ATLAS DDM Requirements for 2007

Dario Barberis

CERN & Genoa University/INFN

ATLAS 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)

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

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 (T0D1)
    - 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

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 and 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!