HPCx Service Report
December 2005

1 Introduction

This report covers the period from 1 December 2005 at 0800 to 1 January 2006 at 0800, a service month of 744 hours.

Utilisation continues to be high at nearly 83%, despite the holiday period. Capability usage was nearly 57% of the total, a new record for the service, beating the previous record of 48.3% last month.

2 Usage

2.1 Availability

Incidents

During this month, there were 13 incidents. As a result of problems in the transition to Phase 2A, they included five 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	5
2	3
3	5
4	0

The attributions for the SEV 1 incidents were as follows:

SEV1	*Incidents*	*MTBF*
IBM	2.0	366
Site	3.0	244
External	0.0	∞
Overall	5.0	146

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

Failure	*Site*	*IBM*	*External*	*Reason*
05.158	0%	100%	0%	GPFS failure on /hpcx/home
05.159	0%	100%	0%	GPFS failure on /hpcx/home
05.162	100%	0%	0%	Site firewall failure
05.163	100%	0%	0%	Site firewall failure
05.164	100%	0%	0%	Capability region drained in error

Serviceability

There was a total of 9.4 hours of scheduled downtime this month.

Attribution	*UDT*	*Serviceability*
IBM	12:35	98.3
Site	17:13	97.7
External	0:00	100.0
Overall	29:48	95.9

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 (Other)*	*AUs charged*	*%age of charged AUs*
e03	191090	220	921055	22.2%
e04	38091	175	166613	4.0%
e05	153159	469	712989	17.2%
e06	26785	19	113129	2.7%
e07	112	0	539	0.0%
e08	3193	0	15372	0.4%
e10	496	0	2389	0.1%
e11	22616	0	108883	2.6%
e14	41460	31	199755	4.8%
e17	6602	8	31824	0.8%
e20	32959	34	158845	3.8%
e24	43	0	207	0.0%
e25	590	560	5535	0.1%
e26	602	1	2905	0.1%
e27	0	0	0	0.0%
e28	3	1	17	0.0%
e31	955	4	4615	0.1%
e33	21	0	100	0.0%
EPSRC Total	518776	1522	2444770	58.9%
n01	48553	8	233793	5.6%
n02	65042	11	313193	7.5%
n03	67219	0	259334	6.3%
n04	1401	127	7353	0.2%
NERC Total	182215	145	813673	19.6%
p01	24	1	121	0.0%
PPARC Total	24	1	121	0.0%
c01	6213	68	30242	0.7%
CCLRC Total	6213	68	30242	0.7%

b02	56055	11	235742	5.7%
b03	102197	0	492022	11.9%
BBSRC Total	158252	11	727764	17.5%

x01	13838	0	66623	1.6%
External Total	13838	0	66623	1.6%

z001	13370	55	64637	1.6%
z002	261	1	1263	0.0%
z004	0	1	7	0.0%
z06	4	1	24	0.0%
HPCx Total	13636	59	65930	1.6%

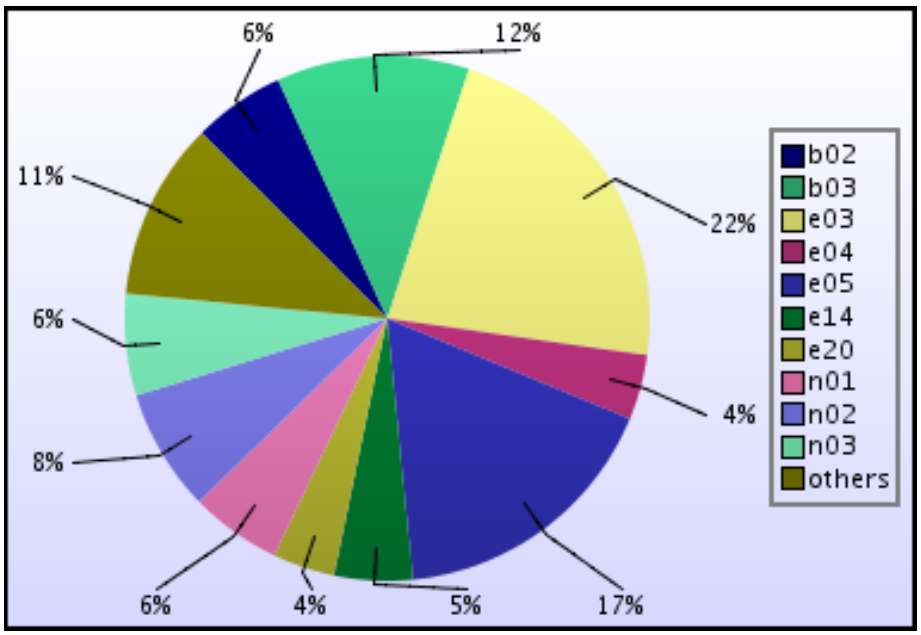

2.3 CPU Usage by Job Type

The figures for *Raw AUs* given here show the number of AUs actually supplied by the system to users' jobs. This month, because of the upgrade, two conversion rates for the AU were used, corresponding to the Linpack benchmark results for the two phases:1 CPU hour = 3.8675 AUs for Phase 2, and 1 CPU hour = 4.81445 AUs for Phase 2A.

Number of Processors	*Raw AUs*	*%age*	*Number of Jobs*
≤32	401939	9.3%	2539
33–64	127223	3.0%	342
65–128	674302	15.7%	385
129–256	647991	15.1%	207
257–512	1140590	26.5%	147
513–1024	1307035	30.4%	50

Utilisation by Region

The system is divided into 3 regions.

Capacity Region (26 nodes, 416 processors, jobs ≤128 processors): 1203464 AUs were used, 81.8% of the total available in the region

Capability region (64 nodes, 1024 processors, jobs >128 processors): 3095616 AUs were used, 85.5% of the total available in the region.

Two further nodes (32 processors) are reserved for interactive parallel jobs.

2.4 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 graph below plots slowdown against run-time (ignoring jobs of less than 5 minutes duration). Despite the continuing heavy utilisation of the system, slowdowns are in general satisfactory.

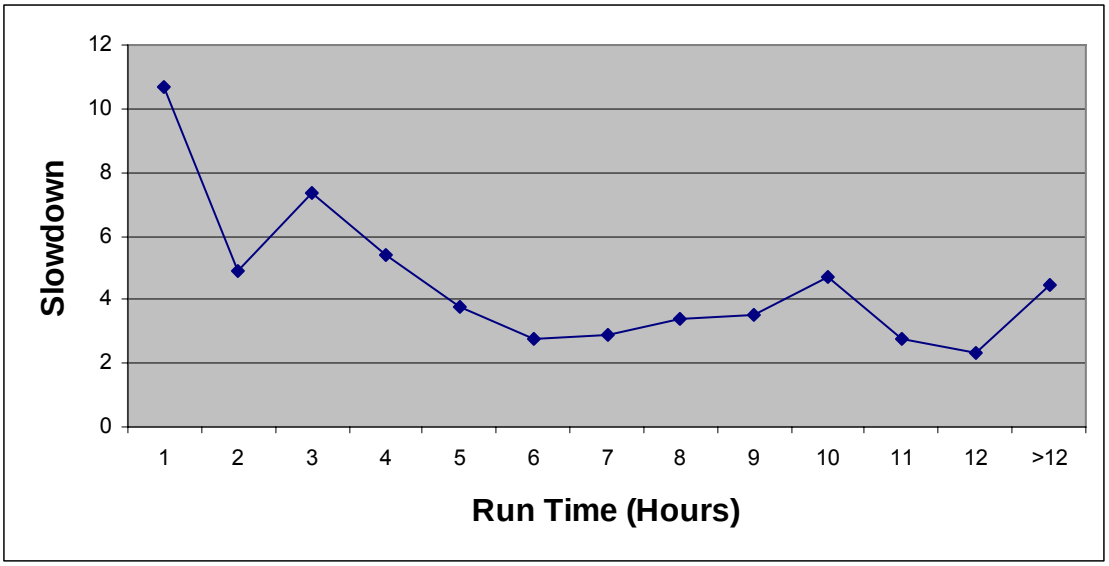

In the graph below, we plot the slowdown figures against the number of processors used and ignoring the development jobs of less than 1 hour. These figures are also satisfactory in general.

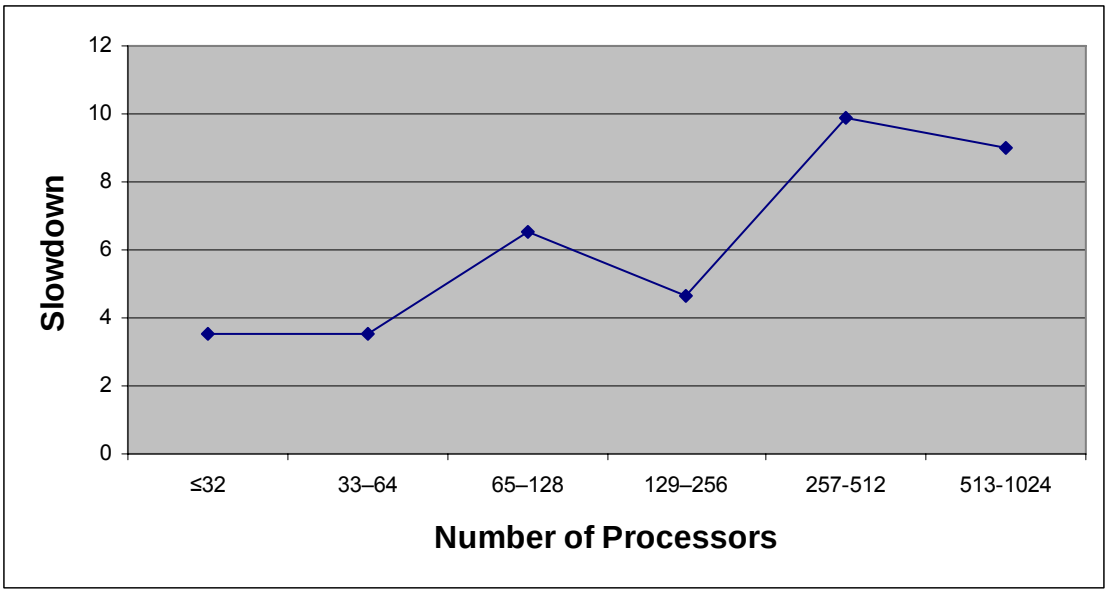

Job wait times

The following table and graph shows the average wait time (in hours) for each class of job. These are also satisfactory in general. The longer wait times for the larger jobs reflect the very heavy capability usage this month. Note that the 9-hour jobs classes have been withdrawn, as they were scarcely ever used; and that job classes using more than 1024 processors are not supported at present.

Job Class	*Category*	*Maximum Number of CPUs*	*Maximum Job length*	*Average wait time*	*Number of Jobs*
par32_1	parallel	32	1	0.7	1636
par32_3	parallel	32	3	5.2	96
par32_6	parallel	32	6	11.0	753
par64_1	parallel	64	1	1.5	180
par64_3	parallel	64	3	9.6	26
par64_6	parallel	64	6	10.4	135
par128_1	parallel	128	1	3.2	203
par128_3	parallel	128	3	23.0	16
par128_6	parallel	128	6	21.5	27
par128_12	parallel	128	12	31.5	139
par256_1	parallel	256	1	6.0	84
par256_3	parallel	256	3	4.9	12
par256_6	parallel	256	6	11.1	81
par256_12	parallel	256	12	15.3	30
par512_1	parallel	512	1	14.1	45
par512_3	parallel	512	3	32.4	29
par512_6	parallel	512	6	19.4	39
par512_12	parallel	512	12	66.1	34
par1024_1	parallel	1024	1	43.3	14

par1024_3	parallel	1024	3	7.5	1
par1024_6	parallel	1024	6	49.0	4
par1024_12	parallel	1024	12	30.3	31
serial_1	serial	1	1	3.3	875
serial_12	serial	1	3	0.3	251
serial_3	serial	1	6	0.2	26
serial_6	serial	1	9	0.0	25
inter32_1	interactive	32	1	0.0	1702
course32_1	parallel	32	1	0.0	14

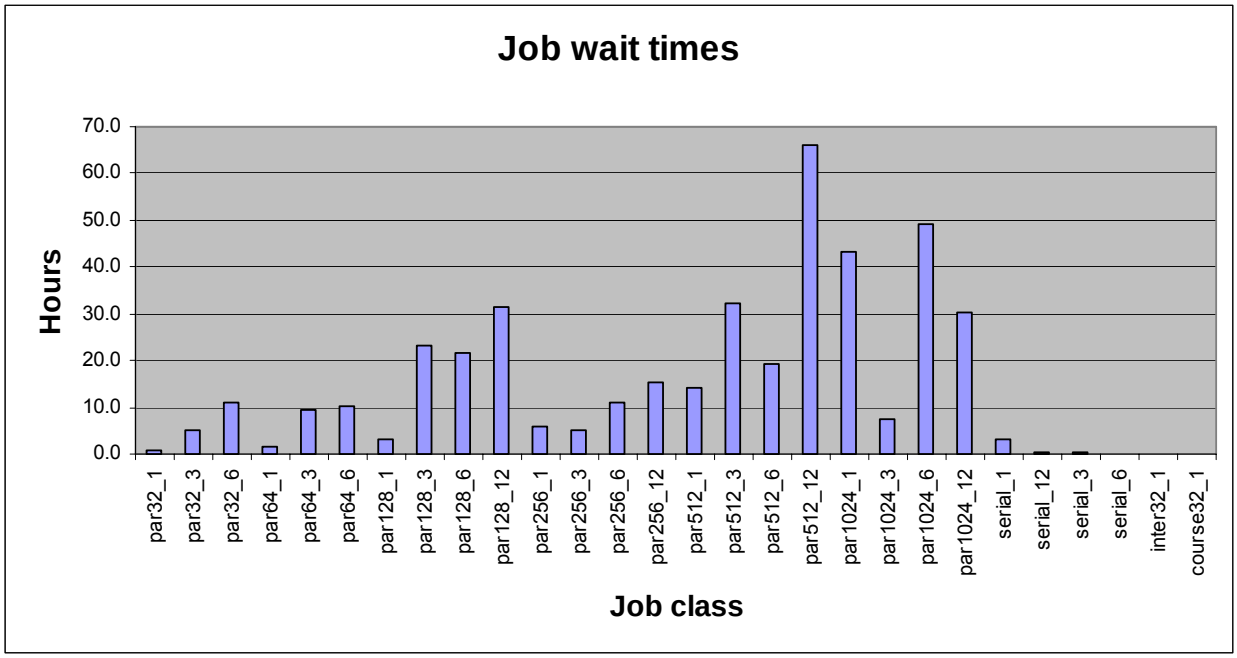

2.5 Disk Occupancy

The quotas shown in these tables reflect the current active quotas on the system, and exclude resource allocated to the project but held in reserve by the PIs.

Home Space

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

Consortium	*Disc Occupancy (Kb)*	*Disc Quota (Kb)*
b02	46,251,488	51,200,000
b03	4,600,160	51,200,000
b04	64	51,200,000
b05	16,802,368	51,200,000
b06	15,353,152	51,200,000
c01	83,896,736	102,400,000
e01	45,075,040	50,006,016
e02	23,633,152	39,760,896
e03	62,321,088	230,412,288
e04	76,314,432	102,400,000
e05	204,767,904	455,680,000
e06	273,628,608	307,200,000
e07	5,638,464	20,480,000
e08	19,760,512	20,480,000
e10	10,148,000	10,240,000
e11	40,582,848	102,400,000
e12	8,976,640	20,480,000
e14	91,621,792	102,400,000
e15	4,150,624	51,200,000
e16	47,392	20,480,000
e17	19,538,752	51,200,000
e18	35,684,256	40,960,000
e19	44,160	40,960,000
e20	53,866,112	61,440,000
e21	97,792	51,200,000
e22	96	10,240,000
e23	96	51,200,000
e24	722,592	51,200,000
e25	5,335,712	51,200,000
e26	18,597,408	20,480,000
e27	3,482,624	20,480,000
e28	115,712	40,960,000
e29	336,448	30,720,000
e30	64	40,960,000
e31	1,418,080	51,200,000

e32	5,435,200	51,200,000
e33	124,608	51,200,000
n01	45,331,232	51,200,000
n02	92,573,088	131,072,000
n03	42,646,912	102,400,000
n04	161,540,480	307,198,976
n05	2,080	10,240,000
p01	37,356,416	40,960,000
x01	38,244,128	51,200,000
x02	8,956,096	20,480,000
z001	250,454,208	276,481,024
z002	42,895,744	49,153,024
z003	256	3,072
z004	73,837,280	102,400,000
z05	4,288,320	30,720,000
z06	49,765,824	51,200,000
z07	20,775,808	30,720,000
z09	15,735,200	51,200,000

Workspace

Consortium	*Disc Occupancy (Kb)*	*Disc Quota (Kb)*
b02	15,104	1,049,600
b03	55,268,736	102,400,000
b04	64	102,400,000
b05	6,372,000	102,400,000
b06	638,272	102,400,000
c01	81,675,168	102,400,000
e01	1,078,806,464	1,177,600,000
e02	8,555,264	10,240,000
e03	10,016	512,000,000
e04	2,638,768,896	3,276,800,000
e05	162,750,816	499,511,296
e06	308,303,552	409,600,000
e07	52,853,760	102,398,976
e08	143,904	1,024,000
e10	291,044,256	307,200,000
e11	2,320,960	102,400,000
e12	743,584	102,400,000
e14	84,925,056	102,400,000
e15	18,676,096	102,400,000
e16	192	61,440,000
e17	834,560	102,400,000
e18	2,101,024	81,920,000
e19	172,914,272	204,800,000
e20	537,040,064	1,024,000,000
e21	1,024	102,400,000

e22	96	20,480,000
e23	96	102,400,000
e24	35,675,008	102,400,000
e25	81,595,744	102,400,000
e26	128	40,960,000
e27	288	40,960,000
e28	19,967,616	81,920,000
e29	941,088	8,192,000
e30	64	81,920,000
e31	73,916,224	102,400,000
e32	495,840	102,400,000
e33	417,600	102,400,000
n01	276,964,384	512,000,000
n02	1,207,410,976	2,036,738,048
n03	32,032	1,026,048
n04	517,182,816	768,000,000
n05	25,564,480	92,160,000
p01	1,022,464	1,024,000
x01	102,397,248	102,400,000
x02	160	20,480,000
z001	331,418,528	409,598,976
z002	297,056	788,480
z003	192	3,072
z004	24,312,256	25,600,000
z05	256	1,024,000
z06	85,763,392	102,400,000
z07	1,728	1,024
z09	22,863,680	102,400,000

2.6 Tape Archive

Consortium	*Usage (Tapes)*	*Quota (Tapes)*	*Files*	*Data (Gb)*
c01	2	2	17	17
e01	38	38	36747	3425
e03	5	5	18797	429
e04	4	14	1260	254
e26	2	2	72	11
n01	114	120	9124	11461
n02	77	170	71615	12166
n04	20	20	71265	2407
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	27	34.6
Technical	44	56.4
In-depth	7	9.0
PMR	0	0.0
TOTAL	78	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	70	89.7
Website	3	3.8
Other/general	5	6.4
TOTAL	78	100.0

Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	59	83.1	75%
Finished within 72 Hours	69	97.2	97%
Finished after 72 Hours	2	2.8	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	25	92.6	97%
Finished after 48 Hours	2	7.4	

Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*	*PMR*
epcc.ed.ac.uk	13	20	4	0
dl.ac.uk	0	2	0	0
Sysadm	13	22	3	0
Other people	1	0	0	0

3.2 Training

Title of Course	*Start Date*	*Length (Days)*	*Place days*	*HPCx User Days*	*HPCx Staff Days*
Optimisation for Power5	08-Dec	1	25	10	1

The Computational Chemistry on HPCx course, which lasted from 30 November to 2 December, was included in the November report.

4 Staffing

4.1 Science Support Staffing

Daresbury Laboratory

Name	*Days*
Ashworth	7.5
Blake	1.6
Bush	14.0
Guest	4.3
Johnstone	8.5
Jones	2.4
Plummer	17.0
Sherwood	2.1
Sunderland	15.0
Thomas	8.5
Pickles	1.6
van Dam	2.1
Total (Days)	84.5
FTEs	4.8

EPCC

Name	*Days*
Simpson	11.7
Booth	11.2
Henty	7.8
Smith	7.1
Bull	5.3
Fisher	4.5
Hein	13.0
Jackson, Adrian	0.4
Pringle	5.5
Reid	12.7
Nazarova	1.2
Trew	4.3
Gray	7.0
D'Mellow	10.5
Hill	11.5

Dobrzelecki	10.8
Helpdesk	1.9
Total (Days)	126.4
FTEs	7.1

Overall Levels

	FTEs
DL	4.8
EPCC	7.1
Total	11.9

4.2 Systems Staffing

Name	*Days*
Andrews	12.8
Blake	0.0
Brown	20.0
Fisher	6.0
Georgeson	12.8
Franks	12.8
Jones	0.0
Shore	15.0
BITD	17.0
Total (days)	96.3
FTEs	5.4

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%	98.3%
Technology MTBF (hours)	200	300	366
Number of AV FTEs	7.5	10	11.9
Number of training days per month	22.5/12	30/12	30/12
Non in-depth queries resolved within 3 days	85%	97%	97.2%
Number of A&M FTEs	3.75	5.75	5.4
A&M serviceability	80%	99.6%	97.7%

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*
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
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
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
n05	2	Non-linear Wave-particle Instabilities in Plasmas	Dr Mervyn Freeman

BBSRC Projects

Code	*Class*	*Title*	*PI*
b02	1	Modelling enzyme catalysis	Dr Adrian J Mulholland
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

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
x02	OHM Ltd	Mr Mark Westwood

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
e01	1	UK Turbulence Consortium	Prof Neil Sandham
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