CSAR Service

Consolidated Management Report

2nd Quarter 2003

Management Summary

This is the consolidated Management report for the second quarter 2003 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 537 to date.

The workload on both the Origin 3000 (Green) and the Cray T3E (Turing) has been fairly evenly spread across the mid- to high-end ranges of PEs during the second quarter of this year, with the T3E running at approximately three-quarters capacity.

CSAR has been granted an 18 month extension of service contract until June 30th 2006. With this extension CfS is implementing a further technology refresh which will introduce a 256 processor Itanium-2 (Madison) based SGI Altix by end September 2003.

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 2003 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	2002/3											
	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March	April	May	June
HPC Services Availability												
Availability in Core Time (% of time)	99.2%	99.75%	98.75%	99.77%	99.25%	99.21%	99.46%	99.73%	100%	99.74%	97.66%	99.25%
Availability out of Core Time (% of time)	100%	100%	99.42%	99.52%	99.57%	100%	99.89%	100.00%	99.81%	99.81%	99.33%	99.9%
Number of Failures in month	0	1	2	1	1	0	3	1	1	1	4	1
Mean Time between failures in 52 week rolling period (hours)	365	381	381	398	417	515	487	487	515	548	461	548
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	<2	<1	<2	<2	<2	<0.5	<1	<2	<3	<1	<2
Administrative Queries - Max Time to resolve 95% of all queries	<2	<0.5	<2	<0.5	<0.5	<0.5	<1	<0.5	<1	<0.5	<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 formulae, based on the relative NPB performance of Turing, Fermat and Green at installation:

   Turing availability x 143/(143+40+233)] + [Fermat availability x 40/ (143+40+233) + Green availability x 233/(143+40+233)]

2. Mean Time Between Failures for Service Credits is formally calculated from Go-Live Date.

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	2002/3											
	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March	April	May	June
HPC Services Availability												
Availability in Core Time (% of time)	0	-0.039	0.039	-0.039	0	0	0	-0.039	-0.058	-0.039	0.078	0
Availability out of Core Time (% of time)	-0.047	-0.047	0	-0.039	-0.039	-0.047	-0.047	-0.047	-0.047	-0.047	0	-0.047
Number of Failures in month	-0.009	-0.008	0	-0.008	-0.008	-0.009	0	-0.008	-0.008	-0.008	0	-0.008
Mean Time between failures in 52 week rolling period (hours)	0	0	0	0	0	-0.008	0	0	-0.008	-0.008	0	-0.008
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	-0.016	0	0	0	-0.019	-0.016	0	0.016	-0.016	0
Administrative Queries - Max Time to resolve 95% of all queries	0	-0.019	0	-0.019	-0.019	-0.019	-0.016	-0.019	-0.016	0	-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:	-0.05	-0.08	-0.01	-0.07	-0.05	-0.06	-0.06	-0.09	-0.09	-0.07	0.00	-0.06
Quarterly Service Credits:			-0.14			-0.19			-0.24			-0.14

Table 3

1.2 No. of Registered Users

The current position at the end of the quarter is that there are 537 registered users of the CSAR Service.

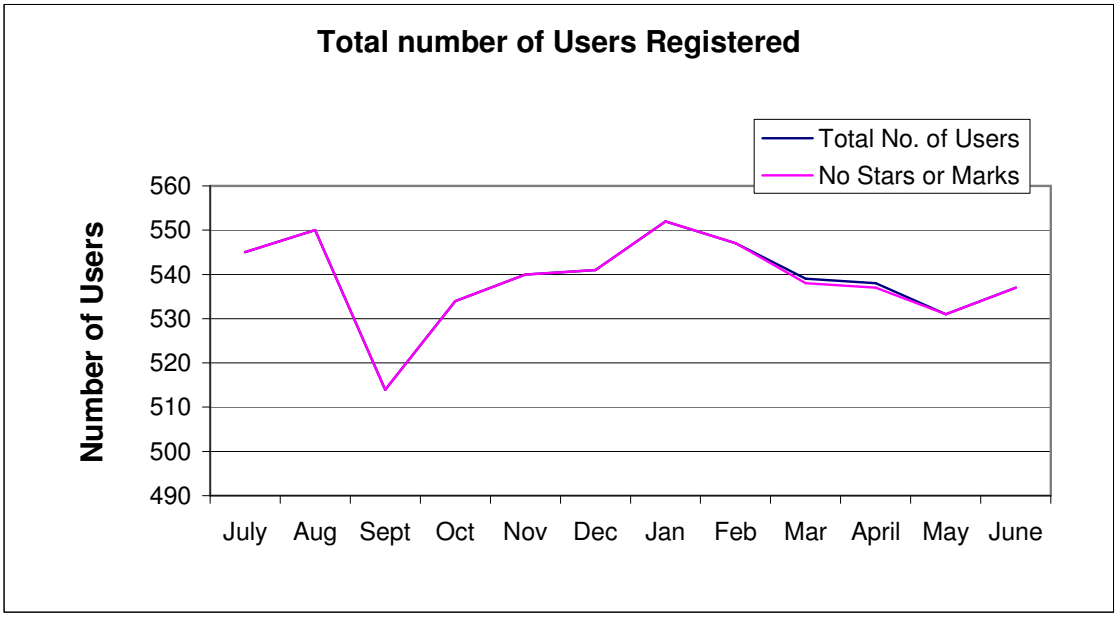

1.3 Service Quality Tokens

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

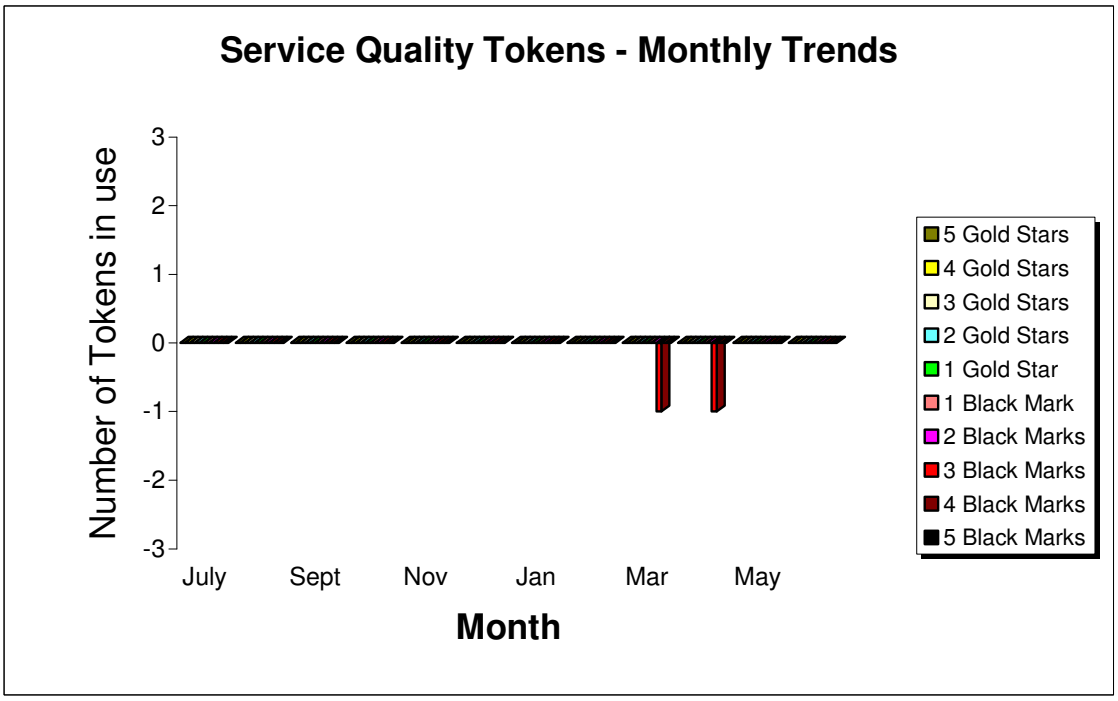

Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

A the end of the quarter there were no black marks or gold stars 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) MPP (T3E) Usage by month, showing usage each month of CPU (T3E PE Elapsed Hours), split by Research Council and giving the equivalent GFLOP-Years as per NPB. The Baseline Capacity is shown by an overlaid horizontal line.

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 Capacity is shown by an overlaid horizontal line.

c) High Performance Disk (T3E) allocated for User Data by month, showing the allocated space each month in GBytes, split by Research Council. The Baseline Capacity (1 Terabyte) is shown by an overlaid horizontal line.

d) 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.

e) HSM/Tape Usage (T3E) 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 PE, CPU, disk and HSM allocations and usage.

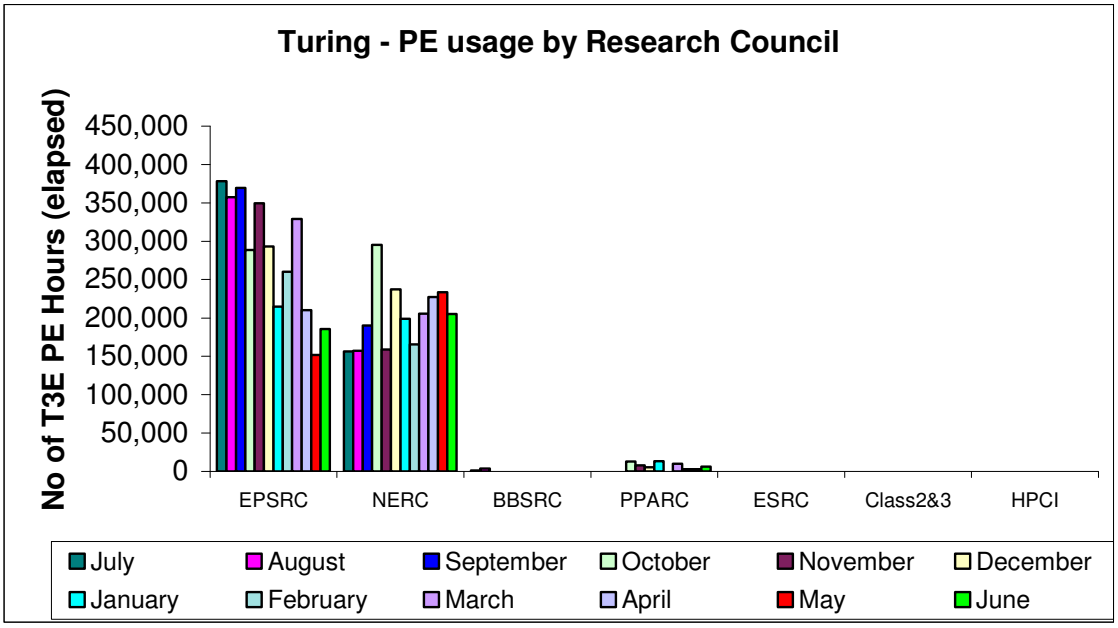

Turing PE usage is shown by Research Council during the last 12 months of service in the above chart.

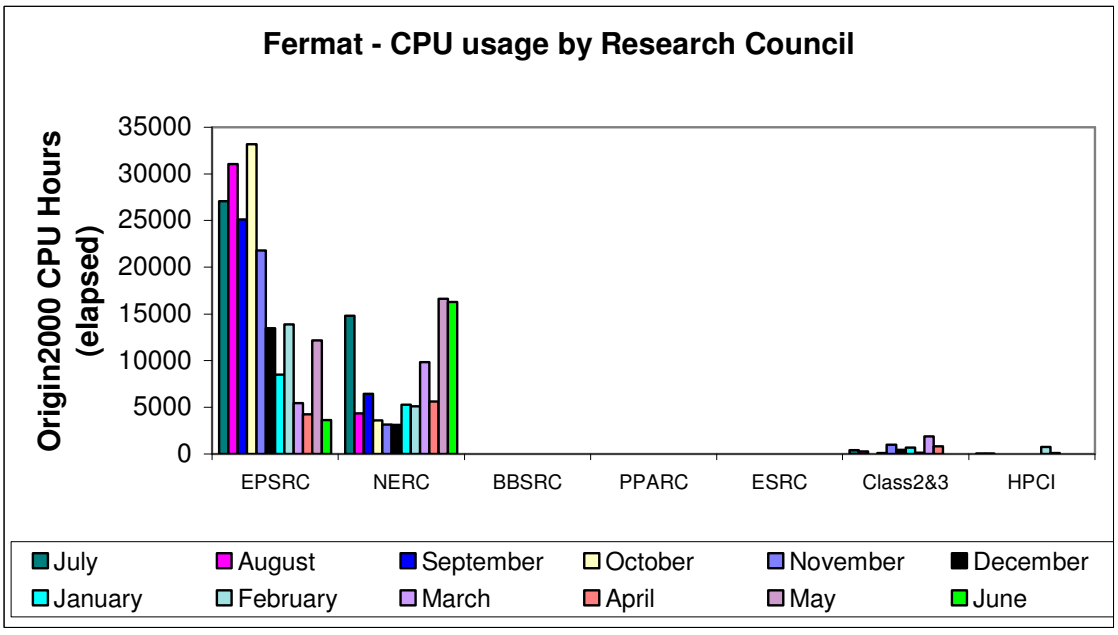

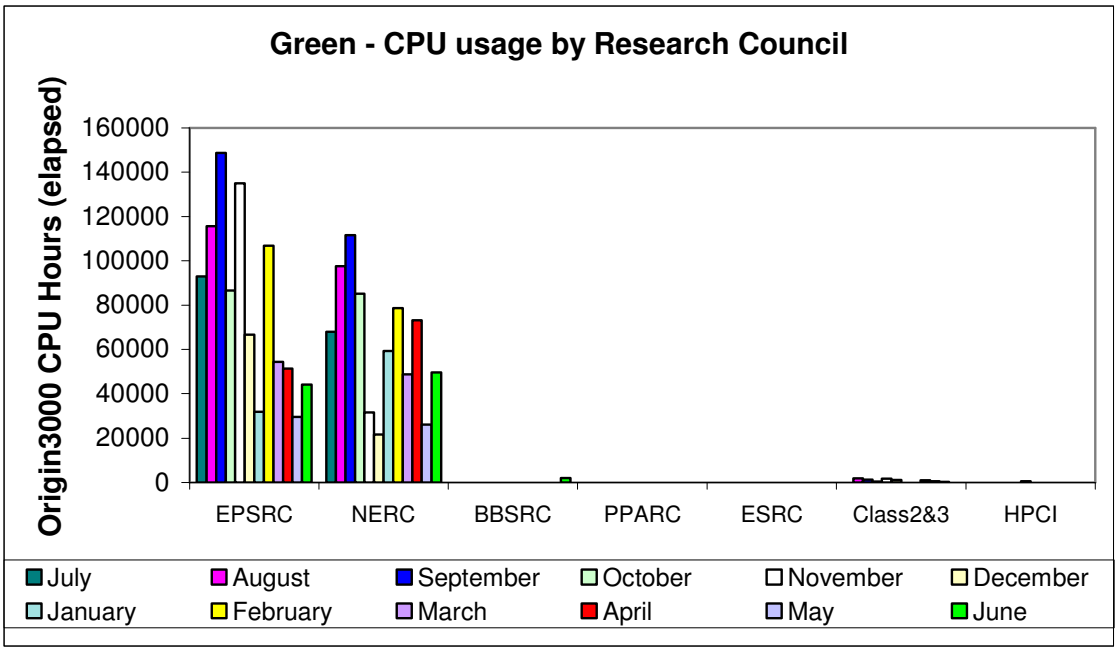

Usage of the two SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service in the above two charts.

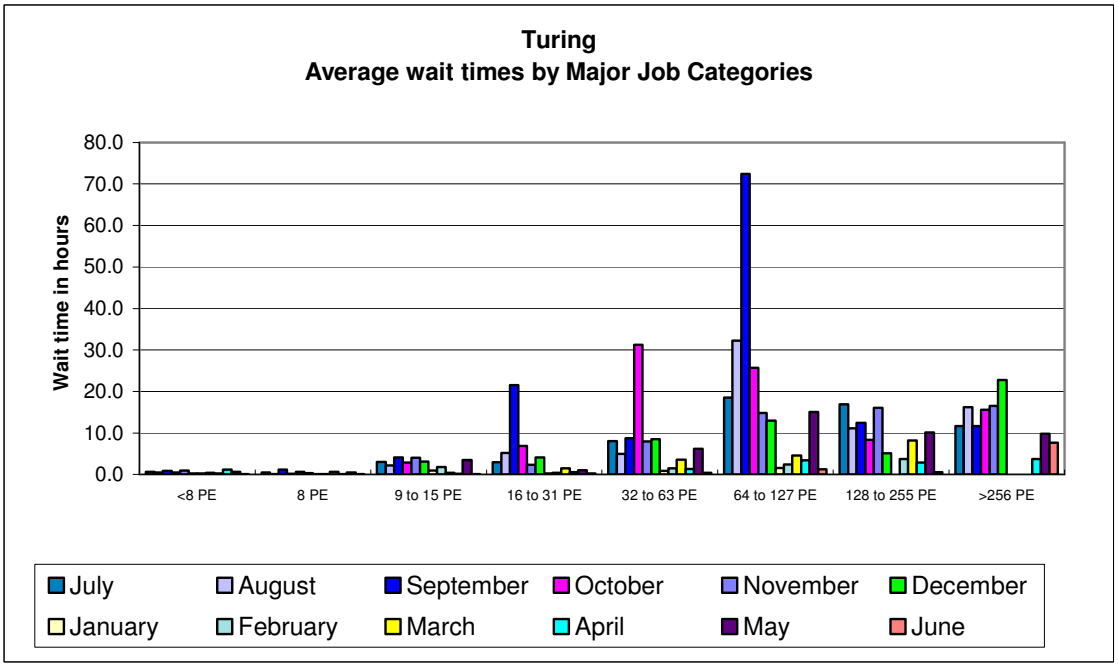

The above chart, and the one below, show the wait time trend in hours on the Turing system.

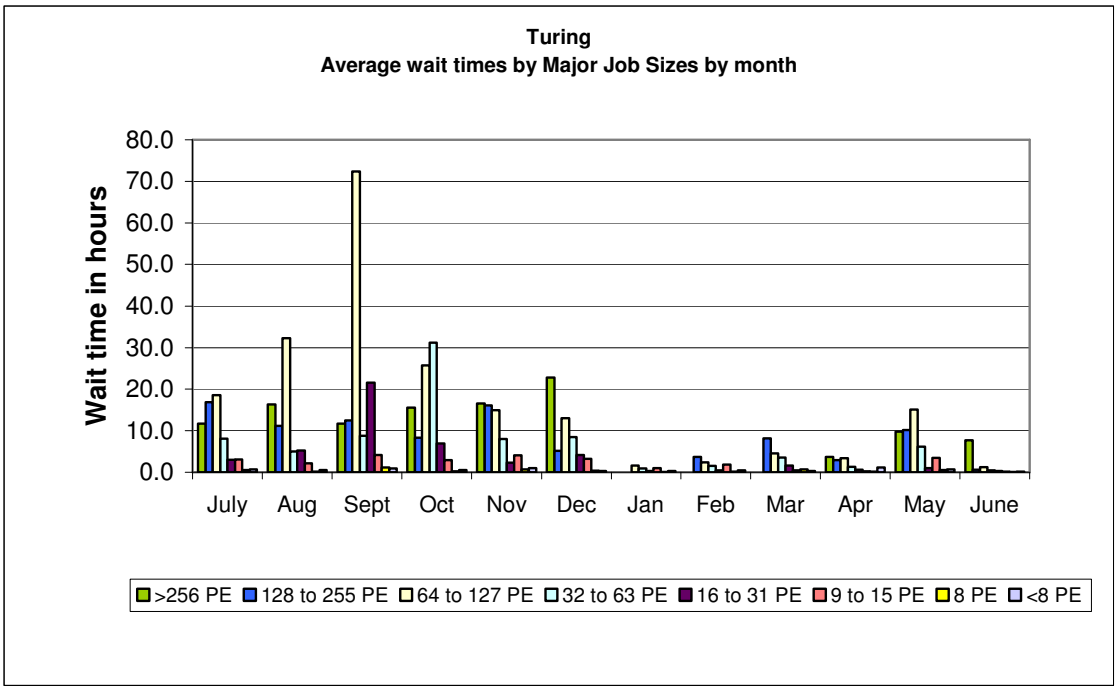

The following two charts show average wait times in hours for the quarter on the Origin 3000 (Green).

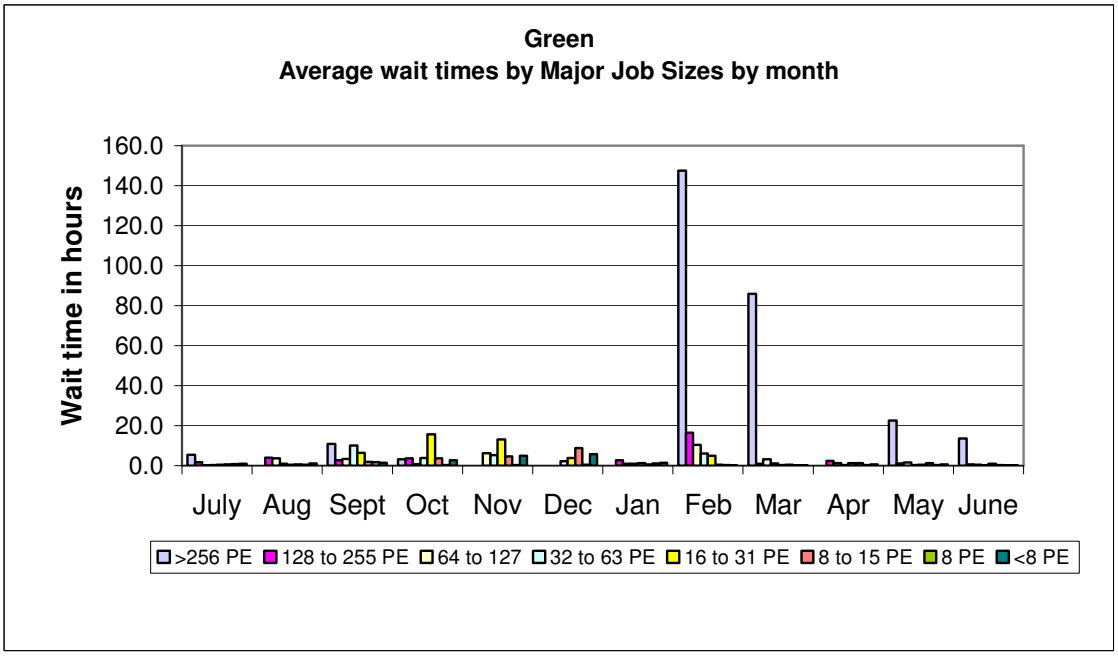

The next three charts show the percentage PE time utilisation by the major job categories on the Turing, Green and Fermat systems for the 2nd quarter 2003.

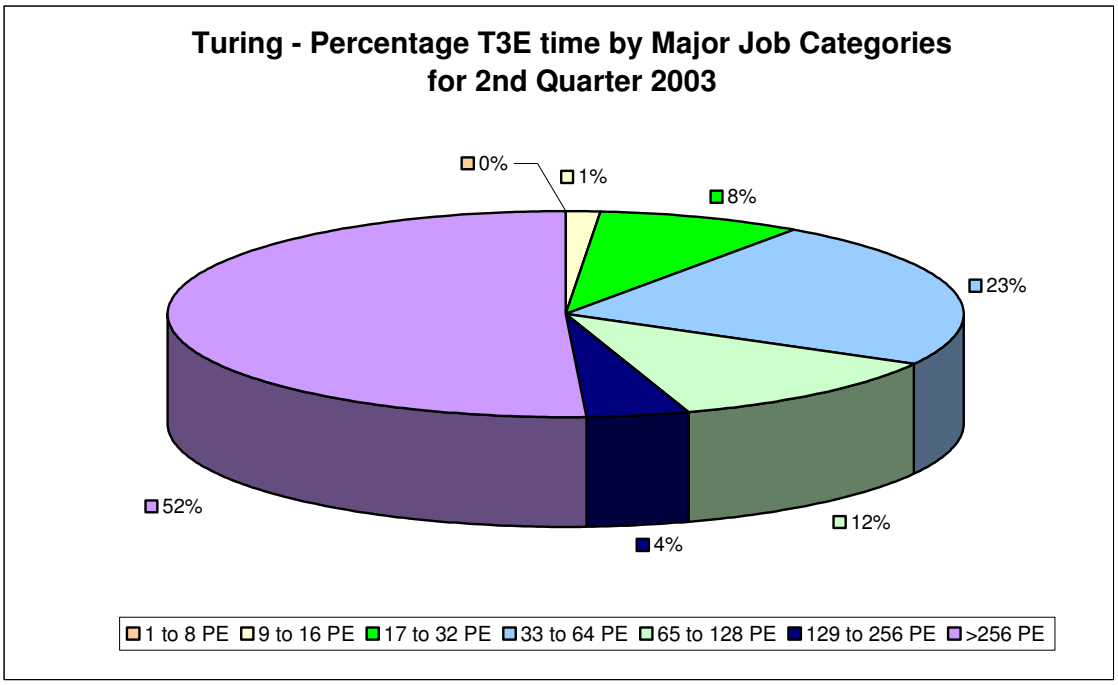

The workload on Turing for the second quarter was predominantly in the greater than 256 PE range.

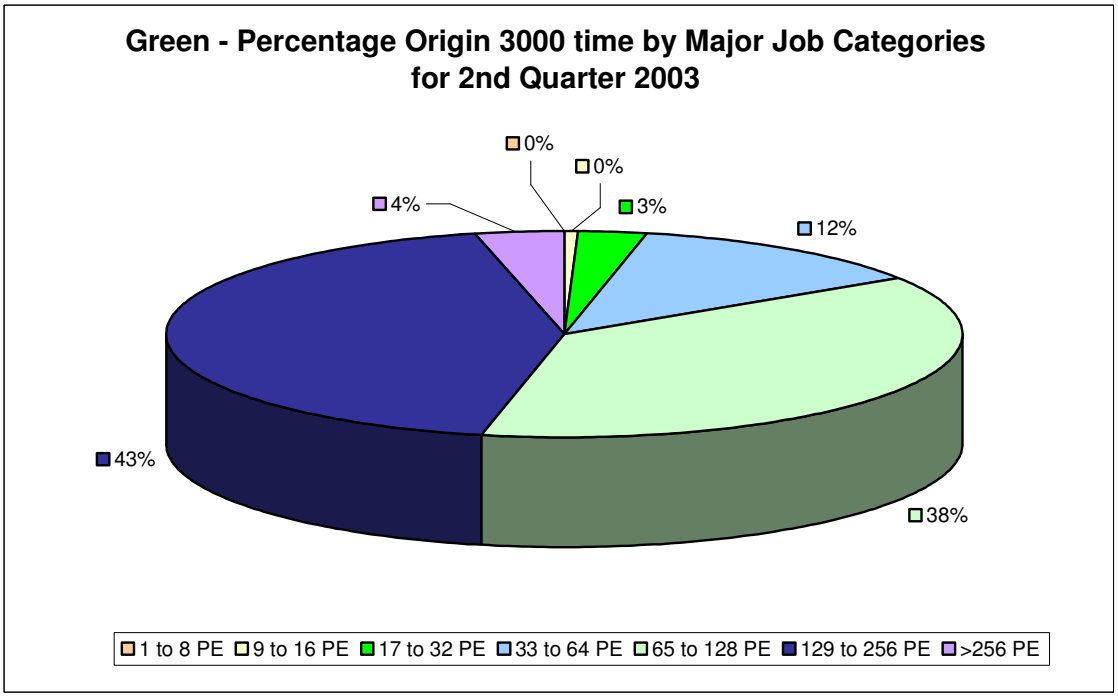

On Green, the 2nd quarter has seen a varied workload, relatively evenly spread across the mid- to high-end ranges of PEs.

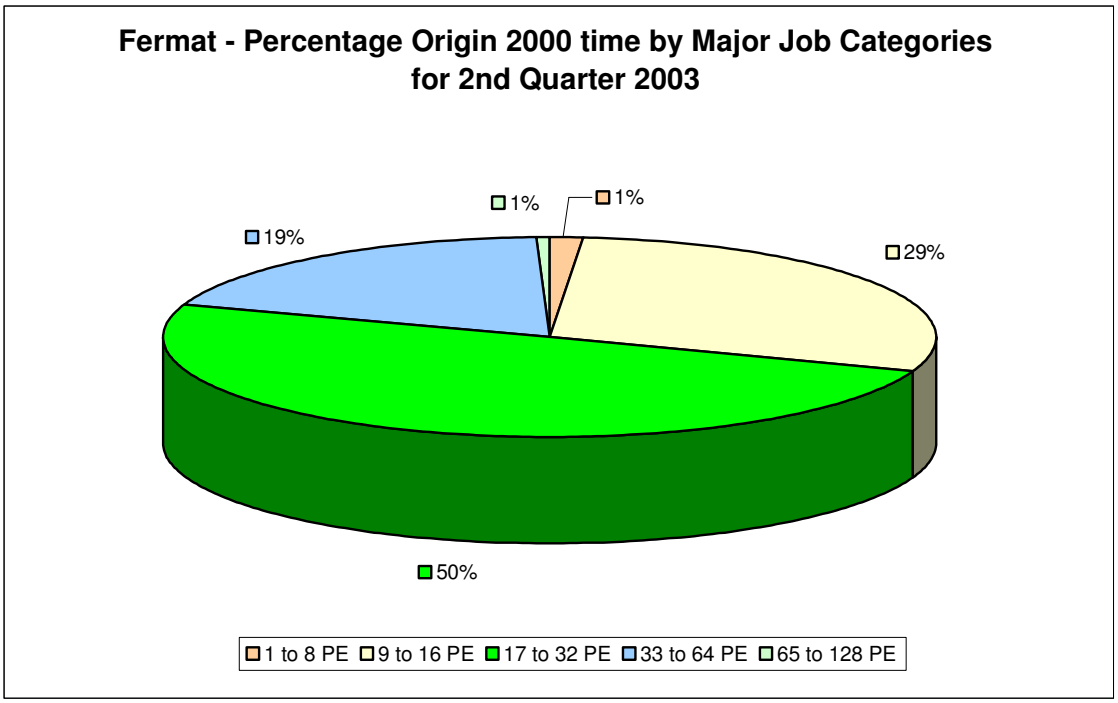

On Fermat the greatest proportion of jobs during the quarter were in the 17 to 32 PE range at 50%.

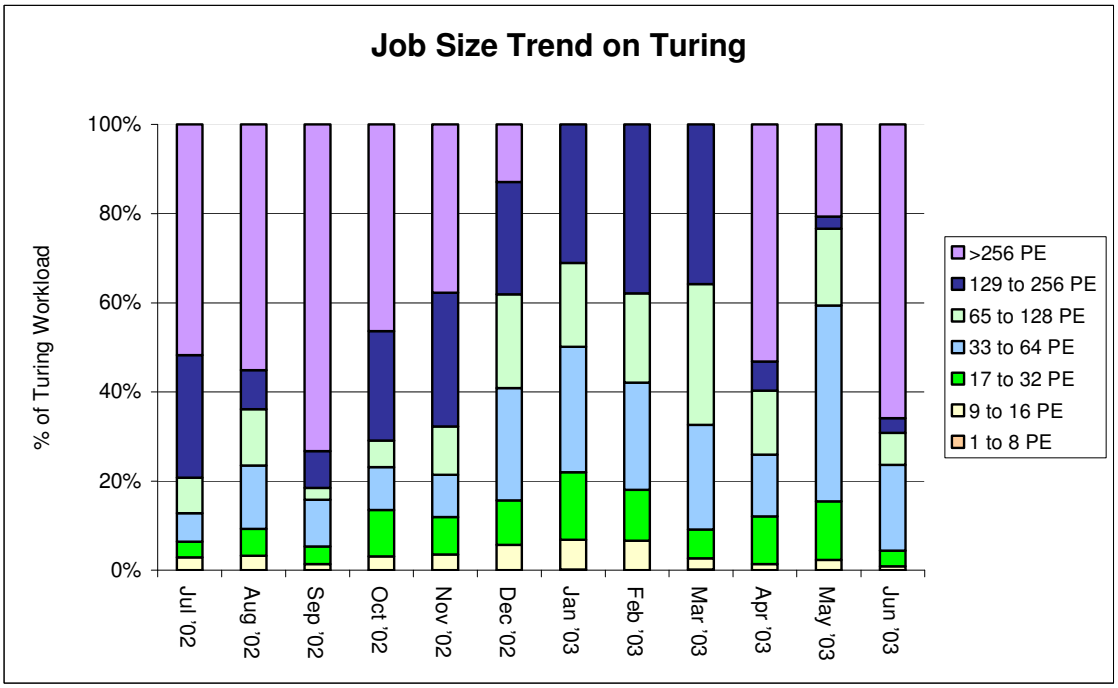

The last few months have seen increasing usage of the T3E for jobs above 256 PEs.

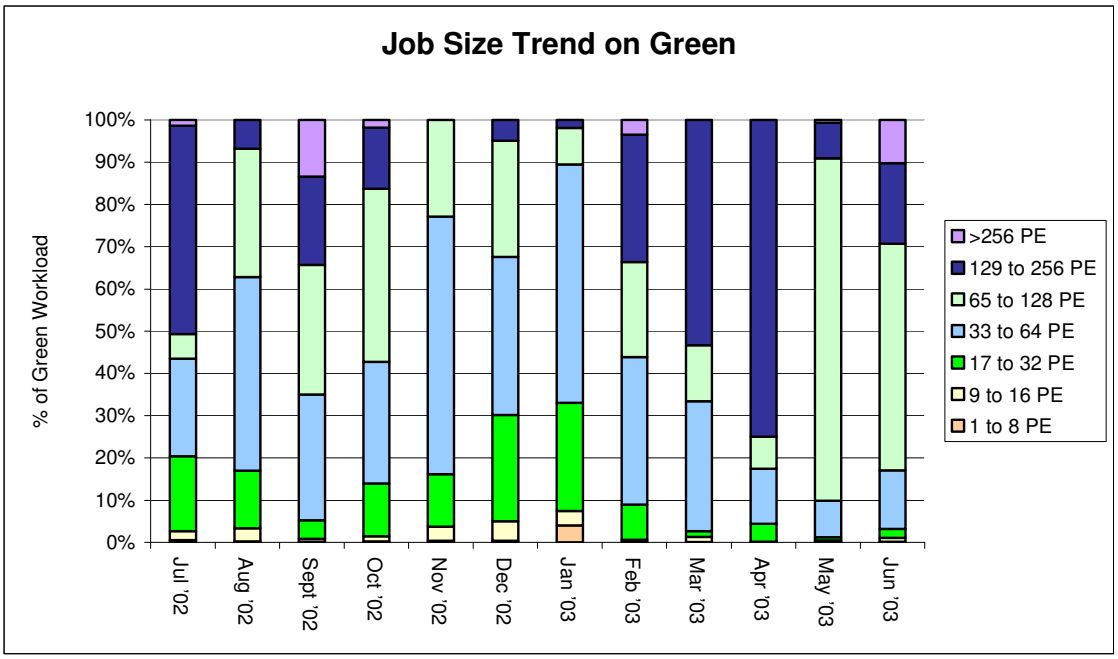

Usage on Green is tending more to the higher-end range of PEs now that all PEs are dedicated to LSF.

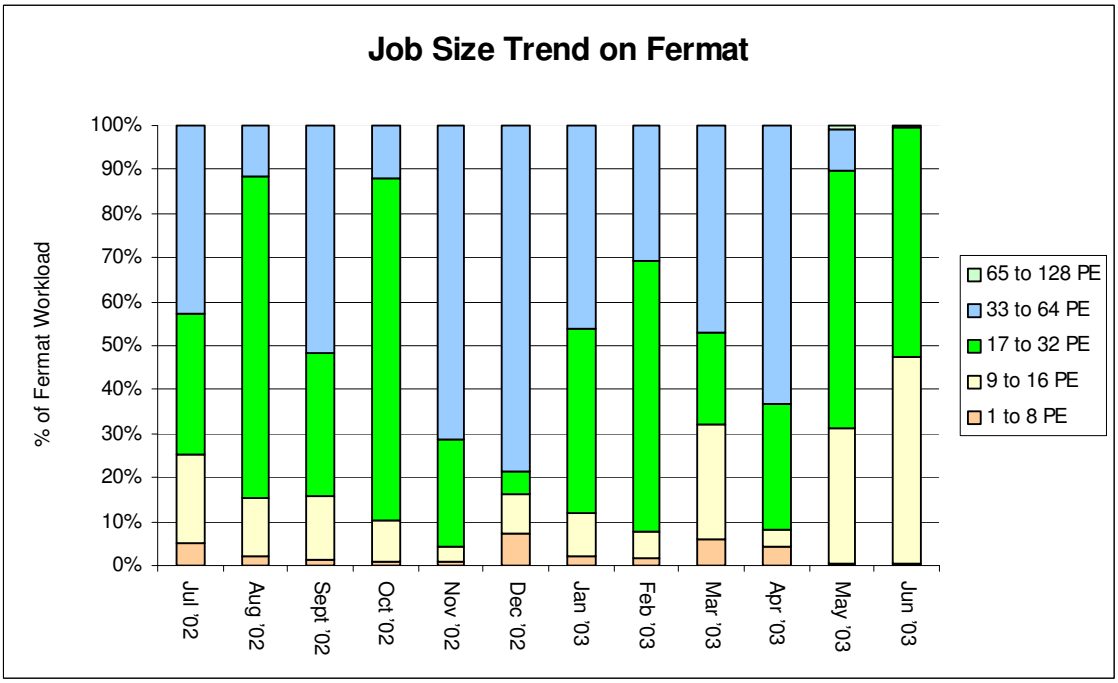

The trend on Fermat currently is that the greatest proportion of the workload is primarily in the mid-range of PEs.

2.2 System Usage Graphs

In all the Usage Charts, the baseline varies dependant on 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.

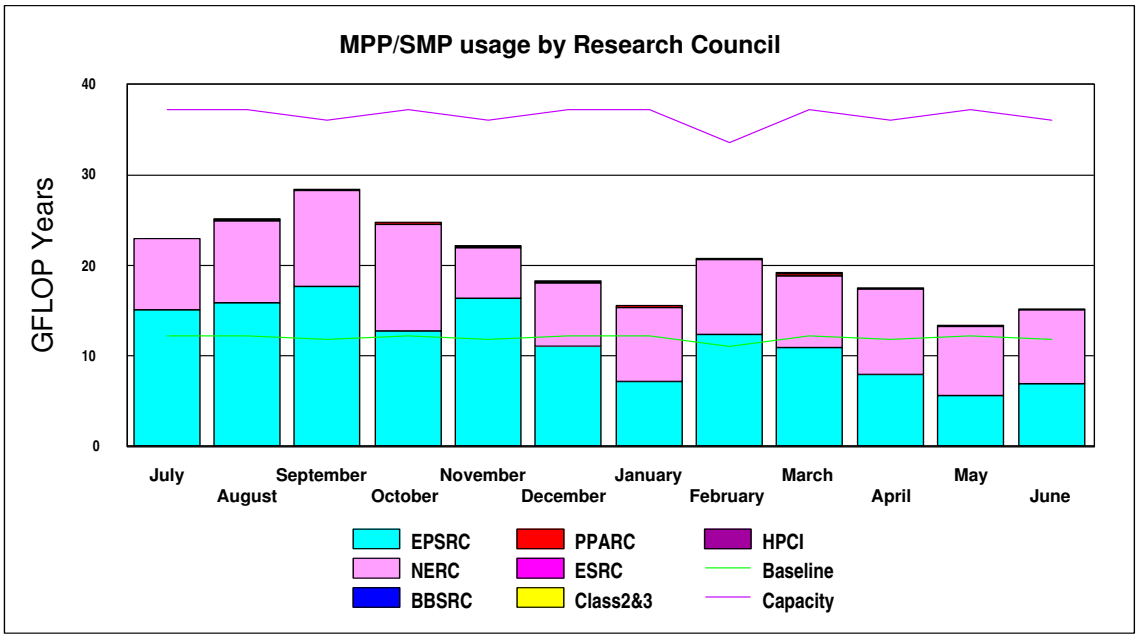

The graph below shows the same service utilisation by major system.

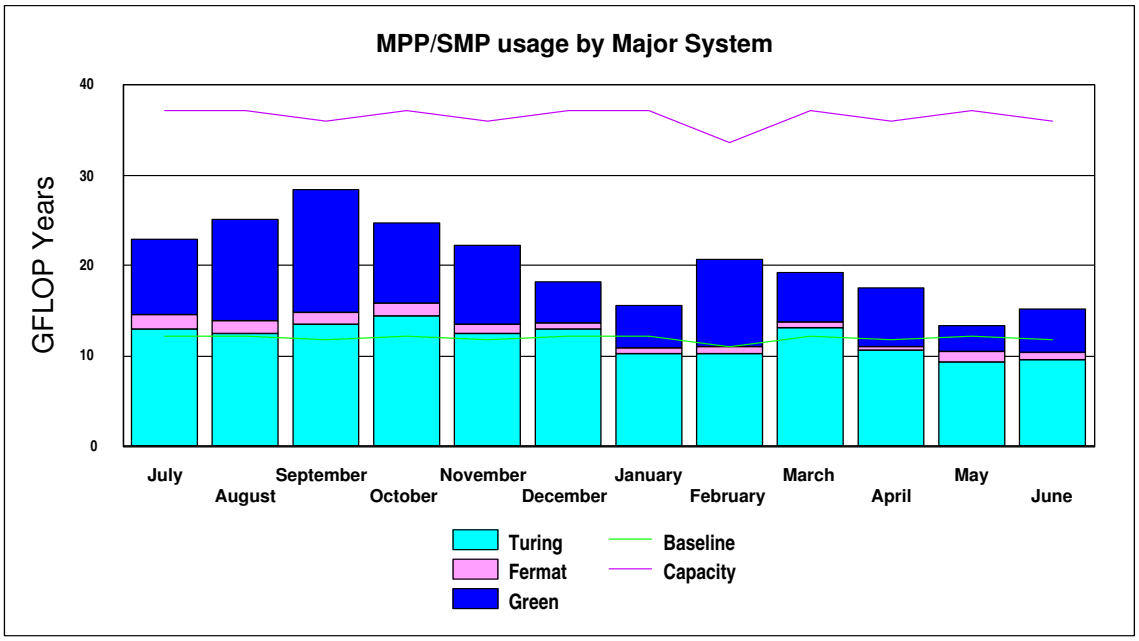

The next chart shows shows the historic interactive usage of the 'baseline' Fermat system (equivalent to 16@250Mhz CPUs) up to the end of February 2003, at which point the interactive usage was transferred to Wren and Fermat became a batch-only system.. Eight of the higher speed 500Mhz CPUs in the Origin 300 system (Wren) deliver the baseline capacity equivalent to that which was previously available on Fermat for interactive usage.

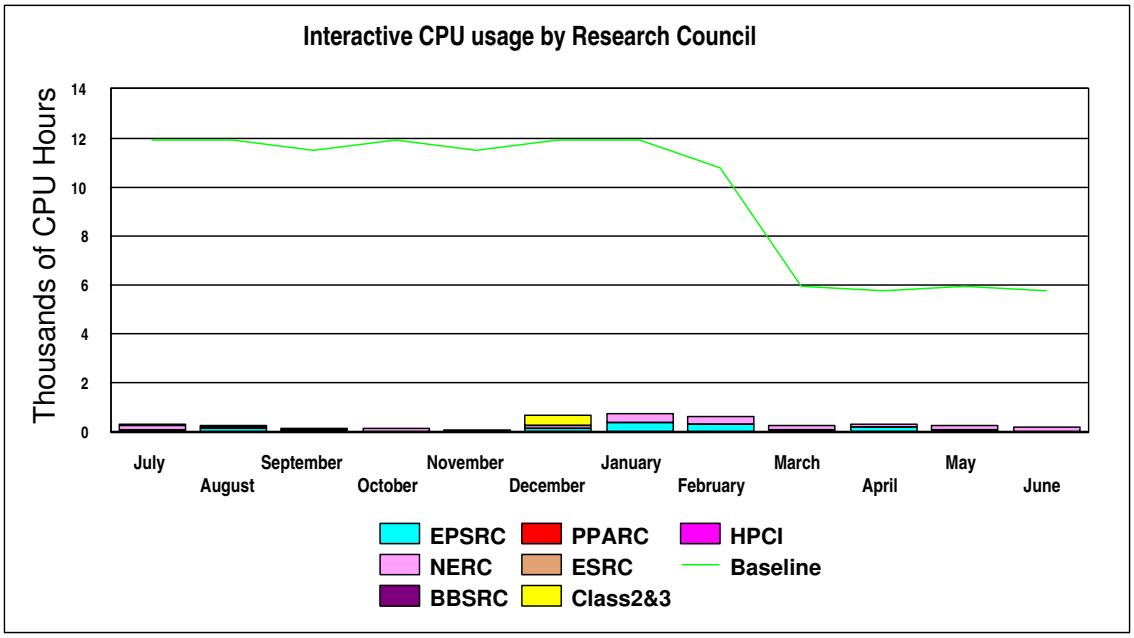

The next series of graphs illustrates the usage of the disk and HSM resources of the system.

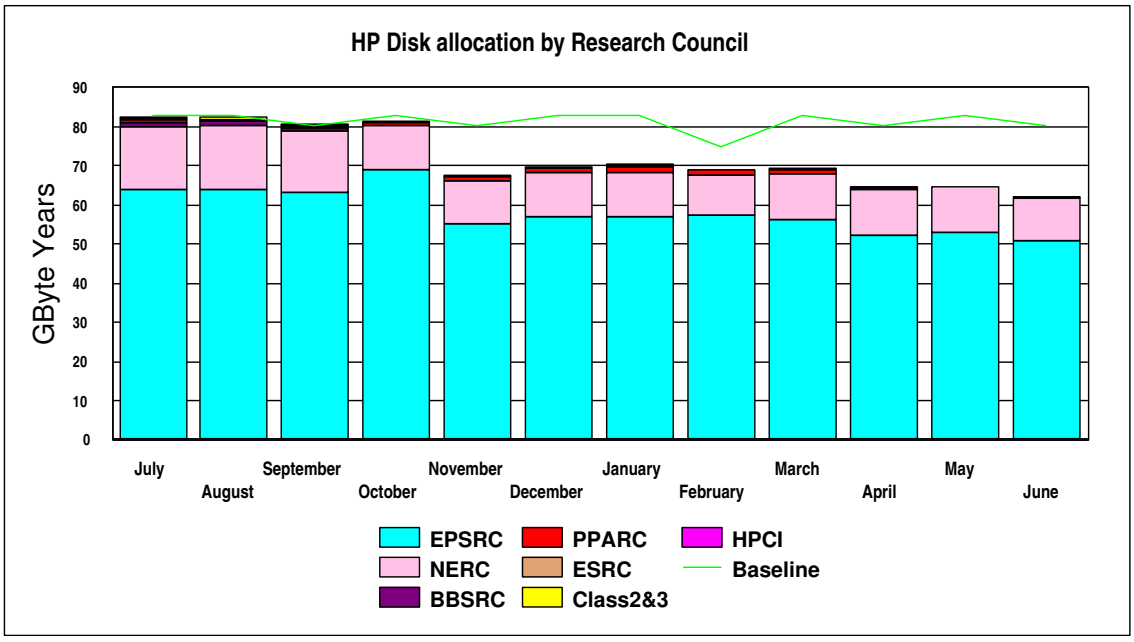

The graph above illustrates the historic allocation of the High Performance Disk on Turing.

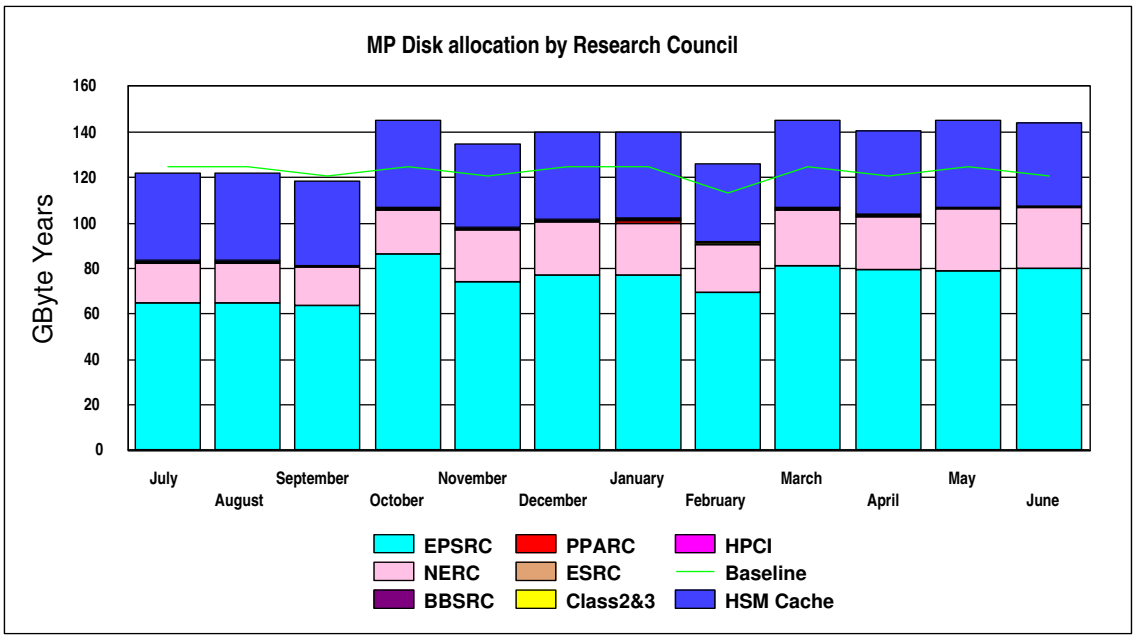

This graph illustrates the historic allocation of the Medium Performance Disk on Fermat.

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

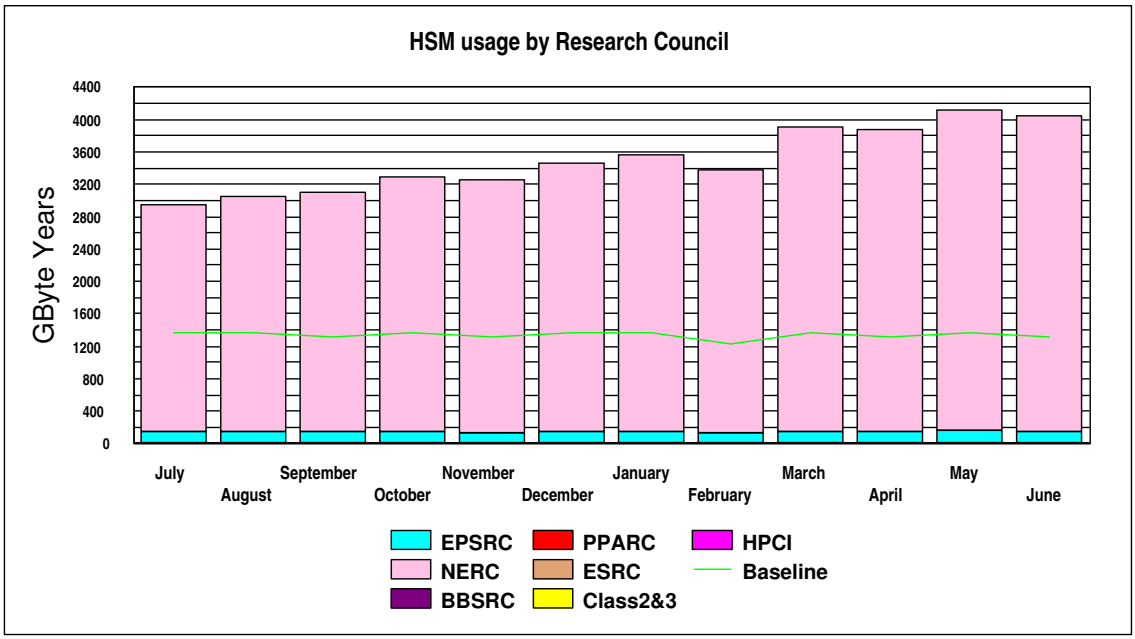

2.2.2 Guest System Usage

There is currently no Guest System usage.

2.3 Service Status, Issues and Plans

Status

The service has been reasonably utilised throughout the second quarter of 2003, with usage exceeding baseline.

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

The four additional fibre-attached tape drives on Fermat have helped to improve the response and reliability of the Data Migration Facility; response times have improved since this addition to the service.

Issues

There are no issues to report for this quarter.

Plans

A 32-PE Altix system (Reynolds) has been introduced as a forerunner to the 256-PE Altix system (Newton) due for installation during September 2003. Further details will be announced as they become available.

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

7 applications for new CSAR projects were received in the second quarter of 2003, requesting a total of 27080 service tokens.

3.2 New Projects

3 new CSAR projects were started with 7744 service tokens being awarded in total.

3.3 Queries

Overall, 203 Class 1, 2 and 3 queries relating to the CSAR service were received by the Helpdesk between 1st April and 30th June 2003.

3.4 Service Quality Tokens

Three black marks were recorded for the poor performance of /hold. Action was taken to address these problems, including the addition of new tape drives. The black marks were subsequently removed.

3.5 Annual Report

The CfS Annual Report is currently being prepared for publication.

3.6 CSAR Focus

Work is underway on the Summer 2003 edition which is due out in July.

3.7 CSAR User Steering Group

The tenth six-monthly meeting of the CSAR User Steering Group was held in June using the Access Grid facilities at Imperial College and the University of Manchester.

The major topics included:

- CSAR service developments – the introduction of the Altix system and the retirement of Turing.
- Issues raised by the user community either via Dr Lois Steenman-Clark, in her capacity as the Chair of the User Liaison Forum, or by the User Survey or other routes.

3.8 CSAR Website

Fiona Cook is managing the creation of the new CSAR website. The frontline support team have attended various courses in preparation and work is now underway. The launch of the new website is intended to coincide with the launch of the Altix system.

3.9 Launch of Altix/Extension to CSAR contract

A letter has been issued to all Principal Investigators this quarter to inform them of the extension to the CSAR contract and the service developments – i.e. the launch of the Altix system and the removal of Turing. The CSAR webpages have been updated accordingly.

4 Scientific Application Support Services

4.1 Training and Education

Kevin Roy and Neil Stringfellow visited the atmospheric group at Cambridge to teach a variety of aspects of using the CSAR Origin systems. The courses included MPI programming, OpenMP programming, serial optimisation, parallel optimization and some of the tools available.

4.2 Consortia Support

4.2.1 Occam Group

Jeff Blundell works with a number of codes the major one being an OpenMP code, q-gcm. This was already a well-optimised code. CSAR staff were involved in analysing this code and providing better solutions where possible to improve both serial and parallel performance.

For this case the work performed involved looking at cache usage and merging loops where appropriate, it also looked at data access patterns and reordering some of the loops. The parallel performance looked at ways of increasing the scalability of some routines which was achieved by a combination of merging some parallel sections, reordering statements to use data from local cache and memory compared to remote cache and memory. More of the code was made parallel particularly 2 routines which were insignificant in the serial case but since they performed worse as the processor count grew, it became a serious performance problem. Also use was made of the automatic parallelizer was made to investigate other regions that were parallel.

Jeff Blundell left with a more optimised and parallelized version of his code with the only bottleneck being the FFT area of his code which involves 2 transposes and an FFT. There are inherent problems in getting this section to scale so he is looking at other methods to achieve this.

4.2.2 Atmospheric Chemistry groups

Dr Glenn Carver and Dr Martyn Chipperfield work on very similar codes in the atmospheric chemistry discipline. They are part of the NERC user base and had made heavy use of the Fujitsu, with the removal of the Fujitsu they had to move to other systems. CSAR staff aided this group in getting codes running, tuned and parallelized for the Origin 3000 and indeed any scalar based MPP machine.

The two groups use codes from a common source and as the science diverged slightly the codes separated and are maintained by the different groups.

The codes had been highly vectorised so in order to get the same level of performance it was necessary to parallelize the codes.

This is where the two groups took different paths. The environmental chemists at Leeds who ran a code called Slimcat were looking for a quick solution to this problem that would give them at least the performance of fuji if not better, for this reason they chose to port to the Origin and use OpenMP.

The environmental chemists at Cambridge use a code called Tomcat and already had a choice; they had an MPI model based on a previous version of the code or they could work with the parts of the code that were common with Slimcat and add in the sections of code that were relevant only to Tomcat.

Slimcat

The code now exhibits good scalability but is limited by the number of height levels they choose to model their part of the atmosphere on, this effectively limits their scalability to 24 processors. However on 24 processors the code scales very well, some tests have been performed using a slightly bigger problem but input datasets for these runs are limited.

Although cost was not the primary concern for this group they have seen that they can achieve 4 times as much science for the same quantity of generic tokens.

Tomcat

The Cambridge group has about 20 staff who work on both the Slimcat and Tomcat models, but whose experience of parallel systems and parallel programming was limited. The second visit addressed this and at their request we ran a 2 1/2 day course on all aspects of parallel code development, from debugging a code to parallelizing a code with OpenMP or MPI visiting serial and parallel optimisation techniques and serial and parallel profiling. There was a further day spent looking at the main code, Tomcat, profiling with Vampir analysing the output and discussing ways to improve the code.

The code has already been tested on model sizes an order of magnitude higher than what they were capable on fuji and can run these simulations quicker, allowing more accurate representation of the earths atmosphere and hence more accurate predictions.

The results from the Slimcat code showing a better performance and price/performance ratio were a big plus. The portability of the new MPI code for the Tomcat group together with some good scalings and the potential for expansion and improvement within the code has provided new opportunities for the group.

4.2.3 Optimisation of R-Matrix software

CSAR Staff have continued to be involved in the optimisation and development of R-Matrix software in collaboration with Queen's University Belfast. There are three areas of the code that we have begun work on

- Redistribution of local matrices to global matrices
- Optimized blas operation for $D = \alpha A \times B + \beta C$
- Development of efficient, functionalized diagonal update subroutine.

Part 1 is a direct follow-up to previous work with this group, which exposed the distribution of arrays from a single processor to a 2-d block cyclically distributed array held on all processors of a Blacs process grid. The consortia use routine pdgemr2d for this purpose, it is in some sense an overextension of the functionality of this routine but allows simple redistribution. A parallel redistribution routine is being developed for this application. There are issues in terms of data files, since a single data file is required for compatibility with previous and subsequent steps in the R-matrix software.

Part 2 is necessary because the normal blas operation

$$C = \alpha A \times B + \beta \times C$$

necessarily overwrites original matrix C. As part of the R-matrix propagation stage, a sub array of the global R-Matrix is sent to DGEMM and the result copied into another sub-array. If the standard DGEMM is used, important data that would remain in sub array C is overwritten, and an additional array D is used, which is expensive in terms of memory and performance. Hence a blas operation of the form

$$D = \alpha A \times B + \beta \times C$$

is required. The implementation of this routine was trivial, but making performance match that of the highly vendor-optimized standard DGEMM is tough.

Part 3 is a more open-ended look at the way that functionality in parallel numerical library routines can be captured to form part of a library of optimised tools. This is extremely useful for any future work, a policy decision has been made to capture as much functionality as possible to reduce programming complexity and to simplify further work in this area (of which there seems to be a great deal).

4.2.4 EROS UK code

CSAR staff are currently performing some in depth optimisation work for the consortium at the University of Glasgow on the EROS UK code. Analysis of this code shows that a dependence on MPI collective communication routine MPI_ALLReduce within a global sum calculation is inhibiting performance

considerably. We are working on a general high-performance replacement to the global sum that uses 1-sided communications.

4.2.5 Netsolve/Gridsolve

Netsolve/Gridsolve is being investigated with the intention of offering a Netsolve service in the future.

4.3 Porting to the Altix

A variety of codes, both third party and users own codes, are being ported by Manchester and SGI staff, and SGI is providing facilities for other code developers to work with Altix machines.

4.3.1 Ported by SGI

The following codes which we consider to be important, given our current and potential user base, have been ported and optimised by SGI:

Castep
SGI report that Castep from Accelrys has been ported and optimised for the Altix/Itanium 2 platform and they have produced timing and scalability figures.

NAMD
SGI report that NAMD, a highly scalable molecular dynamics package from the University of Illinois at Urbana-Champaign, has been ported and optimised for the Altix.

Gaussian
A version of Gaussian is available for the Altix platform, and the software has been ordered.

NWChem
NWChem from Pacific Northwest National Laboratory has been ported to the Altix.

4.3.2 Ported by the University of Manchester

Codes which have been ported by Application/Optimisation team at Manchester:

LB3D
An initial port of the 3D Lattice Boltzmann code which is used in the RealityGrid project has been completed. Initial results show good timings and scaling, and compare favourably with other architectures. The code is currently compiled with full compiler optimization flags (-O3 -ipo), and it is intended to apply further optimisation to this code in the near future.

Tomcat
The Tomcat code has been ported and runs to completion for the test provided.

CPMD
The Car-Parrinello Molecular Dynamics code CPMD has been ported. A performance analysis has also been performed using the different numerical libraries provided by Intel and SGI and also a variety of different compiler options available. The initial results using a small test problem appropriate to the size of the machine indicate that a 16 processor run executes 7 times faster than Turing and 3 and a half times faster than Green. The code also shows good scalability with 16 processors showing a speedup of 14.

DL_POLY
The DL_POLY code written by Bill Smith at Daresbury Laboratories is a major molecular dynamics code used by the UK chemistry and physics community. A new version of this code, DL POLY 3, has been written to be scalable up to hundreds of processors. DL POLY version 2.13 has been ported to the Altix by Neil Stringfellow, but the newer 2.14 version and a port of DL POLY 3 still need to be done.

H2MOL
The H2MOL code from Queen's University of Belfast has been ported. The code show good scaling, but more work is being done on single processor optimisation.

NetCDF
NetCDF has been ported and this should allow improved portability for many codes.

FFTW
FFTW has been ported to the Altix system allowing a working version of a common FFT library.

4.3.3 Ongoing porting work

Larger codes which require more work, or codes for which the porting effort is ongoing:

Unified Model
The unified model versions 4.5 and 5.3 are currently being ported. The size of this code has thrown up a number of issues which may occur elsewhere during the porting process.

Version 5.3 is now ported but shows problems when the compiler options which maintain the highest possible precision are removed. This is currently under investigation.

VASP
SGI have ported VASP and we are now working with SGI to install on Reynolds.

pchan
The code is being ported to the Origin and Altix systems. A version for the Origin systems now exists, but the version for the Altix is awaiting corrections in a library supplied fft routine.

pdns3d
pdns3d has been ported. It runs about three times faster than on Green for the test cases provided.

Polcoms/Shelf
The Polcoms software will be ported as soon as the latest version is received.

4.3.4 Other Software

ScaLAPACK
Porting of the ScaLAPACK parallel numerical linear algebra library is progressing, but problems with the BLACS communication library have hindered progress.

4.4 Other Software Development

Further testing of the use of checkpointing on the Origins is continuing.

The provision of advance reservation and co-allocation facilities in the current Origin batch scheduling system are also being pursued.

5 Collaboration and Conferences

5.1 Events

The following events were attended by SVE staff:

- CARVI Conference, 3-4 April, Vitoria, Spain
  http://w3.euve.org/euve/Carvi/index.htm
  Leng J., Margetts L., John N.W., Leaver G.W., Lever P.G. "Virtual Environments – Improving the Understanding of Engineering Models"

- COMPLAS Conference, 6-10 April, Barcelona, Spain
  http://congress.cimne.upc.es/complas/frontal/default.asp
  Smith, I.M. and Margetts, L. "Portable Parallel Processing for Nonlinear Problems"

- ACME Conference, 24-25 April, University of Strathclyde, UK
  http://mechweb.mecheng.strath.ac.uk/acme2003/index.htm
  Margetts, L. 'Parallel Programming for General Parallel Finite Element Problems'

- SGI Developers Conference, 12-14 May, California, USA

- Cray User Group Summit, 12-16 May, Columbus, Ohio, USA
  http://www.cugoffice.org
  Gibson J and Pettipher, M 'Finite Element Analysis on the Cray MTA-2'
  Pettipher, M. "Managing Supercomputing Resources at the University of Manchester"
  Tate, A. "Optimisation of ScaLAPACK on the Cray X1"

- SGI User Group, 11-13 June, California, USA

- Reality Grid Workshop, 17-18 June, London, UK
  http://www.realitygrid.org/workshop.html

6 Added Value Services

6.1 International Conferences

Fiona Cook organised the exhibition of SVE work at ISC2003 in Heidelberg, Germany, (24-27 June). Fiona attended the exhibition along with Mike Pettipher, Terry Hewitt, James Perrin who gave a demo of the Parallel Support Toolkit and Kaukab Jaffri who gave a demo on resource brokering on the CSAR machines.

Preparations started for the exhibition at SC2003 (Phoenix, USA, November) are progressing well. A proposal for SC Global 2003 on "Collaborative Virtual Design in Engineering", submitted by Lee Margetts, Kevin Roy and Rupert Ford from the University of Manchester, and Simon Bee from the University of Salford, has been accepted.

6.2 Seminars

6.2.1 MRCCS/ESNW Seminar Series

The following seminars, jointly organised by MRCCS and ESNW, have been held during the period:

- The Permis Privilege Management Infrastructure, *David Chadwick*, *University of Salford*
- The Unified Model and Supercomputing at the Met Office, *Paul Selwood, The Met Office*

6.3 Summer School 2003

Planning for the Summer School on HPC in Finite Element Analysis, jointly organised by MRCCS and the National Science Foundation (NSF) of the USA is well underway. Speakers will be from France, Germany and Japan (Earth Simulator) as well as the USA and UK.

6.4 Visualization

The MVC team has completed the first version of the AVS/Express multipipe edition for PC Graphics clusters. Built on top of VRCOs CAVElib API, an AVS Express network can now be run in an immersive environment using a cost effective PC platform. This software will be available to purchase from AVS UK in due course. Meanwhile, please contact MVC for more information or to arrange a demonstration.