ATLAS Software & Computing Readiness for Data-taking

Dario Barberis

(CERN & Genoa University/INFN)

Software release strategy

- We want to keep a "stable" release in production for Tier-0, HLT and online monitoring
  - Only urgent bug fixes go into production caches in this branch
  - A minor release is built every ~2 weeks to consolidate all bug fixes collected in the meantime
- And another "stable" release for simulation production, reprocessing and physics analysis
  - Bug fixes and feature enhancements from Tier-0 and this branch are added to production caches in this series
  - A minor release is built as needed when the production cache grows too much and needs consolidation (usually about once/month)
- Longer-term developments go into "next" (non-production) branch
- Right now we have:
  - Release 15.5.3 with cache 15.5.3.X for Tier-0, HLT and P1 monitoring
  - Release 15.3.1 with cache 15.3.1.Y for Grid productions and analysis (in production)
  - Release 15.6.0 with cache 15.6.0.Z for Grid productions and analysis (in validation)
  - Release 16.0.0 open for long-term developments

Strategy

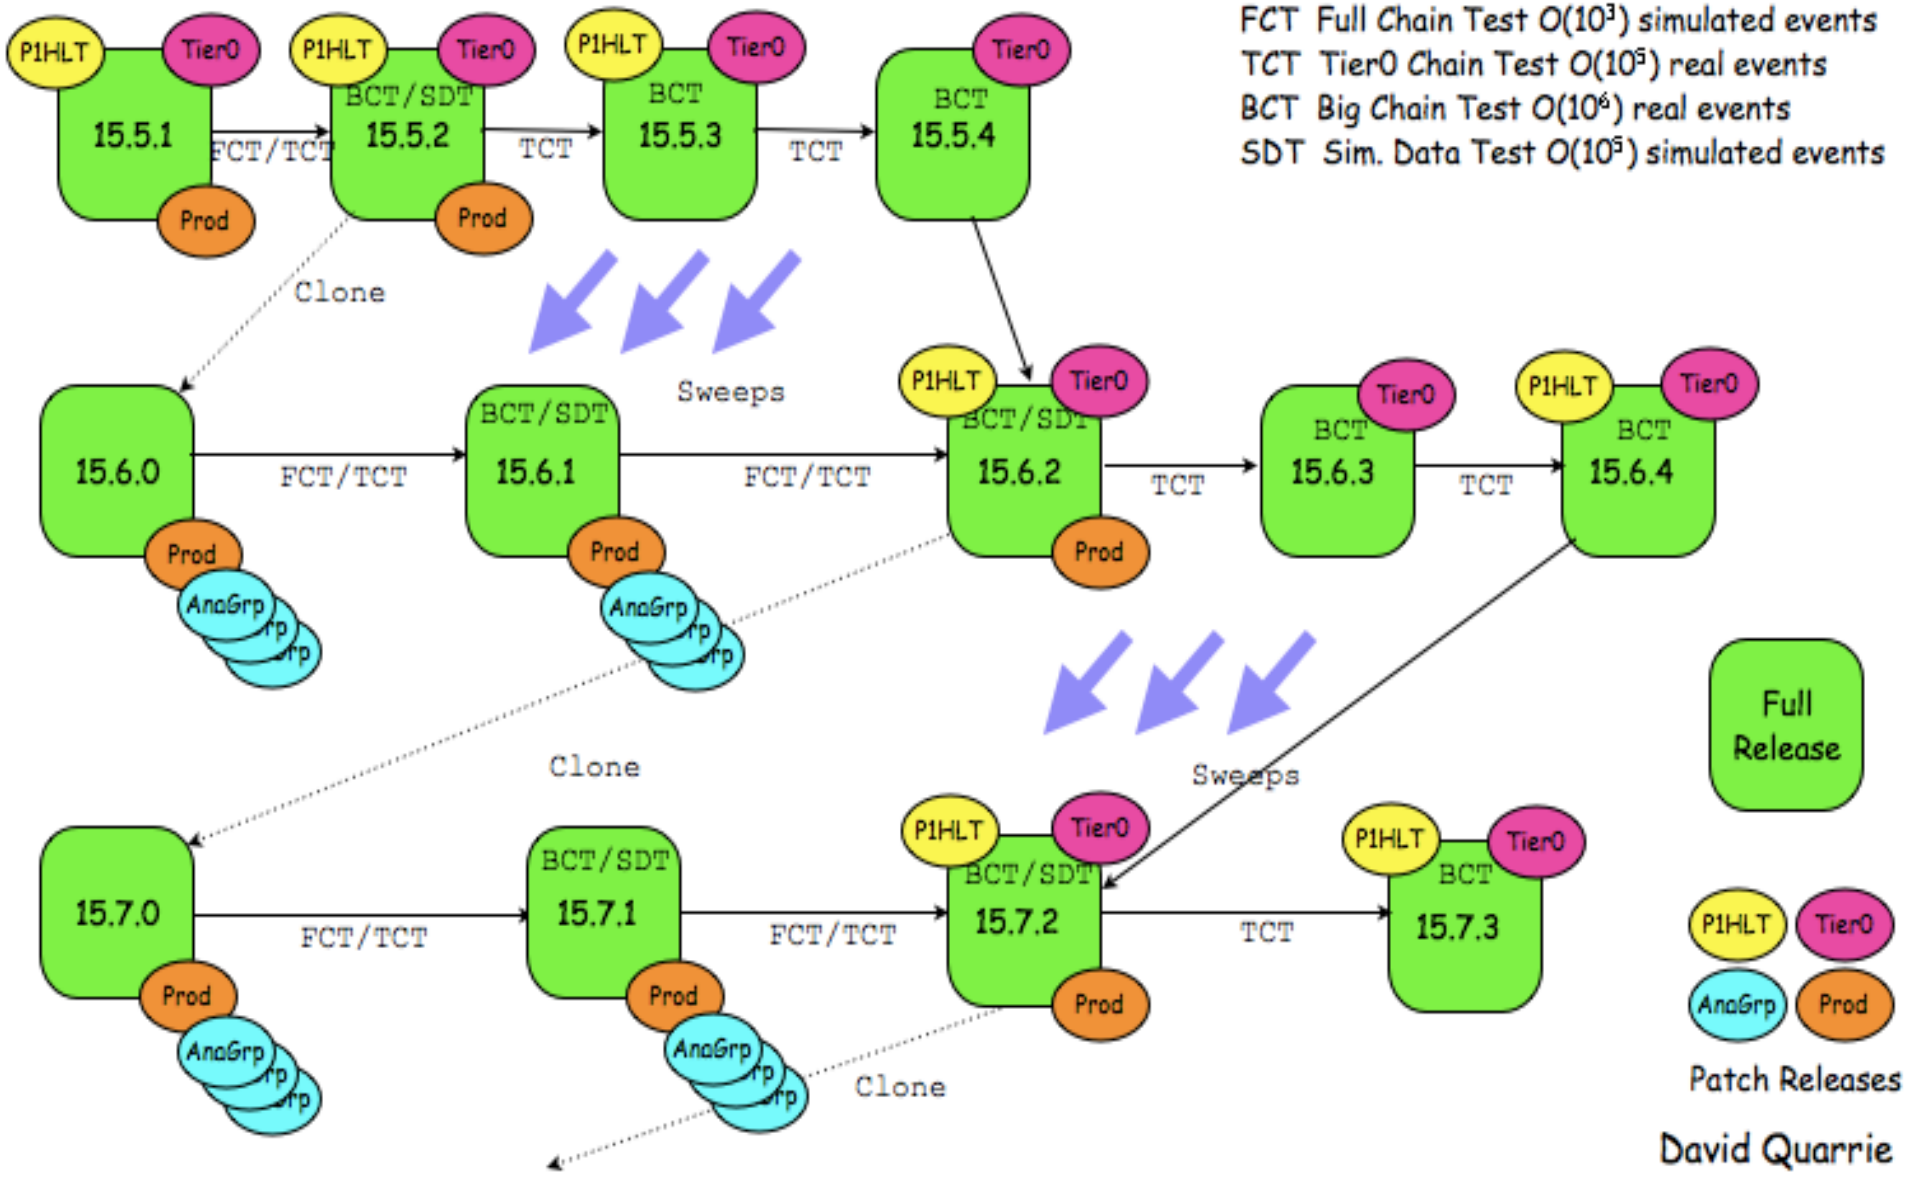
FCT Full Chain Test O(10³) simulated events
TCT Tier0 Chain Test O(10⁵) real events
BCT Big Chain Test O(10⁶) real events
SDT Sim. Data Test O(10⁵) simulated events
P1HLT
Tier0
15.5.1
FCT/TCT
BCT/SDT
15.5.2
TCT
BCT
15.5.3
TCT
BCT
15.5.4
Prod
Clone
Sweeps
15.6.0
FCT/TCT
BCT/SDT
15.6.1
FCT/TCT
BCT/SDT
15.6.2
TCT
BCT
15.6.3
TCT
BCT
15.6.4
AnaGrp
Clone
Sweeps
15.7.0
FCT/TCT
BCT/SDT
15.7.1
FCT/TCT
BCT/SDT
15.7.2
TCT
BCT
15.7.3
Clone
Full Release
Patch Releases
David Quarrie

Platforms

- We build code for SLC4/gcc3.4 and SLC5/gcc4.3
- Current SLC4 builds also run on SL(C)5 machines provided the right set of compatibility libraries is available
- The SLC5/gcc4.3 (32-bit mode) is in validation right now
  - So far all looks good — differences wrt SLC4/gcc3.4 are compatible with different compilers and mathematical libraries
  - We will probably move to SLC5/gcc4.3 at Tier-0 in January unless there is some show-stopper
- SLC5/gcc4.3 seems to save 10-30% CPU time depending on the application
- As soon as all Grid sites move to SL(C)5 we can stop building on SLC4
- We are also validating SLC5/gcc4.3 (64-bit mode) for special applications (high-luminosity pile-up) that need lots of memory

Software performance: memory

- Strong focus on memory usage improvements in the last year
- Lots of work in core, system, combined performance groups
- Continuous monitoring

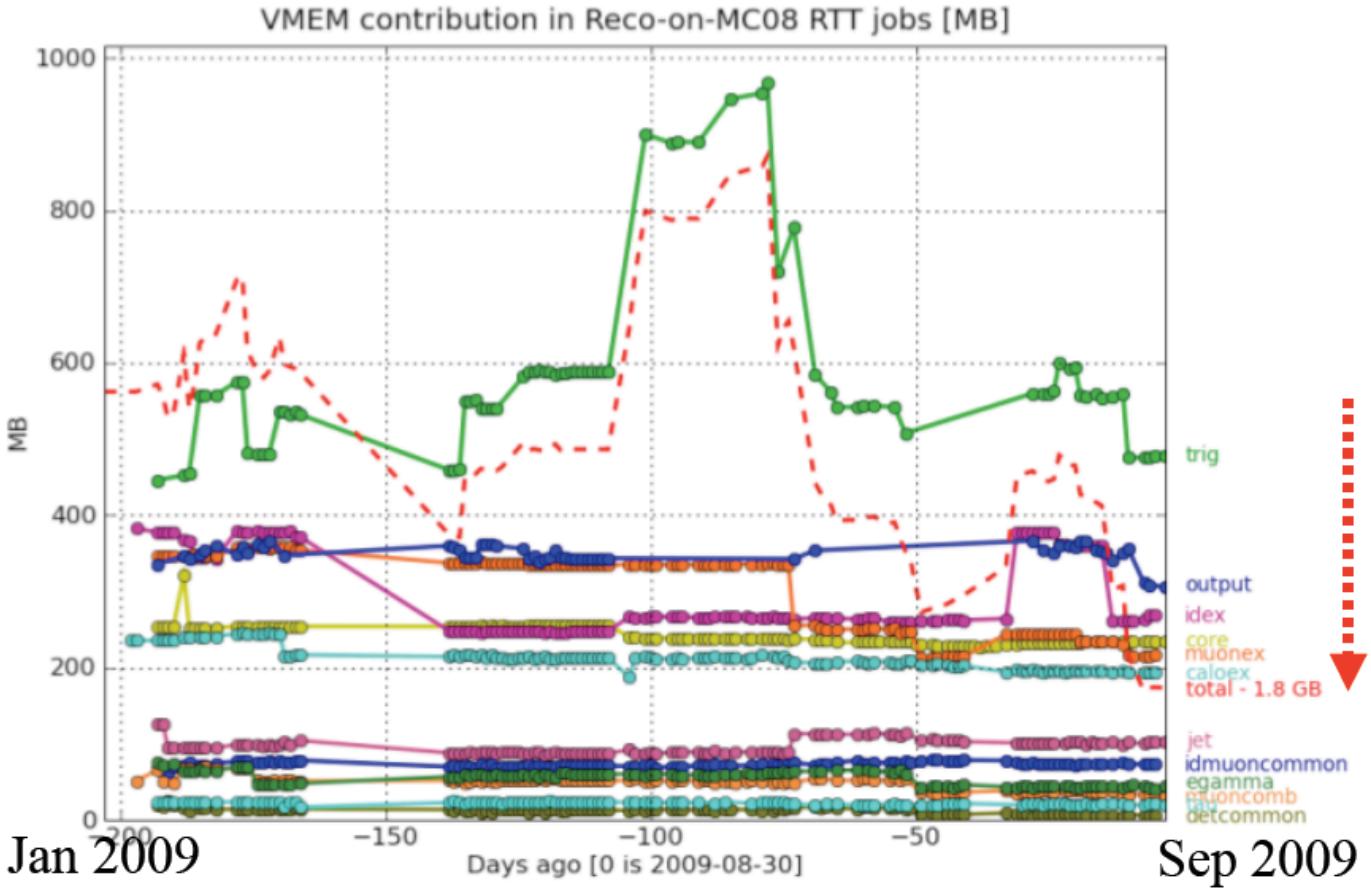

Overall gain 400 MB

(January-September 2009)

Software performance: CPU

- CPU for reconstructing MC t-tbar with release AtlasProduction 15.1.0.2
  - more busy than average ATLAS events
- ID and trigger simulation show significant increase of CPU wrt luminosity, however:
  - ID no special tuning (in any case can be trimmed down overnight with increased threshold)
  - Trigger : loose menu for debugging/tuning (only for MC)

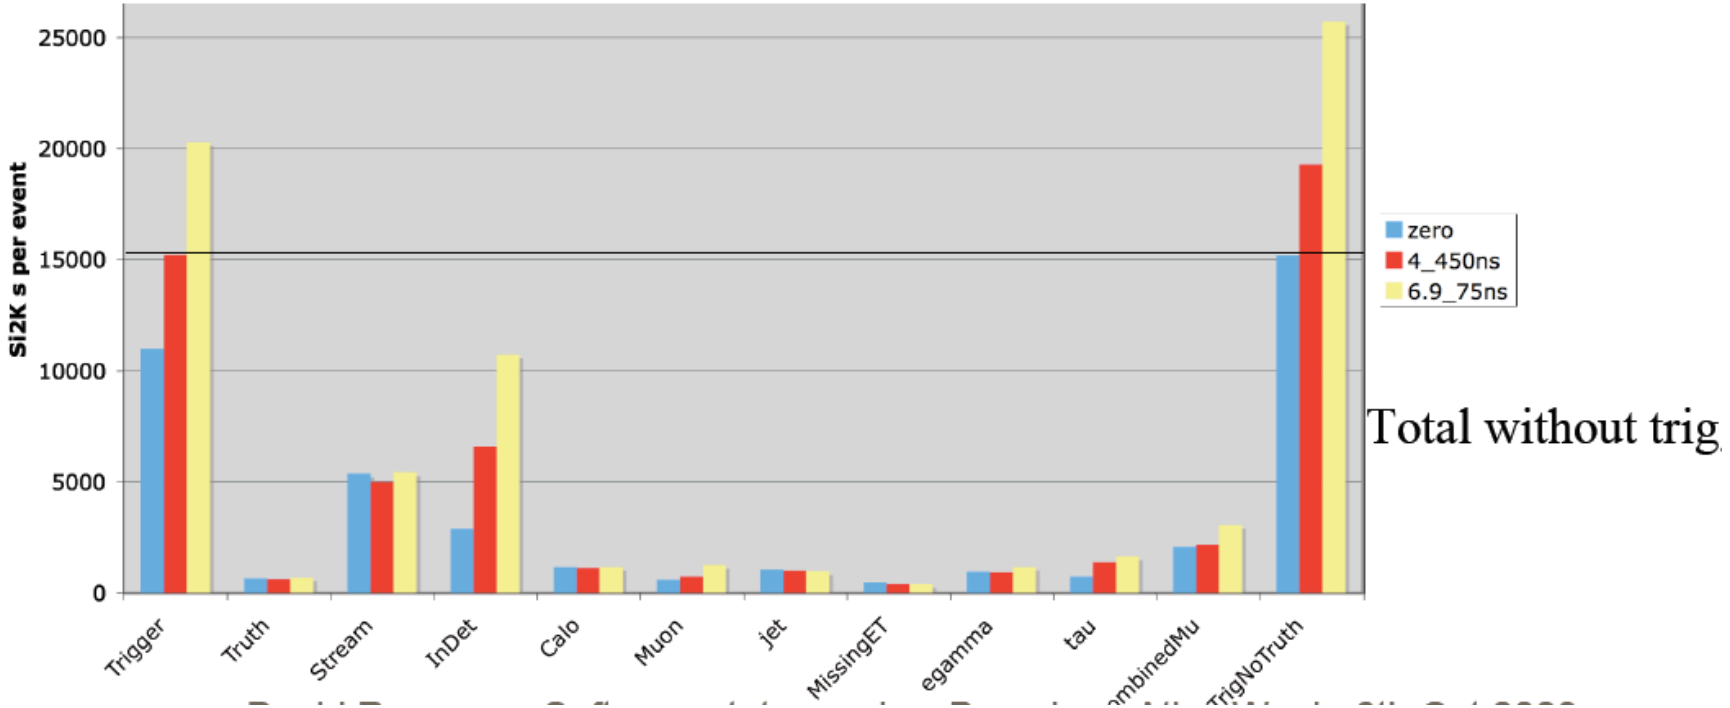

Total without trig

Simulation CPU performance

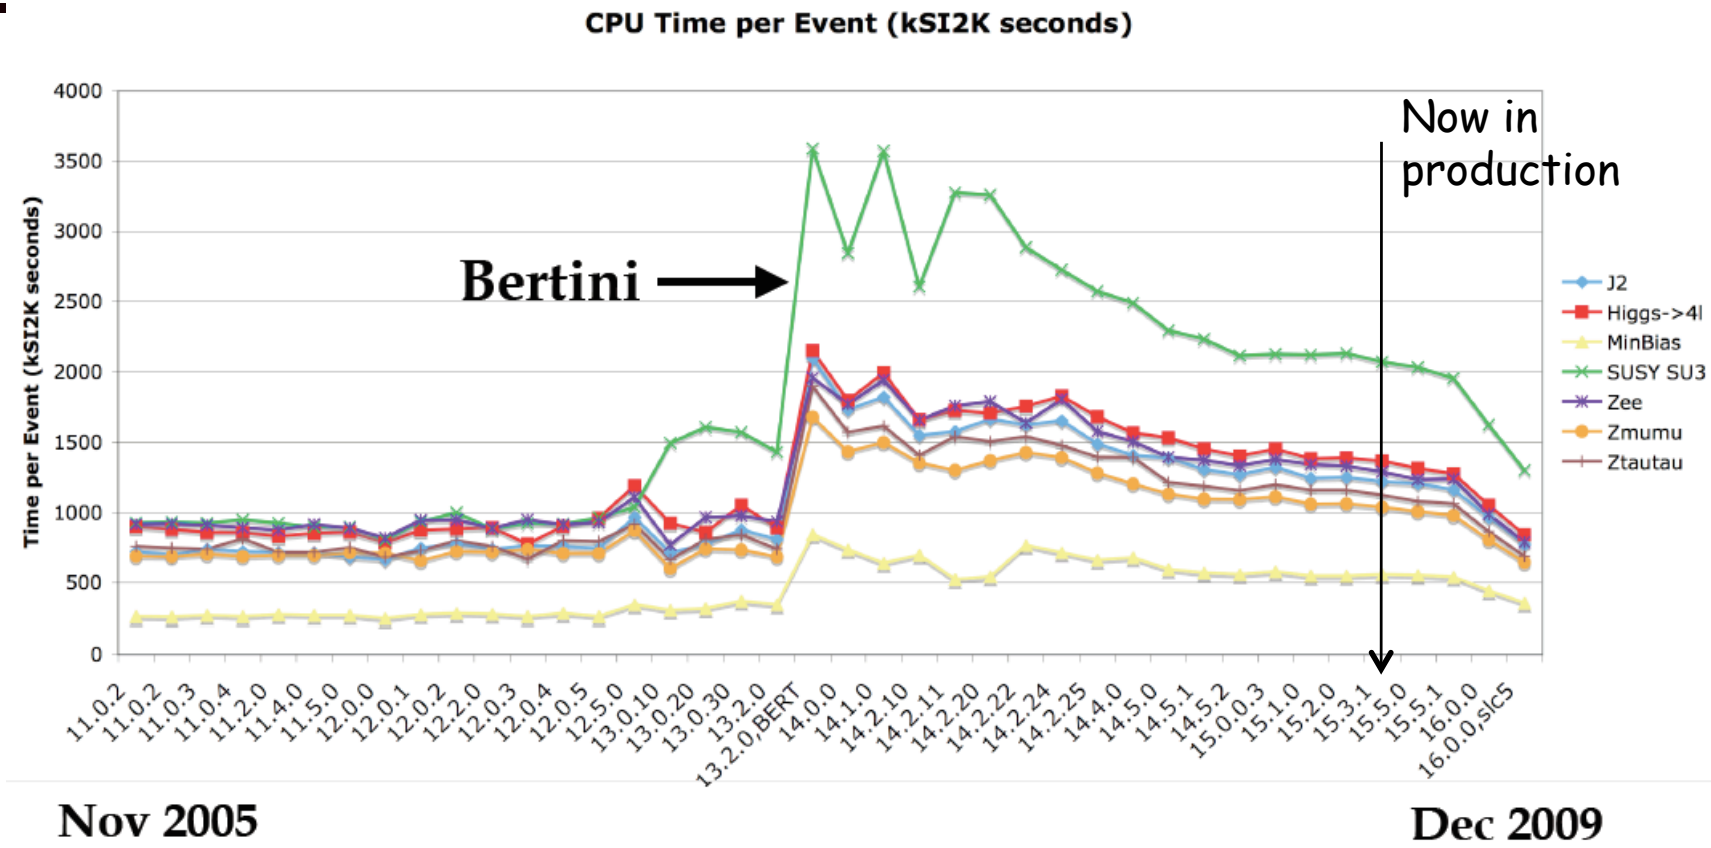

Nov 2005 **Dec 2009**

- About 20% of CPU time is accessing the magnetic field
  - Improved by the "stepper dispatcher" in release 16.0.0
  - Also improved by a new stepper in 15.5.0
- Init time reduced by >50% (9 min -> 4 min)
  - Changes in voxelization (1.5 minutes) and EM physics initialization (3.5 minutes)
- Parameterized showers (7 times faster) are still in validation

Throughput test last month

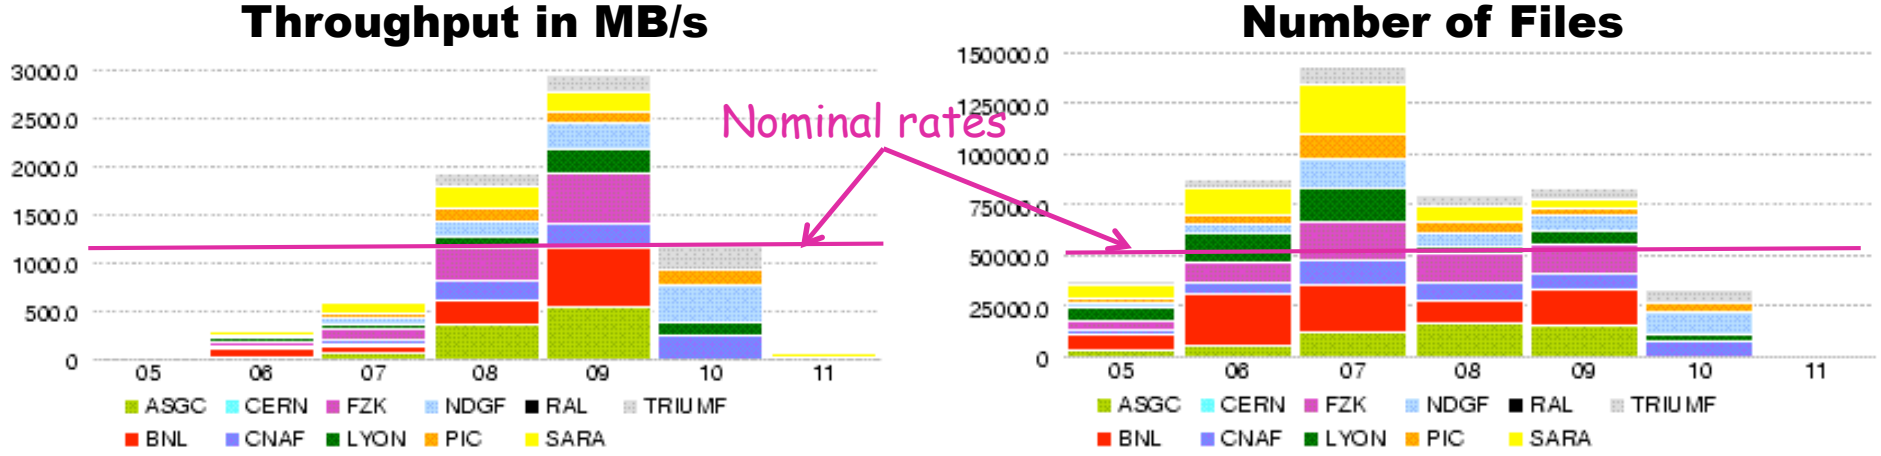

- High throughput tests are run periodically to test links between Tier-0 and Tier-1s, between each pair of Tier-1s, and between each Tier-1 and the associated Tier-2s
  - Last month's test was Tier-0➔Tier-1s to test a new FTS version that supports automatic checksums
    - Phase 1 (5th to 7th) : large number of small files
      - O(1/10) than Nominal Size
    - Phase 2 (7th to 9th) : lower number of large files
      - Nominal Size
    - On the 9th checksum verification was turned on
      - for dCache and DPM destinations
  - Result was partially successful as the new FTS version appears less robust (even if more functional) than the previous one
    - Further development and tests are needed before having it in production

Simulation and reprocessing

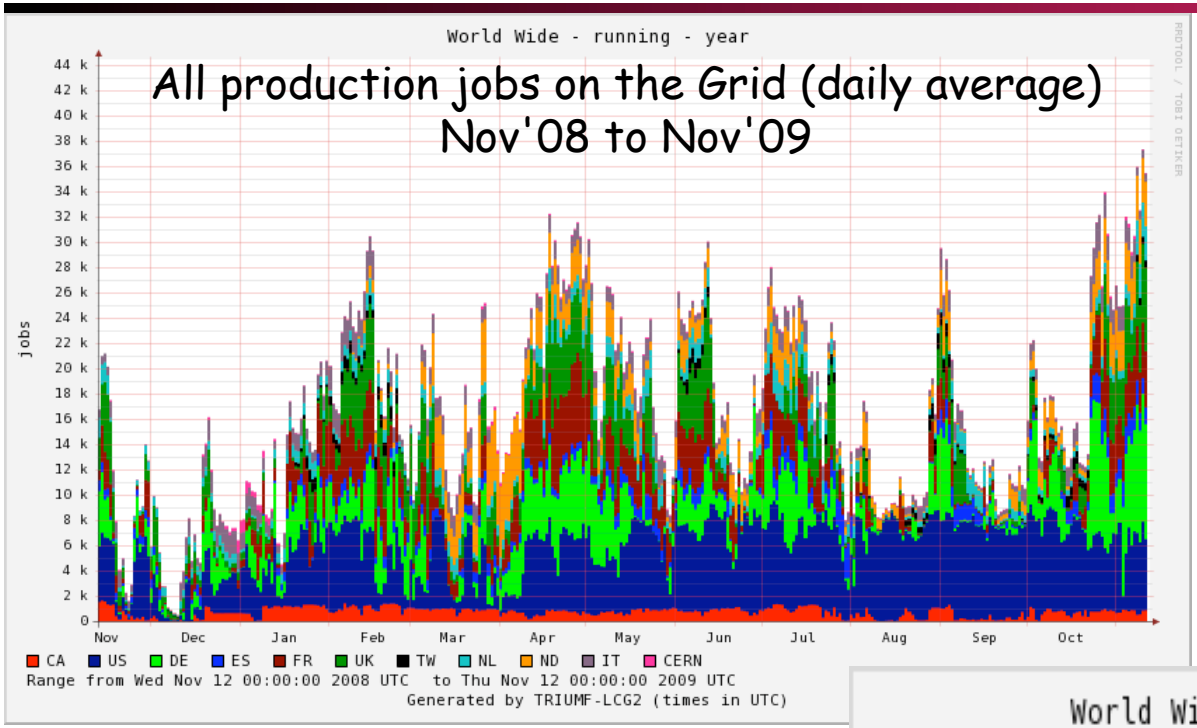

- Simulation production on the Grid started in 2003 and continued with increasing intensity
  - Reached this week 40k simultaneously running jobs at peak usage

- Several reprocessing campaigns were performed this year using 2008 and 2009 cosmics datasets
  - The latest one in September starting from ESD data

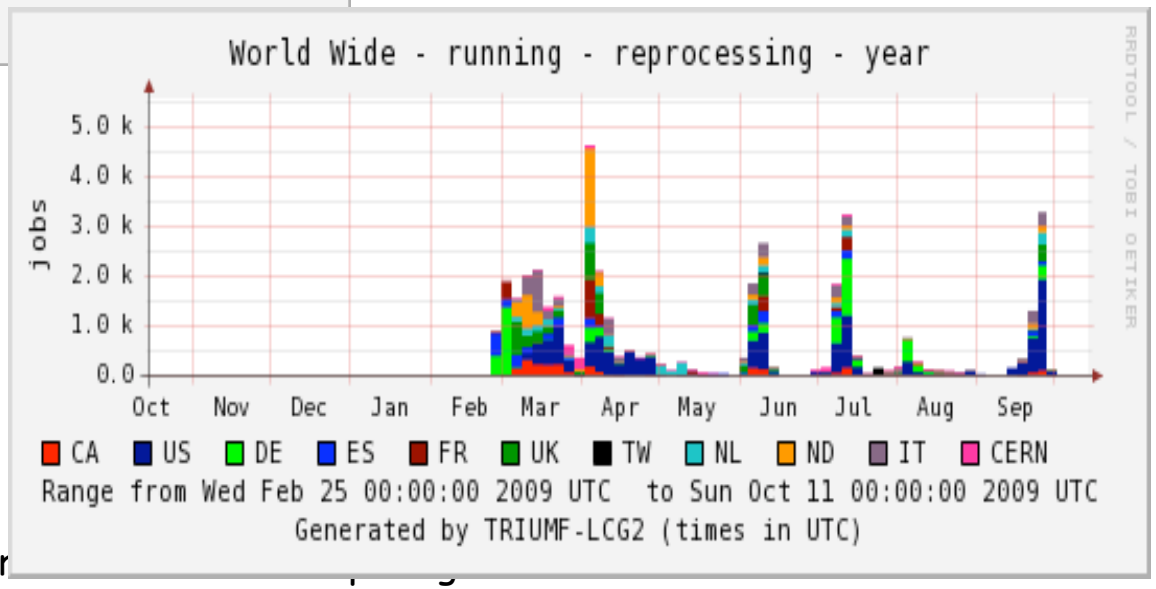

Distributed Analysis Tests

- Several DA tests ran for many years
  - GangaRobot first with single jobs going regularly to each site
  - HammerCloud with intensive cloud (or site) tests: flood of jobs to saturate the CPUs and measure the storage -> CPU transfer rates
- UAT (User Analysis Tests) was run at the end of October — post-mortem discussion next week
- As a result of HC tests several Tier-2 sites improved h/w and performance

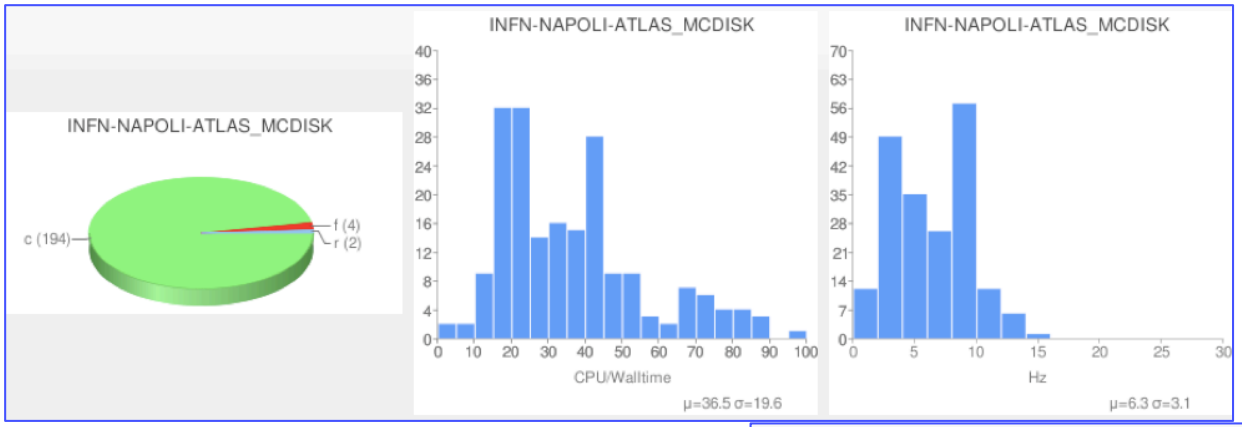

Naples Tier-2 performance in January 2009

Naples Tier-2 performance in August 2009

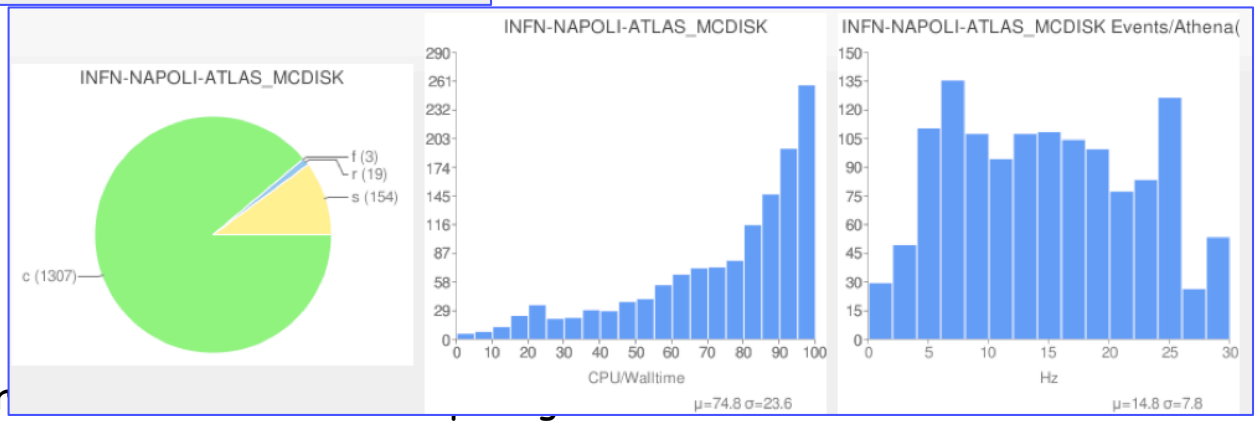

Dar

Conclusions and outlook

- ATLAS has a robust Software & Computing system that can stand the impact of the first LHC collision data
- Nevertheless software development is never finished, as code optimization and improvements are always needed
- The Grid infrastructure, if used carefully and through official tools, can benefit all members of the Collaboration by providing computing power and data storage independently of the geographical location of the collaborators
- Operating this system needs a considerable amount of manpower but, thanks to its distributed nature, operations can be run for most tasks from home institutes

We are eagerly waiting for the first LHC data!