

MANCHESTER
1824

The University of Manchester

IT Services

Supporting your research

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

Welcome to issue #18, September 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/.

New postgraduates

As is our custom each September, we repeat some information in this edition which new research students need to be aware of, such as the variety of computational systems and services on offer. If you are supervising new PhD students, please make sure they see a copy.

Research IT Forum Report

Thanks to all the researchers and IT staff who attended the Research IT Forum on Wednesday 24 September 2014. We would particularly like to thank Dr Caroline Jay – from the School of Computer Science – for her talk on how she would work with the proposed pool of research software engineers (coming on stream as part of the IT transformation programme from around 1 December 2014).

From IT Services, Dr Mike Croucher spoke about an imaginary researcher's journey from arrival at UoM through to her being awarded a Nobel Prize, with illustrations of how she might interact with all four sections of the transformed Research IT support service (within IT Services). Also from IT Services, Dr Simon Hood described the current Computational and Storage Ecosystem we have, with its strengths and weaknesses. And, in the latter case, what he proposes we do to fix them, if work is not already in progress (as it is, for example, with the improvements to the campus wired network).

Travel Fellowships for Supercomputing 2014

Last year, EPSRC funded 20 UK researchers to attend the US Supercomputing conference SC13 and network with US researchers. Competition for the travel fellowships was fierce and one was awarded to Dr Neil Ashton in the School of MACE. This enabled him to establish a new and exciting collaboration with NASA. The trips were funded through a special UK/US HPC Collaboration Network grant awarded to IT's Dr Lee Margetts and colleagues from Imperial College London, the University of Oxford, STFC and NAG Ltd. A second grant application is under review for SC14 in New Orleans (sc14.supercomputing.org/). We have responded to reviewers' comments and expect to hear if we are successful before the next newsletter is published. If you think you might like to apply for a travel fellowship this year, please let Lee know (lee.margetts@manchester.ac.uk). He will contact you if the grant is successful and provide details of how to apply.

Bids to the SSI for Help with Research Software

The Software Sustainability Institute (SSI) – part of which is here at UoM – is once again calling for bids from people who need help further developing their research software. Full details can be found on the SSI's extensive web site:

www.software.ac.uk/news/2014-09-24-need-free-help-your-research-software-try-institutes-open-call.

Some of you may know that the SSI's slogan is *Better Software, Better Research*. For a deeper understanding of what lies behind the headline see SSI co-investigator Prof Carole Goble's recent publication, reproduced on the SSI web site here:

www.software.ac.uk/resources/publications/better-software-better-research.

N8 High Performance Computing – do you need access to more compute power?

As well as access to the Computational Shared Facility (CSF – see below), UoM researchers also have access to a Tier 2 regional High Performance Computing facility, ideal for when the scale of your research outgrows other resources. N8 HPC is an EPSRC funded initiative that aims to seed engagement between industry and academia around research using e-infrastructure, as well as developing skills in the use of e-infrastructure across the N8 partnership of which Manchester is a member.

The infrastructure consists of a "Tier 2" facility, which is available for use on an equal share basis by researchers from all 8 partner sites. The facility is an SGI HPC cluster with 5312 cores. Each node has a compute capability of 320 GigaFLOPS.

Research groups can apply for free access to the facility at any time through a lightweight peer review process. Projects do not need to be funded by EPSRC and do not require an industry component to their research application although this is preferred. Application decisions are swift and assistance is available locally to make the most of the facility.

Please visit the N8 HPC website (www.n8hpc.org.uk) for more information, but if you would like to discuss your requirements or if you need assistance with your application then please get in touch (ITS-Research@manchester.ac.uk).

Research Computing Infrastructure

IT Services develops and administers the University Ecosystem for Computationally-Intensive Research that provides computational resources (both batch and

interactive) and data/file storage for researchers. Below we outline the main components of the Ecosystem. To find out more, please visit ri.itservices.manchester.ac.uk/ or contact the Research Infrastructure team at its-ri-team@manchester.ac.uk.

Computational Shared Facility (CSF)

The CSF is the University flagship high performance/throughput computing facility. It is managed by IT Services for the use of University research groups. It is built on a shared model: all compute nodes are procured using contributions of funds from research groups; the cost of infrastructure such as login nodes, file servers and network equipment is paid for by the University. University schools may also contribute so that all members of that school then have access to the facility. There is no barrier to entry: all contributing research groups from The University may use the CSF – there is no review process. The CSF currently has almost 6000 processor cores and it is tightly integrated with the Research Data Service (see below) forming a complete compute/storage research computing package ideal for those whose work involves large quantities of data. Support for users of the CSF is provided by IT Services with particular focus on helping those new to using HPC/HTC systems become productive; we are happy to provide face-to-face help for small groups and individuals. The CSF originated to provide a professionally managed alternative to individual research groups building and running their own small clusters. For this reason, i.e., to accommodate the (differing) needs of research groups, the service provided is very flexible and responsive: it is an agile service.

Free-at-the-point-of-use Computational Resource

The CSF (above) is contribution-based. We do, however, have a *very limited* amount of free-at-the-point-of-use computational resource to offer. For more information, please use the contact details at the top of this article.

Research Data Service (RDS) – Isilon

IT Services' Research Data Service provides centrally-hosted and administered data storage for research staff and students at The University. Storage provided by this service is accessible from desktop and laptop machines on campus and may also be accessed from on-campus research computing systems, in particular the CSF. Files stored on this service can be considered secure. For example, files corrupted or accidentally deleted can be recovered for up to 28 days. If you have research data storage requirements, you can find out more at www.rds.itservices.manchester.ac.uk/ or by contacting your faculty IT research support team, or the Research Infrastructure team (its-ri-team@manchester.ac.uk). To apply for storage, please use the form at remedy.manchester.ac.uk/cgi-bin/sr.cgi?scid=142.

iCSF (also known as *Incline*)

Incline – the INteractive Computational LINux Environment – is a service designed specifically for interactive and GUI-based computationally-intensive work, such as code development and testing, with applications such as Stata, R and MATLAB. Users of Incline have access to the same home directory (and other RDS/Isilon-based filesystems) as they have on the CSF. Incline is an ideal starting point for those new to remote use of shared computational facilities.

zCSF (also known as *Zrek*) and hCSF (also known as *Hydra*)

There are two specialist branches of the CSF: Zrek hosts emerging technologies and (GP)GPUs including Nvidia K40 and K20 cards, Intel Xeon Phi cards, and both Maxeller and Nallatech FPGAs. Hydra is designed specifically for those with very high memory and/or data input/output requirements.

Workshop: the new Azure Machine Learning Cloud Service

Microsoft Research invite you to participate in a one-day workshop about the new Azure Machine Learning cloud service that makes creating, testing, evaluating and publishing predictive analytics easy -

azure.microsoft.com/en-us/services/machine-learning/.

The event is free to attend and will be held at Microsoft Research Cambridge (UK) on 15 October 2014. It is a hands-on workshop aimed at academic researchers interested in using machine learning for their projects and teaching data science courses.

If interested, please register online at aka.ms/AzureMLCam.

Participants will be provided with access to Azure Machine Learning during the day, and for three months afterwards to continue to develop their work. The cloud service is web browser-based, supports R code, has R and Python APIs, and can be used from participant's own laptops running Linux, Mac and/or Microsoft Windows.

The aim of the workshop is to provide you with: an understanding of Azure Machine Learning and how it can be used; hands-on experience designing, building, training, evaluating and deploying predictive models; an opportunity to discuss your current and future needs with Microsoft and Microsoft Research.

Microsoft are also soliciting proposals for Azure Machine Learning Awards for data science courses, research project, and shared workspaces for research collaboration. The next deadline is 15 November 2014, and then every two months – apply at research.microsoft.com/en-us/projects/azure/ml.aspx.

Microsoft Research is now supporting over 300 projects worldwide as part of its Azure for Research programme (www.azure4research.com). You can also apply for one of our general bi-monthly awards, the next deadline for projects is 15 October 2014. Please submit your project proposals for up to 180,000 core hours of compute and 20TB of data over 12 months at www.microsoftazurepass.com/research.

Note that a general *Azure in Academia Workshop* is also likely to be held in Manchester in the not too distant future: more details when we have them.

SciProgramming Email List

This is a listserv email list intended for Manchester researchers who develop software. It is explained further here:

https://www.applications.itservices.manchester.ac.uk/show_content.php?id=180.

The web page also contains a link to the *Research Software: Made in Manchester* site; our thanks to Dr Mike Croucher in IT (EPS) for setting this up.

Windows 7 Managed Desktop

Answers to your feedback about the Windows 7 Managed Desktop, including:

- Administrative Rights
- Recent changes
- Future changes.

www.staffnet.manchester.ac.uk/news/display/?id=12745

Python Notes for MATLAB Users

Professor Bill Lionheart and student Sophia Coban recently introduced Python as an option for the School of Mathematics undergraduate module *Problem Solving by Computer*. As part of this introduction they have developed a set of lecture notes that teaches the basics of Python for MATLAB users. For more details please see www.walkingrandomly.com/?p=5568.

If you would like advice on how to port MATLAB code to Python, feel free to contact the EPS Applications team at applicationsupport-eps@manchester.ac.uk.

Condor

Condor is free-at-the-point-of-use to any researcher in any faculty. It provides High Throughput style computing (only – there are no fast interconnects or shared storage facilities, such as can be found on, e.g., the CSF or N8 HPC). As well as bespoke code, Condor can run any Linux application that has a batch mode.

Although we have some active professorial users, by far the main user community is PhD students. Graduate (PGR) students like the service as they can get more or less instant access without having to go through supervisor or application approval.

The number of simultaneous users of the system typically varies from around 8 to 20. The total number of processor cores varies during the day: circa 128 during normal office hours, 2500–3500 out of office hours¹.

Apart from a backbone of some 128 cores, we make use of existing PCs in (some of) the EPS Faculty teaching clusters (over night and at weekends, to avoid impacting undergraduate use – we also use Linux rather than Windows as the underlying system to get the full Condor features and further disambiguate it from the undergraduate teaching environment).

IT staff provide training courses and one-to-one support, as well as a code optimization service. We have added a GUI – called *DropAndCompute* – as well as an active web-based monitoring tool that is clickable down to a fine level of detail per job. We have also developed a tool that automatically checks all users' jobs have been adequately described in their submission files. Full details of the Condor system can be found here: condor.eps.manchester.ac.uk.

Subject-specific Email Lists

Please note that these email lists are subject specific discussion forums, rather than a way of contacting IT. For the latter, please see the contact emails in the next section.

finite-elements@listserv.manchester.ac.uk

fortran@listserv.manchester.ac.uk

sciprogramming@listserv.manchester.ac.uk

university-gpuclub@listserv.manchester.ac.uk

¹ Unfortunately, we have an issue currently in that we have to wait for the Windows side of the teaching cluster PCs to download all their applications before we can switch these systems into Linux/Condor overnight and at weekends. We will restore the full service as soon as we possibly can.

Faculty Contact Emails

Humanities:

Hum-ITResearch@manchester.ac.uk.

Engineering and Physical Sciences:

EPS-ITResearch@manchester.ac.uk.

Next edition

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

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

Medical and Human Sciences:

MHS-ITResearch@manchester.ac.uk.

Life Sciences:

FLS-ITResearch@manchester.ac.uk.