



# ATLAS Quarterly Report and Plans

Dario Barberis

CERN & Genoa University/INFN



# Outline

- Software release strategy for LHC turn-on
- Data distribution tests
- Simulation production
- SRM v2 storage classes
- CCRC'08-2 and FDR-2 reports later this week in CCRC Workshop
- Plan of activities until the end of 2008



# Software Releases

- Release 14.0.10 was deployed on 9 April
  - M7 Cosmics (May)
- Release 14.1.0 was deployed on 1 May
  - FDR-2 Tier-0 processing and monitoring
    - Using validated 14.1.0.Y (bug-fix cache built daily for rapid turn-around)
  - Global cosmics and initial single beam running
  - Baseline turn-on simulation production
- Release 14.2.0 built last week
  - Initial colliding beam running
  - Essential to focus on robustness and technical performance (cpu/memory) within the 14.2.X branch
  - "Completion" release 14.2.10 will be built in ~2 weeks after the first round of validation
  - Full deployment mid-July
- General strategy:
  - 3 open bug-fix projects in TagCollector
    - AtlasP1HLT for online bug-fix patches, built if/when needed
    - AtlasPoint1 for Tier-0 reconstruction and monitoring bug-fixes
    - AtlasProduction for simulation production on the Grid
  - All patches to the current release are included (if applicable) into the future release development
- Release 15.0.0 (baseline for 2009 running) foreseen for late Autumn 2008




# Data Distribution Tests

- The Throughput Tests (TT) continue (a few days/month) until all data paths are shown to perform at nominal rates
  - This includes:
    - a) Tier-0 → Tier-1s → Tier-2s for real data distribution
    - b) Tier-2 → Tier-1 → Tier-1s → Tier-2s for simulation production
    - c) Tier-1 ↔ Tier-1 for reprocessing output data
- Test a) is now OK almost everywhere
  - Run again several times, now with SRM v2 end-points
- Test b) is part of simulation production since a long time
- Test c) was run in all combinations last month
- The Functional Test (FT) is also run in the background approximately once/month in an automatic way
  - The FT consists in low rate tests of all data flows, including performance measurements of the completion of dataset subscriptions
  - The FT is run in the background, without requiring any special attention from site managers
  - It checks the response of the ATLAS DDM and Grid m/w components as experienced by most end users




# Example of Tier-1-Tier-1 Transfer

Cumulative data transfer rates  
between all ATLAS Tier-1s  
Only SRM v2.2 end points



Fraction of non-Transferred Datasets Thu 15th 2pm



Fraction of delayed (>24 hours)  
datasets during transfer between all  
ATLAS Tier-1s

Total datasets sample: 629 datasets


Source Dataset Distribution:

- INFN-T1\_DATADISK: 35
- TAIWAN-LCG2\_DATADISK: 28
- IN2P3-CC\_DATADISK: 78
- FZK-LCG2\_DATADISK: 71
- PIC\_DATADISK: 36
- BNL-OSG2\_DATADISK: 174
- TRIUMF-LCG2\_DATADISK: 28
- NDGF-T1\_DATADISK: 26
- SARA-MATRIX\_DATADISK: 101
- RAL-LCG2\_DATADISK: 52



# Distributed Simulation Production (1)

- Simulation production continues all the time on the 3 Grids (EGEE, OSG and NorduGrid)
  - The rate is limited by the needs and by the availability of data storage more than by resources
- Validation of simulation and reconstruction with release 14 is still in progress






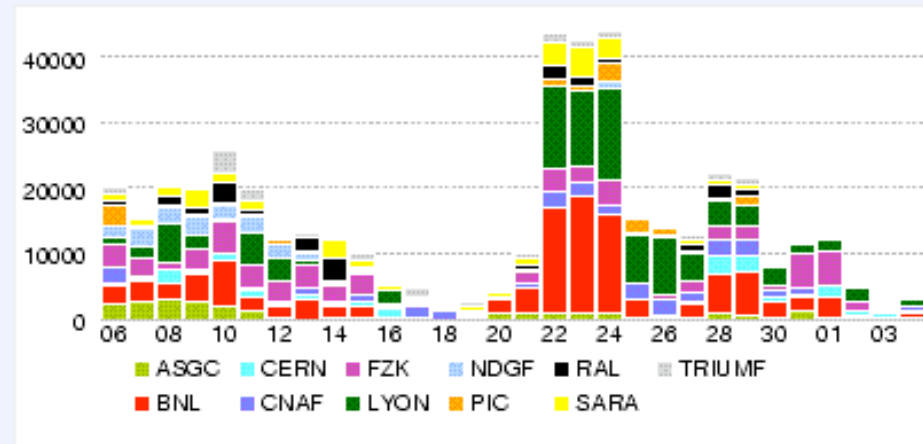
# Distributed Simulation Production (2)

Inter-cloud transfers of simulation production from 6 May to 4 June 2008


Throughput (MB/s)



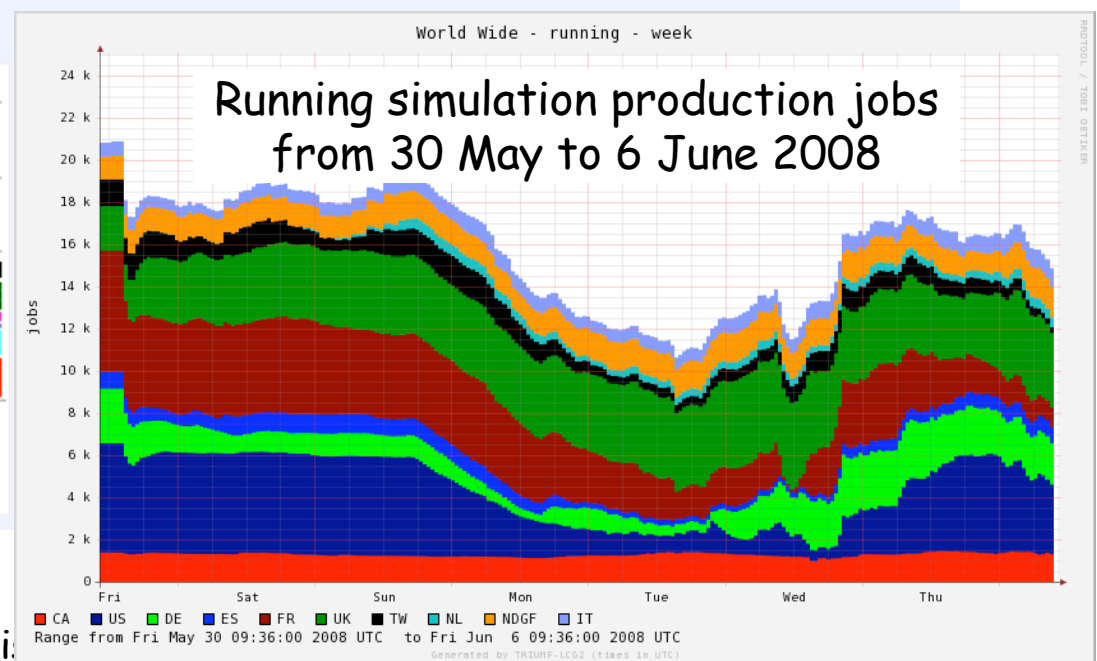
Data Transferred (GBytes)



Completed File Transfers

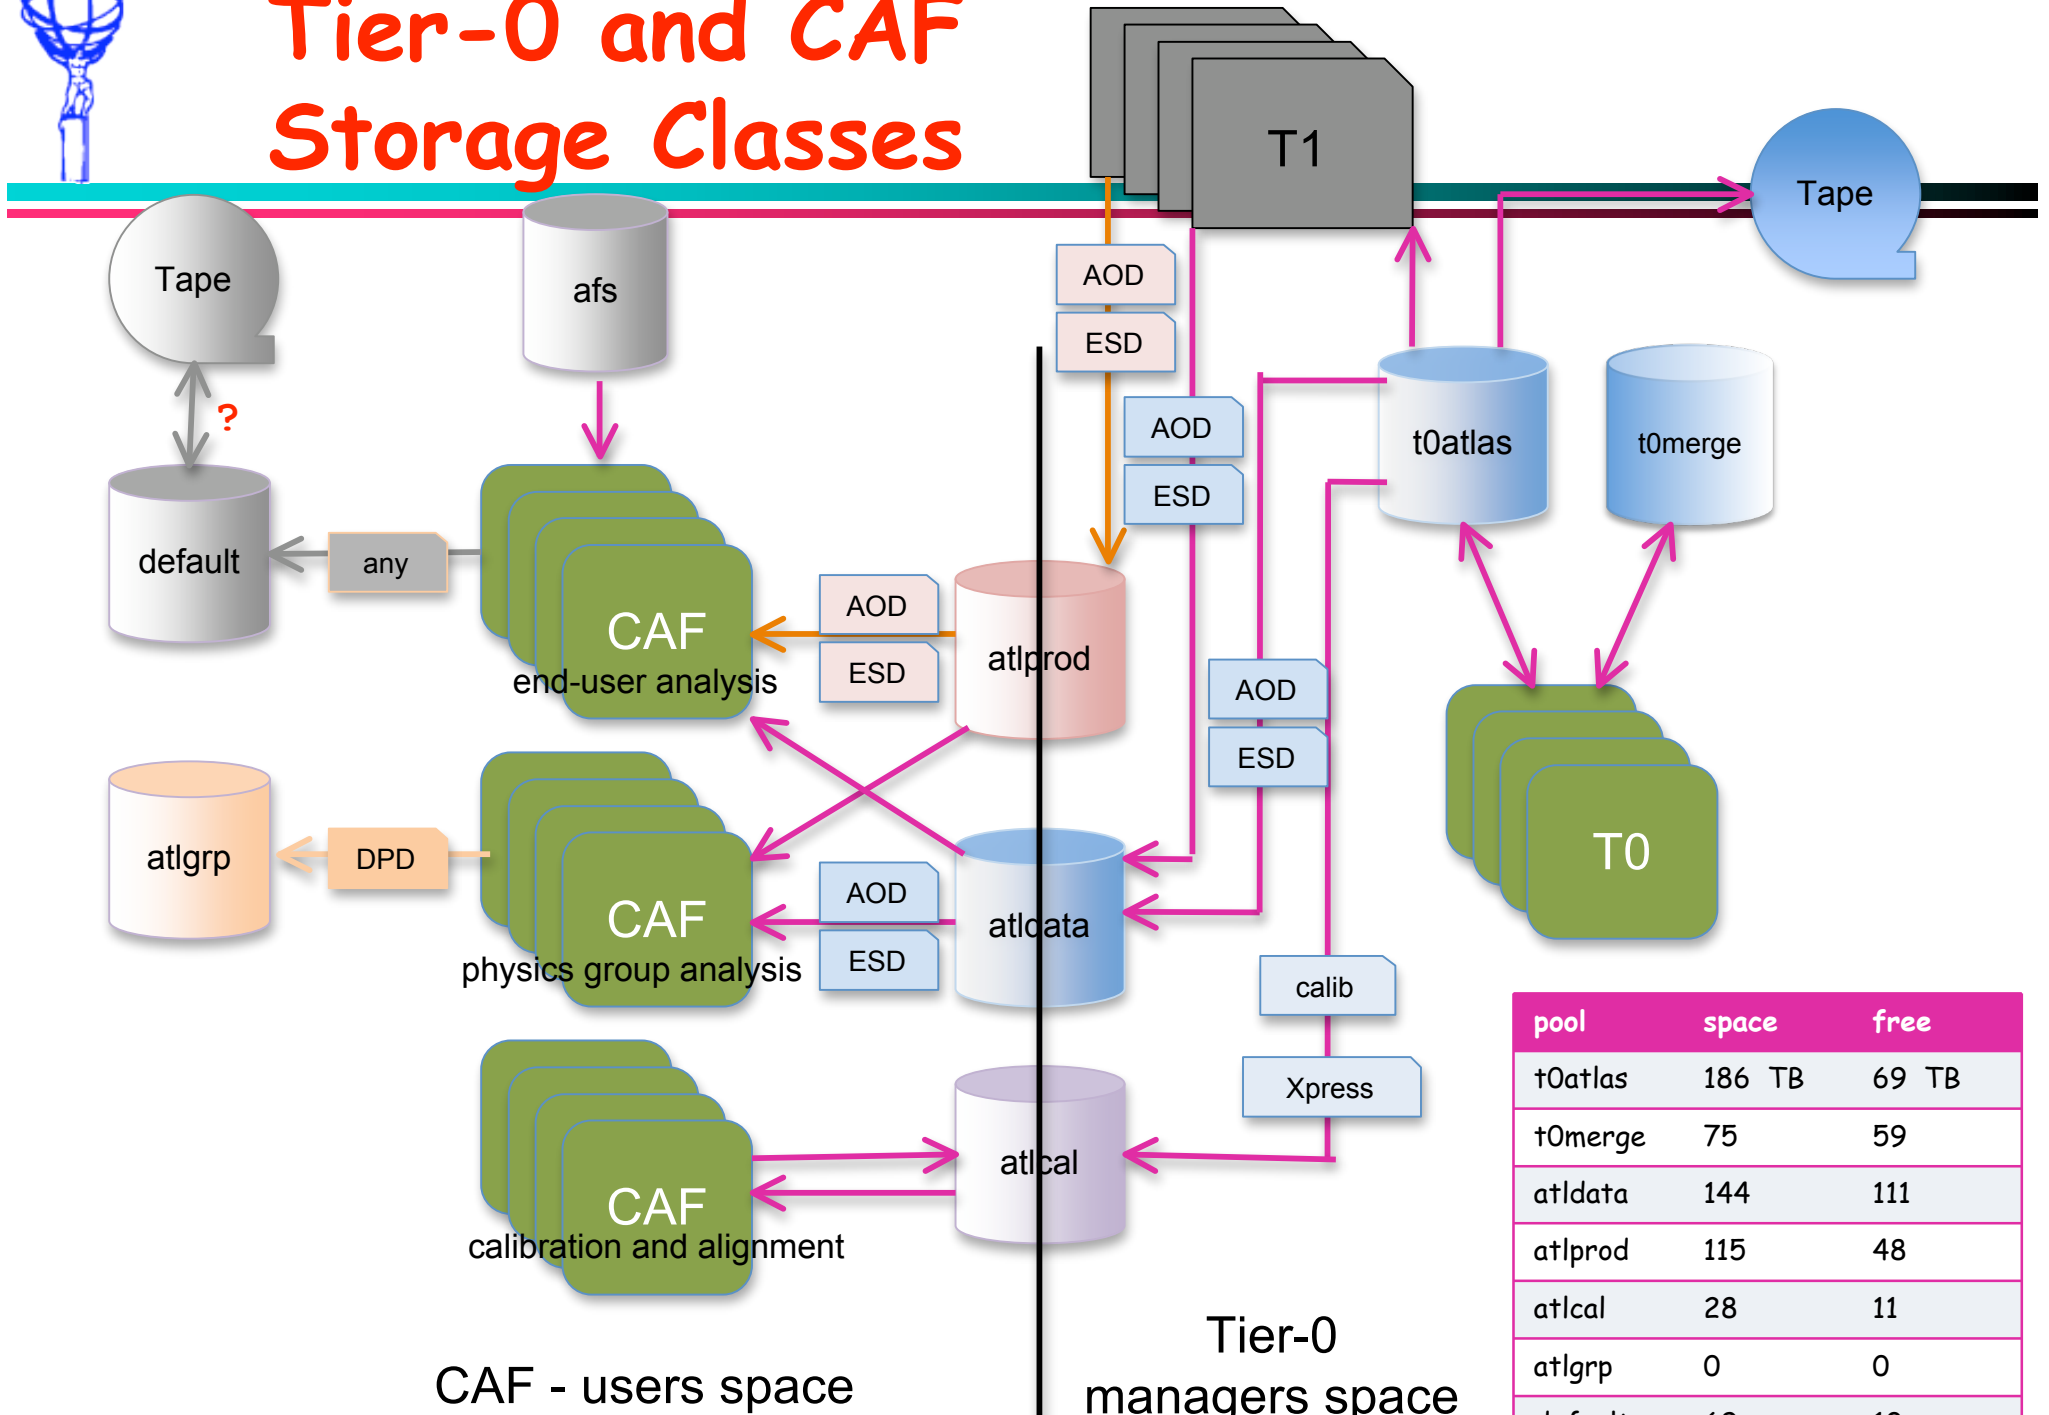


World Wide - running - week





# Tier-0 and CAF Storage Classes

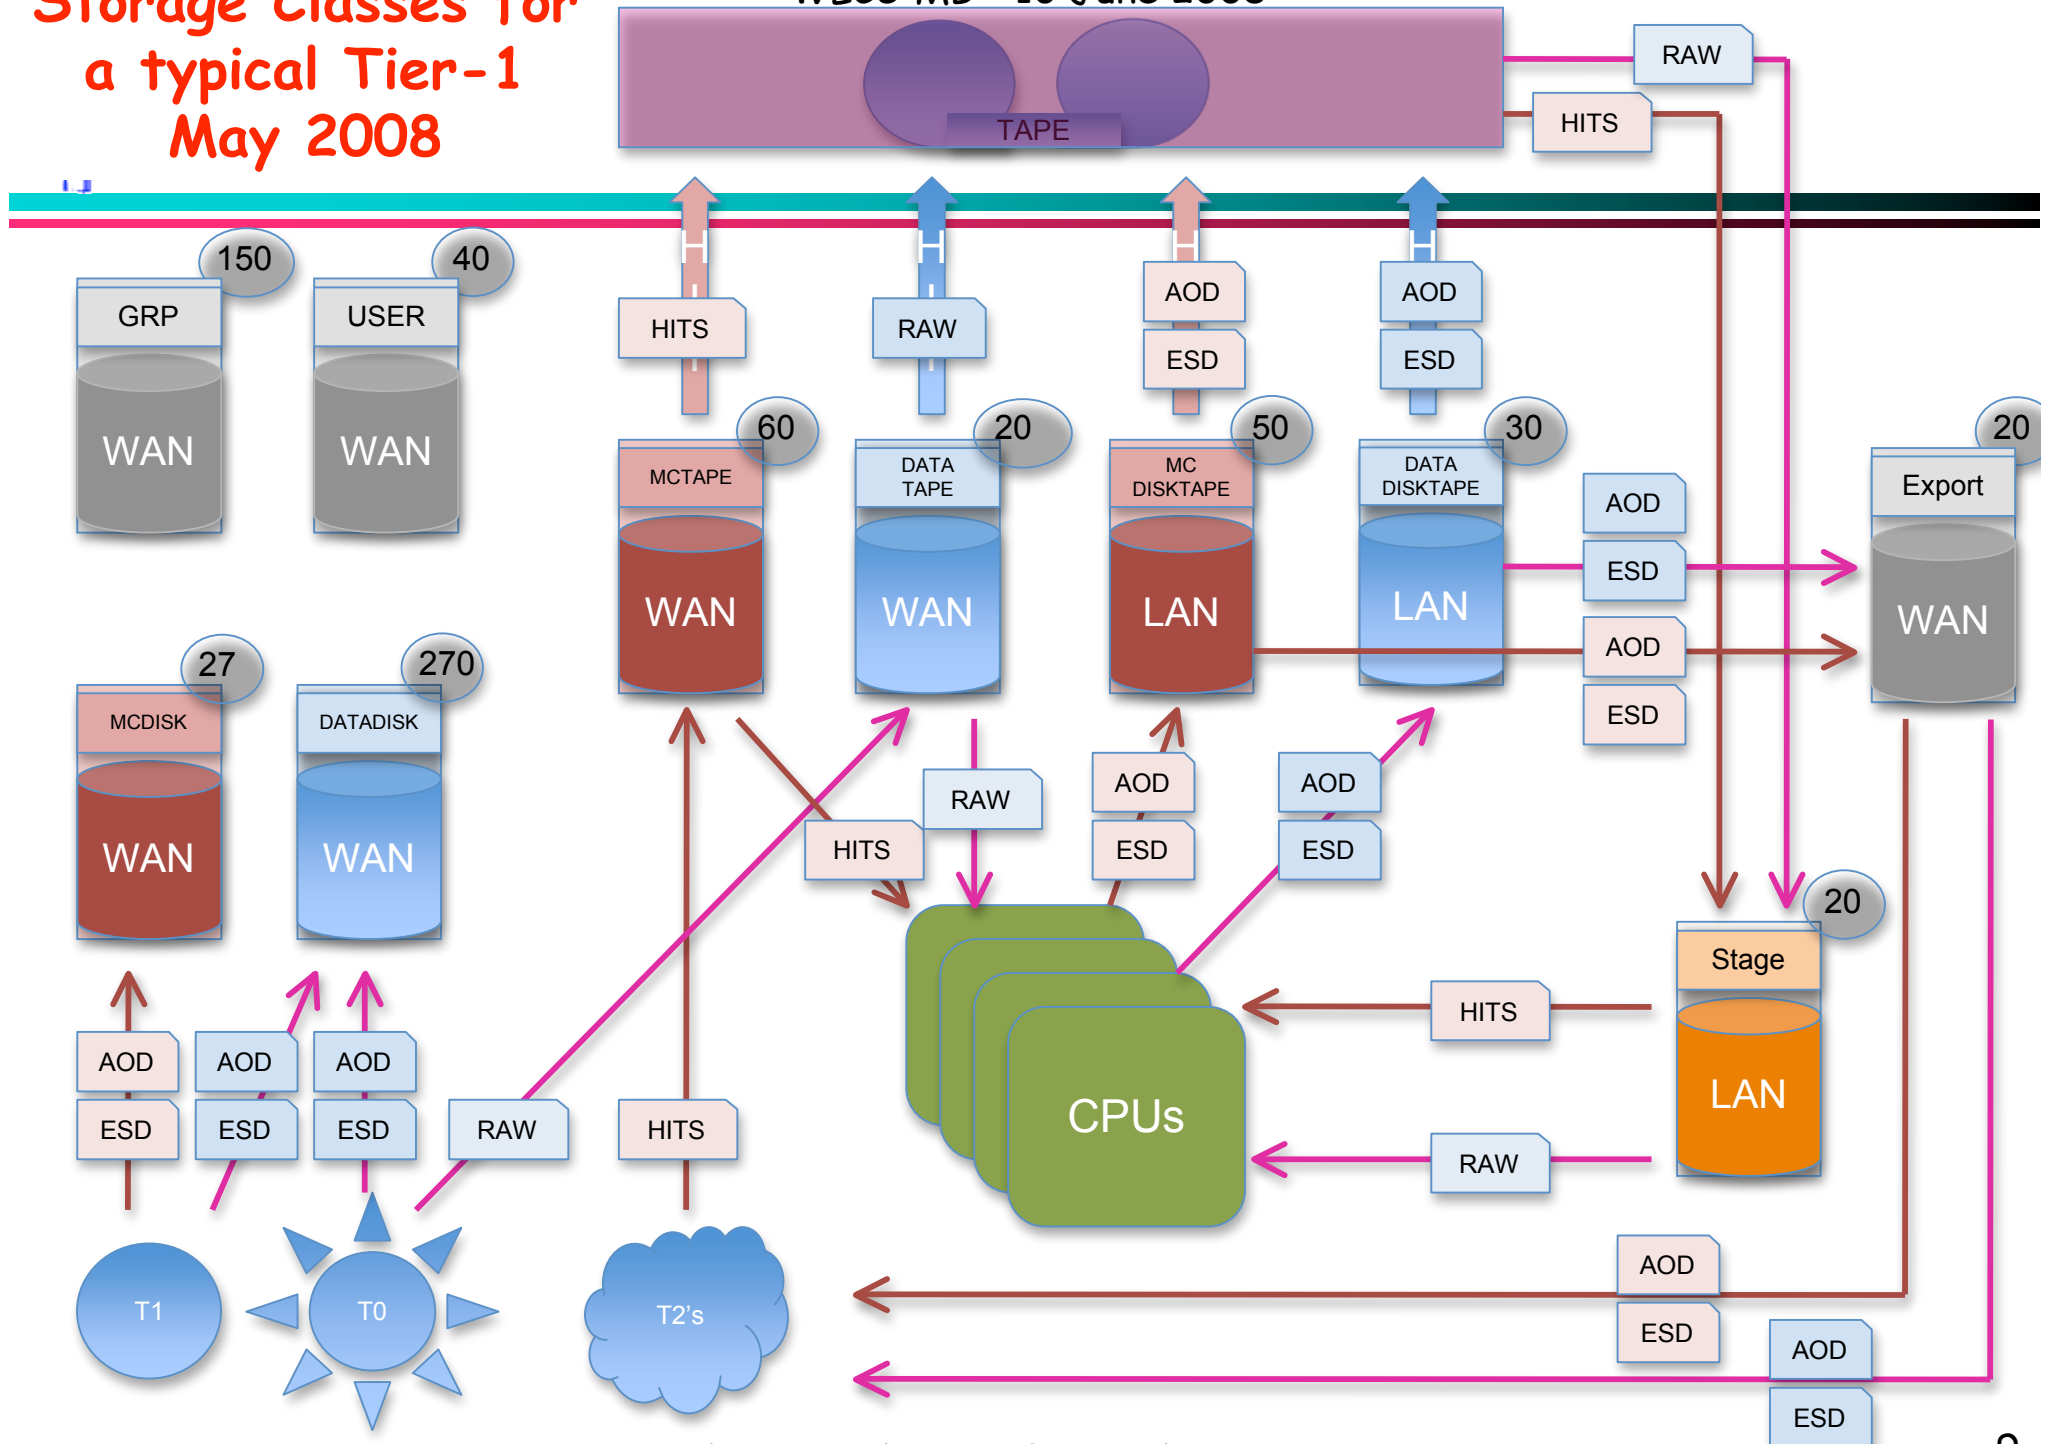


pool	space	free
t0atlas	186 TB	69 TB
t0merge	75	59
atldata	144	111
atlprod	115	48
atlcal	28	11
atlgrp	0	0
default	69	10



# Storage Classes for a typical Tier-1 May 2008

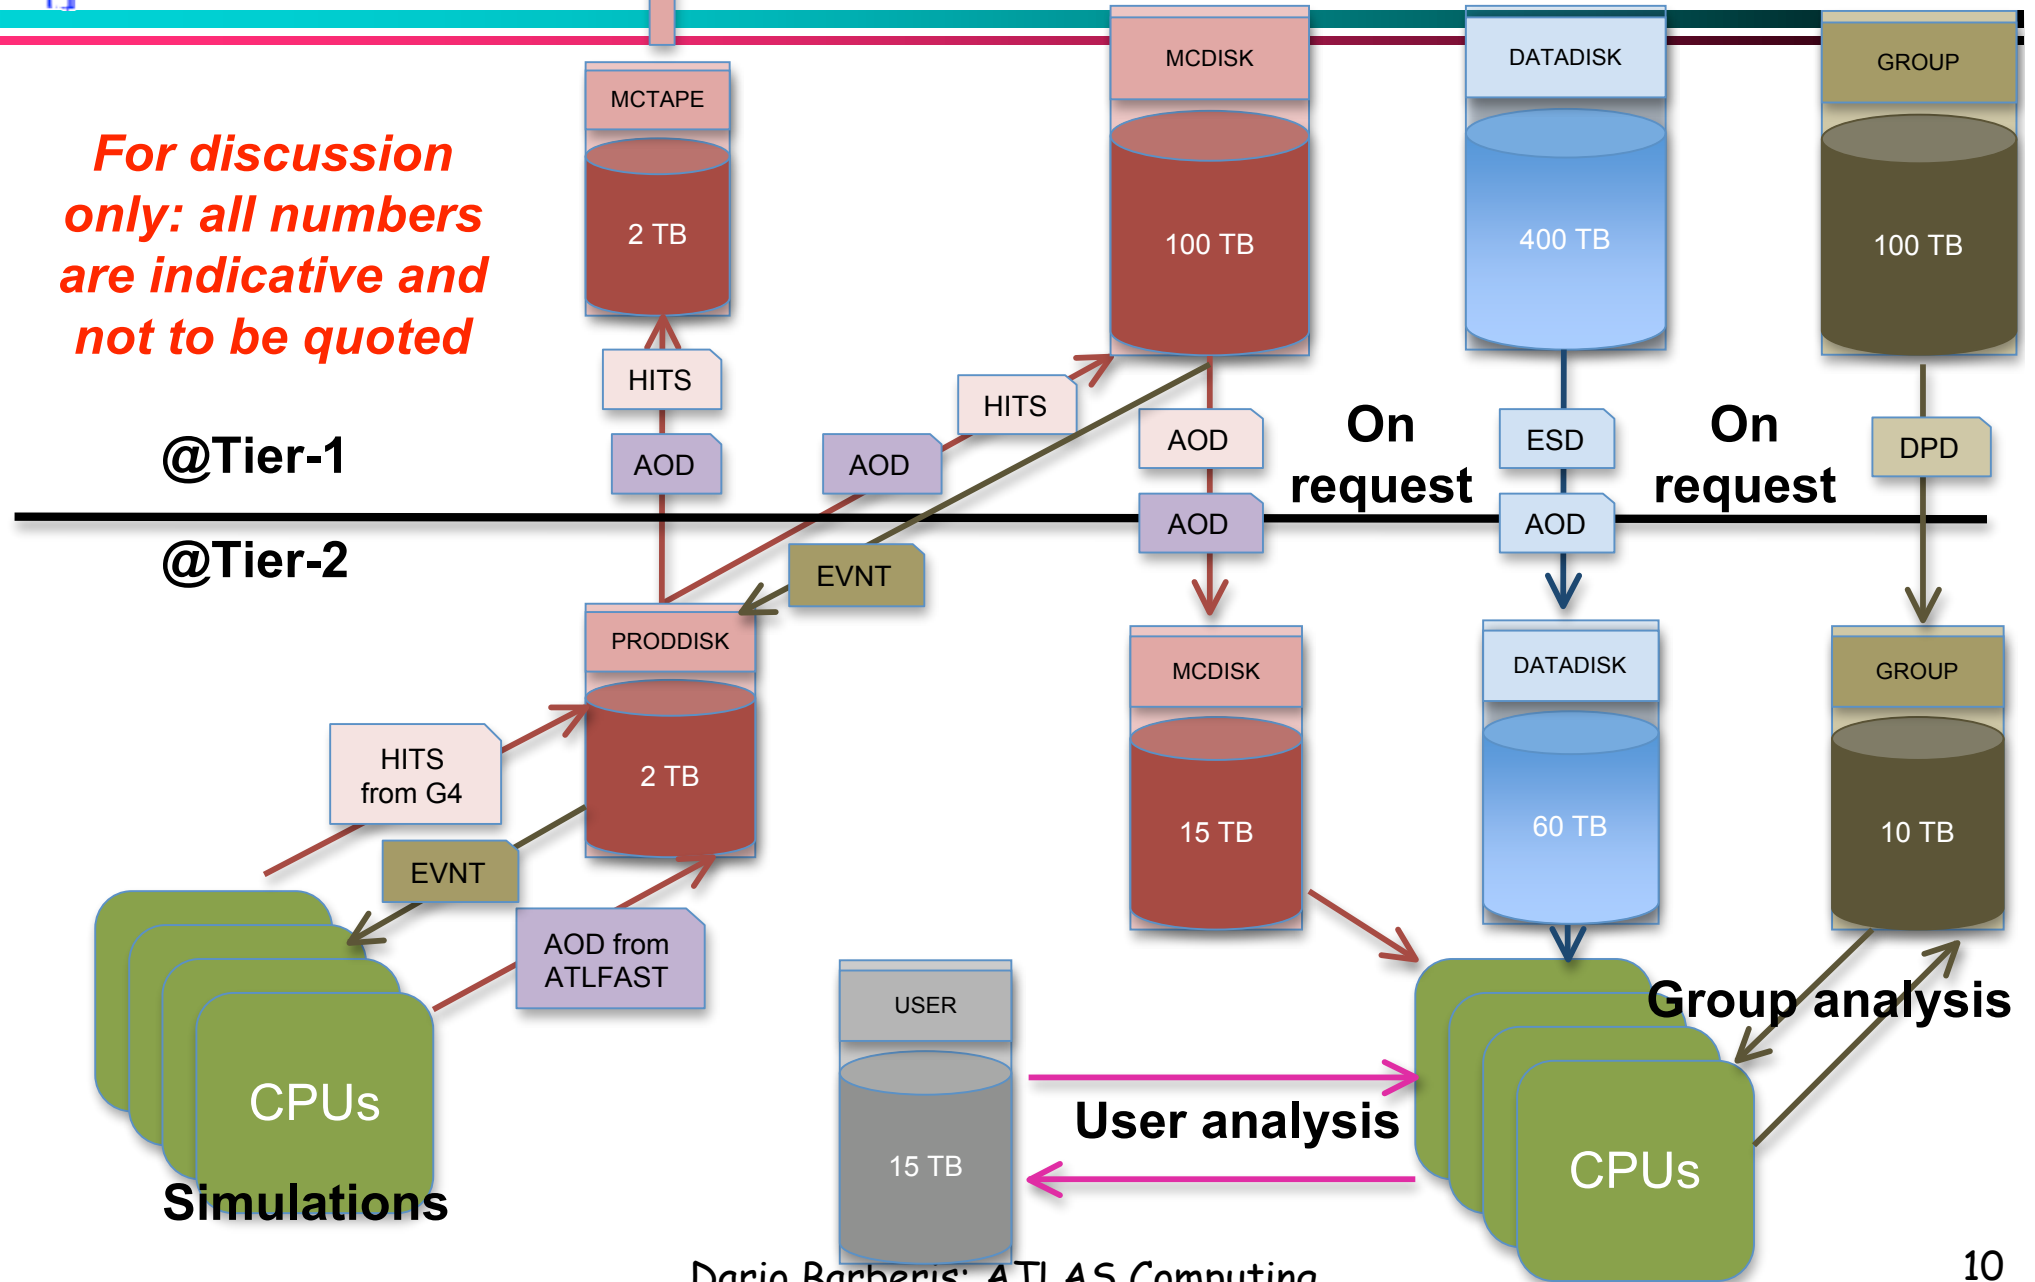
WLCG MB - 10 June 2008





# Storage for Tier-2's

*For discussion only: all numbers are indicative and not to be quoted*



@Tier-1

@Tier-2

On request

On request

Group analysis

User analysis

Simulations

CPUs

CPUs



# Evolution of the storage model

- We tried to implement a reasonable set of storage classes (and storage tokens) for CCRC'08-2 and FDR-2, using only SRM v2.2 end points
  - Not all our original model can be implemented, as the T1D1 class is not working as in the SRM 2.2 specs, neither in dCache nor in Castor
- We are therefore moving towards a more conservative approach:
  - Treat T1D0 and T0D1 as completely separate storage instances
  - Copy data twice in case they have to go to disk and tape at destination, or trigger an internal copy from disk to tape
    - This is the case of samples of RAW and ESD produced at Tier-0, and all simulation output from Tier-2s
- In this way we lose the T1D1 functionality but retain control of which datasets are on disk
  - More control comes with more active data management on our side
- Conclusion: the definition of ATLAS storage classes for 2008 is still in progress



# Plans

- Software releases:
  - 14.X.Y releases
    - Bug fixes only for HLT/Tier-0 and Grid operations
  - Release 15.0.0
    - Late autumn 2008. Include feedback from 2008 running. Base release for 2009 operations.
- Cosmic runs:
  - M7-M8
    - Continuing in June-July 2008
  - Continuous mode
    - In parallel with initial LHC operations
- Collision data:
  - Ready to go from mid-July for what concerns ATLAS S&C