

UCL Computational Resource Allocation Group (CRAG)

MEETING MINUTES

19th April 2013

In Attendance:

1. Prof Nik Kaltsoyannis (Chair) - *Molecular Quantum Dynamics and Electronic Structure*
2. Dr Vincent Plagnol – *Next Generation Sequencing*
3. Dr Ben Waugh - *High Energy Physics*
4. Dr Bruno Silva - *Research Computing Platforms Team Leader (Service Lead), ISD*
5. Dr Andrew Martin - *Bioinformatics and Computational Biology*
6. Dr Nicholas Achilleos - *Astrophysics and Remote Sensing*
7. Clare Gryce - *Head of Research Computing and Facilitating Services, ISD*
8. Thomas Jones - *Research Platforms Team Leader (Infrastructure Lead), ISD*
9. Jo Lampard - *Senior Research IT Services Facilitator, ISD*

Apologies:

1. Prof Dario Alfe - *Thomas Young Centre (Materials Science)*

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 15th March 2013

The Group approved the Minutes of the March 15th 2013 meeting, with the following corrections:

Item 5: The second sentence should read “It was suggested that BCS remove projects that **were** being given priority access”.

Item 9: Michail Stamatakis is from the Department of Chemical Engineering, not Chemistry.

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.

There were none for the current month.

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

(Doc: Cfi_Access_Application_v0 6)

The discussion of the following applications was postponed until after agenda item 7, in the light of the review of the IRIDIS access policy.

The following applications were approved:

- David Fonz
- Gary Davies

The following applications were discussed:

- John Brodholt

As John had several collaborators from the University of Leeds who would also require accounts on IRIDIS, it was decided that they should be asked to apply for a UCL username to allow them to be registered on IRIDIS with a UCL email address. This would enable them to be accounted for as UCL users within the Cfi system. Once this had been done, the application would be approved. As the job size could potentially be anywhere from 8 – 128 cores, their usage would need to be carefully monitored in the first instance in order to prevent the smaller jobs from flooding the high quality service partition. In the long term, this type of job would ideally be directed to the lower quality partition by the job scheduler (CG to investigate whether this would be technically possible (see new Action 122).

- Kristen Marino

The jobs were large (requiring more than 100 cores) so would be suitable for IRIDIS; however, GROMACS is available on Emerald and the jobs should run well on GPUs. BCS to speak to Kristen to see if this is a possibility; if not, her application for IRIDIS would be approved (see new Action 123).

5. Legion usage report for March. Report available <http://feynman.rits-isd.ucl.ac.uk:8888>

BCS presented the Legion usage statistics for March 2013; slowdown is now calculated using requested time and not run time. He pointed out that the graph display was still not ideal as very small run times contaminate the data and lead to problems with the scale. It was noted that TYC jobs still showed very short wait times which could be due to backfill, but the group felt that the actual reason for the apparent anomaly should be established; in contrast, two users (uccajab and zccamku) still showed very long wait times for their jobs. BCS to investigate (see new action 124).

6. Enquiry to current Emerald account holders regarding their use of the system - *discussion (CG)*

CG referred to the last usage statistics graphs, and emphasised that the difference between Bristol's very high usage and UCL's low usage was due to only a very small number of very active users in the Computer Science department at Bristol. She explained that there was a need to justify requests for further capital expenditure in the light of under-usage of existing facilities. She had therefore emailed all of the current UCL Emerald users to find out if there were any factors blocking them from making more use of the system, and had received several replies. She monitors all tickets raised by the Emerald helpdesk system and intervenes when it appears to be necessary; the helpdesk has been prompt and helpful so there is no real problem with support. However, some of the replies to her email would seem to indicate that UCL users were not always contacting the helpdesk as soon as they experienced an issue. For example, Dave Wright had contacted her that morning to complain that there is a limit on the number of jobs he can submit at a time; CG had suggested that he contact the Emerald helpdesk to ask if the limit can be increased or removed – she will back up the request if necessary.

It was felt that there were multiple issues; a problem with communicating Emerald's benefits to potential users, and also UCL's existing users need to be given additional support. CG mentioned that the Research Software Development team will be undertaking a GPU-based project next term in order to increase their skills in this area. The possibility of holding a CUDA workshop was raised; Mike Giles runs a CUDA/NVIDIA summer school and an online course. NK felt that Emerald usage might pick up given sufficient time, but thought that a one-day workshop would be useful in the meantime. BCS pointed out that the GPU users' email list is separate from the one for Legion users, and that perhaps the two should be merged. VP asked if there was a facility for "support at your desk" to get users up and running with porting code to GPU architecture. In the long term, CG thought that the Research Software Development team would be able to help with this; in the meantime, she will get in touch with Istvan Reguly (the Cfl's GPU support officer,

who has already collaborated with Serge Guillas to get his code running on Emerald) to arrange a visit to VP's group.

Research Computing will send out an email to Legion users in the next couple of weeks highlighting the availability of Emerald, the applications already installed, and offering pre-booked desktop support visits; the possibility of holding another GPU event to be discussed next week in a Research Computing group meeting (see new Action 125).

7. IRIDIS access policy review (CG)

CG explained that the next version of IRIDIS (4) is currently under test at Southampton, and the hardware is due to be signed off in the next couple of weeks. This will result in the total number of cores being increased from 4,000 to 12,000, along with UCL's share of the machine being increased to 50% (i.e. 6,000 cores - around the same size as Legion). This will be online by mid-May; as Legion is highly contended, it would be desirable to get as many UCL users onto IRIDIS as soon as possible. The intention is that the 12,000 cores will be divided into 2 logical partitions, one with 4,000 and the other with 8,000 cores. The smaller partition will be easier to gain access to (with usage only limited to each institution's fair share), and is not expected to have fast turnaround, i.e. it will be a low quality service. The larger partition will have much more restricted access, with each institution deciding who should be allowed on and for how long. One institution wants to have a very small number of approved 'high quality' users on the high quality service, with maintenance of the list of approved users being under their control. They can then decide to add/remove users at any desired interval, enabling them to have consistently fast job turnaround.

After a lengthy discussion, it was decided that UCL should adopt a similar approach, with existing IRIDIS users being automatically migrated in due course to the higher quality service. Any active (within the last 6 months) Legion user will automatically be granted access to the lower quality IRIDIS service on a "free-for-all" basis. This would ease the existing pressure on Legion, and users would have an easier route to try out IRIDIS. Obviously users cannot be compelled to migrate to IRIDIS if they require access to software which is only licensed on Legion, but BCS pointed out that a common software stack is being worked on which will limit the number of such cases. NK suggested that 3 months after existing Legion users have been given an account on IRIDIS, they should be informed that any job requiring $>x$ cores (with the value of x still to be determined!) will no longer run on Legion, and should be run on IRIDIS instead. VP asked whether the IRIDIS application form should be removed from the website? CG said that before the next meeting, the Research Computing team would consider how a single account application form could be designed (but still retaining the option of applying for exceptional access), and also how automatic IRIDIS account creation for existing Legion users could be implemented (the current process involves ssh keys and is very slow to setup). They would also start the

process of identifying exceptional IRIDIS users, to be migrated to the high quality service as soon as possible (see new action 126).

8. AOB

Miracle have now used up all of their allocation. BCS's team will be dealing with data migration and closing external accounts; a generous closedown period has been allowed for this.

9. Next Meeting Date and Agenda

- 10th May 2013, 1pm – 3pm (venue to be confirmed).
-

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 (Chair)
5. Review of Legion and IRIDIS usage statistics

New items:

6. Review of draft policy for Legion Priority Queue access (BCS/CG)
7. Update on IRIDIS expansion (CG)

LIST OF CURRENT ACTIONS

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

	Actions	Status	Owner
76	Peter Harrison request for additional Unity resource	<p>(27/01/2012): To follow up with requestors regarding additional information about check-pointing and problem decomposition.</p> <p>(24/02/2012): Update from DG, Peter Harrison has been working closely with RC team and adapting code where possible. CRAG agreed to extend 72hr wall-clock to 10 days on Unity. CLOSED to be REVISTED at a later date</p> <p>(24/05/2012): DG to contact the user regarding progress of his work.</p> <p>(26/06/2012): Re-assigned to BCS</p> <p>(25/07/2012): BCS to speak to contact Peter Harrison.</p> <p>(10/09/2012): PH requests continued use of Unity and will provide review for next CRAG meeting.</p> <p>(12/10/2012): CRAG approved Peter Harrison's indefinite request with proviso that Legion is cited in research papers. BCS to provide update to CRAG in 3 months.</p> <p>(14/12/2012): ONGOING – update next month</p> <p>(11/1/2013): ONGOING – update next month</p> <p>(8/2/2013): Re-visit in 3 months.</p> <p>(19/4/2013): BCS has now spoken to Peter, and he no longer requires the resource.</p> <p>CLOSED</p>	BCS
91	Establish policy for requesting Priority CP hours	<p>(10/09/2012): CG to circulate draft policy paper to CRAG members and inform Serge Guillas that his request is under review. All CRAG members to report back on implementation of Priority Queue.</p> <p>(12/10/2012): TJ to investigate implementation of Priority Access using 'Projects' method as discussed and agreed by group. Gold Accounting Software – RC to investigate by further testing. NA to provide local Miracle users for testing.</p> <p>(9/11/2012): Still pending. Meanwhile, TJ to set up priority access for Miracle jobs as previously agreed using same set up on Miracle as for Harvest project (Serge Guillas).</p> <p>(14/12/2012): Done for Miracle jobs; TJ to present Gold</p>	TJ

		<p>accounting software information at the next CRAG.</p> <p>(14/12/2012): Gold accounting software installed. Client and lustre upgrade still pending. TJ to report back at next CRAG.</p> <p>(8/2/2013): TJs team are testing, looking into SGE and thinking about how to implement it. TJ to write up and report at next CRAG.</p> <p>(15/3/2013): Gold accounting software installed, configured and deployed on 18/3/2013.</p> <p>(19/4/2013): Agenda item for the next meeting to review and approve the draft policy for access to the Priority Queue.</p> <p>ONGOING</p>	
96	Record of Cfl applications	<p>(9/11/2012): It was agreed that a spreadsheet record of all Cfl requests, including reasons for rejection where appropriate, should be maintained.</p> <p>11/1/2013 – CG to maintain list of usage and report to CRAG every three months.</p> <p>(8/2/2013): ONGOING – update next month</p> <p>(15/3/2013): ONGOING</p> <p>(19/4/2013): CG has set up the spreadsheet with BCS to maintain</p> <p>ONGOING</p>	CG/BCS
101	Cfl usage statistics	<p>(14/12/2012): CG to clarify whether STFC's share of Emerald is included in the under-utilisation figures.</p> <p>(11/1/2013): ONGOING</p> <p>(8/2/2013): ONGOING</p> <p>(15/3/2013): ONGOING</p> <p>(19/4/2013): STFC's share of Emerald is now explicitly shown in the usage data</p> <p>CLOSED</p>	CG
103	Cfl IRIDIS job classes	<p>(14/12/2012): CG to obtain data on IRIDIS job class distribution.</p> <p>(11/1/2013): CG to present job clusters from Cfl</p> <p>(8/2/2013): ONGOING</p> <p>(15/3/2013): ONGOING</p> <p>(19/4/2013): All available data is now included in the</p>	CG

		<p>quarterly report; CG is examining the possibility at obtaining monthly data. IRIDIS to be added onto the standing agenda item for Legion usage.</p> <p>CLOSED</p>	
110	Consortia mailing list	<p>(11/1/2013): BCS to advertise Consortia mailing lists on the Research Computing webpage with an archive and search facility for consortium members.</p> <p>(8/2/2013): The archives are not searchable at the moment; TJ to request that the lists are transferred to Mailman which will allow this. It may not be possible to have a single cross-searchable archive for all the lists. TJ will find out if non-members can have read-only access. BCS will add links for each archive to the web pages if necessary.</p> <p>(8/2/2013): BS to replace/update mailing list on consortia webpage and send mailman links to consortia leaders to view all messages.</p> <p>(15/03/2013): Mailman mailing lists added to consortia webpage. Mailman mailing list links not sent yet. ONGOING</p> <p>(19/4/2013): Consortium leaders still need to be sent information about their respective mailing lists; TJ to do this.</p> <p>ONGOING</p>	TJ and BCS
113	Change of NGS Consortium Leader	<p>(8/2/2013): TJ to look at the NGS mailman list and queues for VP (if necessary).</p> <p>(15/03/2013): ONGOING</p> <p>(19/04/2013): DONE</p> <p>CLOSED</p>	TJ
114	Legion scratch quotas	<p>(8/2/2013): BCS to implement increased Legion scratch quotas of double the existing default allocation (200Gb).</p> <p>(15/03/2013): ONGOING</p> <p>(19/04/2013): The recent issue with NFS over the preceding weekend had prevented TJ from switching scratch quotas back on; the information has been updated, but the database needs to be moved to another server in case that was the cause of the problem. Once done, TJ will be implementing the increase to 200Gb default allocation next week.</p> <p>ONGOING</p>	BCS/TJ

115	Amend stats graph so that zero-utilising Consortia are still shown.	(15/03/2013): ONGOING (19/04/2013): ONGOING	BCS
116	Request for Additional Resources Form	(15/03/2013): Amend form to state for how long users would require the additional resource (19/04/2013): BCS has now amended the form with the additional fields. CLOSED	BCS
117	Dr. Oleg L. Polyansky	(15/03/2013): BCS to liaise with Oleg L. Polyansky to establish limit for backfill pending further steer on IRIDIS access from Cfl Research Policy group. (19/04/2013): BCS has spoken to Oleg and his colleague Lorenzo; the backfill limit has been removed as there is a large variance in runtime. To be reviewed after one month. ONGOING	BCS
118	Dr Furio Cora	(15/3/2013): BCS to liaise with Furio Cora and suggest re-application based on use of Legion. (19/04/2013): BCS had emailed Furio to inform him that the application had been refused, and asked him to contact him for further advice. He had not yet received a reply. Closed owing to change of policy at Cfl. CLOSED	BCS
119	Legion usage Report available http://feynman.rits-isd.ucl.ac.uk:8888	(15/3/2013): BS to present graph of slowdown per user using both requested time and run time. (19/04/2013): ONGOING	BCS
120	Role of Consortia leaders and Account Application Policies and Processes	(Doc: RC_Consortia_and_account_application_process_v1.0.pdf) (15/03/2013): BS re-format document to present to the Research Computing Working Group for approval. (19/04/2013): BCS to reformat the document; CG to add approval of the document as an agenda item for the next RCWG meeting. ONGOING	BCS/CG

121	Michail Stamatakis UCL Chemical Engineering, £30k funding for computing infrastructure	<p>(15/3/2013): BS to report back to CRAG on:</p> <ul style="list-style-type: none"> • Costings • Integration • Hosting <p>(19/04/2013) BCS has obtained hardware costings as part of the process for the next upgrade and is to meet with Michail in the next few days. TJ also needs to get costings for hosting. They have not yet agreed on whether there will be a common software stack.</p> <p>ONGOING</p>	BCS
122	John Brodholt – IRIDIS application	(19/04/2013): BCS to inform John that his Leeds collaborators will be required to apply for UCL usernames before the application can be approved. CG to investigate the possibility of the IRIDIS job scheduler automatically diverting jobs only requiring a small number of cores to the low quality service partition.	BCS/CG
123	Kristen Marino – IRIDIS application	(19/04/2013): BCS to speak to Kristen to see if her GROMACS jobs are suitable for running on Emerald.	BCS
124	Legion usage - job wait time anomalies	(19/04/2013): BCS to investigate the very short wait times for the TYC consortium, and the long wait times for users uccajab and zccamku.	BCS
125	Increasing Emerald support and awareness	(19/04/2013): CG to contact Istvan Reguly (to arrange a visit to VP's group. An email to be sent out an email to Legion users highlighting the applications already installed on Emerald and offering pre-booked desktop support visits. The possibility of holding a one-day CUDA Workshop GPU event to be discussed.	CG/BCS /JL
126	New Legion/IRIDIS access policy	(19/04/2013): The Research Computing team to produce a new draft access policy and bring to the next meeting for discussion/approval. A single account application form for Legion and IRIDIS to be designed. The possibility of automatic IRIDIS account creation for existing Legion users to be investigated. Exceptional IRIDIS users to be migrated to the high quality service to be identified as soon as possible.	CG/BCS