CSAR Service

Consolidated Management Report

2nd Quarter 2006

Management Summary

This is the final consolidated Management report for the CSAR Service at the point of closure after 7½ years of successful service.

The number of users at the end of the quarter stood at a total of 409.

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 this final quarter.

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 2nd 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 July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	Mar	April	May	June
HPC Services Availability												
Availability in Core Time (% of time)	98.75%	99.04%	97.85%	98.50%	98.33%	96.99%	99.25%	94.82%	98.53%	99.47%	88.47%	96.53%
Availability out of Core Time (% of time)	99.45%	99.4%	96.82%	99.63%	99.12%	99.3%	99.40%	99.79%	98.46%	99.43%	95.04%	99.32%
Number of Failures in month	4	5	6	1	4	4	2	2	7	2	4	3
Mean Time between failures in 52 week rolling period (hours)	2430	1078	625	596	461	379	353	330	260	244	217	203
Help Desk												
Non In-depth Queries - Max Time to resolve 50% of all queries	<0.25	<0.25	<0.25	<0.25	<0.5	<0.25	<0.5	<0.25	<0.25	<0.25	<0.25	<0.25
Non In-depth Queries - Max Time to resolve 95% of all queries	<2	<1	<2	<3	<5	<2	<2	<3	<3	<1	5>	<1
Administrative Queries - Max Time to resolve 95% of all queries	5>	<0.5	<0.5	<0.5	<2	<1	<3	<5	<2	<0.5	<2	<0.5
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 mont	2	2	2	2	2	2	2	2	2	2	2	2

Table 2

Notes:

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

   [Fermat availability x 40/ (40+233+343)] + [Green availability x 233/(40+233+343)] + [Newton availability x 343/(40+233+343)]

2. 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 July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	Mar	April	May	June
HPC Services Availability												
Availability in Core Time (% of time)	0.039	0.039	0.078	0.039	0.078	0.078	0	0.115	0.039	0	0.115	0.078
Availability out of Core Time (% of time)	0	0	0.039	-0.039	0.000	0	0	-0.039	0.039	0	0.039	0
Number of Failures in month	0.008	0.0156	0.023	-0.008	0.008	0.008	0	0	0.023	0	0.008	0
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	-0.016	-0.019	-0.016	-0.019	-0.019	-0.019	-0.019	-0.019
Non In-depth Queries - Max Time to resolve 95% of all queries	0	-0.016	0	0.016	0.031	0	0	0.016	0.016	-0.016	0.046	-0.016
Administrative Queries - Max Time to resolve 95% of all queries	0.046	-0.019	-0.019	-0.019	0	-0.016	0.016	0.031	0	-0.019	0	-0.019
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 mont	0	0	0	0	0	0	0	0	0	0	0	0
Monthly Total & overall Service Quality Rating for each period:	0.02	-0.01	0.04	-0.03	0.04	0.01	-0.01	0.04	0.04	-0.04	0.08	0.00
Quarterly Service Credits:			0.05			0.03			0.06			0.04

Table 3

1.2 No. of Registered Users

The position at the end of the quarter is that there were 409 registered users of the CSAR Service.

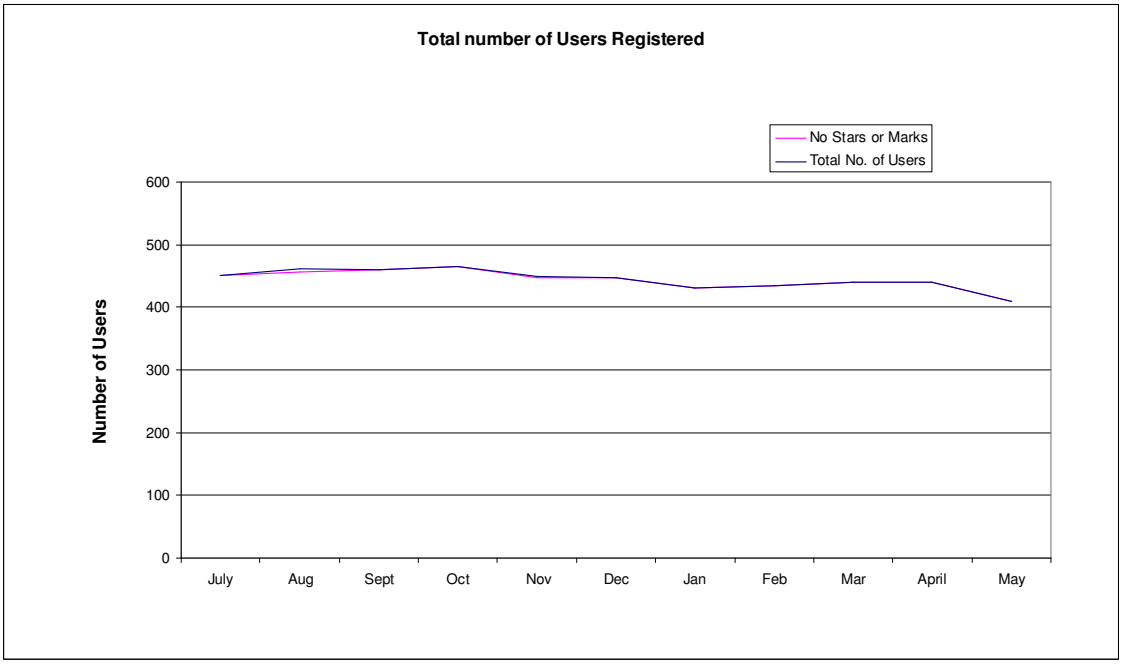

1.3 Service Quality Tokens

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

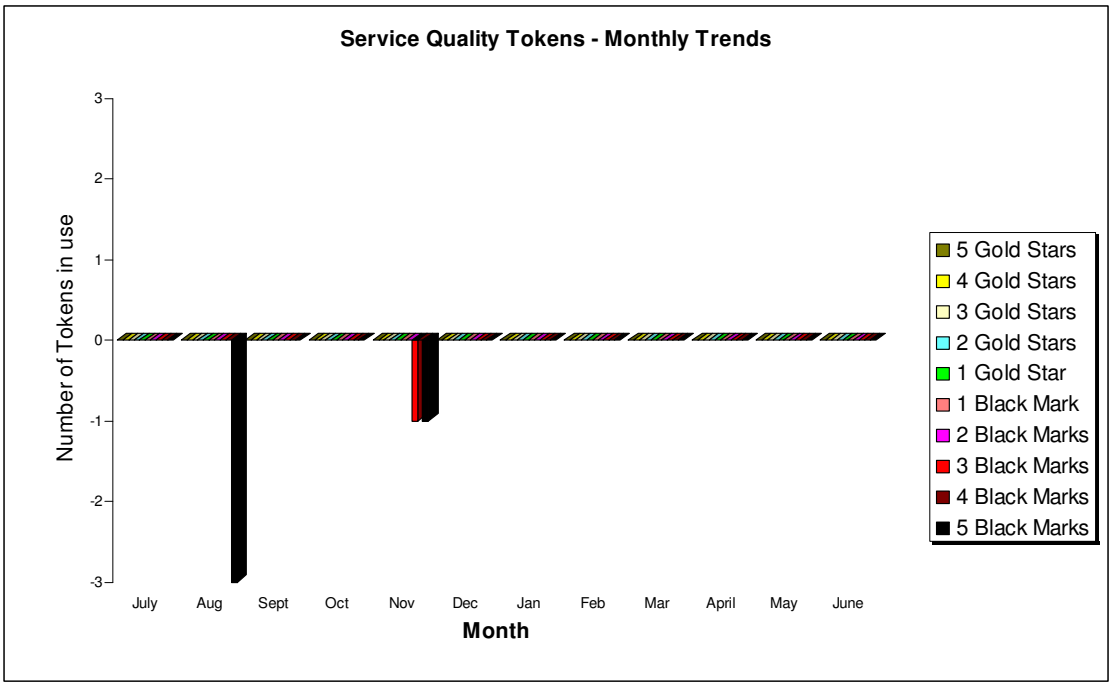

Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

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.

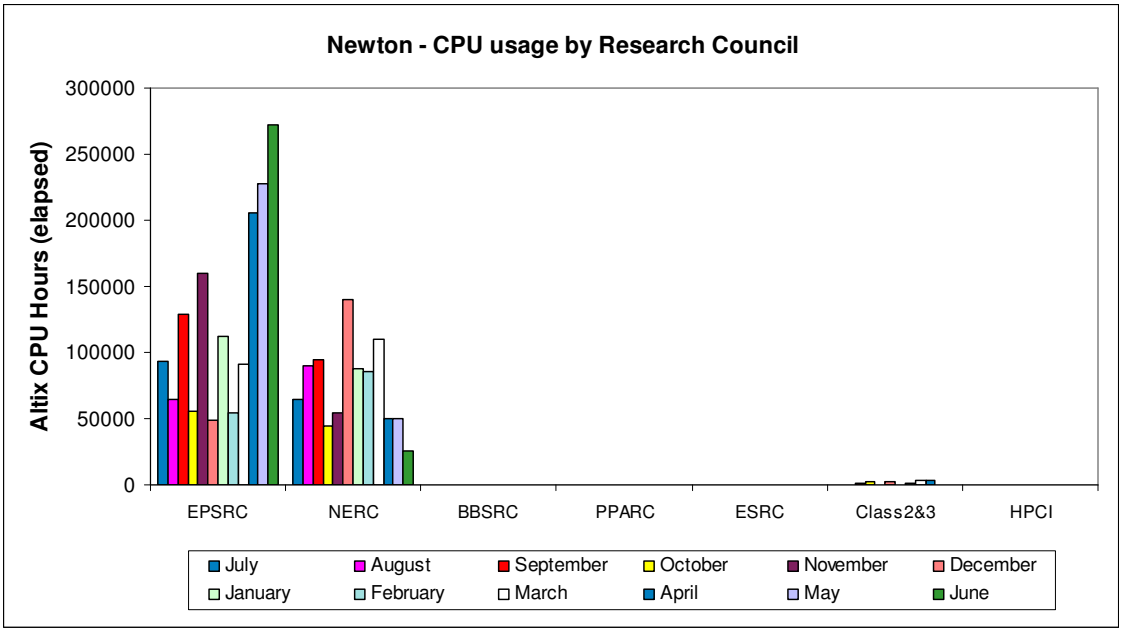

Usage of the SGI Altix 3700 system Newton is shown by Research Council in the above graph.

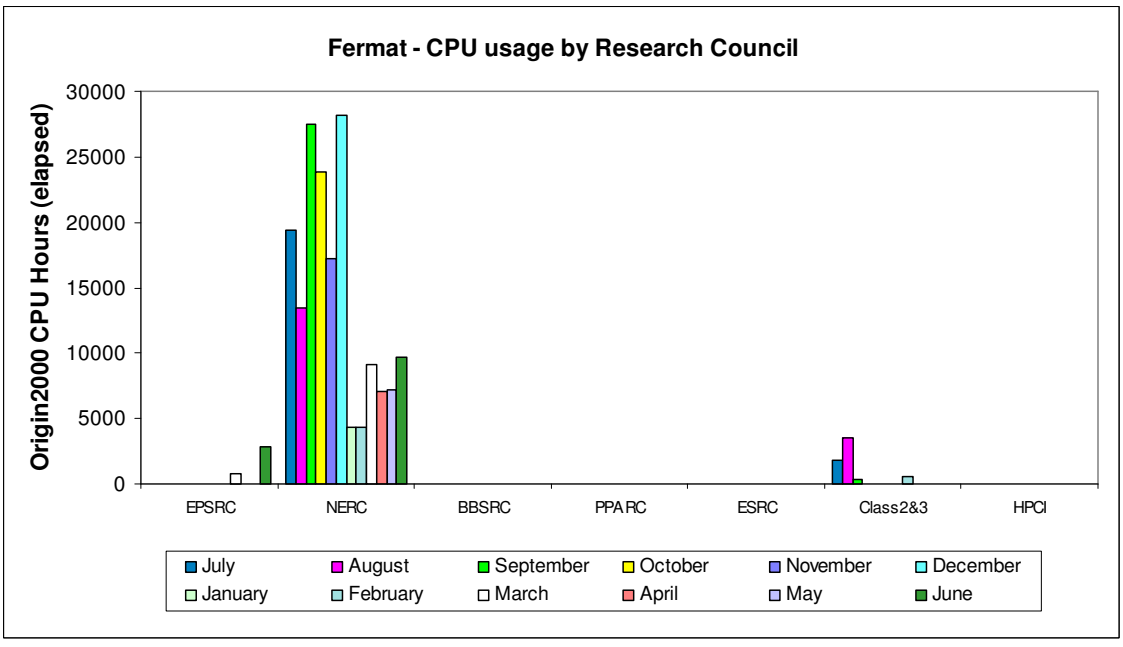

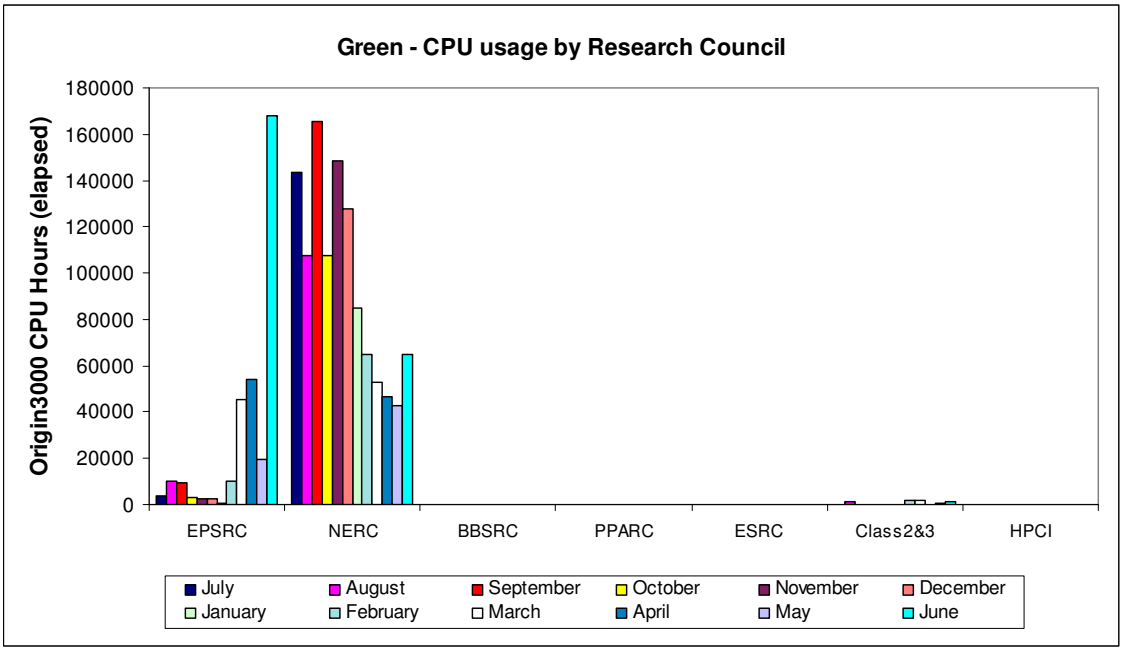

Usage of the two batch SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service in the preceding two charts.

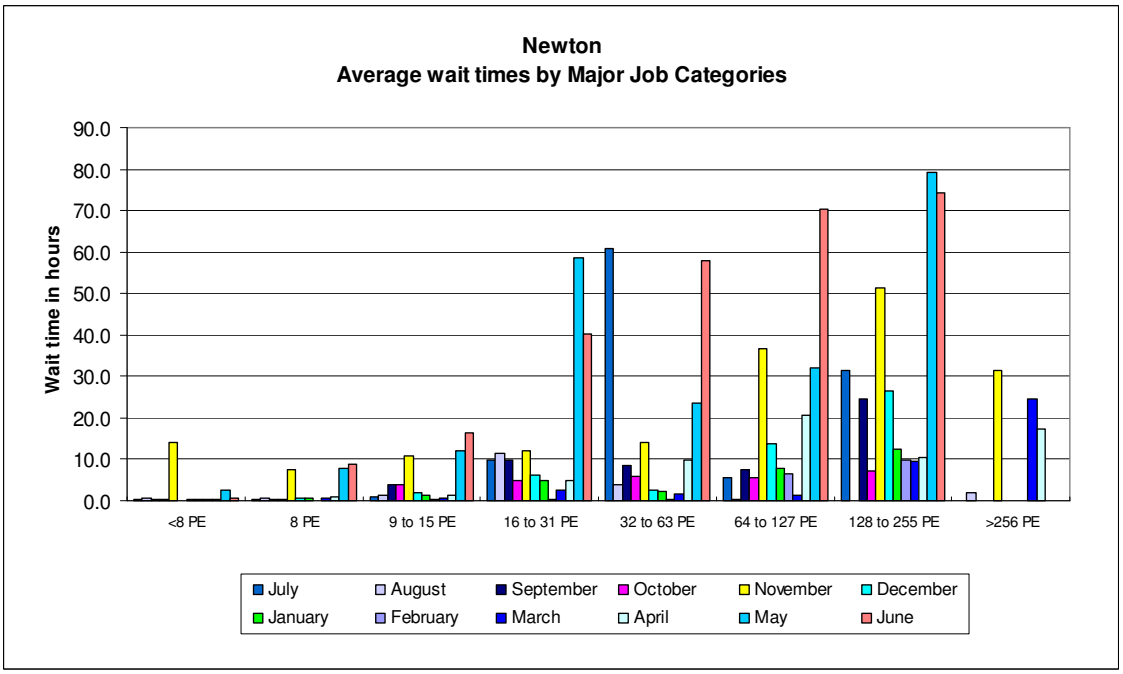

The above chart, and the one below, shows the wait time trend in hours on the Altix 3700 Newton.

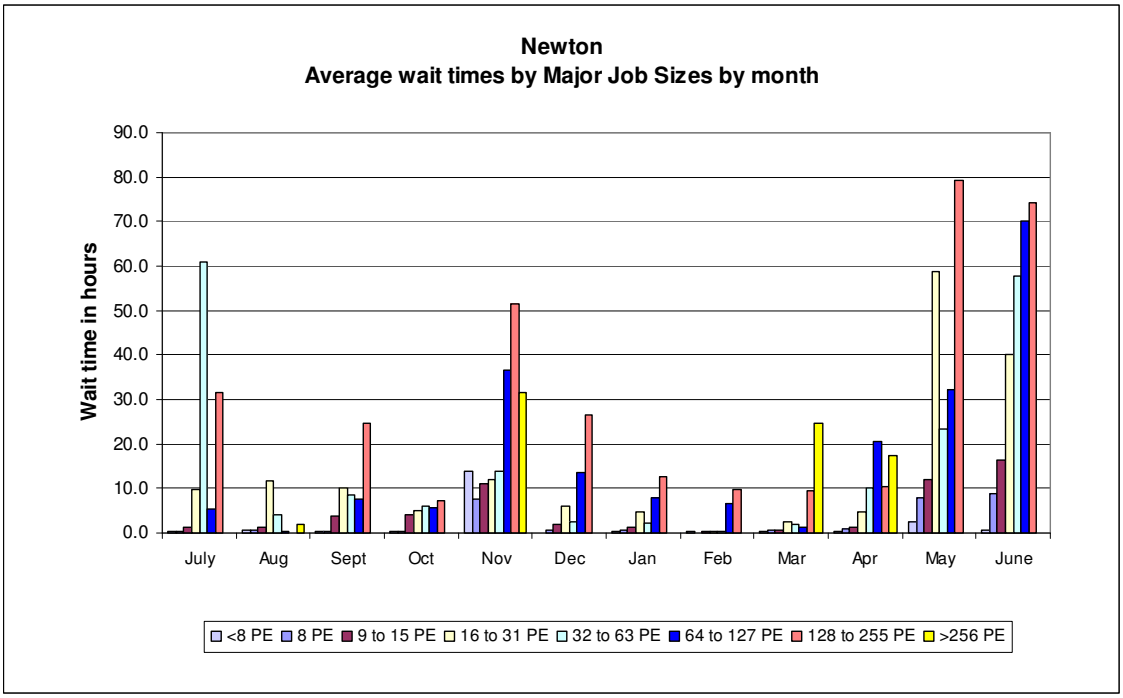

The following two charts show average wait times in hours for the quarter on the Origin 3000 Green.

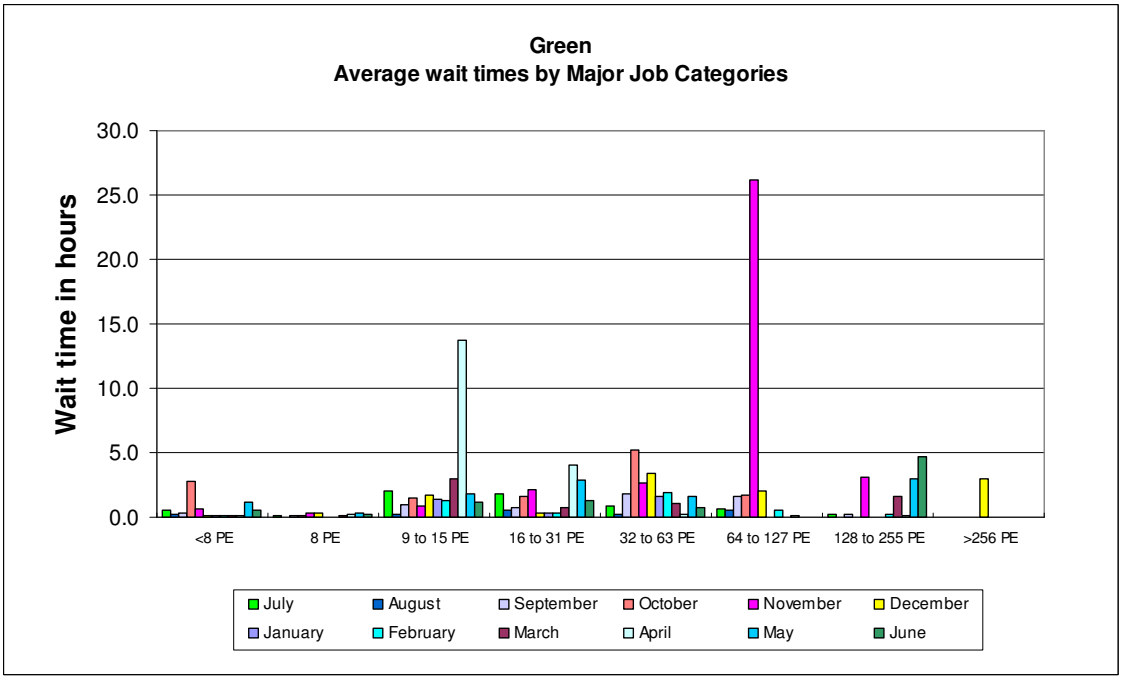

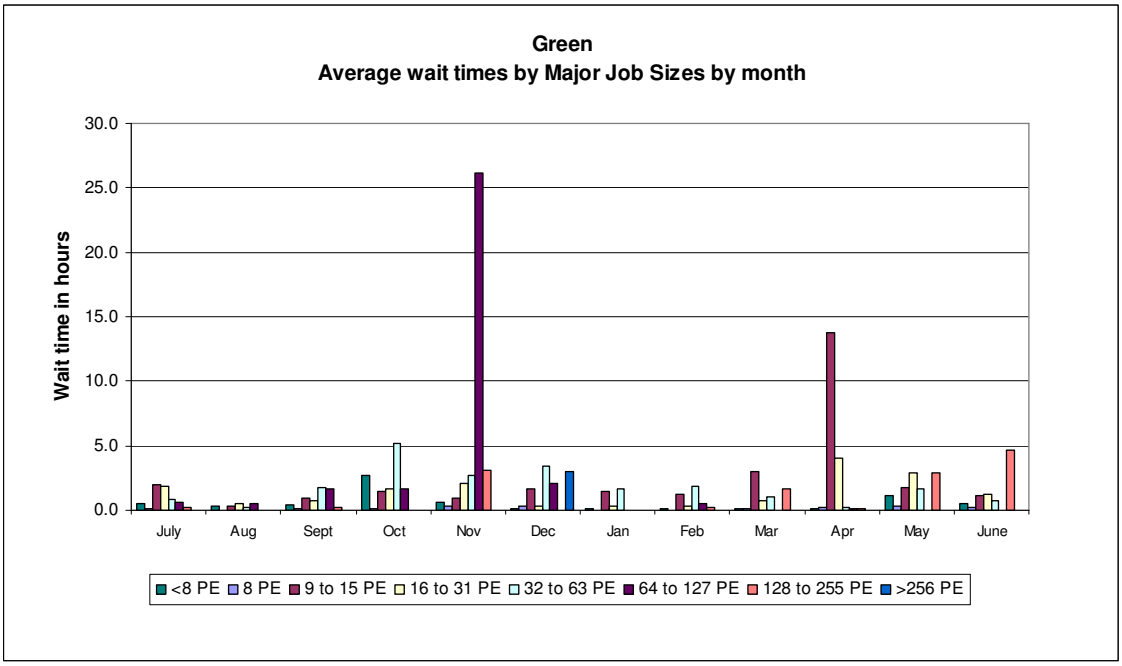

The next series of four charts show the percentage CPU time utilisation by the major job categories on the Green, Fermat and Newton systems for the 2nd quarter 2006.

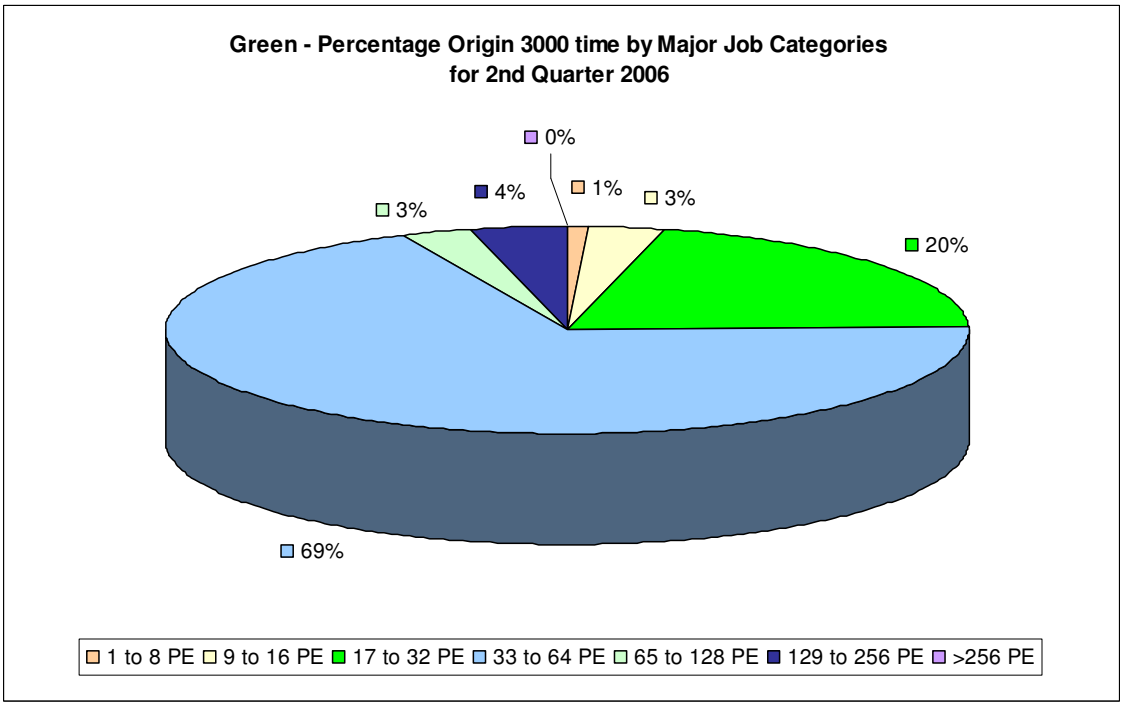

During this quarter the greatest concentration of the workload on Green was in the 33 to 64 PE range.

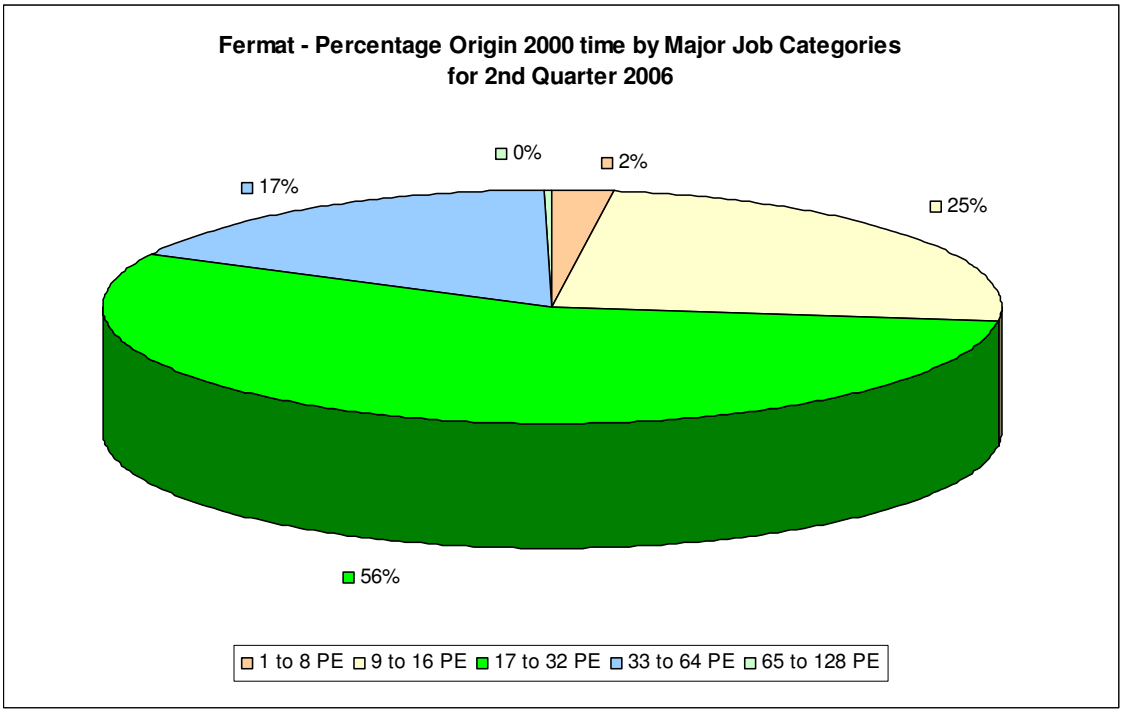

Just over half of the workload on Fermat during this quarter was in the 17 to 32 PE range.

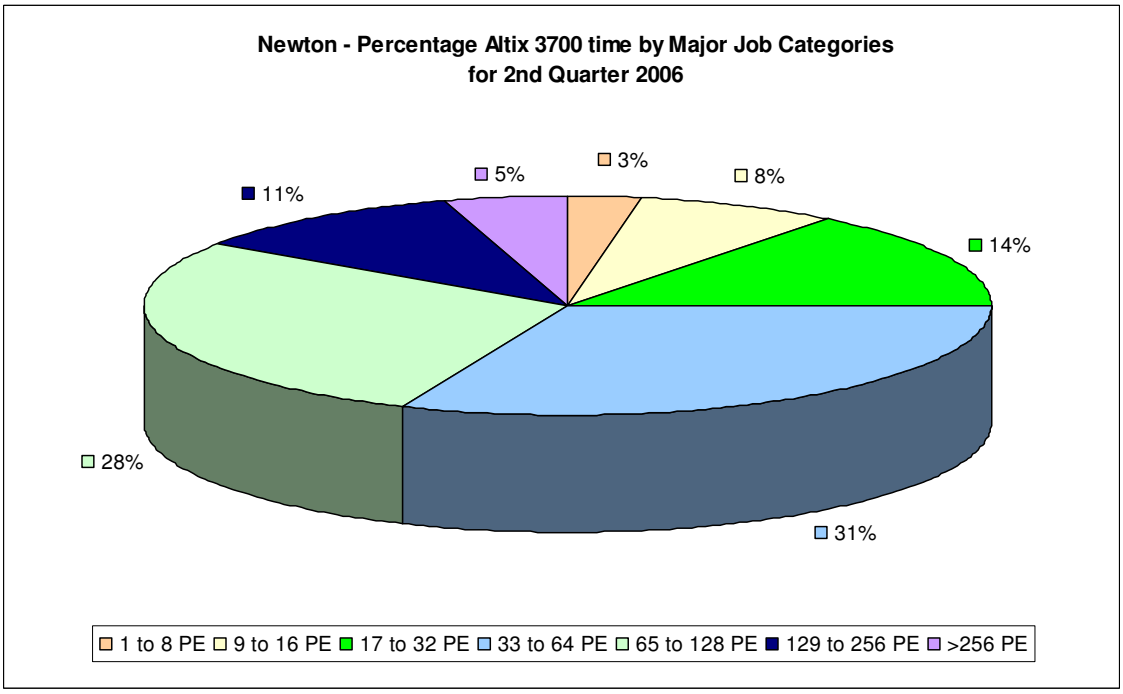

During the second quarter the workload on Newton was spread evenly across a range of PEs.

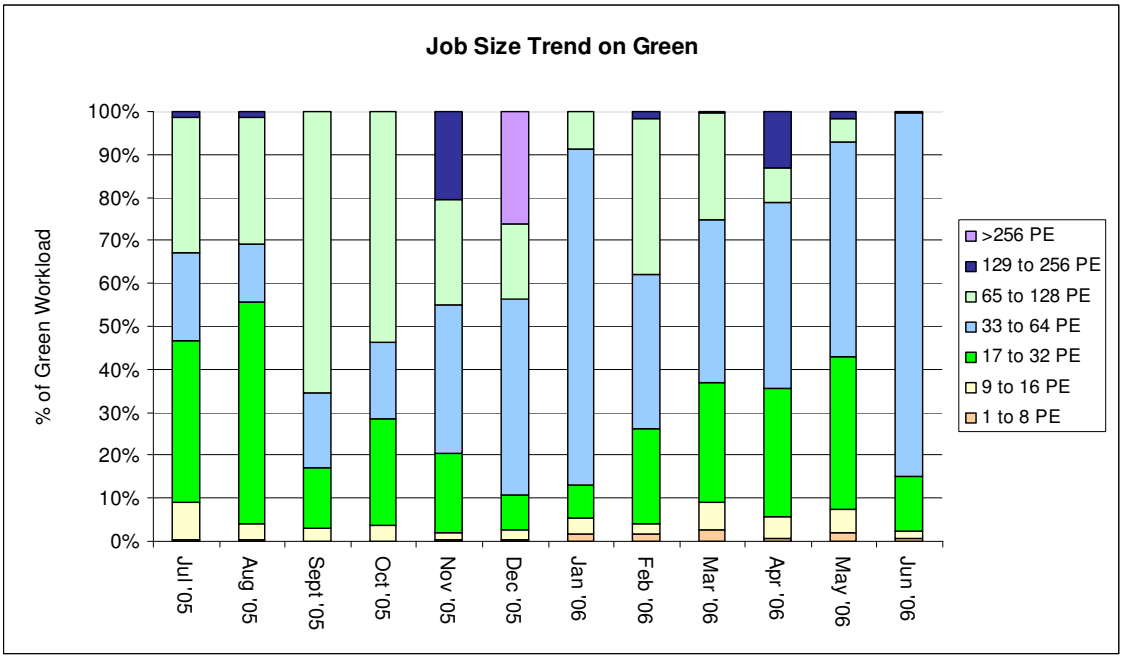

Job size trends on Green tended to be in the mid-range PEs during this quarter.

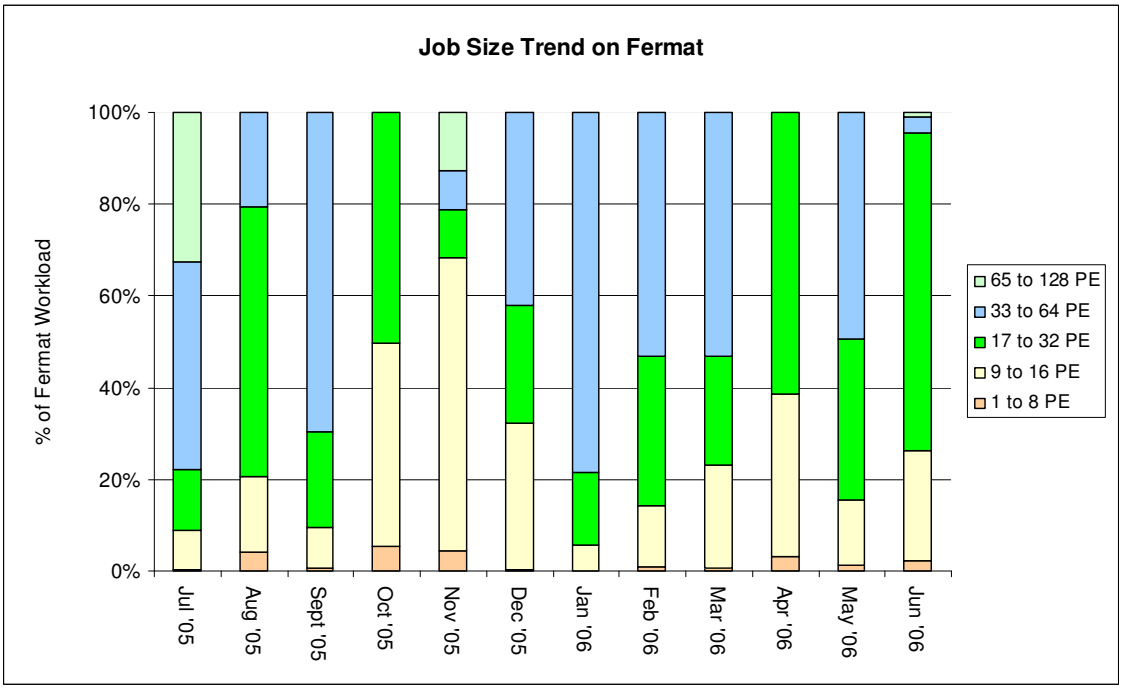

The 17 to 32 PE range was predominantly utilised on Fermat in the second quarter.

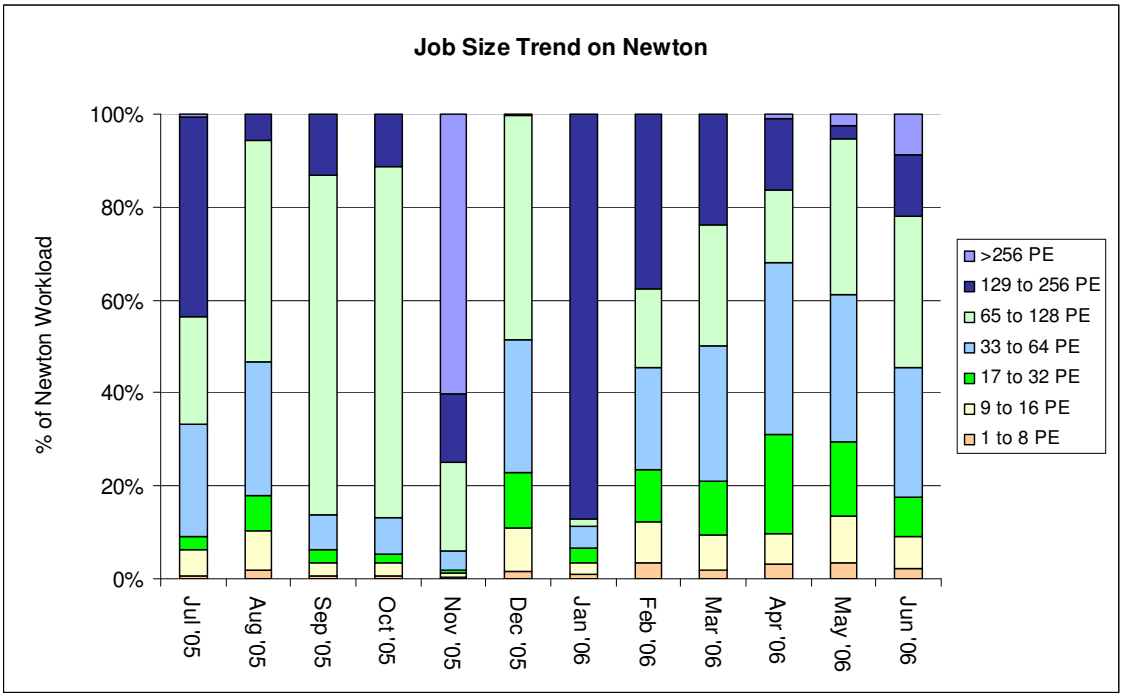

During the second quarter usage on Newton was concentrated mainly in the mid- to high-range PEs.

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.

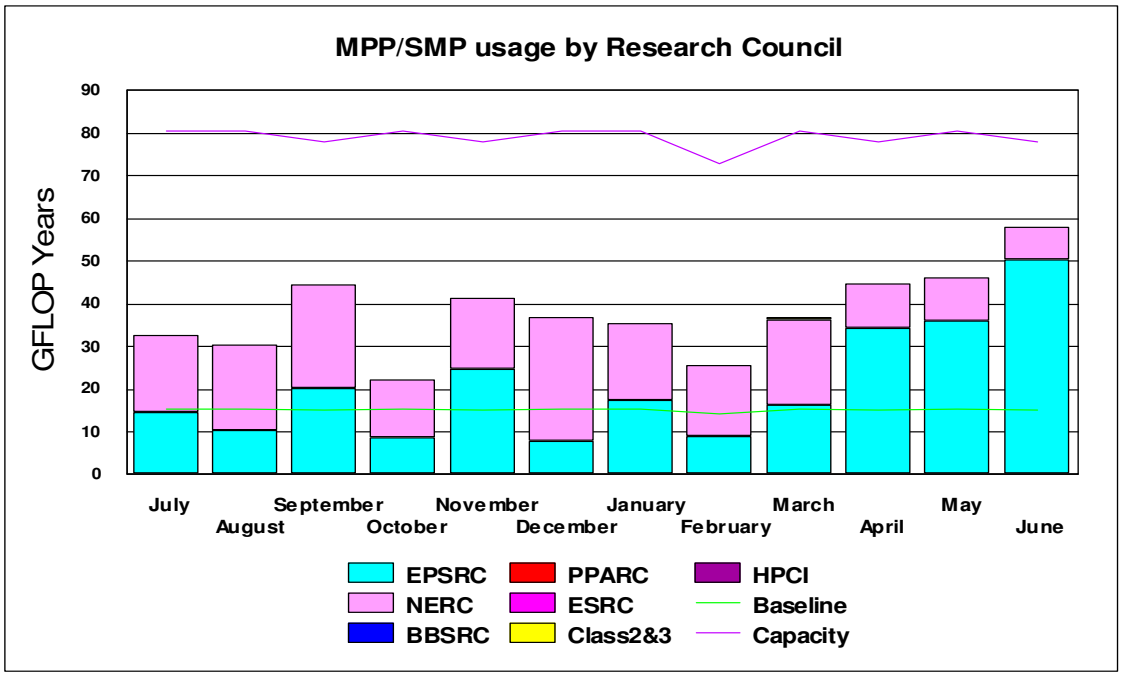

The graph below shows the same service utilisation by major system.

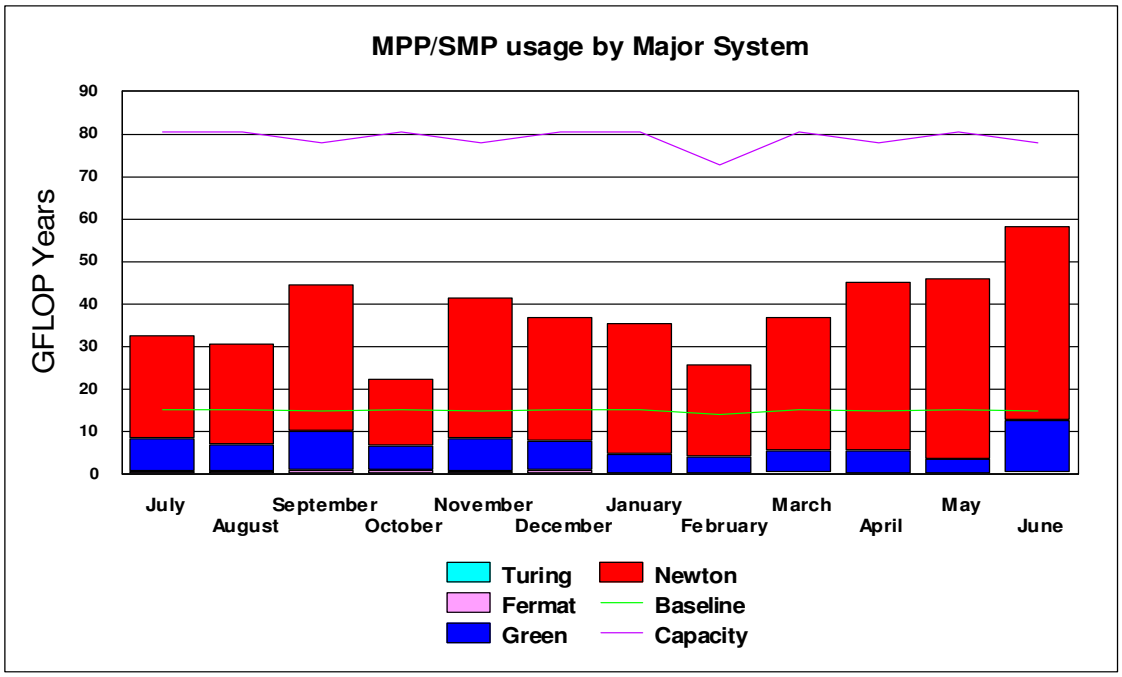

The next chart shows the historic interactive usage of the Origin 300 system Wren. Eight of the higher speed 500Mhz CPUs in Wren deliver the baseline capacity equivalent to that which was previously available on the Origin 3000 system Fermat for interactive usage.

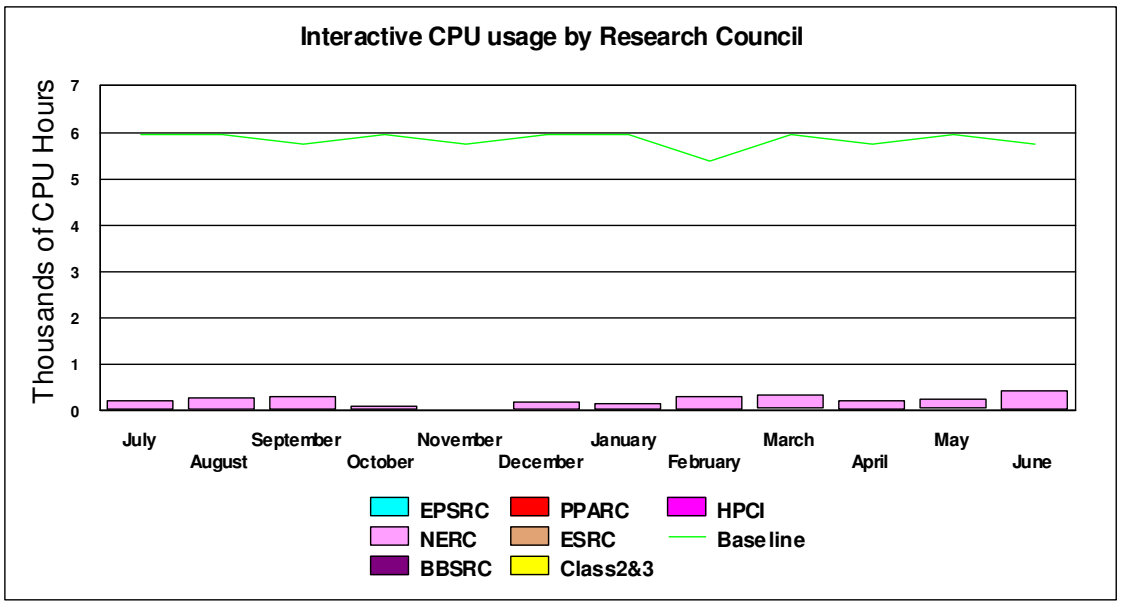

The following two charts detail the historic usage of the Origin 3000 system (Green) and the Altix 3700 system (Newton).

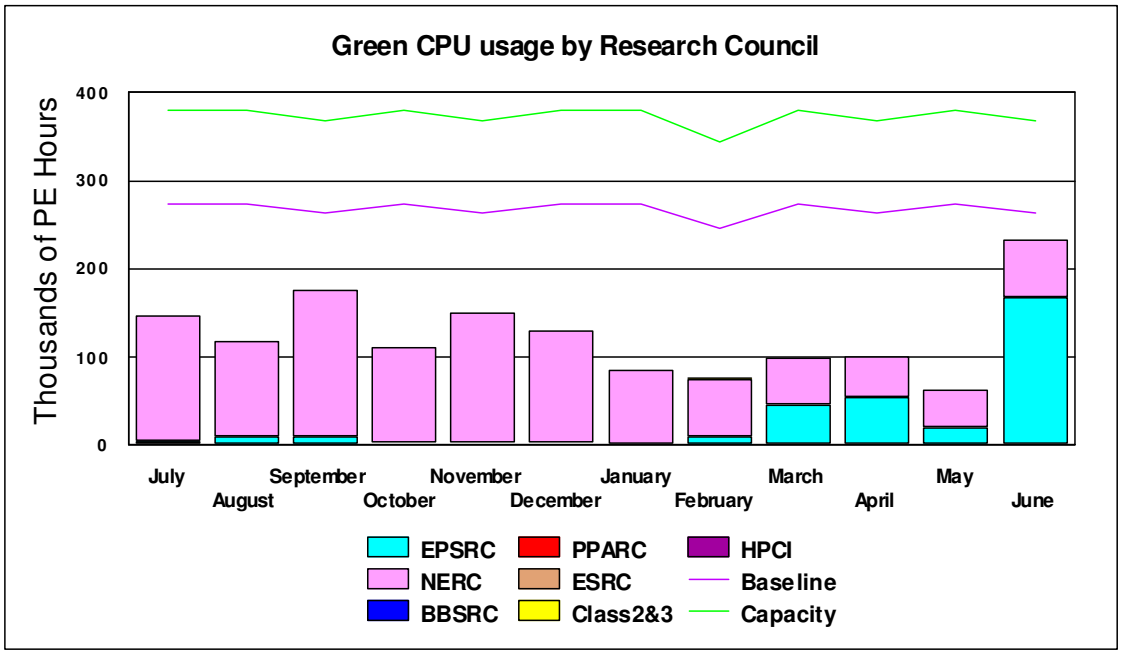

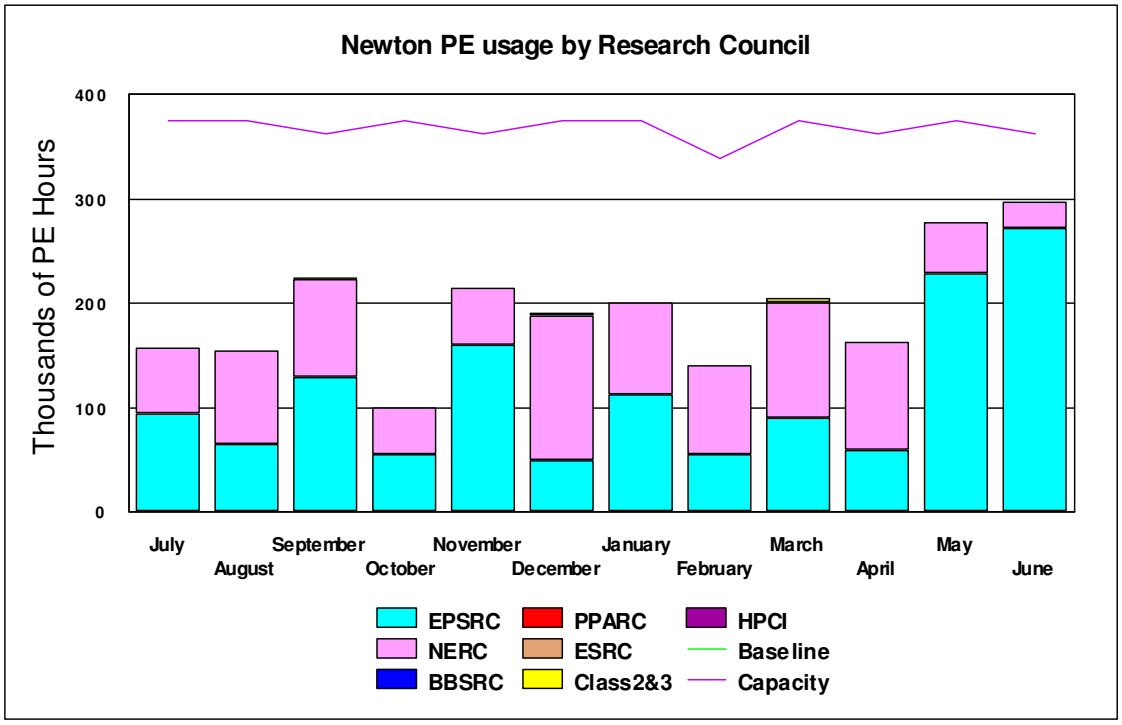

The next series of graphs illustrates the usage of the Medium Performance disk and HSM/tmp resources of the system.

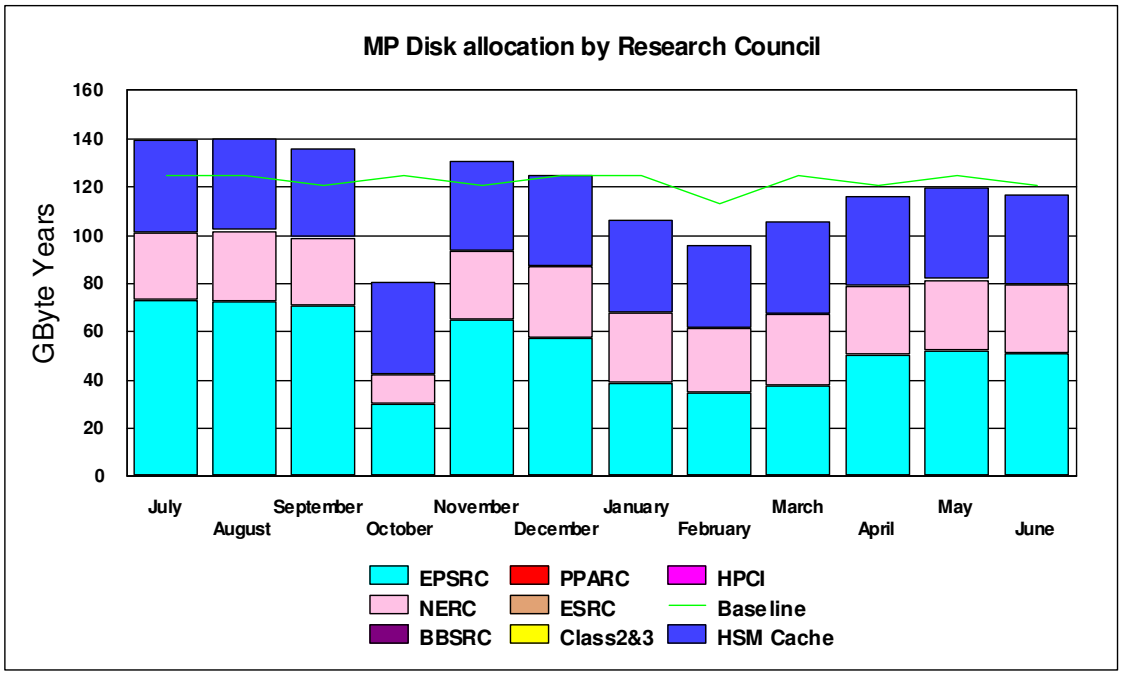

The above graph illustrates the historic allocation of the Medium Performance Disk on Fermat and the SAN.

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.

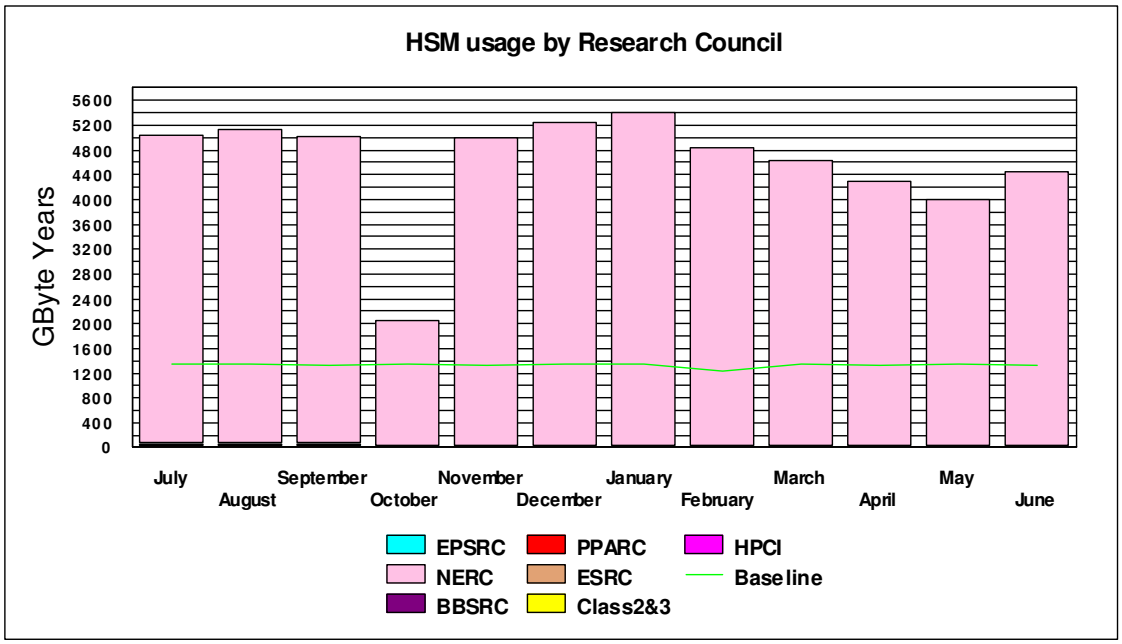

2.2.2 Guest System Usage

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 second quarter of this year.

Service Tokens Refunded: Quarter 2 2006 Usage							
System	Consortia						Total
	cse171	cse174	cse133	cse110	csn003	csn006	
Green 256+ PEs				2.68		1.58	4.26
Green 384+ PEs							0
Newton 128+ PEs		457.42	160.41		1.35		619.18
Newton 192+ PEs	141.16		52.38				193.54
Total Tokens	**141.16**	**457.42**	**212.79**				**816.98**

This is within the CfS Management Board's forecast.

Note that due to the closure of the CSAR service on June 30th, there were no capability incentive discounts for June.

2.4 Service Status, Issues and Plans

Status

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

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

Issues

There are no issues to report for this quarter.

Plans

The CSAR Service ceased operation on Friday 30 June 2006 at the end of its current contract. For projects that run beyond this date, the Research Councils are responsible for providing access to alternative resources. All users who currently have accounts on the CSAR systems are responsible for transferring their data elsewhere. For further details refer to the “Closure of the CSAR Service” page on the CSAR website at http://www.csar.cfs.ac.uk/user_information/service_end/.
.

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

1 new CSAR projects was started, with a total of 500 tokens allocated.

3.3 Finished Projects

11 projects ended.

3.4 Queries

A total of 130 CSAR queries were dealt with:

- 78 non-in-depth
- 52 user registration and admin
- 0 in-depth

3.5 Service Quality Tokens

No black or gold marks were awarded during this period.

3.6 Closure of CSAR

Details of projects and remaining resources were provided to EPSRC and HPCx. Extensive efforts have been made to ensure all users are aware of the service closure and the removal of their data.

4 Scientific Application Support Services

4.1 Training and Education

The following courses were delivered:

- Porting from CSAR to HPCx, 25 May (Jon)
- Share Memory Parallelisation with Open MP, 19 June, for the group in Salisbury (cse110, Leach) (Kevin)

4.2 Consortia Support/Software

Zoe Chaplin and Jon Gibson have continued their work optimising and parallelising the Amazon-SC code on behalf of David Ingram (project cse135). The original code contained many bugs which have all been fixed. Timings on Newton are 970 seconds.

The code has also been parallelised using mpi to divide the time of the run into 'chunks'. The input files we have been given means that each 'chunk' size is about 24s so, even with a timing of 550s on HPCx for 0.1s simulation time, the code is going to take more than a day to run. Some work has also been performed, therefore, on writing dumps and using these to restart the code.

There have been regular progress meetings with the project group, and considerable progress has now been made.

Kevin Roy began work on Tomcat, an OpenMP code for atmospheric chemistry. The code was a heavy consumer of CPU and Martyn Chipperfield of the NCAS consortium was looking to reduce the cost of running the model.

Kevin Roy worked with members of the Jonathon Pitts project cse129 who visited Manchester in May. They were looking at how to get several applications ported and to obtain advice on how to proceed on national HPC systems.

Craig Lucas worked on Fortran data mining codes for project cse131. They are serial codes designed to simplify the performance analysis of the existing heavily templated C++ codes and increase efficiency. They are written in a modular and elegant way such that different data mining algorithms can be switched between by specifying different data-splitting and hashing methods. Thus the output will be a suite of codes. Some further work is required to finalise the routines and test them.

Joanna Leng supported Hengui Zhang with his work on heart simulation under a CSAR class 3 project. Joint submissions to the All Hands Meeting (successful) and to Europar (awaiting outcome) have been made.

Joanna Leng provided flow visualisation support for Gavin Tabor's class 3 project, producing images of one of the data sets using AVS/Express. This compared well to the results provided by Fluent.

Vendel Szeremi has performed a significant amount of work for Philippe Young's new project, cse169, on modelling of advanced materials. This project uses software to generate the large meshes required for the project, so parallelisation with OpenMP and optimisation has been performed to achieve this.

Additionally work has been performed to improve the scalability of the solver used in the ParaFEM library, used in this project. This has involved the use of the mesh partitioner, METIS, and rewriting some of the communications routines to reduce the total amount of communications.

4.2.1 PETSC

Jon Gibson installed PETSC (Portable, Extensible Toolkit for Scientific Computation) onto the Origins. This is a suite of data structures and routines for the scalable (parallel – MPI) solution of scientific applications modelled by partial differential equations.

4.2.2 Exonerate

Jon Gibson installed Exonerate on Newton. Exonerate is a generic tool for pairwise sequence comparison in bioinformatics.

4.2.3 UM

Zoe Chaplin has been testing the latest version of UM 6.1 on Newton.

4.2.4 FFTW

Kevin Roy put the latest version of FFTW, 3.1.1, onto Newton at a user request. This version was required as the change from versions 2 to 3 changed the API, which meant that codes written for version 3 could not use version 2.

4.2.5 LAMMPS

Tim Robinson installed a new version, 12th April release, of LAMMPS onto Newton at a user request.

5 Collaboration and Conferences

5.1 MRCCS Projects

5.1.1 MRCCS Seminars

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

- Professor Bob Stone, "Virtual Environments and Interactive 3D: A Reflection Spanning 20 Years", 28 April
- Professor Blackledge, "Fractional Dynamics and some of its Applications", 26 May
- Dr Richard Henchman, "Solvent free energy from a single simulation via cell theory", 02 June
- Julia Handl, "Nature-inspired Approaches to Clustering", 09 June
- Dr Aladdin Ayesh, "DMU Creative Robots", 16 June
- Professor Min Chen & Dr David Chisnall, "The Use of Volume Scene Graphs in Very Large Dataset Visualization", 23 June

5.2 Events

Events/workshops this quarter included:

- Kevin Roy, Craig Lucas and Mike Pettipher attended CUG 2006 in Lugano, Switzerland, 8-11 May. Kevin and Craig both presented papers. Kevin's, entitled "Parallel Performance Analysis on Cray Systems", examined a tool he developed that made analysis and interpretation of multiple profiles from multiple parallel jobs easier. Craig gave the paper "Symmetric Pivoting in ScaLAPACK". This was a new routine, incorporating a pivoting technique untried in ScaLAPACK, and aimed for a future release. Mike narrowly failed in his bid to become the Vice President of the Cray User Group, but remains a Director at Large for the group.
- Kevin Roy attended SGIUG 2006 in Las Vegas, US and then presented a paper entitled "MPI Collective Operations and Beating Them". The paper examined the performance of collective operations on SGI Altix systems and looked to improve upon them using single-sided communication. Kevin was also accepted onto the SGIUG board, where he will be director of Novel Technologies.
- Craig Lucas attended PARA06 in Umea, Sweden, 18-22 June. He was invited to speak at the mini-symposium on Recent Advances in Dense Linear Algebra and presented the paper "Pivoted Cholesky and QR Updating".
- Joanna Leng organised and held a workshop on the User Requirements of Visualization at the Second International Conference of e-Social Science.
- Joanna Leng organised a workshop with the Institute for the "Protection and Security of the Citizen", ISPRA, in early April.
- Kevin Roy prepared a paper for EuroPVM/MPI in conjunction with the CSAR user Iain Barrass. This paper discussed the parallelisation of a Disease Propagation code.
- Firat Tekiner presented "Attribute Selection Methods for Filtered Attribute Subspace based Bagging with Injected Randomness (FASBIR)" at SCDM 2006, Washington DC.
- Kevin Roy produced a poster consisting of 10 slides on the University of Manchester for the Gelato organisation. Gelato is a worldwide organisation supporting Linux Itanium users.
- Tim Robinson, Firat Tekiner and Kevin Roy attended the EPSRC Seminar day at the University of Manchester Conference Centre, 14 June.
- Tim Robinson attended Good Practice in Academic Intellectual Property & Commercialisation (session 4 – licensing), 24 June.

5.3 Publications

Isiuchi, S., Fujii, M., Robinson, T.W., et al., 'Vibrational overtone spectroscopy of phenol and its deuterated isotopomers', Journal of Physical Chemistry, A 110 (23): 7345-7354, June 15 2006.

Whittley, I.M., Bagnall, A.J., Bull, L, Pettipher, M., Studley, M., Tekiner, F., 'Attribute Selection Methods for Filtered Attribute Subspace based Bagging with Injected Randomness (FASBIR)', In Feature Selection for Data Mining Workshop, Part of the 2006 SIAM Conference on Data Mining, 2006.

Bagnall, A.J., Whittley, I.M., Studley, M., Pettipher, M., Tekiner, F., Bull, L., 'Variance Stabilizing Regression Ensembles for Environmental Models', In IEEE World Congress on Computational Intelligence, 2006.

Bagnall, A.J., Cawley, G., Whittley, I.M., Studley, M., Bull, L., Tekiner, F., Pettipher, M., 'Super Computer Heterogeneous Classifier Meta-Ensembles', International Journal of Data Warehousing and Mining. Accepted, expected to be published Second Quarter 2007.

6 Added Value Services

6.1 Reconfigurable Computing with FPGAs

- Kevin Roy helped organise and coordinate RSSI'06. The Reconfigurable Systems Summer Institute is held annually at NCSA, IL, US. The University of Manchester was a co-host and helped review papers and materials for the conference.
- Kevin created the bylaws for the OpenFGPA organisation. These have now been passed.

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 attended bimonthly technical board meetings.

6.3 ScaLAPACK

Craig Lucas completed a first version of the Pivoted Cholesky Factorization for the ScaLAPACK library. Although the starting point was the serial routine he had written for LAPACK, there was much work needed to parallelise the code and adapt it for the block cyclic data distribution of ScaLAPACK. The code requires some refinement but it has provided the possibility of using symmetric pivoting in ScaLAPACK. Symmetric pivoting requires rows and columns to be swapped at every step of the algorithm, instead of just rows in existing ScaLAPACK routines. Because of the data distribution this requires a large communication overhead involving all processes. Furthermore the routine is one that has been requested by users of the library, thus there will be benefit to the user community.