

# **CSAR Service**

## **Consolidated Management Report**

### **1st Quarter 2002**

#### **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 547 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 86% 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
<b>HPC Services Availability</b>						
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
<b>Fujitsu Service Availability</b>						
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
<b>Help Desk</b>						
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
<b>Others</b>						
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											
	April	May	June	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March
<b>HPC Services Availability</b>												
Availability in Core Time (% of time)	100%	99.70%	99.70%	98.49%	98.49%	98.49%	98.60%	98.60%	100.00%	99.86%	99.73%	99.70%
Availability out of Core Time (% of time)	99.40	99.40	99.40	98.49%	100%	99.40	99.50%	99.50%	98.49%	99.89%	99.85%	99.97%
Number of Failures in month	1	3	3	4	2	2	2	2	4	2	1	2
Mean Time between failures in 52 week rolling period (hours)	674	584	584	438	398	365	365	365	337	350	324	313
<b>Fujitsu Service Availability</b>												
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%
<b>Help Desk</b>												
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	<2	<2	<1	<1	<1	<1	<1	<1	<1	<2	<1
Administrative Queries - Max Time to resolve 95% of all queries	<0.5	<0.5	<0.5	<1	<2	<1	<1	<0.5	<2	<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%
<b>Others</b>												
Normal Media Exchange Requests - average response time	0	<0.5	<0.5	<0.5	0	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5
New User Registration Time (working days)	0	0	0	0	0	0	0	0	0	0	0	0
Management Report Delivery Times (working days)	12	10	10	10	10	10	10	10	10	10	10	10
System Maintenance - no. of sessions taken per system in the month	1	0	0	1	2	2	2	2	2	2	2	2

**Table 2**

**Notes:**

- HPC Services Availability has been calculated using the following formulae, based on the relative NPB performance of Turing, Fermat and Green at installation:  

$$\text{Turing availability} \times 143 / (143 + 40 + 233) + [\text{Fermat availability} \times 40 / (143 + 40 + 233)] + \text{Green availability} \times 233 / (143 + 40 + 233)$$
- 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											
	April	May	June	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March
<b>HPC Services Availability</b>												
Availability in Core Time (% of time)	-0.058	-0.039	-0.039	0.039	0.039	0.039	0.039	0.039	-0.058	-0.039	-0.039	-0.039
Availability out of Core Time (% of time)	0	0	0	0.039	-0.047	0	-0.039	-0.039	0.039	-0.047	-0.047	-0.047
Number of Failures in month	-0.008	0	0	0.008	0	0	0	0	0.008	0	-0.008	0
Mean Time between failures in 52 week rolling period (hours)	-0.008	-0.008	-0.008	0	0	0	0	0	0	0	0	0
<b>Help Desk</b>												
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	0	-0.016	-0.016	-0.016	-0.016	-0.016	-0.016	-0.016	0	-0.016
Administrative Queries - Max Time to resolve 95% of all queries	-0.019	-0.019	-0.019	-0.016	0	-0.016	-0.016	-0.019	0	-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
<b>Others</b>												
Normal Media Exchange Requests - average response time	0	-0.002	-0.002	-0.002	0	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002
New User Registration Time (working days)	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019
Management Report Delivery Times (working days)	0.003	0	0	0	0	0	0	0	0	0	0	0
System Maintenance - no. of sessions taken per system in the month	-0.003	-0.004	-0.004	-0.003	0	0	0	0	0	0	0	0
Monthly Total & overall Service Quality Rating for each period:	-0.05	-0.06	-0.06	0.00	-0.03	-0.02	-0.04	-0.04	-0.04	-0.08	-0.08	-0.07
Quarterly Service Credits:	-0.16			-0.05			-0.11			-0.23		
Annual Service credit	-0.55											

**Table 3**

### 1.2 No. Of Registered Users


The current position at the end of the quarter is that there are 547 registered users of the CSAR Service.

It can be seen from the chart below that although the overall number of users dipped due to a number of projects finishing during February, user levels are once again on the increase.



### 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, three marks and one star 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 gold star	cse013	25/03/02	Good assistance with problem
3 black marks	csn001	15/02/02	Repeated problems accessing /hold

## 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;

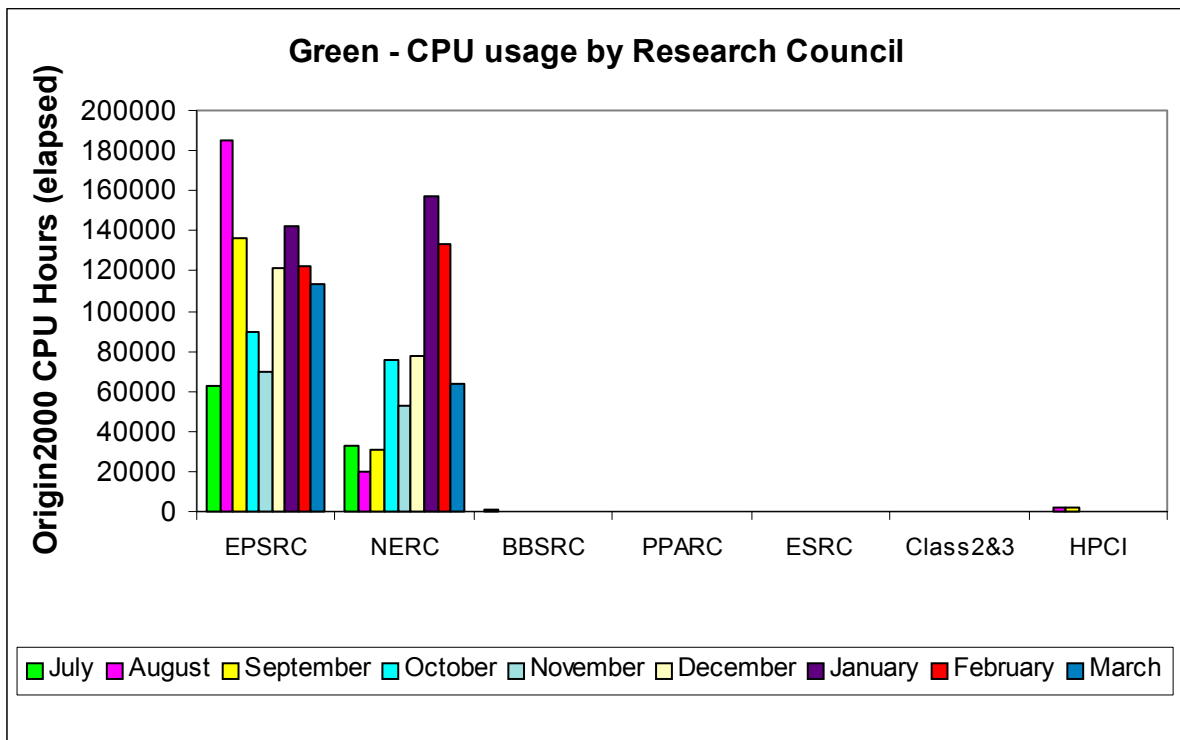

- MPP (T3E) Usage by month, showing usage each month of CPU (T3E PE Elapsed Hours), split by Research Council and giving the equivalent GFLOP-Years as per NPB. The Baseline Capacity is shown by an overlaid horizontal line.
- SMP (Origin) Usage by month, showing usage each month in CPU Hours, split by Research Council and giving the equivalent GFLOP-Years as per NPB. The Baseline Capacity is shown by an overlaid horizontal line.
- 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.
- 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.
- 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.




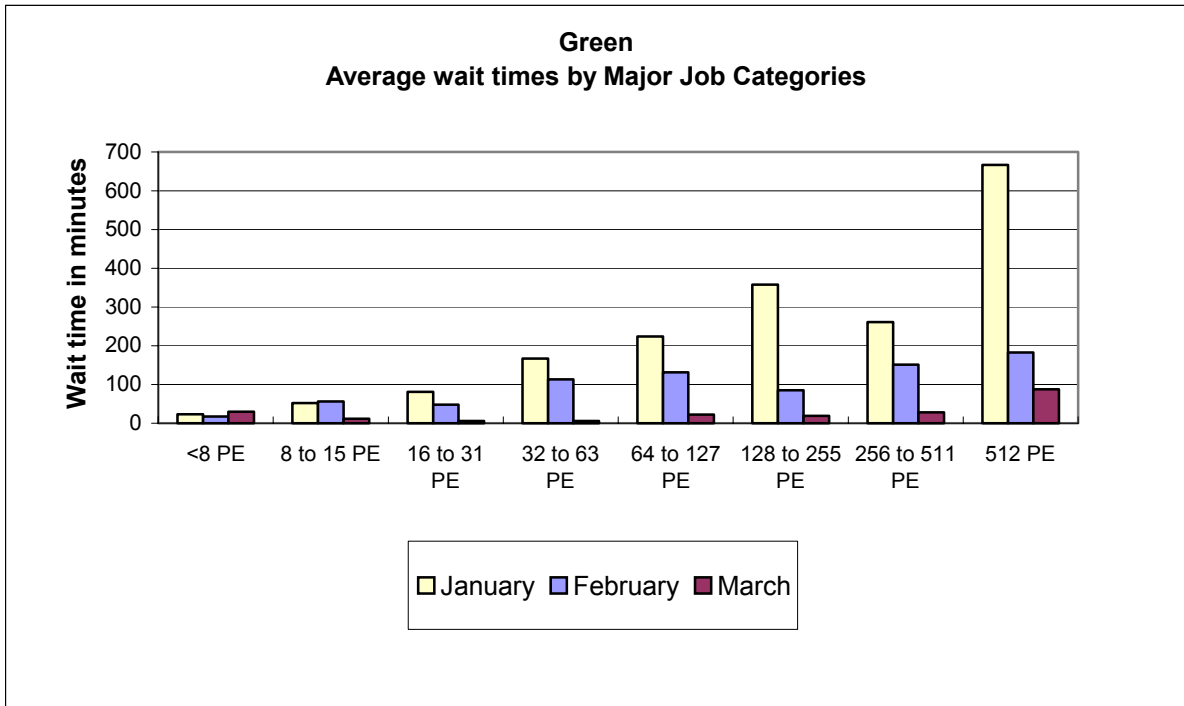
The above chart shows the wait time trend on the Turing system. The quarter has seen job wait times gradually decrease.




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




The next two charts show the percentage PE time utilisation by the major job categories on the Turing and Green systems for the 1<sup>st</sup> quarter 2002.



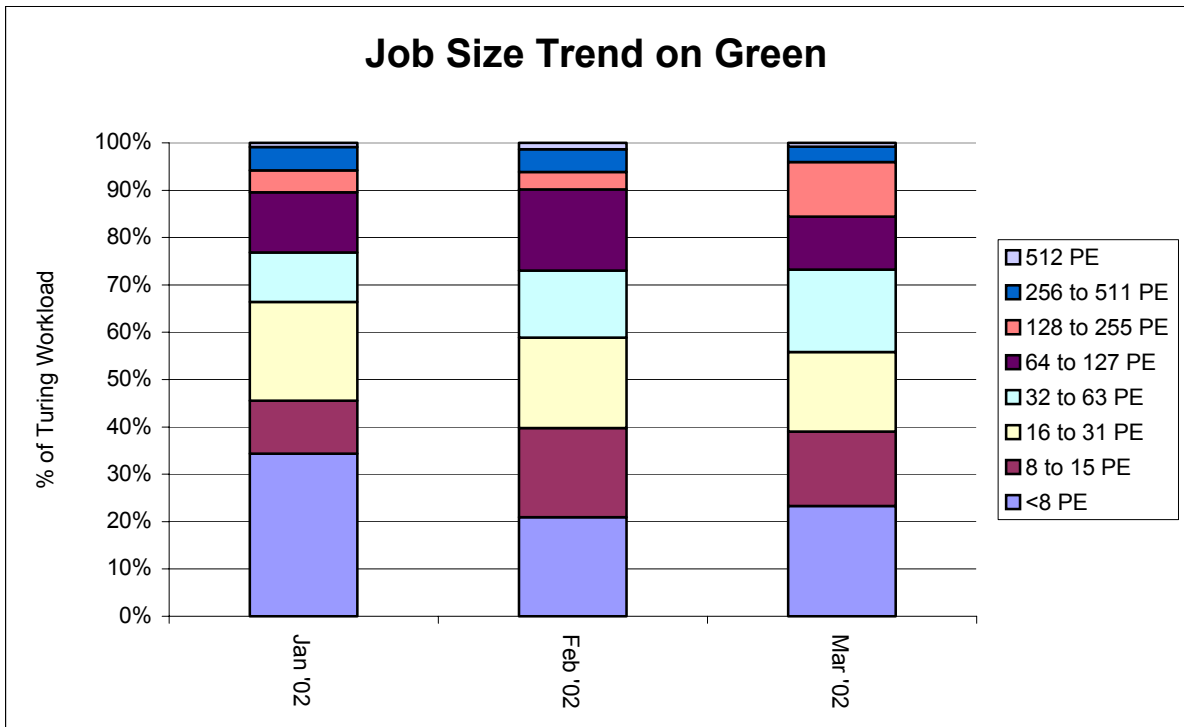
The predominant job size on Turing is large, with the percentage of jobs greater than 64 PEs for the quarter being 86%.



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, as displayed in the chart below, 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.




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



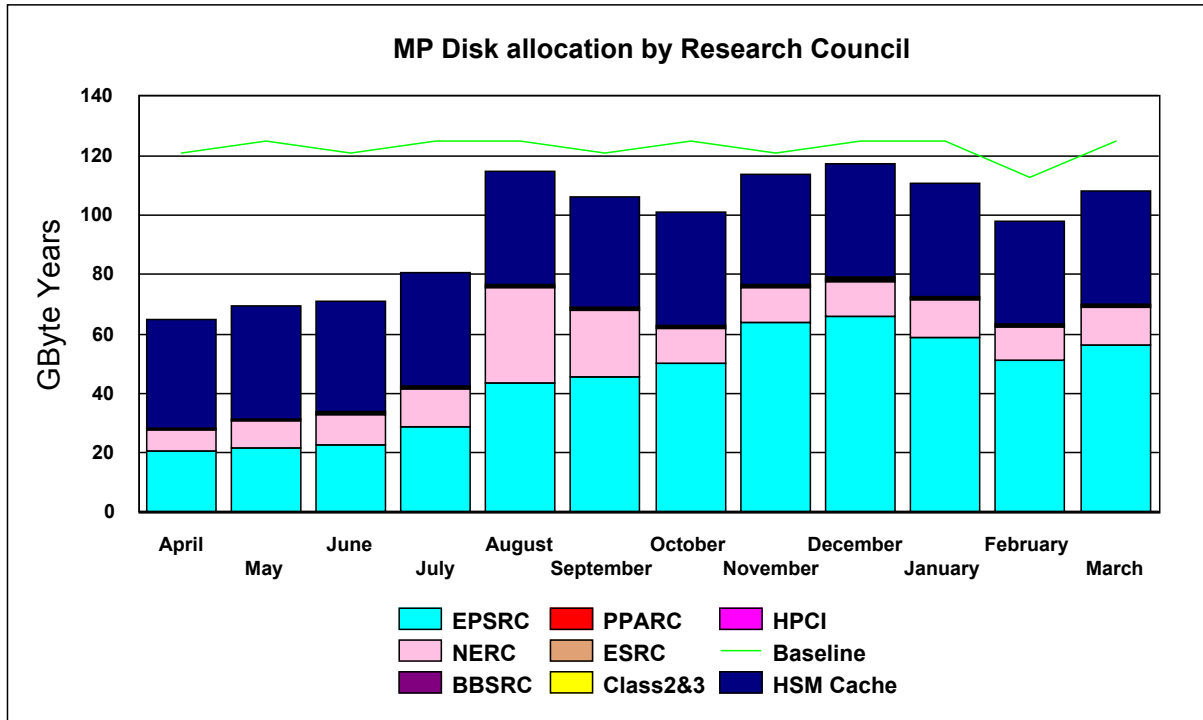
The next chart shows the historic interactive usage of the 'Baseline' Fermat system (equivalent to 16 @250Mhz CPU's).



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

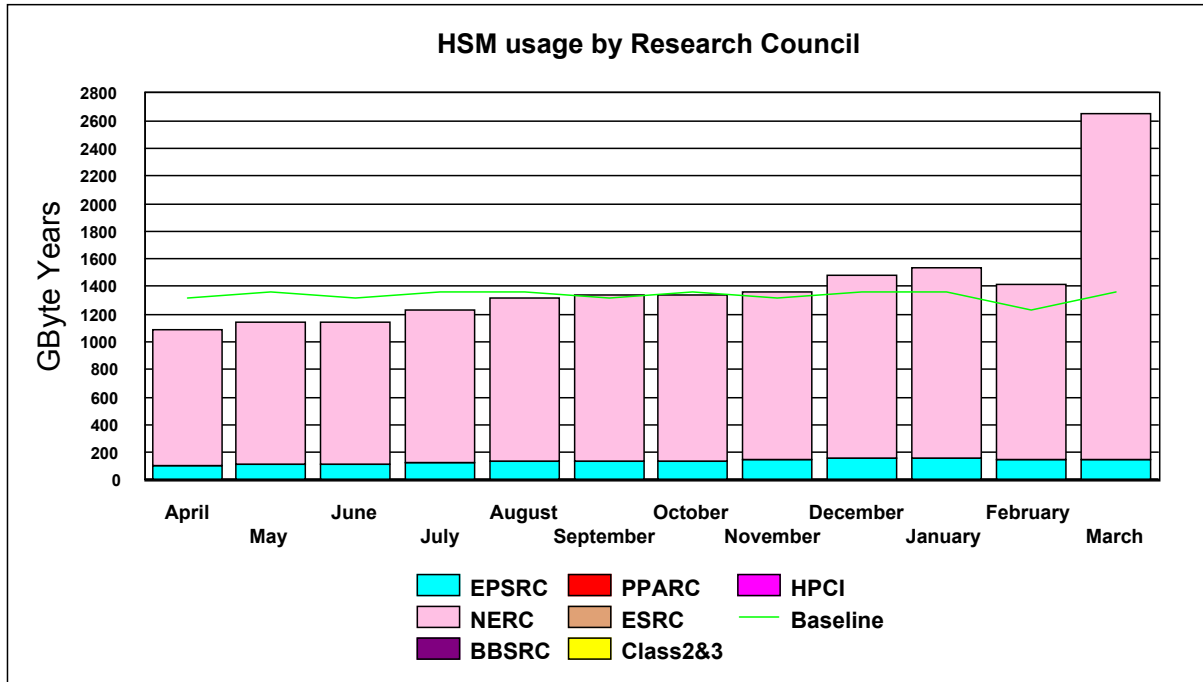


The preceding graph illustrates the historic allocation of the High Performance Disk on Turing, showing baseline disk fully utilised.




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 exceeded the overall Baseline of 16 Terabytes, and now totals 26.54 Terabytes. The primary usage is for NERC and, in total, this amounts to about 25.1 Terabytes of data storage at present. The large jump in HSM usage that is evident in the March figures is due to the correction of a 32 bit arithmetic error in an accounting program. The relevant people were informed of this.



## 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 all three months of the quarter.

## 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.

Following acceptance testing in December, Green has now gone into full production as the technology refresh system from January 1<sup>st</sup> 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 proposed SAN implementation is now in its final stages. The bulk of the equipment is now on site awaiting installation.

### **3 Project Management, Documentation and User Feedback**

This section covers aspects relating to the registration of projects and users, the management of projects and resources, topics associated with documentation and user feedback.

#### **3.1 Project Applications**

11 applications for new CSAR projects were received in the first quarter of 2002.

#### **3.2 New Projects**

5 new CSAR projects were started.

#### **3.3 Queries**

Overall, 209 queries relating to the CSAR service were received by the Helpdesk between 1<sup>st</sup> January and 31<sup>st</sup> March 2002.

#### **3.4 Generic Token Usage**

Implementation of daily checking of project token usage with warning messages being sent to PIs and the Research Councils has taken place with the ability to stop project usage at 100% of token allocation utilisation.

#### **3.5 Documentation**

The CSAR web-site has been updated to reflect the status of the Technology Refresh.

#### **3.6 Annual Report**

The CfS Annual Report is currently being prepared for publication.

#### **3.7 CSAR Focus**

The Spring edition of CSAR Focus is being produced.

### **4 Scientific Application Support Services**

#### **4.1 Training and Education**

The following courses were given at Manchester Computing:

- + Shared Memory Parallelisation with OpenMP
- + Introduction to MPI
- + Advanced MPI
- + Introduction to Access Grid / Node Operators' Tutorial
- + Visualization and Virtual Reality for Scientific Applications

Developments are being made to the Advanced MPI course, to include one-sided communications..

In February courses in MPI and OpenMP were given at the National Supercomputing Centre in Linköping, Sweden. The two day course was well attended and subsequent feedback has been very positive

In addition, an "Introduction to HPC and the CSAR Service" course was given to British Antarctic Survey staff (in Cambridge) in March.

#### **4.2 Consortia Support**



Work on the shmem alternative to ScaLAPACK continues. The original work for a specific consortium proved to have a significant performance gain and should assist the consortium in producing capability codes. The results will be presented at CUG 2002 in Manchester. The work continues in relation to the wider CSAR user community, though its reliance upon ScaLAPACK or similar routines must first be surveyed.

Support was given to a variety of consortia with respect to improving the I/O performance of jobs run on Green. As a result of this activity, performance problems with one particular code were detected. The cause is being investigated by SGI, but a workaround has been provided for the consortium concerned.

Debugging and optimisation is being carried out on the Large Eddy Simulations (LES) codes.

Optimisation is being performed on behalf of the OCCAM consortium.

Some benchmarking has taken place for the computational chemistry code, CASTEP.

## **5 Collaboration and Conferences**

### **5.1 UKHEC Reports**

The case study report "Visualization of Historic Data Bases of Sunspot Observations and Solar Dynamo Simulation" has been completed and posted on the UKHEC web site at <http://www.ukhec.ac.uk/publications/>

### **5.2 MRCCS Projects**

#### **5.2.1 Reality Grid**

Andrew Porter joined to provide programming support on VIPAR. Robin Pinning moved from CSAR applications support to work on grid middleware developments.

#### **5.2.2 E-Science**

CSAR continues to participate in the Engineering Task Force, whose mission is to realise a production UK e-Science Grid. Michael Daw leads the AccessGrid working group and Stephen Pickles leads the Grid User Management working group.

#### **5.2.3 Access Grid**

The Manchester Access Grid is now established as a production facility. An average of two meetings a week are being held which involve the Manchester node. These meetings include the MRCCS Seminar Series, e-Science Engineering Task Force meetings, meetings for a number of e-Science projects and UK Test Cruises. A member of the technical support staff (Andrea Jackson) is now in charge of the day-to-day running of the Manchester node and is node operator for many events.

#### **5.2.4 SAMD Project**

Kevin Roy has been involved in the setting up of the SAMD (Seamless Access to Multiple Datasets) project.

#### **5.2.5 Eurogrid**

The Eurogrid project has continued on schedule, with a project meeting in Warsaw in March. Just prior to the meeting, on the 1st of February, Manchester commenced the operation of the Certificate Authority for the project. Since then, certificates have been issued to over thirty users at sites across Europe. Manchester's next task in the project, a software deliverable, is to write a Resource Broker, due in June 2002 - this work is now well underway.

#### **5.2.6 GRIP**

GRIP, the Grid Interoperability Project, has now started and is proceeding well. Manchester is coordinating publicity for this project, which includes running the website. Since the initial meeting, there has been a project

meeting in Warsaw, and a further technical meeting held on the AccessGrid, which FZ-Juelich (Germany), Manchester, Southampton, and Argonne National Laboratories (USA) attended.

Kaukab Jaffri moved from the CSAR Frontline service to work on this project.

### **5.2.7 International Collaboration**

Mike Pettipher and Stephen Pickles visited SARA in the Netherlands to discuss issues relating to the provision of HPC services based on large Origin systems and to explore the opportunities for collaborative ventures..

## **5.3 Events**

John Brooke, Jon MacLaren and Stephen Pickles attended the fourth Global Grid Forum (GGF4) in Toronto and are active in several areas, including Accounting, Scheduling and Grid User Services. This meeting saw a strengthening of Manchester's activities in Global Grid Forum. Jon MacLaren led a discussion on Grid Resource Economies within the Accounting Group, and has become one of the co-chairs of the proposed Grid Resource Management Working Group (Scheduling Area), which is looking at developing standard protocols for advanced reservation.

Jon MacLaren attended the UK Grid Portals and Grid Services Workshop, at the new National e-Science Institute, in January 2002. He talked about portal development using UNICORE software, which is used in the Eurogrid and GRIP projects.

Michael Daw attended the Access Grid Retreat 2002, held in San Diego on 4-5 March. This event was held to discuss future directions of the Access Grid.

## **6 Added Value Services**

### **6.1 The VIP Laboratory**

Work is underway on developing an implementation of the multi-pipe edition of AVS/Express for use on PC graphics clusters.

There have been visits to the VIP Laboratory by the British Computer Society and the Institute of Electrical Engineers.

### **6.2 Distributed Visualization**

A recent edition of "Tomorrow's World" featured a world first, as a revolutionary new method of using 3D computer graphics in the operating theatre was used by a surgeon for the first time. This development was the result of a collaborative venture between Manchester Visualization Centre and Professor Rory McCloy in the Department of Surgery.

The system uses computer graphics technology to improve the accuracy and effectiveness of cancer surgery by creating 3D images that surgeons can consult during surgery. Professor McCloy used the new system for the first time during an operation carried out on a woman, with a suspected tumour in her pancreas. The new system meant that Professor McCloy was able to refer to a six-foot wide 3D visualisation of the patient's pancreas and surrounding anatomy, projected on to the wall of the operating theatre.

### **6.3 International Conferences**

Preparations are progressing well for the Cray User Group (CUG) conference which is being hosted by the University of Manchester from 20<sup>th</sup> – 24<sup>th</sup> May ([http://www.cugoffice.org/cugoffice\\_realms/CUG\\_S02/index.html](http://www.cugoffice.org/cugoffice_realms/CUG_S02/index.html))

Presentations and exhibitions of CSAR, MRCCS and MVC work are being prepared for the International Supercomputer Conference Europe 2002 in Heidelberg in June (<http://www.supercomp.de>), and for Supercomputing 2002 in Baltimore in November (<http://sc-2002.org>).

#### **6.4 Cray MTA Optimisation**

Work continues on the optimisation of a genomics code for use on the Cray MTA. The results will be presented at the CUG conference in May.

#### **6.5 Seminars**

Most of the Spring programme of seminars was held over the Access Grid. This enabled audiences from remote sites to attend (including Cineca Italy, Montana and Cambridge) and also for speakers to give their talks from remote sites (David Walker from Cardiff and David Harmin from Kentucky).

The Summer programme of events will be available shortly and will continue this trend.