CSAR Service

Consolidated Management Report

2nd Quarter 2001

Management Summary

The system continues to be the flagship HPC facility for UK Academia and Industry, enabling World-Class research and development.

The number of users has grown to a total of 492 to date.

The Cray T3E (Turing) system continued to run almost to full capacity again this quarter. Queue wait times are reducing rapidly with the introduction of the new Fermat and Green resources.

The upgraded Fermat continues to be heavily used.

The Origin 3000 (Green) is currently at 256 PEs and is scheduled to grow to 512 by August 2001, and usage is growing steadily.

The T3E has seen 65% of the Quarters workload at greater than 64 PEs in size.

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 was the year 2000 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	2000/1 July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March	April	May	June
HPC Services Availability												
Availability in Core Time (% of time)	100%	100%	100%	100%	100%	94.90%	99.70%	99.70%	100%	100%	99.70%	99.70%
Availability out of Core Time (% of time)	100%	100%	100%	100%	99.40	98.49%	99.50%	99.40	99.40	99.40	99.40	99.40
Number of Failures in month	0	0	0	0	2	4	1	1	1	1	3	3
Mean Time between failures in 52 week rolling period (hours)	461	626	730	1095	673	584	584	626	674	674	584	584
Fujitsu Service Availability												
Availability in Core Time (% of time)	100%	98.4%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Availability out of Core Time (% of time)	100%	100%	100%	100%	100%	100%	100%	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	<2	<2	<2	<1	<3	<3	<5	<5	<3	<5	<2	<3
Administrative Queries - Max Time to resolve 95% of all queries	<0.5	<2	<2	<0.5	<0.5	<5	<2	<2	<3	<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	0	<0.5	0	<0.5	<0.5	<0.5	<0.5	0	0	<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	12	10	10
System Maintenance - no. of sessions taken per system in the mont	2	2	2	1	2	1	0	2	1	2	0	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 at installation:[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	2000/1 July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March	April	May	June
HPC Services Availability												
Availability in Core Time (% of time)	-0.058	-0.058	-0.058	-0.058	-0.058	0.195	-0.039	-0.039	-0.058	-0.058	-0.039	-0.039
Availability out of Core Time (% of time)	-0.047	-0.047	-0.047	-0.047	0	0	-0.039	0.000	0	0	0	0
Number of Failures in month	-0.009	-0.009	-0.009	-0.009	0	0	-0.008	-0.008	-0.008	-0.008	0	0
Mean Time between failures in 52 week rolling period (hours)	0	-0.008	-0.008	-0.009	-0.008	-0.008	-0.008	-0.008	-0.008	-0.008	-0.008	-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	-0.016	0.016	0.016	0.031	0.031	0.016	0.031	0	0.016
Administrative Queries - Max Time to resolve 95% of all queries	-0.019	0	0	-0.019	-0.019	0.046	0	0	0.016	-0.019	-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	0	-0.002	0	-0.002	-0.002	-0.002	-0.002	0	0	-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	0.003	0	0
System Maintenance - no. of sessions taken per system in the mont	0	0	0	-0.003	0	-0.003	-0.004	0	-0.003	-0.003	-0.004	0
Monthly Total & overall Service Quality Rating for each period:	-0.09	-0.08	-0.08	-0.10	-0.06	0.11	-0.05	-0.03	-0.04	-0.05	-0.06	-0.05
Quarterly Service Credits:			-0.25			-0.05			-0.13			-0.15

Table 3

1.2 No. Of Registered Users

The current position at the end of the quarter is that there are 492 registered users of the CSAR Service.

It can be seen from the chart below that the number of users is relatively stable.

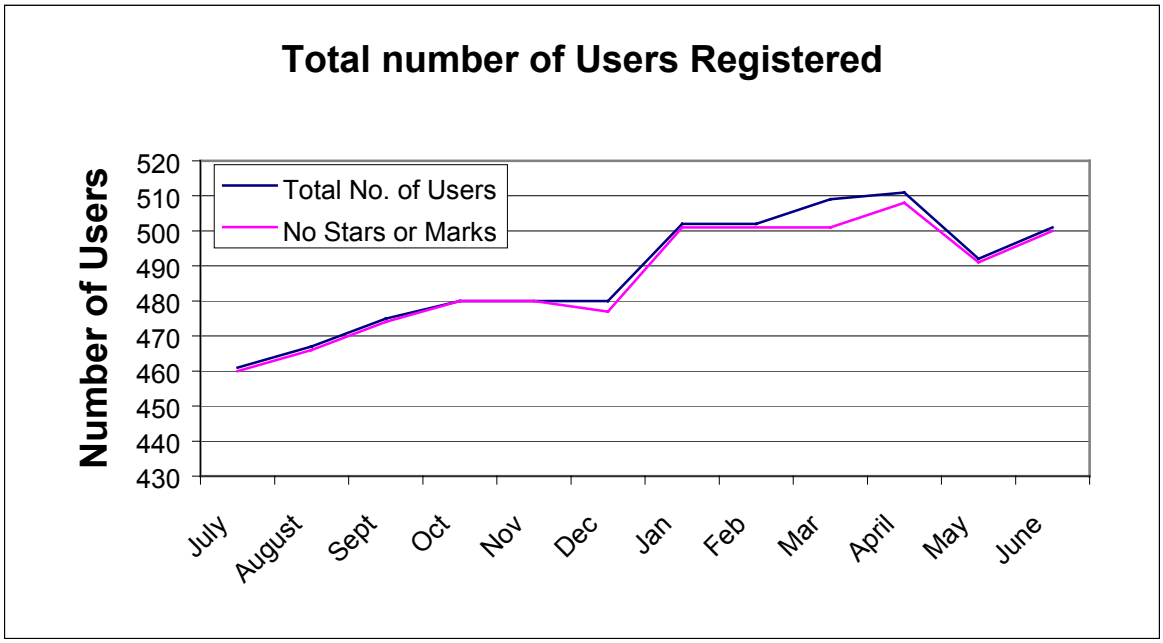

1.3 Service Quality Tokens

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

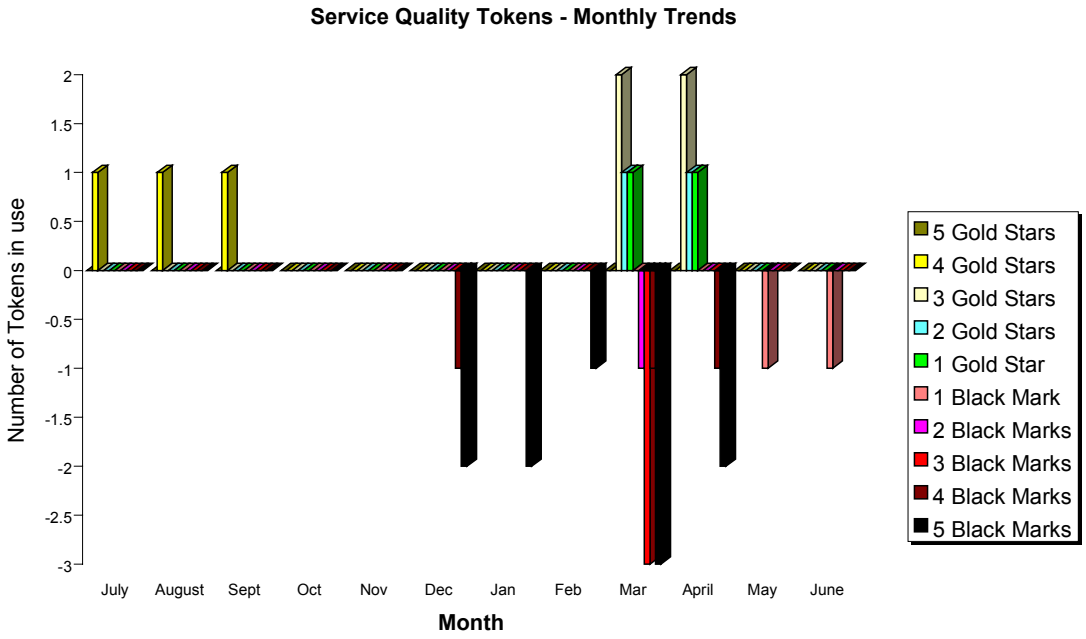

Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

There are, at the end of the quarter, one mark allocated to the service as per the chart below.

SUMMARY OF SERVICE QUALITY TOKEN USAGE

No of Stars or Marks	Consortia	Date Allocated	Reason Given
1 Black Mark	CSN001	12/03/01	Prblems improved though not resolved

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 CPU, PE, disk and HSM allocations and usage.

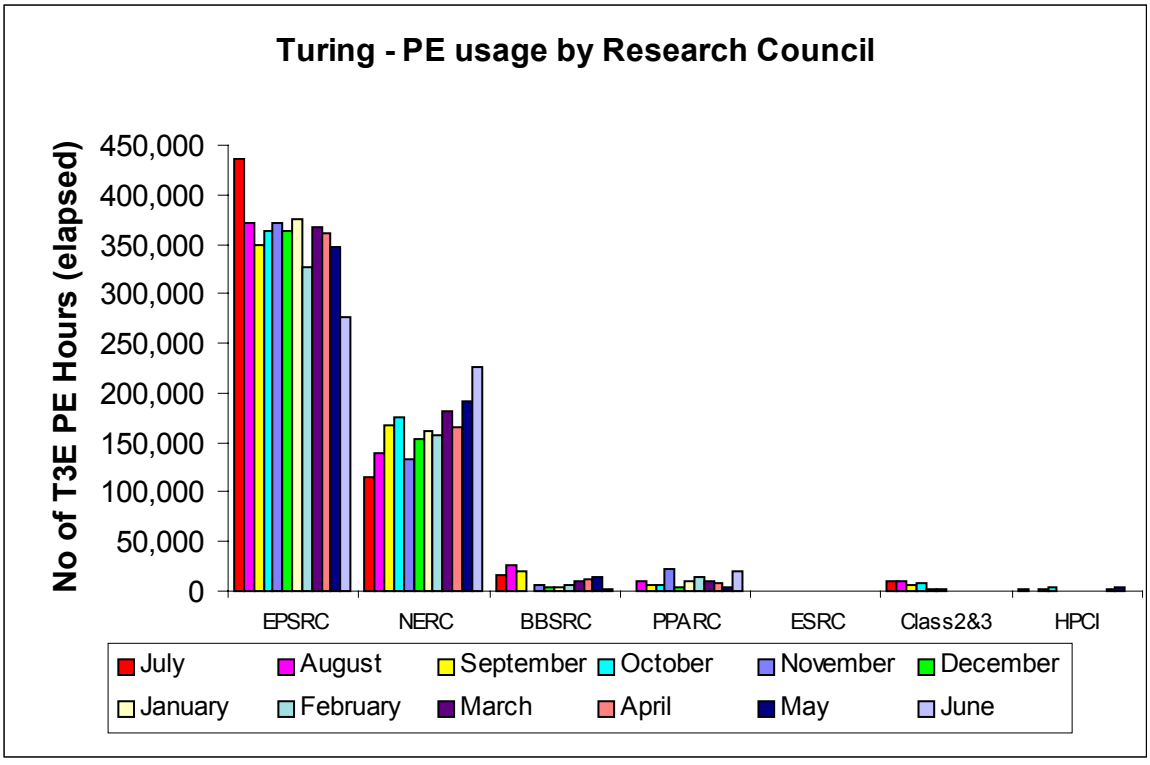

Turing PE usage is shown by Research Council during the last 12 months of service is shown in the above chart.

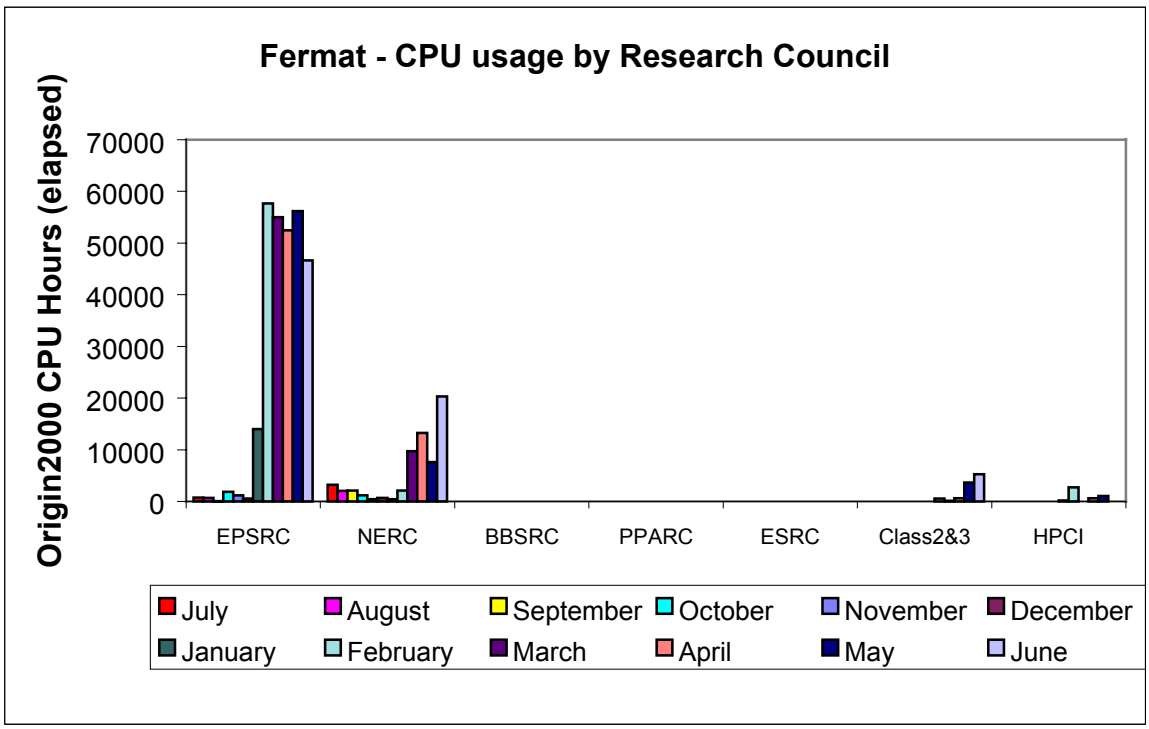

Origin 2000 CPU usage is shown by Research Council during the last 12 months of service is shown in the above chart. CPU usage of the Origin 3000 (Green) will be reported next Quarter.

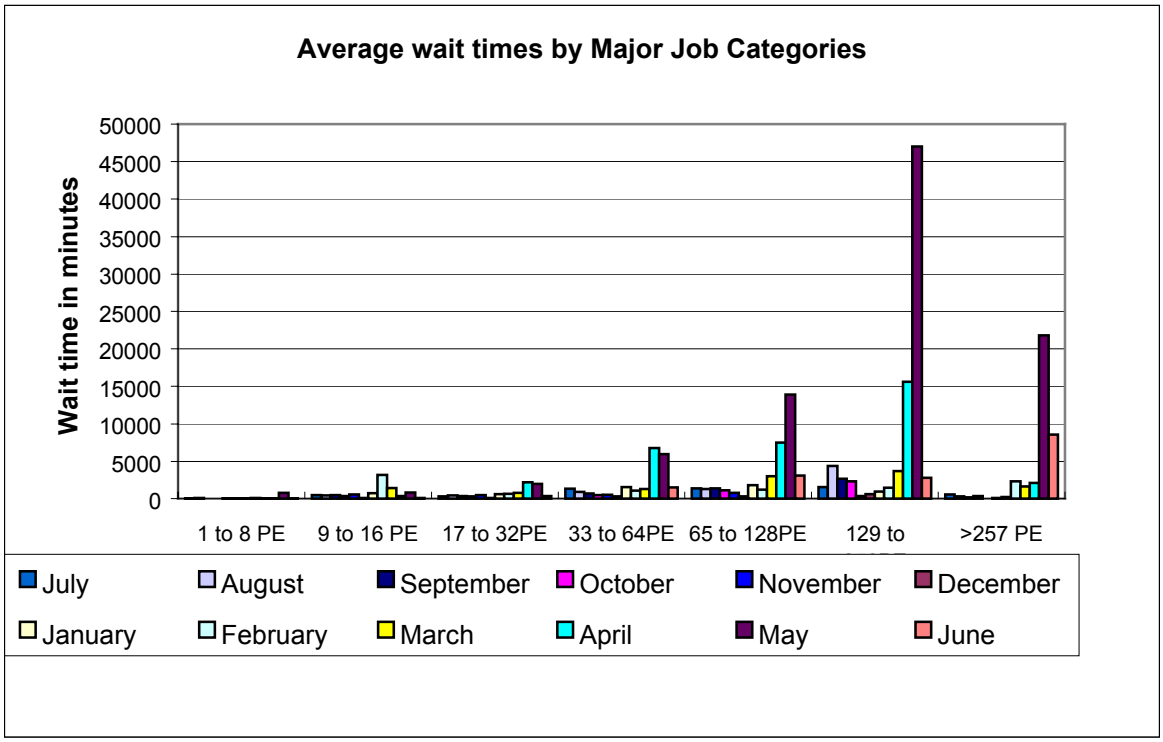

The above chart shows the wait time trend on the Turing system. The trend towards rising wait times is now being reversed for all job sizes.

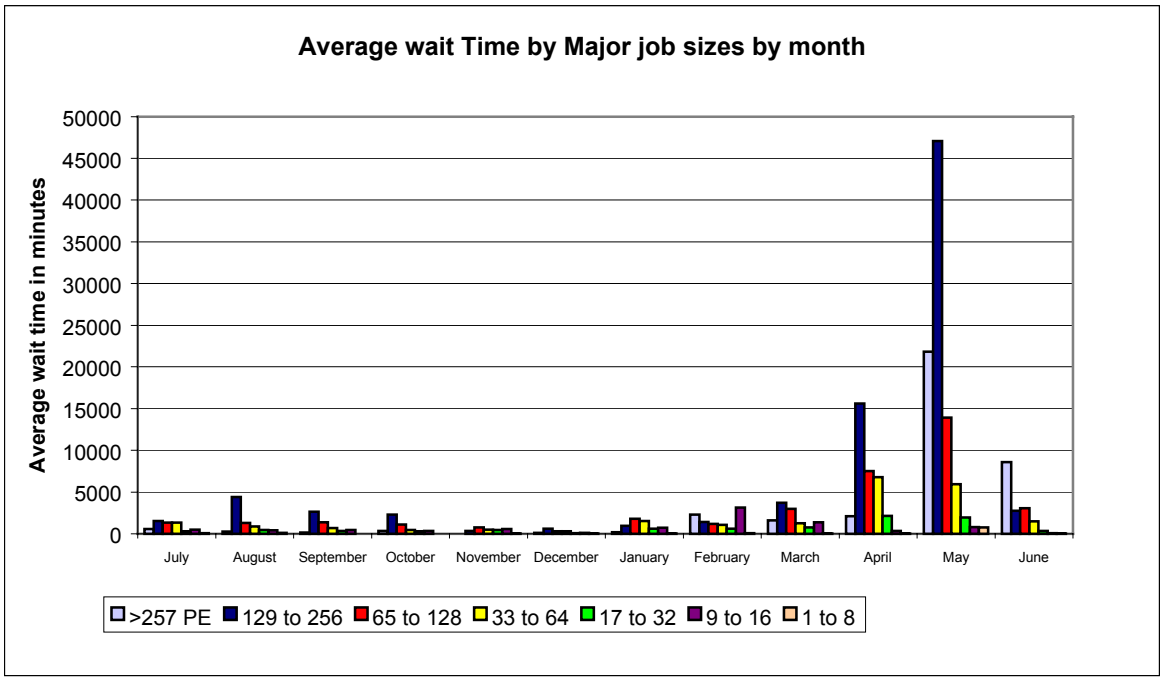

It can be seen from the above graph that the additional capacity added to the service, in the shape of the upgrade to Fermat and the addition of the Origin 3000 (Green), are now reducing the burden on the T3E (Turing) and reducing the queue wait times.

The next chart shows the percentage PE time utilisation by the major job categories on the Turing system for the 2nd quarter 2001.

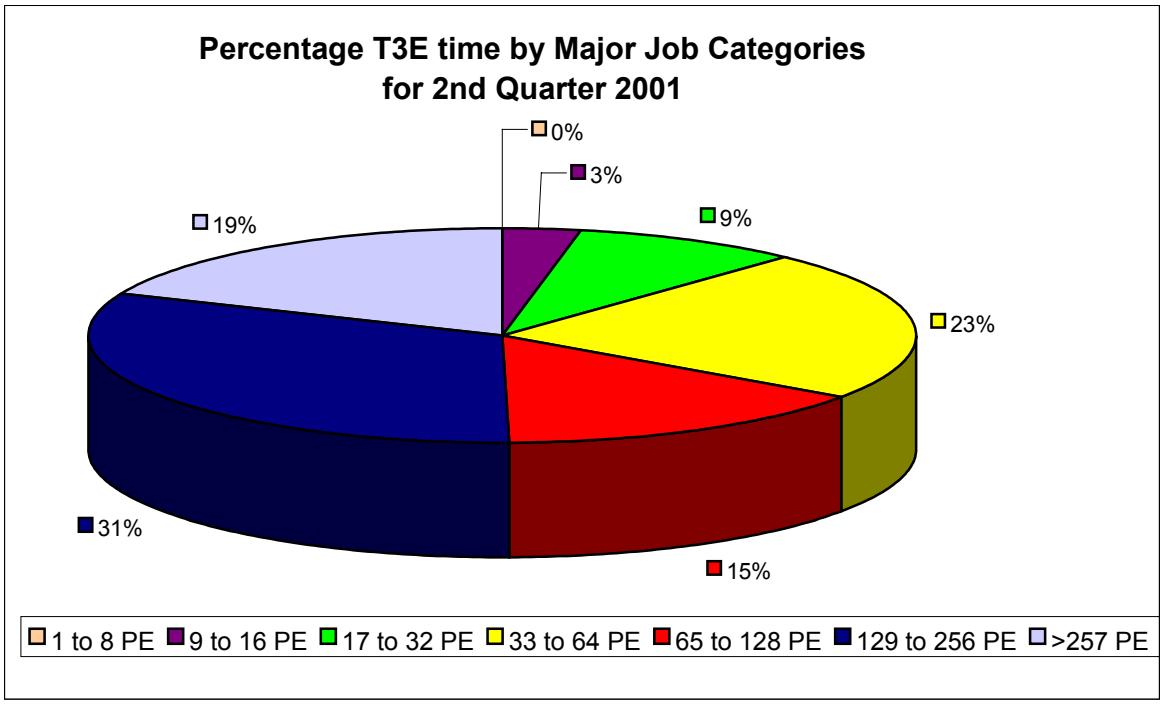

The trend on job size has shifted and now the predominant job size is in the mid range, with the percentage of jobs greater than 64 PEs for the Quarter being 65%.

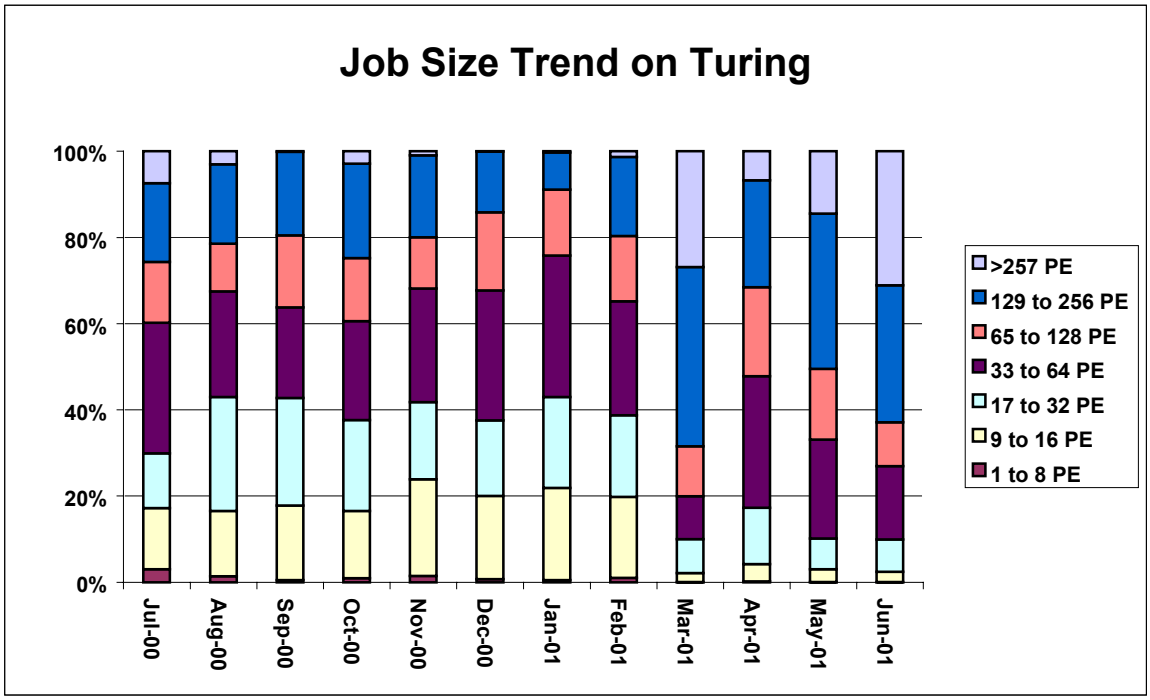

The above chart illustrates the effect that the introduction of the Fermat upgrade, and the Origin 3000/256 have had on the job size profile on Turing, where the use for capability jobs is being promoted

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 CSAR's systems by Research Council for the last 12 months, in particular the strong growth in usage during the past 6 months.

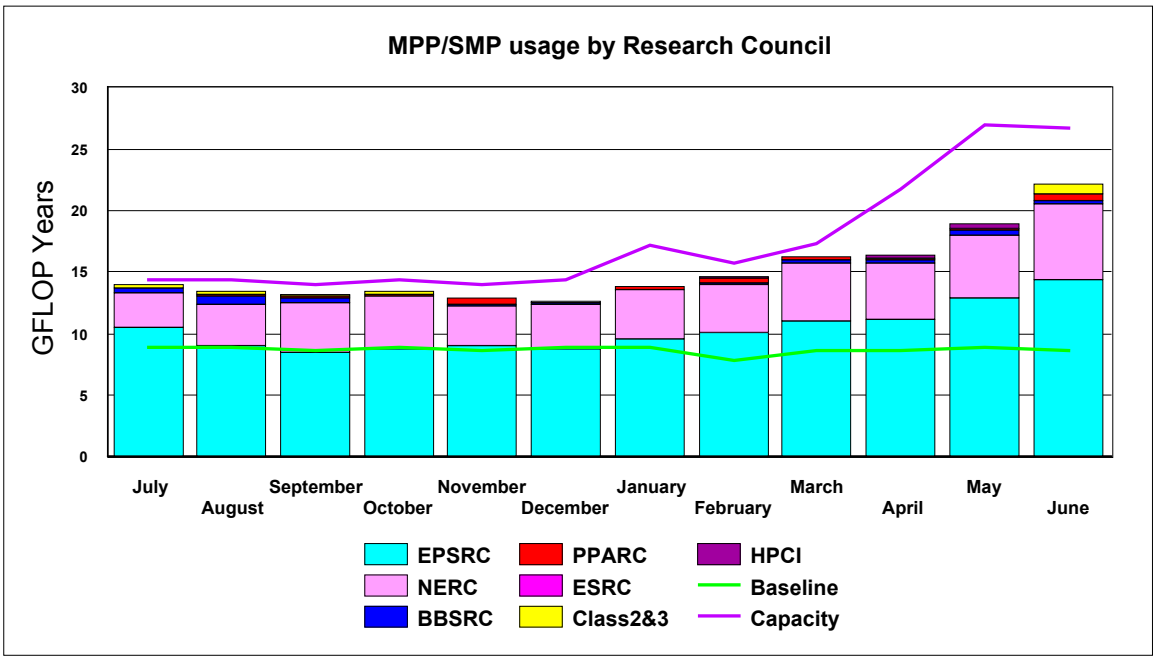

The graph below shows the same service utilisation by major system. .

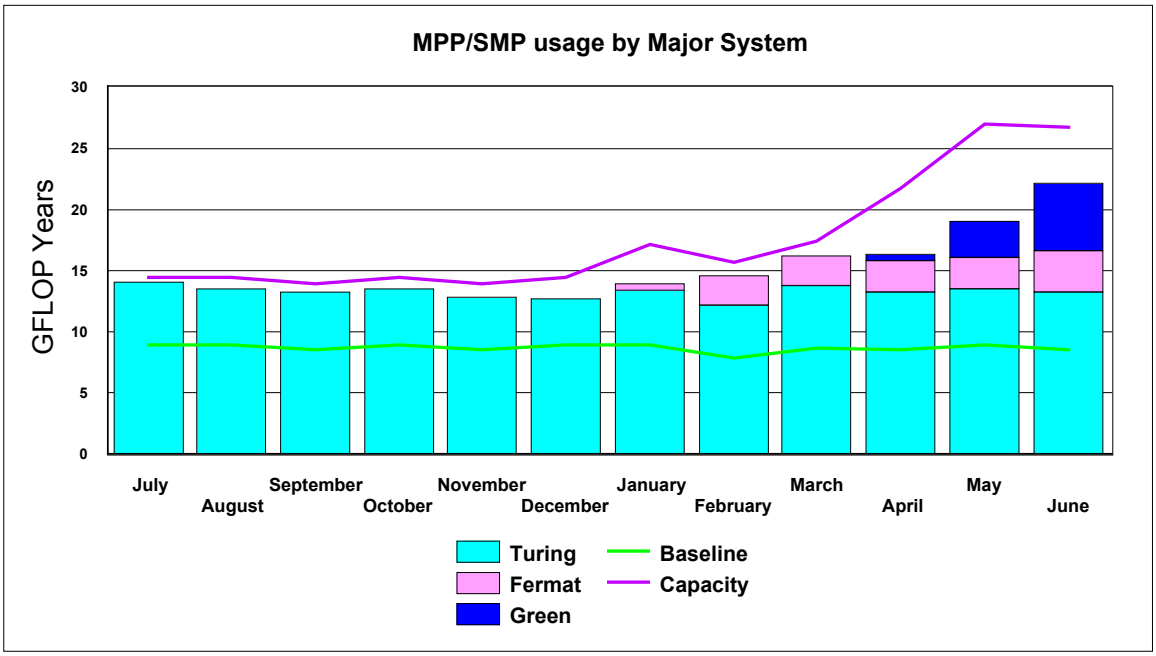

The next chart shows the historic interactive usage of the 'Baseline' Fermat system (equivalent to 16 @250Mhz CPU's).

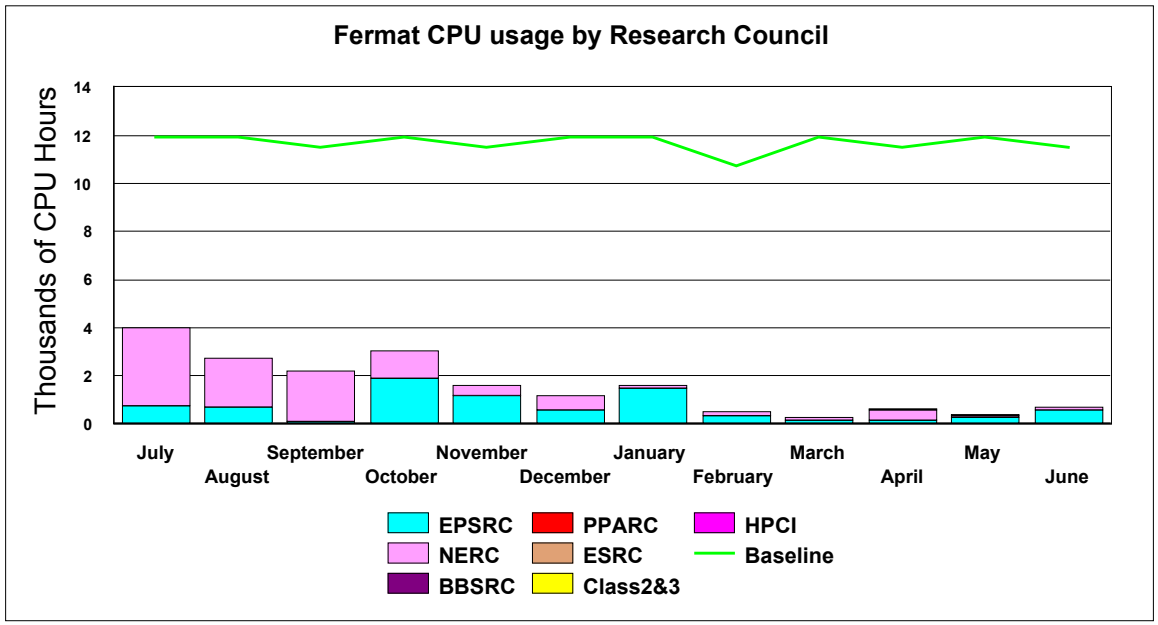

The next series of graphs illustrates the usage of the disk and HSM resources of the system.

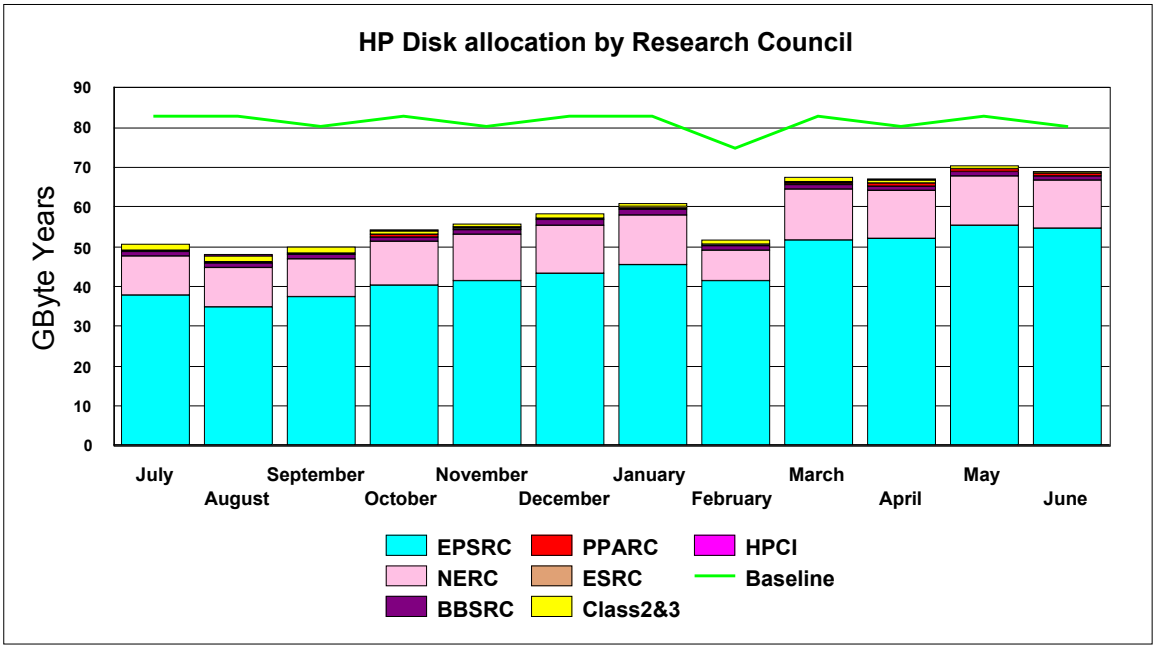

The preceding graph illustrates the historic allocation of the High Performance Disk on Turing, showing slight increases in demand.

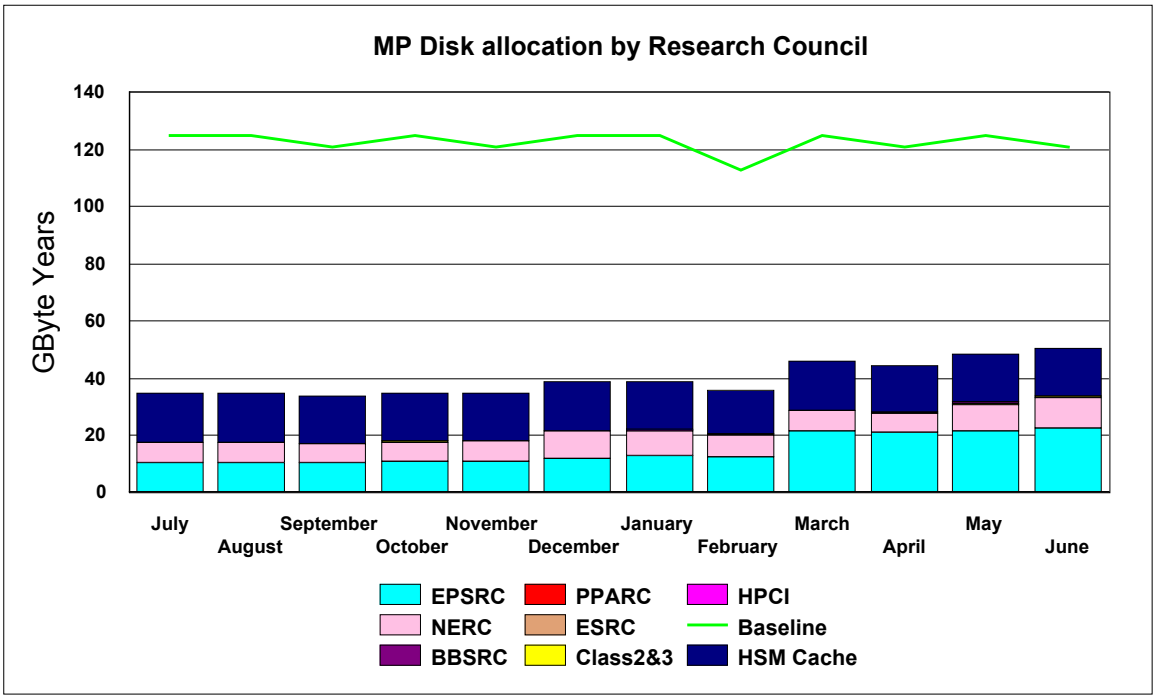

The graph above illustrates the historic allocation of the Medium Performance Disk on Fermat. From July 2000, 200 Gbytes has been used as a data cache for the enhanced HSM system.

The graph below shows the historic HSM usage by Research Council funded projects. The primary usage is for NERC, and equates in total to about 14 Terabytes at present.

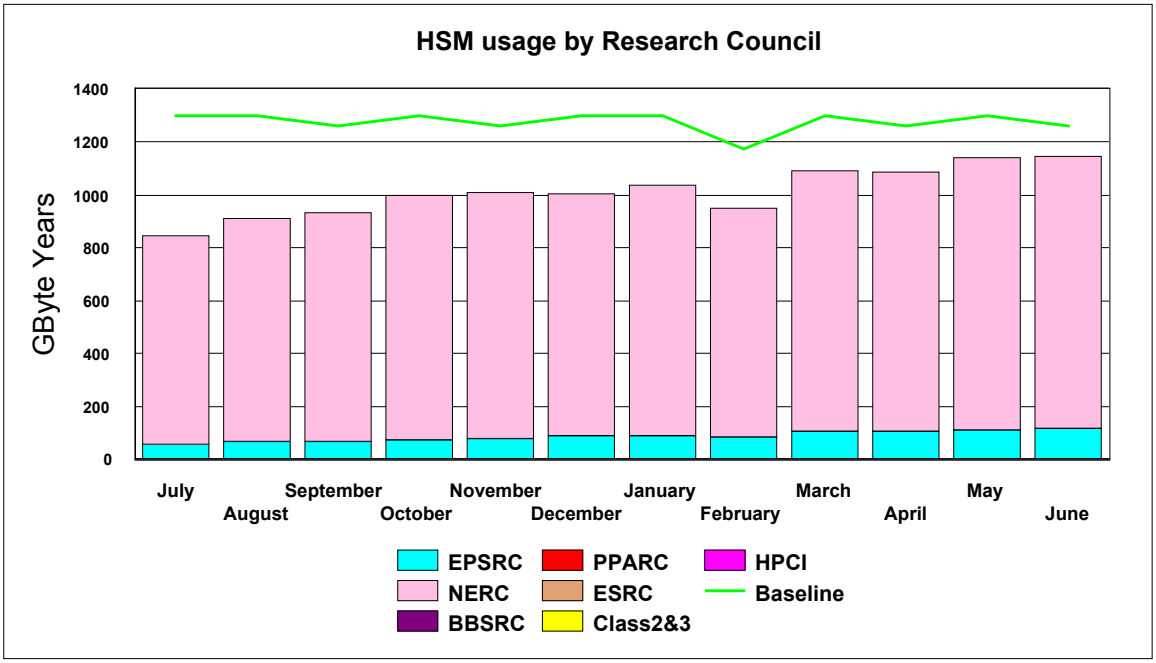

2.2.2 Fujitsu System Usage Graph

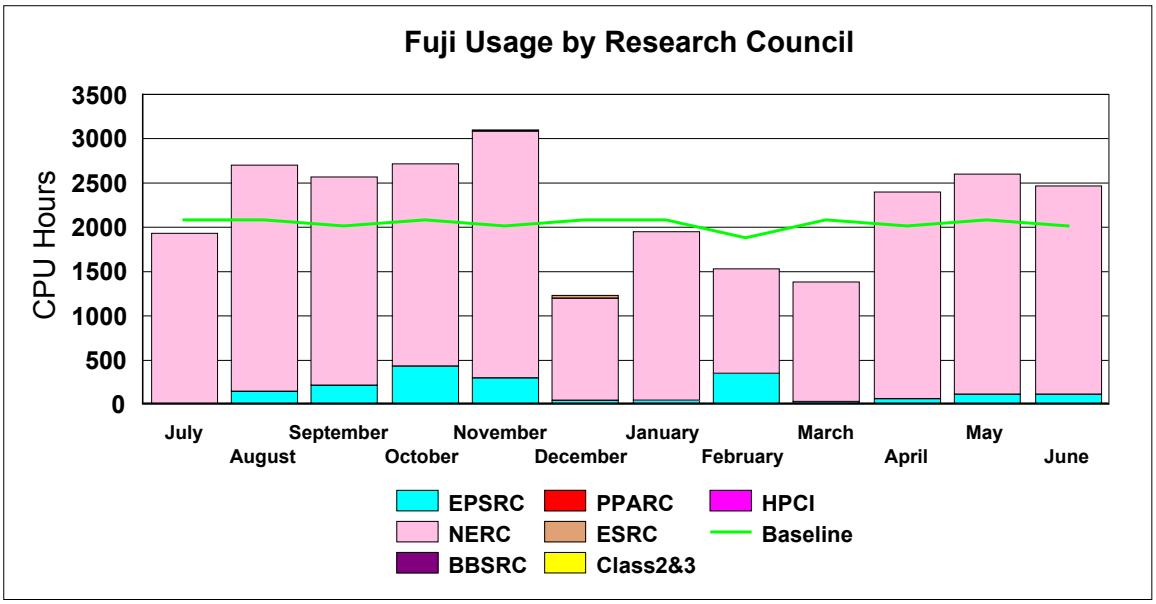

The above graph shows the current CPU usage on the Fujitsu VPP 300 NERC system based at the University of Manchester. The Fujitsu usage exceeded the baseline for the quarter.

2.2.3 Guest System Usage

A Compaq ES40 cluster is now situated at Manchester. A NEC SX4 vector system is available through CSC's Maidstone Data Centre. Neither systems have any current usage and may be withdrawn from service within the next 3 months if there continues to be negligible usage.

2.3 Service Status, Issues and Plans

Status

The service continues to be heavily used.

The IA64 prototype system is now being used for single CPU code optimisation work and is soon to be upgraded to the next phase with production standard Itanium based components.

The new Origin 2000/128 (Fermat) and Origin 3000/256 (Green) are fully installed and released as production batch engines.

Issues

Wait times can still be excessive at times due to the time demands on all the machines, however the additional Fermat and Green resources are starting to reduce the queue times on Turing.

Grant requests can exceed the term of the current CSAR Service and consortia may thus not consider the CSAR Service if their requirements demand a service over a longer period. Enhancements to the CSAR Service are becoming more difficult to justify with only 3.5 years of the current contract term to complete.

NERC have a requirement to hold 12 Terabytes of additional data (ECMWF Data Sets) and the optimum (lower cost) solution would require the HSM system to be upgraded and necessitate a larger disk cache area for the more current information to be held nearline.

Plans

Green will be upgraded to a 512 processor system by the end of July 2001 with, initially, a 256 processor job limit. This however will be increased to allow capability users to run jobs approaching 512 processors in size.

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 Management

Work was carried out in the registration system and on Turing to enable a consortium to allocate quota at a group level. The registration system and Green were modified to allow sub-allocation of green CPU to subprojects. It was agreed at a CfS Board meeting that a 10% reduction will be made in the cost of jobs using 512 or more processors on Turing. This came into effect on 1st June and will operate as a refund of resources to be made on a monthly basis.

3.2 User Registration

Code is being produced to allow CSAR users to submit their postal address upon self-registration.

3.3 Annual Report

The CSAR Annual Report 1999-2000 has been published on the CSAR webpages:

http://www.csar.cfs.ac.uk/admin/reports/index.shtml

3.4 User Survey

The summary report for the Second Annual User Survey has been published on the CSAR webpages:

http://www.csar.cfs.ac.uk/admin/reports/user_surveys/

3.5 User Steering Group Meeting

The 6th User Steering Group Meeting was held in June 2001 at the University of Manchester. It was attended by representatives from the CSAR user community, EPSRC, NERC, and CfS. The major topic of discussion was the Technology Refresh to the CSAR Service.

3.6 CSAR Audit

This review of the CSAR Service has been completed and presented to the CfS Management Board.

3.7 CSAR Focus

The 7th (Summer) edition of CSAR Focus has been produced and will be distributed in early July.

3.8 Documentation

A provisional website has been released for the "HPC on SGI NumaFlex" Summer School (8th - 12th October 2001). This summer school will be focused primarily on topics of immediate relevance to users of Green, in light of the decision to make Green the next CSAR flagship facility. The CSAR Technical Survey was completed by a small number of users. Their views will be taken into account in shaping the CSAR Service.

Additional work was carried out on the Green and Fourier introductory guides.

4 Science Application Support Services

4.1 Consortia Support

Code optimisation for Fuji and Turing has been carried out on behalf of two consortia.

Fine tuning of AMBER and installation of the smp version of AMBER was performed on behalf of another consortium.

4.2 Training and Education

Courses given:

- SGI Origin3000

Developments have been made to the following courses:

- IA-64 course
- Moving from C to C++ (formerly C++ for Scientific Applications)

The course schedule for the first semester of the year 2000/2001 is now being drawn up.

4.3 Service Developments

4.3.1 Fermat

The latest version of MPT (1.5.0.0) is now the default. The default version of the MIPSpro compilers has been upgraded from version 7.3.1.1 to 7.3.1.2. The default version of SCSL was upgraded from 1.1.0.0 to 1.3.0.0. Work was carried out to establish queue limit specifications.

4.3.2 Turing

The popular text editor 'pico' has been installed in /usr/local/unsupported bin. Access to the machine via telnet has been withdrawn in favour of the secure shell, ssh. There has been improvement in job turnaround resulting particularly from work being moved to Fermat and Green, and this is expected to continue as more people move their codes to Green.

4.3.2 Green

Green, the 128 400 MHz processor SGI Origin3000, was installed in April. It was subsequently upgraded to 256 processors with 256 Gbytes memory in May. It, therefore, already has more total memory than Turing and four times the memory per processor of Turing. Green will be upgraded again to 512 processors in July to meet the CSAR Technology Refresh requirements. The latest version of MPT (1.5.5.0) is now the default. AMBER 6.0, Gaussian98 and NAG have been installed.
The default version of SCSL was upgraded from 1.1.0.0 to 1.3.0.0. The default version of the MIPSpro compilers has been upgraded from version 7.3.1.1 to 7.3.1.2.

4.3.4 Fourier

Fourier was upgraded to 32 processors. CSAR staff attended a course on Itanium provided by SGI.

4.3.5 Kelvin

CSAR staff attended a course provided by COMPAQ.

4.3.6 Fuji

Access to the machine via telnet has been withdrawn in favour of the secure shell, ssh.

5 Collaboration and Conferences

5.1 UKHEC Reports

- The UKHEC case studies on VIPAR and the Terra consortium are being finalised before publication on the UKHEC website.
- The VR/Visualization technology watch report has been completed and has been sent to Daresbury for QA.

5.3 UKHEC Newsletter

MRCCS edited the Spring issue and contributed articles on Mini Grids, SG Global and the Access Grid and reviews of Visualisation and Grid conferences.

5.4 MRCCS Projects

5.4.1 EuroGrid

The MRCCS role in this project is to develop a GRID resource broker. Jon MacLaren is currently completing an implementation of UNICORE on Turing. Jon attended the EUROGRID/UNICORE Joint Technical Meeting, Bruhl in May and the EUROGRID Project Meeting, Paris in June. Jon also spoke on the project at the recent National Services Representatives' Committee Meeting in Manchester.

5.4.2 Grid Support

- The development of the Access Grid node at MC is progressing well.
- MRCCS is part of the collaborative venture recently funded to provide Grid support to the UK academic community.
- MRCCS has been awarded funding for the regional e-Science support centre in the North West.

5.4.3 North West Centre for Advanced Virtual Prototyping

MRCCS is part of a consortium from the Universities of Manchester, Salford, Lancaster and UMIST, which has been awarded £1.7M from the OST North West Science Funds to set-up a Research Centre in Advanced Virtual Prototyping. This will result in the appointment of a new member of staff to work on the MRCCS component of the project - an advertisement has been sent to press and has been published on the CSAR and MRCCS websites.

5.5 Events

5.5.1 Events Attended

Kevin Roy and Daniel Kidger attended CUG Summit 2001, May 21-25th, Indian Wells, California. A presentation was given with respect to CUG 2002, which is being hosted by MRCCS.

Jo Leng attended a meeting for VVECC (Visualization and Virtual Reality Community Club) at RAL in May.

Jo Leng attended the NWVRG meeting in May.

Jon MacLaren attended the EuroGlobus Workshop, Lecce, Italy in June.

Jon Gibson presented a talk entitled "The Jodrell Bank Pulsar Search Code and the Move to Real-Time Processing" at the 4th HLRS Metacomputing workshop in Stuttgart, Germany in May.

John Brooke attended and presented at a number of Grid related meetings. This included an invitation to give a demonstration at HPCN 2001on 'Getting Maximum Performance from Transcontinental Networks' involving the Cray T3E system from CSAR and the Universities of Stuttgart and Pittsburgh.

Stephen Pickles presented at the 'IEEE International Symposium on Cluster Computing and the Grid' at Brisbane in May, and visited supercomputing sites in Canberra and Melbourne.

Daniel Kidger is preparing material for an engineering conference in July.

6 Added Value Services

6.1 Joint Projects

The VIP Laboratory

One of the main goals of the VIP Laboratory is to enable high quality visualisation of data obtained through HPC. Where appropriate, an immersive environment using active stereo can be employed. A variety of research groups are currently using the facility, including applications in engineering, medicine, earth sciences, and building planning. Scientists can visualise their data without having to be skilled computer programmers. The software currently available to them includes AVS/Express multipipe edition, REALAX Virtual Reality software, and MUSE Development Environment 2000. Recent visitors include Schlumberger Geoquest, Immerse New Media, Wentworth College and Kubota, Japan. Enhancements have been made including a new MPE demo suite. The VIP Laboratory features in a television article on Virtual Medicine now being broadcast regularly on the Einstein TV Cable network. For details on how to access this facility
contact Dr Nigel John (n.w.john@man.ac.uk).

6.2 Cray MTA Optimisation

Work has started on the optimisation of a genomics code for use on the Cray MTA. This has involved two weeks training in Seattle.

6.3 International Conferences

Organisation of the two European conferences to take place in Manchester in the summer of 2001, Europar and Eurographics, are progressing well. Europar will take place on 28th-31st August 2001, and Eurographics shortly afterwards on 3rd-7th September. Web pages are available for both: http://www.man.ac.uk/europar/ and http://www.eg.org/egorg2001

Europar is the major European Parallel Computing Conference with over 400 attendees. In 2001 it is hosted by University of Manchester and is being organised jointly by Computer Science and CSAR/MRCCS.
Joanna Leng, John Brooke and Kaukab Jaffri of CSAR are on the Local Organising Committee. Kaukab Jaffri has produced the Advance Programme, a copy of which was sent out to CSAR users with the new edition of CSAR Focus. Invited speakers include Jack Dongarra, Ian Foster, Dennis Gannon, Tony Hey, Martin Kersten and Thomas Sterling.

The Eurographics 2001 Advance Programme is now available from the EG2001 website.

CUG 2002 will be held by MRCCS in Manchester. Initial preparatory work has already begun.

6.4 Seminars

The following MRCCS seminar was held in this quarter: The MRCCS Research Seminars Summer Programme is occasional. CSAR has invited all CSAR consortia to give a seminar on current research topics at the seminars on High Performance Computing and Visualisation on Friday afternoons. Such visits could be combined with the free code portability assessment programme.

Dr. Steven Newhouse, Parallel Software Group, CS, Imperial College "Software Environments to Effectively Utilise Federated Computational Communities".

Jo Leng gave a seminar in Lancaster entitled "Visualisation of Spherical Geometries Produced for Geophysical and Astrophysical Research".