

UCL Computational Resource Allocation Group (CRAG)

MEETING MINUTES

24th May 2012

In Attendance:

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

Apologies:

1. Prof Dario Alfe - *Thomas Young Centre (Materials Science)*
2. Dr Nicholas Achilleos - *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 30th March

The Group approved the Minutes of the March 30th 2012 meeting.

2. Recap on overall governance structure for Research Computing

CG provided a verbal summary on the current governance arrangements, for the benefit of new members.

3. Review of revised membership of CRAG

NK noted the revisions to the membership of the CRAG, explaining that new members had been selected to represent the research domains that make greatest use of the Legion service.

4. Review of revised CRAG ToRs as proposed by Research Computing Working Group (RCWG)

NK noted the revisions made to the CRAG Terms of Reference by the Research Computing Working Group, to which the CRAG reports. The new ToR's are as follows:

- i. To monitor the availability and delivery of computational services provided by ISD Research Computing according to agreed Key Performance Indicators (KPIs)
- ii. To monitor the allocation and usage of Research Computing computational and associated data storage services, according to the above policy
- iii. To provide periodic reports on the above and any other matters arising to the Research Computing Working Group, or the RIISG as required

It was noted that formal KPIs are pending development by the RCWG and/or RIISG.

It was noted that data storage services at (2) only include scratch/home and any 'staging' areas coupled to Research Computing compute platforms; i.e. they do not include Research Data Services

It was agreed that NK would speak to the Chair of the RCWG to ask whether the CRAG or the RCWG should draft 'strawman' KPIs for consideration (**See ACTION 86**)

5. Report from Research Computing Working Group

CG gave a verbal report from the last meeting of the RCWG on 27th April, with contributions from NK, and taking questions from group members.

The following points were noted:

- i. The RCWG has agreed a new framework for resource allocation on Legion that will allow research projects/groups to purchase dedicated resources under a 'leased' arrangement, and/or access to a high-priority job queue also on a chargeable basis, in addition to the currently available no-cost 'fairshare' arrangement. It was noted that the former option would not realistically be available until ISD is able to offer a readily scalable HPC service, which in itself is dependent on a project – currently in hand – to establish an offsite Datacentre capability for UCL. This would ensure that the implementation does not negatively impact the majority of researchers who use the service under 'fairshare'. It was noted that some groups may wish to purchase additional filesystem space, as well as or instead of compute and that this facility should be

provided under the scheme. CG reported that the next steps are for the ISD team to consider implications for implementation, and report back to the RCWG.

- ii. The RCWG had noted developments re the Centre for Innovation facilities and development, previously presented to the CRAG, and noted that oversight of the resource allocation policy for the Cfl facilities might be delegated to the RCWG in the medium term. CG noted that it ideally this should be integrated with the local policies overseen by the CRAG.
- iii. The RCWG received a brief from CG on the budget bids that had been prepared and submitted by the ISD Research Computing team, for consideration and ultimate approval by the RIISG. The RCWG had been asked to prioritise the bids and had agreed the following ranking:
 - 1) Legion Component Replacement and/or Warranty Extensions
 - 2) Legion IV
 - 3) Off-site HPC Pilot
 - 4) e-Infrastructure South Cfl
 - 5) RC Platforms Research and Development
 - 6) Software Development Environment
 - 7) HPC Service for Students and Teaching
 - 8) Cloud and Virtualisation for HPC, Research and Development

6. Update on Current Actions

The list of current Actions (below) was updated, and new Actions arising from the RCWG and charged to the CRAG, were added.

7. Requests for Additional Resources

(a) Scratch Quota requests

BCS outlined the requests for extension of Scratch quota that had been received subsequent to the implementation of quotas on /scratch:

Pascale Gerbauld: three-month grace period granted from the date of the meeting (24th of May) and 5TB base quota henceforth.

Francesco Lescai: 21 day grace period and 2TB base quota.

Sergey Yurchenko, Clara Silva, Oleg Polyanski: creation of shared project area with 20TB of shared space between these three users, all under the *Exomol* project.

The group agreed that all requests should be granted by that the users/projects concerned should be advised that the extensions may be subject to revision of withdrawal by the CRAG at any time in the future.

(b) Any other requests

The request received from J Buckeridge was reviewed by the group; it was agreed that further information to justify the lack of checkpointing that was used as the basis for the extension, was required.

8. March and April Service Report

CG presented the Legion, Unity and Condor service reports for the months of March and April.

9. Legion IV Project Update

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

10. AOB

The group discussed presenting to researchers requests for additional resources in order to enable innovative research. The discussion led to the idea of advertising what high quality research has been enabled by Research Computing platforms, and particularly Legion.

11. Next Meeting Date and Agenda

The next meeting of the CRAG are scheduled for:

- 26th June 2012, 1pm – 3pm, Wilkins JBR Meeting Room
- 27th July 2012, 1pm - 3pm, South Wing G14 Committee Room

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.

Clare Gryce
ISD
1st June 2012

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> <p>(24/05/2012): ON HOLD (lack of manpower)</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

		(24/05/2012) Scratch policy implemented and enforced. CLOSED	
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"> 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> <p>(24/05/2012): IN PROGRESS. Consolidated in item 86.</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> <p>(24/05/2012): 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> <p>(24/05/2012): Item is now line of business and is now closed as a CRAG item. CLOSED</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 REVISTED at a later date</p> <p>(24/05/2012): DG to contact the user regarding progress of his work.</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> <p>(25/05/2012): RC to contact Serge regarding his pipeline and install missing components.</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 – consolidated in item 88</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>	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 – consolidated in item 88</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 – consolidated in item 88</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 – consolidated in item 88</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>	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</p>	NK

		resources.	
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 –</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>79 -</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>83 -</p> <p>(24/02/2012): To provide report and analysis on the Serial vs. Parallel parts of the system for top 5 users.</p> <p>84 -</p> <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>85 -</p> <p>Develop simple utilities of transparency for infiniband jobs, a simple map of what is running and where?</p>	BCS