

US ATLAS Facility R&D Notes

<https://indico.cern.ch/event/1201515/>

November 30 - December 2, 2022

Please feel free to add ideas to this document. Capture your ideas and thoughts here about what the facility should consider working on or doing more on. Could be just clarifications or extensions of existing activities or completely new areas we may want to explore.

Related item: Need to add items to the ATLAS Demonstrator in advance of Jan 23 S&C week (See

https://docs.google.com/spreadsheets/d/1Fx_R1WMM_J6OROfy8lqLIWpUejEkO25SHA_BYqXMjpE/edit#gid=0 and talk at <https://indico.cern.ch/event/1202607/>

Alternative Caching Networks

- Explore Varnish as a potential replacement for squid.
 - Much simpler to deploy
 - Works out of box
 - Does not need any disk storage
 - Cheaper to run
 - Works great with 2GB RAM vs 32GB for Squid
 - No disk storage
 - Can afford to run separate Varnishes for Frontier and CVMFS
 - Much simpler to monitor
 - No need for SNMP
 - Reports all the metrics correctly vs Squid that rounds things to 0
 - Varnish is a reverse proxy which is arguably what we actually need, not a forward proxy. Often random requests get proxied and reported as misses.

Improve container image delivery to facility

- Harbor proxy cache?
- cvmfs?

Preparing for high-speed analysis

- Do we know how to build a storage infrastructure that can support analysis rates of factors of 10-1000 faster than the current analysis tools allow?

Identifying infrastructure bottlenecks

- Can we run load tests or construct other example workloads that will accurately mimic planned infrastructure use-cases for HL-LHC?
- Do we have the monitoring capabilities to identify and characterize the bottlenecks that may exist?

Power efficient computing

- Establish arm processor cluster testbed. Using current and upcoming ATLAS workloads, measure the productivity(number of events simulated/events reconstructed etc)/watt. From a facility point of view, understanding if we can use lower power CPU hardware is important.

Heterogeneous storage systems for efficient analysis data access

- Use Object Stores (DAOS in particular) - RNTuple to measure performance using analysis benchmarks. See ["Exploring Object Stores for High-Energy Physics Data Storage"](#)
 - Collaborate with ATLAS Software data management team (Peter VG, Serhan Mete et al.) to write DAOD_PHYS/lite in RNTuple format.
 - use RNTUPLE formatted data on posix filesystem, dCache (xrootd interface) and OS to determine processing rate. (w/o corrections)

Developing tiered storage

- Gain insight about data on site: usage pattern analysis, data volume, breakdown by user/system metadata
- Map data category and usage to storage classes (e.g., frozen on tape, etc.)
- Data movement layer

Revisiting Data Lake Dynamics

- Hiro's study says only 1% of the site's storage is used in a day.. and provides guidance for I/O relevant to data lakes
- Q - could this be a way to reduce spinning disk capacity costs?

- A data lake with VP queues and xcaches at diskless sites - studies of cache sizes, testing scale, infrastructure performance for various configurations (nodes, disks, disk types)

VP at scale

- Scale up VP queues across US facilities
- Perform sustained performance test over a week (TBD)
- Needed measurements?
 - network bandwidth before / after
 - CPU efficiencies
 - Job efficiencies
 - Job startup time
- How to coordinate?
- Need to write down a series of steps - implication for Run3 and potentially Run4.

Global cold storage across US Facilities

- Investigate an Object Storage service used by US Facilities for storing low usage, small files
-