CSAR Service

Consolidated Management Report

3rd 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 519 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) now has 512 CPUs, 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 Oct	Nov	Dec	Jan	Feb	March	April	May	June	July	Aug	Sept
HPC Services Availability												
Availability in Core Time (% of time)	100%	100%	94.90%	99.70%	99.70%	100%	100%	99.70%	99.70%	98.49%	98.49%	98.49%
Availability out of Core Time (% of time)	100%	99.40	98.49%	99.50%	99.40	99.40	99.40	99.40	99.40	98.49%	100%	99.40
Number of Failures in month	0	2	4	1	1	1	1	3	3	4	2	2
Mean Time between failures in 52 week rolling period (hours)	1095	673	584	584	626	674	674	584	584	438	398	365
Fujitsu Service Availability												
Availability in Core Time (% of time)	100%	100%	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	<1	<3	<3	<5	<5	<3	<5	<2	<2	<1	<1	<1
Administrative Queries - Max Time to resolve 95% of all queries	<0.5	<0.5	<5	<2	<2	<3	<0.5	<0.5	<0.5	<1	<2	<1
Help Desk Telephone - % of calls answered within 2 minutes	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Others												
Normal Media Exchange Requests - average response time	0	<0.5	<0.5	<0.5	<0.5	0	0	<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	12	10	10	10	10	10
System Maintenance - no. of sessions taken per system in the mon	1	2	1	0	2	1	2	0	0	1	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 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 Oct	Nov	Dec	Jan	Feb	March	April	May	June	July	Aug	Sept
HPC Services Availability												
Availability in Core Time (% of time)	-0.058	-0.058	0.195	-0.039	-0.039	-0.058	-0.058	-0.039	-0.039	0.039	0.039	0.039
Availability out of Core Time (% of time)	-0.047	0	0.039	-0.039	0	0	0	0	0	0.039	-0.047	0
Number of Failures in month	-0.009	0	0.008	-0.008	-0.008	-0.008	-0.008	0	0	0.008	0	0
Mean Time between failures in 52 week rolling period (hours)	-0.009	-0.008	-0.008	-0.008	-0.008	-0.008	-0.008	-0.008	-0.008	0	0	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	0.031	0.031	0.016	0.031	0	0	-0.016	-0.016	-0.016
Administrative Queries - Max Time to resolve 95% of all queries	-0.019	-0.019	0.046	0	0	0.016	-0.019	-0.019	-0.019	-0.016	0.016	-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	-0.002	-0.002	-0.002	-0.002	0	0	-0.002	-0.002	-0.002	0	-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	0.003	0	0	0	0	0
System Maintenance - no. of sessions taken per system in the mont	-0.003	0	-0.003	-0.004	0	-0.003	-0.003	-0.004	-0.004	-0.003	0	0
Monthly Total & overall Service Quality Rating for each period:	-0.10	-0.06	0.12	-0.05	-0.03	-0.04	-0.05	-0.06	-0.06	0.00	-0.02	-0.02
Quarterly Service Credits:			-0.03			-0.13			-0.18			-0.04

Table 3

1.2 No. Of Registered Users

The current position at the end of the quarter is that there are 519 registered users of the CSAR Service.

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

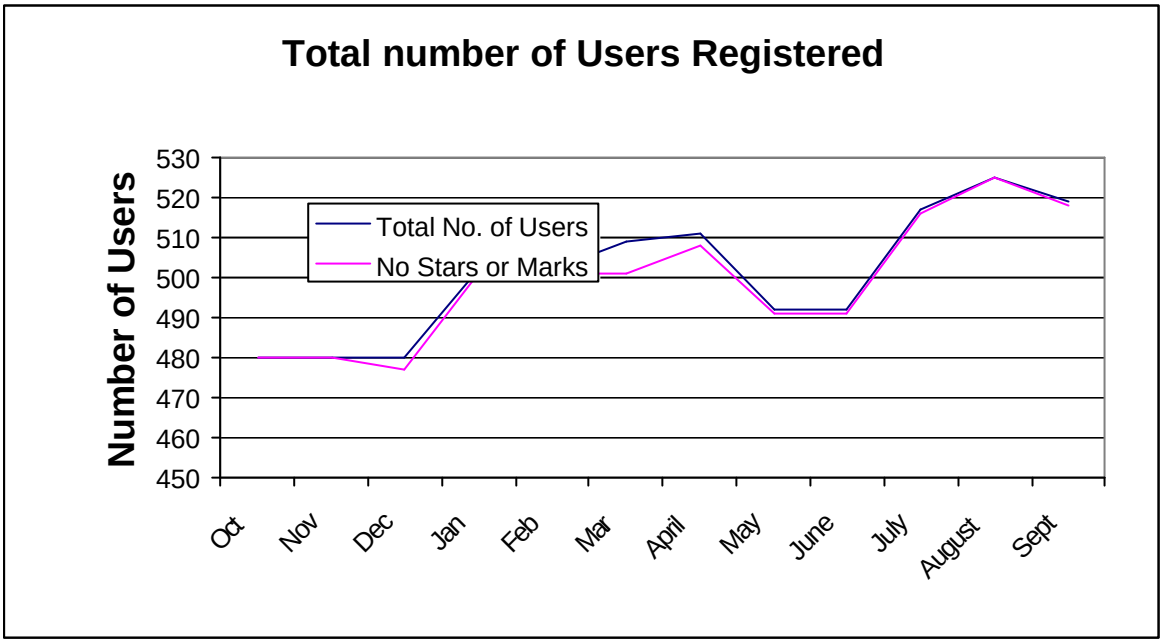

1.3 Service Quality Tokens

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

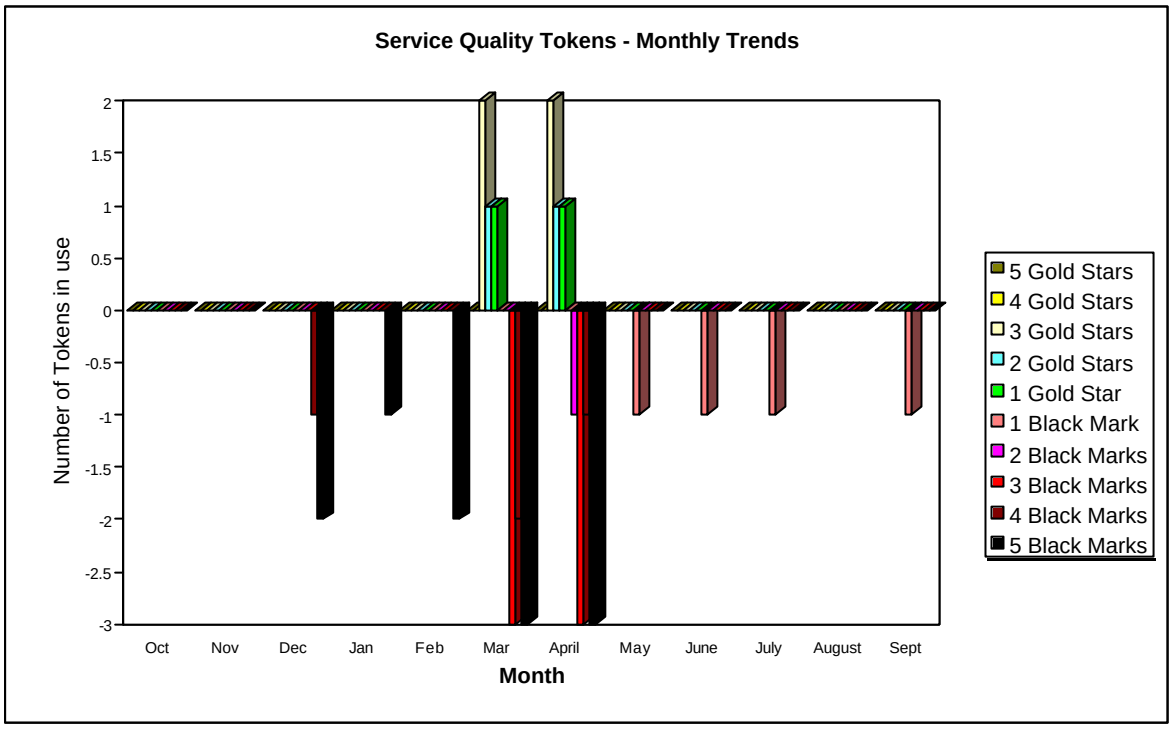

Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

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

No of Stars or Marks	Consortia	Date Allocated	Reason Given
1 Black Mark	cse006	25/09/01	Login problems continuing

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.

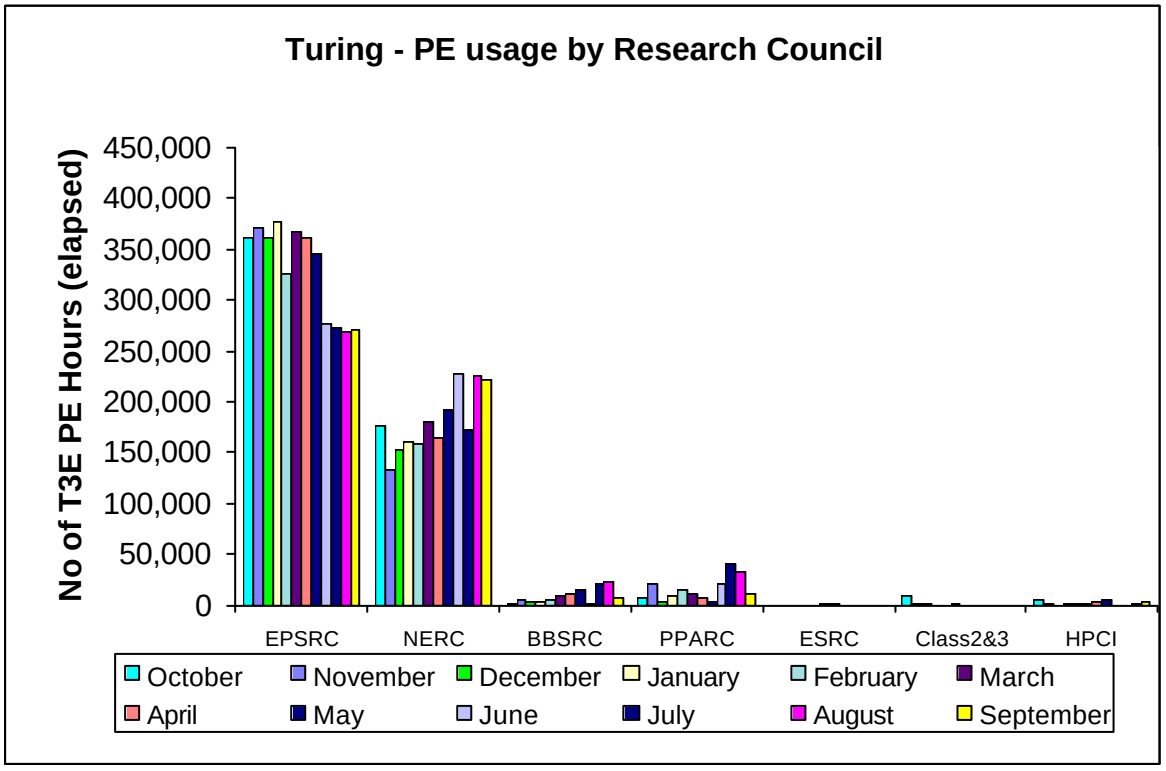

Turing PE usage is shown by Research Council during the last 12 months of service is shown in the above chart.

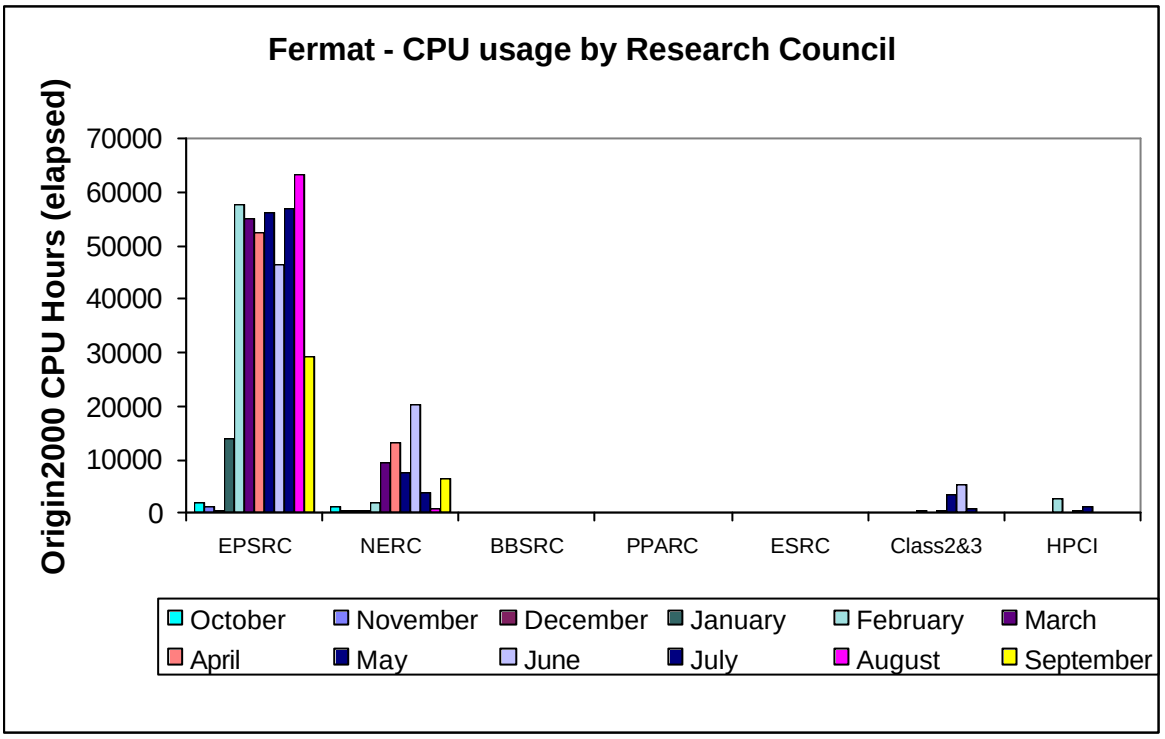

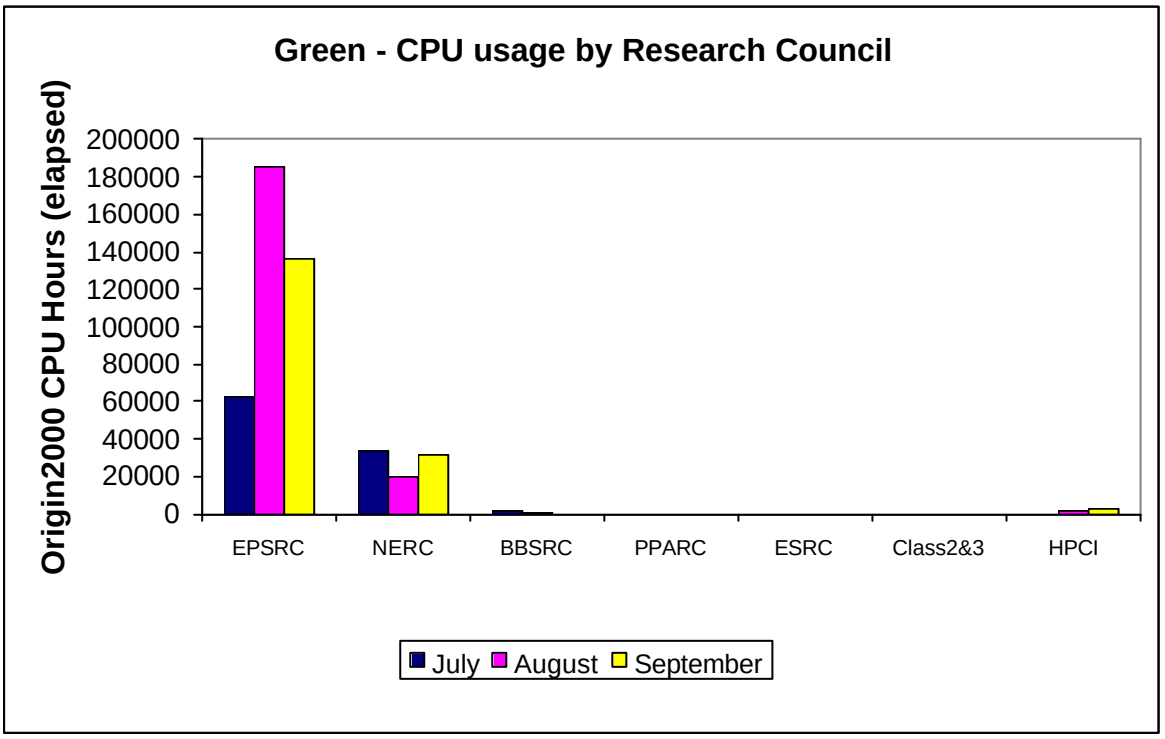

Usage of the two SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service is shown in the above charts.

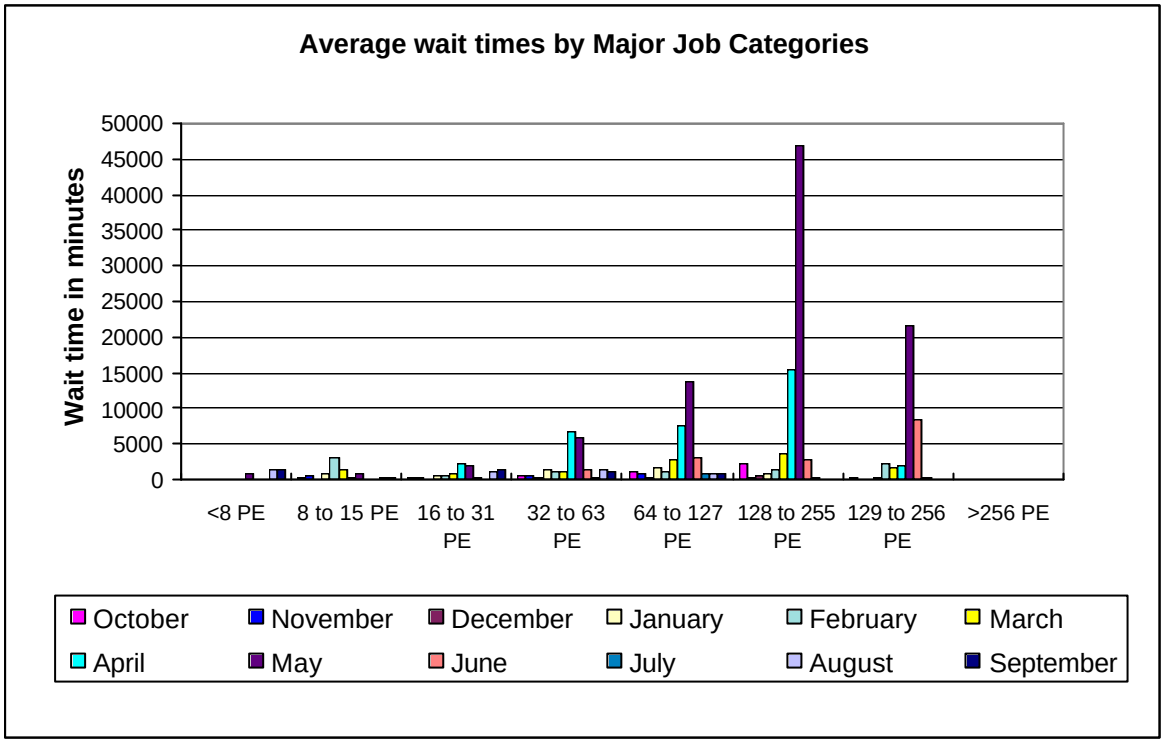

The above chart shows the wait time trend on the Turing system. The trend towards rising wait times is now being reversed for all job sizes.

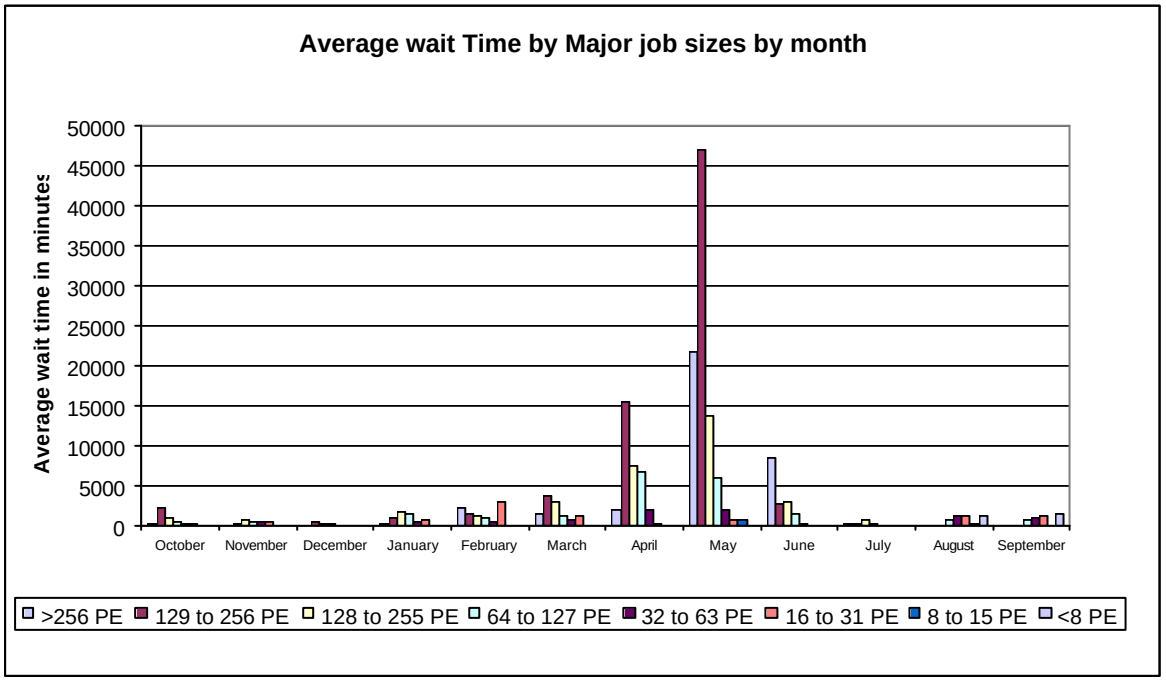

It can be seen from the above graph that the additional capacity added to the service, in the shape of the upgrade to Fermat and the addition of the Origin 3000 (Green), are now reducing the burden on the T3E (Turing) and reducing the queue wait times.

The next chart shows the percentage PE time utilisation by the major job categories on the Turing system for the 3rd 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%.

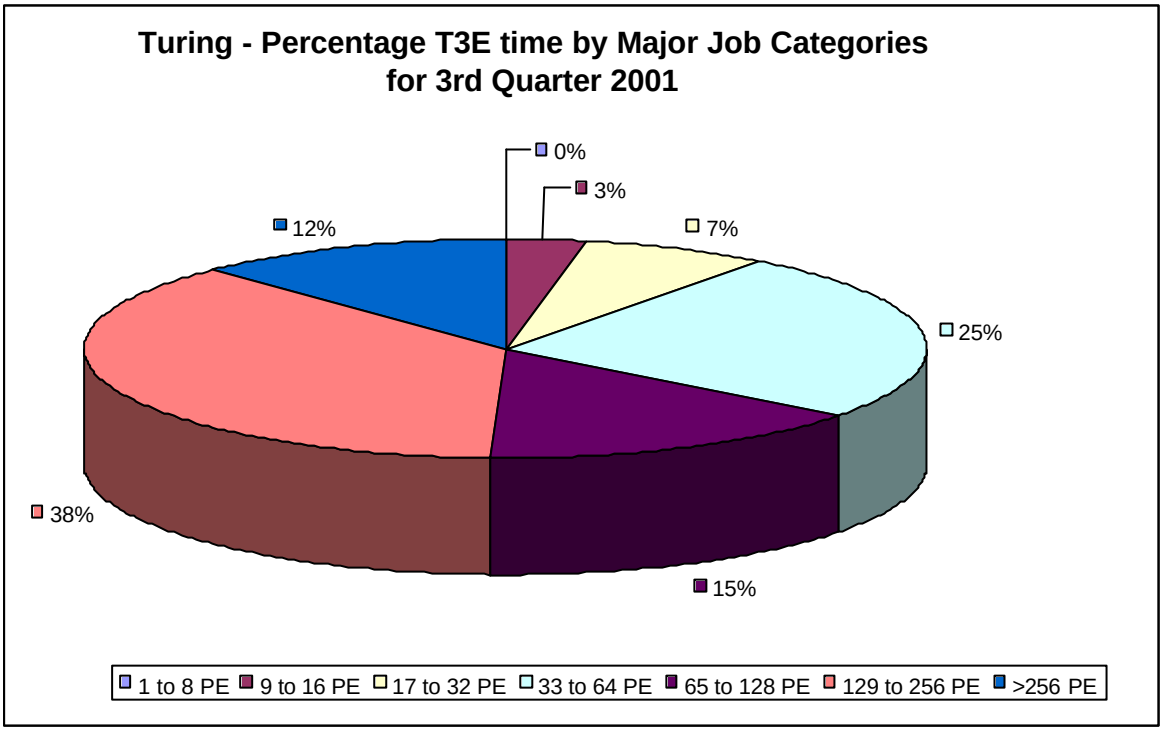

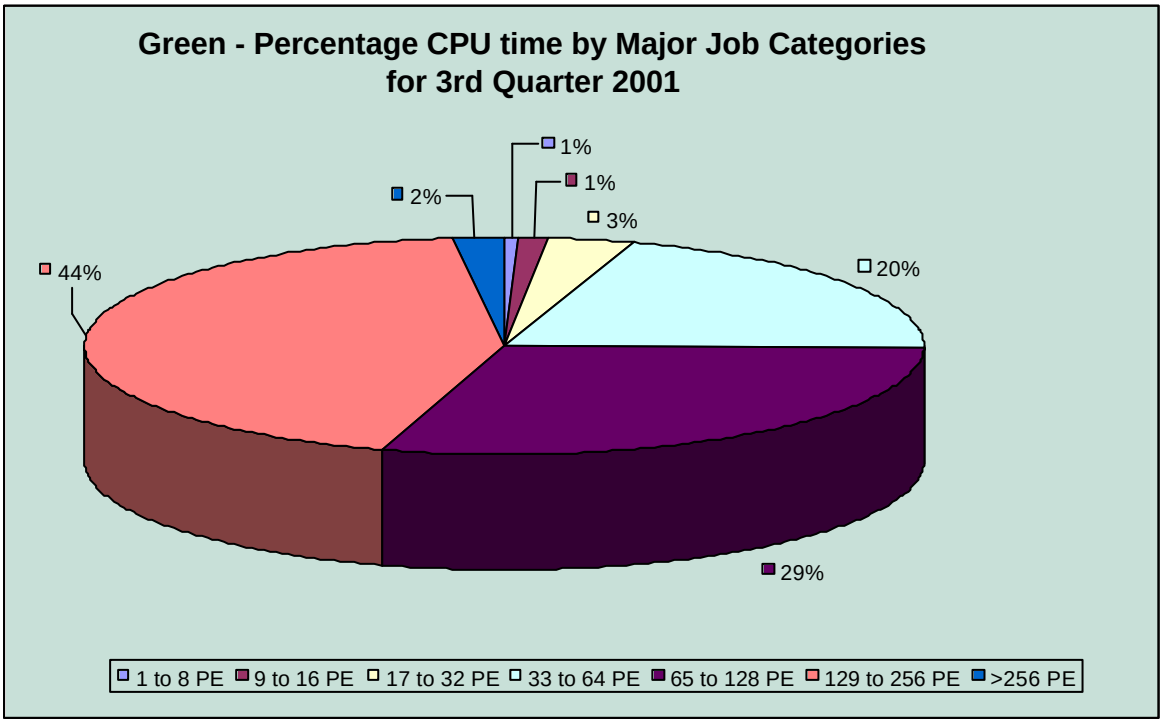

On Green, the percentage of CPU utilisation used by jobs requiring greater than 64 CPUs for the quarter was 75%.

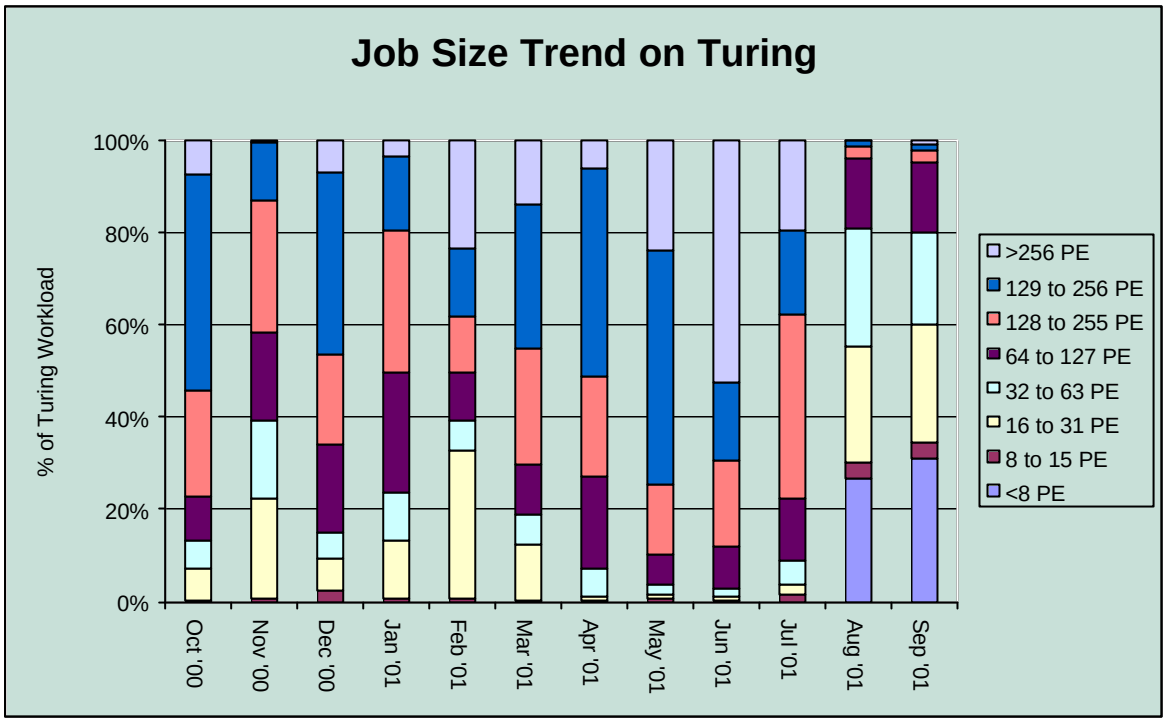

The above chart illustrates the effect that the introduction of the Fermat upgrade, and the Origin 3000/256 have had on the job size profile on Turing. Many of the larger jobs have now moved to Green, leaving a wide spread of workloads on Turing

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 year's 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.

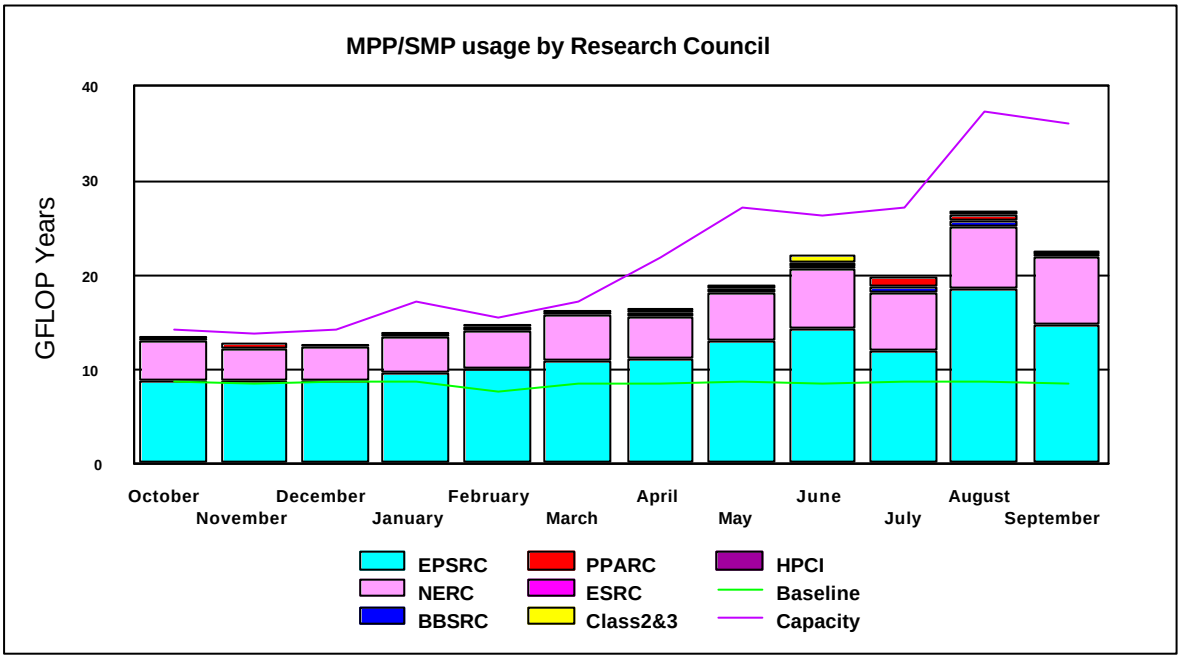

The graph below shows the same service utilisation by major system. .

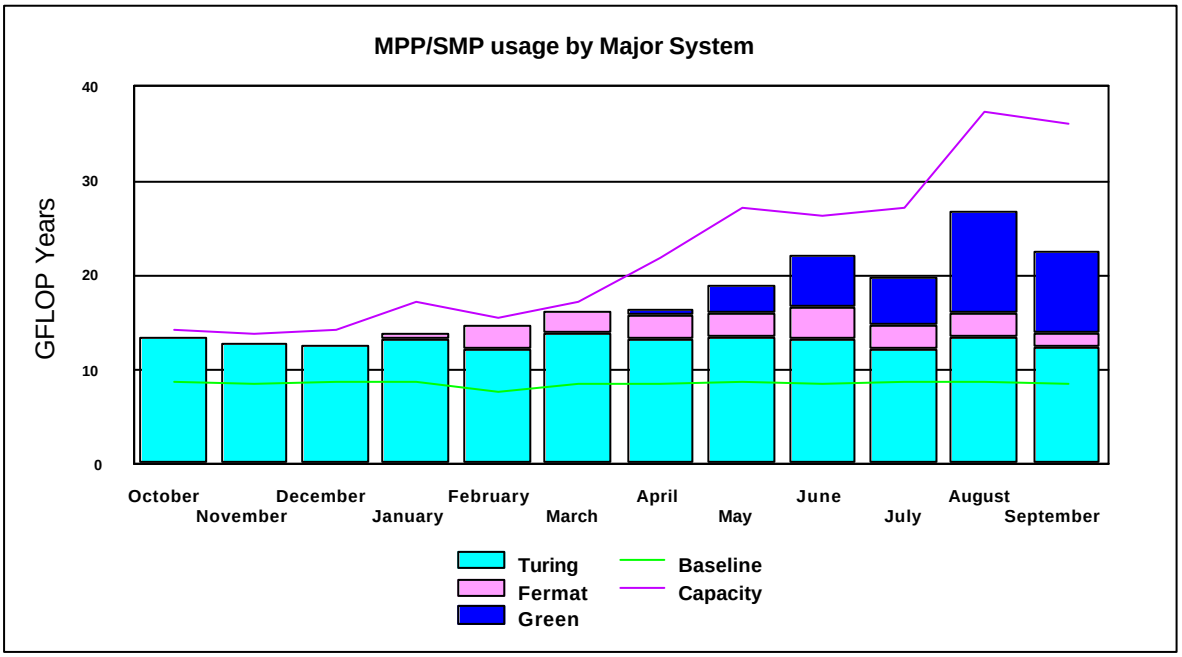

The next chart shows the historic interactive usage of the 'Baseline' Fermat system (equivalent to 16 @250Mhz CPU's).

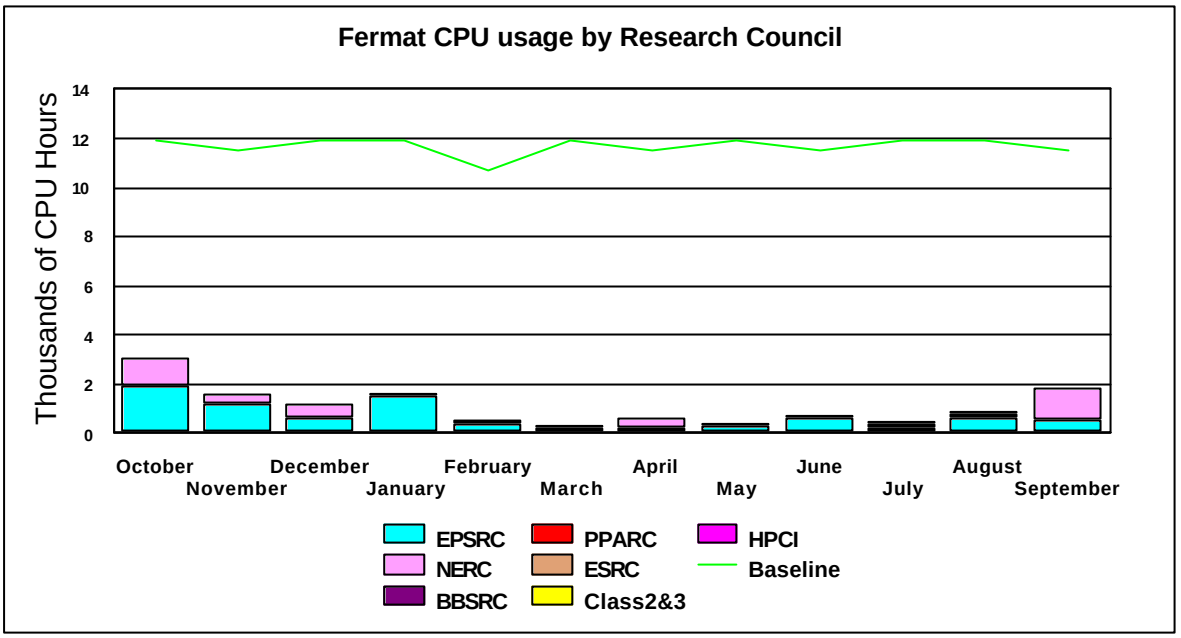

The next series of graphs illustrates the usage of the disk and HSM resources of the system.

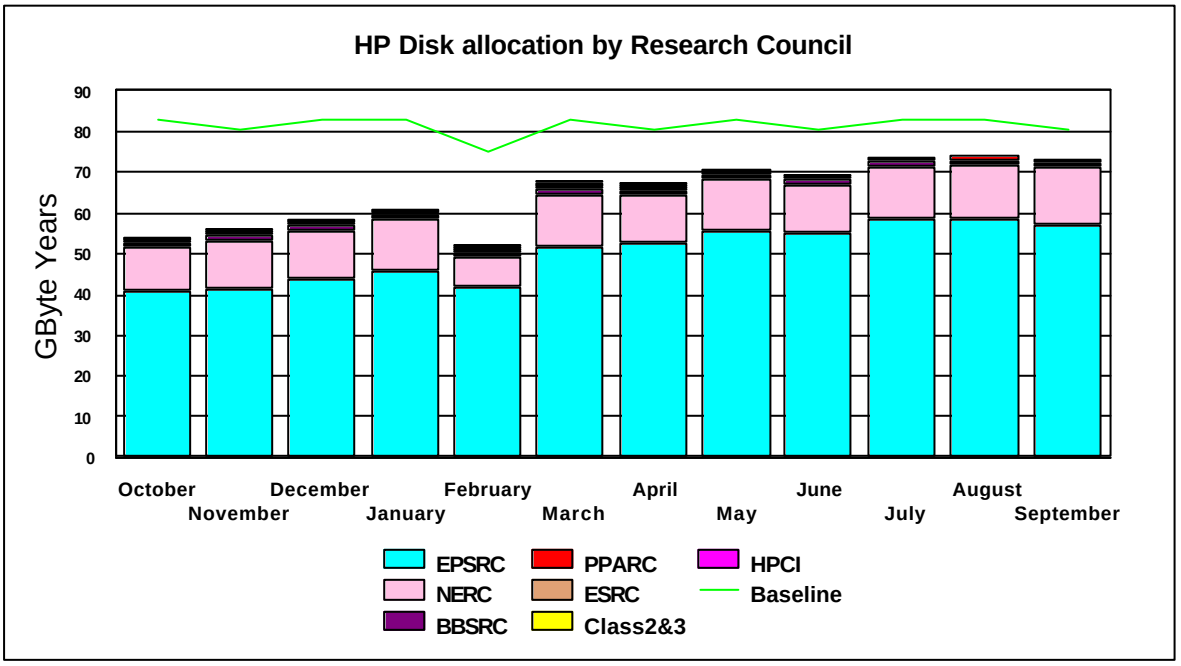

The preceding graph illustrates the historic allocation of the High Performance Disk on Turing, showing slight increases in demand.

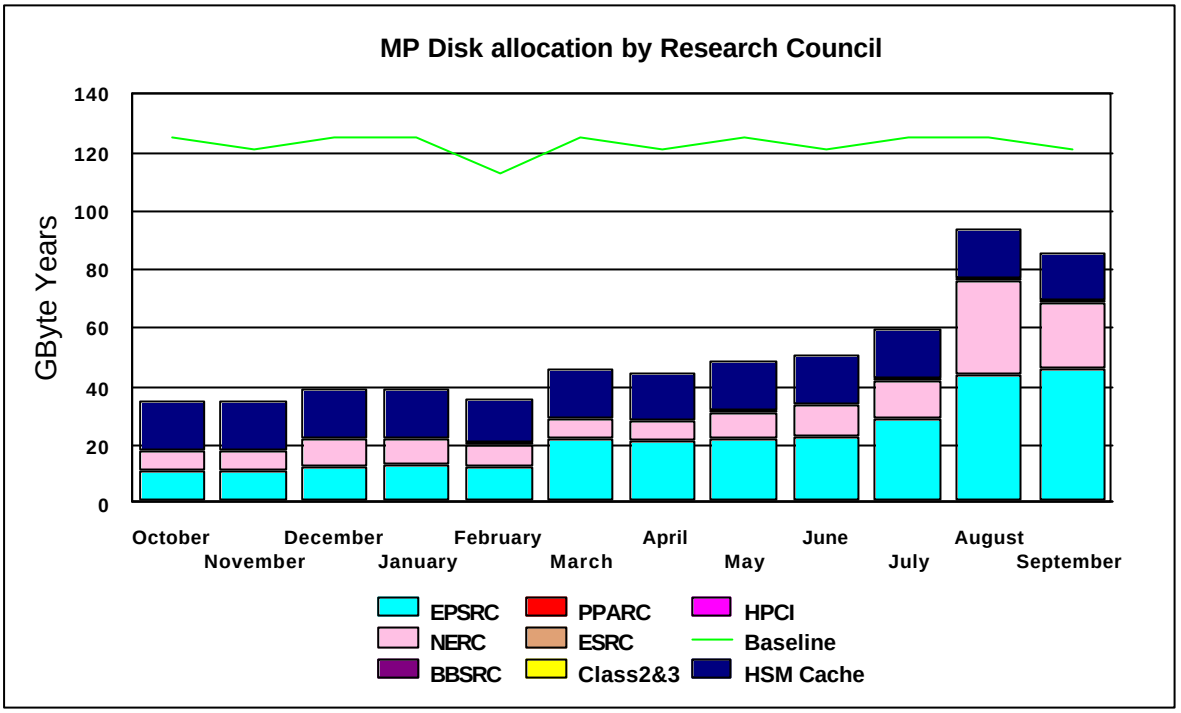

The graph above illustrates the historic allocation of the Medium Performance Disk on Fermat. From July 2000, 200 Gbytes has been used as a data cache for the enhanced HSM system.

The graph below shows the historic HSM usage by Research Council funded projects, which has just reached the overall Baseline of 16 Terabytes. The primary usage is for NERC and, in total, this amounts to about 14 Terabytes of data storage at present.

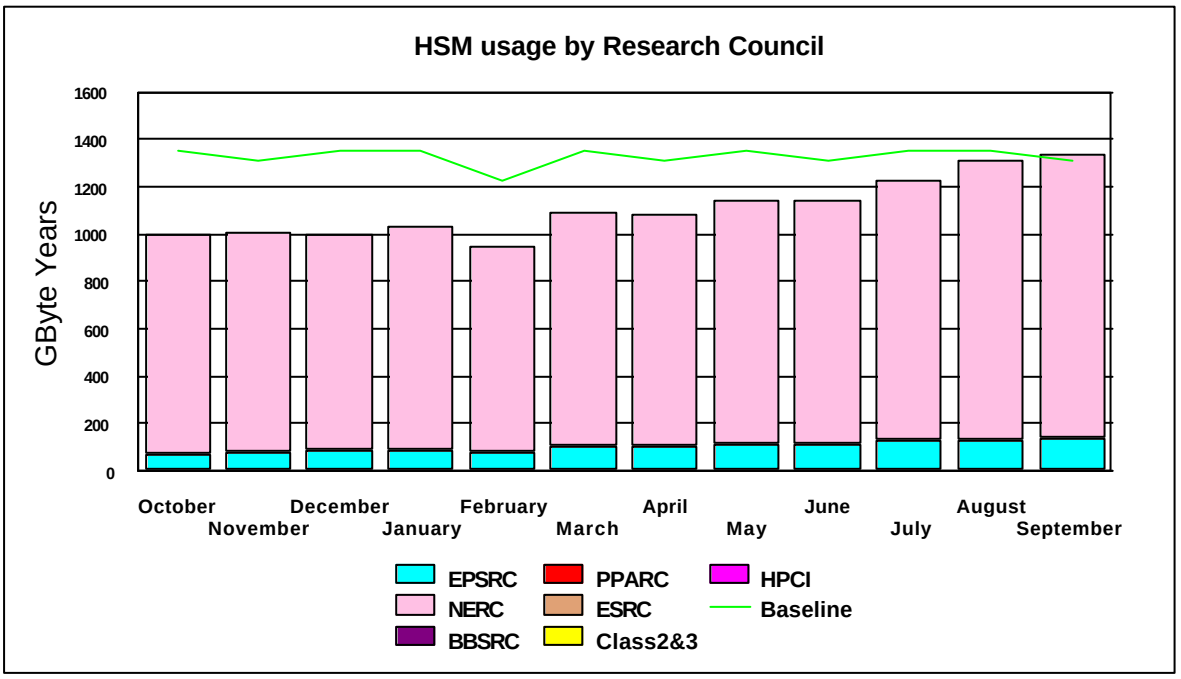

2.2.2 Fujitsu System Usage Graph

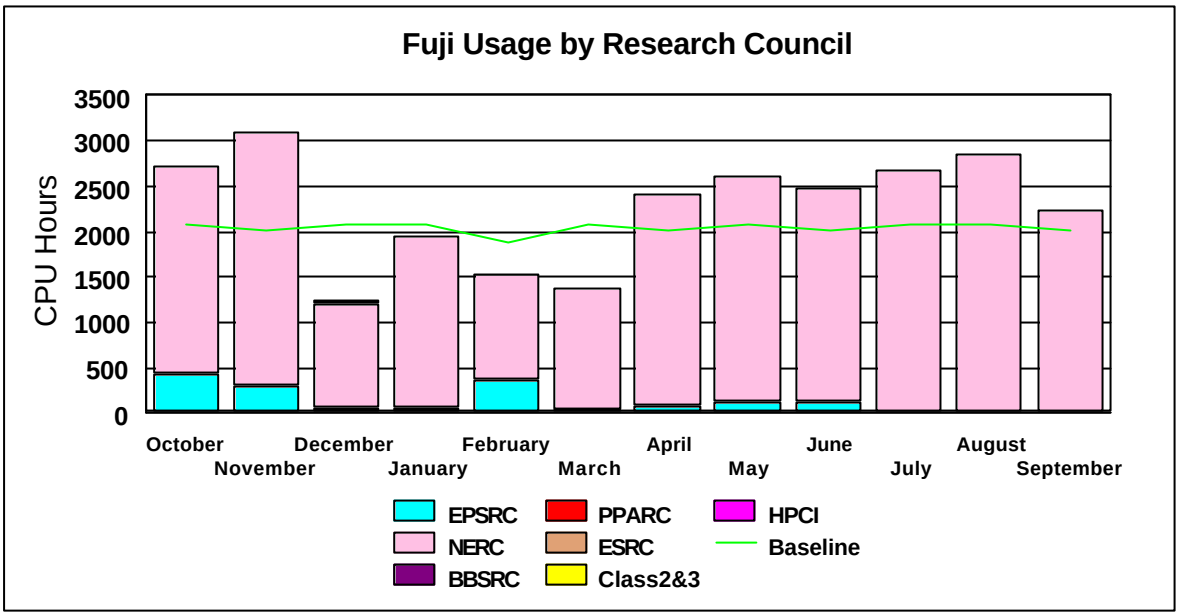

The above graph shows the current CPU usage on the Fujitsu VPP 300 NERC system based at the University of Manchester. The Fujitsu usage exceeded the baseline for the quarter.

2.2.3 Guest System Usage

A Compaq ES40 cluster is now situated at Manchester. An 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 3000/512 (Green) is fully installed and released as production batch engines.

Issues

Wait times can still be significant at times due to the time demands on all the machines, however the additional Fermat and Green resources are have substantially reduced the queuing 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 a little over 3 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 formally acceptance tested during November.

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 Proscribed Users

All users who have not agreed to the conditions of use have been denied access to the CSAR systems.

3.2 Documentation

A new section documenting the newly installed Bio-Informatics codes has been made available on the CSAR web pages.

3.3 Publicity

New CSAR and MRCCS leaflets were produced prior to the Europar and Eurographics conferences. Posters on CSAR, MRCCS and MVC work are now being prepared.

4 Science Application Support Services

4.1 Summer/Autumn School

The High Performance Computing Summer School "HPC on Intel Itanium" , scheduled for 2-6 July, was cancelled as a result of the decision to provide a MIPS based rather than Itanium based system for the CSAR Technology Refresh. A three-day course on use of the Origin3000 systems was provided instead.

Preparations for the HPC & SGI NUMAflex Autumn School (8-12 October) are in progress.

4.2 Courses

A 2-day training course was held in Manchester, entitled “Visualization & Virtual Reality for Scientific Applications”.

The Autumn 2001 course schedule has been produced and can be viewed on the website (www.csar.cfs.ac.uk/courses/).

Developments have been made to the OpenMP course.

4.3 Service Developments

4.3.1 Green

Green was upgraded to 512 processors with 512 Gbytes memory in July. It therefore now has more than double the total memory of Turing and four times the memory per processor of Turing.

4.4 Software

4.4.1 Bio-Informatics Codes

The following Bio-Informatics codes have been installed on both *Fermat* and *Green*:

FASTA
BLAST
HMMER
Wise2

FASTA and BLAST have been optimised for SGI Origin systems by SGI. Further details are available at: www.csar.cfs.ac.uk/software/bioinf.shtml.

4.4.2 Globus

Globus 1.1.4 has been installed on the CSAR systems *Turing* and *Fermat*, and the local system *Kilburn*. Installation on *Green* will follow.

4.4.3 Totalview

The Totalview debugger is being installed on Fermat. It will be available for debugging codes using up to 64 processors.

4.4.4 Vampir

The performance analysis tool Vampir is being installed on Fermat, for use in conjunction with the associated trace facility Vampirtrace to be installed on Green. These will provide additional details on communication performance of codes running on Green.

4.4.5 Message Passing Toolkit (MPT)

The latest version of MPT (1.4.0.2) is now the default

5 Collaboration and Conferences

5.1 UKHEC Reports

The three UKHEC reports, "Visualization of Spherical Geometries Produced by Large Scale Simulation on the Cray T3E at Manchester, UK", "VIPAR: A Case Study for UKHEC" and "Scientific Applications of Visualization, Virtual Reality and High Performance Visualization Computers" with "Scientific Examples of Virtual Reality and Visualization Applications" have all been refereed and are now available from the UKHEC web site:

`http://www.ukhec.ac.uk/publications/`

5.2 UKHEC Newsletter

MRCCS contributed a number of articles to and edited issue 3, published in July.

5.4 MRCCS Projects

5.4.1 EuroGrid

A specification for the resource broker has been written by Manchester with some input from FECIT. This was Manchester's first deliverable for the project.

5.4.2 Access Grid

Our Access Grid node is now fully operational. We have had successful demonstrations of the node to about 30 people at the EuroPar Conference and 50 people at the Reality Centre Special Interest Group (RCSIG) Conference, each of which involved remote sites from around the globe. We participated in the SC Global Production Institute (hosted by Boston University) and ANL, Chicago over the Access Grid. This was a week of training in Access Grid use for operators, speakers and producers. It was particularly aimed at SC Global 2001, at which we are hosting 2 events via the Access grid (on the subjects of Global Supercomputing and Solar-Terrestrial Physics). We will also be participating in events hosted by other remote sites

5.4.3 The Reality Grid

The major proposal for a GRID test-bed, called the RealityGrid has been successful. The objective is to enable the realistic modelling of complex solid and fluid structures at the meso and nano-scale levels, and or the discovery of new materials. High performance computing and visualization are critical to this test-bed. The project involves collaboration between a number of high profile groups from different universities and has substantial industrial backing.
This will result in four new appointments for MC and two for CNC.

5.4.4 E-Science

MRCCS was successful in EPSRC e-Science pilot projects, and will host the North West regional e-Science center. UoM will also be involved in the proposed UK Grid Support Center.

5.4.5 North West Centre for Advanced Virtual Prototyping

Meetings have been held with CNC and Salford to identify the first stages of the project. The MRCCS position has been advertised, but unfortunately nobody has been appointed. It will be re-advertised shortly.

5.5 Events

5.5.1 Events Attended

EuroGrid meeting in Paris in June.

Global Grid Forum (GGF2) in Washington DC.

Mike Daw attended the Internet2/UKERNA event in London in September at which he presented outlines of the work we have done under funding from this project, namely Global Supercomputing and Access grid. The 7 other sites that have current projects also made presentations.

At the second UKHEC annual seminar in Edinburgh in September, Jon MacLaren spoke about the various grid-related activities at Manchester, focussing on the EuroGrid project. Stephen Pickles chaired one of the sessions and was a member of the discussion panel.

Stephen Pickles, John Brooke and Michael Kramer authored poster "Processing astronomical data on the Grid in real time", presented by John Brooke at the HPC Town Meeting on 26 September.

5.5.2 Publications

S.M. Pickles, J.M. Brooke, F. Costen, E. Gabriel, M. Mueller, M. Resch
"Metacomputing across intercontinental networks"
Future Generation Computer Systems
17 (2001) 911-918.

UKQCD Collaboration: C.R. Allton, S.P. Booth, K.C. Bowler, J. Garden, A. Hart, D. Hepburn, A.C. Irving, B. Joo, R.D. Kenway, C.M. Maynard, C. McNeile, C. Michael, S.M. Pickles, J.C. Sexton, K.J. Sharkey, Z. Sroczynski, M. Talevi, M. Teper, H. Wittig
"Effects of non-perturbatively improved dynamical fermions in QCD at fixed lattice spacing" (hep-lat/0107021)

6 Added Value Services

6.1 The VIP Laboratory

Contact Nigel John

6.2 International Conferences

The two European conferences organized by MRCCS in Manchester this summer, Europar and Eurographics, were very successful.

Preparations have started for the Cray User Group (CUG) conference to be held in Manchester in May 2002.

6.3 Seminars

The seminar program for the Autumn 2001 term is now available: www.man.ac.uk/mrccs/seminars/

Douglas Johnson from the OHIO Supercomputer Center gave a seminar entitled:
"Cluster Computing and Computational Grids at The Ohio Supercomputer Center" during his stay at Manchester Computing in July.

Jarek Nabrzyski from Pozan supercomputer center gave an MRCCS seminar entitled “PIONIER – Polish Optical Internet: Advanced Applications, Services and Technologies for the Information Society.”

6.4 Summer Student Program

Eight students were involved in the ten-week summer student program. The projects involved various topics in both HPC and visualisation, including the development of the first UK Access Grid.