
ATLAS Concerns for the future

— Andrej, Torre, Ale, Eric, Simone —

Manpower (Brainpower!)

- Expertise coming from T1s and T2s!
 - a large overlap between site operations and experiment support
 - we cannot afford to lose it
- Guarantee the stability of resources in spite of shrinking funding for manpower at sites - how to address?
- Clouds, today at 15%, more in the future
 - even grid clusters are difficult to maintain
 - how well will the clouds be supported and how much the experiments need to micromanage?
- Career possibilities are limited for experts in computing and software
 - especially for the physicists
 - difficult to motivate people to contribute on long and more permanent basis

Sites deployment and evolution

- Till recently : uniform service provisioning for Tier-2
- But even now: what should new sites deploy? ...not very clear!
 - which CE? how to configure it?
 - DPM? dCache? ... what makes sense long term?
 - new technologies: ObjectStores? http federations?
 - deployment models: OSG, EGI ...
- Possible future:
 - compute intensive sites, cached storage -- lower operational cost
 - data intensive sites: providing permanent storage with good connectivity - higher operational cost
- How to adjust the deployment model?

Storage persistency

- Compute resources are “easy”:
 - volatile: if a site goes down, there is no permanent damage
 - migration to new technologies (OS, services) is “transparent”/short term operational activity
- Storage resources are much more difficult:
 - site instabilities cause partial unavailability of permanent data
 - deployment of new storage technology is a long term operational activity
- Decommissioning of storage resources is expensive:
 - long migration of unique data - months, years
 - potential loss of unique data -- happened to ATLAS on few sites in the past
- How to ensure a stable storage?
 - MoU - required?, enough ?
 - long term commitment? level of support? planned funding?

Upgrade studies considerations for Run-3/4

- Memory consumption:
 - 4 to 8GB of RSS
- Reconstruction time:
 - x15 as compared to MC15 at $\mu=25$
- Need for dedicated resources:
 - most of the sites are not able to run them out of the box
 - often competing with other heavy requests (HLT reprocessing, Heavy-Ion reprocessing)
 - we need to find a way to run the upgrade jobs on part of ATLAS resources without disruption of regular activities

Job resource requirements

- RSS (PSS) instead of VMEM implemented in WMS
 - some sites not ready yet (LRMS with cgroups needed)
- Various activities:
 - high-memory vs usual 2GB/core jobs
 - high-I/O vs low I/O jobs
 - multi-core (and MPI) jobs vs single-core jobs
 - EventService (single event processing) vs fixed no-of-events jobs
- Those workloads can vary a lot (campaigns)
 - sites with static partitioning not effective - eg. we cannot use all resources for mcore (not more than 60%)
 - should the sites be pushed to be more dynamic, or should general-purpose vs limited-functionality sites be introduced?

Networking

- Need for efficient network performance and connectivity monitoring
 - ESNet provides central and full monitoring overview - in discussion to extend it to Geant
- Packet loss is critical for high RTT
 - needs to be propagated to the network operators
 - Crucial in view of concentration of the storage and remote access
- How to ensure sufficient network quality?
 - measurement
 - procedures and propagation to NRENs
- Some countries have much better networking infrastructure than the others
 - are the less performant networks sufficient for the new paradigm?

Common Software

- Common software is, as ever, essential and difficult
- Open source community points the way: practical, bottom-up, community driven
- ATLAS supports and participates in the community initiative following this model, the HEP Software Foundation
 - HSF startup team leadership & membership
 - Contributed [HEP Software & Computing Knowledge Base](#), [websites & newsletter](#), EC2 platform for HSF websites, ...
 - Investigating common project possibilities
 - Starting to leverage expertise via HSF, notably on git migration
- But HSF participation is (persistently) very low and consequently progress on building the HSF into what the optimists imagined 18 months ago is very slow
- How to address? How to build up the effort?
- Institutions capable of making/supporting contributions should do so
- Individuals should be encouraged and supported to participate
- cf. Benedikt's talk on Wed for more on HSF