CSAR Service

Consolidated Management Report

1st Quarter 2000

Management Summary

The system is the flagship HPC facility for UK Academia and Industry, enabling them to continue their World-Class research and development.

The number of users has grown to a total of 427 to date.

The Turing machine has now been upgraded by 200 PEs making it the single largest machine in the UK.

The allocation of the Service Quality Tokens over the quarter has provided a useful measure of user satisfaction with the current position showing no black marks outstanding.

The system continues has been highly utilised this quarter.

The STK silo underwent an upgrade this quarter with an additional 4 tape drives being added subject to final approval of the CCN by the Research Councils.

Anouncements for the 2 week summer school on Linux supercomputing and the 6th Cray/SGI MPP workshop, both being held in Manchester in September 2000, were made.

CfS is currently 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 was fair (see Table 3), i.e. Yellow measured against the CPARS performance targets.

CSAR Service - Service Quality Report - Actual Performance Achievement

Service Quality Measure	April	May	June	July	Aug.	Sept	Oct	Nov	Dec	Jan	2000 Feb	March
HPC Services Availability												
Availability in Core Time (% of time)	97.10%	98.50%	99.70%	99.70%	100%	100%	100%	100%	100%	96.11%	95.0%	99.70%
Availability out of Core Time (% of time)	98.10%	99.71%	99.40%	99.40%	99.40%	99.5%	100%	100%	99.70%	98.52%	100%	99.5%
Number of Failures in month	1	3	2	2	1	1	0	0	1	4	1	2
Mean Time between failures in 52 week rolling period (hours)	480	453	395	391	416	437	486	534	563	230	515	486
Fujitsu Service Availability												
Availability in Core Time (% of time)	N/A	N/A	N/A	N/A	N/A	N/A	N/A	99.30%	100%	100%	100%	100%
Availability out of Core Time (% of time)	N/A	N/A	N/A	N/A	N/A	N/A	N/A	100%	100%	100%	100%	100%
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	<3	<3	<2	<2	<1	<3	<2	<1	<1	<1	<2
Administrative Queries - Max Time to resolve 95% of all queries	<2	<2	<1	<1	<1	<1	<2	<1	<0.5	<0.5	<1	<2
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	0	0	0	0	0	0	0	0
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 month	0	1	2	2	2	1	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 and Fermat:[Turing availability x 122 / (122 + 3.5)] + [Fermat availability x 3.5 / (122 + 3.5)]

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	April	May	June	July	Aug.	Sept	Oct	Nov	Dec	Jan	2000 Feb	March
HPC Services Availability												
Availability in Core Time (% of time)	0.078	0.039	-0.039	-0.039	-0.058	-0.058	-0.058	-0.058	-0.058	0.195	0.195	-0.039
Availability out of Core Time (% of time)	0.039	-0.039	0	0	0	-0.039	-0.047	-0.047	-0.039	0.000	-0.047	-0.039
Number of Failures in month	-0.008	0	0	0	-0.008	-0.008	-0.009	-0.009	-0.008	0.008	-0.008	0
Mean Time between failures in 52 week rolling period (hours)	0	0	0	0	0	0	0	-0.008	-0.008	0.008	-0.008	0
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	0.016	0	0	-0.016	0.016	0	-0.016	-0.016	-0.016	0
Administrative Queries - Max Time to resolve 95% of all queries	0	0	-0.016	-0.016	-0.016	-0.016	0	-0.016	-0.019	-0.019	-0.016	0
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	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
New User Registration Time (working days)	0	0	0	0	0	0	-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 month	-0.004	-0.003	0	0	0	-0.003	0	0	0	0	0	0
Monthly Total & overall Service Quality Rating for each period:	0.03	-0.01	-0.03	-0.04	-0.05	-0.08	-0.07	-0.09	-0.09	0.07	0.03	-0.06
Quarterly Service Credits:			0.00			-0.17			-0.25			0.04

Table 3

1.2 Service Quality Tokens

The current position at the end of the quarter is that of the 427 registered users of the CSAR Service none have any outstanding Service Quality Tokens as per the guidelines of the User Steering Group. See below:

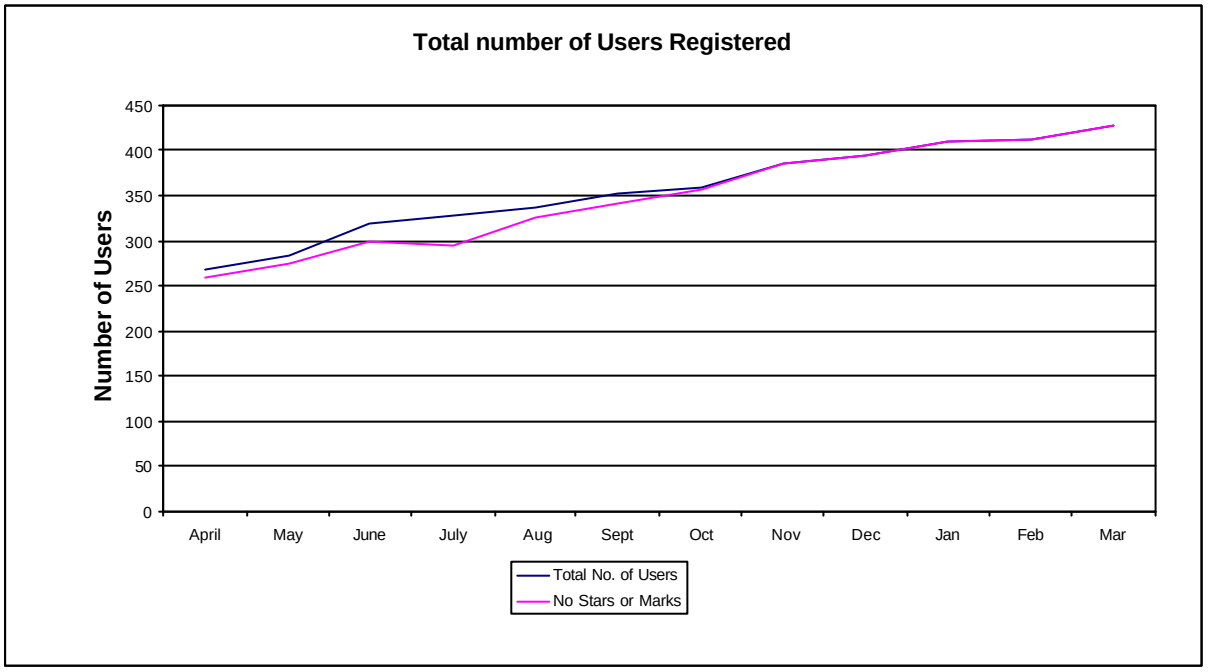

The area graph below illustrates the monthly usage trend of Service Quality Tokens:

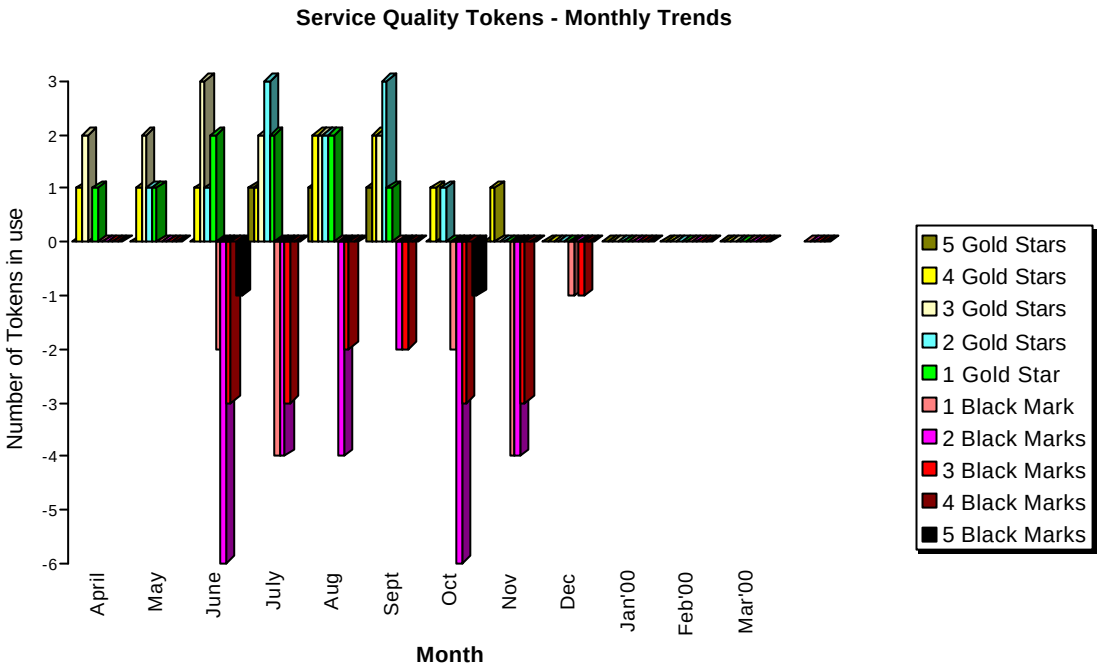

Over the course of the quarter the position is that as a management tool the Service Quality Tokens have enabled the user to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

There are, at the end of the quarter, no outstanding stars or marks, nor have there been either stars or marks placed against the service at any point during the quarter.

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 mo nth of CPU (T3E PE Elapsed Hours), split by Research Council and giving the equivalent GFLOP-Years as per NPB. The Baseline Capacity (103 GFLOP-Years) 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 (3.5 GFLOP-Years) 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 current monthly CPU, PE, disk and HSM allocations and usage.

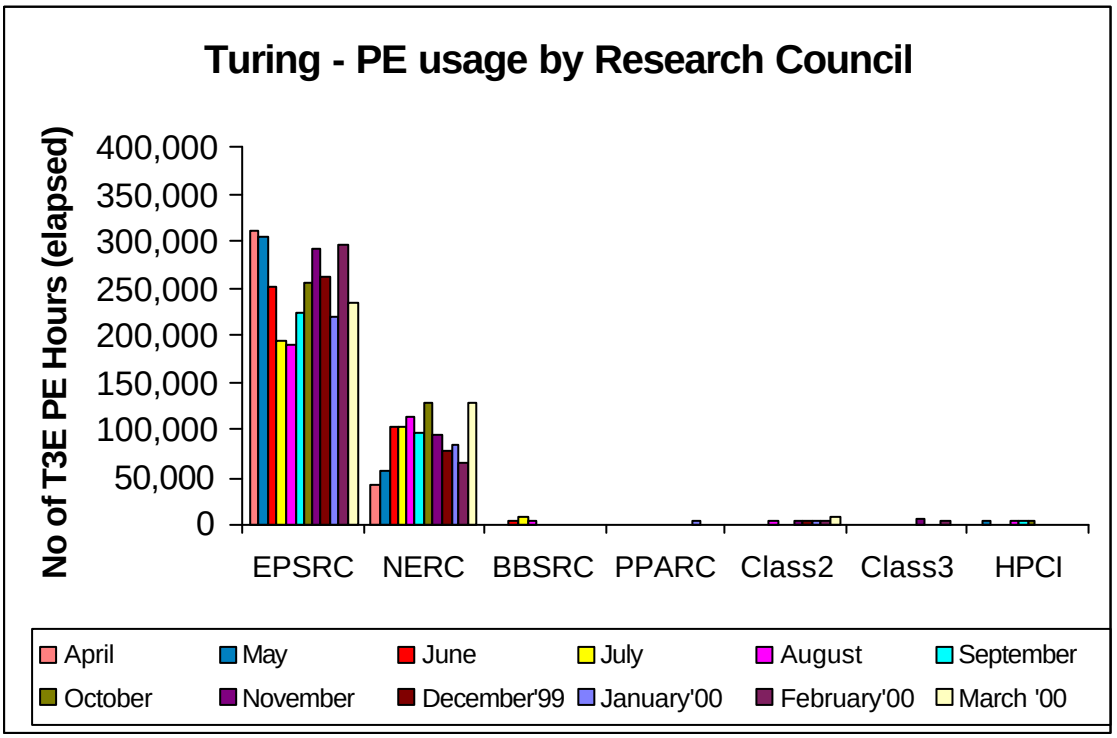

Turing PE usage is shown by Research Council during the months of service to date in the above chart.

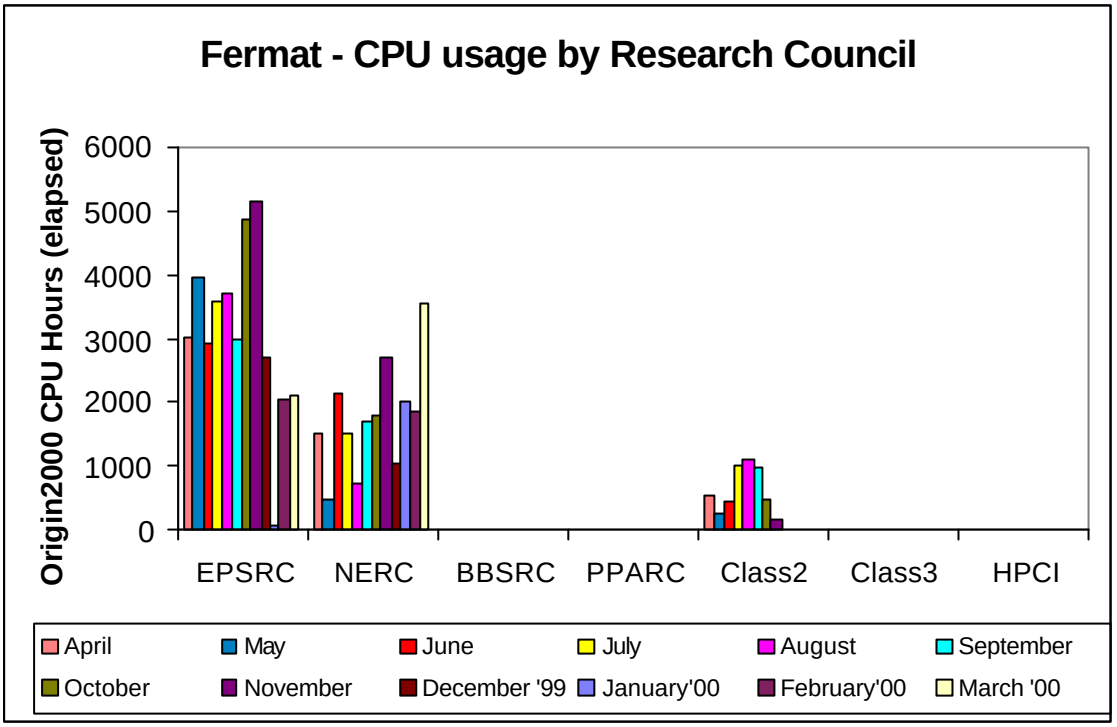

Origin2000 CPU usage is shown by Research Council during the months of service to date in the above chart.

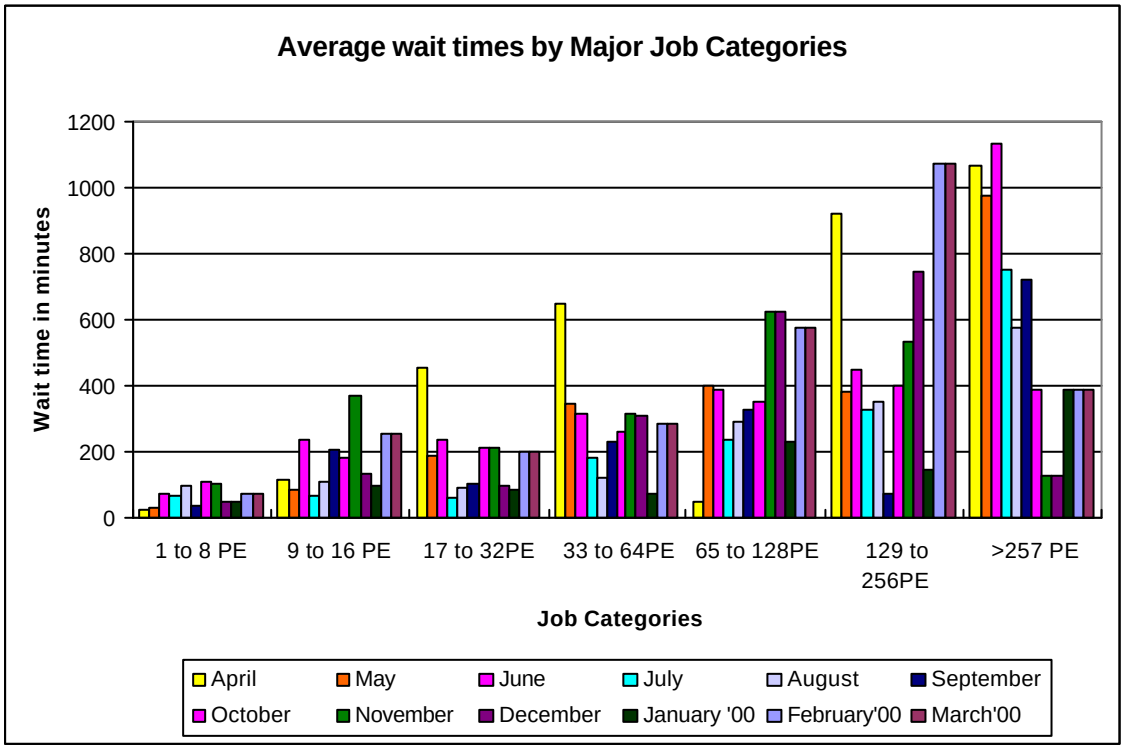

The chart below shows the average wait time trend over the months from January to date.

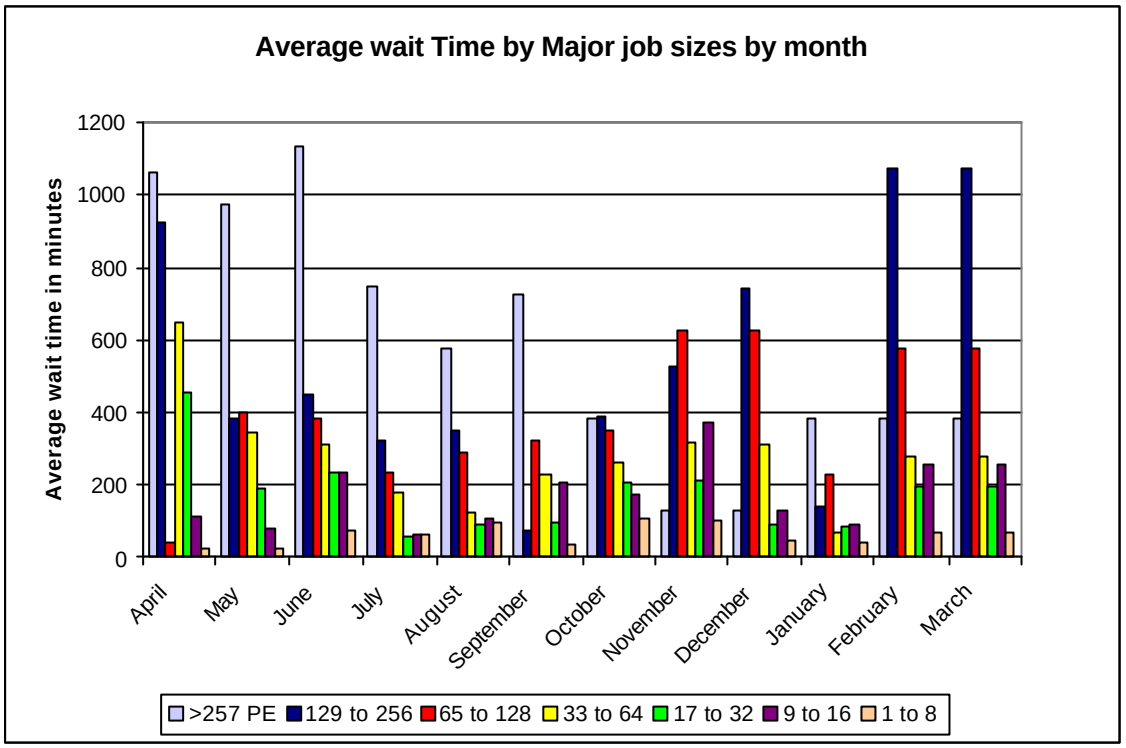

It can be seen from the above graph that enhancements to the scheduling on Turing have reduced the average wait times however attention must be paid to ensure sufficient head room exists in the system to prevent wait times from rising

The next chart shows the percentage PE time utilisation by the major job categories on the Turing system for the final quarter in 2000.

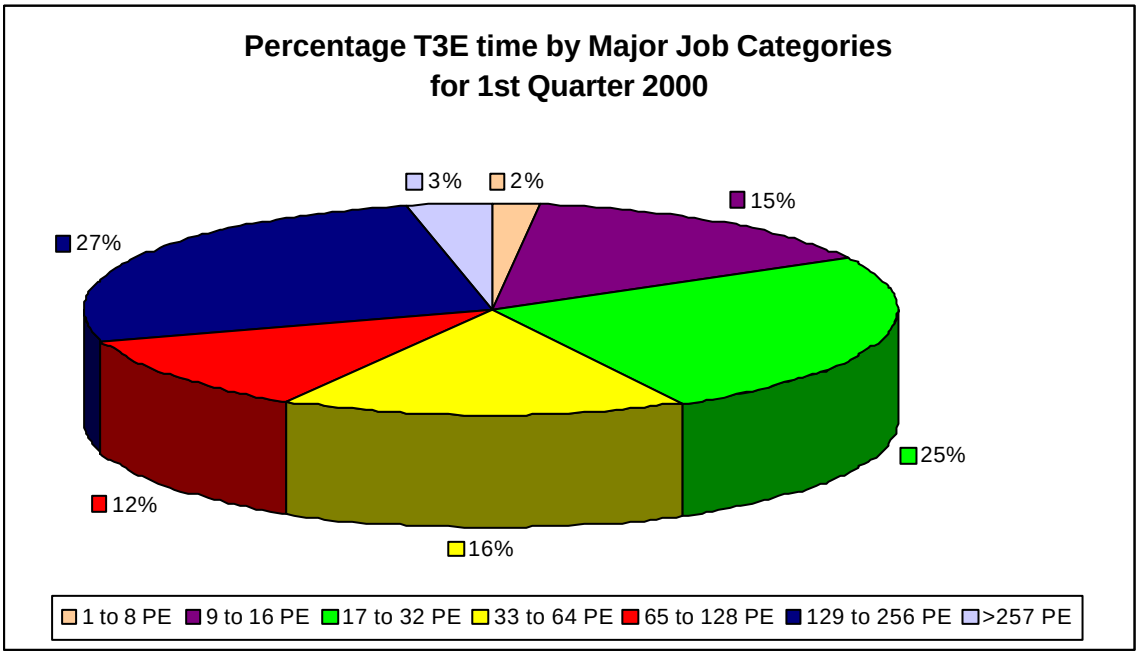

2.2 Baseline System Usage Graphs

In all the Usage Charts, the baseline varies dependant on the number of days in each month, within a 365-day year. The reduced Baseline in November 1998 represents half a month.

The graph below shows the PE hour's utilisation on Turing by Research Council from November 1998.

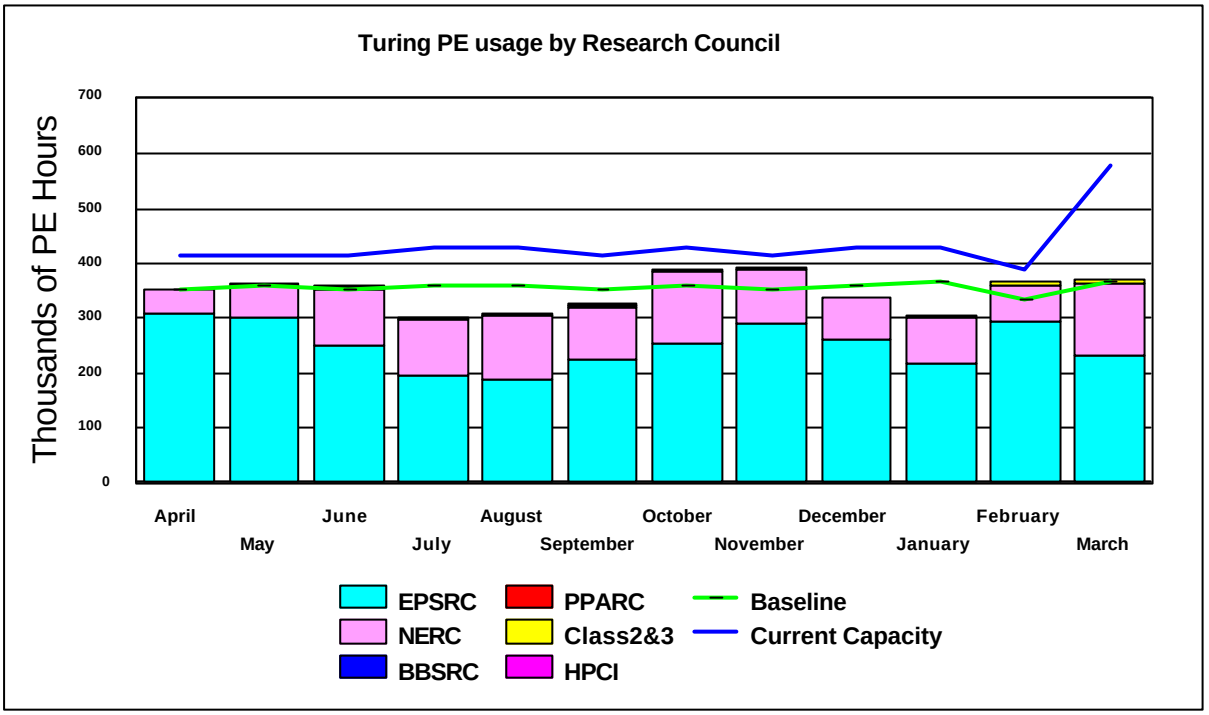

The graph below shows the historic CPU usage on Fermat by Research Council from November.

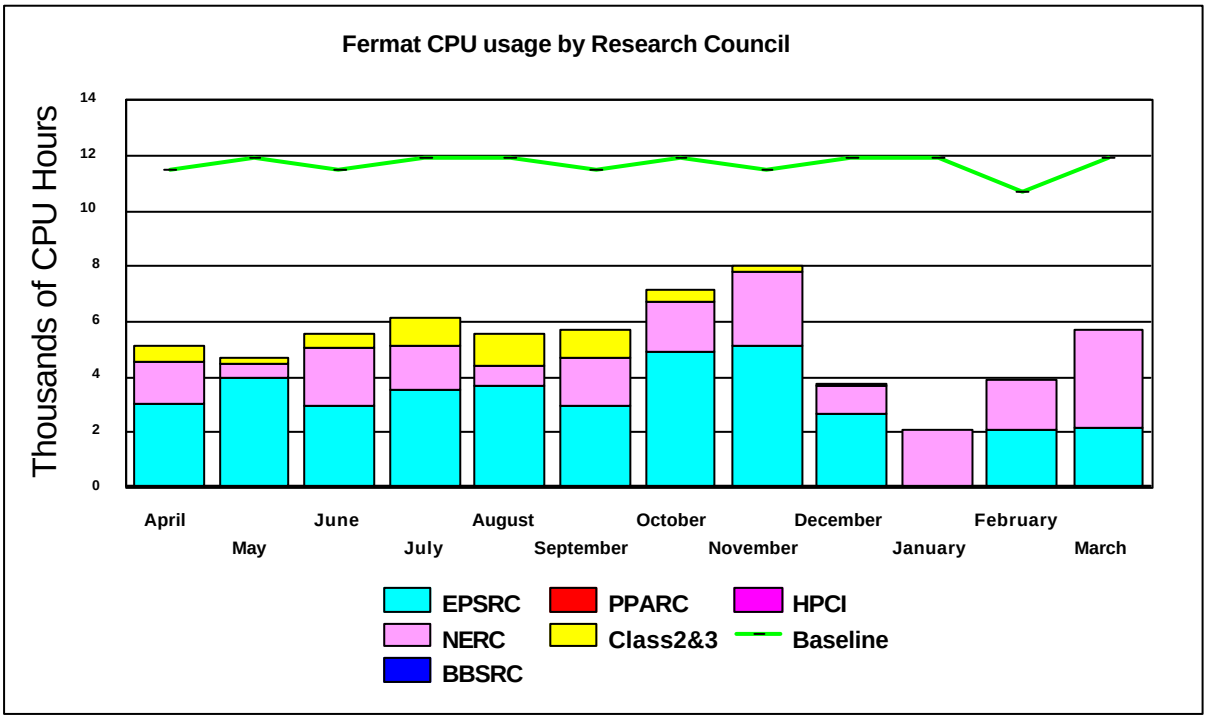

The next series of graphs illustrates the usage of the disk and HSM resources of the system.

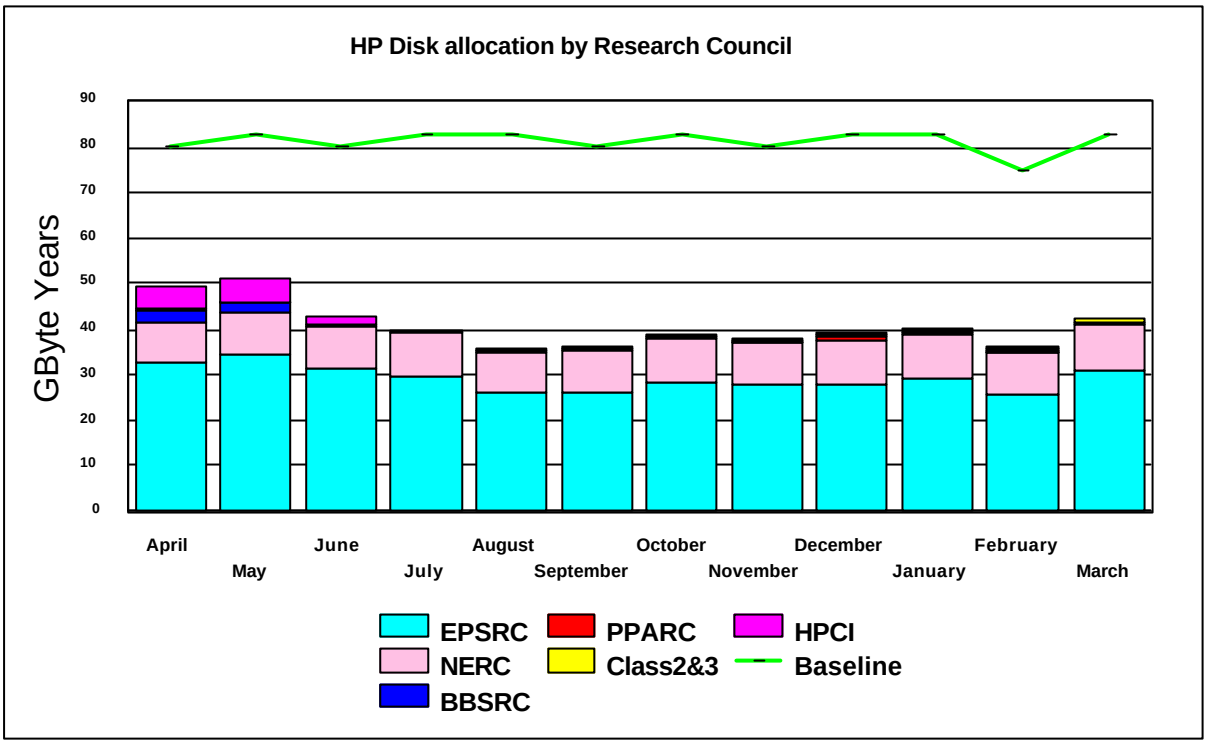

The preceding graph illustrates the historic allocation of the High Performance Disk on Turing.

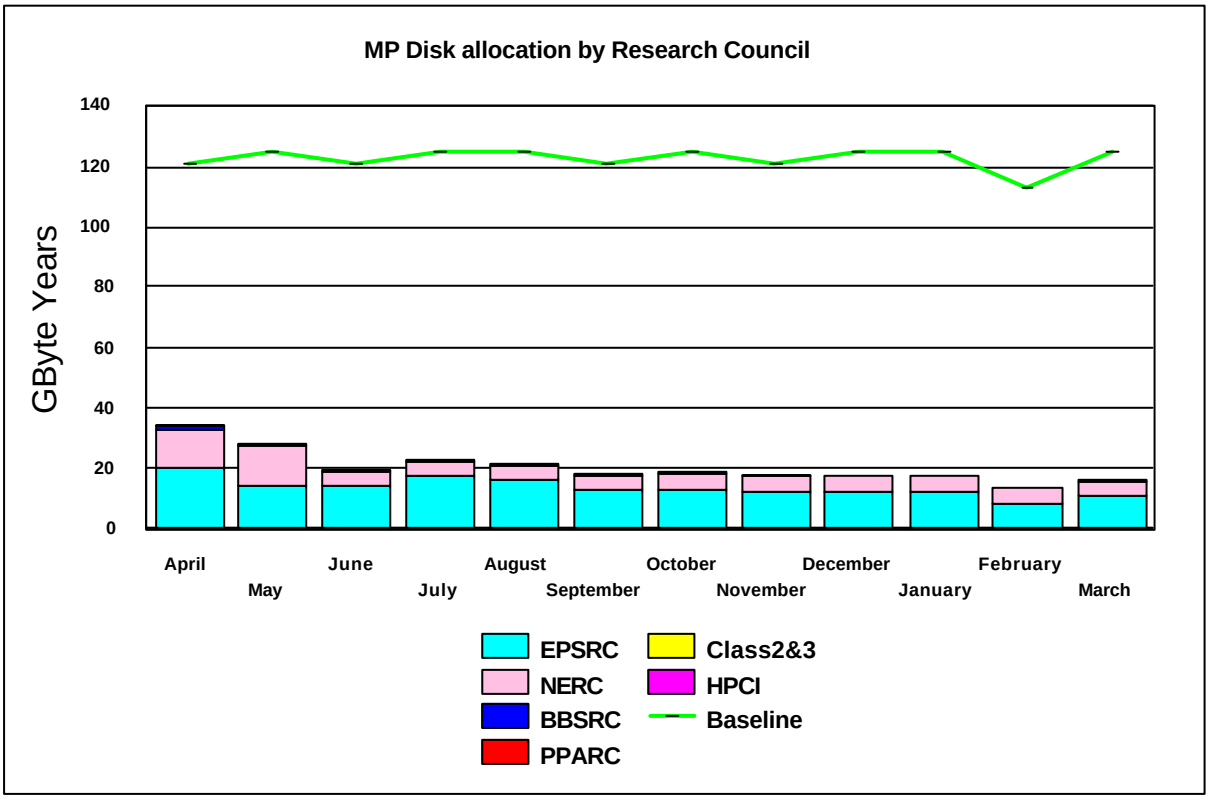

The graph above illustrates the historic allocation of the Medium Performance Disk on Fermat.

The graph below shows the historic HSM usage by Research Council funded projects. The primary usage is for NERC.

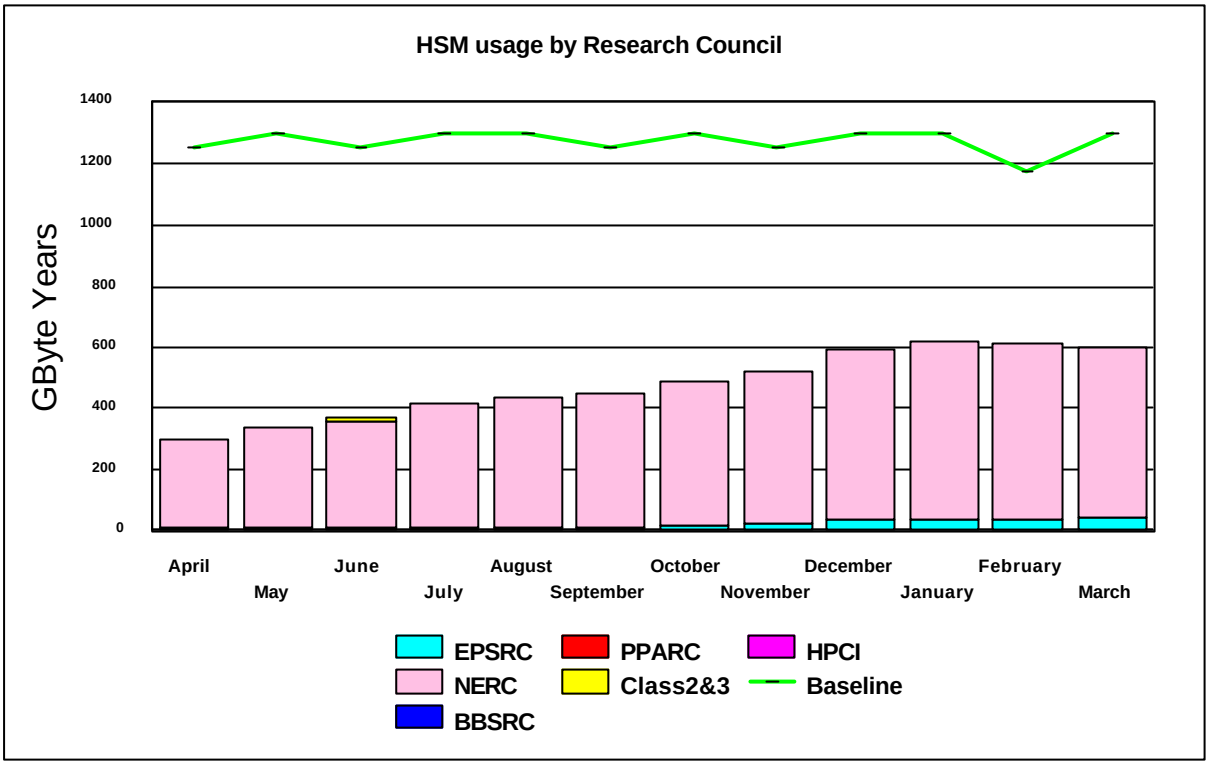

2.2.1 Fujitsu System Usage Graph

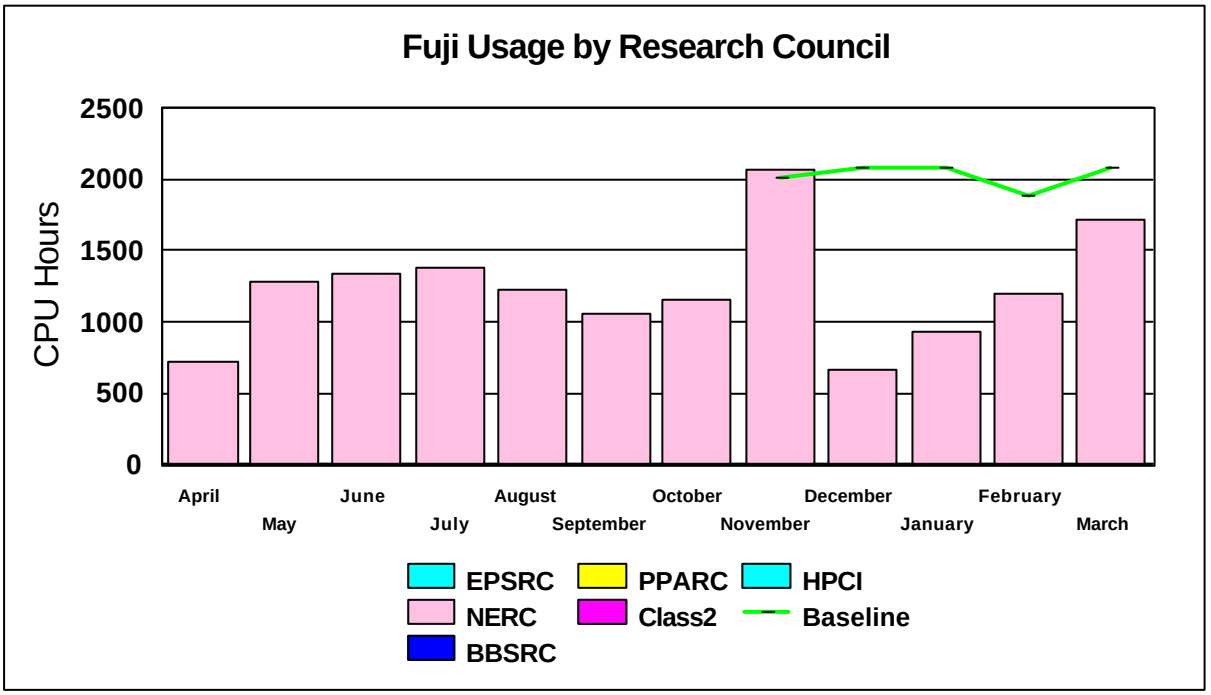

The above graph shows the current CPU usage on the Fujitsu VPP 300 NERC system based at the University of Manchester.

The system has been fully integrated into the CSAR service from November, and was 100% available both in core and out of core time during December.

2.2.1 Guest System Usage Graphs

The CSC J90 Guest system has been withdrawn from service in agreement with EPSRC and NERC.

2.3 Service Status, Issues and Plans

Status

The service has run above baseline in two of the three months in the quarter.

Issues

Turing was successfully upgraded by 200 PEs in this quarter.
The system is now capable of running jobs upto 700 PEs in size.

Plans

The STK silo was upgraded in March by the addition of four 9840 tape drives. A larger hold area is to be created in accordance with the relevant CCN which will bring about enhancements to the HSM system.

3 Science Application Support Services

3.1 Consortia Support

Work is proceeding with the parallelisation/optimisation of CASTEP on behalf of the UKCP consortium. The academic release of 4.2 is on the point of being released on CD to relevant CSAR users. The visualization work for the Terra group has made substantial progress and the Terra group have commented very favourably on the results. This work is being disseminated to the UK community via a case-study report for UKHEC and forms the centrepiece of an article on visualization in the UKHEC newsletter.

Work on coupling ocean-atmosphere models has been made available to the Occam group who are evaluating its potential for production use. International links with other groups running coupled models have been made via a visit and presentation to CERFACS and this expertise is available to CSAR users.

The offer of a free evaluation of applications by the Applications/Optimisation team has been taken up by several groups and relevant results from this excercise form the basis of articles to the UKHEC newsletter and will feature in the next edition of CSAR focus.

Work on parallelising a code of major importance for consortium CSE030, "HPC for Complex Fluids" has recently been handed back to the user for evaluation. This work involved parallelising a code completely from scratch and was funded by support tokens. The exercise has been very valuable in developing the CSAR team's ability to negotiate large-scale support via the tokens scheme and has shown that this scheme can provide a way in which users can evaluate the progress of support effort and have a set of objectives by which this progress can be judged.

The experience gained from support work is being used to advise the Research Councils on the support needs of new project applications. It is clear that for the majority of projects, sufficent support tokens need to be built into the initial application and that where this has been done, projects have reaped considerable benefit.

3.2 Collaboration with ‘HPCI Centres’

MRCCS is one of the partners in the UKHEC consortium which provides support and information of current developments in high-end computing. This has enabled CSAR to continue its fruitful collaboration with the former HPCI centres. One very useful result has been the exchange of information/expertise with Daresbury over matters of cluster computing and the newly emerging GRID technology (see also Section 6).

Speakers from both EPCC and Daresbury gave seminars in the MRCCS series and this allowed valuable informal contact between the major HPC centres to maximise expertise and avoid wasteful duplication of effort.

3.3 Conferences and Meetings

The following conferences have been attended

* "Second International Workshop on Next Generation Climate Models for Advanced High Performance Facilities", February 23-25, Meteopole Toulouse. A CSAR paper was given on "Using Metacomputing to Process Scientific Data", reporting on the SC99 experiments.

* A presentation on current High End Computing in the UK was given to meetings in Bologna and Sweden.

* A presentation of the CSAR service was given as a seminar to CERFACS. At this visit there were major discussions on coupling climate models, with particular relevance to OASIS which was developed at CERFACS. The CSAR expertise in intercontinental metacomputing was shared with relevant groups at CERFACS and the intention is to continue and develop this collaboration.

* "AAAI 2000 Spring Symposium Series, Smart Graphics", March 20-22 Stanford, CA CSAR paper presentedon "A case study of 3D Visualization for Computational Data in a Spherical System"

* "Isthmus 2000, Eunis 2000, EGrid Forum", Poznan (Poland), 10-13 April

3.4 Training and Education

The CSAR service has continued to make a wide range of courses available. The courses have proved useful and relevant (as assessed by the users on the course questionnaires). We are still concerned about attendance not being as high as we would like and some interesting courses had to be cancelled through lack of sufficient numbers. We have asked CSAR users informally about this and one problem seems to be that of getting funding for accomodation to visit Manchester for the duration of the course. Provision may be needed for this in research grants along with the usual provision for subsistence for conference attendance. We have also offered to give courses at users sites and the course on Fujitsu Programming was given at Lancaster to benefit several new Fujitsu users.

We have completed a contract to provide training and support for the Met Office by adapting our materials on optimisation and debugging for users of the Unified Model. These materials have been passed to the Met Office but we have the right to use this material to train users of the Unified Model on the CSAR service. This offer will be made to the UGAMP consortium but other CSAR users who may be interested are encouraged to contact CSAR.

3.5 Service Developments

Work has proceeded on provision of the secure shell on all CSAR facilities and this service should be available shortly. It will be announced via the usual channels.

CSAR applications have worked with systems managers at CSC to investigate and improve the performance of applications running on command PEs. This work has resulted in recommendations which are curretnly being tested prior to implementation and announcement.

5 User Registration and Project Management Services

This section covers aspects relating to the registration of projects and users, and the management of projects and resources.

5.1 Quarterly usage reports by email

A new feature is the production of quarterly emails providing a summary of actual and planned usage in the previous quarter. The first of these emails went out to project administrators prior to Easter. It is intended to provide a brief summary to assist in the monitoring of usage, and to serve as a reminder to keep individual capacity plans up to date. It is too early to report on any feedback, or any improvement in the accuracy of individual capacity plans.

In the last report it was stated that an email would be sent out to advise projects when they are able to retrade their resources (three months after the previous retrade) and to remind them to review their capacity plans. However, in reality most projects retrade only when necessary, for example when they have exhausted the bulk of their resources, and so this proposal was deemed ineffective in practice. The quarterly email serves a similar purpose as well as providing more information.

5.2 Conditions of use

The conditions of use of the CSAR systems are to be revised imminently in order to comply with current UK Department of Trade and Industry regulations on access to HPC systems. In essence these regulations prohibit the use of high performance computers for any research, development, manufacture and procurement of weapons

of mass destruction and their delivery systems, in support of internal repression or international aggression, or any other restricted usage as may be listed from time to time. New users will need to agree to these conditions before they are permitted to register, and all existing users will be asked to confirm that they agree to the conditions. It is expected that in general users will not have difficulty in complying with these conditions; any problems will be dealt with on a case-by-case basis.

6 Value-Added Services

6.1 Joint projects

The development of the Multi-pipe utility for AVS and SGI multi-pipe graphics has met all its deadlines with a first release to a company in December 2000. Discussions are under way about distribution of the software, and about a phase II development.

6.2 Research Liaison

The work with Stuttgart and Pittsburgh continues and is reported elsewhere.
The collaborative project with DL and EPCC submitted to EPSRC for Core HPC support is reported elsewhere.

6.4 New Applications Initiative

The multi-pipe capable module for AVS/Express has now been made into a commercial product by AVS Inc, which will be demonstrated to Research Council users at the Visualisation and Virtual Environments Computational Club at Salford University on Friday, 14th January. This exciting product allows users to rapidly develop post simulation visualisations on a workstation or low powered machine, and then migrate them into a fully immersive environment such as a Reality Centre or CAVE by the substitution of a single AVS module.

CfS have delivered computational and development resources to a research project intended to couple major supercomputing centres together. The scientific problems tackled included a 2,000,000 particle molecular dynamics simulation, a flow simulation of a spacecraft re-entering the atmosphere, and a large-scale analysis of radio astronomy data.

6.3 Databases

We have not been able to discuss with EPSRC how best to populate the database, but we now have some material to kick-start it anyway. We will discuss with EPSRC how to promote its use.

6.4 Adding value to other major research facilities

The National Centre for Virtual Environments at Salford University have started a research project into computational steering, coupling their Virtual Reality environments with simulations running on CSAR services. CfS have been active in our assistance, offering networking and applications design consultancy, training, and software product support. The N.C.V.E. are planning computational steering applications in the areas of medical imaging and mechanical engineering (collision detection).
As mentioned in section 6.4, CfS have undertaken computational grid applications development. One anticipated benefit is to allow the Jodrell Bank radio telescopes to reconstruct signals in near real time (i.e. while the observation is taking place), and so make better use of the observation time.

6.5 ASCI prototyping capability

Recent developments associated with the ASCI project indicate that the original outline plans for an ASCI prototyping capability are no longer adequate. CfS has been working to define a more concrete set of proposals

associated with the capability, in order to maximise its usefulness to UK academia. These plans were presented to the CSAR User Steering Group on 15th December 2000.

CfS is now intending to take these plans forward by presenting a Contract Change Note (CCN) within the next month.

MRCCS

MRCCS Seminars
The first series of seminars organised by MRCCS on topics in High Performance Computing and Visualisation were held on Friday afternoons in the Department of Computer Science at the University of Manchester. A wide range of topics was covered. Speakers came from around the UK. Attendance was predominantly from the North West of England including people from Daresbury Laboratory, and the Universities of Salford and Leeds. The series for the second semester is currently being organised.

Presentations
W T Hewitt presented a lecture entitled "HPC in the UK" at the ARCADE meeting in Bologna.

Summer School 2000
The second MRCCS Summer School will take place in early September and will be on the subject 'Supercomputing with Linux'.

EU Projects
We are partners in two projects that are being submitted by 17th January closing date under Framework V.

International Conferences
Since the last report we have secured Euro-par 2001 in Manchester, 28th – 31st August 2001. Professor John Gurd will be the programme chair.

A meeting of the EG2001 national organising meeting has taken place to establish roles etc. Mr W T Hewitt, Dr N John and Dr I Herman (CWI, NL) are the conference chairs; Dr A Chalmers (Bristol) and Ms Theresa-Marie Rhyne (EPA, USA) are the international programme chairs. The conference themes will include computational steering/visual supercomputing.

A similar meeting for Euro-Par 2001 will be held late January/early February.

Summer Students
Planning for the 2000 summer student programme has started.