

# UCL Computational Resource Allocation Group (CRAG)

## MEETING MINUTES

11<sup>th</sup> January 2013

### In Attendance:

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

### Apologies:

1. Prof Andrew Philips - *Epidemiology*
2. 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 14<sup>th</sup> December 2012**

The Group approved the Minutes of the December 14<sup>th</sup> 2012 meeting.

### **2. Update on status of current Actions**

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

Action 98 passed from CG to BCS to publicize the use of Legion@UCL and new scheduler information graph.

**3. Any requests for additional resources: (a) scratch quota requests, (b) any other request.**

Scratch Quota requests

The following requests were approved:

- Abil Aliev

Other requests:

The following request was discussed:

- Isaac Sugden

Regarding IS's application to access longer queues on Unity cluster. It was decided that this application be placed on hold, until the user provided further details as to why this job specifically needed to be completed on Unity. BCS would seek further details from user (see new Action 106).

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

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

The following applications were submitted and approved:

- None

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

BCS presented the Legion usage stats for December 2012 (it was noted that viewing them in IE is not recommended). NK expressed concern that the stats for the Built Environment and Medical Imaging showed long wait times in relation to run time; he had also noticed two particular users: uccaek and uccaalo whose jobs shared this characteristic. It was decided to add the consortia to per-user graphs to help further identify any correlation between users and consortia, and any repeated patterns. (see new Action 107).

It was decided that the CRAG monthly users report list users per consortium. (see new Action 108).

**6. Verbal update on Legion IV project (*Bruno Silva*)**

BCS gave a brief overview of the current status of Legion IV project and explained that the Legion IV project would end with the Lustre upgrade. CG explained there

was still budget associated with the Legion IV project and funds could possibly be directed to:

- An off-site scalable facility
- Improve existing Legion structure, i.e., improve the core switch for Legion.

## **7. Role of Consortia leaders - discussion (Chair)**

(Doc: RC\_Consortia\_ToTs\_v1.doc)

There was discussion regarding the need for consortia and what would be expected of consortia in relation to granting access to the Research Computing facilities. The the CRAG Terms of Reference for Research Computing Consortia were discussed and some preliminary modifications were made. CG explained the importance of reporting on research themes and how, for example, the EPSRC use research themed categories to allocate funding. It was also discussed whether to adhere to the existing structure of consortia or to create a new structure mapping it with UCL departments, as this would indicate where outreach is required. It was decided to amalgamate CRAG Terms of Reference for Research Computing Consortia with the Legion Re-Application Document and draft to be presented at the next meeting. (see new Action 109). It was decided that the Consortia mailing lists be advertised on the Research Computing webpage with an archive and search facility for consortium members. (see new Action 110).

## **8. AOB**

CG informed the CRAG that responsibility for meetings would now be handed over to BCS; CG would however still attend future meetings, where relevant.

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

- 8<sup>th</sup> February 2012 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. Any requests for additional resources: (a) scratch quota requests, (b) any other request.
4. Review of any Centre for Innovation (Cfi) access requests (Chair)
5. Reporting data/stats to be made available to users
6. Proposal for a policy allowing Teaching and Learning usage of Legion by UG students and taught PG students – discussion (Bruno Silva)

## 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><b>(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.</b></p>	BCS

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><b>(11/1/2013): Re-visit in 3 months.</b></p>	BCS
88	Consolidated metrics and information for reporting and users	<p>(24/05/2012): Merge items 70,79,83,84,85 and report at next meeting</p> <p><b>70 – (26/06/2012): ONGOING</b></p> <ol style="list-style-type: none"> <li>1) Total node utilisation – The amount of time a node is being utilised by any job vs. total up-time (simple binary measure per node) summed over all nodes.</li> <li>2) Used/requested memory – The ratio of effectively utilised memory to requested memory.</li> <li>3) Requested memory vs. Requested cores plotted over time – Simultaneous (side by side) plot of CPU and memory utilisation over time, possibly normalised for each component.</li> </ol> <p>Points 1 and 3 have now been implemented as database entries and the RC team is now working on presenting these as a report. Implementation of Point 2 requires further investigation.</p> <p><b>25/07/12 – Database entries are now ready and will be presented at next meeting.</b></p> <p><b>14/12/2012 – ONGOING; Item 2) cannot be implemented as SGE has a bug in the current version that reports actual memory utilised by a job inaccurately.</b></p> <p><b>83 - (26/06/2012)</b></p> <p>To provide report and analysis on the Serial vs. Parallel parts</p>	BCS

		<p>of the system for top 5 users.</p> <p>25/07/2012 - DONE</p> <p><b>It was agreed that a discussion needs to be had at a later date, in the form of a presentation at a future meeting, about what information is made available to users and what information is made available to the group – add item to agenda.</b></p> <p><b>85 - (26/06/2012): ONGOING</b></p> <p>Develop simple utilities of transparency for infiniband jobs, a simple map of what is running and where.</p> <p>25/07/2012 – ONGOING</p> <p><b>14/12/2012 - ONGOING</b></p> <p>(12/10/2012): Note that in future discussion in meetings re usage should be 'by exception', with all members to review reports using online tool (link has already been circulated) prior to each meeting. Next step will be to extract data from scheduler for live up to minute reporting.</p> <p>(9/11/2012): Academic members to review reports using online tool forward comments as to which information should NOT be available to users</p> <p><b>14/12/2012 - DONE</b></p> <p>(9/11/2012): BCS to anonymise data for viewing service metrics</p> <p>14/12/2012 – DONE</p> <p>(14/12/2012): Action reduced to Infiniband utilities.</p> <p><b>CLOSED</b></p>	
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><b>(14/12/2012): Gold accounting software installed. Client</b></p>	TJ

		<b>and lustre upgrade still pending. TJ to report back at next CRAG</b>	
94	User viewable 'live' service info	(9/11/2012): To add core service statistics and view of present user windows of opportunities and number of nodes available to website.  <b>(14/12/2012): Service statistics added to webpage. View of present window of opportunities is not necessary owing to the fact that backfill actually takes care of automatically filling scheduling gaps. – ONGOING</b>  <b>CLOSED</b>	BCS
96	Record of Cfl applications	(9/11/2012): It was agreed that a spreadsheet record of all Cfl requests, including reasons for rejection where appropriate, should be maintained.  <b>11/1/2013 – CG to maintain list of usage and report to CRAG every three months.</b>	CG
98	Publicising use of Legion@UCL and new scheduler information graph	(14/12/2012): CG to publicise the location on the website which advises how to acknowledge use of Legion (Legion@UCL) and the new graph of scheduler information, by emailing the research-computing list.  <b>(11/1/2013): BCS to inform users of availability of Legion stats and web location advising how to acknowledge use of Legion (Legion@UCL) and the new graph of scheduler information.</b>	BCS
99	Cfl GPU applications developer	<b>(14/12/2012): There is a dedicated GPU developer to help with applications on Emerald; Legion users to be advised of this.</b> <b>CLOSED</b>	CG
100	Legion jobs showing unusually long wait times	<b>(14/12/2012): BCS to contact the Legion user who had a job with a very large wait time to run-time ratio (uccaeak) to see if there was a problem; also to check December's statistics to see if the pattern was repeated for the SysBioMed consortium.</b> <b>CLOSED</b>	BCS
101	Cfl usage statistics	<b>(14/12/2012): CG to clarify whether STFC's share of Emerald is included in the under-utilisation figures.</b>  <b>(11/1/2013): ONGOING</b>	CG
102	Cfl Emerald applications	<b>(14/12/2012): BCS to circulate a list of the applications which are available on Emerald to the research-computing mailing list.</b>	BCS

		<b>CLOSED</b>	
103	Cfl IRIDIS job classes	(14/12/2012): CG to obtain data on IRIDIS job class distribution. <b>(11/1/2013): CG to present job clusters from Cfl</b>	CG
104	Legion 7-day Gaussian queue	<b>(14/12/2012): BCS to email the research-computing list advising Legion users that they need to reapply for access to the 7-day Gaussian queue before the end of February, if still required, or access will be revoked.</b> <b>CLOSED</b>	BCS
105	Legion Teaching and Learning use	<b>(14/12/2012): BCS to bring a proposal for a Legion Teaching &amp; Learning policy for discussion at the next meeting.</b> <b>(11/1/2013): ONGOING</b>	BCS
106	Request for Additional Resources	<b>(11/1/2013): BCS to liaise with Isaac Sugden for further information as to why he requires use of Unity.</b>	BCS
107	Legion usage report for December.	<b>(11/1/2013): BCS to further investigate long wait times for users:</b> <ul style="list-style-type: none"> <li>• ucceak</li> <li>• uccaalo</li> </ul> <b>To add consortia to stats graph:</b> <ul style="list-style-type: none"> <li>• Built Environment</li> <li>• Medical Imaging</li> </ul>	BCS
108	Consolidated metrics and information for reporting and users	<b>(11/1/2013): BCS to add to monthly report, the users listed per consortium.</b>	BCS
109	Role of Consortia leaders.	<b>(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.</b>	CG



110	Consortia mailing list	<b>(11/1/2013): BCS to advertise</b> Consortia mailing lists on the Research Computing webpage with an archive and search facility for consortium members	
-----	------------------------	---	--