

CSAR Service

Consolidated Management Report

1st Quarter 2003

Management Summary

This is the consolidated Management report for the first quarter 2003 of the CSAR HPC facility for UK Academia and Industry, which enables World-Class research and development.

The number of users has grown to a total of 539 to date.

The workload on both the Origin 3000 (Green) and the Cray T3E (Turing) has been fairly evenly spread across the mid- to high-end ranges of PEs during the first quarter of this year, with the T3E continuing to run to almost full capacity.

CSAR has been granted an 18 month extension of service contract until June 30th 2006. With this extension CfS is implementing a further technology refresh which will introduce a 256 processor Itanium-2 (Madison) based SGI Altix by end September 2003.

Interactive provision on Fermat was removed at the end of February, with Wren taking over as the primary interactive system. Fermat has now become a dedicated batch system.

The Fujitsu VPP 300/8 system Fuji was removed at the end of March in accordance with the wishes of NERC.

LSF, with CPUsets, is now in full production usage on Fermat and Green, with usage of these systems growing steadily.

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 for the 1st Quarter 2003 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	2002/3											
	April	May	June	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March
HPC Services Availability												
Availability in Core Time (% of time)	96.17%	96.08%	97.66%	99.2%	99.75%	98.75%	99.77%	99.25%	99.21%	99.46%	99.73%	100%
Availability out of Core Time (% of time)	97.75%	99.90%	99%	100%	100%	99.42%	99.52%	99.57%	100%	99.89%	100.00%	99.81%
Number of Failures in month	2	1	4	0	1	2	1	1	0	3	1	1
Mean Time between failures in 52 week rolling period (hours)	302	324	313	365	381	381	398	417	515	487	487	515
Fujitsu Service Availability												
Availability in Core Time (% of time)	96.89%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Availability out of Core Time (% of time)	98.92%	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	<2	<5	<2	<2	<1	<2	<2	<2	<0.5	<1	<2
Administrative Queries - Max Time to resolve 95% of all queries	<2	<3	<5	<2	<0.5	<2	<0.5	<0.5	<0.5	<1	<0.5	<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	<0.5	<0.5	<0.5	<0.5	<0.5	<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)	10	10	10	10	10	10	10	10	10	10	10	10
System Maintenance - no. of sessions taken per system in the month	2	2	2	2	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	2002/3											
	April	May	June	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	March
HPC Services Availability												
Availability in Core Time (% of time)	0.078	0.078	0.078	0	-0.039	0.039	-0.039	0	0	0	-0.039	-0.058
Availability out of Core Time (% of time)	0.039	-0.047	0.000	-0.047	-0.047	0	-0.039	-0.039	-0.047	-0.047	-0.047	-0.047
Number of Failures in month	0	-0.008	0.000	-0.009	-0.008	0	-0.008	-0.008	-0.009	0	-0.008	-0.008
Mean Time between failures in 52 week rolling period (hours)	0	0	0	0	0	0	0	0	-0.008	0	0	-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	-0.016	0	0.031	0	0	-0.016	0	0	0	-0.019	-0.016	0
Administrative Queries - Max Time to resolve 95% of all queries	0	0.016	0.031	0	-0.019	0	-0.019	-0.019	-0.019	-0.016	-0.019	-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	-0.002	-0.002	-0.002	-0.002	-0.002	-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	0	0	0	0	0	0	0	0	0	0	0
System Maintenance - no. of sessions taken per system in the month	0	0	0	0	0	0	0	0	0	0	0	0
Monthly Total & overall Service Quality Rating for each period:	0.03	0.00	0.05	-0.05	-0.08	-0.01	-0.07	-0.05	-0.06	-0.06	-0.09	-0.09
Quarterly Service Credits:	0.08			-0.14			-0.19			-0.24		

Table 3

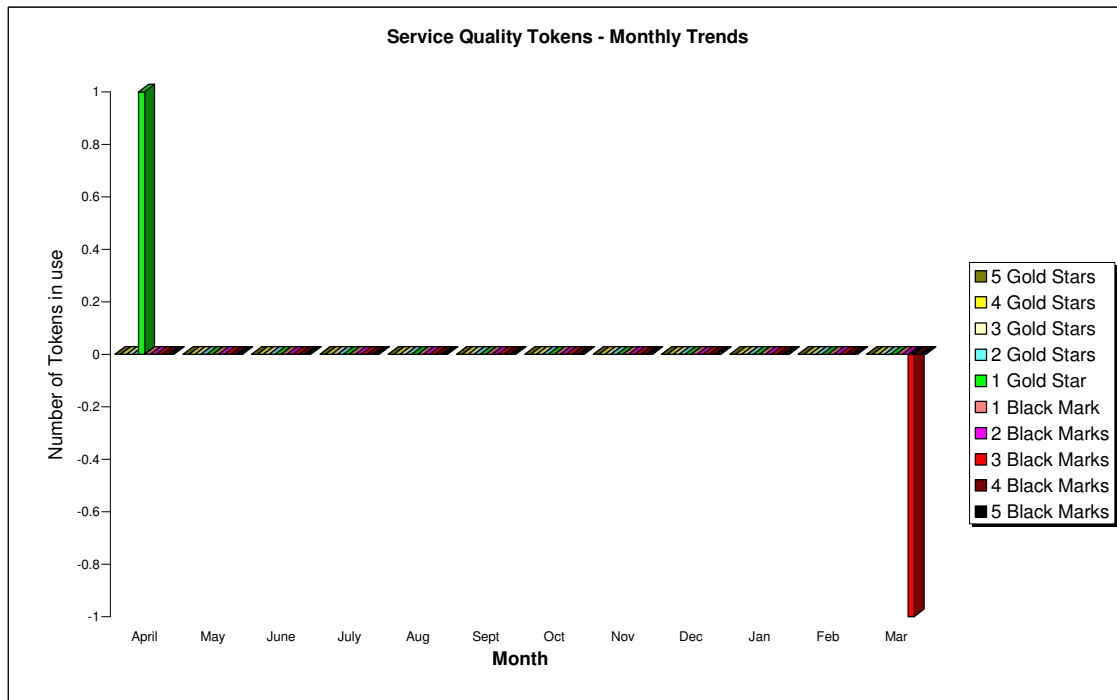
1.2 No. of Registered Users

The current position at the end of the quarter is that there are 539 registered users of the CSAR Service.



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.

At the end of the quarter there were three 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
3 black marks	csn001	27/03/03	Problems with access to /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:


- 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.

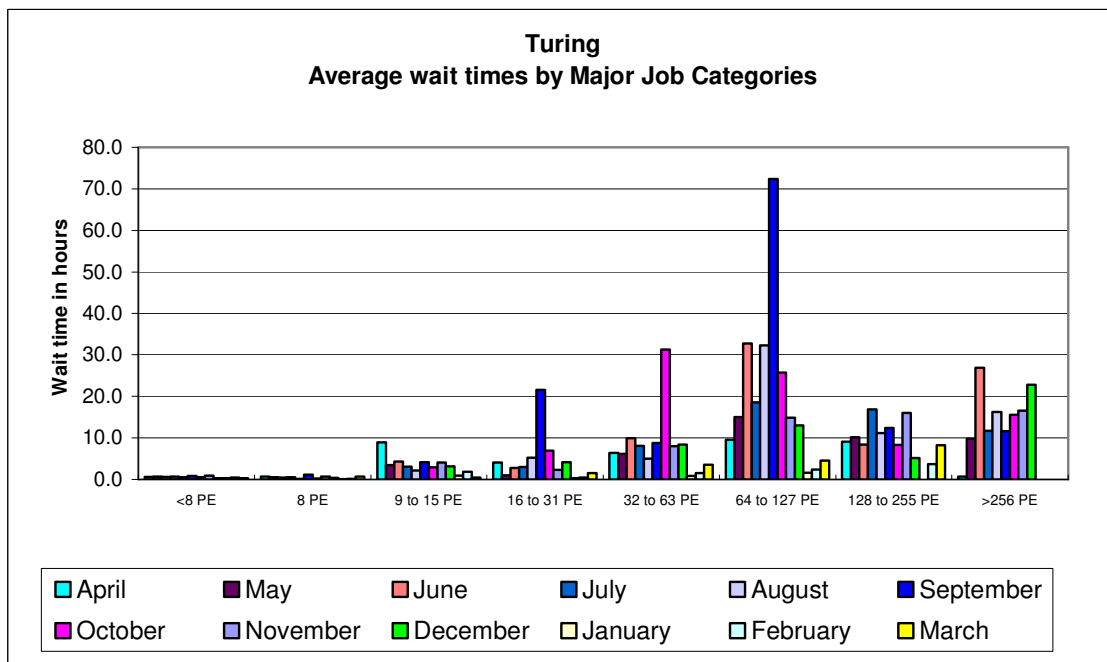


Turing PE usage is shown by Research Council during the last 12 months of service in the above chart.







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 two charts.




The above chart, and the one below, show the wait time trend in hours on the Turing system.




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



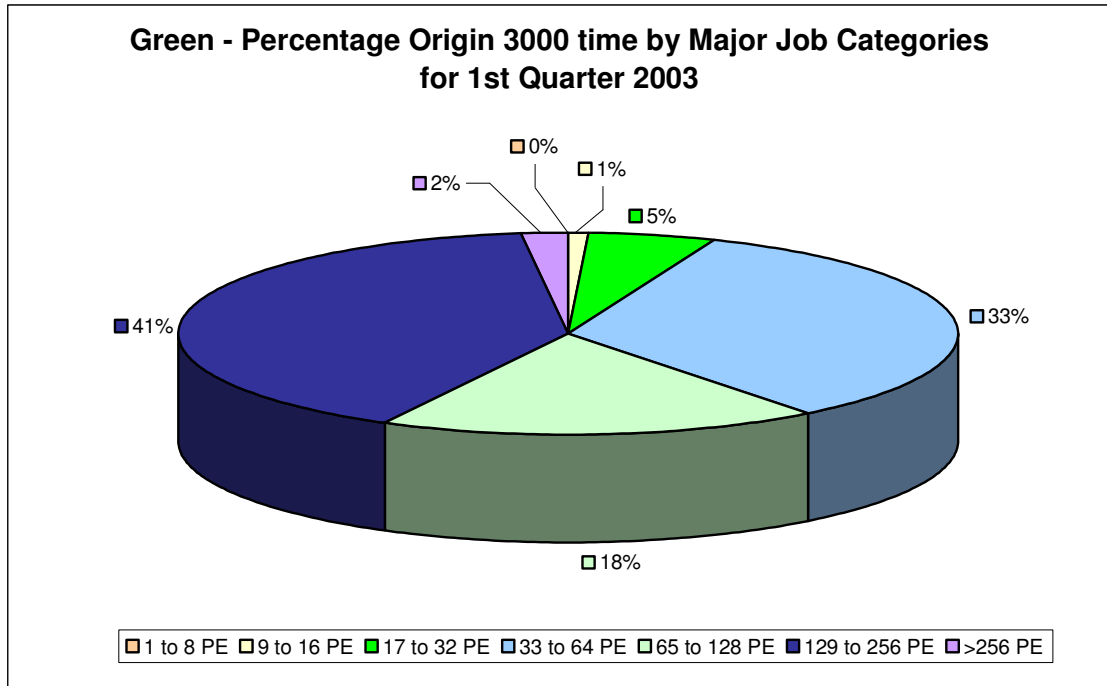


The two "spikes" shown in wait times for Green during February and March were due to two large jobs submitted which had to wait for a high load of smaller jobs to finish.


The next three charts show the percentage PE time utilisation by the major job categories on the Turing, Green and Fermat systems for the 1st quarter 2003.




The workload on Turing for the first quarter was fairly evenly spread among the mid- to high-end ranges of PEs.




On Green, the 1st quarter has seen a varied workload, fairly evenly spread across the mid- to high-end ranges of PEs.




On Fermat the greatest proportion of jobs during the quarter were in the 17 to 64 PE range at 83%.



The last few months have seen decreasing usage of the T3E for jobs above 256 and increasing usage in the upper mid-range PEs.



Usage on Green is tending more to the higher-end range of PEs now that all PEs are dedicated to LSF.




The trend on Fermat currently is that the workload is relatively evenly spread across the range of PEs.

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

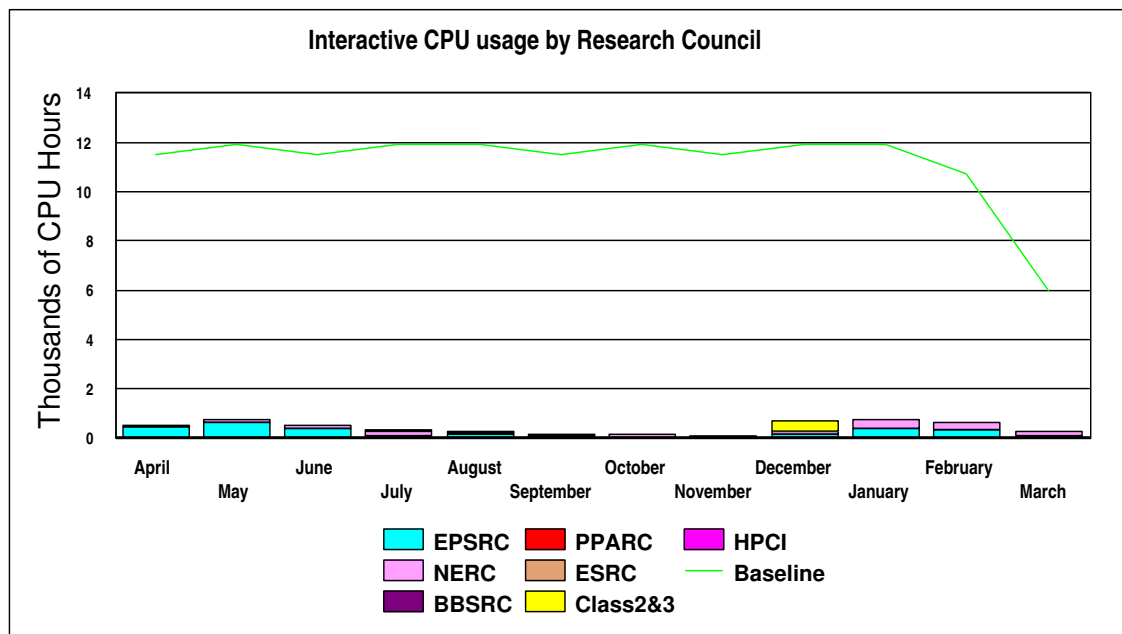
The graph below shows the Gflop year's utilisation on CSAR's systems by Research Council for the last 12 months.




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 CPUs) up to the end of February 2003, at which point the interactive usage was transferred to Wren and Fermat became a batch-only system.. Eight of the higher speed 500Mhz CPUs in the Origin 300 system (Wren) deliver the baseline capacity equivalent to that which was previously available on Fermat for interactive usage.



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




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

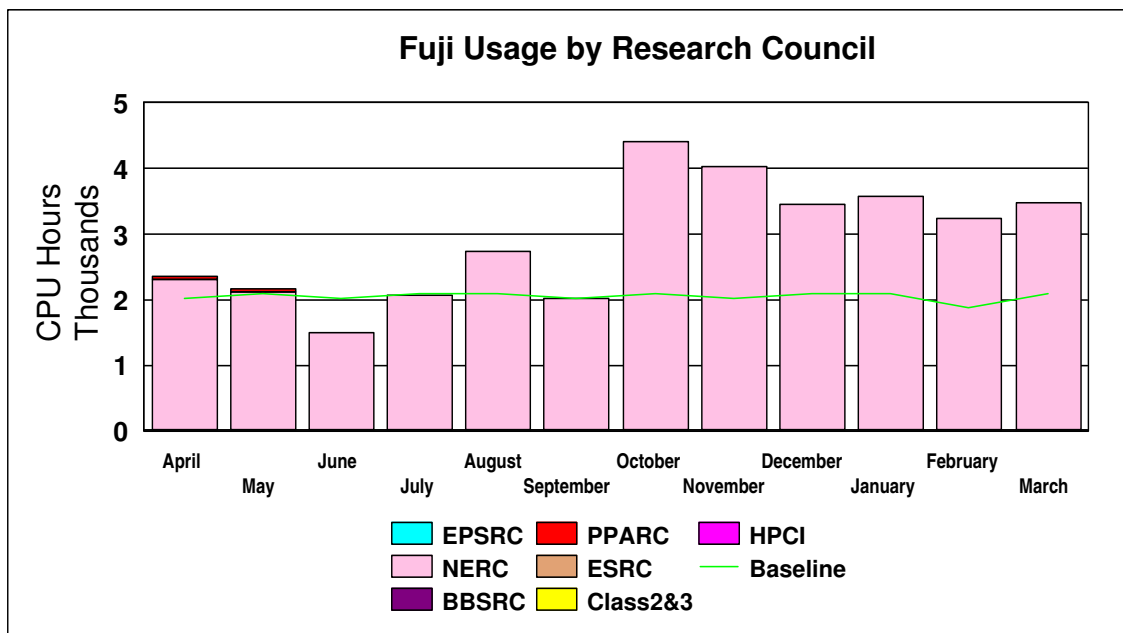


The graph above illustrates the historic allocation of the Medium Performance Disk on Fermat.

The graph below shows the historic HSM usage by Research Council funded projects, which has exceeded the overall Baseline of 16 Terabytes, and now totals 46 Terabytes.



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 was variable during the quarter.

2.2.3 Guest System Usage

The Compaq ES40 cluster has now been withdrawn from service due to continued zero usage demand.

2.3 Service Status, Issues and Plans

Status

The workload on the machines is showing a fairly balanced spread of work across all major systems.

Wren has now taken over from Fermat as the interactive system, leaving Fermat to be a dedicated batch system.

The end of the quarter saw the decommissioning of the Fujitsu VPP 300/8 system Fuji, in accordance with the wishes of NERC.

Issues

There are no issues to report for this quarter.

Plans

Plans are now underway to introduce a 32-PE Altix system into the service by 30th June 2003. This will be followed by a 256-PE Altix system which is to be introduced by September 30th 2003. Further details will be announced as they become available.

It is also the intention of CfS to further upgrade the Silo configuration by the addition of a second tape silo and 4 additional fibre-attached tape drives. These will be used to improve the response and reliability of the Data Migration Facility and to automate remote copying of file system backups.

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 2003, requesting a total of 70975 service tokens.

3.2 New Projects

4 new CSAR projects were started with 10670 service tokens being awarded in total.

3.3 Queries

Overall, 212 Class 1, 2 and 3 queries relating to the CSAR service were received by the Helpdesk between 1st January and 31st March 2003.

3.4 Service Quality Tokens

Three black marks have been recorded for the poor performance of /hold. Action is being taken to address these problems, both in terms of additional tape drives and in an additional on-line area of storage.

3.5 Annual Report

The CfS Annual Report is currently being prepared for publication.

3.6 CSAR Focus

The Winter edition of CSAR Focus has been published and distributed. Planning has started for the Summer 2003 edition.

3.7 CSAR User Steering Group

The ninth six-monthly meeting of the CSAR User Steering Group was held early in January using the Access Grid facilities at Imperial College and the University of Manchester.

The major topics included:

- o CSAR service developments, particularly LSF and the possibility of installing an SGI SN IPF system.
- o Issues raised by the user community either via Dr Lois Steenman-Clark, in her capacity as the Chair of the User Liaison Forum, or by the User Survey or other routes.
- o CSAR and HPC(X). Both CfS and EPSRC described how the services would operate together.

3.8 LSF

The LSF user guide and status pages have been modified as required during the last three months.

3.9 Audit

An audit of the CSAR service took place this quarter.

4 Scientific Application Support Services

4.1 Training and Education

The following CSAR training courses were delivered this quarter:

- o Introduction to High Performance Computing
- o Introduction to MPI
- o Advanced MPI
- o Fortran 90
- o Share Memory Parallelisation with OpenMP

4.2 Consortia Support

Adrian Tate continues to work closely with one consortia and plans to create an alternative to their matrix scatter routine. Adrian has discussed the possibility of this same group becoming involved in some experimental Netsolve based work. This project is now very likely, and we expect a variety of machines and a mixture of Belfast and Manchester staff to be involved in this work over the summer.

Kevin Roy and Neil Stringfellow visited the atmospheric chemistry group at Cambridge to help them parallelise their codes in order that they could migrate from Fuji to the Origin machines.

4.3 Applications Software

CPMD has been installed on the Origins and Turing.

Installation of the PETSc library on the Origins and Turing is underway.

Gaussian O3 is being obtained for the Origins.

Zoe Chaplin is porting the Unified Model to the SGI Origins. The mesoscale version of the model is now up and running and bit reproducing.

4.4 Other Software

In order to allow the most flexible job scheduling service, substantial testing is being performed on the checkpoint/restart facility on the Origins in conjunction with LSF.

MIPSpro 7.4 compilers have been installed on the Origin systems.

LSF is now fully installed and is operational on all the Origins. As of 7th January, it is the primary batch system for the Origins, with 540 of the 556 available batch CPUs now devoted to LSF.

4.5 Benchmarking

We are developing a suite of benchmarking codes for the CSAR systems – choosing a suitable set of benchmarks (both standard benchmarks, e.g. NPB, and also application/user codes, e.g. NAMD) with the aim of running them comprehensively on Turing, Green and the forthcoming Altix, and publishing this information (for the application codes) as part of our software pages on the CSAR site. It is intended to use this as a basis for a suite of UK academic benchmark codes.

4.6 "Top Ten" Scripts

A set of scripts has been created to record the usage of the three main CSAR systems by executable rather than by user/project etc. This keeps the logs of commands issued by batch and interactive jobs lasting longer than ten minutes, and also summarises these logs periodically to show the top ten commands/executables by total cpu-hours.

5 Collaboration and Conferences

5.1 MRCCS Projects

5.1.1 Reality Grid

The RealityGrid project demonstrated trans-European computational steering and visualization using CSAR systems at SGI's VizSummit held in Paris on February 3-5 2003. The demonstration, an evolution of that which took place at SC2002 in Baltimore, utilised 128 processors on CSAR's Origin 3800, Green, to run a Lattice-Boltzmann simulation which was being steered from a laptop and visualized in real time on an SGI Onyx in Paris.

5.1.2 Access Grid

Usage of the Access Grid at Manchester remains very high.

Recommendations from the videoconferencing report mentioned in the last Quarterly Report went to the JISC Committee for the Support of Research (JCSR). All were approved for funding. When a JISC circular is eventually released, Manchester will hope to be high on the list of successful applicants. The report may be viewed at http://umbriel.dcs.gla.ac.uk/Nesc/general/technical_papers/UKeS-2002-04.html.

The next big event for Access Grid will be SC Global 03. This year, SC Global will be much more integrated within the Supercomputing conference. Further information about SC Global (including submission deadlines, etc) may be viewed at <http://www.scglobal.org>.

5.1.3 GRIP

GRIP has taken part in two events in the last quarter:

- o Donal Fellows attended the Globus WORLD conference, San Diego, US, 13-17 January 2003, where he presented a GRIP resource broker poster.
- o John Brooke and Jon MacLaren chaired a number of workshops/BOF sessions at the Global Grid Forum 7 at Tokyo, Japan, 4-7 March 2003.

Jon MacLaren is no longer working on the GRIP/EUROGRID project. GRIP/EUROGRID has recruited a new member of staff, Kevin Garwood, who started work on the project in February and is based in the Computer Science department. He is working on the ontology.

5.1.4 UKHEC

The UKHEC project formally ended on 31st December 2002. This successful collaboration between MRCCS, DL and EPCC provided valuable support to the high end computing community over the 3 years of its term.

Three events took place this quarter under this project:

- o Data Management for High End Computing, 10th January 2003 – CLRC Rutherford Appleton Laboratory.
- o Visualization for High End Computing, 6th February 2003 – University of Manchester.
- o Workshop on Parallel Finite Element Analysis, 13th-14th March 2003 – University of Manchester.

MRCCS attended all of these seminars and hosted the latter two.

5.1.5 Advanced Virtual Prototyping /Finite Elements

There is progress in integrating MPI-based programs, specifically the parallel Finite Element Analysis codes under development at the University of Manchester into the Virtual Prototyping environment. Lee Margetts is also investigating the possibility of staging an Access Grid demo of Virtual Prototyping for SGGlobal 2003.

5.1.6 Co-array Fortran Development

Adrian Tate is working on Co-array Fortran based development. From this project we expect to co-sponsor a Co-array Fortran workshop later this year.

5.2 Events

The following events were attended by SVE staff:

- o UKHEC seminar: Data Management (Rutherford Appleton Laboratory, 10th January)
- o IBM seminar (Manchester Computing, 17th January)
- o Dell-sponsored HPC seminar (Manchester Computing, 22nd January)
- o UKHEC Review (Daresbury Laboratory, 20th February)
- o "Modelling with terascale computers" (RI, London, 21st February), CSAR exhibited their work at this meeting
- o Materials Modelling event (UMIST, 18th March)
- o Zoe Chaplin attended the 5th International Workshop on Next Generation Climate Models for Advanced High Performance Computing Facilities in Rome from 3-5 March. This workshop was mainly attended by Japanese Climate Modellers who are using the Earth Simulator, but representatives from a number of other major climate centres were there as well, together with a number of computer vendors

6 Added Value Services

6.1 International Conferences

Preparations are underway for the exhibitions at ISC2003 (Heidelberg, Germany, June) and SC2003 (Phoenix, USA, November), and for a presentation at CUG 2003.

6.2 Seminars

6.2.1 MRCCS/ESNW Seminar Series

The following seminars, jointly organised by MRCCS and ESNW, have been held during the period:

- o ViSiCADE – an Integrated Interactive Virtual Prototyping Environment, *Dr M Ashworth, Daresbury Laboratory.*
- o Kinematics Engines and Constraint Recognition in Virtual Reality, *Dr S Bee, University of Salford.*
- o Stochastic Modelling of Heterogeneity in Geomaterials: The Case for High-End Computing, *Dr M Hicks, University of Manchester.*
- o EU DataGrid and GridPP Authorization and Access Control, *Dr A McNab, University of Manchester.*
- o RealityGrid project, *Dr S Pickles, University of Manchester.*

6.2.2 Virtual Prototyping

Lee Margetts gave a seminar in the Universidad Politécnica de Madrid entitled 'Future Design Spaces – Virtual Prototyping and Distributed Computing'.

6.3 Summer School 2003

Planning for the Summer School on HPC in Finite Element Analysis, jointly organised by MRCCS and the National Science Foundation (NSF) of the USA is well underway. Speakers will be from France, Germany and Japan (Earth Simulator) as well as the USA and UK.

6.4 Visualization

The Manchester Visualization Centre is being presented with a Medal of Achievement in April for work on the Op3D project. MVC has submitted a case study on 3D Volume Visualization in the Operating Room which will officially become part of the Computerworld Honors Collection on Sunday April 6th in San Francisco Hall when the 2003 Collection is formally presented to the Global Archives.

The Parallel Support Toolkit (PST) project commenced on 1st January 2003. This will provide a framework for the execution on development of parallel visualization modules running AVS/Express environment. The work is being done by MVC and is supported by AVS, KGT, JAERI (Japanese Atomic Energy Research Institute). The initial part is a feasibility study (providing parallel isosurfacing for a number of systems) to help secure further funding, for which the three main deliverables have been accomplished on time.