HPCx Service Report
September 2003

1 Introduction

This report covers the period from 1 September 2003 at 0800 to 1 October 2003 at 0800. This gives a service month of 720 hours.

Utilisation continues to be very high, with NERC use at three times last month's figure. Capability use fell back; at this stage we attribute this principally to statistical fluctuations. There were two SEV 1 incidents this month.

2 Usage

2.1 Availability

Incidents

During this month, there were 8 incidents, 2 of which were at SEV 1. The following table indicates the severity levels of the incidents, where SEV 1 is defined as a *Failure* (in contractual terms). The definitions used for severity levels can be found in Appendix A.

Severity	*Number*
1	2
2	3
3	2
4	1

The attributions for the SEV 1 incidents were as follows:

SEV 1	*Incidents*	*MTBF*
IBM	1.75	418
Site	0.25	2928
External	0	∞
Total	2	366

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

Failure	*Site*	*IBM*	*External*	*Reason*
03.125	25%	75%		LL home directory inaccessible
03.129		100%		GPFS failure, cluster wide

Serviceability

Attribution	*UDT*	*Serviceability*
IBM	2:58	99.6%
Site	0:19	99.9%
External	0:00	100.0%
Overall	3:17	99.6%

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	15840.0	0.0	40012	2.3%
e02	47231.5	0.0	122716	7.2%
e03	75064.0	1.1	190066	11.1%
e04	167899.1	27.0	425258	24.9%
e05	194655.7	63.4	496567	29.0%
e06	129790.4	3.0	329909	19.3%
e11	0.0	0.0	1	0.0%
EPSRC Total	630480.7	94.5	1604530	93.8%
n01	29343.1	29.6	74372	4.3%
n02	29.9	0.2	76	0.0%
n03	5591.7	0.0	14158	0.8%
n04	442.1	1.0	1122	0.1%
n05	4522.8	12.8	11484	0.7%
NERC Total	39929.6	43.6	101212	5.9%
p01	97.9	5.7	262	0.0%
PPARC Total	97.9	5.7	262	0.0%

z001	1532.9	0.0	3959	0.2%
z002	4.7	0.0	12	0.0%
z004	232.0	0.0	590	0.0%
z06	98.6	0.0	250	0.0%
HPCx Total	1868.2	0.0	4811	0.2%

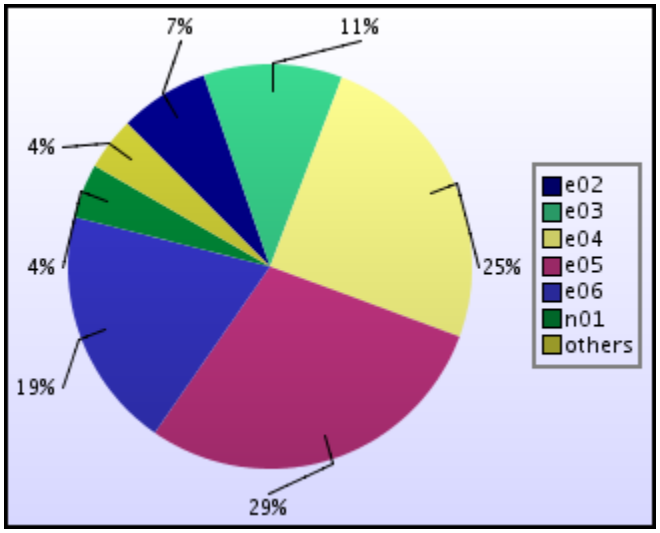

CPU Usage by Job Type

Note: These data will now be presented in AUs.

Number of Processors	*AUs*	*%age*	*Number of Jobs*
8	15867	0.9%	1468
16	20674	1.2%	1806
32	72998	4.3%	611
64	105541	6.2%	875
128	992566	58.0%	619
256	404336	23.6%	281
512	82602	4.8%	37
1024	7873	0.5%	3

Development Region (192 CPUs): a total of 215079 AUs were used, which represents 61.4% utilisation of the maximum available in this region.

Capability Region (1024 CPUs): a total of 1487378 AUs were used, which represents 79.7% utilisation of the maximum available in this region.

2.3 Slowdown and 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)

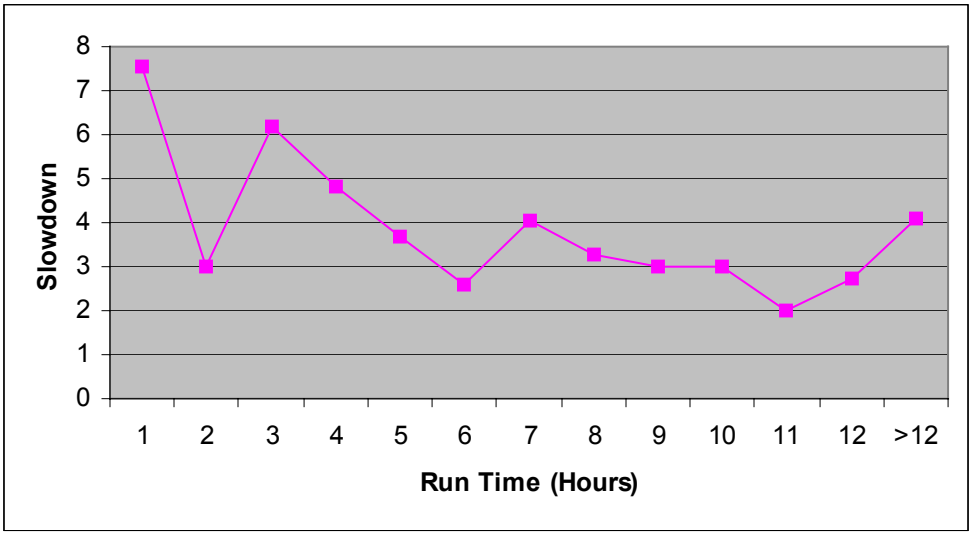

Slowdowns of less than around 10 are usually regarded as reasonable. The above graph plots slowdown against run-time (ignoring jobs of less than 5 minutes duration).

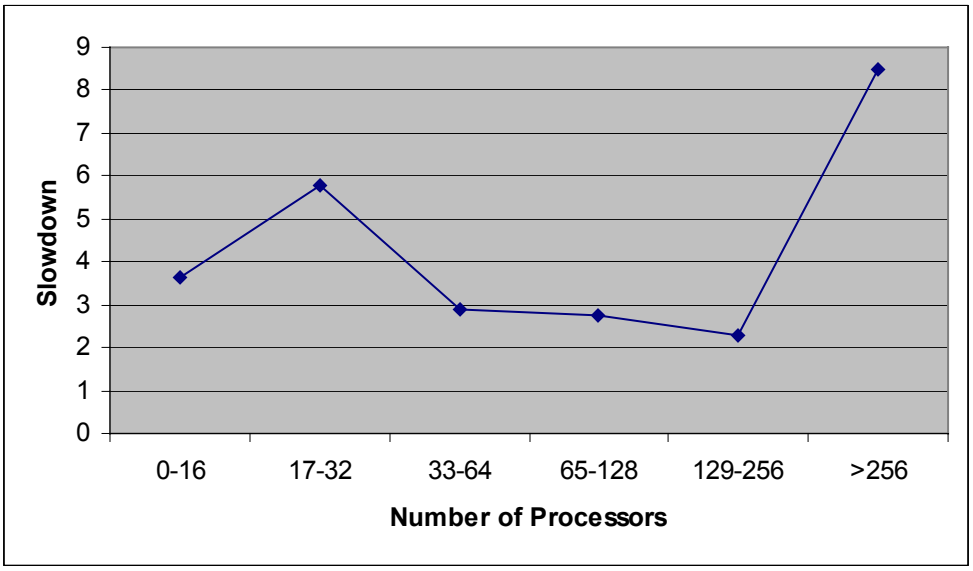

In the graph above, we plot the slowdown figures against the number of processors used and ignoring the development jobs of less than 1 hour.

Job wait times

The following table shows the average wait time (in hours) for each class of job during the month.

Job Class	*Category*	*Number of CPUs*	*Maximum Job length*	*Average wait time*	*Number of Jobs*
par1_1	parallel	8	1	0.2	1310
par1_3	parallel	8	3	4.3	4
par1_6	parallel	8	6	9.8	154
par2_1	parallel	16	1	0.2	1722
par2_3	parallel	16	3	5.5	6
par2_6	parallel	16	6	18.9	78
par4_1	parallel	32	1	0.9	412
par4_3	parallel	32	3	5.0	8
par4_6	parallel	32	6	17.0	191
par8_1	parallel	64	1	0.6	718
par8_3	parallel	64	3	7.7	21
par8_6	parallel	64	6	14.1	136
par16_1	parallel	128	1	1.4	167
par16_3	parallel	128	3	1.2	26
par16_6	parallel	128	6	4.9	55
par16_9	parallel	128	9	29.7	2
par16_12	parallel	128	12	18.2	369
par32_1	parallel	256	1	1.9	164
par32_3	parallel	256	3	2.7	6
par32_6	parallel	256	6	4.4	31
par32_9	parallel	256	9	6.2	5
par32_12	parallel	256	12	10.9	57
par64_1	parallel	512	1	6.1	21
par64_3	parallel	512	3	0.0	0
par64_6	parallel	512	6	8.7	16
par64_12	parallel	512	12	0.0	0
par128_1	parallel	1024	1	8.5	2
par128_3	parallel	1024	3	0.0	0
par128_6	parallel	1024	6	38.1	1
par128_12	parallel	1024	12	0.0	0
serial_1	serial	8	1	0.2	156
serial_6	serial	8	6	0.4	175
serial_12	serial	8	12	0.2	14
inter1_1	interactive	8	1	0.0	1260
inter2_1	interactive	16	1	0.0	304
inter4_1	interactive	32	1	0.0	224

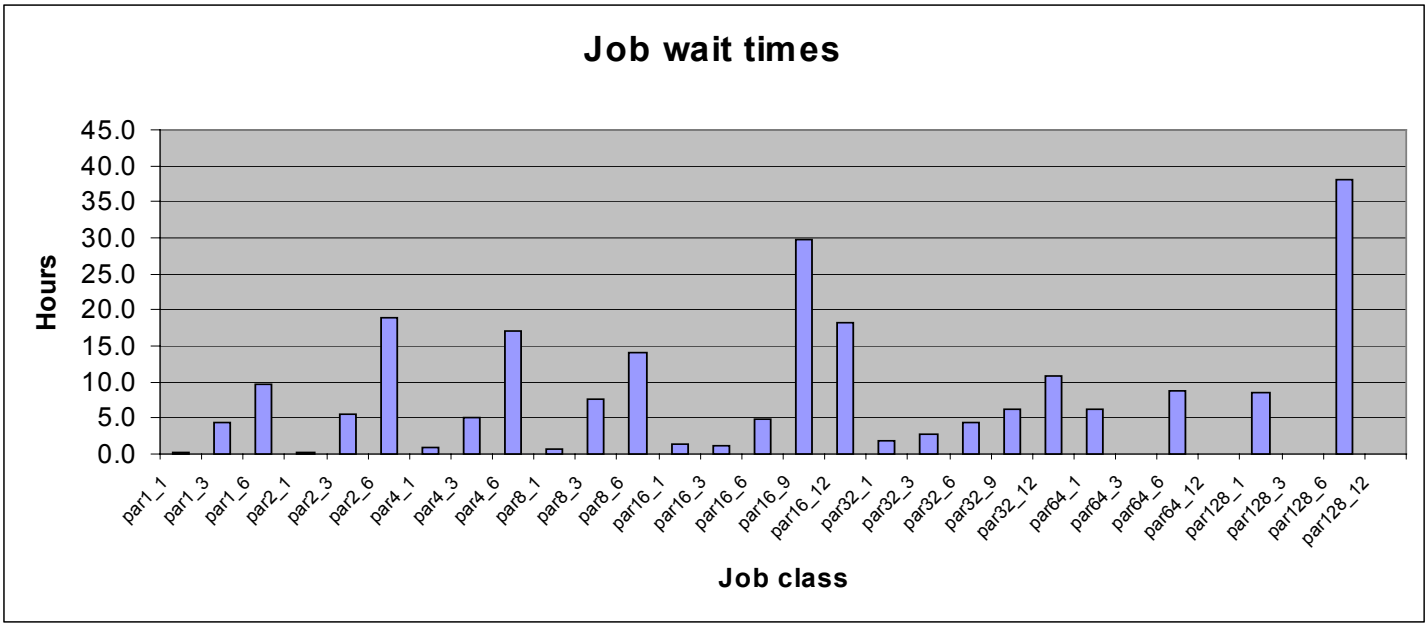

2.4 Disk Occupancy

Home Space

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

Consortium	*Disc Occupancy (Kb)*	*Disc Quota (Kb)*
b01	64	1024000
c01	224	
e01	17566976	50006016
e02	19888768	39760896
e03	74303904	102412288
e04	94796224	102400000
e05	244043936	307200000
e06	93254304	204800000
e07	484288	20480000
e08	902368	20480000
e09	64	
e10	4797216	10240000
e11	440896	
n01	25144608	51200000
n02	5765504	10240000
n03	12237440	
n04	54841184	102400000
n05	97952	10240000
p01	14640800	20480000

y001	64	
y004	175936	
y007	480	
z001	107578016	153602048
z002	23711488	49153024
z003	224	2048
z004	22059104	51200000
z05	256	20480000
z06	22721952	51200000
z07	1433600	10240000

Workspace

Consortium	*Disc Occupancy (Kb)*	*Disc Quota (Kb)*
b01	64	1024000
c01	224	
e01	17566976	50006016
e02	19888768	39760896
e03	74303904	102412288
e04	94796224	102400000
e05	244043936	307200000
e06	93254304	204800000
e07	484288	20480000
e08	902368	20480000
e09	64	
e10	4797216	10240000
e11	440896	
n01	25144608	51200000
n02	5765504	10240000
n03	12237440	
n04	54841184	102400000
n05	97952	10240000
p01	14640800	20480000
y001	64	
y004	175936	
y007	480	
z001	107578016	153602048
z002	23711488	49153024
z003	224	2048
z004	22059104	51200000
z05	256	20480000
z06	22721952	51200000
z07	1433600	10240000

2.5 Mass Store

There is currently no usage of the mass store.

3 Support

3.1 Helpdesk

Classifications

Category	*Number*	*% of all*
Administrative	32	42.1
Technical	41	53.9
In-depth	2	2.6
PMR	1	1.3
TOTAL	76	100

Note: The PMR category indicates in-depth queries that result in Problem Management Reports for IBM.

Service Area	*Number*	*% of all*
Phase 1 platform	49	64.5
Website	11	14.5
Other/general	16	21.1
TOTAL	76	100

Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	59	80.8	75%
Finished within 72 Hours	73	100.0	97%
Finished after 72 Hours	0	0	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	31	96.9	97%
Finished after 48 Hours	1	3.1	

Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*	*PMR*
epcc.ed.ac.uk	23	21	2	0
dl.ac.uk	4	6	0	1
Sysadm	5	14	0	0
Other people	0	0	0	0

3.2 Training

Title of Course	*Start Date*	*Length (Days)*	*HPCx Users*	*HPCx Staff*
Message-Passing Programming	29/9/03	2	9	0

This course was held in Cambridge, hosted by the Cambridge University High-Performance Computing Facility. Apart from the 9 HPCx users, 31 people who were registered by the hosts attended.

4 Staffing

4.1 Science Support Staffing

Daresbury Laboratory

Name	*Days*
Ashworth	16.5
Blake	5.5
Bush	20.0
Guest	5.5
Jones	10.3
Plummer	3.0
Sunderland	22.0
Total (Days)	82.8
FTEs	4.7

EPCC

Name	*Days*
Simpson	13.5
Booth	14.2

Henty	12.8
Bull	8.2
Hein	13.3
Pringle	16.7
Smith	15.8
Anjomshoaa	0.8
Jackson (Adrian)	12.0
Helpdesk	3.0
Total (Days)	110.3
FTEs	6.2

Overall Levels

	FTEs
DL	4.7
EPCC	6.2
Total	10.9

4.2 Systems Staffing

Name	*Days*
Andrews	16.5
Blake	0.0
Brown	22.0
Elwell	12.8
Fisher	18.5
Franks	16.5
Jones	5.1
Shore	5.3
BITD	22.0
Total (days)	118.6
FTEs	6.7

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 which 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 serviceability	80%	99.2%	99.6%
Technology MTBF (hours)	200	300	418
Number of AV FTEs	7.5	10	10.9
Number of training days per month	30/12	40/12	2
Non in-depth queries resolved within 3 days	85%	97%	100.0%
Number of A&M FTEs	3.75	5.75	6.7
A&M serviceability	80%	100%	99.9%

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	*Class*	*Title*	*PI*
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

NERC Projects

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

PPARC Projects

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

BBSRC Projects

Code	*Class*	*Title*	*PI*
b01	2	Quantum Chemistry Studies of the Rusticyanin Protein Crystal	Prof Samar Hasnain

CCLRC Projects

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

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	Dr 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
z05	Outreach Projects	Dr Richard Blake
z06	Application Porting	Dr David Henty
z07	Package Installation	Dr Mike Ashworth