HPCx Service Report
March 2006

1 Introduction

This report covers the period from 1 March 2006 at 0800 to 1 April 2006 at 0800. Taking into account the change to summer time, this is a service month of 743 hours.

Overall utilisation has recovered strongly, to nearly 84%. Indeed, we delivered more than 4.3 millions AUs, the highest monthly number so far. Capability usage was almost 40% of the total.

2 Usage

2.1 Availability

Incidents

During this month, there were 9 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	2
3	5
4	0

The MTBF figures for this month were as follows:

SEV1	*Incidents*	*MTBF*
IBM	1.0	732
Site	0.0	∞
External	1.0	732
Overall	2.0	366

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

Failure	*Site*	*IBM*	*External*	*Reason*
06.028	0%	0%	100%	Network loss at Warrington
06.029	0%	100%	0%	Router failure disconnected 12 nodes from switch

Serviceability

There was a total of 10.4 hours of scheduled downtime this month.

Attribution	*UDT*	*Serviceability*
IBM	1:15	99.8
Site	0:00	100.0
External	11:45	98.4
Overall	13:00	98.2

2.2 CPU Usage by Consortium

Main Service

Consortium	*CPU Hours (Parallel)*	*CPU Hours (Other)*	*AUs charged*	*%age of charged AUs*
e01	55276	79	266502	6.2%
e03	0	0	0	0.0%
e04	26177	32	126181	2.9%
e05	105266	180	507662	11.8%
e06	116569	50	559779	13.0%
e07	4201	0	20226	0.5%
e08	7438	0	35808	0.8%
e10	171	1	827	0.0%
e11	8130	0	39140	0.9%
e14	74018	28	356490	8.3%
e15	683	0	3288	0.1%
e17	102	28	625	0.0%
e18	0	0	0	0.0%
e20	92206	1039	448925	10.5%
e24	64977	0	312827	7.3%
e25	1919	83	9638	0.2%
e26	1731	0	8338	0.2%
e27	3	0	12	0.0%
e28	28699	0	117446	2.7%
e29	468	0	2251	0.1%
e31	1060	20	5200	0.1%
e32	23045	0	77665	1.8%
e35	16	0	79	0.0%
e36	14	0	65	0.0%
e37	1543	0	7429	0.2%
e40	7400	0	35625	0.8%
EPSRC Total	621109	1541	2942029	68.5%
n01	97806	14	470949	11.0%
n02	65449	20	315201	7.3%
n03	44043	242	213206	5.0%
n04	21280	368	104227	2.4%
NERC Total	228579	644	1103583	25.7%

p01	562	0	2708	0.1%
PPARC Total	562	0	2708	0.1%

c01	33371	871961	161176	3.8%
CCLRC Total	33371	871961	161176	3.8%

BBSRC Total	55	0	265	0.0%

x01	2750	100	13724	0.3%
x03	2891	0	13919	0.3%
External Total	5641	100	27643	0.6%

z001	11072	100	53784	1.3%
z004	0	19	94	0.0%
z06	193	1	936	0.0%
HPCx Total	11265	120	54814	1.3%

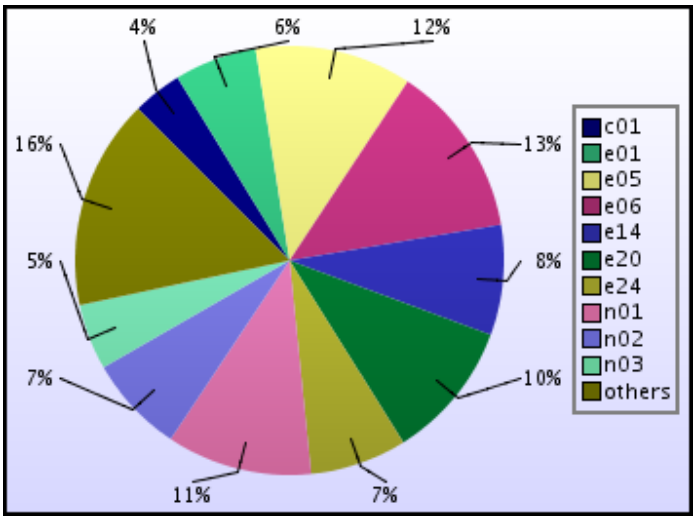

Development Service

Consortium	*CPU Hours (Parallel)*	*CPU Hours (Other)*	*AUs charged*	*%age of charged AUs*
n01	2203	0	10606	2.6%
n02	32661	11	157299	38.3%
n03	50337	0	242344	59.1%
NERC Total	85201	11	410249	100.0%

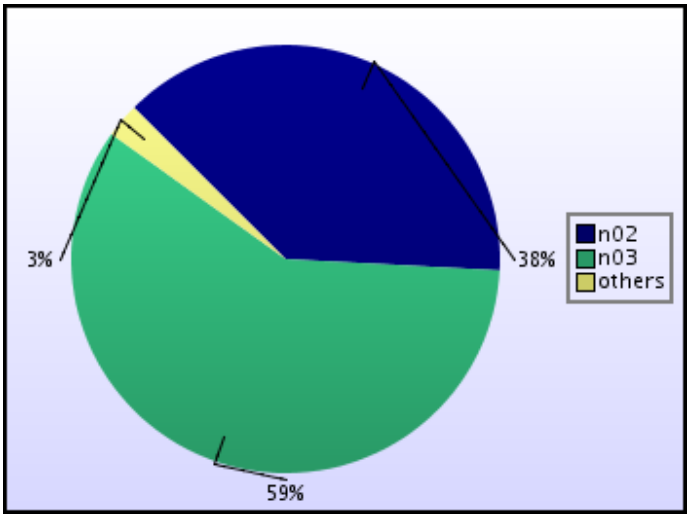

2.3 CPU Usage by Job Type

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	365582	8.4%	2912
33–64	290777	6.7%	671
65–128	785989	18.1%	525
129–256	1221218	28.2%	600
257–512	1458733	33.6%	216
513–1024	213508	4.9%	25

Utilisation by region

Capacity Region (26 nodes, jobs using ≤128 CPUs): a total of 1442348 raw AUs were used; that is 96.9% of the total available in this region

Capability Region (64 nodes, jobs using >128 CPUs): a total of 2893459 raw AUs were used; that is 79.0% of the total available in this region

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

Overall utilisation was 83.5%.

Development Service

Number of processors	*Raw AUs*	*%age*	*Number of jobs*
≤32	263847	64.3%	946
33–64	47879	11.7%	225
65–128	98469	24.0%	204
129–256	0	0.0%	0

Overall utilisation was 60.6%.

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 is in general satisfactory, even though the service is extremely busy. Slowdowns for short jobs on the development service have improved following the scheduling changes which we made in February.

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.

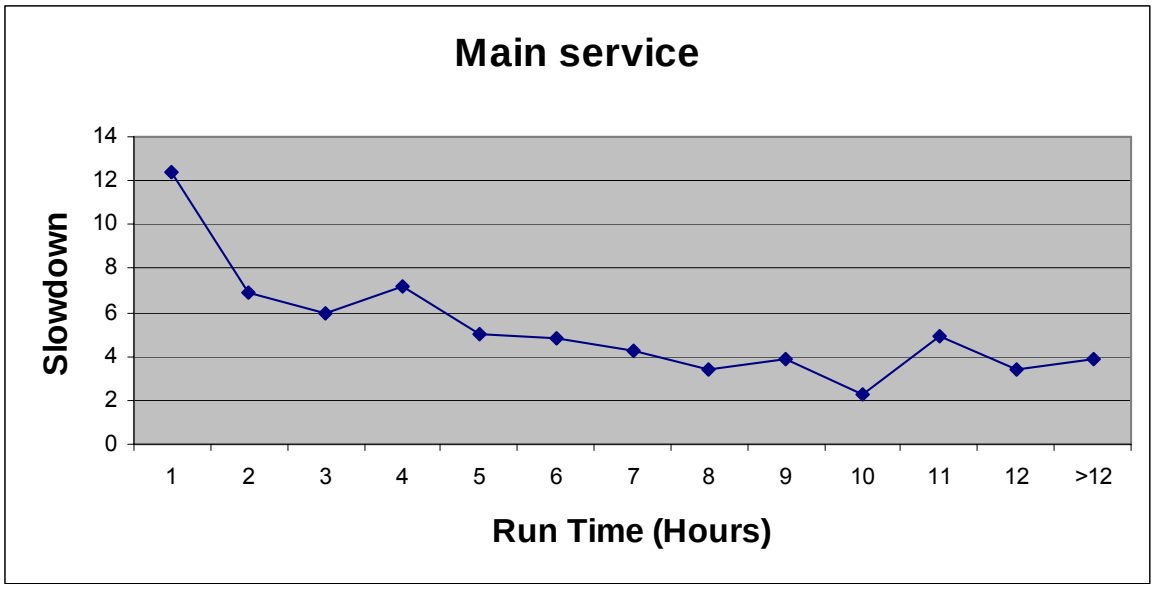

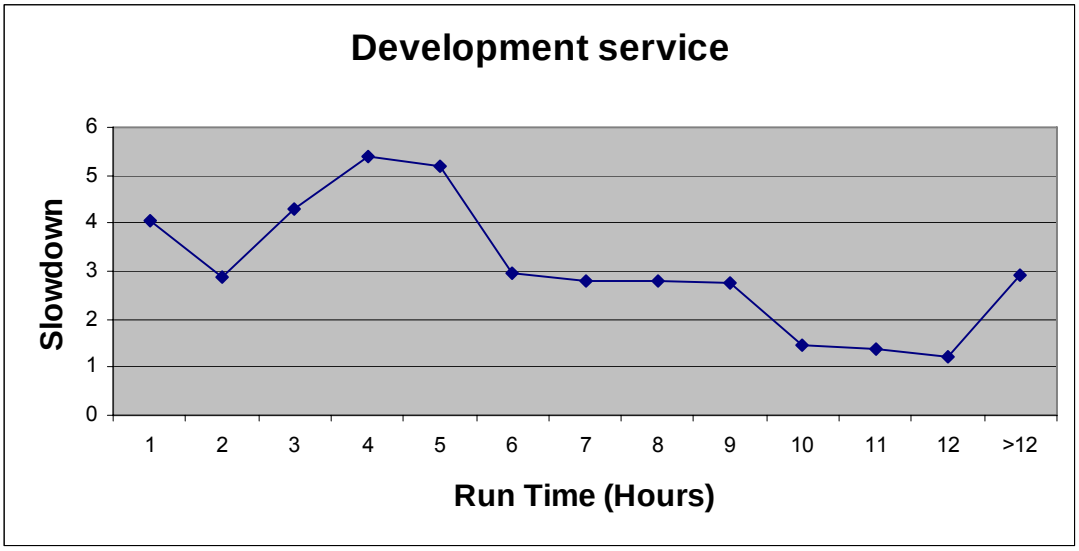

Slowdowns by number of processors

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.

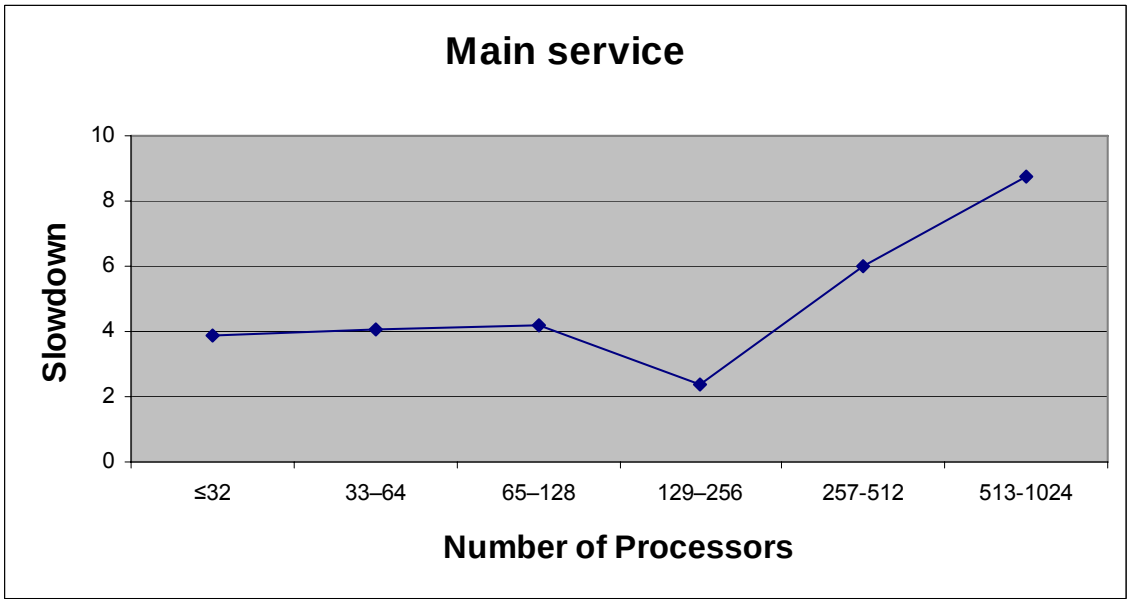

Most of the 513-1024-processor jobs were short ones;

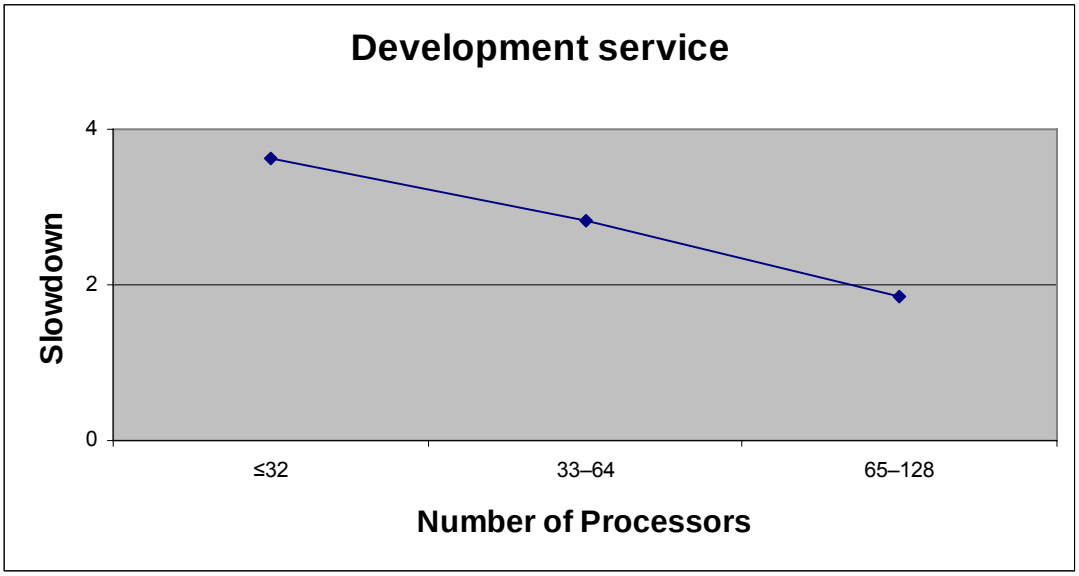

Job wait times

The following table and graph shows the average wait time (in hours) for each class of job on the main service. Wait times for some classes have lengthened, but this is unavoidable, given the exceptionally heavy load on the service. 128-processor jobs, for example, run in the capacity region, where this month the load was 97%; so an average wait time of 46 hours for 12-hour 128-processor jobs is not unexpected.

Job Class	*Category*	*Maximum Number of CPUs*	*Maximum Job length*	*Average wait time*	*Number of Jobs*
par32_1	parallel	32	1	1.3	2157
par32_3	parallel	32	3	9.8	214
par32_6	parallel	32	6	15.2	541
par64_1	parallel	64	1	2.1	408
par64_3	parallel	64	3	11.5	27
par64_6	parallel	64	6	25.9	235
par128_1	parallel	128	1	6.2	306
par128_3	parallel	128	3	26.9	26
par128_6	parallel	128	6	29.1	33
par128_12	parallel	128	12	46.0	160
par256_1	parallel	256	1	2.3	260
par256_3	parallel	256	3	3.7	53
par256_6	parallel	256	6	10.3	165
par256_12	parallel	256	12	11.8	123
par512_1	parallel	512	1	4.5	93
par512_3	parallel	512	3	24.8	19
par512_6	parallel	512	6	43.6	12

par512_12	parallel	512	12	21.3	92
par1024_1	parallel	1024	1	30.6	20
par1024_3	parallel	1024	3	0.0	0
par1024_6	parallel	1024	6	0.0	0
par1024_12	parallel	1024	12	20.1	5
serial_1	serial	1	1	1.3	1373
serial_12	serial	1	12	0.6	121
serial_3	serial	1	3	0.0	11
serial_6	serial	1	6	0.0	83
inter32_1	interactive	32	1	0.0	4474
course16_1	interactive	16	1	0.0	488
course32_1	parallel	32	1	0.0	0

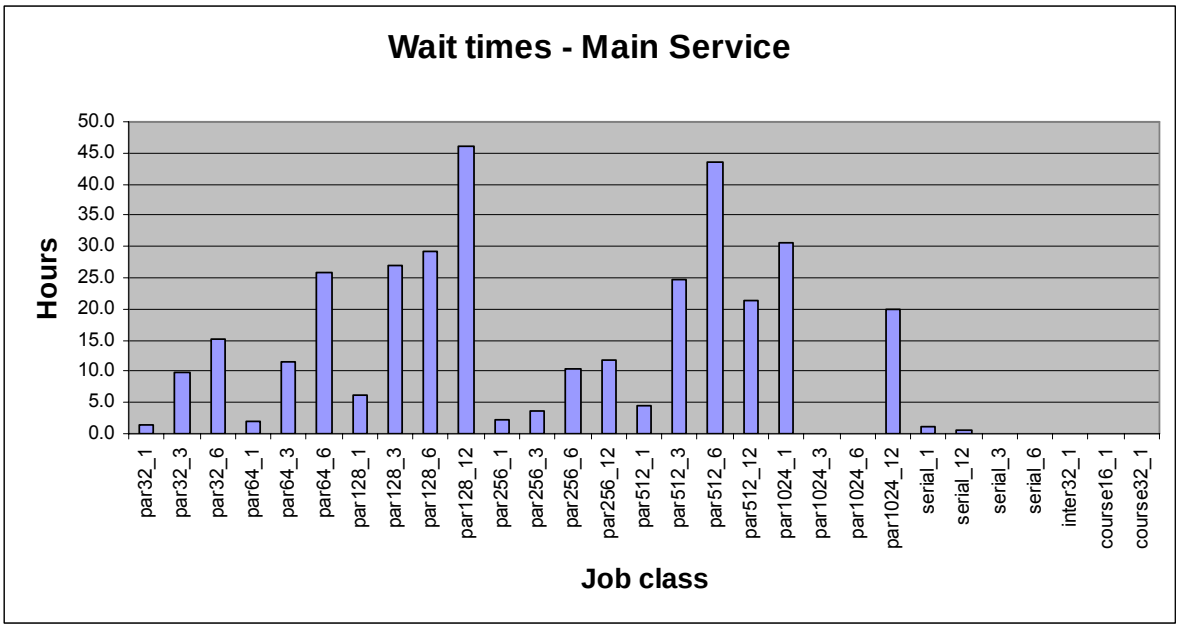

The wait times for the development service are shown below. One class shows a moderately long average wait time; we are consulting with the user groups on this.

Job Class	*Category*	*Maximum Number of CPUs*	*Maximum Job length*	*Average wait time*	*Number of Jobs*
parn16_20m	parallel	16	20 mins	0.0	142
parn16_1	parallel	16	1 hour	0.1	79
parn16_6	parallel	16	6 hours	23.8	50
parn16_12	parallel	16	12 hours	0.0	1
parn32_20m	parallel	32	20 mins	0.1	197
parn32_1	parallel	32	1 hour	0.2	119
parn32_6	parallel	32	6 hours	7.8	207

parn32_12	parallel	32	12 hours	10.7	151
parn64_1	parallel	64	1 hour	1.8	225
parn64_6	parallel	64	6 hours	0.0	0
parn64_12	parallel	64	12 hours	0.0	0
parn128_1	parallel	128	1 hour	1.1	204
parn128_6	parallel	128	6 hours	0.0	0
parn128_12	parallel	128	12 hours	0.0	0
serial_1	serial	1	1 hour	0.1	351
serial_12	serial	1	12 hours	0.0	5

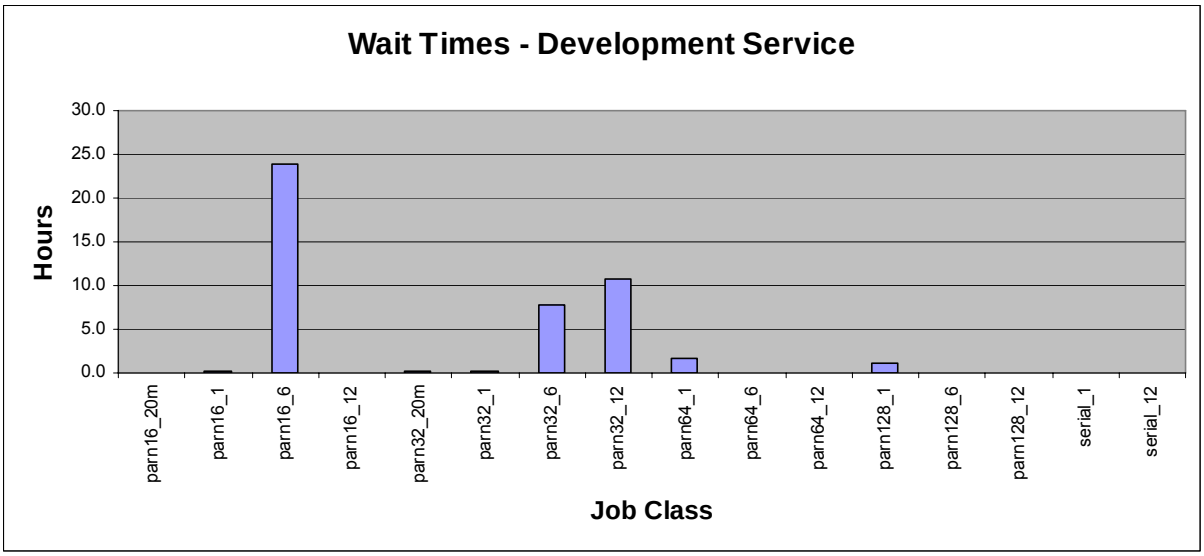

2.5 Disk Occupancy

Home Space

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

Consortium	*Disc Occupancy (Mb)*	*Disc Quota (Mb)*
b02	4106	50,000
b03	4348	50,000
b08	0	50,000
c01	93821	100,000
e01	27993	48,834
e02	23079	38,829
e03	55129	225,012
e04	53407	100,000
e05	208319	445,550
e06	278827	300,000
e07	10679	20,000
e08	49961	50,000
e10	5780	10,000
e11	39646	100,000
e14	91835	100,000
e15	28627	50,000
e16	133	20,000
e17	18275	50,000
e18	38170	40,000
e19	43	40,000
e20	55170	60,000
e21	96	50,000
e22	128	10,000
e23	0	50,000
e24	1075	50,000
e25	5368	50,000
e26	16313	20,000
e27	3401	20,000
e28	19414	40,000
e29	2069	30,000
e30	0	40,000
e31	45198	50,000
e32	31610	50,000
e33	122	50,000
e34	0	50,000
e35	101	100,000
e36	2090	50,000
e37	11644	100,000

e40	1342	50,000
n01	44882	100,000
n02	92336	128,000
n03	47281	100,000
n04	167355	299,999
p01	55280	200,000
x01	41409	50,000
x02	8746	20,000
x03	795	50,000
z001	237581	270,001
z002	44906	48,001
z003	0	3
z004	71699	100,000
z05	4188	30,000
z06	49858	50,000
z07	21002	30,000
z09	15366	50,000

Workspace

Consortium	*Disc Occupancy (Mb)*	*Disc Quota (Mb)*
b02	15	1,025
b03	60381	100,000
b08	0	50,000
c01	80576	100,000
e01	988061	1,150,000
e02	8355	10,000
e03	10	500,000
e04	1305921	3,200,000
e05	211046	487,804
e06	283978	400,000
e07	53148	99,999
e08	141	5,000
e10	284223	300,000
e11	14887	100,000
e14	93432	150,000
e15	18238	100,000
e16	0	60,000
e17	1639	100,000
e18	6498	80,000
e19	168862	200,000
e20	858775	1,000,000
e21	1	100,000
e22	0	20,000
e23	0	100,000
e24	250179	300,000

e25	93090	150,000
e26	0	40,000
e27	0	40,000
e28	37908	80,000
e29	5296	8,000
e30	0	80,000
e31	90874	100,000
e32	484	100,000
e33	1162	100,000
e34	0	100,000
e35	0	200,000
e36	0	50,000
e37	251	150,000
e40	0	100,000
n01	368326	500,000
n02	1412270	1,999,003
n03	60	41,002
n04	654568	750,000
p01	41764	50,000
x01	63093	100,000
x02	0	20,000
x03	167	50,000
z001	382139	399,999
z002	290	770
z003	0	3
z004	23740	25,000
z05	0	1,000
z06	73415	100,000
z07	2	1
z09	22328	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	500,000
n02	154,354	5,210,003

2.6 Tape Archive

Consortium	*Usage (Tapes)*	*Quota (Tapes)*	*Files*	*Data (Gb)*
c01	2	2	17	17
e01	38	38	37407	3616
e03	5	5	18797	429
e04	4	14	1260	254
e14	8	10	19164	178
e26	2	2	516	24
n01	143	160	17387	14784
n02	137	180	83745	17746
n04	21	30	75505	2620
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	60	58.3
Technical	37	35.9
In-depth	5	4.9
PMR	1	1.0
TOTAL	103	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	91	88.3
Website	2	1.9
Other/general	10	9.7
TOTAL	103	100.0

Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	87	89.7	75%
Finished within 72 Hours	97	100.0	97%
Finished after 72 Hours	0	0.0	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	60	100.0	97%
Finished after 48 Hours	0	0.0	

Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*	*PMR*
epcc.ed.ac.uk	52	17	3	0
dl.ac.uk	1	8	0	0
Sysadm	6	12	2	0
Other people	1	0	0	1

3.2 Training

There were no training courses in March. Training courses planned for the next two months include:

- *Fundamental Concepts of High Performance Computing*
  Tuesday, 18 April - Thursday, 20 April
- *Practical Software Development*
  Tuesday, 25 April - Thursday, 28 April
- *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	13.2
Blake	2.3
Bush	21.0
Guest	5.8
Johnstone	9.8
Jones	4.2
Plummer	23.0
Sherwood	2.9
Sunderland	23.0
Thomas	11.5
Pickles	2.1
van Dam	2.0
Total (Days)	120.7
FTEs	6.8

EPCC

Name	*Days*
Simpson	15.8
Booth	21.7
Henty	9.3
Smith	12.6
Bull	5.0
Fisher	9.0
Hein	18.3
Pringle	2.1
Reid	3.7
Stratford	6.2
Nazarova	11.6
Trew	4.3
Gray	9.7
D'Mellow	15.0
Hill	3.9
Johnson	10.8

Helpdesk	0.7
Total (Days)	159.6
FTEs	9.0

Overall Levels

	FTEs
DL	6.8
EPCC	9.0
Total	15.8

4.2 Systems Staffing

Name	*Days*
Andrews	17.3
Blake	0.0
Brown	23.0
Fisher	11.0
Georgeson	15.0
Franks	17.3
Jones	1.1
Shore	16.5
BITD	23.0
Total (days)	124.1
FTEs	7.0

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%	99.8%
Technology MTBF (hours)	200	300	732
Number of AV FTEs	7.5	10	15.8
Number of training days per month	20/12	25/12	0/3
Non in-depth queries resolved within 3 days	85%	97%	100.0%
Number of A&M FTEs	3.75	5.75	7.0
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 Turbulence Consortium	Dr Gary Coleman
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
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
z09		HECToR Benchmarking	Dr Edward Smyth

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.

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

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