

UCL Computational Resource Allocation Group (CRAG)

MEETING MINUTES

24th February 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
- Ben Waugh [BW] – Physics and Astronomy

Apologies:

- Eric Fraga [EF] – Chemical Engineering
- Ian Eames [IE] – Mechanical Engineering
- Bruno Silva (acting CRAG Secretary) [BCS] – ISD, Research Computing
- David Gregory [DG] – ISD, Computational Biology Research Facilitator

1. Minutes of 27th January 2012 meeting

The Group approved the Minutes of the January 27th 2012 meeting.

2. Review of current actions

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

3. Options for reducing variations in usage by user.

The group discussed options for reducing variations in usage by individual users. It was agreed that while the current high levels of usage observed for some projects/individuals were not an intended outcome of the current policy, policies that would forcibly moderate such anomalies were undesirable (e.g. more 'heavyweight' peer review or fixed allocations) as they would be hard to implement and may reduce overall system utilisation.

It was agreed that it was more important to understand what these users are using Legion for i.e. what type/quality of research is being performed.

It was suggested that CG should develop a process and associated policy for eliciting information about research outcomes from 'top' users and seek approval for this from the Research Computing Platforms Working Group.

JB requested that the following information be sought for the top current Legion users:

- 1) Status of users i.e. undergraduate or postgraduate
- 2) What they were researching
- 3) Publications
- 4) Funding

The top five users were as follows:

1st Jan – 20th July 2011

ucbpko
ucfboao
uccaijs
uccajkl
ucaaea

11th July 2011 – 20th February 2012

ucbpmko
zcqsc29
rmjtaph
ucbtcds
uccaea

Date of next meeting

The next meeting of the group is scheduled for the 30th March 2012 (1-2pm).

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>	BCS/TJ

67	<p>Provide users with the means to estimate memory usage by their jobs more accurately (on Legion, but also applicable to Unity):</p> <ul style="list-style-type: none"> • Provide User Test Nodes • Offer scripts to users in which allow them to profile their jobs • Inform users regarding availability of these scripts and useful techniques (once user test nodes and scripts are available). 	<p>(8/03/2011) User test nodes are not yet available but this work has been scheduled.</p> <p>(30/09/2011) User test nodes when busy will return message to command line via qrsh, informing the user to “try later”, nodes are used interactively. Jobs can be given instant access if nodes are available. Users are permitted 2 hour sessions via scheduler which can be double via CPU sets. User test nodes are still under development.</p> <p>(28/10/2011): Testing still pending by selected users, to be completed by December 2011.</p> <p>(25/11/11) depends on implementation of User Test Nodes.</p> <p>(27/01/2012): IN PROGRESS. User test nodes have been implemented. This work to be added to monthly programme of ‘small works’ for RC Admin team.</p> <p>(24/02/2012): COMPLETED</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>	BCS

		<p>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.</p> <p>2) Used/requested memory – The ratio of effectively utilised memory to requested memory.</p> <p>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> <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>	
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>	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>	EF
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>	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 permit 72hr wall-clock usage over 10 days on Unity</p>	DG

77	Uncertainty Quantification Pipeline for Climate Models project	(27/01/2012): To set up appropriate priority queue for project (by June 2012) (24/02/2012): PENDING	BCS
78	Epidemiology use of Legion Nov and Dec 2011	(27/01/2012): To provide further information about number of projects using Legion during period. (27.2.2012): COMPLETED. Stats provided reviewed projects as follows: Nov-4 projects, Dec-3 projects, Jan-5 projects. Epidemiology major user.	CG
79	Legion usage 17-64 cores	(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. (24/02/2012): IN PROGRESS. Has been added to work planning log.	CG
80	Email re priority of short jobs	(27/01/2012): To email users reminder that shorter jobs are given higher priority. (24.2.2012): Completed	BCS
81	Usage per user graphs – fairshare changes	(27/01/2012): To produce ‘usage per user’ graphs be produced to show to effect of the changes to the fairshare implementation made last year. (24/02/2012): COMPLETED. Report presented at meeting.	OK
82	Condor report – CPU hours used?	(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. (24/02/2012): IN PROGRESS. Has been added to work planning log.	CG
83	Options for reducing variations in usage	(24/02/2012): To provide report and analysis on the Serial vs. Parallel parts of the system for top 5 users.	CG