HPCx Service Report
January 2003

1 Introduction

This report covers the period from 1 January 2003 at 0800 to 1 February 2003 at 0800. This gives a service month of 744 hours.

The early months of service are typically a `bedding-in' period --- nevertheless, the numbers of failures and incidents has again been higher than we would have liked.

2 Usage

2.1 Availability

Incidents

During this month, there were 29 incidents; 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	6
2	16
3	6
4	1

One of the SEV 1 incidents was due to the power supply; we have a waiver for this type of incident until 18 March. The attributions for the other SEV 1 incidents were as follows:

SEV 1	*Incidents*	*MTBF*
IBM	2.5	298
Site	2.5	298
Total	5	149

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

Failure	*Site*	*IBM*	*Reason*
03.009		50%	Undocumented dependency on CWS by GPFS
	50%		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 the responsibility of the technology supplier
03.023	100%		Uninterrupted access to the external network is site responsibility

Availability

Attribution	*Wallclock Hours Lost*	*Serviceability*
IBM	2.74	99.6%
Site	4.39	99.4%
Overall	7.15	99.0%

Processor hours lost: 9352.5 (out of 952320)
Processor serviceability: 99.0%

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*	*%age*
e01	53795.4	0.1	129109	14.46%
e02	39459.6	0.0	94703	10.61%
e03	20.5	0.1	49	0.01%
e04	3667.8	0.3	8804	0.99%
e05	146208.6	7.9	350920	39.31%
e06	9765.2	0.1	23437	2.63%
EPSRC Total	252917.1	8.5	607021	68.01%

z001	118797.2	11.4	285141	31.95%
z002	38.0	0.1	91	0.01%
z004	140.6	0.1	338	0.04%
HPCx Total	118976	12.0	285570	32.00%

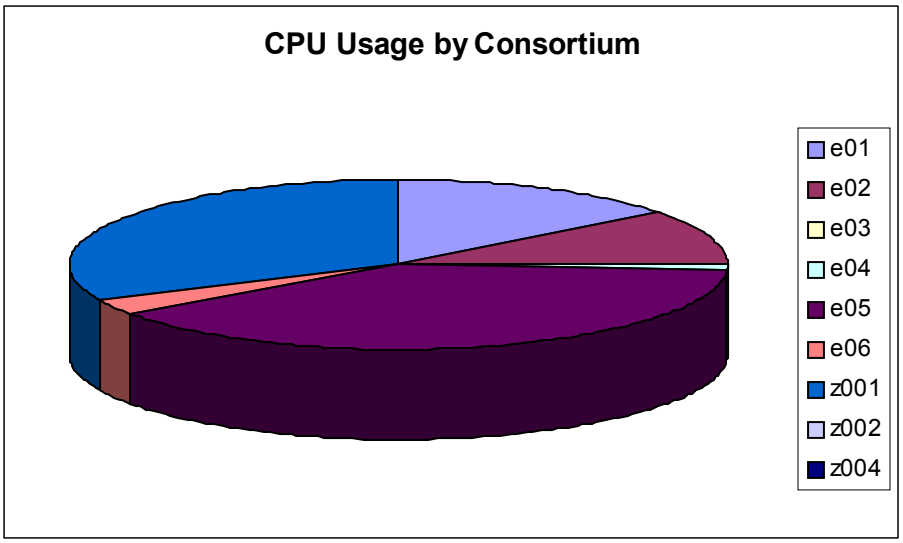

2.3 CPU Usage by Job Type

Number of Processors	*CPU Hours*	*%age*	*Number of Jobs*
8	5070.52	1.36%	1893
16	5089.55	1.37%	538
32	18055.86	4.85%	884
64	37169.94	9.99%	256
128	136310.38	36.65%	428
256	122259.43	32.87%	201
512	42419.56	11.41%	65
1024	5517.75	1.48%	27
1280	5070.52	1.36%	1893

Development Region (192 CPUs): a total of 65387 CPU hours were used which represents 45.77% utilization of the maximum available in this region.

Capability Region (1024 CPUs): a total of 301813 CPU hours were used which represents 39.62% utilization of the maximum available in this region.

Serial Region (64 CPUs): a total of 21 CPU hours were used which represents 0.04% 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)

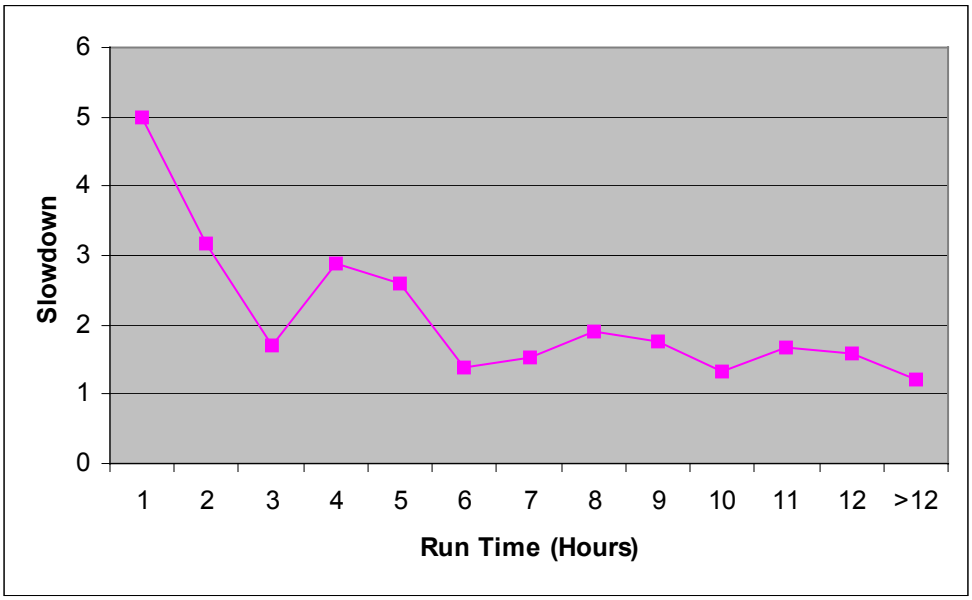

Slowdowns of less than around 10 are usually regarded as reasonable. The above graph plots slowdown against run-time (ignoring jobs of less than 5 minutes duration). The graph demonstrates that only short development jobs had noticeable slowdowns. It is important to note that as many of these jobs were quite short, the actual job wait times were also relatively short. Moreover, during the early months of service, this category has been very busy.

In the following graph, we plot the slowdown figures against the number of processors used and ignoring the development jobs of less than 1 hour. Again, none of these slowdowns exceed 10.

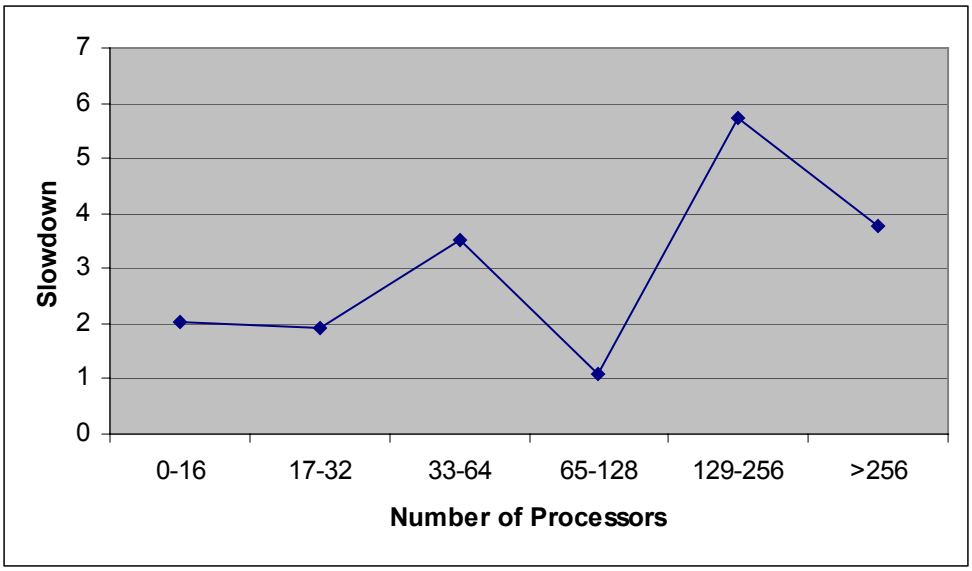

2.5 Disk Occupancy

The disk occupancy has increased steadily during the first few months of service, although no problems have been reported or noticed. However, we do not have detailed records of the disk occupancy for January. We have recently changed the way this data is collected and stored which should allow us to provide details on trends in future reports.

2.6 Mass Store

There is currently no usage of the mass store.

3 Support

3.2 Helpdesk

Classifications

Category	*Number*	*% of all*
Administrative	77	47.2
Technical	78	47.9
In-depth	8	4.9
TOTAL	163	100.0

Service Area	*Number*	*% of all*
Phase 1 platform	125	76.7
Website	16	9.8
Other/general	22	13.5
TOTAL	163	100

Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	133	85.8	75%
Finished within 72 Hours	153	98.7	97%
Finished after 72 Hours	2	1.3	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	75	97.4	97%
Finished after 48 Hours	2	2.6	

Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*
epcc.ed.ac.uk	53	12	5
dl.ac.uk	1	12	1
Sysadm	23	49	2
Other people		5	

3.3 Training

We ran 4 different courses during January. The final one of these was run by IBM staff at EPCC.

Title of Course	*Start Date*	*Length (Days)*	*HPCx Users*	*HPCx Staff*
Object Oriented Programming for HPC	7/1/03	3	2	1
Exploiting the Computational Grid	21/1/03	3	3	1
Introduction to the HPCx Service	27/1/03	1	2	4
Effective Exploitation of the HPCx Service	28/1/03	3	6	11

4 Staffing

4.1 Science Support Staffing

Daresbury Laboratory

Name	*Days*
Ashworth	16.5
Blake	5.5
Bush	20.0
Guest	11.0
Plummer	0.0
Sunderland	21.0
Jones	3.5
Total	77.5
FTEs	4.4

EPCC

Name	*Days*
Simpson	12.2
Booth	11.7
Henty	12.3
Breitmoser	4.5
Bull	6.8
Egbert	3.3
Fisher	19.7
Hare	6.5
Hein	14.6
Jackson, A N	12.5
Jackson, W A	19.0
Johnson	4.2
Murdoch	0.6
Pringle	12.4
Smith	2.8
Stratford	1.7
Helpdesk	0.7
Total (Days)	145.5
FTEs	8.2

Overall FTE Levels

	January
DL	4.4
EPCC	8.2
Total	12.6

4.2 Systems Staffing

Name	*Days*
Andrews	17.6
Blake	13.2
Brown	22.0
Elwell	15.8
Franks	16.5
Jones	3.5
Shore	5.2
Walmsley	21.0
BITD	30.6
Total	145.4
FTEs	8.2

Note: BITD covers a range of bookings from a support department who provide approximately 1 FTE to support computer room operations, electrical and mechanical site services and networking and security. Roughly a dozen staff charge time to the project in amounts that vary from month to month. We believe that it adds no value to report these individual bookings although a full listing can be provided annually if required.

5 Summary of Performance Metrics

Metric	*TSL*	*FSL*	*Monthly Measurement*
Technology Availability	80%	99.2%	99.6%
Technology MTBF (hours)	200	300	298
Number of AV FTEs	7.5	10	12.6
Number of training days per month	30/12	40/12	10
Non in-depth queries resolved within 3 days	85%	97%	98.7%
Number of A&M FTEs	3.75	5.75	8.2
A&M serviceability	80%	100%	99.4%

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