

HPCx Quarterly Report

January-March 2004

1 Introduction

This report covers the period from 1 January 2004 at 0800 to 1 April 2004 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

- During this quarter, utilisation of the service has exceeded 80% for each month, with March being another new record month.
- The job size using the most time remains 128 CPUs although almost 30% of the time was used on capability jobs.
- We expect demand to remain high during the transition to Phase 2 and that the Phase 2 machine will fill up reasonably quickly.
- The number of active research groups is now at 30, which is the maximum agreed with EPSRC for 2004.
- Staffing effort levels for both the Science Support and Systems teams have been high during this quarter in preparation for moving the service onto the Phase 2 hardware.
- The initial Phase 2 system has now grown to 20 Regatta H+ compute frames connected with a pre-General Availability HPS (High Performance Switch). We were able to start running an early-user service on this at the end of March.

- We have now implemented access to the file archive system and a number of user groups have started to make use of this.
- Significant science support effort has been expended on ensuring the transition to the Phase 2 system will be as smooth as possible for users. For example, we:
 - ported a number of applications codes to Phase 2;
 - have installed key tools and libraries;
 - and have produced a document on the differences on Phase 2, which was available for the early users.
- The second HPCx Annual Seminar will be held in Edinburgh on 9 July and will be preceded by a new scaling course on 8 July.
- There will also be a workshop in London on 26 May covering initial user experiences on Phase 2 and including a closed session with NERC users to better understand their needs.
- As part of our role in ATWG, significant effort has gone into characterising the performance of the benchmark codes on both Phase 1 and Phase 2.
- The Software Engineering team has been investigating the performance of the new switch and has been looking at ways to exploit the larger LPARs on Phase 2.
- Outreach has focussed on getting the Life Sciences programme up and running and on preparing for the Industry Day held at the beginning of April.

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:

	January	February	March
Incidents	13	15	13
Failures	0	5	5

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

January

None

February

<i>Failure</i>	<i>Site</i>	<i>IBM</i>	<i>External</i>	<i>Reason</i>
04.020	0%	0%	100%	External networking failure
04.026	100%	0%	0%	Power loss
04.031	100%	0%	0%	LL problems after recovery from power loss
04.033	0%	100%	0%	GPFS loss after switch plane failure
04.039	0%	0%	100%	External networking failure

March

<i>Failure</i>	<i>Site</i>	<i>IBM</i>	<i>External</i>	<i>Reason</i>
4.057	100%	0%	0%	Networking problems
4.063	100%	0%	0%	Networking problems
4.068	0%	100%	0%	LL central manager problem
4.067	100%	0%	0%	Networking problems
4.069	100%	0%	0%	Replace failed network switch

3.1.2 Performance Statistics

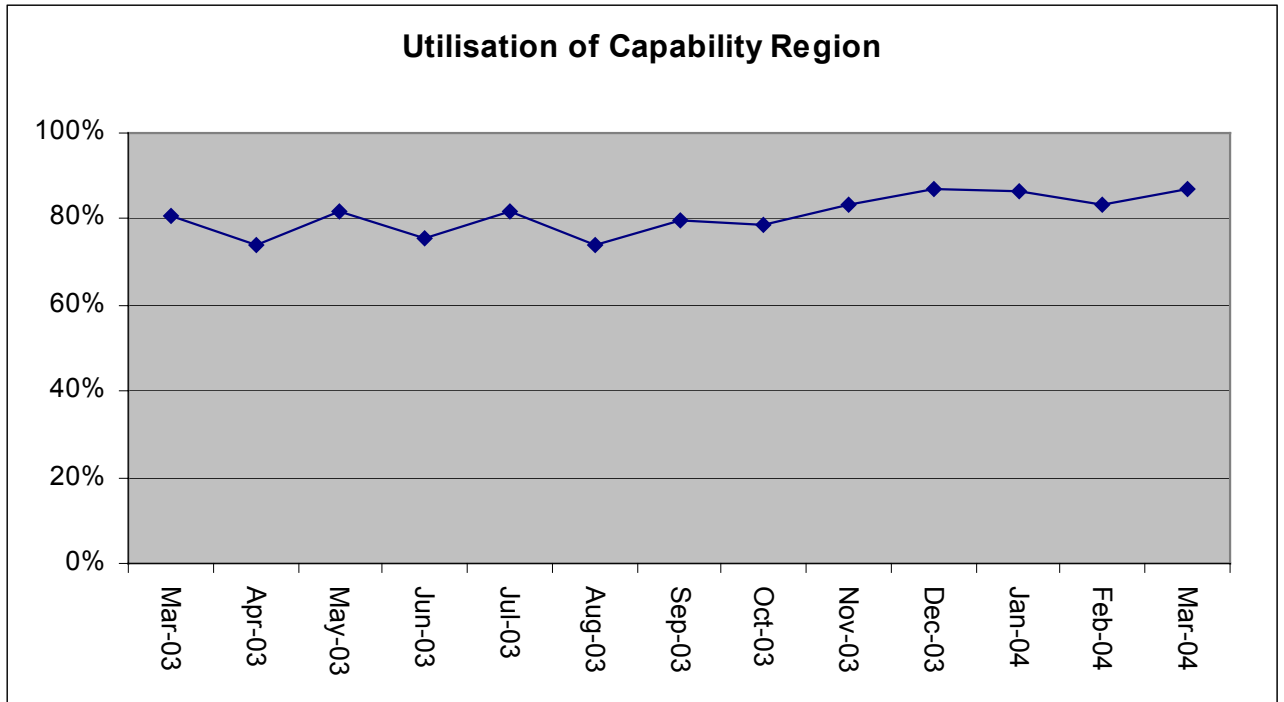
This section uses the definitions agreed in Schedule 7, ie,

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

<i>Attribution</i>	<i>Metric</i>	<i>January</i>	<i>February</i>	<i>March</i>	<i>Quarterly</i>
IBM	Failures	0	2	1	3
	MTBF	∞	366	732	732
	Serviceability	100.0%	99.8%	99.7%	99.8%
Site	Failures	0	1	4	5
	MTBF	∞	732	183	439
	Serviceability	100.0%	97.9%	99.7%	99.2%
External	Failures	0	2	0	2
	MTBF	∞	366	∞	1098
	Serviceability	100.0%	98.9%	100.0%	99.6%
Total	Failures	0	5	5	10
	MTBF	∞	146	146	220
	Serviceability	100.0%	96.6%	99.3%	98.7%

3.2 Capability Utilisation

The monthly utilisation for the 1024-processor capability region is shown in the following graph. This has exceeded 80% each month this quarter, with the figure for March being another record.




3.3 Capacity Planning

Predicted Utilisation

The following graph shows the utilisation since the start of the project and the projected utilisation until November 2004. The scale on the y-axis is AUs per hour, where the peak that HPCx Phase 1 could currently deliver is around 3240 AUs per hour, and Phase around twice that (the red line in the graph). The yellow line 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 graph suggests that during the transition to Phase 2 the service will be substantially overloaded. Even when Phase 2 comes online, it appears that it is likely to fill up quickly.

Numbers of Research Consortia

There are currently 30 research consortia using the HPCx system. This is the maximum number for 2004, agreed between HPCx and EPSRC.

3.4 CPU Usage by Job Size



The above graph shows that the job size using the most time is still 128 CPUs. Capability usage is slightly down, at 28.4%

3.5 AU Usage by Consortium

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


<i>Consortium</i>	<i>January</i>	<i>February</i>	<i>March</i>	<i>Quarterly</i>	<i>%age</i>
e01	117405	221640	80072	419117	7.4%
e02	774031	433964	219	1208214	21.3%
e03	32492	31027	136065	199584	3.5%
e04	22578	37594	74293	134465	2.4%
e05	106613	215693	369777	692083	12.2%
e06	616992	583646	930319	2130957	37.5%
e07	2157		2125	4282	0.1%
e08	47	1053	7690	8790	0.2%
e10		2616	63578	66194	1.2%
e11	79697	7836	17687	105220	1.9%
e12	705	6316	2911	9932	0.2%
e13	42840	18703		61543	1.1%
EPSRC Total	1795557	1560088	1684736	5040381	88.7%

n01	72108	45151	59804	177063	3.1%
n02	18336	90192	462	108990	1.9%
n03	10783	11299	53260	75342	1.3%
n04	5895	6945	47972	60812	1.1%
NERC Total	107122	153587	161498	422207	7.4%

p01	9	416	0	425	0.0%
PPARC Total	9	416	0	425	0.0%

c01	23665	25715	112684	162064	2.9%
CCLRC Total	23665	25715	112684	162064	2.9%

z001	12629	22378	17808	52815	0.9%
z002	0	34	33	67	0.0%
z004	139			139	0.0%
z06	1424	26	139	1589	0.0%
HPCx Total	14192	22438	17980	54610	1.0%



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.

<i>Consortium</i>	<i>AUs Used</i>	<i>AUs Charged</i>	<i>Discount</i>
e01	433213	419116	14096
e10	94562	66193	28368

3.6 Helpdesk

3.6.1 Classifications

<i>Category</i>	<i>Number</i>	<i>% of all</i>
Administrative	76	31.5
Technical	146	60.6
In-depth	16	6.6
PMR	3	1.2
TOTAL	241	100

<i>Service Area</i>	<i>Number</i>	<i>% of all</i>
Phase 1 platform	226	93.8
Website	13	5.4
Other/general	2	0.8
TOTAL	241	100

3.6.2 Performance

<i>All non-indepth queries</i>	<i>Number</i>	<i>%</i>	<i>Target</i>
Finished within 24 Hours	185	83.3	75%
Finished within 72 Hours	221	99.5	97%
Finished after 72 Hours	1	0.5	

<i>Administrative queries</i>	<i>Number</i>	<i>%</i>	<i>Target</i>
Finished within 48 Hours	75	98.7	97%
Finished after 48 Hours	1	1.3	

3.6.3 Experts Handling Queries

<i>Expert</i>	<i>Admin</i>	<i>Technical</i>	<i>In-Depth</i>	<i>PMR</i>
epcc.ed.ac.uk	43	60	9	3
dl.ac.uk	6	25	4	0
Sysadm	27	60	3	0
Other people	0	1	0	0

3.7 Service Quality Tokens

<i>Date</i>	<i>Person</i>	<i>Value</i>	<i>Comment</i>	<i>Status</i>
Mar 28, 2004 12:51:50 AM	Dr Dario Alfe	••••	account locked again (for who knows how long). policy on reset account passwords should really be changed.	This was a Sunday morning. Nonetheless, Dr Alfe's account was unlocked by 8:30 AM.
Feb 26, 2004 9:37:26 PM	Dr Dario Alfe	•••	Also, I can see no reason why my e05da account must have the password reset. This is now happening too often, and I would really like somebody to tell me what's wrong with this account. I have put this request many time to the helpdesk but I had no reply	Accounts are locked after three attempts fail because of a bad password. As far as we can see, this is what has happened on a number of occasions to this account
Feb 26, 2004 9:34:38 PM	Dr Dario Alfe	•••	It is taking an untollerably long amount of time to sort out a simple question:	Dr Alfe sent a query about this at 17:36 on Feb 26. The next morning a staff

			giving me a new password for accessing my e05da account. I believe this kind of requestes should be automated, and satisfied automatically within seconds.	member phoned him and assisted him to log in. He was also sent details of how to log in without a password, using Agent Forwarding. However, we believe that security concerns make it inadvisable to fully automate the issuing of new passwords.
--	--	--	--	--

4 Support

4.1 Applications Support (*Dr David Henty*)

We are progressing well against all the key targets identified in the Annual Plan for 2004. The two main aims for this quarter were to start developing Phase 2 documentation in advance of user service and to kick off the organising committee for this year's Annual Seminar.

4.1.1 Documentation

The main target for this year is to ensure that documentation is updated to include all changes introduced by the Phase 2 service. To this end, we have produced a new document "HPCx: Phase 2 Differences" which was available on the WWW pages in time for early user access.

4.1.2 Technical Reports

Two reports were due this quarter, one on "Mixed-mode applications" and the other on "Performance of communications libraries". We have actually produced three reports

- **HPCxTR0401:** *An LPAR-customized MPI_AllToAllV*, Martin Plummer and Keith Refson.
- **HPCxTR0402:** *Introducing LAPI and Capturing the Performance of its Point-to-Point Communications on HPCx and the Colony SPswitch2*, Christos Kartsaklis.
- **HPCxTR0403:** *Mixed Mode Applications on HPCx*, Jake Duthie, Mark Bull and Lorna Smith.

where we have investigated the performance of two different communications libraries: MPI and LAPI.

4.1.3 Training

In the first three months of 2004 we ran the following four courses:

- **Loughborough, 19-20 January:** *Message-Passing Programming*.
- **Edinburgh, 27-29 January:** *Object-Oriented Programming for HPC*.
- **Edinburgh, 10-12 February:** *Hardware, Compilers and Performance Programming*.
- **Edinburgh, 2-4 March:** *Applied Numerical Algorithms*.

where the course in Loughborough was run on request, and places were reserved for HPCx users on all the other courses. The statistics are summarised below alongside the annual targets (where appropriate).

Metric	Total	Target
Course days	11	30
Number of courses	4	
Different courses	4	12
Different locations	2	4
Student-days for HPCx users	82	
Student-days for HPCx staff	9	
Student-days available for HPCx	190	600

A new course on “Improved Performance Scaling on HPCx” is under development and will be run on 8 July, the day before this year’s Annual Seminar. In response to requests for more work on visualisation, we are looking into porting the Visualisation Tool Kit (VTK) to HPCx. If this is successful, we will include the use of VTK on HPCx in our existing “Scientific Visualisation” course, which already covers the toolkit itself in some detail.

4.1.4 Conferences and Workshops

Organisation of the Second HPCx Annual Seminar, to be held in Edinburgh on Friday 9 July, is well underway and we are already open for registration via the HPCx WWW pages. The organising committee currently meets weekly and is now focusing on the programme for the day.

We are also organising a workshop focusing on initial user experiences on the Phase 2 system. This is being held at University College, London, on Wednesday 26 May. The first part of this meeting will comprise a closed session for us to discuss issues with NERC users of HPCx, an event organised at the request of STAC.

4.1.5 User Group

The third HPCx User Group meeting will take place in Edinburgh on the evening of Thursday 8 July, immediately prior to the buffet reception for the Annual Seminar.

4.1.6 Newsletter

The third issue of Capability Computing is just going into production, with printed copies to be available in time for the Annual Seminar in July.

4.1.7 Packages

We continue to maintain all the package accounts, with recent work including updates to GAUSSIAN and work on a parallel version of FLUENT. Package accounts are currently being migrated to Phase 2 in advance of user service.

4.2 Outreach Activities (*Dr Richard Blake*)

Over the past quarter the outreach activities have progressed in the following areas:

- Workplans have been developed with three of the Lifescience projects; workplans will be developed with the two remaining projects by the end of April.
- A discussion paper highlighting the issues that need to be addressed in providing an industrial service on HPCx has been drafted. This will provide a basis for discussing Quality of Service issues for industrial users at the HPCx Industry Day meeting and a follow-up questionnaire.
- Both Fluent and Abaqus have been ported to HPCX Phase 2 and initial test results reported to industrial customers alongside initial discussions on the required Quality of Service. The ports of the commercial codes took longer than expected.
- The HPCx Industry Day is planned for 5 April 2004. Some 50 attendees are expected, coming approximately equally from industry (including public research labs), academia and HPCx staff. The meeting will seek to explore the opportunities for HPCx/ HECTOR class systems to address key computational science and engineering research challenges in industrial research applications in particular in the materials, chemicals, lifesciences and environment areas.
- A report on lessons learnt from the Teragyroid experiment undertaken last November will be submitted in mid-April.

4.3 Terascaling Applications (*Dr Martyn Guest*)

The work described below covers the period January-March 2004, 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 Phase 2 Acceptance Testing

AIMPRO, CASTEP, DL-POLY, H2MOL and PCHAN have been benchmarked on HPCx and Phase 2 as part of the acceptance test exercise. Results have been reported to the ATWG via the Phase 2 Acceptance Test Report. In many cases, this work has included new profiling activities using Vampir and MPITrace tools.

4.3.2 Computational Materials

AIMPRO

- AIMPRO benchmarking results have been discussed with Patrick Briddon. Analyses using Vampir and MPITrace tools have been carried out and benchmark results compared with the SGI Altix, Newton.

Castep

- CASTEP has been ported and tested on Phase 2 for the acceptance tests. A major rewrite of the customized MPI_AllToAllV has taken place to take account not only of 32-way LPARs but also, very importantly, the distribution of processors to communication links. This work means that the MPI_AllToAllV is now much more generally applicable to SMP-cluster machines. The customized MPI_AllToAllV improves performance noticeably but even more so if process binding is enforced. This has been carried out using the IBM's new VSRAC (Versatile System Resource Allocation and Control) tool.
- Comparison tests of CASTEP between Phase 2 and Newton have also been carried out. With proper process binding the customized CASTEP on Phase 2 is faster than on Newton. However there are optimisation issues on Newton that have not yet been addressed (e.g. replacing MPI_AllToAllV with SHMEM calls).
- A parallel bug in CASTEP has been located and fixed. This related to the distribution and cross-referencing of 'standard' and 'fine' grids among the processors.
- The released CASTEP executables for licensed users have been updated. CASTEP users who require 256-512 processors or more have been asked to seek help and advice with any performance problems.

Vasp

- Two versions of the code have been received and we are currently profiling the code to locate bottlenecks.

4.3.3 Molecular Simulation

Siesta

- The scaling and performance of this electronic structure simulation code is being investigated. Profiling of this code shows that a large proportion of the time is spent in MPI_Bcast calls. Replacing MPI_Bcast with an LPAR-aware version (developed by Stephen Booth within the software development team) leads to a performance improvement of about 10% when running a sample with 720 atoms on 32 processors. We are currently investigating the performance of this code further.

NAMD

- This molecular dynamics code makes very efficient use of the HPCx architecture up to 1024 processors, given a large enough problem size. We have been investigating the performance of this code on the Phase 2 system, which shows an expected level of performance improvement. This work has been presented at two meetings:
Joachim Hein presented this work at, (1) a tutorial on computational chemistry codes at ScicomP9, Bologna, 24th-26th March 2004; and (2) at a RealityGrid (e10 consortium) access grid meeting on the 29th March 2004.

Gromacs

- This molecular dynamics code is an alternative to NAMD. As a result of a request from a user, we are currently porting this application to HPCx.

DL_POLY

- The current version of the DL_POLY 3 has been downloaded, compiled and checked for correctness. The code has been profiled using xprofiler and Vampir (uninstrumented code). The code has extensive Vampir instrumentation; however, this is too large for HPCx to process. We are currently instrumenting the most expensive routines only.

4.3.4 Atomic and Molecular

H2MOL

- As part of the acceptance testing, we have been investigating the performance of the H2MOL code. Profiling of this code reveals that the code is more load-imbalanced on Phase 2 than on Phase 1. This may be partly due to the increase in logical partition size from 8 to 32 processors, increasing the complexity of the memory hierarchy.

4.3.5 Molecular Electronic Structure

GAMESS-UK

- A new parallel SCF solver for GAMESS-UK has been written. As far as possible the SCF process has been cast into an object-oriented matrix form and a generic matrix manipulation library, written in Fortran90, has been implemented to provide the underlying matrix operations. After an initial serial implementation, a distributed data parallel implementation, based on ScaLAPACK/BLACS/MPI has been written and is now being tested.
- For large systems and large processor counts the data distribution should reduce the memory requirements by approximately a factor of two relative to the replicated data solver. On a short run on a test system of 4000 basis functions the SCF time on 128 nodes is now more than 20% faster than the earlier implementation and larger speedups are expected for longer runs, larger systems and larger processor counts.
- It is intended that the library component of this development can be deployed as a way of introducing ScaLAPACK (and potentially other parallel matrix libraries) into codes with similar requirements.

Molpro

- Andrew Sunderland attended a CCP6 meeting at UCL in January and reported on Eigensolver performance on HPCx. We received a copy of Molpro to look at, and downloaded very large application matrices (n=30,000) from Jonathan Tennyson's research group for testing against various Eigensolvers.

Chemistry Codes on Phase 2

- A number of chemistry codes have been migrated to Phase 2, and performance compared with that obtained on the Phase 1 system. These include GAMESS-UK, NWChem, CPMD, DL_POLY (2 and 3) and CHARMM.

4.3.6 Computational Engineering

Fluent

- Fluent has been installed on HPCx and Phase 2. We are currently testing Fluent in shared memory parallel mode using the DSTL benchmarks.

4.3.7 Environmental Science

POLCOMS

- Some optimisation work has been carried out on the global communications routines, including an improved user interface using Fortran 90 generic functions and internally generated communication tags.

4.3.8 Libraries

- We have investigated the parallel PLapack Eigensolvers QR and MRRR (Multiple Relatively Robust Representations - recently installed on HPCx) and compared these with the corresponding ScaLAPACK solvers on HPCx and Phase 2 for symmetric, dense matrices. This investigation has been carried out using actual application datasets from CRYSTAL (standard Eigenproblem) and PLATO (generalised Eigenproblem). The lower latency of the HPS switch on Phase 2 means that runs on large numbers of processors (256 and more) run significantly faster (approx. a factor of two).
- Andrew Sunderland presented Eigensolver timings from HPCx at Scicomp9, Bologna, comparing ScaLAPACK, PLapack, Block-Jacobi and PeIGS parallel diagonaliser routines. The presentation focused on terascaling issues for the parallel diagonalization stages in Crystal, AIMPRO and PRMAT.
- Eigensolver benchmarks have been run on Newton for Phase 2 vs. Newton comparison.

4.3.9 Tools

TotalView

- TotalView has been installed on Phase 2.

Vampir

- Vampir has been installed on Phase 2.

Paraver

- There have been stability and license server problems with Paraver on HPCx. We are currently in contact with CEPBA to obtain a new version for installation on HPCx and Phase 2.

4.3.10 Technical Assessments

- In response to grant applications we now carry out some Technical Assessments of submitted codes by prospective users. Two such

assessments have been carried out this quarter, including a dialogue with the users and the production of an agreed (between us and the users) Assessment Report.

4.3.11 Consortium Meetings and Presentations

- Andrew Sunderland gave a presentation entitled "A Guide to Eigensolvers on HPCx" at the CCP6 Consortium meeting at UCL on 23rd January 2004.
- Mike Ashworth, Mark Bull and Martyn Guest visited IBM Poughkeepsie (January 27th - 29th) for discussions on Phase 2 & HPS performance on HPCx applications.
- Andrew Sunderland and Elena Breitmoser gave a presentation entitled "Parallel Eigensolver Performance on the HPCx System" at Scicomp9, Bologna, 24th-26th March 2004.
- Joachim Hein presented his work on CPMD at a tutorial on "High Performance Computing in Computational Chemistry" at ScicomP9, Bologna, 24th-26th March 2004.
- Joachim Hein also presented work on CPMD at a RealityGrid (e10 consortium) access grid meeting on the 29th March 2004.

4.4 Software Engineering (*Dr Stephen Booth*)

The main focus of the Software Engineering team in this quarter has been preparations for the Phase 2 upgrade. Mostly this has involved investigations of the Federation communication hardware and looking into ways of exploiting shared memory programming.

4.4.1 Low Level Communications

A good understanding of the low level communication performance is vital to the efficient use of systems like HPCx.

- We have been running a large number of low-level communication benchmarks on the Phase 2 system to explore the performance of the Federation switch. These results have been presented internally and a performance report for users is being prepared.
- We have delivered a report on LAPI performance on HPCx Phase 1. [HPCxTR0402](#) "*Introducing LAPI and Capturing the Performance of its Point-to-Point Communications on HPCx and the Colony SPswitch2*", Christos Kartsaklis

4.4.2 Shared Memory Techniques

Shared memory programming techniques are expected to be of increased importance on the Phase 2 system as this system will be constructed out of 32 processor SMP nodes rather than the 8 processor nodes used in phase-1. We have therefore been investigating the use of shared memory techniques on HPCx. Initially we have been concentrating on mixed-mode programming where applications are written to use a combination of OpenMP and MPI.

- We have delivered a technical report on the use and performance of mixed mode programming on HPCx.
 - [HPCxTR0403](#) "*Mixed Mode Applications on HPCx*", Jake Duthie, Mark Bull, Lorna Smith
- Lorna Smith presented a talk on this work at the IBM System Scientific Computing User Group (SCICOMP9) in Bologna, Italy on the 23rd – 26th March, 2004. "Mixed mode programming on a clustered p690 system", L.Smith, J.M.Bull, J.Duthie <http://www.spcicomp.org/ScicomP9/>

4.4.3 Grid Computing

We have been taking an active part in the process to set up the National Grid Service (NGS). This is a production Globus grid made up from the National Supercomputer services (HPCx and CSAR) and the new JISC data and compute clusters. Mostly this activity has been to attend the NGS management meetings as the current software revision on HPCx seems to be compatible with that required by the NGS.

4.4.4 Data Handling

The HPCx administration website (the SAF) has been extended to provide project managers with additional information about the Tape archive. In addition to the number of tapes in use by a Project the SAF can now also report the number of files stored and the amount of data in Gbytes.

4.4.5 General Terascaling Techniques

We have started to investigate methods for optimizing global exchange operations on HPCx. Global exchange operations are where every processor needs to send a message to every other processor. This kind of communication operation is very important for many applications, as parallel FFTs typically require global exchanges.

Global exchange operations generate large numbers of small to medium sized messages and message latency is therefore often important. We are developing an implementation that uses a shared memory segment to combine messages from within an LPAR allowing the data to be sent using a smaller number of larger messages.

4.4.6 Systems Programming

In addition to the improved reporting for the Tape archive the SAF has also been extended in the following ways:

- Improved database structure for holding use data. This should improve the speed of report generation.
- Functions to automatically transfer user account records from the Phase 1 to Phase 2 system.
- Improved postal address format in the database.
- Improved project reports for project managers including more detailed information about the use of capability incentives.

4.5 Operations and Systems (*Mr Mike Brown*)

4.5.1 Staffing

No change in staffing levels, although the coverage pattern has been adjusted to make better use of resources. Coverage still remains substantially in excess of the "core hours" contractual requirement, although such additional coverage remains under review.

There has been the need to apply a lot of additional effort because of the phased build-up of the Phase 2 migration system. The additional workload has been considerable, and has required substantial amounts of staff effort both within, and outside "normal working hours".

4.5.2 Test and Development System

The test and development system has continued to be of critical importance in maintaining the service upon the Phase 1 production platform, and the need for such a facility in support of the Phase 2 service is very great.

4.5.3 Maintenance Sessions

Regular maintenance sessions are still being taken on the Phase 1 system, although some have been taken because of plant services requirements.

4.5.4 Reliability/Stability

The reliability of the system continues to be acceptable, although problems still remain with switch adapters.

4.5.5 File Archive

The initial version of the file archive system (under TSM) went into early-user service in January), and has proven to be stable.

The second phase of the product roll-out was initiated in March, with users invited to comment upon the user interface.

4.5.6 Phase 2 Migration

With the agreement of all parties, the Phase 2 migration was accelerated and by the end of March 20 compute frames (plus 2 for I/O) were in an "early user" fitness-for-purpose trial that is set to run for at least ten days.

Substantial infrastructure upgrade works were required to the computer room, which were undertaken largely during normal user service without incident, and credit is due to the Estates Management team at Daresbury Laboratory.

4.6 Staffing

<i>AV</i>	<i>January</i>	<i>February</i>	<i>March</i>
DL	5.7	5.7	6.2
EPCC	8.6	8.7	9.9
Total	14.2	14.4	16.1

<i>Systems</i>	5.9	5.8	6.6
----------------	-----	-----	-----

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%	100.0%	99.8%	99.7%
Technology MTBF (hours)	200	300	∞	366	732
Number of AV FTEs	7.5	10	14.2	14.4	16.1
Number of training days per month	30/12	40/12	5/1	8/2	11/3
Non in-depth queries resolved within 3 days	85%	97%	100.0%	98.7%	100.0%
Number of A&M FTEs	3.75	5.75	5.9	5.8	6.6
A&M serviceability	80%	100%	100.0%	96.6%	99.3%

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

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

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
e08	2	Organic Solid State	Prof Sarah Price
e09	2	Molecular Properties and their Geometry	Prof Peter Taylor
e10	1	Reality Grid	Prof Peter Coveney
e11	1	Bond making and breaking at surfaces	Prof Sir David A King
e12	1	Parallel programs for the simulation of complex fluids	Dr Mark R Wilson
e13	1	TeraGyroid project	Dr Richard J Blake
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

NERC Projects

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
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

PPARC Projects

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
p01	1	Atomic Physics and Astrophysics	Prof Alan Hibbert

BBSRC Projects

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
b01	2	Quantum Chemistry Studies of the Rusticyanin Protein Crystal	Prof Samar Hasnain
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

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
c01	1	Daresbury Laboratory Facilities Agreement Consortium	Dr Richard J Blake

Externally-funded Projects

<i>Code</i>	<i>Title</i>	<i>PI</i>
x01	HPC-Europa	Dr J-C Desplat

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 Porting	Dr David Henty
z07	Package Installation	Dr Mike Ashworth