

ATLAS Tier3 workshop at the OSG all-hand meeting



ATLAS
EXPERIMENT

2009-12-06, 10:03 CET
Run 141749, Event 405315

FNAL
8-11 March 2010

ATLAS Tier3 workshop

Massimo Lamanna (CERN IT/ES)

What is a Tier3?

- Working definition
 - “Non pledged resources”
 - “Analysis facilities” at your University/Institute/...
- Tier3 level
 - The name suggests that it is another layer continuing the hierarchy after Tier0, Tier1s, Tier2s...
 - Probably truly misleading...
 - Qualitative difference here:
 - **Final analysis vs simulation and reconstruction**
 - **Local control vs ATLAS central control**
 - **Operation load more on local resources (i.e. people) than on the central team (i.e. other people)**

What is a Tier3?

- Comments:
 - No concept of size (small Tier3 vs big Tier2...)
 - Tier3s can serve (and be controlled by) a subset of the ATLAS collaboration (local or regional users).
 - At this stage there are no constraints of quality of service offered to the collaboration (or services at all, since a Tier3 could be used only by the physicists at the corresponding university)
- Non-pledged resources does not mean uncontrolled or incoherent
 - Need to provide a **coherent** model (across ATLAS)
 - Small set of template to be followed while setting up a Tier3 for ATLAS users.

Tier3: interesting features

- Key characteristics (issues, interesting problems)
 - Operations
 - Must be simple (for the local team)
 - Must not affect the rest of the system (hence central operations)
 - Data management
 - Again simplicity
 - Different access pattern (analysis)
 - I/O bound, iterative/interactive
 - More ROOT-based analysis (PROOF?)
 - Truly local usage
 - “Performances”

Tier3

- Of course the recipe *Tier3 = (small) Tier2* could make sense in several cases
- But in several other cases:
 - Too heavy for small new sites
 - “Human cost”
 - The new model is appealing for small Tier2-like centre as well
- In all cases:
 - Got the first collisions! The focus is more and more on doing the analysis than supporting computing facilities ;)

What is the difference between a Tier2 and a Tier3?

- Independently from the magic recipes for building and operating a successful ATLAS Tier3, we should consider the following directions:
 - **Maintaining grid services vs using grid clients**
 - **Tier3 as an independent layer** (with respect of the T0/T1/T2 infrastructure)
- With the goal to guarantee:
 - Simpler and smoother (central) operations
 - Development point of view: Tier3s can be seen as an incubator of new solutions for the whole system

Important points

- Need of an [ATLAS Tier3 model](#)
 - Build on top of the US ATLAS Tier3 initiative and experience
 - Copy the good ideas!
 - Make it ATLAS-wide!
 - Status:
 - Working groups
 - Prototype sites

ATLAS Tier3 workshop

- 25-26 January 2010 at CERN
 - <http://indico.cern.ch/conferenceDisplay.py?confId=7705>
 - More than 70 participants from ATLAS
- 13 presentations only on plans/experience
 - more than 1 per cloud, here the granularity is more "country"
 - Typically T3 is a single experiment facility
 - Notable exceptions: **DESY** and **Lyon** analysis facilities (NAF and LAF)
- All main aspects of the problem discussed
 - Data management (file, sw, cond files, dbs), job distribution (DorDA, gcp, others), BROCF

→ 6 working groups

- **Distributed storage (Lustre/Xrootd/GPFS)**

- Rob Gardner (Chicago) and Santiago Gonzalez Hoz (Valencia)

- **DDM-Tier3 link**

- Simone Campana (CERN)

- **Tier 3 Support**


- Dan van der Ster (CERN)

- **PROOF Working Group**

- Wolfgang Ehrenfeld (DESY) and Neng Yu (Wiscosin)

- **Software / Conditions data Working Group**

- Alessandro de Salvo (INFN Roma) and Asoka da Silva



Status report
In this workshop



Please note!

Project metrics

- We would like to demonstrate one (a few) viable models for ATLAS Tier3
- The criteria are:
 - 1) **Performance** (analysis work)
 - 2) **Maintainability** (small local system management effort)
 - 3) **Sustainability** (small and understood implication on central operations)
- N.B. Importance: $3 > 2 > 1$
- Testbed!
 - Set up of an ATLAS Tier3 community which can

How to attach a number to it?

- **Performance**

- Analysis jobs

- Reliability $[(\text{tot-err})/\text{err}]$
 - Efficiency [CPU/Elapsed]
 - HammerCloud can do the job (at least for PanDA – need a local PF? Is a cron enough?)

- **Maintainability** (local effort)

- Measure the effort ???

- Stability over several weeks (HC)
 - Feedback (questionnaire)

- **Sustainability** (global effort)

- Experience of the central team

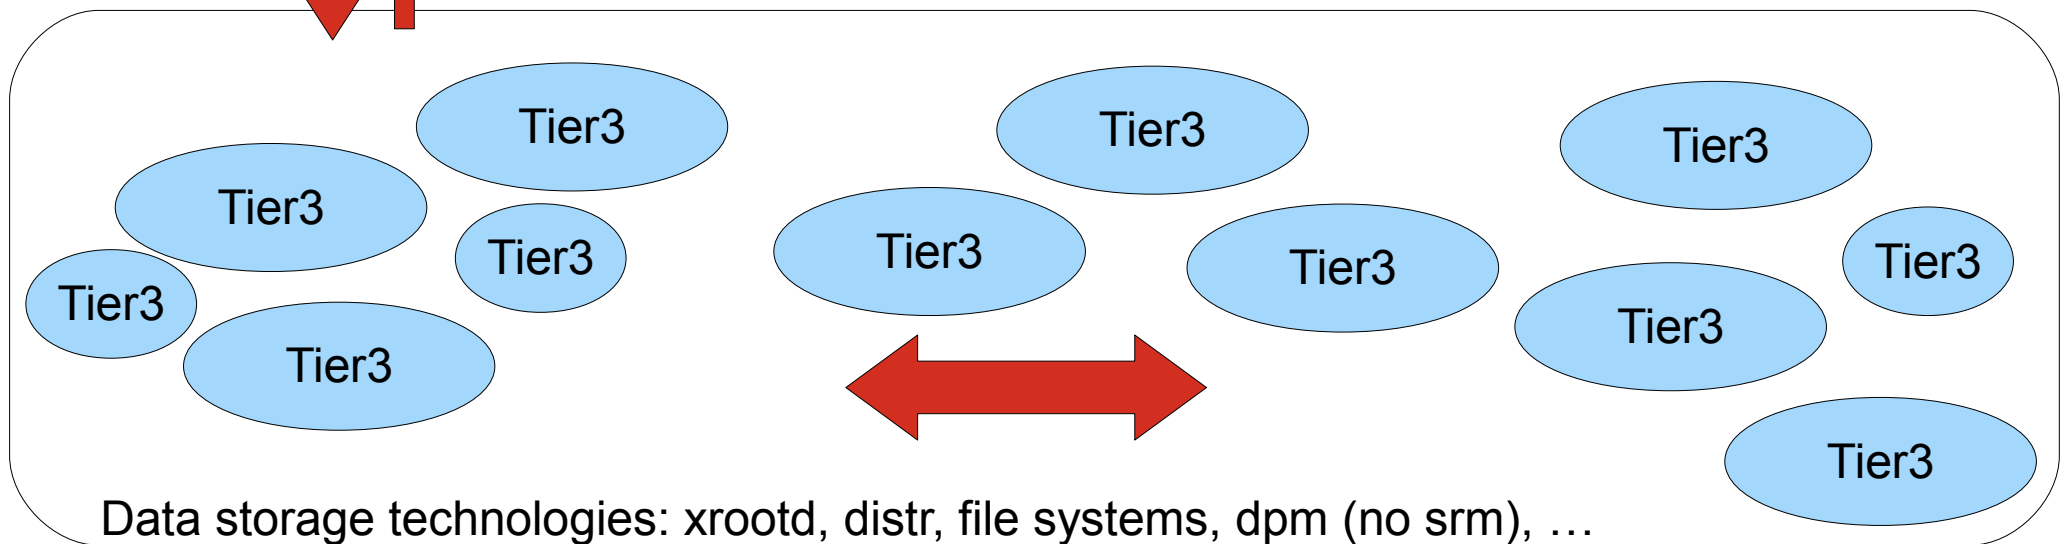
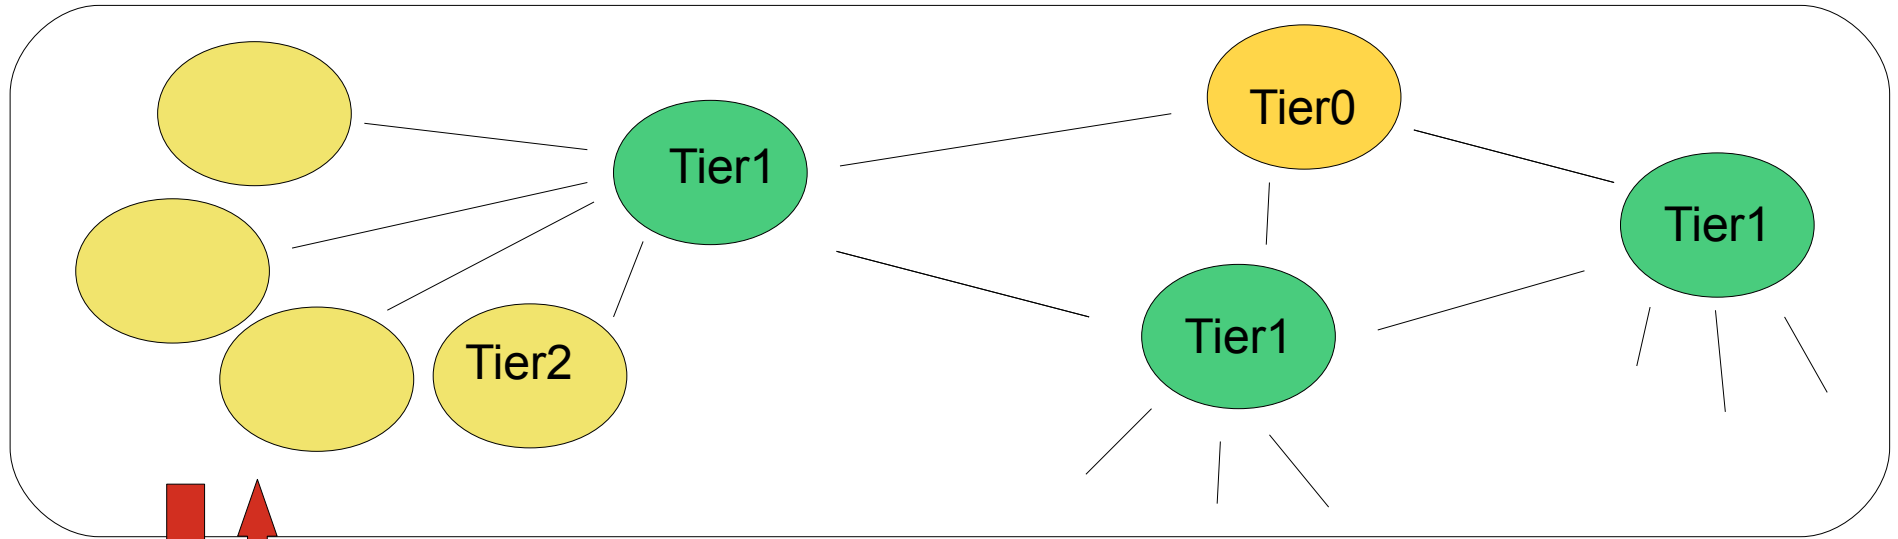
Personal notes

- I am in favour to start with the most simplified model:
 - Each Tier3 download the data they want and they have full responsibility (consistency, deletion, etc...)
 - The catalogue is the Tier3 “file system”
 - At this point this is not “advertised”
 - Compatible with PanDA (recent development) and local submission (includes Ganga with the local back-end, say LSF)
 - Most of the central load will be DAST oriented (User Support)
 - Assuming the Tier3 community is “closer” to the final users one than to the grid folks

Boring?

- Tier3-Tier3 data exchange
 - Possible across xrootd implementations
 - Desirable! Useful? Does it work for us?!?
- PROOF analysis
 - At least for ntuple scan... Unclear situation (investment) for Athena analysis (at least unclear to me)
 - Desirable! Useful? Does it work for us?!?
- For these things (and others):
 - Indispensable to have ATLAS first-hand experience!

System boundaries



Next stop?

- ATLAS SW week: CERN (April 19-23)

