
ADC Overview


Andrej Filipcic, Alessandro Di Girolamo
Nurcan Ozturk

GridKa and ADC F2F Meeting, May 23-24, 2017

Distributed Production in 2017

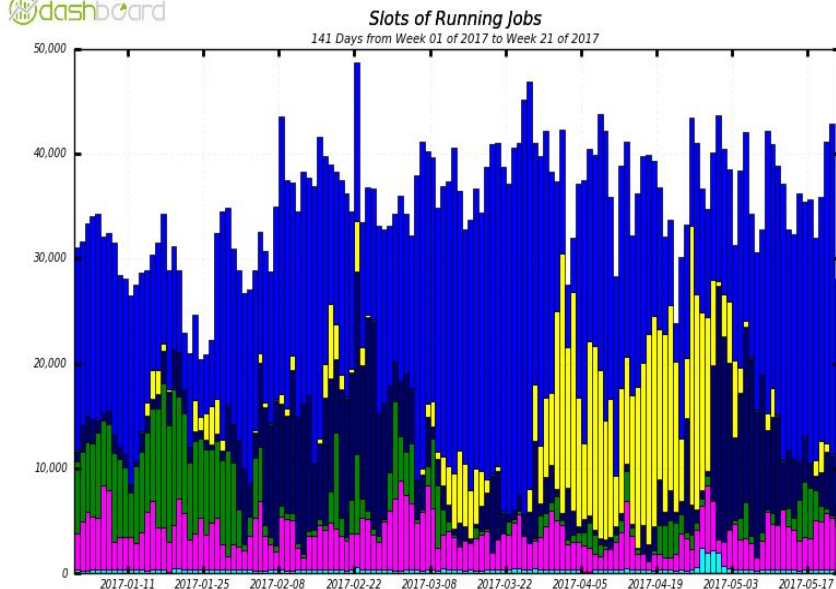


Slots of Running Jobs
141 Days from Week 01 of 2017 to Week 21 of 2017




- Full MC+data derivations for Moriond and spring conferences completed in early 2017.
- Reprocessing of data15+data16 completed in March-April.
- Derivations on reprocessed data is about to start.
- MC16 campaign running in full speed.
- Analysis continue to use ~30k slots on average.

DE cloud jobs - per activity, per site





■ MC Simulation ■ Data Processing ■ MC Reconstruction ■ Group Production ■ Analysis
■ Others



■ FZK-LCG2 ■ LRZ-LMU ■ UNI-FREIBURG ■ DESY-HH ■ CSCS-LCG2
■ PRAGUE LCG2 ■ MPPMU ■ DESY-ZN ■ GOEGRID ■ WUPPERTALPROD
■ CYFRONET-LCG2 ■ IEPSAS-KOSICE ■ PSNC ■ FMPIH-UNIBA ■ UNI-SIEGEN-HEP
■ HEPHY-UIBK ■ TUDRESDEN-ZIH ■ UNI-DORTMUND ■ UNI-BONN

DE Transfers (source)



Production Plans in 2017/2018

- Data taking/processing, open-ended derivation production.
- 3-step MC16 production campaign will continue:
 - **MC16a, b:** From Jan 2017, 2015-2016 mu profile, b: mu 40-60
 - **MC16c, d:** From May 2017, 2017 trigger menu, expected mu profile, d: true mu profile
 - **MC16e:** In 2018, expected mu and trigger menu for 2018
- Heavy Ion reprocessing by the end of the year.
- Fast reprocessing (Release 21+conditions update) of 2017 data will happen. No MC reprocessing.
- Release 22 is for Run-3, only validation will run.

Hot Topics for ADC in 2017

- Premixing: digi+reco is too I/O heavy - split in
 - Minbias + neutrino signal digitization -> RDO
 - Overlay minbias RDO + signal HITS and reconstruct
 - 50-100M event RDO minbias to be used with reco
 - Dataset distribution to be changed
- Fast-chain
 - Simulation
 - Digitization
 - Reconstruction
- New monitoring effort provided by CERN-IT
 - Current ATLAS job and DDM dashboards will be migrated to Kibana/Graphana
 - Active involvement by ATLAS on what needs to be migrated/new features to be added, etc
- Understanding workflow and software performance
 - E.g. CPU efficiency of ATLAS jobs was degraded in the last few years - site/activity/ncore dependence

WFMS

- Main topics:
 - **Harvester** - being developed for resource provisioning for timely optimization of CPU allocation among various resource types (grid/HPCs/Clouds). First version for HPCs is deployed at NERSC and Titan, functional tests ongoing. Version for the grid usage is under development.
 - **Event Service** - fine grained event level processing instead of file level processing. Active development/commissioning/monitoring effort.
 - **Predictive Analytics**- anomaly detection; overload and misconfigurations in the system, malfunctioning computing resources, malicious activities
- Some of recent changes:
 - [Global shares](#) implemented - to better control resources needed by each activity
 - Time-To-Complete (TTC) estimation - for tasks and chain of tasks
 - Several monitoring updates - [MC16a campaign report](#), etc
- Analysis/production batch shares
 - Readiness on moving to only one proxy

DDM

- Main topics:
 - Automating lifetime model
 - Archiving via tarball to put on tape
 - Deletions based on space reporting from Rucio values
 - How to reduce data transfers - optimize workflows
 - Closeness matrix to be made more dynamic
 - Rucio and event service - how Rucio would scale for that
 - How to use Object Store more
- Some of recent changes:
 - Groupdisk decommissioning
 - SRM decommissioning
 - Storage dumps available in Hadoop that can be compared with sites' own records

HPCs and Clouds

- SuperMuc, Draco/Hydra, Piz Daint
 - How much resources can we get?
 - Simul only?
- ATLAS@home
 - Very stable for MC simulation
 - Large contribution from MPI, WLCG performance-studies cluster, Prague
 - Accounting can be provided if significant resources are allocated

Containers - Singularity and Docker


- ATLAS is moving to towards full containerization for production and analysis
 - simplifies site maintenance and centralizes deployment
- Singularity - for production
 - Tuned for batch systems and HPCs, easy to install, one rpm, straightforward config
- Docker - for software development
 - More difficult to deploy, usage foreseen on few proactive sites
- Large scale singularity deployment started with all modern OS sites in March, several sites already using it; SIGNET, ARNES, HPC2N, LRZ, MWT2, ...
- Under evaluation:
 - Full job containers
 - Partial pilot execution inside the container (eg transfers, payload execution)
- WLCG is supporting the container vision
 - To be followed up in Manchester

Infrastructure

- [Lightweight sites](#)
- Diskless (cache-only) sites, first discussion on [policy and procedure](#) last week
- direct/remote IO
- Batch systems
- Storage access for other experiments

Site Availability and Performance (ASAP)

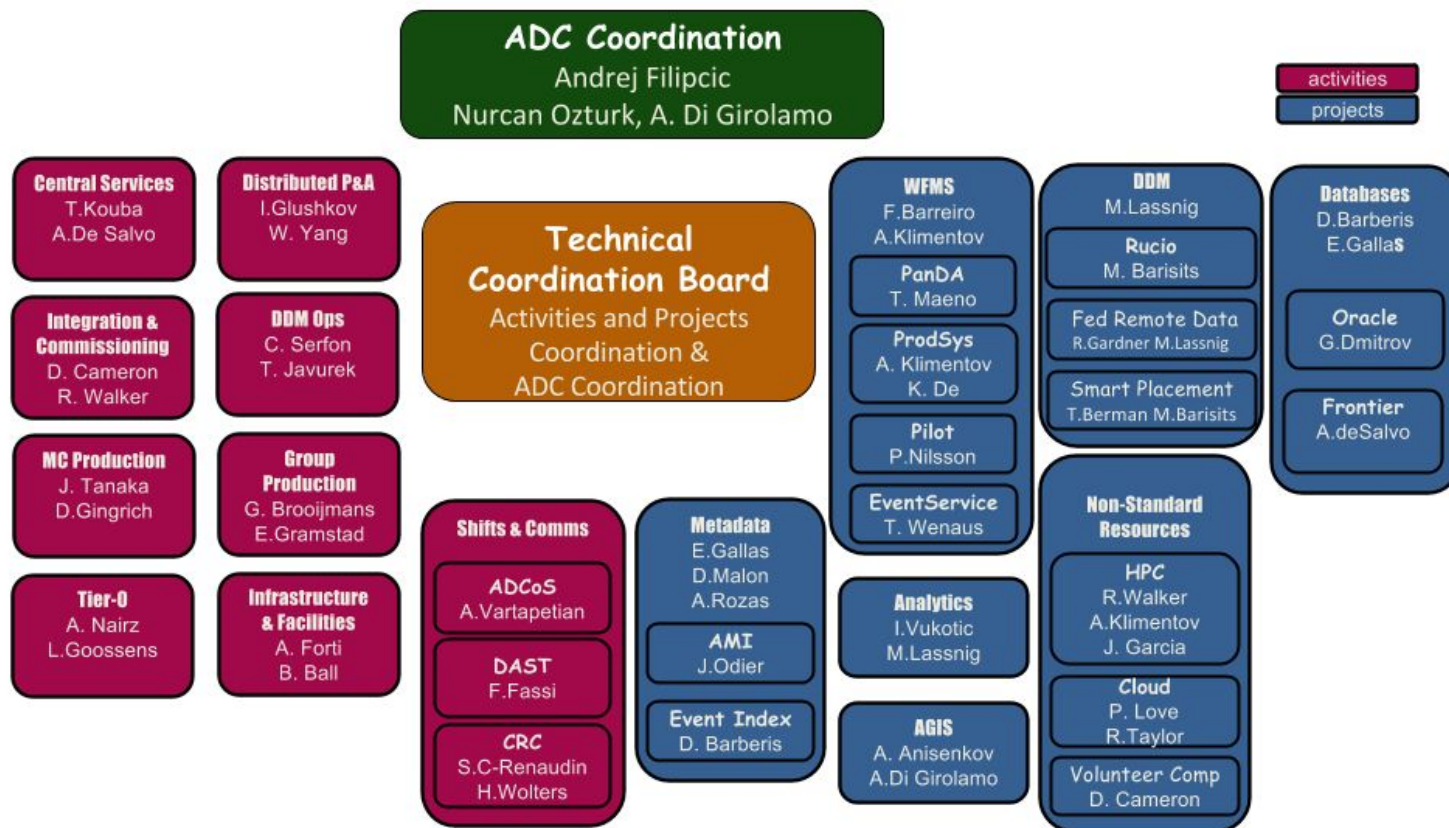
- Important to review site availabilities and performance, a [new Kibana monitoring](#) available.
- Currently production and analysis combined, in the search field they can be separated; data.id:1394 for production and data.id:782 for analysis.
- From a quick look, TUDeresden-ZIH and UNI-BONN is in critical in last 6 months.



Tape Usage

- How we can use more tape in ATLAS? For instance reading inputs directly from tape for derivation production.
- Two tape stress tests have been conducted, first with couple of Tier1 sites in January. then with all Tier1 sites in February with a [report](#) presented.
- Tape performance of FZK was not very good in these tests due to parallel activity - competing with CMS and dcache problem, however this weekend FZK showed a very good performance (distribution of reprocessed data to Tier1 tapes). Any idea what changed?
 - 210 TB delivered in 3 days, 160 TB in the first day (which is disk buffer)
- We need to understand from our side (DDM or FTS) how to protect the buffer getting full.

Communication with ADC - Projects and Activities



ADC Meetings

	Monday	Tuesday	Wednesday	Thursday	Friday
9 am	Morning meeting	Morning meeting	Morning meeting	Morning meeting	Morning meeting
9:30 am		DDM Ops	Tier0 News	Central Services	
10:00 am			Software News	DDM Dev	
3 pm	S&C Mgmt meeting		DPA meeting		
4 pm	Technical Coordination Board	ADC Weekly		WFMS	
5 pm		Infrastructure & Facilities	Event Service Analytics (in rotation)	ProdSys CREM	

Other meetings:

WLCG Ops (weekly, CRC attends)
 WLCG Ops Coord (monthly)
 ATLAS-IT (monthly)
 ADC&CERN-IT
 ATLAS Computing Workflow
 Performance Understanding
 (every 3-4 weeks)
 ADC&CERN-IT new monitoring meeting
 (every 3-
 ADCoS, CRC, DAST (weekly)

Communications, Shifts, Manpower, Operations

- Long term commitment from German cloud to ADC operations, many thanks!
 - DPA, analysis, HammerCloud, monitoring, shifts
- DE cloud mailing list and GGUS tickets working fine for site/service issues
- Any suggestions on communications and for better integration of DE cloud people in ADC?