The University of Manchester CIR Campus Ecosystem: Émerging Technology Support

What is the Ecosystem?

The Computationally-Intensive Research Ecosystem is a response to feedback asking for an integrated system of infrastructure designed to address all aspects of research groups' computational work and requirements. It comprises:

- traditional batch computational facilities;
- a facility for interactive computation, e.g., for development work;
- high-capacity, resilient storage;
- computational facilities and storage linked via a dedicated, secure, fast network for data-transfer;
- a virtual machine service for research groups;
- a cluster for sharing emerging tech. hardware.

Local Versus Centrally-Run Infrastructure

The RI Team administer IT infrastructure for Before 2010, many small "beowulf" HPC clusters existed computationally-intensive research (CIR) — many of the facilities mentioned here: The CSF, Redqueen, The on campus. Some were well-run by academics and postgrads; others were not. All took time to administer which was better spent on research; most had many iCSF/Incline and Zrek; also the Research Virtual Desktop wasted "spare" CPU cycles. Since then, most such Service, SSH gateway and SSHFS service. In addition, we are the business owner of the Research Data Storage beowulfs have been decommissioned and financial Service, Research Data Network and the Research Virtual contributions have been made to the CSF instead — "HPC" infrastructure has been centralised. Academics Machine Service. If you are interested in finding out more now have access to a shared, professionally-run, campus about any of these services, please email service, with all the benefits that brings — where all "spare" CPU cycles may be used by others. its-ri-team@manchester.ac.uk

Following the success of this strategy, academics are now encouraged to make use of, and contribute to - buy into - other centralised research infrastructure run by IT Services, including the emerging technology service introduced here.



Computational Components of the Ecosystem

The CSF (aka Danzek)

The Computational Shared Facility is the University flagship "HPC" cluster with approximately 5500 CPU cores. It is used for a wide variety of work including HPC, HTC, and jobs requiring up to 512 GB RAM and/or fast/parallel filesystem

Danzek contains 34 GPUs in 17 nodes — some of these nodes are connected *via* Infiniband.

The zCSF (aka Zrek)

Zrek is the ecosystem's main emerging technology focus see box The zCSF, aka Zrek.

The Interactive CSF — the iCSF (aka Incline)

The iCSF, is designed specifically for interactive and GUIbased computationally-intensive work. It is ideal for code development work involving multiple very short runs. Incline is commonly used with the Research Virtual Desktop Service (which uses NX/X2GO for X11 compression).

Incline contains three GPUs which are used for visualisation work.

Other Computational Components

The remaining computational components do not directly support emerging technology, but for completeness we mention:

- **Hydra** for very high memory and big data work;
- Redqueen for non-pooled nodes;
- Condor for high-throughput of short jobs.

The Research Infrastructure Team

Alternatively, please visit our Web site

http://ri.itservices.manchester.ac.uk/

Michael Smith (FLS) Kadmon Job Queue

Storage and Data Transfer within the Ecosystem

Computational facilities within the ecosystem access common filesystems:

- each user has the same home-directory everywhere;
- large shared-areas for data storage are available to research groups free-at-the-point-of-use as part of the RDS (see below);
- a fast, dedicated network (the RDN, see below) links the RDS and the computational facilities.

The Research Data Storage Service (RDS)

IT Services provides centrally-hosted and administered data storage for research staff and students - the RDS. (Filesystems are replicated and snapshotted.)

The storage provided by this service is accessible from desktop and laptop machines on campus and may als accessed from on-campus research computing syster (including the CSF and the iCSF).

To find out more about the RDS, please visit

http://www.rds.itservices.manchester.a

The Research Data Network (RDN)

Many users of CIR have large quantities of data which be moved from experimental instrument to RDS and/c RDS to computational cluster. This requirement is sa by the RDN which connects all nodes on all facilities the ecosystem to the RDS using fast dedicated hardw a secure network.



The zCSF, aka Zrek

The zCSF — aka Zrek — is a loose-knit computational cluster the purpose of which is to bring together miscellaneous servers procured by research groups at The University which host emerging technology such as Intel Xeon Phi cards, FPGAs and the latest computational GPUs such as the Nvidia K40.

The cluster provides a managed environment for these nodes: moving system administration overhead from research groups to IT Services; making it easier to share and exploit the technology; and ensuring that users of each node can benefit from Zrek being part of the UoM CIR Ecosystem.

Zrek currently comprises the following nodes:

Merlion

Maxeler FPGA.

Xenomorph

Two Xeon Phi 7120P cards. Each card has 61 cores (244 threads) and 16 GB of

Besso

RAM

Nvidia K40 card (2880 cores and 12 GB RAM).

Kaiju

Two Nvidia K20 cards (each 2496 cores and 5 GB RAM).

Namazu

Altera FPGA.

Accessing and Contributing to Zrek

Should you wish to make use of the technology contained within Zrek, or contribute hardware, please contact us at

its-ri-team@manchester.ac.uk.

m so be ms	Ecosystem Gateways — Working from Office, Home and Barcelona
	Off-campus access to the Ecosystem is supported by a variety of globally-accessible services:
c.uk/	 SSH Gateway University staff and postgrads can login to the gateway from anywhere in the world and then hop to any component of the campus ecosystem.
n must or from tisfied within vare on	 SSHFS RDS storage is not accessible off-campus. The SSHFS service offers off-campus access to RDS shares used on ecosystem computational facilities.
	• Research Virtual Desktop Service This allows users to: access the CSF (Danzek), the iCSF (Incline) and Zrek from off-campus; do interactive/GUI-based work over relatively slow connections; and re-connect to the same desktop session from office, home and elsewhere.