

UCL Computational Resource Allocation Group (CRAG)

MEETING MINUTES

25th July 2012

In Attendance:

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

Apologies:

1. Clare Gryce - Head of Research Computing and Facilitating Services, ISD
2. Thomas Jones - Research Platforms Team Leader (Infrastructure Lead), ISD

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 26th June

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

Item 3

It was agreed that the allocation of accounts be added to the agenda for the next meeting.

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.

(a) Scratch Quota requests

The following requests were approved:

Abil Aliev

(b) Any other requests

None

4. Legion Service Stats for June

100% service availability was reported, with 85% computing core availability, of which 98% was utilized. There were a total of 127 active users over this period. The lower availability of computing cores had to do with the fact that two CUs were taken out at the KLB for datacenter maintenance reasons.

5. Legion Usage Stats for June (using new reporting framework)

The reports were presented it was noted that one user consumed around 23% of all cluster resources. Concerns were raised around the backfill of serial jobs on Legion, particularly around the maximum wall clock time allowed for this. RC was asked to provide information to the CRAG regarding the present configuration of backfill and job mapping to compute nodes. RC should also provide information about job characteristics of top user.

BCS will provide web reports in advance of the next meeting.

6. Review of regular Unity and Condor Reports

BCS explained that Unity stats are not available this month owing to technical problems with the machine which are currently being resolved.

It was agreed at the next meeting that July and August statistics will be seen in conjunction.

7. Date of next meeting

TBC – second week of September.

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>ON HOLD – 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>(26/06/2012): ON HOLD (lack of manpower)</p>	CG/BW
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 & 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. ?? is going to start speaking to users. Going to make Condor useable from Legion itself, and should be available by next year. Will be a hybrid system.</p>	BCS
76	Peter Harrison request for additional Unity resource	(27/01/2012): To follow up with requestors regarding additional information about check-pointing and problem decomposition.	DG/BCS

		<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>	
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>(26/6/2012): Link publications on RC website and provide list of Consortia publications for standard Legion use.</p> <p>(25/07/2012): Now have a complete list from main users of Legion. To be discussed with CG when she returns from leave.</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>70 – (26/06/2012): ONGOING</p> <ol style="list-style-type: none"> 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. 2) Used/requested memory – The ratio of effectively utilised memory to requested memory. 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. <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>25/07/12 – Database entries are now ready and will be presented at next meeting.</p> <p>83 - (26/06/2012)</p> <p>To provide report and analysis on the Serial vs. Parallel parts of the system for top 5 users.</p> <p>25/07/2012 - DONE</p> <p>It was agreed that a discussion needs to be had at a later date, in the form of a presentation at a future meeting,</p>	BCS

		<p>about what information is made available to users and what information is made available to the group – add item to agenda.</p> <p>85 - (26/06/2012):</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>	
89	Researcher volunteers for RITs Website Working Group	<p>(26/06/2012): VP and DW to source Researcher volunteers to provide input on new RITs website.</p> <p>25/07/12 – Volunteers have been sourced.</p>	CG
90	Provide information to CRAG about backfill policy and job characteristics of top user.	<p>(25/07/2012) – Action requested</p> <p>this was requested to clarify whether the job distribution observed on July was due to backfill settings and user job profile.</p>	BCS