

HPCx Service Report

November 2005

1 Introduction

This report covers the period from 1 November 2005 at 0800 to 1 December 2005 at 0800, a service month of 720 hours.

This month, the service moved to Phase 2A, increasing the number of AUs available to users. Despite this, there was only one SEV1 incident. Overall utilisation remained high, at nearly 84%, and capability use was at its highest level ever: more than 48%.

2 Usage

2.1 Availability

Incidents

During this month, there were 10 incidents, one 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.

<i>Severity</i>	<i>Number</i>
1	1
2	1
3	8
4	0

The attributions for the SEV 1 incident were as follows:

<i>SEV1</i>	<i>Incidents</i>	<i>MTBF</i>
IBM	0.0	∞
Site	1.0	732
External	0.0	∞
<i>Overall</i>	1.0	732

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

<i>Failure</i>	<i>Site</i>	<i>IBM</i>	<i>External</i>	<i>Reason</i>
05.156	100%	0%	0%	16 nodes locked by memory overflow

Serviceability

There was a total of 15.8 hours of scheduled downtime this month.

<i>Attribution</i>	<i>UDT</i>	<i>Serviceability</i>
IBM	0:00	100.0
Site	3:35	99.5
External	0:00	100.0
<i>Overall</i>	3:35	99.5

2.2 CPU Usage by Consortium

The PIs and titles for the various consortia are listed in Appendix B.


<i>Consortium</i>	<i>CPU Hours (Parallel)</i>	<i>CPU Hours (Other)</i>	<i>AUs charged</i>	<i>%age of charged AUs</i>
e03	197160	34	808120	23.1%
e04	15166	50	60295	1.7%
e05	195497	436	718702	20.5%
e06	56726	70	229161	6.5%
e08	6129	0	24164	0.7%
e10	128	0	493	0.0%
e11	763	3	2960	0.1%
e14	14342	25	59191	1.7%
e15	51	0	197	0.0%
e17	18253	12	72830	2.1%
e18	2595	0	10036	0.3%
e19	9720	0	40204	1.1%
e20	51026	11	197962	5.6%
e24	296	11	1271	0.0%
e25	259	100	1650	0.0%
e26	387	0	1863	0.1%
e27	5786	0	22376	0.6%
e28	2	1008	3908	0.1%
e29	748	2	2900	0.1%
e31	1	8	31	0.0%
e32	4893	0	22905	0.7%
z09	2610	0	12567	0.4%
<i>EPSRC Total</i>	<i>582538</i>	<i>1770</i>	<i>2293788</i>	<i>65.4%</i>
n01	105315	26	422913	12.1%
n02	87050	28	342948	9.8%
n03	55246	6	222532	6.3%
n04	7513	83	31258	0.9%
<i>NERC Total</i>	<i>255123</i>	<i>142</i>	<i>1019650</i>	<i>29.1%</i>
p01	1157	22	4832	0.1%
<i>PPARC Total</i>	<i>1157</i>	<i>22</i>	<i>4832</i>	<i>0.1%</i>

c01	18354	410	83156	2.4%
<i>CCLRC Total</i>	18354	410	83156	2.4%

b02	4344	0	18417	0.5%
<i>BBSRC Total</i>	4344	0	18417	0.5%

x01	530	0	2048	0.1%
<i>External Total</i>	530	0	2048	0.1%

z001	17195	61	75807	2.2%
z002	57	0	263	0.0%
z004	0	1	5	0.0%
z06	1531	1	7190	0.2%
<i>HPCx Total</i>	18783	63	83266	2.4%



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.

<i>Number of Processors</i>	<i>Raw AUs</i>	<i>%age</i>	<i>Number of Jobs</i>
≤32	348910	9.7%	2827
33–64	262131	7.3%	785
65–128	651229	18.2%	755
129–256	592044	16.5%	313
257–512	951652	26.5%	148
513–1024	779351	21.7%	58
>1024	1200	0.0%	2

The regions of the Phase 2 and Phase 2A systems are configured differently.

Phase 2:

Development region: 9 frames, 288 processors; ran jobs ≤64 processors

Production region: 40 frames, 1280 processors; ran jobs >64 processors

Interactive parallel region: 1 frame, 32 processors.

Phase 2A:

Capacity region: 26 nodes, 416 processors; runs jobs ≤128 processors

Capability region: 64 nodes, 1024 processors; runs jobs >128 processors

Interactive parallel region: 2 nodes, 32 processors

It can be seen that 128-processor jobs have been moved to the smaller region. The intention is to maintain the high throughput of capability jobs despite the decrease in size of the Capability region. We will be closely monitoring the results of the new configuration.

Combining usage for the old Development and new Capacity regions, we find a total of 875075 AUs, 84.7% of the total available in these regions.

Combining usage for the old Production and the new Capability regions, we find a total of 3324610 AUs, 85.6% of the total available in these regions.

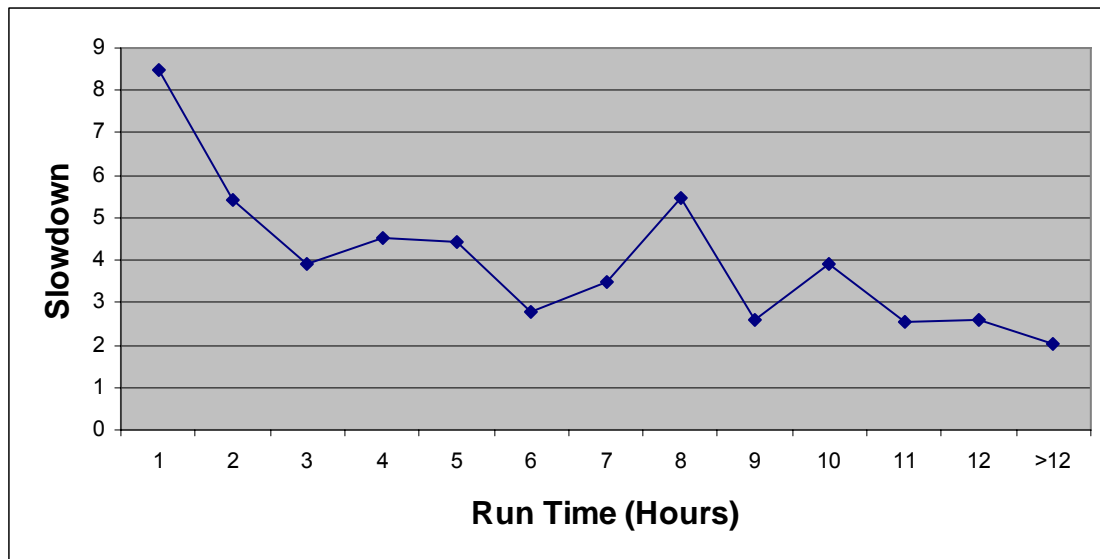
2.4 Slowdown and Job Wait Times

Slowdown


Slowdown is a widely used measure of the relative wait times of different classes of jobs. It is defined as:

$$\text{Slowdown} = (\text{job run time} + \text{job wait time}) / (\text{job run time})$$

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). In general, despite the heavy utilisation of the service, the pattern of slowdowns continues to be 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.




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, although a few classes are waiting for moderately long periods. Note jobs using more than 1024 processors will not run on Phase 2A. The parx- and pary- job classes were special ones for use during the Grid experiments at SC|05.

Job Class	Category	Maximum Number of CPUs	Maximum Job length	Average wait time	Number of Jobs
par32_1	parallel	32	1	1.2	2099
par32_3	parallel	32	3	2.5	178
par32_6	parallel	32	6	9.8	530
par64_1	parallel	64	1	5.1	521
par64_3	parallel	64	3	11.4	49
par64_6	parallel	64	6	11.2	211
par128_1	parallel	128	1	1.8	430
par128_3	parallel	128	3	3.9	82
par128_6	parallel	128	6	6.4	82
par128_9	parallel	128	9	6.2	2
par128_12	parallel	128	12	33.4	137
par256_1	parallel	256	1	2.1	178
par256_3	parallel	256	3	4.0	28
par256_6	parallel	256	6	5.8	38
par256_9	parallel	256	9	8.7	2
par256_12	parallel	256	12	19.4	67
par512_1	parallel	512	1	3.7	67
par512_3	parallel	512	3	3.9	22
par512_6	parallel	512	6	11.7	20

par512_9	parallel	512	9	0.0	0
par512_12	parallel	512	12	9.4	39
par1024_1	parallel	1024	1	18.4	24
par1024_3	parallel	1024	3	11.5	9
par1024_6	parallel	1024	6	39.5	8
par1024_9	parallel	1024	9	0.0	0
par1024_12	parallel	1024	12	8.0	42
par1280_1	parallel	1280	1	35.0	2
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	0.7	1044
serial_12	serial	1	3	0.1	37
serial_3	serial	1	6	0.0	44
serial_6	serial	1	9	0.0	129
serial_9	serial	1	12	0.0	2
inter32_1	interactive	32	1	0.0	3291
course32_1	parallel	32	1	0.0	42
parx256_6	parallel	256	6	1.7	4
parx256_72	parallel	256	72	0.0	8
pary32_6	parallel	32	6	0.0	8
pary32_72	parallel	32	72	13.7	3



2.5 Disk Occupancy

Home Space

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

<i>Consortium</i>	<i>Disc Occupancy (Kb)</i>	<i>Disc Quota (Kb)</i>
b02	39,317,984	51,200,000
b03	5,184	51,200,000
b04	64	51,200,000
b05	16,802,368	51,200,000
b06	15,353,152	51,200,000
c01	67,487,840	102,400,000
e01	45,074,944	50,006,016
e02	23,633,152	39,760,896
e03	53,432,832	230,412,288
e04	76,091,456	102,400,000
e05	203,840,544	419,840,000
e06	272,480,736	307,200,000
e07	5,232,608	20,480,000
e08	19,292,320	20,480,000
e10	10,097,984	10,240,000
e11	38,214,720	102,400,000
e12	8,976,640	20,480,000
e14	99,697,120	102,400,000
e15	4,150,624	51,200,000
e16	47,392	20,480,000
e17	18,511,232	51,200,000
e18	35,684,256	40,960,000
e19	44,160	40,960,000
e20	47,907,968	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	36,437,600	51,200,000
e26	17,944,352	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	320	102,400
n01	45,155,232	51,200,000
n02	92,884,576	131,072,000

n03	30,297,984	102,400,000
n04	143,029,088	307,198,976
n05	2,080	10,240,000
p01	35,863,584	40,960,000
x01	32,352,992	51,200,000
x02	8,956,096	20,480,000
z001	226,942,176	235,521,024
z002	42,968,288	49,153,024
z003	256	3,072
z004	73,384,768	102,400,000
z05	4,288,320	30,720,000
z06	50,734,368	51,200,000
z07	20,766,912	30,720,000
z09	15,735,200	51,200,000

Workspace

<i>Consortium</i>	<i>Disc Occupancy (Kb)</i>	<i>Disc Quota (Kb)</i>
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	79,465,408	102,400,000
e01	1,064,197,824	1,177,600,000
e02	8,555,264	10,240,000
e03	10,016	512,000,000
e04	1,871,219,808	2,252,800,000
e05	153,971,072	294,711,296
e06	347,623,936	409,600,000
e07	52,853,760	102,398,976
e08	143,904	1,024,000
e10	290,676,800	307,200,000
e11	192	102,400,000
e12	743,584	102,400,000
e14	67,718,368	102,400,000
e15	18,676,096	102,400,000
e16	192	61,440,000
e17	815,104	102,400,000
e18	2,101,024	81,920,000
e19	172,914,272	204,800,000
e20	499,309,280	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	42,861,312	102,400,000

e26	128	40,960,000
e27	288	40,960,000
e28	19,954,176	81,920,000
e29	941,088	8,192,000
e30	64	81,920,000
e31	8,038,464	102,400,000
e32	495,840	102,400,000
e33	160	204,800
n01	255,348,704	512,000,000
n02	1,507,218,016	2,016,258,048
n03	550,880	1,026,048
n04	452,889,568	768,000,000
n05	25,564,480	92,160,000
p01	1,022,464	1,024,000
x01	97,642,720	102,400,000
x02	160	20,480,000
z001	391,968,864	409,598,976
z002	297,024	788,480
z003	192	3,072
z004	24,312,256	25,600,000
z05	256	1,024,000
z06	82,202,848	102,400,000
z07	1,728	1,024
z09	22,863,680	102,400,000

2.6 Tape Archive

<i>Consortium</i>	<i>Usage (Tapes)</i>	<i>Quota (Tapes)</i>	<i>Files</i>	<i>Data (Gb)</i>
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	100	100	5929	10022
n02	70	70	67158	11023
n04	20	20	71265	2407
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. The quotas for n01 and n02 have both been increased since the end of November.

3 Support

3.1 Helpdesk

Classifications

<i>Category</i>	<i>Number</i>	<i>% of all</i>
Administrative	69	46.0
Technical	74	49.3
In-depth	7	4.7
PMR	0	0.0
TOTAL	150	100.0

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

<i>Service Area</i>	<i>Number</i>	<i>% of all</i>
Phase 2 platform	124	82.7
Website	10	6.7
Other/general	16	10.7
TOTAL	150	100.0

Performance

<i>All non-indepth queries</i>	<i>Number</i>	<i>%</i>	<i>Target</i>
Finished within 24 Hours	127	88.8	75%
Finished within 72 Hours	143	100.0	97%
Finished after 72 Hours	0	0.0	

<i>Administrative queries</i>	<i>Number</i>	<i>%</i>	<i>Target</i>
Finished within 48 Hours	67	97.1	97%
Finished after 48 Hours	2	2.9	

Experts Handling Queries

<i>Expert</i>	<i>Admin</i>	<i>Technical</i>	<i>In-Depth</i>	<i>PMR</i>
epcc.ed.ac.uk	60	33	7	0
dl.ac.uk	2	14	0	0
Sysadm	6	27	0	0
Other people	1	0	0	0

3.2 Training

<i>Title of Course</i>	<i>Start Date</i>	<i>Length (Days)</i>	<i>Place days</i>	<i>HPCx User Days</i>	<i>HPCx Staff Days</i>
Scientific Visualisation using VTK	23-Nov	1	25	12	0
Parallel IO using MPI-IO	24-Nov	1	25	13	4
Computational Chemistry on HPCx (at Daresbury)	30-Nov	3	60	45	9

4 Staffing

4.1 Science Support Staffing

Daresbury Laboratory

<i>Name</i>	<i>Days</i>
Ashworth	10.4
Blake	2.2
Bush	22.0
Guest	5.5
Johnstone	11.0
Jones	4.3
Plummer	22.0
Sherwood	2.8
Sunderland	18.0
Thomas	6.0
Pickles	2.2
van Dam	2.8
Total (Days)	109.0
FTEs	6.1

EPCC

<i>Name</i>	<i>Days</i>
Simpson	18.5
Booth	22.2
Henty	13.2
Smith	19.5
Bull	7.8
Fisher	6.5
Hein	24.8
Jackson, Adrian	8.2
Pringle	6.6
Reid	21.5
Stratford	2.2
Nazarova	7.5
Trew	5.5
Gray	13.6
D'Mellow	19.7
Hill	17.6
Dobrzelecki	11.3

Helpdesk	3.0
Training	5.5
Total (Days)	234.7
FTEs	13.2

Overall Levels

	<i>FTEs</i>
DL	6.1
EPCC	13.2
Total	19.4

4.2 Systems Staffing

<i>Name</i>	<i>Days</i>
Andrews	16.5
Blake	0.0
Brown	24.0
Fisher	8.0
Georgeson	16.5
Franks	16.5
Jones	0.0
Shore	15.8
BITD	22.0
Total (days)	119.3
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

<i>Metric</i>	<i>TSL</i>	<i>FSL</i>	<i>Monthly Measurement</i>
Technology serviceability	80%	99.2%	100.0%
Technology MTBF (hours)	200	300	∞
Number of AV FTEs	7.5	10	19.4
Number of training days per month	22.5/12	30/12	29/11
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%	99.6%	99.5%

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

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
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 Decanometer 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

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
p01	1	Atomic Physics and Astrophysics	Prof Alan Hibbert

NERC Projects

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
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

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
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

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
c01	1	Daresbury Laboratory Facilities Agreement Consortium	Dr Richard J Blake

Externally-funded Projects

<i>Code</i>	<i>Title</i>	<i>PI</i>
x01	HPC-Europa	Dr J-C Desplat
x02	OHM Ltd	Mr Mark Westwood

HPCx Projects

<i>Code</i>	<i>Title</i>	<i>PI</i>
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

<i>Code</i>	<i>Class</i>	<i>Title</i>	<i>PI</i>
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