UNIVERSITY OF BIRMINGHAM | BEAR
BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

Installing and curating software for heterogeneous compute environments

Dr Andrew Edmondson

Advanced Research Computing, University of Birmingham

BEAR services


BlueBEAR HPC


BEAR Research Data Store


BEAR Cloud


CaStLeS (powered by BEAR)


BEAR Gitlab


BEAR Data Transfer


BEAR Archive


BEAR Research Data Network

BEAR Software


Exclusive BEAR


BEAR Database Service

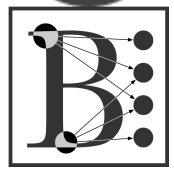
Baskerville Tier 2 HPC

https://intranet.birmingham.ac.uk/bear/services


UNIVERSITY OF BIRMINGHAM | BEAR BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

CPU and OS

2012
- BlueBEAR 2 – Sandy Bridge – SL6
- Single, homogeneous cluster

2016
- Added BlueBEAR Haswell – EL7 – heterogeneous cluster
- Launched BEAR Cloud – Haswell – EL7 / U1604

2017
- Broadwell – EL7
- NVIDIA P100s

CPU and OS

2018
- POWER9 – EL7
- NVIDIA V100s

2019
- Cascade Lake – EL7

2020
- Upgrade EL7 => EL8 and U1604 => U2004
- Decommissioned Sandy Bridge nodes

CPU and OS

2021
- BlueBEAR Ice Lake – EL8 – NVIDIA A30 and A100
- Baskerville Tier 2

2022
- ?

2023
- ?

Baskerville

- 46 Ice Lake nodes with 512G RAM and quad NVIDIA A100 GPUs
- Total of 184 A100 GPUs
- Novel accelerator technologies:
  - NextSilicon ExCALIBUR project on Baskerville - https://excalibur.ac.uk/projects/novel-hardware-software-architecture-testbed/
  - More planned to come

- https://www.baskerville.ac.uk/
- Come and see us in the Research Zone!

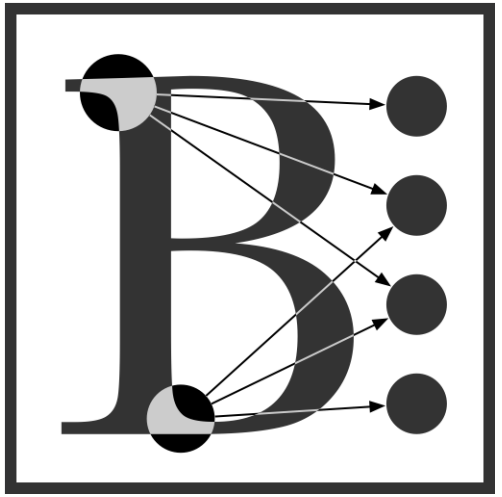

UNIVERSITY OF BIRMINGHAM | BEAR BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

Simon Thompson

Our infrastructure story from a homogeneous Sandy Bridge cluster to a multi-architecture, multi-OS suite of services: “The Thompson Era”

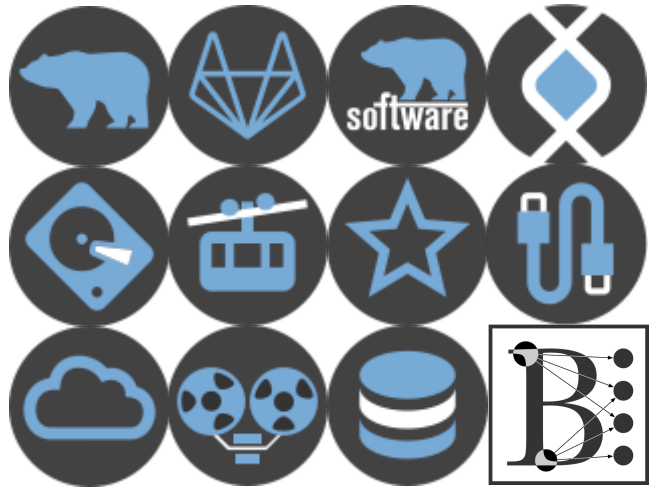

Who will be next?

https://www.jobs.ac.uk/job/CKY408/research-computing-infrastructure-architect

Closing date 15th December

UNIVERSITY OF BIRMINGHAM | BEAR BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

Users of BEAR's compute services

- Wide variety of researchers
  - Traditional HPC users and non-traditional
  - Staff, PhD, PGT, UG

- Slurm-based BlueBEAR HPC (SSH)
- OpenStack BEAR Cloud VMs (SSH/x2go)
- Open OnDemand BEAR Portal
  - Jupyter/Rstudio/MATLAB/…
  - BlueBEAR GUI

UNIVERSITY OF BIRMINGHAM | BEAR BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

The challenge we face

- Wide variety of people using BEAR's compute services
- 1,181 unique applications/libraries
- Six different CPU/OS combinations (today)
- 3,122 distinct modules
- Upgrade OS => rebuild software…
- Add new CPU or GPU => rebuild software…
- And new software is being requested all the time…

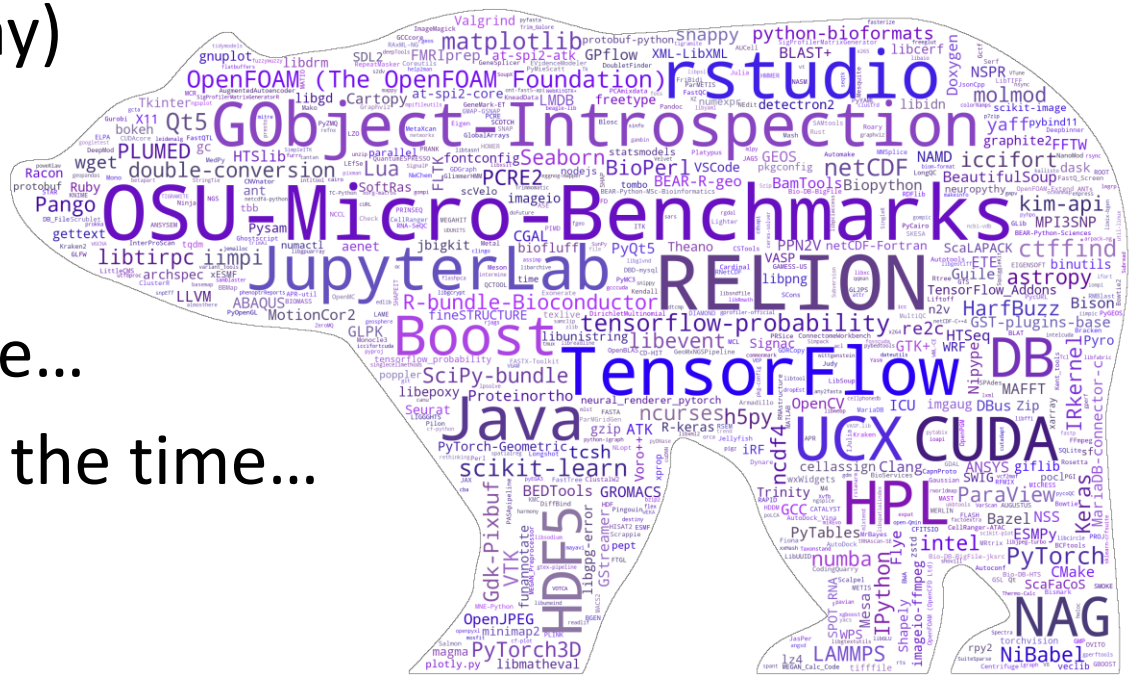

UNIVERSITY OF BIRMINGHAM | BEAR
BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

Installing and curating software

- The rest of this talk:
  - EasyBuild: Installing software efficiently
  - Documenting what's installed
  - ReFrame: Testing and finding problems
  - Monitoring and visualising usage
  - Deprecating and removing old software
  - Problems of a changing environment

EasyBuild

"EasyBuild is a software build and installation framework that allows you to manage (scientific) software on High Performance Computing (HPC) systems in an efficient way."

https://easybuilders.github.io/easybuild/

UNIVERSITY OF BIRMINGHAM | BEAR BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

EasyBuild

- We started in 2016 – when Haswell arrived
- We contributed lots upstream
  - POWER9 support
  - New applications/versions
  - Testing
  - Dr Simon Branford is now an official EasyBuild Maintainer
- We've benefited hugely from other people's work
- See https://github.com/bear-rsg/easybuild-easyconfigs

EasyBuild

- REALLY helpful for installing the same thing for multiple arch/OS
- E.g. our 2019a environment has:
  i. EL7-cascadelake
  ii. EL7-haswell
  iii. EL7-power9
  iv. EL7-sandybridge
  v. Ubuntu16.04-haswell
  vi. EL8-cas
  vii. EL8-has
  viii. EL8-p9
  ix. U20-has

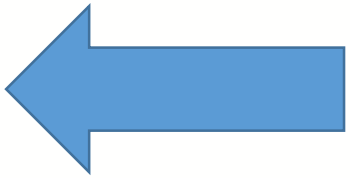

Upgraded OS
Changed naming convention
Sandy Bridge decommissioned

Documentation

- We used to maintain our user-facing docs manually…
  - Out of date
  - Time consuming

- Now we use https://bear-apps.bham.ac.uk
  - Django-based website
  - API for generating docs from modules
  - “Spend time to save time…”

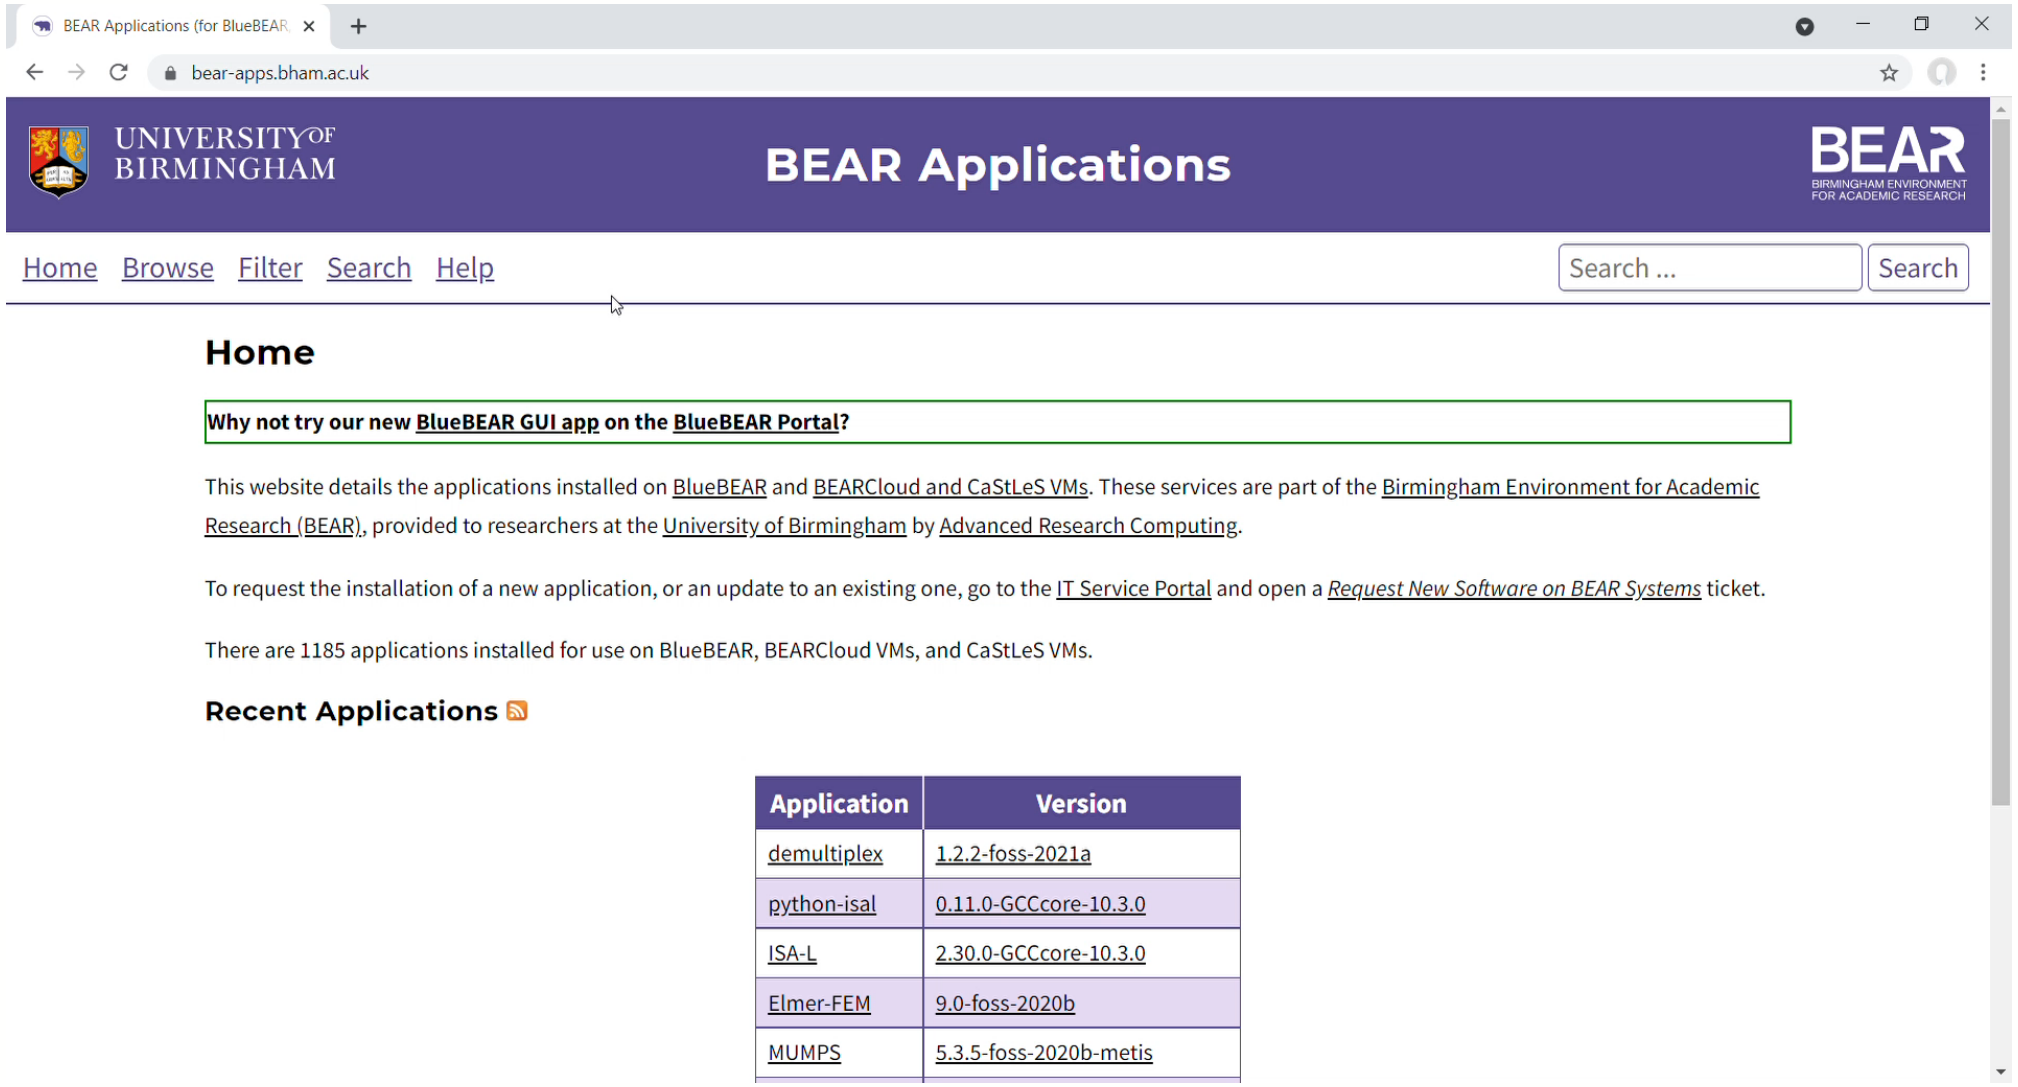
BEAR Applications (for BlueBEAR
bear-apps.bham.ac.uk
UNIVERSITY OF BIRMINGHAM
BEAR Applications
BEAR
BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH
Home Browse Filter Search Help
Search ... Search
Home
Why not try our new BlueBEAR GUI app on the BlueBEAR Portal?
This website details the applications installed on BlueBEAR and BEARCloud and CaStLeS VMs. These services are part of the Birmingham Environment for Academic Research (BEAR), provided to researchers at the University of Birmingham by Advanced Research Computing.
To request the installation of a new application, or an update to an existing one, go to the IT Service Portal and open a Request New Software on BEAR Systems ticket.
There are 1185 applications installed for use on BlueBEAR, BEARCloud VMs, and CaStLeS VMs.
Recent Applications
Application
Version
demultiplex
1.2.2-foss-2021a
python-isal
0.11.0-GCCcore-10.3.0
ISA-L
2.30.0-GCCcore-10.3.0
Elmer-FEM
9.0-foss-2020b
MUMPS
5.3.5-foss-2020b-metis

ReFrame – HPC testing

"ReFrame is a powerful framework for writing system regression tests and benchmarks, specifically targeted to HPC systems…"

https://reframe-hpc.readthedocs.io

- Nightly tests – can detect all kinds of problems with software, systems, filesystems, package updates, ...
- Single-node stress tests
- Methods to generate a single test that runs on all subsets in a group of nodes – e.g. to detect MPI/networking problems.

Monitoring usage

- Imagine you're upgrading the OS, and you know lots of software won't work any more (due to major glibc changes, for example)…
  - Do you have to rebuild and fix everything?
- Imagine you've been running the same cluster for 10 years
  - Do you have to keep supporting all the old software?
- What are people using? Do we want to ask them to change? Can we deprecate old software? Can we delete old software?

Monitoring usage

- We wrote a Django-based website (for admin use)
  - Log all module loads to disk (using module headers)
  - Ingest all module loads into the database via cron
  - “Spend time to save time…”

Modules DB

bearadmin.bham.ac.uk/modules/index

UNIVERSITY OF BIRMINGHAM

BEAR Admin

BEAR
BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

Modules | Deprecated | Users | Top Modules

Search ... | Search | Logout

BEAR Apps Versions

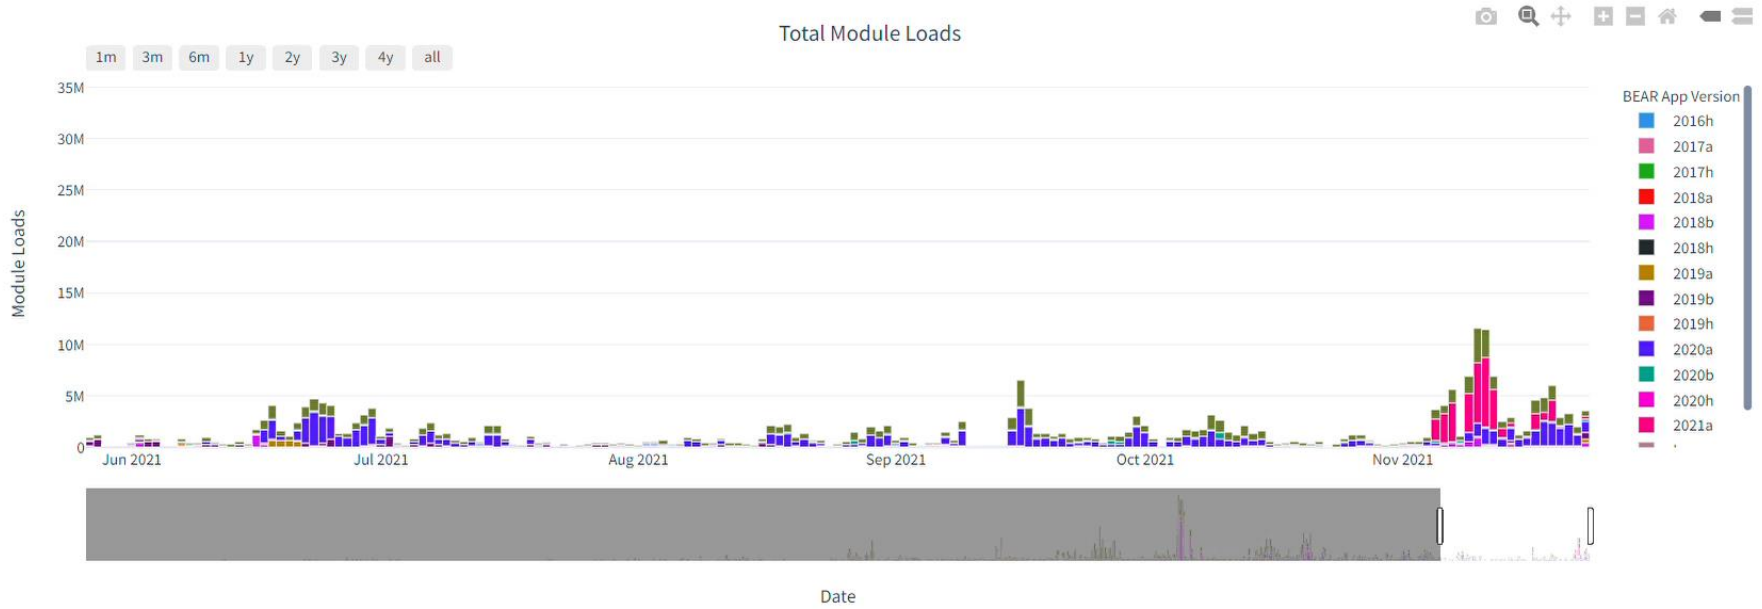

Top

Problems of a changing environment

Example: EL7 to EL8 and Ubuntu 16.04 to Ubuntu 20.04 OS upgrades

- Mid 2020 – “New OS coming”
  - EL7 final scheduled release was 7.9
  - EL8 needed to support latest CUDA on POWER9, for TensorFlow 2.3 etc.
  - EL8 planned for Baskerville, and support for future hardware
  - U1604 out of support in 2021
- Dec 2020 – CentOS support model changed!
  - Change of plans! Switched from CentOS Stream 8 to RedHat 8.
  - Upgraded GPU nodes only
- April 2021
  - Upgraded remaining compute nodes to RedHat 8
  - Upgraded all VMs to CentOS Stream 8, or Ubuntu 2004

Problems of a changing environment

Example: EL7 to EL8 and Ubuntu 16.04 to Ubuntu 20.04 OS upgrades

- Virtually everything needed to be re-installed
  - System library changes – e.g. glibc
- Many old versions wouldn't re-install on EL8 – simple patches but time-consuming!
- Deprecated everything installed before 2018
  - Found who was using old software (using our modules database)
  - Talked to them and helped them to move to new versions, or installs of the old versions with newer compilers, libraries etc.
  - New policy and automatic messaging (next slides)

Deprecating and removing old software

Applications Support and Retention Policy

- Applications installed on the BEAR systems are organised by the year of installation.
- All mention of years in this policy refer to installation date.
- Applications from the current year and the previous two years will usually be supported.
- **Supported applications may be removed at short notice if they can no longer run on the BEAR systems.**
- Applications installed more than two years ago will usually be deprecated.
- **All deprecated applications will warn the user that they are deprecated, each time they are loaded.**
- Any deprecated application may be removed without warning.
- After an application has been deprecated it may become unsupported, instead of being removed.

https://bear-apps.bham.ac.uk/help

UNIVERSITY OF BIRMINGHAM | BEAR BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

Deprecating and removing old software

```
$ module load ScaLAPACK/2.0.2-gompi-2018a-OpenBLAS-0.2.20
------------------------------ WARNING --------------------------------
This module is deprecated:
 * ScaLAPACK/2.0.2-gompi-2018a-OpenBLAS-0.2.20

Please see https://intranet.birmingham.ac.uk/bear-apps-versions for
details of the 2020 OS upgrade for BEAR systems.

This module will NOT be available after the OS upgrade.

Please see https://bear-apps.bham.ac.uk for alternative versions.
-----------------------------------------------------------------------
```

UNIVERSITY OF BIRMINGHAM | BEAR BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

Questions?

UNIVERSITY OF BIRMINGHAM | BEAR BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH

Come and lead our architecture, infrastructure and systems team!

https://www.jobs.ac.uk/job/CKY408/research-computing-infrastructure-architect

Closing date 15th December

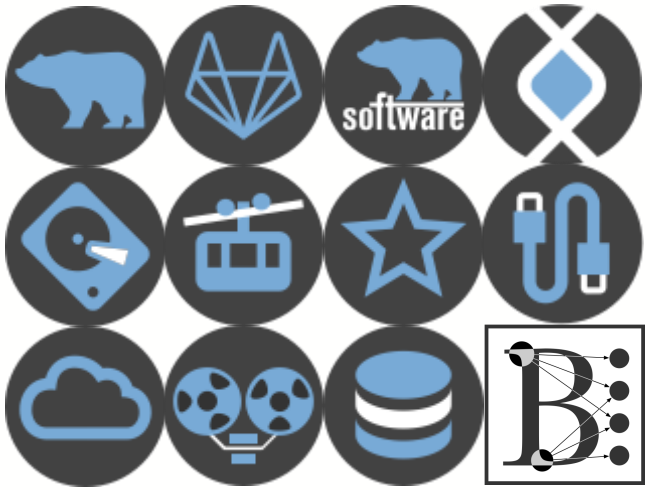

UNIVERSITY OF BIRMINGHAM | BEAR BIRMINGHAM ENVIRONMENT FOR ACADEMIC RESEARCH