HPCx Quarterly Report
July – September 2005

1 Introduction

This report covers the period from 1 July 2005 at 0800 to 1 October 2005 at 0800.

The next section summarises the main points of the service for this quarter. Section 3 gives details of the usage of the service, including failures, serviceability, CPU usage, helpdesk statistics and service quality tokens. A summary table of the key performance metrics is given in the final section. The Appendices define the incident severity levels and list the current HPCx projects.

2 Executive Summary

- EPSRC have recently approved the upgrade of HPCx to a Power 5 system. The new system will be built up alongside the existing system during October and should go into service towards the end of November.

- The Annual Plan for 2005 was approved at the Oversight Committee meeting in August.

- Although overall utilisation was lower than the previous quarter, both capability usage (37%) and NERC usage (25%) have been high.

- The capacity-planning look-forward suggests that the HPCx system may be very busy over the next few months, as the e05 consortium have a lot of time to use before their end-date of April 2006.

- There are currently 37 research consortia using HPCx and 5 new EPSRC consortia are due to start in October. Although there are 3 EPSRC consortia that have reached their end-dates, we are clearly approaching the current maximum of 40 research consortia.

- After the successful talk at this year's Edinburgh International Science Festival, we have been accepted for next year's Festival. David Henty will

give a talk titled *Supercomputers: rise of the machines.* We are also investigating participation in next year's National Science Week.

- The latest HPCx newsletter is being mailed out to 3500 people. It includes three research articles from users and a *Stop Press* announcement of the Power 5 upgrade.

- Preparations are well under way for the SPICE and DYNOME experiments at SC2005 in Seattle. This has required work both on Grid-enabling codes and in updating of Grid software (port-forwarding and MPICH-G2).

- One of this quarter's HPCx Technical Reports, *On the Performance of Molecular Dynamics Applications on Current High-End Systems*, was published in the August 2005 issue of *Philosophical Transactions of the Royal Society of London Series A; Theme Issue on Scientific Grid Computing*.

- The summer provided an opportunity to run a good number of HPCx courses, bringing the training activity back on schedule for the year. The programme for the final quarter sees the development and delivery of four days of completely new training material for HPCx users.

- Organisation of The Third Annual HPCx Seminar, *Capability Science on HPCx*, is progressing well. The event will run at Daresbury on 5th December, immediately preceding the 16th Machine Evaluation Workshop.

- Martyn Guest has written a paper based on the ten consortium visits carried out so far this year. A presentation of the findings to the STAC meeting in September was well received and has demonstrated the value of these visits.

- This quarter has seen significant optimisation work on a variety of codes in different areas including: SCAM, GAMESS-UK, VASP and PDVR3DRZ. We have also ported to HPCx a number of other codes in support of the Life Sciences outreach activity.

- Significant testing of the new xlf9 compiler was performed on a wide range of codes prior to its roll-out to users in September. No problems were found and some codes obtained performance improvements. A presentation on the new features in the compiler was made at the HPCx User Group via Access Grid.

- As part of our investigation into future portability of performance, we have ported and profiled a variety of HPCx codes to BlueGene. These include the HPCx implementation test codes and other key HPCx applications. A presentation on this has been made available on the HPCx web site.

3 Usage Statistics

3.1 Availability

3.1.1 Failures

The monthly numbers of incidents and failures (SEV 1 incidents) are shown in the table below:

	July	*August*	*September*
Incidents	13	10	7
Failures	3	1	0

The number of non-failure incidents in July and August was under-reported in the monthly reports as a result of a spreadsheet error.

The following tables give more details on the attribution of the failures:

July

Failure	*Site*	*IBM*	*External*	*Reason*
05.100	0%	0%	100%	External network failure
05.106	50%	50%	0%	Login node crashed/reloaded (twice)
05.107	50%	50%	0%	Login node crashed/reloaded (twice)

Failures 5.106 and 5.107 each represent a pair of events: at 19:16 and 22:26 on 12 August, and at 10:25 and 11:10 the next day.

August

Failure	*Site*	*IBM*	*External*	*Reason*
05.119	100%	0%	0%	Network firewall is a site responsibility

September

There were no failures in September

3.1.2 Performance Statistics

This section uses the definitions agreed in Schedule 7, ie,

- MTBF = (24 x 30.5)/(number of failures in month)
- Serviceability (%) = 100 x (WCT – SDT – UDT) / (WCT– SDT)

Attribution	*Metric*	*July*	*August*	*September*	*Quarterly*
IBM	Failures	1	0	0	1
	MTBF	732	∞	∞	2196.0
	Serviceability	99.7%	100.0%	100.0%	99.9%
Site	Failures	1	1	0	2
	MTBF	732	732	∞	1098.0
	Serviceability	99.7%	100.0%	100.0%	99.9%
External	Failures	1	0	0	1
	MTBF	732	∞	∞	2196.0
	Serviceability	99.1%	100.0%	100.0%	99.7%
Total	Failures	3	1	0	4
	MTBF	244	732	∞	549.0
	Serviceability	98.5%	100.0%	100.0%	99.5%

3.2 Capability Utilisation

Overall utilisation for the quarter was 67.9%, with capacity utilisation at 37.2% of the total.

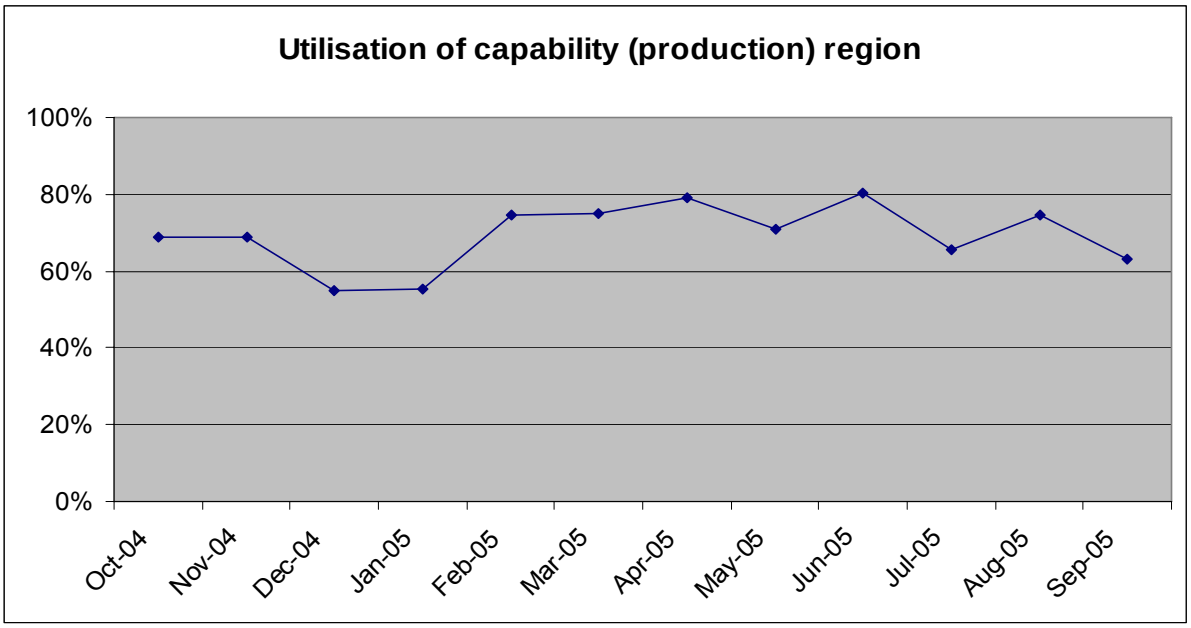

3.3 Capacity Planning

Predicted Utilisation

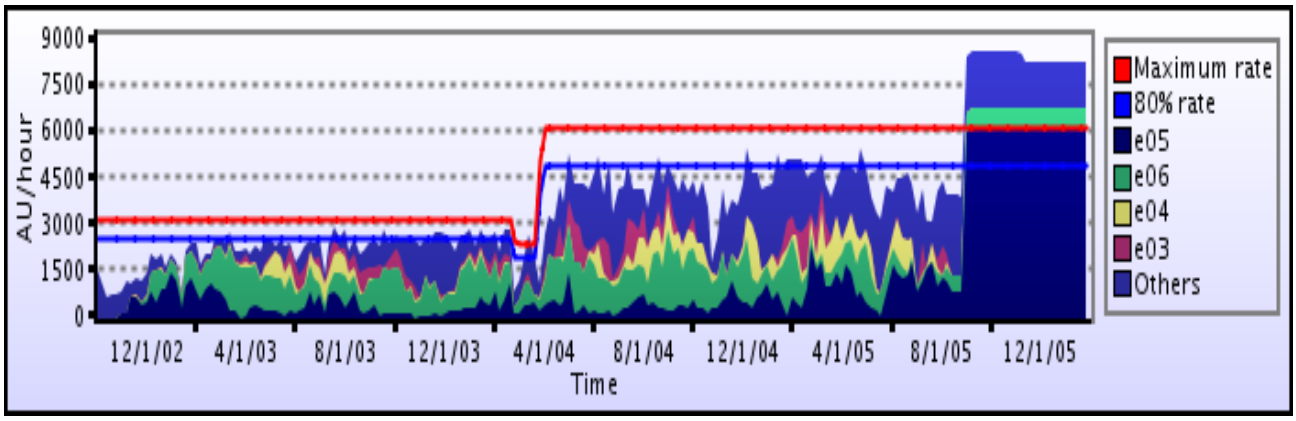

The graph above shows the utilisation since the start of the project and the projected utilisation until the end of 2005. The scale on the *y*-axis is AUs per hour, where the peak that HPCx Phase 1 could deliver was around 3240 AUs per hour, and Phase 2 6188 AUs per hour (the upper red line in the graph). The lower line (in blue) corresponds to the more practicable 80% level.

The graph assumes that each project will use its remaining allocation pro rata with its usage profile from the SAF, which is often simply that on the original application form.

The large block of time to the right of the graph corresponds to the large amount of time unused by the e05 Materials Chemistry consortium, which is currently scheduled to close at the end of April 2006.

Numbers of Research Consortia

There are currently 37 research consortia using the HPCx system. Five other projects have now been closed.

In addition, there are two externally funded projects.

3.4 CPU Usage by Job Size

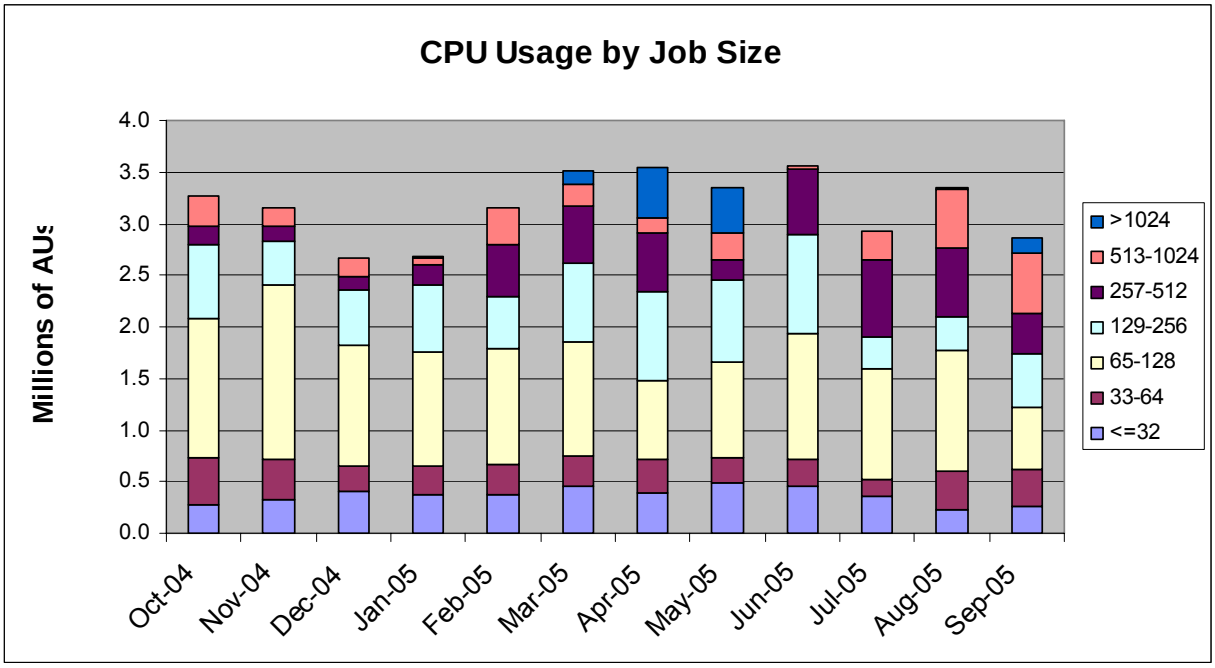

3.5 AU Usage by Consortium

The PIs and titles for the various consortia are listed in Appendix B.

Consortium	*July*	*August*	*September*	*Quarterly*	*%age*
e02	170988	87980	97932	356900	4.1%
e03	320	963	186932	188215	2.1%
e04	266529	612817	119088	998434	11.4%
e05	363762	1152764	953233	2469759	28.1%
e06	707838	274118	117296	1099252	12.5%
e07	878	2		880	0.0%
e08	44257	8247	7549	60053	0.7%
e10	16099	1	10655	26755	0.3%
e11	17707	3784	2394	23885	0.3%
e14	36305	9862	39862	86029	1.0%
e15	4	727	247	978	0.0%
e17	1415	5269	8489	15173	0.2%
e18	23093	28131	23644	74868	0.9%
e19	31208	45429		76637	0.9%
e20	47998	44712	146508	239218	2.7%
e24	2769	13077	20666	36512	0.4%
e25	18	31	1	50	0.0%
e26	1	2473	179	2653	0.0%
e27	536	159	2544	3239	0.0%
e28	10	21008	15474	36492	0.4%
e29	4395	1	0	4396	0.1%
e31	2	3737	1197	4936	0.1%
z09	20631	63713	116485	200829	2.3%
EPSRC Total	1756765	2379003	1870375	6006143	68.5%
n01	476859	263707	260145	1000711	11.4%
n02	328069	330883	148324	807276	9.2%
n03	42829	112457	103255	258541	2.9%
n04	118167	20627	26129	164923	1.9%
NERC Total	965924	727673	537853	2231450	25.4%
p01	3999	6	3244	7249	0.1%
PPARC Total	3999	6	3244	7249	0.1%

c01	132889	61614	49541	244044	2.8%
CCLRC Total	132889	61614	49541	244044	2.8%

b02	17	4	3602	3623	0.0%
b03	12067	38751	17895	68713	0.8%
BBSRC Total	12083	38756	21497	72336	0.8%

x01	7064	14252	31781	53097	0.6%
x02				0	0.0%
External Total	7064	14252	31781	53097	0.6%

z001	57574	56550	26542	140666	1.6%
z002	828			828	0.0%
z004	2252	15562	54	17868	0.2%
z06	4	430	89	523	0.0%
HPCx Total	60658	72542	26684	159884	1.8%

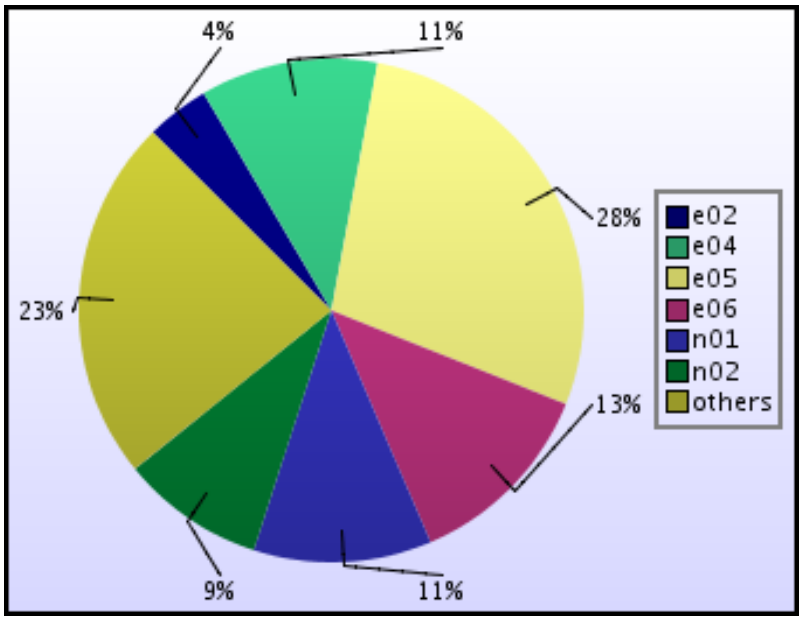

3.5.1 Discounts

There are now a number of user codes that have qualified for capability discounts. The following table shows the discounts that were awarded during the last quarter.

Consortium	*AUs Used*	*AUs Charged*	*Discount*
e04	1009103	998434	10668
e05	2842133	2469758	372375

3.6 Helpdesk

3.6.1 Classifications

Category	*Number*	*% of all*
Administrative	159	54.3
Technical	111	37.9
In-depth	23	7.8
PMR	0	0.0
TOTAL	293	100.0

Service Area	*Number*	*% of all*
Phase 2 platform	235	80.2
Website	7	2.4
Other/general	51	17.4
TOTAL	293	100.0

3.6.2 Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	227	84.1	75%
Finished within 72 Hours	269	99.6	97%
Finished after 72 Hours	1	0.4	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	155	97.5	97%
Finished after 48 Hours	4	2.5	

3.6.3 Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*	*PMR*
epcc.ed.ac.uk	131	49	15	0
dl.ac.uk	2	23	5	0
Sysadm	23	39	3	0
Other people	3	0	0	0

3.7 Service Quality Tokens

No quality tokens were set this quarter.

4 Support

4.1 Applications Support *(Dr David Henty)*

4.1.1 Documentation

This has been another stable quarter in terms of the hardware and software configuration of the machine, so again no major documentation changes have been required. However, some general maintenance has taken place: KOJAK, a suite of performance analysis tools, was installed and the User Guide updated accordingly; a presentation outlining the F2K features available under xlf9, presented at the User Group, is now available on the WWW.

4.1.2 Technical Reports

A total of four reports were planned for Q3 in the following areas:

a) Parallel Performance Analysis Tools
b) Comparison of Molecular Dynamics Codes for the Life Sciences.
c) Use of Scalable Eigensolvers: a Case Study
d) Data Handling on the Grid for Capability Computing

where **a)** is a report postponed from Q2.

We have produced the following three reports this quarter:

- **HPCxTR0508**: *Performance and Profiling of the LAMMPS Code on HPCx,* Fiona Reid and Lorna Smith.
- **HPCxTR0509**: *On the Performance of Molecular Dynamics Applications on* **Current High-End Systems, Joachim Hein, Fiona Reid, Lorna Smith, Ian** Bush, Martyn Guest and Paul Sherwood.
- **HPCxTR0510**: *Use of Scalable Eigensolvers: a Case Study*, Andy Sunderland, Gavin Pringle and Ian Bush.

Reports **08**, **09** and **10** correspond directly to the titles **a)**, **b)** and **c)**. It was decided to illustrate the use of parallel performance analysis tools via a practical case study, and the LAMMPS report demonstrates how MPITrace and VAMPIR were used on HPCx to understand the performance characteristics of a major software package on large numbers of processors. Report **09** is a substantial piece of work that was also published in the August 2005 issue of *Philosophical*

Transactions of the Royal Society of London Series A: Theme Issue on Scientific Grid Computing. Report **10** describes the optimisations made by the Terascaling Team to the parallel eigensolver stage of several application codes (including GAMESS-UK, PRMAT, PLATO, SIESTA and VASP). The parallel performance improvements obtained on HPCx are detailed in the report.

The work for report **d)** is well under way, but has taken longer than anticipated to complete. For this study we are undertaking a survey of data transfer between HPCx and other nodes on the National Grid Service. There was some delay involved in obtaining accounts on the other NGS nodes, but they have now all been set up and the report will be competed in Q4.

There are a total of twelve reports due for this year. The profile in the 2005 Annual Plan had ten reports due by the end of Q3, so we are currently on target.

4.1.3 Training

In Q3 of 2005 we ran the following six courses:

- **EPCC / Access Grid, 20 July**: *Task Farming on HPCx*
- **University of Leeds, 26 July**: *Using the HPCx Service*
- **University of Leeds, 27-29 July**: *Message-Passing Programming using MPI.*
- **EPCC, 22 August:** *Fundamental Concepts of HPC*
- **EPCC, 23-24 August:** *Object-Oriented Programming for HPC*
- **EPCC, 29 August – 1 September:** *Practical Software Development*

The courses at Leeds, which hosts an NGS node, were organised following comments received at the Life Sciences workshop, where it was mentioned that it was a convenient location for many users. This was confirmed by very good attendance figures for both courses.

Statistics are summarised below alongside annual targets (where appropriate):

Metric	Total	Target
Course days	21	30
Number of courses	11	12
Different locations	4	4
Student-days for HPCx users	324	
Student-days for HPCx staff	47	
Student-days available for HPCx	465	600

As predicted in the Q2 report, when the figures appeared low, the Q3 training schedule has brought this activity back on target for the year.

The following courses are scheduled for Q4:

- **EPCC, 25– 27 October**: *Message-Passing Programming using MPI*
- **EPCC, 24 November:** *Using the Visualisation Toolkit VTK*
- **Daresbury, 30 Nov–2 Dec**: *Computational Chemistry on HPCx*
- **EPCC, 1-day course:** *Parallel IO using MPI-2*
- **Daresbury, 1-day course:** *Compiling for the Power5 Processor*

where the dates for the last two courses are currently being finalised. With this programme, we will meet all our training targets for 2005.

4.1.4 Workshops and Conferences

The second HPCx workshop for this year is *Progress in Environmental Engineering: Is There a Role for High-End Computing?,* to be held in Daresbury on 21 November.

The Third Annual HPCx Seminar, *Capability Science on HPCx*, is taking place on 5 December at Daresbury, timed to directly precede the *16th Machine Evaluation Workshop*. The programme for the HPCx event is almost complete, and we are already taking registrations on the WWW.

4.1.5 User Group

A user group meeting was held over Access Grid on 20 July. A total of five sites participated, with some two dozen people taking part. Two presentations were given, on Task Farming and new features of the xlf9 compiler, followed by a general discussion. Minutes and copies of the presentations are available from the *events* section of the HPCx WWW site.

Despite some minor technical difficulties, this was a much more successful meeting than when we first used the Access Grid for a user group in April 2003. We plan to continue to hold one meeting per year in this manner.

4.1.6 Newsletter

The sixth issue of Capability Computing, *Modelling fusion on HPCx*, was produced in September as planned, and 3500 copies are currently being mailed out. We will also take copies to upcoming events such as Supercomputing and the Annual Seminar. Fortunately, there was just enough time to include a stop-press article announcing the Power5 upgrade to HPCx.

4.1.7 Packages

More than 50 pieces of software are now centrally supported under the packages mechanism, the most recent additions being KOJAK and UNICORE.

4.2 Outreach Activities *(Dr Richard Blake)*

4.2.1 Outreach to Lifesciences

Molecular Dynamics Codes

The work has been reported in more detail in the TeraScaling report.

Integrative Biology Project

Michael Holden has now successfully ported the Cardiac Arrhythmia Research Package (CARP) to HPCx, and the problems reported previously with PETSC have been solved.

Michael has also ported the Sheffield Cardiac Arrhythmia Model (SCAM). This is a C and OpenMP code that models heart tissue in a similar way to CARP but using different techniques. After porting the code to HPCx it was analysed and profiled. Further scope for parallelisation was identified along with some serial code optimisations. The effect of these changes was to reduce the program run time by up to 75%.

Genome Searches

Lorna Smith has been in contact with a group studying “Genome Wide Searches for RNA Localisation Motifs”. This group utilises two codes: AMBER and a shared memory code called MC-SYM. The group have estimated that their calculations would take 50 years executing serially, so are obviously keen to consider using HPCx. A small amount of time has been granted from Outreach, and we are working with them to investigate task farming MC-SYM and AMBER.

QDVE Project

The HPCx-based part of the project is drawing to a close. Marcus Durrant has now used most of his HPCx resource and has also left BBSRC JIC to work at Northumbria University. Collaboration has been re-established and we are starting to working with him to adapt the software for grid systems (e.g. NGS) and regional grid systems.

Computational Enzymology

A report on our experiences with the GAMESS-UK/CHARMM replica path capability has been passed to the Bristol group and the code, incorporating the bug fixes and enhancement is now being tested by that group on HPCx. There has been a considerable hiatus on this activity in that the PDRA who previously was working on the project has re-assigned to the Dynome project (see below). A planed visit to NIH has had to be rescheduled for November, to coincide with SC05 and the dynome demonstrator.

Dynome Demonstrator

HPCx staff supported this team (led by Jon Essex) in the preparation and planning of the bid to run a US/UK demonstration of HPC/Grid resources applied to large-scale biomolecular simulation and distributed data analysis (making us of BioSimGrid tools). The current status is that the project was partially funded and is preparing a somewhat scaled-down demonstration for SC 2005, working jointly with NAMD development team, Klaus Schulten on the scientific side and Rick Kufrin on the technical side. HPCx staff are now assisting with the experiments which will look into the execution of the grid-enabled NAMD (NAMD-G).

SPICE Demonstrator

Work has focussed on updating NAMD to support the pore simulation sub-project, to implement MPI-CH G2 for the arterial flow and turbulence modelling projects and to enable UK-Light access to HPCx.

CCP15 – a Collaborative Computation Project for Biomolecular Simulation

BBSRC has funded CCP15 to provided support to the biomolecular simulation community. The Project will seek to provide a representative voice for interacting with the Research Councils, coordinate research activities across cluster systems, high-end systems and e-science focussing in particular on training.

4.2.2 Other Outreach Activities

Financial Modelling

We reported previously on a group of financial modelers who had been given access to HPCx, as part of the outreach activity. This group has been able to demonstrate that they can make effective use of over a thousand processors. We will continue to work with them to develop a proposal for larger amounts of time.

LAser Materials Processing Simulation CONsortia (LAMPs-CON)

Dr Markus Gross (Cambridge University) visited EPCC on the 1st September to port and test the performance of his code on HPCx, for a future grant proposal. Joachim Hein worked with Markus and the code was successfully ported and tested. The proposal has been completed and we are currently finishing the technical assessment; the proposal will then be submitted.

Free Electron Laser Modelling

HPCx have been working with CCLRC's ASTEC staff to port a Free Electron Laser modelling code – GENESIS – to the Phase 2 system. 3D time dependent simulations have been undertaken to investigate the photon pulse quality of systems such as the 4th Generation light-source. The parallel MPI version of the code yielded immediate speed-ups of 100 on 256 processors. This compute capability makes it possible to simulate a Regenerative Amplifier FEL to saturation, in time dependent mode. A full Class 1 application is in preparation.

Industry

Astra Zeneca and Daresbury Laboratory have signed a three-year collaboration starting in October 2005 to design improved drug formulations. This includes a £50K budget for computer time. Commercial pricing for Fluent is still under discussion. Visits to industrial end-users and their suppliers are under way. Recent discussions have been held with Unilever Research Port Sunlight, Chimatica, BNFL and Scienomics.

4.2.3 Improve Public Awareness

Science Festival

We have again been accepted to present a talk at next year's Edinburgh International Science Festival. The talk will be given by David Henty, and is titled: Supercomputers: rise of the machines.

Cardiovascular Conference

Gavin Pringle presented a talk titled: HPC facilities in the UK and internationally at the Cardiovascular Haemodynamics and Modelling conference on the 25th – 27th September, 2005.

Science Week

Lorna Smith and Tracy Peet will attend a briefing session in Livingston in October, to consider the possibility of organising an event for the BA National Science week next year.

SRS 25th Anniversary

Daresbury Laboratory hosted some 80 VIPs around the SRS Annual User Meeting with some 40 attendees participating in a tour of the HPCx facility.

4.2.4 Visualisation

A high-performance compute, data and visualisation system is out to tender with responses expected in mid-October.

CCLRC is involved in a JISC UK Visualisation Support Network. Activities include generic training activities and access to the visualisation system at Daresbury Laboratory will be contributed as a resource. HPCx will focus on visualisation demonstrations of large data sets in collaboration with a number of Consortia.

4.3 Terascaling Applications *(Dr Martyn Guest)*

4.3.1 Capability Science

Consortium Visits

Martyn Guest wrote a paper based on the visit reports carried out so far this year and made a presentation of the findings to the STAC meeting on 28th Sep 2005. Further visits are planned for the final quarter.

Capability Incentives

The following codes have gained capability incentives in this quarter:

- PDVR3DRZ
- CP2K

A full list of codes with capability incentives can be found at:

http://www.hpcx.ac.uk/support/documentation/capability.html

4.3.2 Quantum Chemistry Codes

GAMESS-UK
Ian Bush's work on GAMESS-UK has revolved around two fronts. Firstly the implementation of the Restricted Hartree-Fock (RHF) version of the code that uses shared memory segments to store large, replicated objects has been completed, and the performance characteristics have been investigated. The use of shared memory segments requires a critical region in the Fock matrix build, and this could diversely affect performance. In practice it was found that though the effect was large for small molecules, for large ones the performance degradation was negligible, and in fact in some cases there was a very small improvement. This is probably due to improved Level 3 cache usage, as only one copy of the relevant matrices is now stored instead of up to 32.

The resulting code was used for calculations on a number of clusters taken from the isocitrate lyase protein, and calculations on molecules with over 12,000 basis functions were demonstrated. These results have been presented to the American Chemical Society. Prior to the work described above using shared memory segments, these calculations were possible but only 12 processors per node could be used before reaching the memory limit. The new code, which allows the full 32 processors per node to be utilised, shows a performance improvement of over a factor of two.

The second work on GAMESS-UK has concentrated on developing methods to help study why convergence problems are occurring, driven by experiences with calculations on the Rusticyanin molecule. We have implemented facilities to print the frontier orbitals and Mulliken and Lowdin populations at each iteration and also changes in these quantities. This is not as straightforward as might be thought due to the manipulation of distributed objects.

4.3.3 Computational Materials Codes

CASTEP
Martin Plummer has tested CASTEP using the xlf9.1 compiler to ensure a problem-free transition to general use of this compiler. Benchmark tests showed the new compiler either did not affect or slightly improved (~5%) performance. Following the previous quarter's work on optimising CASTEP restarts, a section of the code which writes data to disk was rewritten to streamline the MPI and avoid occasional problems with hanging reported by users (these problems were more systematic on other clustered systems but could in principle affect HPCx for very large jobs). The latest release of CASTEP (version 3.2) was made available by the CASTEP Developers Group and this is currently undergoing tests and customisation/optimisation for HPCx prior to its official release. Martin has also been working with Bartosz Dobrzelecki (EPCC) to characterise and optimise the performance of the path-integral quantum-MD intelligent task-farming routines.

CRYSTAL
Optimisation of the serial integral evaluation in CRYSTAL had introduced incompatibilities with the parallel code. Ian Bush has resolved these issues.

VASP
Kenton D'Mellow has investigated and profiled two versions of VASP on HPCx, in an attempt to increase scalability. VASP's scalability is in part determined by its collective communications, particularly MPI_Alltoall and MPI_Alltoallv (see earlier reports). An optimised version of MPI_Alltoall is available for 64bit applications, so in an effort to utilise this, he has been working on producing a stable 64-bit executable for version 4.6. In addition to this, a planned approach to MPI_Alltoallv had been devised by Stephen Booth (SPB), which significantly cuts overhead especially at small message sizes (<2k). This has been ported to Fortran for incorporation into VASP and other codes. With these optimised communications it is anticipated that a new optimum parallelisation value NPAR (as described by Gavin Pringle's preliminary report on VASP optimisation) may be reached, and allow significantly better scaling. In addition to this, it is anticipated that general code performance will improve at 64-bit. Current work is focused on extensively benchmarking this.

SIESTA
The work done on the SIESTA code in 2004 is being updated for a Journal publication.

4.3.4 Engineering Codes

UKAAC Consortium
Andrew Sunderland has completed an analysis of the parallel performance of ROTORMBMGP on HPCx and on a selection of other HEC platforms. A paper entitled *"Parallel performance of a UKAAC helicopter code on HPCx and other large-scale facilities"* has been submitted to the Parallel CFD 05 Proceedings for publication.

Fluent
Fluent on HPCx has been upgraded to version 6.2.16.

4.3.5 Physics Codes

Atomic & Molecular Physics Consortium
Andrew Sunderland has carried out further testing of the 'uniform energy grid' code and an xlf9.1 optimised binary has been supplied to users.

CENTORI
Joachim Hein has continued to work with the group from Culham to use an alternative parallelisation strategy to improve the scaling of the code. He has worked on the parallelisation of cadence, which is a simpler cut down code for CENTORI. He has experimented with replacing the present solver with another solver found in the ScaLAPACK library - results to date are inconclusive. For CENTORI the xlf9 compiler leads to a performance improvement of more than 20%.

ChemReact
PDVR3DRZ has obtained a gold capability incentive award, see above. This is due to the BLACS sub-context harnessing prototype supplied by the Andrew Sunderland.

4.3.6 Life Science Codes

DL_POLY
Ian Bush has compared the performance of the FFT in DL_POLY 3 with that provided by IBM in the PESSL library. The FFT in DL_POLY is slower than IBM's, but it fits directly onto the data distribution used in DL_POLY and therefore avoids an expensive data re-distribution, so in practice the two methods are comparable in time up to around 256 processors, after which the implementation in DL_POLY is markedly better dues to its intrinsic better scaling. Comparisons with FFTW have also been started.

Martin Plummer has successfully combined the DL_POLY2 code with a Perl script written by Andrew Sunderland to create a task-farming version which can

perform several simulations simultaneously. If required this code can be further adapted to be guided by a 'master' program to control the various farmed-out simulations. Martin is liaising with Ilian Todorov (DL) to allow similar adaptation of DLPOLY_3 (although this code is designed for larger molecular systems than DL_POLY2).

LAMMPS
A new million atom system has been supplied by Chris Greenwell. This system can run on both LAMMPS 2001 and 2005 allowing a direct (or as near as we can get) comparison between the two versions of the code. The benchmarks have been tested on HPCx and the performance and scaling investigated. The 2005 code is at least twice as fast as 2001 code. Installation of the 2005 code requires the inclusion of a number of patches - these will not apply on HPCx therefore the patching was carried out on Linux and the source then copied over to HPCx.

Peter Coveney/Chris Greenwell have invited Fiona Reid to be involved in the "Benchmarks for atomistic mineral simulations: a working group" which will be funded by the National Institute for Environmental eScience. There is likely to be a 3 day workshop/meeting for this in Cambridge during March/April 2006.

The LAMMPS technical report (reported in previous reports, but not made public) is now publicly available on the HPCx website (HPCxTR0508).

The Phil. Trans. paper reported previously (HPCxTR0509) is also now available on the HPCx website.

NAMD
NAMD is to be utilised in one of the SC demonstrator projects (see the Software engineering report). However NAMD version 2.5 would not run the users configuration across a number of platforms including HPCx. This problem has been fixed in version 2.6b1, which has since been installed on the service. However the user still experienced problems. Removing a number of assertions in the makefile enabled us to run the user-supplied configuration files. However the user has since reported new problems which are presently under investigation.

The performance of the NAMD version 2.6b1 on a standard benchmark has been assessed. The new version gives a performance improvement of about 12% when compared to the older version 2.5. The recent change of the C++ compiler from vac6 to xlc7 gives an additional improvement of 17%.

4.3.7 Compilers

All members of the Terascaling Team carried out tests of the new Fortran compiler, xlf9, prior to its roll out to users, as this was a major upgrade to the compiler. A number of codes were studied and no problems were found. For some codes a small performance improvement was found.

4.3.8 Tools

Paraver has been upgraded to version 3.4.

4.3.9 POWER5 benchmarking

IBM provided access over two days to a 16-processor single p5-575 node for benchmarking purposes. We ran the five acceptance test codes (AIMPRO, CASTEP, DL_POLY, H2MOL and PCHAN) plus a number of low-level kernels. The last in order to provide deeper understanding of the implications of the revised memory architecture compared to the p690. Results were reported to EPSRC as part of the scientific and technical case for the Phase2a upgrade.

4.3.10 Publications

Martin Plummer is coordinating a Terascaling Team submission to the Materials Chemistry Consortium (e05) special edition of the Journal of Materials Chemistry entitled High Performance Computing. Contributing authors are Joachim Hein, Martyn Guest, Kenton d'Mellow and Ian Bush. The paper describes the optimisation of a range of Materials Chemistry codes in order to exploit large-scale capability resources, such as HPCx. The codes included are CASTEP, CRYSTAL, DL_POLY, GAMESS-UK, NAMD, SIESTA and VASP. The article will be submitted in the next quarter (by the end of 2005) and we plan to include initial results from Phase2a in a final section to make the article as up-to-date as possible.

"Parallel performance of a UKAAC helicopter code on HPCx and other large-scale facilities", A.G. Sunderland, D.R. Emerson, and C.B. Allen, submitted to Parallel CFD 05 Proceedings.

4.3.11 Presentations

An Overview of HPCx was provided as part of the presentation*, The UK's e-Science Programme; Perspectives on Science Delivery,* J. Gordon and M.F. Guest, CAS2K5 Workshop, Annecy, 11-15th September, 2005

4.4 Software Engineering *(Dr Stephen Booth)*

4.4.1 General Terascaling and Optimisation techniques

We have been investigating the deployment of our planned AlltoAllv library real application codes. We are currently investigating the use of this library in the VASP code. As part of this investigation we have also developed Fortran bindings for this package.

4.4.2 Advanced data handling and Grid computing

The Globus installation on HPCx has been updated to version 2.4.3. This is the most recent version of the Globus-2 software. This version includes the MPI Globus flavours that will be required by the SPICE project for demonstrations at Supercomputing 2005. The SPICE project involves closely coupled meta-computing between HPCx and machines in the US TeraGrid. This is particularly difficult to achieve with HPCx as the compute nodes of HPCx are on a private network and are not directly connected to the internet. To work around this problem we have been developing a new port forwarding system to allow jobs running on the HPCx backend processors to make and receive arbitrary socket connection. This system is based on the standard socks protocols and requires no source code modification to application. The only requirement is that an additional library is specified when the application is built. This should make the new system more useful and easier to use than the port forwarder we developed for SC-2003. A prototype version of this system is now complete and being tested.

A number of research groups now want to transfer large amounts of data to and from HPCx. These include:

- Integrative Biology
- HiGEM

We have started work on an investigation of the relative performance of the different mechanisms of transferring data to and from HPCx. This will include and investigation of data transfers over the new UK-Light connection. A report on this work should be available next quarter.

Future trends

We have investigated the portability and performance of a number of important application codes to the BlueGene system, working with the developers when needed. These include:

- OCCAM
- CASTEP
- DLPOLY3
- H2MOL
- NAMD
- LAMMPS
- LUDWIG
- PCHAN
- MDCASK

In addition we have been running some low-level benchmarks such as STREAMS. Much of this work including performance comparisons between HPCx and BlueGene on a number of important codes was presented at a joint QCDOC/BlueGene workshop in Edinburgh and is available top HPCx users here: *http://www.hpcx.ac.uk/research/hpc/presentations/Hein_bluegene.pdf*

4.5 Operations and Systems *(Mr Mike Brown)*

4.5.1 Phase2A upgrade

Much effort has been expended in planning the upgrade from the existing Phase2 service to the POWER5-based Phase2A service.

There has been a close interaction with IBM on both the implementation and on the transition design, with a principal aim being that of making the migration as seamless, and risk-free, as possible, whilst minimising downtime.

The required alterations to the room's power distribution system were undertaken in September, with the delivery of Phase2A equipment planned for early October.

4.5.2 Incidents

SEV1 incidents (contractual failures) numbered 4 this quarter, unchanged from Q2. The number of non-failure incidents, at 26, was significantly lower than the rest of the year.

4.5.3 Outreach

A meeting of the IBM UK HPC User Group was held at the IBM plant in Montpellier during September.

This presented an opportunity to view some of the HPC manufacturing capability, as well as to meet senior figures from IBM p-series support.

4.6 Staffing

AV	*July*	*August*	*September*
DL	5.8	4.9	5.7
EPCC	8.6	10.0	8.4
Total	14.4	14.9	14.1

Systems	5.2	5.4	6.0

5 Summary of Performance Metrics

Metric	*TSL*	*FSL*	*July*	*August*	*September*
Technology serviceability	80%	99.2%	99.7%	100.0%	100.0%
Technology MTBF (hours)	200	300	732	∞	∞
Number of AV FTEs	7.5	10	14.4	14.9	14.1
Number of training days per month	22.5/12	30/12	13/7	20/8	20/9
Non in-depth queries resolved within 3 days	85%	97%	99.3%	100.0%	100.0%
Number of A&M FTEs	3.75	5.75	5.2	5.4	6.0
A&M serviceability	80%	99.6%	99.7%	100.0%	100.0%

Colour	*Meaning*
	Exceeds FSL
	Between TSL and FSL
	Below TSL

Note 1: The number of training days is reported as a running total since the start of the year.

Note 2: The above table includes the revised FSL targets for *training days* and *A&M serviceability*, which have been agreed with EPSRC.

Appendix A: Incident Severity Levels

SEV 1 — anything that comprises a FAILURE as defined in the contract with EPSRC.

SEV 2 — NON-FATAL incidents that typically cause immediate termination of a user application, but not the entire user service.

The service may be so degraded (or liable to collapse completely) that a controlled, but unplanned (and often very short-notice) shutdown is required or unplanned downtime subsequent to the next planned reload is necessary.

This category includes unrecovered disc errors where damage to filesystems may occur if the service was allowed to continue in operation; incidents when although the service can continue in operation in a degraded state until the next reload, downtime at less than 24 hours notice is required to fix or investigate the problem; and incidents whereby the throughput of user work is affected (typically by the unrecovered disabling of a portion of the system) even though no subsequent unplanned downtime results.

SEV 3 — NON-FATAL incidents that typically cause immediate termination of a user application, but the service is able to continue in operation until the next planned reload or re-configuration.

SEV 4 — NON-FATAL recoverable incidents that typically include the loss of a storage device, or a peripheral component, but the service is able to continue in operation largely unaffected, and typically the component may be replaced without any future loss of service.

Appendix B: Projects

B.1 Current Projects

EPSRC Projects

Code	*Class*	*Title*	*PI*
e02	1	Ab-initio simulation of covalently bonded materials	Dr Patrick Briddon
e03	1	Multi-photon, electron collisions and BEC HPC consortium	Prof Ken Taylor
e04	1	Chemreact Computing Consortium	Prof Jonathon Tennyson
e05	1	Materials Chemistry using Terascaling Computing	Prof Richard Catlow
e06	1	UK Car-Parrinello Consortium	Prof Paul Madden
e07	2	Turbulent Plasma Transport in Tokamaks	Dr Colin M Roach
e08	2	Organic Solid State	Prof Sarah Price
e10	1	Reality Grid	Prof Peter Coveney
e11	1	Bond making and breaking at surfaces	Prof Sir David A King
e14	1	Blade and Cavity Noise	Prof Neil Sandham
e15	2	CSAR/HPCx Collaboration	Dr Mike Pettipher
e16	1	Cardiac virtual tissues	Prof Arun V Holden
e17	1	Integrative Biology	Dr David Gavaghan
e18	1	DARP: Highly swept leading edge separations	Prof Michael A Leschziner
e19	1	Edinburgh Soft Matter and Statistical Physics Group	Prof Michael E Cates
e20	1	UK Applied Aerodynamics Consortium	Dr Ken Badcock
e21	1	Intrinsic Parameter Fluctuations in Decananometer MOSFETs	Prof Asen M Asenov
e22	1	Preconditioners for finite element problems	Prof David J Silvester
e23	1	Exploitation of Switched Lightpaths for e-Science Applications	Prof Peter Clarke

e24	1	DEISA - Distributed European Infrastructure for Supercomputing Applications	Dr David Henty
e25	1	Turbulent vortex motion in stratified flows	Dr Gary Coleman
e26	1	Simulation of Radioprobing	Dr Charlie Laughton
e27	1	SPICE	Prof Peter V Coveney
e28	1	Towards the Dynome	Dr Jonathan W Essex
e29	1	Free-surface-piercing circular cylinders	Dr Eldad Avital
e30	1	Metal/Oxide Interfaces at the Atomic Level	Dr Nora de Leeuw
e31	1	Lateral Straining of Wall-Bounded Turbulence	Dr Gary N Coleman
e32	1	Rapid Prototyping of Usable Grid Middleware	Prof Peter V Coveney
z09		HECToR Benchmarking	Dr Edward Smyth

PPARC Projects

Code	*Class*	*Title*	*PI*
p01	1	Atomic Physics and Astrophysics	Prof Alan Hibbert

NERC Projects

Code	*Class*	*Title*	*PI*
n01	1	Large-Scale Long-Term Ocean Circulation	Dr David Webb
n02	1	NCAS	Prof Alan J Thorpe
n03	1	Computational Mineral Physics Consortium	Dr John Brodholt
n04	1	Shelf Seas Consortium	Dr Roger Proctor
n05	2	Non-linear Wave-particle Instabilities in Plasmas	Dr Mervyn Freeman

BBSRC Projects

Code	*Class*	*Title*	*PI*
b02	1	Modelling enzyme catalysis	Dr Adrian J Mulholland
b03	1	Towards a virtual outer membrane	Prof Mark S Sansom
b04	1	Life sciences software development	Dr Jo L Dicks
b05	1	Virtual forced evolution of catalytic transition metal complexes	Dr Marcus Durrant
b06	2	Biomolecular computational chemistry	Prof Jonathan D Hirst

CCLRC Projects

Code	*Class*	*Title*	*PI*
c01	1	Daresbury Laboratory Facilities Agreement Consortium	Dr Richard J Blake

Externally-funded Projects

Code	*Title*	*PI*
x01	HPC-Europa	Dr J-C Desplat
x02	OHM Ltd	Mr Mark Westwood

HPCx Projects

Code	*Title*	*PI*
z001	HPCx Support	Dr Alan Simpson
z002	Systems and Operations	Mr Mike Brown
z003	Test Project	Dr Denis Nicole
z004	HPCx Training	Dr David Henty
z05	Outreach Projects	Dr Richard Blake
z06	Application Porting	Dr David Henty
z07	Package Installation	Dr Mike Ashworth

B.2 Former Projects

Code	*Class*	*Title*	*PI*
b01	2	Quantum Chemistry Studies of the Rusticyanin Protein Crystal	Prof Samar Hasnain
e01	1	UK Turbulence Consortium	Prof Neil Sandham
e09	2	Molecular Properties and their Geometry	Prof Peter Taylor
e12	1	Parallel programs for the simulation of complex fluids	Dr Mark R Wilson
e13	1	TeraGyroid project	Dr Richard J Blake