HPCx Service Report
December 2002

1 Introduction

This report covers the period from the official start of full service on 9 December 2002 at 0800 to 1 January 2003 at 0800. This gives a service month of 552:00:00 hours. The early months of service are typically a `bedding-in' period --- nevertheless, the numbers of failures and incidents has been higher than we would have liked.

2 Usage

During this month, there were 42 active users from 7 different consortia.

2.1 Availability

Incidents

During this month, there were 16 incidents, of which 15 were attributed to IBM and 1 was attributed to A&M. The following table indicates the severity levels of the incidents, where severity level 1 is defined as a *Failure* (in contractual terms). The definitions used for severity levels can be found in Appendix A.

Severity	*Number*
1	5
2	3
3	8
4	0

The following table gives more details on the Severity 1 incidents:

Severity 1 Incident	*Reason*
02.012	HACMP collapsed. Logins lost
02.013	System idled to restart switch
02.014	16 LPARs fenced out
02.015	16 LPARs idled due to switch table entry problems
02.023	24 LPARs idled due to switch table entry problems

Note: the partial switch failure (which took out a complete plane) is not counted here as a "failure", but the knock-on effect is, ie, HACMP fell over when the system was idled in order to restart.

The overall MTBF is 110 hours.

Availability

Wallclock time lost to Severity 1 incidents: 15:04 hours
Overall serviceability: 97.3%

CPU Time lost: 4211:36 (out of 706560)
Processor serviceability: 99.4%

2.2 CPU Usage by Consortium

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

Consortium	*CPU Hours (Parallel)*	*CPU Hours (Serial)*	*AUs*	*Percentage Utilisation*
z001	151931:29:36	0:31:01	364637	87.45%
z002	134:06:00	0:12:11	322	0.08%
HPCx Total	152065:35:36	0:43:12	364959	87.53%
y004	469:04:00	0:00:00	1126	0.27%
y006	4148:39:28	0:00:00	9957	2.39%
Early User Total	4617:43:28	0:00:00	11083	2.66%
e01	2031:01:36	0:00:00	4874	1.17%
e02	14029:36:40	0:00:00	33671	8.08%
e04	996:30:08	1:33:42	2395	0.57%
EPSRC Total	17057:08:24	1:33:42	40940	9.82%

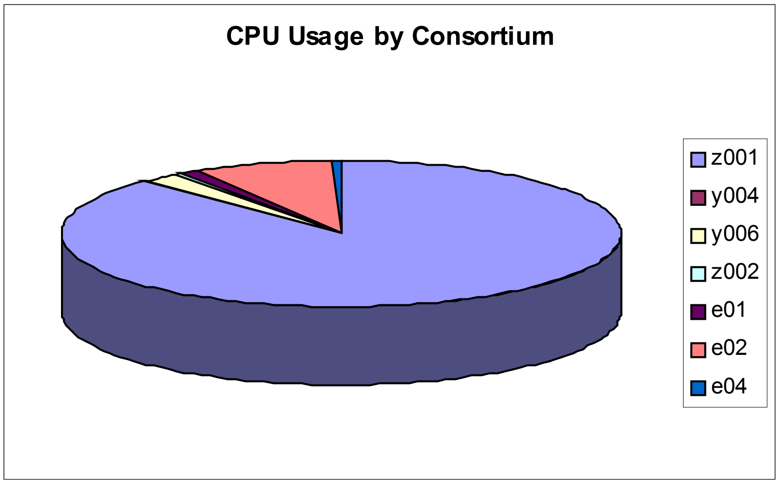

Most of the user groups are still in development and so most of the usage was by the HPCx science support staff. This time has been used on preparing a couple of key applications for production runs, as discussed in Appendix C.

2.3 CPU Usage by Type

Development Region (192 CPUs)

Number of CPUs Used	*CPU Hours*	*Percentage of Total Utilisation*	*Number of Jobs*
8	1930:23:36	1.11%	876
16	673:09:52	0.39%	518
32	3449:40:48	1.99%	392
64	8333:11:28	4.80%	261
Total	14386:25:44		2047

This represents 13.57% utilization of the maximum available in this region.

Capability Region (1024 CPUs)

Number of CPUs Used	*CPU Hours*	*Percentage of Total Utilisation*	*Number of Jobs*
128	10993:31:44	6.33%	80
256	67894:36:48	39.08%	261
512	68609:16:48	39.49%	75
1024	7435:39:44	4.28%	33
Total	154933:05:04		449

This represents 27.41% utilization of the maximum available in this region.

Serial Region (64 CPUs)

	CPU Hours	*Number of Jobs*
Serial	2:16:54	120

This represents 0.01% utilization of the maximum available in this region.

2.4 Job Wait Times

Slowdown

Slowdown is a widely used measure of the relative wait times of different classes of jobs. It is defined as:

Slowdown = (job run time + job wait time) / (job run time)

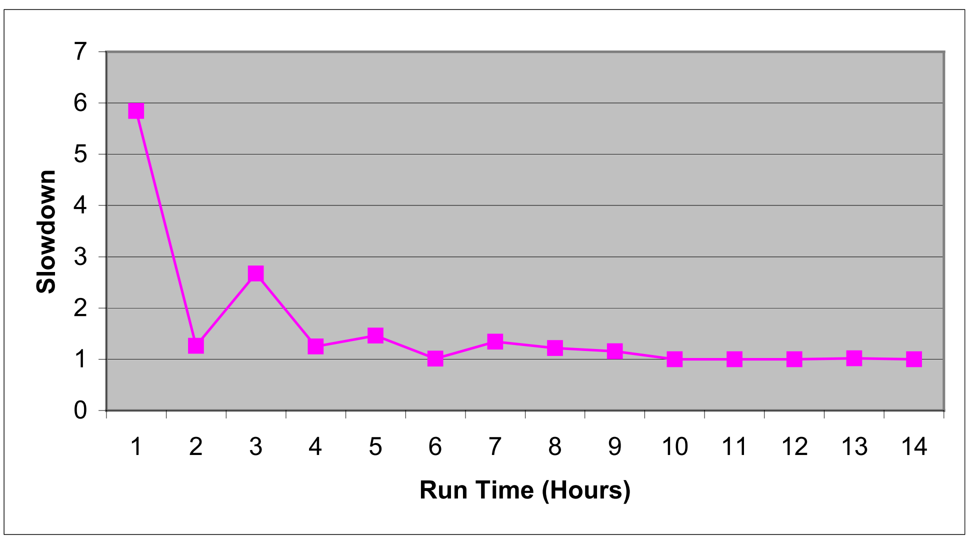

Slowdowns of less than around 10 are usually regarded as reasonable. The above graph shows that only short development jobs (of less than 1 hour) had slowdowns of more than 3. It is important to note that as many of these jobs were very short, the actual job wait times were also short. Moreover, during this first month of service, this category was very busy.

In the following graph, we plot the slowdown figures against the number of processors used and ignoring these development jobs of less than 1 hour.

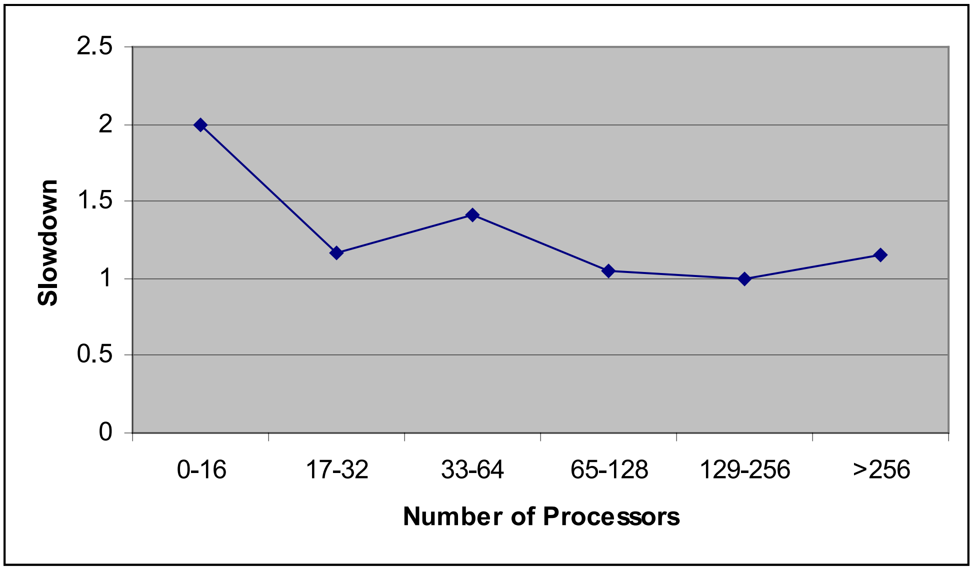

2.5 Disk Occupancy

Home Space

Home space is the part of the disk space which is regularly backed up.

Consortium or Sub-consortium	*Disc Occupancy (KB)*
e01	2872160
e02	1706944
e03	12224
e04	30452832
e05	14162464
e06	96
y001	64
y002	161574432
y003	288
y004	4441760
y005	190272
y006	2477888
y007	1682752
z001	138218208
z002	3622208
z002-a	201920
z003	192

Workspace

Consortium or Sub-consortium	*Disc Occupancy (Kb)*
e01	234253920
e02	480
e03	224
e04	95360
e05	11717344
e06	96
y001	64
y002	800
y003	128
y004	256
y005	160
y006	128
y007	128
z001	86119648
z003	128

2.6 Mass Store

There is currently no usage of the mass store.

3 Support

3.2 Helpdesk

Classifications

Category	*Number*	*% of all*
Administrative	30	36.6
Technical	46	56.1
In-depth	6	7.3
TOTAL	82	100.0

Service Area	*Number*	*% of all*
Phase 1 platform	49	59.8
Website	13	15.9
Other/general	20	24.4
TOTAL	82	100.0

Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	68	89.5	75%
Finished within 72 Hours	76	100.0	97%
Finished after 72 Hours	0	0.0	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	29	96.7	97%
Finished after 48 Hours	1	3.3	

Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*
epcc.ed.ac.uk	21	14	2
dl.ac.uk		7	
Sysadm	9	22	4
Other people		3	

3.3 Training

No training courses were run in December although 2 will run in January.

We ran the following courses in October/November and they were made available both to prospective HPCx users and to HPCx staff.

Title of Course	*Length (days)*	*HPCx Users*	*HPCx Staff*
Fundamental Concepts of HPC	3	2	2
Practical Software Development	3	0	2
Shared Memory Programming	3	0	1
Message Passing Programming	3	4	1
Specialised Programming Models	3	0	1
Introduction to the HPCx Service	1	6	3

4 Staffing

4.1 Science Support Staffing

Daresbury Laboratory

Name	*Days*
Ashworth	11.3
Blake	4.8
Bush	17.5
Guest	9.5
Plummer	13.0
Sunderland	9.5
Total	65.5
FTEs	3.7

EPCC

Name	*Days*
Simpson	5.0
Booth	10.5
Henty	9.2
Breitmoser	0.7
Bull	7.0
Egbert	5.8
Fisher	12.7
Hare	8.7
Hein	10.7
Jackson	10.8
Johnson	8.2
Murdoch	0.5
Pringle	7.7
Smith	2.7
Stratford	2.2
Helpdesk	0.5
Total (Days)	102.7
FTEs	5.8

Overall FTE Levels

	December
DL	3.7
EPCC	5.8
Total	9.5

The slightly low staffing figures were due to annual leave being taken at Christmas.

4.2 Systems Staffing

Name	*Days*
Andrews	12.0
Blake	4.8
Brown	26.7
Elwell	16.0
Franks	12.0
James	8.0
Shore	19.0
Walmsley	18.0
Total	116.5
FTEs	6.6

5 Summary of Performance Metrics

Metric	*TSL*	*FSL*	*Monthly Measurement*
Technology Availability	80%	99.2%	97.3%
Technology MTBF (hours)	200	300	110
Number of AV FTEs	7.5	10	9.5
Number of training days per year	30	40	0
Non in-depth queries resolved within 3 days	85%	97%	100.0%
Number of A&M FTEs	3.75	5.75	6.6
A&M serviceability	80%	100%	100.0%

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

Code	*Title*	*PI*
e01	UK Turbulence Consortium	Prof Neil Sandham
e02	Ab-initio simulation of covalently bonded materials	Dr Patrick Briddon
e03	Multi-photon, electron collisions and BEC HPC consortium	Prof Ken Taylor
e04	Chemreact Computing Consortium	Prof Jonathon Tennyson
e05	Materials Chemistry using Terascaling Computing	Prof Richard Catlow
e06	UK Car-Parrinello Consortium	Prof Paul Madden

Early User Projects

Code	*Title*	*PI*
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

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

Appendix C: HPCx Usage under z001

1. UKTC's Incompressible Channel Code
A world leading Reynolds Number of 1440

Academic Involvement: Professor Sandham and UKTC
z001 (epsrcj): 57844:33:20

Time has been used trying to improve the performance of the UKTC's incompressible channel code. The initial porting from the Cray T3E went well but the performance of the code, both single processor and communications, was not good. The original intention was to run a world-leading calculation at a Reynolds number of 1440, but the projected run times have been too prohibitive. The consortium has reluctantly moved to lower Reynolds number calculations (Re = 720) to look at acoustic data. This is certainly internationally competitive research but far short of what the original goal was.

At a consortium meeting last November, the UKTC asked Dr. Johnstone to see what improvements could be made to the code's performance to see if they could be in a position to tackle the higher Reynolds number case. If little improvement could be made the consortium would abandon that challenge. Much of the time has therefore been spent trying to enable this world-leading calculation and the code's single processor performance has been improved significantly, reducing the run time from 100s per time step to 65s per time step. There remains, however, a communication bottleneck that prevents the code scaling very well and the test runs have been done on the minimum number of processors required, which for this case is 256.

Recent runs have been targeted at initial "production" work that would allow the consortium to make an informed decision about whether to proceed with the full-scale simulation. At this stage, the computational effort required to perform the full simulation is not well known and it is necessary to run the simulation for an unknown period until the transients begin to settle down and useful statistical information can be extracted. This is where the project currently stands. Professor Sandham has been informed of recent progress and results based on these initial calculations and has decided that the consortium will attempt the full simulation.

2. The Simulation of Biological Systems (UKCP and the Materials Chemistry Consortia)

Academic Involvement: UKCP/Materials Chemistry
z001(ijb) 61137:24:40. z001(hpcxmp) 20941:29:12

This work has been undertaken to establish the simulation of biological systems using the CRYSTAL periodic HF/DFT code. This is the target of both the UKCP and Materials Chemistry consortia, so that establishing the principle is an important first step. The initial work is a collaboration between CSE, the SRD (who have excellent data for this system) and Imperial College and will be presented by Montanari at the SRRTNet meeting in Frascati during Feb.

To date a full SCF calculation has been performed on crystalline crambin with 1284 atoms per cell and a rich basis set of 14,000 Gaussian type orbitals. To our knowledge this is by far the largest periodic Gaussian orbital calculation that has been performed to date. Previous calculations are exemplified by the work of Challacombe and Schwegler at the Minnesota Supercomputer Institute (J. Chem. Phys. 106 (13) 5526 (1997) and involved just 3836 basis functions.