

UCL Computational Resource Allocation Group (CRAG)

MEETING MINUTES

30th March 2012

In Attendance:

- John Brodholt [JB] – (Chair) Earth Sciences
- Clare Gryce [CG] – ISD, Research Computing
- Thomas Jones [TJ] – ISD, Research Computing
- Nik Kaltsoyannis [NK] – Chemistry
- Owain Kenway [OK] – ISD, Research Computing
- Andrew Martin [AM] – Structural and Molecular Biology
- Bruno Silva [BCS] – ISD, Research Computing

Apologies:

- Eric Fraga [EF] – Chemical Engineering
- Ian Eames [IE] – Mechanical Engineering
- David Gregory [DG] – ISD, Computational Biology Research Facilitator
- Ben Waugh [BW] – Physics and Astronomy

1. Minutes of 24th February 2012 meeting

The Group approved the Minutes of the February 24th 2012 meeting.

2. Review of current actions

Action list was reviewed and updated with relevant comments
(See Actions list below)

3. Presentation by ISD on current fairshare implementation (see Action 75) – Thomas Jones/William Hay

WH presented the current fair-share function. It was suggested that a queue reporting diagram or text output be devised to view what jobs and how many are running on all types of node available. **AGREED** (see Action 84 below)

4. Brief verbal report on current projects – Clare Gryce

CG reported on the Legion IV project as approved by the Research Computing Working Group and the RIISG. The Project comprises the following phases:

- Replace current high-performance (Lustre) file system hardware – technical design is in progress and a Project Board is being convened, to include senior academic representation from both physical sciences and biosciences.
- An expansion of the on-site Legion compute service will follow. A proportion of the budget is being allocated to provide 'seed funding' for an offsite HPC facility in 2013.

5. Brief verbal report from RC Forum meeting – John Brodholt

The RC Forum meetings started again with reporting from RC team and the CRAG regarding governance changes, progress of future projects and future ones. A representative from the Miracle Astrophysics consortium stated concerns about prioritisation of jobs on Legion as this is a paying consortium. It was clarified that the agreement guaranteed share, and not priority access.

6. Matters arising from Feb reports – Chair/All

It was noted that the number of current users had increased to 130. It was proposed that node utilisation data be sought in order to compile proposal in order to purchase a new machine.

Date of next meetings

The next meetings of the group are scheduled for:

- 24th May 2012, 1pm - 3pm, Edgeworth Room (Drayton B16), Drayton house, 30 Gordon St, London, WC1H1H 0AX.
- 26th June 2012, 1pm – 3pm, Wilkins JBR Meeting Room, Wilkins Building Gower St, London, WC1E 6BT,
- 27th July 2012, 1pm - 3pm, South Wing G14 Committee Room, UCL Main Campus.

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>IN PROGRESS – 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>	CG/BW
56	Scratch space policy to be enforced appropriately, users informed.	<p>(11/03/11): Policies to be enforced after Lustre upgrade as quotas may be used to retrieve usage information quickly.</p> <p>(8/07/2011): Scratch will be deleted entirely by the end of August, leaving some time for users to clear away their data. Policies will be enforced from then on.</p> <p>(30/09/2011): Scratch now deleted, new policy to be agreed – technical development proceeding</p> <p>(28/10/2011): Implementation of agreed policy has started.</p> <p>(25/11/11): Inform users that jobs will be stopped if over quota, programme for this to be written and implemented</p> <p>(27/01/2012): IN PROGRESS. Scripts have been completed and are under test prior to informing users of new policy.</p> <p>(24/02/2012): IN PROGRESS. Scripts to monitor usage have been written and tested. Next steps are to roll out usage and provide information and advice to users on website include exceeding quota allowance - users exceeding allocation will be given 14 days to reduce quota.</p> <p>(30/3/2012): IN PROGRESS. An Email has been sent to users explaining the new quota regime and advising how they should submit requests for extra quota. RC Team authorised to use discretion re User quota requests and escalate to CRAG where appropriate.</p>	BCS/TJ

70	<p>10/06/2011</p> <p>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>	<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"> 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>(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>	BCS
----	---	---	-----

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>	BCS
74	25.11.11 GAMS licence	<p>(25.11.11) Discuss GAMS license for Legion</p> <p>(27/01/2012): IN PROGRESS. License model is perpetual, fixed number of users managed via group (not license manager). Chemical Engineering to progress and provide further feedback on requirements.</p> <p>(24/02/2012): IN PROGRESS. No update available</p> <p>(30/03/2012): IN PROGRESS. BCS to contact EF to establish who is now responsible for this.</p>	BCS
75	Presentation to CRAG on Fairshare implementation	<p>(27/01/2012): RC Team to give presentation to next CRAG meeting explaining current Legion Fairshare implementation.</p> <p>(24/02/2012): PENDING. Postponed to March meeting.</p> <p>(30/03/2012): COMPLETED</p>	CG/TJ
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 REVISTED at a later date</p>	DG
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>	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>	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>	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>	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>	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>	BS