

## Enabling Data Analytics as a Service For Large Scale Facilities

K. Woods, R. Clegg, N. Cook, R. Millward, Tessella Ltd, Abingdon, UK  
F. Barnsley, C. Jones, STFC/RAL, Didcot, UK

### Context of our Challenge

Advanced scientific research facilities are generating increasing volumes of data, and need increasingly advanced computing infrastructure, products and services to **interpret** and **manage** the data obtained during research.

For many scientists making use of advanced high performance computing infrastructure is a challenge. Without appropriate **experience and training**, exploitation of such resources can be daunting.

### The Ada Lovelace Centre (ALC)

The ALC is an integrated, cross-disciplinary data intensive centre fostering **better exploitation of research** carried out at large scale UK facilities.

ALC and Tessella have developed three new products, **Piezo**, the **VMM** and **DMS**, designed to abstract away the complexity of advanced infrastructure to free scientists from the complex mechanics of managing datasets and environments, enabling them to focus on the **analysis and interpretation** of their data.

### Piezo

**Piezo** creates **on-demand single-user Spark clusters**. **Spark** jobs run as **Docker** containers to ensure consistency and repeatability between different execution platforms. **Kubernetes** is used to orchestrate creation, management and tear-down of the cluster.




Figure 1: High-level Piezo component architecture

A **Harbor** repository is used to manage Docker containers. Underlying data storage utilises **Ceph**, accessed via an **S3** interface, to decouple processing from the Spark implementation.

Scientists are able to monitor progress and performance through a dedicated **Prometheus** interface, with more detailed job-specific information available via the Spark UI.

**Piezo** is a powerful general mechanism built on Kubernetes. We have demonstrated its capabilities through Spark but other frameworks could be injected to achieve different scientific use cases.

### Virtual Machine Manager (VMM)

The **VMM** utility manages compute resource scaling in a cloud. Built on **OpenStack** and **libcloud** it provides an automated, cross-cloud mechanism for managing pools of virtual machines. Each pool contains a configurable number of identical virtual machines pre-provisioned with an analysis environment and the tools researchers require to analyse their experimental data be that general-purpose machines, machines with many cores, machines with large amounts of memory, GPUs, etc.




Figure 2: VMM conceptual model

To perform an analysis, researchers select the dataset and the type of virtual machine most appropriate to their needs. The **VMM** will **allocate a machine** of the selected type and **automatically mount** the dataset without further intervention from the researcher. Because the virtual machines are pre-provisioned they are quick to load and be ready for use.

### Data Movement System (DMS)

During data analysis data must be moved to where it is needed. As data volumes grow, transfer of data between archives and compute clusters becomes increasingly inefficient. **DMS** addresses this problem by transferring only data that is needed in an optimised, **'just-in-time'** manner.




Figure 3: Logical structure of the DMS

The **DMS** uses a **FUSE** client to provide a folder-based view into underlying datasets but, crucially, does not transfer any files until requested to do so. **DMS** supports 2-way file transfer- archived data can be transferred to a compute cluster and processed data can be saved back to the archive.

**DMS** also features **data transfer adapters** allowing facilities to configure data transfer patterns best suited to their own local circumstances. The adapter mechanism is open, allowing facilities to create custom adapters based on user demand.

We acknowledge funding from UK Research and Innovation | STFC (UK SBS IT 18160).