	10	601	88	35	71.5%
× IN2P3-LPC	0	19	0	348	<-x2->	13	316	105	0	100%
× IN2P3-CPPM	0	33	0	948	480	16	340	79	0	100%
× TOKYO-LCG2	0	19	0	679	249	16	153	75	0	100%
× IN2P3-LAPP	0	75	0	814	421	5	362	75	0	100%
× GRIF	0	103	0	3441	1894	46	1121	70	4	94.6%
× IN2P3-LPSC	0	19	0	404	192	8	132	42	0	100%
× IN2P3-CC-T2	0	0	0	0	1	77	0	0	0	-
× RO-02-NIPNE	0	19	0	273	133	4	112	0	0	-
total	*0*	*365*	*0*	*15897*	*10150*	*684*	*3373*	*4407*	*164*	*96.4%*

CRITICAL | WARNING | NORMAL | GOOD | NO_ACTIVITY

https://twiki.cern.ch/twiki/bin/view/Atlas/PandaErrorCodes

Monitoring production: monitoring your site

http://dashb-atlas-prodsys-test.cern.ch/dashboard/request.py/overview

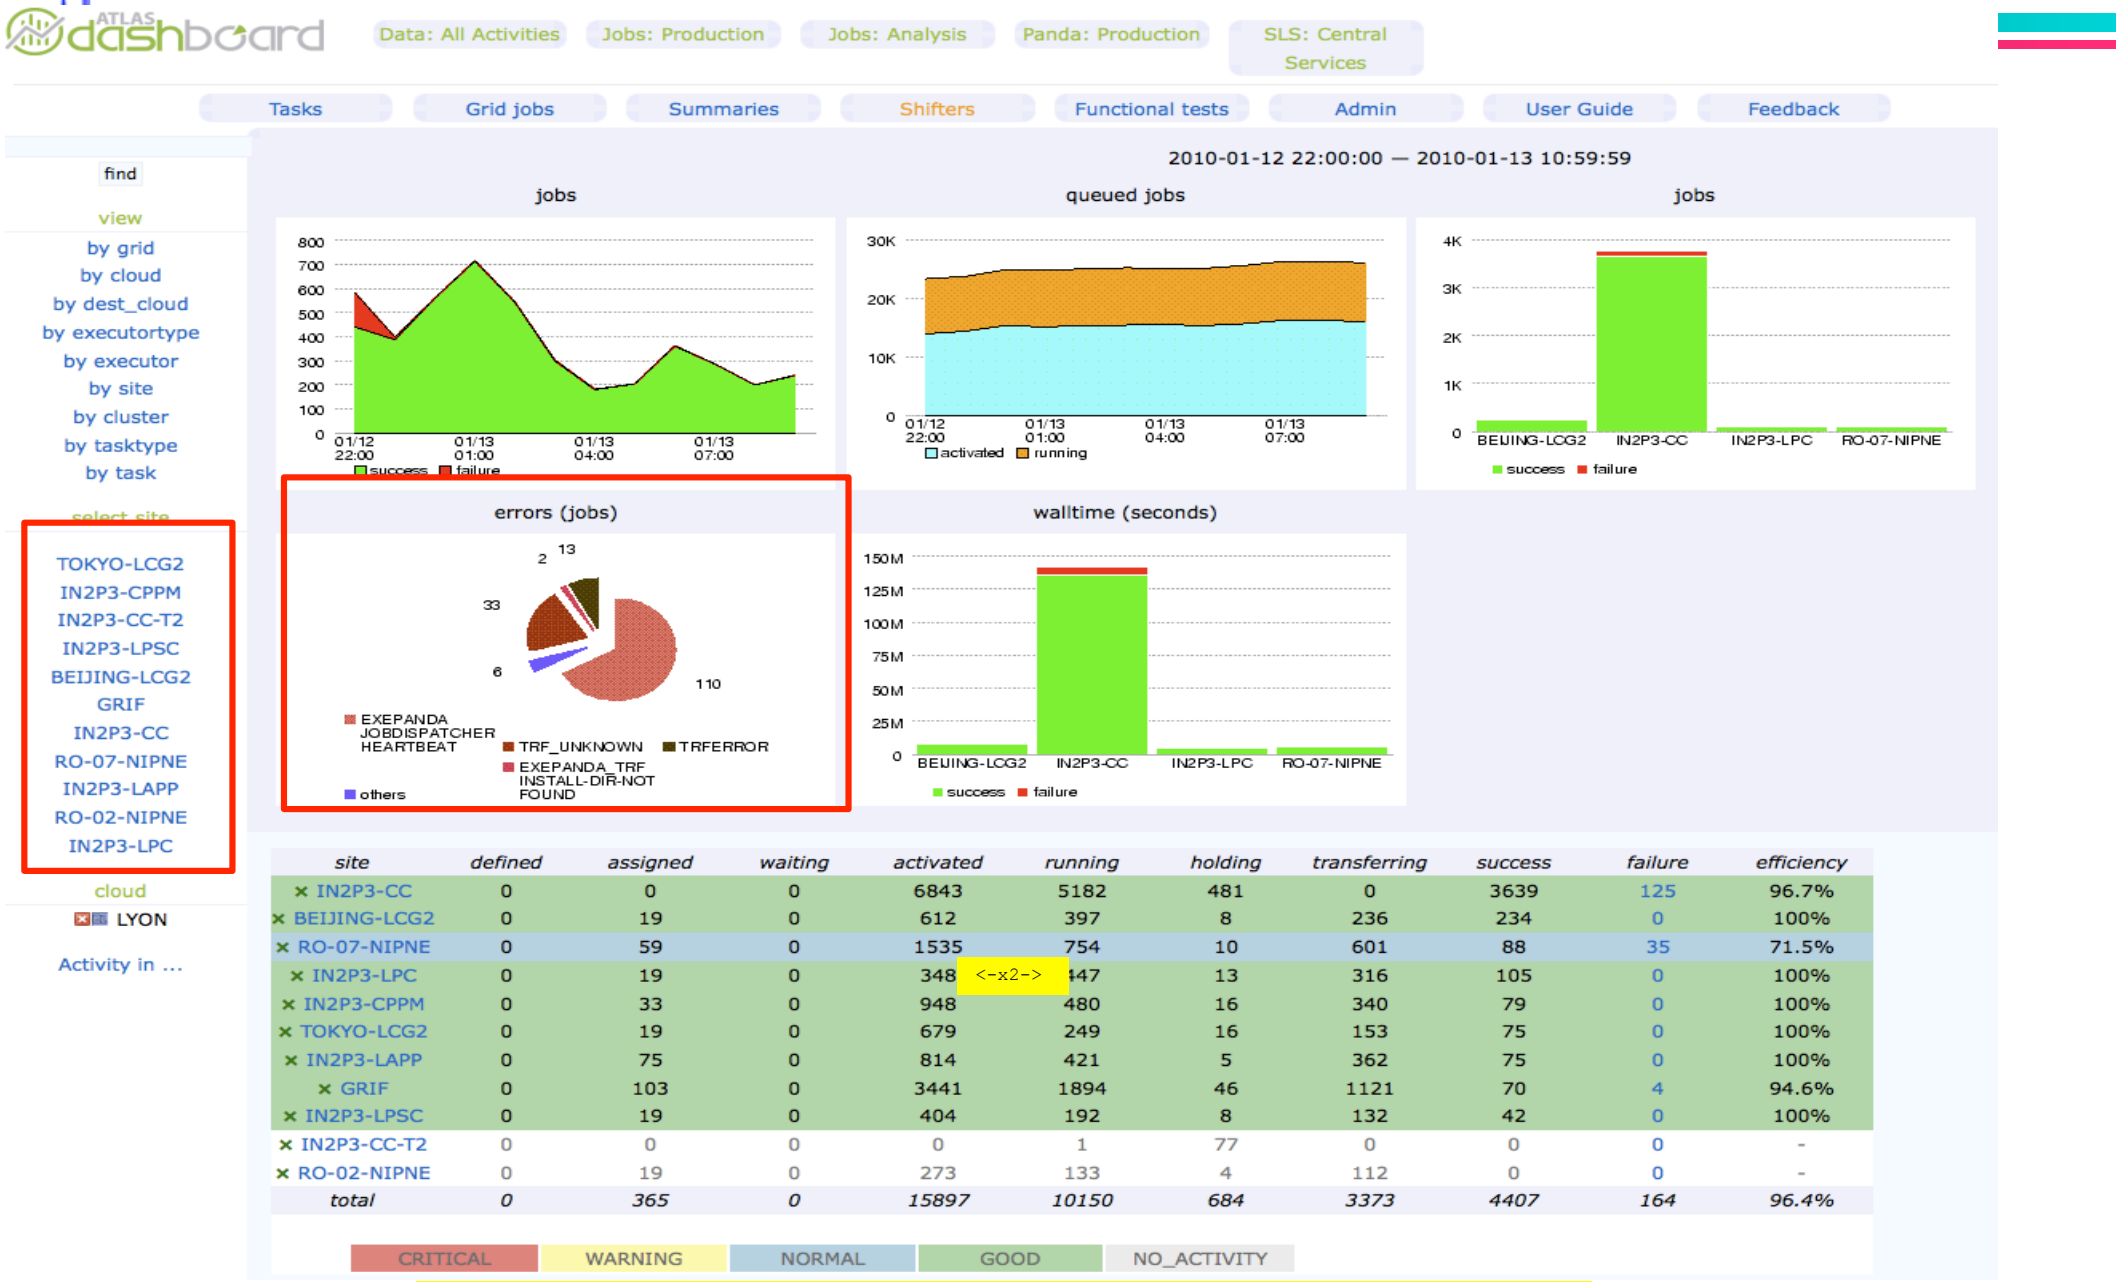

site	defined	assigned	waiting	activated	running	holding	transferring	success	failure	efficiency
IN2P3-CC	0	0	0	6843	5182	481	0	3639	125	96.7%
BEIJING-LCG2	0	19	0	612	397	8	236	234	0	100%
RO-07-NIPNE	0	59	0	1535	754	10	601	88	35	71.5%
IN2P3-LPC	0	19	0	348	<-x2->	13	316	105	0	100%
IN2P3-CPPM	0	33	0	948	480	16	340	79	0	100%
TOKYO-LCG2	0	19	0	679	249	16	153	75	0	100%
IN2P3-LAPP	0	75	0	814	421	5	362	75	0	100%
GRIF	0	103	0	3441	1894	46	1121	70	4	94.6%
IN2P3-LPSC	0	19	0	404	192	8	132	42	0	100%
IN2P3-CC-T2	0	0	0	0	1	77	0	0	0	-
RO-02-NIPNE	0	19	0	273	133	4	112	0	0	-
total	*0*	*365*	*0*	*15897*	*10150*	*684*	*3373*	*4407*	*164*	*96.4%*

CRITICAL | WARNING | NORMAL | GOOD | NO_ACTIVITY

https://twiki.cern.ch/twiki/bin/view/Atlas/PandaErrorCodes

Monitoring production: monitoring your site

http://dashb-atlas-prodsys-test.cern.ch/dashboard/request.py/overview

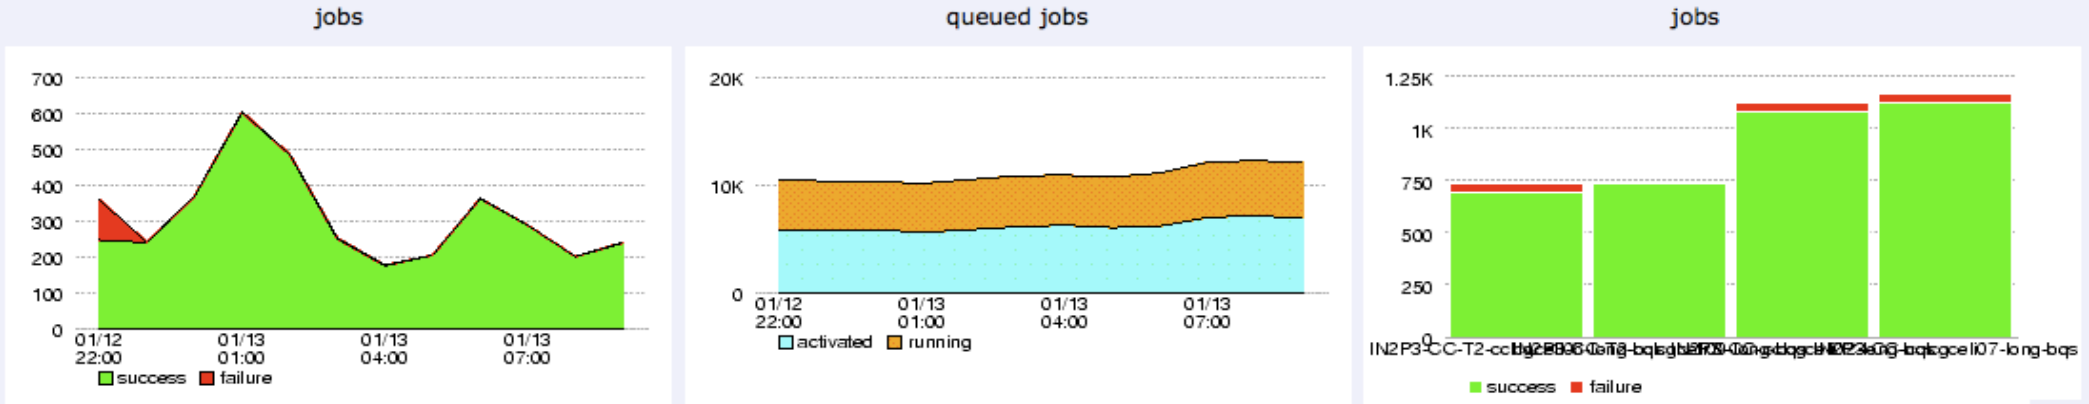

Going further by clicking on the failures link

Many relevant info contained in these statistics

cluster	defined	assigned	waiting	activated	running	holding	transferring	success	failure	efficiency
× IN2P3-CC-cclcgceli07-long-bqs	0	0	0	0	2010	13	0	1124	42	96.4%

type	endpoint	status (SAM)	time
CE	cclcgceli02.in2p3.fr	ok	2010-01-13 09:49:02 UTC
CE	cclcgceli07.in2p3.fr	ok	2010-01-13 09:53:18 UTC

SAM CE tests results

Monitoring production: monitoring your site

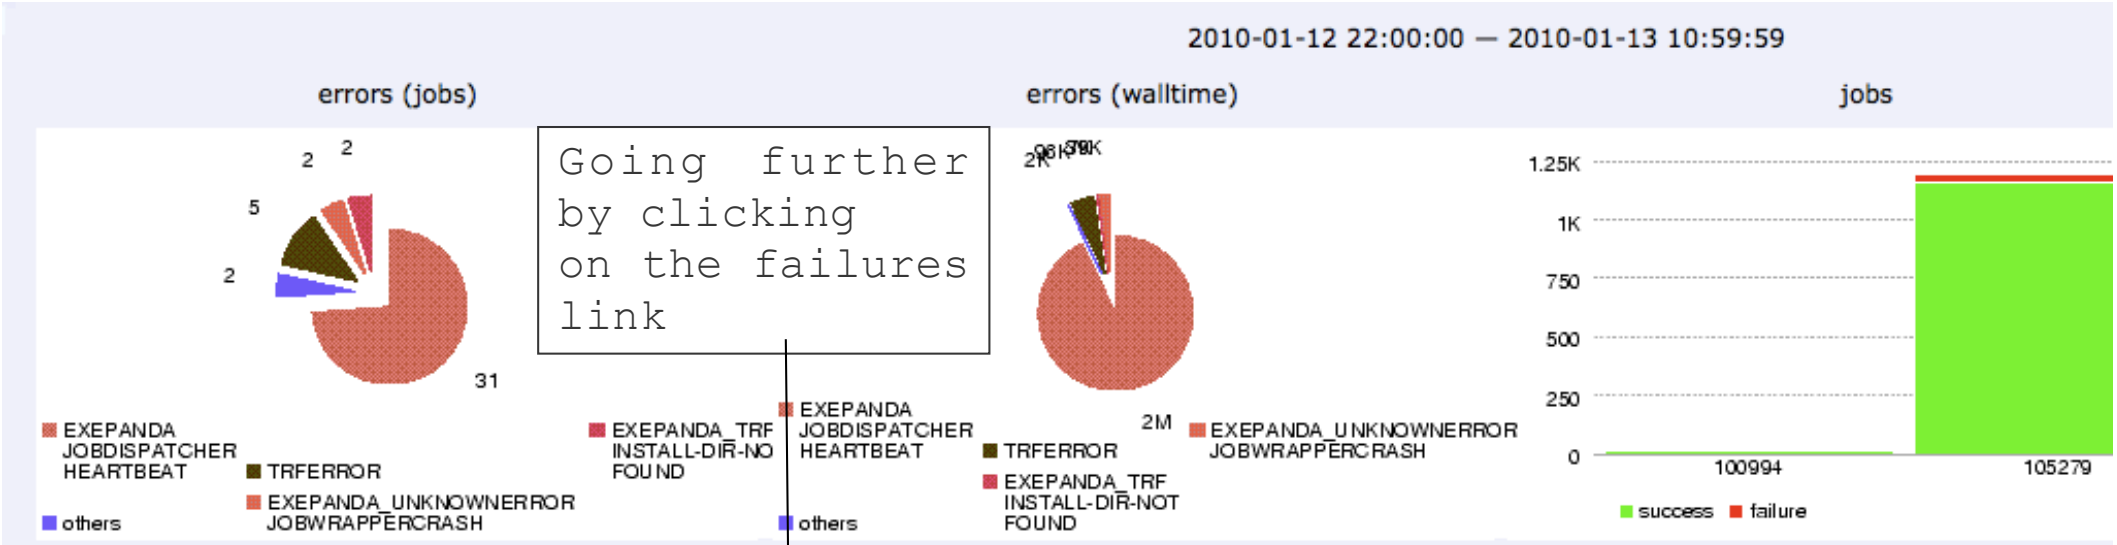

restrict to: EXEPANDA_JOBDISPATCHER_HEARTBEAT (31), TRFERROR (5), EXEPANDA_TRF_INSTALL-DIR-NOT-FOUND (2),

task	1	2	3	others
105279 (42) •	EXEPAND...EARTBEAT (31)	TRFERROR (5)	EXEPAND...OT-FOUND (2)	others (4)

The magic blue dot: click to see if the task is running in other clusters. useful to disentangle site/task issues

site	defined	assigned	waiting	activated	running	holding	transferring	success	failure	efficiency
× IN2P3-CC	0	0	0	0	0	0	0	1155	42	96.5%
total	*0*	*0*	*0*	*0*	*0*	*0*	*0*	*1155*	*42*	*96.5%*

Monitoring production: monitoring your site

this error (jobs)

most common error messages

message (click to expand)	jobs
lost heartbeat : 2010-01-12 22:42:37	1
lost heartbeat : 2010-01-12 22:32:24	1
lost heartbeat : 2010-01-12 22:36:13	1
lost heartbeat : 2010-01-12 22:32:13	1

text/csv

jobexeid	jobdeffk	taskfk	jobname	error	message
66575162	61077288	105279	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.e507_s624_tid105279._072293.job	EXEPANDA_JOBDISPATCHER_HEARTB	lost heartbeat : 2010-01-12 2

error text:
lost heartbeat : 2010-01-12 22:32:33
jobexeid: 66575162
supervisor:
infoexecutor:
creationtime:
attemptnr: 1
errorcode:
partnr:
endtime: 2010-01-12 22:32:33+00:00
modificationtime: 2010-01-13 04:39:16+00:00
nevents:
starttime: 2010-01-12 06:41:04+00:00
execluster: IN2P3-CC-cclcgceli07-long-bqs
processing host: ccwl0133
facilityid: 1040933136 (click for logs)
software: 15.3.1

Job execution details

Access to full job details (link to Panda)

66575146	61077272	105279	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.e507_s624_tid105279._072277.job	EXEPANDA_JOBDISPATCHER_HEARTB	lost heartbeat : 2010-01-12 2
66574545	61077162	105279	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.e507_s624_tid105279._072167.job	EXEPANDA_JOBDISPATCHER_HEARTB	lost heartbeat : 2010-01-12 2
66574499	61077116	105279	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.e507_s624_tid105279._072121.job	EXEPANDA_JOBDISPATCHER_HEARTB	lost heartbeat : 2010-01-12 2
66574498	61077115	105279	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.e507_s624_tid105279._072120.job	EXEPANDA_JOBDISPATCHER_HEARTB	lost heartbeat : 2010-01-12 2
66574435	61077052	105279	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.e507_s624_tid105279._072057.job	EXEPANDA_JOBDISPATCHER_HEARTB	lost heartbeat : 2010-01-12 2

Monitoring production: monitoring your site

Panda job information

Jobs: 1040933136

Click for help

PandaID, Owner, Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1040933136 borut.kersevan@ijs.si	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.e507_s624_tid105279._072293.job #1	failed	2010-01-12 04:13	2:27:25	15:51:29	01-12 22:32	FR/LYON, production	130
	Error details: jobDispatcher: lost heartbeat : 2010-01-12 22:32:33							
	In: mc09_7TeV.105831.JF35_herwig_jet_filter.evgen.EVNT.e507_tid104135_00 **Out:** mc09_7TeV.105831.JF35_herwig_jet_filter.simul.HITS.e507_s624_tid105279_07							

Job 1040933136 details

Job outputs were registered in the IN2P3-CC_MCDISK,IN2P3-CC_DATADISK,IN2P3-CC_DATATAPE,IN2P3-CC_MCTAPE storage element

4 files for job 1040933136:

Filename	Type	Status	Dataset
DBRelease-7.5.1.tar.gz guid=d5ae0e1e-2968-4644-a28c-2540074f7d0	input	ready	ddo.000001.Atlas.Ideal.DBRelease.v070501
EVNT.104135._000255.pool.root.1 guid=64287E86-07F0-DE11-9162-001E4F18BD9A	input	ready	mc09_7TeV.105831.JF35_herwig_jet_filter.evgen.EVNT.e507_tid104135_00
log.105279._072293.job.log.tgz.1 guid=719055ef-b2b3-4337-b81e-cc497f29e636	log	failed	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.log.e507_s624_tid105279_07 (destination block: ..._sub04943187)
HITS.105279._072293.pool.root.1 (no guid)	output	failed	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.HITS.e507_s624_tid105279_07 (destination block: ..._sub04943186)

Show associated task 105279
Show recent jobs for task 105279

Transformation tags: e507 s624 Interpret tags and show transformation configuration

Find and view log files

No job log extract found for job PandaID=1040933136

Look for logging monitor records for job 1040933136

Job failure category: time

File lookup error

GUID 719055ef-b2b3-4337-b81e-cc497f29e636
File lookup failed.
Returned: [LFC][lfc_getreplica][] lfc-prod.in2p3.fr: 719055ef-b2b3-4337-b81e-cc497f29e636: No such file or directory lcg_lr: No such file or directory

Cannot find PFNs from guid 719055ef-b2b3-4337-b81e-cc497f29e636
Error looking up file in IN2P3-CC_MCTAPE replica catalog.
Lookup output:

Job lost heartbeat, no time to store log

JobSpecs for job 1040933136:

PandaID	1040933136
jobDefinitionID	61077288
schedulerID	atlasprod-FR
pilotID	http://cclcgatlas04.in2p3.fr:43210/autopilot/logs/factory/production/2010-01-12/cclcgceli07.in2p3.fr_2119_jobmanager-bqs-long/4966752.6.out BatchID: lcg0112035108-09988 Batch system type: BQS
creationTime	2010-01-12 04:13:39
creationHost	voatlas23.cern.ch
modificationTime	2010-01-13 04:39:16
modificationHost	ccwl0133
AtlasRelease	Atlas-15.3.1
transformation	csc_atlasG4_trf.py
homepackage	AtlasProduction/15.3.1.8
prodSeriesLabel	pandatest
prodSourceLabel	managed
prodUserID	borut.kersevan@ijs.si
assignedPriority	130
currentPriority	130
attemptNr	1
maxAttempt	--
jobStatus	failed
jobName	mc09_7TeV.105831.JF35_herwig_jet_filter.simul.e507_s624_tid105279._072293.job

Forbidden

You don't have permission to access /autopilot/logs/factory/production/2010-01-12/cclcgceli07.in2p3.fr_2119_jobmanager-bqs-long/4966752.6.out on this server.

Additionally, a 403 Forbidden error was encountered while trying to use an ErrorDocument to handle the request.

Apache/2.0.52 (Red Hat) Server at cclcgatlas04.in2p3.fr Port 43210

Monitoring production: monitoring your site

- Once accessed to job's final information one has -usually- enough tools to understand de errors.
- Site contacts may want to cross-check and spot persistent vs. transient problems, for instance:
  - Problems with I/O (lcg-cp/lcg-cr)
    - Disentangle transient vs. persistent problems:
      - timing out because an ATLAS pool is overloaded ?
      - Missing files ?
        - Disentangle unaccessible/lost/missing data at the site by browsing the LFC

Monitoring production: monitoring your tasks

- Monitoring duality: site efficiency vs. task efficiency
  - Snapshot of task progress and problematic jobs per cloud

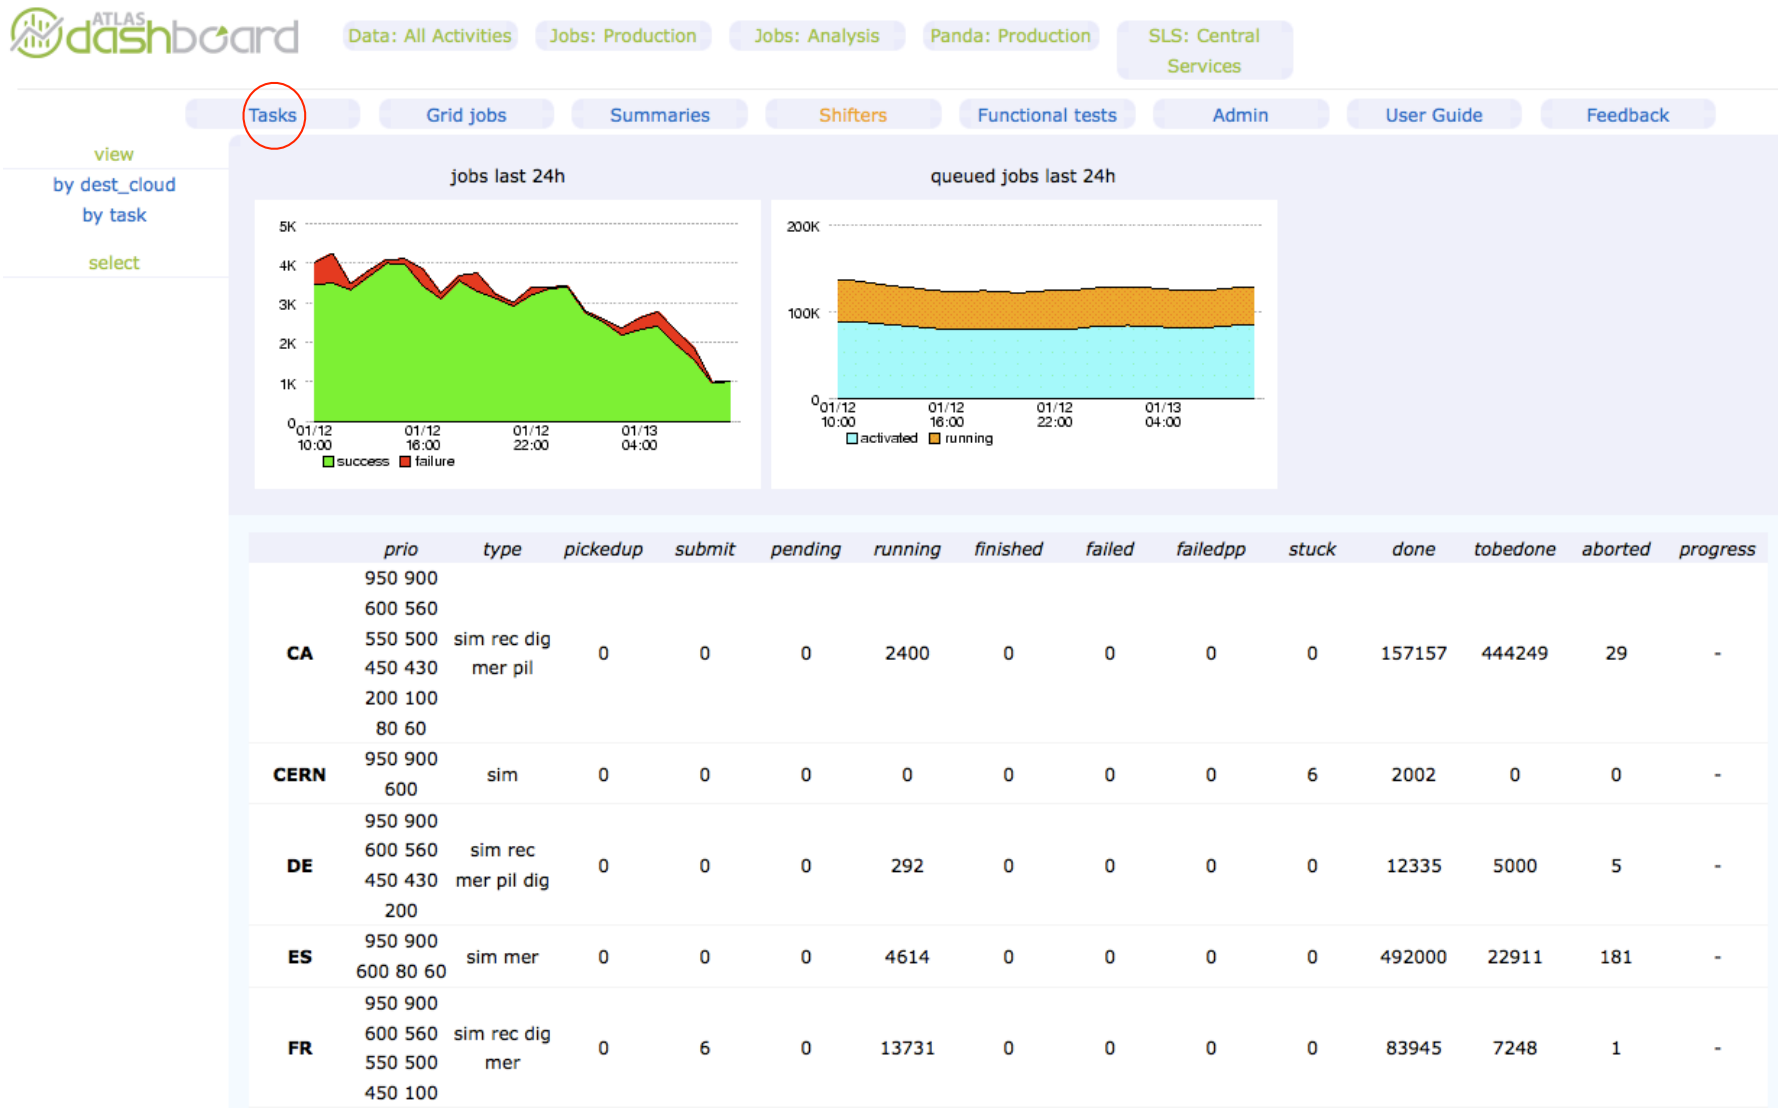

	prio	type	pickedup	submit	pending	running	finished	failed	failedpp	stuck	done	tobedone	aborted	progress
CA	950 900 600 560 550 500 450 430 200 100 80 60	sim rec dig mer pil	0	0	0	2400	0	0	0	0	157157	444249	29	-
CERN	950 900 600	sim	0	0	0	0	0	0	0	6	2002	0	0	-
DE	950 900 600 560 450 430 200	sim rec mer pil dig	0	0	0	292	0	0	0	0	12335	5000	5	-
ES	950 900 600 80 60	sim mer	0	0	0	4614	0	0	0	0	492000	22911	181	-
FR	950 900 600 560 550 500 450 100	sim rec dig mer	0	6	0	13731	0	0	0	0	83945	7248	1	-

Monitoring production: monitoring your tasks

	prio	type	pickedup	submit	pending	running	finished	failed	failedpp	stuck	done	tobedone	aborted	progress
FR	600 560 550 500 450 100	sim rec dig mer	0	6	0	13731	0	0	0	0	83945	7248	1	-
105438	950	simul				26					20			20%
105449	900	simul				9					1			10%
100994	600	simul									1689		1	99.9%
105369	560	reco									102	514		16.6%
105358	550	digit				10					621			69%
105318	500	reco									10	190		5%
105321	500	reco									10	1983		0.5%
105421	450	merge										10		0%
105429	450	merge										10		0%
105406	450	merge										10		0%
105409	450	merge										10		0%
105273	100	simul				1045					16327	586		81.6%
105279	100	simul		6		12641					65165	3935		65.2%

Tasks assigned to LYON cloud
Now one can go to Panda to look for task details:

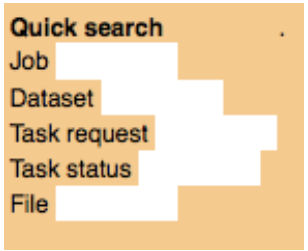

Monitoring production: monitoring your tasks

- http://panda.cern.ch:25980/server/pandamon/query?mode=listtask

Tasks (Jobs) Requests Statistics. (Tue Jan 12 13:41:13 2010 EDT)

ListTaskReqByID ListTaskReqByState

Tasks												Jobs					
total	testing	holding	pending	submitting	rejected	submitted	running	done	finished	failed	aborted	total	submitted	running	done	Last Request	Last Check
9181	0	0	10	80	0	0	76	8409	606	0	0	9248266	256182	2438196	6553888	Jan 11 11:58	Jan 12 12:16

Tasks Requests (only active requests are listed. Aborted requests are not included. The whole list is here)

Task name	Task ID	Req Jobs	Done Jobs	Total events	Prio	Grid	State	Timestamp
valid1.105001.pythia_minbias.recon.e465_s681_r1047	105384	10	2	50000	990	panda@us	submitting	Jan 11 11:58
valid1.105001.pythia_minbias.digit.e465_s681	105383	200	10	50000	990	panda@us	running	Jan 11 11:57
valid1.105200.T1_McAtNlo_Jimmy.recon.e380_s680_r1046	105382	66	1	50000	990	panda@us	submitting	Jan 11 11:55
valid1.105200.T1_McAtNlo_Jimmy.digit.e380_s680	105381	1000	333	50000	990	panda@us	running	Jan 11 11:55
valid1.105200.T1_McAtNlo_Jimmy.recon.e380_s593_r1046	105380	200	200	50000	990	panda@ca	done	Jan 11 11:54
mc09_900GeV.108311.pythia_sdiff_DW.recon.e504_s655_s657_d256_r1022	105379	10	0	1000000	600	panda@nl	submitting	Jan 9 14:25
mc09_900GeV.108311.pythia_sdiff_DW.recon.e504_s655_s657_d257_r1023	105378	10	0	1000000	640	panda@nl	submitting	Jan 9 14:24
data09_900GeV.00141534.physics_CosmicMuons.merge.r988_p62	105377	576	576	9524000	850	panda@us	done	Jan 9 08:31
data09_2TeV.00142308.physics_BPTX.merge.r988_p62	105376	159	159	1930000	850	panda@us	done	Jan 9 08:24
data09_1beam.00140536.express_express.merge.r988_p62	105375	3	3	28000	850	panda@us	done	Jan 9 08:03
data09_1beam.00140477.express_express.merge.r988_p62	105374	25	25	242000	850	panda@us	done	Jan 9 08:03
mc09_900GeV.105004.pythia_ddiff.recon.e466_s678_s679_d274_r1045	105373	798	786	900000	560	panda@it	submitting	Jan 8 22:16
mc09_900GeV.105004.pythia_ddiff.recon.e466_s678_s679_d274_r1045	105372	0	0	100000	560	panda	pending	Jan 8 22:16
mc09_900GeV.105004.pythia_ddiff.recon.e466_s676_s677_d273_r1044	105371	0	0	900000	560	panda	pending	Jan 8 22:15
mc09_900GeV.105004.pythia_ddiff.recon.e466_s676_s677_d273_r1044	105370	0	0	100000	560	panda	pending	Jan 8 22:15
mc09_900GeV.105004.pythia_ddiff.recon.e466_s674_s675_d272_r1043	105369	616	102	900000	560	panda@fr	submitting	Jan 8 22:15
mc09_900GeV.105004.pythia_ddiff.recon.e466_s674_s675_d272_r1043	105368	0	0	100000	560	panda	pending	Jan 8 22:15
mc09_900GeV.105003.pythia_sdiff.recon.e466_s674_s675_d272_r1043	105367	900	883	900000	560	panda@es	running	Jan 8 22:14
mc09_900GeV.105003.pythia_sdiff.recon.e466_s674_s675_d272_r1043	105366	0	0	100000	560	panda	pending	Jan 8 22:14
mc09_900GeV.105003.pythia_sdiff.recon.e466_s676_s677_d273_r1044	105365	900	897	900000	560	panda@de	running	Jan 8 22:14
mc09_900GeV.105003.pythia_sdiff.recon.e466_s676_s677_d273_r1044	105364	0	0	100000	560	panda	pending	Jan 8 22:14
mc09_900GeV.105003.pythia_sdiff.recon.e466_s678_s679_d274_r1045	105363	0	0	900000	560	panda	pending	Jan 8 22:13
mc09_900GeV.105003.pythia_sdiff.recon.e466_s678_s679_d274_r1045	105362	0	0	100000	560	panda	pending	Jan 8 22:13
mc09_900GeV.105001.pythia_minbias.recon.e500_s678_s679_d274_r1045	105361	1000	949	1000000	560	panda@uk	running	Jan 8 22:12
mc09_900GeV.105001.pythia_minbias.recon.e500_s676_s677_d273_r1044	105360	1	0	1000000	560	panda@nl	submitting	Jan 8 22:11
mc09_900GeV.105001.pythia_minbias.recon.e500_s674_s675_d272_r1043	105359	41	10	1000000	560	panda@ca	submitting	Jan 8 22:10
mc09_900GeV.105004.pythia_ddiff.digit.e466_s674_s675_d272	105358	900	616	900000	550	panda@fr	running	Jan 8 22:08
mc09_900GeV.105004.pythia_ddiff.digit.e466_s674_s675_d272	105357	10	0	100000	550	panda@us	submitting	Jan 8 22:08
mc09_900GeV.105004.pythia_ddiff.digit.e466_s676_s677_d273	105356	10	0	900000	550	panda@ca	submitting	Jan 8 22:08
mc09_900GeV.105004.pythia_ddiff.digit.e466_s676_s677_d273	105355	10	0	100000	550	panda@us	submitting	Jan 8 22:08
mc09_900GeV.105004.pythia_ddiff.digit.e466_s678_s679_d274	105354	900	831	900000	550	panda@it	running	Jan 8 22:07
mc09_900GeV.105004.pythia_ddiff.digit.e466_s678_s679_d274	105353	10	0	100000	550	panda@us	submitting	Jan 8 22:07
mc09_900GeV.105003.pythia_sdiff.digit.e466_s678_s679_d274	105352	10	0	900000	550	panda@us	submitting	Jan 8 22:06
mc09_900GeV.105003.pythia_sdiff.digit.e466_s678_s679_d274	105351	10	0	100000	550	panda@us	submitting	Jan 8 22:06

Monitoring production: monitoring your tasks

- http://panda.cern.ch:25980/server/pandamon/query?mode=listtask

Task (mc09_900GeV.105003.pythia_sdiff.recon.e466_s676_s677_d273_r1044) Request Parameters

Input Attributes

Project	mc09_900GeV
Input dataset	mc09_900GeV.105003.pythia_sdiff.digit.e466_s676_s677_d273
Transformation	Reco_trf.py
Transformation Version	15.5.4.10
Trasnfromation Cache	AtlasTier0
Transformation Parameters	--athenaopts,--ignoreerrors,--omitvalidation,DBRelease,autoConfiguration,beamType,conditionsTag,e
Values	NONE,ALL,NONE,8.3.1,everything,NONE,OFLCOND-DR-BS900-ANom-03,NONE,ATLAS-GEO-08-02-00,
Total Number Of Output Events	900000
First File Number in Input Dataset	101
Total Input Files	900
Number Of Events for Output File	1000
CPU per Event	100
Memory Usage	2000
Grid Flavour	panda@de
E-mail	borut.kersevan@ijs.si

Output Atributes

Output Task Name :	**mc09_900GeV.105003.pythia_sdiff.recon.e466_s676_s677_d273_r1044**
Output Datasets	TAG_COMM.AOD.ESD.HIST
Priority	560
SW Release :	**15.5.4**
State	running
Request Time :	**Fri Jan 8 22:14:05 2010**
Comments	MC09 production 20 - URGENT
Project Mode	default
Modification Time :	**Never**
Modified by	None
Request Id :	105365
Total_req_jobs :	**900**
Total_done_jobs :	**897**
Total_avail_jobs :	**900**
Physics group :	**physics**
Queue :	**default**
Ctag :	**r1044**
Bug report :	0

AddRequest UpdateRequest

ATLAS production: Panda Job states

- Typical cloud view in Panda

FR Sites	Job Nodes	Jobs	Latest	Pilot Nodes	defined	assigned	waiting	activated	sent	running	holding	transferring	finished	failed	tot	trf	other
Site Name	2482	810	01-12 13:39	1834	0	279	0	18478	0	9069	92	2322 / 0	7471	810	10%	2%	8%
BEIJING	108	6	01-12 13:39	104	0	19	0	701	0	391	10	171 / 0	405	6	1%	0%	1%
CPPM	82	0	01-12 13:39	112	0	37	0	798	0	475	11	192 / 0	596	0	0%	0%	0%
GRIF-IRFU	155	0	01-12 13:39	209	0	76	0	2262	0	1304	9	547 / 0	1224	0	0%	0%	0%
GRIF-LAL	178	253	01-12 13:39	188	0	19	0	2183	0	1021	2	27 / 0	518	253	33%	0%	33%
GRIF-LPNHE	51	0	01-12 13:39	61	0	19	0	210	0	146	2	57 / 0	188	0	0%	0%	0%
IN2P3-LPSC	29	0	01-12 13:39	34	0	19	0	414	0	191	0	115 / 0	141	0	0%	0%	0%
LAPP	83	0	01-12 13:38	104	0	33	0	722	0	416	3	197 / 0	485	0	0%	0%	0%
LPC	135	1	01-12 13:38	123	0	19	0	341	0	398	3	296 / 0	349	1	0%	0%	0%
LYON	780	325	01-12 13:39	0	0	0	0	6672	0	2472	22	0 / 0	1739	325	16%	0%	15%
Lyon-T2	708	24	01-12 13:39	649	0	0	0	2631	0	1168	11	0 / 0	1157	24	2%	1%	1%
ROMANIA02	30	1	01-12 13:37	34	0	19	0	82	0	100	0	42 / 0	27	1	4%	0%	4%
ROMANIA07	108	200	01-12 13:39	160	0	0	0	978	0	743	10	520 / 0	419	200	32%	22%	10%
TOKYO	35	0	01-12 13:39	56	0	19	0	484	0	244	9	158 / 0	223	0	0%	0%	0%

- defined : job-record inserted in PandaDB
- assigned : dispatchDBlock is subscribed to site
- waiting : input files are not ready
- activated: waiting for pilot requests
- sent : sent to a worker node
- running : running on a worker node
- holding : adding output files to DQ2 datasets
- transferring : output files are moving from T2 to T1
- finished : completed successfully
- failed : failed due to errors

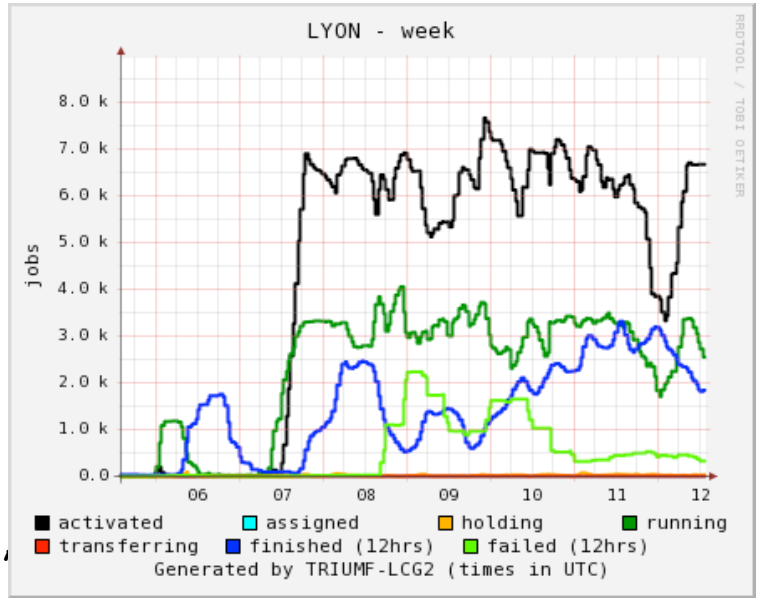

I.Ueda, X. Espinal:

Why is my site not running jobs ?

- First check if site is excluded from production
  - Shifters take sites offline when massive failure are found
    - Site status can be found here (look for status column offline/online)
      - http://panda.cern.ch:25980/server/pandamon/query?dash=clouds
- Check if there are defined jobs in Panda.
  - No defined jobs means nothing to be run.
- Check if there are assigned jobs. If there are assigned jobs means site should run production unless:
  - Input data blocks are not arriving at the site
    - High number of jobs in waiting state. When found, need investigation as usually mean that the input data is not found. Panda brokering is modulated for waiting jobs.
- Check if pilots are running
  - If there are activated jobs and site is not running production, possibly means that pilots are not arriving, so cannot pull jobs from Panda.
- Check if transferring jobs number is high
  - Meaning that outputs cannot be aggregated to T1. brokerage has a protection against this
- Check there is enough disk space
  - lack of free disk space block brokering

Data processing: cloud stats per activity

- Repro-Panglia
  - http://gridinfo.triumf.ca/panglia/graph-generator/?

Panda Graph Generator

Cloud FR

Task reprocessing

ALL
reprocessing
reco
pile
evgen
simul
merge

Time year

Size large

Generated Graph output

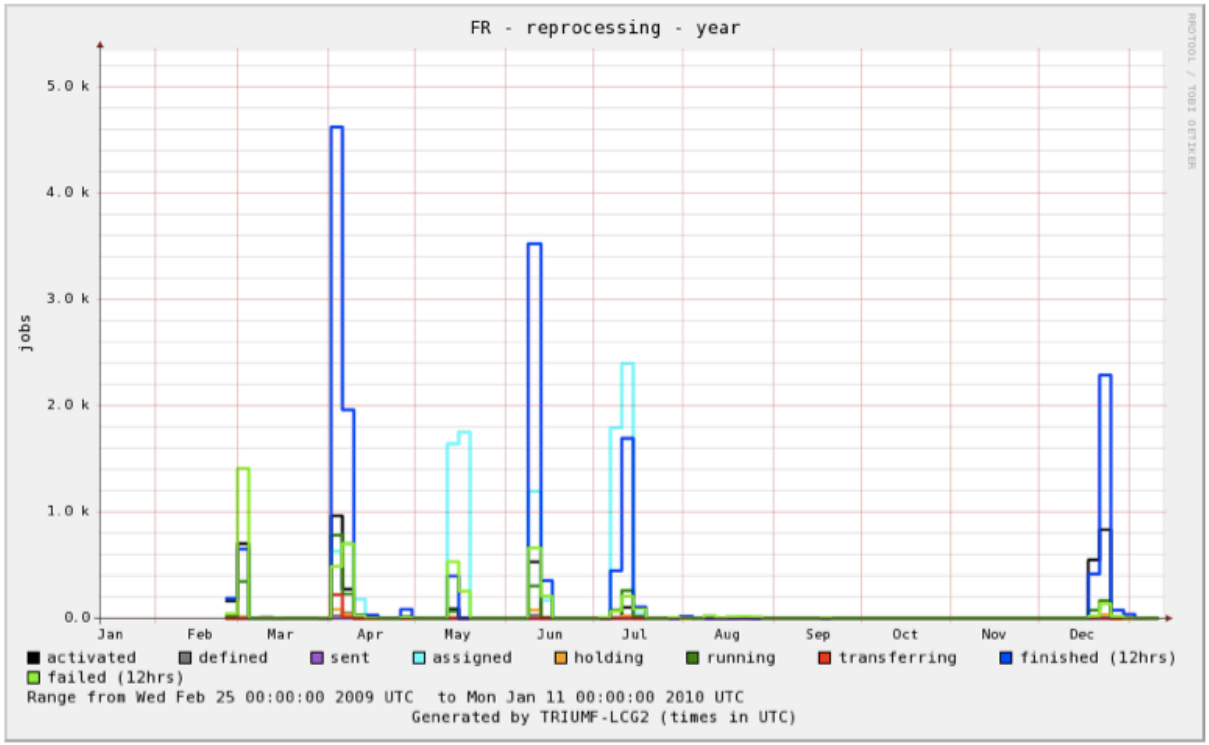

Link for this image:
http://gridinfo.triumf.ca/panglia/graph-generator/graph.php?SITE=FR&TIME=year&SIZE=large&TASK=reprocessing

Simulation Functional tests

- ATLAS is running a continuous load of simulation FT
  - Well known working jobs, stable release
  - Shifters usually look at the stats to knwo if a site is healthy.
  - Soon the stats will be inserted in the site exclusion policy metrics
  - Good way to know if your site has some endemic problem (not Athena related)
- Efficiency should be high by definition, over 90% eff.

Simulation Functional tests

ATLAS dashboard

Data: All Activities | Jobs: Production | Jobs: Analysis | Panda: Production | SLS: Central Services

Tasks | Grid jobs | Summaries | Shifters | Functional tests | Admin | User Guide | Feedback

find

view
by grid
by cloud
by dest_cloud
by executortype
by executor
by site
by cluster
by tasktype
by task

select cloud
CNAF
FZK
PIC
TRIUMF
CERN
LYON
BNL
ASGC
SARA
None
RAL
NDGF

task-flag
functional test

Activity in ...

2010-01-08 11:00:00 — 2010-01-13 11:59:59

jobs

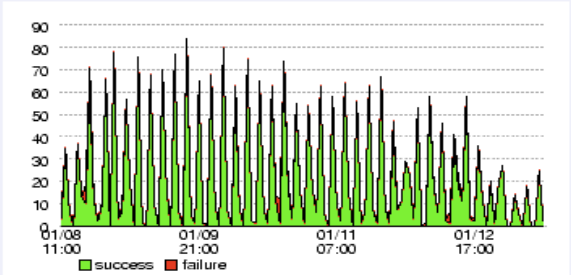

queued jobs

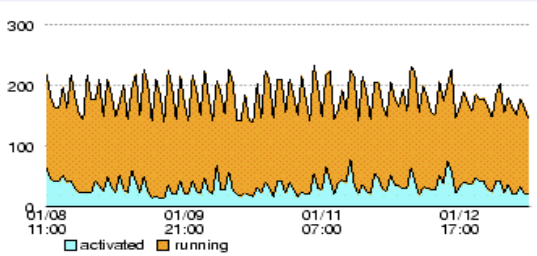

jobs

errors (jobs)

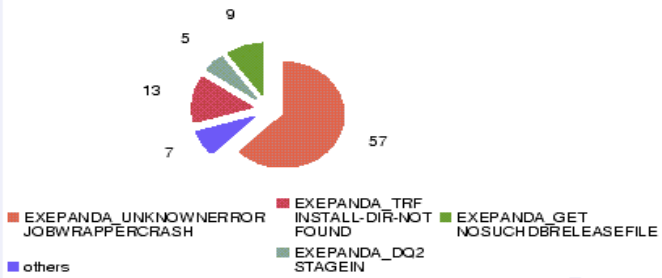

walltime (seconds)

cloud	defined	assigned	waiting	activated	running	holding	transferring	success	failure	efficiency
CERN	0	0	0	0	10	0	0	383	0	100%
CNAF	0	0	0	0	4	0	301	260	65	80%
LYON	0	0	0	0	0	0	0	324	0	100%
BNL	0	0	0	0	0	2	10	310	0	100%
ASGC	0	0	0	0	0	3	0	271	0	100%
PIC	0	0	0	0	24	9	161	218	2	99.1%
None	0	0	0	11	2	1	3	189	22	89.6%
FZK	0	0	0	1	48	0	225	204	0	100%
SARA	0	0	0	8	2	0	59	164	0	100%
RAL	0	0	0	0	0	2	10	139	0	100%
TRIUMF	0	0	0	0	21	7	241	126	2	98.4%
NDGF	0	0	0	0	14	0	0	29	0	100%
total	*0*	*0*	*0*	*20*	*125*	*24*	*1010*	*2617*	*91*	*96.6%*

CRITICAL | WARNING | NORMAL | GOOD | NO_ACTIVITY

Simulation Functional tests

2010-01-08 11:00:00 — 2010-01-13 11:59:59

jobs

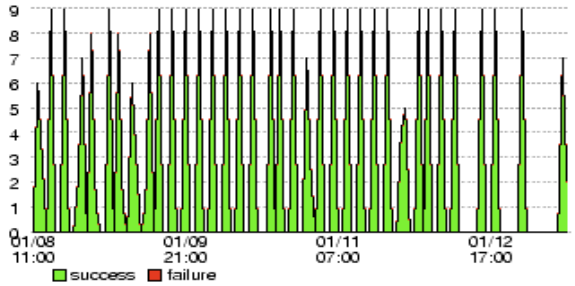

queued jobs

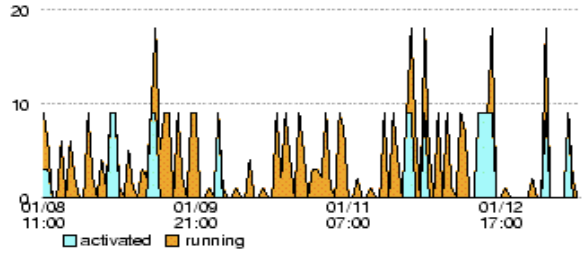

jobs

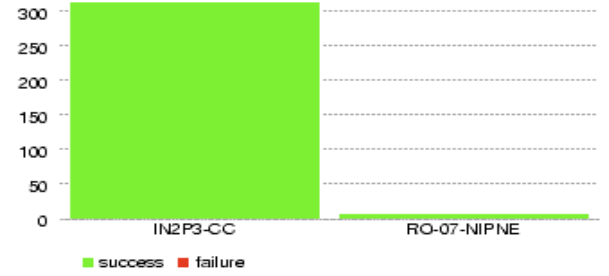

errors (jobs)

walltime (seconds)

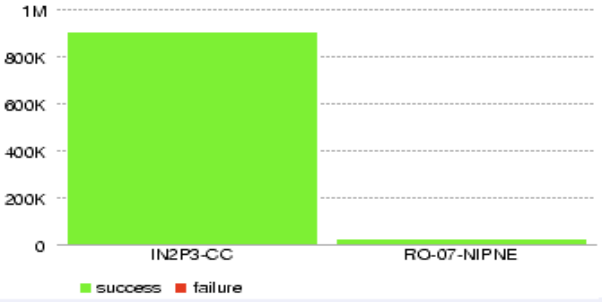

You are watching functional tests jobs. The plot below shows where they ran in the last two weeks. Click the plot to get to the error detail. (large plot)

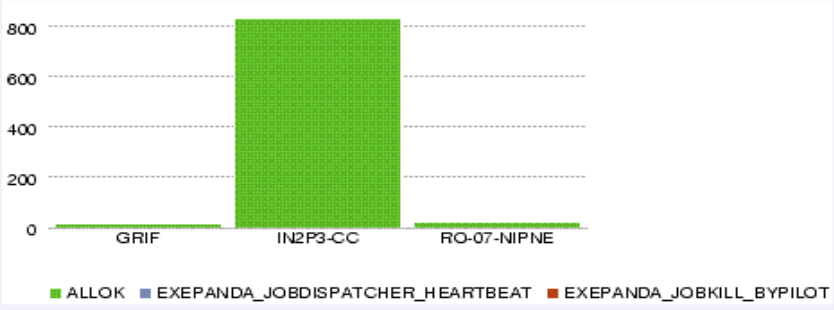

site	defined	assigned	waiting	activated	running	holding	transferring	success	failure	efficiency
× IN2P3-CC	0	0	0	0	0	0	0	315	0	100%
× RO-07-NIPNE	0	0	0	0	0	0	0	9	0	100%
total	*0*	*0*	*0*	*0*	*0*	*0*	*0*	*324*	*0*	*100%*

Centralized tests of basic functions

- SAM tests
  - wLCG SAM tests (ops VO)
  - ATLAS-VO specific SAM tests
    - http://dashb-sam-atlas.cern.ch/dashboard/request.py/historicalsmryview
    - https://lcg-sam.cern.ch:8443/sam/sam.py
  - We will use the ATLAS SAM tests on SRM in the site exclusion policy (see the later slides)
    - So, please check the results before asking for putting the site back into the system

Monitoring: Centralized tests and critical services

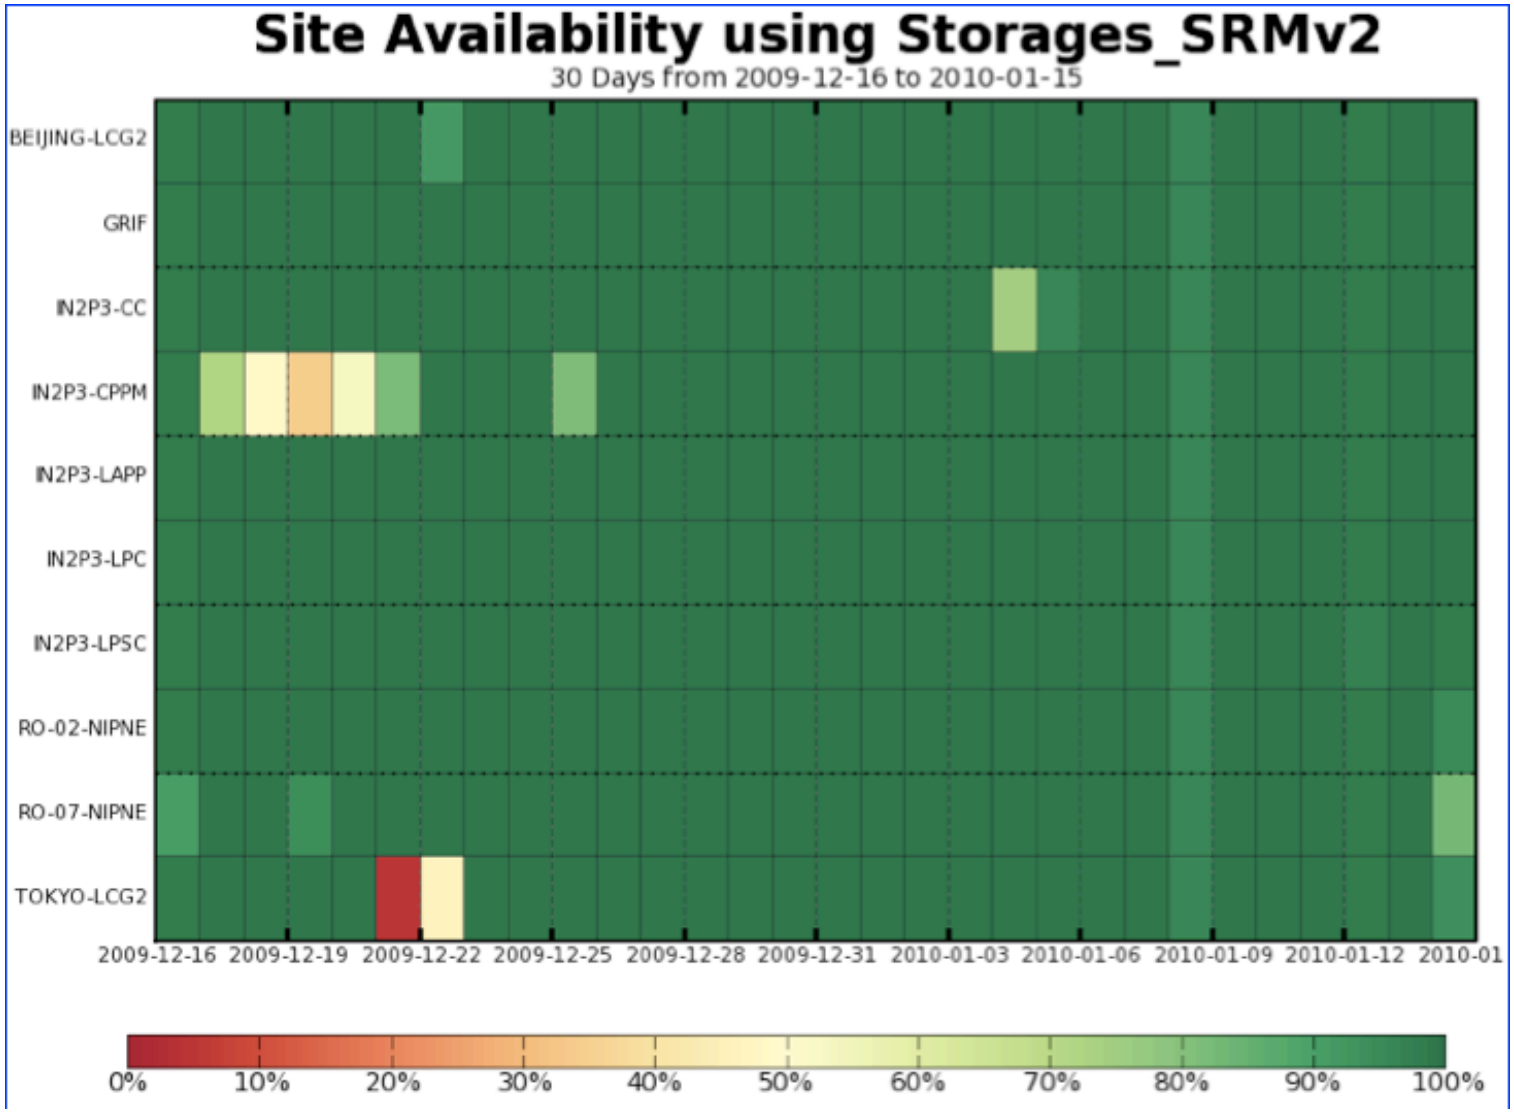

Pilot Factories

Monitoring distributed analysis

- As we saw there are two ways that users run jobs in your clusters:
  - Panda backend (pilot based)
    - Easily traceable with Panda monitoring
  - WMS backend
    - Not easily traceable
- There are continuous tests where performance of both backends can be examined:
  - Gangarobot
  - Hammercloud
- Site exclusion policy will gather GR and GC stats and compute site availability based on the results
  - Sites not achieving targets will be excluded from DA activity

GR and HC

- http://gangarobot.cern.ch/hc/

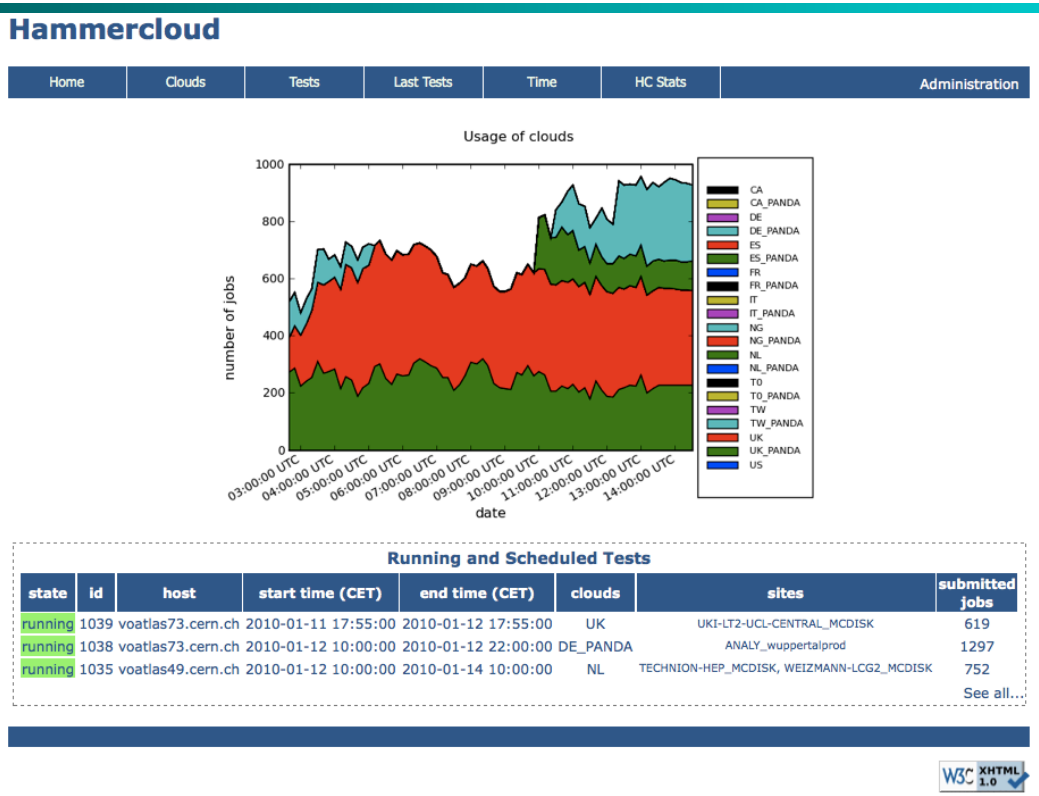

- Daily Gangarobot tests
  - http://gangarobot.cern.ch/

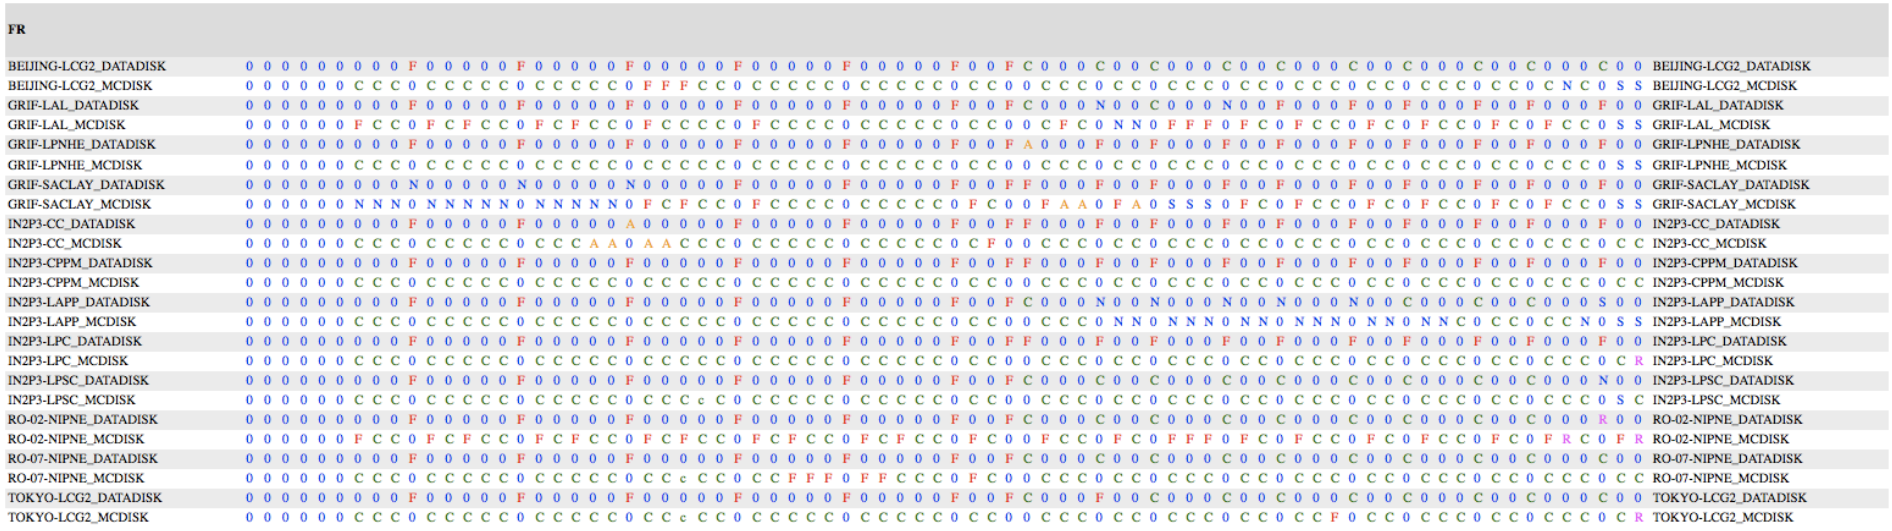

GR and HC

Hammercloud

Home | Clouds | Tests | Last Tests | Time | HC Stats | Administration

all the tests

state	id	host	clouds	start time (CET)	end time (CET)	sites	submitted jobs
running	1039	voatlas73.cern.ch	UK	2010-01-11 17:55:00	2010-01-12 17:55:00	1	619
running	1038	voatlas73.cern.ch	DE_PANDA	2010-01-12 10:00:00	2010-01-12 22:00:00	1	1297
completed	1037	voatlas73.cern.ch	DE_PANDA	2010-01-11 17:00:00	2010-01-12 05:01:04	12	3556
completed	1036	voatlas73.cern.ch	UK_PANDA	2010-01-11 15:31:00	2010-01-12 14:31:38	11	13744
running	1035	voatlas49.cern.ch	NL	2010-01-12 10:00:00	2010-01-14 10:00:00	2	752
completed	1034	voatlas73.cern.ch	UK	2010-01-09 15:57:00	2010-01-09 23:57:01	1	113
completed	1033	voatlas49.cern.ch	NL	2010-01-09 15:00:00	2010-01-09 15:01:50	1	0
completed	1032	voatlas49.cern.ch	FR_PANDA	2010-01-09 18:00:00	2010-01-10 18:00:06	13	845
completed	1031	voatlas49.cern.ch	FR_PANDA	2010-01-08 17:00:00	2010-01-09 17:00:09	13	845
completed	1030	voatlas73.cern.ch	UK_PANDA	2010-01-08 11:00:00	2010-01-08 23:01:29	2	4984
completed	1029	voatlas49.cern.ch	NL	2010-01-08 12:00:00	2010-01-08 12:01:51	1	0
completed	1028	voatlas49.cern.ch	TW_PANDA	2010-01-08 10:46:00	2010-01-08 10:47:49	2	0
completed	1027	voatlas49.cern.ch	TW_PANDA	2010-01-08 10:42:00	2010-01-08 10:44:39	2	0
completed	1026	voatlas49.cern.ch	ES_PANDA	2010-01-08 09:20:00	2010-01-09 09:15:09	7	2110
completed	1025	voatlas49.cern.ch	IT_PANDA	2010-01-07 11:04:00	2010-01-08 10:04:10	4	1486
completed	1023	voatlas49.cern.ch	IT_PANDA	2010-01-07 11:04:00	2010-01-08 10:04:10	4	1655
completed	1022	voatlas73.cern.ch	DE	2009-12-30 22:00:00	2009-12-31 02:00:04	12	619
completed	1021	voatlas73.cern.ch	DE	2009-12-30 19:15:00	2009-12-30 23:15:01	12	647
completed	1020	voatlas73.cern.ch	DE_PANDA	2009-12-30 19:20:00	2009-12-30 23:20:04	12	909
completed	1019	voatlas49.cern.ch	UK	2009-12-25 09:19:00	2009-12-26 09:19:15	2	5253
completed	1018	voatlas49.cern.ch	UK	2009-12-24 09:19:00	2009-12-25 09:19:12	2	5171
completed	1017	voatlas49.cern.ch	UK	2009-12-23 09:19:00	2009-12-24 09:19:18	2	4820
draft	1016	voatlas49.cern.ch	UK	2009-12-23 08:56:00	2009-12-24 08:56:00	2	0
completed	1015	voatlas49.cern.ch	UK_PANDA	2009-12-21 16:30:00	2009-12-22 16:31:15	2	4908
completed	1014	voatlas49.cern.ch	IT_PANDA	2009-12-19 16:03:00	2009-12-20 15:00:05	1	104
completed	1012	voatlas73.cern.ch	DE	2009-12-18 14:30:00	2009-12-18 20:30:22	2	454
completed	1011	voatlas73.cern.ch	DE_PANDA	2009-12-18 10:30:00	2009-12-18 22:30:08	1	1225
completed	1010	voatlas73.cern.ch	DE_PANDA	2009-12-17 18:30:00	2009-12-18 06:33:58	13	34886
completed	1009	voatlas73.cern.ch	DE	2009-12-17 17:30:00	2009-12-17 22:30:32	2	407
completed	1006	voatlas49.cern.ch	FR_PANDA	2009-12-18 16:00:00	2009-12-19 04:00:00	4	260

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 Next »

Want to modify running tests? Contact dan or johannes to create an account.

W3C XHTML 1.0

GR and HC

Summary

state	id	host	clouds	start time (CET)	end time (CET)	submitted jobs
completed	1032	voatlas49.cern.ch	FR_PANDA	2010-01-09 18:00:00	2010-01-10 18:00:06	845

Input type: PANDA
Output DS: user09.JohannesElmsheuser.ganga.sitetest_squid_55DPD.FR.20091209.1.[sitename]
Input DS Patterns: data09_cos.*RPCwBeam.merge.DPD_MUONCOMM.f*_m*
Ganga Job Template: /data/gangarobot/hammercloud/inputfiles/CosmicsAnalysis/CosmicsAnalysis_panda.tpl
Athena User Area: /data/gangarobot/hammercloud/inputfiles/CosmicsAnalysis/CosmicsAnalysis_1550.tar.gz
Athena Option file: /data/gangarobot/hammercloud/inputfiles/CosmicsAnalysis/MuAnalysis_topOptions_FRONTIER_SERVER_CLEARED.py

View Test Directory (for debugging)

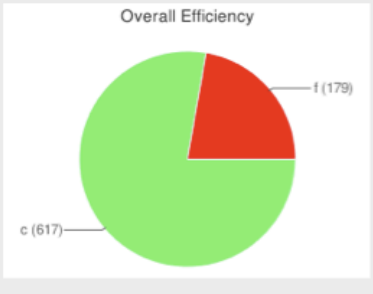

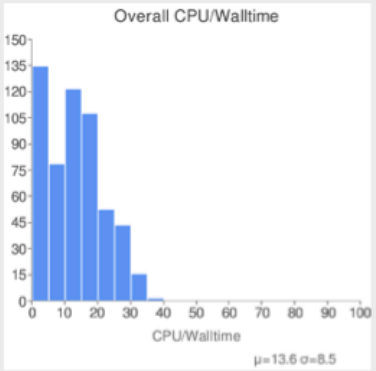

Sites

Site	Submitted jobs	Running jobs	Completed jobs	Failed jobs	Num datasets per bulk	Min queue depth	Max running jobs	Resubmit enabled	Resubmit force
ANALY_LYON_DCACHE	0	1	0	64	55	55	55	no	no
ANALY_LPC	0	0	64	1	55	55	55	no	no
ANALY_LYON	0	0	0	65	55	55	55	no	no
ANALY_LAPP	0	0	65	0	55	55	55	no	no
ANALY_BEIJING	0	1	64	0	55	55	55	no	no
ANALY_ROMANIA02	45	0	18	2	55	55	55	no	no
ANALY_TOKYO	0	2	63	0	55	55	55	no	no
ANALY_GRIF-IRFU	0	0	64	1	55	55	55	no	no
ANALY_LPSC	0	0	65	0	55	55	55	no	no
ANALY_GRIF-LPNHE	0	0	65	0	55	55	55	no	no
ANALY_ROMANIA07	0	0	62	3	55	55	55	no	no
ANALY_CPPM	0	0	65	0	55	55	55	no	no
ANALY_GRIF-LAL	0	0	22	43	55	55	55	no	no
Total	45	4	617	179					

Boosting and tuning your site for distributed analysis

- And the best, site resp can schedule its own tests for site/cloud debugging or explore new configs to boost performance
  - Using HammerCloud scheduler, choosing backend and access patterns:
    - Panda/WMS
    - Copy-To-WN: the WN script pre-copies the input files prior to starting athena.
    - FileStager: use athena's FileStager module to copy the input files in a background thread.
    - Direct: use the local protocol to directly open the input files
      - ➡ To determine which data access method to use at a given site, pay attention to the CPU/Wallclock results during HammerCloud tests. In general the access method which uses the most CPU will be the optimal at the site.

Boosting and tuning your site for distributed analysis

- Collect your stats and publish them here:
  - https://twiki.cern.ch/twiki/bin/view/Atlas/HammerCloudDataAccess

FR cloud

Most recent tests are 662 (Panda/Copy-to-WN), 652/656 (FS) and 653/657 (direct)

Site	Data Access Method	Reason (comment and/or link to HC tests)	in Ganga IS	in Panda IS	in AGIS
BEIJING	?	insufficient recent tests. (60% w/ FS though)			
GRIF-LAL	FileStager	FS 83%			
GRIF-LPNHE	?	Copy 52% no testing w/ FS			
GRIF-SACLAY/IRFU	FileStager	70% FS vs. 51% Copy			
IN2P3-CPPM	unclear	Copy 33% vs 32% FS			
IN2P3-LAPP	FileStager	Copy 58% vs. 82% FS			
IN2P3-LPC	unclear	16% copy vs. 11% FS			
IN2P3-LPSC	?	67% copy vs. no FS testing?			
RO-02-NIPNE	FileStager	FS 49% vs. Copy 34%			
RO-07-NIPNE	?	insufficient recent tests			
TOKYO	FileStager	Copy 52% vs. FS 72% vs. direct 38%			

Updated by DanielVanDerSter on 12 Oct 2009

Scheduling and configuring HC tests

Hammercloud

You are connected as xavier, click here to Logout!

Home | Clouds | Tests | Last Tests | Time | HC Stats | Administration

Ask for an account !

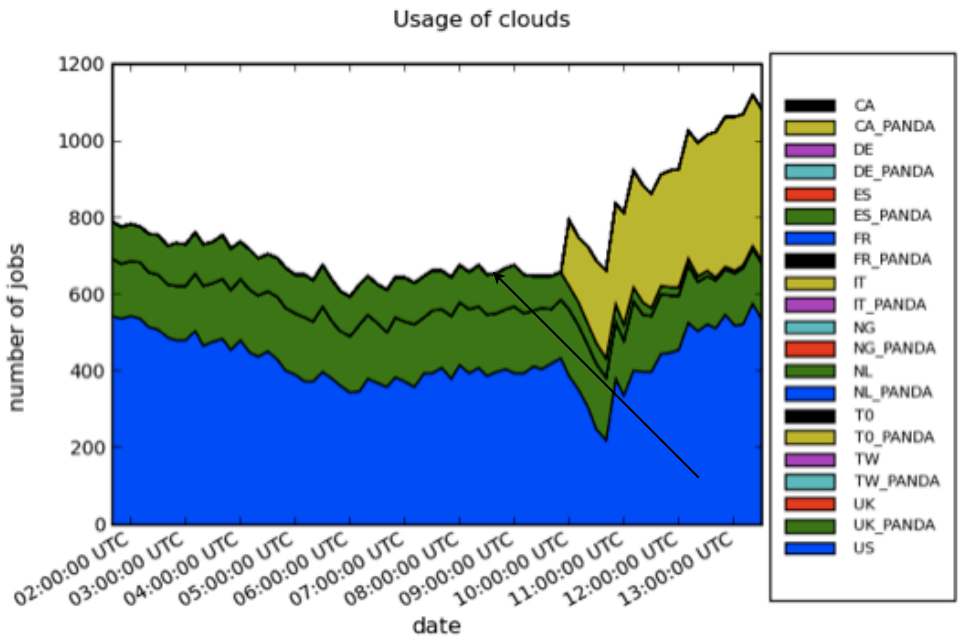

Running and Scheduled Tests

state	id	host	start time (CET)	end time (CET)	clouds	sites	submitted jobs
running	1043	voatlas49.cern.ch	2010-01-13 10:00:00	2010-01-14 22:01:00	IT	INFN-FRASCATI_MCDISK, INFN-NAPOLI-ATLAS_MCDISK, INFN-ROMA1_MCDISK, 1 more...	1027
running	1041	voatlas73.cern.ch	2010-01-12 17:30:00	2010-01-13 17:30:00	US	ANALY_MWT2	7015
running	1040	voatlas49.cern.ch	2010-01-12 18:00:00	2010-01-14 18:00:00	UK_PANDA	ANALY_MANC	2449
running	1035	voatlas49.cern.ch	2010-01-12 10:00:00	2010-01-14 10:00:00	NL	TECHNION-HEP_MCDISK, WEIZMANN-LCG2_MCDISK	986

See all...

Scheduling and configuring HC tests

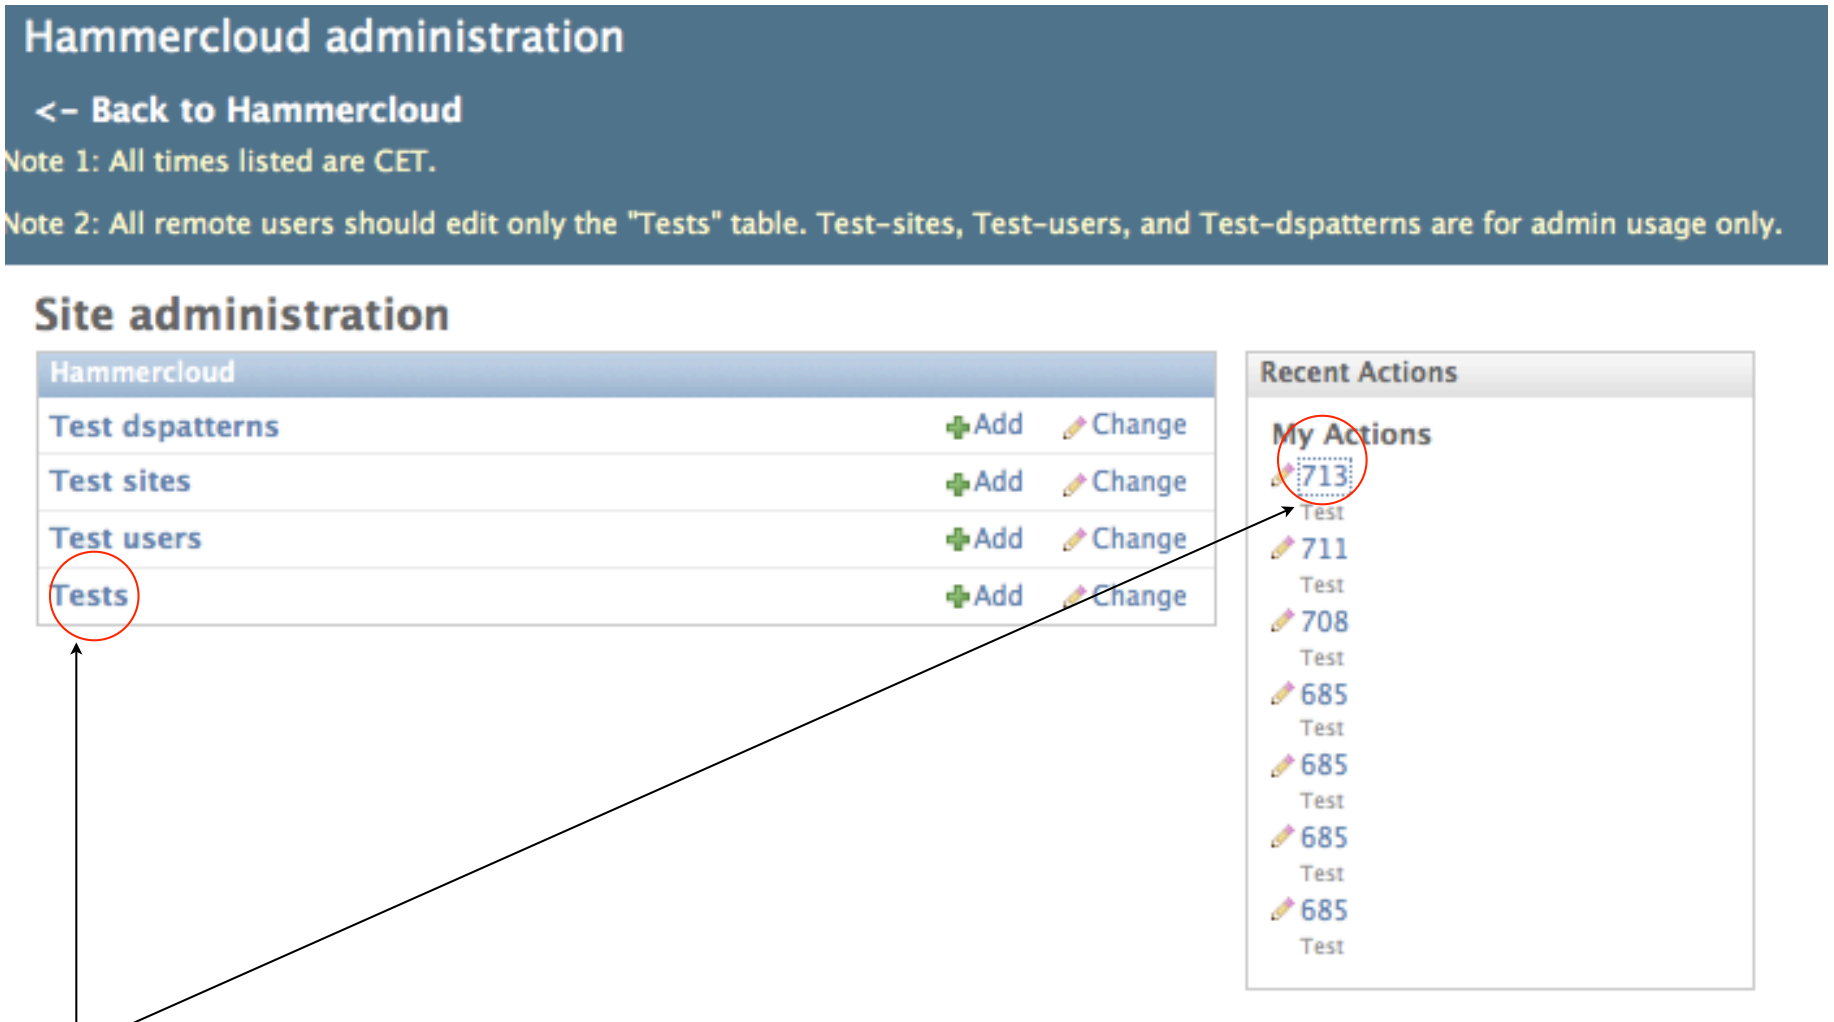

Easiest way is to edit and modify an already scheduled test and adapt it for your site

Scheduling and configuring HC tests

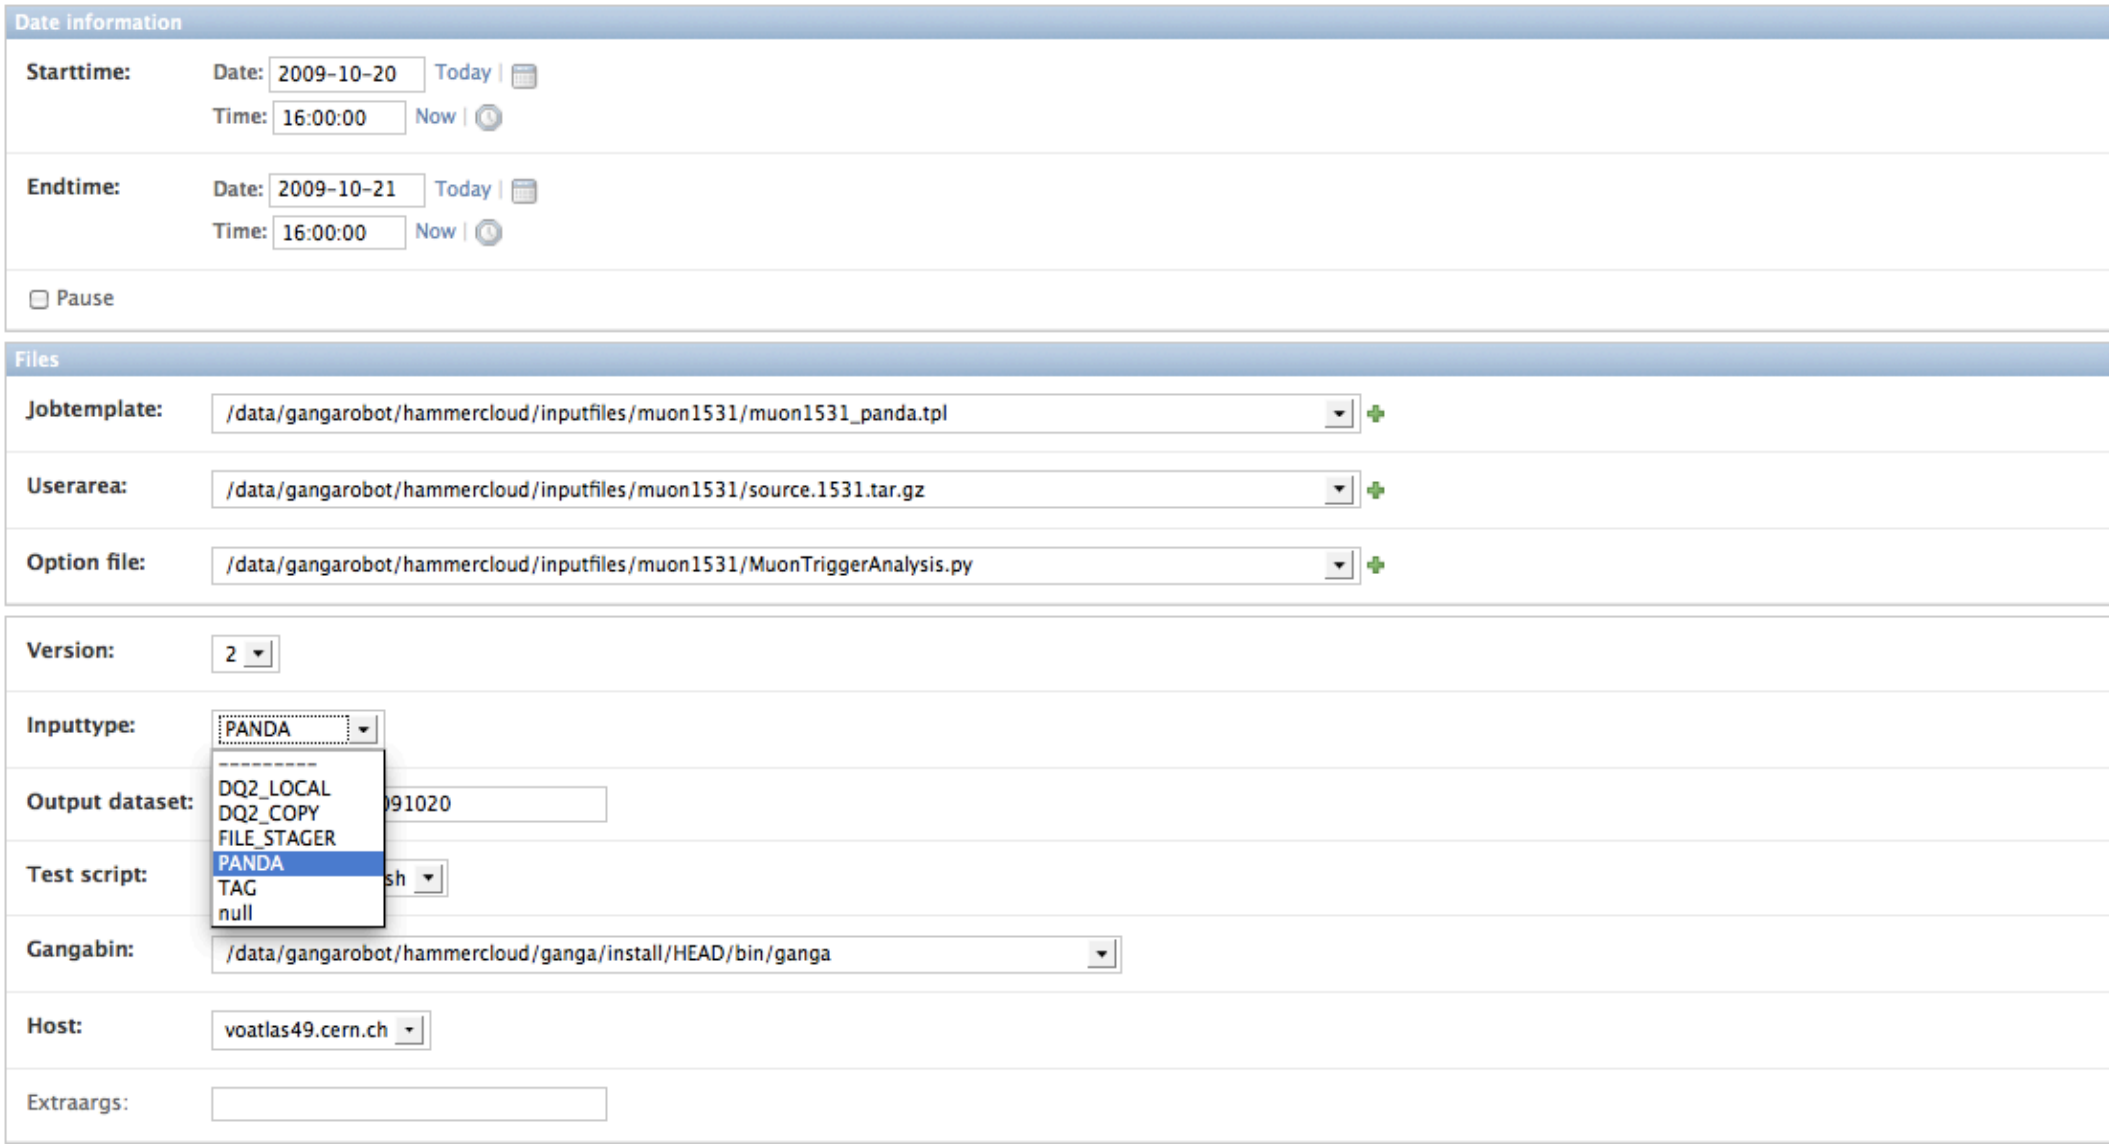

Scheduling and configuring HC tests

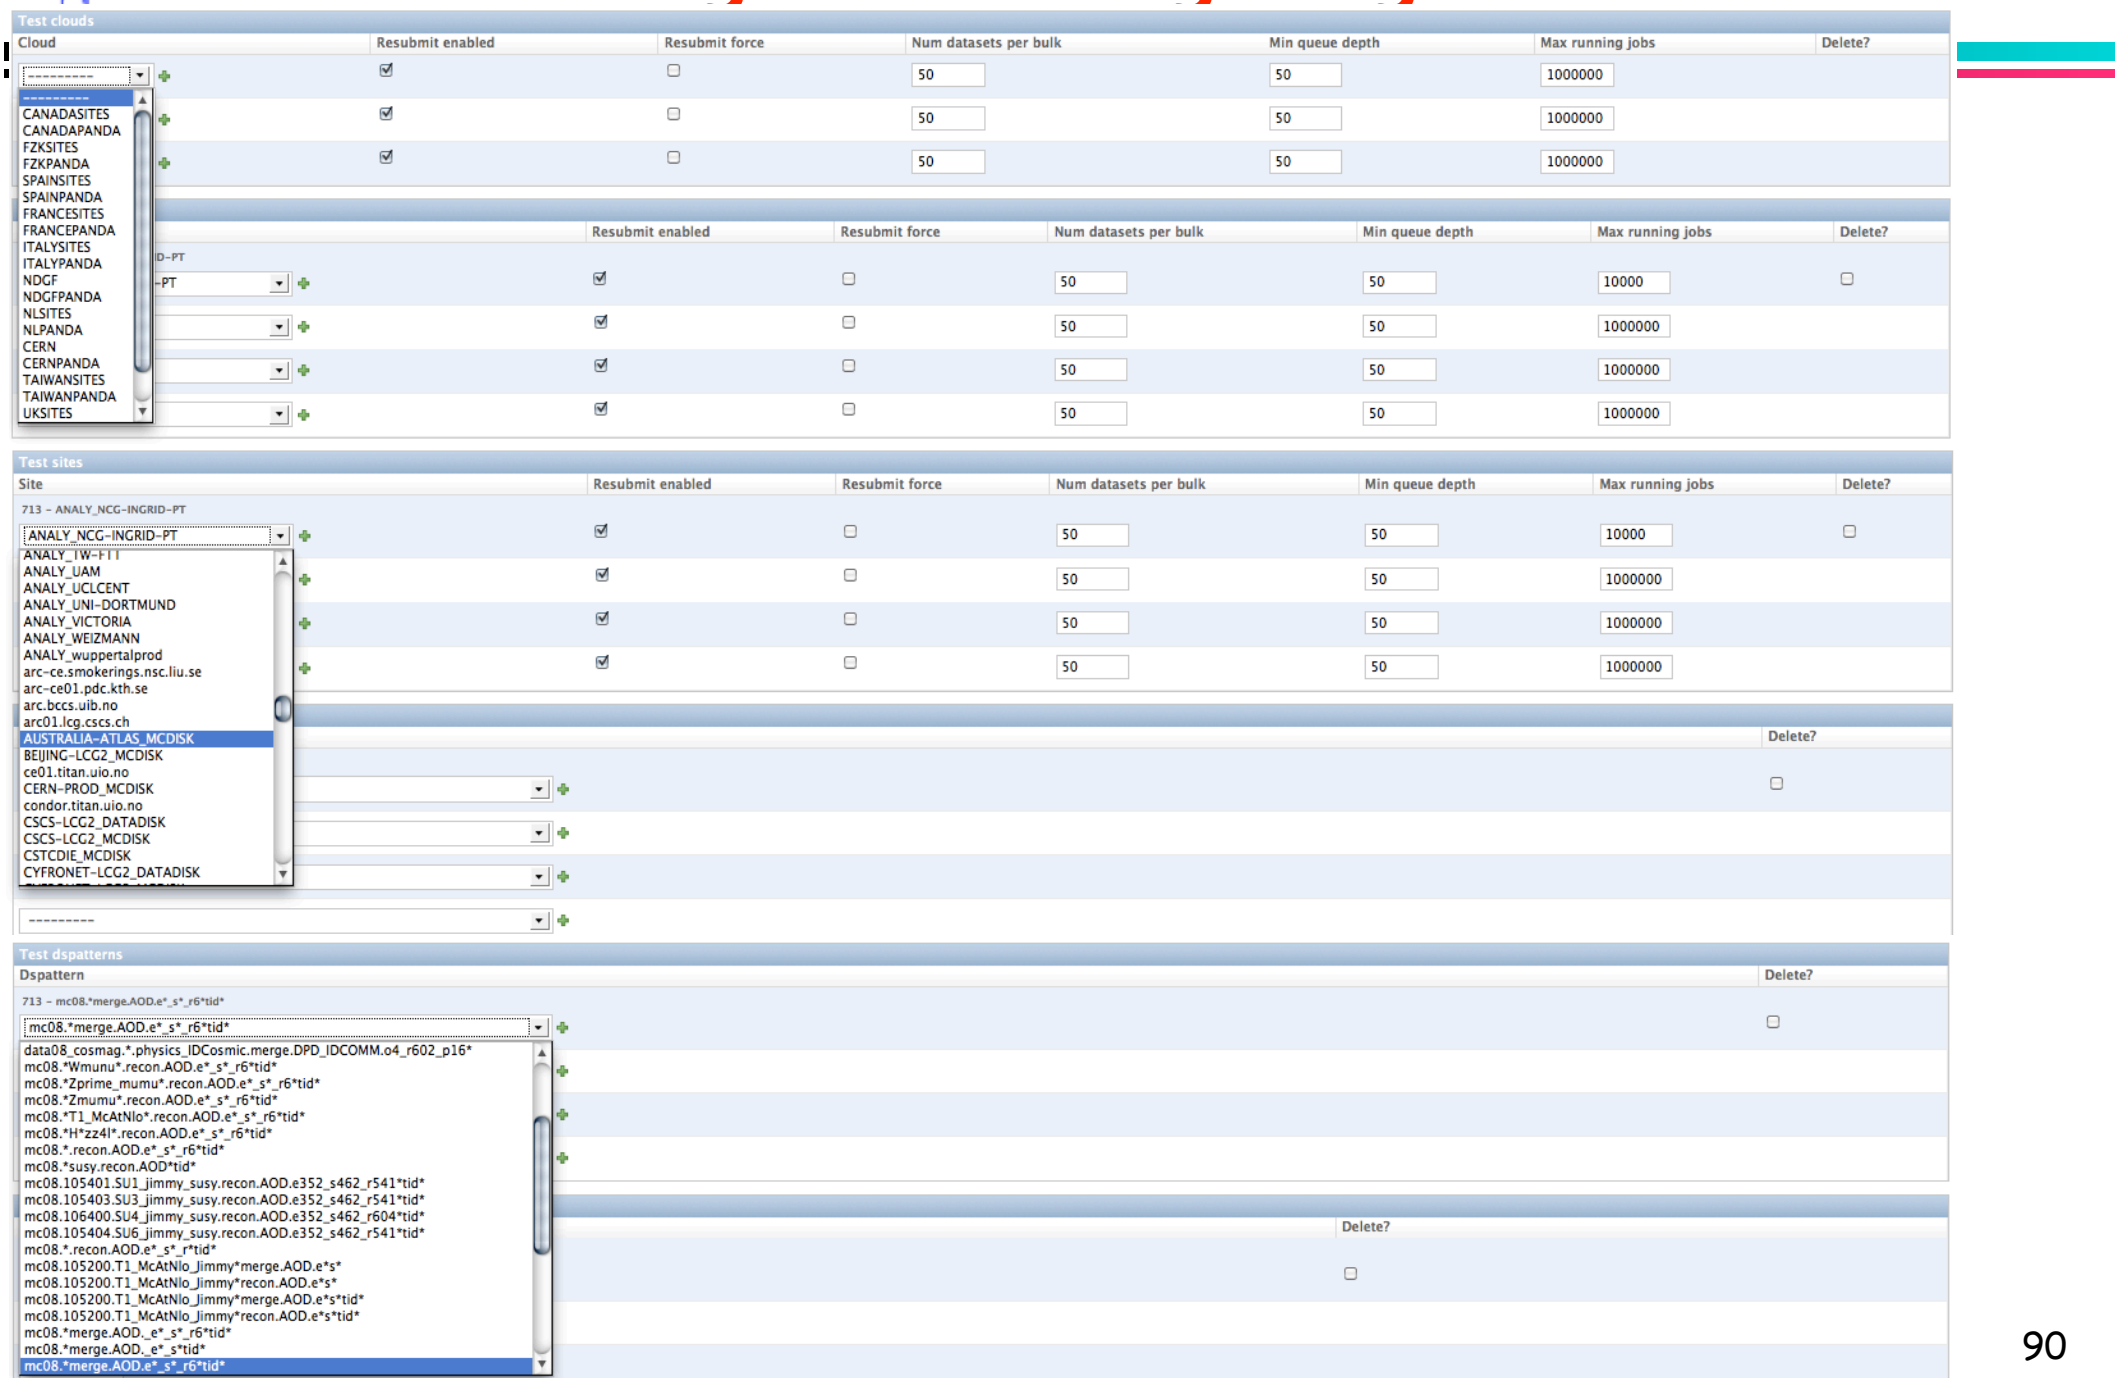

Scheduling and configuring HC tests

By state
All
draft
unapproved
tobescheduled
scheduled
submitting
running
completed
error

Action: --------- Go

Clone multiple tests
Clone single test
Delete selected tests
Send selected tests for approval

Id		time	State	Clouds	Sites	Inputtyp
104		9, 2010, 9:15 a.m.	draft	ES_PANDA	ANALY_IFAE, ANALY_IFIC, 6 more...	PANDA
104		14, 2010, 10:01 p.m.	running	IT	INFN-FRASCATI_MCDISK, INFN-NAPOLI-ATLAS_MCDISK, 2 more...	FILE_STAGER
1042	Jan. 12, 2010, 6 p.m.	Jan. 13, 2010, 6 a.m.	completed	US	ANALY_MWT2	PANDA
1041	Jan. 12, 2010, 5:30 p.m.	Jan. 13, 2010, 5:30 p.m.	running	US	ANALY_MWT2	PANDA
1040	Jan. 12, 2010, 6 p.m.	Jan. 14, 2010, 6 p.m.	running	UK_PANDA	ANALY_MANC	PANDA
1039	Jan. 11, 2010, 5:55 p.m.	Jan. 12, 2010, 5:55 p.m.	completed	UK	UKI-LT2-UCL-CENTRAL_MCDISK	FILE_STAGER
1038	Jan. 12, 2010, 10 a.m.	Jan. 12, 2010, 10 p.m.	completed	DE_PANDA	ANALY_wuppertalprod	PANDA
1037	Jan. 11, 2010, 5 p.m.	Jan. 12, 2010, 5:01 a.m.	completed	DE_PANDA	ANALY_CSCS, ANALY_CYF, 11 more...	PANDA
1036	Jan. 11, 2010, 3:31 p.m.	Jan. 12, 2010, 2:31 p.m.	completed	UK_PANDA	ANALY_BHAM, ANALY_CAM, 10 more...	PANDA
1035	Jan. 12, 2010, 10 a.m.	Jan. 14, 2010, 10 a.m.	running	NL	TECHNION-HEP_MCDISK, WEIZMANN-LCG2_MCDISK, 1 more...	FILE_STAGER
1034	Jan. 9, 2010, 3:57 p.m.	Jan. 9, 2010, 11:57 p.m.	completed	UK	UKI-LT2-UCL-CENTRAL_MCDISK	FILE_STAGER
1033	Jan. 9, 2010, 3 p.m.	Jan. 9, 2010, 3:01 p.m.	completed	NL	IL-TAU-HEP_SCRATCHDISK	FILE_STAGER
1032	Jan. 9, 2010, 6 p.m.	Jan. 10, 2010, 6 p.m.	completed	FR_PANDA	ANALY_BEIJING, ANALY_CPPM, 12 more...	PANDA
1031	Jan. 8, 2010, 5 p.m.	Jan. 9, 2010, 5 p.m.	completed	FR_PANDA	ANALY_BEIJING, ANALY_CPPM, 12 more...	PANDA
1030	Jan. 8, 2010, 11 a.m.	Jan. 8, 2010, 11:01 p.m.	completed	UK_PANDA	ANALY_GLASGOW, ANALY_QMUL, 1 more...	PANDA
1029	Jan. 8, 2010, noon	Jan. 8, 2010, 12:01 p.m.	completed	NL	IL-TAU-HEP_SCRATCHDISK	FILE_STAGER
1028	Jan. 8, 2010, 10:46 a.m.	Jan. 8, 2010, 10:47 a.m.	completed	TW_PANDA	ANALY_TAIWAN, ANALY_TW-FTT, 1 more...	PANDA
1027	Jan. 8, 2010, 10:42 a.m.	Jan. 8, 2010, 10:44 a.m.	completed	TW_PANDA	ANALY_TAIWAN, ANALY_TW-FTT, 1 more...	PANDA
1026	Jan. 8, 2010, 9:20 a.m.	Jan. 9, 2010, 9:15 a.m.	completed	ES_PANDA	ANALY_IFAE, ANALY_IFIC, 6 more...	PANDA
1025	Jan. 7, 2010, 11:04 a.m.	Jan. 8, 2010, 10:04 a.m.	completed	IT_PANDA	ANALY_INFN-FRASCATI, ANALY_INFN-MILANO-ATLASC, 3 more...	PANDA
1023	Jan. 7, 2010, 11:04 a.m.	Jan. 8, 2010, 10:04 a.m.	completed	IT_PANDA	ANALY_INFN-FRASCATI, ANALY_INFN-MILANO-ATLASC, 3 more...	PANDA
1022	Dec. 30, 2009, 10 p.m.	Dec. 31, 2009, 2 a.m.	completed	DE	CSCS-LCG2_DATADISK, CYFRONET-LCG2_DATADISK, 11 more...	DQ2_LOCAL
1021	Dec. 30, 2009, 7:15 p.m.	Dec. 30, 2009, 11:15 p.m.	completed	DE	CSCS-LCG2_DATADISK, CYFRONET-LCG2_DATADISK, 11 more...	DQ2_LOCAL
1020	Dec. 30, 2009, 7:20 p.m.	Dec. 30, 2009, 11:20 p.m.	completed	DE_PANDA	ANALY_CSCS, ANALY_CYF, 11 more...	PANDA
1019	Dec. 25, 2009, 9:19 a.m.	Dec. 26, 2009, 9:19 a.m.	completed	UK	UKI-LT2-QMUL_MCDISK, UKI-LT2-RHUL_MCDISK, 1 more...	FILE_STAGER
1018	Dec. 24, 2009, 9:19 a.m.	Dec. 25, 2009, 9:19 a.m.	completed	UK	UKI-LT2-QMUL_MCDISK, UKI-LT2-RHUL_MCDISK, 1 more...	FILE_STAGER
1017	Dec. 23, 2009, 9:19 a.m.	Dec. 24, 2009, 9:19 a.m.	completed	UK	UKI-LT2-QMUL_MCDISK, UKI-LT2-RHUL_MCDISK, 1 more...	FILE_STAGER
1016	Dec. 23, 2009, 8:56 a.m.	Dec. 24, 2009, 8:56 a.m.	draft	UK	UKI-LT2-QMUL_MCDISK, UKI-LT2-RHUL_MCDISK, 1 more...	FILE_STAGER

Monitoring DA in Panda

- Basically same structure as for production/reprocessing
  - http://panda.cern.ch:25980/server/pandamon/query?dash=analysis
- Stats exclusively for panda backend in DA tools
- Make sure your factories feed the farms:
  - Special role for Panda based analysis: /atlas/Role=pilot

Monitoring DA in Panda

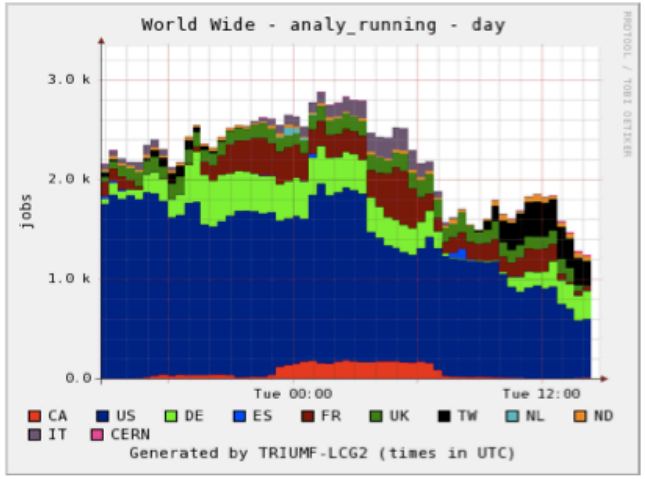

Analysis Summary By Cloud

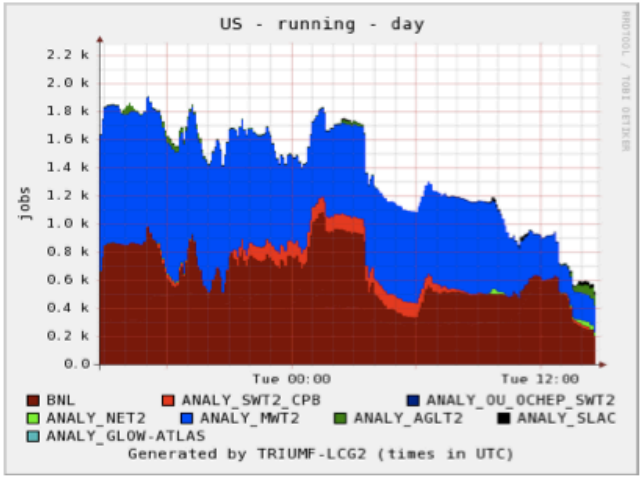

Analysis Summary By Site

Analysis job summary, last 24 hours (Details: errors, nodes) pathena analysis queue status

Cloud Information	Job Nodes	Jobs	Latest	Pilot Nodes	defined	assigned	waiting	activated	sent	running	holding	transferring	finished	failed	tot	trf	other
Overall Analysis	5804	27072	01-12 14:31	18445	1645 / 0	1 / 0	0	18966 / 0	1 / 0	1228 / 0	141 / 0	116 / 0	42077 / 0	27072 / 0	39%	13%	26%
CA	221	50	01-12 14:31	622	14	0	0	244	1	3	20	0 / 0	2296	50	2%	0%	2%
CERN	18	4	01-12 14:29	1091	0	0	0	0	0	1	8	0 / 0	5	4	44%	44%	0%
DE	1494	929	01-12 14:31	4015	44	1	0	331	0	266	52	60 / 0	4753	929	16%	6%	10%
ES	201	465	01-12 14:30	1022	0	0	0	1	0	1	0	0 / 0	273	465	63%	35%	28%
FR	1100	2443	01-12 14:30	3649	211	0	0	77	0	91	10	0 / 0	2016	2443	55%	29%	25%
IT	104	7	01-12 14:22	78	2	0	0	109	0	1	0	0 / 0	808	7	1%	1%	0%
ND	185	988	01-12 14:31	242	0	0	0	11697	0	42	18	56 / 0	1907	988	34%	10%	24%
NL	59	662	01-12 14:24	378	65	0	0	61	0	0	1	0 / 0	38	662	95%	34%	61%
TW	225	606	01-12 14:31	1176	36	0	0	30	0	235	4	0 / 0	242	606	71%	24%	48%
UK	1391	430	01-12 14:28	2663	84	0	0	84	0	0	0	0 / 0	13223	430	3%	0%	3%
US	806	20488	01-12 14:31	3509	1189	0	0	6332	0	588	28	0 / 0	16516	20488	55%	17%	39%

ATLAS releases

- Releases are automatically installed by the ATLAS-LCG installation machinery
  - But some time the installation fails and your site can be missing a needed release
    - Errors found in simulation or analysis jobs
  - Site resp can trigger release installation and validation here:
    - https://atlas-install.roma1.infn.it/atlas_install/protected/rai.php

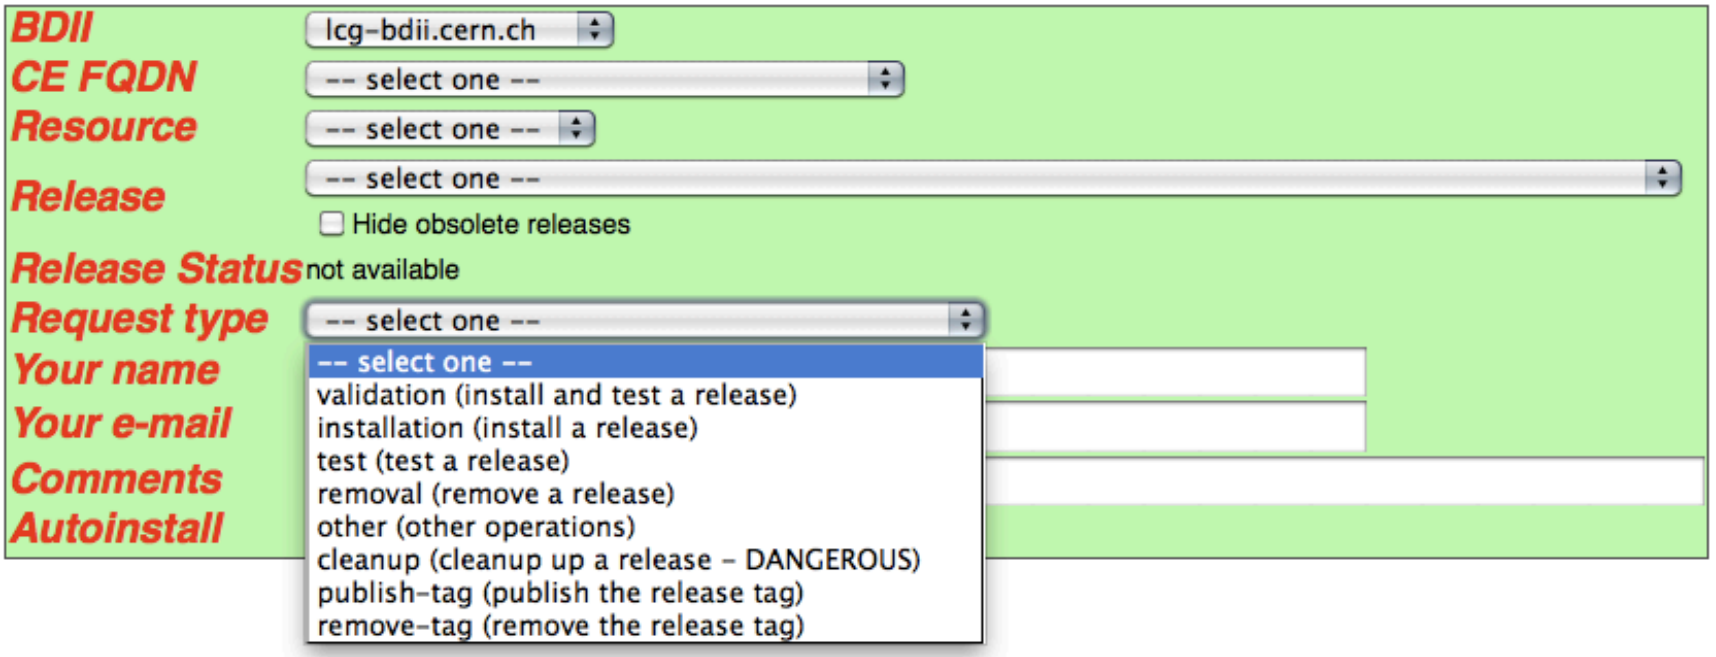

ATLAS releases

- ATLAS offline software release status:
  - ➡ http://atlas-computing.web.cern.ch/atlas-computing/projects/releases/status/

Atlas Offline Software Release Status

The Nightly Build Summary page summarizes the nightly builds including the special purpose and migration nightly builds. Maintained by Atlas.Release@cern.ch

Latest Production Releases

Release	Patch Release	Usage
15.6.1	AtlasProduction 15.6.1.6	Physics validation and SLC5 compatibility tests.
15.5.4	AtlasTier0 15.5.4.11	Collisions and Cosmics Dec'09 and reprocessing
	AtlasTier0 15.5.4.20	Trigger reprocessing
15.5.3	AtlasP1HLT 15.5.3.10	Cosmic data taking Oct-Nov'09
15.3.1	AtlasProduction 15.3.1.12	Physics validation and SLC5 compatibility tests.
	AtlasProduction 15.3.1.20	Special cache for ESD to ESD reprocessing.
14.2.25	AtlasProduction 14.2.25.10	Simulation production.
	AtlasProduction 14.2.25.11	Overlay reconstruction.

Legacy Production Releases

Release	Patch Release	Usage
15.5.0	AtlasP1HLT 15.5.0.2	Cosmics data taking.
	AtlasTier0 15.5.0.2	Tier0 reconstruction of combined cosmics.
15.4.0	AtlasP1HLT 15.4.0.6	Cosmics data taking.
	AtlasTier0 15.4.0.4	Tier0 reconstruction of combined cosmics.
	AtlasTier0 15.4.0.2	Tier0 reconstruction of combined cosmics.
15.3.0	AtlasProduction 15.3.0.2	Validation of reconstruction - reprocessing of the mc08 and processing of the mc09 samples.
15.2.0	AtlasProduction 15.2.0.1	Physics validation
	AtlasP1HLT 15.2.0.10	Cosmics data taking.
	AtlasTier0 15.2.0.14	Tier0 reconstruction of combined cosmics.
	AtlasTier0 15.2.0.16	Tier0 reconstruction of combined cosmics.
15.1.0	AtlasProduction 15.1.0.8	Validation
15.0.1	AtlasTier0 15.0.1.8	Validation re-reprocessing of 2008 cosmics data and early comissioning runs.
15.0.0	AtlasProduction 15.0.0.4	MC09 Event generation and Geant4.9.2 preliminary validation.
	AtlasTier0 15.0.0.7	Validation re-reprocessing of 2008 cosmics data and early comissioning runs.
	AtlasP1HLT 15.0.0.1	TDAQ Technical Run.
14.5.2	AtlasProduction 14.5.2.6	Cosmics re-reprocessing.
	AtlasProduction 14.5.2.12	Cosmics simulation.
14.5.1	AtlasProduction 14.5.1.6	LUCID & Cosmics Simulation.
14.5.0	AtlasProduction 14.5.0.7	Reprocessing of cosmics data at the Tier0's.
	AtlasTier0 14.5.0.8	HLT reprocessing.

Release Status and Plans

CondDB access

Introduction

- Software in Athena is completely intependent of database technology
  - The COOL and CORAL packages (developed by the LCG Applications Area) interface ATLAS software to different technologies
    - Oracle, SQLite, MySQL, FroNTier/Squid etc.
- The geometry and most of conditions data are stored in the Oracle database
  - COOL provides tools to define Intervals of Validity (IoV) for each record, and a hierarchical tagging of folders
- Some payloads (LAr calibrations and InDet alignments) are too large to be stored efficiently in Oracle
  - They are stored as POOL files and referenced from the Oracle database
- The master Oracle offline database, and the master copy of POOL files, are at CERN
  - The contents of the Oracle database is replicated in real time (~15 min max delay) to all 10 Tier-1s, which have read-only copies
  - The POOL files are replicated to all Tier-1s (and soon Tier-2s) by DDM

DB access use case and technologies

- Major use cases:
  - Simulation production
    - Tier-1s, Tier-2s, Tier-3s
  - Tier-0 processing
    - Tier-0
  - Reprocessing
    - Tier-1s and some Tier-2s
  - Calibration and alignments
    - Tier-1s, CAF, calibrations centres
  - Group and user analysis
    - Everywhere
- Available technologies:
  - Direct access to Oracle databases
    - Good for jobs running at CERN and Tier-1s but possible overload problems
  - (Conditions) DB release
    - Best for production tasks needing fixed conditions
  - FroNTier/Squid (access to Oracle databases through web caches)
    - FroNTier cache in front of Oracle server helps with load problems and local (Squid) cache solves latency problems for jobs running at Tier-2/3s
  - DB-on-demand (SQLite extraction of DB data)
    - Useful to run on "disconnected laptops"

Recommended strategies

- Simulation production and reprocessing tasks
  - The (C)DB release has been in use satisfactorily for several years
    - The DB release is an extraction of the needed constants (SQLite DB and POOL files) into a tar file that is copied to the local worker node and accessed locally
    - To reduce local network traffic, the Conditions DB release for reprocessing separates SQLite files with the conditions for a given run, so that only the required files are copied
  - Direct access to Oracle databases is also functional but much less efficient in CPU usage, especially when contacting remote Oracle servers
- Tier-0 processing
  - Tier-0 jobs are released only after a hand-shake with the Oracle database to make sure that conditions data for the run to be processed are available
  - Job submission is throttled to avoid overloading the Castor servers
    - This protects the Oracle server from overloads too
    - POOL files are stored on afs at CERN so there is no access problem
  - In case of need (foreseen downtime of Oracle at CERN) a CDB release can be used as back-up solution

Recommended strategies

- Calibration and alignment tasks
  - Tasks run on Tier-0, CAF or Tier-1s: use direct access to the Oracle databases
  - Tasks run on Tier-2s and calibration centres: access Oracle databases through FroNTier/Squid caches
- Group and user analysis tasks
  - For reasons of uniformity and portability, these jobs should access Oracle databases through FroNTier/Squid caches
    - In this way jobs can run anywhere with no configuration changes
    - POOL files will now be replicated also to Tier-2s
  - Jobs running at Tier-1s can always access Oracle databases directly
    - This possibility is functional but slower for remote jobs
  - People wishing to run in disconnected mode ("travelling laptop") can use the DB-on-demand tool to take a snapshot of DB information and POOL files they need and use it locally
    - Care should be taken to avoid proliferation of these snapshots with different contents

Need for an efficient remote access

- CondDB data is stored at the Tier-1 ORACLE RACs
  - Data is replicated via ORACLE streams technology from Tier-0 (CERN) to all Tier-1s
    - Having problems in a single Tier-1 can degrade the CondDB data replication
      - ➡ Which could happen during reprocessing campaigns and was already observed
- Found that accessing ORACLE from WAN is 20 times slower than local access
  - Long resolution times when accessing CondDB from remote sites
    - http://indico.cern.ch/getFile.py/access?contribId=9&sessionId=12&resId=1&materialId=slides&confId=50976
- Users do also need an efficient access to the DB:
  - Running against Event Data requires access to conditions data that is NOT included with the DB release.
    - The software distribution scheme only distributes the conditions data needed to process the simulated data.
    - Not all the conditions data is in Oracle. Most calibrations data are in POOL files which are referenced from within Oracle and must be locally on the site running the job (too large to fit in Oracle DB: expected 500GB/year)
      - Pool files would need to be copied now to T2s

Remote oracle access

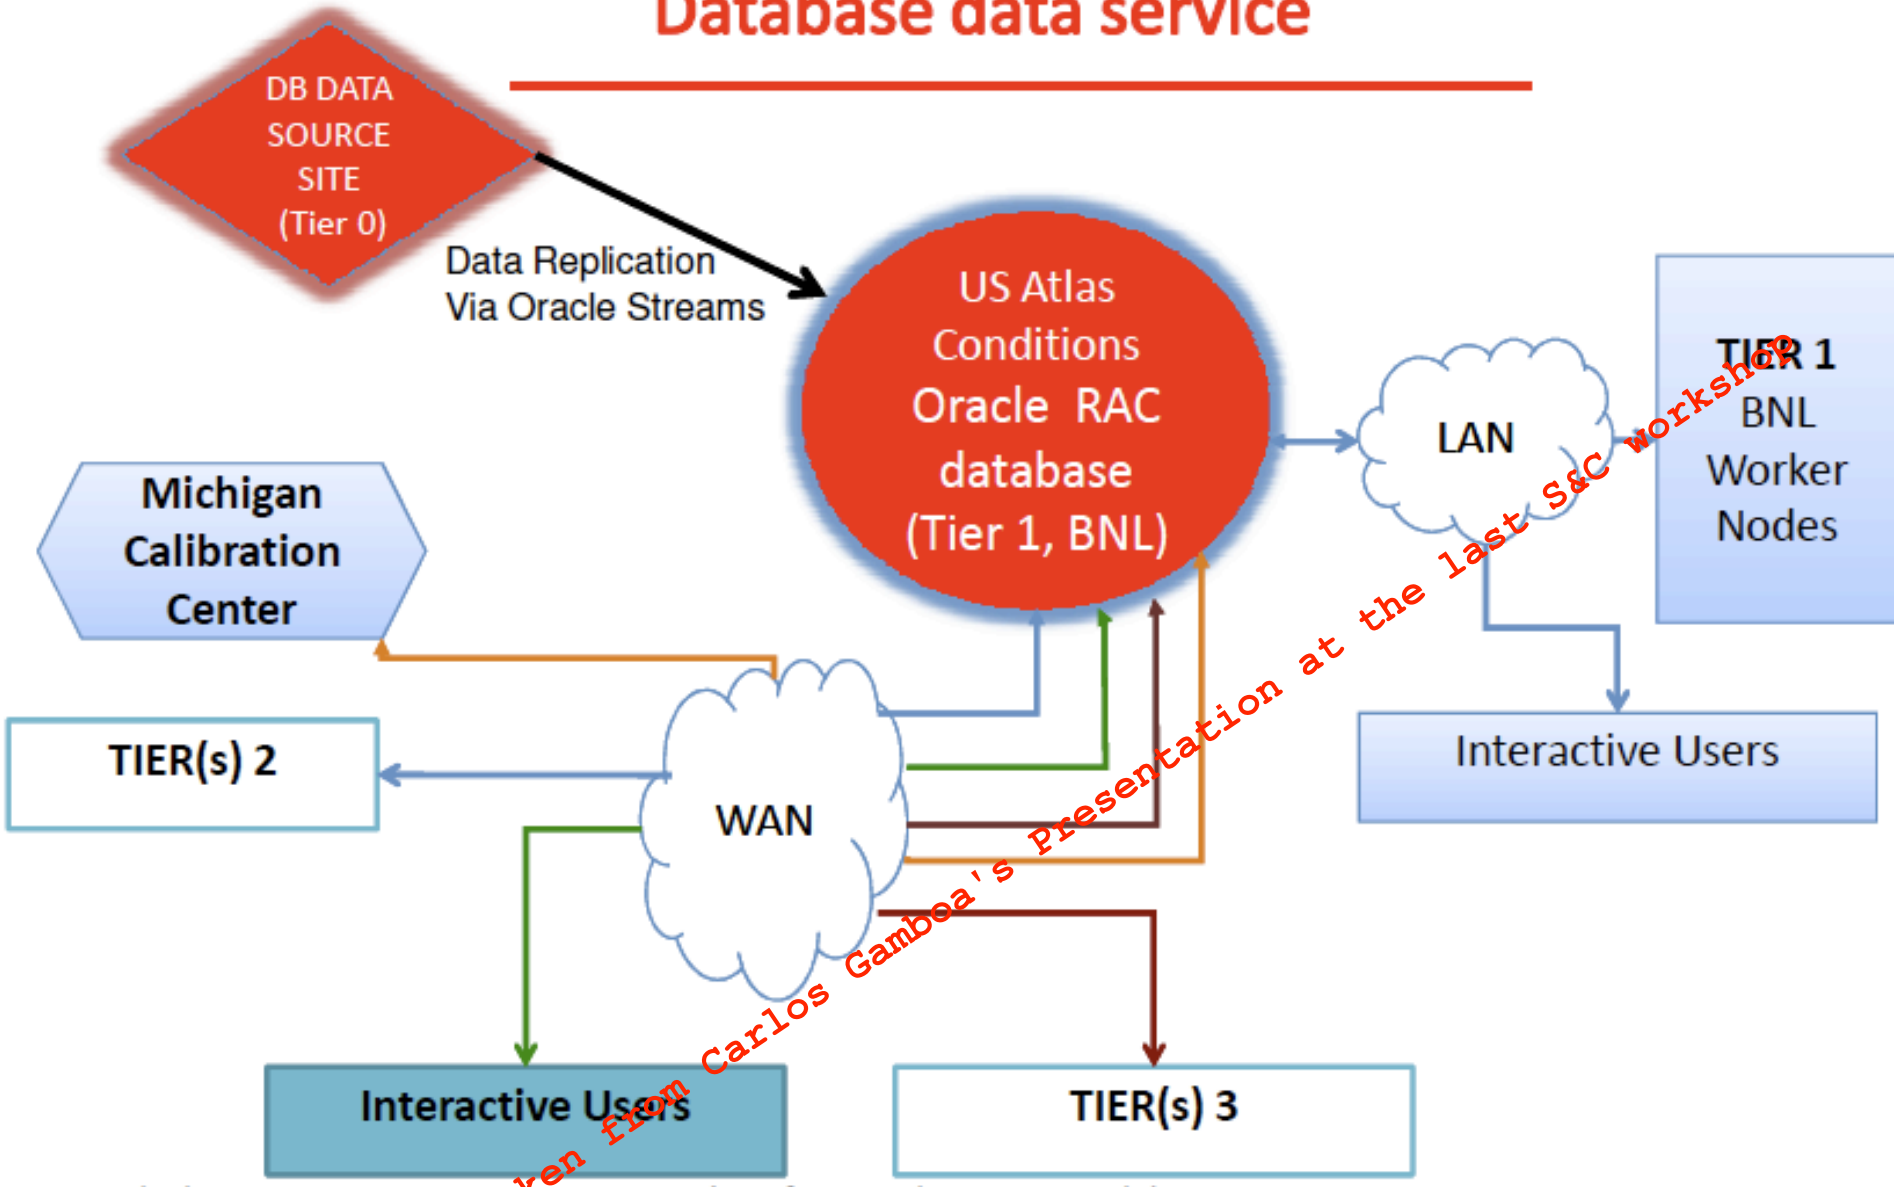

Squid/FroNTier

- Having one (or some FroNTier) and squid servers at T1 and T2 would provide much efficient remote CondDB access
  - FroNTier make caching of ORACLE queries
  - Squids do local caching connecting with Frontier when new queries are placed

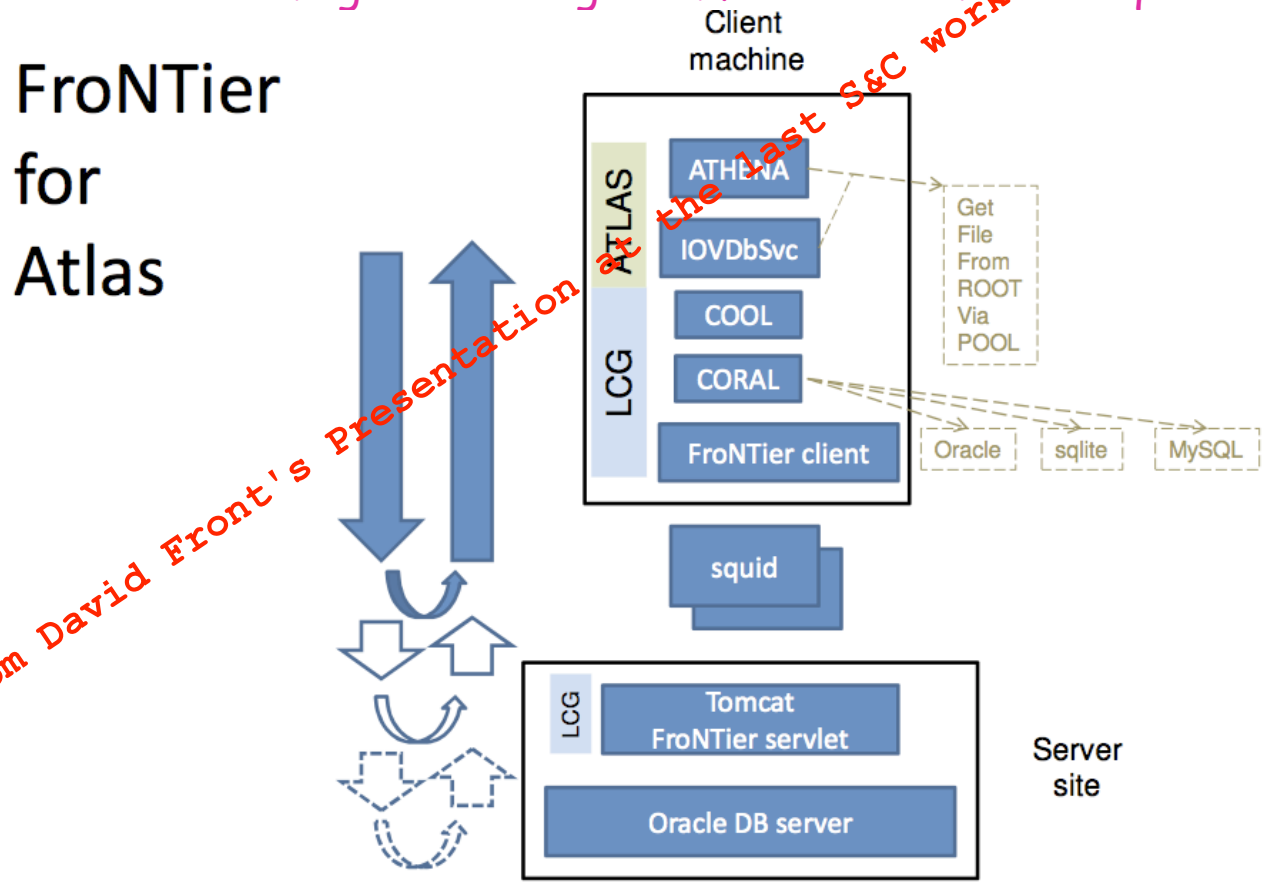

Taken from David Front's Presentation at the last S&C workshop

Site configuration

- Frontier Servers
  - A site can be configured to have primary and failover Frontier servers. Therefore the T2 Squid needs to permit caching from either. This does not set which Frontier will be used; just which can be used. Since the Squid will be configured to only accept certain clients, i.e. the WNs, it is not necessary to restrict which sites can be cached. This simplifies the configuration, and allows the squid to be used for other purposes (fetch-crl is good one).
- The current list of ATLAS frontier servers is:-
  - http://atlassq1-fzk.gridka.de:8021/fzk
  - http://atlasfrontier1.cern.ch:?/?
  - http://squid-frontier.usatlas.bnl.gov:23128/frontieratbnl * currently no access control so will work also direct from WNs
  - http://atlfrontier.pic.es:3128/pic-frontier
  - http://lcgft-atlas.gridpp.rl.ac.uk:3128/frontierATLAS
  - http://frontier.triumf.ca:3128/ATLAS_frontier
- FZK and TRIUMF only allow authorized clients, so please send your squid server names/IPs to avoid disappointment.
  - We need some assignment of clouds to primary and backup Frontier servers. CERN Frontier will be used for some scaling tests.
- The frontier client can be configured to failover to a backup Squid proxy, in case of a problem with the primary. The backup must accept connections from the remote WN`s, so this is a little bit of bookkeeping work to setup, but I think it`s worth it for redundancy. Each site should pair up with a network near partner and exchange WN IP ranges(probably NATs). Some T2's may choose not to deploy the Squid at this time, perhaps due to not having the hardware available. In which case, they should agree with a network-near T2 to use theirs, and presumably the same backup.

Site configuration

- The boxes correspond to:
- Squid server installed and configured. Local test successful.
- T1 Frontier launchpad accepts T2 Squid, WN, local desktop IP ranges
- The client can failover to directly access the Frontier launchpad. If this happens a lot it could be a problem, but I think we should allow it.
- Run the fnget test from each client group, with and without the http_proxy set.
- Same for back-up Squid.
- Local script in $VO_ATLAS_SW_DIR/local/setup.sh
- to configure FRONTIER_SERVER env, override liddcap.so, add libshift, LD_PRELOAD for dpm
- creates PFC to CD flat files in HOTDISK (no HOTDISK, no tick!)
- It is created by the SW install tool. Green tick here means green on status
- Panda schedconfig so that this script is sourced.
- Single panda test jobs accessing Db. Either Ganga robot, or manual readreal submission(below) will do.

CA Show...

DE Show...

ES Show...

FR Hide

FR	Catherine&Ghita						
Site	**Squid**	**T1 auth**	**fnget test**	**backup**	**setup.sh**	**schedconfig**	**1 job**
BEIJING-LCG2	http://atlassq.ihep.ac.cn:3128						
GRIF-IRFU							
GRIF-LAL	http://grid15.lal.in2p3.fr:3128	✓	✓				
GRIF-LPNHE							
GRIF-SACLAY							
IN2P3-CC							
IN2P3-CC-T2							
IN2P3-CPPM							
IN2P3-LAPP	http://lapp-squid01.in2p3.fr:3128	✓	✓				
IN2P3-LPC							
IN2P3-LPSC							
NIPNE							
TOKYO-LCG2	http://conddb-px01.icepp.jp:3128	✓	✓				

IT Show...

ND Show...

NL Show...

TW Show...

UK Show...

US Show...

Clouds	Frontier	Backup
CA	TRIUMF	BNL
US,TW	BNL	TRIUMF
DE,NL,ND	FZK	PIC
ES,FR,IT	PIC	FZK
UK	RAL	PIC

Criticality of services

MoU SLA

- T1

Service	*Maximum delay in responding to operational problems*			*Average availability[2] measured on an annual basis*	
	Service interruption	Degradation of the capacity of the service by more than 50%	Degradation of the capacity of the service by more than 20%	During accelerator operation	At all other times
Acceptance of data from the Tier-0 Centre during accelerator operation	12 hours	12 hours	24 hours	99%	n/a
Networking service to the Tier-0 Centre during accelerator operation	12 hours	24 hours	48 hours	98%	n/a
Data-intensive analysis services, including networking to Tier-0, Tier-1 Centres outwith accelerator operation	24 hours	48 hours	48 hours	n/a	98%
All other services[3] – prime service hours	2 hour	2 hour	4 hours	98%	98%
All other services[3] – outwith prime service hours[6]	24 hours	48 hours	48 hours	97%	97%

http://lcg.web.cern.ch/LCG/MoU/MoU_basics/Annex3_min_membership.doc

MoU SLA

- T2

Service	*Maximum delay in responding to operational problems*		*Average availability[2] measured on an annual basis*
	Prime time	*Other periods*	
End-user analysis facility	2 hours	72 hours	95%
Other services[3]	12 hours	72 hours	95%

Site exclusion policy

The seed...

- ATLAS is composed by 80+ sites
- ATLAS is running a broad range of computing activities
- ADC-experts and shifters' duty is to spot problems and communicate it to sites
- Operations were swamped by minor things that ATLAS site people could cure

A bit of context...

- ATLAS asked clouds two years ago to setup an specialized support group at every cloud
  - Composed by ATLAS contacts at the sites (atlas-cloud-support-<cloud>)
    - Usually DDM issues were routed to this groups
- Two months ago ATLAS asked to form cloud-squads, taking profit from the previous support group but asking more implication to the daily activities and take care of ongoing site problems that need specific actions from site people, f.i. missing/lost/corrupted files
  - That squads were formed within Savannah infrastructure
- Due to the not-infinite manpower of ATLAS operations and the start of data taking, ADC operations must make sure things are working and protect experts and shifters from overkills
  - Site exclusion policy started to be defined
- The goals of site exclusion are:
  - Sites need to be stable and provide quality of service to the whole collaboration
  - Operations people should concentrate on major problems and left to the sites the solution and the validation of site-specific problems

Site exclusion procedure

- Once a site problem or degradation is noticed, shifters do report as usual:
  - File GGUS and eLog
- Shifters then pre-evaluate the severity checking a set of defined metrics per activity and contact ADC experts if site qualifies for it
  - ADC expert will decide eventual site exclusion/blacklisting
- When site is excluded or blacklisted, the cloud-squad is contacted. The ticket will contain:
  - Link to the GGUS
  - Information about site status:
    - Explicitly telling if site is excluded from some ADC activity or blacklisted
- Once these steps has been followed by ADC people, responsibility is passed to the site. From then on they are responsible for:
  - Acknowledgement and fix of the problem
  - Guarantee site performance (succeed in wLCG and ATLAS SAM tests)
  - Answer back GGUS and cloud-squad ticket so ATLAS can do assessment

Site exclusion procedure (2)

- Assessment is based on metrics using the ADC testing activities:
  - DDMFT/SRMFT, AFT and DPFT
- Once the metrics are passed, site is set back into the ADC operation activities.
- Shifters/experts do log the exclusion/blacklisting in ATLAS downtime Google Calendar
  - ... and also log inclusion in testing activities once the site claim they are ready

Site exclusion metrics

- Site degradation:
  - Data Transfer (DT), site excluded if:
    - DDMFT/SRMFT efficiency <80% (48h)
    - Data Placement efficiency <80% (48h)
  - Distributed Analysis (DA), site excluded if:
    - AFT performance is <90% in the last 24h
      - AFT not yet ready, interim metrics:
        - GangaRobot/Panda efficiency <50% in the last 12h.
- Site failure:
  - Data Transfer (DT), site excluded if:
    - DDMFT/SRMFT efficiency is 0% (12h)
    - Data Placement efficiency is 0% (12h)
- Site blacklisting:
  - AFT fail for 5+ days
  - SRMFT fails for 3+ days
- Site excluded from Data Processing (DP) if blacklisted or excluded in DT

Site recovery metrics

- Site recovery metrics after being excluded from service activities:
  - DT: DDMFT/SRMFT running stable with efficiency >90% for 24h
  - DA: AFT running stable with efficiency >90% for 24h
  - DP: DPFT running stable with efficiency > 90% for 24h
- After blacklisting, ADC activities can be resumed only after two days of successful functional tests in DT, DA and DPFT.

Site exclusion example

Test/days of running	0	1	2	3	4	5
AFT	+	-	-	-	-	-
DDMFT	+	+	-	+	+	+
SRMFT	+	+	-	+	+	+

Activities/days of running	0	1	2	3	4	5
DA	on	on	off	off	off	off
DT	on	on	on	off	on	off
DP	on	on	on	off	on	off

→ Blacklisted AFT fails 5+ days

Summary of responsibilities

Activity	Problem routing	Exclusion/acceptance responsible(s)	Site fixing and verification	ATLAS certification
Data Transfers	Shifters file GGUS/eLog and notify experts. Experts contact cloud-squad.	Experts	cloud-squads	Shifters and Experts
Distributed Analysis	Shifters file GGUS/eLog. Notify experts if site should be blacklisted.	Handled by AFT and DAST team	cloud-squads	Handled by AFT and DAST team
Data Processing	Shifters file GGUS/eLog. Shifters/experts contact cloud-squad.	Cloud, T1: experts and DP team T2s: shifters	cloud-squads	Shifters and Experts

http://dashb-sam-atlas.cern.ch/dashboard/request.py/historicalsmryview

Impact measure

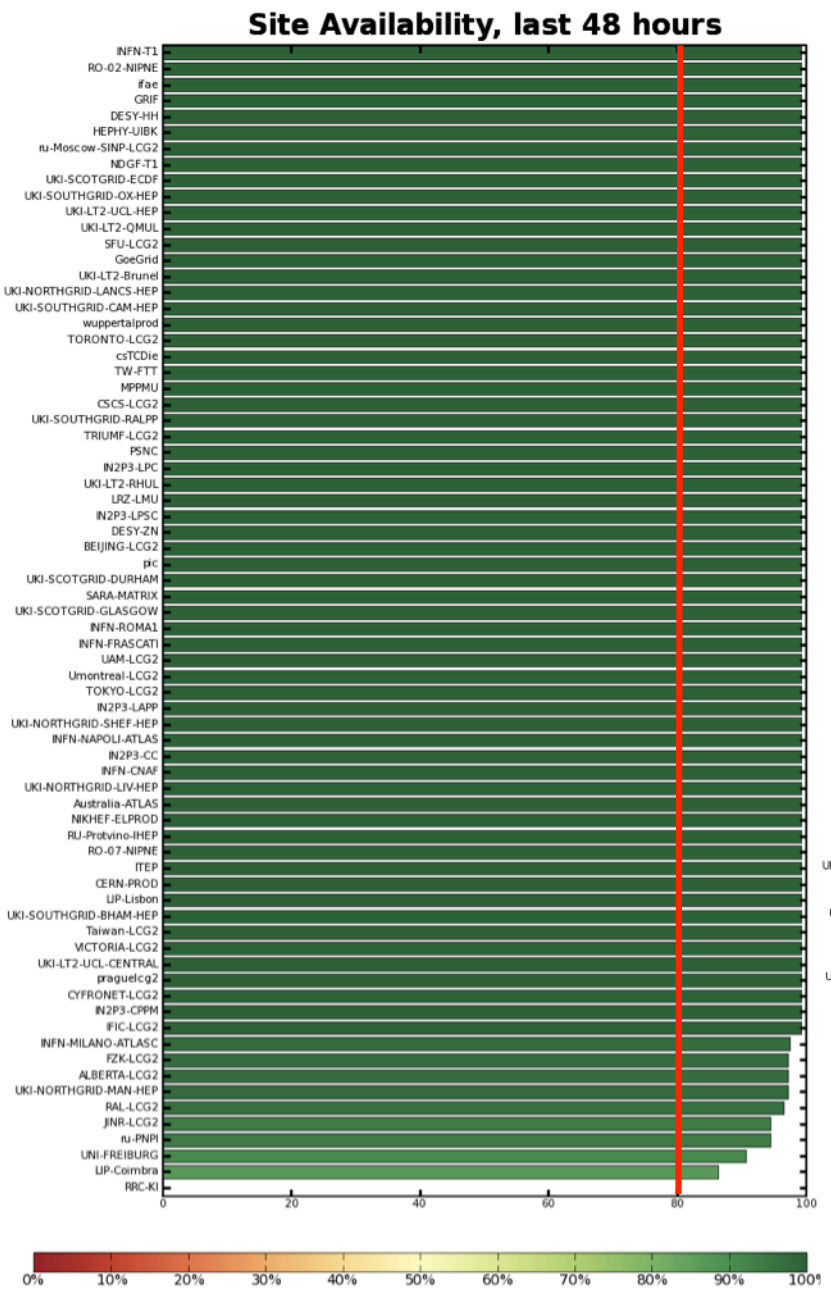

SRMv2 ATLAS SAM TESTS

DDMFT (24h): http://atladcops.cern.ch:8000/drmon/ftmon_T1-T2_matrix_day.html

Impact measure (2)

Last subscription: 30 Nov 10:47:13 | Last FC checked: 30 Nov 10:42:31 | Last transfer: 30 Nov 10:41:58

DDM FT

ASGC	TW FTT* (444/368)												
BNL	AGLT2* (535/530)	MWT2* (356/348)	NET2* (295/295)	SLACXRD* (304/304)	SWT2* (292/292)	WISC* (295/295)							
CNAF	FRASCATI* (62/55)	MILANO* (186/166)	NAPOLI* (186/166)	ROMA1* (185/172)									
FZK	CSCS* (150/150)	CYF* (150/150)	DESY HH* (442/442)	DESY ZN* (442/442)	GOEGRID* (133/133)	HEPHY UIBK* (87/87)	LRZ* (142/142)	MPPMU* (142/142)	PRAGUE* (44/44)	PSNC*	UNI* (150/150)	WUP* (149/149)	
LYON	BEIJING* (100/100)	CPPM* (50/50)	LAL* (240/240)	LAPP* (151/151)	LPC* (151/151)	LPNHE* (501/501)	NIPNE_02* (100/100)	NIPNE_07* (100/100)	SACLAY* (259/259)	TOKYO* (974/831)			
NDGFT	SE SNIC T2_DATADISK (89/89)	SIGNET* (44/44)											
PIC	IFAE* (187/187)	IFIC* (364/364)	LIP COIMBRA_DATADISK (53/36)	LIP LISBON_DATADISK (53/53)	NCG INGRID PT_DATADISK (44/39)	UAM* (187/170)							
RAL	BHAM* (35/35)	CAM* (35/35)	GLASGOW* (354/354)	LANCS* (247/247)	LIV* (105/87)	OXF* (105/93)	QMUL* (142/142)	RALPP* (179/158)	RHUL* (212/175)	SHEF* (35/35)	UKI LT2 UCL CENTRAL_DATADISK (35/35)		
SARA	CSTCDIE* (51/51)	IHEP* (151/151)	IL TAU HEP_DATADISK (51/51)	JINR* (201/201)	PNPI* (151/151)	RU MOSCOW FIAN LCG2_DATADISK (101/101)	RU MOSCOW MEPHI LCG2_DATADISK (101/101)	TECHNION HEP_DATADISK (51/51)	TR 10 ULAKBIM* (201/201)	WEIZMANN LCG2_DATADISK (51/51)			
TRIUMF	ALBERTA* (202/202)	AUSTRALIA ATLAS_DATADISK (40/40)	CA ALBERTA WESTGRID T2_DATADISK (41/0)	SFU* (578/577)	TORON* (5/5)	VICTORIA* (23/23)							

http://dashb-sam-atlas.cern.ch/dashboard/request.py/historicalsmryview

Impact measure (4)

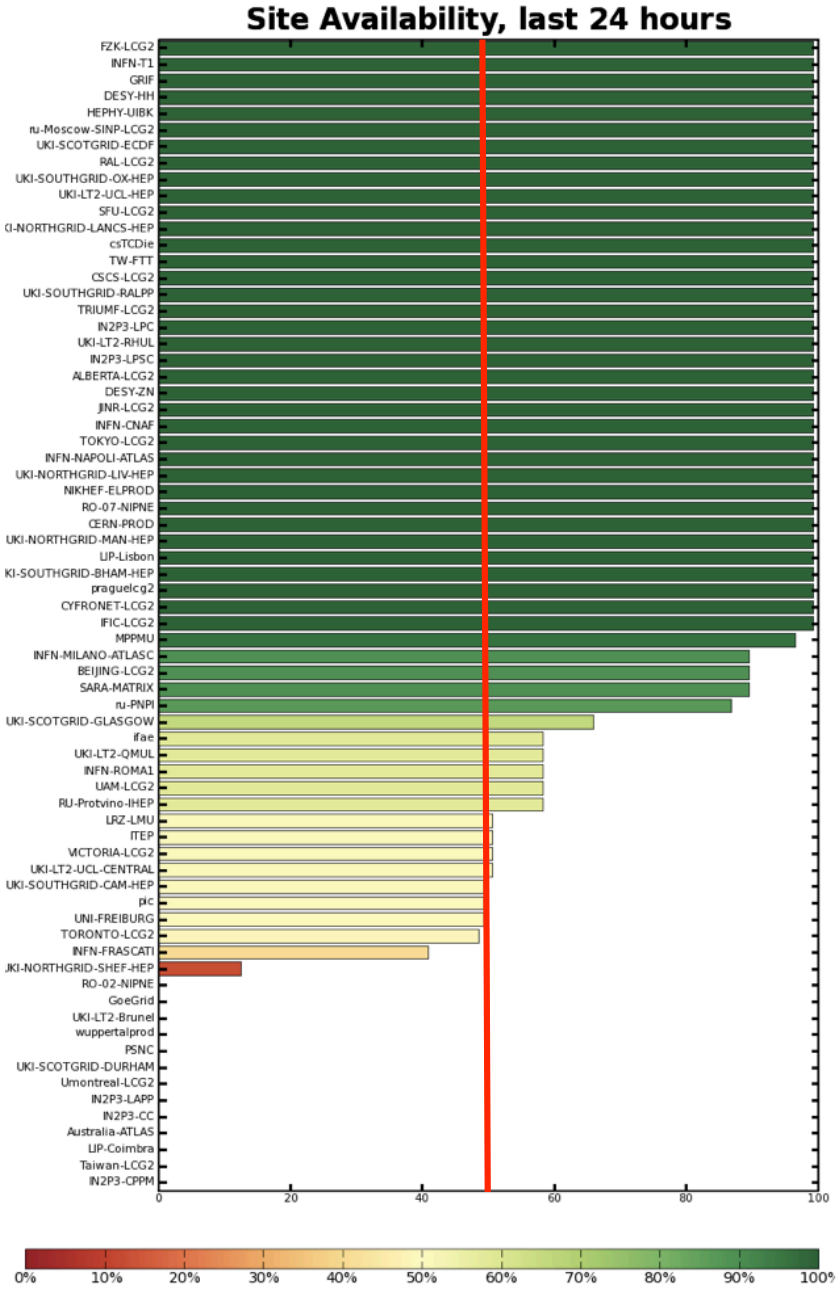

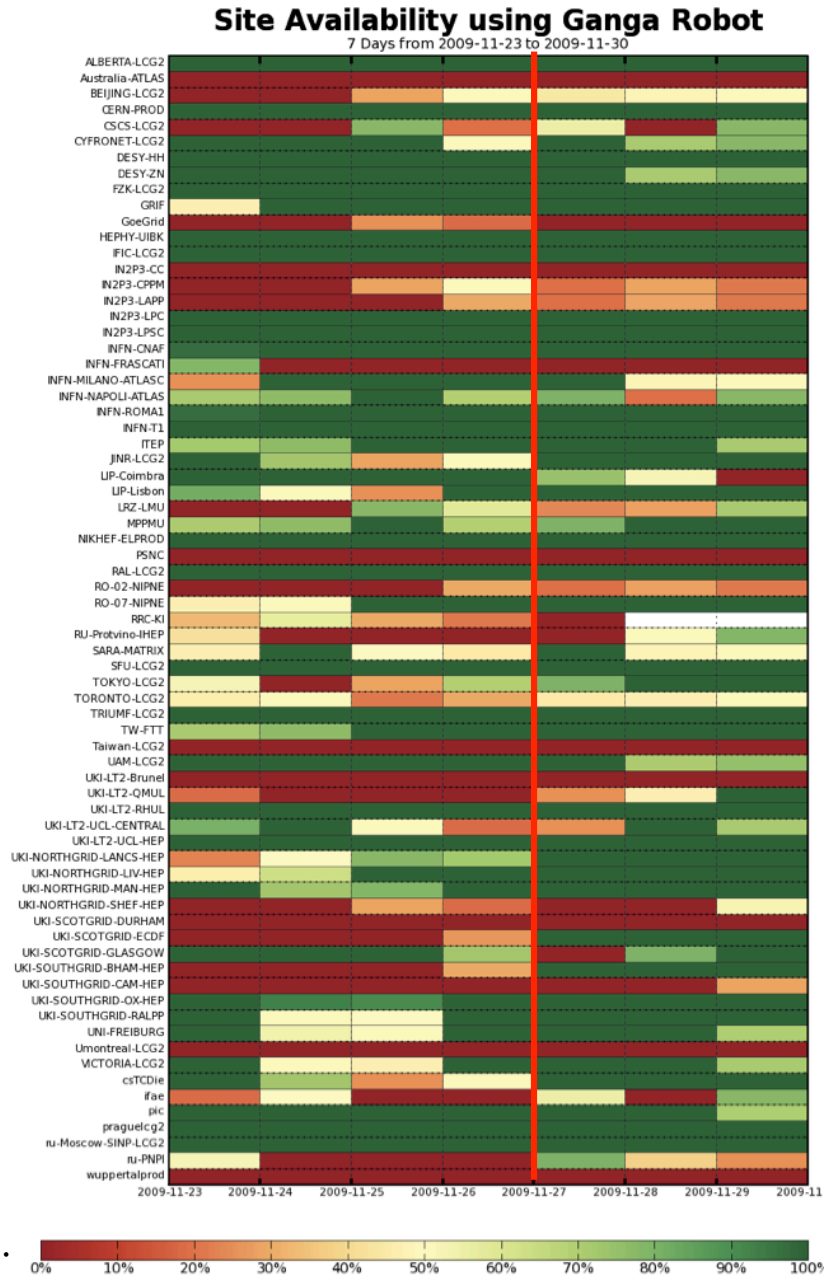

Ganga Robot (WMS)

Conclusions

- Primary goal of this policy is provide quality of service to the ATLAS community
- Excluding/blacklisting site is not a punishment but a protection against disappointment by the ATLAS scientific community
  - Ensure good behavior of data flow and user analysis
- Very simple things are tested:
  - Basic storage testing and basic analysis jobs
- Blacklisting and data transfer exclusion will be supervised by ADC experts
- Complete Analysis Functional Tests still to be implemented
  - Interim period covered by Gangarobot tests (Panda and WMS backends)
- cloud-squads become an important entity at cloud level
- Site exclusion document available at:
  - https://twiki.cern.ch/twiki/pub/Atlas/AtlasDistributedComputing/SiteExclusion_v8.pdf

Site Exclusion Policy Status

- Working group created to work on the SEP
- Most important part is to gather metrics, and represent it altogether in a single page
- We decided to use the dashboard's SSB
  - Matrix of every metric (SRM-SAM tests, DDMFT and SAM-GR) together with maintenance column and a final column where to have a binary decision about sites
- The plan is to have this tool ready by the LHC re-startup (15th February)
- Minimal version already in place
- The goal of this HLV monitoring is that:
  - Shifters and expert will spot immediately site that are not achieving the metrics
  - Admins/contacts at the sites will know their status as seen by ATLAS Distributed Computing
  - Possible automation in the future

ATLAS Site Status Board

http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteview?

Site Name	SAM					Maintenance in SAM
	LFC	LFC_L	SRMv2	CE	FTS	
AGLT2	n/a	n/a	ok	error	n/a	up
ALBERTA-LCG2	n/a	n/a	error	error	n/a	up
Australia-ATLAS	n/a	n/a	ok	ok	n/a	up
BEIJING-LCG2	ok	ok	ok	ok	n/a	up
BU_ATLAS_Tier2	n/a	n/a	ok	error	n/a	up
CERN-PROD	error	ok	ok	ok	ok	up
CSCS-LCG2	n/a	n/a	ok	ok	n/a	up
CYFRONET-LCG2	n/a	n/a	ok	error	n/a	up
DESY-HH	n/a	n/a	ok	ok	n/a	up
DESY-ZN	n/a	n/a	ok	ok	n/a	up
FZK-LCG2	ok	ok	ok	ok	ok	up
GRIF	n/a	n/a	ok	ok	n/a	up
GoeGrid	n/a	n/a	ok	ok	n/a	up
HEPHY-UIBK	n/a	n/a	ok	error	n/a	up
IFIC-LCG2	n/a	n/a	ok	ok	n/a	up
IN2P3-CC	ok	ok	ok	ok	ok	up
IN2P3-CPPM	n/a	n/a	ok	ok	n/a	up
IN2P3-LAPP	n/a	n/a	ok	ok	n/a	up
IN2P3-LPC	n/a	n/a	ok	ok	n/a	up
IN2P3-LPSC	n/a	n/a	ok	ok	n/a	up
INFN-CNAF	n/a	n/a	n/a	warn	n/a	up
INFN-FRASCATI	n/a	n/a	ok	ok	n/a	up
INFN-MILANO	n/a	n/a	n/a	n/a	n/a	up
INFN-MILANO-ATLASC	n/a	n/a	ok	maint	n/a	All services in maint
INFN-NAPOLI-ATLAS	n/a	n/a	ok	ok	n/a	up

Site Name	SAM 24h average	SAM 48h average
AGLT2	99	99
ALBERTA-LCG2	0	0
Australia-ATLAS	99	98
BEIJING-LCG2	99	99
BU_ATLAS_Tier2	99	99
CERN-PROD	99	99
CSCS-LCG2	99	99
CYFRONET-LCG2	99	99
DESY-HH	99	99
DESY-ZN	99	99
FZK-LCG2	99	95
GRIF	99	99
GoeGrid	99	99
HEPHY-UIBK	99	97
IFIC-LCG2	99	99
IN2P3-CC	99	99
IN2P3-CPPM	99	99
IN2P3-LAPP	99	99
IN2P3-LPC	99	99
IN2P3-LPSC	99	99
INFN-CNAF	99	99
INFN-FRASCATI	92	83
INFN-MILANO-ATLASC	99	99
INFN-NAPOLI-ATLAS	99	99
INFN-ROMA1	99	99

Missing Pieces

- Datasets being produced
  - average file size, number of files can be found in the prodsys
  - the datasets recently created can be found in dq2
    - dq2-list-dataset-by-creationdate [--site=] --younger
- Subscriptions
  - Monitoring to see how many datasets are subscribed to a site and how much volume
    - dq2-list-subscription-site shows a list of subscriptions
    - dq2-list-files (or dq2-ls -f) shows the size of each dataset
  - Site Exclusion in T0 export
- Deletion
  - Better monitoring

Missing Pieces

- LFC/SE
  - Consistency
  - ACL
- Operations
  - Too many manual operations, need some more automitization...
- anything else?
- contributions are welcome (suggestions, tool development, ...).

Feedback from the Squad?

Feedback from the Sites?