

# UCL Computational Resource Allocation Group (CRAG)

## MEETING MINUTES

10<sup>th</sup> May 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. Clare Gryce - *Head of Research Computing and Facilitating Services, ISD*
6. Thomas Jones - *Research Platforms Team Leader (Infrastructure Lead), ISD*
7. Jo Lampard - *Senior Research IT Services Facilitator, ISD*
8. Tom Couch - *Senior Research IT Services Facilitator, ISD*
9. Prof Dario Alfe - *Thomas Young Centre (Materials Science)*

### Apologies:

1. Dr Andrew Martin - *Bioinformatics and Computational Biology*
2. Dr Nicholas Achilleos - *Astrophysics and Remote Sensing*

*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 19<sup>th</sup> April 2013**

The Group approved the Minutes of the April 19<sup>th</sup> 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

There were none for the current month.

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

*(Doc: Cfi\_Access\_Application\_v0.7.doc)*

The following application was approved pending further information:

- David Wright

Although David Wright owns an EMERALD account, an application was still made to use IRIDIS for GROMACS and Ambertools (MD code). GROMACS may have very limited GPU acceleration (update: version 4.6 has been written with GPU acceleration in mind). BCS to enquire from David Wright as to why this application is not suitable for EMERALD (**see new Action 127**).

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

BCS presented the Legion usage statistics for April 2013. It was noted that TYC jobs still showed very short wait times. BCS explained that this was due to a large number of users in that consortium, all of whom consumed a small amount of CPU time. Because priority is assigned to projects (effectively individual users) this resulted in a large amount of CPU time for the consortium and low wait times; in contrast, two users (uccajab and zccamku) still showed very long wait times for their jobs. BCS investigated this and found that user uccajab was requesting 4GB of RAM per process and 64 core jobs, which requires 16 nodes of type W to be reserved for his jobs. User zccamku was requesting 7 core jobs, and this was causing him long wait times because of the single-node nature of his jobs – there are no reservations for this type of job. BCS noted that user ucablpe also had large wait times owing to the nature of his jobs – mainly serial with very short run times.

## 6. Report back to the CRAG on the Research Computing working Group (RCWG) meeting of 02.05.2013 (*Chair*)

NK reported that the CRAG proposal for future Cfl Iridis Access policy document had been approved by the RCWG. There was a very strong steer from Professor John Brodholt, Dept of Earth Sciences each consortium be assigned a resource quota, managed by consortia leaders. CG highlighted that this would be very difficult to implement. NK indicated that as IRIDIS would be undergoing an extension, quota assignment was not a priority. BCS would re-design the Cfl Proposal for Access to Facilities application form for review by July. It was agreed this would be a useful process to undertake prior to migration. **(see new Action 128).**

## **7. Update on IRIDIS expansion (CG)**

CG stated that there were no updates to report but would be meeting with the IRIDIS team soon to discuss the following:

- At what stage are they with migration?
- Where are they with the scheduling idea proposed by UCL?
- How easy will it be to transfer UCL users onto the local partition?

## **8. Review of draft policy for Legion Priority Queue access (BCS/CG)**

BCS informed all that there were two main considerations with regards to the policy:

- Type of priority access
- Mechanism on how priority access is given

NK explained how the Miracle consortium distributed priority access and would this model be a viable option. TJ explained that some work would have to be done in Gold to redistribute priority. BCS explained that another option could be to use preemption by checkpointing running jobs and interrupting them to allow priority jobs to start.

The pilot project for Michail Stamatakis, UCL Chemical Engineering had been agreed of £35k for 4 node servers of 12 cores and priority access would be given as followed:

- Full access to machines. When not in use, allow others to use these machines
- Allow access to the software stack via virtual machines, pending outcome of common software stack tests

BCS enquired as to how long priority access should be given? NK suggested guaranteeing a minimum time of three years, which would be reviewed and re-negotiated accordingly. BCS to inform Michail Stamatakis of minimum priority access time (Action 121).

NK highlighted the importance of a cost recovery mechanism and suggested tying priority access to mechanism, in order to evaluate how much priority could be bought for a specific amount of money. NK also suggested working out CPU cost per hour while still maintaining enough access for fair share. NK highlighted that this costing would be extremely helpful when approaching funding bodies. It was suggested that the pricing model could be based on existing cases:  
3 years x costing of support + price of CPU time divided by 3 years or an alternative formula: cost of hardware divided by the number of CPUs x 3 years  
BCS and TJ to liaise with Miracle and HECToR teams, to understand how they cost priority access. **(see new Action 129)**.

NK stated that current priority should be given to the Michail Stamatakis priority access project. It was agreed to review the policy document after its completion.

## **9. AOB**

None

## **10. Next Meeting Date and Agenda**

14<sup>th</sup> June 2013 from 1pm – 3pm, Cruciform Foyer 102 Seminar Room 2,  
Cruciform Building, GOWER STREET, LONDON, WC1E 6BT.

### **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 usage statistics,
6. Review of IRIDIS and Emerald usage statistics

## LIST OF CURRENT ACTIONS

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

	<b>Actions</b>	<b>Status</b>	<b>Owner</b>
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 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><b>ONGOING</b></p>	TJ
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 spread sheet with BCS to maintain.</p> <p><b>(10/5/2013): CG to hand over to BCS.</b></p>	BCS

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.</p> <p>(19/4/2013): Consortium leaders still need to be sent information about their respective mailing lists; TJ to do this.</p> <p><b>(10/5/2013): TJ reported problems/issues re mailing lists to mailing team.</b></p>	TJ and BCS
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><b>(10/5/2013): To go live 14.05.2013</b></p>	BCS/TJ
115	Amend stats graph so that zero-utilising Consortia are still shown. Legion usage Report available <a href="http://feynman.rits-isd.ucl.ac.uk:8888">http://feynman.rits-isd.ucl.ac.uk:8888</a>	<p>(15/3/2013): BS to present graph of slowdown per user using both requested time and run time.</p> <p>(19/04/2013): ONGOING</p> <p><b>(10/5/2013): ONGOING</b></p>	BCS
117	Dr. Oleg L. Polyansky	<p>(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.</p> <p>(19/04/2013): BCS has spoken to Oleg and his colleague</p>	BCS

		<p>Lorenzo; the backfill limit has been removed as there is a large variance in runtime. To be reviewed after one month.</p> <p><b>(10/5/2013): BCS to contact Oleg regarding access to IRIDIS, to also acquire usage requirement statistics and details of total amount of resource required. In order for RC to identify approximately the number of jobs to be submitted by Oleg. To be reviewed after one month.</b></p>	
119	<p>Legion usage Report available</p> <p><a href="http://feynman.rits-isd.ucl.ac.uk:8888">http://feynman.rits-isd.ucl.ac.uk:8888</a></p>	<p>(15/3/2013): BS to present graph of slowdown per user using both requested time and run time.</p> <p>(19/04/2013): ONGOING</p> <p><b>(10/5/2013): ONGOING</b></p>	BCS
120	<p>Role of Consortia leaders and Account Application Policies and Processes</p>	<p>(Doc: RC_Consortia_and_account_application_process_v1.0.pdf)</p> <p>(15/03/2013): BS re-format document to present to the Research Computing Working Group for approval.</p> <p>(19/04/2013): BCS to reformat the document; CG to add approval of the document as an agenda item for the next RCWG meeting.</p> <p>ONGOING</p> <p><b>CLOSED</b></p>	BCS/CG
121	<p>Michail Stamatakis UCL Chemical Engineering, £35k funding for computing infrastructure</p>	<p>(15/3/2013): BS to report back to CRAG on:</p> <ul style="list-style-type: none"> <li>• Costing</li> <li>• Integration</li> <li>• Hosting</li> </ul> <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><b>(10/5/2013): CG to obtain true operational cost i.e. overheads, data resources which must be passed on. To note that this is not offered as a service but as a pilot project. To be reviewed in one month.</b></p> <p><b>(10/5/2013): BCS to inform Michail Stamatakis of three minimum priority access time.</b></p>	BCS
122	<p>John Brodholt – IRIDIS application</p>	<p>(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</p>	BCS/CG

		diverting jobs only requiring a small number of cores to the low quality service partition. <b>(10/5/2013): ONGOING</b>	
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. <b>(10/5/2013): BCS to contact Kristen and establish if her GROMACS jobs are successfully running on EMERALD, also suggest that Kristen contact GPU specialist at Cfl.</b>	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. <b>(10/5/2013): BCS to contact uccajab and zccamku to establish and understand their jobs, perhaps suggest compacting jobs.</b>	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.  <b>(10/5/2013): Handed over to JL to circulate email, organise a CUDA (GPU) workshop for UCL Genetics Institute (VP) as well as organise further GPU workshops for UCL depts..</b>	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.  <b>(10/5/2013): ONGOING</b>	CG/BCS
127	David Wright IRIDIS application	<b>(10/5/2013): BCS to investigate why this application is not suitable for EMERALD.</b>	BCS



128	Cfl Proposal for Acces to Facilities application form	<b>(10/5/2013): BCS to re-design the Cfl Proposal for Access to Facilities application form for review by July.</b>	BCS
129	Legion Priority Queue access costing	<b>(10/5/2013): BCS and TJ to liaise with Miracle and HECToR teams to understand how they cost priority access.</b>	BCS/TJ