



ATLAS Distributed Computing: Commissioning & Integration, Clouds & opportunistic resources, and more...

Alessandro Di Girolamo
on behalf of the ADC community



IT-SDC: Support for Distributed Computing

Commissioning & Integration: ProdSys2

- Developed new distributed production framework
 - DEFT: enable production manager to define requests & translate the requests into chains of task
 - JEDI: generate the job definition, babysit user analysis jobs
 - Panda: manage the execution of the jobs on the distributed infrastructure
- MC, Group, Reprocessing coordinators are already implementing their some of the workflows in ProdSys2
- Several MC tasks already running on ProdSys2
 - Validation ongoing with respect to ProdSys1
- Reprocessing campaign:
 - Bulk with ProdSys1, a part with ProdSys2 to crosscheck

ProdSys2: some immediate benefits

- Data loss: faster reproduction capability
- Automatic transient data handling:
 - minimize the amount of transient data
- Dynamic dimensioning of jobs
 - based on scout jobs
 - a bit farer in time the Event Service
- Auto-rescaling of failing jobs
- Log files merging (and storing to tape)

Commissioning & Integration: Rucio

□ Rucio integration:

- File renaming to Rucio convention **DONE**
- LFC -> to Rucio migration **DONE**
- DQ2 -> Rucio **ONGOING**



Commissioning & Integration: Rucio

□ Rucio integration:

- File renaming to Rucio convention **DONE**
- LFC -> to Rucio migration **DONE**
- DQ2 -> Rucio **ONGOING**

- Need to complete successfully the Rucio Full Chain Test
- Migration of DQ2 entities in Rucio in background
- new DQ2 Rucio aware clients
 - Move activities to Rucio backend
 - First tests with HC Analysis Functional Tests

Rucio commissioning

- Since ~1 year: Rucio continuous internal validation test
 - Daemons stability, validation of releases, DB, etc.
- Since mid-May, ongoing: Rucio Full Chain Test
 - validate all the Rucio blocks in “real life”
 - Stress tests the Rucio building blocks
 - Simulate the Tier0 export workflow with test data
 - Not Grid infrastructure stress test!
 - Integrate & commission missing functionalities, e.g.:
 - Interaction with tape systems
 - Quota management
 - Tier0-Tier1s-Tier2s transfer “hops”

Foreseen for 6 weeks, extended for few weeks more

Data access

- Rucio enable usage of different (from SRM) protocols
 - Which doesn't mean we can completely get rid of SRM yet
 - E.g. TAPE recall , quotas

- WebDAV can be used today for:
 - See log files directly on browser,
 - download output on your laptop
 - More activities under testing (next slide)

- Xrootd can be used for direct data access in root/Athena
 - evaluating also WebDAV, but not productionready yet

WebDAV commissioning and performance evaluation

- Functional tests on all ATLAS site to verify the webDAV setup

- Reproduce and compare FAX testing for webDAV with HammerCloud
 - improve the results by fixing the various bugs

- Test deletion performances

WebDAV commissioning and performance evaluation: issues

- On many sites, few or even no file could be opened
 - dCache sites are being updated to version 2.6.31
 - On dpm sites some disk nodes are failing: contact admin
 - WebDAV http server instabilities on some dpm sites: dpm developers
- Some jobs crashing: Davix developer on the issue:
 - Error message “error: Payload out of memory”
- Problem of proxy not supported on most sites: cream team
- Still problem of compilations on some sites
- Some site haven't yet opened WebDAV doors for remote access: ATLAS request

FAX: Federated data Access through Xrootd

- 56 sites, 91% of the data available (target to 95%)
 - In progress NDGF, Nikhef, some other Tier2s
- In production
 - failover in case of data access issues
 - some Tier3 and opportunistic CPUs
- Under testing
 - “overflow” use case: WAN data access from sites with available CPUs but not data
- Future
 - Event Service
- Stress testing in progress

xrootd, FAX, HTTP policy & plans

ADC has agreed the following policy regarding xrootd and FAX at T1s and T2Ds:

In the Feb 2014 ATLAS S&C Week ADC Operations session it was agreed as policy that T1s and T2Ds are to offer xrootd access to their storage, where the storage technology allows it. ADC furthermore asks and encourages sites not yet in the FAX federation to take the modest additional step beyond supporting xrootd of joining FAX. If there are technical issues, then please let ADC know.

- We intend to demonstrate WAN data access at scale (~10% of data access) in DC14, utilizing the technology available today (xrootd, FAX)
 - Consequently, timescale for installation is in time for pre-DC14 testing
- We intend to explore and possibly utilize HTTP as technology for federating storages and enabling WAN data access
 - Compare xrootd, http for WAN access during 2014
- Also will put HTTP in production (e.g. downloads/dq2-get) sooner as they solve long standing issues impacting users
 - Does our data have to be ATLAS-only read protected?? Disables caching
- Therefore we ask sites to enable HTTP access via WebDAV on same timescale.

□ FAX: We ask sites to deploy FAX by DC14, i.e. June, including commissioning and testing.

□ HTTP/WebDAV: ATLAS has use cases for which intend to use http/webDAV as primary protocol for data access/transfer: we ask sites to deploy it with the same timeline of FAX, with slightly lower priority respect to FAX.

! Feel free to discuss with ATLAS in case of any problem or concern

11

Slide already presented in several meetings!
STILL VALID!

The above is from ATLAS SW and Computing week 24-28 February 2014.
ATLAS requested it also during WLCG OpsCoord meeting :
<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsMinutes140306#ATLAS>



ATLAS Tier0 evolution

□ CERN LSF batch system tests

- Successfully tested (at 2x the foreseen std workload) setup with dedicated master & resources
 - 2 local users with different shares. Still to be tested with Grid users/jobs

□ Future Storage infrastructure: model to be finalized

- EOS for data from P1
 - EOS JBOD safer than CASTOR RAID, multiple copies possible
 - RAW and derived data available for some time for quick local access
- CASTOR as “cold storage” : tape backend
 - Replica on CASTOR through a third-party transfer

□ Spill over to Tier1s still to be defined



MultiCore

- All sites are requested to implement M CORE queues, dynamic provisioning is recommended

- ATLAS will need in the next months on M CORE 30-50% of resources
 - G4 simul will run on M CORE with new releases
 - Part of the DC14 reconstruction will run on M CORE when validated
 - Analysis not yet foreseen on M CORE

- will not run single core jobs on M CORE payload

Frontier & CVMFS

□ Frontier

- Frontier launchpads migrated to CERN AI (support by ATLAS CS)
 - Common work ATLAS-CMS

□ CVMFS

- Switch for the central services to CVMFS 2.1 in August.
 - No problem expected from the sites: should all have CVMFS 2.1(.19) by then
- HPC and opportunistic non-ATLAS resources need to access the external libraries from CVMFS, mainly via parrot.
 - Problems with multi-repositories and parrot being addressed by the developers - beta version available.



Squid

- For Frontier access today recommend the Frontier squid RPM (using squid 2.x)
 - we are actively looking to work with sites to test running a standard squid 3.x (for both/either CVMFS/Frontier)
 - @CERN: ATLAS Squids being moved to CERN IT managed service
- Large files on CVMFS:
 - ATLAS aware of the possible stress on site caches provoked by this small number of large files: actively working to remove them


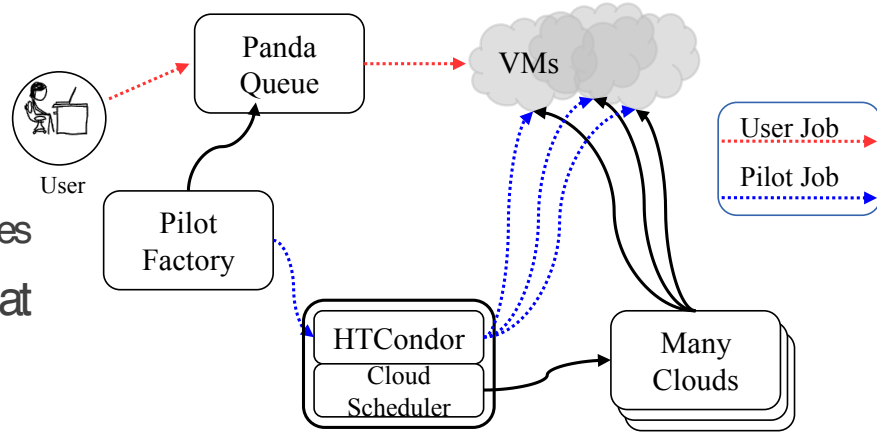
RRC-KI-T1 Russian proto-Tier 1

- Disk storage and computing resources commissioned for several ATLAS activities
- Tape system: site still working to be able to provide to the experiment a stable reliable service
- Networking: 10 Gbps protected bandwidth to LHCOPN (since couple of months), 3 Gbps of R&N connectivity, full peering with all Russian ATLAS T2 sites at 1-10 Gbps.
- RRC-KI-T1 is effectively used by ATLAS as site without pledges:
 - production: running in "multi-cloud processing" mode since almost one year.
 - analysis: running since couple of months. Data shipped through dynamic data placement (PD2P)
 - no primary/custodial data
- waiting for the site to sign the MoU



Cloud resources

- Cloud Operation team setup:
 - consolidate the solutions adopted
 - Evaluate new possibilities
 - Ease the integration of new resources
- 1 image and contextualization that can run anywhere
- Using Shoal for dynamic Squid "Federation" ([github](#))
- Cloud Scheduler for VM provisioning and management ([website](#))
 - 2M jobs completed in 2014
- Evaluating Vacuum and Vcycle ([website](#))



Cloud resources: Sim@P1

Sim@P1 enable ATLAS TDAQ HLT farm (LHC Point1) to ATLAS MonteCarlo production since July 2013

□ Collaborative effort of BNL, ATLAS and CERN-IT/SDC

□ Based on OpenStack (Folsom)

□ 1.3k nodes: up to 2.6k VMs □ 20.5k cores

- TDAQ □ Sim@P1: 6 Hz VM instantiation rate, 1.5/2Hz job flow. Sim@P1 □ TDAQ: 10 min!

□ VM: in 2013 based on CernVM 2.6.0 (SLC 5.8), now on “golden image” as other ADC clouds.

- Investigating the possibility of using CernVM 3.3.0

□ SingleCore & MultiCore:

- today static, plan to try HTCondor dynamic partitioning

□ Agreeing with ATLAS management operations during data taking:

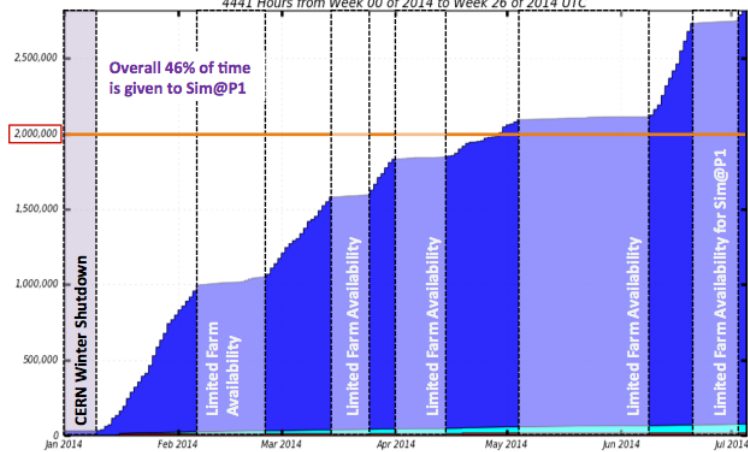
- reasonable to use Sim@P1 for LHC stops at least 24 hours long



Cloud resources: Sim@P1 – cont'




Completed jobs Cumulative CERN-P1(_MCORE)
4441 Hours from Week 00 of 2014 to Week 26 of 2014 UTC



■ MC Simulation (2,744,916) ■ Others (58,388) **Jan 1, 2014 – Jul 4, 2014 Total: 2.82M**



CPU consumption Good Jobs in seconds
4442 Hours from Week 00 of 2014 to Week 26 of 2014 UTC



Jan 1, 2014 – Jul 4, 2014

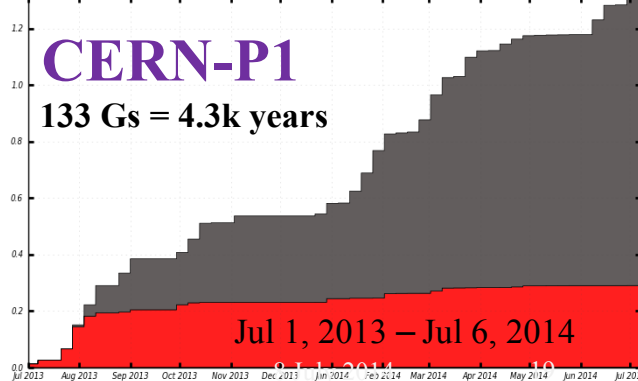
■ MC Simulation (1,015,848,591.126) ■ Extra Production (8,242,070,480)

Total: 1,024,090,661,606, Average Rate: 64,036/s

CERN-P1(_MCORE)

- ~7% of all ATLAS resources for MC production
- as one Tier1 as overall CPU/WALLClock delivered over the past year (but running only EvGen& Simul)

8903 Hours from Week 26 of 2013 to Week 27 of 2014 UTC




■ simul (102,025,335,459) ■ evgen (28,938,272,139) ■ gangarobot-pft (66,247,724) ■ hammercloud (40,257,187)

HPC

- HPC in: DE, NO, SW, CH, US
- Over the past 6 months contributing to ATLAS with 10-20k cores
- Quite different world respect to the Grid, and very variegated
 - no external connectivity from/to the WN
 - Pilot “outside” the HPC itself (edge node)
 - X86 but also non-intel (bluegene, for now only event generation)
- ARC-CE provides perfect gatekeeper
 - data in/out handled, no grid MW or ext connectivity needed on WNs
- Lots of possible cycle “available” in the future if:
 - optimize/automate our backfilling strategy
 - SW optimization: e.g. Disk access, GPUs...

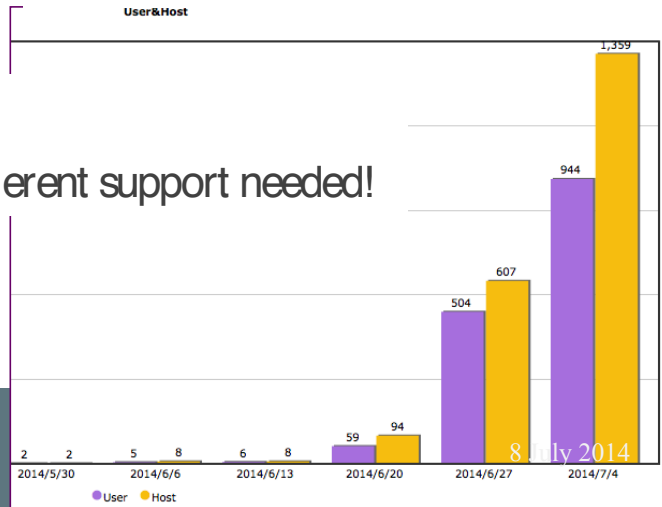
BOINC – ATLAS@Home



- Server with ARC-CE and Boinc server + ATLAS@Home app
- BOINC PandaResource
 - Low priority MC simul short jobs

□ ~1k users:

- Complete different support needed!

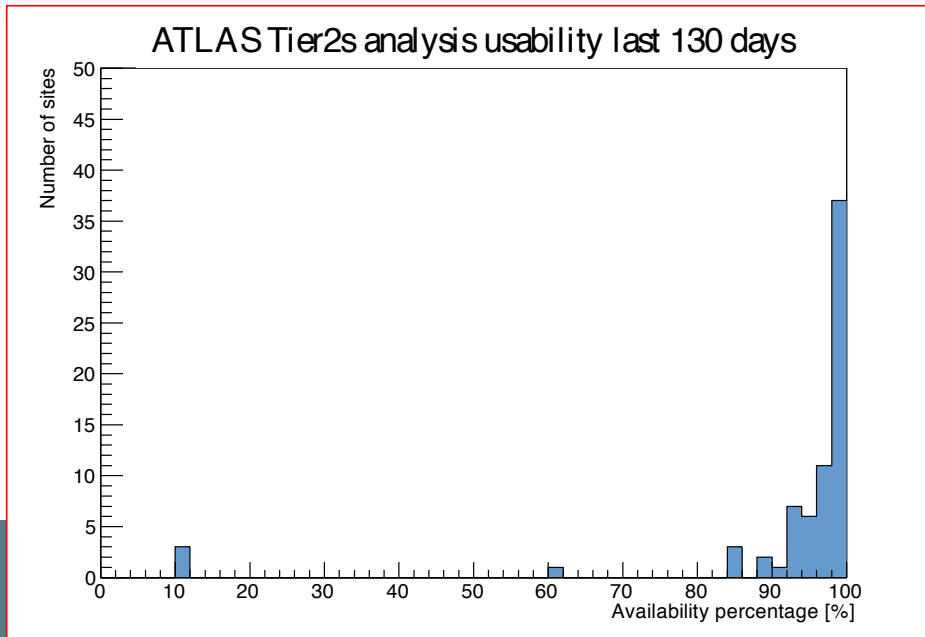


AGIS – ATLAS Grid Information System

- Concept: resources "provided by" are "used by" ATLAS entities
 - decoupling the services from the ADC frameworks
- Organized in plugins, based on Django, exposing in REST ...
- Uniform naming convention and topology:
 - Services, workload management, data management, monitoring: structured and connected.
- reliable ATLAS Central Service since more than 2 years:
 - always evolving! Ongoing examples:
 - dynamic queues: expose to Panda all the "possibilities" offered by the various sites simplifying the life of the site admin and of the shifters
 - re-organization of storage protocols
- Discussion/evaluation ongoing with CMS
 - useful also for ATLAS having other experiments using and contributing to it!

ATLAS Site usability

- Automatic way to categorize ATLAS sites in place since few months:
 - Based on usability of the site for analysis over the past 30 days
 - Recomputation possibility - similar to SAM procedure
 - [ABCD categorization](#)
 - Useful to have status of the sites over the past month, and also to keep history over longer period: 6 months in the histogram here




ATLAS Site usability

□ Evaluating the possibility to expose to WLCG as site availability/reliability the ATLAS usability

□ Why:

- SAM tests to measure ATLAS availability/reliability for WLCG sites are often not representing well enough the real ATLAS usability: these new metrics can complement the atomic tests to provide more realistic availability and reliability numbers
- resources are accounted to ATLAS but not all of them are really usable!

□ Still few technical details to sort out before discussing with WLCG MB



Conclusions

□ Opportunistic resources

- Able to use (almost) whatever available:
 - integration work is a lot
 - consolidating solutions

□ Site usability - resource availability

- Working to expose to the scrutiny group and WLCG MB (avail/rel reports) sites usability as seen by ATLAS

□ Many new technologies and frameworks for Run2

- Commissioning ongoing
- No showstoppers

Run2: (almost) ready to take off!



Back Up



..

□ ..

• ..

• ..

□ ..

□ ..

□ ..

□ ..

□ ..



..

□ ..

• ..

• ..

□ ..

□ ..

□ ..

□ ..

□ ..



■

□ ..
□ ..

□ ..
□ ..

□ ..
• ,
• .

□ .
• .
□ ..
• ..

!
•
• ..
! ..