

UCL Computational Resource Allocation Group (CRAG)

MEETING MINUTES

17th January 2014

In Attendance:

1. Prof Nik Kaltsoyannis (Chair) – *Molecular Quantum Dynamics and Electronic Structure*
2. Prof Dario Alfe – *Thomas Young Centre (Materials Science)*
3. Dr Simon Kuhn – *Engineering Sciences*
4. Dr Sergey Yurchenko – *TBC*
5. Dr Bruno Silva – *Research Computing Platforms Team Leader (Service Lead), ISD*
6. Jo Lampard - *Senior Research IT Services Facilitator, ISD*
7. Dr Tom Couch – *Research IT Services Facilitator, ISD*

Apologies:

1. Dr Nicholas Achilleos – *Astrophysics and Remote Sensing*
2. Dr Vincent Plagnol – *Next Generation Sequencing*
3. Thomas Jones – *Research Platforms Team Leader (Infrastructure Lead), ISD*
4. Clare Gryce – *Head of Research Computing and Facilitation Services, ISD*

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

1. Approval of Minutes of last meeting on 13th Decemeber 2013

The Group approved the Minutes of the December 13th 2013 meeting.

2. Update on status of current Actions

The list of current Actions (below) was updated, and new Actions arising were added.

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

Eugenio Piasini's request for additional storage on Legion was approved by the group. EP's storage to be increased to 1TB for three months.

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

No requests were submitted.

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

These statistics refer to the last month before the functional share policy was implemented so the CRAG decided that there is little point in looking at these stats now. They will be revisited for comparison with next month's stats, post change in share policy. A cursory inspection suggests nothing unusual.

6. Review of IRIDIS and EMERALD usage statistics

The group note that UCL is using close to its IRIDIS allocation but that underuse is still an issue for EMERALD. JL points out that NVIDIA have been invited to run a CUDA training course next term which may stimulate interest in GPU computing.

With respect to the slowdown graphs presented at the end of the report, the group queried why there is any slowdown on EMERALD if machine is only 43% used? Also, why does it appear that UCL's slowdown is much higher than Oxford's on IRIDIS?

ACTION: BCS to ask Tim Metcalf for an explanation of these stats supported by evidence (see new action 149).

7. General process and policy for priority access and leasing of existing HPC resources draft review

The CRAG approves the document. The group then discussed the options for promoting this new policy. NK expressed that the policy document should be made available to researchers. BCS advised caution over promoting the policy too much given the limited data centre space.

ACTION: BCS to report back to next CRAG meeting with a plan for promoting the policy (see action 140).

8. Suggestions by the Department of Statistical Science for improvement of the Legion access process

BCS explained that statistical science is considering using legion as their sole computational resource. This would allow them to get rid of high spec workstations under researchers' desks and replace them with something more lightweight. They have approached RCPS to see if this possible. Statistical science are effectively asking for a departmental reserve on Legion, but a sticking point is that they are concerned that the current application process for setting up and renewing Legion accounts is too cumbersome. BCS made the case for a new application process which would allow PIs in this department to add users to their projects and to take responsibility for providing information on the benefits gained from use on an annual basis.

NK pointed out that this effectively turns a small job for a large number of people into a large job for a few people. The group agreed that the information gained as a result of the application and renewal process is important for justifying the existence of the service and supporting proposals for future upgrades, and that there is an onus on the users of this

'free' service to provide that information.

BCS suggested that another option is to allow a simplified access process for the reserved departmental hardware whilst access to the rest of the Legion cluster would require the usual application process.

The group agreed that the new online application and renewal process may address many of the concerns that statistical science have which may lead them to reconsider their objections.

ACTION: BCS will advise Statistical Science that access to the centrally funded pool will require them to use the standard application and renewal process, but that a departmental reserve may have its own application process (see new action 150).

The group then discussed the new online application forms. NK suggests the following changes:

- the application form should capture data on a per project basis, expanding to allow users to enter data about multiple projects if appropriate
- the renewal form should make it clear that researchers only need to complete the project details if it's a new project.

SY also suggests that an example of a completed form be made available to guide users as to the amount of detail to include for each section of the form.

ACTION: The new application and renewal forms are approved (subject to the above amendments) and should be implemented (see action 145).

9. Nodes of type T (high memory node) policy regarding highly threaded jobs

BCS: currently these machines have 32 cores which is a unique characteristic on top of their large memory. Should these nodes be available for high threading jobs as well as high memory jobs?

NK: Would these jobs block the jobs of users requiring high memory, and wouldn't using two 16 core nodes be an acceptable alternative?

BCS: Two 16 core nodes would not necessarily be an alternative. Currently jobs requesting more than 48GB of RAM go to the fat nodes. 32 core jobs with low ram are effectively treated as backfill on the T nodes and are restricted to a maximum wall clock time of 12 hours.

The group decided to maintain the status quo and that researchers who require 32 core nodes for an extended period of time should make a formal request to the CRAG for an exception to the current 12 hour limit.

10. Discussion: Building a simple KPI which reflects user experience with wait times

There is a need for a KPI for wait times on Legion which can be used to demonstrate the need for increased capacity on Legion as use increases over time. Clare Gryce has previously suggested mean slowdown with deviation for this purpose.

The group discussed the ease with which this measure could be calculated and how meaningful it would be. It was agreed that this measure should be calculated for each job type on a monthly basis as a trial to see if any meaningful trends can be observed.

ACTION: 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 (new action 151).

11.AOB

None

12.Next meeting date and agenda

14th February 2014 from 1pm – 3pm, Venue: Room 103, 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:

- none

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>	BCS

		<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>	
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>	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>	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 & 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>	BCS
140	General policy proposal for priority access to Research Computing resources	<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>	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>	BCS

		(13/12/2013): ONGOING (17/01/2014): ONGOING	
142	Legion usage statistics	(11/10/2013): <ul style="list-style-type: none"> • Add new node classes, • identify users by names and user id • investigate the long wait times for Social and Historical Sciences (22/11/2013): <ul style="list-style-type: none"> • Add new node classes - DONE • identify users by names and user id - DONE • investigate the long wait times for Social and Historical Sciences - ONGOING (13/12/2013): ONGOING Investigate the long wait times for Social and Historical Sciences - ONGOING until SHS jobs are submitted (17/01/2014): CLOSED (possibility of revisiting if problem recurs)	BCS
143	IRIDIS user feedback	(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 12 th December. (13/12/2013): ONGOING (17/01/2014): ONGOING	OK
144	Interactive computing requirements document	(22/11/2013): Work with e-learning department to create a requirements document for an interactive computing platform to support research based teaching. (13/12/2013): ONGOING (17/01/2014): CRAG has already made sufficient decisions in this area. CLOSED	CG
145	Web mock-up of new application form	(22/11/2013): Implement changes to form: <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. 	OK/BCS

		<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 	
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>	BCS
147	Legion usage statistics	<p>(13/12/2013): Investigate unusual slowdown behaviour for zcqs34 and ucesjw0</p> <p>(13/12/2013): During the meeting OK discovered that the users' jobs were restricted to nodes of type T,U,Z (either by the user or as a result of the resources requested) and that they were submitting large arrays of jobs so they experienced a miniature version of the array job wait artefact issue.</p> <p>(17/01/2014): User managed to submit a zero core job which failed instantly after 145 hour queue time. CRAG understands unusual slowdown behaviour. CLOSED</p>	OK
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>	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? 	BCS

150	Statistical science legion access query	(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.	BCS
151	KPI for legion wait times	(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.	BCS