

CSAR Service

Consolidated Management Report

3rd Quarter 2004

Management Summary

This is the consolidated Management report for the third quarter 2004 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 473 to date.

The workload on both the Origin 3000 Green and the Altix 3700 Newton has been fairly evenly spread across the mid- to high-end ranges of PEs during the third quarter of this year.

CSAR has been granted an 18-month extension of service contract until June 30th 2006. With this extension CfS implemented a further technology refresh introducing a 256 processor Itanium-2 (Madison) based SGI Altix, now enhanced to 512 processors with 1 Terabyte of memory.

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 3rd quarter 2004 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	2003/4											
	Oct	Nov	Dec	Jan	Feb	March	April	May	June	July	Aug	Sept
HPC Services Availability												
Availability in Core Time (% of time)	98.84%	98.95%	98.75%	97.49%	98.16%	98.51%	89.39%	94.21%	97.49%	97.97%	100%	99.52%
Availability out of Core Time (% of time)	99.28%	97.74%	98.3%	98.88%	97.9%	99.48%	91.90%	99.73%	97.85%	100%	99.2%	99.80%
Number of Failures in month	4	3	5	5	4	3	5	3	4	2	2	2
Mean Time between failures in 52 week rolling period (hours)	365	337	283	265	243	383	314	280	245	235	222	211
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	<2	<1	<1	<2	<2	<2	<2	<3	<2	<2	<1	<2
Administrative Queries - Max Time to resolve 95% of all queries	<1	<1	<1	<0.5	<0.5	<0.5	<0.5	<0.5	<1	<0.5	<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 formula, based on the relative NPB performance of Fermat, Green and Newton at installation:

$$[\text{Fermat availability} \times 40 / (40+233+343)] + [\text{Green availability} \times 233 / (40+233+343)] + [\text{Newton availability} \times 343 / (40+233+343)]$$
- Mean Time Between Failures for Service Credits is formally calculated from a rolling 12-month period.

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	2003/4											
	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	0.039	0.078	0.078	0.039	0.195	0.195	0.078	0.078	-0.058	-0.039
Availability out of Core Time (% of time)	-0.039	0.078	0.078	0	0.078	0	0.039	0	0.078	-0.047	0	-0.047
Number of Failures in month	0.008	0	0.0004	0.0004	0.008	0	0.0004	0.0	0.008	0	0	0
Mean Time between failures in 52 week rolling period (hours)	0	0	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002
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	-0.016	-0.016	0	0	0	0	0.016	0	0	-0.016	0
Administrative Queries - Max Time to resolve 95% of all queries	-0.01551	-0.016	-0.016	-0.019	-0.019	-0.019	0	-0.019	-0.016	-0.019	-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	0.02	0.02	0.01	0.05	-0.01	0.09	0.07	0.05	-0.02	-0.07	-0.07
Quarterly Service Credits:	0.02			0.05			0.21			-0.15		

Table 3

Overall, system availability has steadily improved since the previous Quarter.


1.2 No. of Registered Users

The current position at the end of the quarter is that there are 473 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 one user had allocated four black marks to the service:

SUMMARY OF SERVICE QUALITY TOKEN USAGE

No of Stars or Marks	Consortia	Date Allocated	Reason Given
4	cs3025	17/09/04	Long delay in help with porting software

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) SMP (Altix/Origin) Usage by month, showing usage each month of CPU (GFLOP-Years as per NPB), split by Research Council and by system. Overlaid horizontal lines show the overall Capacities.
- 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 and overall Capacity are shown by overlaid horizontal lines.
- c) 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.
- d) HSM/Tape Usage by month, showing the volumes held each in GBytes, split by Research Council. The Baseline Capacity (16 Terabytes) available will be shown by an overlaid horizontal line.


2.1 Service Usage Charts

The graphs below show recent monthly CPU, disk and HSM allocations and usage.




Usage of the SGI Altix 3700 system Newton is shown by Research Council in the above graph.







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




The above chart, and the one below, shows the wait time trend in hours on the Altix 3700 Newton.




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







The next series of four charts show the percentage CPU time utilisation by the major job categories on the Green, Fermat and Newton systems for the 3rd quarter 2004.




On Green, the 65 to 128 PE range has seen the greatest percentage of workload during this quarter.



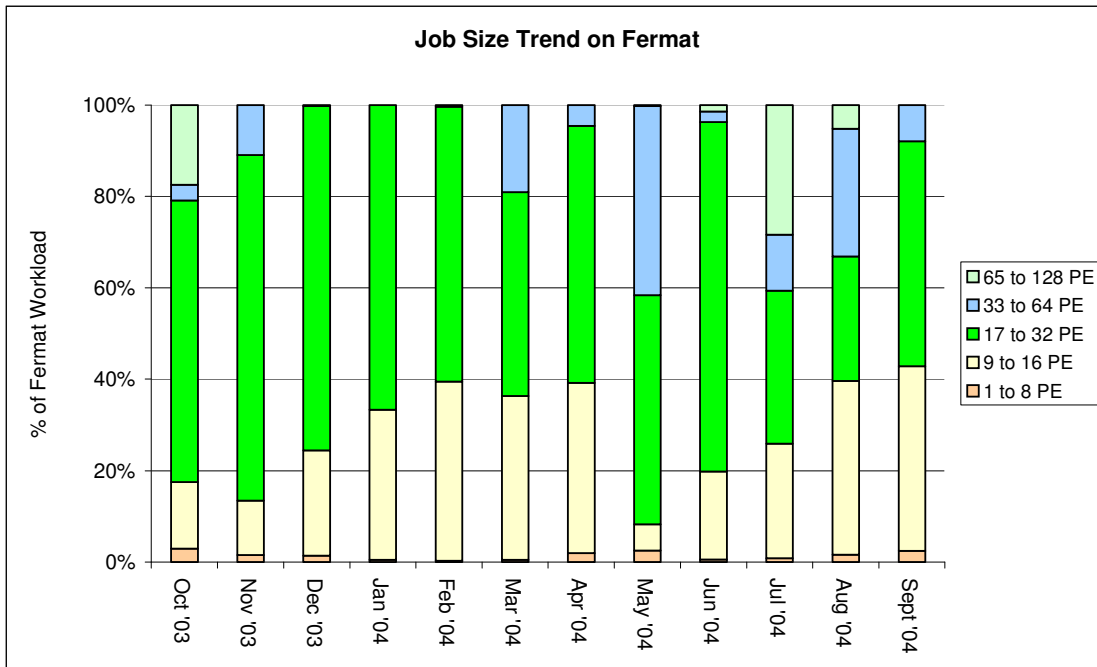
On Fermat there was a varied spread of work across the entire range of PEs.



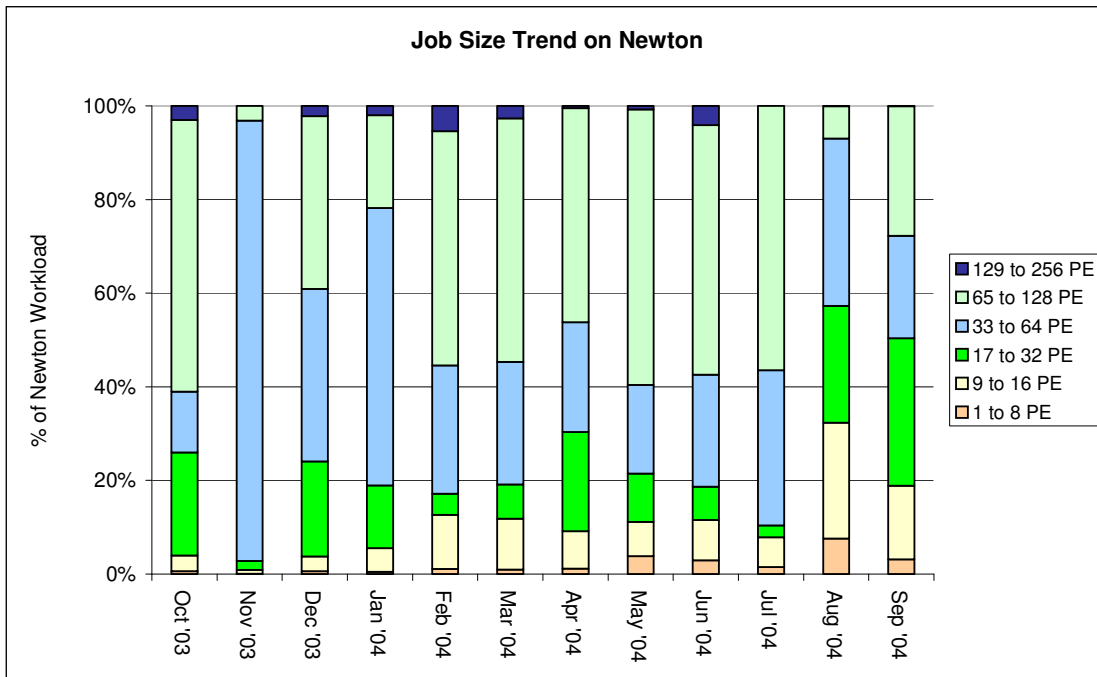
There was a good spread of work across Newton during the third quarter.



Usage on Green tended more to be in the mid-range of PEs during the 3rd quarter.



The workload on Fermat for this quarter was also spread fairly evenly across the mid-range of PEs.




During the third quarter Newton's usage was reasonably spread across the machine.

2.2 System Usage Graphs


In all the Usage Charts, the baseline varies dependant upon the number of days in each month, within a 365-day year.

2.2.1 Baseline System

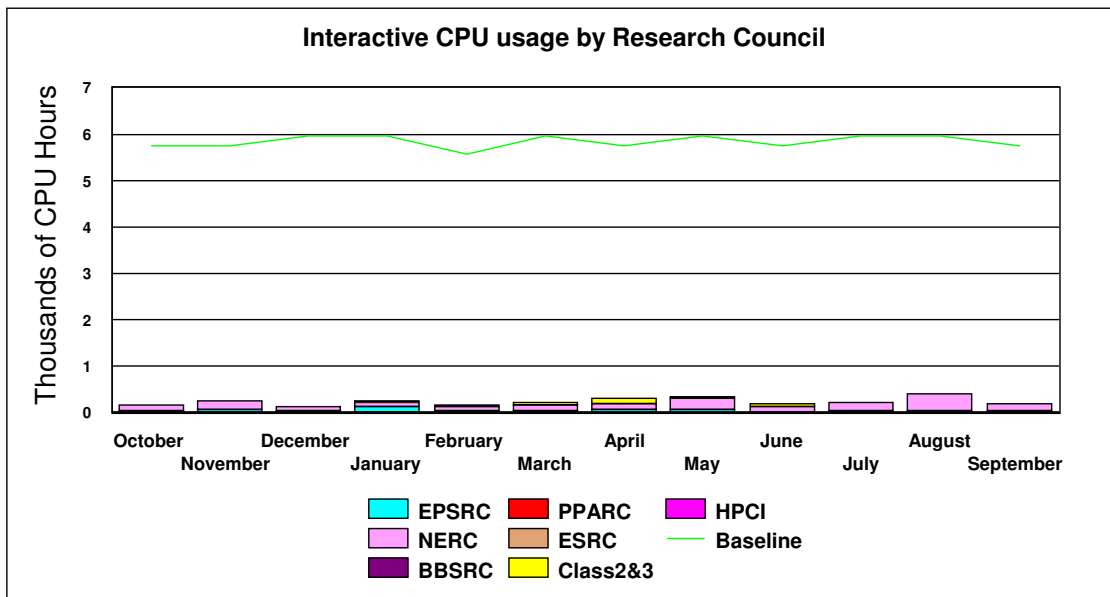
The graph below shows the Gflop Years utilisation on the CSAR systems by Research Council for the last 12 months.




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




The next chart shows the historic interactive usage of the Origin 300 system Wren. Eight of the higher speed 500Mhz CPUs in Wren deliver the baseline capacity equivalent to that which was previously available on the Origin 3000 system Fermat for interactive usage.




The following two charts detail the historic usage of the Origin 3000 system (Green) and the Altix 3700 system (Newton).



NERC is by far the largest user of Green.




The next series of graphs illustrates the usage of the Medium Performance disk and HSM/tmp resources of the system.



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

The next graph shows the historic HSM usage by Research Council funded projects, which has exceeded the overall Baseline of 16 Terabytes, and now totals about 48 Terabytes.



2.2.2 Guest System Usage

There is currently no Guest System usage.

2.3 Capability Incentives

Capability incentives were historically given on the T3E system Turing for jobs of 512 PEs and above. In July 2003 it was announced that discounts for capability jobs available on all CSAR systems had been approved to include the SGI Origin 3000 system Green and the SGI Altix 3700 system Newton.

These capability incentives were agreed with the Research Councils to encourage capability usage of the national supercomputers for greater scientific achievement, and offer the following discounts:

System	No of Processors	Discount
newton	192+ CPUs	15% discount
newton	128+ CPUs	10% discount
green	384+ CPUs	15% discount
green	256+ CPUs	10% discount

Discounts are given in the form of refunded Service Tokens.

Changes in usage patterns will be monitored and, subject to review, CfS reserve the right to change the incentives at any future date.

The following table displays the capability incentive discounts granted during the third quarter of this year.

Service Tokens Refunded: Quarter 3 2004 Usage							
System	Consortia						Total
	cse076	cse133	csn003	csn066			
Green 256+ PEs				6.69			6.69
Green 384+ PEs			3.39				3.39
Newton 128+ PEs	342.52	12.59					355.11
Newton 192+ PEs							0
Total Tokens							365.19

This is within the CfS Management Board's forecast.

2.4 Service Status, Issues and Plans

Status

The service has been reasonably utilised throughout the third quarter of 2004, with usage exceeding baseline.

During the quarter there was a relatively balanced spread of work across all major systems.

A three-phase upgrade to the Altix system Newton was completed during this quarter. During July an additional 128 processor node was added to the system, with each processor in the new node having 4GB memory, 0.5TB in total.. The second and third phases were completed during August, with the memory on the original 256 processor system being upgraded to 2GB per processor and a further 128 processors added to the new node, these being the faster 1.5GHz processor chips. This upgrade means that the Newton system totals 512 processors with 1 Terabyte of memory.

Issues

The increased stability gained after having completed extensive operating system upgrades over the past few months mean that there are no issues to report for the third quarter.

Plans

It is planned to combine the 4 x 64 processor original Newton nodes into one 256 processor Single System Image. This work is expected to be carried out during the last quarter of the year, with further details to be announced 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

10 applications for new CSAR projects were received, requesting a total of 370,830 service tokens.

3.2 New Projects

7 new CSAR projects were started with 139,324 service tokens being awarded in total.

3.3 Finished Projects

6 projects finished

3.4 Queries

A total of 152 CSAR queries were dealt with:

- o 109 non-in-depth
- o 38 user registration and admin
- o 5 in-depth:

3.5 Service Quality Tokens

Four black marks were received due an unacceptably long delay in help with porting software. This was resolved after discussion with the research group concerned and agreement was made for progress with this work.

3.6 CSAR Focus

The Summer/Autumn edition has been completed and distributed.

3.7 CSAR Website

The new CSAR website is now live. Notification was sent to all users upon its completion advising of the details of the new site.

3.8 User Steering Group

The 12th CSAR User Steering Group meeting was held on 12th July 2004.

4 Scientific Application Support Services

4.1 Training and Education

No training courses were supplied during this quarter.

4.2 Consortia Visits

Kevin Roy attended the Materials Chemistry consortium meeting in London and gave a presentation about the developments in the CSAR service. The feedback from the meeting was very valuable. Kevin also visited other researchers at the same site.

Neil Stringfellow and Craig Lucas attended the Turbulence Consortium meeting in the New Forest. Brief details about the service and about the consortium optimisation work were presented. Again the feedback was invaluable, as were discussions about further work and consortium proposals.

4.3 Consortia Support/Software

The primary work has been in porting and optimising codes for the SGI Altix service, Newton. Additional testing has been performed on the new 1.5GHz processors.

4.3.1 NWChem

Further testing has shown that good performance can be achieved using the local /tmp disks. Performance on the 1.5GHz, with larger cache, has shown generally about 15-20% performance improvement over the 1.3GHz processors. CSAR is aware that further development work is being carried out with the Global Arrays software on the Altix, particularly with respect to the use of multiple partitions. This may lead to further performance improvements.

4.3.2 VASP

Further work has been performed in rebuilding the latest version of VASP (4.6.21) with the latest compiler (version 8.1). This version uses SHMEM and shows about 20% better performance than the standard version on small numbers of processors. It also scales better resulting in about 40% improvement on 128 processors. Additional compiler bugs have been reported. Due to compiler issues, the 7.1 compiler may still be the best choice in some circumstances. We are working with SGI and Intel to resolve outstanding issues.

4.3.3 Tomcat

Additional optimisation work has been done on the Tomcat code, on behalf of the NCAS consortium. Improvements in the advection section mean the code now scales to 64 processors. Load balancing in the chemistry section is ongoing.

4.3.4 Pchan

Further work was carried out on this code on behalf of the Turbulence consortium. This has included:

- removal of asynchronous and memory copies to get the best out of the single copy features of MPT
- investigating the use of MPI 1-sided. This may result in more MPI calls because MPI-2 functions in MPT do not allow non-contiguous datatypes
- needed MPI_TYPE_CREATE_RESIZED, because unimplemented MPI-2 routines meant that complicated MPI could not be replaced with a single call
- investigations in the use of MPI-I/O indicate improvements, but care is required when using the SAN disks
- it was noted (and reported to SGI/Intel) that the very useful EXTERNAL32 feature is not currently implemented.

4.3.5 PolComs

When received, the Wave modelling section of the code performed poorly compared to the rest. The target was to improve performance by a factor of 10!

This was achieved on 16 processor runs and load balancing optimisations improved the scalability. It has involved rewriting most of the wave section of the code. There is ongoing work to verify the numerical accuracy of optimisations. This will allow the group to make better use of all machines.

4.3.6 Class 3 Project

Work has proceeded in the parallelisation of a code (LENS) on behalf of a Class 3 project. This has involved more work than initially expected.

4.3.7 Visualisation Support

Visualisation work is required for the Celebration of Engineering event in November. It is intended to build on this work to assist other CSAR users in their use of visualisation.

5 Collaboration and Conferences

5.1 MRCCS Projects

5.1.1 Advanced Virtual Prototyping Research Centre

The SC Global showcase entitled “Collaborative Finite Element Analysis” (Lee Margetts et al.) was accepted for the Supercomputing Conference to be held in November 2004. Work has proceeded in preparation for this event

Preparations are also underway to hold a workshop on virtual prototyping in 2005.

5.1.2 Celebration of UK Engineering Research

Similarly preparations have been made for EPSRC’s Celebration of UK Engineering Research to be held on Wednesday 17th November 2004 in London. This is part of the International Review of Engineering, taking place throughout the week; it is the sixth such review, and only the second in Engineering. Demonstrations are being developed using a stereo projection facility (hired from the University of Birmingham).

5.1.3 NetSolve

NetSolve is an RPC based client/agent/server system that allows remote access of both hardware and software components. There are now developments to use Globus, which is likely to help in the use of CSAR systems.

5.2 National Grid Service

The formal start of service, announced at the All Hands Meeting, was on 1st September 2004.

A course on SRB, organised by the Worldwide Universities Network (WUN), was held at Manchester.

An experimental OGSA DAI service is in preparation.

The Reality Grid performed molecular dynamics simulations on the Teragrid and the NGS systems, demonstrating steering, job spawning and migration.

5.3 Access Grid Support Centre

A new Access Grid facility has been installed at the University of Manchester. This will be used in the forthcoming SC Global 2005.

Manchester staff are actively involved in the preparations for SC Global. Paul Kuchar is the European test co-ordinator.

6 Added Value Services

6.1 International Conferences

Preparations are well underway for SC2004 at Pittsburgh in November, at which the University of Manchester will have an exhibition stand.

6.2 Visits

There have been numerous visits to the University and by our staff, particularly involving computer vendors and software developers, including Platform Computing (LSF) and HKS (Abaqus).

6.3 Visualisation

Progress is being made in the collaboration between Access Grid technology and Manchester Visualisation Centre, to create the first UK installation of a passive stereoscopic access grid node. Consisting of six Christie Digital LX32 projectors, an 8x2 metre real-laced passive stereoscopic surface, driven by three dedicated nVidia QuadroFX 3000G graphics cards. High-end scientific visualisation employing the standard commercial packages, such as AVS and Amira, will be available within the Access Grid.