

Data Intensive Research - Enabling and Optimising flexible 'big data' Workflows

Ian Corner

CSIRO, 1 Wilfrane Cres, Yarralumla, Australia, ian.corner@csiro.au

LENGTH OF EXTENDED ABSTRACT

The extended abstract should be no longer than 2 pages using this template (including text, tables, illustrations and appendices). A third page can be used for author bios.

DELIVERY

SHORT ABSTRACT [Maximum 300 words]

As data growth and proliferation continues to outpace research grade infrastructure, do we need a new approach to the problem? Are we 'big data' ready? Or do we only have 'lots of data?'

For 'big data' to have a future it needs to be:

- discoverable,
- related,
- well connected, and
- easily mapped to existing/future workflows.

Can we transition from having 'lots of data' into 'big data' while:

- reducing costs,
- improving data management practices,
- accelerating workflows, and
- opening up new workflow possibilities?

Data only speaks when it is analysed. Analysis expresses relationship through workflow. 'Big data' requires optimised workflow while preserving relationship. Research is often ad-hoc so workflows must remain flexible. What data management frameworks and optimised infrastructure do we require to accelerate workflows and maintain relationship while reducing overheads?

The CSIRO strategy is based on:

1. Increasing bandwidth of data to compute by building our disk arrays not for 'data storage' but for 'high speed data cache',
2. Transforming the use of peta-scale tape libraries from 'backup' (as an afterthought) to 'data storage with integrated protection',
3. Escaping the monolithic data problem by containerising datasets to achieve a flexible mechanism for connecting data to workflows, and
4. Allowing throughput optimisations by using pre-defined data categories to communicate access patterns and protection regimes to the infrastructure.

This talk looks at what can be achieved today, based on existing technologies and real eResearch workflows. It provides a scalable plan to survive and prosper, as datasets and workflows continue to grow.

ABOUT THE PROJECT [EXTENDED ABSTRACT]

CSIRO, the Commonwealth Scientific and Industrial Research Organisation, is Australia's national science agency and one of the largest and most diverse research agencies in the world.

STACC PROJECT - CAPABILITY ENABLEMENT

The **Storage and Compute Cloud (STACC)** Project was endorsed in October 2012, to develop and deliver internal research cloud services and supporting tools, facilities, technologies and frameworks. The project deliverables were focused towards providing increased science efficiency and productivity through access to more flexible and higher performing technology.

CAPABILITIES DEVELOPED AND DELIVERED DURING THE STACC PROJECT INCLUDE:

- Access to reusable, scalable and shared resources;
- Development of a range of storage and compute acceleration tools and systems, delivering accelerated research workflows;
- Provision of underpinning resources, facilities and frameworks;
- Delivery of self-service tools and automation development, including provision of cloud orchestration tools and technology; and
- Rapid and reliable business continuity response and protection of research data assets.

SERVICES

Support - the service delivery team works with researchers to understand their workflow, storage and compute requirements thereby ensuring data and workflows are structured in a way that optimises processing while reducing infrastructure foot print.

High Performance Storage Facility (HPSF) - The high performance storage facility enables users of data intensive workflows access to non-blocking data storage. The facility may be scaled up or down to meet the changing requirements of the scientist's workflow without significant re-investment or reworking of infrastructure. Parallel workflows are supported through either policy based scheduling of IO or categories into dedicated pools of underlying storage.

Long Term Storage Facility (LTSF) - As the storage capacity of the HPSF is limited in capacity the LTSF provides the 'general storage' solution, allowing current datasets i.e. those being processed in the immediate future to be cached on the HPSF. In addition, LTSF treats the issue of data protection by enabling datasets to be stored at two different geographic locations.

Optimised Storage Management (OSM) - High bandwidth dedicate storage appliance deployed on a dataset by dataset basis (between 1TB and 250TB in capacity). Appliance is dedicated to a project or small working group of researchers whose activities revolve around a single science dataset. Within the appliance, data is catagorised into subsets permitting the research to communicate the significance, performance or otherwise of the larger dataset at points within their workflow. The appliance integrates both HPSF and LTSF for optimal performance and minimal footprint.

Optimised Compute Appliances - (OCA) - Readily deployable compute appliances with integrated security and automatic connection to the researcher's storage appliance. Appliances are deployed either as permanent or on-demand and integrate automatically with the researcher's current datasets and codebases.

Optimised Compute Framework - (OCF) - The optimised compute framework links the researcher to the infrastructure in a manner which optimises science workflows. The framework ensures researchers can develop, share and peer review new codes in a distributed manner. The framework ensures that as more compute appliances are brought online existing code is readily accessible and executable in a scalable manner. As compute appliances are shut down and deleted the framework protects code from loss.

Clustered Inter-process Communication - Acceleration of research workflows generally requires the ability to run related tasks in parallel utilising low latency interconnects between compute appliances. This facility depends on advanced protocols such as RDMA over IB.

Logical Collection Management - Management of large sets of 'large sets of related data' requires a containerised approach. Logical Collection Manager provides an automated approach to ensuring related data in the form of 1,000s to 100,000,000s of files may be handled as a single object. The service embeds protection against silent data loss and silent data corruption. To facilitate integration within the science workflow Logical Collection Management integrates self-provisioning of datasets from within the science workflow.

CLIENTS

STACC services and capabilities have been progressively delivered to 'early adopters' across CSIRO research areas - with approximately 1.5PB of data storage with closely coupled compute services, delivered to 85 groups/projects across CSIRO. The 'STACC' project is ending in June 2014, with the services moving to a fully operational model under the IM&T Scientific Computing group.

ABOUT THE AUTHOR

Ian Corner has worked in the IT industry since 1987, with a focus on business oriented outcomes in commercial, engineering, government and science organisations.