



ATLAS Computing: from Development to Operations

Dario Barberis

CERN & Genoa University/INFN



Outline

- Major operations in 2007
- Computing Model in a nutshell
- Component testing
 - Tier-0 processing
 - Data distribution
 - Distributed production and reconstruction of simulated data
 - Distributed analysis
 - Reprocessing
 - Calibration Data Challenge
 - Data streaming
- Integration testing
 - Full Dress Rehearsal
- Requirements on Grid Tools



Experiment operations in 2007

- The Software & Computing infrastructure must support general ATLAS operations in 2007:
 - Simulation production for physics and detector studies
 - Cosmic-ray data-taking with detector setups of increasing complexity throughout the year
 - Start of "real" data-taking, at low energy, in November 2007
- In addition, the S&C system has to be fully commissioned
 - Shift from development-centric towards operation-centric
 - Test components of increasing complexity
 - Component integration towards the full system test ("Full Dress Rehearsal") in Summer (early Autumn) 2007
- This is what we call since last year "Computing System Commissioning" (CSC)

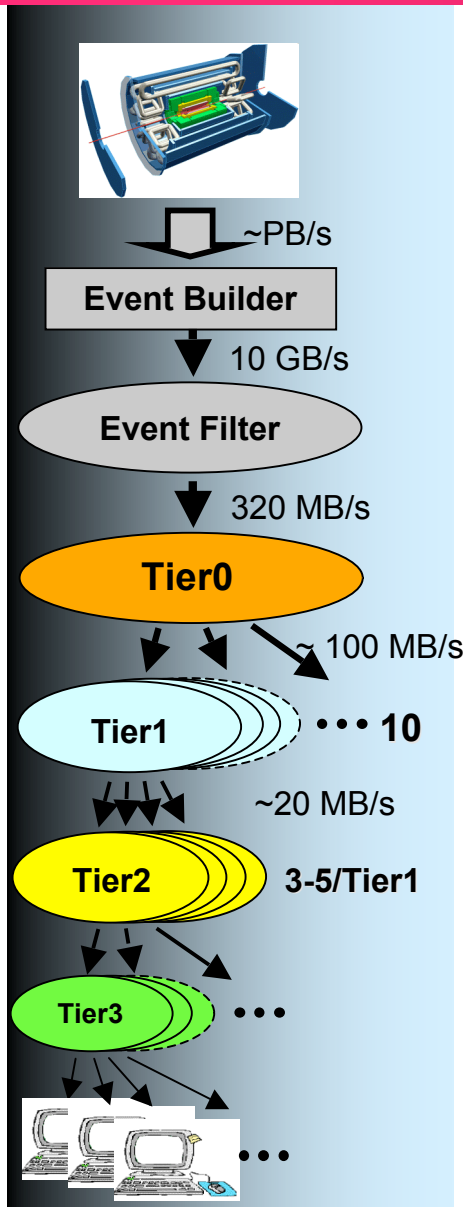


Computing Model: central operations

- Tier-0:
 - Copy RAW data to Castor tape for archival
 - Copy RAW data to Tier-1s for storage and subsequent reprocessing
 - Run first-pass calibration/alignment (within 24 hrs)
 - Run first-pass reconstruction (within 48 hrs)
 - Distribute reconstruction output (ESDs, AODs & TAGs) to Tier-1s
- Tier-1s:
 - Store and take care of a fraction of RAW data (forever)
 - Run "slow" calibration/alignment procedures
 - Rerun reconstruction with better calib/align and/or algorithms
 - Distribute reconstruction output (AODs, TAGs, part of ESDs) to Tier-2s
 - Keep current versions of ESDs and AODs on disk for analysis
- Tier-2s:
 - Run simulation (and calibration/alignment when appropriate)
 - Keep current versions of AODs on disk for analysis



Data replication and distribution



In order to provide a reasonable level of data access for analysis, it is necessary to replicate the ESD, AOD and TAGs to Tier-1s and Tier-2s.

RAW:

- Original data at Tier-0
- Complete replica distributed among all Tier-1
 - Randomized dataset to make reprocessing more efficient

ESD:

- ESDs produced by primary reconstruction reside at Tier-0 and are exported to 2 Tier-1s
- Subsequent versions of ESDs, produced at Tier-1s (each one processing its own RAW), are stored locally and replicated to another Tier-1, to have globally 2 copies on disk

AOD:

- Completely replicated at each Tier-1
- Partially replicated to Tier-2s (depending on each Tier-2 size) so as to have at least a complete set in the Tier-2s associated to each Tier-1
 - Every Tier-2 specifies which datasets are most interesting for their reference community; the rest are distributed according to capacity

TAG:

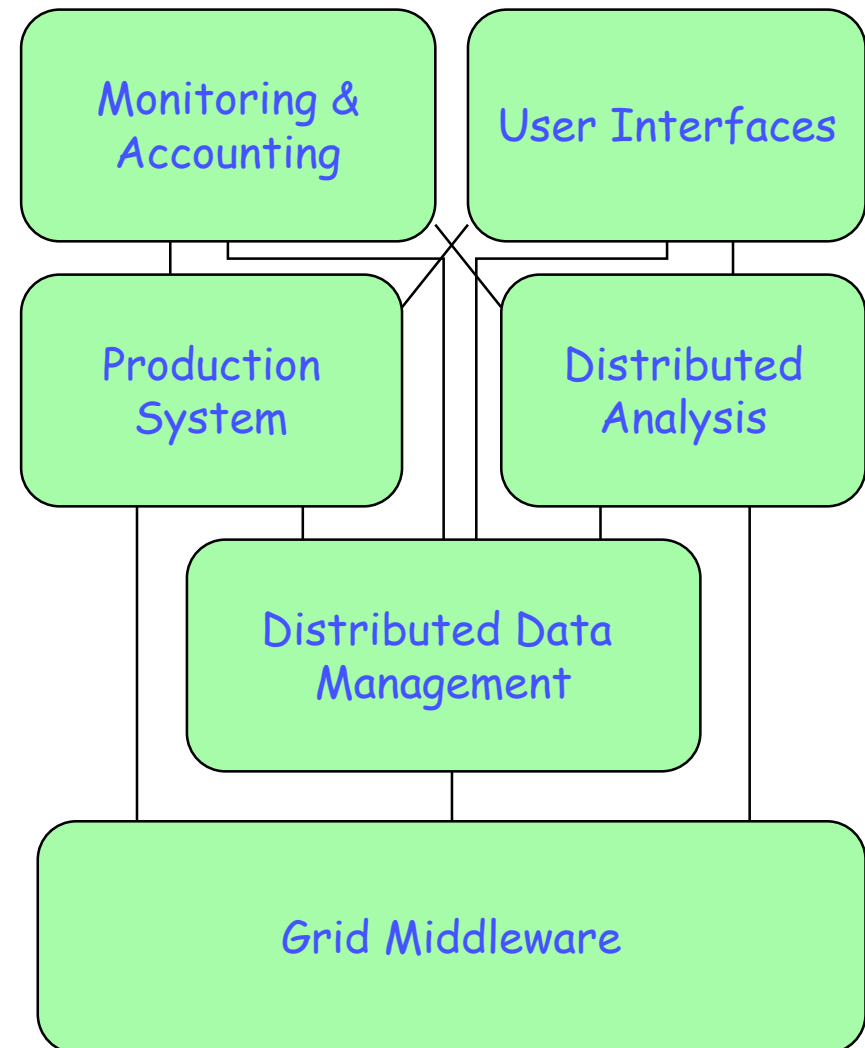
- TAG files or databases are replicated to all Tier-1s (Root/Oracle)
- Partial replicas of the TAG will be distributed to Tier-2 as Root files
 - Each Tier-2 will have at least all Root files of the TAGs that correspond to the AODs stored there

Samples of events of all types can be stored anywhere, compatibly with available disk capacity, for particular analysis studies or for software (algorithm) development.



ATLAS Grid Architecture

- The ATLAS Grid architecture has to interface to 3 middleware stacks: gLite/EGEE, OSG, NG/ARC
- It is based on 4 main components:
 - Distributed Data Management (DDM)
 - Distributed Production System (ProdSys)
 - Distributed Analysis (DA)
 - Monitoring and Accounting
- DDM is the central link between all components
 - As data access is needed for any processing and analysis step!
- Development and deployment activities are still needed throughout 2007
- During 2007 we also have to move from support for pure simulation production operations to the full range of services specified in the Computing Model
 - Including placing data (datasets) of each type in the correct location and sending analysis jobs to the locations of their input data





CSC tests: Tier-0 processing

- Two rounds of Tier-0 processing tests are foreseen in 1H-2007:
 - February 2007 onwards: Tier-0 tests 2007/Phase 1
 - integration with data transfer from the online output buffers (SFOs)
 - first prototype of off-line Data Quality monitoring integrated
 - more sophisticated calibration scenarios exercised
 - first prototype of TO operator interface
 - strategy in place for ATLAS software updates
 - first experiments with tape recall
 - May 2007: Tier-0 tests 2007/Phase 2
 - integration with real SFO hardware completed
 - first production version of off-line Data Quality monitoring in place
 - all expected calibration scenarios exercised
 - first production version of Tier-0 operator interface in place
 - all relevant tape-recall scenarios exercised
 - End of May: integration with Data Streaming tests
 - See later slides



CSC tests: data distribution

- Several types of data distribution tests were performed in 2006 and will continue this year
- Tier-0 → Tier-1 → Tier-2 distribution tests
 - Following the Computing Model for the distribution of RAW and reconstructed data
 - Will be performed periodically, trying to achieve
 - Stability of the distribution and cataloguing services
 - Nominal rates for sustained periods in the middle of 2007
- Simulated data storage at Tier-1s
 - Collecting simulated data from Tier-2s for storage on disk (and tape) at Tier-1s
 - This is actually a continuous operation as it has to keep in step with the simulation production rate
- Distribution of simulated AOD data to all Tier-1s and Tier-2s
 - Also has to keep going continuously at the same rate as simulation production



CSC tests: simulation production

- ATLAS is expecting to produce fully-simulated events at a rate of up to 30% of the data-taking rate
 - i.e. 60 Hz, or 3M events/day, towards the end of 2007
- Right now we are able to simulate 2-3M events/week
 - Limited by the availability of facilities (CPU and storage) and by our software and middleware stability
- We plan to increase the production rate:
 - By a factor 2 by May-June 2007
 - By another factor 2 by October-November 2007
- According to MoU pledges, this is still a long way lower than nominally available capacities
 - But we know that not all pledged capacities actually exist and are available to us
- On our side we are working on improving our production software quality
- We expect a similar commitment from middleware developers



CSC tests: distributed analysis

- Our distributed analysis framework (*GANGA*) allows job submission to 3 Grid flavours (*EGEE*, *OSG* and *NG*) as well as to the local batch system
- It is now interfaced with the *DDM* system
 - Work is in progress on improving the interfaces to metadata
- Near future plans:
 - Test Posix I/O functionality and performance for sparse event reading with different tools (*GFAL*, *rfio*, *dcap*, *xrootd*) and different back-ends (*DPM*, *dCache*, *Castor SEs*)
- In Spring 2007:
 - Test large-scale concurrent job submission
 - Measure the read performance for concurrent access to the same files by large number of jobs
 - Collect metrics for the number of replicas of each file that will be needed for data analysis as a function of the number of users of a given dataset



CSC tests: reprocessing

- There will be many reprocessing steps of 2007 data in the first half of 2008
 - But as long as 2007 data will (most likely) not be much, we can try to keep the "good" RAW data on disk all the time
- Real reprocessing at Tier-1s (and Tier-0 when not taking data) will only occur in the second half of 2008
- One essential component of the reprocessing framework is the "prestaging" functionality in SRM 2.2
 - If we want to seriously test reprocessing before that is available, we have effectively to implement it ourselves for each SE type
- We therefore decided to defer full reprocessing tests at Tier-1s (including recalling RAW data from tape) until SRM 2.2 with prestaging functionality will be available
 - In the meantime we can nevertheless test the environment at each Tier-1, taking the Tier-0 Management System (TOMS) as example



Calibration Data Challenge

- Scope of the CDC is to deploy and test the full calibration/alignment chain
- Many components are needed for the CDC to have success:
 - Simulation software supporting a misaligned and miscalibrated detector geometry
 - Simulation started already in 2006 with software release 12
 - Release 13 in April will be used to simulate runs with more complex misalignment modes
 - Algorithmic software to calculate the calib/align constants
 - Mostly already available
 - Conditions database to store calib/align constants
 - Already implemented in COOL for almost all detectors
 - Distribution of the Conditions DB to the Tier-1s for remote processing
 - Using the tools developed by the 3D project
 - Reconstruction algorithms that use the parameters stored in COOL
 - Mostly already available; will be completed for release 13
- We currently plan to start the final phase of the CDC as soon as release 13 is available
 - Using initially existing data simulated with release 12
- The more complex modes will be exercised later in the Spring/Summer 2007



CSC tests: data streaming

- NB: in this context, “data streaming” means “separating RAW events into streams according to their trigger signature, at the end of HLT and of Tier-0 processing”, not “providing a continuous data flow between sites”
- The data streaming tests (and even more later the FDR) implement, for the first time, the file structure and organization of our data (luminosity blocks, trigger streams, file merging, etc.)
- They tests effectively assemble and integrate all software components that are needed to set up the full processing chain for real data
 - Including the software to be run on the Tier-0, the dataset creation systems, data and metadata cataloguing etc.
- Data streaming tests started in Autumn 2006 and will continue at increasing levels for the next few months
 - At the end of the May 2007 Tier-0 processing test, we foresee to merge the two tests into a single suite
 - This will be from then on the testbed for the treatment of real data



Full Dress Rehearsal

- The FDR in July-October 2007 will test the functionality and performance of the complete Software & Computing system ahead of the first data-taking period
 - It will progressively integrate the infrastructure prepared and tested in the first half of 2007 in the separate tests described so far
- Once completed, it will allow us to inject simulated events in RAW data format into (the later stages of) the TDAQ system and pass them on to the Tier-0 and beyond, including processing and final data distribution and analysis
- The first phase will start in July 2007 using s/w release 13 and building on the infrastructure already set up by the Data Streaming, Tier-0 and Data distribution tests
- The final phase in September-October will include the full system tests and use release 14 (foreseen for August 2007)
 - The major aim is to have all parts of the system running concurrently and stably at a rate as close as possible to the nominal data-taking rate (200 Hz average) by the end of October
 - In order to test the global computing infrastructure, simulation production and reprocessing will have to run at the same time, including their data distribution, at a rate as close to nominal as possible



Grid Tools: Data Management

- We depend on a number of Grid data management tools:
 - FTS, LFC, SRM, lcg-utils
- These tools must run RELIABLY and with HIGH PERFORMANCE
- Our own DDM infrastructure relies on the performance of the underlying tools
- A working SRM 2.2 is absolutely necessary for our DDM operations as soon as possible this year
 - We are now suffering from many problems with the instabilities of current SRM-1 installations
- A few nominal rates for reference:
 - Data transfer rates:
 - 1 GB/s export rate from CERN to all Tier-1s (each one takes its fraction), 1/3 of which to tape (T1D0) and 2/3 to disk (TOD1)
 - Each Tier-1 will export AODs at a maximum of 20 MB/s per associated Tier-2
 - Each Tier-1 will export reprocessed data at the same rate as ESD/AOD data received from Tier-0 and receive other production at half that rate
 - Fully-simulated data will be only 30% of real data but events are 50% larger: add a factor ~1.5 to the data transfer rate due to reprocessing
 - File registration rates:
 - Up to 30k RAW files/day will be produced by the online system
 - The same order of magnitude applies for reconstructed events
 - Multiply by ~3 for concurrent reprocessing and simulation production
 - O(100k) files/day will be generated by scheduled productions alone
 - We estimate that at least an equivalent number of files will be generated by user activities
- The data transfer and cataloguing middleware must be able to support the above rates continuously from Autumn 2007 onwards



Grid Tools: Workload Management

- As usual, the keywords are STABILITY, RELIABILITY, ROBUSTNESS, PERFORMANCE
 - The experience of the last few years is worrying
- We were forced so far to develop a much thicker interface to Grid tools than our original intention 5 years ago
- We have a stable version of the LCG CE that will soon also provide proper accounting information:
 - Good, we do not need anything more
- For job distribution, we are still using (and testing) different approaches:
 - gLite Resource Broker
 - Used in simulation production on the EGEE Grid
 - Unstable service, but crash program set up to turn it into a real stand-alone service
 - Condor-G
 - Alternative (currently more efficient) way to submit production jobs
 - Condor glide-ins
 - Experimental system still under test, using pilot jobs
 - PanDA
 - OSG developed environment that submits pilot jobs to sites holding input data
 - Under test on the EGEE Grid using Condor-G to distribute the pilot jobs
 - Custom interface to the NG/ARC client to submit production and analysis jobs
- Some time in Spring 2007 we have to reduce the number of options by selecting the most promising tools



Grid Tools: Monitoring & Accounting

- Monitoring tools for job and data management are finally becoming available thanks to the efforts of the LCG/ARDA group
 - The ARDA "dashboard" provides increasingly complete coverage for jobs running on the Grid(s), and now also for data storage and data transfer facilities
- CPU usage accounting tools are presently making rapid progress but are still insufficient to establish "who did what, and where"
 - We need full coverage of job statistics for internal and external accounting
 - We also need different views (by site, activity group, role, time window etc.) for statistically accumulated data
 - And we need the possibility to get individual user's records in case of need
- Storage accounting tools are still primitive
 - Some of the new functionality of SRM 2.2 should help here
- In any case we need at least minimally complete accounting tools by mid-2007 to be able to enforce ATLAS internal policies



Grid Tools: Resource Management

- CPU capacity management:

- Finally after many years a few tools that will allow us to manage the CPU capacities across different activities are being deployed
- Such tools are not needed in case there is a single job queue, but this is not the case for ATLAS, therefore the CEs have to be aware of ATLAS groups and roles as defined in the VOMS database
 - That's why we defined groups and roles in the first place...
- It is more than evident that by the time we start taking data later this year we have to reserve allocations at Tier-1s and Tier-2s for high priority activities, such as detector studies, calibrations etc, and move away from the initial "free for all" model

- Storage allocation management:

- We would like very much not to have to reinvent the wheel ourselves
- ACLs have existed for a long time in the Unix world and this is all we need
 - Of course we can try to use more complex schemes too but this is not our preference
- Again VOMS groups and roles should be used to define access rights to different storage areas
 - But not all necessary tools are on the market right now



Summary 2007 timeline

- Running continuously throughout the year (increasing rates):
 - Simulation production
 - Cosmic ray data-taking (detector commissioning)
- January to June:
 - Data streaming tests
- February through May:
 - Intensive Tier-0 tests
- From February onwards:
 - Data Distribution tests
- From March onwards:
 - Distributed Analysis (intensive tests)
- May to July:
 - Calibration Data Challenge
- June to October:
 - Full Dress Rehearsal
- November:
 - **GO!**



Conclusions

- ATLAS has a comprehensive plan for Computing System Commissioning
 - It builds on the *Data Challenges* and *Service Challenges* performed in the last few years
- The aim is to arrive at the beginning of LHC operation with a well-tested, robust and satisfactorily functional Software & Computing system
 - Priority will be given to testing and integration of the components that are absolutely necessary to operations
 - Many "nice things to have" from now on will have to be thoroughly tested before being integrated into the ATLAS software base
 - Similarly for the Grid middleware tools we use
- Stability and robustness are the keywords for this year



Conclusions

- ATLAS has a comprehensive plan for Computing System Commissioning
 - It builds on the Data Challenges and Service Challenges performed in the last few years
- The aim is to arrive at the beginning of LHC operation with a well-tested, robust and satisfactorily functional Software & Computing system
 - Priority will be given to testing and integration of the components that are absolutely necessary to operations
 - Many "nice things to have" from now on will have to be thoroughly tested before being integrated into the ATLAS software base
 - Similarly for the Grid middleware tools we use
- Stability and robustness are the keywords for this year

I hope to be able to show you real LHC data processed on the Grid
in a year's time!