



RCSI

Leading the world
to better health



Thursday, 22 September 2022



RCSI

RCSI Research IT Service Managed Slurm Cluster

WILLIAM BROWN

Senior System Administrator (Research IT)
Information Technology

The Challenge

Why do we need a compute cluster?

?

HIGH COMPUTATIONAL DEMAND

Research staff are processing large datasets where the compute resources needed exceed what is possible with a desktop computer or a single server.

?

GRANT FUNDING

A single 3-year research grant cannot cover the cost of a compute cluster or a large storage system. It is not practical to try to pool portions of many grants to buy systems or renew them at regular intervals.

?

PERSONAL DATA

We process a lot of identifiable Personal Data in order to deliver precision medicine and in conjunction with clinical trials. This cannot be processed on most external services.

?

FAIR DATA

A single large storage system provides greater safety from data loss and security. Data is easier to find over long retention periods to meet FAIR data requirements from grant awarding bodies.

?

TIME IS PRECIOUS

A 3-year grant cycle does not allow time for a research scientist to gain the IT skills to optimise systems, and there is no point in every scientist becoming an IT expert. Leave IT to trained IT staff and release time for the scientist to do research.

RCSI Research IT Service

ICHEC Kay



Slurm Cluster



Microsoft Azure
Blob Storage




Network Storage



Off-Site Access



OneDrive



RCSI Network



ICHEC KAY

ICHEC – Irish Centre for High-End Computing:

A very large Slurm cluster, where jobs that require very high CPU counts can be run.

Up to 1,600 CPUs per job.

Storage space is limited

Only anonymised data can be processed.

Can be accessed from anywhere.



ONEDRIVE

Every staff member has 1TB storage

Accessible from anywhere

Good for collaboration

Cost included in Office licence

Not high performance



ISILON STORAGE

Highly scalable – currently 1PB

Enterprise class levels of resilience

Multi-Protocol so exact same data accessible from user laptop and Slurm cluster.

Integrated to Active Directory so identity is assured.

Access to data is controlled through Active Directory; easy to audit who has access to any data.

Physically secured.

Encrypted data transfer for sensitive data.



SLURM CLUSTER

Small Slurm cluster

Up to 104 CPUs per job

1.5TB RAM

Bigger than any PC

Software installed by IT on request

Disks encrypted.

Integrated to Active Directory so identity is assured.

Physically secured.

Encrypted data transfer for sensitive data.

CIS Hardened

What is Slurm?



Slurm is an open source cluster management and job scheduling system for Linux clusters.

Slurm is relatively easy to install and configure, with a large global community of users. It is used in the largest compute clusters in the world and in most universities.

The key features are:

- Allocates exclusive or non-exclusive use of resources (CPUs, memory, GPUs, software licences).
- Starts, monitors and if necessary kills tasks on the nodes (servers) allocated to jobs.
- Manages the queue of pending jobs, taking into account resources requested, priorities and allocation of resource to the user.
- Maximises use of costly resources.

A cluster consists of at least:

- Login node
- Control node
- Compute nodes

Predictable and Reliable for the User

The user requests the resources needed for their job, by adding directives in the job script file or on the command line.

The job runs as soon as those resources are available; and is guaranteed those resources.

One user's job is (mostly) not impacted by another user's jobs.

Jobs may run on a single node or across many nodes (though not all software supports this).

Massively parallel compute tasks can be run.

Workflows can be designed to run steps in sequence or parallel as appropriate.

Job steps can have different resource requirements.

User can submit a heap of jobs and go away to do something else.

Bioinformatics software routinely runs for days or weeks.

Successful Delivery

FAIR

FAIR DATA

We know where the data is, and it is safe and secure. At project end it can be moved to a suitable public repository or retained as required.



TRAINING

Staff use a system that is used in other universities aiding mobility both ways.



PREDICTABLE SOFTWARE

IT provide software in a standard way, the same on ICHEC Kay and our local cluster. Many versions can co-exist.



STANDARD PIPELINES

Workflows can be defined so the same process is repeatably applied to different datasets. Aids peer re-analysis.



RAPID PROVISION

Storage can be rapidly provided on confirmation of grant funds; much faster than buying in hardware.



THROUGHPUT

All the compute resources are available 7*24. Limited only by capacity.



DATA LIFECYCLE


All data has an identified owner both on site and when archived; archived data has a retention period defined.




PERSONAL DATA

Auditable security and security-by-design allow processing of Personal Data.

Future

- More integration to other resources:
 - Cloud
 - REDCap 
- Share best practice internally (& externally?)
- CentOS replacement
- Bioinformatics centre: Understand:
 - Each tool better and optimize workflows
 - The data
 - Hardware capability
- SFI Research Infrastructure Programmes (€)
- European Environment for Scientific Software Installations (EESSI)



Thursday, 22 September 2022



RCSI

Thank you

WILLIAM BROWN

Senior Systems Administrator (Research IT)
Information Technology

CONSILIO

1784

MANUQUE