

HPCx Quarterly Report

Jan-Mar 2003

1 Introduction

This report covers the period from 1 January 2003 at 0800 to 1 April 2003 at 0800, although, where appropriate, data is also reported for December 2002.

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. Section 4 includes reports on progress against the Annual Plan from each of the functional teams. A summary table of the key performance metrics is in Section 5. The Appendices define the incident severity levels and list the current HPCx projects.

2 Executive Summary

- Utilisation of the service has risen consistently over the last few months and is now around 80% on the capability region.
- The modal job size is 128 CPUs and more than a third of the utilisation was from jobs of at least 256 CPUs.
- However, the number of incidents remains high and the number of failures even appears to be increasing. Correspondingly, the MTBF figures are significantly short of the agreed targets. However, the serviceability figures have been kept reasonable thanks to the efforts of the Systems team.
- Interactions with the users have been good:
 - the helpdesk is exceeding all its targets;
 - the training programme is well under way;
 - the numbers of AV and A&M FTEs have been high;
 - the outreach to lifesciences has identified a number of projects due to start in July;
 - there is a user group arranged for April and the first newsletter is due in May.
- The Terascaling team has begun work on a good range of applications codes and has had a number of successes.

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:

	December	January	February	March
Incidents	16	29	22	30
Failures	3	5	6	11

The number of incidents has been consistently high and the apparent trend in the number of failures is clearly worrying.

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

December

<i>Failure</i>	<i>Site</i>	<i>IBM</i>	<i>External</i>	<i>Reason</i>
02.012		100%		Loss of primary login server is responsibility of technology supplier
02.013		100%		Loss of primary login server is responsibility of technology supplier
02.023		100%		Accessibility to LPARs is responsibility of the technology supplier

January

<i>Failure</i>	<i>Site</i>	<i>IBM</i>	<i>External</i>	<i>Reason</i>
03.009	50%	50%		Undocumented dependency on CWS by GPFS Site decision to take CWS down during service time
03.013		100%		GPFS is the responsibility of the technology supplier
03.020	100%			Uninterrupted access to the external network is site responsibility
03.021		100%		Accessibility to LPARs is responsibility of the technology supplier
03.023	100%			Uninterrupted access to the external network is site responsibility

February

<i>Failure</i>	<i>Site</i>	<i>IBM</i>	<i>External</i>	<i>Reason</i>
03.035		100%		Undocumented dependency on CWS by GPFS
03.036				Re-occurrence of 03.035 during IBM investigation
03.038			100%	Loss of JANET connection external to site

03.040	100%			Uninterrupted access to the external network is site responsibility
03.043		100%		Availability of GPFS is technology supplier responsibility
03.046		100%		Operation of switch fabric is technology supplier responsibility
03.049		100%		Network interface on system is technology supplier responsibility

March

<i>Failure</i>	<i>Site</i>	<i>IBM</i>	<i>External</i>	<i>Reason</i>
03.053		100%		Availability of GPFS & switch is technology supplier responsibility
03.054		100%		Availability of GPFS & switch is technology supplier responsibility
03.055		100%		Operation of switch fabric is technology supplier responsibility
03.056		100%		Availability of GPFS & switch is technology supplier responsibility
03.057	100%			Uninterrupted access to the external network is site responsibility
03.059		100%		Availability of GPFS & switch is technology supplier responsibility
03.069	100%			Operations error
03.075		100%		Ability to submit jobs is technology supplier responsibility
03.076				Re-occurrence of 03.075 within 12 hours counts as single failure
03.077		100%		Ability to submit jobs is technology supplier responsibility
03.078		100%		Debugging applied under instructions from IBM
03.053		100%		Availability of GPFS & switch is technology supplier responsibility

Notes

These are now the official record of these incidents. There have been a small number of changes since the original monthly reports:

- 02.014 and 02.015 were re-classified as SEV 2 as exactly 10% of the system was down which does not constitute a failure under Schedule 7 of the contract.
- 03.036 was merged with 03.035 as it was a re-occurrence of the same failure within 12 hours.

3.1.2 Performance Statistics

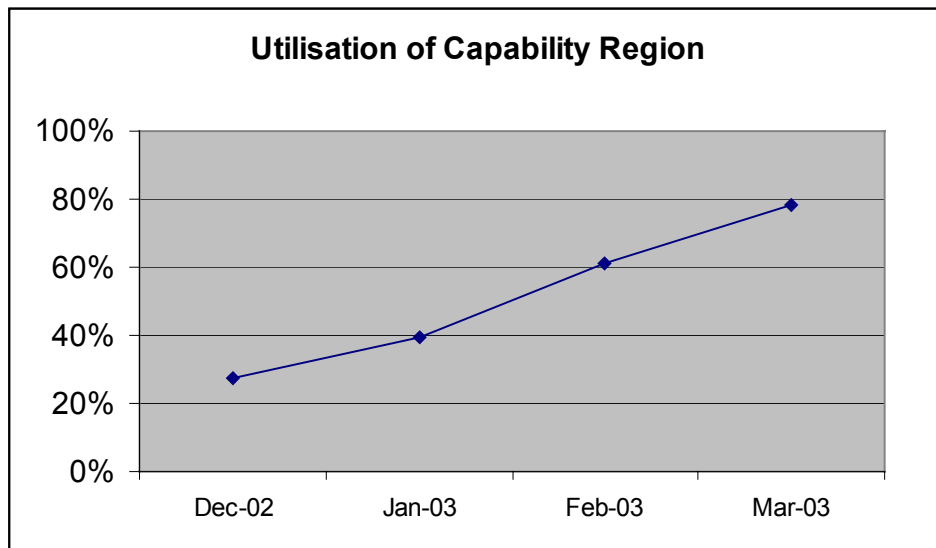
This section uses the definitions agreed in Schedule 7, ie,

- $MTBF = (24 \times 30.5) / (\text{number of failures in month})$
- $\text{Serviceability (\%)} = 100 \times (WCT - SDT - UDT) / (WCT - SDT)$


<i>Attribution</i>	<i>Metric</i>	<i>December</i>	<i>January</i>	<i>February</i>	<i>March</i>	<i>Quarterly</i>
IBM	Failures	3	2.5	4	9	15.5
	MTBF	184	293	183	81	142
	Serviceability	98.9	99.6	98.0	96.8	98.2
Site	Failures	0	2.5	1	2	5.5
	MTBF	∞	293	732	366	399
	Serviceability	100.0	99.4	99.6	99.9	99.6
External	Failures	0	0	1	0	1
	MTBF	∞	∞	732	∞	2196
	Serviceability	100.0	100.0	99.9	100.0	99.9
Total	Failures	3	5	6	11	22
	MTBF	184	146	122	66	100
	Serviceability	98.9	99.0	97.5	96.8	97.8

3.2 Capability Utilisation

The monthly utilisation for the 1024-processor capability region is shown in the graph below. This has increased month on month and was almost 80% in March.



3.3 CPU Usage by Job Size



The above graph shows that the modal job size is currently 128 CPUs. However, the relative utilisation by large job sizes is increasing and, during this quarter, more than a third of the utilisation was from jobs of at least 256 CPUs.

3.4 CPU Usage by Consortium

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


Consortium	December	January	February	March	Quarterly	%age
e01	2031.0	53795.4	44490.2	6212.7	104498.3	7.70%
e02	14029.6	39459.6	58424.1	36649.0	134532.7	9.91%
e03		20.5	49.9	949.4	1019.8	0.08%
e04	996.5	3667.8	5833.6	6580.3	16081.7	1.18%
e05		146208.6	216410.2	351737.4	714356.2	52.61%
e06		9765.2	164842.1	179079.0	353686.3	26.05%
EPSRC Total	17057.1	252917.1	490050.1	581207.8	1324175.0	97.52%

n01			2.0	3997.8	3999.8	0.29%
n02				19.4	19.4	0.00%
n03				218.5	218.5	0.02%
NERC Total			2.0	4235.7	4237.7	0.31%

z001	151931.5	118797.2	18947.8	9128.0	28075.8	2.07%
z002	134.1	38.0	501.3	41.6	542.9	0.04%
z004		140.6			140.6	0.01%
z05				771.0	771.0	0.06%

z06				4.4	4.4	0.00%
HPCx Total	152065.6	118976.0	19449.1	9945.0	29394.1	2.16%

Note: EPSRC have agreed that HPCx usage during January will not count against HPCx's annual allocation and so this usage is not included in the quarterly totals in the table above.



3.5 Helpdesk

3.5.1 Classifications

<i>Category</i>	<i>Number</i>	<i>% of all</i>
Administrative	150	40.4
Technical	199	53.6
In-depth	22	5.9
TOTAL	371	100.0

<i>Service Area</i>	<i>Number</i>	<i>% of all</i>
Phase 1 platform	304	81.9
Website	32	8.6
Other/general	35	9.4
TOTAL	371	100.0

3.5.2 Performance

<i>All non-indepth queries</i>	<i>Number</i>	<i>%</i>	<i>Target</i>
Finished within 24 Hours	286	81.9	75%
Finished within 72 Hours	343	98.3	97%
Finished after 72 Hours	6	1.7	

<i>Administrative queries</i>	<i>Number</i>	<i>%</i>	<i>Target</i>
Finished within 48 Hours	146	97.3	97%
Finished after 48 Hours	4	2.7	

3.5.3 Experts Handling Queries

<i>Expert</i>	<i>Admin</i>	<i>Technical</i>	<i>In-Depth</i>
epcc.ed.ac.uk	86	64	9
dl.ac.uk	6	35	4
Sysadm	58	91	6
Other people		9	3

3.6 Service Quality Tokens

<i>Date</i>	<i>Person</i>	<i>Value</i>	<i>Comment</i>	<i>Status</i>
26-Mar-2003 15:47:55	<i>Prof Neil Sandham</i>	•••	Reporting still not working - as PI, I seem to have no way of comparing allocation to researchers with actual usage	This problem has now been fixed and the user has been informed
05-Mar-2003 17:17:03	<i>Dr Fernando Bresme</i>	**		
25-Feb-2003 11:09:14	<i>Dr David Webb</i>	**	Help with quickly setting up accounts.	
06-Dec-2002 12:23:11	<i>Dr Dario Alfe</i>	****	Fantastic machine and excellent service, almost everything is perfect. It would be everything if I could solve my memory problem. It is somewhat disappointing that a job capable of running on a sun-sparc with 16 processors runs out of memory on HPCx.	The memory problems have been discussed and a solution proposed

4 Support

4.1 Applications Support (*Dr David Henty*)

During the first months of full operational service we have begun the process of building up a good working relationship with users and understanding their requirements. This has, overall, gone extremely well and has benefited significantly from the high quality of the SAF interface to common administrative operations such as account creation and time allocation. The full automation of these tasks over the WWW has meant that there have been few teething problems for users and the applications support team has been investigating important, long-term user issues right from the start of service.

4.1.1 Helpdesk

The helpdesk has operated extremely smoothly, and the statistics above show that we are meeting our targets for answering user queries. Common problems are being identified and additional information placed in the User Guide or the online FAQ as appropriate.

4.1.2 Documentation

The HPCx User Guide is the main source of information and it is continually updated and revised. It is currently at version 1.1, and is available both in online and printable formats. We have also made substantial amounts of IBM documentation available from the WWW pages, and have many local copies to avoid long download times for users.

4.1.3 Technical Reports

Of the eight technical reports identified in the Annual Plan, we are actively working on five of these (*Scaling of Lattice Codes, Java, Single-sided Communications, Fourier Transforms and Parallel IO*) and plan to have them all completed by the next quarterly report. We currently have preliminary versions of the first two available on the WWW, with the work on Lattice Codes having been presented at the ScicomP7 meeting in Germany in March.

4.1.4 Training

Of the two new courses proposed in the Annual Plan, we have already run one of these (DL_POLY on HPCx) and are beginning development of the optimisation course. In collaboration with IBM we were able to put on an in-depth course "Effective Exploitation of the HPCx System" in January, presented by Charles

Grassl and Farid Parpia who are two of IBM's HPC specialists from the US. We have worked closely with the systems team to ensure that those attending HPCx courses have full and easy access to the system via special course accounts that have privileged access to specifically tailored LoadLeveler queues.

The quarterly training statistics were as follows:

Course days	17
Number of courses	7
Different courses	7
Student-days for HPCx users	71
Student-days for HPCx staff	49
Student-days available for HPCx	226

4.1.5 Workshops / Conferences

HPCx made a presentation to the user meeting at Scicomp7, asking for Scicomp9 to be held at EPCC in 2004. We were unsuccessful in that bid (the 2004 European event will be at Cineca in Italy) but did secure Scicomp11 for 2005. To fulfill our commitments to international conferences we will be holding a one-day meeting at Daresbury in early December 2003. In order to maximise attendance it will be held just before the well-established Machine Evaluation Workshop, a strategy that worked well for the 2002 UKHEC Annual Seminar. The first of the two planned Workshops will be held in Q2, and the second in Q3.

4.1.6 User Group

The first HPCx User Group will take place on 24th April over Access Grid, with venues at EPCC, Daresbury and Imperial College pre-booked. At this meeting we expect to ask for volunteers for a chair from the user community, and will canvass opinion regarding possible joint user meetings with the CSAR service.

4.1.7 Newsletter

Production of the first newsletter is well underway and we plan to publish it in May. The second of this year's newsletters will be available six months later in early November, available in time for distribution at Supercomputing 2003.

4.1.8 Packages

The list of currently supported packages, tools and libraries on HPCx includes: Pchan, POLCOMS, Aimpro, Crystal, DLPOLY, GAMESSUK, CASTEP, AMBER, NWChem, CFX, H2MOL, NEWT, Vampir, Paraver and ScaLAPACK. Each has a nominated member of staff to maintain the software with an associated package

account. Optimisation of these packages for HPCx is being undertaken as part of the Terascaling activity.

In addition to the continual work to ensure that each installation is up to date, we have put substantial effort into ensuring that the user interface to ScaLAPACK is as clean as possible. This is complicated on the HPCx system by the many possible compile-time options for software (eg 32 or 64-bit, threaded or unthreaded) and the interaction with the IBM versions of selected ScaLAPACK routines available in the ESSL libraries. However, we understand all the issues and aim to provide the user with a single library that can be used in all circumstances.

4.2 Outreach (*Dr Richard Blake*)

To date there has been significant progress in a number of the outreach activities:

- EPSRC issued a 'Call for Proposals' for applications from the engineering and physical sciences community researchers that are not currently users of HPC facilities. The proposals are currently being reviewed within EPSRC. The HPCx Added Value team worked with a number of applicants (Mark Wilson - Durham and the Plasma Modelling Group at Culham) to port their codes to the HPCx system, explore salacity and advise on resource requirements for the proposal. Effort will be allocated from the Added Value team to support these new projects once a resource plan has been agreed with EPSRC.
- The HPCx service has actively progressed the IBM Lifesciences programme. Working with BBSRC, a number of researchers from the lifesciences community have been invited to submit proposals. These have been reviewed by IBM and funding for staff at a number of university sites and at Daresbury Laboratory and EPCC approved. The projects cover the areas of computational enzymology, molecular dynamics simulations of virtual outer membranes, simulation of the evolution of gene sequences and modelling of the radiation damage. Contracts, workplans and mechanisms for allocating time to these projects are currently being discussed. The effort supported by IBM increases the size of the Added Value team to 12 FTEs for the two years of the service commencing July 2003 thus meeting the performance related aspects of the contract. These two staff members at the centre must clearly work on lifescience projects approved of by IBM and the Research Council providing the computing time. This is a strong boundary condition on the outreach resource.

- HPCx followed up its initial programme of supporting a number of BBSRC projects through the IBM Lifesciences programme by participating in a workshop on 26th-27th of March which has been organised by Oxford University and CLRC to explore a strategy for molecular simulation for the BBSRC community. A strong recommendation from this meeting was the establishment of a CCP to support the broad base of BBSRC computing applications. The CCP may well identify a number of applications requiring access to the highest performance systems which will develop HPC Consortia bids in the fullness of time.
- Plans are being developed for a 'HPCx Industry Day' on 30th September 2003. The meeting will seek to demonstrate the capabilities of HPCx class systems in addressing key computational science and engineering research challenges in industrial research applications in particular in the materials, chemicals, lifescience and environment areas.
- Work has started on reviewing the international portfolio of HPC applications in particular those being considered as part of the Blue Planet project.

4.3 Terascaling Applications (*Dr Martyn Guest*)

The work described below covers the period January-March 2003, and details evaluation and development terascaling activities around application codes, libraries and tools, plus details of staff training, and attendance at Consortium meetings and associated events, including presentations by members of the Terascaling Team.

4.3.1 Computational Materials

Castep

- Set up Newtep licensed-user executable and data files in `/usr/local/packages/castep`. General liaison with A. Wander and K. Refson to broadcast information and comeback to licensed users.
- Set up and re-ran castep4.2 benchmark achieving best ever time of 402s.
- Newtep development: memory profiling of Newtep. Memory optimisation for cache-efficiency. Performed large-scale Newtep initial benchmarks and basic time profiling. Large swathes of redundant memory were removed plus accompanying improvements in serial code and use of communication collectives. Detected a convoluted MPI bug in Newtep. Continued study of structure of code for further optimisation.

- Performed tests of OpenMP/esslsm and mixed mode parallelism. Generally full MPI is better (without overuse of gather/reduce etc) but if code structure does not allow 8 tasks per node then esslsm can improve things with the correct number of threads selected.

Crystal

- Initial implementation and subsequent tuning of the distributed data version of the Crystal code on HPCx.
- Continued to eliminate excess memory usage in Crystal. This involved both tidying the code used to run Crambin, making it more general purpose (updating the main CVS repository) and also developing new methods to eliminate, distribute or calculate on the fly some of the remaining large replicated arrays.
- Performed Benchmarks runs of Crambin at STO-3G, 6-31G and 6-31G* levels on HPCx with comparison run on green.
- With the memory requirements now scaling linearly with the problem size (thus allowing appreciably bigger problems to be addressed), investigations focused on the possibility of running Rusticyanin, which is of interest to protein crystallographers. This is somewhat bigger than Crambin, the number of basis functions for the three different basis sets being (numbers for Crambin, our previous biggest case, are shown below for comparison):

<i>Basis</i>	<i>Rusticyanin</i>	<i>Crambin</i>
STO-3G	13906	3948
6-31G	25420	7194
6-31G**	43750	12354

- Further as this protein contains copper(II) unrestricted runs are required, making the problem bigger still. STO-3G and 6-31G have been demonstrated to fit on the machine (but not yet run to convergence), 6-31G** is too big. 6-31G* will be investigated.
- The studies of rusticyanin showed that at these problem sizes some of the start up routines were taking an unacceptably long time to run. All but one of these has now been parallelized, reducing the start up time to 8 minutes from roughly 50 for 6-31G Rusticyanin. Once the last routine has been parallelized (which is not a trivial job), we expect the start up to take roughly 1 minute.

GW Space

- The present serial and parallel versions of the code have been ported to HPCx, along with the associated libraries, and the scaling of the parallel code investigated. The code presently does not scale very well and is not memory-efficient.

- Although poor, the scaling of the HPCx code appears to be comparable with the scaling on other machines (e.g. CSAR T3E and a Beowulf).
- Currently writing a new more memory-efficient version of the code, in consultation with the authors at York.

4.3.2 Molecular Simulation

DL_POLY

- Optimisation of part of the start up procedure.
- W. Smith and I.J. Bush prepared and presented a DL_POLY training course at Edinburgh

CPMD

- Initial version of the code obtained from NCSA and implemented on HPCx with assistance from Klaus Schulten's group (Illinois)
- Scalability tests performed across 512 CPUs. Larger benchmarks under construction

Plato

- The code has now been successfully ported to the IBM.
- Met with author to discuss the details of the code and obtain actual simulation to time.
- Installed ARPACK. Note, despite ARPACK claims to the contrary, Plato still requires both LAPACK version 2 and LAPACK version 3 to compile (and hence run).
- Investigated relevant libraries, required for the code. This included public libraries LAPACK V3, LAPACK V2 (from ARPACK), BLAS (public and ARPACK), ScaLAPACK and IBM's own numerical libraries (BLAS (again), ESSL, ESSL_SMP, PESSL, PESSL_SMP).
- Currently profiling the code using the 3 out of 4 serial diagonalisation techniques in the original version of the code.

AMBER

- Compiled and run the latest version of the main executable of this Molecular Dynamics package. The code had already been ported to the HPCx machine.
- Scaling (with no optimisation or enhancements to aid tera-scaling) appears to be very poor. Scaling is improved on larger numbers of processors by underfilling LPARs, but the scaling is still poor.
- Currently profiling the code.

4.3.3 Atomic and Molecular

H2MOL

- Investigated performance gains associated with non-blocking message passing. Passed on optimised H2MOL routines to A&M consortium members.

PFARM

- Ported the external region code (PFARM) to HPCx, validated results and tested performance. The code scales fairly well on HPCx, but it is designed for fully distributed memory architectures. PFARM was evaluated further through new large-scale Fe²⁺ problems including 272 target states and 600 channels; previous calculations on the Cray T3E had a maximum of 400 channels.
- Aspects of restructuring to take advantage of the HPCx architecture are currently being tested (e.g. by maximising the usage of local shared memory).
- A new version of PFARM is being developed which will allow switching of eigensolvers between ScaLAPACK and Peigs.
- 64-bit versions of the A&M codes have been installed on HPCx. This raised some issues surrounding the promotion of long integers in 64-bit mode C code.
- The internal region code RMat2 was ported to HPCx.

4.3.4 Molecular Electronic Structure

GAMESS-UK and NWChem

- Both GAMESS-UK and NWChem have been implemented on HPCx, based on the Global Array tools from PNNL and IBM's underlying LAPI communication library.
- Extensive testing and performance evaluation of the LAPI-based implementations, with performance comparisons of HPCx and other high-end platforms.
- Generation of a number of kernel-based benchmarks (e.g. matrix multiply, diagonalisation) to provide a direct comparison between MPI- and LAPI-based applications. These clearly show major performance problems when comparing ScaLAPACK/MPI with the problematic GA/LAPI codes. Ultimately these kernels will be used as demonstrators and targets for IBM's in-house developments, both around colony and Federation.
- Use of the GAMESS-UK code on HPCx was highlighted at the Codes Workshop and Training Course at Daresbury (3-4 April, 2003 - M.F. Guest, P. Sherwood and H. van Dam)

4.3.5 Computational Engineering

CFX

- Installed and tested on HPCx and solutions can be run in parallel under MPI.

PNEWT

- Installed on HPCx and its performance analysed for the given test cases. Other test cases are currently under preparation by the UKTC consortium. NEWT has also been profiled using the Paraver trace tool.
- Analysed and improved memory management in partitioning software for PNEWT on HPCx.

Turbulence Code

- Currently investigating the scaling of a turbulence code for Professor McComb, a potential HPCx user. The code has been successfully ported from an IBM SP3 in Spain to the HPCx system. The initial version supplied produced incorrect results, which limited progress. However this has now been resolved and results agree within rounding errors.
- Initial tests indicate the HPCx system takes about half of the time per iteration compared to the IBM SP3, when running on 32 processors.
- Currently discussing instrumentation of the code with the group for more detailed timing tests and a scalability study.

4.3.6 Environmental Science

POLCOMS

- Installed on HPCx. Initial benchmarking showed poor performance scaling. This was found to be due to some collective calls associated with error checking code, which could be removed.
- Further benchmarking now shows excellent scaling at 1 km grid resolution to 1024 processors.

4.3.7 Libraries

- Converted libraries to versions that contain both 32 and 64 bit versions of the object file so that only one library need be maintained for each package. This was done for a number of libraries: LAPACK, BLACS, ScaLAPACK. A short article on how to do this has been written-up for the web site
- EIGENSOLVERS – Peigs 2.1 and Peigs 3.0 and alternative ScaLAPACK eigensolver routines have been installed on HPCx. Their scaling performance has been analysed for a range of matrix sizes.

- The latest version of Peigs3.0 has now been installed. Currently testing performance, along with other eigensolvers for large-scale application matrices.
- Discussions with Jonathon Tennyson's group about obtaining test matrices which are typical for the atomic and molecular community.
- Installed ARPACK (See PLATO section).
- Investigated the performance of FFTs on HPCx (FFTW, ESSL and PESSL). Currently writing a report that will be made available on web site.

4.3.8 Tools

- IBM MPI Trace tools and the Paraver and Vampir trace tools have been installed on HPCx and made available to users.
- Updated Vampir to version 3.0; currently testing enhanced features.

4.3.9 Staff Training

- Most members of the terascaling team attended an HPCx optimisation course given by IBM at Edinburgh.

4.3.10 Consortium Meetings and Presentations

- M. Plummer spoke at the 10th March UKCP meeting at DL, presenting HPCx, Castep and Newtep benchmarks and progress on Newtep optimization. Fielded questions on HPCx.
- M. Plummer attended the 28 March meeting of the Materials Chemistry Consortium. Fielded questions about HPCx and promoted Newtep on HPCx (task-farming, memory optimization, AllToAll ongoing improvements): the consortium is recommending Newtep over Vasp where the two codes are alternatives.
- I.J. Bush presented work on the Crambin benchmarks at the BioSRRTnet meeting in Frascati, Italy. (February, 2003)
- M.F. Guest spoke at the "Biomolecular Simulation - Developing a Strategy" meeting at Cosener's House, Abingdon, 24-25 March, 2003, with a presentation entitled "HPCx : A New Resource for UK Computational Science. Challenges and Opportunities in Biomolecular Simulation". Lorna Smith attended this meeting. Fielded questions on HPCx, capability computing and software packages.
- M.Ashworth gave a presentation entitled "HPCx: A New Resource for UK Computational Science - Overview and Benchmarks" and J. Hein gave a presentation entitled "A scalability study on a 1280 processor p690 system" both at ScicomP 7, 4-7 March 2003, Goettingen, Germany.

4.4 Software Engineering (*Dr Stephen Booth*)

4.4.1 Low Level Communications

Investigation of the low level communication protocols of HPCx, in particular, MPI, MPI-2 single sided, LAPI and GA-tools.

- We have started work on the development of a set of standard communication benchmarks for Low level communications. So far we have benchmarks for:
 - MPI-1
 - MPI-2 Single sided
 - LAPIThese are being extended and a report comparing the different libraries is under development.
- We have investigated the performance of MPI collective communications on HPCx. The IBM MPI library has not been optimised for clusters of SMPs. We have developed SMP optimised collective routines that can be linked into an existing application to improve the performance of some of the common collective operations.

4.4.2 Grid Integration

- Globus-2 has been installed on HPCx and is undergoing final configuration.

4.4.3 MPI and Mixed Mode Programming

- *Single sided Communications:* The performance of MPI-2 single sided communication on HPCx has been investigated. MPI-2 single sided communications seems to suffer from a large message latency compared with normal MPI communications. A report on this work is in preparation.
- *MPI I/O performance:* A report on MPI I/O on HPCx is in preparation.

4.4.4 General Terascaling Techniques

- We presented a paper at ScicomP7 at Gottingen Germany describing a scalability study of an example code on HPCx.

4.4.5 Java Performance

- A report is available on the HPCx website comparing the performance of Java, Fortran and C. This work was performed as part of UKHEC.

4.4.6 System Administration Functions (SAF)

- As expected we have needed to continue to put significant development effort into the SAF during the first 3 months of the project. Developments have included:
 - Improving the reporting mechanism to handle the increased amounts of data now being produced.
 - Additional charts and information in the reports.
 - Additional output formats for the reports.
 - Improved handling of Disk Quota management.
 - Improved machine status reporting.
 - Allow use accounts to be members of more than one project.
 - Recording disk use history as well as current status.
- Effort on SAF development is expected to be less in the second quarter however we intend to continue development in the following areas:
 - Improved capacity planning support.
 - Additional work on automated reporting.

4.5 Systems and Networking (*Mr Mike Brown*)

4.5.1 Group objectives

- To house, operate, maintain and develop a production quality user service for scientific capability computing on the principal HPCx service platforms so as:
 - to meet all mandatory and agreed desirable requirements specified within the contracts between the Consortium and the Research Councils;
 - to maximise the QoS (Quality of Service) for the user population by being as flexible and responsive as is reasonable to user needs and requests whilst maximising fairness of access to resources;
 - to operate the service in a professional manner commensurate with its world-leading status;
 - to meet all statutory and contractual requirements, as well as accepted sector best-practice standards, as regards security of access, security of data and physical protection of the service whilst

applying appropriate safe working practices as required under Health and Safety regulations.

4.5.2 Modus operandi

- Although *not obligated to do so under contract*, currently we are providing a 7 x 24 service with 16 hours on-site systems support coverage Mon-Fri, and 8 hours Sat-Sun. All other times are covered by an on-call rota. This is to enhance QoS by the shortening of incident response times. Because of the high incident rate, this level of cover has been very hard on the staff concerned at times.
- During each on-site 8-hour shift (12 during the week), there is a designated “Duty Shift Systems Administrator” whose role is to actively monitor the system, and respond to all queries and account management requests that are directed towards the Group during that shift. A detailed shift log is maintained (and preserved) for every shift.
- Standard Operating Procedures are in place (and are constantly being revised and updated)

4.5.3 Planned Enhancements to Service

- Less restrictive memory limits for shared-memory jobs:
 - enhancement to LoadLeveler configuration to support a more dynamic approach to applying memory limits is in design phase
- Support for interactive parallel jobs:
 - to enable effective application development and to provide access to debugging facilities, 4 LPARs (initially) will be configured to support interactive parallel applications. Access to the LPARs will be in shared, non-exclusive mode, and will be under control of an amended LoadLeveler configuration. This facility is under test and an announcement of availability is expected to be made soon
- HSM:
 - a HSM (hierarchical storage management) facility will be made available soon. This will enable data to be migrated from a large cache filesystem to the 24 LTO drive tape-library, with (it is anticipated) similar functionality to DMF as used originally on the Edinburgh CRAY T3D/T3E and latterly on the CSAR T3E services
- Phase 2 system:
 - planning has already started for the future installation of the Phase 2 system, with a Phase 2 Transition Planning Group having been set up between CLRC/UoE and IBM
- The principal issue is how to achieve a seamless (as far as possible) transition between Phase 1 and Phase 2 services, when the following all has to change (*at the SAME time*):

- the processors (p690 “H” to “H+”)
- the entire switch network (“colony” to “federation”)
- the OS across all systems (AIX 5.1 to 5.2)
- the frame partitioning to change from 4 x 8 LPAR to 1 x 32 SMP
- the parallel system support environment (PSSP to CSM)
- all GPFS data has to be recreated - filesystems will not survive the transition from a PSSP-managed system to a CSM-managed one

4.5.4 Key Issues

- Reliability
 - very high number of incidents have occurred on the service since December. While the hardware has generally been reliable, the reliability of the switch (and all that relies upon it, such as GPFS) has been *extremely* poor. 98 incidents since start of service, 26 at Severity 1 (although some of these have been coalesced)
- Lack of development resources (IBM)
 - the HPCx system is the largest p690 cluster in the world, this causes issues in that problems encountered cannot necessarily be replicated in-house by IBM as they have no access to significant development resources
- Lack of development resources (HPCx)
 - there is a clear and pressing need for a separate development system to be available at DL alongside the main machine - that replicates the principal hardware, software and interconnect of the service platform, so that fixes can be adequately (and safely) tested before live application
- Inability to undertake tests upon service platform
 - highlights the lack of an available development resource. The cycle of applying fixes and new configuration features is hampered by the need to avoid ANY risk of service disruption (even minor) because of the contractual framework. Impacts on QoS.
- IBM support structure in the UK
 - this is currently not set up to support large-scale HPC-class systems, with minimal knowledge and experience on large configurations centred around a switch fabric. Deficiencies have been identified, and IBM is proposing to address these, but the early months of the service have suffered as a consequence of a lack of specialised knowledge within the UK.
- Transition to phase 2
 - delivering a seamless transition to phase 2 processor and switching technology will be a challenge.

4.5.5 Successes

- External links:
 - membership of SP-XXL (international group for sites with very large IBM SP-type systems) - meets twice per year
 - membership of IBM UK HPC “user” Group - meets with other large UK IBM HPC-class sites (ECMWF, AWE) on a regular basis to discuss common issues (particularly focussing on IBM support)
- Hardware reliability is well above expectation
 - the bulk of failures to date have occurred with redundant components (disc spindles within RAID arrays, duplicated bulk-power regulators).
 - only ONE MCM (processor module) has been replaced to date
- AIX (within a single box) is as solid and reliable as would be expected from a mature SMP operating system
 - the issues start when trying to maintain a service across a cluster of 166 instances of AIX with a shared global filesystem over an interconnect that may charitably be best described as sub-optimal
- At Daresbury Lab we have commissioned a Computer Room with state-of-the-art support infrastructure
 - a 1.6MVA rotary UPS incorporating a 64-litre V16 2500hp standby diesel generator has just been commissioned
 - power and cooling for Phase 2 is next to be installed

4.6 Staffing

<i>AV</i>	<i>January</i>	<i>February</i>	<i>March</i>
DL	4.4	3.7	4.1
EPCC	8.2	8.0	8.5
Total	12.6	11.7	12.6

<i>Systems</i>	8.2	7.1	7.9
----------------	-----	-----	-----

5 Summary of Performance Metrics

<i>Metric</i>	<i>TSL</i>	<i>FSL</i>	<i>January</i>	<i>February</i>	<i>March</i>
Technology serviceability	80%	99.2%	99.6%	98.0%	96.8%
Technology MTBF (hours)	200	300	293	183	81
Number of AV FTEs	7.5	10	12.6	11.7	12.6
Number of training days per month	30/12	40/12	10/1	17/2	17/3
Non in-depth queries resolved within 3 days	85%	97%	98.7%	98.7%	97.8%
Number of A&M FTEs	3.75	5.75	8.2	7.1	7.9
A&M serviceability	80%	100%	99.4%	99.6%	99.9%

<i>Colour</i>	<i>Meaning</i>
	Exceeds FSL
	Between TSL and FSL
	Below TSL

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: Current Projects

EPSRC Projects

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
e01	1	UK Turbulence Consortium	Prof Neil Sandham
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

NERC Projects

<i>Code</i>	<i>Title</i>	<i>PI</i>
n01	Large-Scale Long-Term Ocean Circulation	Dr David Webb
n02	NCAS	Prof Alan J Thorpe
n03	Computational Mineral Physics Consortium	Dr John Brodholt
n04	Shelf Seas Consortium	Dr Roger Proctor

Early User Projects

<i>Code</i>	<i>Title</i>	<i>PI</i>
y001	Materials	Dr Patrick Briddon
y002	DNS of Turbulent Flow	Prof Neil Sandham
y003	Multi-photon and Electron Collision Processes	Prof Ken Taylor
y004	Materials	Prof Jonathon Tennyson
y005	UKAEA	Dr Tim Hender
y006	UK Car-Parrinello Consortium	Prof David Price
y007	Climate Modelling	Ms Lois Steenman-Clark

HPCx Projects

<i>Code</i>	<i>Title</i>	<i>PI</i>
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 Installation	Dr David Henty