

CSAR Service

Consolidated Management Report

1st Quarter 2006

Management Summary

This is the consolidated Management report for the first quarter 2006 of the CSAR HPC facility for UK Academia and Industry, which enables World-Class research and development.

The number of users stands at a total of 441 to date.

The workload on both the Origin 3000 Green and the Altix 3700 Newton has been fairly evenly spread across the mid- to high-end range of PEs during the first quarter of this year.

CSAR has been granted an 18-month extension of service contract until June 30th 2006. With this extension CfS implemented a further technology refresh introducing a 256 processor Itanium-2 (Madison) based SGI Altix, now enhanced to 512 processors with 1 Terabyte of memory.

CfS remains active in the UK Grid Forum.

Introduction

This Management Report includes a section for each of the main service functions:

1. Service Quality
2. HPC Services
3. Science Applications Support Services
4. Training & Education Services
5. User Registration & New User Services
6. Value-Added Services

Each section includes a status report for the period, including notable achievements and problems, also noteworthy items for the next period.

1 Service Quality

This section covers overall Customer Performance Assessment Ratings (CPARS), HPC System availability and usage, Service Quality Tokens and other information concerning issues, progress and plans for the CSAR Service.

1.1 CPARS

Table 1 gives the measure by which the quality of the CSAR Service is judged. It identifies the metrics and performance targets, with colour coding so that different levels of achievement against targets can be readily identified. Unsatisfactory actual performance will trigger corrective action.

CSAR Service - Service Quality Report - Performance Targets

Service Quality Measure	Performance Targets					
	White	Blue	Green	Yellow	Orange	Red
HPC Services Availability						
Availability in Core Time (% of time)	> 99.9%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Availability out of Core Time (% of time)	> 99.8%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Number of Failures in month	0	1	2 to 3	4	5	> 5
Mean Time between failures in 52 week rolling period (hours)	>750	>500	>300	>200	>150	otherwise
Fujitsu Service Availability						
Availability in Core Time (% of time)	> 99.9%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Availability out of Core Time (% of time)	> 99.8%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Help Desk						
Non In-depth Queries - Max Time to resolve 50% of all queries	< 1/4	< 1/2	< 1	< 2	< 4	4 or more
Non In-depth Queries - Max Time to resolve 95% of all queries	< 1/2	< 1	< 2	< 3	< 5	5 or more
Administrative Queries - Max Time to resolve 95% of all queries	< 1/2	< 1	< 2	< 3	< 5	5 or more
Help Desk Telephone - % of calls answered within 2 minutes	>98%	> 95%	> 90%	> 85%	> 80%	80% or less
Others						
Normal Media Exchange Requests - average response time	< 1/2	< 1	< 2	< 3	< 5	5 or more
New User Registration Time (working days)	< 1/2	< 1	< 2	< 3	< 4	otherwise
Management Report Delivery Times (working days)	< 1	< 5	< 10	< 12	< 15	otherwise
System Maintenance - no. of sessions taken per system in the month	0	1	2	3	4	otherwise

Table 1

Table 2 gives actual performance information for the period. Overall, the CPARS Performance Achievement for the 1st quarter 2006 was satisfactory (see Table 3), i.e. Green measured against the CPARS performance targets.

CSAR Service - Service Quality Report - Actual Performance Achievement

Service Quality Measure	2005/6											
	April	May	June	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	Mar
HPC Services Availability												
Availability in Core Time (% of time)	96.00%	99.50%	97.13%	98.75%	99.04%	97.85%	98.50%	98.33%	96.99%	99.25%	94.82%	98.53%
Availability out of Core Time (% of time)	99.90%	99.54%	99.22%	99.45%	99.4%	96.82%	99.63%	99.12%	99.3%	99.40%	99.79%	98.46%
Number of Failures in month	4	1	2	4	5	6	1	4	4	2	2	7
Mean Time between failures in 52 week rolling period (hours)	227	241	257	786	559	406	394	330	286	271	257	212
Help Desk												
Non In-depth Queries - Max Time to resolve 50% of all queries	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.5	<0.25	<0.5	<0.25	<0.25
Non In-depth Queries - Max Time to resolve 95% of all queries	<2	5>	<1	<2	<1	<2	<3	<5	<2	<2	<3	<3
Administrative Queries - Max Time to resolve 95% of all queries	<1	<0.5	<2	5>	<0.5	<0.5	<0.5	<2	<1	<3	<5	<2
Help Desk Telephone - % of calls answered within 2 minutes	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Others												
Normal Media Exchange Requests - average response time	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5
New User Registration Time (working days)	0	0	0	0	0	0	0	0	0	0	0	0
Management Report Delivery Times (working days)	10	10	10	10	10	10	10	10	10	10	10	10
System Maintenance - no. of sessions taken per system in the month	2	2	2	2	2	2	2	2	2	2	2	2

Table 2

Notes:

- HPC Services Availability has been calculated using the following formula, based on the relative NPB performance of Fermat, Green and Newton at installation:

$$[\text{Fermat availability} \times 40 / (40+233+343)] + [\text{Green availability} \times 233 / (40+233+343)] + [\text{Newton availability} \times 343 / (40+233+343)]$$
- Mean Time Between Failures for Service Credits is formally calculated from a rolling 12-month period.

Table 3 gives Service Credit values for each month to date. These are accounted on a quarterly basis, formally from the Go-Live Date. The values are calculated according to agreed Service Credit Ratings and Weightings.


CSAR Service - Service Quality Report - Service Credits

Service Quality Measure	2005/6											
	April	May	June	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	Mar
HPC Services Availability												
Availability in Core Time (% of time)	0.078	0	0.078	0.039	0.039	0.078	-0.039	0.078	0.078	0	0.115	0.039
Availability out of Core Time (% of time)	-0.047	-0.039	0	0	0	0.039	-0.039	0.000	0	0	-0.039	0.039
Number of Failures in month	0.008	-0.008	0	0.008	0.0156	0.023	-0.008	0.008	0.008	0	0	0.023
Mean Time between failures in 52 week rolling period (hours)	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002
Help Desk												
Non In-depth Queries - Max Time to resolve 50% of all queries	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.016	-0.019	-0.016	-0.019	-0.019
Non In-depth Queries - Max Time to resolve 95% of all queries	0	0.046	-0.016	0	-0.016	0	0.016	0.031	0	0	0.016	0.016
Administrative Queries - Max Time to resolve 95% of all queries	-0.016	-0.019	0	0.046	-0.019	-0.019	-0.019	0	-0.016	0.016	0.031	0
Help Desk Telephone - % of calls answered within 2 minutes	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004
Others												
Normal Media Exchange Requests - average response time	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002
New User Registration Time (working days)	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019
Management Report Delivery Times (working days)	0	0	0	0	0	0	0	0	0	0	0	0
System Maintenance - no. of sessions taken per system in the month	0	0	0	0	0	0	0	0	0	0	0	0
Monthly Total & overall Service Quality Rating for each period:	-0.01	-0.03	0.01	0.02	-0.01	0.04	-0.03	0.04	0.01	-0.01	0.04	0.04
Quarterly Service Credits:	-0.03			0.05			0.03			0.06		

Table 3


1.2 No. of Registered Users

The current position at the end of the quarter is that there are 441 registered users of the CSAR Service.



1.3 Service Quality Tokens

The graph below illustrates the monthly usage trend of Service Quality Tokens:



Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

At the end of the quarter no gold stars or black marks were allocated to the service.

2 HPC Services Usage

Usage information is given in tabular form, and in graphical format. The system usage information covers:


- CPU usage
- User Disk allocation
- HSM/tape usage

This is illustrated in a number of graphs including;


- a) SMP (Altix/Origin) Usage by month, showing usage each month of CPU (GFLOP-Years as per NPB), split by Research Council and by system. Overlaid horizontal lines show the overall Capacities.
- b) SMP (Origin) Usage by month, showing usage each month in CPU Hours, split by Research Council and giving the equivalent GFLOP-Years as per NPB. The Baseline and overall Capacity are shown by overlaid horizontal lines.
- c) Medium Performance Disk (Origin) allocated for User Data by month, showing the allocated space each month in GBytes, split by Research Council. The Baseline Capacity (1.5 Terabytes) is shown by an overlaid horizontal line.
- d) HSM/Tape Usage by month, showing the volumes held each in GBytes, split by Research Council. The Baseline Capacity (16 Terabytes) available will be shown by an overlaid horizontal line.


2.1 Service Usage Charts

The graphs below show recent monthly CPU, disk and HSM allocations and usage.

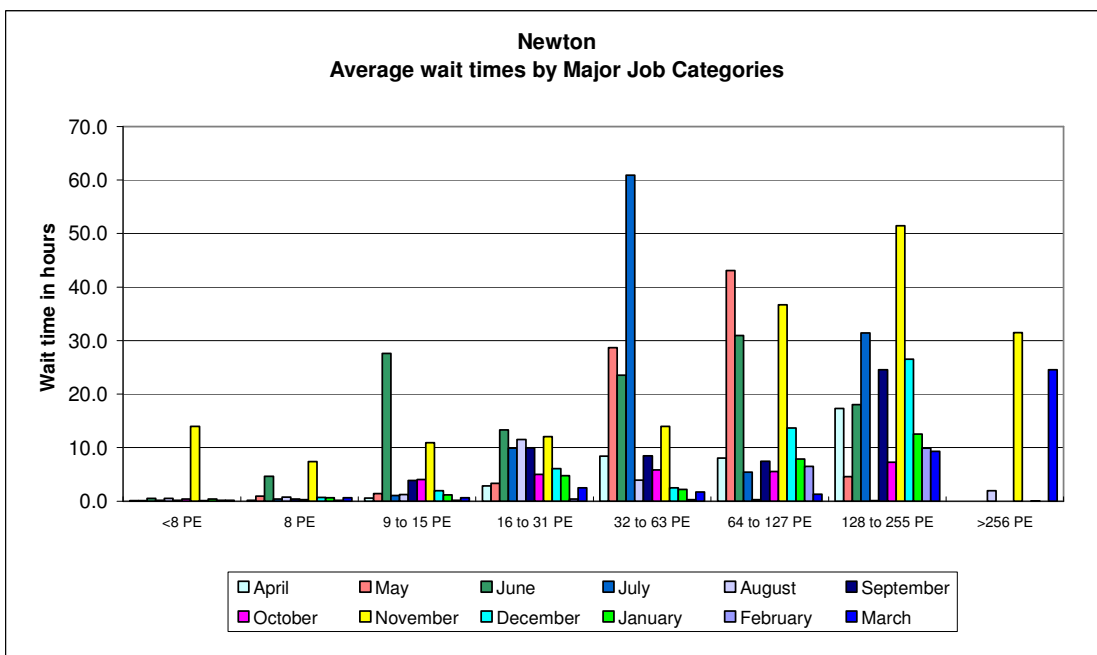


Usage of the SGI Altix 3700 system Newton is shown by Research Council in the above graph.







Usage of the two batch SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service in the preceding two charts.




The above chart, and the one below, shows the wait time trend in hours on the Altix 3700 Newton.




The following two charts show average wait times in hours for the quarter on the Origin 3000 Green.







The next series of four charts show the percentage CPU time utilisation by the major job categories on the Green, Fermat and Newton systems for the 1st quarter 2006.




During this quarter half of the workload on Green was in the 33 to 64 PE range.




The workload on Fermat during this quarter was concentrated primarily in the 33 to 64 PE range.




During the first quarter there a large proportion of the workload on Newton was in the 129 to 256 range of PEs.



Job size trends on Green tended more to be in the mid-range PEs during this quarter.



The 33 to 64 PE range was predominantly utilised on Fermat in the first quarter.




During the first quarter usage on Newton was concentrated mainly in the mid- to high-range PEs.

2.2 System Usage Graphs


In all the Usage Charts, the baseline varies dependant upon the number of days in each month, within a 365-day year.

2.2.1 Baseline System


The graph below shows the Gflop Years utilisation on the CSAR systems by Research Council for the last 12 months.





The graph below shows the same service utilisation by major system.




The next chart shows the historic interactive usage of the Origin 300 system Wren. Eight of the higher speed 500Mhz CPUs in Wren deliver the baseline capacity equivalent to that which was previously available on the Origin 3000 system Ferret for interactive usage.



The following two charts detail the historic usage of the Origin 3000 system (Green) and the Altix 3700 system (Newton).




The next series of graphs illustrates the usage of the Medium Performance disk and HSM/tmp resources of the system.



The above graph illustrates the historic allocation of the Medium Performance Disk on Fermat and the SAN.

The next graph shows the historic HSM usage by Research Council funded projects, which has exceeded the overall Baseline of 16 Terabytes, and now totals about 48 Terabytes.



2.2.2 Guest System Usage

There is currently no Guest System usage.

2.3 Capability Incentives

Capability incentives were historically given on the T3E system Turing for jobs of 512 PEs and above. In July 2003 it was announced that discounts for capability jobs available on all CSAR systems had been approved to include the SGI Origin 3000 system Green and the SGI Altix 3700 system Newton.

These capability incentives were agreed with the Research Councils to encourage capability usage of the national supercomputers for greater scientific achievement, and offer the following discounts:

System	No of Processors	Discount
newton	192+ CPUs	15% discount
newton	128+ CPUs	10% discount
green	384+ CPUs	15% discount
green	256+ CPUs	10% discount

Discounts are given in the form of refunded Service Tokens.

Changes in usage patterns will be monitored and, subject to review, CfS reserve the right to change the incentives at any future date.

The following table displays the capability incentive discounts granted during the first quarter of this year.

Service Tokens Refunded: Quarter 1 2006 Usage							
System		Consortia					Total
		cse171	csn006				
Green	256+						0
Green	384+						0
Newton	128+	28.16	333.79				361.95
Newton	192+	678.1					678.1
Total Tokens		706.26	333.79				1040.05

This is within the CfS Management Board's forecast.

2.4 Service Status, Issues and Plans

Status

The service has been reasonably utilised throughout the first quarter of 2006, with usage exceeding baseline.

There was a relatively balanced spread of work across all major systems throughout this quarter.

Issues

A series of unrelated hardware issues caused several outages on Green during March. These were rectified by the vendor as soon as possible as each occurred.

Plans

There are currently no plans to report for the CSAR service.

3 Project Management, Documentation and User Feedback

This section covers aspects relating to the registration of projects and users, the management of projects and resources, topics associated with documentation and user feedback.

3.1 Project Applications

No applications for new CSAR projects were received.

3.2 New Projects

4 new CSAR projects were started, with a total of 234,919 tokens allocated.

3.3 Finished Projects

3 projects ended.

3.4 Queries

A total of 144 CSAR queries were dealt with:

- o 75 non-in-depth
- o 66 user registration and admin
- o 3 in-depth

3.5 Service Quality Tokens

No black or gold marks were awarded during this period.

3.6 Annual User Survey

The report of the annual user survey was published and can be viewed online at http://www.csar.cfs.ac.uk/project_management/reports/user_survey/2005.pdf.

4 Scientific Application Support Services

4.1 Training and Education

The following courses were delivered:

- Introduction to HPC, 3 March (Jon)
- Shared Memory Parallelisation with OpenMP, 6 March (Zoe)
- Advanced MPI, 9 March (Jon)
- Introduction to MPI, 8 March (Jon and Tim)
- Itanium 2, 23-24 March (Zoe)

4.2 Consortia Support/Software

There has been a focus this quarter on supporting projects that will need to migrate from CSAR to HPCx at the end of the CSAR service. Of particular importance are I/O issues and determining what resources projects can be transferred to HPCx.

Kevin Roy worked with Iain Barrass from the Health Protection Agency (HPA) on optimising his application. Iain visited for a week and Kevin improved the performance of his code by a factor of 4. A technical report was written describing the research, which involves activities such as predicting the spread of avian flu and we are discussing a potential follow-on joint project. In addition, Joanna Leng discussed with Iain how he could improve visualisation of his work, and discussions are in progress regarding submitting a follow-on project. This work is also relevant to a forthcoming visit by UoM staff to the Institute for the Protection of the Citizen, Ispra in Italy.

Jon Gibson and Zoe Chaplin have continued their work optimising and parallelising the Amazon-SC code on behalf of project cse135. Work is likely to be necessary beyond the end of the CSAR contract. There have been regular progress meetings with the project group, and considerable progress has now been made.

Jon Gibson is looking at the optimisation of the cgles code on behalf of Prof. John Williams (Queen Mary), CSAR project cse111.

Lee Margetts provided in-depth advice about parallelising non-trivial Lattice Boltzmann programs, on behalf of Prof. Peter Coveney's project.

Lee Margetts and Joanna Leng supported Henggui Zhang with his work on a CSAR class 3 project, 'Computing functional roles of gene mutation leading to sudden cardiac death'. Joanna is developing applications that allow the production of animations from his time dependant data. It is likely that a class 1 application in heart modelling will result from this project.

Gavin Tabor visited the Manchester site. Joanna Leng discussed visualisation techniques and demonstrated how the Virtual Reality facility could be exploited in his project.

4.2.1 DL_POLY

Installed new releases of DL_POLY on Newton – DL_POLY_2.16 and DL_POLY_3.06 (Tim Robinson).

4.2.2 CASTEP

Installed a new release of Castep on Newton and the Origins – castep-4.0 (Tim Robinson).

4.2.3 AMBER

Compiled the replica exchange method (REM) into parallel sander (part of the AMBER distribution) on Newton (Tim Robinson).

4.2.4 FLUENT

Installed the Fluent software on Newton. We have a demonstration licence valid until the end of April. Fluent is a commercial parallel CFD package. Only one user has been given access to it for a specific project. (Jon Gibson and Lee Margetts).

4.2.5 UM

A considerable amount of time was spent in January and early February beta-testing version 6.1 of the UM. Results have been passed back to the Met. Office, to be incorporated into the full release of the model. (Zoe Chaplin).

4.2.6 Machine Status

Jon Gibson wrote a machine_status script for Newton which runs a number of health checks.

5 Collaboration and Conferences

5.1 MRCCS Projects

5.1.1 MRCCS Seminars

Seminar series of HPC and visualisation related topics (organised by Jon and Lee, chaired by Jon):

- Dr Ian McLuckie, Advanced Integrated Engineering Solutions Ltd, “Faster Real World CAE Solutions – Using Integrated Knowledge Based Systems”, 27 January
- Dr Matthias Heil, University of Manchester, “oomph-lib – The object-oriented multi-physics finite-element library”, 17 January
- Dr Tom Jackson, University of York, “Storage Request Broker (SRB)”, 24 February

5.2 Events

Events/workshops this quarter included:

- CSAR User Steering Group meeting, London/Manchester
- The 2nd Technical Symposium on Reconfigurable Computing with FPGAs, 9-10 February and Workshop on Programming FPGAs, 8 February, University of Manchester (Kevin Roy organised this event; see Added Value Services)
- UMARI sponsored event, the A-Star workshop, aimed at looking into DTI funding into work on intelligent travel. The meeting looked at how this could be applied to airports. Farnborough, 26 February
- BBSRC road show at Manchester on 2 March
- Joanna Leng exhibited HPC and Visualisation activities at the Museum of Science and Industry in Salford, 12 March
- Meet the Scientist workshop organised by EPSRC as part of their program to aid communicating science to the public, 18 January
- NAFEMS/EU funded AUTOSIM Workshop, Barcelona, 17-18 January
- NAFEMS Iberian Steering Group Meeting, Barcelona, 17 January
- Bioinformatics and the Grid meeting
- SOS HPC meeting, Maui
- Visit to NERSC, US

5.3 Publications

Smith, I.M. and Margetts, L., ‘The convergence variability of parallel iterative solvers’, Engineering Computations, Volume 23, Issue 2, pp. 154-165, 2006.

6 Added Value Services

6.1 2nd Symposium on Reconfigurable Computing with FPGAs

The 2nd symposium on FPGAs in HPC was hosted by MC in February. A one-day workshop, in conjunction with SGI, was held prior to the symposium. Both events were hugely successful, with about 80 delegates attending. It is planned that this meeting will become an annual event, hosted at Manchester. Manchester will also participate in the next similar meeting at NCSA, Illinois in July 2006.

6.2 National Grid Service

The CSAR service provides a resource as part of the NGS. Tim Robinson is a member of the NGS Technical Board and attends bimonthly technical board meetings.

6.3 Visualisation

CSAR visualisation work was presented on a portable virtual reality environment at the half day visualisation event at the Museum of Science and Industry in Salford, 12 March.