CSAR Service

Consolidated Management Report

3rd Quarter 2005

Management Summary

This is the consolidated Management report for the third quarter 2005 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 460 to date.

The workload on both the Origin 3000 Green and the Altix 3700 Newton has been fairly evenly spread across the mid-range of PEs during the third 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 3rd quarter 2005 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	2004/5 Oct	Nov	Dec	Jan	Feb	Mar	April	May	June	July	Aug	Sept
HPC Services Availability												
Availability in Core Time (% of time)	97.08%	98.50%	99.50%	97.37%	97.85%	97.85%	96.00%	99.50%	97.13%	98.75%	99.04%	97.85%
Availability out of Core Time (% of time)	98.67%	98.78%	99.2%	99.73%	99.5%	99.80%	99.90%	99.54%	99.22%	99.45%	99.4%	96.82%
Number of Failures in month	3	4	2	3	5	4	4	1	2	4	5	6
Mean Time between failures in 52 week rolling period (hours)	212	208	225	237	231	223	227	241	257	275	241	207
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	<0.25	<0.25	<0.25	<0.25	<0.25
Non In-depth Queries - Max Time to resolve 95% of all queries	<3	<1	<0.5	<0.5	<2	<1	<2	5>	<1	<2	<1	<2
Administrative Queries - Max Time to resolve 95% of all queries	<1	<0.5	<1	<0.5	<2	<1	<1	<0.5	<2	5>	<0.5	<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	2004/5 Oct	Nov	Dec	Jan	Feb	Mar	April	May	June	July	Aug	Sept
HPC Services Availability												
Availability in Core Time (% of time)	0.078	0.039	-0.039	0.078	0.078	0.078	0.078	0	0.078	0.039	0.039	0.078
Availability out of Core Time (% of time)	0	0	0	-0.039	0	-0.047	-0.047	-0.039	0	0	0	0.039
Number of Failures in month	0.008	0.008	0	0.008	0.0156	0.008	0.008	-0.008	0	0.008	0.0156	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	-0.019	-0.019	-0.019	-0.019	-0.019
Non In-depth Queries - Max Time to resolve 95% of all queries	0.016	-0.016	-0.019	-0.019	0	-0.016	0	0.046	-0.016	0	-0.016	0
Administrative Queries - Max Time to resolve 95% of all queries	-0.016	-0.019	-0.016	-0.019	0	-0.016	-0.016	-0.019	0	0.046	-0.019	-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	-0.02	-0.06	-0.02	0.03	-0.02	-0.01	-0.03	0.01	0.02	-0.01	0.04
Quarterly Service Credits:			-0.05			-0.01			-0.03			0.05

Table 3

1.2 No. of Registered Users

The current position at the end of the quarter is that there are 460 registered users of the CSAR Service.

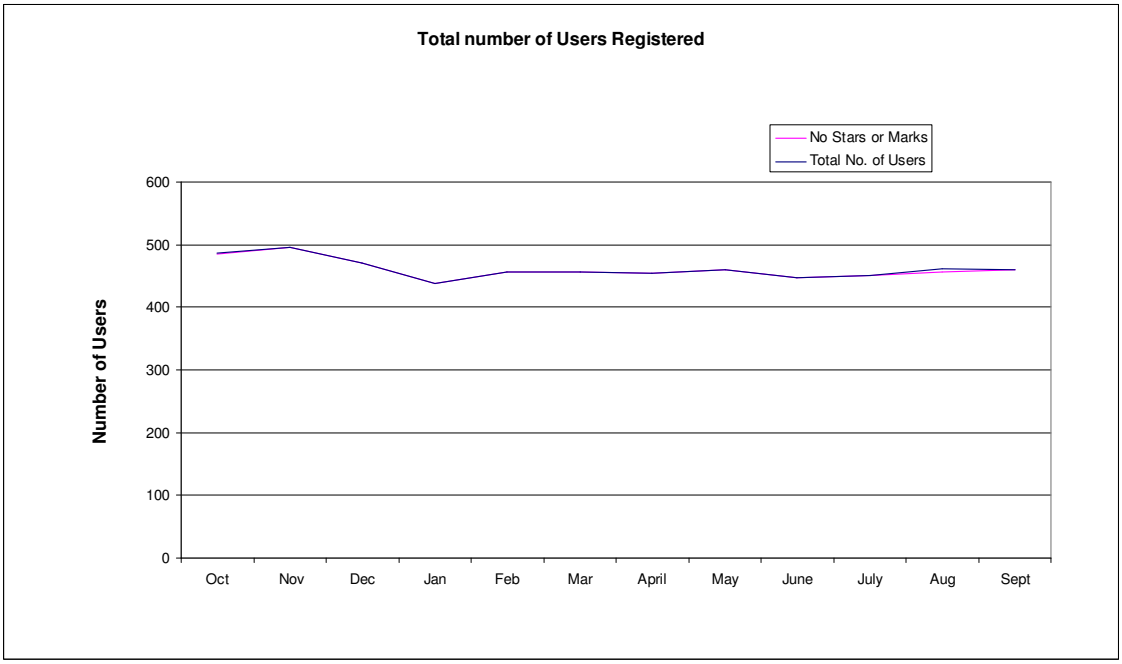

1.3 Service Quality Tokens

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

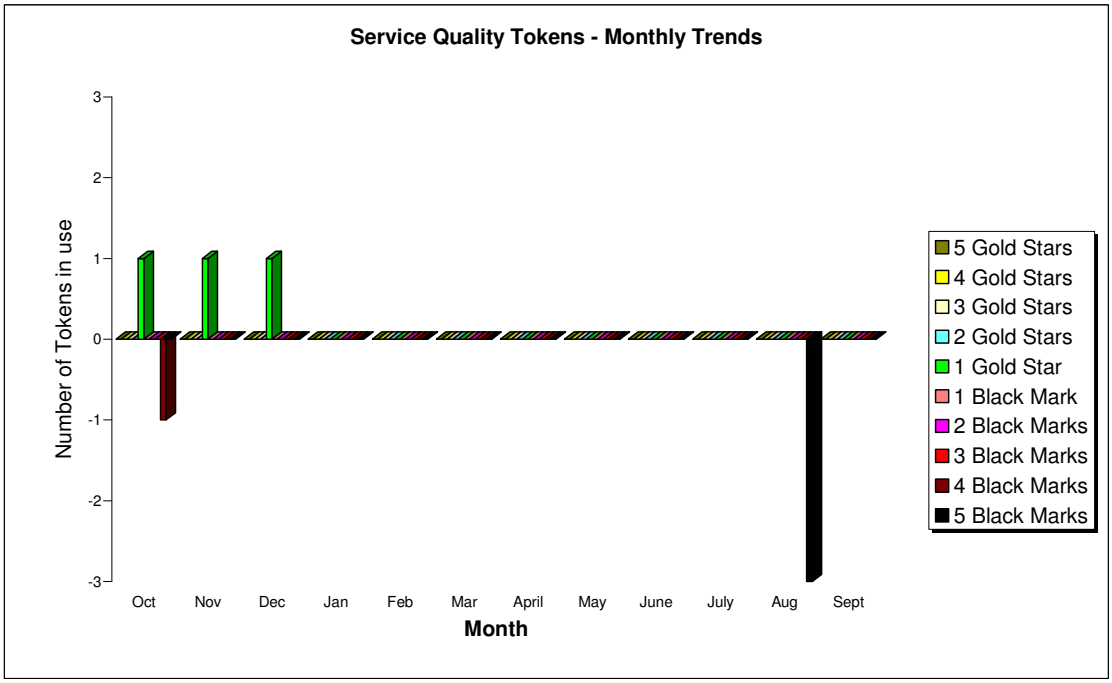

Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

At the end of the quarter no gold stars or black marks had been 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.

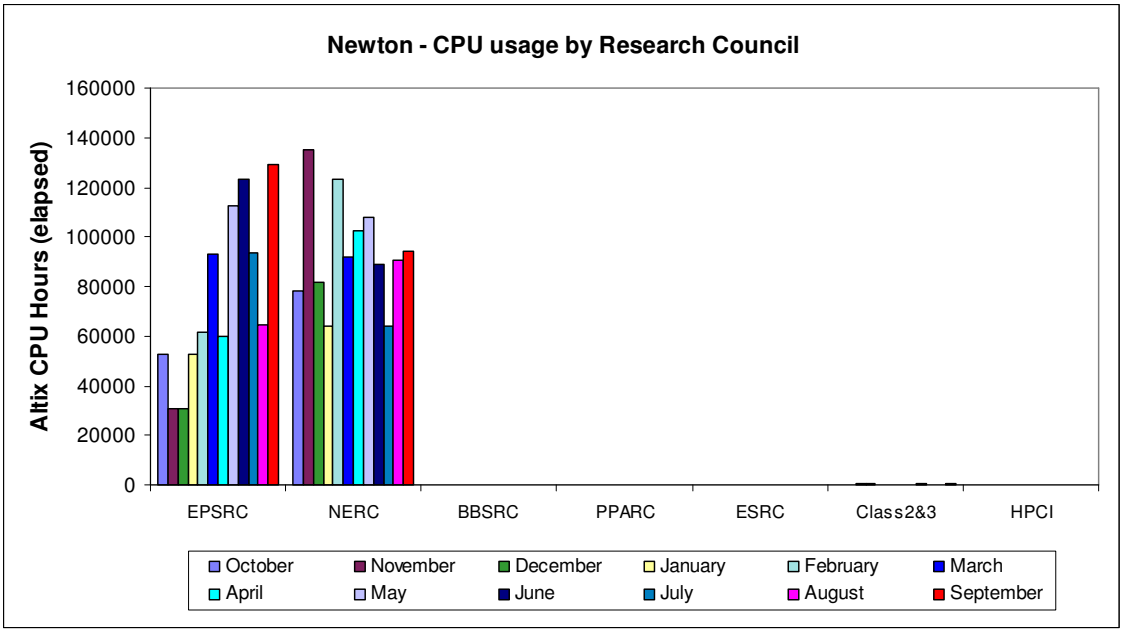

Usage of the SGI Altix 3700 system Newton is shown by Research Council in the above graph.

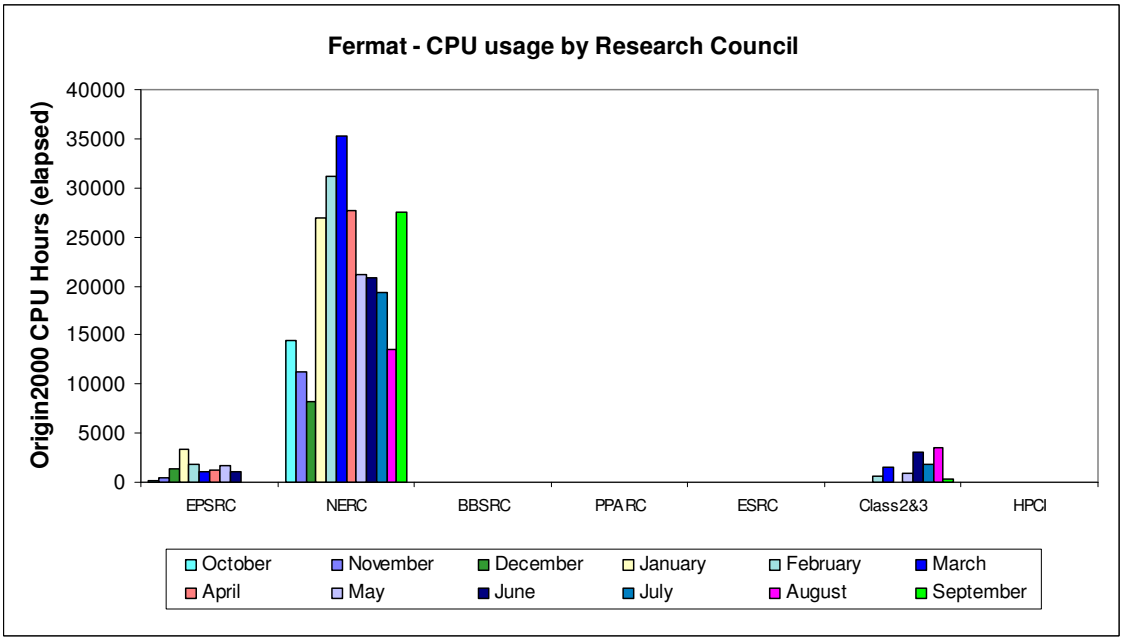

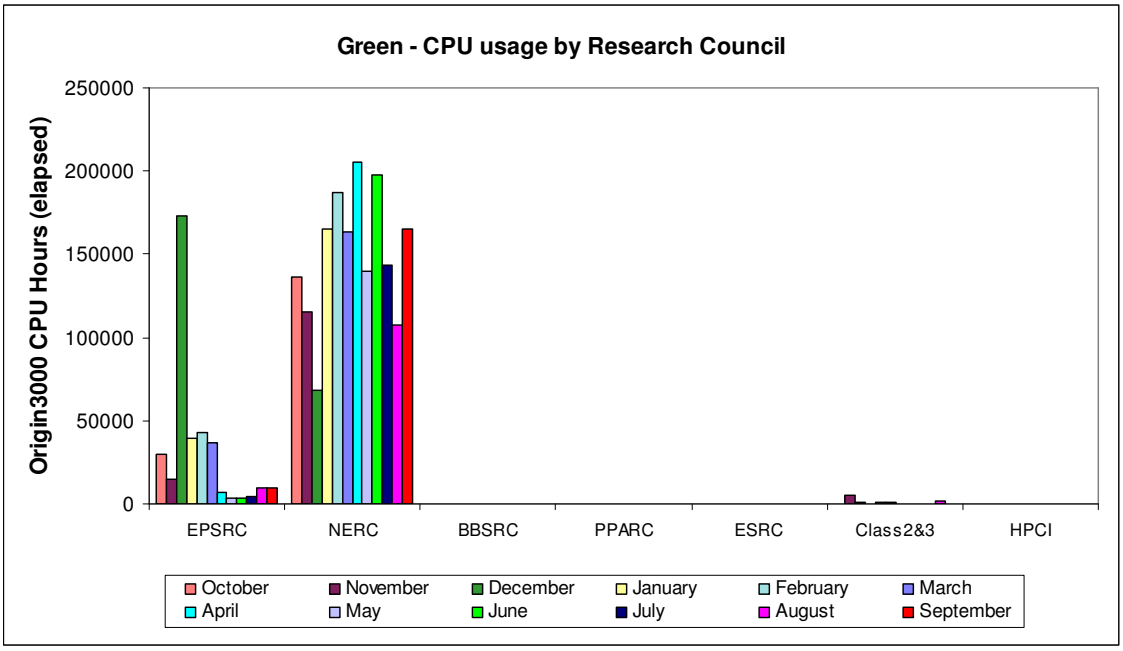

Usage of the two batch SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service in the preceding two charts.

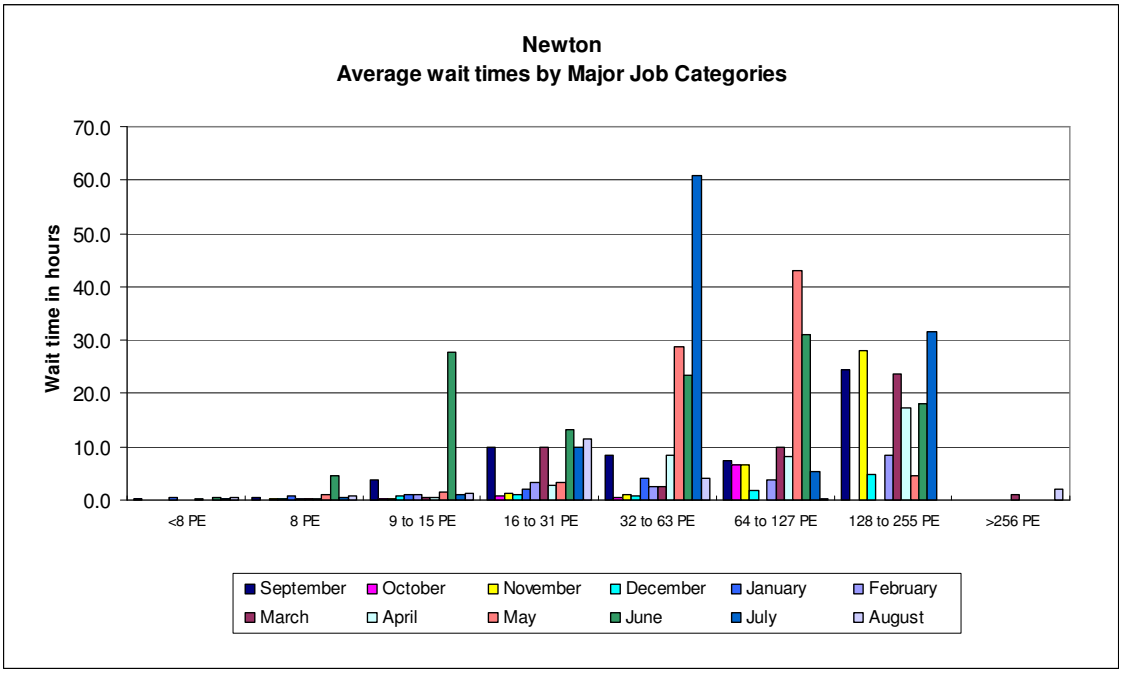

The above chart, and the one below, shows the wait time trend in hours on the Altix 3700 Newton.

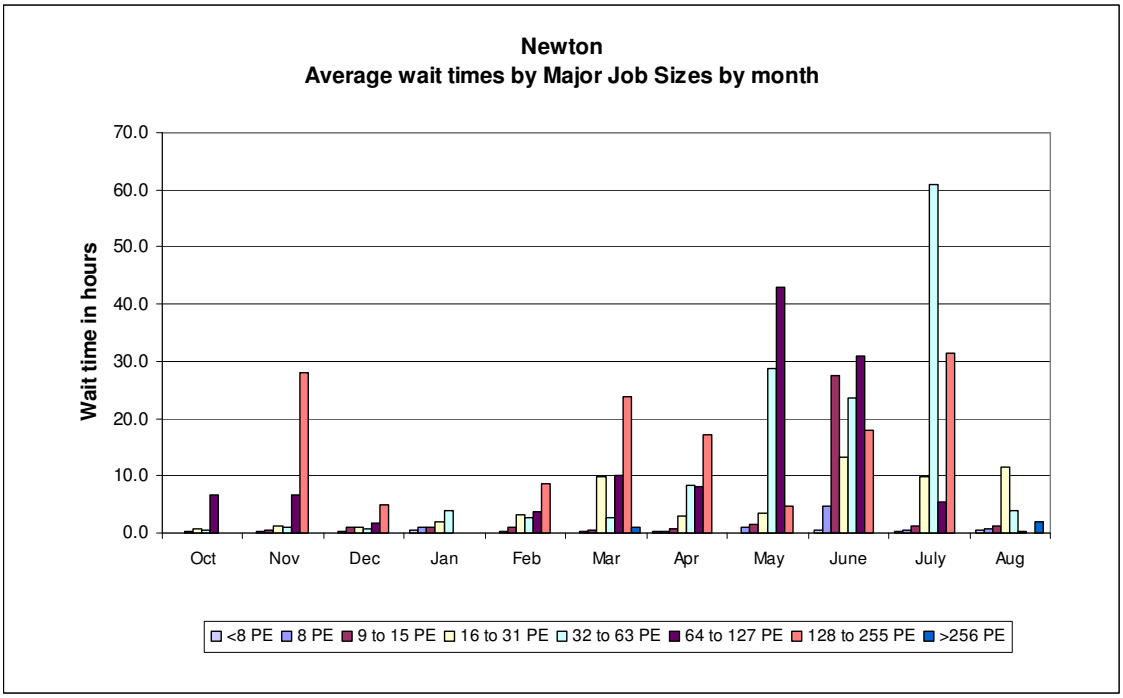

The following two charts show average wait times in hours for the quarter on the Origin 3000 Green.

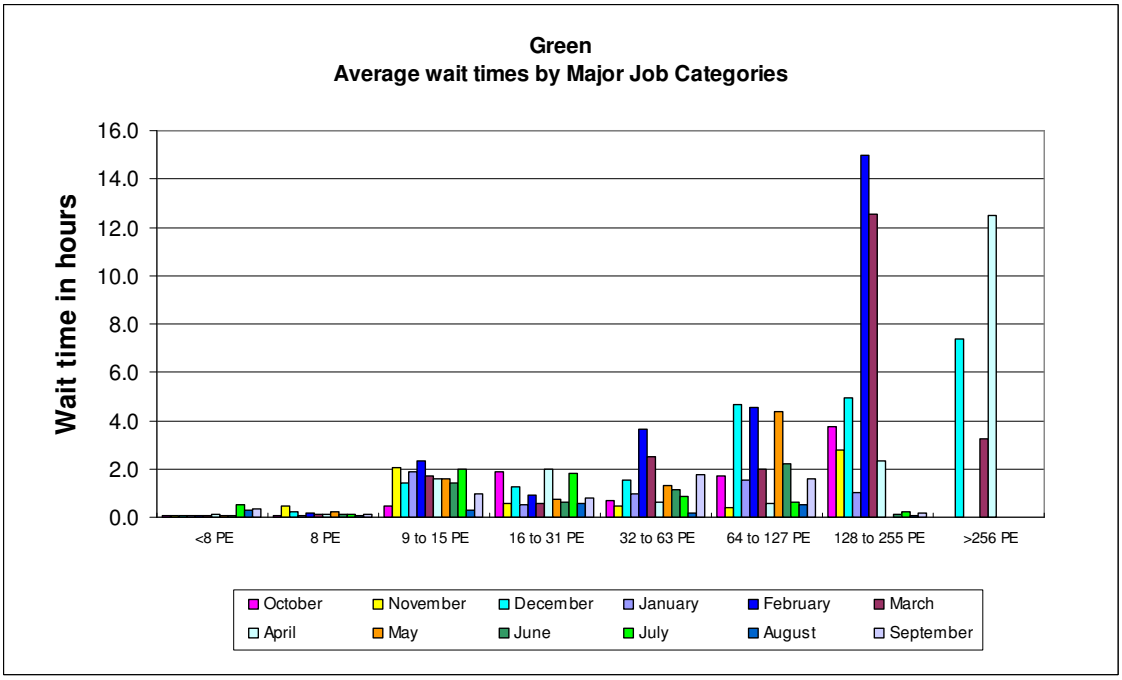

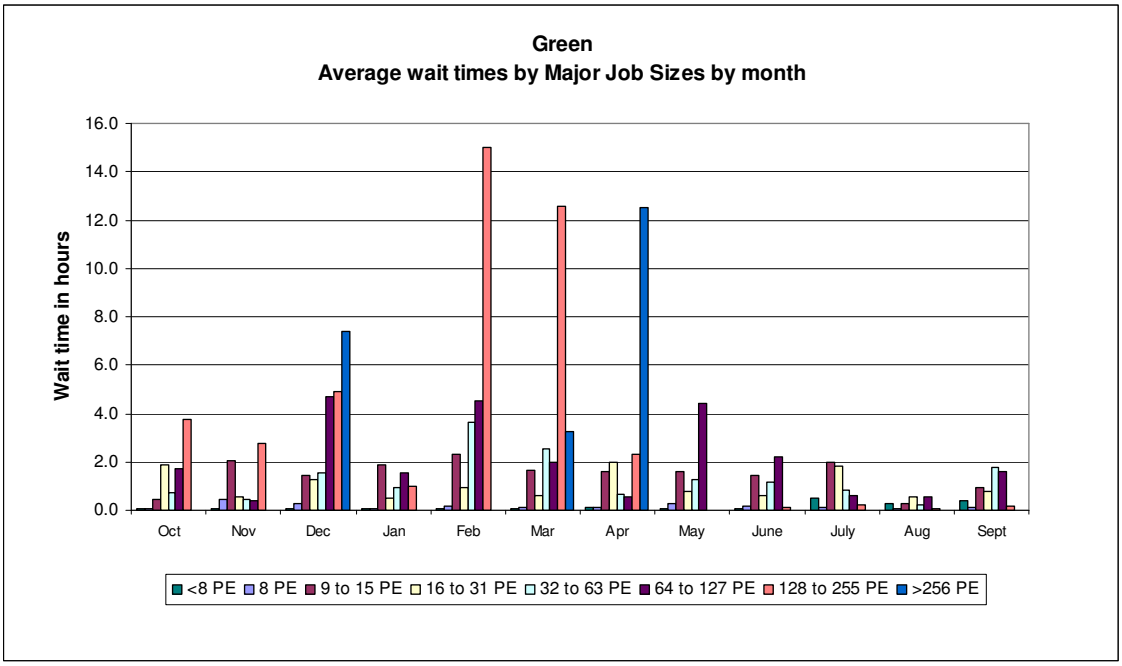

The next series of four charts show the percentage CPU time utilisation by the major job categories on the Green, Fermat and Newton systems for the 3rd quarter 2005.

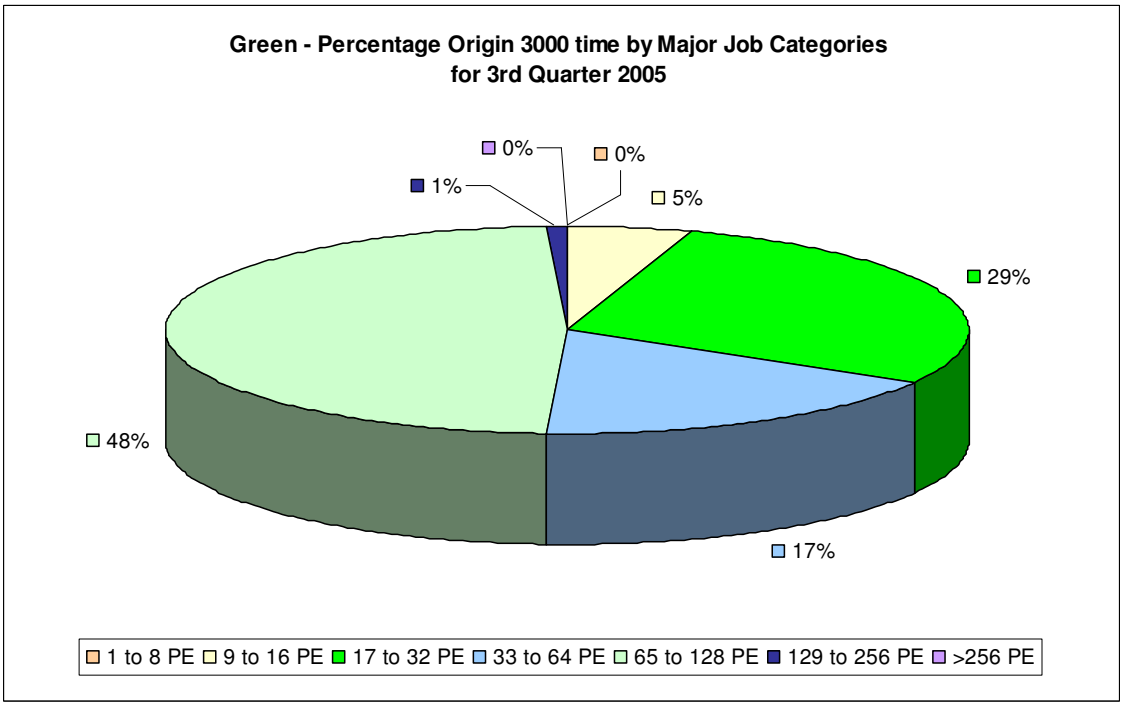

During this quarter there has been a good spread of work across the mid-range PEs on Green.

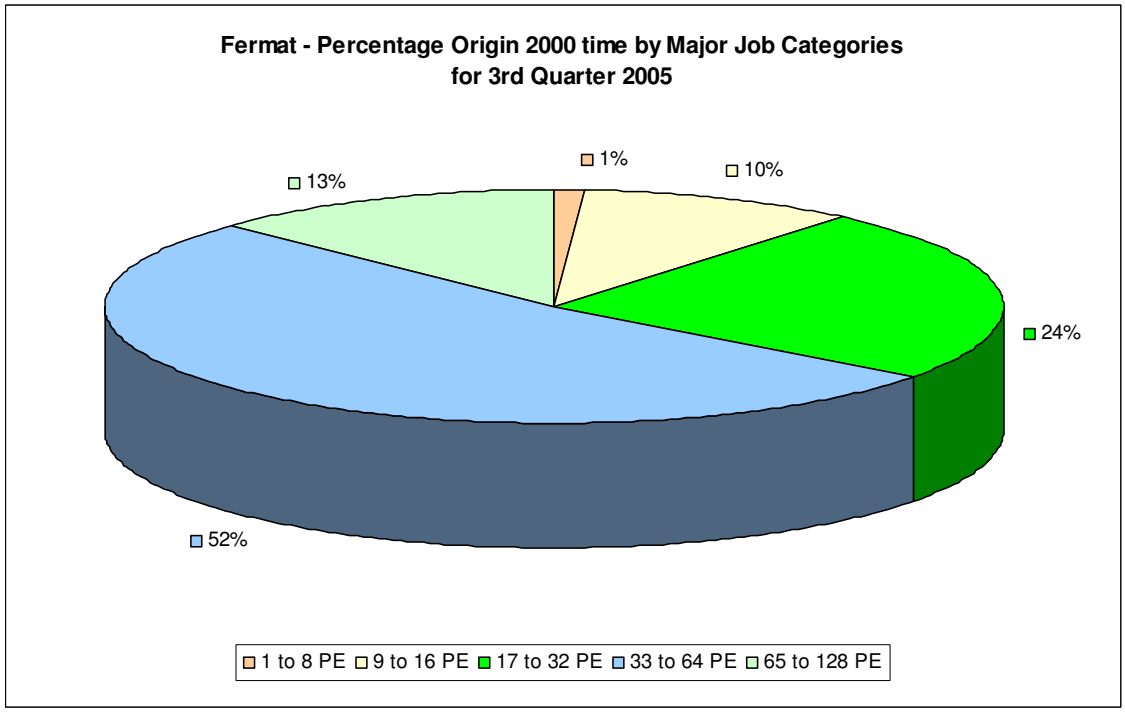

The workload on Fermat during this quarter was concentrated primarily in the 33 to 64 PE range.

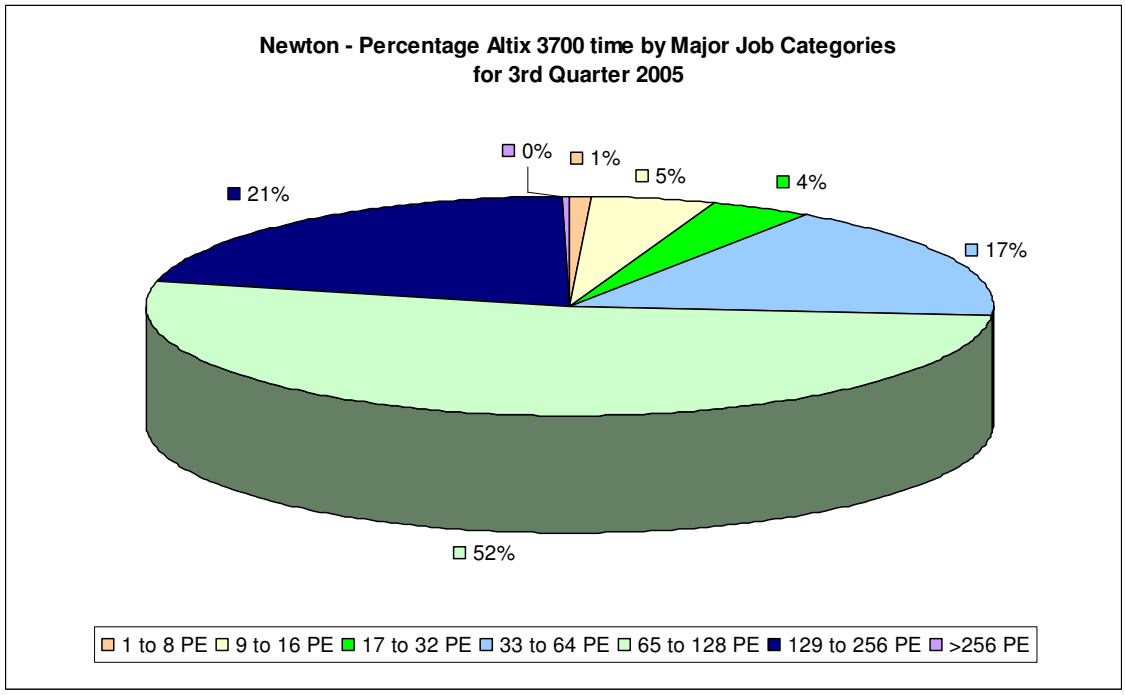

During the third quarter, the majority of work on Newton was in the 65 to 128 PE range.

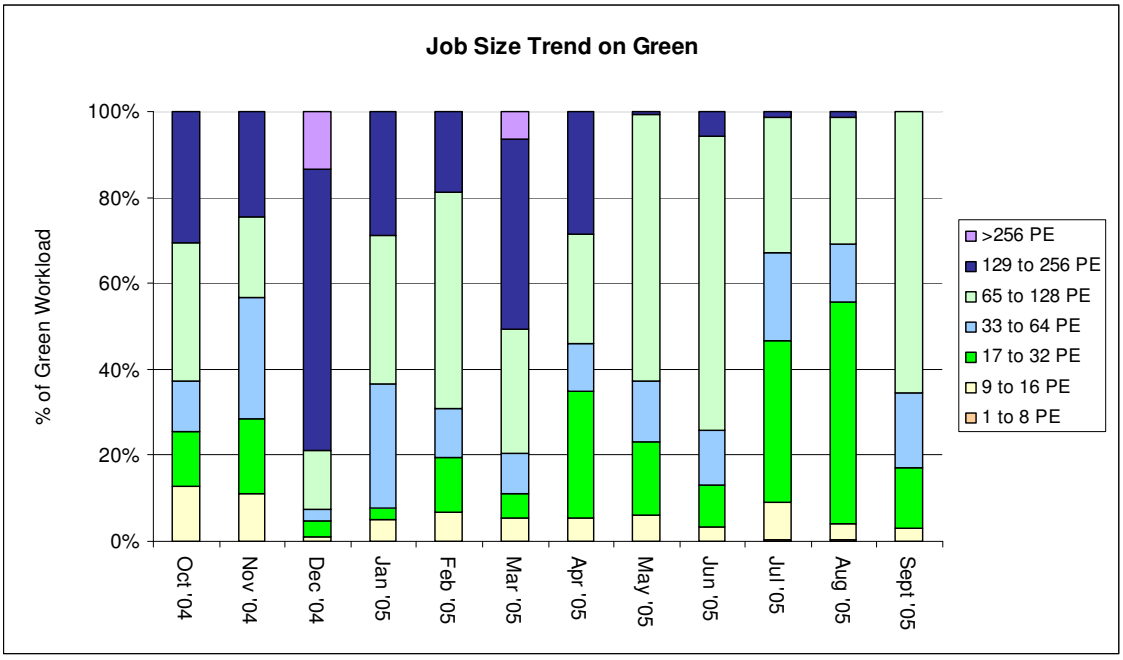

Usage on Green tended more to be in the mid-range PEs during this quarter.

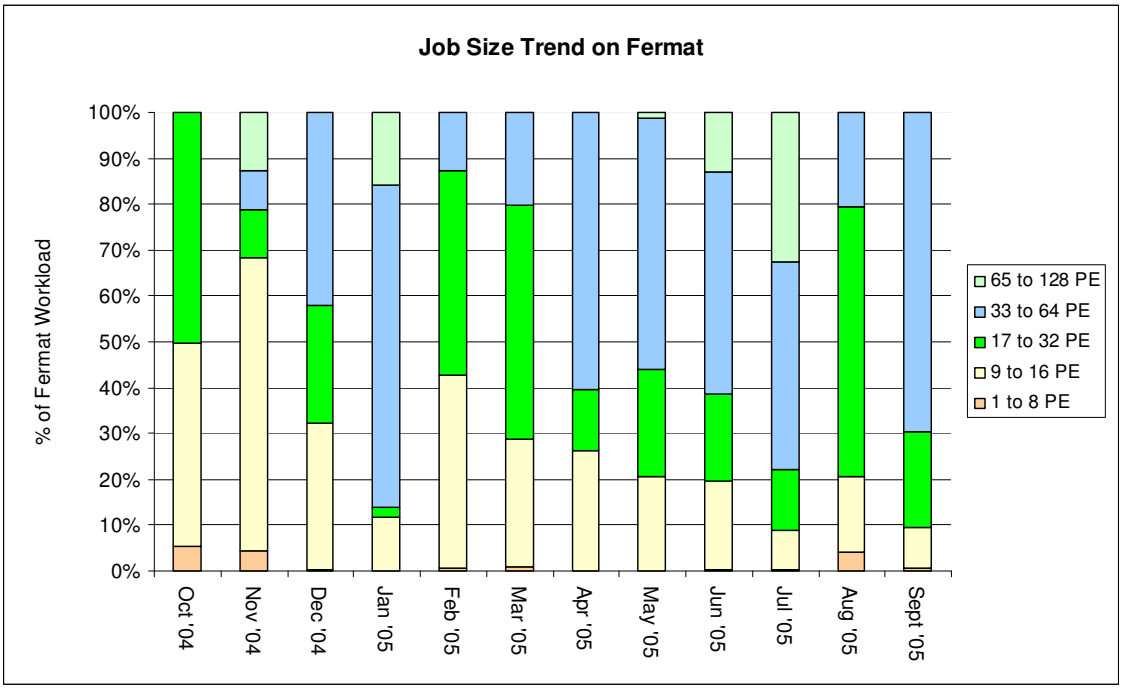

The workload on Fermat for this quarter was spread fairly evenly across the mid-range PEs.

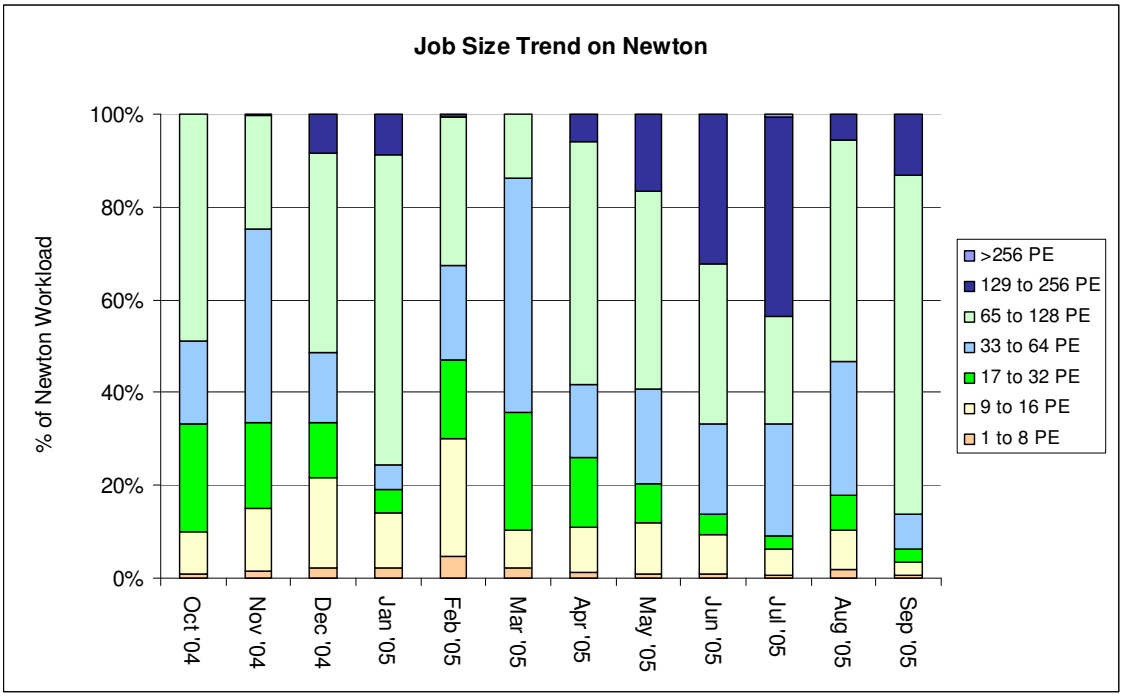

During the third quarter usage on Newton was concentrated mainly in the mid- to high-range PEs.

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.

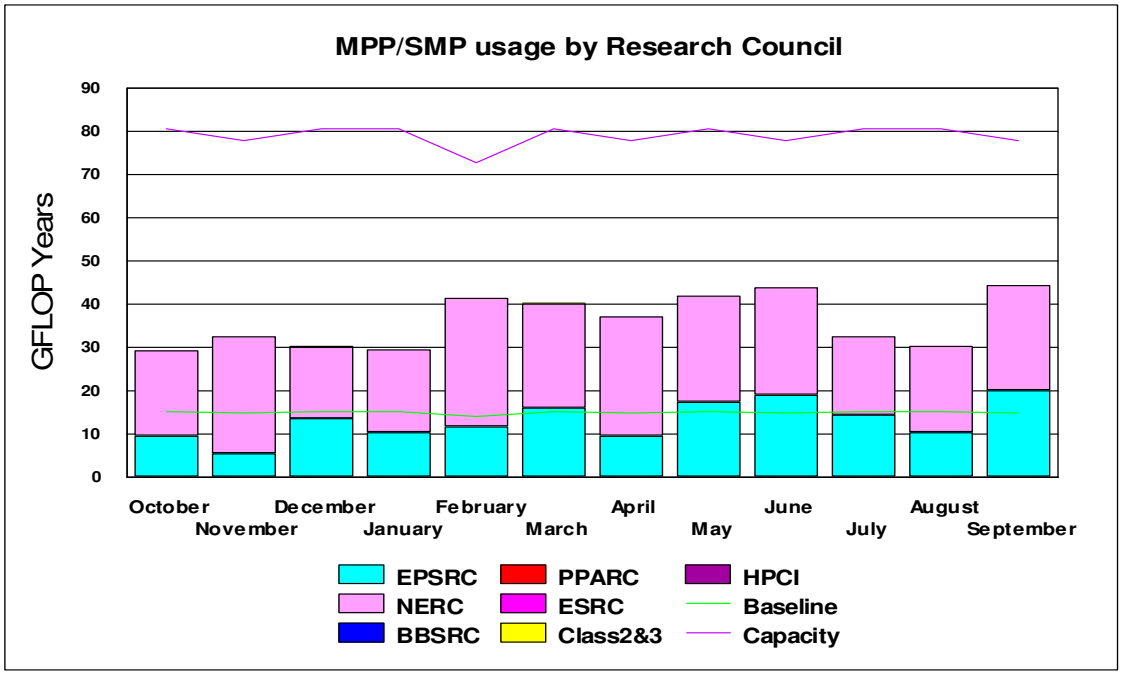

The graph below shows the same service utilisation by major system.

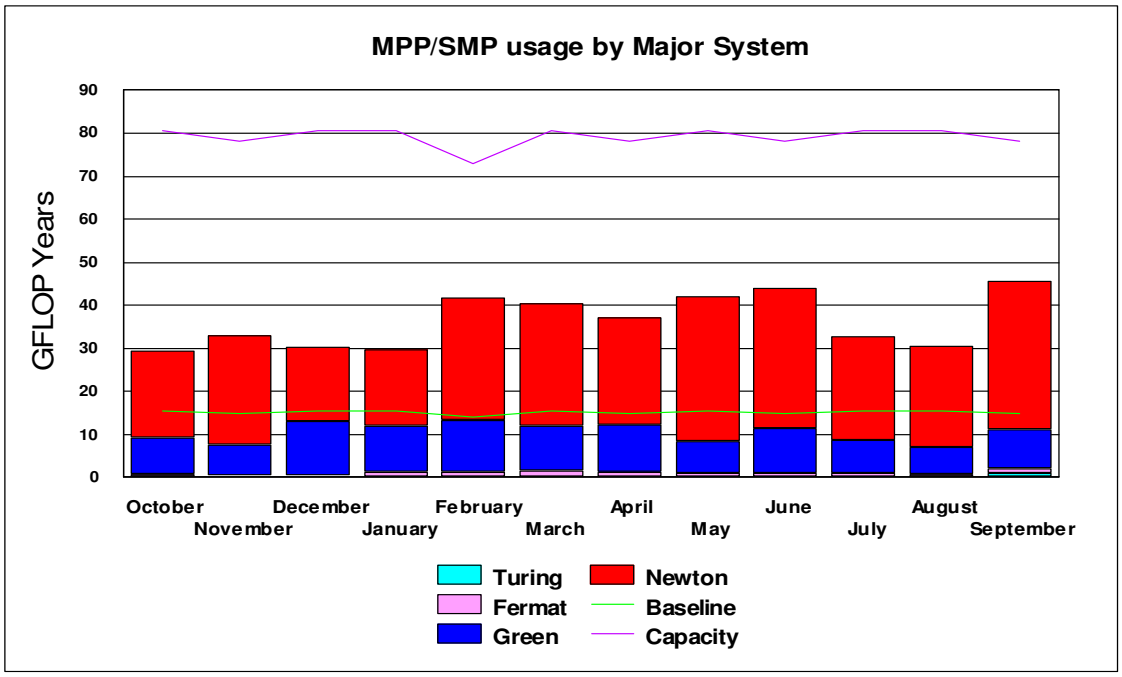

The next chart shows the historic interactive usage of the Origin 300 system Wren. Eight of the higher speed 500Mhz CPUs in Wren deliver the baseline capacity equivalent to that which was previously available on the Origin 3000 system Fermat for interactive usage.

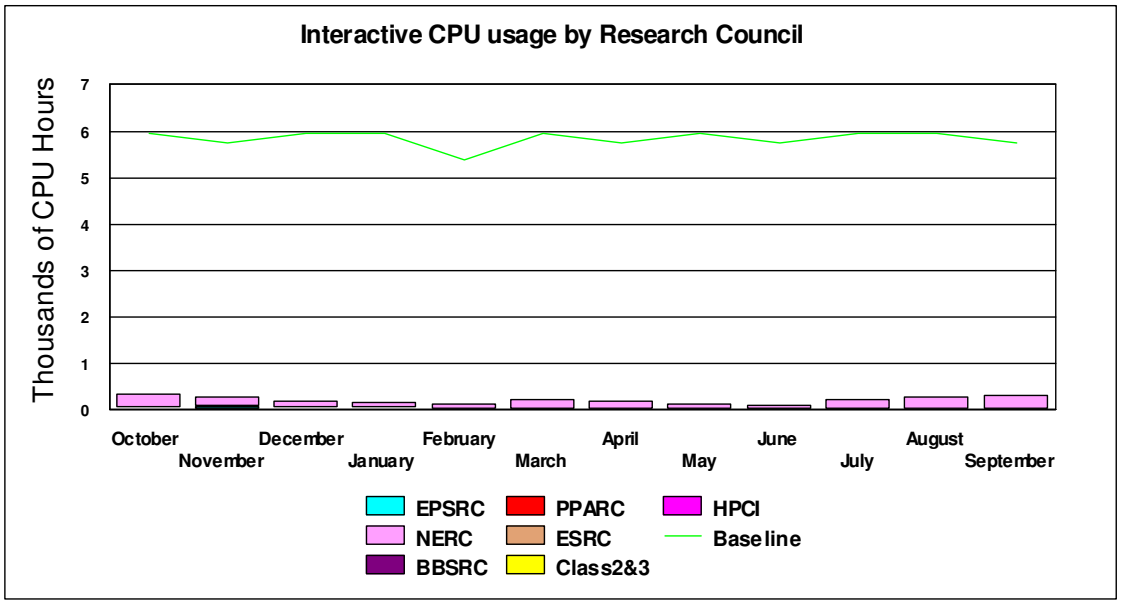

The following two charts detail the historic usage of the Origin 3000 system (Green) and the Altix 3700 system (Newton).

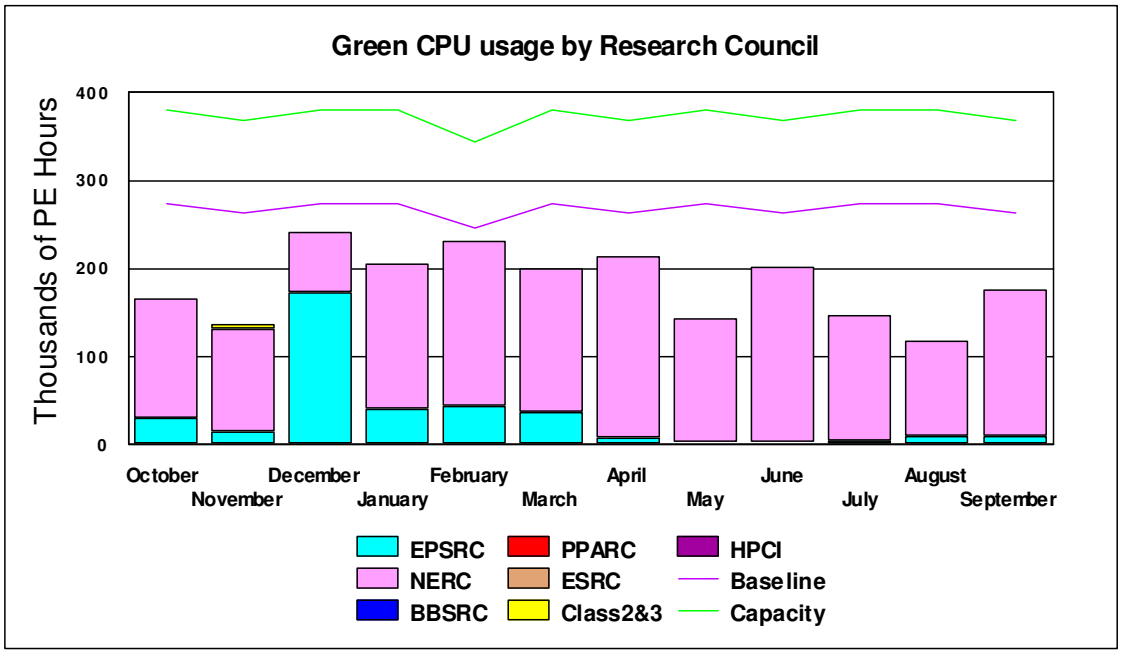

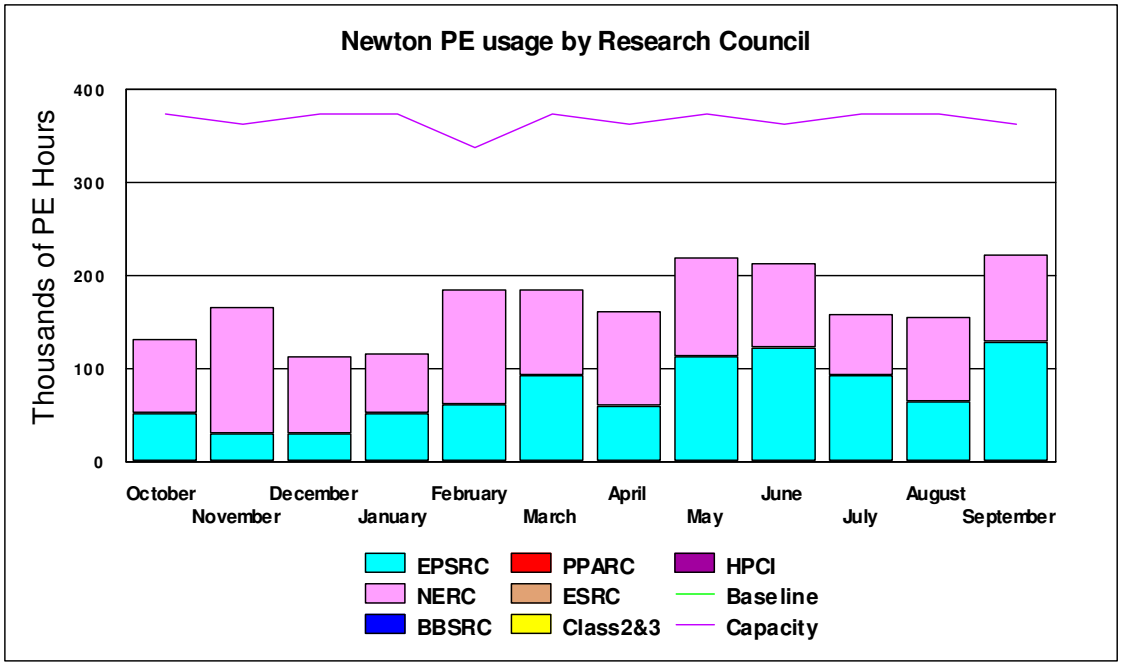

The next series of graphs illustrates the usage of the Medium Performance disk and HSM/tmp resources of the system.

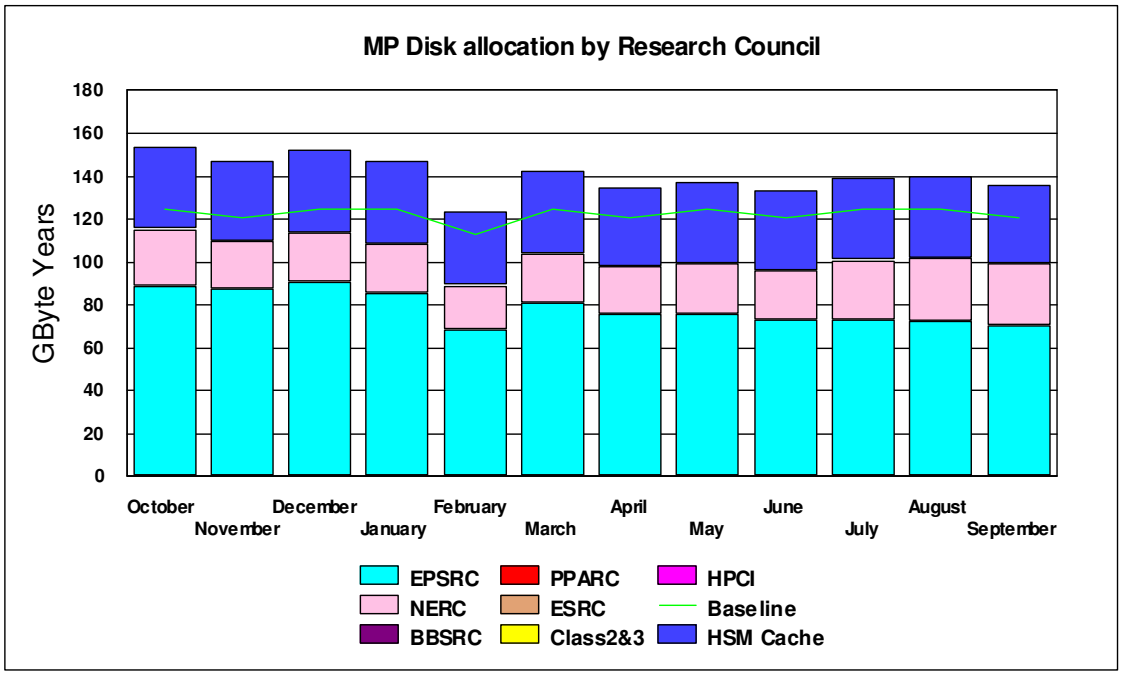

The above graph illustrates the historic allocation of the Medium Performance Disk on Fermat and the SAN.

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.

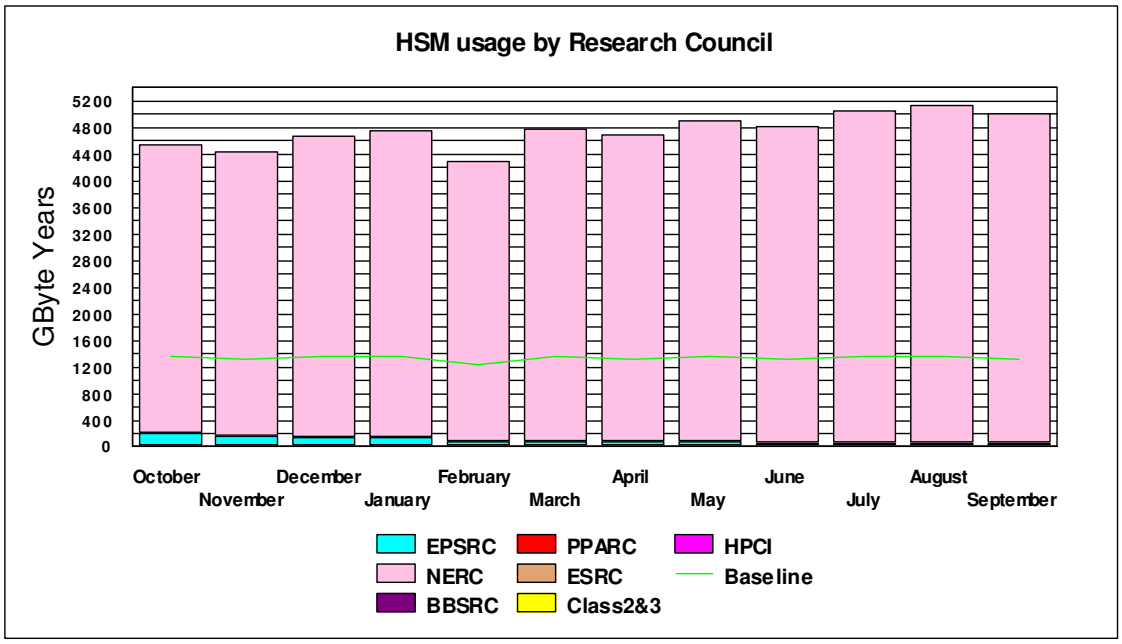

2.2.2 Guest System Usage

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 third quarter of this year.

Service Tokens Refunded: Quarter 3 2005 Usage							
System	Consortia						Total
	csn001	cse075	cse121	cse086	cse133		
Green 256+ PEs							0
Green 384+ PEs							0
Newton 128+ PEs	539.14	108.64	1331.63		227.45		2206.86
Newton 192+ PEs		0.01		534.94			534.95
Total Tokens	**539.14**	**108.65**	**1331.63**	**534.94**	**227.45**		**2741.81**

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 third quarter of 2005, with usage exceeding baseline.

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

Towards the end of July the two 256PE Newton nodes were not only successfully combined into one physical 512 processor node, but at the same time the system was converted into a 512 PE Single System Image. This greatly increases the largest job size that can be run on Newton, and is expected to bring additional stability to the system.

Issues

There were several issues with stability on Newton, most notably throughout September, the majority of which were finally tracked down to being a rogue user code which, every time it crashed, also crashed the system. Remedial work was undertaken by the applications support team in conjunction with SGI to ascertain the exact problem with the code.

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 project applications were received.

3.2 New Projects

4 new projects were started, with a total of 1600 tokens allocated.

3.3 Finished Projects

5 projects finished

3.4 Queries

A total of 107 CSAR queries were dealt with:

- 71 non-in-depth
- 31 user registration and admin
- 5 in-depth:

3.5 Service Quality Tokens

No gold stars or black marks were awarded during this period.

3.6 CSAR Focus

The 14th edition of CSAR Focus was published and is also available online at www.csar.cfs.ac.uk/about/csarfocus/focus14/focus14.pdf.

4 Scientific Application Support Services

4.1 Training and Education

The following courses were delivered:

- Kevin Roy delivered most of the 4 day course "Programming and Optimisation on the Cray XD1". Much development work was done on the presentation.
- Jon Gibson delivered the course "Introduction to MPI",as part of the Cray XD1 training course at Aston University. Some development work was done on the presentation and exercises.
- Craig Lucas delivered a new course "Introduction to ScaLAPACK on the XD1", again as part of the Cray XD1 training course at Aston University. Much development work was done on the presentation.

4.2 Consortia Support/Software

4.2.1 MPICH-G2

MPICH-G2 1.2.6 was compiled on top of the MPI flavour of Globus on Newton and Green. This was built with vendor MPT for intramachine communication and TCP for intermachine communication. MPICH-G2 was tested by successfully running ring-type programs within Newton and across Newton and Green. Web pages were written to provide instructions for using MPICH-G2 on Newton (http://www.csar.cfs.ac.uk/user_information/grid/grid-middleware.shtml).

4.2.2 Tomcat

Optimisation work on Tomcat code for Glenn Carver of Cambridge was completed, work is continuing on documentation and write-ups.

4.2.3 Reality Grid

Tim Robinson worked closely with the RealityGrid consortium, incorporating the UK-arm (SPICE project, using NAMD with visualisation and steering) and the US partners' projects VORTONICS and Nektar. This work has been at a number of levels, ranging from software (source compilation, installation and testing of NAMD 2.6b1 on Newton and Green), to middleware (MPICH-G2 installation and testing on Newton and Green), to the facilitation of network requirements (UKLight connections on Newton and Green), to the practical issues of cross-site coordination (implementing policies for advanced reservations) and documentation.

4.2.4 h2mol

Work continued with Ken Taylor's group to get large h2mol jobs to run. (Ken Taylor and Dan Dundas).

4.2.5 Globus

Globus 4.0.1 pre-web services MPI flavours, both vendor cc and gcc, were installed on Newton.

4.2.6 Namd

NAMD 2.6b1 was installed on Newton. This upgrade worked for RealityGrid where 2.5 had failed; names, it doesn't crash when using more than 64 processors, and doesn't lead to meaningless/incorrect simulations when run on less than 64 processors.

4.2.7 Siesta

SIESTA 1.3 was installed on Newton.

4.2.8 PUM

Testing was finished of the PUM v6.0 on Newton.

4.2.9 Amazon SC

Work has continued for David Ingram's group at MMU. This has proved more problematic than originally anticipated due to the method they are using to solve their equations.

There is now an MPI version of their code available for use, but some modifications required to the data are currently awaited from the group. Unfortunately the group have not been able to achieve much with the addition of OpenMP to the code, but hope to improve this by rewriting their matrix solver routines..

4.2.10 FFTW

FFTW 3.0.1 was installed on the Origins.

4.2.11 Pam Crash

Pam Crash was installed on Newton.

4.2.12 Python

Python 2.4.1 was installed on Newton.

4.2.13 Gromacs

Installation and testing was carried out on single- and double-precision versions of GROMACS 3.2.1 on Newton.

4.2.14 Beagle

The Beagle Generic Programming Package was installed on Newton.

4.2.15 Boost

The C++ Boost libraries were maintained on Newton.

4.2.16 Visualisation

Visualisation support was given to several Class 3 projects.

5 Collaboration and Conferences

5.1 MRCCS Projects

5.1.1 MRCCS Seminar Series

The Access Grid based seminar series on HPC and visualisation related topics continued on Fridays, including:

- Peter Hofstee, IBM – Cell Processors: Motivation, Architecture, Design, Programming and Applications

5.2 Events

Events attended include:

- FPGA conference at NCSA, 11th-13th July
- Liquid Computing web-based seminar, 13th July
- First meeting of the NAFEMS Iberian Steering Group, 18th July
- CCP5 Conference: "Computational techniques and applications for materials modelling", Keele University, 30th August-2nd September
- CAS2K5, Annecy, France, 11th-15th September
- NAFEMS seminar: "Quality and Reliability of CFD Simulations II: Industrial Challenges", Nottingham, 13th September
- All Hands Meeting, Nottingham, 19th-22nd September
- Cray Technical Workshop, Lugano, 20th-22nd September

5.3 Publications

Margetts, L., Smith, I.M, Leng, J. 'Simulating Dinosaur Trackway Formation', COMPLAS, Barcelona, September 2005

Margetts, L., Ford, R., Porter, A. and Haines, R. 'RapidFire: A novel grid enabled application for interactive finite element analysis', TCN CAE. Margetts was awarded EU funding to attend, and as a result is now a member of the Marie Curie Training Network in the field of scientific computing (2005-2007). Details can be found at http://www.eu4ax.net .

6 Added Value Services

6.1 CSAR

An online form was produced for requesting advanced reservations for the CSAR machines. This enabled cross-site runs at iGrid and is being used to facilitate testing for SC2005.

6.2 SC2005

Preparations continued for Manchester Computing's exhibition at this year's Supercomputing Conference, to be held in Seattle in November.