HPCx Service Report
April 2006

1 Introduction

This report covers the period from 1 April 2006 at 0800 to 1 May 2006 at 0800, a service month of 744 hours.

After last month's record delivery of AUs, utilisation in April fell to around 50%. The main causes of this were a large drop in utilisation by three major EPSRC projects, and a fall in utilisation by the NERC projects, who in March had been given large allocations of time and encouraged to use them quickly.

2 Usage

2.1 Availability

Incidents

During this month, there were 10 incidents, 1 of which was 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	1
2	4
3	5
4	0

The MTBF figures for this month were as follows:

SEV1	*Incidents*	*MTBF*
IBM	1.0	732
Site	0.0	∞
External	0.0	∞
Overall	1.0	732

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

Failure	*Site*	*IBM*	*External*	*Reason*
06.036	0%	100%	0%	Maintenance session over-run

Serviceability

There was a total of 63 hours of scheduled downtime this month; this was required in order to upgrade the switch in preparation for the move to Phase 3.

Attribution	*UDT*	*Serviceability*
IBM	2:00	99.7
Site	0:00	100.0
External	0:00	100.0
Overall	2:00	99.7

2.2 CPU Usage by Consortium

Main Service

Consortium	*CPU Hours (Parallel)*	*CPU Hours (Other)*	*AUs charged*	*%age of charged AUs*
e01	39804	165	192431	8.5%
e03	12	0	58	0.0%
e05	49367	169	238486	10.6%
e06	61407	0	295643	13.1%
e07	5195	0	25009	1.1%
e08	9873	0	47531	2.1%
e14	58285	5	280635	12.5%
e15	113	0	543	0.0%
e17	4905	31	23766	1.1%
e18	4664	0	22453	1.0%
e19	0	0	0	0.0%
e20	15610	0	75153	3.3%
e24	4696	0	22608	1.0%
e25	5221	338	26762	1.2%
e26	0	1	5	0.0%
e27	15594	0	75077	3.3%
e28	21456	0	96908	4.3%
e31	6182	32	29914	1.3%
e32	43443	0	148874	6.6%
e33	1946	0	9370	0.4%
e35	819	20	4042	0.2%
e36	640	0	3080	0.1%
e37	8670	0	41742	1.9%
e40	10328	0	49623	2.2%
EPSRC Total	368228	762	1709714	75.9%
n02	26926	7	129670	5.8%
n03	15597	0	75092	3.3%
n04	13155	24	63446	2.8%
NERC Total	55678	31	268208	11.9%
p01	5022	0	24176	1.1%
PPARC Total	5022	0	24176	1.1%

c01	21925	38	103608	4.6%
CCLRC Total	21925	38	103608	4.6%

b08	12215	0	58810	2.6%
BBSRC Total	12215	0	58810	2.6%

x01	3637	400	19439	0.9%
x03	9898	0	47654	2.1%
External Total	13536	400	67093	3.0%

z001	2300	25	11193	0.5%
z002	49	0	234	0.0%
z05	24	0	114	0.0%
z06	1998	21	9719	0.4%
HPCx Total	4370	46	21260	0.9%

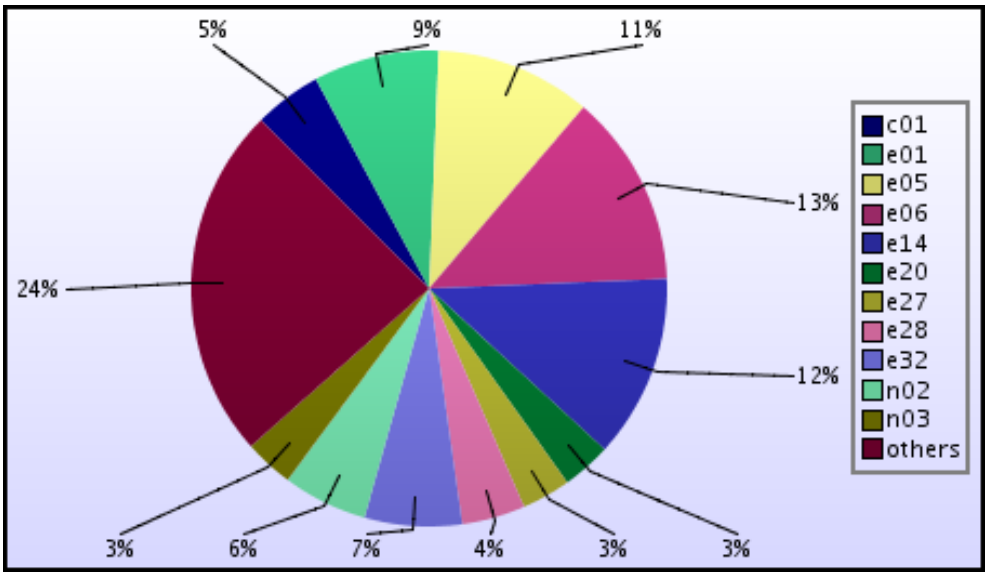

Development Service

Consortium	*CPU Hours (Parallel)*	*CPU Hours (Other)*	*AUs charged*	*%age of charged AUs*
n01	2010	0	9676	2.8%
n02	28077	16	135253	38.9%
n03	42175	0	203049	58.4%
NERC Total	72262	17	347979	100.0%

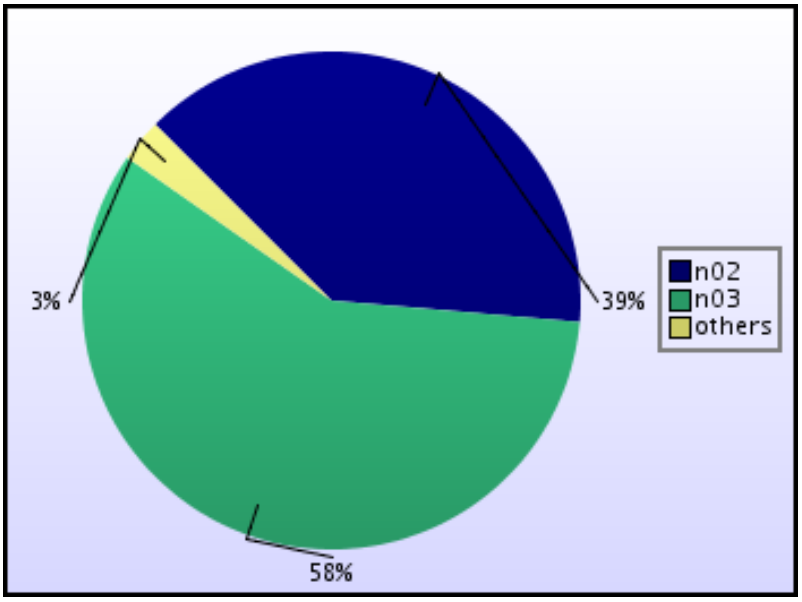

2.3 CPU Usage by Job Type

The figures for *Raw AUs* given here show the number of AUs actually supplied by the system to users' jobs.

Main service

Number of processors	*Raw AUs*	*%age*	*Number of jobs*
≤32	348314	15.0%	2422
33–64	162573	7.0%	419
65–128	535990	23.1%	405
129–256	769873	33.2%	510
257–512	231333	10.0%	208
513–1024	267539	11.6%	141

Utilisation by region

Capacity Region (26 nodes, jobs using ≤128 CPUs): a total of 1046878 raw AUs were used; that is 72.6% of the total available in this region

Capability Region (64 nodes, jobs using >128 CPUs): a total of 1268745 raw AUs were used; that is 35.7% of the total available in this region

The remaining 2 nodes are reserved for interactive-parallel work.

Overall utilisation was 49.7%.

Development Service

Number of processors	*Raw AUs*	*%age*	*Number of jobs*
≤32	213086	61.2%	569
33–64	43909	12.6%	169
65–128	90903	26.1%	232
129–256	0	0.0%	0

Overall utilisation was 57.3%.

2.4 Slowdown and Job Wait Times

Slowdowns

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.

Currently the pattern of slowdowns on both services is satisfactory.

Slowdowns by runtime

The following graphs show the slowdowns recorded for jobs of differing run times, ignoring those which ran for less than 5 minutes.

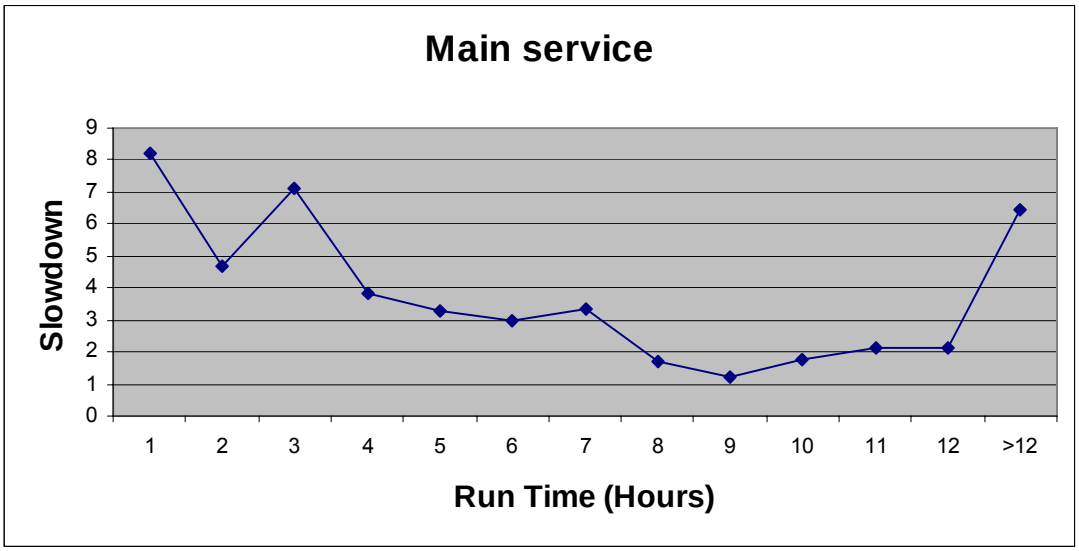

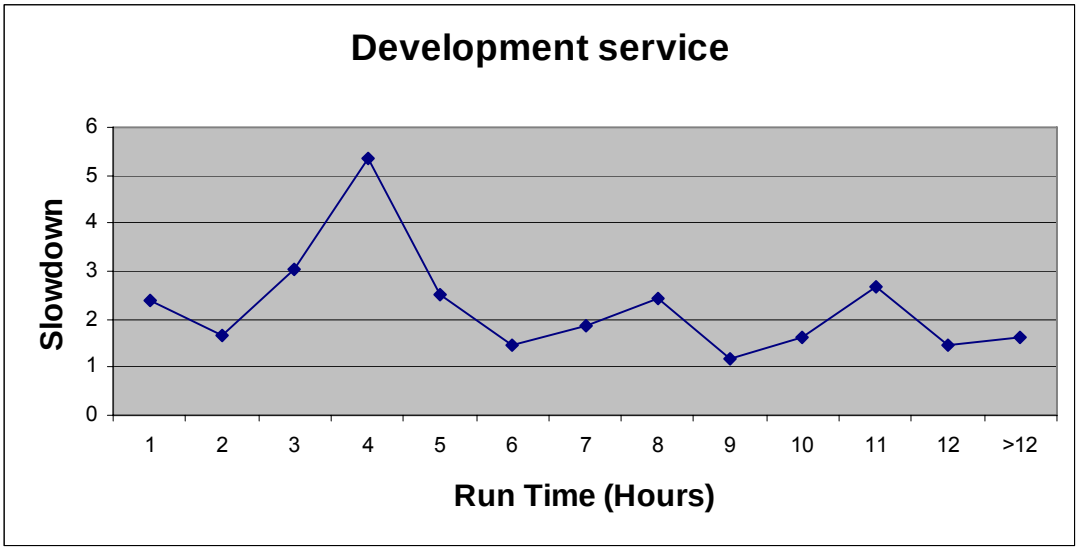

Slowdowns by number of processors

In the graphs below, we plot the slowdown figures against the number of processors used. Jobs which ran for less than 1 hour are ignored.

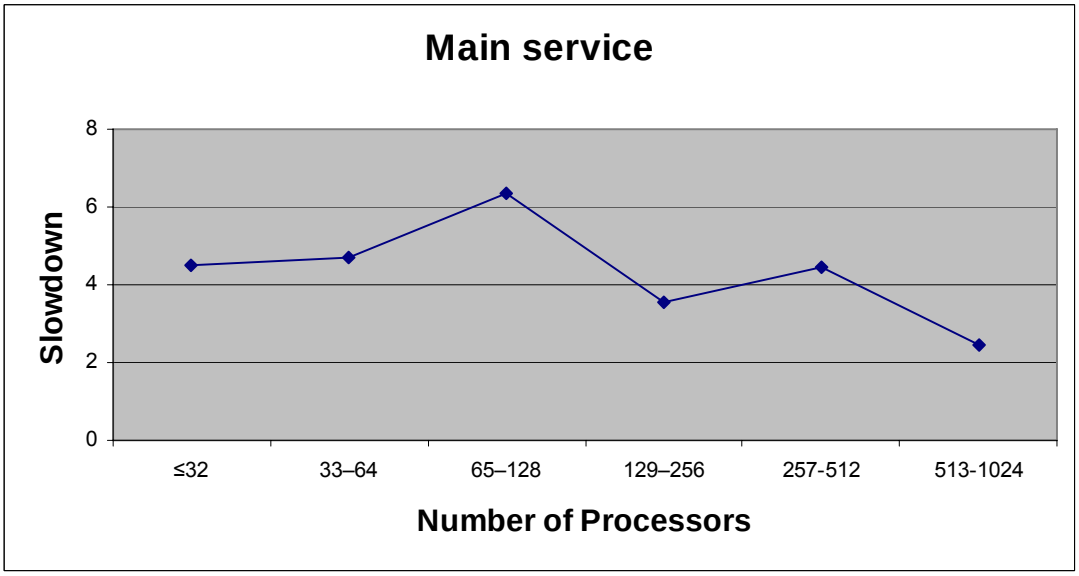

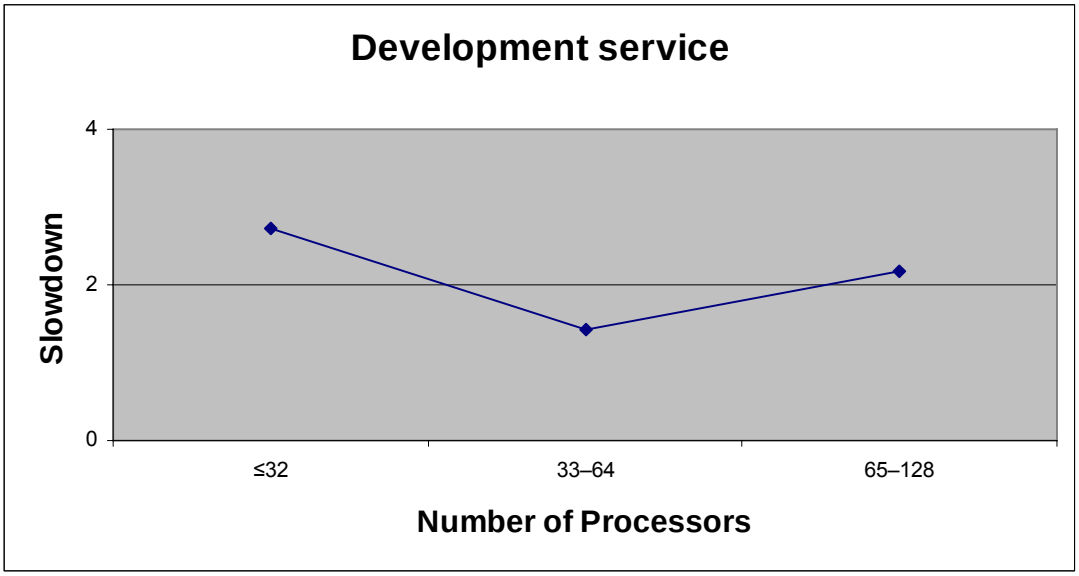

Job wait times

The following table and graph shows the average wait time (in hours) for each class of job on the main service. The somewhat longer wait times for 128-processor jobs are associated with the movement of these jobs into the capacity region, one of the changes at the start of the development service. Utilisation in the capacity region continues to be high (see section 2.3).

Job Class	*Category*	*Maximum Number of CPUs*	*Maximum Job length*	*Average wait time*	*Number of Jobs*
par32_1	parallel	32	1	0.7	1672
par32_3	parallel	32	3	10.0	199
par32_6	parallel	32	6	11.3	551
par64_1	parallel	64	1	0.5	237
par64_3	parallel	64	3	6.0	10
par64_6	parallel	64	6	11.8	172
par128_1	parallel	128	1	2.2	247
par128_3	parallel	128	3	17.7	10
par128_6	parallel	128	6	25.0	52
par128_12	parallel	128	12	35.4	96
par256_1	parallel	256	1	0.7	234
par256_3	parallel	256	3	0.6	29
par256_6	parallel	256	6	6.6	170
par256_12	parallel	256	12	1.0	77
par512_1	parallel	512	1	0.4	172
par512_3	parallel	512	3	2.6	19
par512_6	parallel	512	6	38.8	5
par512_12	parallel	512	12	3.7	12

par1024_1	parallel	1024	1	4.0	124
par1024_3	parallel	1024	3	1.3	6
par1024_6	parallel	1024	6	9.5	11
par1024_12	parallel	1024	12	0.0	0
serial_1	serial	1	1	7.2	257
serial_12	serial	1	12	0.3	185
serial_3	serial	1	3	0.0	2
serial_6	serial	1	6	0.0	12
inter32_1	interactive	32	1	0.0	2042
course16_1	interactive	16	1	0.0	0
course32_1	parallel	32	1	0.0	0

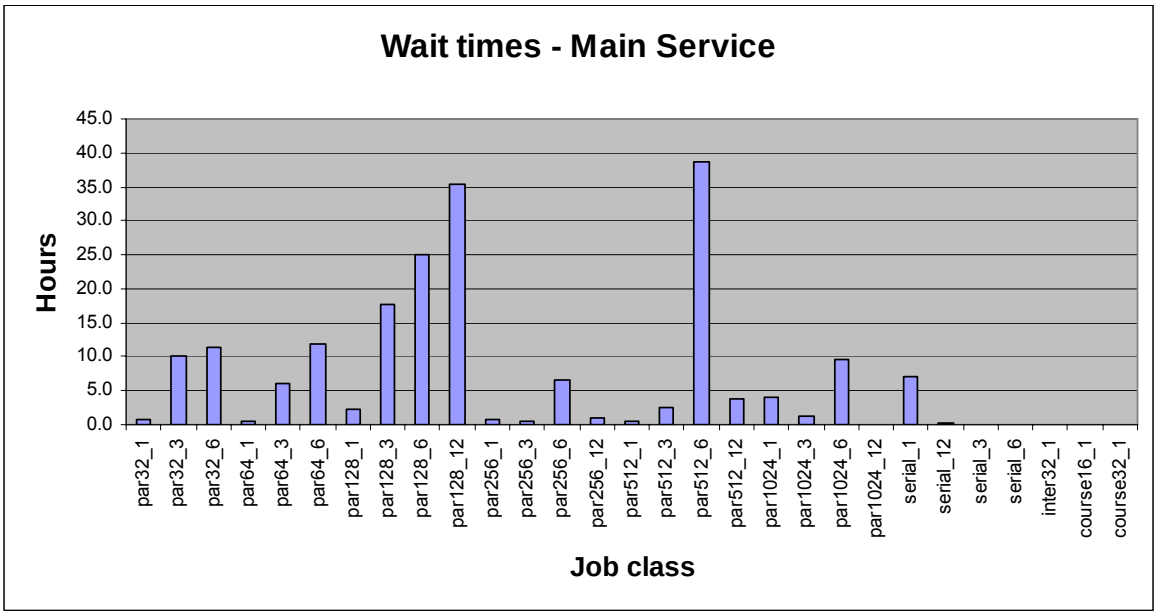

The wait times for the development service are shown below. The very long wait time for the single par128_20m job is because it was submitted immediately before the 63-hour scheduled downtime and did not run until after it.

Job Class	*Category*	*Maximum Number of CPUs*	*Maximum Job length*	*Average wait time*	*Number of Jobs*
parn16_20m	parallel	16	20 mins	0.1	39
parn16_1	parallel	16	1 hour	0.1	94
parn16_6	parallel	16	6 hours	5.5	16
parn16_12	parallel	16	12 hours	15.8	15
parn32_20m	parallel	32	20 mins	0.1	54
parn32_1	parallel	32	1 hour	0.2	100
parn32_6	parallel	32	6 hours	7.6	88
parn32_12	parallel	32	12 hours	6.2	163

parn64_1	parallel	64	1 hour	0.4	169
parn64_6	parallel	64	6 hours	0.0	0
parn64_12	parallel	64	12 hours	0.0	0
parn128_20m	parallel	128	20 mins	90.1	1
parn128_1	parallel	128	1 hour	0.8	231
parn128_6	parallel	128	6 hours	0.0	0
parn128_12	parallel	128	12 hours	0.0	0
serial_1	serial	1	1 hour	0.5	789
serial_12	serial	1	12 hours	0.0	0

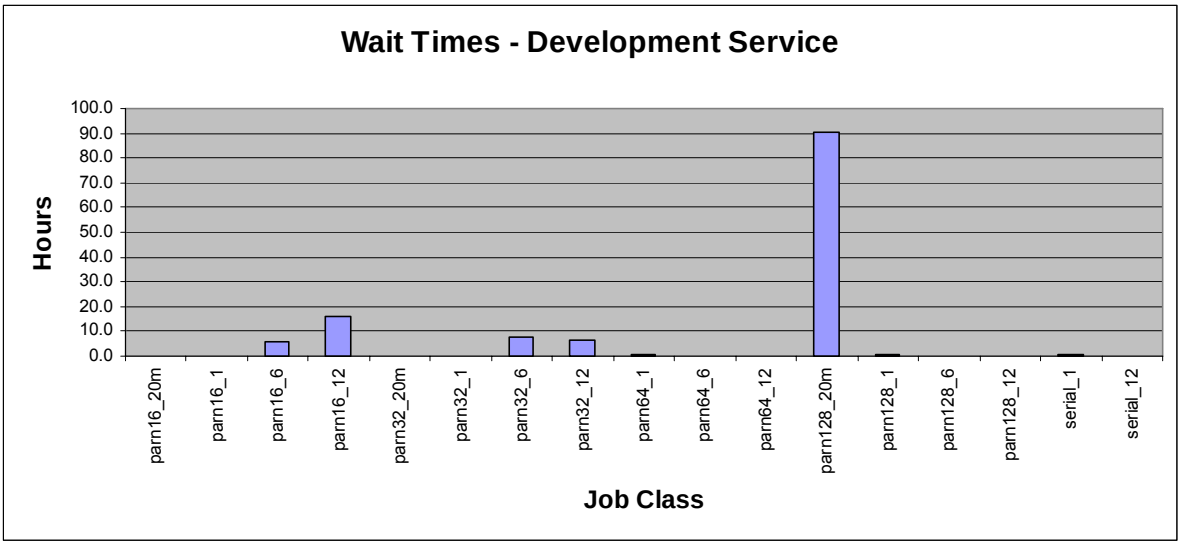

2.5 Disk Occupancy

Home Space

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

Consortium	*Disc Occupancy (Mb)*	*Disc Quota (Mb)*
b02	4,105.7	50,000
b03	4,348.1	50,000
b08	0.2	50,000
c01	97,559.6	100,000
e01	32,243.6	48,829
e02	23,079.3	38,829
e03	55,181.2	225,012
e05	211,740.2	445,550
e06	239,397.8	300,000
e07	10,190.7	20,000
e08	29,427.2	50,000
e10	5,817.2	150,000
e11	39,645.7	100,000
e14	78,113.1	100,000
e15	28,674.2	50,000
e16	132.9	20,000
e17	17,954.7	50,000
e18	37,158.4	40,000
e19	553.7	40,000
e20	58,745.7	60,000
e21	95.5	50,000
e22	127.9	10,000
e23	0.1	50,000
e24	38,054.8	200,000
e25	5,371.6	50,000
e26	12,152.4	20,000
e27	5,403.6	20,000
e28	30,065.2	40,000
e29	2,078.8	30,000
e30	0.1	40,000
e31	42,390.2	50,000
e32	46,895.1	50,000
e33	121.7	50,000
e34	0.1	50,000
e35	193.5	100,000
e36	27,020.6	50,000
e37	30,236.6	100,000
e40	2,527.1	50,000

e41	0.1	100,000
n01	44,288.5	100,000
n02	95,954.7	128,000
n03	54,866.3	100,000
n04	159,615.7	299,999
p01	79,185.3	200,000
x01	43,788.7	50,000
x02	8,746.2	20,000
x03	634.6	50,000
z001	238,576.8	270,001
z002	37,925.9	48,001
z003	0.2	3
z004	71,701.2	100,000
z05	4,188.3	30,000
z06	49,370.0	50,000
z07	21,050.7	30,000
z09	15,366.4	50,000

Workspace

Consortium	*Disc Occupancy (Mb)*	*Disc Quota (Mb)*
b02	14.8	1,025
b03	49,407.8	100,000
b08	5,019.5	50,000
c01	91,902.7	100,000
e01	1,073,600.2	1,149,995
e02	8,354.8	10,000
e03	9.8	500,000
e05	208,649.9	487,804
e06	369,354.2	400,000
e07	57,189.6	99,999
e08	140.8	5,000
e10	284,222.9	400,000
e11	14,886.8	100,000
e14	101,894.5	150,000
e15	18,238.4	100,000
e16	0.2	60,000
e17	1,686.0	100,000
e18	10,434.3	80,000
e19	168,861.6	200,000
e20	895,538.3	1,000,000
e21	1.0	100,000
e22	0.1	20,000
e23	0.1	100,000
e24	28,122.1	500,000
e25	126,477.6	150,000

e26	0.1	40,000
e27	0.3	40,000
e28	46,869.5	80,000
e29	5,296.1	8,000
e30	0.1	80,000
e31	88,585.8	100,000
e32	56,271.6	100,000
e33	1,951.2	100,000
e34	0.1	100,000
e35	0.1	200,000
e36	0.1	50,000
e37	9,740.8	150,000
e40	0.2	100,000
e41	0.1	200,000
n01	411,667.1	500,000
n02	1,642,665.7	1,999,003
n03	60.3	41,002
n04	638,720.8	750,000
p01	41,764.4	50,000
x01	91,810.5	100,000
x02	0.2	20,000
x03	177.6	50,000
z001	385,496.0	399,999
z002	389.3	770
z003	0.2	3
z004	23,739.8	25,000
z05	856.7	20,000
z06	71,211.6	100,000
z07	1.7	1
z09	22,327.8	100,000

Development space

This is the disk space reserved for users of the development service.

Consortium	*Disc Occupancy (Mb)*	*Disc Quota (Mb)*
n01	0.0	500,000
n02	186,186.7	5,210,003

2.6 Tape Archive

Consortium	*Usage (Tapes)*	*Quota (Tapes)*	*Files*	*Data (Gb)*
c01	2	2	17	17
e01	36	38	37270	3428
e03	5	5	18797	429
e14	8	10	19164	178
e26	2	2	545	27
n01	144	160	17553	14852
n02	146	180	86479	18858
n04	30	30	105130	2934
z001	2	10	6189	50
z002	4	4	5810	15
z06	1	3	833	68

Note that a tape is counted in the *Usage* column even if it is only partly occupied.

3 Support

3.1 Helpdesk

Classifications

Category	*Number*	*% of all*
Administrative	47	63.5
Technical	27	36.5
In-depth	0	0.0
PMR	0	0.0
TOTAL	74	100.0

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

Service Area	*Number*	*% of all*
Phase 2 platform	63	85.1
Website	1	1.4
Other/general	10	13.5
TOTAL	74	100.0

Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	63	85.1	75%
Finished within 72 Hours	73	98.6	97%
Finished after 72 Hours	1	1.4	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	47	100.0	97%
Finished after 48 Hours	0	0.0	

Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*	*PMR*
epcc.ed.ac.uk	42	10	0	0
dl.ac.uk	1	5	0	0
Sysadm	4	11	0	0
Other people	0	1	0	0

3.2 Training

Title of Course	*Start Date*	*Length (Days)*	*Place days*	*HPCx User Days*	*HPCx Staff Days*
Fundamental Concepts of HPC	18-Apr	3	75	45	15
Practical Software Development	25-Apr	4	100	24	20

Training courses planned for May include:

- *Shared Memory Programming*
  Tuesday, 2 May - Thursday, 4 May
- *Message Passing Programming*
  Tuesday, 9 May - Thursday, 11 May
- *Parallel Decomposition*
  Wednesday, 24 May - Friday, 26 May
- *Porting Codes from CSAR to HPCx*
  Thursday, 25 May (in conjunction with CSAR)
- *Applied Numerical Algorithms*
  Tuesday, 30 May - Friday, 2 June

4 Staffing

4.1 Science Support Staffing

Daresbury Laboratory

Name	*Days*
Ashworth	7.5
Blake	1.6
Bush	18.0
Guest	4.5
Johnstone	9.0
Jones	3.3
Plummer	18.0
Sherwood	2.3
Sunderland	18.0
Thomas	1.5
Pickles	1.5
van Dam	2.1
Total (Days)	87.2
FTEs	4.9

EPCC

Name	*Days*
Simpson	11.6
Booth	16.8
Henty	11.6
Smith	12.7
Bull	4.5
Fisher	5.0
Hein	5.8
Jackson	4.6
Pringle	1.9
Reid	6.1
Stratford	2.4
Nazarova	3.2
Trew	4.3
Gray	2.1
D'Mellow	14.8
Hill	2.9

Beckett	5.2
Training	7.5
Total (Days)	122.6
FTEs	6.9

Overall Levels

	FTEs
DL	4.9
EPCC	6.9
Total	11.8

4.2 Systems Staffing

Name	*Days*
Andrews	10.5
Blake	0.0
Brown	18.0
Fisher	7.0
Georgeson	13.5
Franks	10.5
Jones	0.8
Shore	13.5
BITD	18.0
Total (days)	91.8
FTEs	5.2

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 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%	99.7%
Technology MTBF (hours)	200	300	732
Number of AV FTEs	7.5	10	11.8
Number of training days per month	20/12	25/12	7/4
Non in-depth queries resolved within 3 days	85%	97%	98.6%
Number of A&M FTEs	3.75	5.75	5.2
A&M serviceability	80%	99.6%	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: Projects

B.1 Current Projects

EPSRC Projects

Code	*Class*	*Title*	*PI*
e01	1	UK Trubulence Consortium	Dr Gary Coleman
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
e10	1	Reality Grid	Prof Peter Coveney
e11	1	Bond making and breaking at surfaces	Prof Sir David A King
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
e18	1	DARP: Highly swept leading edge separations	Prof Michael A Leschziner
e19	1	Edinburgh Soft Matter and Statistical Physics Group	Prof Michael E Cates
e20	1	UK Applied Aerodynamics Consortium	Dr Ken Badcock
e21	1	Intrinsic Parameter Fluctuations in Decananometer MOSFETs	Prof Asen M Asenov
e22	1	Preconditioners for finite element problems	Prof David J Silvester
e23	1	Exploitation of Switched Lightpaths for e-Science Applications	Prof Peter Clarke
e24	1	DEISA - Distributed European Infrastructure for Supercomputing Applications	Dr David Henty
e25	1	Turbulent vortex motion in stratified flows	Dr Gary Coleman
e26	1	Simulation of Radioprobing	Dr Charlie Laughton
e27	1	SPICE	Prof Peter V Coveney

e28	1	Towards the Dynome	Dr Jonathan W Essex
e29	1	Free-surface-piercing circular cylinders	Dr Eldad Avital
e30	1	Metal/Oxide Interfaces at the Atomic Level	Dr Nora de Leeuw
e31	1	Lateral Straining of Wall-Bounded Turbulence	Dr Gary N Coleman
e32	1	Rapid Prototyping of Usable Grid Middleware	Prof Peter V Coveney
e33	1	Engineering Functional Coatings	Prof Roger Smith
e34	1	Dissolution of Bioactive Phosphate Glasses	Dr N de Leeuw
e35	1	Non-adiabatic processes	Dr T Todorov
e36	1	Jets in Cross-Flow	Dr Y Yao
e37	1	LESUK_3	Prof JJ McGuirk
e40	1	Computational Quantum Many-Body Theory	Prof R Needs
e41	1	Flow in Weapon Bays	Dr George N Barakos
z09		HECToR Benchmarking	Dr Edward Smyth

PPARC Projects

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

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

BBSRC Projects

Code	*Class*	*Title*	*PI*
b02	1	Modelling enzyme catalysis	Dr Adrian J Mulholland
b08	1	IntBioSim	Prof M S Sansom

CCLRC Projects

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

Externally-funded Projects

Code	*Title*	*PI*
x01	HPC-Europa	Dr J-C Desplat
x03	IBM	Mr Derrick J Byford

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

Note: The original project e01 ended on 30 April 2005. The new UKTC project started on 1 March 2006. At the request of the PI it was assigned the same code as the old one, and inherited its disk space.

B.2 Former Projects

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

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
e09	2	Molecular Properties and their Geometry	Dr Mark R Wilson
e12	1	Parallel programs for the simulation of complex fluids	Dr Richard J Blake
e13	1	TeraGyroid project	Mr Mark Westwood
x02		OHM Ltd	Dr Mervyn Freeman
n05	2	Non-linear Wave-particle Instabilities in Plasmas	