

# UCL Computational Resource Allocation Group (CRAG)

## MEETING MINUTES

26th June 2012

### In Attendance:

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

### Apologies:

1. Prof Dario Alfe - *Thomas Young Centre (Materials Science)*
2. Dr Andrew Martin - *Bioinformatics and Computational Biology*
3. Dr Dugan Witherick - *Astrophysics and Remote Sensing*

*Note: Minutes below provide a high level summary of decisions taken and actions assigned by the Group. CRAG members are referred to the slide deck for the meeting for further information on matters presented.*

### **1. Approval of Minutes of last meeting on 26<sup>th</sup> June**

The Group approved the Minutes of the June 26<sup>th</sup> 2012 meeting. The following updates were noted:

#### **Item 5.3: Budget bids:**

- 1) Legion Component Replacement and/or Warranty Extensions
- 2) Legion IV
- 3) Off-site HPC Pilot: 1,2 & 3 to be compacted into one bid
- 4) e-Infrastructure South Cfl: To be headed by Prof David Price
- 5) RC Platforms Research and Development: Budget to be cut/reduced
- 6) Software Development Environment: Budget to be cut/reduced

- 7) HPC Service for Students and Teaching: *Budget to be cut/reduced*
- 8) Cloud and Virtualisation for HPC, Research and Development: *Budget to be cut*

There was no final outcome regarding Capital Budget.

**Item 7(b): Any other requests:**

The request received from *J Buckeridge* for allocation was granted.

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

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

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

NK emphasised that requests should not be replied to electronically, as the CRAG is to discuss and approve requests on a case by case basis.

(a) Scratch Quota requests

The following requests were approved:

*Sergey Yurchenko*  
*Sofia Moffoblou*

BS indicated that although users were registered they were not using their allocated accounts. CG suggested RC monitor users and that an agenda item be added to discuss the allocation of accounts.

(b) Any other requests

The following request was approved:

*David Jones*: priority access of 3 days (100 cores over 3 days) for completion and submission to run until the end of July 2012.

**4. Proposals for Research Computing KPIs**

CG presented the paper proposing RC KPIs which would measure the success of the group in delivering services. Both KPIs and SLAs would be approved by RCWG and signed off by RIISG. CG defined both KPI and SLA as follows:

KPI (Key Performance Indicator) – indicates whether RC is actually performing the services that it should be. KPIs could be tracked over time and measured.

SLA (Service Level Agreement) – forms the basis of agreement between ISD and the User Community.

The following KPIs were agreed:

1. Service Availability (for all compute services) - SLA
2. CPU/Node Availability (for all cluster services) - SLA
3. Mean Time Between Service Interrupt
6. Number of user Departments to include:
  - Number of active users
  - Number of active departments
  - Number of CPUs used

The following would be incorporated into General Reports:

4. User Satisfaction Rating (for each Research Computing platform service)
5. Number of peer reviewed papers published
7. Service utilisation
8. Wait time/slow down (agreed not to aggregate)
9. Total associated research funding (breakdown by funding body)

VP suggested a report on the responsiveness of Legion, CG would consider this accordingly. It was agreed that this document be re-submitted at a later stage.

## **5. Presentation of new reporting framework**

CG presented the new reporting framework, taking questions from group members.

## **6. May service report**

BS presented the Legion, Unity and Condor service reports for the month of May.

## **7. Legion IV update**

CG gave a verbal report on the progress of the Legion IV project.

## **8. AOB**

CG explained that the RITs webpage would be undergoing a major overhaul incorporated within the re-design of the internal ISD website. CG required volunteer Researchers to sit on a working group and provide input from a Researcher's perspective on content, presentation and navigation.

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

The next meeting of the CRAG are scheduled for:

- 25<sup>th</sup> July 2012, 1pm – 3pm, Room 103, ROBERTS BUILDING, Torrington Place, London, WC1E 7JE

It was agreed that hard copies of papers would not be provided and that all papers should be provided electronically in advance so that members would print their own paper copies should they wish.

## LIST OF CURRENT ACTIONS

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

	Actions	Status	Owner
10	HEP and RC to decouple the running of the middleware from Legion and run this from HEP Central.	<p><b>ON HOLD</b> – meetings are scheduled between IS-RC and HEP</p> <p>(13/05/2011): Meeting to be scheduled.</p> <p>(8/07/2011): There has been progress on the technical side and RC is waiting for feedback from relevant users.</p> <p>(30/09/2011): Meeting to be scheduled.</p> <p>(28/10/2011): IN PROGRESS, HEP to resource manpower.</p> <p>(27/01/2012): IN PROGRESS with HEP</p> <p>(24/02/2012): IN PROGRESS with HEP</p> <p><b>(26/06/2012): ON HOLD (lack of manpower)</b></p>	CG/BW
70	10/06/2011 Investigate potential memory usage statistics and metrics (e.g. used vs. requested) to measure how efficiently users are requesting resources on Legion and Unity.	<p>(8/07/2011) This item has been merged with item 68 Look into CPU utilisation metrics and statistics in conjunction with this. This will both allow reporting to the CRAG and provide users with useful information for correct estimation.</p> <p>(30/09/2011): Metrics have been created. Technical implementation will start once these metrics are reviewed. Re-send metrics for evaluation in next meeting.</p> <p>(28/10/2011): RC to provide proposals and circulate to CRAG mailing lists - Done. Implementation has started.</p> <p>Statistics would be gathered in the following manner: Weekly sampling at least (higher time resolution may be needed for data to be meaningful). Separate statistics for each node class. Statistics done over all users (monitoring on a per user basis may be done for support reasons).</p> <p>(27/01/2012): IN PROGRESS. The RC Sys Admin team is working on obtaining the relevant data and writing this to a database.</p> <p>(24/02/2012): IN PROGRESS. Agreed that the following statistics were previously proposed to the CRAG and approved.</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</li> </ol>	BCS

		<p>time – Simultaneous (side by side) plot of CPU and memory utilisation over time, possibly normalised for each component.</p> <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>(30/03/2012): IN PROGRESS. BCS working on framework for storage of system data and generation of ad-hoc and more flexible reporting. To be presented at next meeting.</p> <p><b>(24/05/2012): IN PROGRESS. Consolidated in item 86.</b></p>	
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><b>(26/06/2012): IN PROGRESS. Shared storage areas have been mounted on Condor machines, applications require placement.</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><b>(26/06/2012): Re-assigned to BCS</b></p>	DG/BCS

77	Uncertainty Quantification Pipeline for Climate Models project	<p>(27/01/2012): To set up appropriate priority queue for project (by June 2012)</p> <p>(24/02/2012): PENDING</p> <p>(25/05/2012): RC to contact Serge regarding his pipeline and install missing components</p> <p><b>(26/06/2012): Pipeline completed and priority queue established. CLOSED</b></p>	BCS
79	Legion usage 17-64 cores	<p>(27/01/2012): Stats showing usage by job class 17 - 64 cores to be broken down to show usage by 17 – 32 and 33 – 64.</p> <p>(24/02/2012): IN PROGRESS. Has been added to work planning log.</p> <p>(30/03/2012): Stats showing usage by job class to be broken down to show usage by 79 – 65, 66 - 128 and 129 – 156.</p> <p>(24/05/2012): IN PROGRESS</p> <p><b>(26/06/2012): – consolidated in item 88</b></p>	TJ
82	Condor report – CPU hours used?	<p>(27/01/2012): To look into whether it is possible to get a report on total CPU hours used by Condor jobs, as part of the regular Condor report.</p> <p>(24/02/2012): IN PROGRESS. Has been added to work planning log.</p> <p>(30/03/2012): NO FURTHER PROGRESS</p> <p>(24/05/2012): NO FURTHER PROGRESS</p> <p><b>(26/06/2012): – consolidated in item 88</b></p>	TJ
83	Options for reducing variations in usage	<p>(24/02/2012): To provide report and analysis on the Serial vs. Parallel parts of the system for top 5 users.</p> <p>(30/03/2012): NO FURTHER PROGRESS</p> <p>(24/05/2012): NO FURTHER PROGRESS</p> <p><b>(26/06/2012): – consolidated in item 88</b></p>	CG
84	Queue Conveyor	<p>(30/03/2012): Devise queue reporting diagram/list to determine whether jobs are running on:</p> <p>Node types W, X, Y, Z and GPU</p> <p>(24/05/2012): NO FURTHER PROGRESS</p> <p><b>(26/06/2012): – consolidated in item 88</b></p>	TJ/WH

85	Utilities of Transparency	<p>Develop simple utilities of transparency for infiniband jobs, a simple map of what is running and where?</p> <p>(24/05/2012): NO FURTHER PROGRESS</p> <p><b>(26/06/2012): – consolidated in item 88</b></p>	BS
86	Draft KPI's	<p>(24/4/2012) – Action opened: NK to speak to RCWG Chair whether CRAG or RCWG should draft 'strawman' KPIs.</p> <p><b>(26/06/2012): CLOSED</b></p>	NK
87	Reports on key publication reports to be used for promoting Legion	<p>(24/4/2012) – Action opened: NK to email Consortium Leaders asking them to solicit lists of key publications that have been enabled via use of Legion, for use on RC website and to encourage more requests for 'additional resources.'</p> <p><b>(26/6/2012): Link publications on RC website and provide list of Consortia publications for standard Legion use.</b></p>	NK
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>79 - (26/06/2012): COMPLETED</b></p> <p>(27/01/2012): Stats showing usage by job class 17 - 64 cores to be broken down to show usage by 17 – 32 and 33 – 64.</p> <p>(24/02/2012): IN PROGRESS. Has been added to work planning log.</p> <p>(30/03/2012): Stats showing usage by job class to be broken down to show usage by 79 – 65, 66 - 128 and 129 – 156</p> <p><b>83 - (26/06/2012): PENDING - can be created from 84</b></p> <p>(24/02/2012): To provide report and analysis on the Serial vs. Parallel parts of the system for top 5 users.</p>	BCS



		<p><b>84 - (26/06/2012): IMPLEMENTED</b></p> <p>(30/03/2012): Devise queue reporting diagram/list to determine whether jobs are running on: Node types W, X, Y, Z and GPU</p> <p><b>85 - (26/06/2012): PENDING – Information dependant</b></p> <p>Develop simple utilities of transparency for infiniband jobs, a simple map of what is running and where.</p>	
89	Researcher volunteers for RITs Website Working Group	<b>(26/06/2012): VP and DW to source Researcher volunteers to provide input on new RITs website.</b>	CG