HPCx Service Report
October 2005

1 Introduction

This report covers the period from 1 October 2005 at 0800 to 1 November 2005 at 0800. Taking into account the end of summer time, this is a service month of 745 hours.

This month we delivered more than 3.9 million AUs to users, more than any month so far. Overall utilisation was more than 86%, which is also a record. Some 40% of this work was at capability scale, and NERC usage was 30% of the total. There was one SEV1 incident, which was caused by an external network problem.

2 Usage

2.1 Availability

Incidents

During this month, there were 19 incidents, only one of which was at SEV 1; most of the rest were minor problems. 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	3
3	15
4	0

The attributions for the SEV 1 incident were as follows:

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

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

Failure	*Site*	*IBM*	*External*	*Reason*
05.142	0%	0%	100%	Loss of external network

Serviceability

There was a total of 40 minutes of scheduled downtime this month.

Attribution	*UDT*	*Serviceability*
IBM	0:00	100.0
Site	0:00	100.0
External	0:05	99.9
Overall	0:05	99.9

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	135006	43	522302	13.8%
e04	9621	5	36494	1.0%
e05	411226	278	1398584	36.9%
e06	49907	417	188986	5.0%
e08	3312	0	12808	0.3%
e10	10916	45	42390	1.1%
e14	16365	99	63672	1.7%
e17	6176	0	23886	0.6%
e18	11033	0	42669	1.1%
e20	50933	55	197195	5.2%
e24	331	10	1320	0.0%
e25	509	7	1997	0.1%
e26	0	2	8	0.0%
e27	331	0	1280	0.0%
e28	4975	0	19240	0.5%
e31	33	23	217	0.0%
EPSRC Total	710673	983	2553049	67.3%

n01	147275	32	569708	15.0%
n02	102189	21	395295	10.4%
n03	43154	75	167188	4.4%
n04	1521	14	5937	0.2%
NERC Total	294139	141	1138128	30.0%

p01	1055	36	4219	0.1%
PPARC Total	1055	36	4219	0.1%

c01	9334	658	38640	1.0%
CCLRC Total	9334	658	38640	1.0%

b02	6606	0	25547	0.7%
BBSRC Total	6606	0	25547	0.7%

x01	2037	0	7880	0.2%
External Total	2037	0	7880	0.2%

z001	5991	40	23326	0.6%
z002	0	0	1	0.0%
z004	327	0	1263	0.0%
z06	120	37	607	0.0%
HPCx Total	6439	77	25196	0.7%

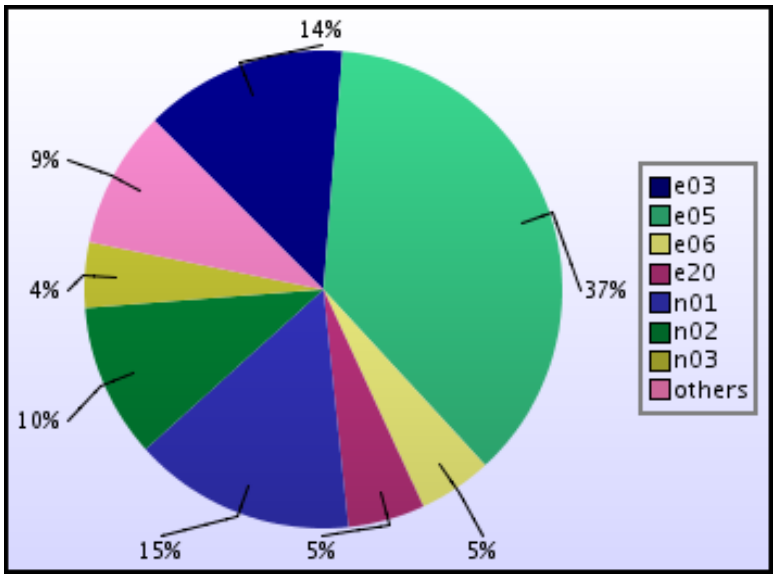

2.3 CPU Usage by Job Type

The figures for *Raw AUs* given here show the number of AUs actually supplied by the system to users' jobs. It uses the conversion rate for the AU which corresponds to the results of the Linpack benchmark running on the new platform; that is, 1 CPU hour = 3.8675 AUs.

Number of Processors	*Raw AUs*	*%age*	*Number of Jobs*
≤32	481681	12.1%	3191
33–64	243905	6.1%	717
65–128	903394	22.7%	871
129–256	769966	19.3%	283
257–512	927810	23.3%	120
513–1024	657860	16.5%	45
>1024	0	0.0%	0

The system is divided into three regions.

Development Region (9 frames, jobs using ≤64 CPUs): a total of 725586 raw AUs were used; that is 87.4% of the total available in this region

Production Region (40 frames, jobs using >64 CPUs): a total of 3259030 raw AUs were used; that is 88.4% of the total available in this region

The remaining frame is 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). Considering the exceptionally high utilisation this month, these figures are good.

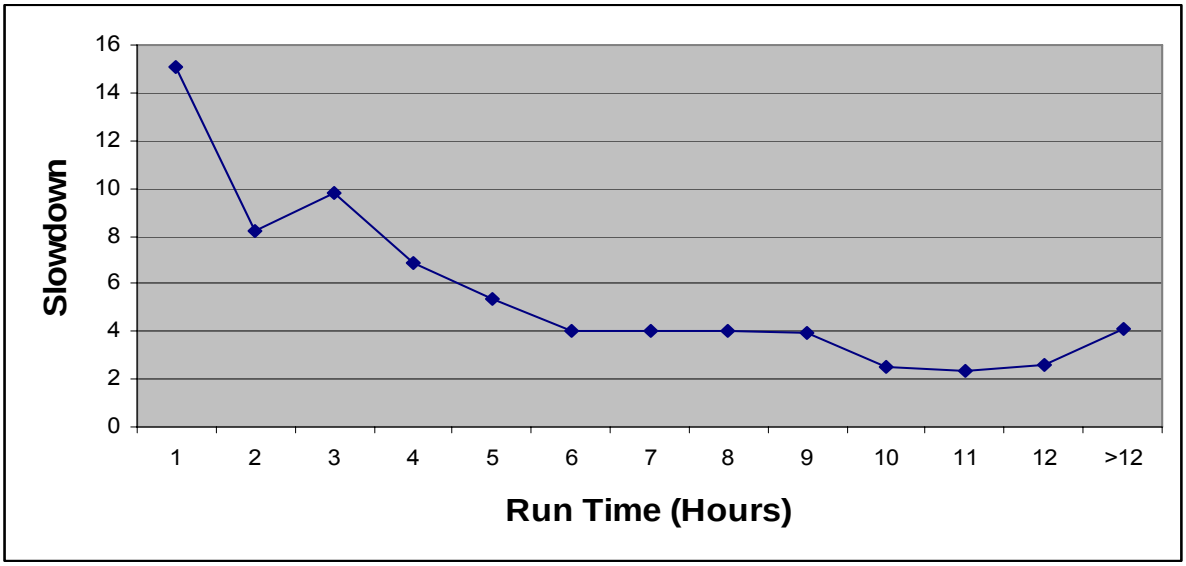

In the graph below, we plot the slowdown figures against the number of processors used and ignoring the development jobs of less than 1 hour. These figures are remarkably good, considering the high utilisation.

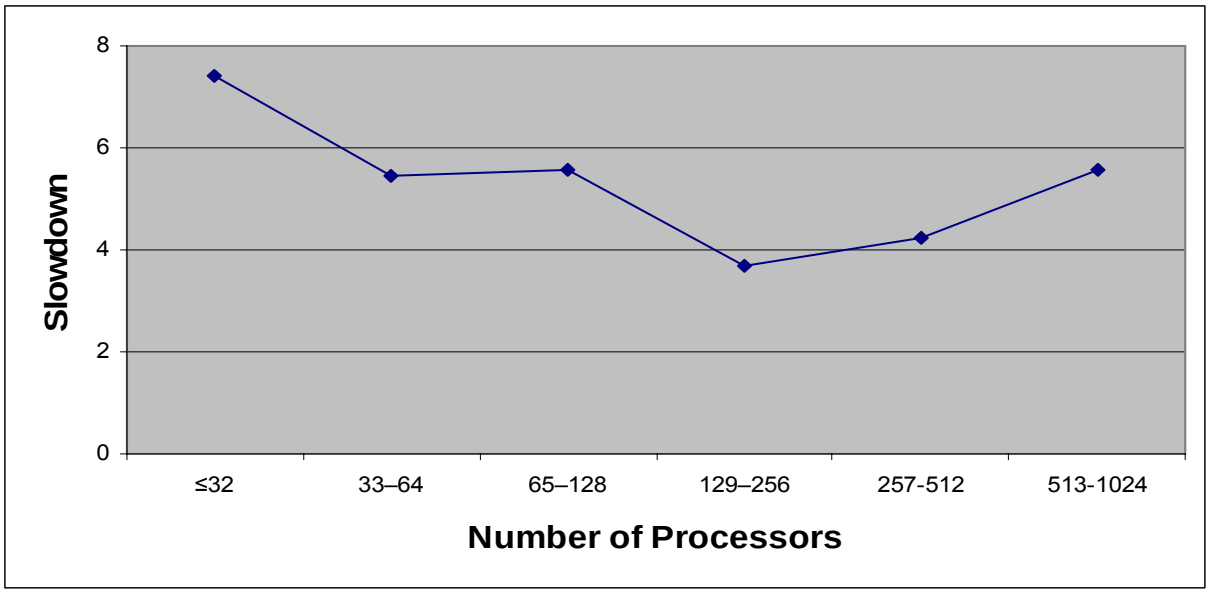

Job wait times

The following table and graph shows the average wait time (in hours) for each class of job. (The *parx512_1*, *parx512_72* and *parx768_72* classes were created for the reserved work relating to the SPICE exercise at SC|05.) These figures are also satisfactory in general, although several spikes reflect the very heavy load on the system this month.

Job Class	*Category*	*Maximum Number of CPUs*	*Maximum Job length*	*Average wait time*	*Number of Jobs*
par32_1	parallel	32	1	1.5	2148
par32_3	parallel	32	3	14.1	147
par32_6	parallel	32	6	19.7	882
par64_1	parallel	64	1	2.9	381
par64_3	parallel	64	3	19.1	39
par64_6	parallel	64	6	16.5	297
par128_1	parallel	128	1	2.8	428
par128_3	parallel	128	3	4.7	79
par128_6	parallel	128	6	6.0	138
par128_9	parallel	128	9	0.0	0
par128_12	parallel	128	12	41.8	199
par256_1	parallel	256	1	4.6	142
par256_3	parallel	256	3	6.6	23
par256_6	parallel	256	6	13.5	27
par256_9	parallel	256	9	14.0	6
par256_12	parallel	256	12	17.5	83
par512_1	parallel	512	1	13.0	39
par512_3	parallel	512	3	31.9	6
par512_6	parallel	512	6	20.5	16

par512_9	parallel	512	9	0.0	0
par512_12	parallel	512	12	11.8	59
par1024_1	parallel	1024	1	48.4	8
par1024_3	parallel	1024	3	46.4	1
par1024_6	parallel	1024	6	0.0	0
par1024_9	parallel	1024	9	0.0	0
par1024_12	parallel	1024	12	16.4	36
par1280_1	parallel	1280	1	0.0	0
par1280_3	parallel	1280	3	0.0	0
par1280_6	parallel	1280	6	0.0	0
par1280_9	parallel	1280	9	0.0	0
par1280_12	parallel	1280	12	0.0	0
serial_1	serial	1	1	0.5	907
serial_12	serial	1	3	0.3	31
serial_3	serial	1	6	0.0	9
serial_6	serial	1	9	0.1	73
serial_9	serial	1	12	0.0	16
inter32_1	interactive	32	1	0.0	3839
course32_1	parallel	32	1	0.0	0
parx512_6	parallel	512	6	0.4	40
parx512_72	parallel	512	72	0.0	1
parx768_72	parallel	768	72	0.0	2

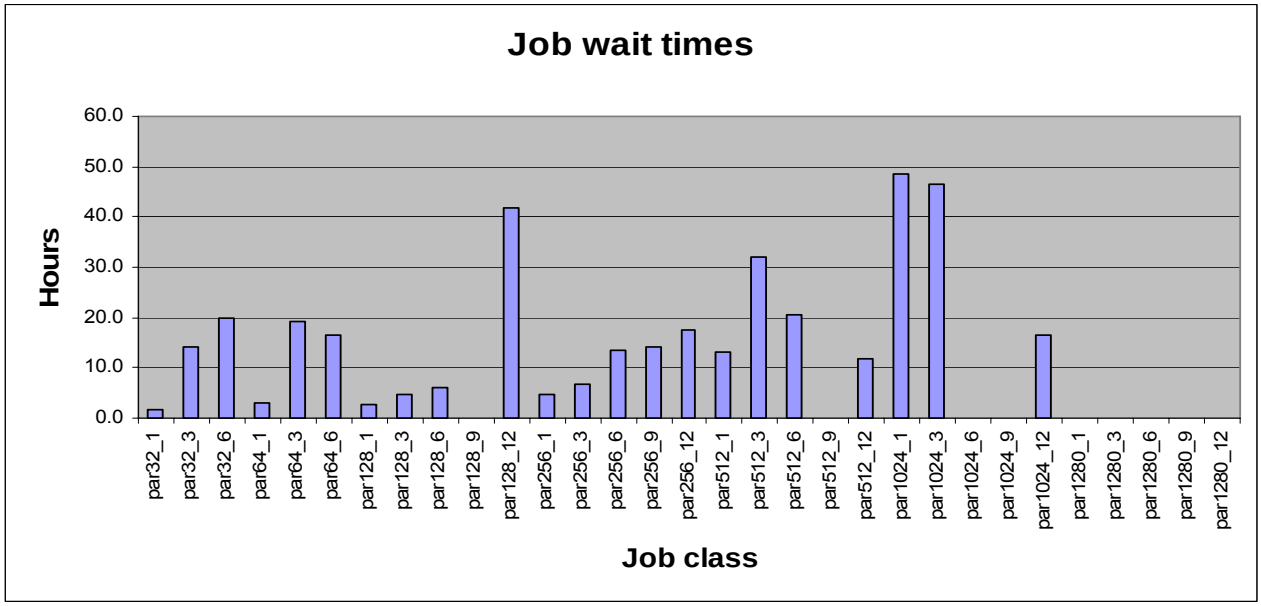

2.5 Disk Occupancy

Home Space

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

Consortium	*Disc Occupancy (Kb)*	*Disc Quota (Kb)*
b02	32,217,152	51,200,000
b03	5,152	51,200,000
b04	64	51,200,000
b05	16,802,368	51,200,000
b06	15,353,152	51,200,000
c01	101,319,680	102,400,000
e01	45,074,848	50,006,016
e02	23,569,056	39,760,896
e03	123,137,312	230,412,288
e04	98,535,904	102,400,000
e05	213,998,368	383,692,800
e06	272,329,792	307,200,000
e07	6,172,768	20,480,000
e08	18,574,400	20,480,000
e10	6,942,432	10,240,000
e11	38,172,608	102,400,000
e12	8,976,640	20,480,000
e14	83,888,608	102,400,000
e15	4,084,864	51,200,000
e16	47,392	20,480,000
e17	19,140,256	51,200,000
e18	40,354,464	40,960,000
e19	43,744	40,960,000
e20	51,861,888	61,440,000
e21	97,792	51,200,000
e22	96	10,240,000
e23	96	51,200,000
e24	708,480	51,200,000
e25	5,324,640	51,200,000
e26	18,970,048	20,480,000
e27	437,408	20,480,000
e28	160	40,960,000
e29	2,155,936	30,720,000
e30	64	40,960,000
e31	1,418,080	51,200,000
e32	96	51,200,000
e33	192	102,400
n01	46,440,192	51,200,000
n02	96,276,800	131,072,000

n03	30,096,864	102,400,000
n04	134,772,352	307,198,976
n05	2,080	10,240,000
p01	32,686,944	40,960,000
x01	32,200,256	51,200,000
x02	8,956,096	20,480,000
z001	209,632,704	235,521,024
z002	43,040,960	49,153,024
z003	256	3,072
z004	73,371,840	102,400,000
z05	4,288,320	30,720,000
z06	50,321,696	51,200,000
z07	13,197,120	30,720,000
z09	8,628,288	51,200,000

Workspace

Consortium	*Disc Occupancy (Kb)*	*Disc Quota (Kb)*
b02	15,104	1,049,600
b03	22,950,688	102,400,000
b04	64	102,400,000
b05	6,372,000	102,400,000
b06	638,272	102,400,000
c01	83,603,008	102,400,000
e01	1,083,966,048	1,177,600,000
e02	8,555,264	10,240,000
e03	10,016	512,000,000
e04	1,810,377,024	2,252,800,000
e05	157,077,984	273,924,096
e06	309,396,672	409,600,000
e07	52,853,760	102,398,976
e08	139,232	1,024,000
e10	291,202,848	307,200,000
e11	192	102,400,000
e12	743,584	102,400,000
e14	67,718,368	102,400,000
e15	18,251,584	102,400,000
e16	192	61,440,000
e17	768	102,400,000
e18	160	81,920,000
e19	172,772,544	204,800,000
e20	470,833,024	1,024,000,000
e21	1,024	102,400,000
e22	96	20,480,000
e23	96	102,400,000
e24	35,650,976	102,400,000

e25	3,483,264	102,400,000
e26	128	40,960,000
e27	224	40,960,000
e28	18,420,416	81,920,000
e29	128	8,192,000
e30	64	81,920,000
e31	14,094,016	102,400,000
e32	96	102,400,000
e33	128	204,800
n01	344,804,224	512,000,000
n02	1,301,614,944	1,504,257,024
n03	293,472	1,026,048
n04	458,741,984	768,000,000
n05	25,564,480	92,160,000
p01	1,022,464	1,024,000
x01	94,794,464	102,400,000
x02	160	20,480,000
z001	363,473,120	409,598,976
z002	297,024	788,480
z003	192	3,072
z004	24,309,472	25,600,000
z05	256	1,024,000
z06	40,905,312	102,400,000
z07	1,696	1,024
z09	26,277,344	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	84	100	4327	8319
n02	50	50	75329	8753
n04	19	20	64038	2102
z001	2	2	6189	50
z002	3	4	1619	11
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	49	50.5
Technical	42	43.3
In-depth	5	5.2
PMR	1	1.0
TOTAL	97	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	84	86.6
Website	4	4.1
Other/general	9	9.3
TOTAL	97	100.0

Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	76	83.5	75%
Finished within 72 Hours	90	98.9	97%
Finished after 72 Hours	1	1.1	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	47	95.9	97%
Finished after 48 Hours	2	4.1	

Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*	*PMR*
epcc.ed.ac.uk	39	15	2	0
dl.ac.uk	0	7	1	1
Sysadm	10	20	2	0
Other people	0	0	0	0

3.2 Training

Title of Course	Start *Date*	*Length (Days)*	*Place days*	*HPCx User Days*	*HPCx Staff Days*
Message-Passing Programming with MPI	25-Oct	3	75	72	1

4 Staffing

4.1 Science Support Staffing

Daresbury Laboratory

Name	*Days*
Ashworth	12.7
Blake	1.9
Bush	21.0
Guest	4.3
Johnstone	10.5
Jones	4.0
Plummer	22.0
Sherwood	2.8
Sunderland	22.0
Thomas	11.0
Pickles	2.1
van Dam	2.8
Total (Days)	116.9
FTEs	6.6

EPCC

Name	*Days*
Simpson	12.0
Booth	18.9
Henty	9.6
Smith	15.2
Bull	2.5
Fisher	8.5
Hein	16.6
Jackson, Adrian	5.1
Pringle	4.0
Reid	13.5
Nazarova	4.6
Trew	5.3
Gray	8.6
D'Mellow	15.9
Hill	18.8
Dobrzelecki	10.8
Helpdesk	4.3
Total (Days)	174.1
FTEs	9.8

Overall Levels

	FTEs
DL	6.6
EPCC	9.8
Total	16.4

4.2 Systems Staffing

Name	*Days*
Andrews	16.1
Blake	0.0
Brown	21.0
Fisher	8.0
Georgeson	16.5
Franks	15.0
Jones	0.0
Shore	8.3
BITD	22.0
Total (days)	106.9
FTEs	6.0

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%	100.0%
Technology MTBF (hours)	200	300	∞
Number of AV FTEs	7.5	10	16.4
Number of training days per month	22.5/12	30/12	23/10
Non in-depth queries resolved within 3 days	85%	97%	98.9%
Number of A&M FTEs	3.75	5.75	6.0
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*
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
e33	1	Engineering Functional Coatings	Prof Roger Smith
e32	1	Rapid Prototyping of Usable Grid Middleware	Prof Peter V Coveney
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