HECToR Annual Report 2009

01 January – 31 December 2009

Issue: 1.0

Contents

1 Introduction ... 4
2 Executive Summary ... 4
3 Quantitative Metrics ... 6
3.1 Reliability ... 6
3.1.1 Performance Statistics ... 6
3.2 HECToR Utilisation ... 7
3.2.1 Overall Utilisation ... 7
3.2.2 XT Utilisation by AU ... 7
3.2.3 XT Utilisation by Job Size ... 8
3.2.4 XT Number of jobs per queue ... 8
3.2.5 Usage by Consortium ... 9
3.2.6 HECToR Usage by Application Area ... 13
3.3 Performance Metrics ... 14
4 Helpdesk ... 15
4.1.1 Helpdesk Targets ... 15
4.1.2 Queries by Service Metric ... 15
4.1.3 Queries by Category ... 15
4.1.4 Queries by Handler Category ... 16
4.1.5 Helpdesk General ... 16
5 Systems Hardware ... 17
5.1 HECToR Technology Changes ... 17
5.1.1 CLE2.1 Upgrade ... 17
5.1.2 Phase 2A Quad-Core Upgrade ... 17
5.1.3 HECToR Archive Solution ... 18
5.1.4 HECToR Phase 2B ... 19
5.1.5 PBS v10 ... 19
5.2 Severity-1 Incidents ... 20
5.2.1 Technology Failures ... 20
5.2.2 Service Provision Failures ... 21
5.2.3 External Failures ... 21
5.3 Single Node Failures ... 22
6 Cray Centre of Excellence ... 24
6.1 Executive Summary ... 24
6.2 Applications Support ... 24
6.2.1 Cray UK Focus ... 24
6.2.2 CoE Steering Committee ... 26
6.2.3 Other Applications ... 27
6.2.4 Phase 2 Development ... 29
6.3 Workshops and Training Events ... 29
6.4 Seminars ... 30
7 The HECToR Computational Science and Engineering (CSE) Support Service ... 31
7.1 Overview of the CSE Service ... 31
7.2 Highlights of 2009 ... 31
7.3 The CSE Helpdesk ... 32
7.4 Training ... 35
7.5 Outreach ... 36
7.6 Other Core CSE Activities ... 36
7.7 Distributed CSE ... 36
7.8 Completed DCSE Projects ... 37

7.8.1 OCEANS 2025 (NEMO) 37
7.8.2 Improving the Parallelisation of CASINO 37
7.8.3 Enhanced scalability of the domain decomposition within CP2K 38
7.8.4 Performance enhancements for the GloMAP aerosol model Part 1 38
7.9 Ongoing DCSE Projects 39
7.9.1 Efficient Parallel Algorithms for ChemShell 39
7.9.2 Cloud and Aerosol Research on Massively-parallel Architectures 39
7.9.3 Parallel Algorithms for the Materials Modelling code CRYSTAL 40
7.9.4 Better Parallelism for the electron-atom scattering codes PRMAT 40
7.9.5 Porting and Optimisation of Code_Saturne® on HECToR (XT4/X2) 40
7.9.6 WRF code optimisation for Meso-scale Process Studies (WOMPS) 41
7.9.7 Multigrid improvements to the Citcom code 41
7.9.8 Hybrid time-dependent density functional theory in CASTEP Part 1 42
7.9.9 Performance enhancements for the GloMAP aerosol model Part 2 42
7.9.10 Parallelisation of CABARET 43
7.9.11 Scaling Turbulence Applications to thousands of cores 43
7.9.12 Direct Numerical Simulations (DNS) of turbulent fluid flows 43
7.9.13 Implementation of established algorithms to extend HELIUM 44
7.9.14 Massive Remote Batch Visualizer – MRBV 44
7.9.15 Improvements in I/O for DL_POLY_3 44
7.9.16 Dynamic load balancing and rigid body dynamics for DL_POLY_3 45
7.9.17 HPC Driven Next-Generation Modelling of the World's Oceans (ICOM) 45
7.9.18 Improving Load Balancing and Parallel Partitioning in ICOM 45
7.9.19 HECToR Distributed CSE Support for OpenFOAM 45
7.9.20 Improving the performance of GWW 46
7.9.21 Fast Fourier Transformations for Gyrokinetic Plasma Simulations 46
7.9.22 SPRINTing with HECToR 46
7.9.23 Improving scalability of CP2K on multi-core systems 46
7.9.24 Scalability Optimization for Large Scale In-Silico Simulations of Cardiac Bioelectric Activity 47
7.10 Plans for 2010 48
Appendix A: Terminology 49
Appendix B: Projects on HECToR 50

1 Introduction

This report covers the period from 1 Jan 2009 at 08:00 to 1 Jan 2010 at 08:00.

The next section of this report contains an executive summary for the year.

Section 3 summarises service availability, performance and utilisation statistics for the year. Section 4 covers the Helpdesk statistics. Systems support is covered in Section 5, with the work of the Cray Centre of Excellence described in Section 6 and the Computational Science and Engineering (CSE) Support provided by NAG covered in Section 7.

The Appendices define some of the terminology and incident severity levels and list the current HECToR projects together with their overall utilisation profile to date.

This report and the additional SAFE report are available to view online at http://www.hector.ac.uk/about-us/reports/annual/2009.php

2 Executive Summary

2009 saw a number of key changes to the HECToR Service, with upgrades to both the operating system and the arrival of quad-core.

The highlights of the service over the year included:

- Helpdesk statistics for the year were excellent. A total of 3717 queries were handled with targets exceeded in almost all areas.
- 2009 has been a busy year for the CSE service. The CSE helpdesk has continued to be busy and to receive positive feedback from users. Training courses have been a big success and are now being run at a variety of locations.
- The Distributed CSE programme is proving very popular and the projects completed so far will lead to savings of several million pounds through improved code efficiency.
- The Cray Centre of Excellence (CoE) for HECToR continued to deliver targeted high-end application support to a number of user groups. The Centre was heavily involved, along with the wider Cray applications team, in preparing users for the transition to quad-core and for the CLE 2.1 upgrade.
- The CoE has also participated extensively in the Cray Centre of Excellence Steering Committee which aims to target the Centre's activities to important UK HPC applications and users. An additional activity has been the Cray UK Focus which has consisted of a series of UK-wide workshops, the aim of which is to place a strong emphasis on scaling of UK applications to very high core counts.
- Average utilisation for the year on the Cray XT4 was 59%. Utilisation peaked in July during the 'Open Access' phase of the Phase 2A upgrade. A Capability Incentive scheme for discounting large jobs was implemented in 4Q09. Initial results from this appear to be positive. A Low Priority Access mechanism was also enabled on HECToR on 16th December. While preliminary results are encouraging, further data is required in order to ascertain the success of this initiative.

- The EPSRC Resource Allocation Panel (RAP) call was announced in 4Q09. The initial meeting of the panel will be in February 2010. The BBSRC initiative 'Widening the BBSRC HPC User Base' also promises new and novel research on HECToR.
- In 2009 there were 37 technology-attributed service failures compared to 46 in 2008. The volume of software failures has decreased considerably from 32 in 2008 to 6 in 2009; the upgrade to CLE2.1 appears to have had a positive impact. In addition to the technology failures, there were 8 site and external failures. The overall MTBF was 195 in 2009, an improvement on the 172 in 2008.
- The X2 was very reliable in 2009. Charging on the Cray X2 was suspended in February to improve utilisation. The overall utilisation for the year was 44%.
- The HECToR operating system was upgraded to CLE2.1 on 6th May. The upgrade which had been rescheduled from March was completed as planned.
- The HECToR Phase 2A upgrade to quad-core was completed in June. Work commenced on the two-phase upgrade on 8th June and by June 29th HECToR was available to users with 5664 quad-core nodes. The acceptance tests were completed in July.
- The HECToR website had a major transition in 2009. The style, layout and content of the site were updated. Prospective users of the service are now presented with a simplified view of the service. Case studies covering a diverse range of research were kindly provided by HECToR users to illustrate examples of HECToR use.
- The HECToR Archive solution was made available for early access users in November. No major issues were encountered and the archive is now available for all users. To date over 31TB of data has been archived.
- The HPCx service closes on 31st Jan 2010. In relation to this a number of projects have recently moved resources from HPCx to HECToR. Additional HECToR introductory training has also been scheduled in order to support any new users.
- The Contract Change Notice for HECToR Phase 2B was signed in August. Detailed planning has commenced for the upgrade to a Cray XT6 in 1H10, followed by an upgrade to the Gemini interconnect in late 2010. An early user service is expected to be online remotely in the U.S. in March 2010.

3 Quantitative Metrics

3.1 Reliability

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

	1Q09	*2Q09*	*3Q09*	*4Q09*	*2009*
Incidents	94	87	106	60	347
Failures	16	8	13	8	45

The incidents above are primarily related to single node failures. Details on both the service failures and single node fails in 2009 can be found in Section 5.

3.1.1 Performance Statistics

- MTBF = (732)/(number of failures in a month)

  Quarterly MTBF = (3x732)/(number of failures in a quarter)

  Annual MTBF = (12*732)/(number of failures in a year)

Attribution	*Metric*	*1Q09*	*2Q09*	*3Q09*	*4Q09*	*2009*
Technology	Failures	14	8	9	6	37
	MTBF	157	275	244	366	237
Service Provision	Failures	2	0	2.5	1	5.5
	MTBF	1098	∞	878	2196	1597
External	Failures	0	0	1.5	1	2.5
	MTBF	∞	∞	1464	2196	3514
Overall	Failures	16	8	13	8	45
	MTBF	137	275	169	275	195

3.2 HECToR Utilisation

3.2.1 Overall Utilisation

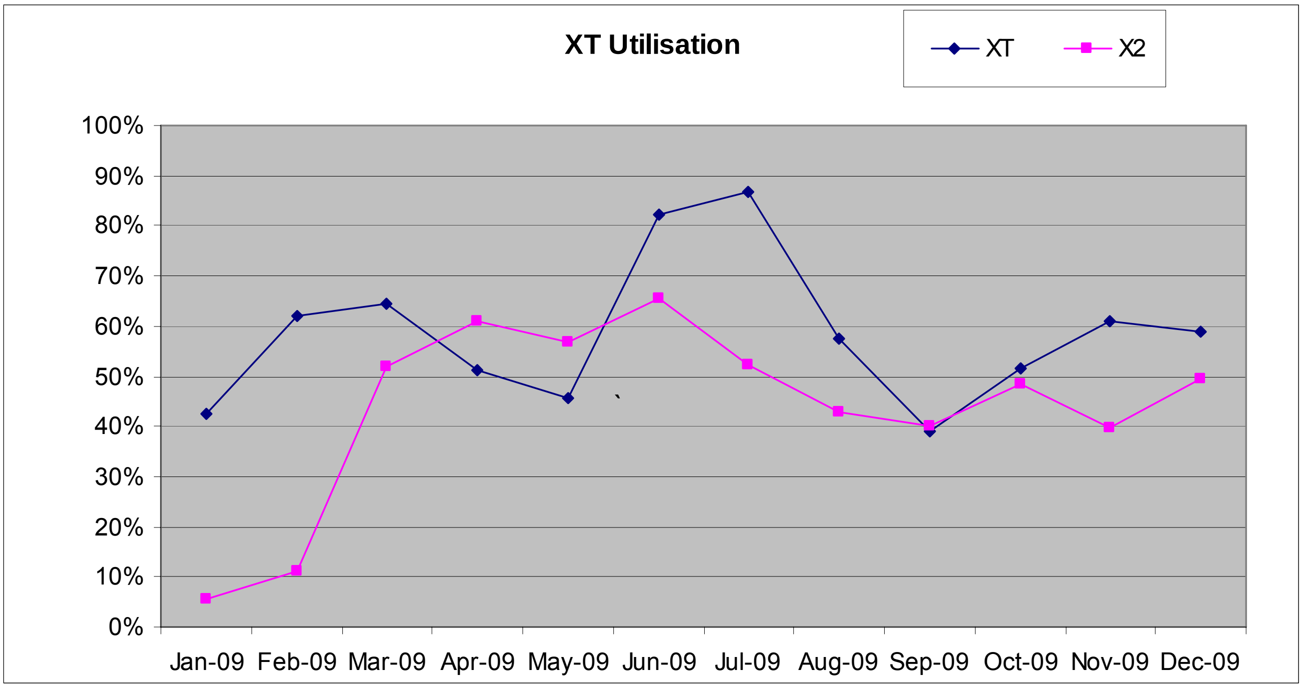

3.2.2 XT Utilisation by AU

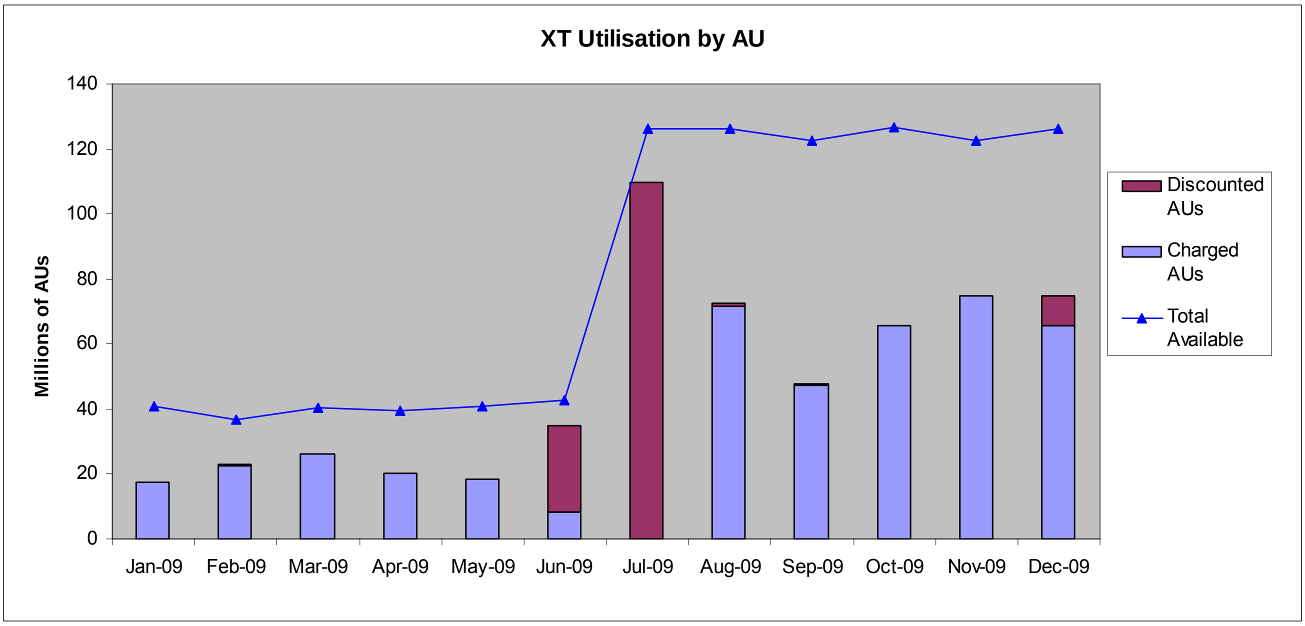

The above graph shows the impact on the overall size of the service due to the upgrade to quad-core processors in Phase 2A. The effect of the open access period in June/July and the introduction of the low priority queue in December can also be seen.

The percentage utilisation has not changed dramatically post quad-core, however the quantity of research conducted by the users has. As an example, in February and November utilisation was approximately 60%. At an AU level, three times as many were used in November.

3.2.3 XT Utilisation by Job Size

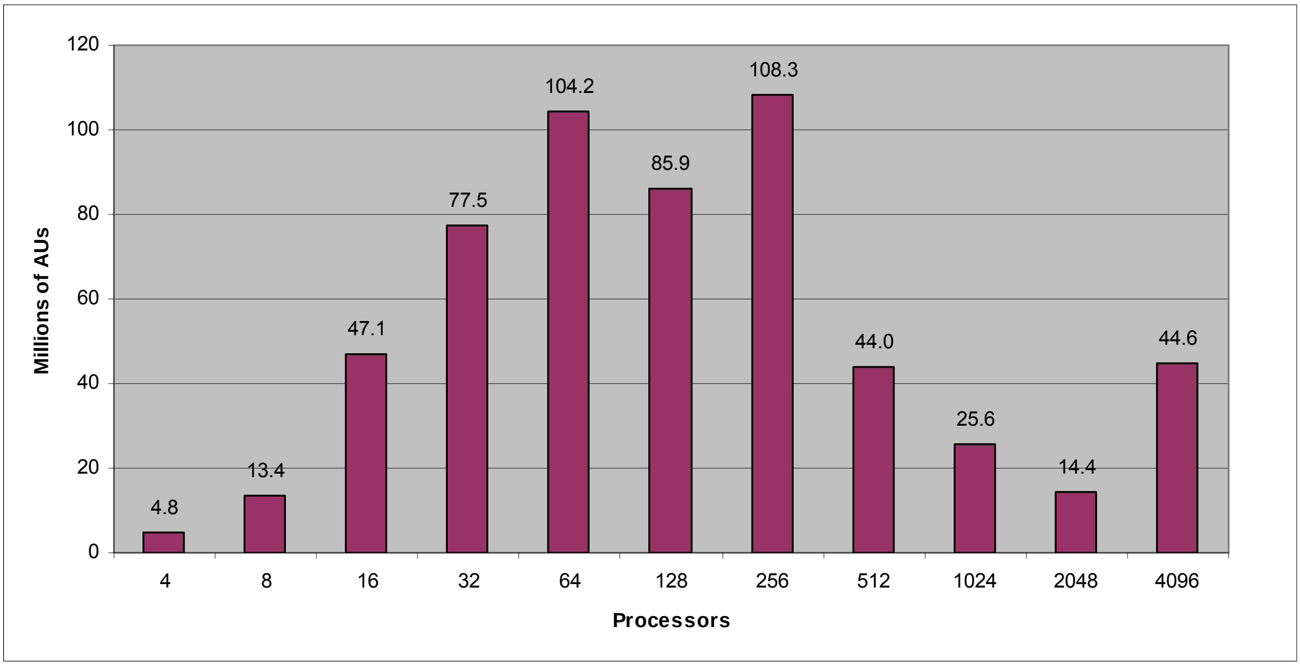

3.2.4 XT Number of jobs per queue

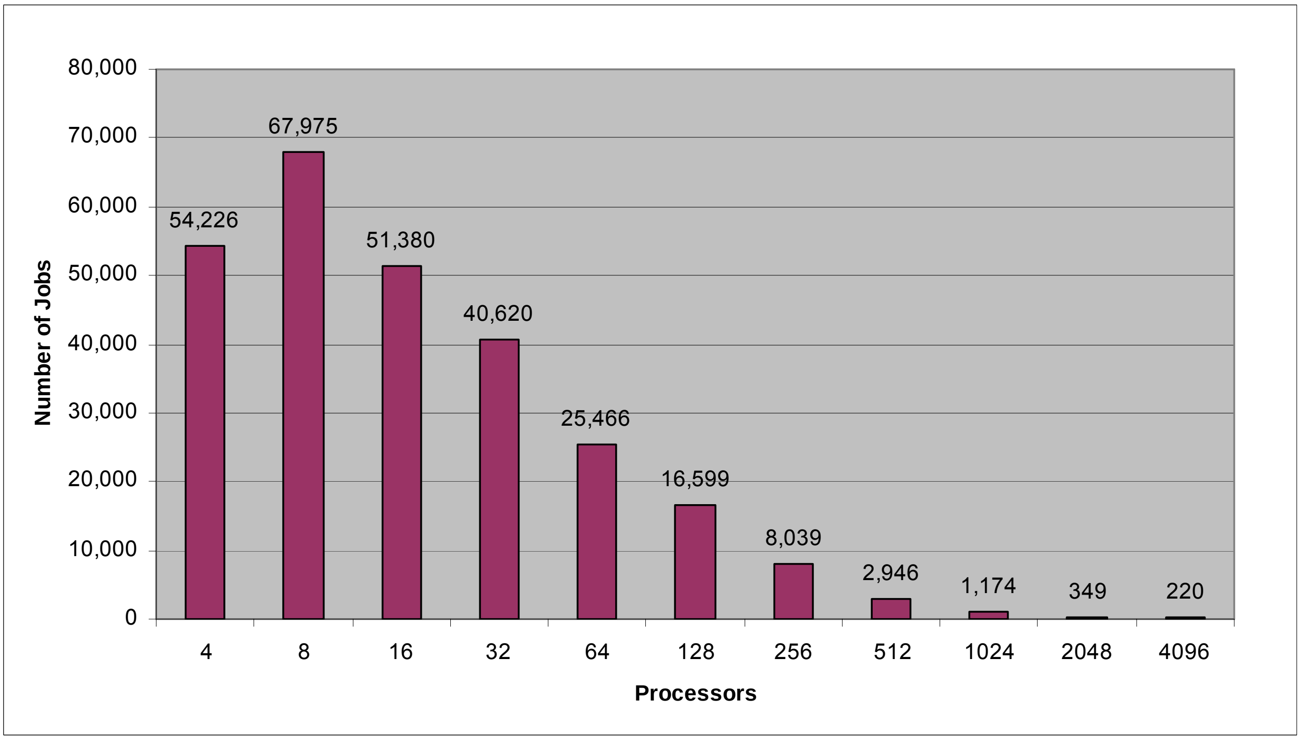

3.2.5 Usage by Consortium

3.2.5.1 XT Usage by Consortium

A total of 991,033,704 AUs were available during this period.

Project	AUs	Number of Jobs	%age of Use	Utilisation
y01	11,110	115	0.00%	0.00%
y02	88	13	0.00%	0.00%
y03	2,337	3873	0.00%	0.00%
y04	33	9	0.00%	0.00%
y05	11,673,742	2410	2.00%	1.18%
y06	126	9707	0.00%	0.00%
y07	482	132	0.00%	0.00%
y11	0	1	0.00%	0.00%
z01	1,739,359	5241	0.30%	0.18%
z02	57,177	1280	0.01%	0.01%
z03	3,173,731	14028	0.54%	0.32%
z06	1,549	10	0.00%	0.00%
Internal Total	**16,659,735**	**36819**	**2.85%**	**1.68%**
c01	8,680,822	8042	1.48%	0.88%
e01	912,118	1917	0.16%	0.09%
e05	42,979,737	16525	7.35%	4.34%
e10	202,073	346	0.03%	0.02%
e24	46,529,014	5708	7.95%	4.69%
e34	19,230	69	0.00%	0.00%
e35	1,975,793	397	0.34%	0.20%
e42	16,343,717	6316	2.79%	1.65%
e59	509,625	64	0.09%	0.05%
e63	13,972,128	3292	2.39%	1.41%
e68	41,077,135	8111	7.02%	4.14%
e69	54,113	61	0.01%	0.01%
e70	733,127	481	0.13%	0.07%
e71	225,425	394	0.04%	0.02%
e72	433,850	195	0.07%	0.04%
e74	0	2	0.00%	0.00%
e75	3,847,638	451	0.66%	0.39%
e76	25,258,510	351	4.32%	2.55%
e77	19,684	109	0.00%	0.00%
e79	17	18	0.00%	0.00%
e81	24,044	195	0.00%	0.00%
e82	10,625	123	0.00%	0.00%
e84	20,486	336	0.00%	0.00%
e85	4,951,375	276	0.85%	0.50%
e88	42,014	11	0.01%	0.00%
e89	77,951,488	26192	13.32%	7.87%

Project	AUs	Number of Jobs	%age of Use	Utilisation
e90	86,511	746	0.01%	0.01%
e92	23,577	31	0.00%	0.00%
e93	64,992	225	0.01%	0.01%
e94	106,473	151	0.02%	0.01%
e96	0	7	0.00%	0.00%
e97	110,364	147	0.02%	0.01%
e98	1,168	85	0.00%	0.00%
e99	50,952	57	0.01%	0.01%
e100	1,560	13	0.00%	0.00%
e101	438,514	488	0.07%	0.04%
e102	3,378,758	1684	0.58%	0.34%
e104	2,654	11	0.00%	0.00%
e105	252,770	153	0.04%	0.03%
e107	140,014	283	0.02%	0.01%
e108	27,421	75	0.00%	0.00%
e109	100,828	38	0.02%	0.01%
e110	7,604,487	1822	1.30%	0.77%
e111	100,470	172	0.02%	0.01%
e112	107,594	100	0.02%	0.01%
e113	15,502	182	0.00%	0.00%
e116	19,708	179	0.00%	0.00%
e117	247,664	72	0.04%	0.02%
e119	0	7	0.00%	0.00%
e121	345,279	316	0.06%	0.03%
e122	21,778,119	1767	3.72%	2.20%
e124	4,226,229	1141	0.72%	0.43%
e125	8,235,501	93	1.41%	0.83%
e126	1,232,387	221	0.21%	0.12%
e127	144	12	0.00%	0.00%
e132	10,230	154	0.00%	0.00%
e134	473	4	0.00%	0.00%
e135	19,677	181	0.00%	0.00%
e136	42,787	37	0.01%	0.00%
e137	190,349	78	0.03%	0.02%
u02	0	5	0.00%	0.00%
u03	11,400	152	0.00%	0.00%
u10	451,751	1210	0.08%	0.05%
EPSRC Total	**336,200,097**	**92081**	**57.47%**	**33.92%**
n01	10,303,575	11780	1.76%	1.04%
n02	68,912,319	106315	11.78%	6.95%
n03	100,752,275	25182	17.22%	10.17%
n04	27,371,966	8744	4.68%	2.76%
NERC Total	**207,340,136**	**152021**	**35.44%**	**20.92%**
b01	300,795	85	0.05%	0.03%
b08	910,000	6738	0.16%	0.09%
b09	2,323	5	0.00%	0.00%
b10	2,765	484	0.00%	0.00%

Project	AUs	Number of Jobs	%age of Use	Utilisation
b11	0	25	0.00%	0.00%
BBSRC Total	**1,215,884**	**7337**	**0.21%**	**0.12%**
p01	622	34	0.00%	0.00%
STFC Total	**622**	**34**	**0.00%**	**0.00%**
T01	962,925	1111	0.16%	0.10%
x01	2,345,714	3062	0.40%	0.24%
x02	2	3	0.00%	0.00%
x04	943,342	302	0.16%	0.10%
x05	9,870	311	0.00%	0.00%
External Total	**4,261,853**	**4789**	**0.73%**	**0.43%**
d03	2,151,259	1762	0.37%	0.22%
d04	811,258	4125	0.14%	0.08%
d05	291,515	359	0.05%	0.03%
d07	1,707,983	216	0.29%	0.17%
d08	25,647	31	0.00%	0.00%
d09	1,052,368	407	0.18%	0.11%
d10	12,240,051	21	2.09%	1.24%
d11	2,374	16	0.00%	0.00%
d14	1,046,244	512	0.18%	0.11%
d15	61	38	0.00%	0.00%
DirectorsTime Total	**19,328,761**	**7487**	**3.30%**	**1.95%**
Total	**585,007,088**	**300568**	**100.00%**	**59.03%**

3.2.5.2 X2 Usage by Consortium

A total of 18,837,504 AUs were available during this period.

Project	AUs	Number of Jobs	%age of use	Utilisation
y02	5,558	240	0.07%	0.03%
y03	204	141	0.00%	0.00%
y05	1,143	55	0.01%	0.01%
z01	1,909	110	0.02%	0.01%
z02	6	1	0.00%	0.00%
z03	80,235	277	0.97%	0.43%
Internal Total	**89,054**	**824**	**1.08%**	**0.47%**
e01	2,994,260	1150	36.16%	15.90%
e05	2,522,298	986	30.46%	13.39%
e10	33,213	41	0.40%	0.18%
e24	997,268	289	12.04%	5.29%
e42	139	5	0.00%	0.00%
e69	10,195	18	0.12%	0.05%
e75	285,314	892	3.45%	1.51%
e85	66,732	59	0.81%	0.35%
e89	1,223,280	178	14.77%	6.49%
e99	1	2	0.00%	0.00%
e102	13	20	0.00%	0.00%
e110	14	20	0.00%	0.00%
e123	1,465	31	0.02%	0.01%
EPSRC Total	**8,134,194**	**3691**	**98.24%**	**43.18%**
n01	0	2	0.00%	0.00%
n02	32,607	337	0.39%	0.17%
n03	1	1	0.00%	0.00%
NERC Total	**32,608**	**340**	**0.39%**	**0.17%**
x01	1	1	0.00%	0.00%
External Total	**1**	**1**	**0.00%**	**0.00%**
d04	23,858	481	0.29%	0.13%
DirectorsTime Total	**23,858**	**481**	**0.29%**	**0.13%**
Total	**8,279,714**	**5337**	**100.00%**	**43.95%**

3.2.6 HECToR Usage by Application Area

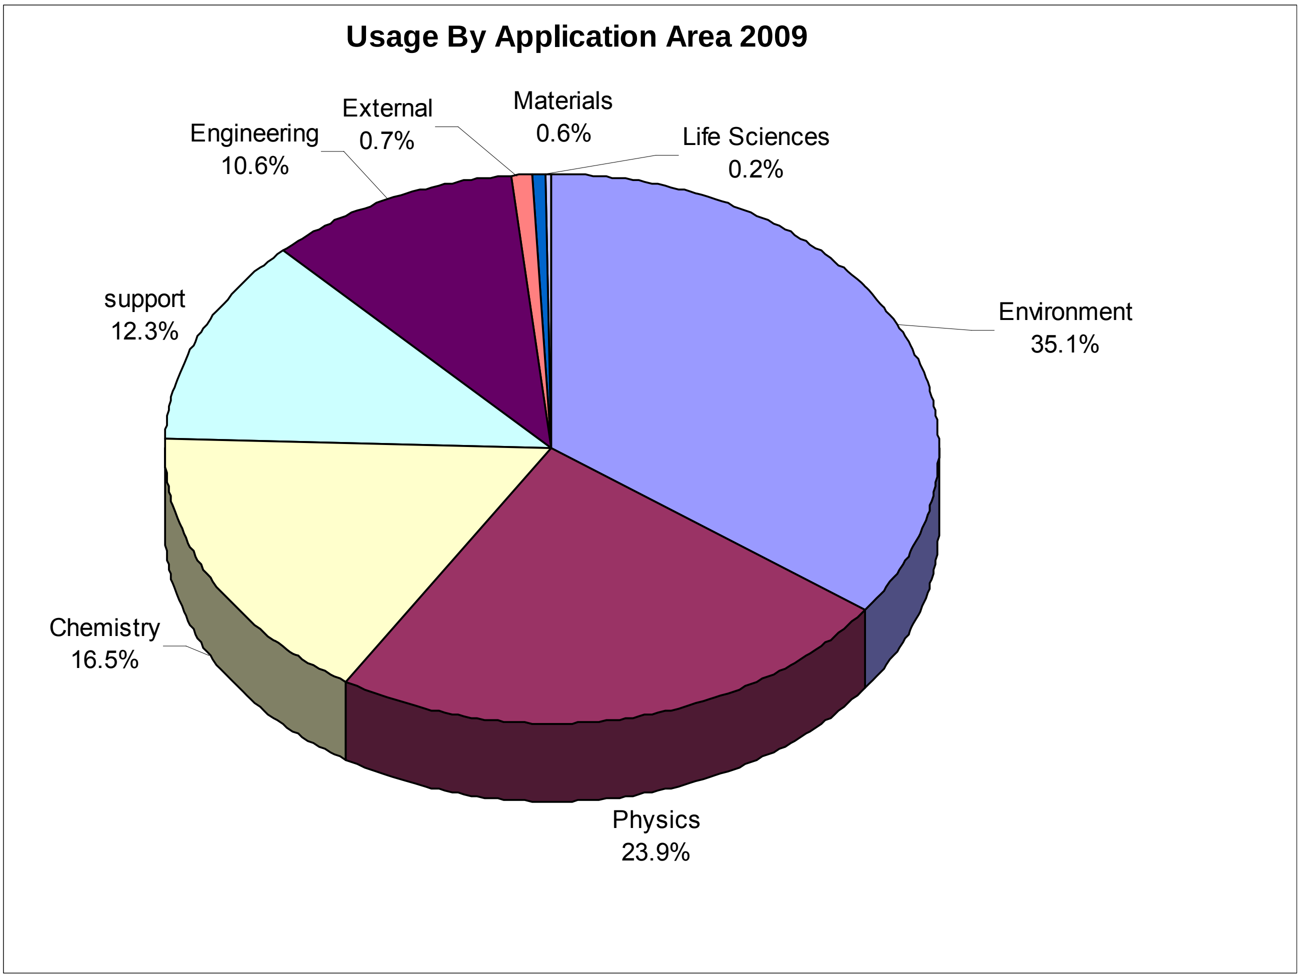

3.3 Performance Metrics

Metric	TSL	FSL	January	February	March	April	May	June	July	August	September	October	November	December	Annual Average
Technology reliability (%)	85.0%	98.5%	98.4%	97.8%	98.4%	99.1%	98.8%	98.2%	99.4%	96.6%	99.3%	100.0%	98.1%	98.9%	98.6%
Technology MTBF (hours)	100	126.4	146.4	122.0	146.4	244.0	366.0	244.0	183.0	183.0	146.4	∞	183.0	183.0	237.4
Technology Throughput, hours/year	7000	8367	8472	8308	8433	8534	8495	7346	8124	8307	8198	8604	8099	8012	8366.3
Capability jobs completion rate	70%	90%	N/A(*)	100.0%	100.0%	100.0%	100.0%	97.7%	96.3%	100.0%	100.0%	100.0%	100.0%	100.0%	99.4%
Non in-depth queries resolved within 1 day (%)	85%	97%	100.0%	100.0%	99.7%	100.0%	99.0%	100.0%	100.0%	100.0%	98.3%	98.8%	100.0%	100.0%	99.6%
Number of SP FTEs	7.3	8.0	8.2	8.2	8.1	8.6	8.9	10.1	9.0	8.4	9.3	8.9	9.2	8.4	8.8
SP Serviceability (%)	80.0%	99.0%	99.4%	98.5%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	99.5%	100.0%	100.0%	93.9%	99.3%

Colour coding:

Exceeds FSL	(green)
Between TSL and FSL	(amber)
Below TSL	(red)

** There were no capability jobs in January.*

4 Helpdesk

A total of 3718 queries with a specified service metric were completed in this period.

4.1.1 Helpdesk Targets

Metric	Pass	Total	Fraction	Target
All queries finished in 1 day	3056	3072	99.5%	97.0%
Admin queries finished in 1 day	2770	2781	99.6%	97.0%
Queries assigned in 30 min	3687	3711	99.4%	97.0%
Technical assessments in 10 days	93	104	89.4%	97.0%

4.1.2 Queries by Service Metric

Service Metric	Queries	Percentage
Automatic	1820	49.0%
Admin	961	25.8%
In-depth	542	14.6%
Technical	291	7.8%
Technical assessment class-1A	67	1.8%
Technical assessment class- 2A	30	0.8%
Technical assessment class-1B	5	0.1%
Technical assessment class- 2B	2	0.1%

4.1.3 Queries by Category

Query Category	Queries	Percentage
New User	515	13.9%
Set group quotas	337	9.1%
Set user quotas	332	8.9%
3rd Party Software	267	7.2%
New Password	253	6.8%
Access to HECToR	235	6.3%
None	232	6.2%
User behaviour	213	5.7%
Node Failure	182	4.9%
Add to group	167	4.5%
Disk, tapes, resources	146	3.9%
Compilers and system software	131	3.5%
New Group	109	2.9%
Batch system and queues	107	2.9%
Other	83	2.2%
User programs	82	2.2%
Login, passwords and ssh	80	2.2%
SAFE	55	1.5%
Join Project	50	1.3%
Static website	27	0.7%
Courses	27	0.7%
Remove account	22	0.6%
Update account	16	0.4%
Performance and scaling	12	0.3%

Query Category	Queries	Percentage
Create certificate	8	0.2%
Grid	7	0.2%
Delete from group	7	0.2%
Archive	6	0.2%
Network	4	0.1%
Delete from project	3	0.1%
Porting	2	0.1%
Delete Certificate	1	0.0%

4.1.4 Queries by Handler Category

Handlers	Admin	Automatic	In-depth	Technical	Technical assessment	Total	%age
USL	883		143	154		1180	31.74%
OSG	71	1820	39	115		2045	55.00%
CSE	1		282	1	104	388	10.44%
Cray Systems	6		78	20		104	2.80%
Other				1		1	0.03%

4.1.5 Helpdesk General

The HECToR website was updated in 1Q09 to facilitate easier navigation to the FAQ and DCSE areas. The internal style guide has now been updated for all the support teams and the site checked to ensure all html validates against the W3C standard.

In 2Q09 there was a further change to the website in parallel with the upgrade to quad-core nodes. The CSE and User Support teams worked well together to ensure there was a seamless transition of the web content. There was a new look to the website along with changes to the HECToR technical specifications, user guide and frequently asked questions.

Finally, in 3Q09 the website was given a new ‘public facing’ image. Prospective users of the service are now presented with a simplified view of the service. These changes were made in line with ideas which had been discussed by the Strategic Management Board.
Case studies were kindly provided by HECToR users to illustrate examples of HECToR use. The sample case studies published so far cover a diverse range of subject areas including;

- Virtual palaeontology – Dinosaur gait reconstruction
- Cancer research
- Reducing Aircraft Noise Pollution
- Improving the model of Cardiac Arrhythmia
- Gene Analysis using Parallel ’R’
- Quantum Chemistry using CP2K
- Breakthroughs in fluid mixing using simulated fractal-generated flows

5 Systems Hardware

5.1 HECToR Technology Changes

2009 was a very busy period for the service, with a number of key technology changes taking place.

5.1.1 CLE2.1 Upgrade

The HECToR operating system was upgraded to CLE2.1 in May. Lessons learned from a previous attempt to complete the upgrade in March were implemented, resulting in a faster, more efficient upgrade.

Initially, it looked as if the upgrade to CLE 2.1 had worked perfectly, with the system passing its regression tests and being returned to service some hours early. However, within 24 hours problems were identified with a small number of user codes. A key issue was that the Unified Model would not recompile. This was not a problem with CLE 2.1 per se, rather it was caused by a change in the korn shell (ksh). Although the Unified Model did form part of the regression tests run by the CSE team, this was a preconfigured test case used for benchmarking which did not fully explore the "features" of the Makefile and hence did not uncover the problems. Some user-owned versions of codes would not run (including CASTEP and WRF). It transpired that these had not been recompiled despite repeated mailings to users on this.

A post-mortem was conducted by the service. Key recommendations included:

- More comprehensive regression testing
- More direct communications between the CSE team and key groups
- Access to the TDS for key consortia prior to major upgrades

These recommendations were adopted for the quad-core upgrade.

The operating system upgrade to CLE2.1 appears to have reduced the number of software related failures.

5.1.2 Phase 2A Quad-Core Upgrade

The HECToR Phase 2A upgrade to quad-core was completed in June. Work commenced on the two-phase upgrade on 8^{th} June and, by 29^{th} June, the 5664 quad-core nodes of HECToR were available to users. A test Linpack run demonstrated a maximum performance of 174 TFlops. This would have placed HECToR at theoretical sixteenth place in the June edition of the Top500.

As well as the upgrade to quad-core processors, all 22,656 memory DIMMs were replaced, and all 60 cabinet blowers had rework completed on them. In total 28,320 components were swapped.

The hardware upgrade and testing went to plan with no major issues encountered. Preventative measures were taken by the team to ensure that the upgrade did not impact the live service. This planning was worthwhile as there was no impact on the main service as a result of the upgrade work which was taking place in parallel. Cray staff were also dedicated to supporting either the upgrade or the main service, and not both at the same time.

The upgrade to quad-core processors took place in two parts. During the first part, half of HECToR's dual-core processors were available to users (2832 processors). During the second part of the upgrade the same number of quad-core processors were made available for users (11,328 cores). This happened on 19th June; at this stage user accounting was disabled. Within 8 hours of the service moving to quad-core processors, there were over 100 quad-core jobs running. This was very positive and showed that the users had adapted quickly and changed scripts as required. The service was soon running at full capacity. On 29th June the remaining 2832 processors were brought online for users as quad-core.

The HECToR Phase 2A acceptance tests were completed on 7th July. The 10-day availability trial was completed as planned on 18th July.

Following on from the upgrade there was an increase in the number of single-node failures in July as a result of faulty DIMMS (See Section 5.3). This was to be expected and failure rates quickly dropped back to lower levels. In summary, the upgrade was very successful providing an enhanced service with minimal disruption.

5.1.3 HECToR Archive Solution

The ability to retain very large volumes of historic data is critical to the research being conducted by certain groups on HECToR (e.g. ocean and climate modelling). Work commenced on the implementation of the HECToR Archive solution in 3Q09.

The solution is an archive mechanism provided by CRAY based on Symantec's Enterprise NetBackup. The archive has an initial capacity of 1300*800GB Tapes (1.02PB). At this time the typical daily archive volume is 1TB.

Two methods of archive are supported

- Daily scheduled archive from special directories within the users work space.
- User initiated on-demand archive via a simple command line interface

The archive completed acceptance tests on November 5th. A number of groups were identified for initial user trials as shown in the table below.

Project	PI
n01	Thomas Anderson/ Beverly de Cuevas
n02	Lois Steenman-Clark
n04	Richard Proctor
e01	Gary Coleman
e85	N. Swaminathan

Early access to the archive commenced on November 18th. Some minor bugs were identified by the early access users which have subsequently been resolved. The archive was made available for all users in late December.

5.1.4 HECToR Phase 2B

The Contract Change Notice for HECToR Phase 2B was signed in August.

A 20-cabinet Cray XT6 system will be installed, with an early-user service available in 1/2Q10 2010. The system will comprise 44,544 cores (464 blades x 8 chips x 12 cores), delivering an estimated peak performance of 338 TFlops. Approximately 30 cabinets of Cray XT4 will also be retained. The Cray XT4 will formally provide the main service at this stage.

This will be followed by an upgrade to the Gemini network (by the end of 2010). The Cray XT6/Gemini system will then provide the main service.

5.1.5 PBS v10

Job reservations are important to users as they can be used to support demonstrations, training courses and urgent computing. Reservations are available in PBS v10 which was installed on HECToR on 5th August. In practice we have discovered a number of problems with enabling reservations. Work is ongoing to resolve these problems. The third-party provider Altair has been engaged and has been working onsite with the service. The root cause of the problem has now been identified and we plan to have a workaround available in 1Q10.

5.2 Severity-1 Incidents

5.2.1 Technology Failures

Cray supplied technology has been responsible for 37 severity-1 incidents in 2009.

This is a breakdown of the failure categories:

- Fifteen compute or service blade related failures
- Eight High Speed Network failures
- Two cabinet blower failures
- Three storage related failures
- Six system software related incidents
- Two instances of a late return to service following scheduled maintenance
- One procedural interrupt

The quad core upgrade in June provided an opportunity to install a new version of firmware which runs on the compute and service blades. The new firmware resolved an issue which had accounted for at least a third of all blade-related failures in 2009 and this has had a positive impact. The remaining blade failures were due to power and control-related issues.

High Speed Network failures are typically caused by single link failures. The Gemini interconnect fabric that will be installed on the Cray XT6 later in 2010 will include features that prevent these kind of failures from impacting the system. A cabinet blower rework program was also implemented in June during the quad-core upgrade to improve the reliability and longevity of these devices.

Two of the storage-related failures resulted in new versions of firmware being installed in the respective storage subsystems and the third problem was a wiring related issue that has also been resolved.

Six software related incidents occurred in 2009. Two of these issues are still under investigation but the rest have been resolved either by configuration changes or following the introduction of version 2.1 of the Cray Linux Environment (CLE) Operating system in May.

The two occasions where scheduled maintenance sessions overran accounted for a combined 42 minutes of delay to the service. A procedural issue which resulted in a SEV-1 incident has been resolved by modifying site procedures.

So far, reliability on quad-core seems slightly better than the long term average on dual-core.

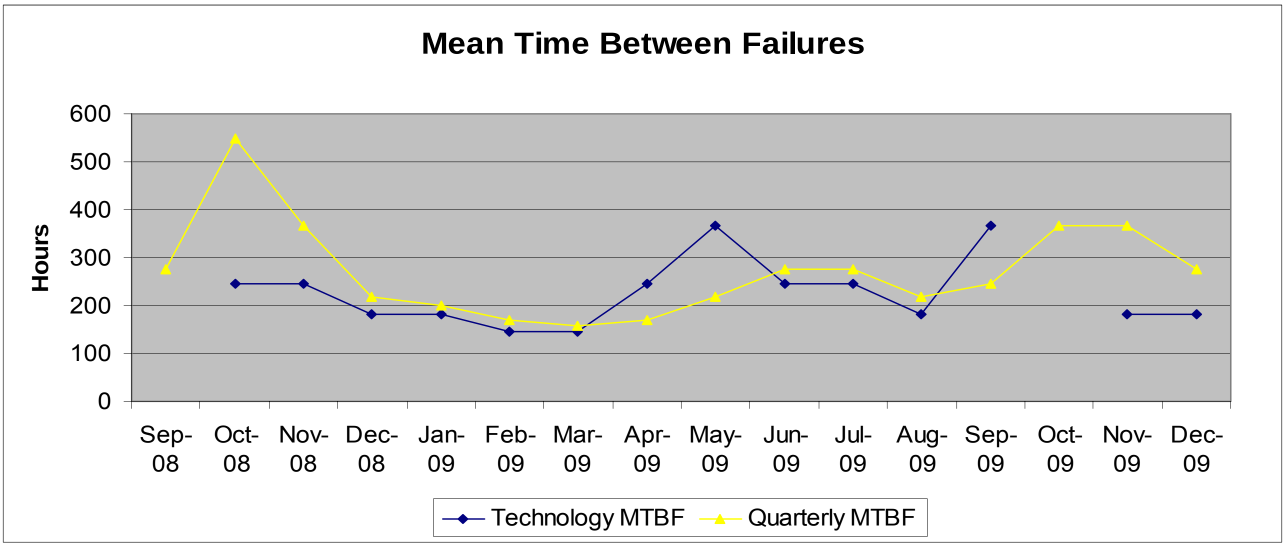

5.2.2 Service Provision Failures

There were 5.5 service provision failures during 2009. Two failures were attributed to filesystem problems on the hardware hosting the website and SAFE. In both of these cases user logins and jobs on the main service were not affected. Changes to the filesystem configuration are now in place and should now prevent this from happening again.

One failure was as a result of a process error in the lead-up to a maintenance slot; the relevant procedures have now been updated.

Plant failures accounted for 2.5 failures. In one incident a valve failed. The second failure was shared between the site and external parties. A severe fault on the external power network triggered a plant failure. The final plant failure was attributed to a UPS failure. During a maintenance period a critical component in one of the UPS modules failed. The component was replaced within hours but there were subsequent configuration parameter problems that prevented the unit from being put back on-line. The service is investigating how the effect of such a problem could be minimised in the future by configuring the UPS units in parallel, but there are unit compatibility and switching safety issues to be addressed

5.2.3 External Failures

There were 2.5 failures in 2009 as a result of external causes. As mentioned above, one failure was shared with the site when an external power fault triggered a plant failure.

Two failures were attributed to Linux security vulnerabilities. During these outages users were prevented from logging in to HECToR, although any jobs in the queues continued running. The associated risk of using an open source operating system has been highlighted to Cray. It is unlikely that this will change however.

5.3 Single Node Failures

The overall volume of single node failure in 2009 peaked in July as a result of the quad-core upgrade but quickly returned to more typical levels. There was a total of 250 node fails in 2009, compared to 298 in 2008, which equates to a 16% improvement year on year.

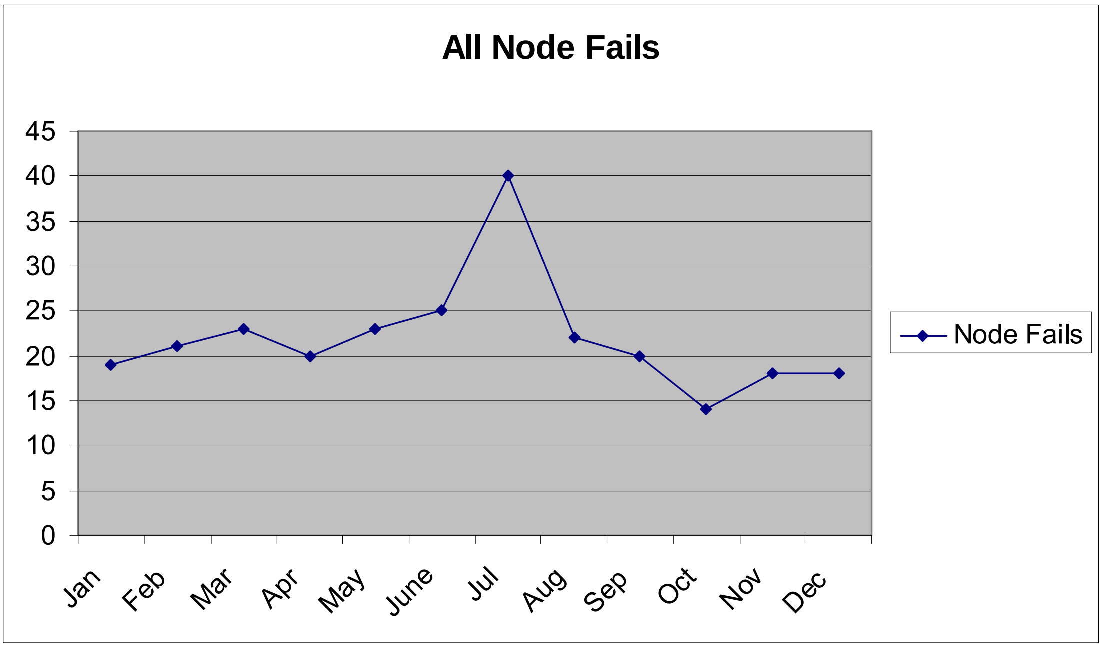

Single node failures can be responsible for the loss of individual user applications. These errors can be broken down into three main categories:

(1) Un-correctable Memory Errors (UME's)
(2) AMD Opteron cache errors
(3) Cyclic Redundancy Check errors from High Speed Network packets

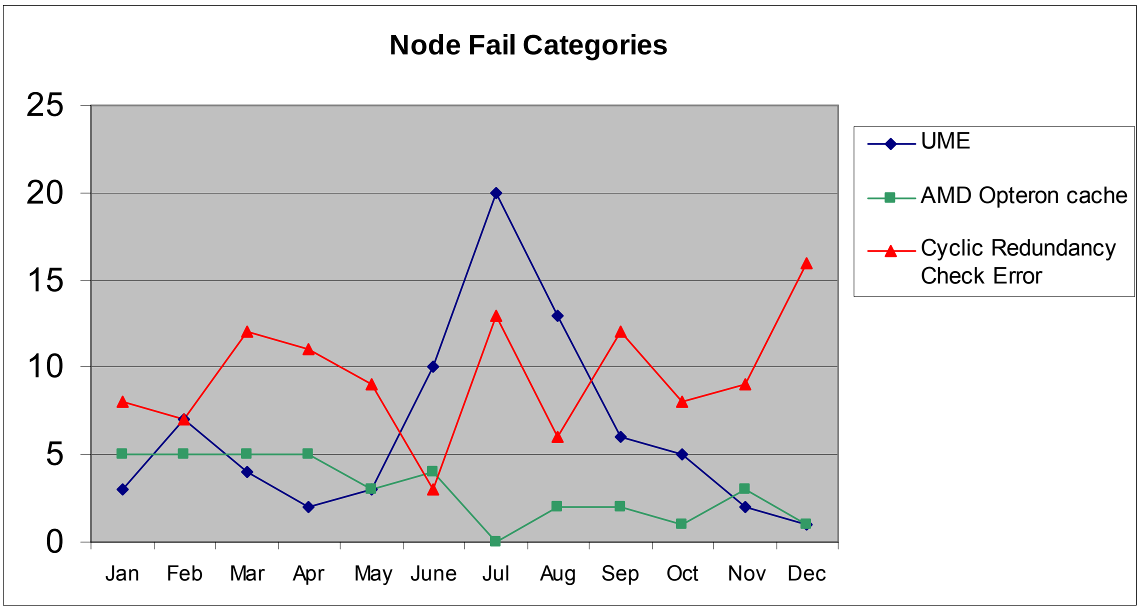

With 5664 compute nodes and 22,848 memory DIMMs in the system we will inevitably see hardware failures in these components. The upgrade to quad-core processors and larger memory in June has resulted in reduced levels for both the un-correctable memory errors (UME's) and cache errors in the AMD Opteron processors, though there was an expected temporary increase due to `infant mortality' immediately after the upgrade occurred.

The rate of Cyclic Redundancy Check (CRC) errors in the High Speed Network which result in the failure of the node receiving a bad packet are continuing at about the same rate as in 2008. These errors can be triggered by either hardware or software events. New code designed to address some categories of software-triggered CRC error events will be installed for testing on the HECToR Test and Development System (TDS) in the near future.

6 Cray Centre of Excellence

6.1 Executive Summary

The year 2009 saw the Cray Centre of Excellence for HECToR continue to deliver targeted high-end application support to a number of user groups. The Centre was heavily involved, along with the wider Cray applications team, in preparing users for the transition to quad-core and the CLE 2.1 upgrade. This involved both improving scaling of applications to take advantage of the greater number of cores, and single-CPU optimisations for taking advantage of the wider SSE units on the quad-core nodes. The Centre has also participated extensively in the Cray Centre of Excellence Steering Committee which aims to target the Centre's activities to important UK HPC applications and users. An additional activity has been the Cray UK Focus which has consisted of a series of UK-wide workshops, the aim of which is to place a strong emphasis on scaling of UK applications to very high core counts. As part of these activities, and as part of the Centre's wider optimisation efforts, a number of applications have been run on the Petaflop jaguarpf system at ORNL, some of these scaling in excess of 100,000 cores.

In other notable areas, the Centre continues to participate in user training and workshops, both in the UK, and as part of the wider HPC community. They also participated in the distributed CSE process and Performance Working Groups; provided guidance and support for postgraduate students studying relevant topics in High Performance Computing; and prepared and delivered assessments on the impact of possible HECToR Phase 2B configurations on the UK's HPC application base. In addition to the Centre's core staff, several members of Cray's Application team were involved in HECToR activities where such specialist input was needed.

6.2 Applications Support

6.2.1 Cray UK Focus

The Cray UK Focus is an initiative led by John Levesque, the director of the Cray Centre of Excellence, and Andy Mason, Sales Director for Europe, which has aimed to put a very strong focus on application scaling using experience gained from extreme scaling on the jaguarpf system at ORNL. The Cray UK Focus activities have consisted of a series of workshops around the UK where the CoE has identified candidate applications for scaling to very high core counts. The workshops also consisted of a series of presentations on how scaling to hundreds of thousands of cores was achieved on jaguarpf. Through the Cray UK Focus, several applications have been examined including SBLI, LUDWIG, DLPOLY, SENGA, Helium and CASINO.

- **SBLI-Shock Boundary Layer Interaction**: SBLI is a code originally developed by Professor Neil Sandham at the University of Southampton for studying turbulent flows. The Cray CoE for HECToR has been running the SBLI code on jaguarpf to very high core counts to examine the scaling behavior of the code. In order to get to these very high core counts, some scaling optimizations were performed.

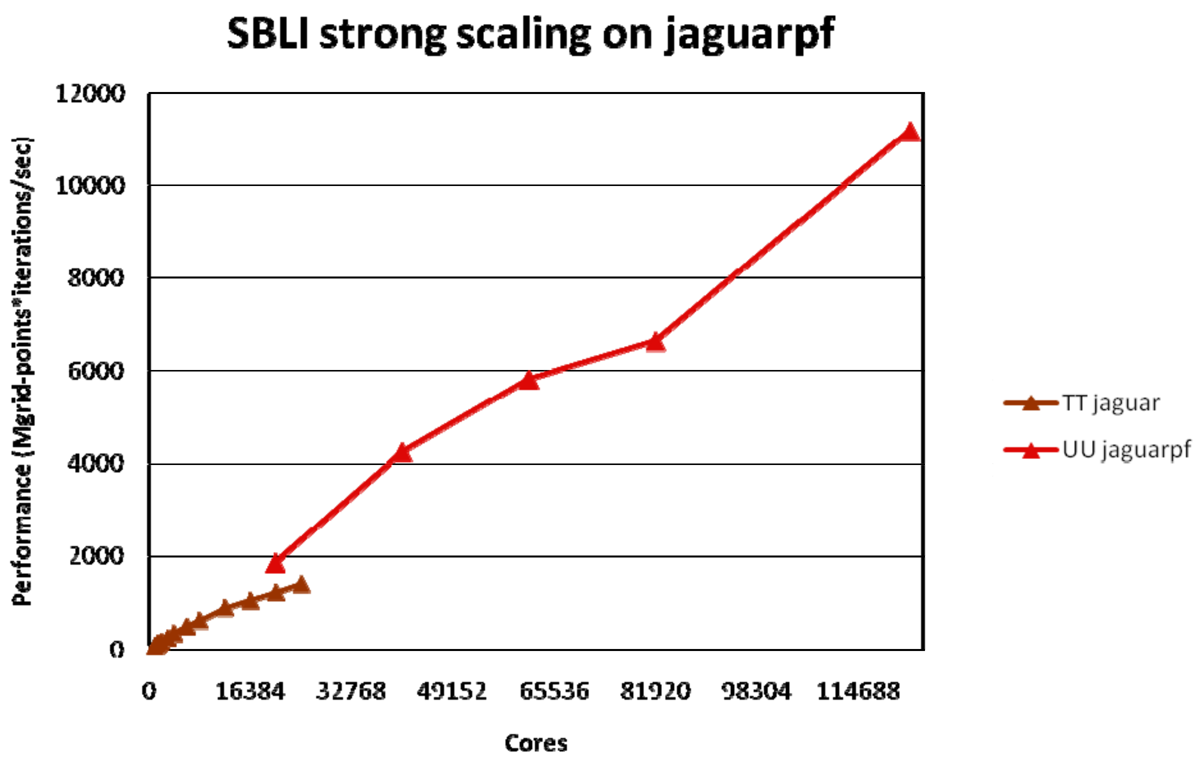

- **Ludwig**: Following on from work done for the Capability Challenge project, several very large runs of Ludwig were performed on the jaguarpf system. Two versions of the code were run at scale, and through these runs it was identified that a particular optimisation used in the code becomes very useful at high core counts.

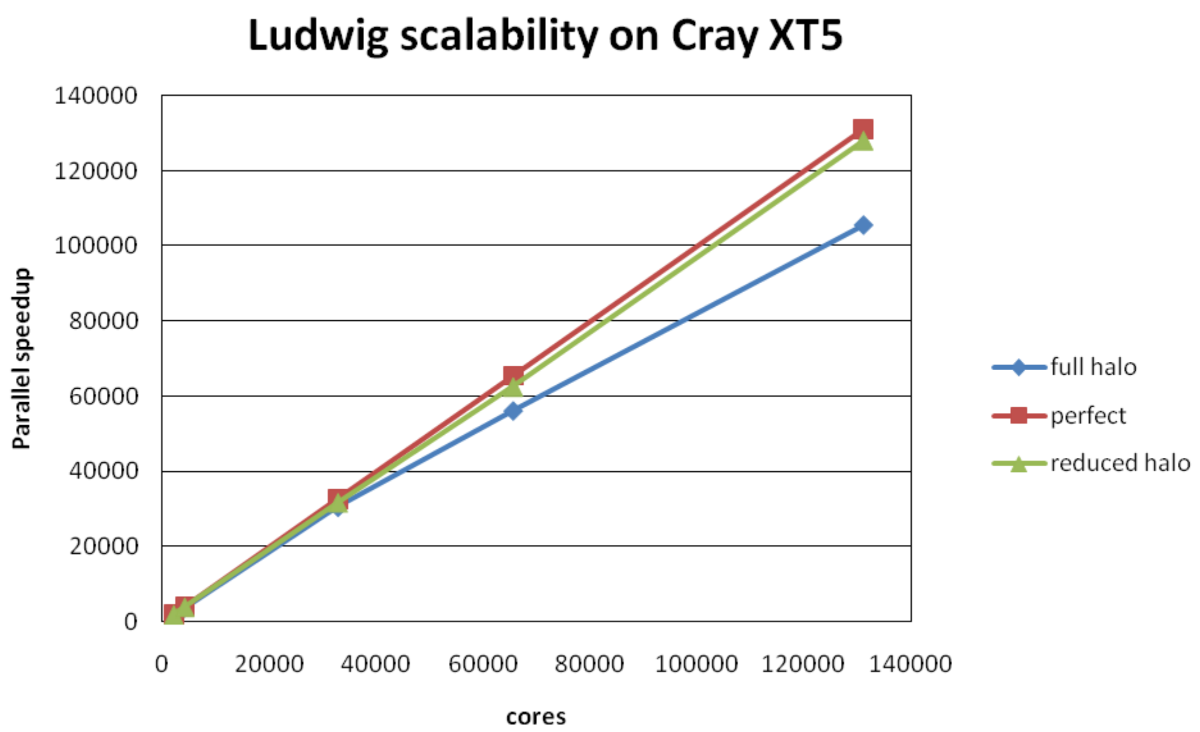

- **DLPOLY**: DL-Poly was run at significant scale which identified a number of scaling issues within the initial configurations. No issues were found in the scaling of the main compute sections of the code.
- **Helium**: Helium was tested at a size far exceeding the capability of HECToR and found no significant problems although the size of complexity of the I/O was the subject of further improvements performed under the CoE Steering Committee Process.
- **Casino:** Casino had never seen many problems on HECToR but when this was run on much larger systems a number of features were observed, in particular the time taken to read in initial data sets grew exponentially.

6.2.2 CoE Steering Committee

This year saw the creation of the CoE Steering Committee, a joint venture between Cray, HPCX, and the UK academic community. The Committee, chaired by Richard Kenway, aims to provide some focus to the Centre's activities by targeting specific user groups which represent important HPC applications in the UK. The initial CoE Steering Committee meeting was 8th July in Edinburgh. Attendees from a number of different user groups participated and after some evaluation, four projects were identified to carry forward: SENGA, HELIUM, HiPSTAR, and HiGEM. The CoE has been actively working on these applications in recent weeks with the aim of providing the users with an Cray XT optimised version of their code, and making the necessary changes to the applications so that the users could return the science which was discussed at the Committee meeting.

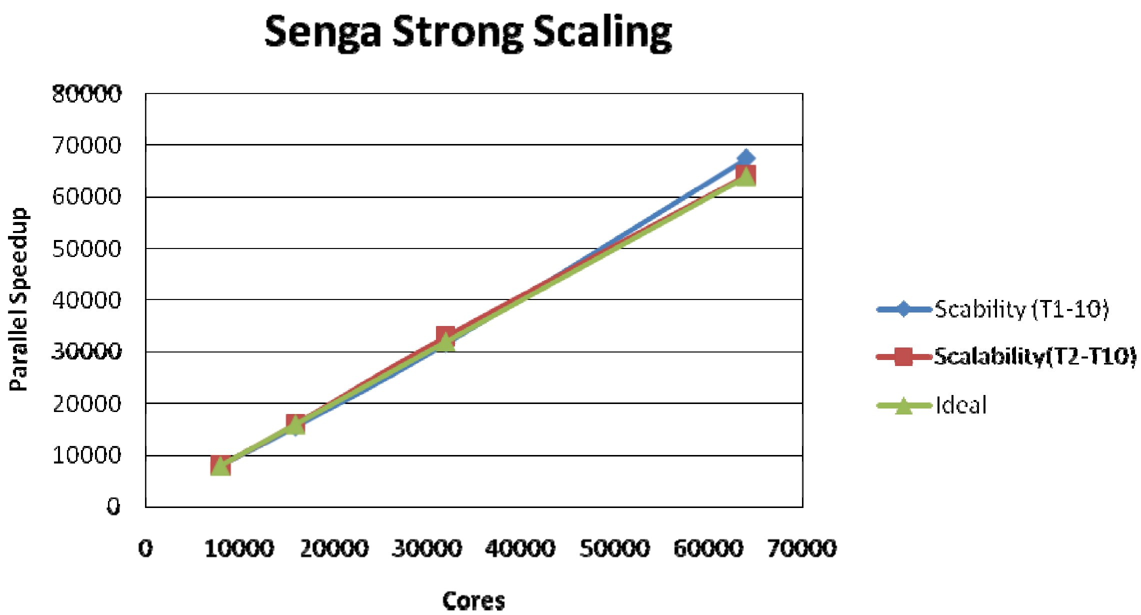

The DNS combustion application SENGA was optimised for Stewart Cant of the University of Cambridge. SENGA is an important application in the field of combustion research, both here in the UK, and internationally. The Cray CoE first produced some very large scaling runs of SENGA on the jaguarpf system to find out the scaling potential of the application. As can be seen in the figure above, the application was scaled to 64,000 cores of jaguarpf, which far exceeded the known scaling behaviour of the code.

In addition a number of routines in this application have been optimised for cache reuse leading to a factor of three in the speedup of the targeted routines.

The HiGEM project aims to test the skill of climate models in predicting decadal scale environmental phenomena like El Nino and help improve confidence in longer term predictions. The project will use a specific version of the Unified Model which will incorporate an I/O server that moves a significant bottleneck away from the compute nodes. The Centre of Excellence is assisting users at the National Centre for Atmospheric Sciences (part of the University of Reading) to tune the IO server and the UM in order to gain improved performance that will allow more ensemble members to run in the allocated time. The results of these simulations will form evidence that will be presented to the next International Panel on Climate Change.

HELIUM integrates the time-dependent Schrödinger equation for few-electron laser-atom interactions. It is developed by Queen's University, Belfast and scales very well to over 16,000 cores on HECToR. Through collaboration with the Centre of Excellence work has focused on separating the I/O from the computation to allow Helium to scale to the limits of the next phase of HECToR and beyond. By dedicating a small number of processors to handle writing checkpoint information to disk the highly scalable computation sections can continue uninterrupted. Checkpoints which previously took approximately 100 seconds on 32,000 cores on jaguarpf are reduced to less than 1 second by dedicating 125 processors as I/O servers

HiPSTAR is a CFD simulation of sub sonic jet engines developed at the University of Southampton. The project will perform a direct numerical simulation of fluid flow around a jet engine to further understand the noise generating mechanisms. A study of this nature has only become possible with the increased computing power available on HECToR. The Centre of Excellence has tuned the serial and parallel sections of the code, creating an 8% overall improvement on smaller numbers of processors and increasing HiPSTAR's ability to scale up to beyond 2048 cores.

6.2.3 Other Applications

The Cray CoE has engaged with a number of other projects within the HECToR project, including:

- **OpenFOAM**: OpenFOAM is an open source multi-physics package for continuum mechanics which is used both in academia and industry for CFD and solid mechanics simulations. The Cray CoE for HECToR, in conjunction with EPCC, has done a HECToR port of OpenFOAM to HECToR. The CoE is also supporting a dCSE project at EPCC to make OpenFOAM available supported software on the HECToR system. Scaling results for OpenFOAM, obtained on an in-house Cray XT5, are shown below and, as far as we are aware, are the largest OpenFOAM runs done on any platform.

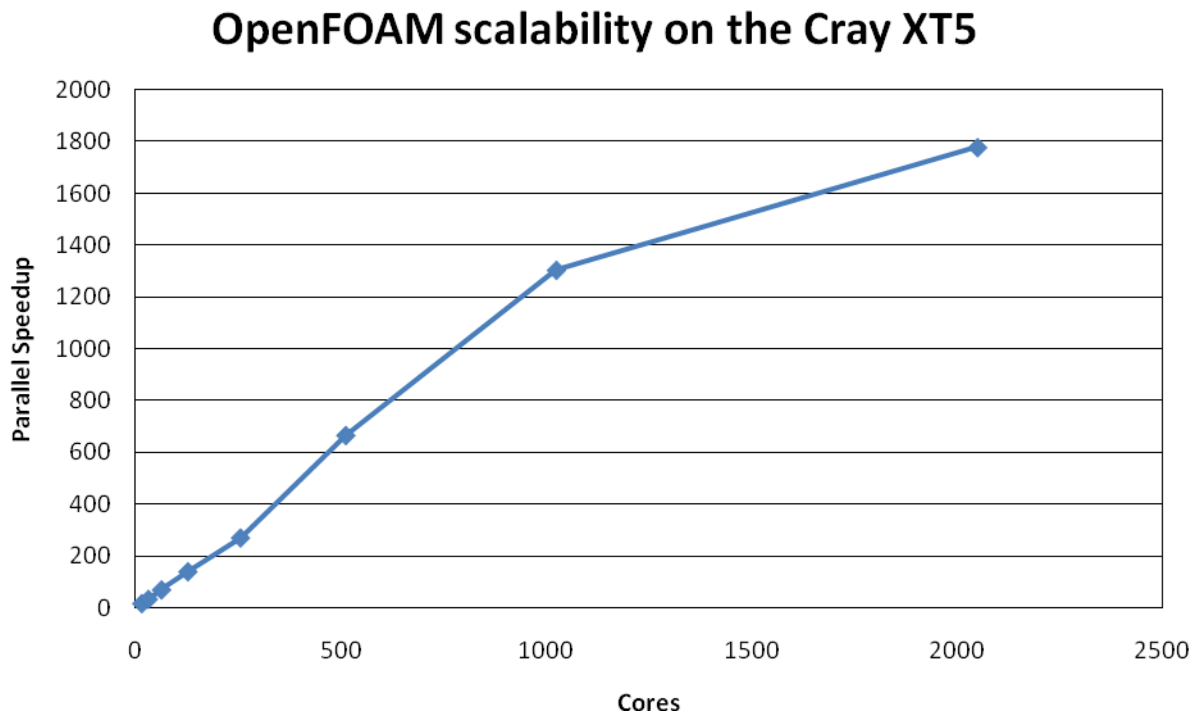

- **BQCD**: BQCD is a QCD application which is part of the DEISA benchmarking suite. The Cray CoE assisted a DEISA user in running BQCD using the full HECToR system, and have since followed up this work on jaguarpf where we have ran both the pure MPI version of the code along with the hybrid OMP/MPI version of the code. As can be seen in the figure below, the hybrid version of the code does very well at scale, and is good validation of the hybrid OMP/MPI approach on the Cray XT.

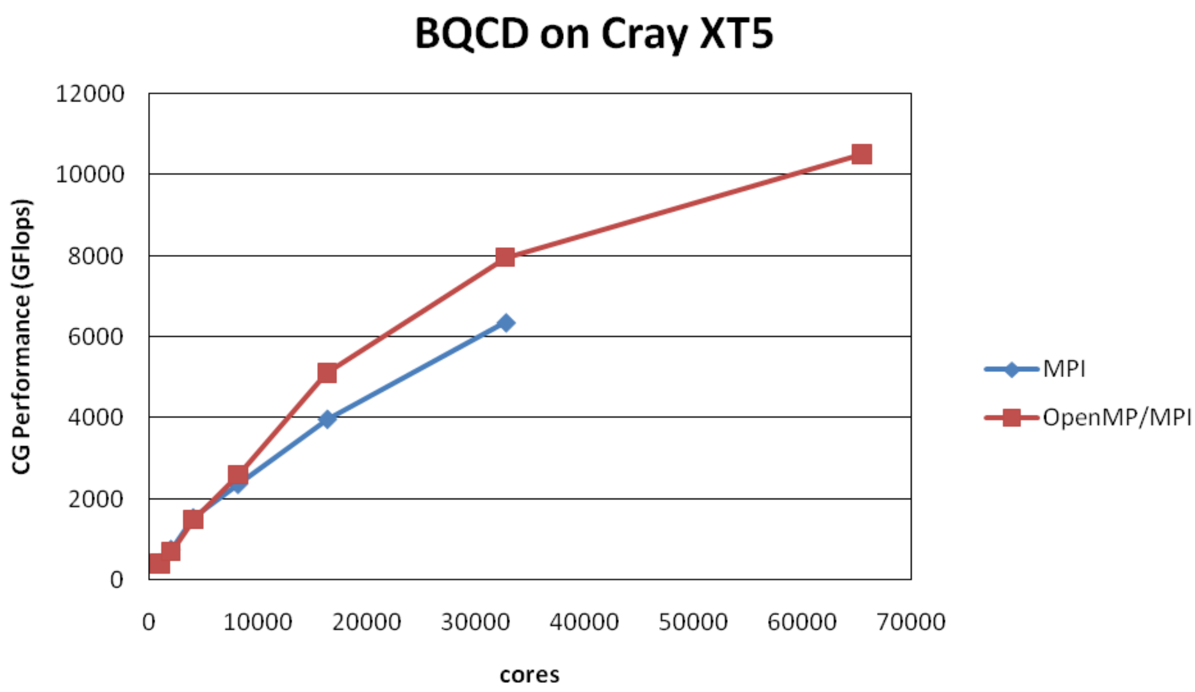

- **TAU**: TAU is a CFD code developed by DLR in Germany and used by Airbus for simulation external flows. TAU is used extensively by Airbus for simulation aerodynamics of aircraft and was used in the design of the Airbus 380. The Cray CoE performed a number of scaling runs of TAU on the HECToR system, running the code to 8,000 cores, which is in excess of the scaling that DLR thought was possible with the provided dataset.

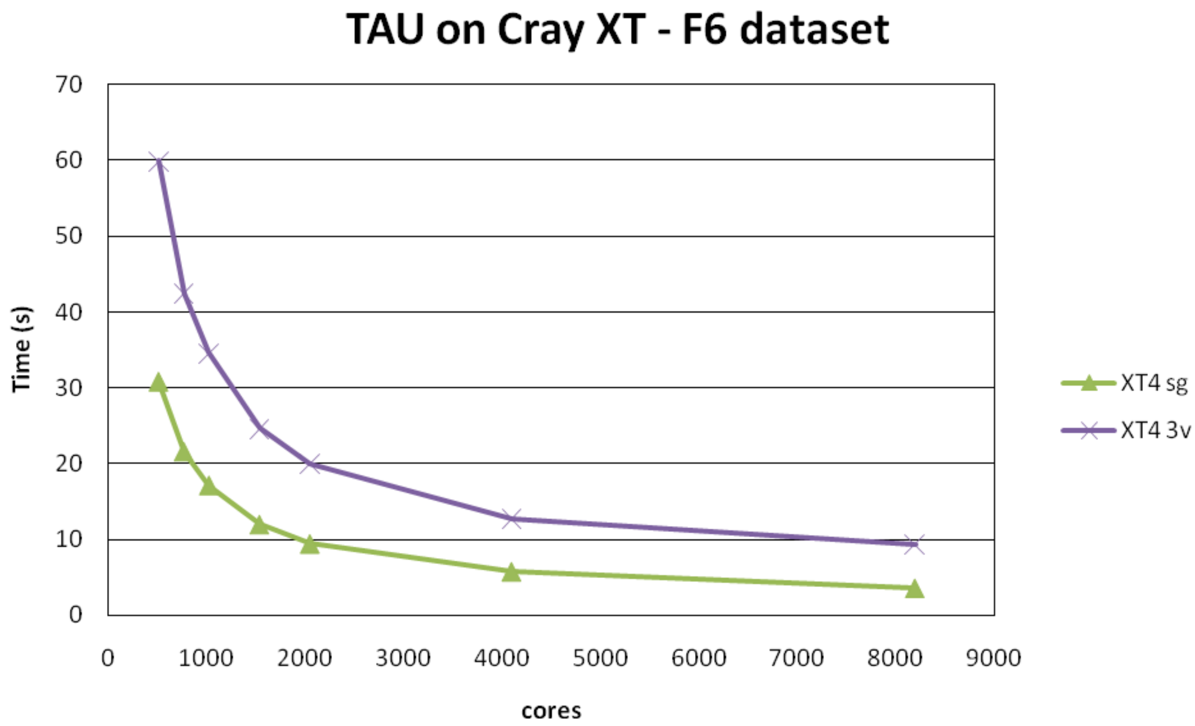

6.2.4 Phase 2 Development

The CoE continued work on analysing application performance on quad-core processors to analyse the performance characteristics when increasing the number of cores; these applications use production codes and datasets. In addition the CoE has continued this work for Phase 2B of HECToR both in terms of involvement of selecting benchmarks for the Phase 2B system and preparing performance estimates.

6.3 Workshops and Training Events

The Cray CoE conducted a number of training events in 2009:

- On the 22nd April the Cray CoE gave a presentation titled “Cray and the Quad-core Experience” at the HECToR Town User Meeting in London. The focus of this presentation was Cray’s experience in upgrading systems and applications to quad-core, and giving initial results for selected HECToR applications on quad-core.
- On the 17th-18th July the Cray CoE, with assistance from Cray technical team, gave a Cray XT quad-core workshop at NAG in Oxford.
- There were four UK Focus meetings in 2009 with meetings held in Edinburgh and Oxford on the 2nd & 3rd June and 8th & 9th September. The focus of these workshops

was to expose the UK community to the issues in scaling applications to the largest systems in the world. As part of this activity a number of UK codes were evaluated on jaguarpf which currently has over 200,000 cores.

- On the 8th July the CoE gave a presentation to the CoE Steering Committee outlining progress to date and current project status.

6.4 Seminars

The Cray CoE has given a number of external seminars of interest to the UK HPC community:

- HLRS Cray XT5m Launch Workshop: The HECToR CoE gave a presentation “CFD Applications at Scale” at the HLRS Cray XT5m Launch in Stuttgart on 18th March. The presentation covered CFD applications at scale using examples from the HECToR and jaguarpf systems and also covered the use of PGAS languages in CFD.
- NAIS Launch and Workshop: The HECToR CoE gave a presentation “Experience with PGAS Languages” at the NAIS Launch and Workshop in Edinburgh on 19th-21st October. The focus of this presentation was the use of PGAS languages in HPC applications.

7 The HECToR Computational Science and Engineering (CSE) Support Service

7.1 Overview of the CSE Service

The Computational Science and Engineering (CSE) service exists to help the user community to make the best use of the HECToR hardware by providing training, web-based resources, and assistance with porting, optimisation and tuning of software. The service is provided by the Numerical Algorithms Group Ltd (NAG), a not-for-profit Company with offices in Oxford and Manchester, and 40 years experience developing mathematical and statistical software. The *Core Team*, made up entirely of NAG staff, responds to in-depth software problems reported by users via the HECToR helpdesk, processes Technical Assessments related to applications for access to HECToR, runs a range of training courses and maintains a range of good practice guides and reference material as part of the service website, and undertakes various outreach activities. The *Distributed Team*, made up of a mixture of NAG staff and staff employed via a contract with a third party, provides dedicated resources for projects to enhance specific applications or support groups of users for periods of between six months and two years.

7.2 Highlights of 2009

2009 has been a busy year for the CSE service. Our helpdesk has continued to be busy and to receive extremely positive feedback from users. Our training courses have been a big success and are now being run at a variety of locations. The Distributed CSE programme is proving extremely popular and the projects completed so far will lead to savings of several million pounds through improved code efficiency.

The big event for the service as a whole during 2009 was the upgrade of the HECToR machine to new quad-core processors. The CSE service helped users to make this transition in several ways: advice and assistance from the core team; case studies on the HECToR website; revised and new training courses; and Distributed CSE effort. The lessons learned during this upgrade will be applied during the introduction of the CRAY XT6 in 2010.

7.3 The CSE Helpdesk

The Core CSE team handles queries from users forwarded by the service helpdesk, carries out technical assessments of applications for HECToR time, undertakes various outreach activities and runs the training courses. The latter two activities are handled separately in this report.

The queries received by the CSE team vary from straightforward requests for advice to requests for assistance in porting, tuning etc. Some queries are resolved straight away while others develop into small projects lasting weeks or even months. The team resolves most queries but if, after investigation, they are found to be connected to system issues, they will be re-assigned to the Service Provider (UoE HPCx) or to Cray. The following table shows how many queries were opened and closed each month. The peak in October was related to requests for Technical Assessments of proposals submitted to the *Software Development for HPC* call.

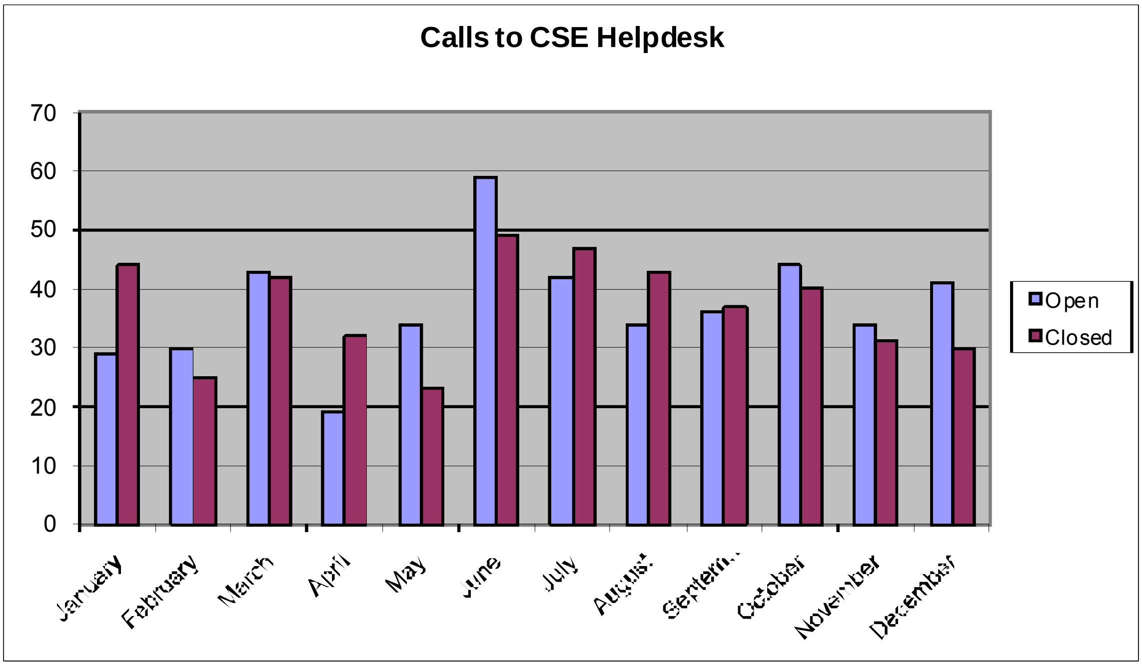

In cases where a technical query (as opposed to a request for a Technical Assessment) is resolved by the team (strictly speaking where the query is closed within the SAFE system by the CSE team), the user is invited to fill in a questionnaire giving feedback about his or her experience and satisfaction with the outcome. This year, 106 have been returned. The responses to the questions are summarised as follows:

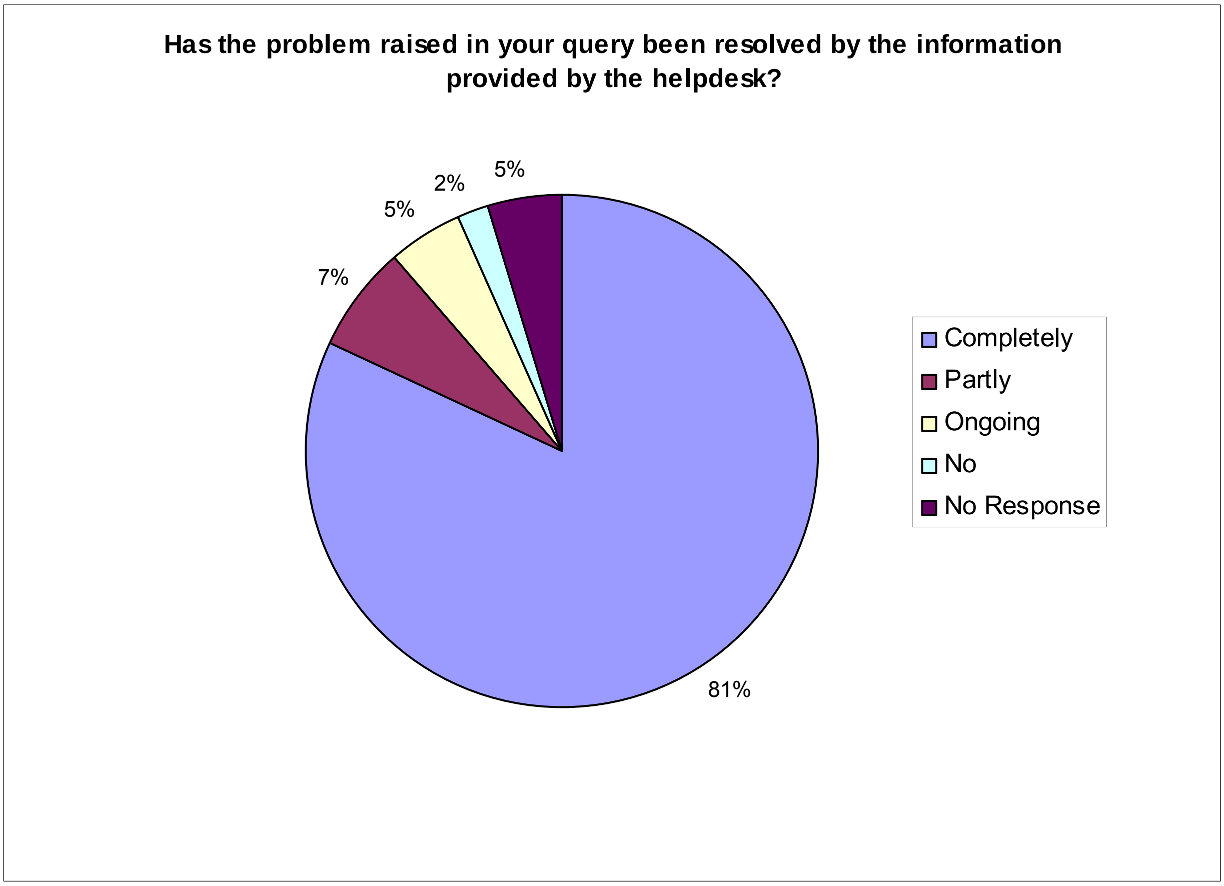

Has the problem raised in your query been resolved by the information provided by the helpdesk?
7%
5%
2%
5%
81%
Completely
Partly
Ongoing
No
No Response

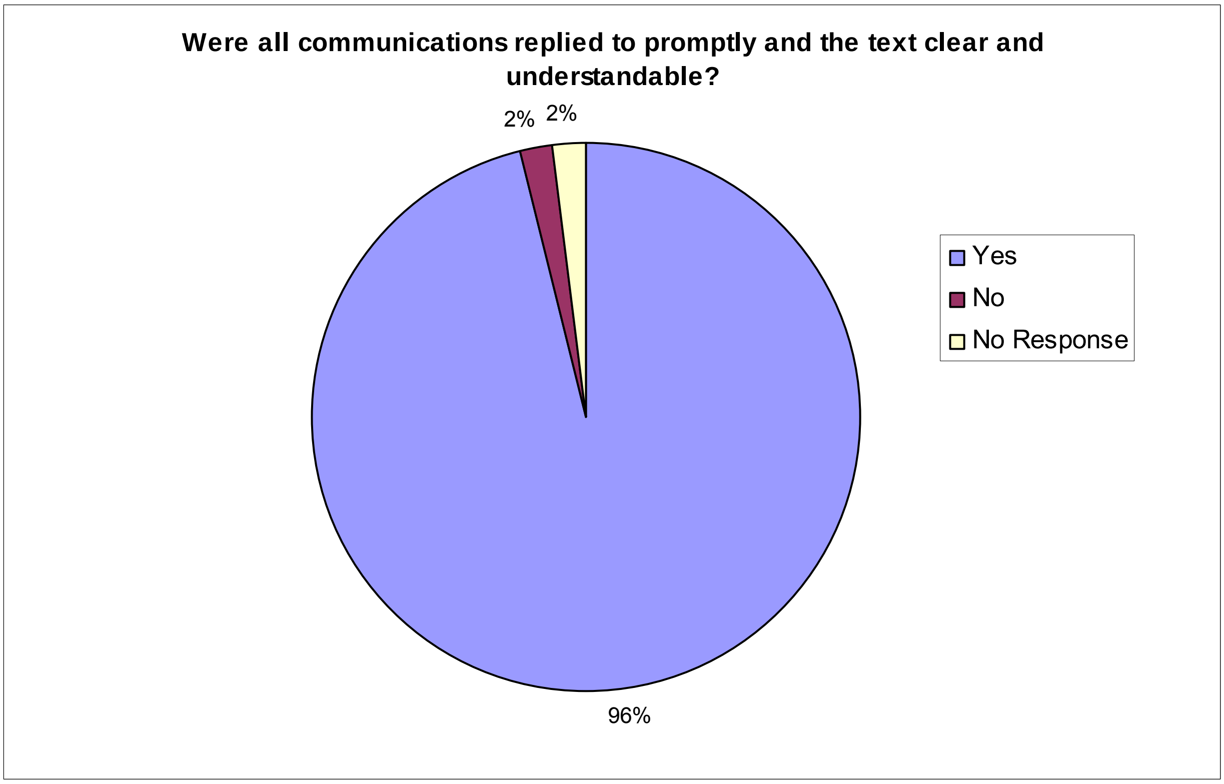

Were all communications replied to promptly and the text clear and understandable?
2%
2%
96%
Yes
No
No Response

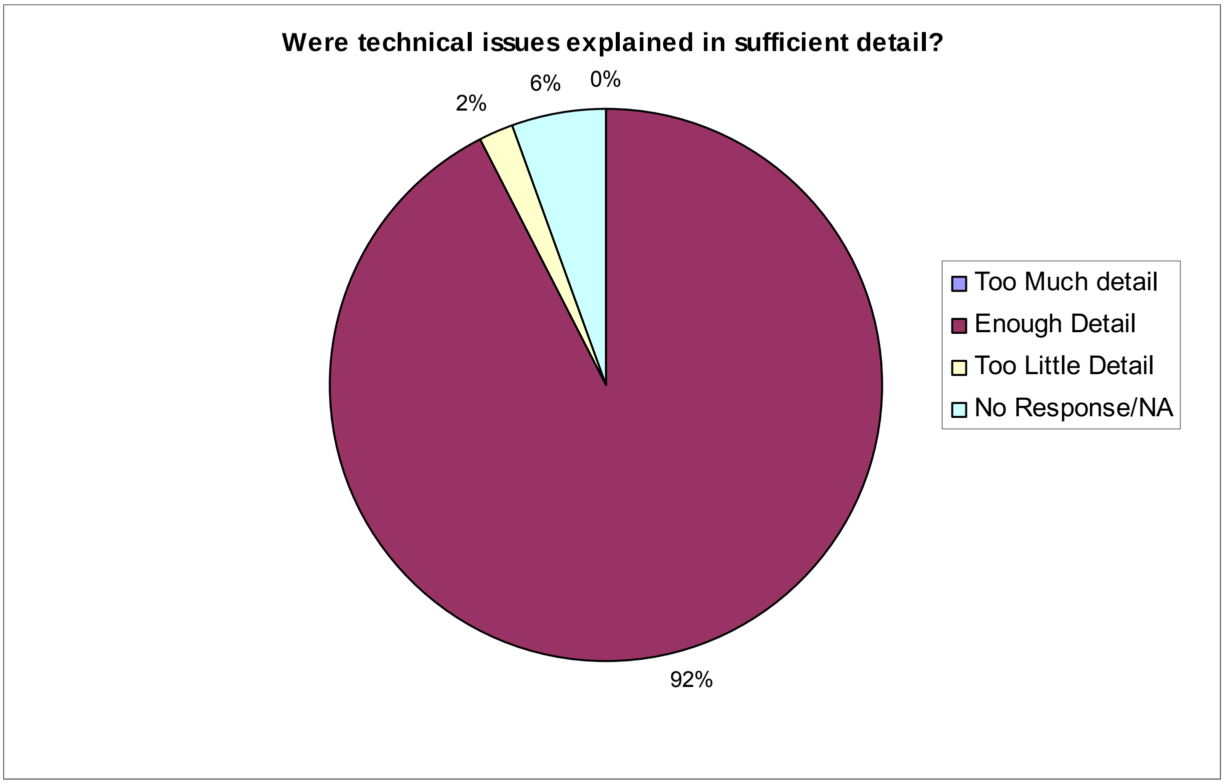

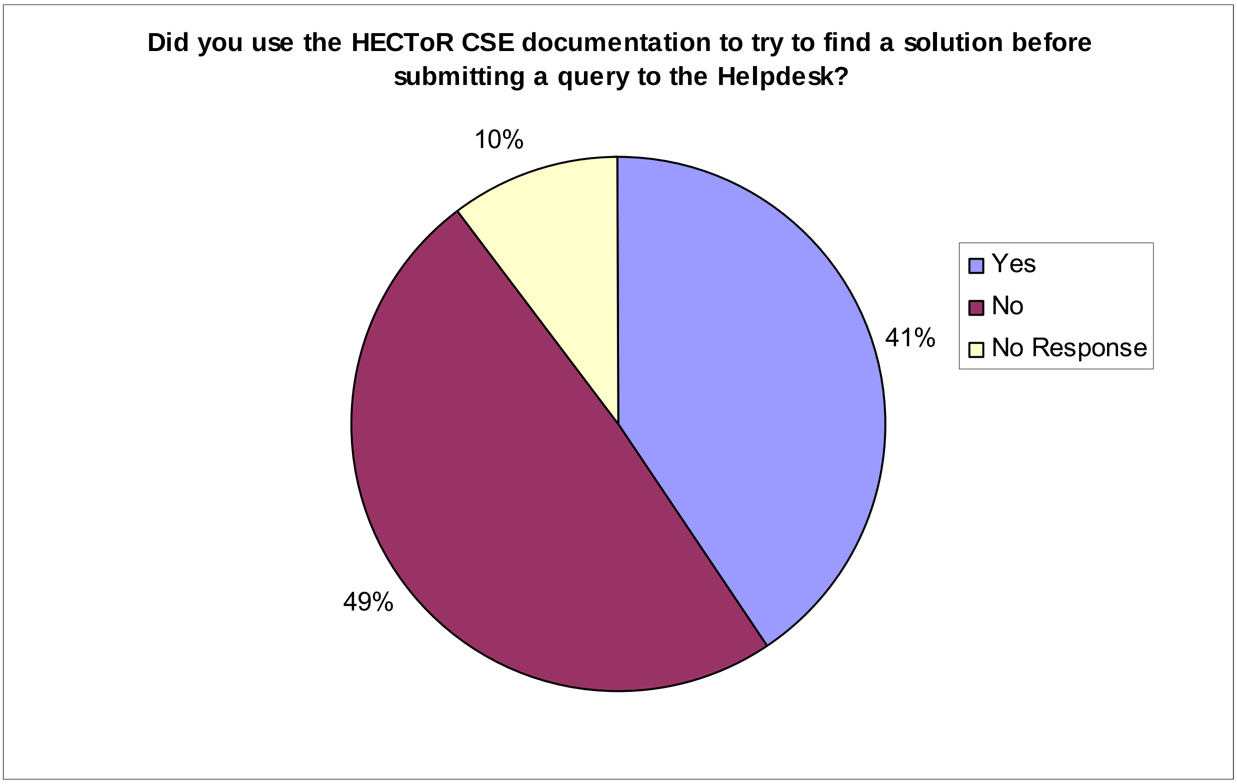

On the relatively rare occasions (less than 2% of responses) that a negative response is received in feedback the CSE team will attempt to understand the reasons behind the response and, if necessary, change or improve their procedures. Most feedback is, however, extremely positive.

7.4 Training

During this period NAG has offered a range of training course, varying from HECToR-specific courses (e.g. *Programming the X2 Vector System),* through more general HPC courses (e.g. *Parallel Programming with MPI*), to courses teaching good practice in software engineering (e.g. *Best Practice in HPC Software Development*). Most courses have been run by NAG staff, although we have run two courses in collaboration with the University of Warwick (*Introduction to HPC* and *Core Algorithms for High Performance Scientific Computing*), and one with STFC (*DL_Poly*).

Take-up for the courses this year has been extremely good, largely because many of the courses have been delivered at users' institutions. We have delivered training at Imperial College, Bristol, Southampton, Warwick, Cambridge, Belfast, Exeter, and Oxford, as well as the NAG offices in Oxford and Manchester. A total of 78 days of training have been offered and delivered, slightly above our contractual level, and there have been 443 registrations.

The complete list of courses offered in 2009 is as follows. The courses range in duration from one day to a week, and complementary courses are often scheduled together for convenience. Many of these courses were run several times in different locations.

- *Introduction to HECToR*
- *Tools and Techniques for Optimising Parallel Codes*
- *Introduction to High Performance Computing*
- *Programming the X2 Vector System*
- *Introduction to MPI*
- *Advanced MPI*
- *OpenMP and Mixed-Mode Programming*
- *Parallel Programming with MPI*
- *Best Practice in HPC Software Development*
- *Core Algorithms for High Performance Scientific Computing*
- *DL_POLY*
- *Parallel I/O*
- *Quad Core Training*

In addition, we ran a training course at a NERC summer school with two streams – one for beginners and one for advanced users.

A number of these courses were new and developed in response to user requests. The five day *Introduction to High Performance Computing* course was developed at the request of EPSRC for recipients of its HPC Software Development grants.

7.5 Outreach

The CSE team has attended a number of national and international meetings to promote HECToR and to learn from experiences of other HPC users and support services. In addition, members of the team make regular visits to HECToR users to discuss their experiences with the service, identify possible training needs or to help them to apply for distributed CSE projects.

Internationally we have visited The National Energy Research Scientific Computing Center (NERSC), Oak Ridge National Laboratory, and Sandia National Laboratories. This helped us to understand the impact of the move from dual- to quad-core, and going forward will help us to manage the transition of the new CRAY XT6 system from the SeaStar to Gemini interconnect. Members of the CSE team have been given accounts on *Jaguar*, currently number 1 on the Top 500 list, to allow them to investigate scaling of codes way beyond what HECToR currently supports, in anticipation of the increase in cores at phase 2B.

Outreach also includes the active dissemination of CSE success stories to the UK and international community (e.g. CSE support processes, new widely applicable optimisation methods, or useful routines/libraries developed under CSE/DCSE work).

7.6 Other Core CSE Activities

The CSE team publishes a newsletter approximately every six weeks which contains information about upcoming training, deadlines for Distributed CSE applications, tips and hints for programming HECToR, and news of changes to the service, particularly the development environment.

The team has produced a number of case studies detailing how to adapt particular codes to quad-core, and comparing how different compilers perform on a basket of applications.

7.7 Distributed CSE

The aim of the Distributed CSE (DCSE) programme is to provide researchers with resources to enable them to:

- port their codes onto HECToR, in particular to work with new codes or to enable previously unsupported features in existing codes;
- improve the performance of their codes on HECToR;
- re-factor their codes to improve long-term maintainability;
- take advantage of algorithmic improvements in the field of high-performance computing.
-

DCSE is designed to complement the usual instruments for funding research which often cannot be used to support software development and maintenance of this kind. The necessary staff can be provided by NAG, an appropriate third party, recruited specially, or seconded from an existing project at the host institution. Where necessary the CSE service will provide special training to equip the DCSE team member with the skills necessary to perform his or her project. A DCSE project can be allocated up to four years of effort, and can last for up to two years.

Applications for DCSE support are considered by an independent review panel, chaired by Professor Chris Bischof of the University of Aachen. The panel recommends which applications should be funded wholly or partially, and NAG then works to agree a contract with the applicant which covers the mechanism for resourcing the project, timescales and a programme of work. The panel met four times during 2009, and considered a total of thirty-four applications requesting 620 months of effort. Of those, twenty-three have been recommended to receive a total of 284 months of effort.

Four DCSE projects ended during 2009. At the end of December twenty-three projects were running with five more scheduled to start in January. Since the start of the service the DCSE panel has allocated almost thirty-seven person-years of effort to the community. In September NAG organised a workshop for all DCSE participants in Oxford, slides from which are available on the HECToR website.

7.8 Completed DCSE Projects

7.8.1 OCEANS 2025 (NEMO)

This project was approved for 6 months effort at the September 2007 call and ran from 01/03/08 - 30/04/09. Dr Andrew Coward of the National Oceanography Centre at the University of Southampton (NOCS) proposed the work with Fiona Reid from EPCC performing the software development.

The NEMO code is used in oceanographic research, operational oceanography seasonal forecast and climate studies. As well as being widely used by several NERC research groups it is also one of the main ocean modelling tools in use by the UK Met Office.

Investigations into grid layout and the removal of land-only cells performed under this DCSE project resulted in significant reductions to the AU usage for a given simulation on HECToR, by as much as 25% for larger processor counts, and a corresponding reduced runtime. NEMO has also been converted to use netCDF 4.0 for its main output files resulting in a reduction in the output file size of up to 3.55 times relative to the original netCDF 3.x code. It is expected that a production length research simulation should benefit in further reduced run time due to fewer I/O bottlenecks resulting from the reduced I/O data sizes.

The NOCS group has used around 6M AUs running NEMO in 2008. Reducing the wall clock time of NEMO by up to 25% could result in a saving in notional cost of AUs by as much as £400k for the remainder of the HECToR service, for only 6 months of person effort.

7.8.2 Improving the Parallelisation of CASINO

This project was approved for 12 months effort at the January 2008 call. The project was supervised by Professor Dario Alfè of the Department of Earth Sciences at University College London (UCL). Lucian Anton from NAG performed the software developments, with the project running 10/03/08 - 31/08/09.

The CASINO code performs Quantum Monte Carlo (QMC) calculations on condensed matter. Such calculations are of rapidly growing importance, because they are generally more accurate than calculations based on other techniques such as density functional theory (DFT) or Hartree-Fock.

The objectives of this DCSE project were to enable the CASINO code to use multicore processors more effectively and therefore be capable of modelling more complex physical systems with greater computational efficiency.

Shared memory techniques were introduced to allow larger models to be computed with greater efficiency by enabling multiple MPI processes on a single node to share a common data set in memory, thus reducing the number of nodes needed for a given simulation.

Further work including introducing hierarchical parallelism with OpenMP and I/O optimisations improved the scalability of the code, enabling CASINO to run 60-80% faster for simulations using more than 10,000 cores.

It is estimated that this DCSE work saved around 12M AUs for a one year research project on HECToR, representing a saving in notional cost of AUs by as much as £760k, with several millions of pounds of savings when applied to future research on HECToR and other supercomputers which are also used to run CASINO.

7.8.3 Enhanced scalability of the domain decomposition within CP2K

This project was approved for 6 months effort at the April 2008 call. The project was proposed by Dr Ben Slater of the Department of Chemistry at UCL. Iain Bethune from EPCC performed the work during 01/08/08 – 27/09/09.

The CP2K code is used to perform atomistic and molecular simulations of solid state, liquid, molecular and biological systems. It facilitates the use of different methods such as density functional theory (DFT) using a mixed Gaussian and plane waves approach, and classical pair and many-body potentials.

The objectives of this project were to improve the MPI communications and load balancing, particularly for the Fast Fourier Transforms within the code.

The overall performance gains from this DCSE work show a speedup of 30% on 256 cores for a generally representative benchmark case. An even greater speedup up of 300% on 1024 cores was made possible for larger test cases. These improvements are now in the official release version of CP2K and are being used by CP2K users worldwide. Researchers on HECToR alone currently use in excess of 10M AUs per year.

7.8.4 Performance enhancements for the GloMAP aerosol model Part 1

This project was approved for 6 months effort at the August 2008 call. The project was supervised by Dr Graham Mann of the School of Earth and Environment at the University of Leeds. Mark Richardson from NAG performed the software development, with the project running 17/10/08 – 31/07/09.

GloMAP is used to simulate global aerosol processes in climate modelling. It is a combination of a chemical transport model TOMCAT, which has algorithms for integrating the gas phase chemistry, deposition and advection of the trace gas and aerosol species around the globe and also a size-resolved aerosol microphysics module to simulate processes such as nucleation, coagulation, condensation and cloud-processing.

The objectives of this project were to improve the single node optimisation of the code, re-factor part of the main code structure and generally improve the performance of the current MPI implementation.

GloMAP is typically used with 64 cores or less on both the vector and scalar parts of HECToR. The performance gains from this DCSE project show a 16.5% improvement when the code is run on 64 cores. The GloMAP code is used by at least twenty HECToR users who are currently benefiting from the AUs saved by using the enhanced version of the code to increase their scientific output.

7.9 Ongoing DCSE Projects

7.9.1 Efficient Parallel Algorithms for ChemShell

The ChemShell chemical environment code uses an embedded cluster approach to the study of heterogeneous catalysis. Accurate characterisation of transition states for extended systems presents a major challenge to practical calculations but with recent algorithmic developments and HECToR this is now possible.

This aims of this project are to implement several efficient parallel algorithms within ChemShell. The project was approved for 9 months effort at the January 2008 call. The project Principal Investigators are Dr Paul Sherwood at Daresbury Laboratory of STFC and Professor Richard Catlow from the Department of Chemistry at UCL. Thomas Keal of Daresbury Laboratories is performing the software development. So far the following objectives have been achieved:

- Installation of the standard (baseline) parallel ChemShell code running on HECToR. ChemShell has been adapted to allow task farmed parallelism using split communicators.
- A speed-up factor of 2 has been achieved on a standard test calculation.
- An interface between the parallel implementation of DL-FIND has been developed such that the separate tasks within ChemShell are given independent work to perform.
- The parallel optimisation for transition-state determination has been improved by a factor of 2 for 1024 cores.

7.9.2 Cloud and Aerosol Research on Massively-parallel Architectures

This project was approved for 12 months support at the January 2008 call. Leading the project is Dr Paul Connolly of the School of Earth, Atmospheric and Environmental Sciences at the University of Manchester, with Jon Gibson from NAG carrying out the software development.

The overall aim of the CARMA DCSE is to improve the computational efficiency and operation of two key numerical cloud models that are of significant importance to the UK atmospheric science community. These models will work in synergy to improve our understanding of the impacts of micro-scale processes on the evolution and life-cycle of clouds.

Within the proposed objectives of the project the following has been carried out so far:

- A major improvement to the pressure solver in the Met Office Large Eddy Model (LEM) has been implemented. This was done by replacing the original Fast Fourier Transform (FFT) method which required periodic boundary conditions with the bi-conjugate gradient method. Jacobi preconditioning is also used for this one dimensional data decomposition.
- This newly implemented iterative solver has been fully tested and validated from known solutions.

7.9.3 Parallel Algorithms for the Materials Modelling code CRYSTAL

This project was approved for 24 months effort at the January 2008 call. The CRYSTAL code is unique in that it can be used to perform Hartree-Fock (HF), Density Functional Theory (DFT) or HF-DFT hybrid calculations with periodic boundary conditions. The code is under continuous development in order to extend its applicability to emerging problems in materials modelling and to ensure its efficient use on modern computer architectures.

Leading this project is Professor Nic Harrison of UCL, assisted by Stanko Tomic from Daresbury Laboratory. To date the following have been achieved:

- The memory bottleneck in the code which is caused by a replicated data structure of the HF and density matrix has been removed by replacing those data structures with their irreducible representations. Results show speedups in multiples of problem size/ (2 x degrees of symmetry of the system).
- Implementation of the new MPP_BLOCK command provides better control of the block size of distributed data. Reducing the block size from the default value of 96 to 64 or 32 achieves more than a 10% speed-up in the diagonalisation part of the code and a 20% speed-up in the back and similarity transform can be demonstrated on 3584 cores.

The improvements are essential to future applications of the code and will benefit a large user base on HECToR and on other supercomputers used by UK researchers.

7.9.4 Better Parallelism for the electron-atom scattering codes PRMAT

This project was approved for 12 months effort at the January 2008 call. Leading the project is Dr Martin Plummer of Daresbury Laboratory, with development being carried out by Andy Sunderland and Cliff Noble. So far, the following progress has been made:

- A greater degree of parallelism has been implemented at different levels within the code to enable it to scale more efficiently with an increasing number of processors.
- Greater flexibility and more automation regarding the configuration of the parallelism for a particular target run now make the code far easier to use.
- Splitting the initial diagonalisation stage of a run into sub-groups has resulted in strong scaling performance (around 80% or more efficiency).
- The work was presented at the Cray User Group meeting (CUG 2009) in Atlanta, May 2009.

7.9.5 Porting and Optimisation of Code_Saturne® on HECToR (XT4/X2)

Code_Saturne® is an open-source finite volume CFD code, developed (and used) by EDF. It solves the Navier-Stokes equations on arbitrary unstructured grids. There are many users of this extremely versatile code and it features in the LES UK consortium computational resources on HECToR.

This project was approved for 18 months effort at the April 2008 call. Leading the project is Professor David Emerson of Daresbury Laboratory. Charles Moulinec and Zhi Shang of Daresbury Laboratories are performing the software implementation. The work is split into four distinct phases:

1. Improvement of the parallelisation for the post-processing and restarting process.
2. Implementation of enhanced parallel partitioning which will enable a full simulation on HECToR, i.e. pre/post-processing and solver all in parallel.
3. Optimisation of Code_Saturne® for multi-core platforms.
4. Specific optimisation for the vector X2 machine.

The following has been carried out so far:

- Initially the two most popular software packages were used for testing and data partitioning, namely ParMETIS and PT-scotch. However, the Zoltan library has now been implemented to facilitate easier use of ParMETIS.
- Meshes for arbitrary geometries are being used for testing with around 200 million cells.

Further work will help improve the I/O situation for post-processing and restarting via the implementation of MPI-IO. This will also improve the scalability of the code.

7.9.6 WRF code optimisation for Meso-scale Process Studies (WOMPS)

The Weather Research and Forecast (WRF) Model is an internationally used medium-to-global-scale model for forecasting the weather and for process studies in the atmosphere. This project was approved for 10 months effort at the April 2008 call. To date the following has been completed:

- Guidelines drawn up to enable WRF users to choose optimal configurations of domain structures to enable efficient use of cache and available processors.
- Optimal compiler options, optimisation of key (numerically/IO intensive) routines and more efficient configurations are now available.
- A 25% increase in performance has been demonstrated with a parallel efficiency of 75% on 1024 cores across a range of jobs.
- Work in progress will further improve the performance and produce a report explaining the advantages and disadvantages of each optimisation approach tried.

Leading the work is Dr Alan Gadian of the School of Earth and Environmental at the University of Leeds, with Andrew Porter of Daresbury Laboratory carrying out the development.

7.9.7 Multigrid improvements to the Citcom code

The Citcom code is part of one of the most widely used suite of computational codes in solid-earth geodynamics. The scientific aim of the project is for the code to be able to model faults due to cracking in materials over a very large length scale range.

This project was approved for 12 months effort at the September 2008 call. Leading the project is Dr Jeroen van Hunen of the Department of Earth Sciences at the University of Durham, with Sarfraz Nadeem from NAG doing the software development. This project is beneficial to both the earth sciences and civil engineering communities due to the versatile nature of this multi-grid finite element code.

The work performed so far has included several parallel code enhancements targeting;

- Better multi-grid cycles;
- Implementation of better performance for the Gauss-Seidel red-black asynchronous communications method by the use of a boundary smoother for the Jacobi method;
- Improving mesh refinement algorithms and coarse grid operators for parallelism.

7.9.8 Hybrid time-dependent density functional theory in CASTEP Part 1

This project was approved for 10.5 months effort at the September 2008 call. The main aim of the work is to implement time-dependent density functional theory (TD-DFT) in the CASTEP code in a manner which scales to thousands of cores on HECToR.

Leading the project is Dr Keith Refson of the Rutherford Appleton Laboratory (STFC), aided by Dominik Jochym. To date the following has been achieved:

- Set up a new TD-DFT module with subroutines to compute the operation of A and B matrices on response wavefunctions.
- Implemented a serial iterative eigensolver to generate lowest few eigenvalues and eigenvectors using a library routine.
- Calculation of the excitation energies of small molecular systems at the Adiabatic Local Density Approximation (ALDA), or gradient-corrected versions thereof, level of theory.
- Implementation of a parallel distributed solver. Plane-waves, bands and k-points can be distributed across processors.
- Perform parallel tests with debugging and validation examples on larger systems.

7.9.9 Performance enhancements for the GloMAP aerosol model Part 2

This project continues on from the previous GloMAP DCSE. This work was approved for 4 months of support at the March 2009 round and started in August 2009. Dr Graham Mann of the School of Earth and Environment at the University of Leeds is again supervising the project and Mark Richardson of NAG is continuing to work with Dr Mann's group.

The aims of this project are to re-factor GLOMAP and TOMCAT in order to further parallelise the codes via a hybrid OpenMP-MPI method. This will take better advantage of the multicore nature of HECToR Phase 2 and beyond by using OpenMP on the intra-node communication and MPI to communicate between nodes, thus reducing the communications overheads per core considerably.

7.9.10 Parallelisation of CABARET

This project was approved for 12 months effort at the December 2008 round. It was proposed by Dr Sergey Karabasov of the University of Cambridge Engineering Department. The work involves the full parallelisation of the CABARET (Compact Accurate Boundary Adjusting high Resolution Technique) code, which is a serial finite volume Navier-Stokes CFD application for turbulent flows. The code was limited by problem size in its serial form. It is also a very versatile code which is capable of handling arbitrary boundary conditions, however due to this feature it also requires an irregular method of domain decomposition.

This is the first DCSE project to involve the complete parallelisation of a serial code. Phil Ridley of NAG has been working on this since March 2009. To date the following has been achieved;

- Implementation of an automated domain decomposition method using METIS.
- Implementation of a parallel version of CABARET which runs with an irregularly decomposed hexahedral grid.
- Full testing and validation has been performed. Simulations can now be performed in a few minutes that would otherwise have taken several days with the serial code.
- For the strong scaling test case of 100000 cells this sustains 60% parallel efficiency on 256 cores.

Dr Karabasov is now able to investigate problem sizes at least 100 times larger than previously and will be applying for more time on HECToR to use the code for these bigger simulations.

7.9.11 Scaling Turbulence Applications to thousands of cores

The aim of this work is to improve the parallel performance of the CFD code EBL. It was approved for 12 months funding at the December 2008 round. The project was proposed by Dr Gary Coleman of the School of Engineering Sciences at the University of Southampton. Dr Coleman was an original developer of the code twenty years ago when it was initially developed for vector machines.

David Scott of EPCC has been working on the project since April 2009. The EBL code models pressure driven flow over a single surface using spectral decomposition. The proposed work centers on the implementation of a 2-D decomposition for the FFTs which replaces the current 1-D method and will enable the code to scale more efficiently. The code has been re-factored to allow the implementation of the 2-D decomposition and this is currently being added to the code.

7.9.12 Direct Numerical Simulations (DNS) of turbulent fluid flows

This work was approved at the December 2008 round for 16 months of support and the proposed project is of a similar theme to the EBL project (above) but the implementation is different. Knowledge transfer between the two projects has been ongoing in order to help share individual experiences.

Professor J. C. Vassilicos of the Department of Aeronautics at Imperial College London is supervising the project and Ning Li of NAG is doing the software development. The Incompact3d code is based on an implicit finite difference solver which can handle flow which passes through fractal generated geometries. The work carried out so far has improved the performance and scalability of the code so that it can handle computational domains with improved efficiency. A library for multi-dimensional FFT decompositions using shared memory has been developed by Dr Li which can easily be used by other codes. Further work will be to improve the I/O performance by the use of MPI-IO and multi-core performance by the addition of OpenMP.

7.9.13 Implementation of established algorithms to extend HELIUM

This is a large software development project which was approved for 44 months of effort at the December 2008 round. The HELIUM code solves the time-dependent Schrodinger equation for a two-electron atom or ion exposed to intense laser light. The proposed work involves; code extensions for crossed laser fields and 2D Bessel transformations, hybrid MPI-OpenMP to resolve memory usage issues, transformation of HELIUM results from spherical to cylindrical coordinates and an extension of the techniques to address atoms and ions with more than 2 electrons.

The project is supervised by Professor Ken Taylor of the Department of Applied Mathematics & Theoretical Physics, Queen's University Belfast (QUB), with the work being carried out by a number of existing members of his team. The overall outcomes of this project will enable the HELIUM code to maintain the ability to produce world class science and bring further exposure to HECToR.

7.9.14 Massive Remote Batch Visualizer – MRBV

This proposal was approved for 10 months of support at the December 2008 round. The aim of the project is to successfully port AVS Express to HECToR to facilitate the handling of large data sets more easily, rather than users being required to transfer their data to their local machines for rendering.

The development will mainly be concerned with the optimisation of the parallel I/O module. Dr Martin Turner of Research Computing Services at the University of Manchester is supervising the work. George Leaver who is also based in the same department has been performing the implementation. AVS Express is widely used by researchers across several departments within the University of Manchester. A large proportion of these are also HECToR users and they will benefit greatly by the outcomes of the project.

7.9.15 Improvements in I/O for DL_POLY_3

This project was approved for 6 months of support for optimising the I/O section of the UK's flagship classical molecular dynamics (MD) modelling code DL_POLY_3. Dr Ilian Todorov of Daresbury is supervising the work and Ian Bush of NAG has been carrying out the implementation. The following points give a brief overview of the work;

- A single MD timestep takes 0.7s but the I/O takes 1000 times longer.
- An improved algorithm for I/O has been implemented using 64 dedicated writing processes for each simulation. Data is sorted before the writing process and the I/O is now 50 times quicker.

- A netCDF version of the algorithm has also been developed and will be implemented during the remainder of the project.

7.9.16 Dynamic load balancing and rigid body dynamics for DL_POLY_3

The aim of this work is to add functionality to DL_POLY_3 which is already available in DL_POLY_2. Although DL_POLY_3 code is more scalable than DL_POLY_2 it doesn't have all of its functionality. The outcome of this project will be to achieve a scalable version of DL_POLY with the extra scientific functionality. This project was proposed by Dr Ilian Todorov and was approved for 18 months of support at the March 2009 round, with Lawrence Ellision of Daresbury Laboratory doing the implementation.

7.9.17 HPC Driven Next-Generation Modelling of the World's Oceans (ICOM)

The Imperial College Ocean Model (ICOM) is a three-dimensional non-hydrostatic parallel ocean model which uses control volume finite element discretisation methods on meshes which may be unstructured in all three dimensions and which may also adapt to optimally resolve the solution dynamics.

This project was approved for 12 months of support at the April 2008 round. Dr Mike Ashworth of Daresbury Laboratory is supervising the work with Xiaohu Guo also based at Daresbury Laboratory performing the implementation. The objectives are:

- An algorithm for optimal renumbering of the mesh nodes.
- Improvements to mesh adaptivity performance.
- Better I/O by implementing a subset of writers approach.
- General code refactoring and compiler optimisation.

7.9.18 Improving Load Balancing and Parallel Partitioning in ICOM

This project was approved for 12 months of support at the September 2009 round. Dr Jon Hill of Imperial College London is leading the work with Paul Woodhams from NAG performing the implementation. The main objective of this work is to improve the scalability of ICOM through the introduction of optimised load balancing. This will be achieved by incorporating the Zoltan library within the code so that it may be used as an alternative to METIS for the irregular domain decomposition.

7.9.19 HECToR Distributed CSE Support for OpenFOAM

OpenFOAM is a widely used and versatile open-source CFD toolbox. The aim of this project is to install and test some requested features of the package on HECToR. This work was approved for 12 months of support at the December 2008 round. Paul Graham of EPCC is supervising the project with Gavin Pringle, also of EPCC, doing the software development.

So far the appropriate modules within the toolbox have been benchmarked and the results have been put on HECToR's third-party package Wiki. The code is widely used amongst the CFD groups at Imperial College London and Daresbury Laboratory. It is now available for HECToR users.

7.9.20 Improving the performance of GWW

The GWW code is a unique application which performs Density Functional Theory (DFT) calculations based on a localised Wannier method. GWW is built around the more widely known Quantum Espresso (QE) code. This project was proposed by Professor Merlyne De Souza from the Department of Electronic and Electrical Engineering at the University of Sheffield, and was approved for 16 months support at the March 2009 round. Professor De Souza is a relatively new user to HPC but realises the potential it will bring by enabling her group to perform new science that would otherwise not be possible.

The objectives of the work are to improve the performance of the all-to-all communications and the Fast Fourier Transforms (FFTs) within the QE code and also implement improved I/O within GWW. Damien Casterman who is based with Professor De Souza is working on the implementation of a 2-dimensional decomposition for the FFTs and the improvement of I/O for GWW. Work to improve the all-to-all data communication is being carried out by Iain Bethune of EPCC.

7.9.21 Fast Fourier Transformations for Gyrokinetic Plasma Simulations

This project was proposed by Dr Colin Roach of the UK Atomic Energy Authority. It was approved for 6 months funding at the March 2009 round. The code in question is GS2 which is used to study low-frequency turbulence in magnetized plasma, often to assess the microstability of plasmas produced in the laboratory and to calculate key properties of the turbulence which results from instabilities. GS2 can also be used to simulate turbulence in plasmas which occur in nature, such as in astrophysical and magnetospheric systems.

The objective of the work is to improve the performance of the Fast Fourier Transforms (FFTs) through more efficient usage of SSE3 instructions available on HECToR's AMD processors. Joachim Hein from EPCC is replacing calls to the external library FFTW2 by calls to FFTW3. This will enable the code to spend at least 40% less time performing the FFTs which is a major part of the total computational time.

7.9.22 SPRINTing with HECToR

SPRINT is an add-on to the R statistical package. SPRINT allows users to write C code to manipulate R objects directly for the more computationally intensive tasks. It is intended to facilitate easy access to the application of HPC for statistical analysis. R is used worldwide and has many contributors and it is available on HECToR as a centrally installed module. SPRINT is used extensively by the Division of Pathway Medicine (DPM) at the University of Edinburgh in order to enable genomic analysis on large data sets through the use of parallel processing.

This project was proposed by Mr Terry Sloan of EPCC who has been working in close collaboration with DPM to develop SPRINT. The project was approved for 6 months support at the June 2009 round and Savvas Petrou who is also based at EPCC has been working on the implementation of a scalable Pearson's correlation statistical function.

7.9.23 Improving scalability of CP2K on multi-core systems

This work was approved for 6 months support at the June 2009 round and follows on from the previous successful CP2K project. It was proposed by Dr Ben Slater of the Department of Chemistry at UCL with Iain Bethune from EPCC continuing his work on the project.

The work involves the implementation of multi-core improvement to the FFTs and more efficiency in the packing and unpacking of MPI buffers for the all-to-all communications. The expected performance improvement from this work will be a 30% speed up in CP2K for up to 1024 processing cores.

7.9.24 Scalability Optimization for Large Scale In-Silico Simulations of Cardiac Bioelectric Activity

The Cardiac Arrhythmia Research Package (CARP) is a software package used to model the bioelectric activity of the heart. It has been developed in C by Dr Gernot Plank from the Oxford e-Research Centre at the University of Oxford. The CARP code is based on an irregular finite element model for the solution of several sets of time dependant non-linear parabolic and elliptic partial differential equations. The PETSc package is used extensively for the linear solvers and data management.

This project was approved for support at the September 2009 round and the main aims are to improve the performance of CARP by improving two key areas which are commonly the source of scalability issues in irregular parallel data decompositions, namely:

- Improvements to load balancing by the application of METIS for the data decomposition.
- Better I/O by implementing MPI-I/O and ordering within non-blocking communication calls.

In addition to these points, there will be extra work targeting improved preconditioning for the equations solvers within PETSc along with more general code re-factoring and compiler optimisation. Lawrence Mitchell from EPCC began working on the project in December.

7.10 Plans for 2010

The big event of 2010 will be the upgrade of the HECToR system from the existing CRAY XT4 to an CRAY XT6. Assisting users in making this transition will be a key focus of the next year.

- **Training:** We will continue to run courses at a range of locations, and will be updating our existing *Quad-Core* course to deal with the new dual-processor twelve-core Magny-Cours based nodes in the CRAY XT6.
- **Outreach:** The CSE service will continue its programme of visits to HECToR users, and will work with other HPC services to share experiences about issues such as data management, resilience, new languages and new architectures. In particular, we will work with other early adopters of the CRAY XT6 to help to determine how best to exploit the novel features of the new system.
- **Distributed CSE:** We will strive to maintain the success of the Distributed CSE Programme and we will be organising another workshop to allow those working on the projects to share each others' experiences. We will be offering support to the early users of the CRAY XT6 and recipients of *Capability Challenge* grants.

Once the CRAY XT6 is available, and the CSE service has had an opportunity to gain some familiarity with it, we will organise a workshop for users to help prepare them to migrate their research to the new machine. We look forward to a successful 2010 for HECToR.

Appendix A: Terminology

TSL : Threshold Service Level

FSL : Full Service Level

SDT : Scheduled Down Time

UDT : Unscheduled Down Time

WCT : Wall Clock Time

MTBF : Mean Time Between Failures = 732/Number of Failures

SP : Service Provision

SP Serviceability% = 100*(WCT-SDT-UDT(SP))/(WCT-SDT)

Technology Reliability % = 100*(1-(UDT(Technology)/(WCT-SDT))

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 file systems 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 on HECToR

Code	Title	Funding Body	Class	Area	PI	Total AUs allocated	AUs used	AUs left
EPSRC Projects								
c01	Support of EPSRC/STFC SLA	EPSRC	Class1a	support	Dr Richard Blake	30,803,723	14,562,885	16,240,838
e01	UK Turbulence Consortium	EPSRC	Class1a	Engineering	Dr Gary N Coleman	483,969,876	2,554,805	481,415,071
e05	Materials Chemistry HPC Consortium	EPSRC	Class1a	Chemistry	Prof C Richard A Catlow	1,129,327,228	41,649,079	1,087,618,149
e10	GENIUS	EPSRC	Class1a	Chemistry	Prof Peter Coveney	9,257,856	5,561,423	3,696,433
e19	Edinburgh Soft Matter and Statistical Physics Group	EPSRC	Class1a	Physics	Prof Michael Cates	1	0	1
e24	DEISA	EPSRC	Class1a	support	Mrs Alison Kennedy	61,365,221	50,226,549	11,138,672
e34	Hydrogen vacancy distribution in magnesium hydride	EPSRC	Class2a	Chemistry	Prof Nora de Leeuw	100,000	34,124	65,876
e35	Non-adiabatic processes	EPSRC	Class1a	Materials	Dr Tchavdar Todorov	11,246,869	1,825,246	9,421,623
e42	Computational Combustion for Engineering Applications	EPSRC	Class1a	Engineering	Prof Kai Luo	32,000,001	19,022,735	12,977,266
e59	Turbulence in Breaking Gravity Waves	EPSRC	Class1a	Engineering	Prof Ian P Castro	708,922	440,752	268,170
e63	UK Applied Aerodynamics Consortium 2	EPSRC	Class1a	Engineering	Dr Nick Hills	15,904,086	10,807,853	5,096,233
e68	Hydrogenation Reactions at Metal Surfaces	EPSRC	Class1a	Chemistry	Dr Angelos Michaelides	50,000,000	31,519,545	18,480,455
e69	Simulations of a Subsonic Cylindrical Cavity Flow	EPSRC	Class2a	Engineering	Dr Aldo Rona	125,001	125,273	-272
e70	Computation of Electron Transfer Properties	EPSRC	Class1a	Chemistry	Dr Jochen Blumberger	960,000	477,401	482,599
e71	Simulating the control of calcite crystallisation	EPSRC	Class1a	Chemistry	Prof John Harding	40,403,522	40,203,133	200,389
e72	Ultrascalable Modelling of Materials	EPSRC	Class2a	Materials	Dr Lee Margetts	8,622,547	8,717,872	-95,325
e75	Terascale DNS of Turbulence	EPSRC	Class1a	Engineering	Prof Christos Vassilicos	27,881,306	27,844,966	36,341
e76	HELIUM Developments	EPSRC	Class1a	Physics	Prof Ken Taylor	42,521,798	22,115,473	20,406,325
e77	Porting of DFT/GW Codes	EPSRC	Class2a	Engineering	Prof Maria Merlyne DeSouza	160,000	60,676	99,324

Code	Title	Funding Body	Class	Area	PI	Total AUs allocated	AUs used	AUs left
e81	e-Collision experiments using HPC	EPSRC	Class2a	Physics	Prof NS Scott	257,095	11,106	245,989
e82	ONETEP: linear-scaling method on High Performance Computers	EPSRC	Class2a	Materials	Dr Peter Haynes	100,000	101,352	-1,352
e84	Vortical Mode Interactions	EPSRC	Class1a	Engineering	Dr Tamer Zaki	9,600,000	19,405	9,580,595
e85	Study of Interacting Turbulent Flames	EPSRC	Class1a	Engineering	Dr N Swaminathan	5,588,610	2,050,132	3,538,478
e89	Support for UK Car-Parrinello Consortium	EPSRC	Class1a	Physics	Dr Matt Probert	360,000,001	80,599,767	279,400,234
e90	Network modelling of wireless cities	EPSRC	Class2a	Engineering	Prof Jonathan M Pitts	100,000	58,605	41,395
e92	Dynamo Action In Compressible Convection	EPSRC	Class1a	Physics	Mr Paul Bushby	4,075,000	74,433	4,000,567
e93	ACE - Architecture Exercise	EPSRC	Service	support	Dr Alan Gray	750,000	579,968	170,032
e96	Materials Property Relationships	EPSRC	Class2a	Materials	Dr Shoufeng Yang	100,000	0	100,000
e98	Non-linear magnetohydrodynamic modelling of tokamak plasmas	EPSRC	Class2a	Physics	Mr Ian T Chapman	100,000	26,287	73,713
e100	Large scale MD and quantum embedding for biological systems	EPSRC	Class2a	Materials	Prof Zheng X Guo	100,000	27	99,973
e101	Optimization of HPCx LES code	EPSRC	Class2a	Engineering	Prof Michael Leschziner	641,009	390,356	250,653
e102	Numerical investigation of aerofoil noise	EPSRC	Class1a	Engineering	Dr Richard D Sandberg	6,484,191	3,115,728	3,368,463
e103	Micromagnetic simulations on HPC architectures	EPSRC	Class2a	Engineering	Dr Hans Fangohr	100,000	0	100,000
e104	Fluid-Mechanical Models applied to Heart Failure	EPSRC	Class1a	Physics	Dr Nicolas Smiths	2,400,000	2,606	2,397,394
e105	Joint Euler/Lagrange Method for Multi-Scale Problems	EPSRC	Class1a	Engineering	Dr Andreas M Kempf	1,300,000	272,453	1,027,547
e106	Numerical Simulation of Multiphase Flow: From Mesocales to	EPSRC	Class1a	Engineering	Prof Kai Luo	3,650,000	0	3,650,000
e107	Parallel Brain Surgery Simulation	EPSRC	Class1a	Life Sciences	Dr Stephane P. A. Bordas	6,000,000	119,764	5,880,236
e108	Unsteady Propeller Noise	EPSRC	Class2a	Engineering	Dr Sergey Karabasov	100,000	57,750	42,250
e110	Computational Aeroacoustics Consortium	EPSRC	Class1a	Engineering	Prof Paul Tucker	39,100,000	7,105,306	31,994,694
e112	Assessment of the ONETEP code	EPSRC	Class2a	Materials	Mr Andrew J Scott	100,000	102,500	-2,500
e113	[dCSE] MRBV - Massive Remove Batch Visualizer	EPSRC	Class2b	support	Dr Martin Turner	85,440	15,484	69,956

Code	Title	Funding Body	Class	Area	PI	Total AUs allocated	AUs used	AUs left
e114	[dCSE] OpenFOAM	EPSRC	Class2b	support	Mr Paul Graham	100,000	0	100,000
e115	Multiscale Modelling of Biological Systems	EPSRC	Class2a	Chemistry	Prof Jonathan W Essex	100,000	0	100,000
e116	Scaling and Benchmarking Spectral Codes	EPSRC	Class2a	Physics	Dr Benson Muite	100,000	19,133	80,867
e117	Getting started on HECTOR	EPSRC	Class2a	Chemistry	Dr Carmen Domene	100,000	247,664	-147,664
e118	Adaptive coupled radiation-transport and fluids modelling.	EPSRC	Class2a	Engineering	Prof Christopher Pain	100,000	0	100,000
e119	Nanoscale Energy Transportation	EPSRC	Class2a	Materials	Dr Dongsheng Wen	100,000	0	100,000
e120	[dCSE] FF Transformations for plasma simulations	EPSRC	Class2b	Physics	Dr Colin M Roach	200,000	0	200,000
e121	[dCSE] Improving Performance using Wannier functions	EPSRC	Class2b	Physics	Prof Maria Merlyne DeSouza	1,700,000	409,959	1,290,041
e122	Multiscale Modelling of Magnetised Plasma Turbulence	EPSRC	Class1a	Physics	Dr Colin M Roach	65,000,000	11,137,523	53,862,477
e123	Finger-jets and turbulent structures	EPSRC	Class2a	Engineering	Dr David Ingram	15,040	0	15,040
e124	Compressible Axisymmetric Flows	EPSRC	Class1a	Engineering	Dr Richard D Sandberg	15,000,000	1,760,244	13,239,756
e125	Fermion Monte Carlo in Slater Determinant spaces	EPSRC	Class2a	Chemistry	Dr Ali Alavi	100,000	252,826	-152,826
e126	Clean Coal Combustion: Burning Issues of Syngas Burning	EPSRC	Class1a	Engineering	Dr Xi Jiang	9,984,000	945,564	9,038,436
e127	Alternative drag-reduction strategies	EPSRC	Class1a	Engineering	Prof Michael Leschziner	7,000,000	144	6,999,856
e128	Rate-Controlled Constrained Equilibrium	EPSRC	Class1a	Engineering	Dr Stelios Rigopoulos	6,230,000	0	6,230,000
e129	Novel Hybrid LES-RANS schemes [ICL]	EPSRC	Class1a	Engineering	Prof Michael Leschziner	7,500,000	0	7,500,000
e130	Novel hybrid LES-RANS schemes [MAN]	EPSRC	Class1a	Engineering	Prof Dominique Laurence	10,500,000	0	10,500,000
e131	Direct Simulation of a Pure Plume impinging on a density surface	EPSRC	Class2a	Environment	Dr Maarten V Reeuwijk	200,000	0	200,000
e132	Parallel Version of a Design Sensitivity Tensegrity Code	EPSRC	Class2a	Engineering	Prof Rod Smallwood	200,000	10,184	189,816
e133	Implementation of Established Algorithms to Extend HELIUM	EPSRC	Class2b	Physics	Prof Ken Taylor	400,000	0	400,000
e134	Numerical Simulation of Turbomachinery Flows	EPSRC	Class2a	Engineering	Dr Francesco Montomoli	200,000	473	199,527
e135	DNS of unsteady turbulent flow over a smooth or a rough surface	EPSRC	Class2a	Engineering	Dr Shuisheng He	200,000	19,675	180,325

Code	Title	Funding Body	Class	Area	PI	Total AUs allocated	AUs used	AUs left
e136	Modelling the UK Wind Power Resource	EPSRC	Class2a	Environment	Dr Gareth Harrison	200,000	42,787	157,213
e137	Turbulent Pipe Flow	EPSRC	Class2a	Engineering	Prof Dwight Barkley	200,000	190,349	9,651
e138	[dCSE] Naturally occurring magnetic mineral systems on HECToR	EPSRC	Class2b	Environment	Prof Wyn Williams	400,000	0	400,000
e139	Scalability Optimization for Largescale in-silico Simulations	EPSRC	Class2b	Life Sciences	Dr Gernot Plank	400,000	80	399,920
e141	A numerical study of turbulent manoeuvering-body wakes	EPSRC	Class1a	Engineering	Dr Gary N Coleman	16,350,000	0	16,350,000
e142	A Study of Doped Semiconducting Nanowires	EPSRC	Class2a	Materials	Mr Arash A Mostofi	200,000	0	200,000
e143	Numerical Investigation of Jet Noise	EPSRC	Class1a	Engineering	Dr Anurag Agarwal	2	0	2
e144	Numerical Simulation of Rotating Stall and Surge	EPSRC	Class1a	Engineering	Dr Mehdi Vahdati	1	0	1
e145	UK-SHEC Consortium	EPSRC	Class1a	Chemistry	Dr T.J. Mays	1	0	1
e146	G protein-coupled receptor dynamics	EPSRC	Class2a	Life Sciences	Dr Irina Tikhonova	199,680	0	199,680
u02	Materials simulation using AIMPRO	EPSRC	Early use	Materials	Dr Patrick R Briddon	4,000,000	3,080,443	919,557
u03	DNS of NACA-0012 aerofoil at Mach 0.4	EPSRC	Early use	Engineering	Dr Gary N Coleman	2,500,000	2,309,233	190,767
u10	Turbulent Plasma Transport in Tokamaks	EPSRC	Early use	Physics	Dr Colin M Roach	2,500,000	2,538,943	-38,943
y08	Testing	EPSRC	Early use	support	Dr David Jenkins	1,000	0	1,000
NERC Projects								
n01	Global Ocean Modelling Consortium	NERC	Class1a	Environment	Dr Thomas Anderson	19,343,840	16,705,496	2,638,344
n02	NCAS (National Centre for Atmospheric Science)	NERC	Class1a	Environment	Dr Lois Steenman-Clark	121,032,034	87,448,243	33,583,791
n03	Computational Mineral Physics Consortium	NERC	Class1a	Environment	Prof John P Brodholt	150,618,316	126,914,367	23,703,949
n04	Shelf Seas Consortium	NERC	Class1a	Environment	Dr Roger Proctor	39,009,435	26,557,147	12,452,288
u07	NCAS	NERC	Early use	Environment	Dr Lois Steenman-Clark	2,000,000	131,438	1,868,562
BBSRC Projects								
b08	Int BioSim	BBSRC	Class1a	Life Sciences	Mr Mark M Sansom	866,000	909,998	-43,998

Code	Title	Funding Body	Class	Area	PI	Total AUs allocated	AUs used	AUs left
b09	Circadian Clock	BBSRC	Class1a	Materials	Prof Andrew A Millar	2,000,000	1,667	1,998,333
b10	SPRINTing with HECToR [dCSE]	BBSRC	Class2b	Life Sciences	Mr Terry Sloan	400,000	2,709	397,291
b100	Widening the BBSRC HPC User Base	BBSRC	Class1a	Life Sciences	Dr Michael Ball	10,000,000	0	10,000,000
b11	ExeterBioSeq	BBSRC	Class2a	Life Sciences	Prof Richard ffrench-Constant	200,000	0	200,000
STFC Projects								
p01	Atomic Physics for APARC	STFC	Class1a	Physics	Dr Penny Scott	3,020,000	0	3,020,000
External Projects								
x01	HPC-Europa	External	Class1a	External	Dr Judy Hardy	5,046,434	2,140,405	2,906,029
x02	BlueArc (TDS)	External	Service	support	Mr M W Brown	1,000	0	1,000
x03	Prospect FS	External	Class1a	External	Mr Davy Virdee	384,000	0	384,000
x04	Futuretec	External	Class1a	External	Mr Davy Virdee	1,191,200	892,001	299,199
x05	FIOS	External	Class1a	Engineering	Mr Davy Virdee	460,800	13,269	447,531
T01	NIMES: New Improved Muds from Environmental Sources.	External	Class1a	Environment	Dr Chris Greenwell	4,113,669	907,733	3,205,936
Director's Time								
y09	Director's Time	Directors' Time	Service	External	Prof Arthur S Trew	29,685,133	82,538	764,170
d01	UKQCD-DT	Directors' Time	Early use	Physics	Dr Jonathan Flynn	3,968,929	3,968,926	3
d03	EUFORIA	Directors' Time	Service	Physics	Mr Adrian Jackson	2,200,000	1,164,046	1,035,954
d04	MSc Projects	Directors' Time	Service	External	Dr David Henty	93,500	58,151	35,349
d09	ICHEC	Directors' Time	Class2a	Materials	Dr Jean-Christophe Desplat	1,000,000	1,014,928	-14,928
d10	d10	Directors' Time	Service	Physics	Prof Ken Taylor	10,000,000	10,426,109	-426,109
d11	NAIS	Directors' Time	Service	Engineering	Prof Mark Ainsworth	416,667	2,375	414,292
d12	CoE HiGEM	Directors' Time	Service	Environment	Dr Len L C Shaffrey	10,000,000	0	10,000,000

Code	Title	Funding Body	Class	Area	PI	Total AUs allocated	AUs used	AUs left
d13	CoE SENGA	Directors' Time	Service	Engineering	Dr Stewart Cant	10,000,000	0	10,000,000
d14	CoE HiPSTAR	Directors' Time	Service	Engineering	Dr Richard D Sandberg	1,000,000	1,046,153	-46,153
d15	HPC-GAP	Directors' Time	Service	Engineering	Dr David Henty	1,000	51	949