ATLAS computing model notes

Snapshot as submitted white paper from S. Campana, A. Klimentov, E. Lancon, T. Wenaus, Jan 23 2017

The challenges of delivering HL-LHC scale computing in less than ten years under tightly constrained costs have been well outlined in a number of white papers. We briefly describe here some considerations and components in how ATLAS is addressing these challenges in its computing model development and planning, particularly the distributed computing aspects. The work described is a mixture of existing capability either deployed or in commissioning targeted at Run 2 and Run 3; new projects at early prototyping stages; and existing or planned R&D activities extending and broadening the existing elements with a longer Run 3 to Run 4 view.

We see some basic precepts:

1. *Storage* is the most constraining factor, as the component that is both most expensive and least amenable to leveraging non-owned resources (ie non-owned storage resources basically don't exist), making this the largest target for development and optimization.
2. *Networking* has been the fundamental enabler for LHC computing thus far and will remain so with a particularly tight coupling to 1). Move as much as possible from file based data access to provisioning exactly the data needed over the network.
3. *Processing* provides the greatest opportunity for leveraging non-owned resources. The dynamics and flexibility required to make maximal use of opportunistic resources are a principal driver for both processing and storage design.
4. *Intelligence* in the form of automated decision-making and dynamic optimization via system-wide, globally aware services leveraging comprehensive metadata and associated analytics is vital to all the above. It is also key to operational sustainability.
5. *Empowering analysis* as the ultimate deliverable should increasingly drive S&C deliverables. Once aggregate scale and capability requirements are met, more focus can go towards the 'last mile' to the physicist performing an analysis, developing tools and services that improve the impedance matching between the analyst's work and the computing she uses. With capability delivered, focus on quality and the user experience.

Our approaches follow a few principal themes and elements:

1. Removing partitions and static structures in the processing and data handling that are increasingly unnecessary and constraining.
    a. In the processing sites: Static hierarchical tiers are replaced by dynamic capability-driven data- and network-aware management of processing resources.
    b. In the storage: Federations dissolve the partitions between small sites into more efficient and economical aggregates providing simpler higher level views and interactions. Consolidating data access protocols also reduces fragmentation and reduces the number of dependencies in the stack. Reducing the number of data access protocols to optimise data access, diminish the number of dependencies in the stack, and reduce the operational costs
    c. In data and processing workflows: Event level processing dissolves file level granularity in processing and managing inputs and outputs, in favor of adaptable fine grained workflows
    d. In the event processing framework: athenaMT dissolves event level granularity allowing algorithms to independently and adiabatically introduce MT, with convergence between online and offline through native support for high level trigger use cases (e.g. event views)
2. Ever more extensive and comprehensive system data gathering supporting intelligent automation and optimization.
3. End to end integration along a number of axes
    a. Physicist to physics result: From the physicist initiator of either production or analysis, through the systems performing the work, to the delivered results and resulting publication, with accompanying data curation
    b. Processing task to consolidated delivered outputs: From a received processing request, through intelligent prioritized brokerage that is dynamically aware of available processing and data resources, to automated

dialogues with those processing resources to shape and provide appropriate processing slots, to managing retries and full (or adequate) completion, to consolidating, curating and providing access to the results.
4. Encouraging a different organisation of the facilities, with storage aggregated in a smaller number of larger centers.

We build on key components that we expect to be with us through Run 3 and probably longer:
- PanDA workload management serving production (via Prodsys2) and analysis workflows
- Rucio data management
- ROOT/xrootd based (local and) remote data access
- HTCondor based grid resource access
- Information system

Missing, incomplete, needed components:
- Standard, uniform means of accessing HPC resources
- Large scale operationally sustainable utilisation of cloud resources

Some of the projects and R&D directions embodying our approaches:
- Fine grained processing friendly to opportunistic resources: Event service
    - allocate work in fine grained ~10min chunks, and if outputs reside on volatile storage, upload with similar granularity (e.g. to object store)
        - resilient against the opportunistic slot vanishing
        - amortizes data movement across the processing time; no post-job peaking
    - decouple flows of produced data from where the processing is done
    - complete tasks across a range of processing resources, while consolidating the outputs
    - output uploads quasi real time and decoupled from processing for maximal efficiency and flexibility; data movement is concurrent with processing, not following on after
- Fine grained asynchronous input data flows: Event streaming service (ESS)
    - just as the event service quantizes outputs into ~10min chunks, the ESS will do the same for the inputs
    - pre-fetching the data corresponding to event ranges assigned by the event service
        - processing works with pre-fetched data so no WAN latency inefficiencies
    - Friendly to networks, friendly to applications consuming WAN resident data, loosens the coupling between processing and data locality
    - be able to better leverage processing resources without local data, if they can consume and deliver data remotely without workflow inefficiencies
    - lessens the dependence on effective caching: caching contributes value to the extent it has a useful cache hit rate, but isn't crucial to making WAN access work, because the access is async
        - caching does have guaranteed value for concurrent usage of same/adjacent data within a cluster, which has a good hit rate built-in
    - establish an extensible distributed data service that can naturally accommodate progressively more sophisticated CDN like functionality, and virtual data-on-demand
        - V1: purely client-driven pre-reads of needed events/branches
        - V1.5: use access stats to drive caching/replication
        - V2: add a data marshalling capability on the server side to deliver only the data needed by the client
        - V3: virtual data?
- Active, flexible, data-rich management of processing resources: Harvester
    - Common layer dynamically provisioning resources in tight collaboration with the central intelligence managing the processing of work in the queue with its requirements & priorities: panda + jedi + harvester
    - Common across resource types (HPC, grid, clouds, …) encapsulating the (substantial!) differences in interfacing resource provisioning to the workload management system via resource specific plugins
    - In terms of dynamic control, data gathering/monitoring, and uniformity across resource types, it integrates resource provisioning much more tightly and powerfully with workflow management than in the past

        - More/better automation and more/better uniformity, towards minimizing operational workloads as platforms proliferate
- Comprehensive, extensible information system: AGIS
    - ATLAS originated and continues to develop the AGIS information system that consolidates site information flexibly and extensibly
    - Uniform source and interface for information allowing characterizing and automating site configurations as the parameter count grows with each successive tool and service supporting site characteristics
    - now being extended beyond ATLAS and LHC
    - a generic json data structure representing information stored by AGIS or added/extended by other info systems and sources will keep information system clients independent of AGIS itself
        - anything that can populate the data structure can be the root info system (which may be composite across >1 system)
    - Further extensions in characterizing sites for workload management purposes, down to the WN level, are planned for Harvester (e.g. WN level network bandwidths)
- Managing and leveraging metadata: Data characterization and curation project and data knowledge base
    - A new group and project with comprehensive responsibility for event-related metadata, including new development of a ‘data knowledge base’ system for flexibly and extensibly hosting the metadata that will accrue in expanding system knowledge and intelligence
    - Compute, network, storage optimization all predicated on intelligent, informed usage: acquiring (meta)data and analyzing it and applying it
- Storage federations: R&D to federate T1/Tn centers within one national cloud using different technologies
    - Evaluate how to federate Russian Federation Tier-1s and Tns sites using dCache, EOS technologies
        - Optionally evaluate federation of dCache and EOS sites with a protosystem as a high level SW in the future
    - Place ALICE and ATLAS data and distribute them randomly within federation
    - Run synthetic tests and experiments specific applications
        - Event selection (ALICE) - I/O intensive
        - TRT reconstruction (ATLAS) - CPU intensive