HPCx Quarterly Report
October-December 2003

1 Introduction

This report covers the period from 1 October 2003 at 0800 to 1 January 2004 at 0800.

The next section summarises the main points of the service for this quarter. Section 3 gives details of the usage of the service, including failures, serviceability, CPU usage, helpdesk statistics and service quality tokens. A summary table of the key performance metrics is given in the final section. The Appendices define the incident severity levels and list the current HPCx projects.

2 Executive Summary

- During this quarter, utilisation of the service has again increased to an average of 83% on the capability region.
- The job size using the most time remains 128 CPUs but the next largest was 1024 CPUs! More than 35% of the time was used on capability jobs.
- We expect demand to remain high during the first half of 2004 and that the Phase 2 machine will fill up reasonably quickly.
- We have now agreed to increase the maximum number of concurrent research groups to 30, as EPSRC have agreed to fund an additional post for 2004. The maximum may increase again to 35, if additional funding is forthcoming.
- There was a cluster of failures during October but since then the system has been remarkably stable. Overall, the MTBF and serviceability figures exceeded the Full Service Levels.
- The first part of the 'Early Access' Phase 2 system has now successfully passed its implementation tests. We now have a total of 8 Regatta H+ compute frames connected with a pre-General Availability HPS (High Performance Switch, previously referred to as 'Federation').

- We have agreed a phased transition plan to a full Phase 2 system, which will minimise the total amount of down-time. EPSRC have agreed to meet £50K of associated costs.

- We have now developed a proposal for archiving access, which we believe meets users needs; this was announced at the User Group meeting and received a favourable response. We are aiming for beta test access for users during January with full roll-out later in 1Q04.

- The TeraGyroid project was successfully demonstrated at SC2003 and won the HPC Challenge Award for Most Innovative Data-Intensive Application.

- The First HPCx Annual Seminar was held at Daresbury on 10 December, immediately preceding the Machine Evaluation Workshop. The Annual Seminar was well attended and had high quality speakers on a good variety of topics. A successful User Group meeting was held following the seminar.

- The Terascaling team produced a paper looking at the efficiency on HPCx of ten user application codes from a variety of different scientific areas. This paper focused on serial efficiency as the limiting factor on overall achieved efficiency. This suite of codes demonstrated efficiencies of 10% to 20%.

- We have recently developed a Performance Optimisation course to help users maximise the efficiency of their codes. The first run of this was held at Daresbury on the day before the Annual Seminar. This run went very well and attracted around a dozen HPCx users.

3 Usage Statistics

3.1 Availability

3.1.1 Failures

The monthly numbers of incidents and failures (SEV 1 incidents) are shown in the table below:

	October	November	December
Incidents	18	7	13
Failures	11	1	1

October was a very poor month with clusters of incidents related to GPFS, the firewall and JANET; since then, the system and infrastructure have been very reliable.

The following tables give more details on the attribution of the failures:

October

Failure	*Site*	*IBM*	*External*	*Reason*
03.132	100%			Firewall reset
03.135	100%			Firewall reset
03.136			100%	JANET problems at MCC
03.137			100%	JANET problems at MCC
03.138	100%			Firewall reset
03.141		100%		GPFS down after job ran out of swap
03.142			100%	JANET problems at MCC
03.143	25%	75%		GPFS down after job ran out of swap
03.144		100%		Switch plane 1 down/crashed GPFS
03.145/03.146	25%	75%		GPFS down after job ran out of swap
03.149	100%			Router failed

November

Failure	*Site*	*IBM*	*External*	*Reason*
3.152	50%	50%	0%	Virtual shared disk problem on one of the IO nodes

December

Failure	*Site*	*IBM*	*External*	*Reason*
3.164	100%	0%	0%	Access to network failure

3.1.2 Performance Statistics

This section uses the definitions agreed in Schedule 7, ie,

- MTBF = (24 x 30.5)/(number of failures in month)
- Serviceability (%) = 100 x (WCT – SDT – UDT) / (WCT– SDT)

Attribution	*Metric*	*October*	*November*	*Dectember*	*Quarterly*
IBM	Failures	3.5	0.5	1	5
	MTBF	209	1464	732	439
	Serviceability	98.2	99.9	99.7	99.3
Site	Failures	4.5	0.5	0	5
	MTBF	163	1464	∞	439
	Serviceability	98.8	99.9	100.0	99.6
External	Failures	3	0	0	3
	MTBF	244	∞	∞	732
	Serviceability	98.3	100.0	100.0	99.4
Total	Failures	11	1	1	13
	MTBF	67	732	732	169
	Serviceability	95.2	99.8	99.7	98.2

3.2 Capability Utilisation

The monthly utilisation for the 1024-processor capability region is shown in the graph below. This has averaged more than 80% for the second half of the year and the last 2 months were the highest on record.

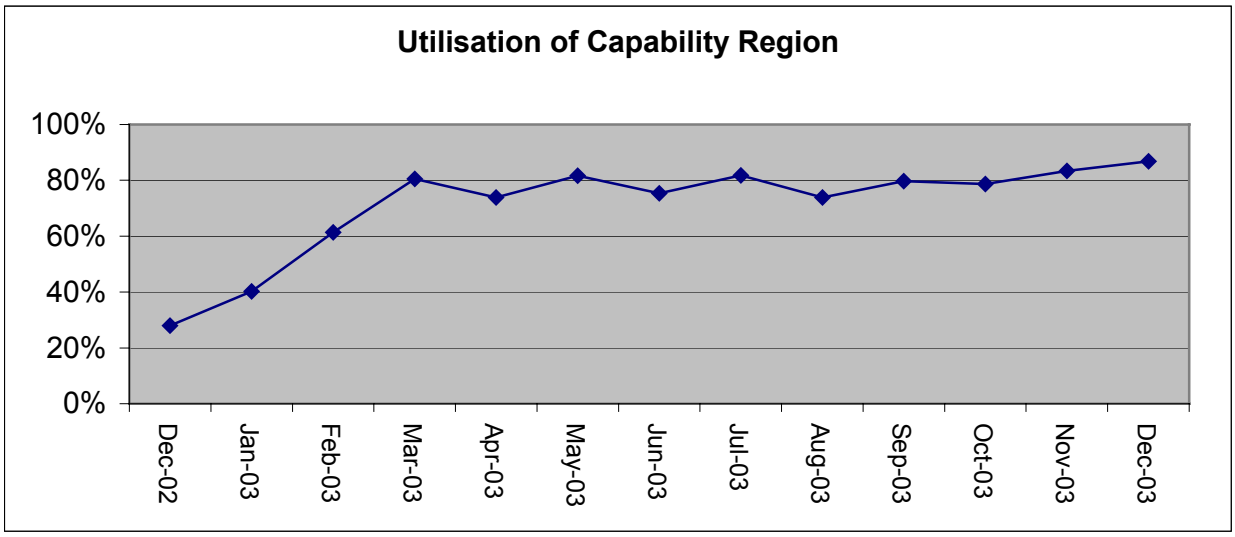

3.3 Capacity Planning

Predicted Utilisation

The following graph shows the utilisation since the start of the project and the projected utilisation until November 2004. The scale on the *y*-axis is AUs per hour, where the peak that HPCx Phase 1 could currently deliver is around 3240 AUs per hour (the red line in the graph). However, the practical maximum is probably around 75% of this, i.e., 2430 AUs per hour, which is shown as a yellow line.

The graph assumes that each project will use its remaining allocation pro rata with its usage profile from the SAF, which is often simply that on the original application form.

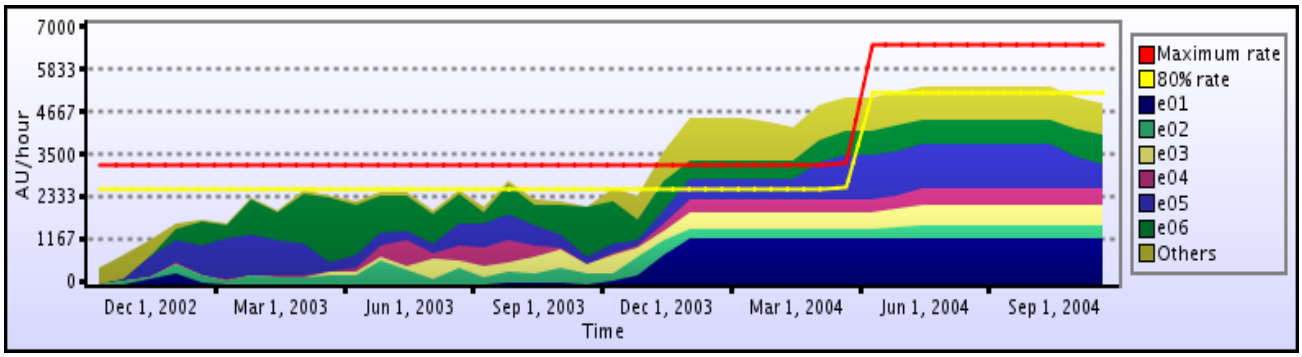

The graph suggests that demand until the start of Phase 2 appears to exceed significantly the practical maximum. It appears likely that, even without new grants coming on to the system, demand on the Phase 1 system will exceed supply. Even when Phase 2 comes online, it appears that it is likely to fill up quickly.

Numbers of Research Consortia

There are currently 26 research consortia using the HPCx system. We have agreed to increase the maximum number of consortia to 30 for 2004, as EPSRC have underwritten an additional post.

3.4 CPU Usage by Job Size

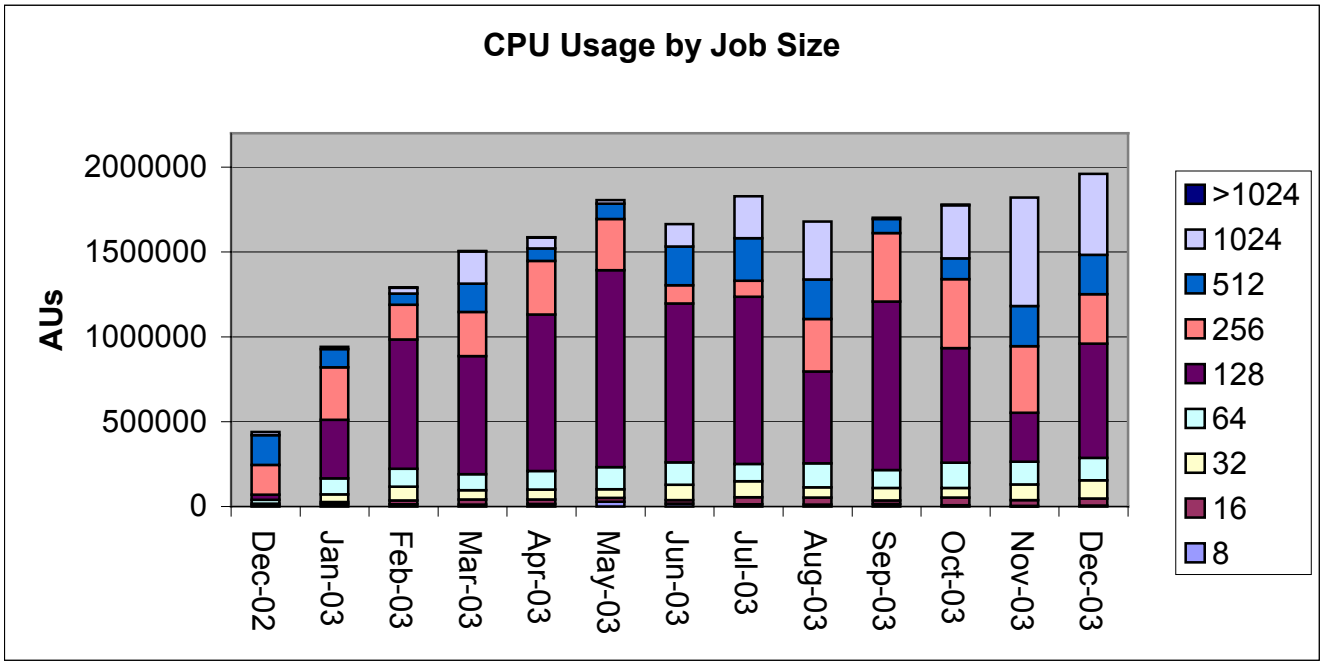

The above graph shows that the job size using the most time is still 128 CPUs. However, during this quarter, utilisation by jobs of at least 256 CPUs has increased to 56% of the total and capability utilisation has increased to more than 36%.

3.5 AU Usage by Consortium

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

Consortium	*October*	*November*	*December*	*Quarterly*	*%age*
e01	69084	49477	138349	256909	4.6%
e02	216737	239540	179955	636232	11.4%
e03	363842	211903	329973	905718	16.3%
e04	209027	17132	43229	269388	4.8%
e05	362957	187852	143708	694517	12.5%
e06	463604	838120	678580	1980304	35.6%
e07		124	1309	1433	0.0%
e08		1121	155	1276	0.0%
e09				13239	0.2%
e10	1413	8	11818	13239	0.2%
e11	2		214335	214337	3.9%

e12		43	43	85	0.0%
e13	5887	229429	54052	289369	5.2%
EPSRC Total	1692553	1774749	1795506	5262807	94.5%

n01	15141	2306	87119	104566	1.9%
n02	3993	9853	25200	39046	0.7%
n03	53956	30384	20840	105180	1.9%
n04	4023	41	4963	9027	0.2%
n05	1012			1012	0.0%
NERC Total	78125	42584	138122	258830	4.7%

p01	843	772	140	1756	0.0%
PPARC Total	843	772	140	1756	0.0%

c01	1462	1532	9626	12620	0.2%
CCLRC Total	1462	1532	9626	12620	0.2%

z001	8002	4246	11374	23622	0.4%
z002	31	40	1	72	0.0%
z004	269			269	0.0%
z06	158	1460	698	2316	0.0%
HPCx Total	8460	5746	12073	26279	0.4%

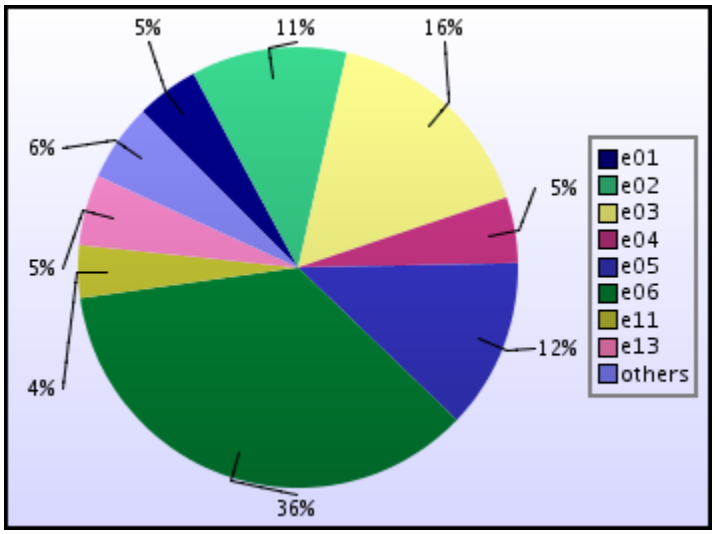

3.5.1 Discounts

There are now a number of user codes that have qualified for capability discounts. The following table shows the discounts that were awarded during the last quarter.

Consortium	*AUs Used*	*AUs Charged*	*Discount*
e01	266177	256909	9268
e10	18748	13239	5509

3.6 Helpdesk

3.6.1 Classifications

Category	*Number*	*% of all*
Administrative	85	36.0
Technical	131	55.5
In-depth	16	6.8
PMR	4	1.7
TOTAL	236	100.0

Service Area	*Number*	*% of all*
Phase 1 platform	200	84.7
Website	21	8.9
Other/general	15	6.4
TOTAL	236	100.0

3.6.2 Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	185	85.6	75%
Finished within 72 Hours	216	100.0	97%
Finished after 72 Hours	0	0.0	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	83	97.6	97%
Finished after 48 Hours	2	2.4	

3.6.3 Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*	*PMR*
epcc.ed.ac.uk	61	52	11	0
dl.ac.uk	2	20	0	1
Sysadm	22	58	4	3
Other people	0	1	1	0

3.7 Service Quality Tokens

Date	*Person*	*Value*	*Comment*	*Status*
Oct 27, 2003 6:33:07 PM	*Dr Changman Moon*	* * * * *		

4 Support

Details of the current status of science support can be found in the *HPCx Annual Report: 2003.*

4.1 Staffing

AV	*October*	*November*	*December*
DL	4.7	4.9	3.9
EPCC	6.5	8.1	6.8
Total	11.2	13.0	10.8

Systems	7.9	6.5	5.4

5 Summary of Performance Metrics

Metric	*TSL*	*FSL*	*October*	*November*	*December*
Technology serviceability	80%	99.2%	98.2%	99.9%	100.0%
Technology MTBF (hours)	200	300	209	1460	∞
Number of AV FTEs	7.5	10	11.2	13.0	10.8
Number of training days per month	30/12	40/12	40/10	49/11	50/12
Non in-depth queries resolved within 3 days	85%	97%	100.0%	100.0%	100.0%
Number of A&M FTEs	3.75	5.75	7.9	6.5	5.4
A&M serviceability	80%	100%	98.8%	99.9%	99.7%

Colour	*Meaning*
	Exceeds FSL
	Between TSL and FSL
	Below TSL

Note: The number of training days is reported as a running total since the start of the year.

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: Current Projects

EPSRC Projects

Code	*Class*	*Title*	*PI*
e01	1	UK Turbulence Consortium	Prof Neil Sandham
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
e09	2	Molecular Properties and their Geometry	Prof Peter Taylor
e10	1	Reality Grid	Prof Peter Coveney
e11	1	Bond making and breaking at surfaces	Prof Sir David A King
e12	1	Parallel programs for the simulation of complex fluids	Dr Mark R Wilson
e13	1	TeraGyroid project	Dr Richard J Blake
e14	1	Blade and Cavity Noise	Prof Neil Sandham

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

PPARC Projects

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

BBSRC Projects

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

CCLRC Projects

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

Early User Projects

Code	*Title*	*PI*
y001	Materials	Dr Patrick Briddon
y002	DNS of Turbulent Flow	Prof Neil Sandham
y003	Multi-photon and Electron Collision Processes	Prof Ken Taylor
y004	Materials	Prof Jonathon Tennyson
y005	UKAEA	Dr Tim Hender
y006	UK Car-Parrinello Consortium	Prof David Price
y007	Climate Modelling	Dr Lois Steenman-Clark

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