

Bede Utilisation at the University of Manchester



George Leaver

George Leaver

As a member of the Research Infrastructure team, George provisions the University's computationally intensive research systems, in particular the Linux HPC clusters and GPU systems available to research groups from all faculties. This covers system administration, application support, software development (multicore and GPU) and training in the use of the University's computational systems and GPU programming.

What is Your Role in Supporting Bede?

I am one of the first people to see applications for Bede at the University of Manchester. I look over the user's requirements and, in conjunction with other colleagues, decide if Bede is the right platform for the project. If it is, I will recommend that a project and user account is created.

Once a project has been added to Bede I can help researchers to access the system, as well as directing them to other support that is available within the university. I also encourage them to join the Bede channel on Slack and signpost them to the Bede user group and its associated mailing list.

What capabilities does Bede add, relative to your on-premise high-performance computing platforms?

Bede uses IBM Power9 CPUs and NVIDIA GPUs alongside high-speed interconnects. This novel architecture is very different to the high-performance platforms that Manchester has, which are

mainly x86-based and ideally suited to CPU computing tasks.

Bede makes extensive use of its GPUs, 4 Tesla v100s per compute node. When combined with the NVLink interconnects, it enables Bede to move tensor outputs to system memory whilst continuing with new calculations. This makes Bede ideally suited to parallel tasks and machine learning and artificial intelligence workloads.

What projects and tasks have you seen utilise Bede?

Alongside the support project, essentially a testing platform for each of the N8 universities using Bede, Manchester has 11 projects using Bede. These include computational fluid dynamics, natural language processing, plasma acceleration, finite element methods, Coronavirus research, airborne biological particle classification, atomistic simulation of various materials and protein interaction for drug design.

Have you been surprised by the variety of projects you've seen apply for time/access to the system?

We have a mixture of existing GPU users who are using Bede to run larger, multi-node GPU jobs that can't be accommodated on the university platforms, and new users that have not previously done much with GPU codes.

It is pleasing to see projects from a variety of disciplines from across the University faculties utilising Bede.

What's the best way for researchers to access the system?

The N8 CIR website has plenty of information about the system, from technical details, research software engineering support and details of the application process. You can find it on the following links: <https://n8cir.org.uk/supporting-research/facilities/bede/>