CSAR Service

Consolidated Management Report

3rd 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 514 to date.

The Cray T3E (Turing) system continued to run almost to full capacity again this quarter.

The Origin 3000 (Green) now has 512 CPUs, and usage is growing steadily. 55% of the usage has been greater than 64 CPUs in size.

The T3E has seen 84% of the Quarter's workload at greater than 64 PEs in size.

CfS remains active in the UK Grid Forum.

The recently installed facilities are now almost fully integrated into the CSAR service. Wren is now running the interactive service in parallel with Fermat, the SAN is in full production use, and LSF is gradually being expanded so as to fully replace NQE in January.

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

<u>Table 1</u> 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

	Performance Targets								
Service Quality Measure	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.

	2001/2											
Service Quality Measure	Oct	Nov	Dec	Jan	Feb	March	April	May	June	July	Aug	Sept
HPC Services Availability												
Availability in Core Time (% of time)	98.60%	98.60%	100.00%	99.86%	99.73%	99.70%	96.17%	96.08%	97.66%	99.2%	99.75%	98.75%
Availability out of Core Time (% of time)	99.50%	99.50%	98.49%	99.89%	99.85%	99.97%	97.75%	99.90%	99%	100%	100%	99.42%
Number of Failures in month	2	2	4	2	1	2	2	1	4	0	1	2
Mean Time between failures in 52 week rolling period (hours)	365	365	337	350	324	313	302	324	313	365	381	381
Fujitsu Service Availability												
Availability in Core Time (% of time)	100%	100%	100%	100%	100%	100%	96.89%	100%	100%	100%	100%	100%
Availability out of Core Time (% of time)	100%	100%	100%	100%	100%	100%	98.92%	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	<1	<1	<1	<2	<1	<1	<2	<5	<2	<2	<1
Administrative Queries - Max Time to resolve 95% of all queries	<1	<0.5	<2	<0.5	<1	<2	<2	<3	<5	<2	<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	<0.5	<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 mon	2	2	2	2	2	2	2	2	2	2	2	2

Notes:

Table 2

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

Turing availability x 143/(143+40+233)] + [Fermat availability x 40/(143+40+233) + Green availability x 233/(143+40+233)]

2. Mean Time Between Failures for Service Credits is formally calculated from Go-Live Date.

<u>Table 3</u> 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

	2001/2											
Service Quality Measure	Oct	Nov	Dec	Jan	Feb	March	April	May	June	July	Aug	Sept
HPC Services Availability												
Availability in Core Time (% of time)	0.039	0.039	-0.058	-0.039	-0.039	-0.039	0.078	0.078	0.078	0	-0.039	0.039
Availability out of Core Time (% of time)	-0.039	-0.039	0.000	-0.047	-0.047	-0.047	0.039	-0.047	0.000	-0.047	-0.047	0
Number of Failures in month	0	0	0.008	0	-0.008	0	0	-0.008	0.000	-0.009	-0.008	0
Mean Time between failures in 52 week rolling period (hours)	0	0	0	0	0	0	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	-0.016	-0.016	-0.016	-0.016	0	-0.016	-0.016	0	0.031	0	0	-0.016
Administrative Queries - Max Time to resolve 95% of all queries	-0.016	-0.019	0	-0.019	-0.016	0	0	0.016	0.031	0	-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	-0.002	-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 mon	0	0	0	0	0	0	0	0	0	0	0	0
Monthly Total & overall Service Quality Bating for each period:	-0.04	-0.04	-0.05	-0.08	-0.08	-0.07	0.03	0.00	0.05	-0.05	-0.08	-0.01
	2.01	2.01	2.00	2.00	2.00	2.07	2.00	2.00	0.00		2.00	2.01
Quarterly Service Credits:			-0.13			-0.23			0.08	I		-0.14



1.2 No. of Registered Users

The current position at the end of the quarter is that there are 514 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 are no black marks or gold stars allocated to the service.

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.
- 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 in hours 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 is now managing the queue wait times on the T3E (Turing) to an acceptable level.



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



Green wait times are higher due to fuller loading of the system typically requiring 24 hours of processing.

The next three charts show the percentage PE time utilisation by the major job categories on the Turing, Green and Fermat systems for the 3rd quarter 2002.



The predominant job size on Turing is large, with the percentage of jobs greater than 64 PEs for the quarter being 84%.



On Green, the percentage of CPU utilisation used by jobs requiring greater than 64 CPUs for the quarter was 55%.



On Fermat the greatest proportion of jobs during the quarter were in the 17 to 64 PE range at 81%.



From the above chart it can be seen that the workload on the T3E has now stabilised.



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.

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 32 Terabytes. The primary usage is for NERC and, in total, this amounts to about 31 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 was variable during the quarter.

2.2.3 Guest System Usage

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

2.3 Service Status, Issues and Plans

Status

The SAN is now fully integrated and in production use. Wren is now running interactive work in parallel with Fermat; the latter will be removed from interactive use in January. At this point Fermat will become a batch-only machine.

Issues

The latter half of this quarter has seen fewer large jobs run due to the impact of the LSF installation. This has resulted in two batch partitions, one LSF and one NQE, existing on the machines at the same time. However, this has provided the user base with an extended testing period.

Plans

The LSF migration will continue in such a manner that the partition size for LSF will be gradually increased and the NQE partition decreased, until only a small legacy partition for NQE remains.

It is planned for Wren to fully replace Fermat as the interactive machine in January 2003.

It is also planned for a transformer maintenance to be carried out in the computer room during December, further details will be provided nearer the time.

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**

9 applications for new CSAR projects were received in the third quarter of 2002.

3.2 New Projects

3 new CSAR projects were started.

3.3 Queries

Overall, 197 queries relating to the CSAR service were received by the Helpdesk between 1st July and 30th September 2002.

3.4 Annual Report

The CfS Annual Report is currently being prepared for publication.

3.5 CSAR Focus

The Autumn edition of CSAR Focus is currently being prepared for publication.

3.6 SAN

Usage reports are being updated to reflect the new resources that have been made available due to the SAN installation. The resource calculator and application form have been updated. All Fermat users have been given accounts on Wren and all projects have been given an initial small allocation of Wren CPU time.

3.7 Web Pages

The CSAR web pages have been considerably updated as part of the Web Accessibility Initiative, in order to improve access to people with disabilities.

4 Scientific Application Support Services

4.1 Training and Education

The late 2002 course schedule is now underway. The type and content of the courses has been revised to emphasise making best use of the CSAR systems.

4.2 Consortia Support

Work continues in the field of parallel numerical library optimisation. Work for one consortium is expected to conclude soon and the results will be presented at the conference Parallel Matrix Algorithms and Applications in Neuchatel, Switzerland, in November of this year. Further work is likely with the same consortium and similar work with other consortia is being pursued. This work has also generated significant interest from a computer vendor.

Other work has included:

- testing of the unified model with LSF
- assistance with a finite element code for a class 2 project
- debugging a large software package

4.3 Applications Software

All of the main application software has been moved over to the SAN. The access method has been improved via the use of modules.

4.3.1 Unified Model

We have installed the Unified Model version 5.3 on Turing and are porting it to the Origins.

4.3.2 GAUSSIAN98

A new version of this software, Rev. 11.3, has been installed on the Origins.

4.3.3 LAMMPS

We are installing the latest version on the CSAR machines.

4.3.4 NAMD

We have installed version 2.4 on the Origins.

4.3.5 MATLAB

We have installed release 12.1 on Wren.

4.3.6 Grid Software

Globus has been installed on Wren. Globus job managers have been installed so that jobs can be submitted to the LSF queues on the Origins. Globus now supports job limits. Unicore and Globus now accept each other's certificates.

4.4 LSF

Extensive testing of LSF has been conducted using a number of codes. A conversion script has been made available to several trial users that converts their existing NQE scripts to LSF, and valuable feedback has been obtained from them. As LSF will be running with dynamic cpusets, further tests have been performed to see what impact these will have on the system and to see how effectively jobs will run in them.

5 Collaboration and Conferences

5.1 MRCCS Projects

5.1.1 Reality Grid

At the UK e-Science All-Hands Meeting in Sheffield on the 2nd - 4th September, a demonstration was performed that used Manchester's Bezier (SGI Onyx) to visualize data produced in a multi-centre computational steering experiment. The simulation, a parallel Lattice-Boltzmann code which simulated the phase separation of two fluids, was run on eight processors of an Onyx at Queen Mary, University of London. The visualization software performed parallel volume rendering and isosurfacing using the vtk library (C++ wrapped by Tk/Tcl) on eight processors of a second SGI Onyx at the University of Manchester. The steering and application-launching software was a Java GUI to the Unicore client, run on a laptop at the North West Regional e-Science Centre's booth in Sheffield. The connection between the laptop and the Onyx in Manchester had a bandwidth of approximately 10 MBps and used VizServer to compress and transfer the images.

5.1.2 Access Grid

The status of the Access Grid at Manchester is much the same as reported in the last Quarterly Report (http://www.csar.cfs.ac.uk/admin/reports/service_quarterly/second2002.pdf). In brief, usage is high and users are varied. An interesting example of this is a performance art event that was held in October - 'navigating gravity' - in conjunction with the University of Florida and including an audience in Australia. More details can be found on http://www.navigatinggravity.net/. Global development on Access Grid technology is imminent with the release of AG2.0 due early next year. Manchester hopes to have a high profile involvement in this effort. The e-Science Core Programme commissioned a comparative report on videoconferencing technologies to help determine a roadmap for the future. Michael Daw managed this effort, which involved many experts from the fields of H.323/H.320 videoconferencing and VRVS as well as Access Grid. This report should soon be publicly available and should form the basis for new UK funding to support distributed collaboration.

5.1.3 GRIP

GRIP, the Grid Interoperability Project is proceeding well. Manchester is coordinating publicity for this project, which includes running the website.

5.1.4 SAMD Project

Work has now been completed on SAMD, the grid application to manipulate econometric datasets on HPC resources, and a report is currently being written.

5.1.5 North West Centre for Advanced Virtual Prototyping

One new area being explored involves the integration of the access grid functionality into the virtual reality framework being developed for the advanced virtual prototyping research centre. This could be of great benefit for the distributed, collaborative working environment - one of the objectives of this project.

5.1.6 Finite Elements

Summer student work into visualisation of finite element data using AVS Express has led to closer collaboration with the School of Engineering. Ongoing collaborative efforts will further develop the visualisation work and bring it into everyday usage by the Geotechnical Engineering Group. It is expected that this will lead to future FEA visualisation projects.

The use of the Cray MTA-2 for the efficient running of finite element codes is being investigated.

5.2 Events

Kevin Roy attended EWOMP 2002, the European Workshop on OpenMP, in Rome, Italy, from September 18th-20th.

The annual EPSRC HPC User Meeting, also known as the "Town Meeting", was held in London on the 30th September. The Computation for Science Consortium responsible for CSAR was well represented, with Terry Hewitt, Mike Pettipher, Andy Jones, and John Brooke (all UoM); John Rawlins and Helen Heath (both CSC); and Andy Grant and Crispin Keable (both SGI) all attending. A total of around 125 attendees also included user representatives, the Research Councils, HPCx and manufacturers.

6 Added Value Services

6.1 International Conferences

We are preparing for SC2002 in November at Baltimore, where we will have an exhibition stand.

6.2 Seminars

The autumn programme of seminars will have subjects relating to HPC application areas, grid computing and clusters. The programme will continue to bring in speakers from around the UK and worldwide. Most of the seminars will be offered to a worldwide audience via the Access Grid, with the speakers themselves often using this technology to save time, energy and expense by presenting remotely.

6.3 Visualisation

In addition to the event described under the reality grid project (Section 5.1.1), on July 21st, the Sunday before the Global Grid Forum in Edinburgh, a demonstration was run to show the new features in VizServer. This was part of Manchester's collaboration with SGI and entitled "Visualization on the Grid". Bezier, the new SGI Onyx at Manchester, was used to run both interactive medical and astrophysics visualizations; the visualization software used was volumizer. Glasgow used a networked PC to view the results. VizServer allowed both Manchester and Glasgow in turn to see and interact with the visualization, altering the view and other parameters.

6.4 Summer Student Program

Two students had placements here over the summer. One student developed tools for visualising finite element data for Manchester University's engineering department and the other developed introductory training materials for UNIRAS.