

# UCL Computational Resource Allocation Group (CRAG)

## MEETING MINUTES

15<sup>th</sup> March 2013

### In Attendance:

1. Prof Nik Kaltsoyannis (Chair) - *Molecular Quantum Dynamics and Electronic Structure*
2. Dr Vincent Plagnol – *Next Generation Sequencing*
3. Prof Dario Alfe - *Thomas Young Centre (Materials Science)*
4. Dr Ben Waugh - *High Energy Physics*
5. Dr Bruno Silva - *Research Computing Platforms Team Leader (Service Lead), ISD*
6. Dr Andrew Martin - *Bioinformatics and Computational Biology*

### Apologies:

1. Dr Nicholas Achilleos - *Astrophysics and Remote Sensing*
2. Clare Gryce - *Head of Research Computing and Facilitating Services, ISD*
3. Thomas Jones - *Research Platforms Team Leader (Infrastructure Lead), ISD*
4. Jo Lampard - *Senior Research IT Services Facilitator, 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 8<sup>th</sup> February 2013**

The Group approved the Minutes of the February 8<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.**

A request for an extension of maximum wall clock time to 10 days on Unity from Alison Wright was approved, with a concession of three months.

It was decided to amend the Request for Additional Resources form, in order for users to stipulate the length of time for which they would require the additional resource. (see new Action 116)

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

*(Doc: Cfi\_Access\_Application\_v0 6)*

The following application was discussed:

- Abil Aliev

NK indicated that “waiting long time for use on Legion” was not an acceptable reason for access to Emerald facilities. AM informed all that GROMACS worked well with GPU. Abil’s case was aimed at obtaining a very short term results for a research paper for submission in April 2013 and a student paper which was to be completed by March 2013. It was decided that Abil Aliev be given priority use of Legion for one month, and if long time GPU use was required then Abil Aliev was to submit another, better argued request.

- Dr. Oleg L. Polyansky

NK highlighted that this was a request for use on IRIDIS, however this machine is for capability computing using codes that scale well above 128 cores, and suggested getting steer from the Cfi Research Policy group. NK also suggested offering special access to Legion to obtain the resource required for these jobs. BW suggested gathering a variety of small jobs from this user to gain a view of the requirements and then offer backfill. It was decided that this request would be rejected and BCS would liaise with user to establish limit for backfill pending further steer. (see new Action 117).

- Dr. Furio Cora

NK highlighted that this was a request to access Gaussian code on IRIDIS, however Gaussian is not available on IRIDIS. There was also a request for use of Crystal on IRIDIS. This request was rejected as the case appeared very short term, more as a stop-gap while additional Hector resources were put into place. NK indicated that “waiting long time for use on Legion” was not an acceptable reason for access to Cfi facilities. It was agreed that BCS contact user and suggest user make case for long term use of Crystal on Legion rather than IRIDIS. (see new Action 118).

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

BCS presented the Legion usage statistics for February 2013. It was suggested that BCS remove projects that were being given priority access. It was highlighted that TYC did not show any waiting times, it was agreed that if this pattern was repeated in the March 2013 statistics it would be investigated by the CRAG accordingly. It was agreed that “slowdown” be calculated using requested time and not run time, as this would present a fairer view. (see new Action 119).

**6. Role of Consortia leaders and Account Application Policies and Processes - discussion (Chair)**

(Doc: [RC\\_Consortia\\_and\\_account\\_application\\_process\\_v1.0.pdf](#))

It was agreed that the section titled Consortia Terms of Reference be amended and that the document then be presented to the Research Computing Working Group for review and approval. (see new Action 120)

**7. Teaching and Learning Cluster proposal and discussion**

BCS informed the group that requirements gathering had been completed and results suggest that Linux machines were required to teach Unix. BCS suggested allocation of space on Legion for this type of training. £20k had been allocated for Legion hardware and extra kit could be bought for HPC training courses. BCS suggested increasing numbers of login nodes, using these to replace Socrates and Plato for Unix training. NK suggested that £20k be not only spent of hardware but also software to aid teaching.

AM suggested using an open source version such as [www.nomachine.com](http://www.nomachine.com) for NX training, highlighting that no downloading of software was required, only a web server i.e. linking one to many machines. It was agreed that a machine was required that offered UNIX interface for X Windows desktop to gain a Linux interface. It was concluded that Research Computing would provide interactive computing resource and parallel computing training which would be undertaken on Legion via node requests.

**8. CRAG Mailing list access rules**

It was highlighted that everyone was able to view and access the CRAG mailing list via [CRAG@ucl.ac.uk](mailto:CRAG@ucl.ac.uk) and that this should remain as such.

## **9. AOB**

BCS informed all that Michail Stamatakis, UCL Department of Chemical Engineering, indicated availability of £30k funding for computing infrastructure. BCS proposed that hardware be purchased which then could be hosted near Legion, additional nodes could also be loaded and some nodes reserved in the infiniband area to offer priority. It was agreed that BCS report back to the CRAG on how this could be achieved with costings, integration and hosting etc. (See new action 121)

Postpone next CRAG meeting on 12<sup>th</sup> April 2013 to 19<sup>th</sup> April 2013.

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

- 19<sup>th</sup> April 2013, 1pm – 3pm, Cruciform Foyer 201 Seminar Room 3, 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

## LIST OF CURRENT ACTIONS

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

	Actions	Status	Owner
73	Devise proactive strategy to inform users of the availability of Condor.	<p>(25/11/11): Email consortium leaders and speak to members of the Genomics community (Jacky Pallas perhaps) regarding Condor to move this forward.</p> <p>(27/01/2012): Pending completion of on-going works to improve usability of Condor service.</p> <p>(24/02/2012): IN PROGRESS. Usability works still in progress; OK to liaise with ISD Datacentre Services to create shared storage for users to install applications (e.g. R &amp; Python) via Research Computing.</p> <p>(30/03/2012): IN PROGRESS. Technical issues regarding mounting of shared storage currently undergoing testing.</p> <p>(24/05/2012): IN PROGRESS. Technical issues regarding mounting of shared storage currently undergoing testing</p> <p>(26/06/2012): IN PROGRESS. Shared storage areas have been mounted on Condor machines, applications require placement.</p> <p>(25/07/12): Can now mount and start moving users. Going to make Condor useable from Legion itself, and should be available by next year. Will be a hybrid system.</p> <p>(10/09/2012): Modification to Condor successful. Users will be contacted to move their web jobs for Condor.</p> <p>(12/10/2012): We have found unexpected problems mounting the K drive – we are investigating what the problem is. Nevertheless, Condor is usable.</p> <p>(9/11/2012): Further testing of K drive to be undertaken and results to be reported at next meeting.</p> <p>(14/12/2012): After further testing, the problem appears to be with either DNS or firewall; some nodes in the Bartlett were removed as job submission was not possible. A temporary solution has been implemented (not using institutional filestore).</p> <p>(11/1/2013): Test launched using RC bespoke file store to remove bad connections across UCL and nodes removed from list. Resume communications with users.</p> <p><b>(8/2/2013): Issues with storage identified; some machines had to be added to server access list while those which would never work (around 200, leaving</b></p>	BCS

		<p>more than 1,000 available) were removed. The problem was not random so now considered resolved. Users still need to be able to mount the drive which allows them access to all tools; Owain Kenway is working on this for the first group of users and will contact them. There are plans to scrap Condor and have similar workloads on Legion; TJ's team have created a virtual machine "Legion node" on the new desktop. Update to be given at next meeting.</p> <p><b>CLOSED</b></p>	
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 REVISITED 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><b>(8/2/2013): Re-visit in 3 months.</b></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><b>(15/3/2013): Gold accounting software installed, configured and deployed on 18/3/2013.</b></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><b>(15/3/2013): ONGOING</b></p>	CG
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><b>(15/3/2013): ONGOING</b></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><b>(15/3/2013): ONGOING</b></p>	CG
105	Legion Teaching and Learning use	<p>(14/12/2012): BCS to bring a proposal for a Legion Teaching &amp; Learning policy for discussion at the next meeting.</p> <p>(11/1/2013): ONGOING</p> <p>(8/2/2013): BCS has put forward a bid for funding; if successful, this might mean purchasing new equipment or re-purposing some of the older racks from Wolfson House, the latter solution being preferred if the Legion upgrade goes ahead. The exact number of nodes is to be discussed, but would give a sizeable cluster. ONGOING</p> <p>(15/03/2013): Proposal for service presented and approved. CLOSED</p>	BCS

107	Legion usage report for December.	<p>(11/1/2013): BCS to further investigate long wait times for users:</p> <ul style="list-style-type: none"> <li>• ucceak</li> <li>• uccaalo</li> </ul> <p>To add consortia to stats graph:</p> <ul style="list-style-type: none"> <li>• Built Environment</li> <li>• Medical Imaging</li> </ul> <p><b>(8/2/2013): The long wait times for uccaalo did not recur last month. uccaalo's jobs require 10Gb per process, requiring a particular cluster of nodes hence incurring longer wait times; BCS explained that if the full memory on a node is required, the entire CU needs to be cleared of other jobs first. Working across nodes is not possible if it is an MPI job (Infiniband has a higher contention CU to CU). This is assuming that he is not using the superqueue. Now that the 7 day queue has been removed, the maximum wait time will be 2 days. BCS expects to see this happening every month – the user is not complaining, but he will continue to pro-actively monitor. CLOSED. (Action reduced to adding Consortia to stats graph).</b></p> <p><b>BCS to amend stats graph so that zero-utilising Consortia are still shown. ONGOING</b></p> <p><b>CLOSED</b></p>	BCS
109	Role of Consortia leaders.	<p>(11/1/2013): CG to amalgamate CRAG Terms of Reference for Research Computing Consortia with the Legion Re-Application Document and present draft at next meeting.</p> <p><b>(8/2/2013): On the agenda for this meeting and was discussed. Amendments to be circulated as soon as possible and taken to the next meeting. ONGOING</b></p> <p><b>(15/3/2013): AGENDA ITEM 6 - CLOSED</b></p>	CG  JL
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-</p>	TJ and BCS



		<p>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><b>(15/03/2013): Mailman mailing lists added to consortia webpage. Mailman mailing list links not sent yet. ONGOING</b></p>	
111	Cfl EMERALD access request by Dimitry Kuzmin	<p>(8/2/2013): BCS to seek clarification concerning the scientific benefit to his work and typical job sizes; response to be circulated to the group for approval so that access can be granted before the next meeting.</p> <p>(8/2/2013): BS to contact Dimitri Kuzmin.</p> <p><b>(15/03/2013): Dmitri has not responded to contact. CLOSED</b></p>	BCS, Group
112	Consolidated metrics and information for reporting and users	<p>(8/2/2013): BCS to investigate anomalies in wait/CPU time, and also to produce a scrollable graph as some of the information is not visible.</p> <p><b>(8/2/2013): Scrollable graph completed.</b></p> <p><b>(15/03/2013): Anomalies were no longer present. CLOSED</b></p>	BCS
113	Change of NGS Consortium Leader	<p><b>(8/2/2013): TJ to look at the NGS mailman list and queues for VP (if necessary).</b></p> <p><b>(15/03/2013): ONGOING</b></p>	TJ
114	Legion scratch quotas	<p><b>(8/2/2013): BCS to implement increased Legion scratch quotas of double the existing default allocation (200Gb).</b></p> <p><b>(15/03/2013): ONGOING</b></p>	BCS
115	Amend stats graph so that zero-utilising Consortia are still shown.	<p><b>(15/3/2013): ONGOING</b></p>	BCS
116	Request for Additional Resources Form	<p><b>(15/3/2013): Amend form to state for how long users would require the additional resource</b></p>	BCS

117	Dr. Oleg L. Polyansky	<b>(15/3/2013): BS to liaise with Oleg L. Polyansky to establish limit for backfill pending further steer on Iridis access from Cfl Research Policy group.</b>	BCS
118	Dr Furio Cora	<b>(15/3/2013): BS to liaise with Furio Cora and suggest re-application based on use of Legion.</b>	BCS
119	Legion usage Report available <a href="http://feynman.rits-isd.ucl.ac.uk:8888">http://feynman.rits-isd.ucl.ac.uk:8888</a>	<b>(15/3/2013): BS to present graph of slowdown per user using both requested time and run time.</b>	BCS
120	Role of Consortia leaders and Account Application Policies and Processes	(Doc: RC_Consortia_and_account_application_process_v1.0.pdf) <b>(15/3/2013): BS re-format document to present to the Research Computing Working Group for approval.</b>	BCS
121	Michail Stamatakis UCL Chemistry, £30k funding for computing infrastructure	<b>(15/3/2013): BS to report back to CRAG on:</b> <ul style="list-style-type: none"> <li>• <b>Costings</b></li> <li>• <b>Integration</b></li> <li>• <b>Hosting</b></li> </ul>	BCS