MANCHESTER 1824
The University of Manchester

IT Services
Supporting your research
A newsletter from the IT Services for Research (ITSfR) group

Welcome to issue #16, July 2014, in a regular series of updates on IT research support and facilities for postgraduates, researchers and academics of The University of Manchester.

Your main Faculty contact email addresses (for research computing support) are given on the last page; emailing your research computing queries this way will also automatically raise a job for you in the *Remedy* tracking system.

Routine work should continue to be submitted through the IT Service Desk on +44 (0) 161 306 5544 or via online submission at: www.manchester.ac.uk/servicedesk.

Web Page: IT Services for Research

The web site for our support areas and all contact details is available here: www.its.manchester.ac.uk/research/.

Research Infrastructure Update

For further information about any of the items below, please contact the IT Services Research Infrastructure team at its-ri-team@manchester.ac.uk.

CSF milestone: over 6000 cores

The Computational Shared Facility has reached another milestone: When the hardware currently on order is added to the system, we will have about 6300 CPU cores in production.

Big Memory Cluster: Hydra Phase 2

Phase 2 of Hydra, the Ecosystem big memory cluster is now in production; Hydra now comprises 10 * 512 GB memory compute nodes each connected to a fast scratch filesystem – Lustre over Infiniband.

Hydra is built on the same shared-facility model as the CSF, with initial contributions from FLS and MHS.

RDS/Isilon and Campus Network Improvements

Usage of the Research Data Storage service is growing quickly, with total filestore in use currently standing at about 300 TB and growing at around 40 TB per month.

We have just taken delivery of 512 TB of new storage to add to the existing facility. This is expected to be in production in a few weeks.

To help research groups take advantage of Isilon for their primary storage eon office and lab-based machines, the campus network is being upgraded.

Several buildings now have a "top of building" connection of 10 Gb/s (i.e., a factor of 10 upgrade). More buildings will be upgraded in the coming year.

IT Services is now in a position to upgrade in-building networks so that speeds of 1 Gb/s will be available to individual labs and offices where a case can be made. It is hoped that the work required for this to be available will start in the coming months.

If you are interested in making greater use of Isilon but are concerned about available network bandwidth, please email your faculty support team.

The CIR Ecosystem: Academic Investment Reaches 2 Million Pounds

As of July 18, the total investment in the Computationally Intensive Research (CIR) Ecosystem by University research groups exceeded two million pounds!

The Ecosystem now provides:

- Batch computing facilities which meet high-throughput and high-performance (parallel) needs
- An interactive computational facility with up to 2 TB (2000 GB) of memory on a compute node
- Large amounts of resilient storage
- Users with the same home-directory on all computational facilities

For a poster giving an overview of the Ecosystem, please visit http://ri.itservices.manchester.ac.uk/ecosystem-poster.

The zCSF (aka Zrek): Emerging Technology

Zrek, the emerging technology cluster in the CIR Ecosystem, now contains six Nvidia K20 GPUs and one K40, following further contributions from researchers in Pharmacy, CEAS and Humanities. Researchers with GPU-optimized simulations who are interested in evaluation access to Zrek should contact us at the above

Next Image-based Modelling Club Meeting

Registration is still open until 8th August 2014.

The University of Manchester Image-based Modelling Club are organising a workshop focusing on "Young Researchers" at the Hartree Centre, Daresbury Laboratory (Warrington) on Tuesday 9th September. Coach transport will be provided to and from the North campus. Anyone is welcome to attend.

In the morning, we will hear about the >£100M investment in world class supercomputing and visualisation facilities at the Hartree Centre. These are available to Manchester researchers through collaborative projects and European research grants. We will also enjoy a visualisation show case featuring recent Manchester Image-based Modelling research in one of the Hartree Centre's immersive stereo visualisation suites.

Lunch will be provided during a poster session. In the afternoon, we will hold a hands-on workshop to discuss the issues and challenges of image-based modelling. This will inform future Image-based Modelling Club events and help progress the use of this technique at Manchester.

For further details and registration, please visit http://www.imc.manchester.ac.uk/events/ibm-workshop

Greater Manchester Data Dive

Mary McDerby and Richard Mervin (Research Applications and Collaboration, IT Services) were members of the team that won first prize at the GM Data Dive on 22nd July organised by Manchester Informatics. The team also included staff from Manchester City Council and Zeyuan Zheng, a PhD student from Computer Science.

Amongst the attendees were people from areas such as the NHS, Manchester City Council, UKDS, SMEs (e.g., DataCentred, Reality Mine etc), as well as University of Manchester staff and postgraduate students.

The data dive/hackathon was organised as part of a programme of events to support and the thriving Big Data community across Greater Manchester.

The team scooped first place by addressing the following problem:

Using NHS prescription data they pulled out the top ten medications prescribed in the Greater Manchester area. Concentrating on medication used to treat clinical depression they mapped the data onto a grid of Greater Manchester in order to assess the use of such medication in low income areas. Data around patient behaviour in the context of exercise and healthy eating was then mapped onto the visualization.

The problem was complex and required the use of data from many sources e.g., NHS, Office of National Statistics, Social, and Ordinance Survey. Significant amounts of data processing and interpretation was required, as well as rapid prototyping and experimentation, in order to visualize and analyse the data.

All the teams present found the day very interesting and Manchester Informatics will be following this up with a similar event at the University in the autumn.

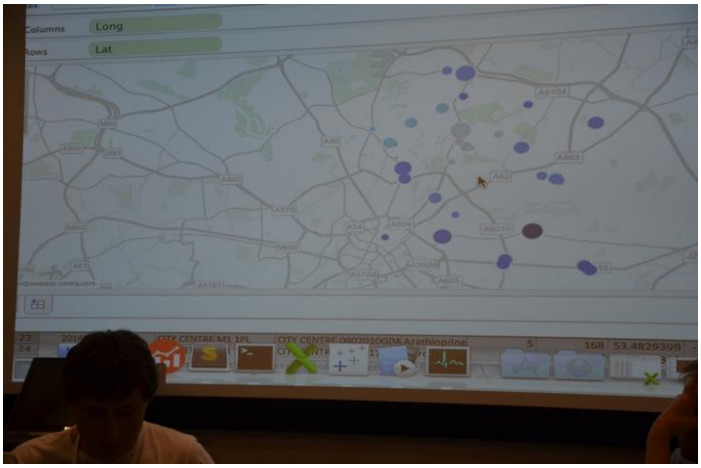

AGM for Research Software Engineers

The second AGM for Research Software Engineers (RSEs) will take place on5-16th September in London. Save the date!

This year, the meeting will be combined with a hackday, so that RSEsget an opportunity not only to meet, but to work together too.

More details about the workshop will be announced in the coming weeks. Keep an eye on the Software Sustainability Institute website, http://www.software.ac.uk/, for the latest news.

Beat the September rush and get your IT requests in now

As with many University departments, the Start of Year period is the busiest time of the year for the IT Service Desk. However, it is possible to take advantage of the (relative) quiet of the summer period by sending in requests for some things you might need once the students arrive. It's one less thing to worry about in September.

If you know you would like any of the following services for next academic year, fill in the relevant form as soon as you can:

- New shared network drive
- New shared email accounts (known as "Role accounts")
- Software installation
- Request redistributed IT equipment
- Setup of a new computer
- IT equipment disposal
- New network socket installation

If you can think of something else you might need setting up, send a message to the Service Desk to discuss this.

Condor pool usage exceeds 20 million hours

Our Condor high-throughput computing pool recently clocked up 20 million hours. Currently the clock says 20,010228, which is around 2,284 computational years. That's a lot of computational research results.

This has been achieved at low cost, as it mainly piggy-backs on teaching cluster PCs (there is a small always on backbone of around 150 cores contributed by EPS, The University and also one or two PIs).

We have rough figures from a couple of cloud computing providers and 20 million hours could easily have cost millions of pounds.

http://condor.eps.manchester.ac.uk

Advanced Optimisation Tutorial

An advanced programming optimization tutorial was recently published 'GEMM: From Pure C to SSE Optimized Micro Kernels'

https://apfel.mathematik.uni-ulm.de/~lehn/sghpc/gemm/index.html

This starts off with a very simple implementation of matrix-matrix multiplication in C and shows how to get maximum performance from the CPU using techniques such as loop unrolling, SSE instrinsics and assembler.

Good reading for anyone who'd like to know some details of advanced CPU optimisation.

Faculty Contact Emails

Humanities:
Hum-ITResearch@manchester.ac.uk.

Engineering and Physical Sciences:
EPS-ITResearch@manchester.ac.uk.

Medical and Human Sciences:
MHS-ITResearch@manchester.ac.uk.

Life Sciences:
FLS-ITResearch@manchester.ac.uk.

Next edition

The next edition of this newsletter will be circulated at the end of August 2014.

If you have any news to contribute, please contact Mike Croucher: michael.croucher@manchester.ac.uk before 25 August 2014.