Data Management

Alastair Dewhurst

Motivation

- The growth in storage requirements is only part of the story.
- We will continue to want to run over “all data” in a timely manner.
- Storage capacity is growing faster than performance.
  - New types of storage such as SSD & NVMe do have better performance.
- Managing storage is the most labour intensive part of the Grid.
- R&D in other software will lead to more data intensive applications.

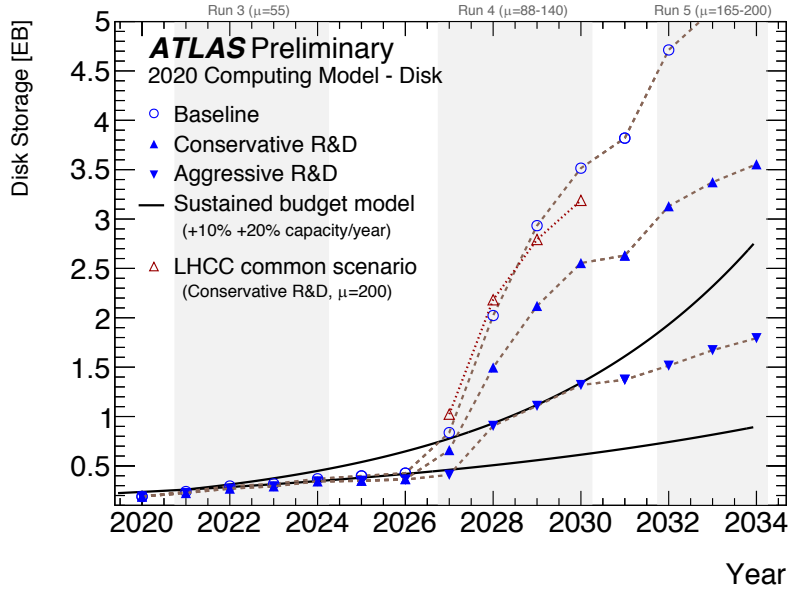

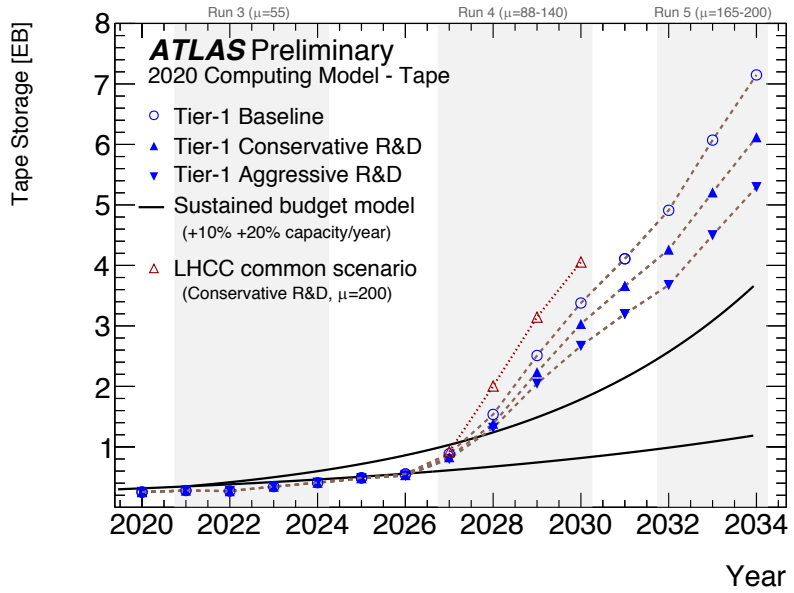

Swift-HEP Proposal

- 0.5FTE for 30 months
  - 12 Project Months at RAL
  - 3 Project Months at Brunel (+3 funded by Brunel)
- Tasks:
  - Build a Data Lake
  - Implement Quality of Service
  - Setup Disk-Less sites
  - Build an analysis facility

WP1: Data Management

D1.1: Setup UK data lake
D1.2: Implement QoS info
D1.3: Rec on data access
D1.4: Analysis Facility

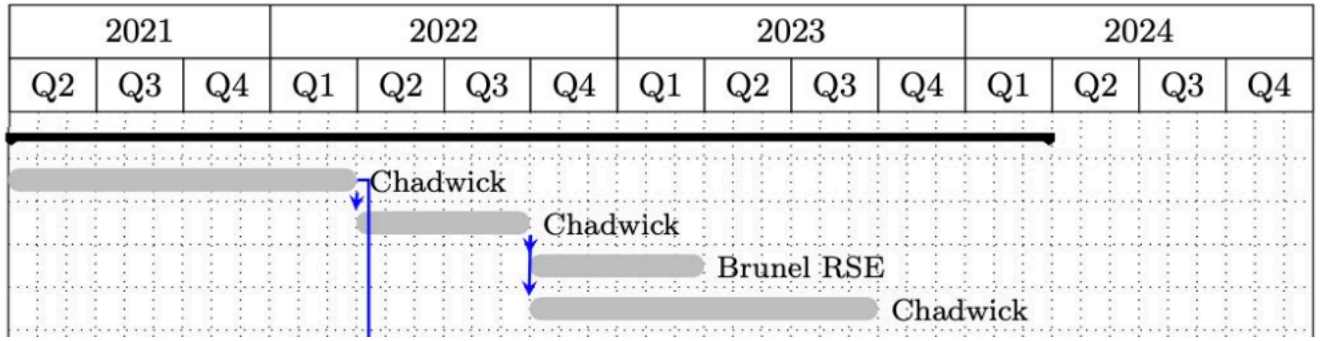

Building a data Lake

- The Data Lake concept moves a lot of the data management from the experiment to the infrastructure.
- will provide the framework for the other improvements
- The UK is likely to be part of large
- RAL Tier-1 already runs a Rucio and FTS instance.
  - Additional funding from IRIS and EGI-ACE to develop and support these services.

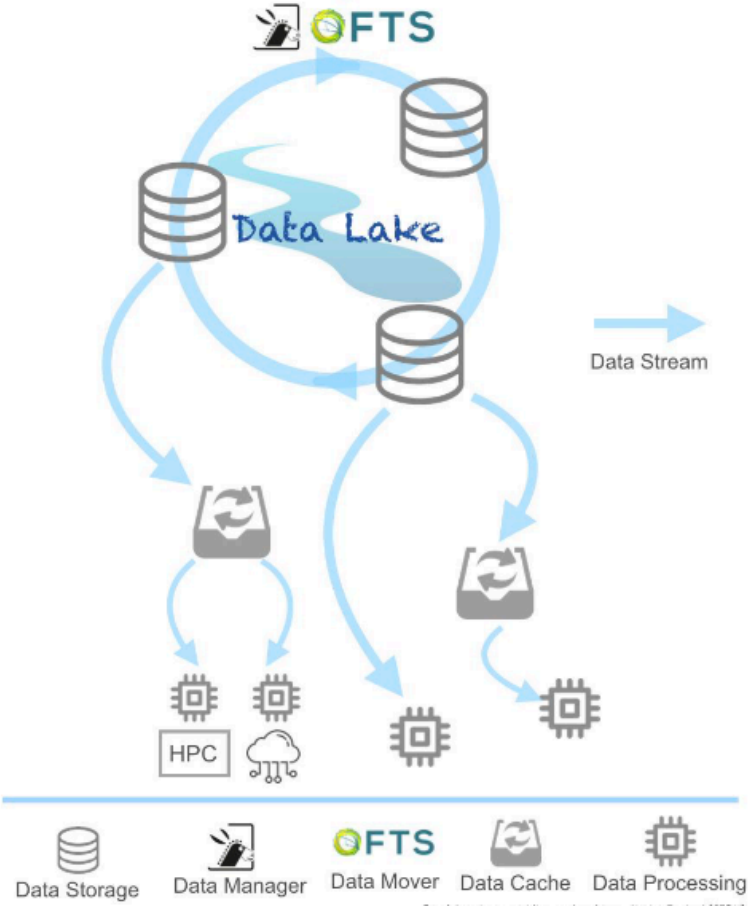

https://wlcg-docs.web.cern.ch/technical_documents/HL-LHC%20Review%20-%20April2020/HL-LHC%20April%202020%20review%20-%20DOMA%20-%20Main%20Document.pdf

Implementing Quality of Service

- We will need to make greater use of cheaper (i.e. less performant) storage.
- We will also need to make greater use of faster storage.
  - We observe better success and job efficiency for jobs using SSD backed storage.
- Sites will need to setup new storage endpoints.
- We will need to ensure we can move data between them in a timely manner.

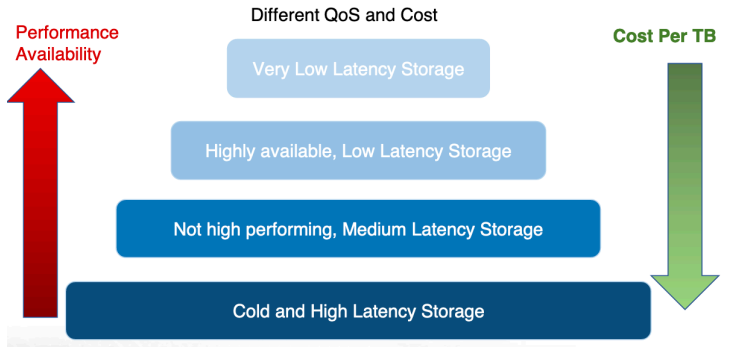

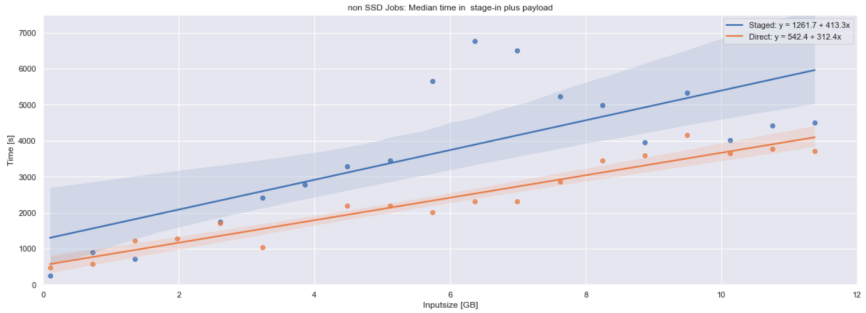

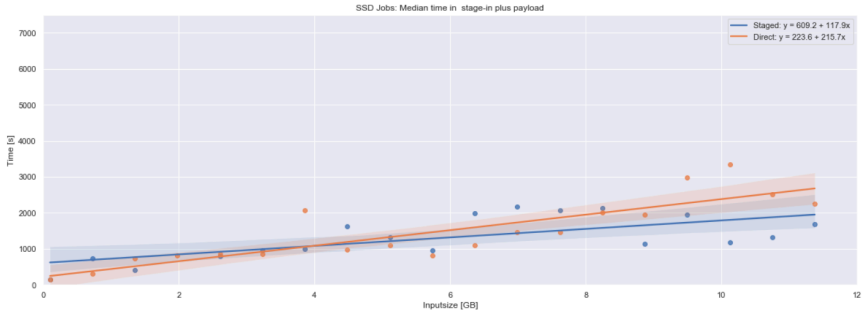

Implementing disk-less sites

- Managing Storage is labour intensive
  - Sites cannot afford a lot of resilience in their storage systems and if something breaks it needs to be fixed.
  - Upgrades and changes need to be done much more carefully because data needs to be kept.
  - For the experiment, each endpoint is additional effort to support.
- GridPP made the decision to consolidate the UK storage into 5 Tier-2s (and the Tier-1).
  - Other sites will still provide CPU.
- Significant work already going on within GridPP to figure out how best to access storage efficiently remotely:
  - Directly, via a transparent cache, via volatile storage.

Building an Analysis Facility

- The amount of data users need to analyze will grow.
  - Turn around times need to stay the same.
  - Error rates need to reduce.
- How we achieve this very much depends on the work in other packages and how users do their analysis.
  - Might be a significant growth in interactive work?
- In September 2020, RAL received funding to build a data analysis centre for LSST
  - 9PB S3 storage and GPU and CPU nodes on OpenStack Cloud.

Summary

- UK is in a strong position going into the start of Swift-HEP.
- We are aiming to provide both more resources and better performance with the same the same ongoing effort & hardware budget.
- There is a significant amount of complimentary work going on.

Science and Technology Facilities Council

Backup

Build a data lake

D1.1 Setup a UK data-lake prototype with four sites and compare metrics with the current system (Month 1-12)

The first task is to set up a UK data-lake. This will build on the DOMA prototype, with the intention that there will be one data lake per region/country. The UK data-lake will use the Rucio instance that is already available at RAL (and being developed with IRIS funds). In this first stage the data-lake would closely follow the generic WLCG plan[1]. This can be broken into three roughly equal steps.

- Configure the core sites that have permanent data storage. To ensure adequate test coverage we will use at least four sites, preferably using different storage technologies.
- Configure additional sites to access data. As well as the core sites, we will configure at least one site with state-less storage (i.e. a cache), one site with no local storage at all and one external site such as a Cloud Compute or an HPC to access the data-lake.
- Create a series of metrics which will measure: Overall time for [data analysis] tasks to complete, storage used, bandwidth used, CPU efficiency of tasks and operational burden (sites and experiment). This will allow comparison with the current infrastructure as well as being able to measure the improvements that come from subsequent tasks.

Implement Quality of Service

D1.2 Implement additional Quality of Service information in the data management (Month 13 - 18)

UK institutes run a variety of different storage technologies, each with their own strengths and weaknesses. In the past the only difference in Quality of Service was between Disk and Tape. We want to add in two major new features:

- High performance (i.e. SSD) storage. Some tasks (e.g. analysis) are becoming increasingly I/O bound and it makes sense to temporarily place the data on high performance storage. Not only does Rucio need to understand the concept of a high performance storage endpoint it also needs to be able to shunt the data into and out of it quickly enough to support the data intensive workflows.
- Reliability of storage (i.e. possibility of losing data). We aim to mitigate the risk of data loss by adding the concept of site reliability to Rucio. For example, as data ages the impact of file loss reduces so after 1 month moving data between new hardware and some that is 5 years old (with a higher chance of failure) could make sense.

Disk-less sites

D1.3 Produce recommendations to sites on how to optimise data access and deploy stateless storage (Months 19-24)

The UK benefits from excellent low latency network connectivity between institutions. This provides significant opportunities when it comes to data access. In many cases computing resources at one site will have excellent access to storage at others. For other sites a state-less storage service may significantly improve CPU efficiency with a minimal increase in operational effort.

Analysis Facility

D1.4 Creation of a [Virtual] Analysis Facility (Months 19-30)

As the amount of data grows, the more likely it is that something will go wrong when trying to process it all. To process datasets in their entirety users can spend a long time chasing up a small number of problematic jobs. The idea of an analysis facility is to have dedicated sites optimised to run user analysis jobs. Within the data-lake model, an analysis facility could actually be made up of multiple sites. This task will be heavily influenced by the earlier work being done in both the data management and workflow management parts of the work packages. This task will provide support to the work described in WP5 to ensure the analysis services developed have a solid infrastructure to rely on. The main goal of an analysis facility is to increase the reliability of the resources used, for example by only using sites with more reliable storage or pre-caching the data for faster access when jobs start. It will also provide tools that make things easier for users, such as easier access (via a web browser) to their job output data or relevant logs. The increased reliability will directly translate in faster analysis turn-around time and more efficient use of existing grid resources.

This work will be further exploited in WP5.1 and will foster a collaboration with the ServiceX team within IRIS-HEP in the USA.

UKRI Science and Technology Facilities Council

Questions?