CSAR Service

Consolidated Management Report

2nd Quarter 2004

Management Summary

This is the consolidated Management report for the second quarter 2004 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 475 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 second quarter of this year.

All systems, Altix and Origins, had their operating systems upgraded during the latter part of the second quarter, to assist with ongoing improvements to the service.

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 planned to be 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 2nd quarter 2004 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	2003/4 July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March	April	May	June
HPC Services Availability												
Availability in Core Time (% of time)	98.83%	98.95%	96.62%	98.84%	98.95%	98.75%	97.49%	98.16%	98.51%	89.39%	94.21%	97.49%
Availability out of Core Time (% of time)	99.57%	100%	98.48%	99.28%	97.74%	98.3%	98.88%	97.9%	99.48%	91.90%	99.73%	97.85%
Number of Failures in month	2	2	4	4	3	5	5	4	3	5	3	4
Mean Time between failures in 52 week rolling period (hours)	487	461	417	365	337	283	265	243	286	246	224	202
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	<1	<0.5	<5	<2	<1	<1	<2	<2	<2	<2	<3	<2
Administrative Queries - Max Time to resolve 95% of all queries	<0.5	<1	<1	<1	<1	<1	<0.5	<0.5	<0.5	<0.5	<0.5	<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	2003/4 July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March	April	May	June
HPC Services Availability												
Availability in Core Time (% of time)	0.039	0.039	0.078	0.039	0.039	0.039	0.078	0.078	0.039	0.195	0.195	0.078
Availability out of Core Time (% of time)	-0.039	-0.047	0.078	-0.039	0.078	0.078	0	0.078	0	0.039	0	0.078
Number of Failures in month	0	0	0.008	0.008	0	0.0004	0.0004	0.008	0	0.0004	0.0	0.008
Mean Time between failures in 52 week rolling period (hours)	0	0	0	0	0	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	-0.019	0.0312	0	-0.016	-0.016	0	0	0	0	0.016	0
Administrative Queries - Max Time to resolve 95% of all queries	-0.019	-0.016	-0.01551	-0.015514	-0.016	-0.016	-0.019	-0.019	-0.019	0	-0.019	-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:	-0.04	-0.04	0.07	-0.03	0.02	0.02	0.01	0.05	-0.01	0.09	0.07	0.05
Quarterly Service Credits:			-0.01			0.02			0.05			0.21

Table 3

1.2 No. of Registered Users

The current position at the end of the quarter is that there are 475 registered users of the CSAR Service.

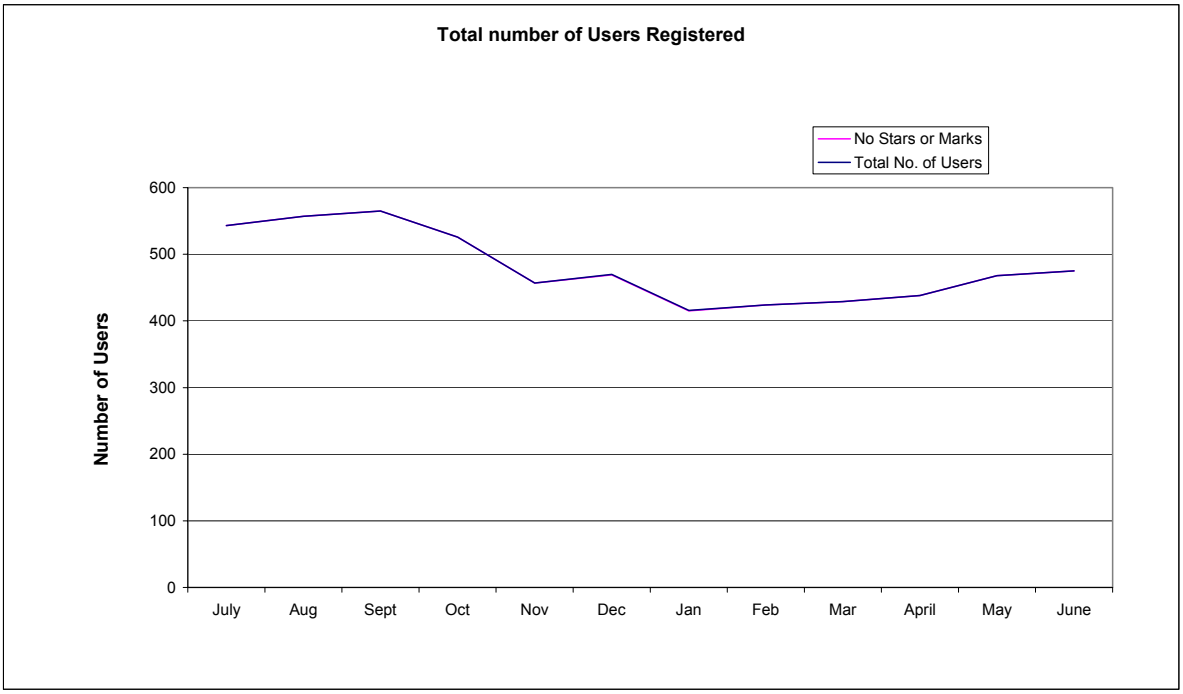

1.3 Service Quality Tokens

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

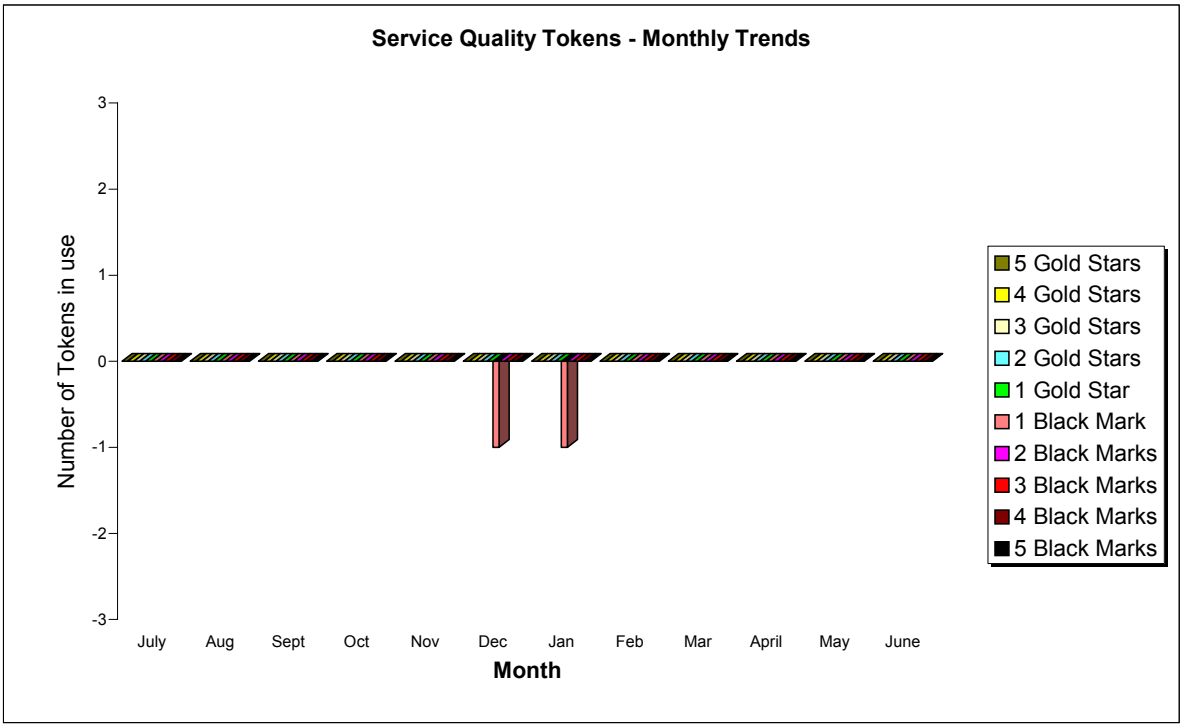

Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

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

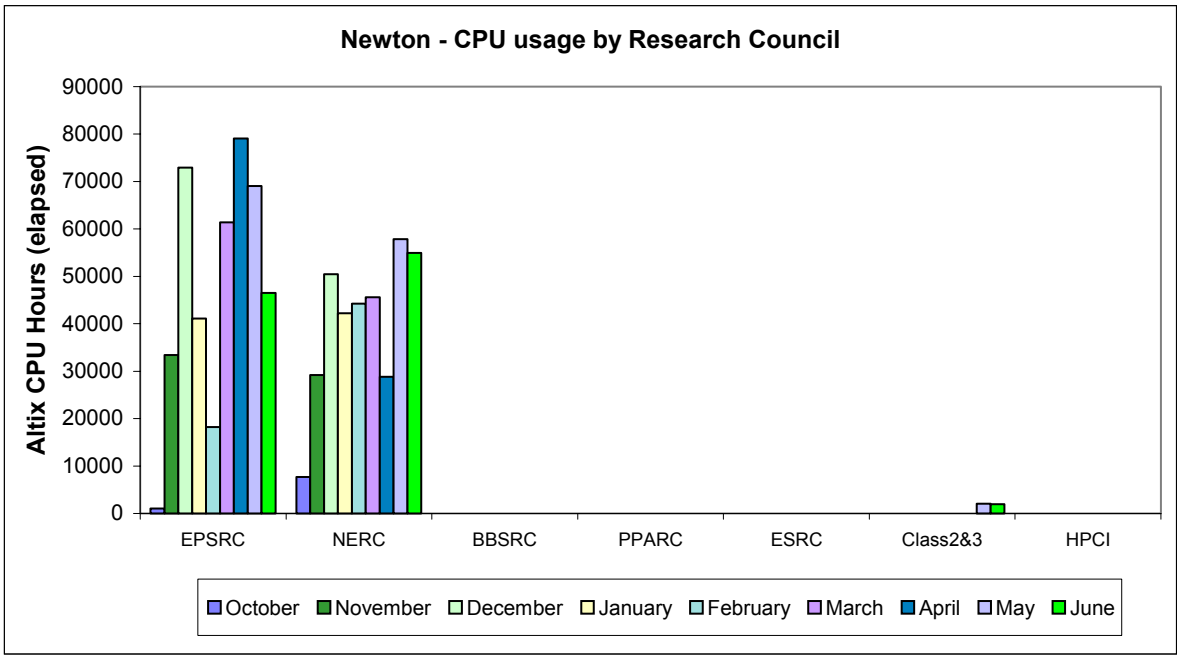

Usage of the SGI Altix 3700 system Newton is shown by Research Council in the above graph, from its introduction in October 2003.

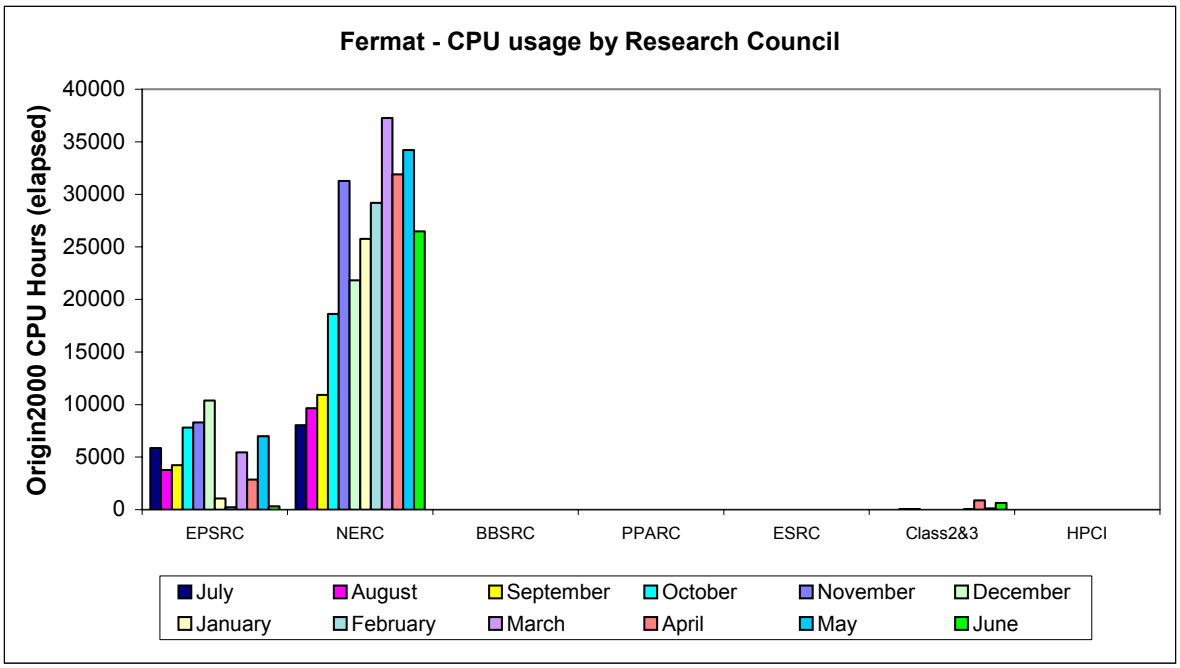

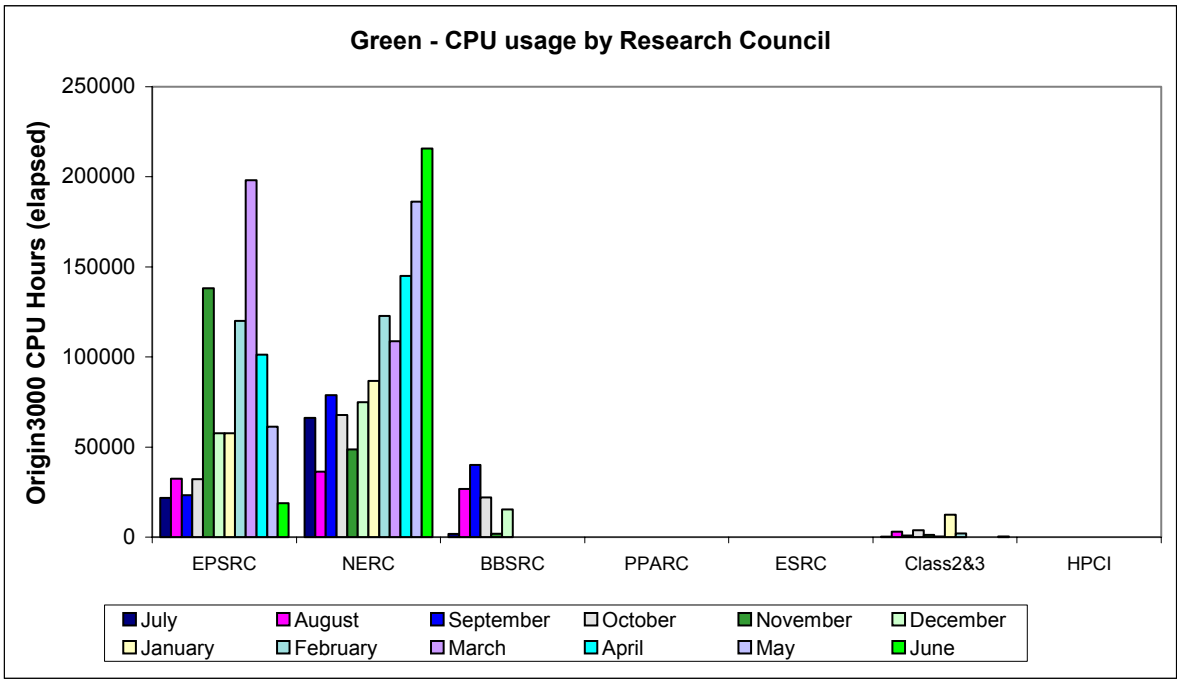

Usage of the two batch SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service in the preceding two charts.

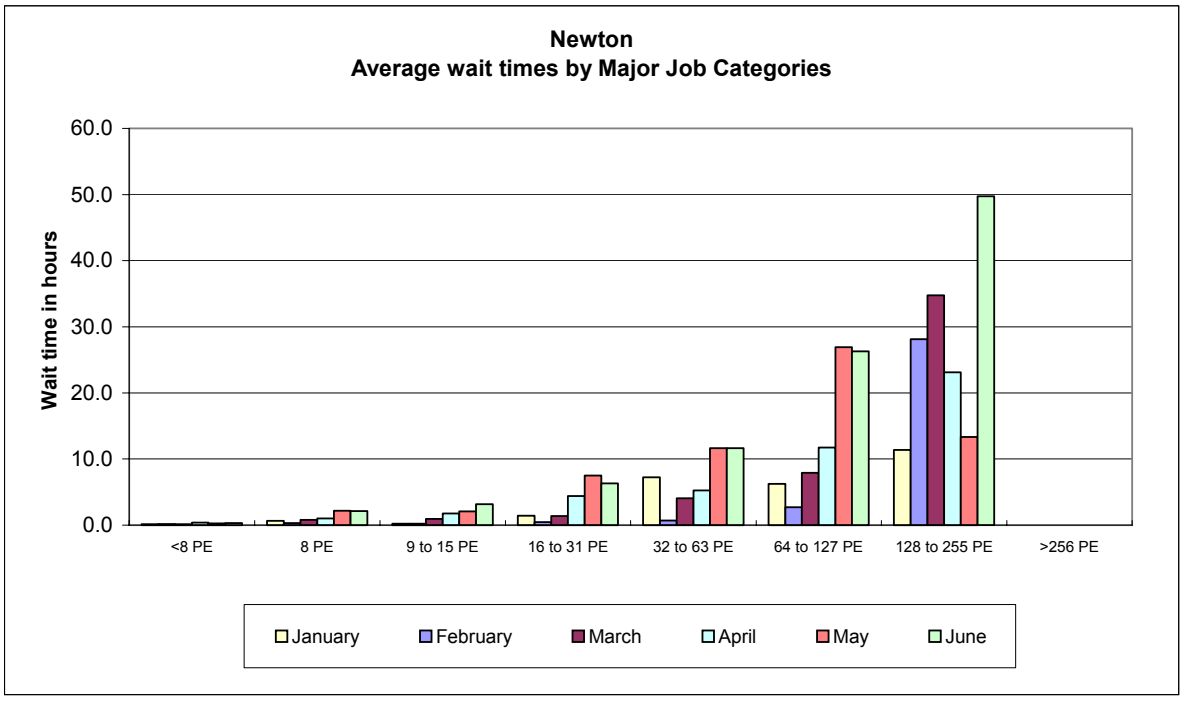

The above chart, and the one below, shows the wait time trend in hours on the Altix 3700 Newton.

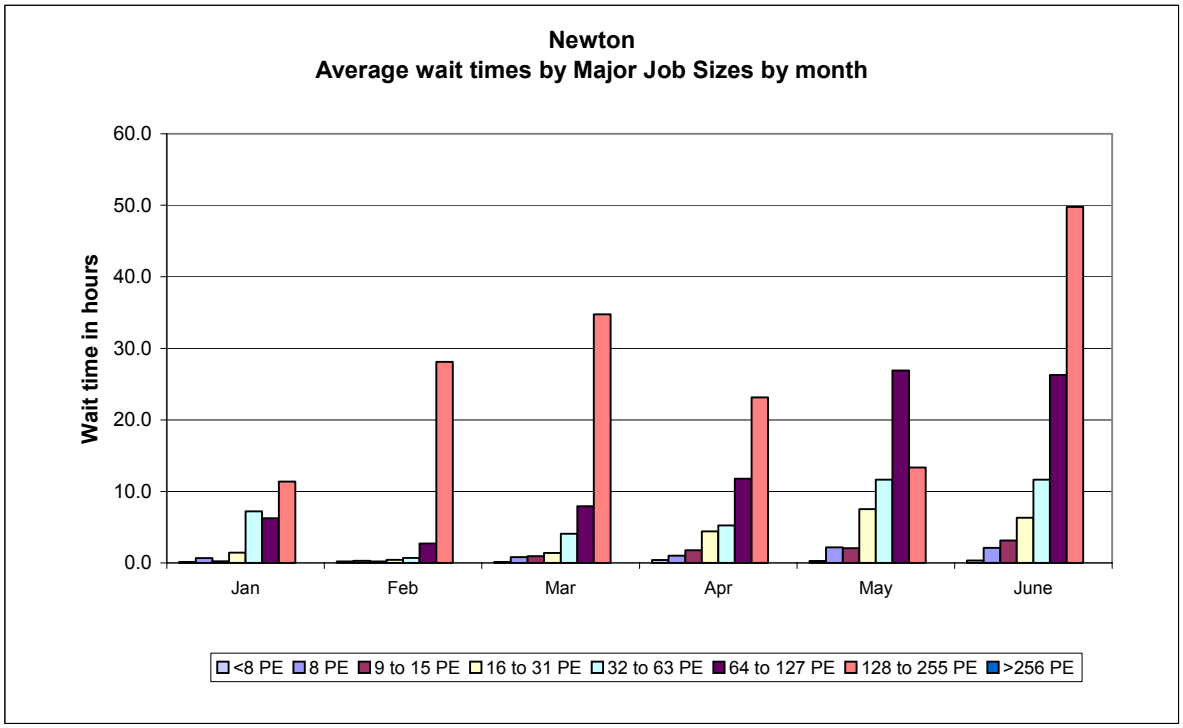

The following two charts show average wait times in hours for the quarter on the Origin 3000 Green.

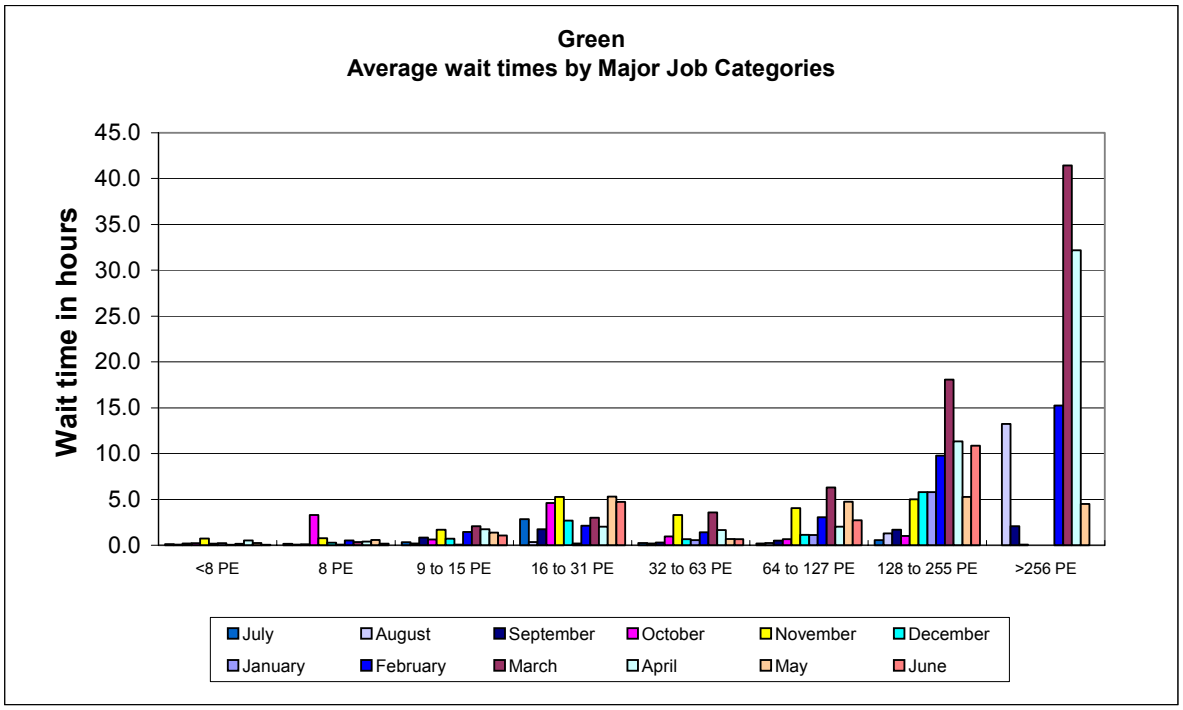

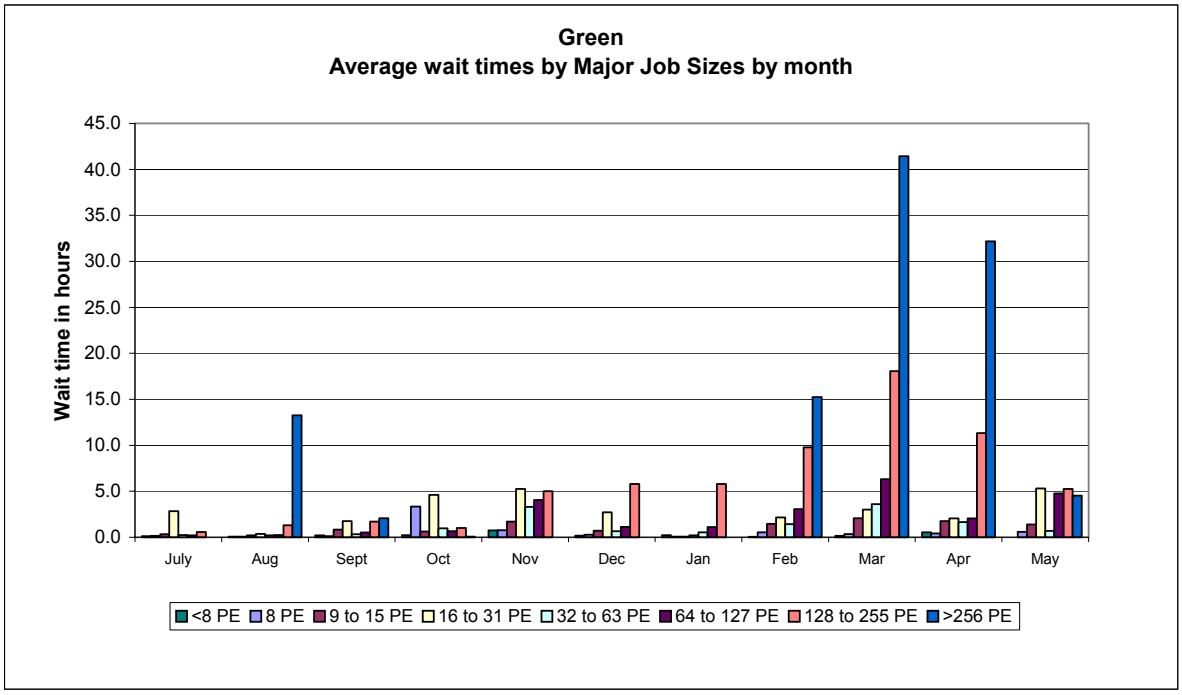

The next series of four charts show the percentage CPU time utilisation by the major job categories on the Green, Fermat and Newton systems for the 2nd quarter 2004.

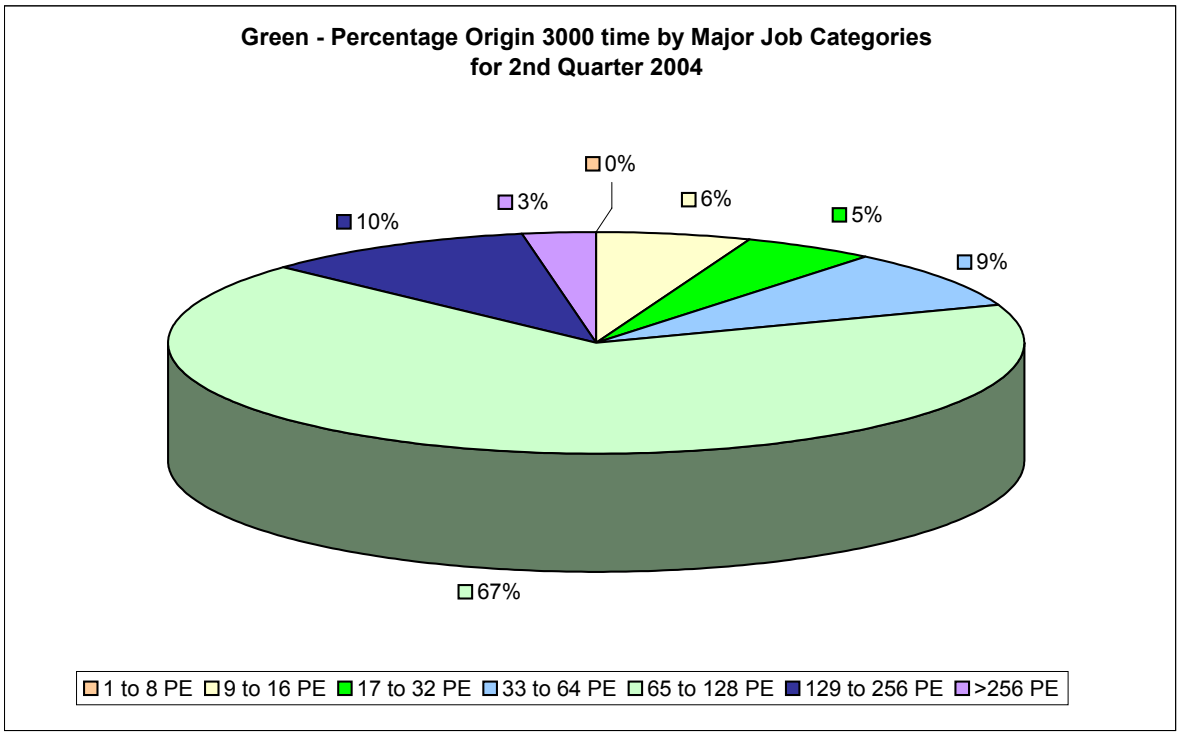

On Green, the 65 to 128 PE range has seen the greatest percentage of workload during this quarter, and significant amounts of Capability computing above 128 processors.

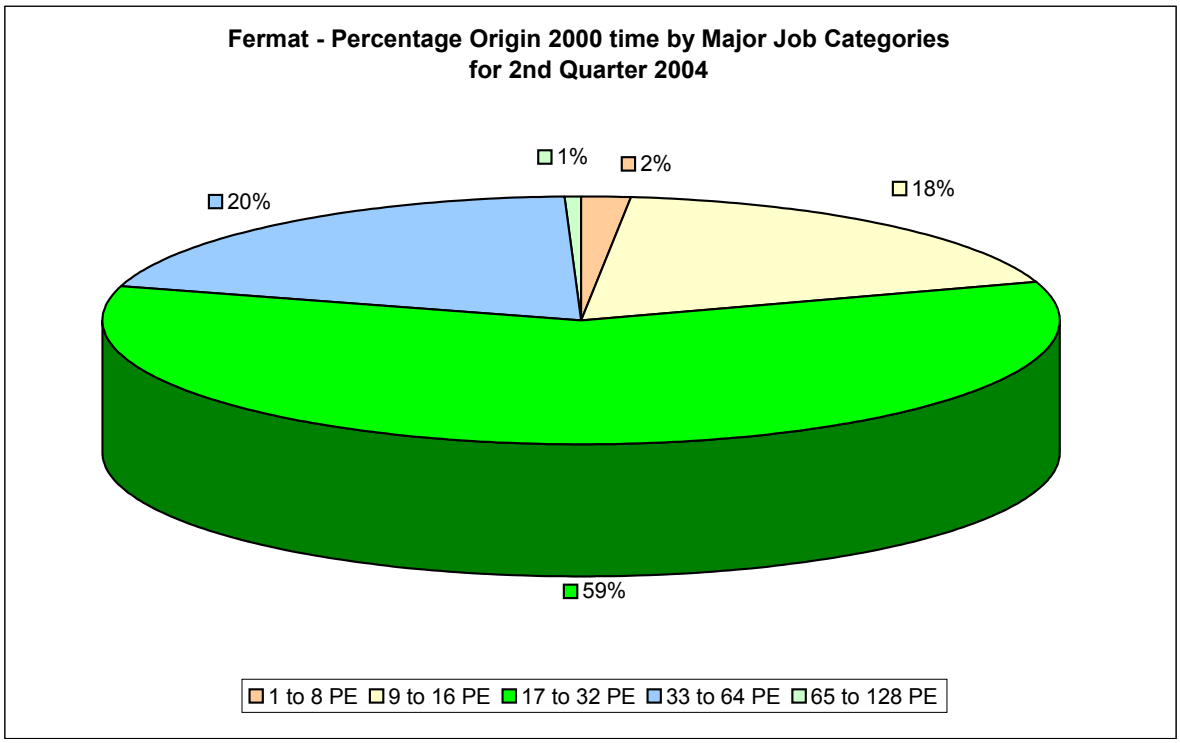

On Fermat the highest concentration of work was in the 17 to 32 PE range at 59%.

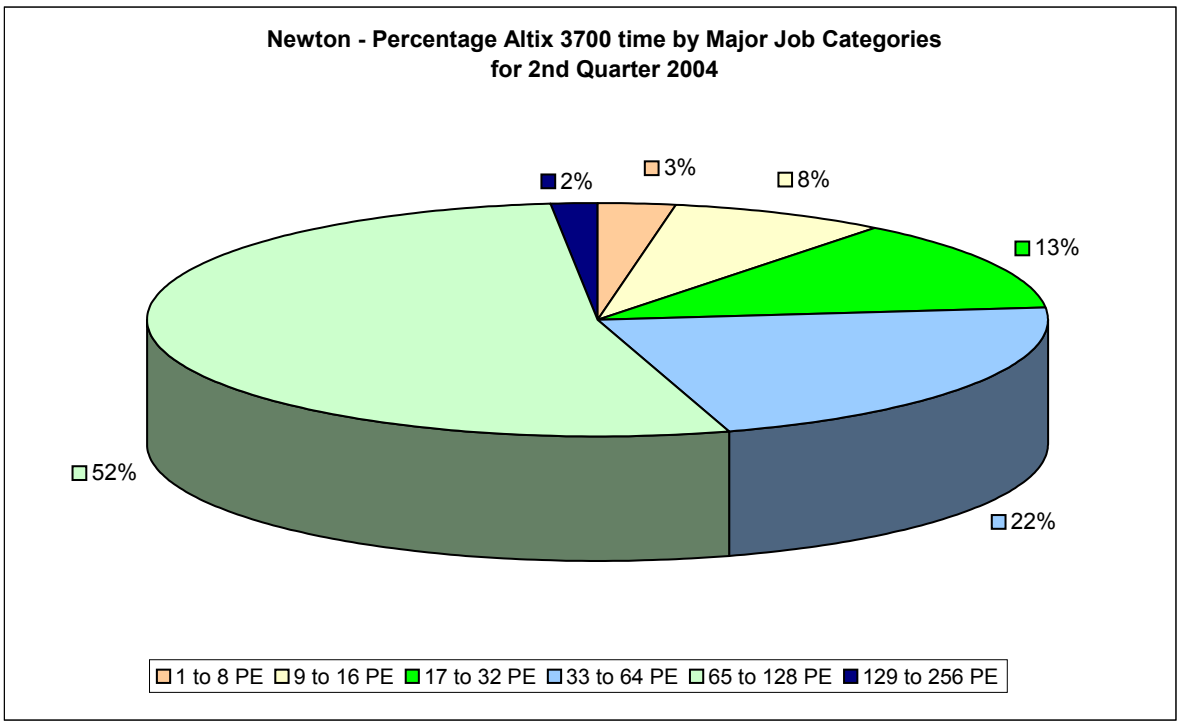

Newton's workload during the second quarter was focused largely in the 65 to 128 PE range, with significant volumes of application development work.

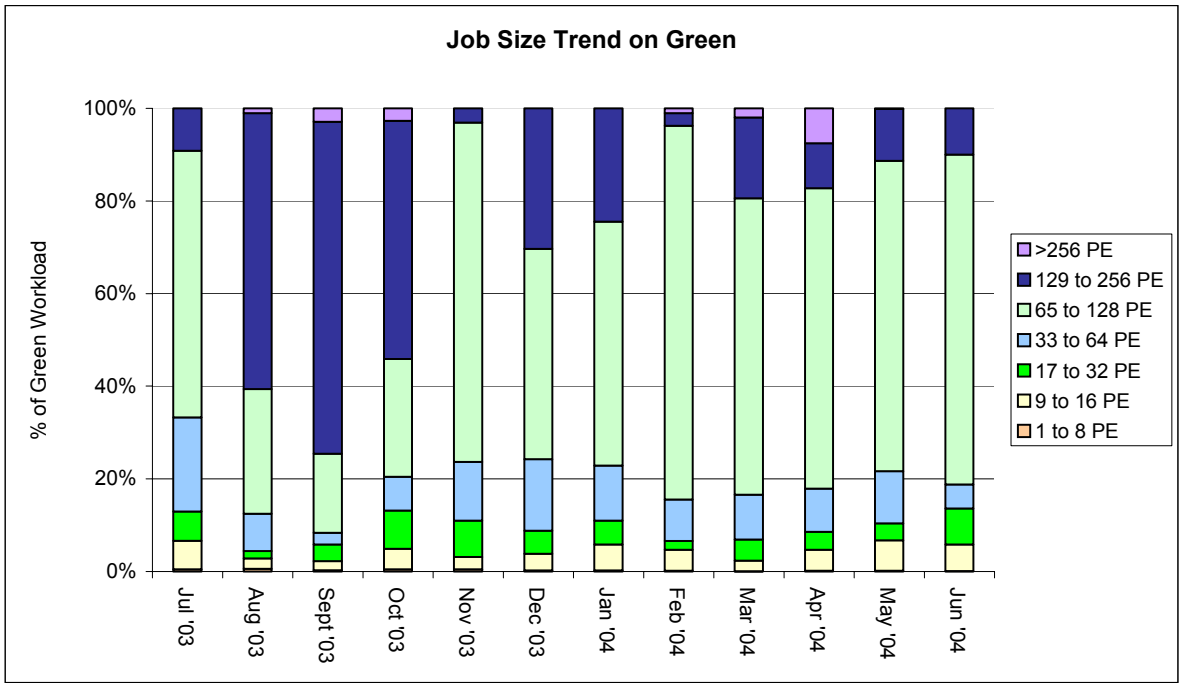

Usage on Green tended more to be in the mid-range of PEs during the 2nd quarter.

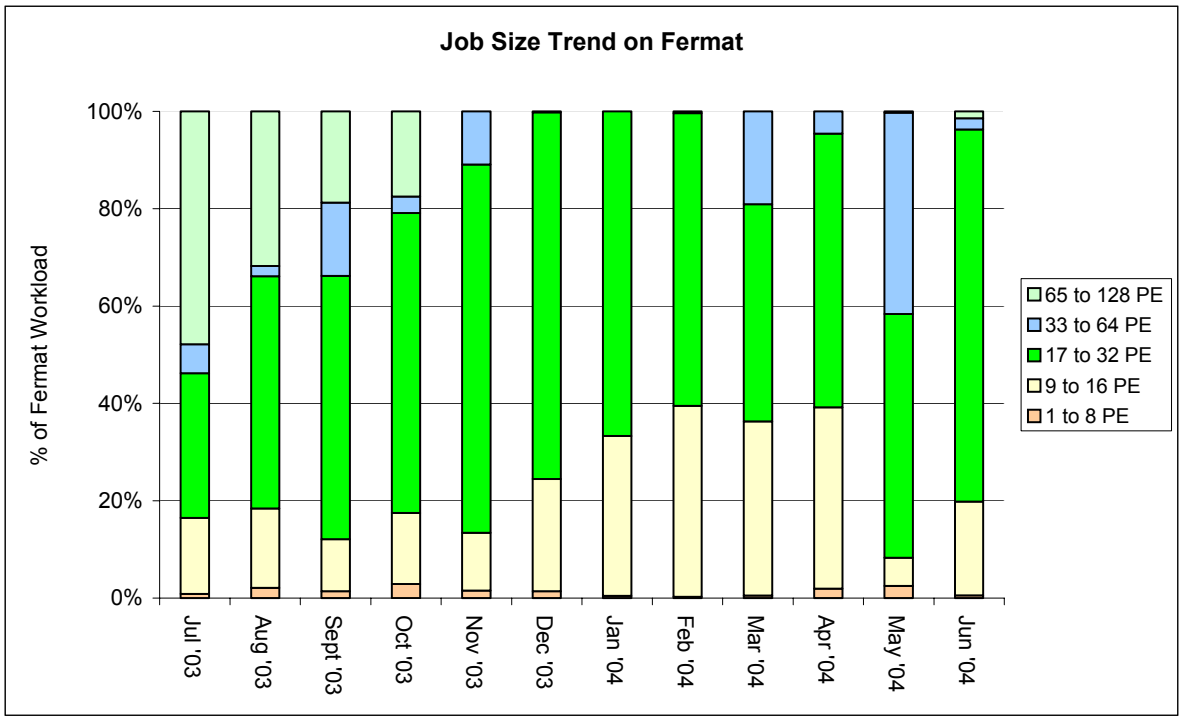

The workload on Fermat for this quarter was also spread fairly evenly across the mid-range of PEs.

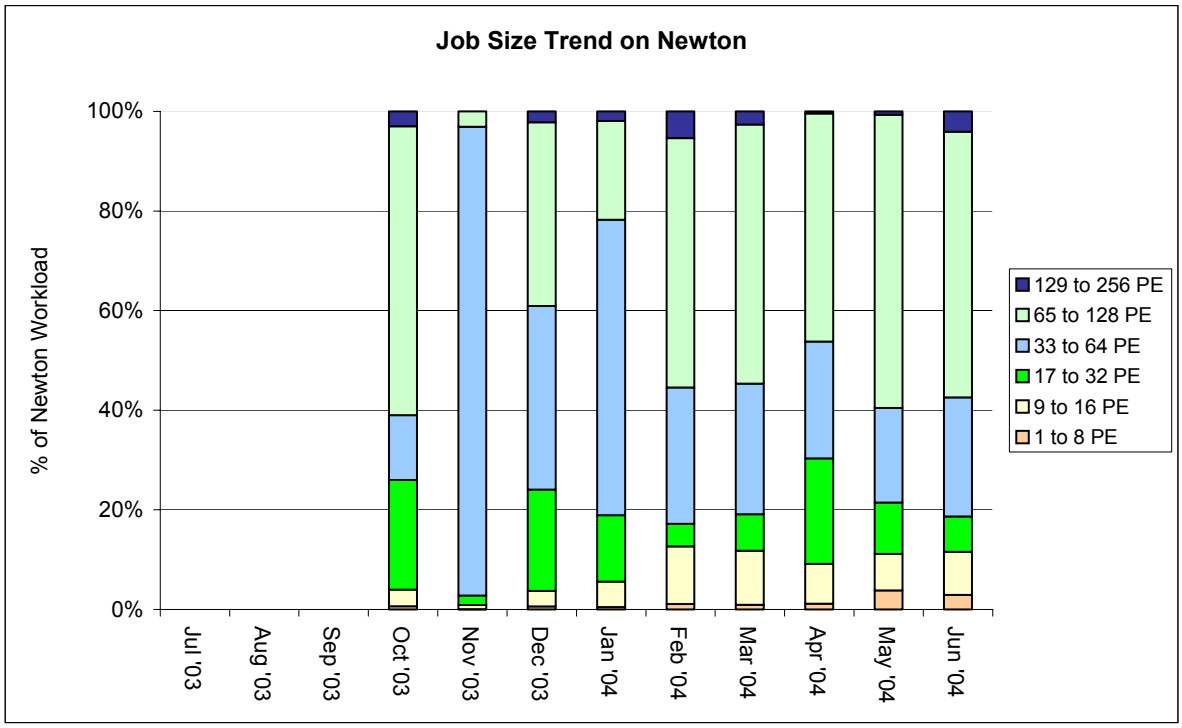

During the second quarter Newton's usage was reasonably spread across the machine, with a slight trend towards increasing job sizes.

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.

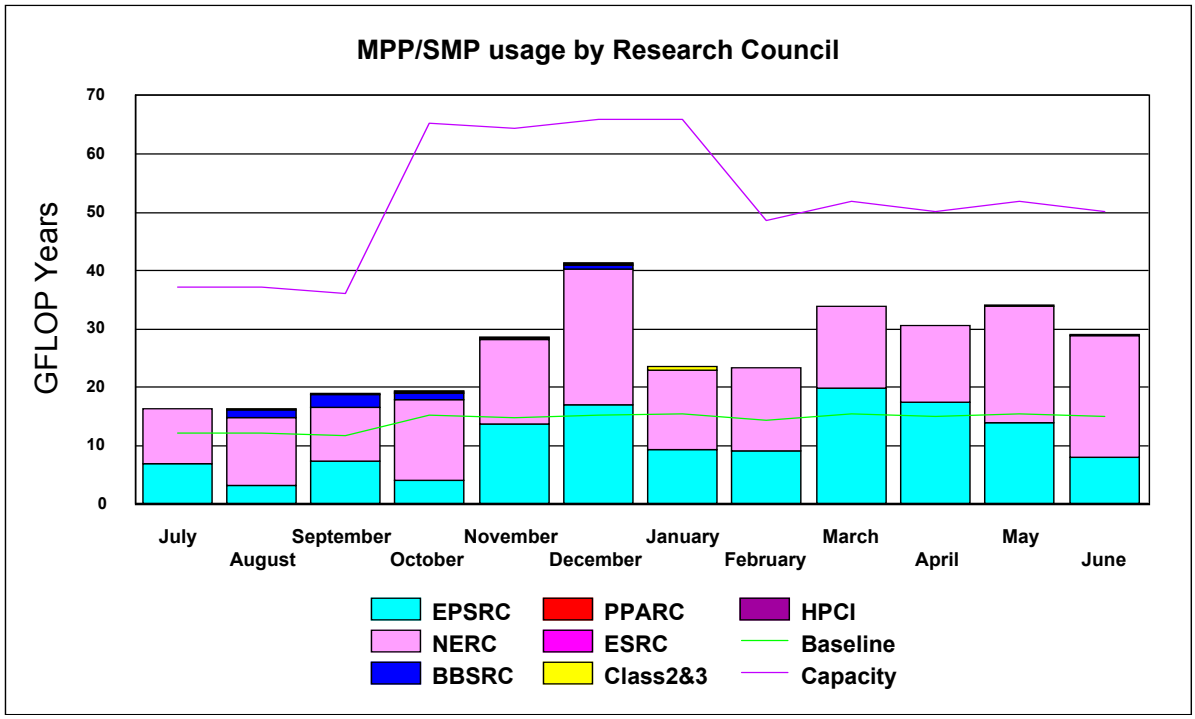

The graph below shows the same service utilisation by major system.

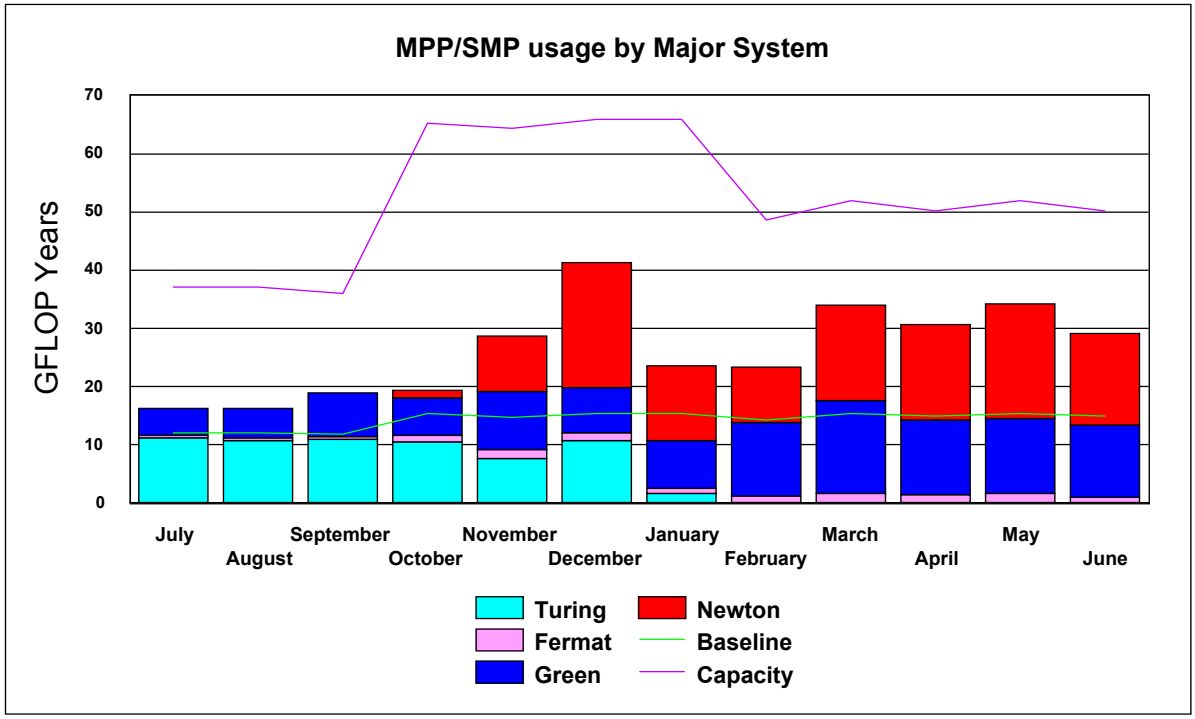

The next chart shows the historic interactive usage of the Origin 300 system Wren. Eight of the higher speed 500Mhz CPUs in Wren deliver the baseline capacity equivalent to that which was previously available on the Origin 3000 system Fermat for interactive usage.

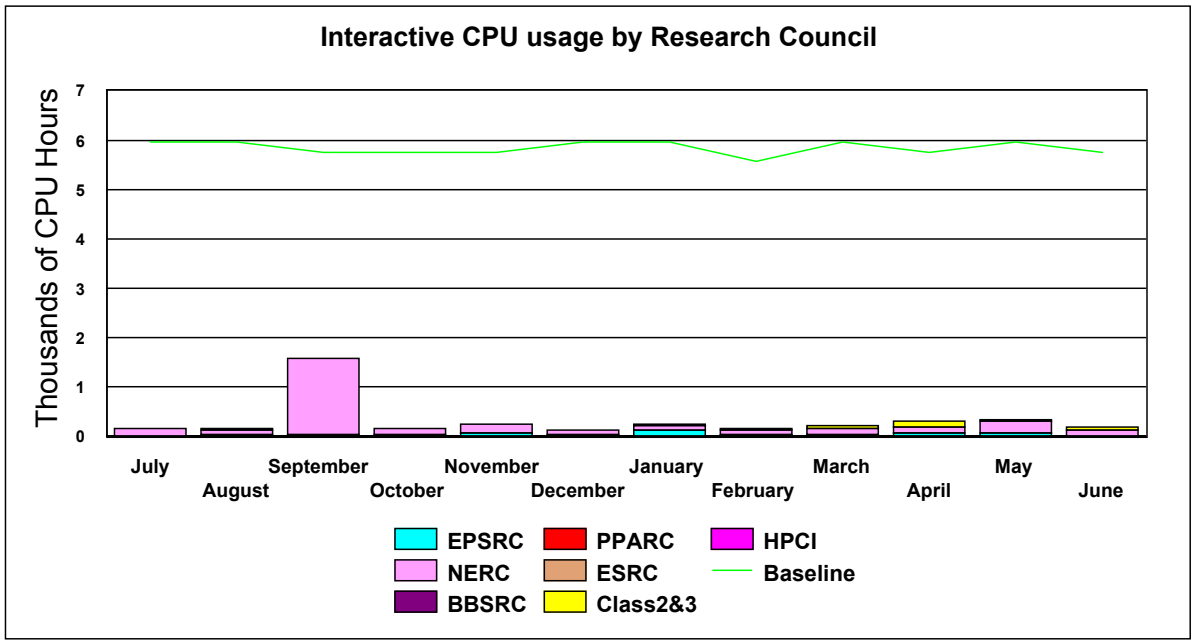

The following two charts detail the historic usage of the Origin 3000 system (Green) and the Altix 3700 system (Newton).

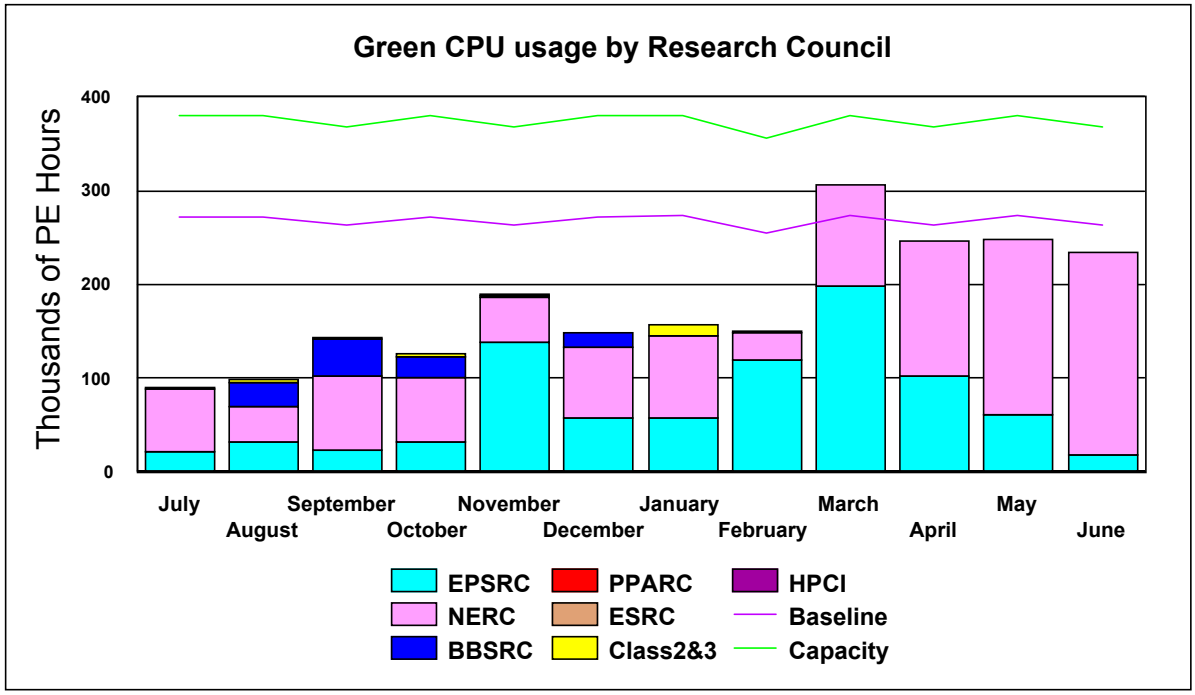

NERC is by far the largest user of Green.

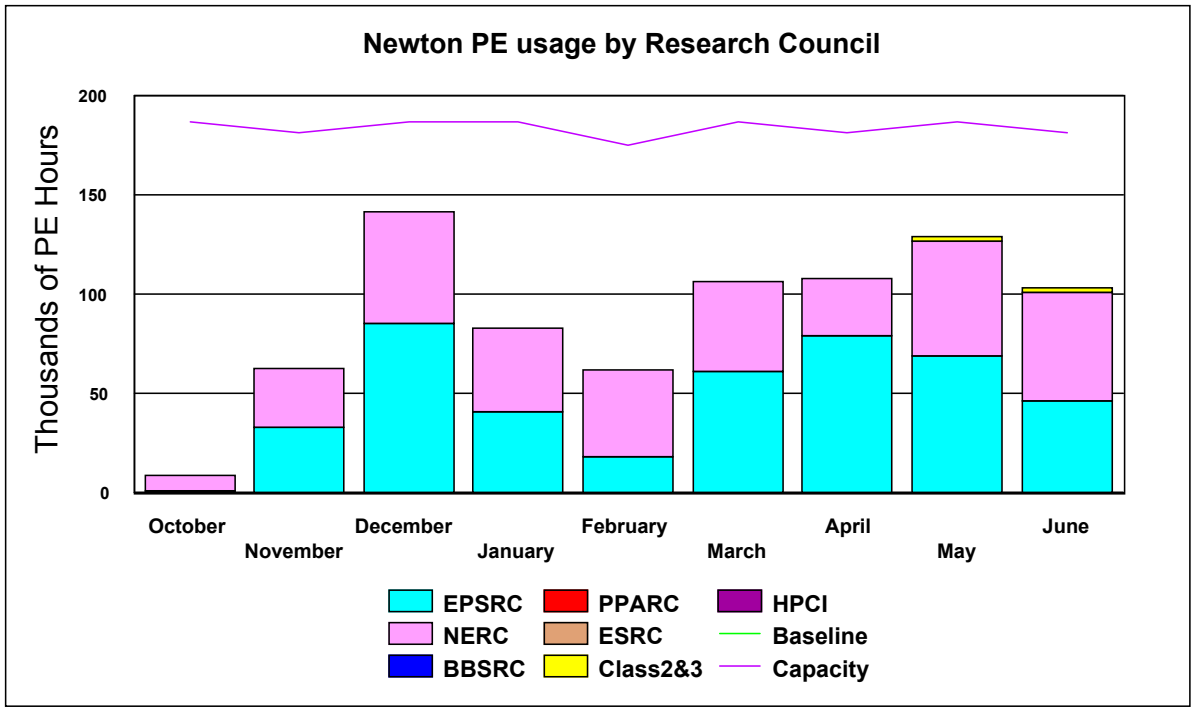

The next series of graphs illustrates the usage of the Medium Performance disk and HSM resources of the system.

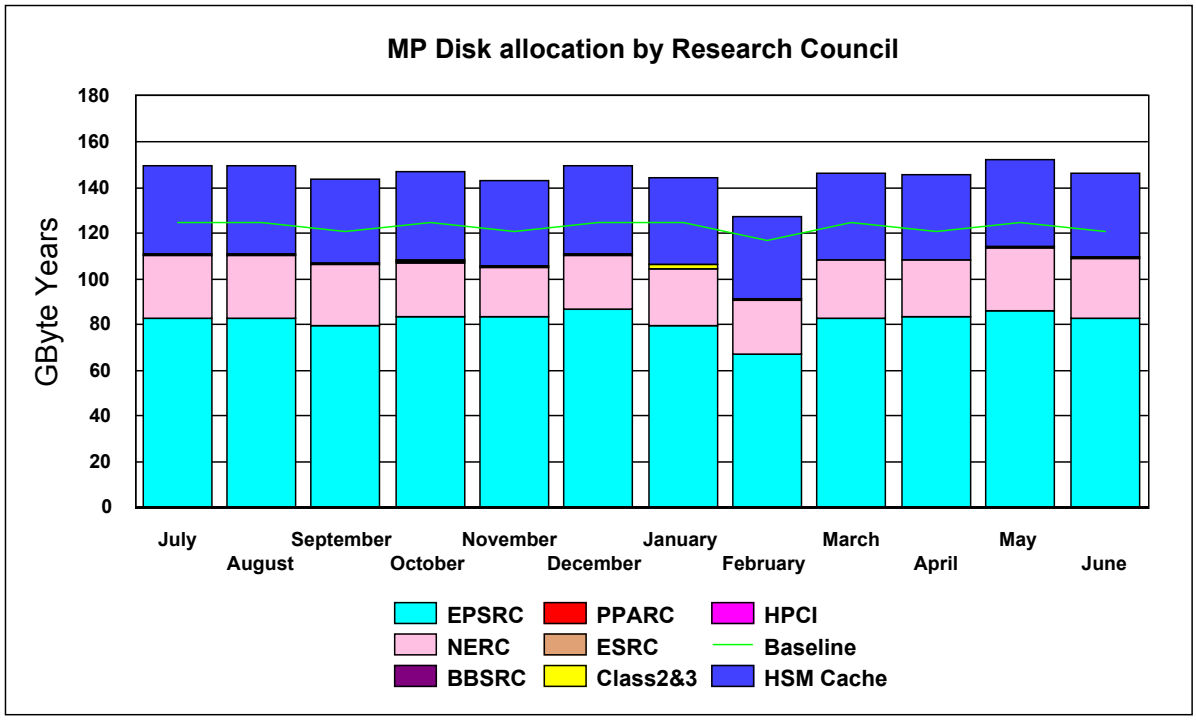

The above graph illustrates the historic allocation of the Medium Performance Disk on Fermat and the SAN.

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.

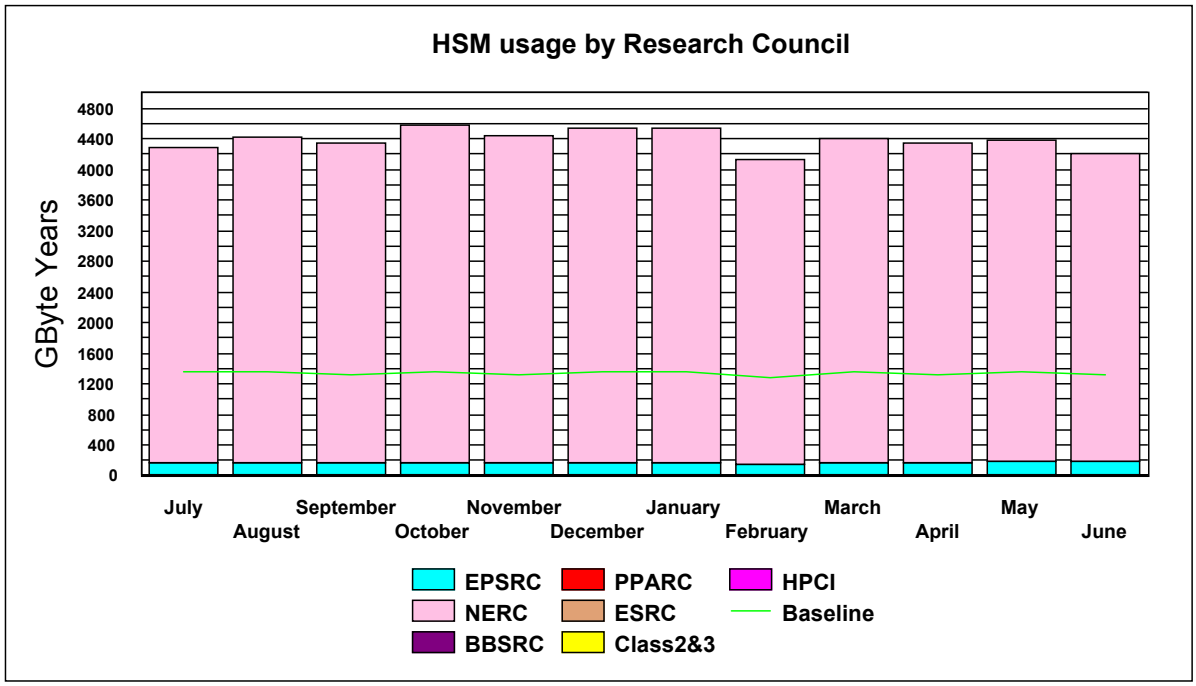

2.2.2 Guest System Usage

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 aproved 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 2004 Usage							
System	Consortia						Total
	cse085	cse076	cse133				
Green 256+ PEs							0
Green 384+ PEs							0
Newton 128+ PEs	98.82	388.23	9.94				496.99
Newton 192+ PEs							0
Total Tokens							**496.99**

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 second quarter of 2004, with usage exceeding baseline.

During the quarter there was a relatively balanced spread of work across all major systems.

Both the Propack operating system of the Altix and the Irix operating system of the Origins were upgraded at the beginning of June, in order to address reliability issues encountered during the quarter.

Issues

Some reliability issues were encountered during the second quarter of 2004, primarily affecting the Altix system Newton. Considerable effort was put into ascertaining the root cause of the problems, the majority of which were subsequently identified as operating system bugs, mainly relating to CXFS and the SAN. In order to rectify as many of these issues as possible, both the Irix operating system of the Origins and the Propack operating system of the Altix were upgraded to the latest stable release, and subsequently supplementary patches to the Altix Propack operating system were applied to target additional issues.. As a result of these efforts, all systems were significantly stabilised and overall job throughput has been improved. All systems in the CSAR service are being carefully monitored to ensure maximum reliability and uptime.

Plans

The Altix system Newton is due to be expanded in a three-phase operation. The first phase will see an additional 128 processor node added to the system, each processor having 4GB memory, 0.5TB memory in total. The second phase will see the memory on the original 256 processor system upgraded to 2GB per processor. Both first and second phases are intended to be carried out during the second half of July. The third and final phase of the expansion will be to add a further 128 processors to Newton, making it a 512 processor system in total. This final phase is currently planned for the second half of the year, with more details to be announced nearer the time.

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

13 applications for new CSAR projects were received, requesting a total of 429,511 service tokens.

3.2 New Projects

4 new CSAR projects were started with 65,274 service tokens being awarded in total.

3.3 Finished Projects

2 projects finished

3.4 Call for new applications to HPC

The call for new applications in HPC resulted in 8 applications to use CSAR resources, of which 5 were successful. These include research into the modelling of large scale communication networks, spacecraft force modelling for system design and supercomputing data mining.

3.5 Queries

A total of 181 CSAR queries were dealt with:

- 125 non-in-depth
- 46 user registration and admin
- 10 in-depth:

3.6 Service Quality Tokens

Two black marks were received when a job failed because the disk on which the home directory resided was full. This appeared to arise when a number of jobs generated significant output simultaneously (the disk typically had about 200GB free space). This was, therefore, seen as a very rare event, but additional disk currently being made available will also help to address such problems in the future.

3.7 CSAR Focus

The Winter/Spring edition has been completed and distributed. The Summer/Autumn edition is now being produced.

3.8 CSAR Website

The new CSAR website is nearing completion; it is hoped that the site will be available within the next few weeks. Notification will be sent to all users nearer to the completion date advising when the new site is due to go live.

3.9 User Steering Group

The 12th CSAR User Steering Group meeting is scheduled for 12th July 2004.

4 Scientific Application Support Services

4.1 Training and Education

The following courses have been given:

- Performance and Development Tools

4.2 Consortia Support/Software

Work has continued particularly on porting and optimising codes both for the SGI Altix service, Newton, and for the SGI Origin systems. There continues to be significant activity in this area

4.2.1 Castep

Benchmarking results on Castep were provided on behalf of the two groups: the Materials Chemistry consortium, and a group who were putting together a proposal.

4.2.2 NWChem

The problems with NWChem were due to high demands on I/O to the SAN; this could be alleviated using local /tmp disks. An alternative was developed which uses memory rather than disk to store the temporary files; this proved successful in reducing wallclock times to that of using local /tmp disks.

4.2.3 CPMD

The latest version of CPMD (3.9.1) has been installed and tested on Newton. This has been compiled with the latest versions of the Intel Compilers (8.0), MKL maths libraries (7.0 beta release) and MPT MPI libraries (1.9.1). A problem that occurred with some calculations in the previous release (3.7.2) no longer occurs with this new release. CPMD continues to perform well on the Altix.

4.2.4 VASP

Further work has been performed in rebuilding the latest version of VASP with the latest compilers. Additional compiler bugs have been reported. Problems were also observed when running across multiple partitions on Newton and using CXFS. This was found to result from a change in the buffering defaults in later versions of the compiler. Due to a couple of compiler issues, we currently recommend using the 7.1 compiler with VASP. Further optimisations are being performed.

4.2.5 Gaussian 03

The supported version of Gaussian installed on the CSAR systems (Altix and Origins) is Gaussian 03 Rev B.04. The previous version, Gaussian 98, is still available on the Origins, but is no longer supported. On the Altix, G03 can only be successfully compiled with a specific version of the Intel compiler (v7.1.008) although much newer versions of the Intel compiler are available. This continues to be investigated by Gaussian and Intel. In addition, G03 has also been linked to the latest SCSL maths libraries installed (currently 1.5.1.0). Some limited testing has shown G03 to be around 5.5x faster on a single Newton processor compared to a single Green processor.

4.2.6 Tomcat

Additional optimisation work has been done with the Tomcat code, on behalf of the NCAS consortium. This has been done to improve the communication in the advection section of the code. Work will continue to address load balancing issues with the chemistry section.

4.2.7 Stream

Parallel performance analysis has been performed on this code, used in a CFD project, resulting in the identification of bottlenecks in the code and in defining how to solve them.

4.2.8 Pchan

Further work was carried out on this code on behalf of the Turbulence Consortium, to enable it to scale well and to allow the possibility of using the expanded memory of Newton to solve large problems. This work is ongoing.

4.2.9 Epidemiology project

Researchers from a project awarded under the first call for new applications to HC visited Manchester to discuss how best to proceed in developing their codes for CSAR.

4.2.10 Class 3 Project

Discussions have been held with the researchers of a new Class 3 project, modelling recovery after damage in single word reading, prior to the parallelisation of a neural network code they wish to use.

4.2.11 Other consortia porting/optimisation

Debugging assistance was given to one researcher who visited Manchester. This resulted in the identification and reporting of a bug in TotalView.

4.2.12 Visualisation support

The work, on behalf of one consortium, in producing high quality images for publications from data generated on the CSAR systems was completed.

4.2.13 Totalview

TotalView software and licenses were updates.

4.3 DDT

We have been investigating the potential use of Streamline's DDT debugger on Newton and Wren. We had highlighted a number of problems with it, which have since been fixed. Recommendations on its purchase for the CSAR service are forthcoming.

4.5 Checkpointing on Green

Testing by CSAR has found significant problems with checkpointing on Green. This is currently under investigation by SGI.

5 Collaboration and Conferences

5.1 MRCCS Projects

5.1.1 Advanced Virtual Prototyping Research Centre

Lee Margetts gave a presentation to the SGI User Group held in Florida in May titled "Exposing the Computational Kernel: Harnessing the Power of HPC for Virtual Prototyping".

The SCGlobal showcase titled "Collaborative Finite Element Analysis" (Lee Margetts et al.) was accepted for the SuperComputing Conference to be held in November 2004. Unfortunately the joint MRCCS/Earth Simulator tutorial proposal (Mike Pettipher and Lee Margetts) titled "Multi-platform Scientific Application Development for High Performance Computers" that was submitted for the same event was unsuccessful. However, the Earth Simulator team have confirmed the intention to continue collaboration and resubmit for 2005. The tutorial may be given as an MRCCS workshop early in 2005.

5.1.2 NetSolve

NetSolve is an RPC based client/agent/server system that allows on to remotely access both hardware and software components. CSAR have installed version 2.0 of the software on our systems and are investigating its use, with a view to being able to launch MPI codes using it.

5.2 Events

Conferences/events attended included:

- Cray User Group
- SGI User Group
- SCC/HTPC Conference in Manchester
- Reality Grid annual workshop
- HPCx industry day
- Cray Technical Workshop in Bologna
- VECPAR

Manchester staff presented at many of these events, giving in total 10 presentations. Manchester also plays a key organising role in both the Cray and SGI user groups.

5.3 National Grid Service

Staff have been involved in setting up the National Grid Service, which started service at the beginning of April 2004. This includes both a data node at Manchester and the CSAR service. The work has included system management, running acceptance benchmarks and managing the website.

5.4 Access Grid Support Centre

The Access Grid Support Centre (AGSC), managed by UKERNA and run by the University of Manchester, was launched on 15 June 2004 in order to provide support for users and potential users. Support covers procurement advice, help on AGSC services and general troubleshooting as well as in-depth training through a series of courses. The AGSC also runs quality assurance tests to aim for improvements in the perceived quality for users and, in addition, a number of services designed to improve the Access Grid experience.

6 Added Value Services

6.1 International Conferences

The International Supercomputing Conference in Heidelberg was attended by Mike Pettipher, Robert Haines and Penny Richardson. A booth representing the University of Manchester illustrated the work done by the CSAR service. This year's conference theme was Applications, Architectures and Trends. Mike Pettipher attended a tutorial on Performance in HPC: Evaluation, Modelling, Benchmarking and Prediction.

Stephen Pickles presented a plenary talk about the TeraGyroid experiments at SC2003, and received the ISC award for the best paper on "Integrated Data and Information Management".

Preparations are being made for SC2004 in Pittsburgh in November, at which the University of Manchester will have an exhibition stand.

6.2 Visits

Manchester represented the UK on behalf of the DTI and hosted a visit from advisors to the French government who were visiting the UK to help formulate a French national HPC strategy. The visitors took away many ideas from the CSAR service and from the wider UK HPC perspective to influence their recommendations.

6.3 Visualisation

An exciting new collaboration between Access Grid technology and Manchester Visualisation Centre is currently being undertaken to create the first UK installation of passive stereoscopic access grid node. Consisting of six Christie Digital LX32 projectors, an 8x2 metre real-laced passive stereoscopic surface, drive by three dedicated nVidia QuadroFX 3000G graphics cards. High-end scientific visualisation employing the standard commercial packages, such as AVS and Amira, will be available within the Access Grid.