Reflections on a career in computing, from teletypes and card decks to a modern University Research Computing service

or
A random(-ish) walk through IT

Paul Hatton
Honorary Senior Research Fellow
The University of Birmingham, U.K.
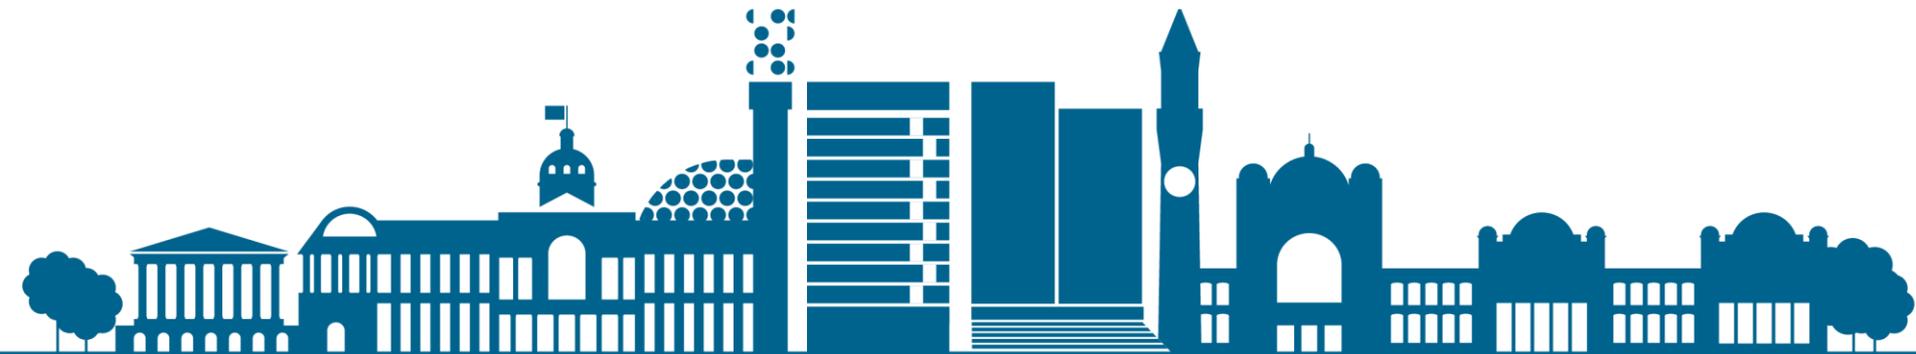
Brief Career History

☐ 1972 - Tried maths at Cambridge

☐ 1972 - Changed to Engineering (↰ #1)

☐ 1975 – Diploma in Computer Science

☐ 1976 – Postgraduate work at Whittle Labs (↰ #2)

☐ 1980 – Rolls-Royce Aerospace, Derby (↰ #3)

☐ 1987 – University of Birmingham, Chemical Engineering (↰ #4)

☐ 1990 – University of Birmingham, Academic Computing Service (↰ #5, 6 …..)

Diploma in Computer Science (1975)

☐ Weekly IBM360 system maintenance (Wednesday afternoons)

Whittle Labs, Cambridge (1976)

Initially mathematical modelling of fan blades ......
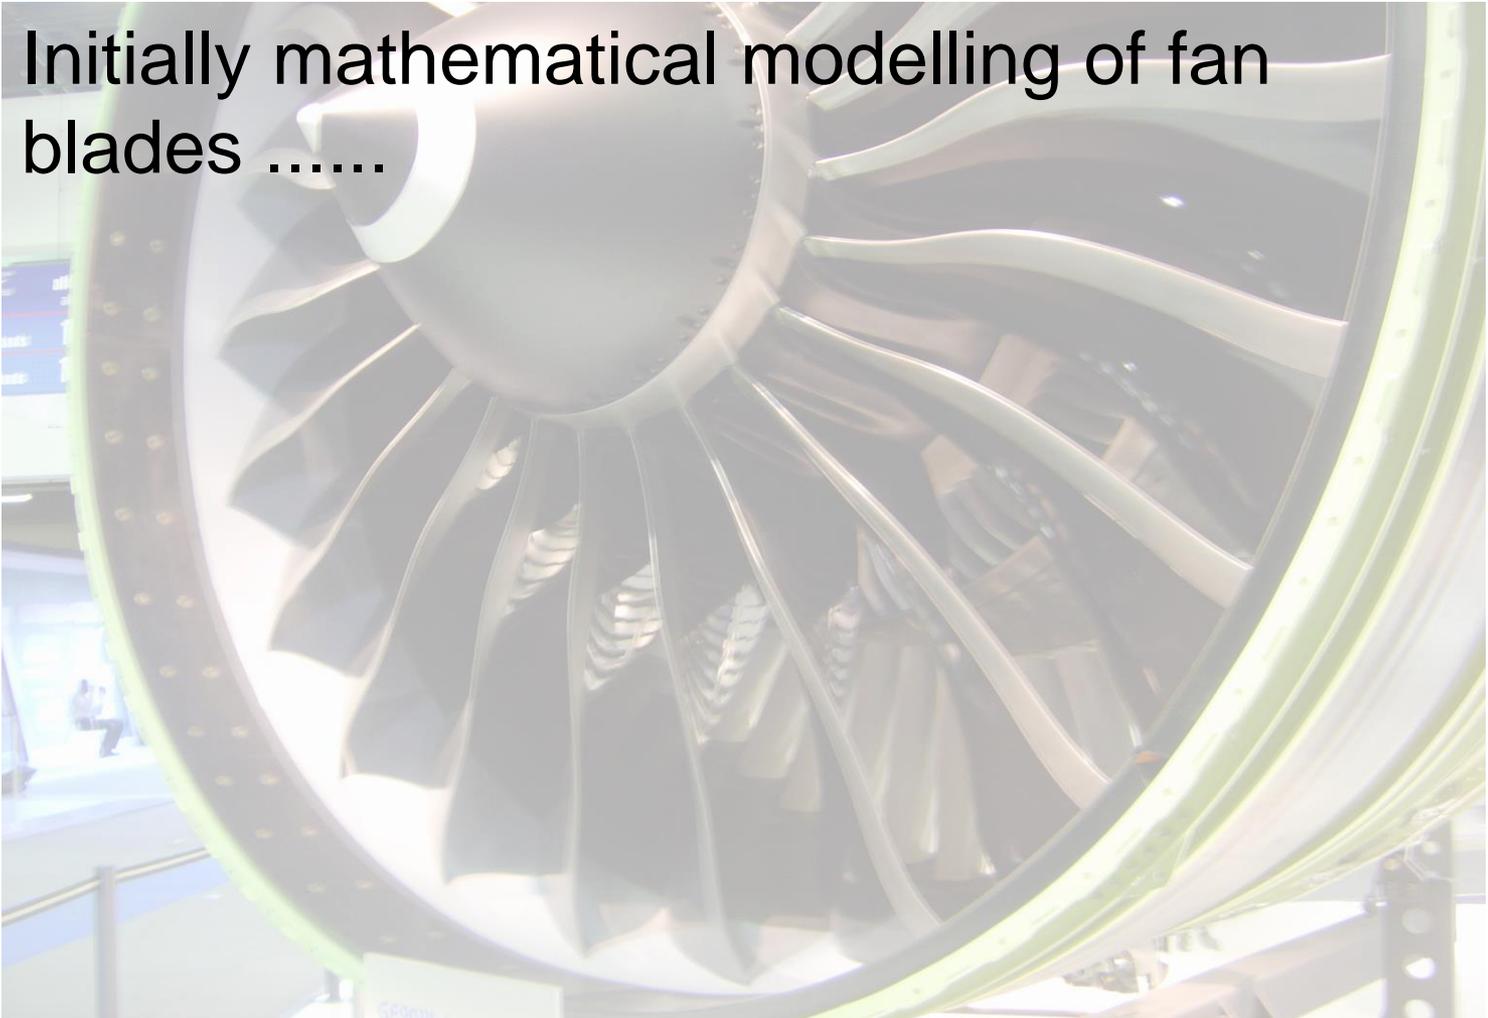
Whittle Labs, Cambridge (1976)

☐ but computers seemed far more interesting:

Phoenix
(TSO that worked)

> HELP GOD
**Deities must be invoked directly and not via Phoenix MVS.**

Rolls-Royce, Derby (1980)

Initially mathematical modelling of fan blades ......
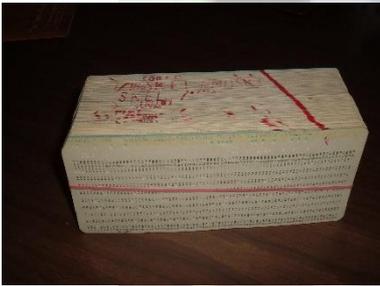
Rolls-Royce, Derby (1980)

☐ Moved to rig testing – real-time data acquisition from compressor rigs
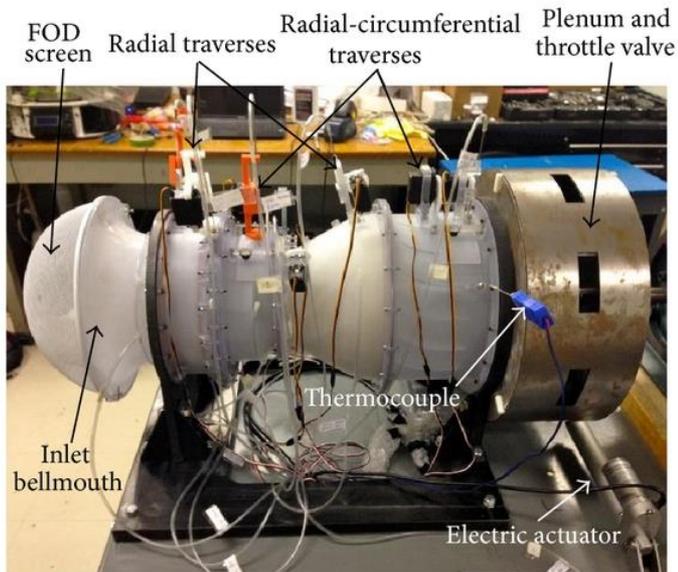
University of Birmingham – Chemical Engineering (1987)

buoyancy-driven 2-phase flow

Combined mathematical modelling, novel instrumentation, real-time data acquisition

University of Birmingham – join Academic Computing Service (1990)

All down to

or, more specifically

*The research grant application process.*

Nostalgia's not what it used to be (#1)

It all started with an IBM 3090 back in 1990

Nostalgia's not what it used to be (#2)

It all started with an IBM 3090 back in 1990
... with a vector facility
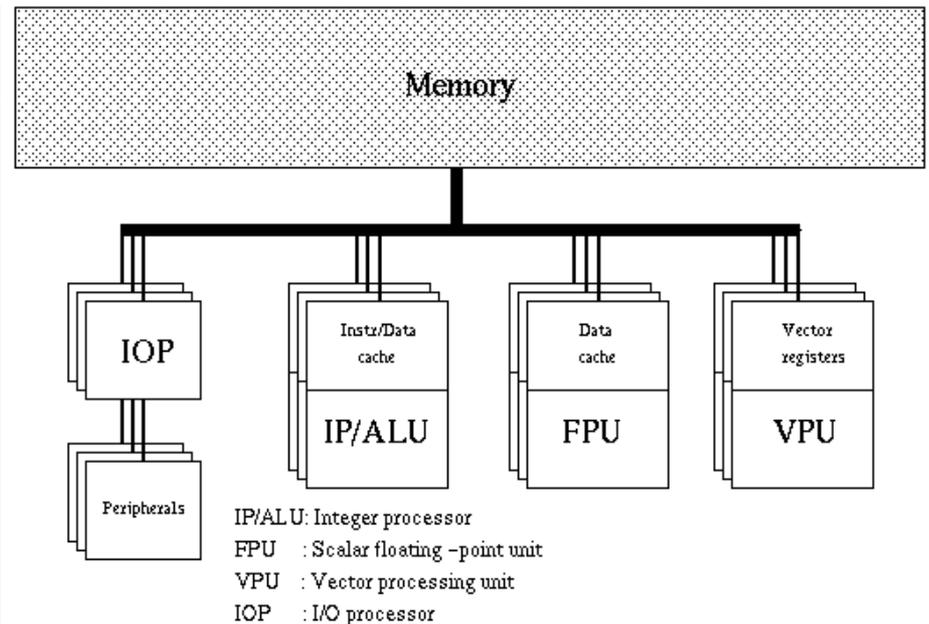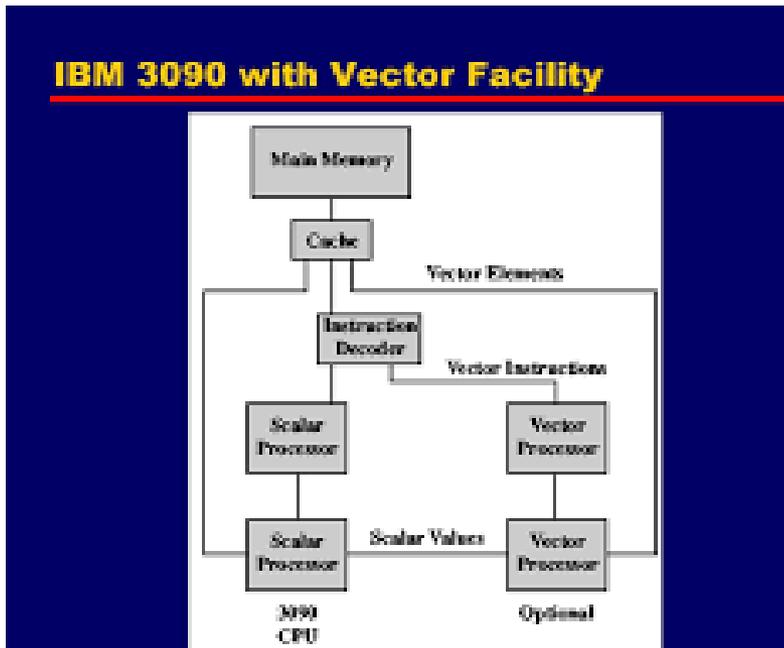
Nostalgia's not what it used to be (#3)

It all started with an IBM 3090 back in 1990
... with a vector facility ... and HSM ...

COMPUTERWORLD
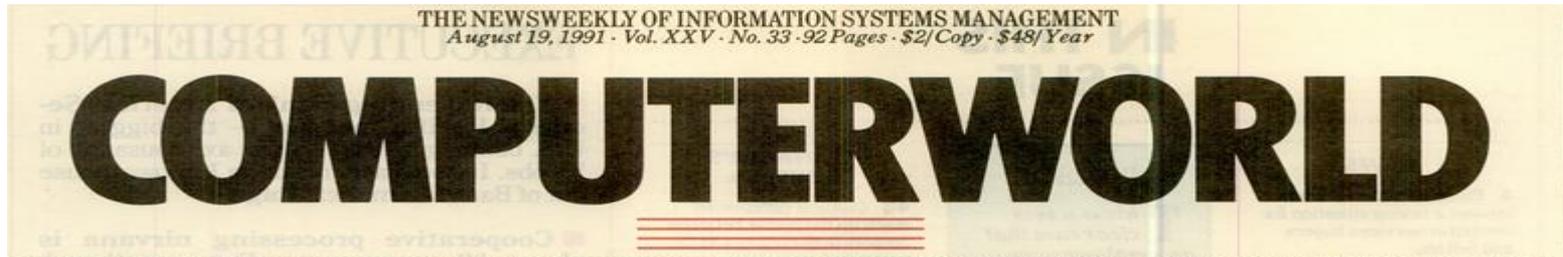
That was prior to 1987, when the company started using IBM's Hierarchical Storage Manager (HSM) for storage pooling. Since then, "we've improved disk utilization by about 15%," says Hansen, who oversees three job entry subsystem complexes of IBM 3090 Class 600 and higher machines.

In early 1990, Hansen boosted efficiency even further by adding IBM's System Managed Storage (SMS — see story page 51). With these products, Ford has been able to freeze — and even decrease — its magnetic disk capacity.

Nostalgia's not what it used to be (#4)

It all started with an IBM 3090 back in 1990
... with a vector facility ... and HSM ...
… and real disc drives …
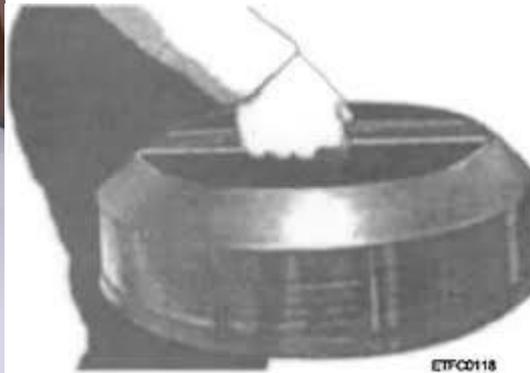
Times change ....
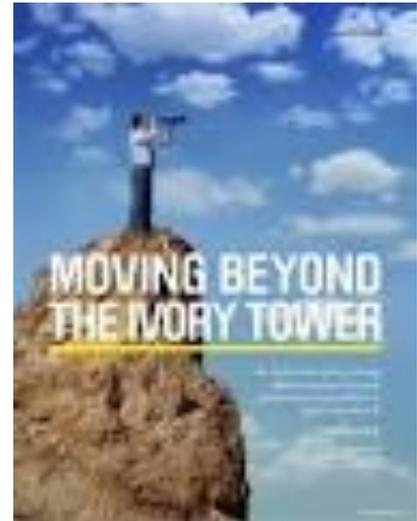
A little more nostalgia (1995)

- 2 x Compaq Alphaservers
- ISDU (Information Service Digital Unix) service
- VAX/VMS (we had to learn a new language)
- One for interactive work (isdugp, EV5 processors), one for batch Compute Intensive (isduci, upgraded to EV6 processors)
- We ended up as a bleeding-edge site for Compaq's cluster filestore
- isdugp service OK, isduci service unreliable due to flaky filestore
- no demand for parallel processing

More recently (early-2003)

- Funded from SRIF-1
- HP Central Applications Service (CAPPS)
- 6 x HP J6700 servers each with 2 x PA-8700 processors at 750 MHz, max 16 GB/server
- HP/UX (we had to learn a new language)
- Funded as a general applications service, not specialist HPC
- Added researcher-funded (Archaeology) 1TB filestore, at great expense to the group – our first example of joint IT Services/researcher funding
- Unreliable filestore – basic NFS mounts to servers and IT Services SAN
- *Users liked it when it worked but constant, and justified, complaints about poor availability*
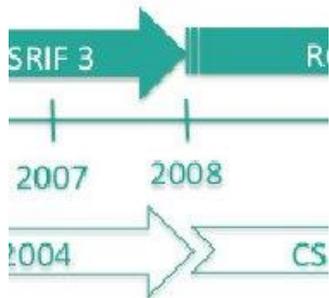
SRIF-3 (2005)

- We don't need to discuss the limitations of the SRIF funding model (do we)?
- Our first foray into 'proper' HPC (2006)
- My first foray into 'proper' procurement
- My first foray into the HPC-SIG
- Tried to be conservative, not bleeding-edge; we needed to provide a stable service
  - Selected dual-core AMD Opteron processors – quad-core processors were overdue and considered too great a risk by us
  - Selected GPFS as a supported cluster filesystem
- Scientific Linux – we had to learn a new dialect (from HP-UX), not a new language
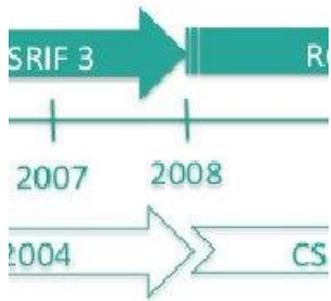
SRIF-3

Didn't work out as hoped:

- No funding released by Birmingham for additional staff
- GPFS flaky – poor performance and reliability - and much buck-passing between integrators and IBM
- As part of the procurement, integrators placed a specialist on site, who showed little interest in providing a stable service but loved tinkering
- *Users liked it when it worked but constant, and justified, complaints about poor availability*
- We couldn't carry on like this ......

a BEAR is conceived ... (2008-ish)

We had to go back to basics since:

- [ ] We had a poor track record in providing a stable HPC or applications service
- [ ] We had little support from the majority of the researchers
- [ ] Many groups were running their own clusters with little incentive from the University to change
- [ ] The SRIF-3 cluster was heading towards end-of-life and there was little enthusiasm from the University to fund anything more than a nominal replacement; any additional staff posts were excluded from this nominal funding
- [ ] etc. ....

We had a novel idea of actually asking researchers what they needed, rather than us giving them what we thought they wanted

What do researchers do?

Some or all of:

- ☐ Chase funding
- ☐ Learn the tools of the trade
- ☐ Collaborate
- ☐ Generate and consume data
- ☐ Compute
  - – Other project-specific resources
- ☐ Interpret
- ☐ Disseminate

And, of course, the broader context
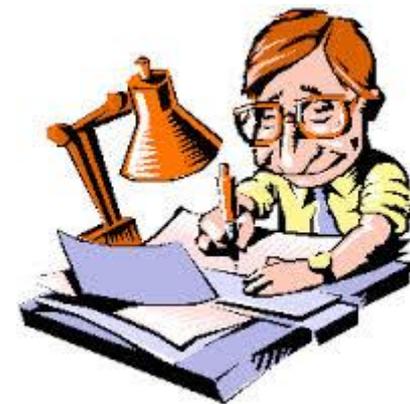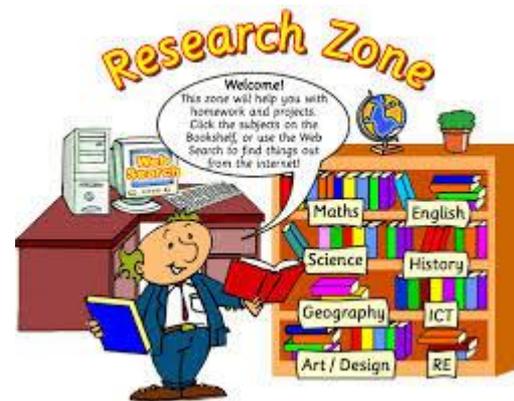
What are IT Services expected to do?

Support a wide range of researchers with different expertise and requirements

Lack of Linux expertise does not mean that research projects are not worthy of centrally provided resources
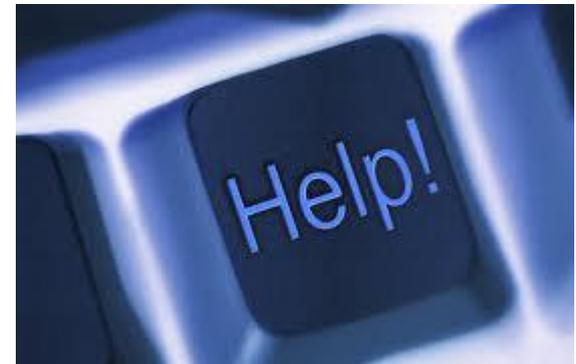
BEAR (Birmingham Environment for Academic Research)

- ☐ In American Indian mythology, the Bear is the font of all wisdom and the source of all knowledge
- ☐ HPC cluster named BlueBEAR, the spiritual original of the Grizzly Bear
- ☐ The vision was to establish a Research Environment with multiple complimentary services .....

BEAR Mk 1
(a BEAR cub)

- ☐ University released (modest) 4 year recurrent University funding from 2012
- ☐ We went to Procurement for BEAR, not solely - or even primarily - for an HPC service
- ☐ The procurement made it clear that this was not a technology demonstrator but was for a stable and reliable service with a partner, not just a supplier
  - – Partnership weighted more heavily than raw hardware in the procurement
- ☐ Set up as a partnership with integrator, Tier-1, visualisation (local and remote) and conferencing partners
- ☐ We started to win back the doubters and build a reputation on campus as a trusted service provider

BEAR Mk 1
(a BEAR cub)

At this time we could offer:

☐ Linux-based HPC (on a good day)

☐ large (at the time) storage for HPC
  – but only local to HPC
  – data had to be transferred to/from HPC with sftp/WinSCP

☐ visualisation centre with active stereo and head tracking

☐ scalable collaborative conferencing
  – IOCOM Visimeet, based on Access Grid
  – think Skype on steroids

☐ **a developing vision based on experience so far**

What does a BlueBEAR cluster look like?

BEAR developments (a growing-up BEAR)

By 2016 the existing BlueBEAR cluster was heading towards end-of-life so another procurement:

- [ ] With the new-found support from the research community the University released more recurrent funds
- [ ] We needed a long-term partnership to develop the existing services
- [ ] Went to procurement for an 8-year Research Computing framework (4+2+2) years to give a long-term relationship
  - Getting this through procurement was quite a challenge
- [ ] Continuity with existing services was essential
- [ ] Resulted in a strengthening of our existing partnerships to give a firm basis for future developments

It seemed a good idea at the time ....

- ☐ Windows HPC
- ☐ Hadoop
- ☐ Batch render service (Blender/Maya/3ds Max)
- ☐ .... And probably others

BEAR
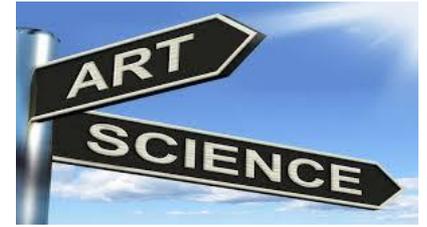
The University expects the services to serve a wide range of users, not just the usual HPC disciplines.

Future plans

Any questions?