CSAR Service

Consolidated Management Report

4th 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 554 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. 79% of the usage has been greater than 128 CPUs in size.

The T3E has seen 83% of the Quarter's 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.

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	2001/2 Jan	Feb	March	April	May	June	July	Aug	Sept	Oct	Nov	Dec
HPC Services Availability												
Availability in Core Time (% of time)	99.70%	99.70%	100%	100%	99.70%	99.70%	98.49%	98.49%	98.49%	98.60%	98.60%	100.00%
Availability out of Core Time (% of time)	99.50%	99.40	99.40	99.40	99.40	99.40	98.49%	100%	99.40	99.50%	99.50%	98.49%
Number of Failures in month	1	1	1	1	3	3	4	2	2	2	2	4
Mean Time between failures in 52 week rolling period (hours)	584	626	674	674	584	584	438	398	365	365	365	337
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	<5	<5	<3	<5	<2	<2	<1	<1	<1	<1	<1	<1
Administrative Queries - Max Time to resolve 95% of all queries	<2	<2	<3	<0.5	<0.5	<0.5	<1	<2	<1	<1	<0.5	<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	0	<0.5	<0.5	<0.5	0	<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	12	10	10	10	10	10	10	10	10
System Maintenance - no. of sessions taken per system in the mont	0	2	1	1	0	0	1	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 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	2001/2 Jan	Feb	March	April	May	June	July	Aug	Sept	Oct	Nov	Dec
HPC Services Availability												
Availability in Core Time (% of time)	-0.039	-0.039	-0.058	-0.058	-0.039	-0.039	0.039	0.039	0.039	0.039	0.039	-0.058
Availability out of Core Time (% of time)	-0.039	0	0	0	0	0	0.039	-0.047	0	-0.039	-0.039	0.039
Number of Failures in month	-0.008	-0.008	-0.008	-0.008	0	0	0.008	0	0	0	0	0
Mean Time between failures in 52 week rolling period (hours)	-0.008	-0.008	-0.008	-0.008	-0.008	-0.008	0	0	0	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	0.031	0.031	0.016	0.031	0	0	-0.016	-0.016	-0.016	-0.016	-0.016	-0.016
Administrative Queries - Max Time to resolve 95% of all queries	0	0	0.016	-0.019	-0.019	-0.019	-0.016	0	-0.016	-0.016	-0.019	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	0	-0.002	-0.002	-0.002	0	-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	0.003	0	0	0	0	0	0	0	0
System Maintenance - no. of sessions taken per system in the mont	-0.004	0	-0.003	-0.003	-0.004	-0.004	-0.003	0	0	0	0	0
Monthly Total & overall Service Quality Rating for each period:	-0.05	-0.03	-0.04	-0.05	-0.06	-0.06	0.00	-0.03	-0.02	-0.04	-0.04	-0.04
Quarterly Service Credits:			-0.15			-0.08			-0.09			-0.12
Annual Service credit												-0.45

Table 3

1.2 No. Of Registered Users

The current position at the end of the quarter is that there are 554 registered users of the CSAR Service.

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

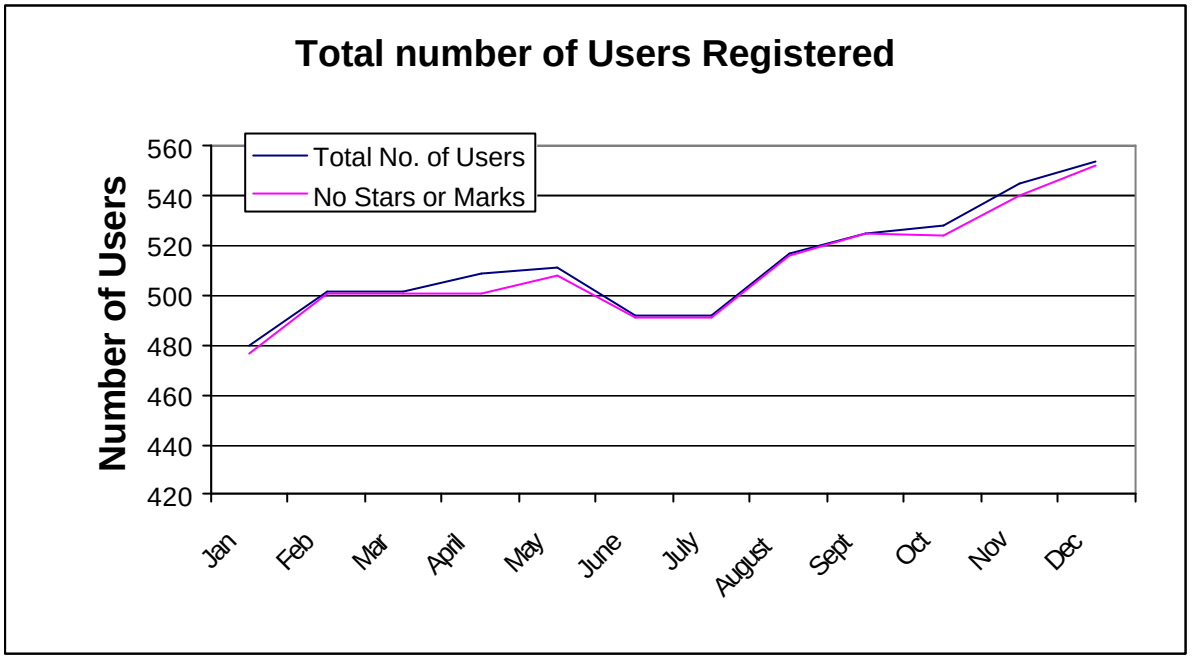

1.3 Service Quality Tokens

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

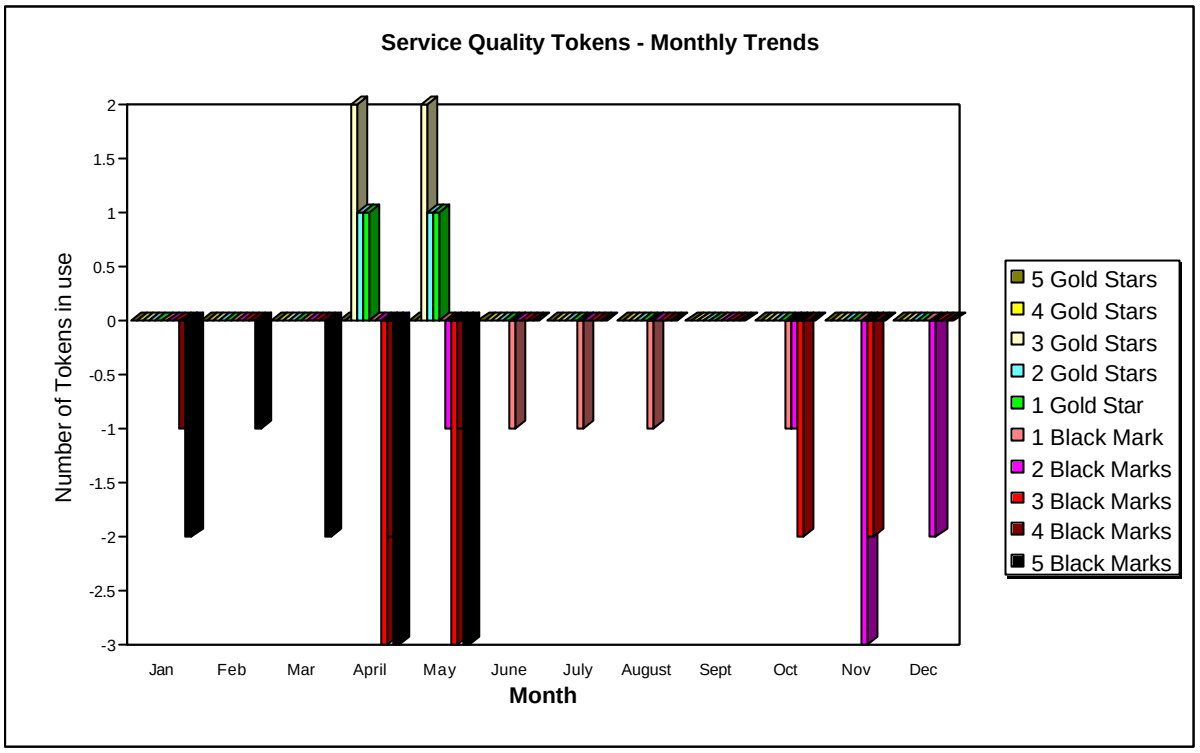

Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

There are, at the end of the quarter, four marks 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
2 Black Marks	cse085	22/11/01	Deterioration of interactive response on Fermat
2 Black Marks	cse085	09/11/01	Problems running jobs on Green

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.

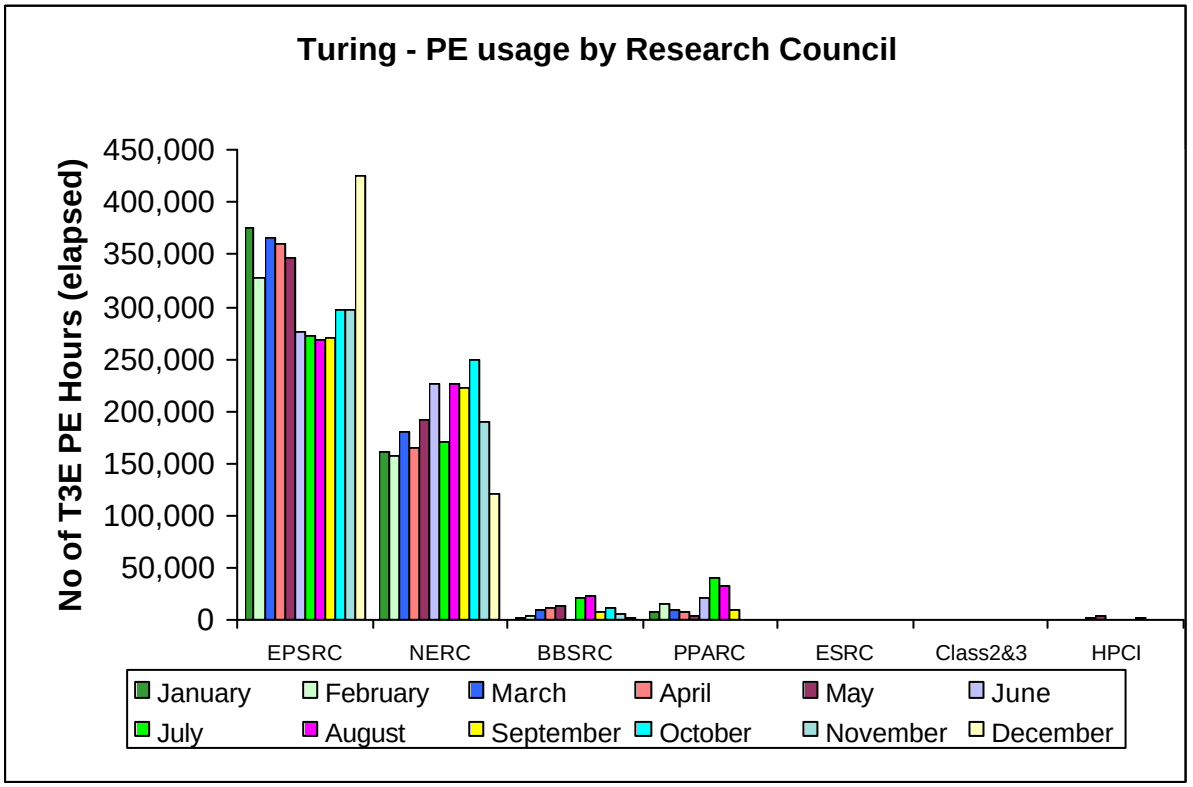

Turing PE usage is shown by Research Council during the last 12 months of service in the above chart.

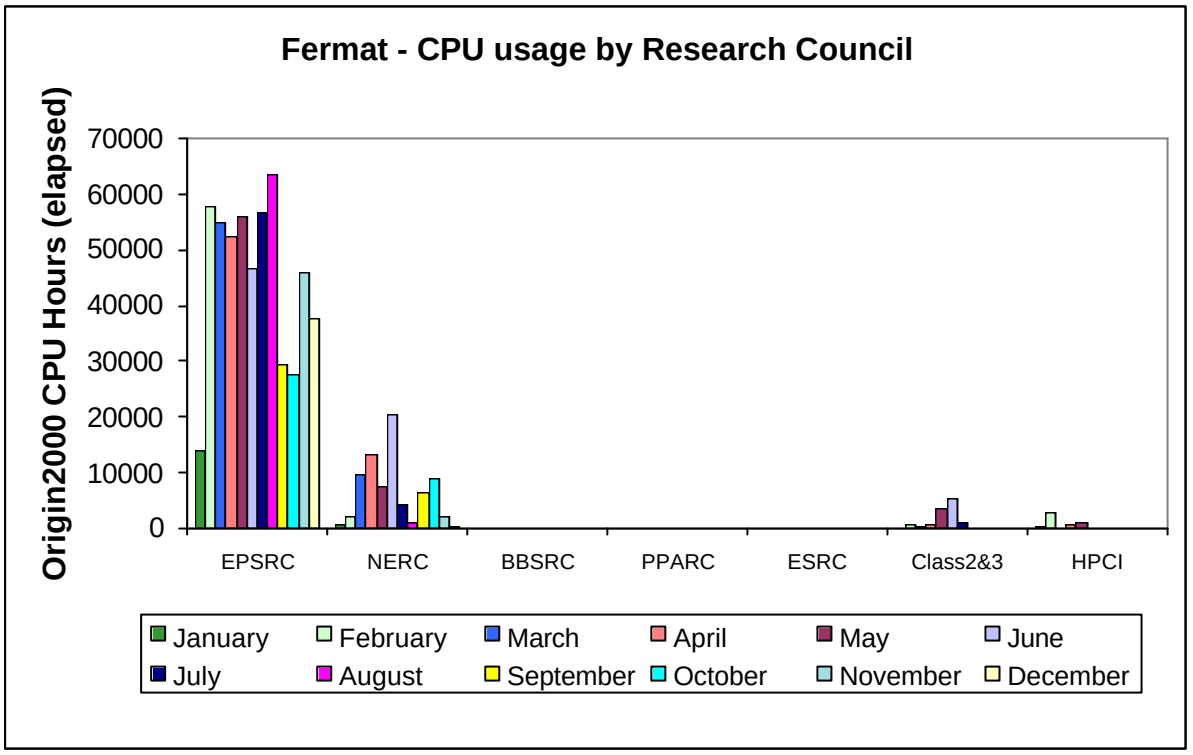

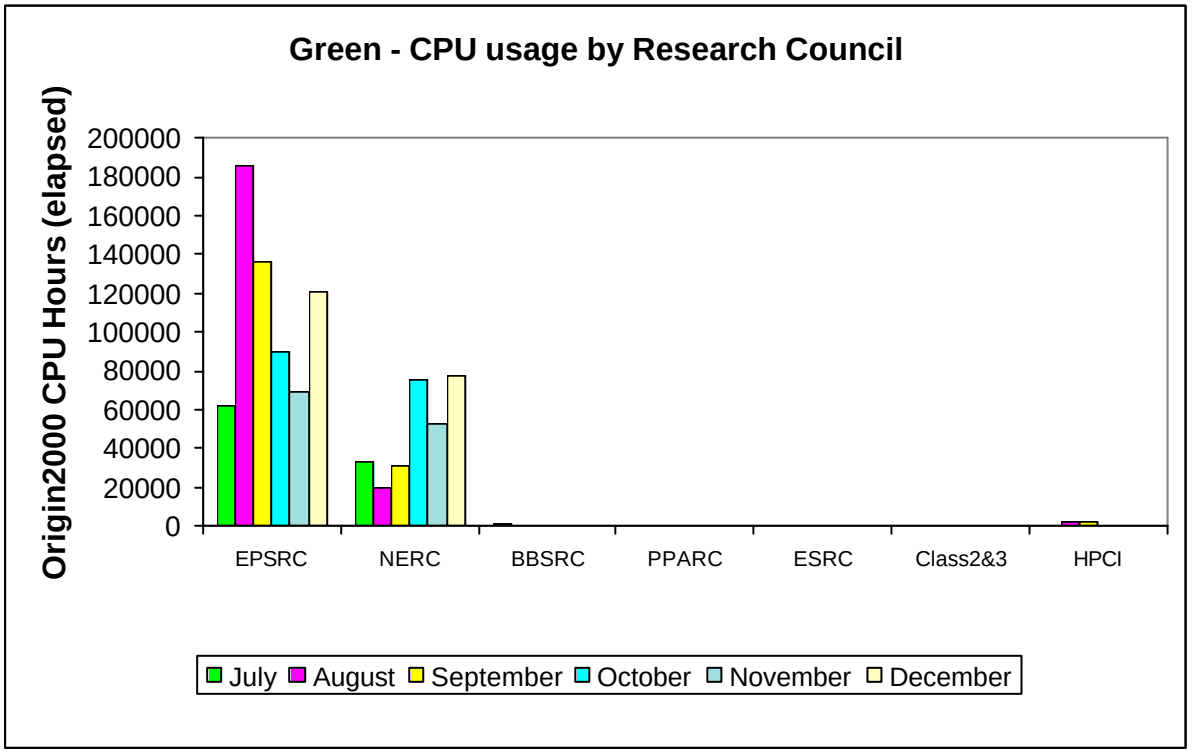

Usage of the two SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service in the above charts.

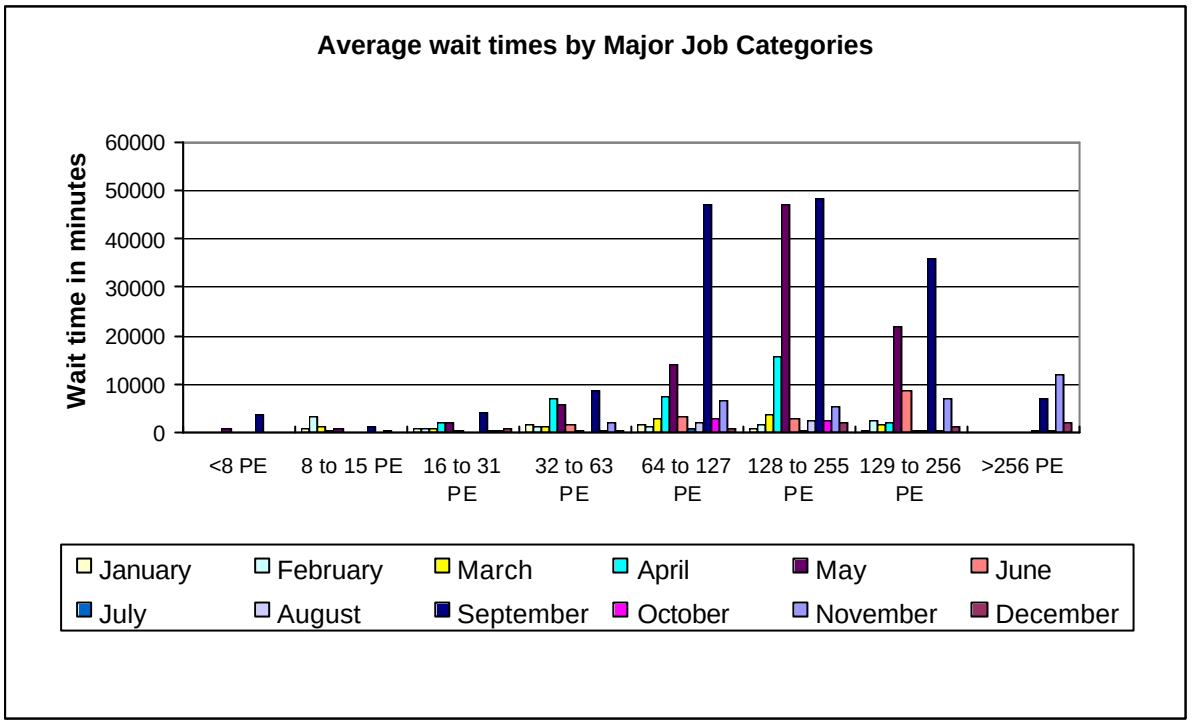

The above chart shows the wait time trend on the Turing system. The quarter has seen job wait times gradually decrease.

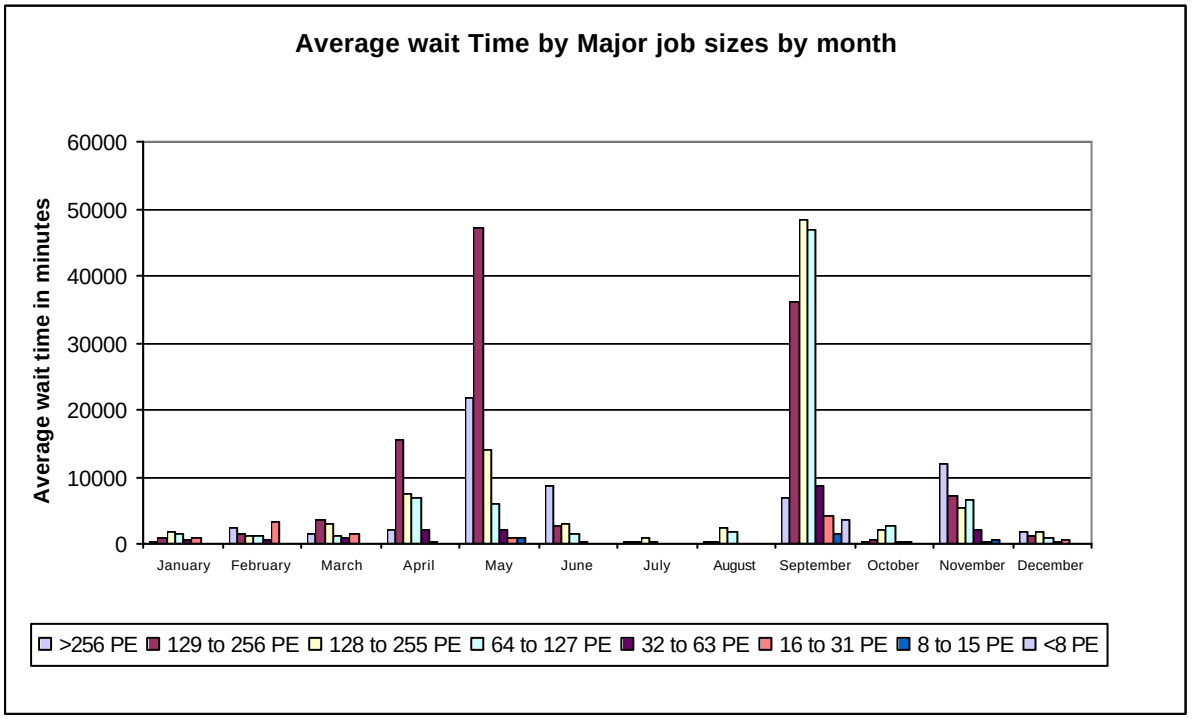

It can be seen from the above graph that the additional capacity added to the service, in the shape of the upgrade to Fermat and the addition of the Origin 3000 (Green), are now reducing the burden on the T3E (Turing) and reducing the queue wait times.

The next chart shows the percentage PE time utilisation by the major job categories on the Turing system for the 4th quarter 2001. The predominant job size is large, with the percentage of jobs greater than 64 PEs for the quarter being 83%.

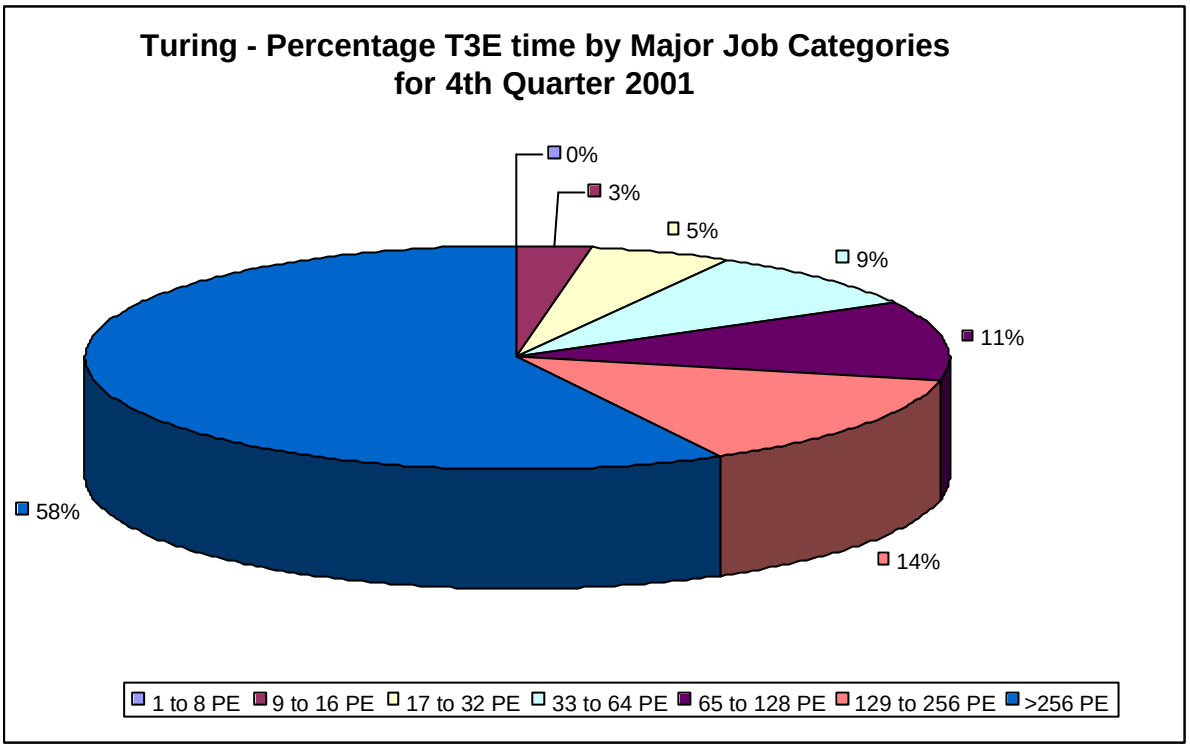

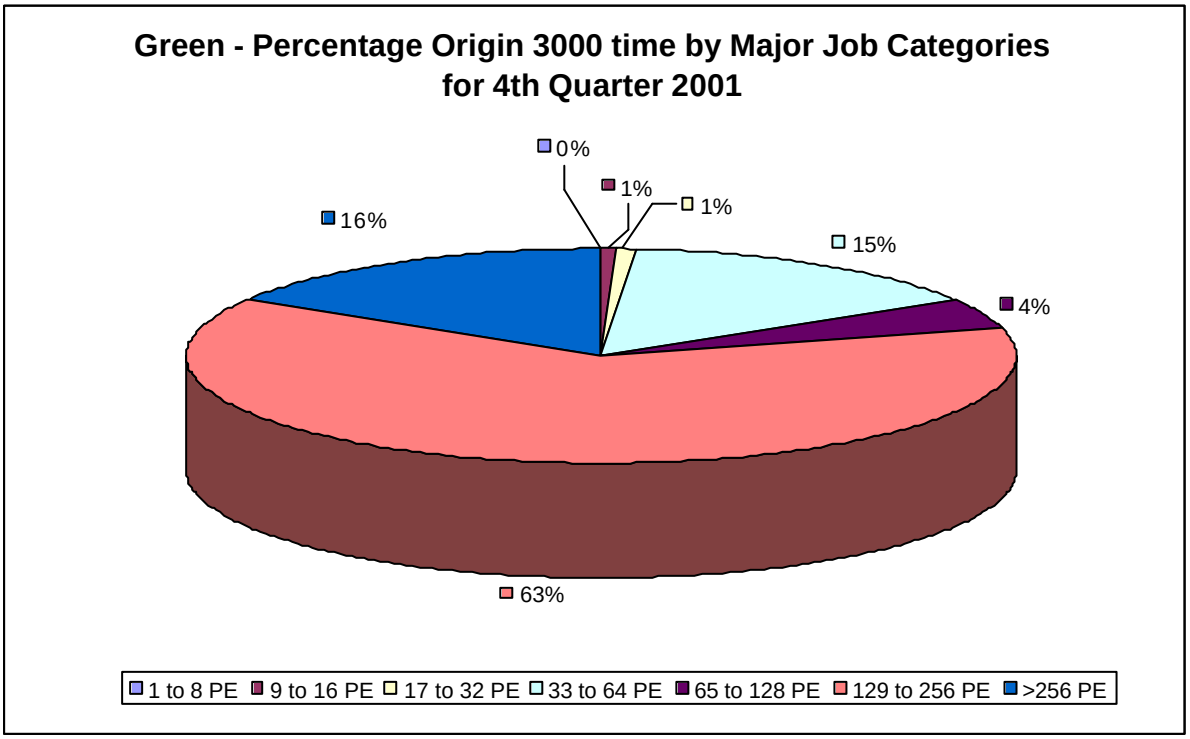

On Green, the percentage of CPU utilisation used by jobs requiring greater than 128 CPUs for the quarter was 79%.

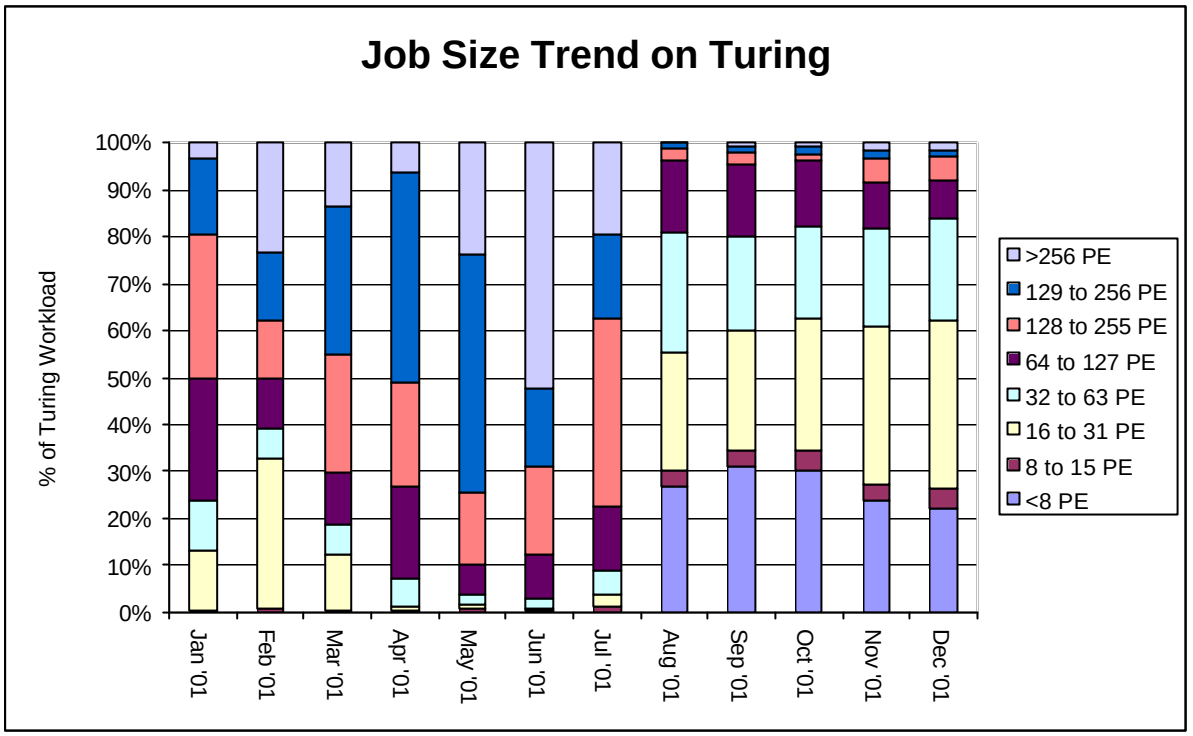

The above chart illustrates the effect that the introduction of the Fermat upgrade and the Origin 3000 (Green) have had on the job size profile on Turing. Many of the larger jobs have now moved to Green, leaving work of the smaller sizes on Turing.

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.

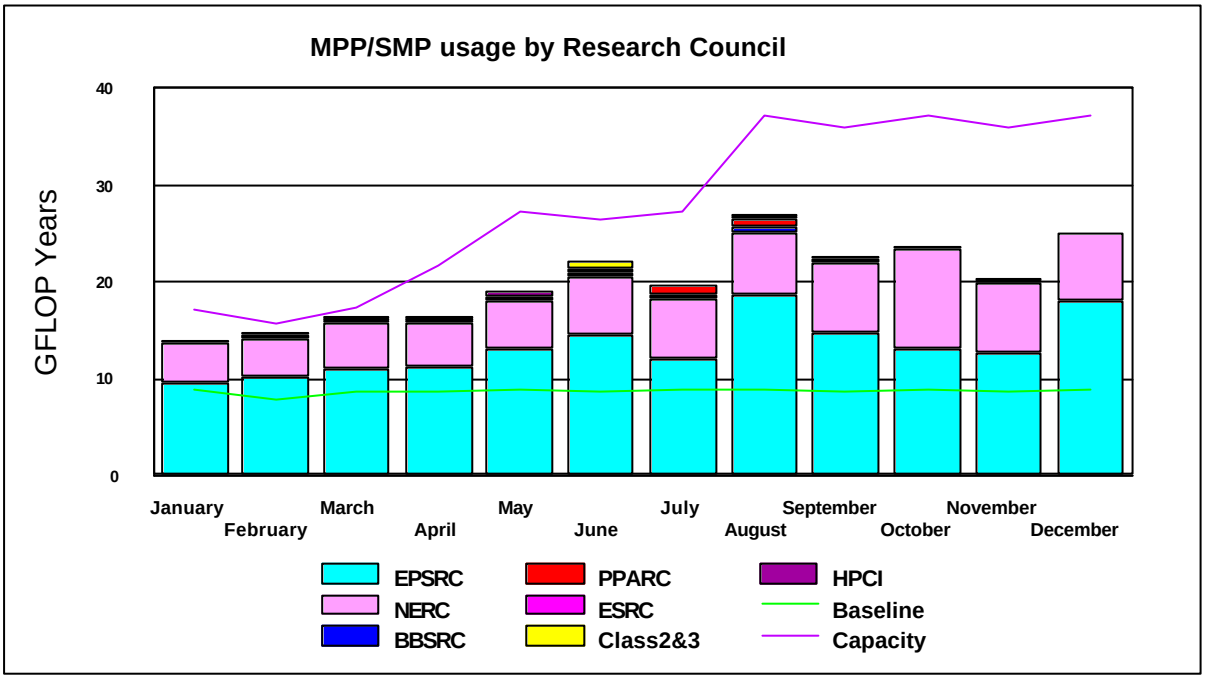

The graph below shows the same service utilisation by major system. .

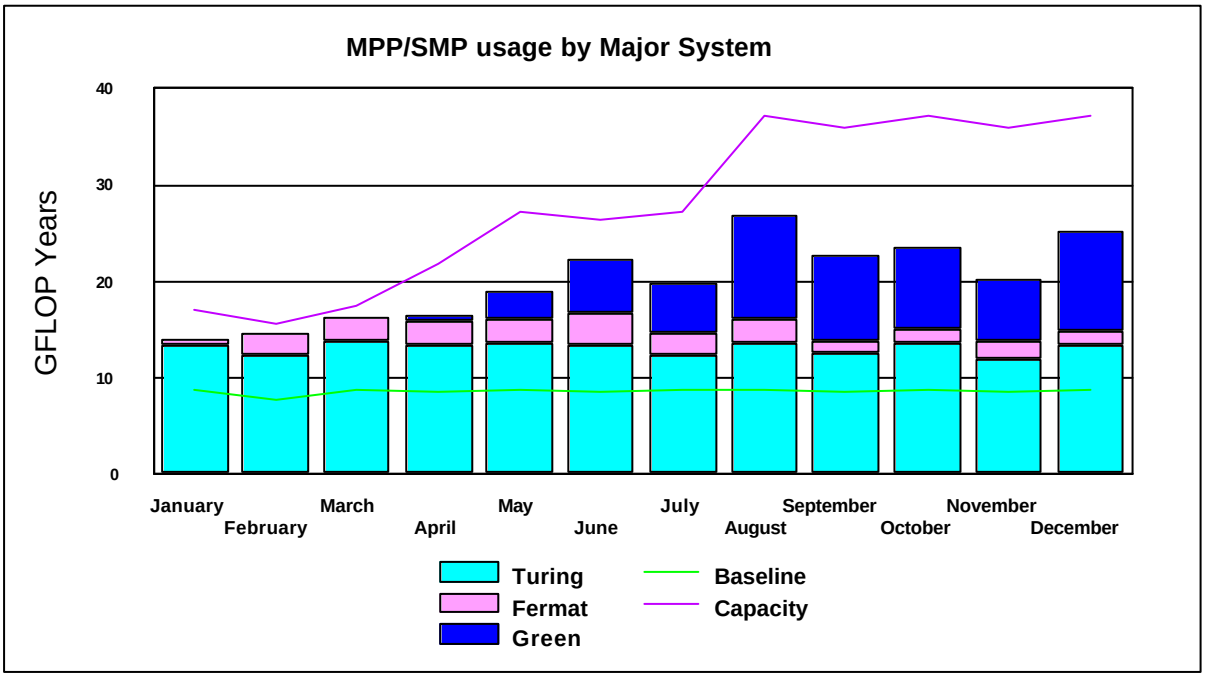

The next chart shows the historic interactive usage of the 'Baseline' Fermat system (equivalent to 16 @250Mhz CPU's).

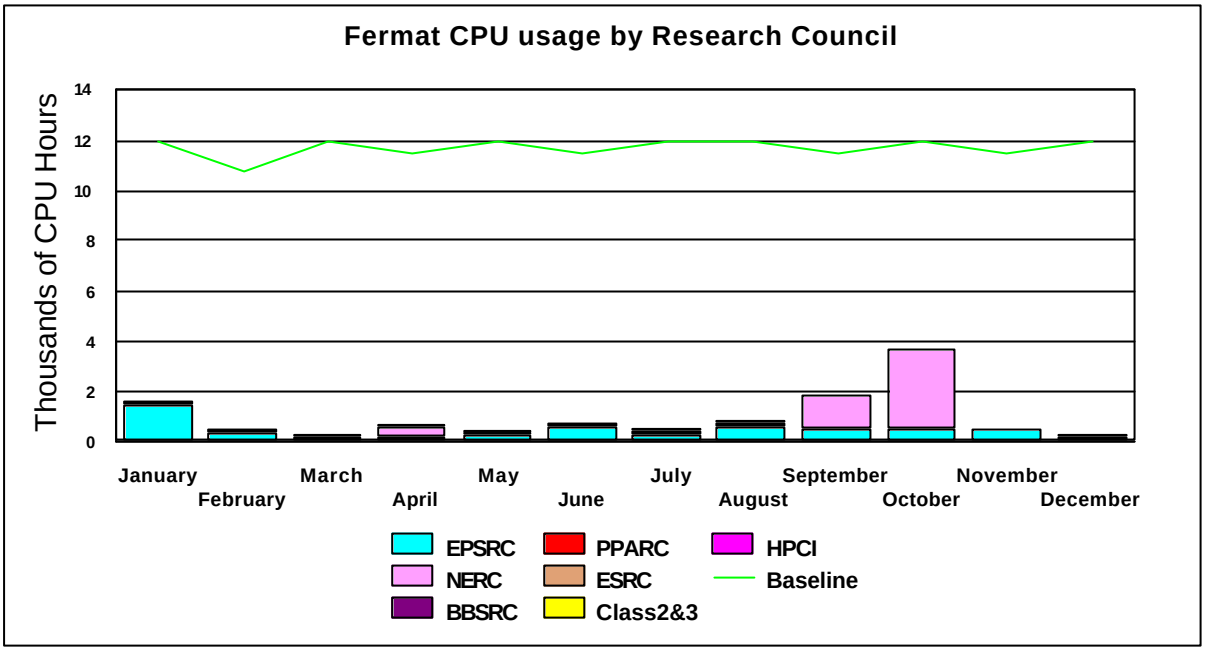

The next series of graphs illustrates the usage of the disk and HSM resources of the system.

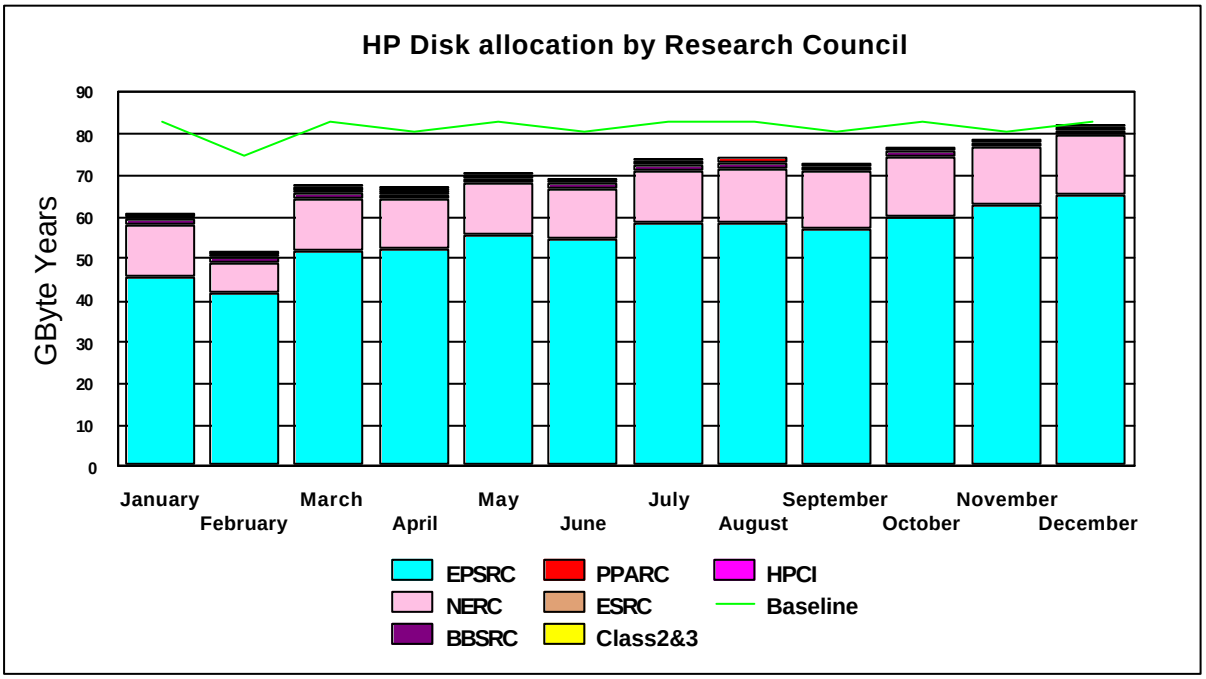

The preceding graph illustrates the historic allocation of the High Performance Disk on Turing, showing baseline disk fully utilised.

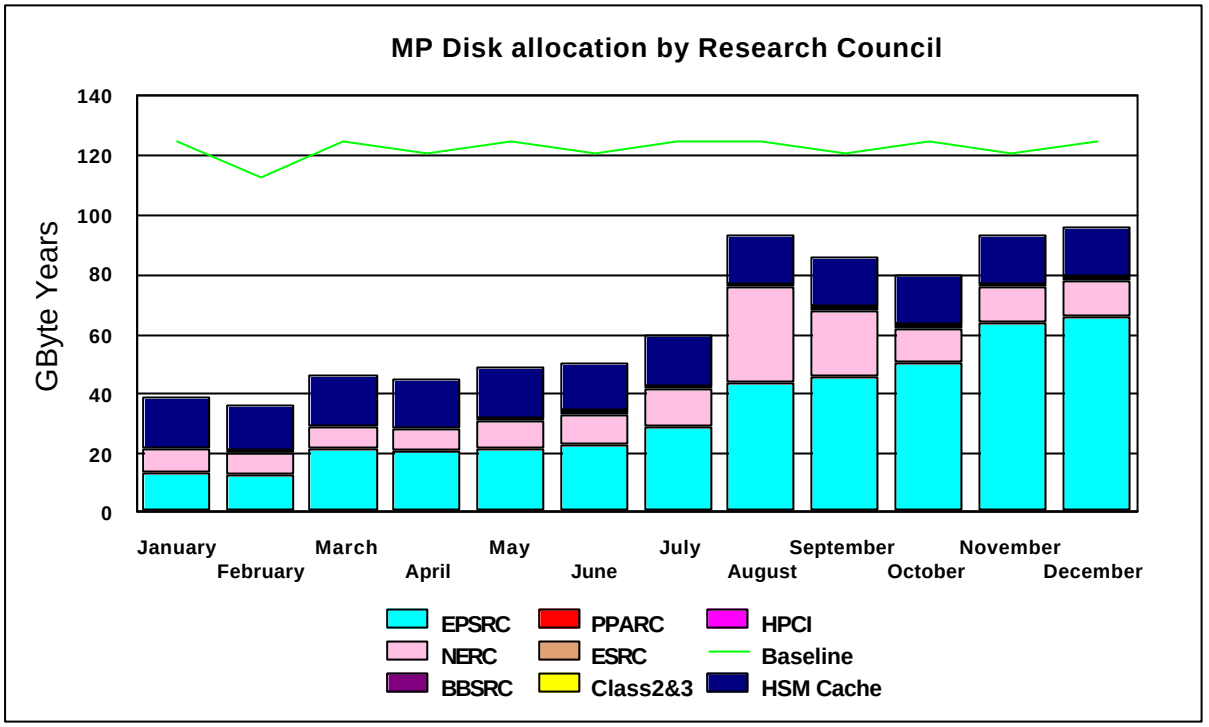

The graph above illustrates the historic allocation of the Medium Performance Disk on Fermat. From July 2000, 200 Gbytes has been used as a data cache for the enhanced HSM system.

The graph below shows the historic HSM usage by Research Council funded projects, which has now exceeded the overall Baseline of 16 Terabytes, and now totals 17.84 Terabytes. The primary usage is for NERC and, in total, this amounts to about 15.9 Terabytes of data storage at present.

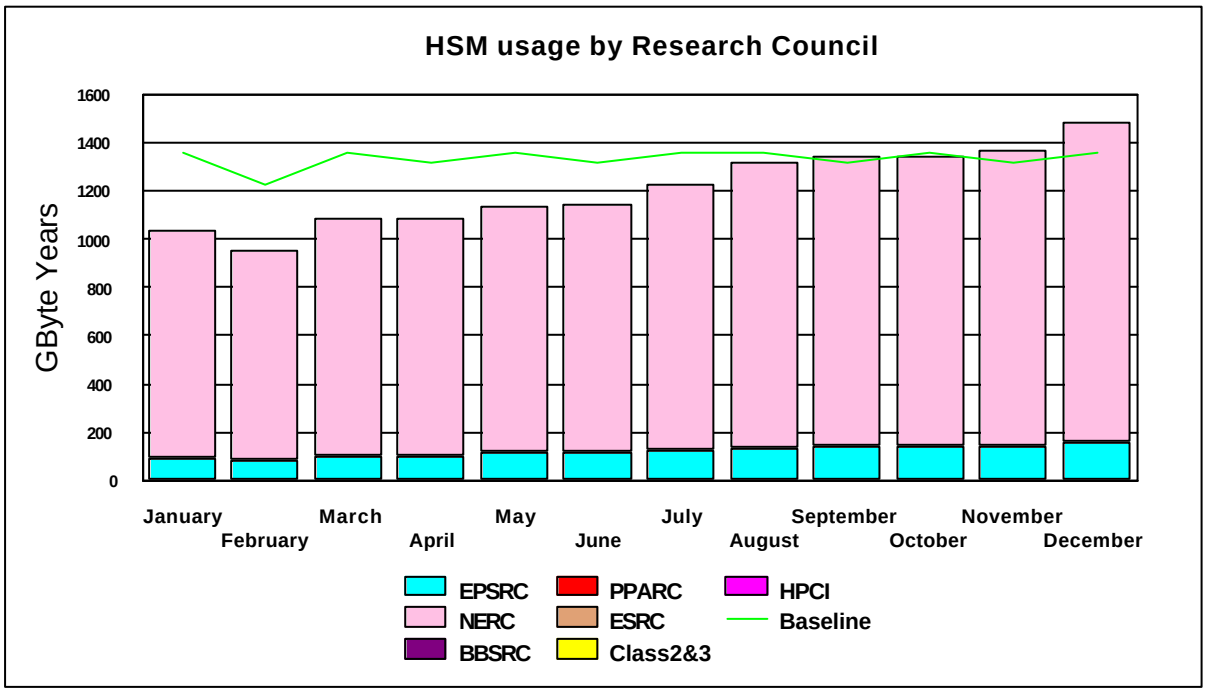

2.2.2 Fujitsu System Usage Graph

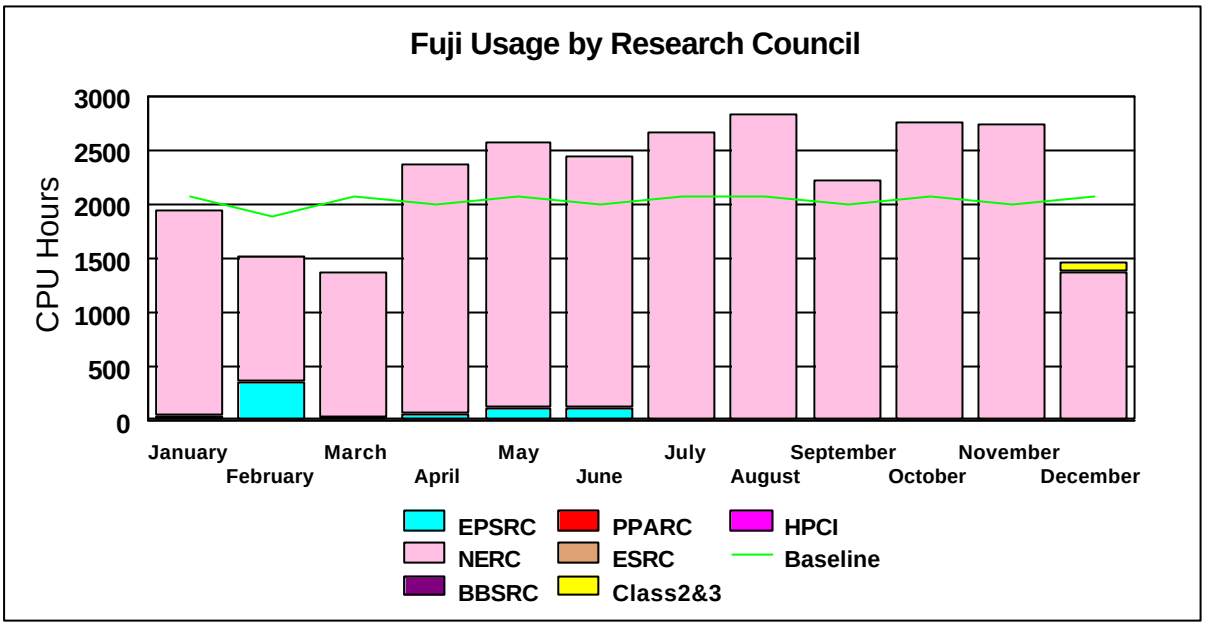

The above graph shows the current CPU usage on the Fujitsu VPP 300 NERC system based at the University of Manchester. The Fujitsu usage exceeded the baseline in two of the three months of the quarter. December also shows Class 2 usage by Vekstein.

2.2.3 Guest System Usage

The Compaq ES40 cluster still shows no guest usage, and the NEC SX4 is being withdrawn from service due to continued zero usage demand.

2.3 Service Status, Issues and Plans

Status

The service continues to be heavily used.

The percentage capability work on Turing and Green has been high for the quarter.

Green underwent acceptance testing in December, with a view to going into full production as the technology refresh system from January 1st 2002.

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 have substantially reduced the queuing times on Turing.

CfS still awaits a decision from NERC regarding the 12 Terabytes of additional data (the ECMWF dataset) that may require storage on the system.

Plans

The implementation of the proposed SAN has been delayed due to the proposed upgrade of Green to 2048 CPUs pending a decision by the Research Councils.

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 Publicity

Numerous posters and leaflets about the CSAR (and MRCCS and MVC) services have been produced for use at SC2001 and elsewhere.

3.2 CSAR User Survey 2001

The annual CSAR User Survey was carried out during December using an on-line form on the CSAR web pages. A report is now being prepared.

3.3 CSAR User Steering Group Meeting

The seventh 6-monthly meeting of the CSAR User Steering Group was held in late December. The meeting was held jointly at Imperial College and at the University of Manchester, using the Access Grids at both sites to allow full participation of all participants at both sites. The chair of the User Liaison Forum (Dr Lois Steenman-Clark) had previously requested feedback from all CSAR users and this was presented at the meeting. This was a particularly successful meeting - the twin site approach and the request for feedback just prior to the meeting were seen to be factors in this success.

4 Science Application Support Services

4.1 Training and Education

4.1.1 Autumn School

A one week workshop covering the NUMAflex architecture and other specific aspects of the Origin systems was given from 8th - 12th October 2001. Discussions with two of the participants have resulted in longer term optimization work with the corresponding consortia. One of the invited speakers was from the SARA Supercomputing Centre – we have now been invited to visit SARA.

4.1.2 Courses

In addition to the in-depth workshop on using the Origin systems, the following courses were given at Manchester:

- Moving from C to C++
- Introduction to AVS/Express
- Introduction to High Performance Computing
- Using the CSAR Service
- Shared Memory Parallelisation with OpenMP
- Introduction to MPI
- Advanced MPI
- Fortran 90

An MPI course was given at Jodrell Bank Observatory.

We have been invited to give some HPC training courses at a National HPC Centre in Sweden in February 2002. This is seen as a further opportunity to collaborate with other international HPC Centres.

The course schedule for next semester has been prepared: http://www.csar.cfs.ac.uk/using/courses/index.shtml.

4.2 Consortia Support

New work has started on producing a SHMEM alternative to some ScaLAPACK routines, used very heavily by one consortium. Preliminary work on some example code has shown a very significant performance gain. These modified routines will now be incorporated into the real code for a full evaluation. If appropriate the modified routines could be made available to any other users with similar requirements.

4.3 Software

The performance analysis tool VAMPIR and the parallel debugger Totalview have now been installed on *Fermat/Green.*

The molecular dynamics package NAMD, has been installed on Green for evaluation by CSAR staff. If deemed appropriate it will be released to the CSAR community.

5 Collaboration and Conferences

5.1 UKHEC Reports

A case study was performed on the visualisation of historic databases of sunspots – to aid computational scientists who are producing computational models of the sun. A UKHEC report is being prepared.

5.2 MRCCS Projects

5.3 EuroGrid

The key goal of the EUROGRID Project is to allow researchers to easily and securely access high-performance computing resources in a seamless fashion, whether these resources are located in their own department or on another continent. The project is building upon the UNICORE software system, developed by Pallas GmbH and is at the forefront of the new field of "GRID Computing". Manchester is specifically concerned with developing a generic resource broker that can match the requirements of a user's job with the resources that are available on the European HPC Grid. Additionally, Manchester is operating a Certification Authority for the project, and providing some HPC resources for test-bed activities.

The EUROGRID project came of age in December 2001 when the project was reviewed by the European Commission,. the review meeting was hosted by Manchester Computing at the University of Manchester "Chancellors" Conference Centre. Representatives attended from all eleven of the European organisations involved in the project, as well as the European reviewing team, led by Kyriakos Baxevanidis. The review itself went very well.

5.4 GRIP - GRid Interoperability Project

The first generation of Grid computing has created several interesting Grid technologies. However, as standardisation is always a step behind, these technologies can often be incompatible. The recently started Grid Interoperability Project, GRIP, will address this problem by developing techniques for integrating Globus and UNICORE, two of the most important and widely used GRID Computing solutions.

GRIP is a two year project, which started on January 2002, and is being funded under the European Commission's IST 5th Framework Programme. The other partners of the project are Pallas GmbH, Forschungszentrum Jülich

GmbH, DWD, FECIT, and the University of Southampton; Argonne National Laboratories USA, who developed Globus. The project will fund one person full-time at Manchester. Specifically, Manchester will attempt to enhance the EUROGRID resource broker so that it can broker for Globus resources as well as UNICORE resources.

5.2.3 Access Grid

Our Access Grid node is now fully operational - there have been numerous successful demonstrations. We participated in SC Global 2001, at which we hosted 2 events via the Access grid (on the subjects of Global Supercomputing and Solar-Terrestial Physics). We have also participated in events hosted by other remote sites. The CSAR User Steering Group meeting was held jointly at Imperial College and Manchester via the Access Grid. The MRCCS seminar series in 2002 will be held over the Access Grid. It is now used as a standard facility, not used purely to demonstrate the technology.

5.2.4 The Reality Grid

The major proposal for an e-Science pilot project, 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 high profile groups from different universities and has substantial industrial backing.

This will result in four new appointments for MC and two for CNC. The interviews for the MC positions will take place shortly.

5.3.5 E-Science

The e-Science North West (ESNW) has now started. This will provide a framework to support e-Science projects in the North West.

A new member of staff, Mike Jones, has been appointed to provide support in this area.

Mike Daw and Stephen Pickles are each leaders of a working group in the Engineering Task Force, established in the UK e-Science programme.

5.3.6 North West Centre for Advanced Virtual Prototyping

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

5.3 Events

Jo Leng attended the IEEE Visualization 2001 conference, October 21st – 26th October, San Diego.(report at http://ukhec.ac.uk/publications/trips/viz2001.html

Mike Jones and Jon MacLaren attended the third Global Grid Forum (GGF3) in Frascati in October 2001. During this meeting, it was agreed that Manchester would participate in the Applications Working Group Testbed Demo at SC2001 by contributing compute cycles from the CSAR machines. This highly successful demo, led by Ed Seidel's group at the Max Planck institute, showed the "Cactus Worm" migrating from resource to resource while simulating the collision of two black holes.

Neil Stringfellow and Robin Pinning attended the Daresbury Evaluation Workshop in November.

5.3.1 Supercomputing 2001

The University of Manchester presented the work of CSAR, MRCCS and MVC at a Research Booth at SC2001 at Denver from November 9-16. We were in the European village with close collaboration with our counterparts in European HPC, notably HLRS Stuttgart, FZ Juelich and the EuroGrid project.

Highlights of our demonstrations were a series of metacomputing experiments linking the US, Europe and Japan. Together with CSC we ensured that the CSAR service was also able to participate in the demonstrations by the Cactus group and we hosted a live demonstration of Cactus running across a European test-bed that included CSAR. This was presented by, Ed Sediel of the Albert Einstein Institute, based at Potsdam. The Cactus team won a special Gordon Bell award for their work in distributed computing and we are pleased to be able to participate in this ground-breaking work which showed that Globus is reliably deployed on the CSAR service.

5.3.2 SCGlobal

Manchester presented the UK contribution to the distributed conference SCGlobal that was held in parallel with SC2001. For the first time this allowed remote sites to participate in the conference and provide seminars, workshops and meetings via Access Grid technology. The Manchester Access Grid has provided a reliable service since August 2001 and was able to pass the very demanding tests required to host an SC Global Constellation site.

Four workshops were presented:

1. Global Metacomputing.
2. Tele-Radiology
3. Solar-Terrestrial Physics
4. Art on the Grid.

In addition to these events, we also presented to audiences at the Manchester Access Grid relays from other events at SC2001, notably the keynote speech by Craig Ventner of Celera Genomics. The experience gained in this work is now being disseminated to the other UK sites who are starting Access Grid nodes and Web pages are available via the MRCCS Web pages

http://www.man.ac.uk/mrccs

6 Added Value Services

6.1 Joint Projects

6.1.1 The VIP Laboratory

The multi-pipe edition of AVS/Express has undergone extensive development using the VIP Laboratory. In particular, support for the tracking of devices in a virtual environment have now been added to the software.

6.1.2 Distributed Visualization

A project to deliver real time volume rendered medical data to an operating theatre in a local hospital is progressing well. This work involves use of the OpenGL vizserver software. It will be presented at the Medecine meets Virtual Reality conference in California in January 2002.

6.2 International Conferences

Preparations are progressing well for the Cray User Group (CUG) conference to be held in Manchester in May 2002 (http://www.cugoffice.org/cugoffice_realms/CUG_S02/index.html).

6.3 Cray MTA Optimisation

Work continues on the optimisation of a genomics code for use on the Cray MTA.

6.4 Seminars

The seminar program for the Spring 2002 term is now available: www.man.ac.uk/mrccs/seminars/

Most of this series of seminars will be presented via the Access Grid, not only to other UK sites, but to anywhere worldwide with an Access Grid. It is hoped that this will encourage further contact and discussions with other HPC sites and users, both by reducing the time and effort required to attend and by increasing the number of participants.