

# ATLAS S&C Week

June 2013

# ATLAS SW Installation & Validation

- Present system: jobs through WMS on each CE
  - Tagging in AGIS all PandaResource attached to that CE
- New system: jobs through Panda on each PandaResources
- Tagging will be done in AGIS for each PandaResource
- releases tagging is now done with panda jobs since June 11
- Failed jobs:
- [https://atlas-vm-db-03.roma1.infn.it/atlas\\_install/list.php?summary=1&orderby=entries&orderdir=desc&status=failed&maxage=24](https://atlas-vm-db-03.roma1.infn.it/atlas_install/list.php?summary=1&orderby=entries&orderdir=desc&status=failed&maxage=24)
- <http://pandamon.cern.ch/jobinfo?jobtype=install&hours=24&plot=no&cloud=UK>

# Reorganization of PandaQueues

- 1 PandaResource ( PandaSiteID): 1 PandaQueue: N CEs associated
- it is a suggestion, not mandatory (but it will simplify their management in the future)
- PandaResource: site computing resources with specific parameters that cannot be mixed (e.g. SL6/5, memory limits, queue length...)
- Panda brokers jobs based on PandaResource information
- Switcher2 will act on PandaResource (exclude PandaResource in case of CE downtimes)
- All PandaQueues associated with 1 PandaResource *\*must\** have the same parameters. Having multiple PandaQueue behind a PandaResource with only difference the CE association is just historical, not needed:
- in AGIS : many CEs can be associated to one single PandaQueue (as already done for ANALY)
- SL(C)6 migration could be an opportunity to do it!

# Services at Sites: CVMFS

- CVMFS is mandatory for all the ATLAS Site
- CVMFS 2.0.X not optimal with multiple repositories
- CVMFS 2.1.11 looks promising, ATLAS encourages sites to test it on
- Was checked in NorduGrid (on sl6)

# PerfSONAR

- WLCG Operations Coordination PerfSONAR deployment Task Force
- <https://twiki.cern.ch/twiki/bin/view/LCG/PerfsonarDeployment>
- 60% of sites have already deployed PerfSONAR
- PerfSONAR v3.3 (allows easier management of central configuration)
  - Release candidate didn't show any issue up to now, expecting "soon" (2 weeks?) to have it as a production release

# WebDAV & XrootD

- WebDAV
  - ATLAS will use it:
    - For LAN data access (if more efficient than other protocols)
    - for Storage management operations for Rucio migration
  - Timeline: by September all sites are expected to provide WebDAV
- access for Storage management operations.
- Up till now only 7 out of 124 Storage elements are WebDAV accessible (and
- have such record in AGIS) (mostly De sites)
- – [https://atlas-agis.cern.ch/agis/storage\\_element/create/](https://atlas-agis.cern.ch/agis/storage_element/create/)
- XrootD
  - ATLAS will use it:
    - For WAN and LAN data access and transfer (if more efficient than other protocols)
    - As a protocol for federating storages
  - Timeline: to be widely deployed over summer
- [https://twiki.cern.ch/twiki/bin/view/Atlas/StorageSetUp#Requirements\\_on\\_SE\\_functions](https://twiki.cern.ch/twiki/bin/view/Atlas/StorageSetUp#Requirements_on_SE_functions)

# Requirements on SE functions

- **Protocols:** srmv2 , xrootd , WebDAV
- **srmv2** : to store files in SE (to utilize space reservation with space tokens and to stage files from tape via srmBringOnline )
- **xrootd** : for local and remote access, read-only to start with, read/write eventually
- **WebDAV** : initially for Storage Management operations for the Rucio migration during 2013 (rw) , eventually for local and remote access (rw)
- **space reservation** ATLAS uses the space tokens
- **checksum** ATLAS uses Adler32 to verify files

# FTS3

- Successfully tested in FT activity since ~ 9 months (and used)
  - Pilot services (RAL, BNL , ASGC, PIC and CERN)
  - Quick bug fixing, incremental developments
- release in EPEL -> ~2 weeks (EPEL testing)
- FTS3 prod instance in e.g. RAL and BNL could be deployed
- Once available, ATLAS plans to start using FTS3 in the same way they use FTS2 now



# FTS 3

- More protocols supported via gfal2 (srm, gsiftp, http, xroot,SRM-free transfers)
- More db back-ends (ORACLE, MySQL)
- Simplified configuration, staging files from archive, support for multiple replicas, session/connection reuse, improved performance
- Pilot service: CERN, RAL, ASGC, PIC,BNL
- Tested by LHC VOs for ~9 months, ATLAS and
- Tested by other VO's (snoplus.snolab.ca, ams02.cern.ch, vo.paus.pic.es, magic, T2K, NA62)
- Moving towards production release ( "3.0" tag )
- FTS 3 will be in EPEL soon

# Production

- Production not completely full after the urgent winter production, few weeks of lack of tasks
- PRODDISK
  - Got rid of it on most of T1s
  - T2s: removing postponed till the end of the year
- Memory issues
  - Currently:
    - Job memory set to what is requested in the task
    - Pilot limits (ulimit) vmem to job values
    - In case of out-of-memory failure → next retry automatically fixed at 4GB (limited number of queues)
  - In future:
    - Tasks with >4GB/process foreseen
    - 4GB is not enough for all cases, it's needed to decide on the algorithm (increase once by 2GB)
    - Only a handful of sites are configured for >4GB jobs.
    - 233 production queues:
      - 198 has memory set to 0, not clear what it means
      - 3 sites at 6GB, 1 site at 10GB
      - The rest in the 3-4GB range
  - ATLAS needs a clear definition of this value and asks the sites (cloud support) to set it to realistic values

# Production

- Maxtime (walltime)
  - 106 sites have maxtime at 0, not clear what it means (2 days?)
  - Most of the other sites have settings in 1.5-2 days
  - Values for both memory and maxtime should be fixed (0 is confusing: default?, unlimited?)
- Frontier/squid
  - MC simulation runs fine without DBRelease →using Frontier/squid only
- MCP
  - Extensive usage of MCP (65 sites)
  - Issues:
    - Slow ramp-up of high-input jobs (pile), can take days
    - Huge backlogs to consolidate the outputs, many FTS channels needed to increase the slots, but the performance still critical for default 2-day timeout
    - Might need to deploy FTS3 sooner on some T1s
- Prodsys scripts still use ancient job scaling rules.
- For the other tasks, we are waiting for JEDI-alpha

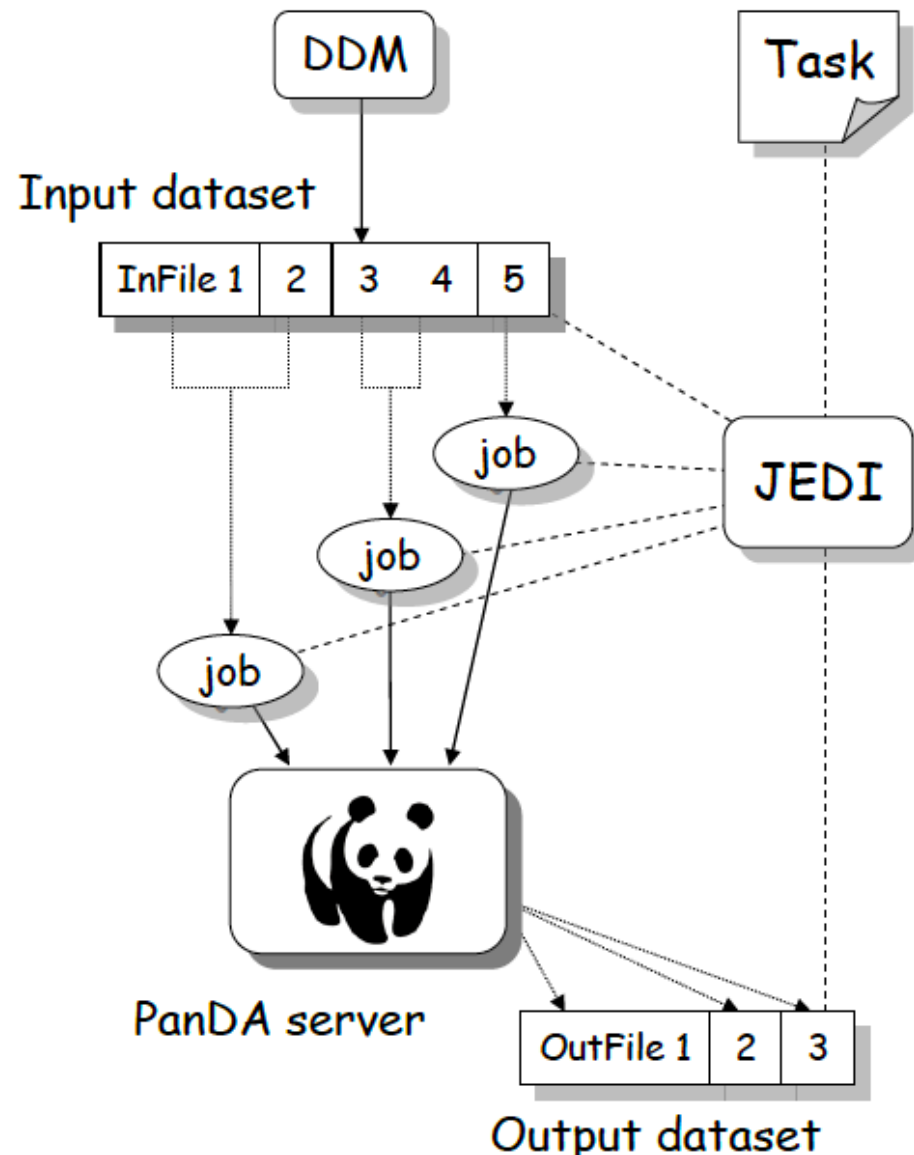
# Production

- Wide MCORE deployment:
  - Need to review the present setup, very complex and involves site actions
- SLC6 validation, and tasks:
  - Just starting, but very little production sites available
  - Most of the sites in testing, need to put them online as soon as possible
- Production jobs will certainly be more complex and demanding

# Production System

- new ProdSys2:
  - Based on tasks rather than jobs:
  - More complex workflows
- new ATLAS Production system consists of two subsystems
  - DEfT (Database Engine for Tasks). Task (meta-tasks, chain tasks definition), upper level I/F for Production Management Team (MC, reprocessing, group production)
  - JEDI (Job Execution and Definition Interface).
    - JEDI will use task formulated by DEfT
    - Integrated with Bamboo and PanDA brokerage
- JEDI- $\alpha$  has been developed
  - Tested for 1.5M jobs/day use case for 1 week
  - The focus now is to put it in production and retire prodsys/bamboo. Timescale: end of June
- DEFT prototype has been developed
- ProdSys2 prototype by the end of 2013.

# Workflow of JEDI



1. Task is submitted to the system
2. Dataset contents are retrieved from DDM
3. JEDI registers output dataset
4. JEDI splits input and generates jobs
5. Jobs are submitted to panda server
6. Files are added to output dataset
7. Task is complete when all input files are processed
8. Output dataset is frozen

Slide from Tadashi Maeno

# SL6 Upgrade

- Sites are requested to upgrade by October
- Exact time decided by sites
- Several issues with upgrade:
  - Very few native SL6 builds of Athena releases → compilation errors
  - Mismatch between OS of worker node and that of submitting host → fix by allowing CMT to get the right configuration:
  - Export CMTUSERCONTEXT=/cvmfs/atlas.cern.ch/repo/tools/slc6/cmt
  - Will be deployed next week
- Peak of user job failures due to compilation issues (many related to ROOT based analyses)
- – Problems handled and solved by DAST
- – But users have to learn the hard way

# RUCIO

- Very good progress in software development
- Core components are ready
- Scale test of core components in preparation for Mid July
- Based on observed workload from production in DQ2
- Functional testing of components successfully done
- Testbed will be provided in September
  - For applications integration
  - For operations to gain experience
- A full scale test of the full system can be foreseen by end 2013



# Rucio naming status

- 30% of all files in CERN LFC follows the Rucio on naming convention (~100M files)
- LOCALGROUPDISK
  - We need to check with sites and ask users if their scripts rely on the old naming convention
  - Users usually rely on the fact that all files belonging to a dataset are in the same directory
  - How do I get the list of files belonging to a dataset?
    - `dq2-ls -pfl <endpoint> <dataset>`
  - for a particular endpoint
    - `dq2-list-files --rucio <dataset>`
  - If pilot already writes to LOCALGROUPDISK then it does it with Rucio naming convention
- The DQ2 Central catalog will refuse old versions (<2.4.0)
  - locally installed DQ2 clients should be updated
  - DQ2 client in cvmfs (under ATLASLocalRootBase) is already pointing to the latest version

# Network

- ATLAS critically depends on the network
- Network problems can be hard to diagnose and slow to fix
- Network Data Access Requirements in Analysis:
  - During busy times – 1 million user analysis and D3PD production jobs per day, user analysis jobs are short jobs
  - 90% of NTUP datasets (by volume) are read 50 times or less
  - Users still read a small fraction of the centrally produced D3PD's (NTUP's)
  - Implies much of the data moved is not read
- ATLAS data in 2015-2016 are more challenging since the trigger rate is increased by a factor of 3
- Network Monitoring for ATLAS
  - Choice of a standard “tool/framework”: perfSONAR
  - Regular network metrics collected and stored in the AGIS and will be used by PANDA in task assignment and in MC production

# Network

- Network awareness will be extended to the middleware
- FTS3 for example
- This will simplify ADC components (e.g. Rucio)
- Software defined networks (virtual circuits) would allow to use network resources more flexibly
- For example loosing the hard partitioning between LHCOPN and LHCONE traffic

# FAX

- Sites join FAX on voluntary basis
- <https://twiki.cern.ch/twiki/bin/view/Atlas/JoiningTheATLASFederation>
- 37 sites already joined
- May be used as fail-over in case of local storage access failures
- Already activated for many US and some UK PandaResources
- ATLAS was running running functional tests
- FAX usage will be introduced in Panda
  - Jobs failover to FAX in case of data access failure
  - will be turned on in US and part of UK
  - Brokering will loose the job-to-data locality

# Opportunistic resources for ATLAS

- HEP now risks to compromise physics because of lack of computing resources.
- Research and commercial clouds
  - Mechanism to cope with peak loads on the Grid
  - Experience Amazon EC2, Helix Nebula
  - Google projectCompute Engine (GCE) preview project :
    - resources for ATLAS for free: ~5M cpu hours, 4000 cores for about 2 month• These are powerful machines with modern CPUs
  - Resources were organized as HTCondor based PanDA queue (Centos 6 based custom built images, with SL5 compatibility libraries to run ATLAS software)
    - Condor head node, proxies are at BNL, output exported to BNL SE
- High-Performance Computers (studies at ANL, LMU, NorduGrid)
  - Each HPC is unique (architecture, HW, OS, “weak” ,limited memory per WN)
  - Unique job submission systems
  - unique security environment
- “ATLAS@home” via BOINC

# HPCs Opportunistic Resources

- HPC offers important and necessary opportunities for HEP
- Competitive (in CPU/\$) with HEP systems
- Problems
  - Very wide spectrum of site policies
    - all HPCs rely on shared FS and have little or no local disk
    - No grid SW on WNs
    - One solution unlikely to fit all