

UCL Computational Resource Allocation Group (CRAG)
Monthly Meeting
Friday 14th February 2014 at 13.00
Room 103, Podium Building, 1 Eversholt Street, London NW1 2DN

Chair:

1. Prof Nik Kaltsoyannis (NK) – *Molecular Quantum Dynamics and Electronic Structure*

Present:

2. Dr Nicholas Achilleos (NA) – *Astrophysics and Remote Sensing*
3. Prof Dario Alfe (DA) – *Thomas Young Centre (Materials Science)*
4. William Hay (WH) – *Datacentre Services, ISD*
5. Thomas Jones – *Research Platforms Team Leader (Infrastructure Lead), ISD*
6. Ian Kirker (IK) - *Research Comp & Facilitating Services, ISD*
7. Jo Lampard (JL) - *Senior Research IT Services Facilitator, ISD*
8. Dr Andrew Martin (AM) - *Structural & Molecular Biology*
9. Dr Vincent Plagnol (VP) – *Next Generation Sequencing*
10. Dr Bruno Silva (BCS) – *Research Computing Platforms Team Leader (Service Lead), ISD*
11. Dr Sergey Yurchenko (SY) – *Atomic, Molecular, Optical and Positron Physics*

Apologies:

12. Clare Gryce (CG) – *Head of Research Computing and Facilitation Services, ISD*
13. Dr Simon Kuhn – *Engineering Sciences*

In attendance:

14. Corrinne Frazzoni – *Administrative Services, ISD (Minutes)*

Note: Minutes below provide a high level summary of decisions taken and actions assigned by the Group.

1. Approval of Minutes of last meeting held on 17th January 2014

The Group approved the Minutes of the last meeting. There were no matters arising.

2. Update on status of current Actions

The list of current Actions (below) was updated. No new Actions were added.

3. Review of any requests for additional resources on local HPC facilities

There were no new requests for additional resources.

4. Review of any Centre for Innovation (Cfi) access requests (Chair)

No requests were submitted.

5. Review of Legion usage statistics <http://feynman.rits-isd.ucl.ac.uk:8888>

This month's statistics show an improvement in the numbers of the user base. This was also the month where a new "functional share" policy was trialled. Comparative analysis of January's statistics with previous months shows that there was a more even distribution of run times across active users in January, higher slowdown for the highest users (which is the desired behaviour) and fewer spikes in slowdown (outliers) for users with lower utilisation of the cluster. The CRAG has decided therefore that it should maintain the new scheduling policy for further analysis in coming months.

6. Review of IRIDIS and EMERALD usage statistics

The group reviewed the IRIDIS and EMERALD usage statistics – it was noted that UCL had higher slowdown on IRIDIS owing to the number of UCL users now using this system. Emerald, has low utilisation. High slowdown numbers are being addressed by the systems administration team at RAL via a mix of user education and scheduler configuration changes. A profile of large job sizes which block computing resources for access to GPU's, and over-estimation of wall clock times appear to be behind the observed slowdown on Emerald. The CRAG will continue monitoring this situation.

7. Update regarding development on new application form

The Research Computing team has received feedback regarding its application process and the perceived high barrier to entry for use of its services. For this reason the CRAG decided that it would be appropriate to review the currently proposed form to make it as simple and easy to use as possible and at the same time collect all information that is deemed essential for the sustainability of the services.

The following modifications were agreed by the group.

Roles and whether funded or not

- PI – assume funded
- Lecturer – can be either funded or unfunded
- Visitor – can be either funded or unfunded – must provide UCL PI name
- PhD Student – assume funded – must provide supervisor name
- Post-Doctoral researcher – assume funded – must provide PI name
- Everyone else – must be funded

Upon selection, the form should automatically display any relevant questions (hide irrelevant ones).

We will not ask for grant funding numbers/codes the first time someone applies, but it will be made clear to them that the account is valid for a year from creation, after which the user will have to provide funding (Award Numbers) and publication output which was generated using Research Computing services to continue having access to the platforms

on offer. We will also remind researchers to add the string Legion@UCL to acknowledgements in their publications for easy search.

The form will request the description of only one valid project title and accompanying description – this will be sufficient for a user to gain access to the services and will be used for categorisation according to consortium.

All services will be provided to a researcher with one single application. There will be one field where the work performed in the project described above should be explained at a technical level, in terms of use of the available platforms, as well as any applications/libraries required.

Funding/collaboration information to be captured at renewal as it will more accurately describe what researchers have produced using the available services.

The publication and funding information required is:

- Award number (we should aim to populate this automatically – we will have to ask researchers to provide this information in the interim. In the future, this could also auto-fill if project code is entered)
- Research outputs in the last year which have used Research Computing platforms and services.

If a researcher is working on a new project at or after re-application, they need to apply again with updated project information and request access to a consortium.

In both new applications and renewal, there will be a commitment to acknowledgement of Legion/Cfl (check box) in order to submit.

Acceptance of T&Cs (check box) in order to submit.

8. Proposal for decommissioning Unity

The group agreed that new account requests to access Unity would be terminated as soon as possible. Current users of the system should be informed that the service will be terminated in three months. Users to be advised that data would be retained for three months maximum after the service is terminated.

9. AOB

Brief discussion of creation of new consortium with Centre for Digital Humanities – the CRAG agreed that this may be useful as there is a growing community of researchers who don't match the current research areas described by the existing consortia, and in particular, the Social and Historical Sciences consortium.

CG requested considering the creation of a minimum threshold for core availability – the CRAG believes that 95% is a reasonable minimum, which may be reviewed periodically.

10. Next meeting date and agenda

Wednesday 12th March 2014 from 13.00-15.00

Venue: Room 104, 1st floor, Podium Building, 1 Eversholt Street, London, NW1 2DN.

Agenda (Items) for the next meeting:

Standing items:

1. Approval of Minutes of last meeting
2. Update on status of current Actions
3. Review of any requests for additional resources on local HPC facilities
4. Review of any Centre for Innovation (Cfi) access requests
5. Review of Legion usage statistics
6. Review of IRIDIS and Emerald usage statistics

New items for next meeting:

- Discussion around the nature of Consortia and their purpose.

LIST OF CURRENTLY APPROVED EXCEPTIONAL REQUESTS

Requesting user	CRAG approval date	details of exception	start date agreed	end date agreed	date removed	Implementation Notes
Francesco Lescai	11/10/2013	5 Terabytes of backed up, node-writeable storage. Will implement as 5 terabytes of scratch, with ongoing work to provide backups to NFS-2	1/11/2013	31/03/2014		Currently only a 5TB quota on Scratch is being granted - we have an issue in Github to provide a backup.
Eugenio Pasini	17/01/2014	Scratch quota increased to 1TB for the requested period	17/1/2014	17/4/2014		

LIST OF CURRENT ACTIONS

Shaded (closed/completed) items will be deleted in the next version.

	Actions	Status	Owner
131	IRIDIS and EMERALD usage statistic	<p>(14/6/2013): BCS to liaise directly with Derek Cross and Timothy Metcalf to seek further clarification on statistics presented for EMERALD</p> <p>(12/7/2013): ONGOING</p> <p>(17/9/2013): ONGOING</p> <p>(11/10/2013): Request Timothy Metcalf to integrate: https://www.emerald.rl.ac.uk/ganglia/graph.php?m=gpu_util_all2_report&z=xxlarge&c=GPU%20Cluster&r=month as a regular monthly report.</p> <p>(22/11/2013): ONGOING, request has been made.</p> <p>(13/12/2013): BCS to add link to stats on next report</p> <p>(17/1/2014): BCS to add link to next report</p> <p>(14/2/2014): CLOSED</p>	BCS

133	Legion users to migrate to IRIDIS	<p>(14/6/2013): Legion users, as those with the highest core count would be migrated to IRIDIS. It was agreed that BCS would identify and provide a list of users of 32 cores and above.</p> <p>(12/7/2013): BCS has created list and will initiate IRIDIS accounts and migrate (with overlap period of 3 months) users accordingly. List to be presented at next meeting.</p> <p>(17/9/2013): BCS to inform users of migration decision and of policy, and to contact individual users (using cores greater than 36).</p> <p>(11/10/2013): Migration successful and on-going, BCS to supply actual numbers on amount of users moved and accounts created.</p> <p>(22/11/2013): 22 new accounts have been created, 14 of which have moved from Legion, migration ongoing. Fuller discussion with BCS at next meeting to discuss date for capping maximum core size on Legion.</p> <p>(13/12/2013): Priority for jobs of 36 cores or more to be given lower weighting on Legion once new application process is in place.</p> <p>(17/1/14): Priority for 36 core jobs to change on 6th of February. No longer dependent upon new application process, although this should be in place by then.</p> <p>(14/2/2014): CLOSED</p>	BCS
134	KLB Power and Cooling	<p>(12/7/2013): TJ to liaise with Simon Marham for an update regarding KLB's power and cooling upgrade work.</p> <p>(17/9/2013): ONGOING</p> <p>(11/10/2013): ONGOING</p> <p>(22/11/2013): Work currently in progress. ONGOING</p> <p>(13/12/2013): CG chasing up. Group expresses deep concern. ONGOING</p> <p>(17/1/14): If nothing happens by next CRAG then consider escalation to higher governance group. ONGOING</p> <p>(14/2/14): Delayed due to safety issues. ONGOING</p>	TJ
135	Review of Legion usage statistics	<p>(12/7/2013): BCS to investigate the unexpected wait time spikes for users with small run times.</p> <p>(17/9/2013): ONGOING</p> <p>(11/10/2013): Standing Agenda Item: Identify (full name &</p>	BCS

		<p>user ID) & contact users with systematic problems, try to resolve problems.</p> <p>(22/11/2013): BCS to investigate whether it is possible to remove jobs from the slowdown graph which are part of arrays that have already started.</p> <p>(13/12/2013): Slowdown statistics for job arrays to be calculated according to start time of first job in array only. Check-pointing jobs also to be treated similarly according to initial start time (except for jobs that fail quickly).</p> <p>(17/01/2014): Pending confirmation. ONGOING</p> <p>(14/2/14): ONGOING</p>	
140	<p>General policy proposal for priority access to Research Computing resources</p>	<p>(17/9/2013): BCS to draft new policy to be presented at next meeting.</p> <p>(11/10/2013): ONGOING</p> <p>(22/11/2013): The group would like an explanation of what the value of the 'C' factor included in the leasing calculations is, and how it was derived.</p> <p>NK suggests that the last paragraph belongs before the section about leasing as it relates to buying hardware.</p> <p>Regarding the access policy for purchased and leased nodes, the group would like to see written down some guarantee of how long owners/leasers would have to wait before they could access their nodes. They would also like to see some consideration of the implications for killing active jobs and how this would be handled.</p> <p>(13/12/2013): BCS to recirculate updated priority access document for next meeting including recommendations for two tier pricing system for immediate/delayed access.</p> <p>(17/01/2014): BCS to report back to next CRAG meeting with a proposal for promoting the new policy.</p> <p>(14/2/14): the proposal was made, and will be implemented as follows: Email to the Research Computing Forum Email to the service mailing lists Information to be provided on website in relevant location (TBD) with "promotional" information</p> <p>ONGOING</p>	BCS

141	Multi-disciplinary research and nature of consortia	<p>(17/9/2013): BCS to provide list of unusual requests for next meeting with Consortia definition and objectives.</p> <p>(11/10/2013): Monitor requests and report to Feb 2014 highlighting any bounced requests by consortia.</p> <p>(22/11/2013): ONGOING</p> <p>(13/12/2013): ONGOING</p> <p>(17/01/2014): ONGOING</p> <p>(14/2/14): Report no monitored requests done, showing a number of cases where applicants had been moved because they misunderstood what the consortia represented. Add discussion to agenda for next meeting. ONGOING</p>	BCS
143	IRIDIS user feedback	<p>(22/11/2013): Email current IRIDIS users to ask them to let us know about strange queuing behaviour. Anecdotal evidence to be taken to Cfl meeting on 12th December.</p> <p>(13/12/2013): ONGOING</p> <p>(17/01/2014): ONGOING</p> <p>(14/2/14): CLOSED</p>	OK
145	Web mock-up of new application form	<p>(22/11/2013): Implement changes to form:</p> <ul style="list-style-type: none"> • make data format easier to analyse • look into possibility of populating renewal form with previous year's publications data from RPS • consider back-end support for hosting the form and associated database. <p>(13/12/2013): OK to update form to include information on platforms and produce final version for approval at next meeting.</p> <p>(17/01/2014): The new forms should be implemented subject to the following changes being made:</p> <ul style="list-style-type: none"> - data to be captured on a per project basis - project data only necessary on renewal form if there is a new project - an example of a completed form should be provided to guide users <p>(14/2/14): Covered in Agenda Item 7. New requirements gathered – implementation has started. ONGOING</p>	OK/BCS

146	Create new consortium for Gatsby Centre	<p>(22/11/2013): Make the necessary arrangements and changes to set up the Gatsby Centre consortium.</p> <p>(13/12/2013): ONGOING</p> <p>Consortium to be added pending new application process implementation</p> <p>(17/01/2014): ONGOING</p> <p>(14/2/14): ONGOING</p>	BCS
148	Proposal of new Scheduling policy	<p>(13/12/2013): Implement functional share policy.</p> <ul style="list-style-type: none"> - To be implemented on 6th January 2014. - CRAG to assess impact at next meeting <p>(17/01/2014): ONGOING – current statistics are not relevant.</p> <p>(14/2/14): CLOSED</p>	BCS/ WH
149	Review of IRIDIS and EMERALD usage statistics	<p>(17/01/2014): BCS to ask Tim Metcalf for an explanation of the following:</p> <ul style="list-style-type: none"> - why is there any slowdown on EMERALD if only 43% is used? - why does it appear that UCL's slowdown is much higher than Oxford's on IRIDIS? <p>(14/2/14): Explanations were given by Timothy Metcalf regarding these questions as reported in the Cfl statistics item in this meeting. CLOSED</p>	BCS
150	Statistical science legion access query	<p>(17/01/2014): BCS to advise statistical science of the CRAG's view that the standard access policy should be followed for centrally funded resources but that a departmental reserve may have its own policy.</p> <p>(14/2/14): Document to send to Stats department is being finalised. ONGOING</p>	BCS
151	KPI for legion wait times	<p>(17/01/2014): After correcting for job arrays, mean slowdown will be calculated for each job type (single core, single node, multi-node etc.) on a monthly basis. The use of this measure will be evaluated at a subsequent CRAG meeting.</p> <p>(14/2/14): This is now being done for senior management reports – will be introduced in coming Legion statistics reports. ONGOING</p>	BCS