



ATLAS

L. Poggioli, LAL

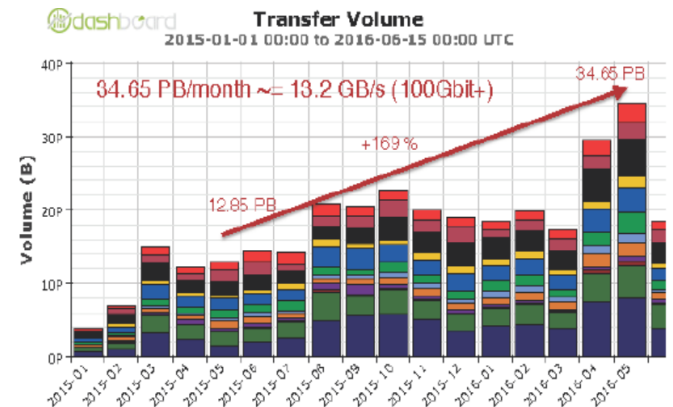


- Status/Concerns
- 2 points for discussion
 - Evolution of the Computing Model
 - Opportunistic resources (Cloud, HPC)

Status

- April 2016 RRB
 - Good usage of CPU
 - HLT farm efficiency improved, baseline 2GB/core
 - Heavy usage of beyond pledges (75% for T2s)
 - T2s & T1s explicitly acknowledged in ATLAS papers
 - Good usage of disk
 - Dataset lifetime, T1/T2 hierarchy flattening
 - Pledges 2017
 - OK (small reduction in T1 CPU under discussion)
 - Effort on CPU wrt Disk
- Run 3
 - Looks OK (AthenaMT for reco & simu)

Concerns: Run 4 (HL-LHC)

- Inputs based on
 - Phase 2 LOI layout, Today's Model, 10kHz HLT



- 1st output
 - Reco dominates (machine learning/coprocessors)
 - AOD storage only (caching, generate on the fly)
- Outcome
 - Run 4 means re-engineering, not evolution only
 - Network is essential

Evolution

- 'Smooth' evolution of the Computing Model
 - Gain in networking -> Dissolve Tiers hierarchy
 - T1s
 - Big data center with storage, throuput, less T1s (?)
 - Distributed storage à la Cloud
 - T2s
 - Less storage, use caching
 - 'Cloudification': eg jobs sent to FR cloud w/o explicit site requirements
 - Less T2s and bigger
- Hunt for opportunistic resources
 - Cloud, HPC, HLT farm, ATLAS@home

T2 storage consolidation (1)

- Discussed at RRB Spring 2016

ATLAS Report to CRSG – April 2016

0 1 5 9 13 17 21 25 29 33 37 41 45 49 53 57 61 65 69 73 0%

The fragmentation of disk end-points at Tier 2s results in a suboptimal use of the storage resources available. Inline with the long term direction of WLCG in preparation for Run-3 and beyond, we will move in the next years toward a model of consolidated or federated disk storage resources at a subset of sites while leveraging CPU cycles at all ATLAS tiers and opportunistic centres. We therefore encourage a reduction of storage capacity at those Tier 2 sites that provide today less than 400 TB of disk space for ATLAS, while we encourage them to increase their CPU capacity.

CRSG Report – April 2016

The wish to consolidate data resulted in a request that T2-sites below 400TB of disk invest in CPU rather than disk. We have expressed concerns about this strategy as this might not fit with the funding schemes of the T2 sites. ATLAS note that the scheme will be implemented gradually and in concert with the funding agencies.

- Concerns 22/71 T2s
- French sites a priori not concerned
- But the trend is there

- Iteration between ATLAS (ICB) and FAs
 - Some FAs or already consolidating their storage
 - Manpower or budget issues (UK, PT)
 - Sites specialization (ie production only) (US SW)
 - Some FAs increasing storage capacity (RO, IL)
 - Possibility to concentrate in 1 site soon (RO)

T2 storage consolidation (2)

- Understand role of small sites
 - Which workflows, DDM endpoints
- Possible evolution of the Model
 - Short (site w/o primary data) vs Long (pure diskless site) term?
 - Storage federation of big sites
- Suggestions
 - 400TB limit is arbitrary (but will increase)
 - Isolated sites -> low performance wrt sites in crowded federations
 - Well managed federations
 - Will decrease central operation load
 - Will maximize storage use for small T2 with disk

Cloud

- Need for extra opportunistic resource
 - Today in ATLAS, Cloud (commercial & academic) only few % total CPU
 - Amazon EC2 large scale test planned
- In Europe
 - HNSciCloud (see PEM's talk)
- In France
 - CC-IN2P3 (Manoulis) OpenStack, Cloud Scheduler, ~400 VMs
 - LPSC (Sha Li et al) Openstack, no manpower

Can we increase the French contribution?

HPC

- Mostly big HPCs centers (US, China*, EU)
 - Today in ATLAS, only few % total CPU
- @T2 level (Cf. mesocentres)
 - LPSC: CIMENT 7k cores 130 Tflops (Chris.' talk)
 - Used in ATLAS for evgen ~100 cores
 - LPNHE (Fred et al): 200 Tflops, 16k cores
 - Used for training and MC calculations

Can we increase contribution to ATLAS?

- @T1 level
 - Modest Top500 > 40

IDRIS & TGCC to be evaluated