The HECToR Service

- an update

22 April 2009

Professor Arthur Trew
Service Director

HECToR

UOE HPCX Ltd
HPCX

RESEARCH COUNCILS UK

Annual utilisation

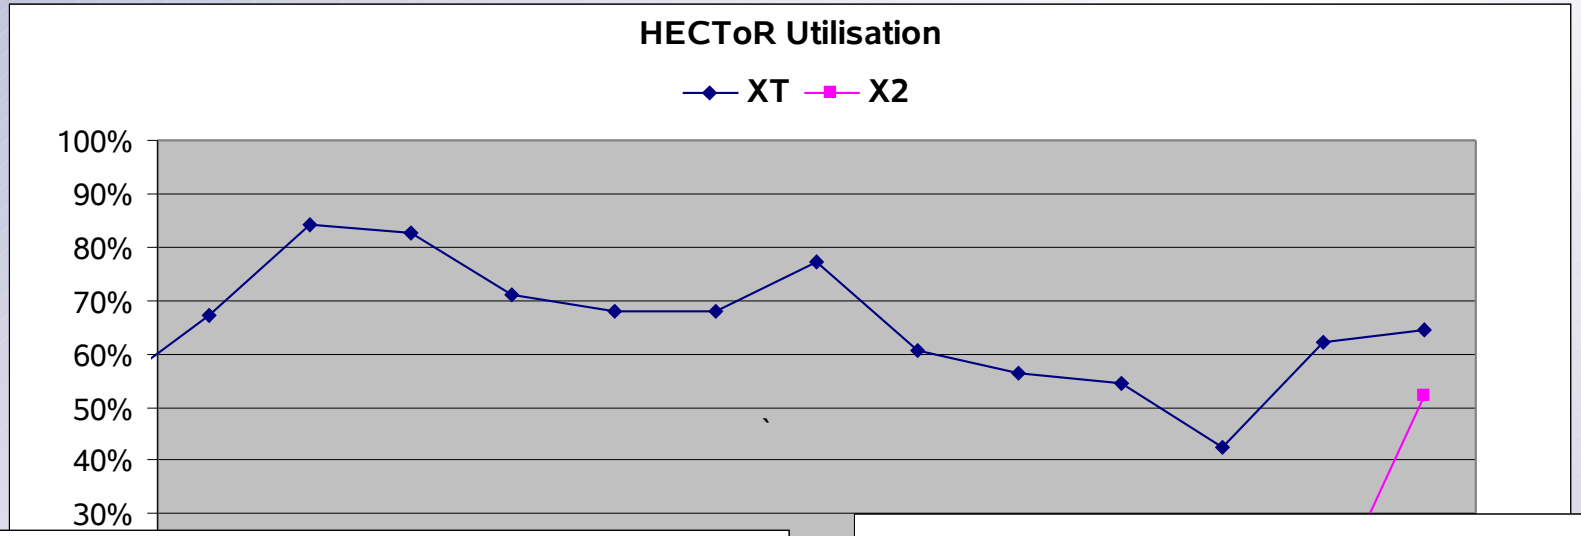

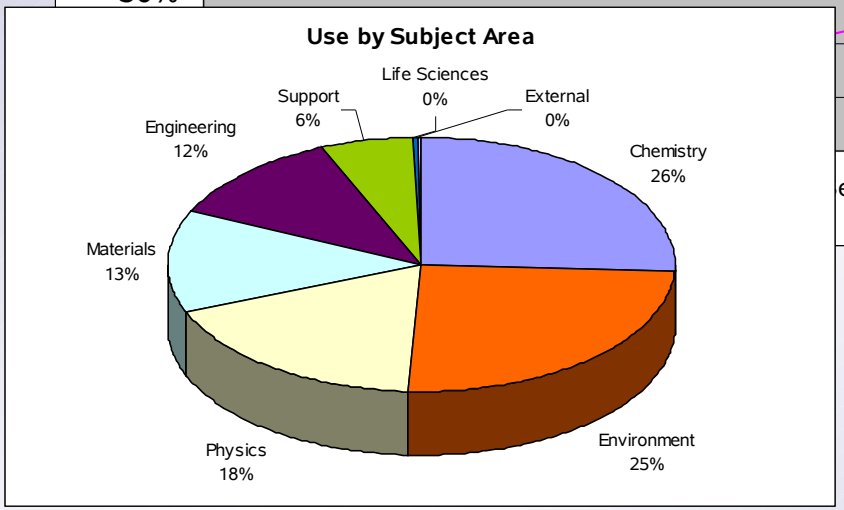

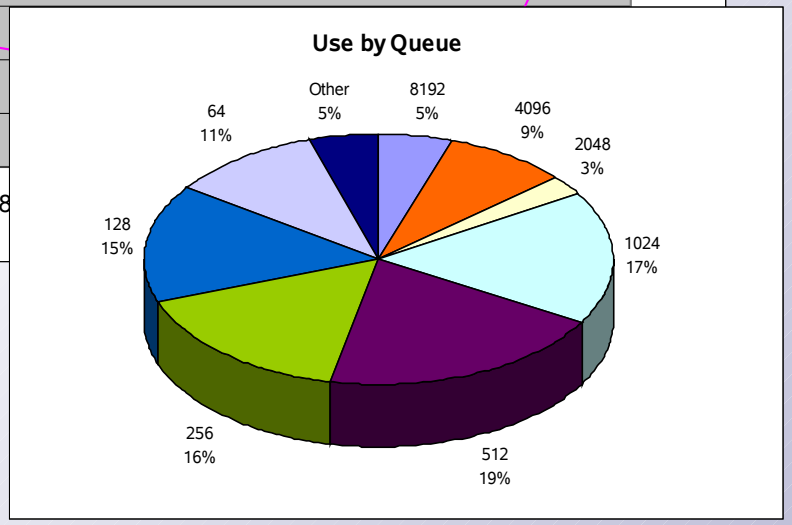

HECToR

UOE HPCX Ltd HPCX

RESEARCH COUNCILS UK

Service Improvements

- Lustre Disk Quotas
  - A fix for the long-standing Lustre 4Tb disk quota limitation was implemented in early October
- HECToR completed conformance tests for joining NGS
  - HECToR is now an NGS Partner
- Job Reservations/Interactive Usage
  - PBS V10 currently under test – available late-May
- Single-processor queues
  - 2 logins nodes allocated, max 4 jobs/node
- CNL 2.1
  - install late-May, required to fix Lustre problems
  - but all codes need recompile

The BIG challenge

- retrofit some form of backup/archive system
  - for little/no money
  - while maintaining the service and data integrity

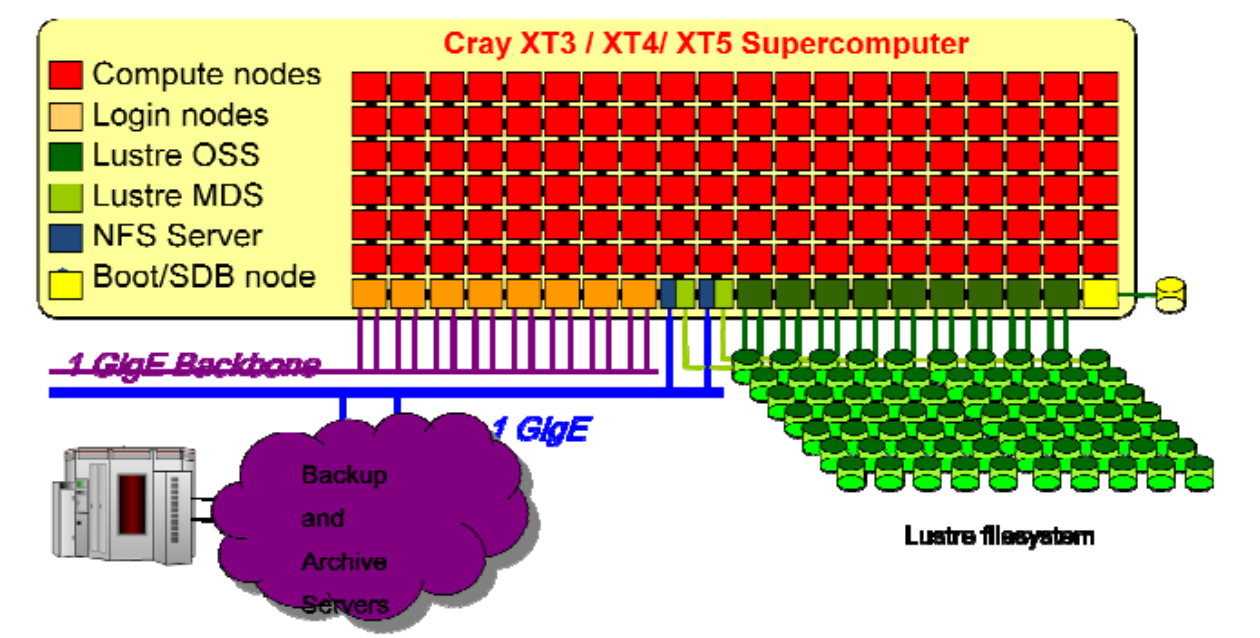

Aim

- provide similar archive functionality as on HPCx

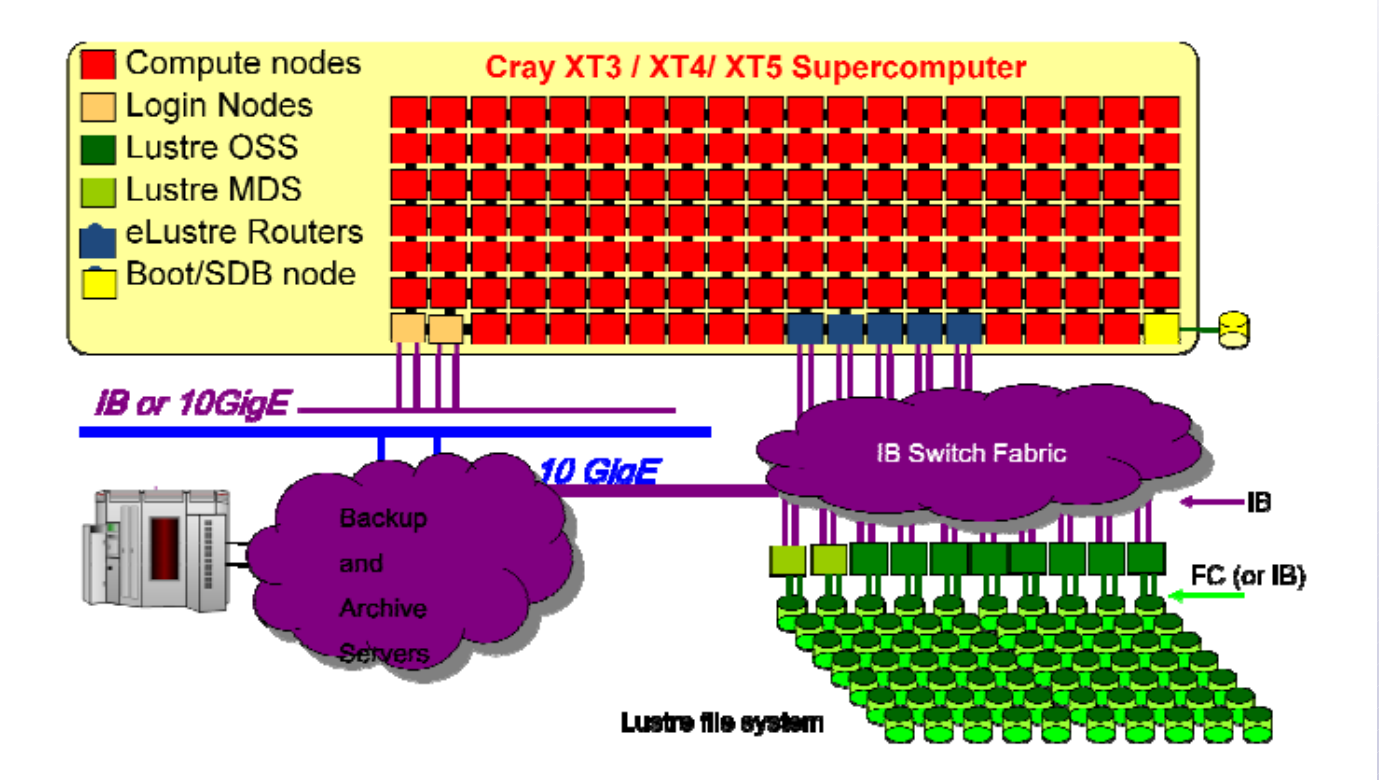

  - additional switch hardware & storage
  - new software to manage file transfers

Where are we?

- our proposal splits the problem into two:
  1. metadata mirroring
     - to (ideally new) disk
     - provides a near real time copy of the primary metadata in case of corruption or unavailability
  2. file archive
     - uses external servers (esDM's)
     - to archive data to tape
     - under test at the moment … the question is can it sustain the required I/O rate?
- user software also under evaluation
- … a solution is still months away

Phase II

- Proposal to upgrade in two Phases
  - Phase IIa: Q2 2009
    - Dual core -> Quad core
      - 200 Tflops peak
    - memory grows from 6GB/node to 8GB/node
    - X2, interconnect, filesystem etc … unaltered
  - Phase IIb: Q3/4 2010
    - XT4 re-sized to 16 cabinets (60 Tflops)
    - “Baker” system installed (360 Tflops)
    - network upgraded with “Gemini”

Phase IIa: when?

- parts on order … their delivery is critical
- two-stage upgrade

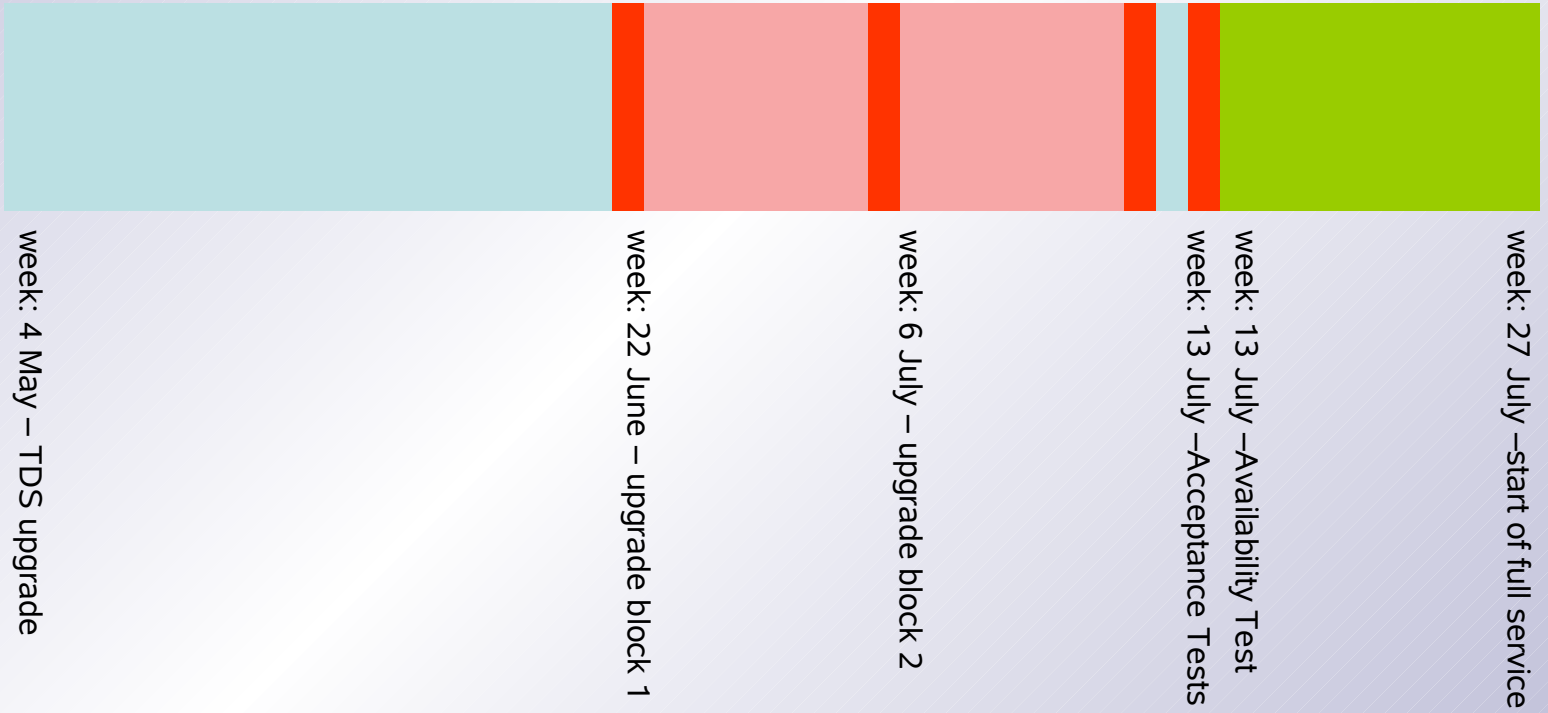

HECToR

UOE HPCX Ltd HPCX

RESEARCH COUNCILS UK

the end of the world …

- ever-increasing clock speed is over
  - 2.8 GHz -> 2.3 GHz Phase IIa
- … can't simply expect the hardware to bail out applications
- multicore *is* the future

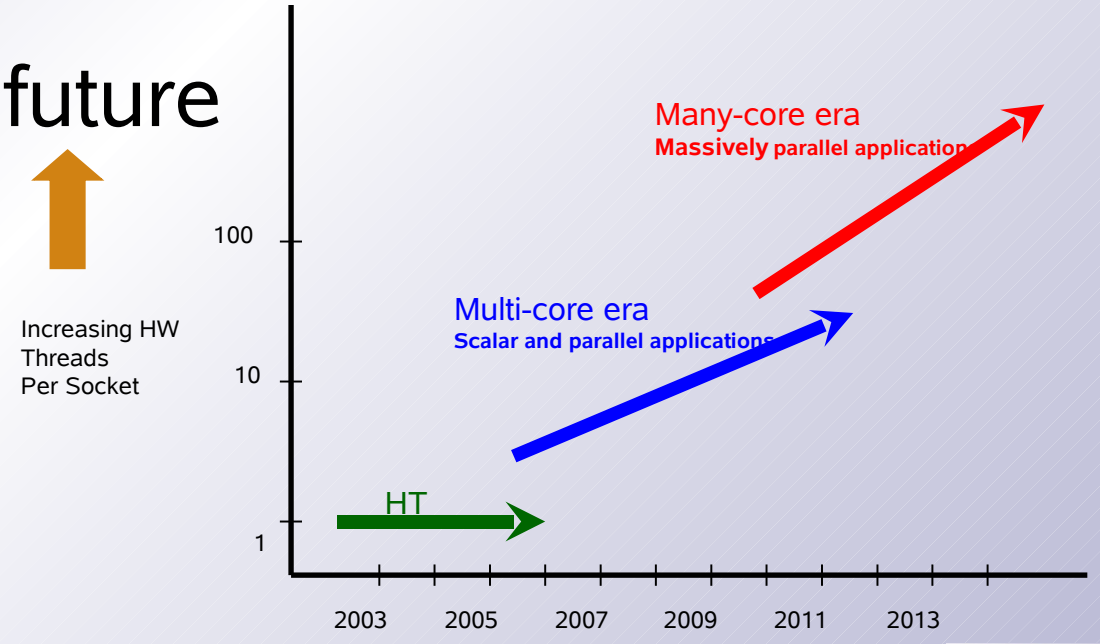

HECToR

UOE HPCX Ltd HPCX

RESEARCH COUNCILS UK

scalability

- scalability is *the* holy grail
- "… but how do I find it?"
- within the HECToR project we have:
  - training & support from NAG
  - opportunity for DCSE projects
  - support from the Cray Centre of Excellence
    - first workshops early in June

NAIS Centre for Numerical Algorithms & Intelligent Software

- Science and Innovation centre
  - Maths at Edinburgh, Strathclyde & Heriot-Watt
  - Informatics at Edinburgh
  - EPCC
- mission: new generation of numerical algorithms
  - written to be scalable & efficient
  - write-once, run anywhere (efficiently)
- looking for applications' collaborations

HECToR

UOE HPCX Ltd HPCX

RESEARCH COUNCILS UK