CSAR Service

Consolidated Management Report

1stQuarter 2005

Management Summary

This is the consolidated Management report for the first quarter 2005 of the CSAR HPC facility for UK Academia and Industry, which enables World-Class research and development.

The number of users has grown to a total of 457 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 ranges 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 1stquarter 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											
	April	May	June	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	Mar
HPC Services Availability												
Availability in Core Time (% of time)	89.39%	94.21%	97.49%	97.97%	100%	99.52%	97.08%	98.50%	99.50%	97.37%	97.85%	97.85%
Availability out of Core Time (% of time)	91.90%	99.73%	97.85%	100%	99.2%	99.80%	98.67%	98.78%	99.2%	99.73%	99.5%	99.80%
Number of Failures in month	5	3	4	2	2	2	3	4	2	3	5	4
Mean Time between failures in 52 week rolling period (hours)	1741	1036	683	607	529	471	396	340	315	284	248	223
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	<2	<3	<2	<2	<1	<2	<3	<1	<0.5	<0.5	<2	<1
Administrative Queries - Max Time to resolve 95% of all queries	<0.5	<0.5	<1	<0.5	<0.5	<1	<1	<0.5	<1	<0.5	<2	<1
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											
	April	May	June	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	Mar
HPC Services Availability												
Availability in Core Time (% of time)	0.195	0.195	0.078	0.078	-0.058	-0.039	0.078	0.039	-0.039	0.078	0.078	0.078
Availability out of Core Time (% of time)	0.039	0	0.078	-0.047	0	-0.047	0	0	0	-0.039	0	-0.047
Number of Failures in month	0.0004	0.0	0.008	0	0	0	0.008	0.008	0	0.008	0.0004	0.008
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	0.016	0	0	-0.016	0	0.016	-0.016	-0.019	-0.019	0	-0.016
Administrative Queries - Max Time to resolve 95% of all queries	0	-0.019	-0.016	-0.019	-0.019	-0.016	-0.016	-0.019	-0.016	-0.019	0	-0.016
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:	0.09	0.07	0.05	-0.02	-0.07	-0.07	0.02	-0.02	-0.06	-0.02	0.02	-0.02
Quarterly Service Credits:	0.09			0.11			-0.12			-0.09		

Table 3

1.2 No. of Registered Users

The current position at the end of the quarter is that there are 457 registered users of the CSAR Service.

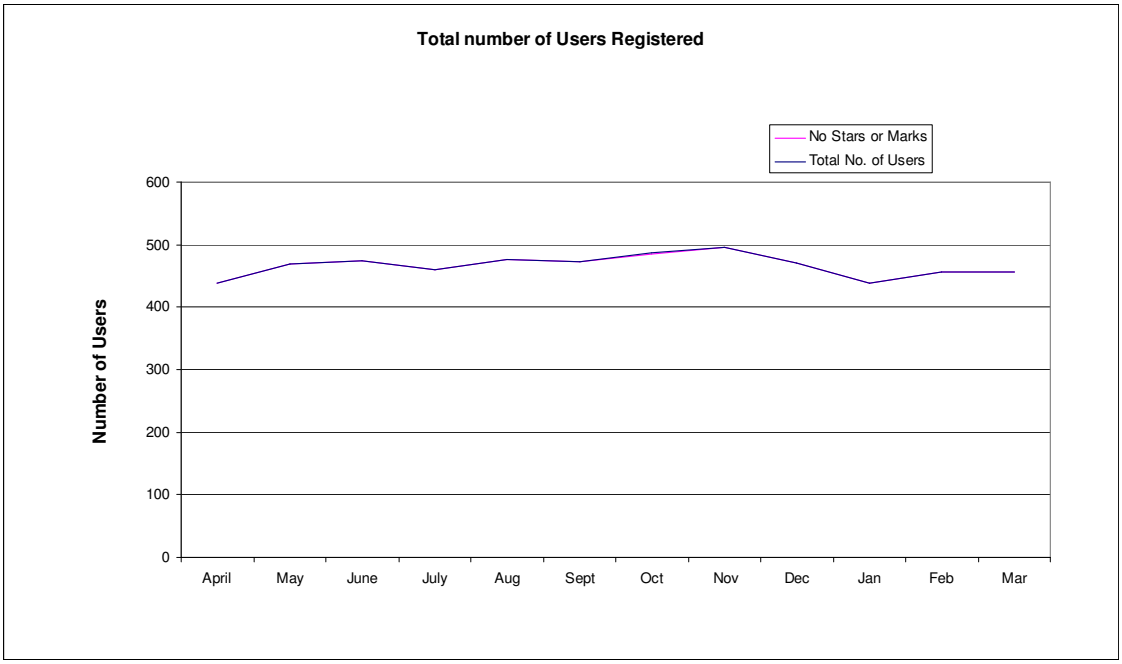

1.3 Service Quality Tokens

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

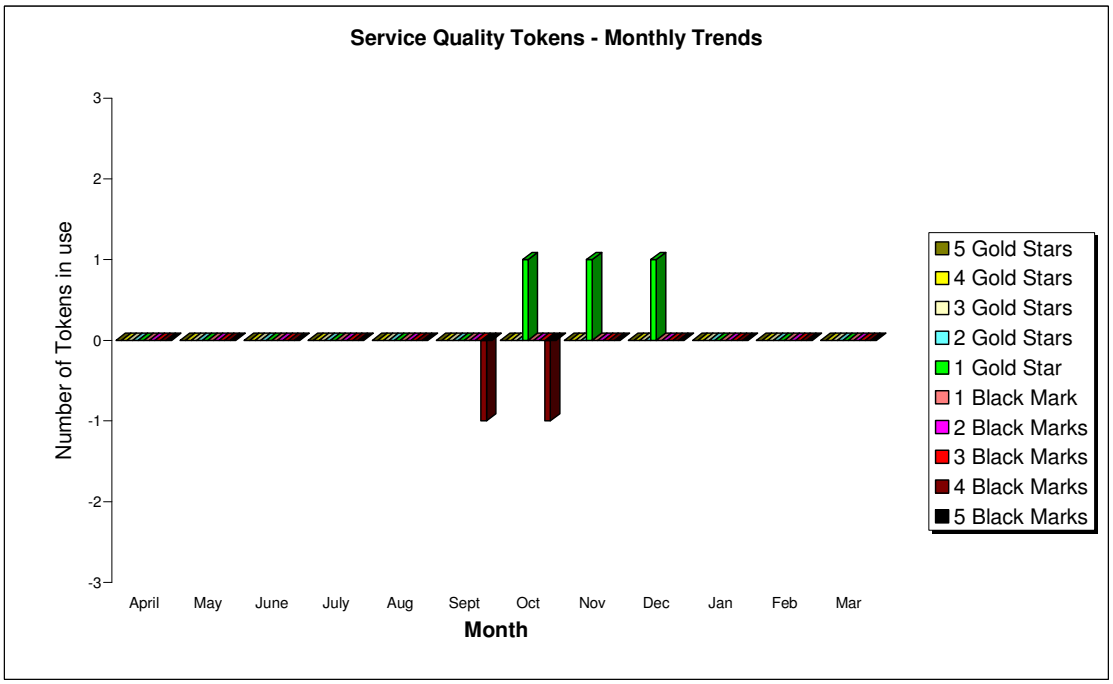

Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

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.

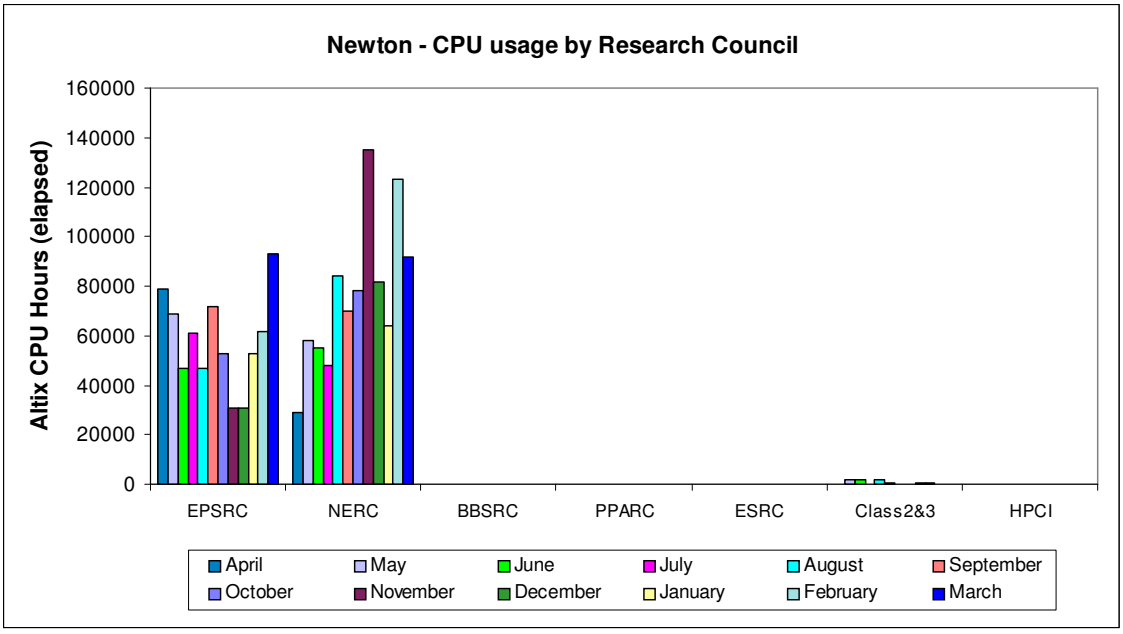

Usage of the SGI Altix 3700 system Newton is shown by Research Council in the above graph.

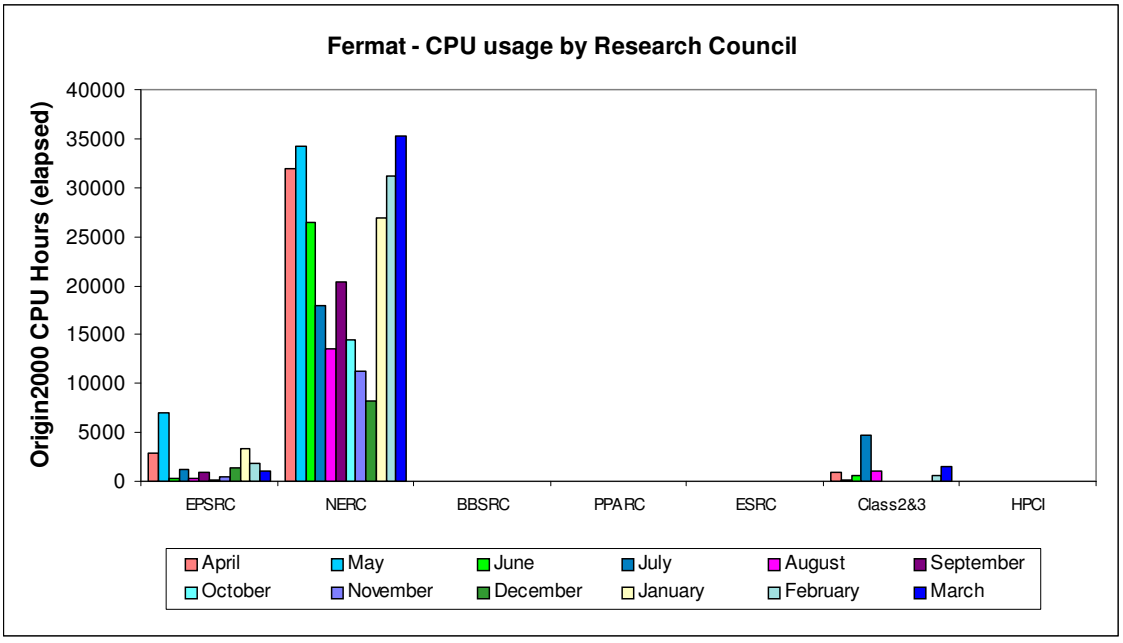

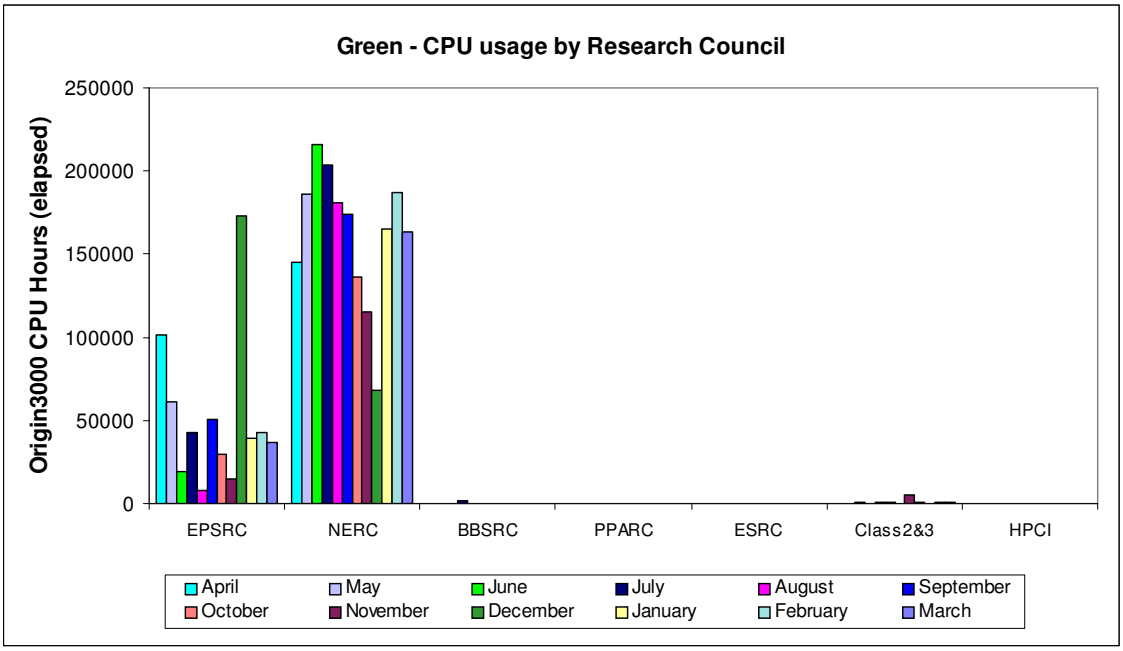

Usage of the two batch SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service in the preceding two charts.

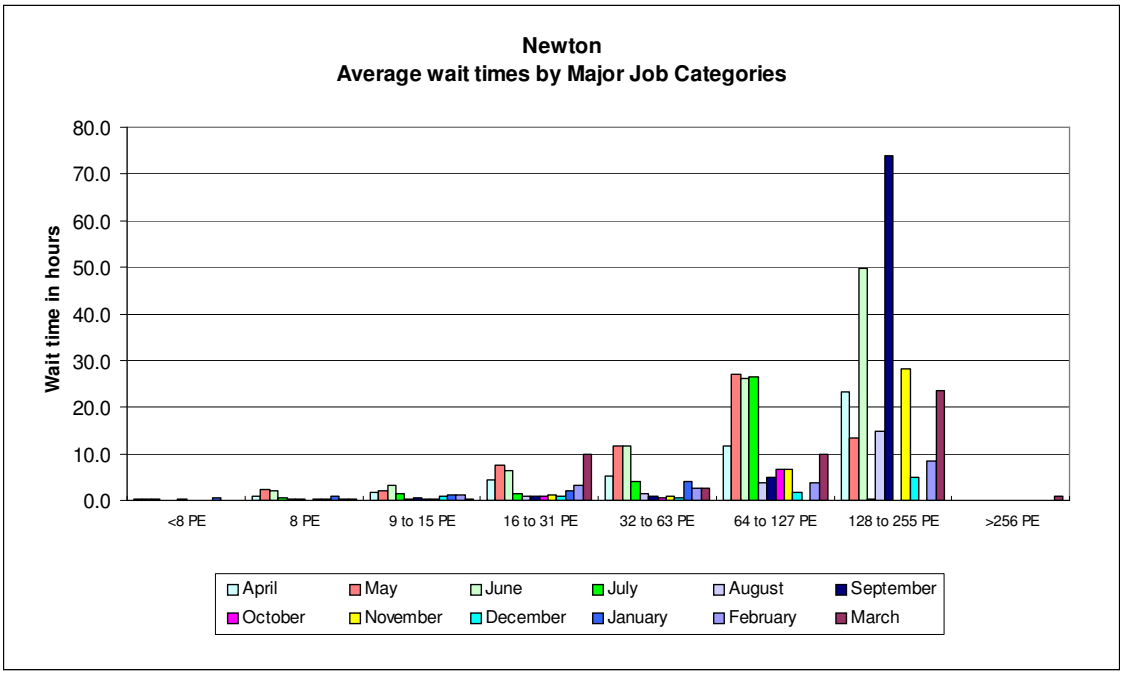

The above chart, and the one below, shows the wait time trend in hours on the Altix 3700 Newton.

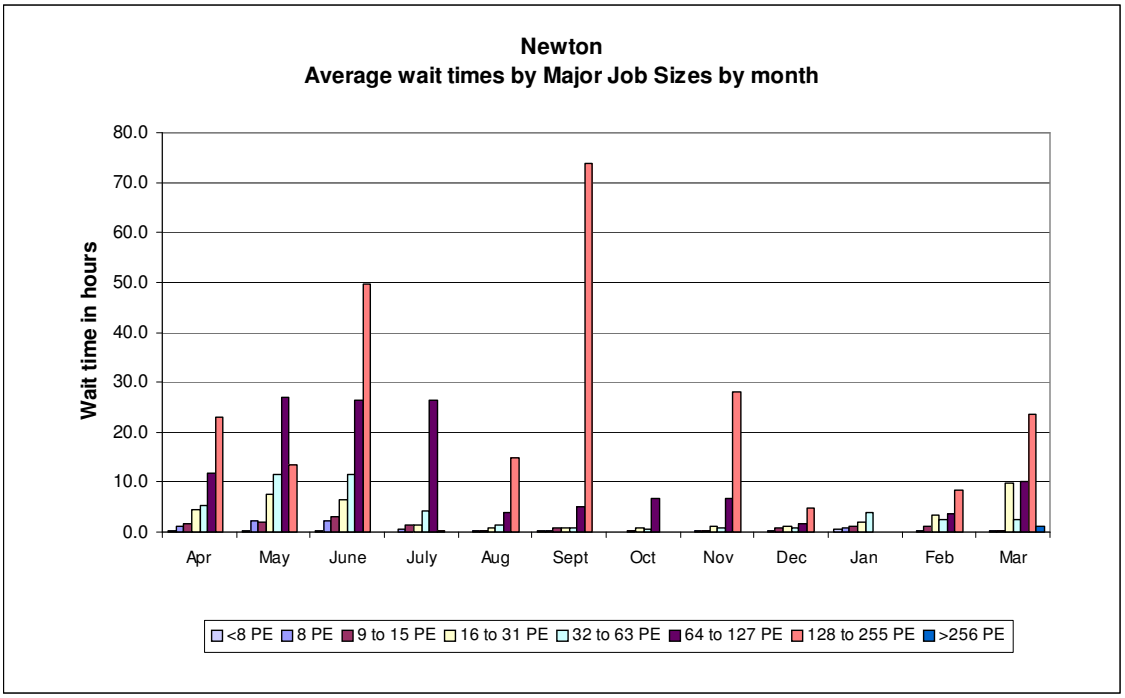

The following two charts show average wait times in hours for the quarter on the Origin 3000 Green.

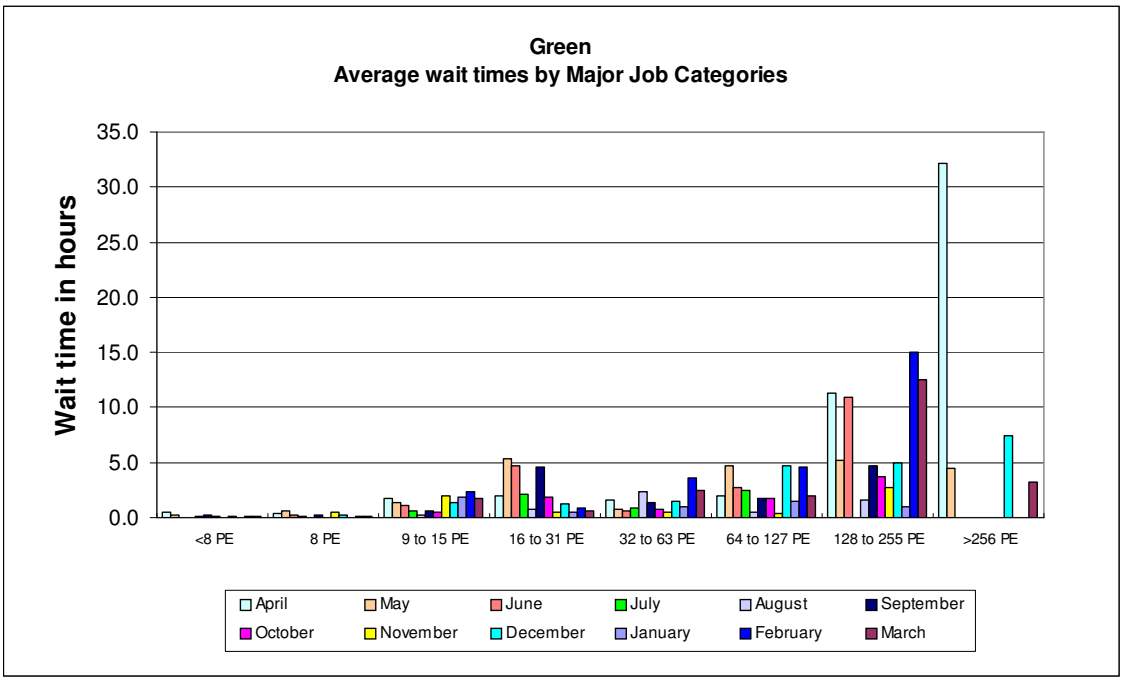

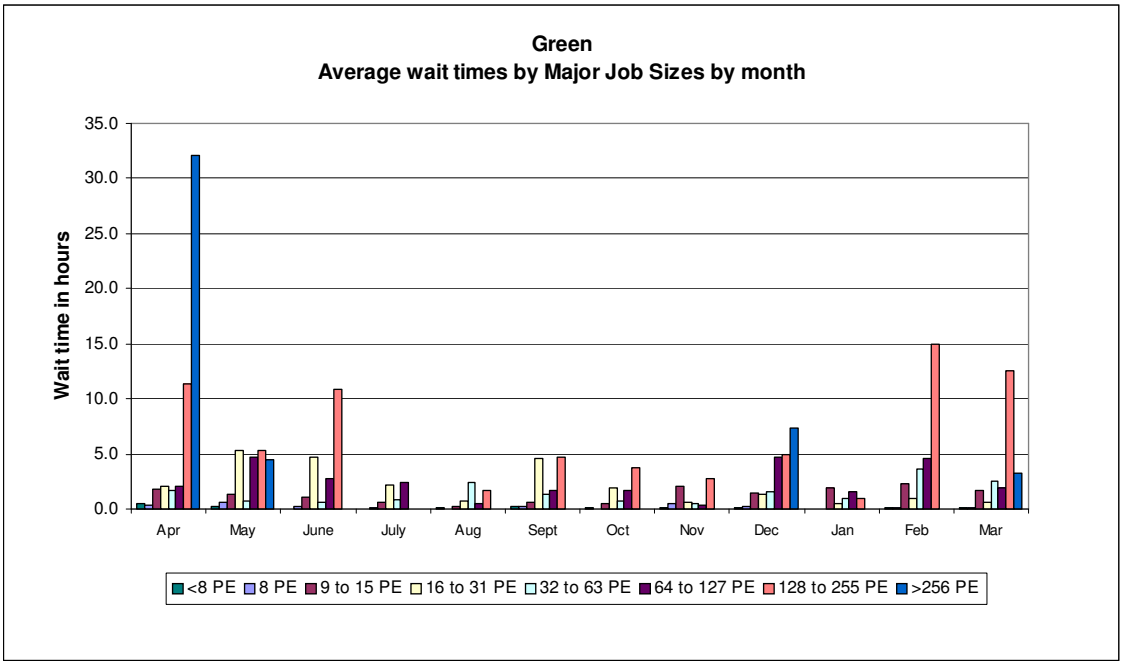

The next series of four charts show the percentage CPU time utilisation by the major job categories on the Green, Fermat and Newton systems for the 1st quarter 2005.

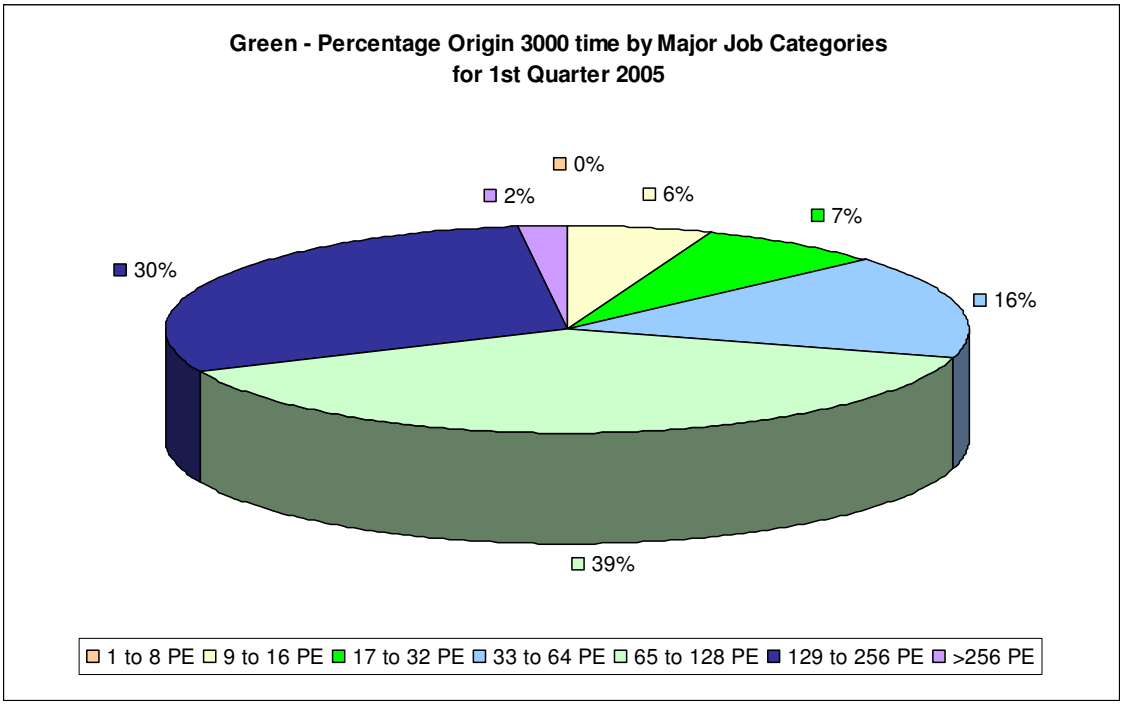

During this quarter there has been a good spread of work across the PE ranges on Green.

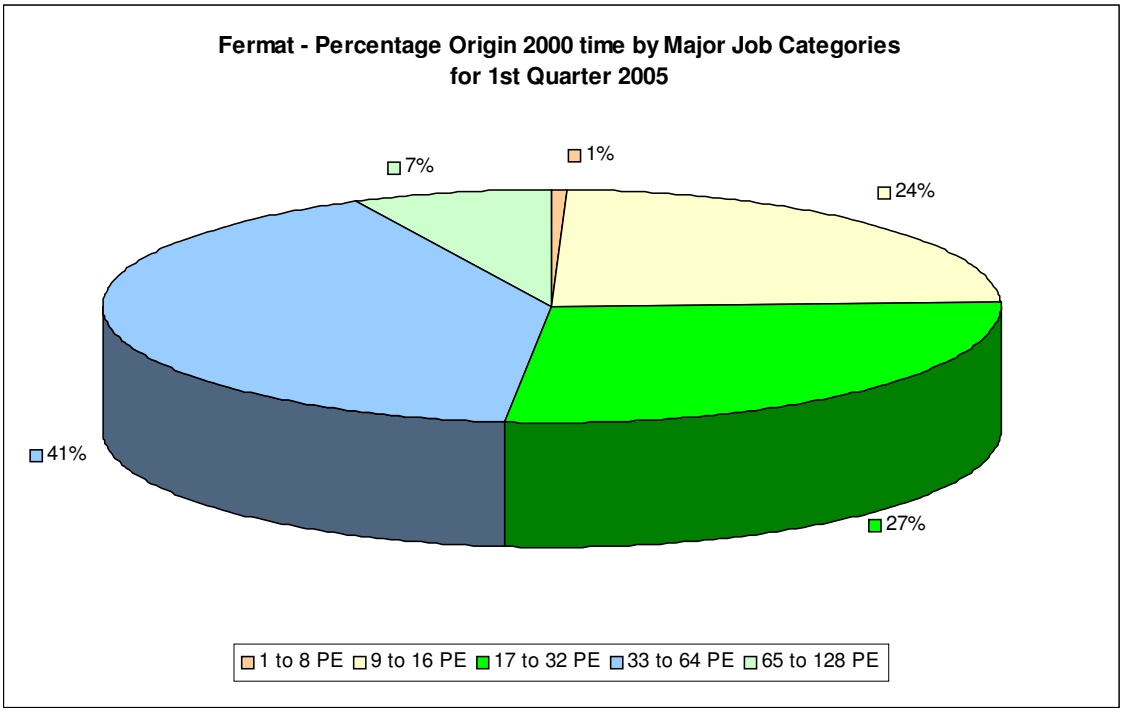

The workload on Fermat during this quarter was concentrated primarily in the mid-range PEs.

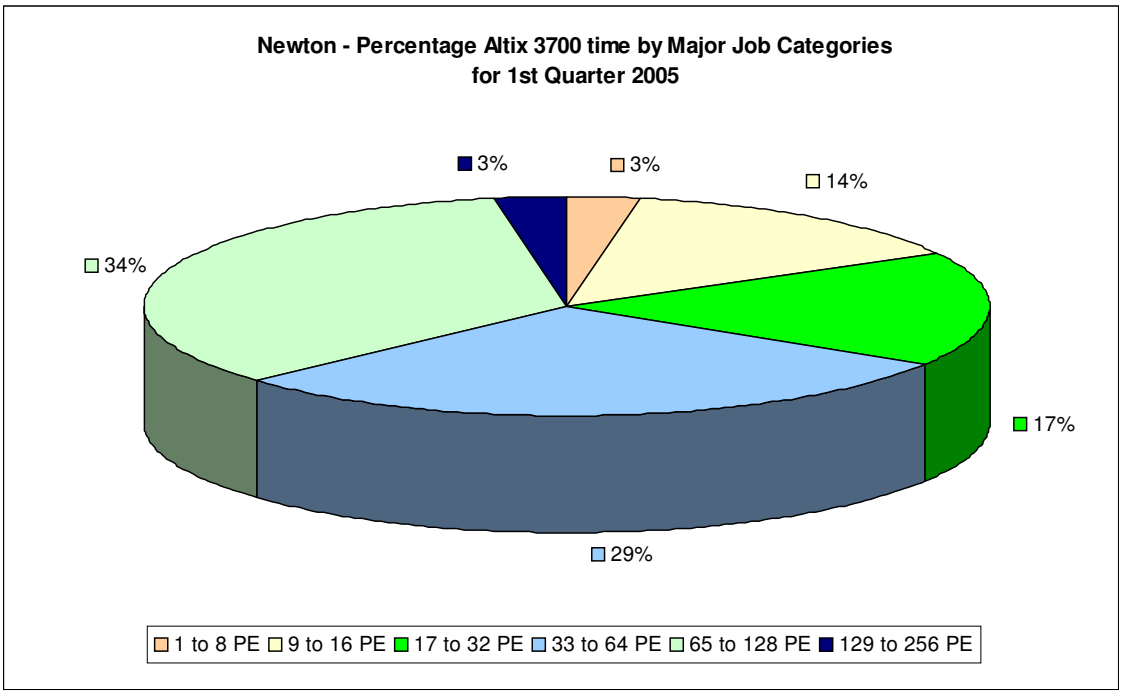

There was a good spread of work across Newton during the first quarter.

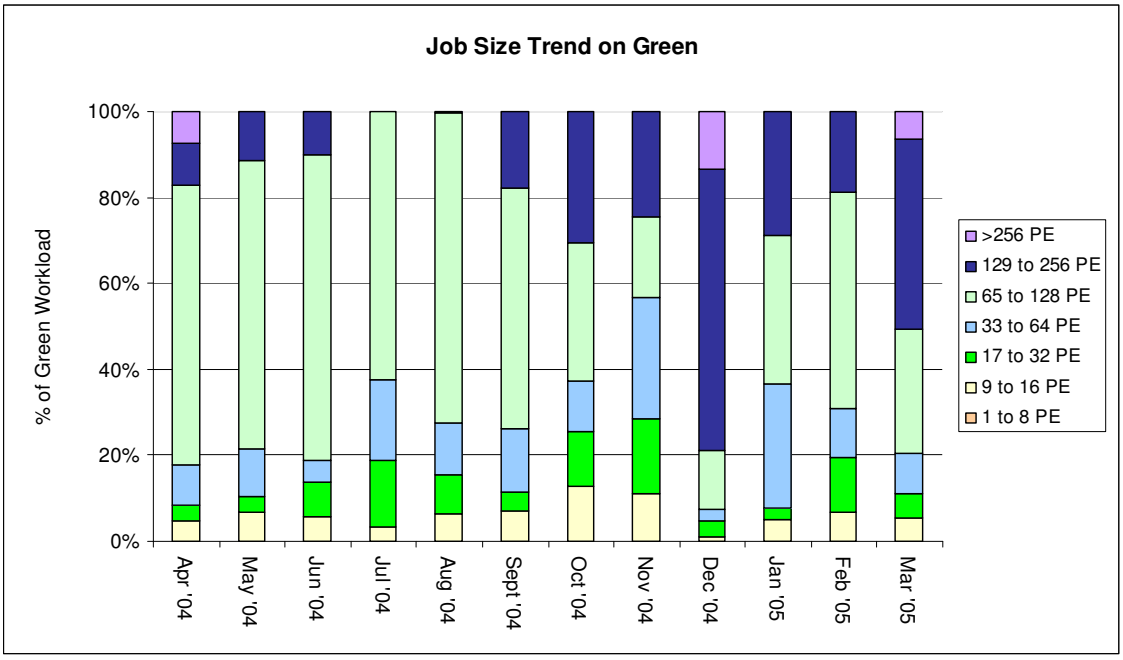

Usage on Green tended more to be in the mid- to high-end range of PEs during the 1st quarter.

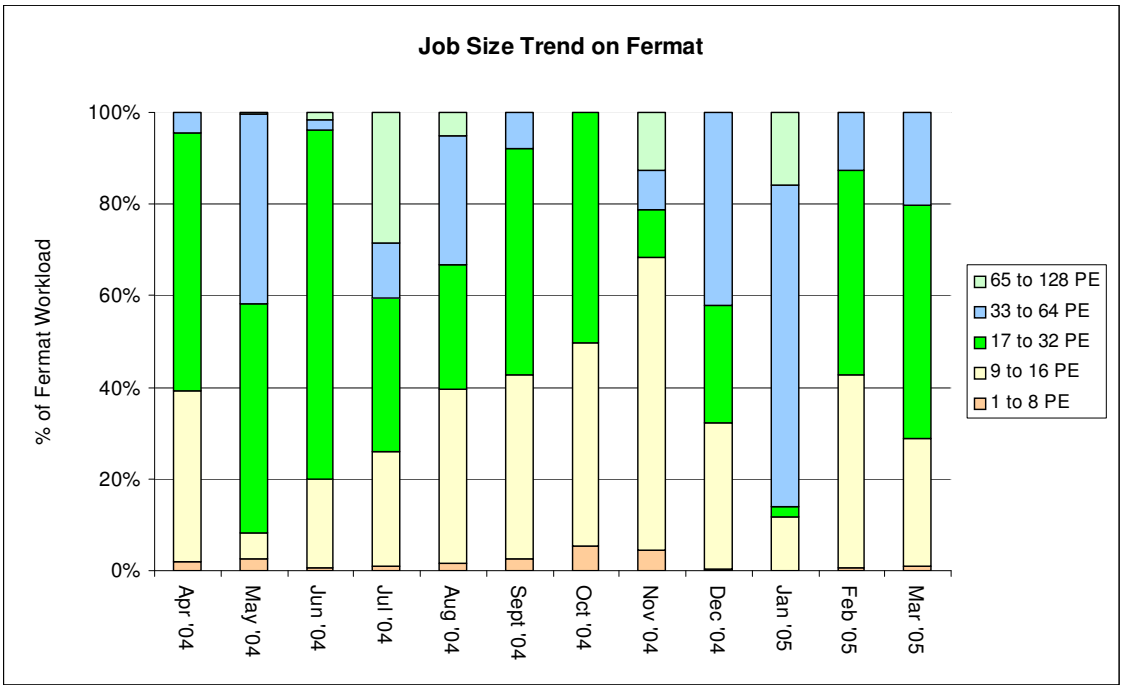

The workload on Fermat for this quarter was spread fairly evenly across the mid-range PEs.

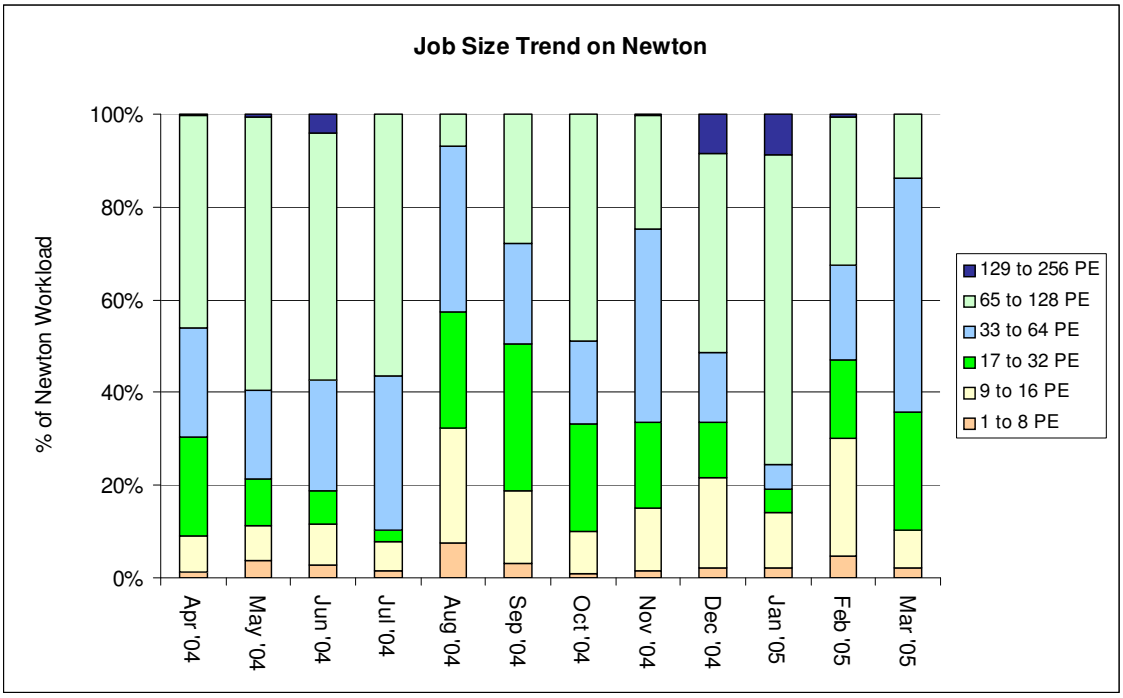

During the first quarter usage on Newton was concentrated mainly in the mid-range PEs.

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.

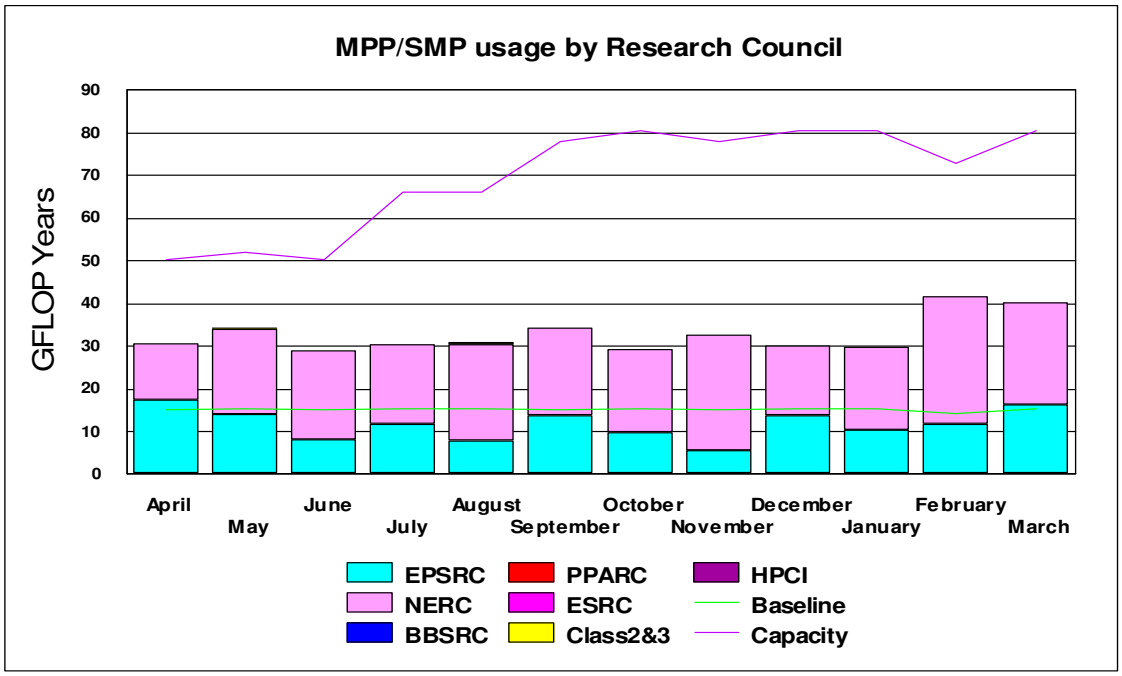

The graph below shows the same service utilisation by major system.

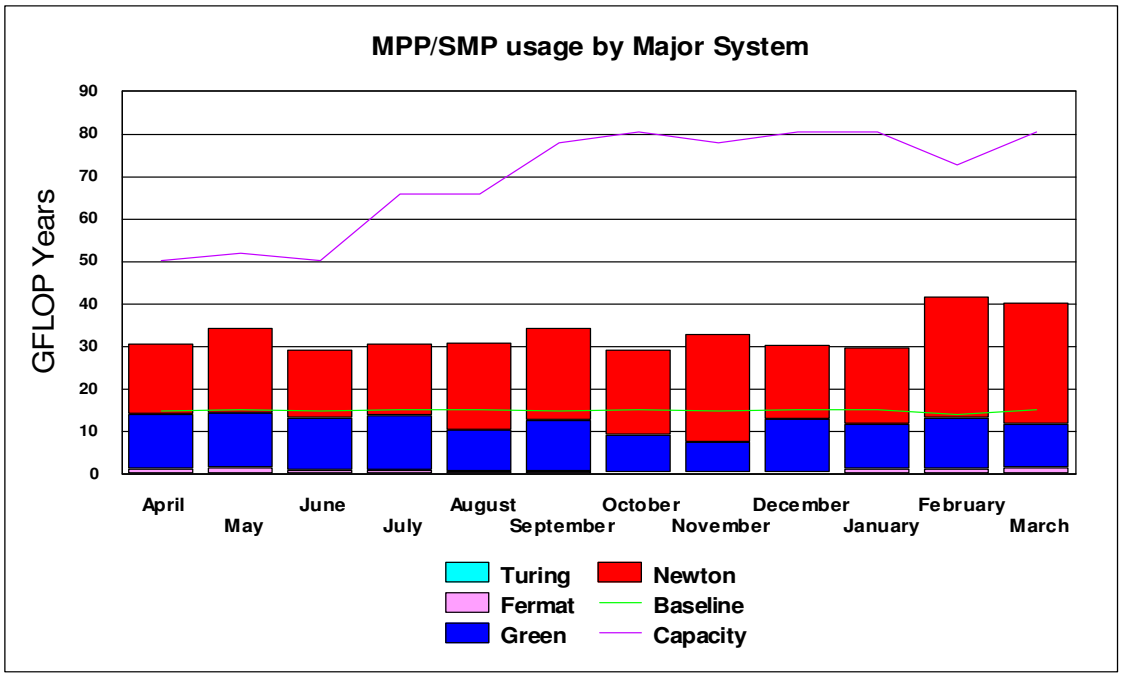

The next chart shows the historic interactive usage of the Origin 300 system Wren. Eight of the higher speed 500Mhz CPUs in Wren deliver the baseline capacity equivalent to that which was previously available on the Origin 3000 system Fermat for interactive usage.

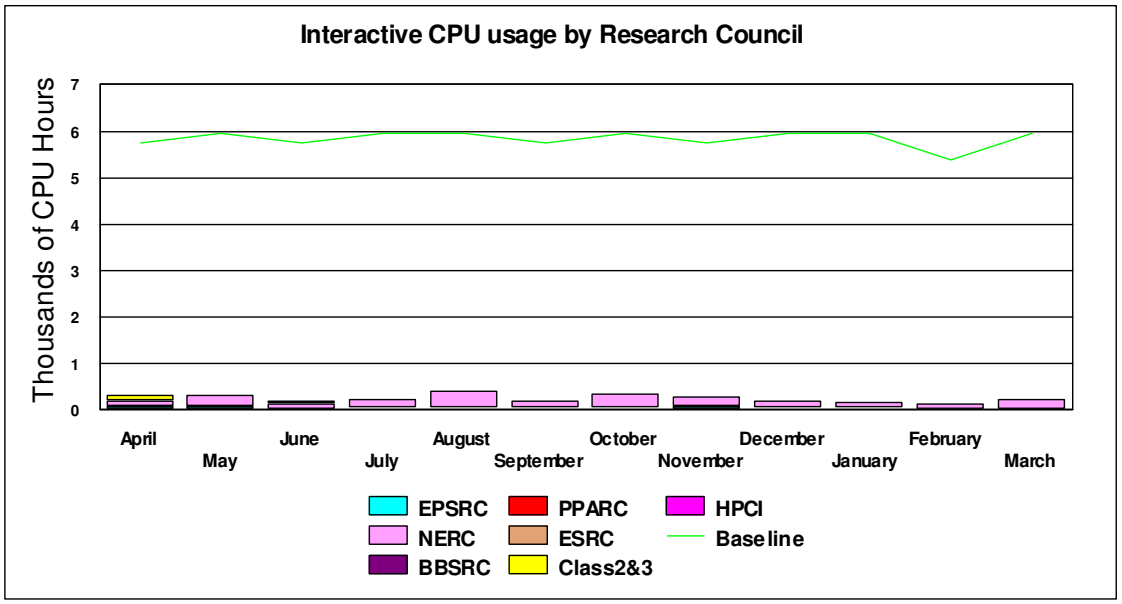

The following two charts detail the historic usage of the Origin 3000 system (Green) and the Altix 3700 system (Newton).

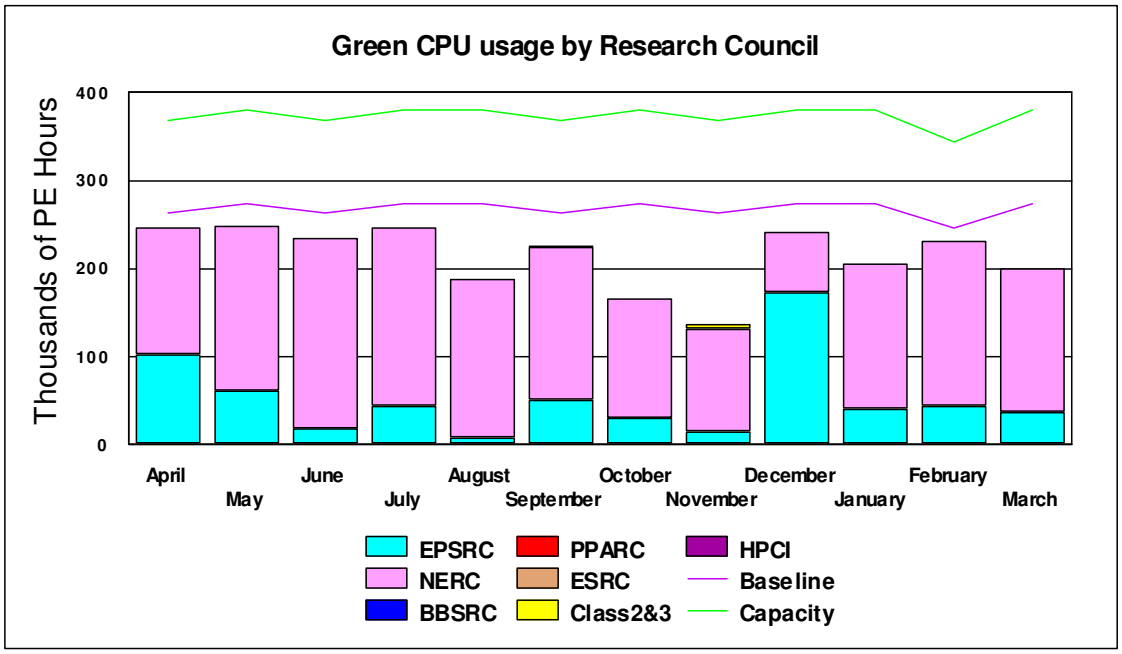

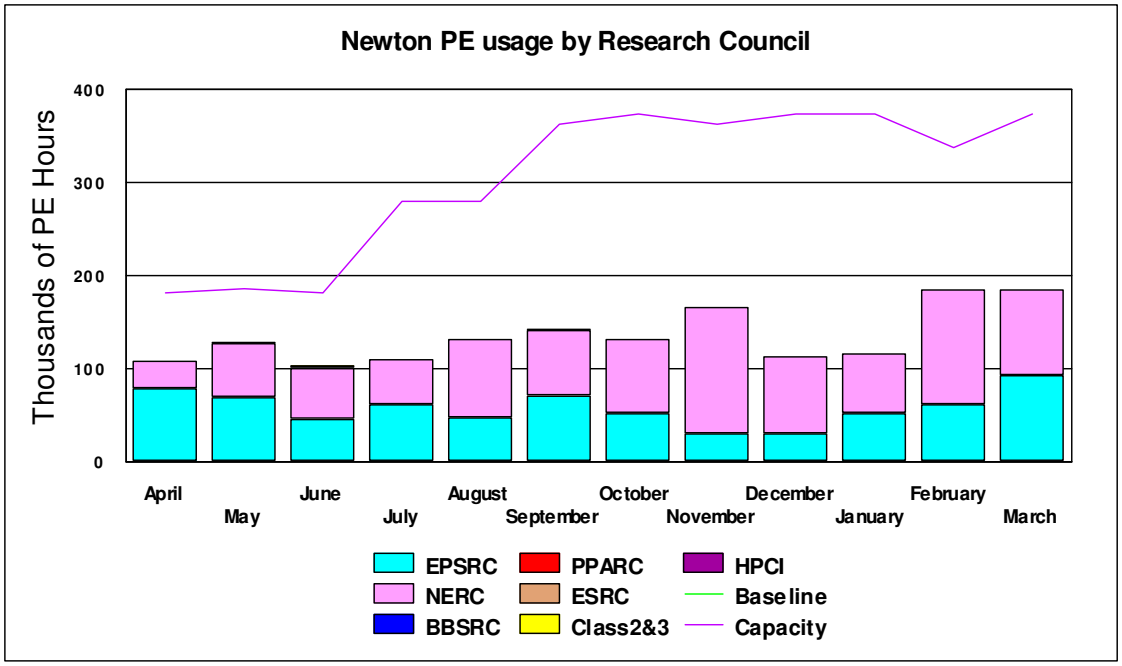

The next series of graphs illustrates the usage of the Medium Performance disk and HSM/tmp resources of the system.

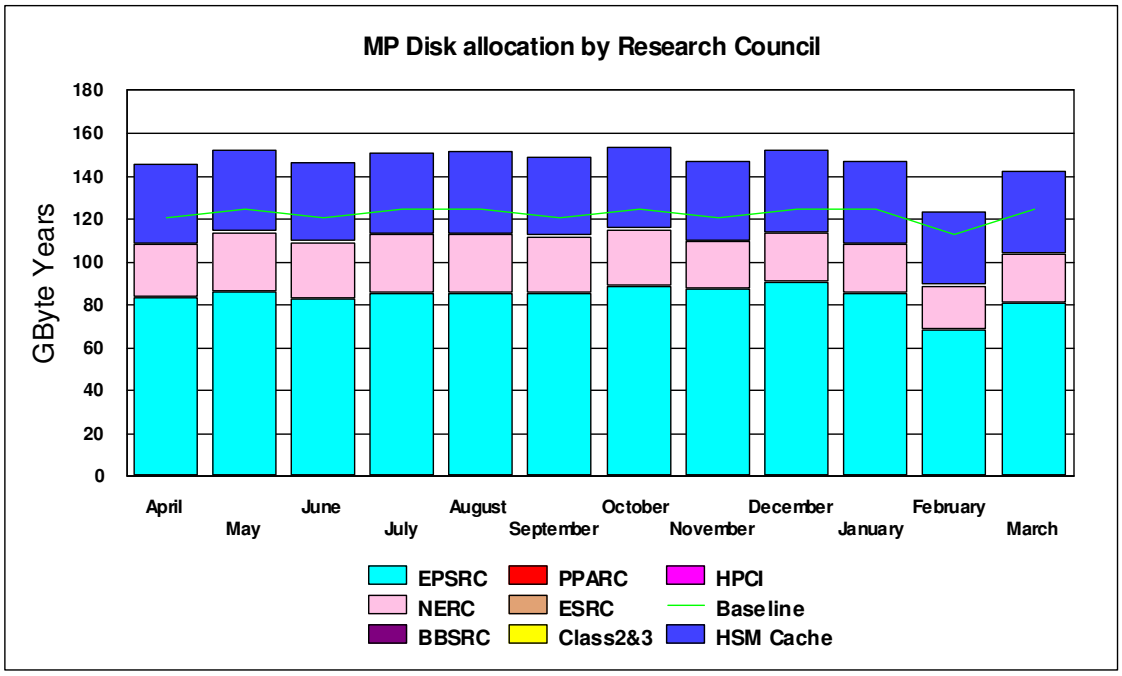

The above graph illustrates the historic allocation of the Medium Performance Disk on Fermat and the SAN.

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.

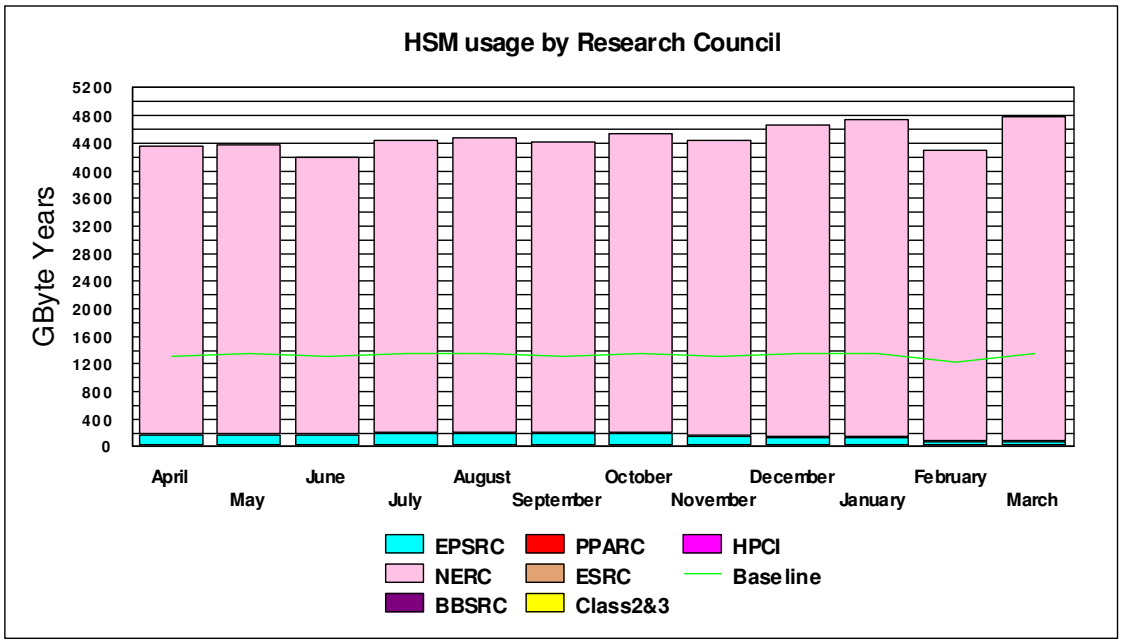

2.2.2 Guest System Usage

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 first quarter of this year.

Service Tokens Refunded: Quarter 1 2005 Usage							
System	Consortia						Total
	cse075	cse076	csn003	csn006	csn015		
Green 256+ PEs			16.14		171.83		187.97
Green 384+ PEs							0
Newton 128+ PEs	27.02	41.1		48.48			116.6
Newton 192+ PEs							0
Total Tokens							**304.57**

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 2005, with usage exceeding baseline.

There was a relatively balanced spread of work across all major systems throughout Quarter 1.

During the quarter, both Newton nodes were moved to a new building, paving way for the two nodes to be combined into one 512 processor physical node. Disruption to the user community was minimised, and the work was completed within the anticipated time scale..

Issues

There were some periods of system downtime during Quarter 1, primarily affecting the Altix system Newton, predominantly due to bugs within the CXFS (SAN) software code. These issues have been urgently addressed by the vendor, and are due to be fixed with the next release of the CXFS software, due to be released and implemented during the first part of the second quarter.

Plans

It is planned to upgrade both Irix and Altix operating systems during the latter half of April, which will at the same time upgrade the version of the CXFS SAN software on both sets of systems. This is a critical upgrade, intended to address the CXFS issues encountered over recent months.

It is also planned to combine the two 256 PE Newton nodes into one 512 processor physical node, expected to be carried during the second quarter this year. This physical combination will mean that the system can additionally at a future date be converted into a 512 PE Single System Image.

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 New Staff

Tim Robinson joined Manchester on 1st January 2005. His background and PhD in Computational Chemistry will complement the team's existing skill base.

3.2 Project Applications

10 applications for new CSAR projects were received, requesting a total of 470,815 service tokens.

3.3 New Projects

4 new CSAR projects were started with 95,579 service tokens being awarded in total.

3.4 Finished Projects

11 projects finished

3.5 Queries

A total of 152 CSAR queries were dealt with:

- 89 non-in-depth
- 55 user registration and admin
- 6 in-depth:

3.6 Service Quality Tokens

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

3.7 CSAR Focus

The 13th edition of CSAR Focus is currently being prepared.

4 Scientific Application Support Services

4.1 Training and Education

- Introduction to HPC and the CSAR Service
- Fortran 90
- Introduction to MPI
- Advanced MPI

4.2 Consortia Support/Software

The primary work has been in porting and optimising codes for the SGI Altix service, Newton, with additional testing on the 1.5GHz processors.

4.2.1 Boost Libraries

The Boost libraries, which provide a number of useful tasks for C++ programmers, have been ported to the CSAR machines.

4.2.2 Data Visualisation

CSAR staff have worked on implementing a library of routines to conveniently write out data to be input into AVS. This software works in parallel, taking advantage of MPI-IO and Parallel HDF.

4.2.3 Lens

Bugs were removed in the recently parallelised version of Lens.

4.2.4 Molpro

The computational chemistry package Molpro, a high level ab initio code for quantum chemistry, has been installed and is currently being evaluated.

4.2.5 Unified Model

Version 6.0 of the Unified Model has been installed and investigated on Newton under the 8.1 compilers. Fixes have been applied to enable it to compile, and these have been passed to the Met. Office.

4.2.6 Development Software

New licenses for the development software TotalView and Intel Trace Tools have been purchased, allowing new versions of these software packages to be installed.

4.2.7 Visualisation Support

A report was compiled for Professor Kai Luo on visualising 3D vectors.

4.2.8 Castep

Castep 3.1.1 was installed on Newton.

4.2.9 New Users to HPC

Discussions have continued with the new applications to HPC consortia, with visits scheduled in the forthcoming quarter. Lee Margetts met Charles Augarde's RA at the ACME conference in Sheffield, where their project was discussed at length.

5 Collaboration and Conferences

5.1 MRCCS Projects

5.1.1 Advanced Virtual Prototyping Research Centre

A conference paper co-authored by Lee Margetts and Joanna Leng was accepted for COMPLAS (Computational Plasticity) in Barcelona, Spain, September 2005.

5.1.2 MRCCS Seminars

The Access Grid based seminar series continued on Fridays with a wide variety of speakers. A number of the invited speakers were interested in submitting research applications for CSAR resources. The event gives such users a great opportunity to network with the CSAR team.

5.1.3 Netsolve

NetSolve is an RPC based client/agent/server system that allows one to remotely access both hardware and software components. There are now developments to use Globus and LSF, which will help in the use of CSAR systems.

5.2 Events

Events attended include:

- Lee Margetts represented CSAR at the ACME (Association for Computational Mechanics in Engineering) annual conference in March, where he gave an invited talk on the National High Performance Computing Services. The visit stimulated such interest that the ACME community are keen to form a new HPC consortium. A proposal is due to be submitted to EPSRC in the coming months.
- Andrew Jones attended the SOS9 Workshop, March 21-23 2005, Davos, Switzerland.

Events organised include:

- A two day Technical Symposium on Reconfigurable Computing with FPGAs was hosted on 21-22 February. The event was a huge success with demand outstripping availability. Cray, SGI and Ohio Supercomputer Centre all sponsored the event, which was organised by Kevin Roy and Carl Ward.
- 10-11 March, Lee Margetts organised a two day International Workshop on Virtual Prototyping. Participants attended from the UK, France, Germany, Spain and Brazil. The workshop was a success with vital contacts being made. Delegates were keen for an annual event.

5.3 Publications

The following paper contributed to by Tim Robinson was published: Rong Z.M., Henry B.R., Robinson, T.W., et al. Absolute intensities of CH stretching overtones in alkenes J. Phys. Chem. A 109 (6): 1033-1041, 2005

6 Added Value Services

6.1 ISC 2005

Preparation work has begun for this year's attendance at the International Supercomputer Conference (ISC). The conference takes place in June (21-24), and this year the event will celebrate its 20th anniversary.

6.2 Visualisation

On March 3 the Visualisation Team had a successful mini-workshop launching the new E-Science North West Passive Stereo Facility. This showed the potential and some uses of a 7-metre stand-alone flat wall stereoscopic visualisation centre, ideal for comparative scientific analysis.

One of the first major funded projects to exploit this new facility is the JISC Virtual Research Environment; SAGE (Stereoscopic Access Grid Environment), which aims to integrate stereoscopic visualisation within the next generation of video conferencing. A key additional part of this project is the construction and recommendation for building small portable passive stereo visualisation facilities, using UK suppliers. This allows the ability to have a cut-down virtual environment centre for the complete price of about £8000, enabling visualisation to occur directly in the classroom or in remote exhibition spaces.

For the high-end visualisation users the VIPL (Visualisation Immersive Projection Laboratory) continues to find new users recently hosting a combined truly immersive experience for sedimentologists using Schlumberger's Inside Reality software compared with Norsk Hydro's HydroVR software. These events used all the features allowing for a fully tracked user interacting and editing large Earth science data-sets employing nearly half of the processing and graphics power of Bezier HPC/HPV. Repeated events are planned as well as cross-cooperative virtual links with the eScience centres at Cardiff and Durham.