



# *Consortium Visits*

Martyn F. Guest

HPCx Terascaling Team



- **Background**
  - Visits advocated as part of the HPCx Annual plan for 2005 (12 visits in total)
- **Main Drivers**
  - Consortium perspective and input on future directions of the service.
  - Issues affecting capability utilisation
- **Visit Schedule and Current Status**
- **Initial Overview of Findings and Headline issues**

Gain a deeper understanding of consortium activities

- Discuss any issues the consortium may have and how we may help - understand and assist with any barriers the consortium faces in achieving its scientific goals.
- Characterise the capability aspirations of each consortium and the performance potential of their codes.
- Understand consortium's longer term goals and their requirements around e.g. HECToR and future Services.

## 1. Scientific Goals & Research Outcomes:

- HPCx should be more involved in helping the consortium to achieve its scientific goals, and together with EPSRC, promote the associated research outcomes and scientific achievements.
- provide ongoing account of consortium activities through joint HPCx/CSAR "capability Science" web site
- together with the consortium, maintain an updated list of publications, presentations, meetings etc.
  - optimum contact for pursuing this, PI, PDRA, ??
  - associated metrics .....

## 2. Capability and Capacity Usage:

- A number of issues affect the balance, some a function of the service itself that can in principle be addressed by HPCx directly, and others which are consortium-specific. Our aim is to understand the balance here.
- Essential to understand these issues in the way they affect not only HPCx utilisation, but machines such as HECToR and future HEC resources. .
- What is the community perception of the service - does the community consider that the major task of HPCx should be to
  - drive capability usage forward across, realistically, a modest number of consortia, and/or
  - progress more modest levels of parallelism (e.g. 128 - 256 CPUs) from a *far greater number of consortium* and their associated codes?

## 2.1 HPCx-specific Issues:

### A. Code Performance and Scalability.

- Computational demands of the codes used by the consortium
- Constraints to scalability within each code - are these a problem in terms of achieving the scientific goals?
- What are the key codes for future work, both on HPCx and HECToR?

### B. Capability Incentive Scheme & its Effectiveness

- Awareness of the capability incentive scheme.
- Which codes would be eligible for discount?
- Do the current levels of discount appear attractive?
- How would the consortium like the scheme to be changed

## 2.1 HPCx-specific Issues:

### C. Throughput, Job Scheduling & Available Resources

- Are you content with the current throughput of your jobs, and if not, what changes would you like to be made?
- Would you favour additional changes that would promote the throughput of capability jobs at the expense of "capacity" jobs
- Are you hindered by current time limits imposed on all jobs?

### D. On the more general nature of the service

- Are there any pre/post processing issues around HPCx utilisation?
- Are there constraints arising from HPCx technology & the associated machine management - memory, disk space etc?

## 2.2 Consortium-specific Issues

### A. Lack of Resources

- Uptake of original allocation, is this on track?
- Does the sub-group structure of consortium constrain the capability usage?
- manpower constraints; is recruitment/retention of staff an issue?

### B. New Users/Applications

- Understand the progress of "new" consortia towards capability status - is this a goal - are consortium staff working on scalability issues or focusing solely on science deliverables?

### C. Virement Issues & Other Computational Resources

- Possible migration from HPCx to the CSAR SGI Altix.
- Use of other computational resources - local, in/outside the UK

## 2.2 Consortium-specific Issues

### D. Scientific Case & Capability Drivers

- Does the existing consortium agenda require more than 128 CPUs?
- Is there a scientific case for moving into the capability regime?

### E. Timing and the Research Cycle

- Current status of consortium's HPCx project? Is there a new bid in the pipeline? how might HPCx help in formulating such a bid?
- Any plans to move to "new problem" regimes within the project? Would the new problem sizes create new capability demands?

## 3 Other Issues and Actions:

- **Training**
  - interest in tailored, application-specific training and/or workshops?
  - Is there a greater incentive for training at the consortium site(s)?

	Consortium		Visit Date	Venue	HPCx Staff
1.	e05	Materials Chemistry	5th April, 3rd May	Royal Institution	I.J. Bush, M. Plummer M.F. Guest, P. Sherwood
2.	e05	Materials Chemistry	11th May	Royal Institution	K. D'Mellow, G. Pringle, P. Sherwood, M.F. Guest
3.	e04	Chemreact Computing	23 <sup>rd</sup> March	UCL	A.G. Sunderland, M.F. Guest
4.	e03	Multi-photon, electron collisions and BEC	27th April	QUB	A.G. Sunderland, M.F. Guest, P. Sherwood
5.	n04	Shelf Seas	22nd July	POL	M. Ashworth
6.	n01	Large-Scale Long-Term Ocean Circulation	11th May	NOC, Southampton	L. Smith, M. Ashworth,
7.	n02	NERC Centres for Atmospheric Science (NCAS)	7th February	University of Reading	L. Smith and M. Bull
8.	e17	Integrative Biology	18th May	RAL	D. Henty, P. Sherwood, L. Smith and M. Bull
9.	e06	UK Car-Parrinello	27th April	QUB	M. Plummer, M.F. Guest, P. Sherwood
10.	e07	Turbulent Plasma Transport in Tokamaks	4 <sup>th</sup> April	Culham	L. Smith, J. Hein and M.F. Guest

- **Interim snapshot based on 10 visits**
  - plus relevant information from the EPSRC HPC visits conducted at the same time by Emma Jones (EPSRC) plus MFG, AGS, IJB etc
- **Capability and Capacity Usage:**
  - The community response to the balance
- **HPCx-specific Issues:**
  - Throughput, Job Scheduling and available resources
  - Other Aspects of the Service
  - Code Performance and Scalability.
- **Consortium-specific Issues**

1. The majority of the work of most HPCx consortia is capacity, and not capability. e03 is the only exception.
2. Many argued that the focus on capability jobs was damaging given the associated handicap imposed on capacity jobs; many requests for an environment that did not penalise capacity jobs:
  - driven by a separate queue mechanism but developed in a way that did not prejudice the capability mission of the service itself.
  - Capacity & capability partitions would share the same file systems & software environment, thus fostering the migration from capacity to capability while providing a closely coupled environment for pre- and post-analysis of capability jobs.

3. Numerous capacity jobs often required to prepare for capability calculations - challenged assumption that capacity jobs would be met through local clusters or mechanisms such as NGS.
4. Closure of the CSAR service & the migration of the associated consortia to HPCx will bring more capacity-based work to HPCx, highlighting the issues raised above.

5. Main criticism of the HPCx was the Memory per processor. This represented a major constraint for many consortia (e03, e04, e05, e06), preventing many calculations. It was repeatedly pointed out that this is significantly less than typical mid-range machines. The Phase2A upgrade will do much to address this criticism.
6. A close second was the poor turnaround for long, 32 processor (or fewer) jobs - "the queues are bad" was repeated on almost every visit. Many of these problems arise from the difficulties to scale codes to 100s of processors, so that jobs are not big enough to be favoured by the scheduler (e03,e01). This constrained the progress of associated capability jobs.

7. Many users would like the existing time limit on all jobs to be increased to at least 24-hours (e05, e04, e07).
8. Many of the consortia were unaware of the incentive scheme for capability codes.
  - Once informed all were interested in using the discounts
  - Current scheme should be broadened to include "odd" CPU counts (n01), and enhanced turnaround as an incentive (e03).
  - Concern that any changes to the capability scheme or scheduling algorithms that reduced the turn around time on 128 processors would be a real problem for users (n02).
9. Website advice about the best number of processors to use with "community codes" for a given size of job.
10. Concern that consortia (e17) who don't have HPC background & experience face too steep a learning curve

11. There were repeated requests for visualisation support in a variety of formats.

- e03 - The lack of a visualisation facility on HPCx causes delays in shipping large data sets either to local facilities or to CSAR.
- e07 - the lack of onboard visualisation capability forces the consortium to download very large files for post-processing on local machines
- e05 - establishing a 'Visualizers Advice' page? Pre- and post-processing codes for preparing and analyzing data files for various materials applications.
- e17 - Users have a requirement to link to high end visualisation services from HPC. The present lack of this capability severely limits their ability to undertake real time exploration and the ability to computationally steer the simulation

12. A number of codes were identified where performance was clearly limiting capability usage:
  - e05 - VASP and Chem-Shell (The initial round of visits left a trail of follow-up actions / meetings, some of which have already taken place.
  - e04 - The PDVR3DRZ code (TAMPA Group) limited by the scalability of ScaLAPACK D&C eigensolver. Subsequently resolved, with the code now in the "Gold Star" category. MOLPRO remains a problem.
  - e07 - CENTORI is presently undergoing major development (EPCC-Culham collaboration) to improve its potential for capability jobs.
13. In some cases constraints to optimising code (e.g. n02 & the UM; e20 - commercial engineering codes, CFX) were outside HPCx control, confining consortia to capacity utilisation and the associated delays on job turn-around.
14. A few instances where the lack of standard numerical libraries such as NAG delayed code development (e03).

## 15. Manpower Constraints

- Code development vs. code exploitation:
  - e05 - too few academic positions for code developers / researchers. Collaborative S/W development e.g., CASTEP, discouraged by system since developers don't get sole credit, and thus appear less competitive for jobs.
  - e03 - difficult to get applications through peer review if one referee is very negative because of lack of appreciation of what code development involves.
  - e01, e04 - Timing of funding for people & computer time - have to succeed at peer review twice.
- e03 - Insufficient appreciation of HPC when it comes to subject (Physics) specific decisions in award of e.g. EPSRC Advanced Fellowships.
- n04 - Staff time (at POL) is main constraint; n01 - loss of key staff has been an issue.
- e17 - ability to exploit capability computing is affected by lack of manpower and expertise to re-architecture codes

## 16. Licences for Commercial Software:

- e20 - Who pays for access to commercial CFD codes? One group could pay for a licence and use the code on HPCx, but other groups who don't have a licence may also want to use it on HPCx - who is responsible for it? Companies would want all the groups to pay for a licence.

## 17. Peer Review Process:

- A better peer review mechanism is needed. The e20 consortium proposal had to be resubmitted because of 1 very critical referee's report so the application process took a whole year.
  - Referees and panels don't always fully understand HPC and don't read the guidelines for refereeing HPC proposals therefore make uninformed critical comments.
  - Should there be a panel just for allocating HPC time?

- Background and current status of the HPCx Consortium Visits
  - 10 of the 12 visits advocated as part of the HPCx Annual plan for 2005 now complete
  - Main Drivers: Consortium perspective and input on future directions of the service; Issues affecting capability utilisation.
- Initial Overview of Findings and Headline issues.
- Consideration of the key issues and recommended actions on completion of the current round of visits in consultation with STAC, HPCx's Oversight Committee and EPSRC