UNIVERSITY OF LIVERPOOL
Research IT | N8 CIR Research Partnership

POWERING DISCOVERY:
THE IMPACT OF DIGITAL RESEARCH INFRASTRUCTURE AT THE UNIVERSITY OF LIVERPOOL

An Impact Narrative Report

Based on Interviews with 6 Researchers and a Survey of 67 Users

Produced as part of the N8 CIR Part-Time Internship Programme

Supervised by:
Dr Philip Murray (School of Management)
Dr Jianping Meng (IT Services)
Dr Benjamin Pietras (IT Services)
Dr Manhui Wang (IT Services)

April 2026

Executive Summary

Digital research infrastructure supports a large share of the research carried out at the University of Liverpool. Two systems do most of this work: the Barkla high-performance computing (HPC) cluster and the HTCondor high-throughput computing (HTC) platform. Together they serve researchers in the physical sciences, engineering, life sciences, health, mathematics, and a growing number of departments in the social sciences and humanities. For most of the researchers who took part in this study, the work they describe could not be carried out without these systems.

This report presents findings from a mixed-methods study conducted between late 2025 and early 2026, as part of the N8 Centre of Excellence in Computationally Intensive Research (N8 CIR) Part-Time Internship Programme. The study combined six in-depth qualitative interviews with active HPC users and a survey of 67 participants. Of these, 49 completed the survey in full and form the basis of the quantitative analysis, whilst a further 18 provided partial responses. Participants were drawn from across the university community.

Across the survey responses and the six interviews, the same finding recurs: for most users the infrastructure is a condition of their research rather than an optional aid. The main quantitative findings were:

- 85.4% of the 48 respondents who answered rated HPC 'Absolutely Essential' (56.2%) or 'Very Essential' (29.2%) to their research.
- 93.8% reported that their results would be impossible or seriously compromised without HPC access.
- 91.6% agreed that HPC had reduced the time taken to produce their research results, the highest scoring of six impact statements.
- Of the 35 respondents who rated the support services, 83% judged them 'Very Effective' or 'Extremely Effective', and none rated them poorly.
- The 49 respondents reported 264 publications, 79 datasets, and 26 software tools attributed to HPC-supported work over five years.
- Barkla 2 supports work in medical imaging, computational fluid dynamics, nuclear engineering, materials discovery, AI in healthcare, and other fields across six broad disciplinary groups.

One theme ran through every interview. HPC has changed not only how quickly research can be done, but the size and type of questions researchers are able to ask. Several described work that would take a week on a desktop machine being completed overnight on Barkla, and problems that could not be attempted at all on local hardware becoming routine. The interviews and survey also record the part the infrastructure plays in winning external funding, supporting national and international collaborations, and producing research with measurable effects in healthcare, aviation, nuclear energy, and materials science.

The study also records clear limits on the current provision. The clusters are over-subscribed, which lengthens queue times; storage quotas constrain some workflows; a steep onboarding curve for less technically experienced users; and demand for GPU resources is rising as more research codes are written to use them. On the evidence gathered here, continued and increased investment in the university's digital research infrastructure directly supports the research community and the university's position as a research intensive Russell Group institution.

1. Introduction

1.1 Context and Background

A large and growing share of research now depends on computation. In fields from aircraft aerodynamics to medical imaging, the questions researchers most want to answer involve data volumes and model sizes that ordinary desktop hardware cannot handle. Problems that were out of reach a generation ago, because the necessary calculations would have taken too long or required too much memory, are now routine on high-performance and high-throughput computing systems. This report examines how that shift plays out at the University of Liverpool.

The University of Liverpool, a founding member of the Russell Group, has invested in its digital research infrastructure (DRI) over many years. The main on-campus HPC system, Barkla (now in its second generation, Barkla 2), gives researchers access to multi-node computing, GPU resources for machine-learning work, and a supported environment for running scientific codes. The HTCondor platform handles high-throughput computing, where the need is to run large numbers of independent jobs rather than one tightly coupled simulation. Beyond the campus, Liverpool researchers also use the national ARCHER2 supercomputer and the N8 Research Partnership's regional Tier-2 facility, Bede.

The university has not previously gathered systematic evidence on how this infrastructure is used, what difference it makes to the research that depends on it, or how clearly that value is conveyed to those who decide on funding. This report addresses that gap, drawing on a survey of 67 users and six interviews with active researchers.

1.2 Aim and Scope of This Report

This report was commissioned as part of the N8 CIR Part-Time Internship Programme and supervised by Dr Philip Murray (School of Management), Dr Jianping Meng (IT Services), Dr Benjamin Pietras (IT Services) and Dr Manhui Wang (IT Services) at the University of Liverpool. It has two aims: to document the effect of Liverpool's digital research infrastructure through the accounts of the researchers who rely on it, and to present that evidence as an impact narrative that can inform decisions about investment, support future funding bids, and contribute to the university's record of research impact.

The study focuses specifically on the Barkla HPC cluster and HTCondor HTC platform as Liverpool's primary DRI assets, with reference to external systems where relevant. It examines impact across multiple dimensions: research quality, speed and scale; research outputs and grant success; skills development; collaboration; and broader societal, environmental, and economic contributions.

1.3 Why This Matters Now

The timing of this report matters. Universities are increasingly asked to show the return on the money spent on research infrastructure. At the same time, the computing demands of research are rising quickly, driven in particular by machine learning and AI workloads that depend heavily on GPUs. Liverpool's recent £3 million upgrade to the Barkla cluster reflects an institutional recognition of this reality. This report provides the evidence base, drawn from the people who use the systems, to inform decisions about further investment.

2. Methodology

2.1 Research Design

The study used a mixed-methods design, pairing a quantitative survey with qualitative interviews. The survey gives a broad picture across the user community and allows the findings to be quantified; the interviews explain what the survey numbers mean in practice, in the words of the researchers themselves. The two strands were designed to support one another: interview themes were checked against the survey results, and survey results were read alongside what interviewees said.

2.2 Survey

A structured online questionnaire was distributed to registered users of the university's HPC and HTC platforms. It collected data on respondent profile (department and career stage); platform and cluster usage; how essential respondents considered HPC to their work; their level of agreement with six statements about HPC's effect on their research; research outputs produced over the previous five years; use of and satisfaction with support and training services; and views on wider societal, environmental, and economic effects. Sixty-seven responses were received. Forty-nine were completed in full and form the basis of the quantitative analysis; the remaining 18 were partial and are drawn on where relevant. One of the 49 respondents did not answer the profile and impact questions, so the figures for those specific questions are based on 48 responses. This is noted in each affected figure, and percentages are calculated on the number who actually answered each question.

2.3 Interviews

Six semi-structured interviews were then carried out, in person, with active users of the university's HPC systems, chosen to cover a range of disciplines, career stages, and patterns of use. The interviews followed the themes raised by the survey in more depth, giving participants room to explain in their own words how the infrastructure supports their work. The participants were:

Participant	Research Area	Role	Primary System
Samuel Ball	Health Data Science / Neural Networks	Lecturer	Barkla
Sebastian Timme	Computational Aerodynamics / CFD	Senior Lecturer	Barkla 2
Dzanis Litskevich	Nuclear Engineering Simulation	Lecturer	Barkla 2
George Darling	Materials Chemistry / Surface Science	Senior Lecturer	Barkla / Archer2
Zhiheng Wang	Computer Science / Digital Twin	MSc Student	Barkla 2
Yalin Zheng	AI in Healthcare / Medical Imaging	Professor	Barkla

Each interview was conducted in person and lasted between 30 and 60 minutes and was recorded and transcribed in full. The questions covered how HPC is used in the participant's research; its effect on the quality, speed, and scale of that research; its part in producing outputs, winning grants, and enabling collaborations; the participant's experience of support and training; and what they felt university leadership should understand about investment in computing infrastructure.

2.4 Analysis

The interview recordings were transcribed and read in full several times to identify recurring points, the main types of impact described, and quotations that represented each theme. Themes were drawn from the interview material and then checked against the survey results. The survey data were analysed using counts and percentages, with each percentage calculated on the number of people who answered that question. Free text survey answers were read and used to add detail to the numbers. The whole analysis was directed at one question: what effect has Liverpool's digital research infrastructure had, and for whom?

3. Research Computing at Liverpool: The Infrastructure

It helps to set out what the infrastructure consists of before turning to its effects. The University of Liverpool runs several DRI platforms, each suited to a different kind of computational work.

Barkla is Liverpool's main on-campus HPC cluster, now in its second generation after a major upgrade completed in 2025. It provides a large pool of CPU compute, including high memory nodes of up to 168 cores and 1.5 TB of RAM, alongside GPU nodes fitted with NVIDIA H100 (80 GB VRAM) and L40 cards for training and evaluating large AI and machine learning models. Barkla is free to use for all registered researchers at the university, which keeps the barrier to large scale computation low.

HTCondor is the university's high-throughput computing platform, designed for researchers who need to run very large numbers of independent jobs rather than a single tightly coupled parallel simulation. It is particularly suited to parameter sweeps, Monte Carlo simulations, and genomics pipelines.

Beyond the campus, Liverpool researchers can apply for time on the national ARCHER2 supercomputer through EPSRC allocation processes and can use the N8 Research Partnership's Tier-2 facility, Bede, at Durham University, along with other specialist systems. These external resources extend what Barkla can do for work that exceeds local capacity. They require applications in advance, however, and are not available for exploratory or development work. As the interviews show, that difference has real day to day consequences for how researchers can work, a point returned to in Section 5.

4. Findings: Who Uses the Infrastructure and How

4.1 Respondent Profile

The survey drew responses from across the university's research community. Engineering and Chemistry were the most represented departments among fully completed responses, with eight respondents each (16.7%). Physics and Computer Science each contributed five respondents (10.4%), and Mathematics four (8.3%). The remaining 18 respondents (37.5%) came from a range of other disciplines: Health Data Science (2), Public Health, Earth and Ocean Sciences, Environmental Sciences, Biochemistry, Pharmacology, Psychology, Economics, Biology, Communication and Media, and Music, with one response each. This breadth of participation reflects the cross-disciplinary reach of the infrastructure.

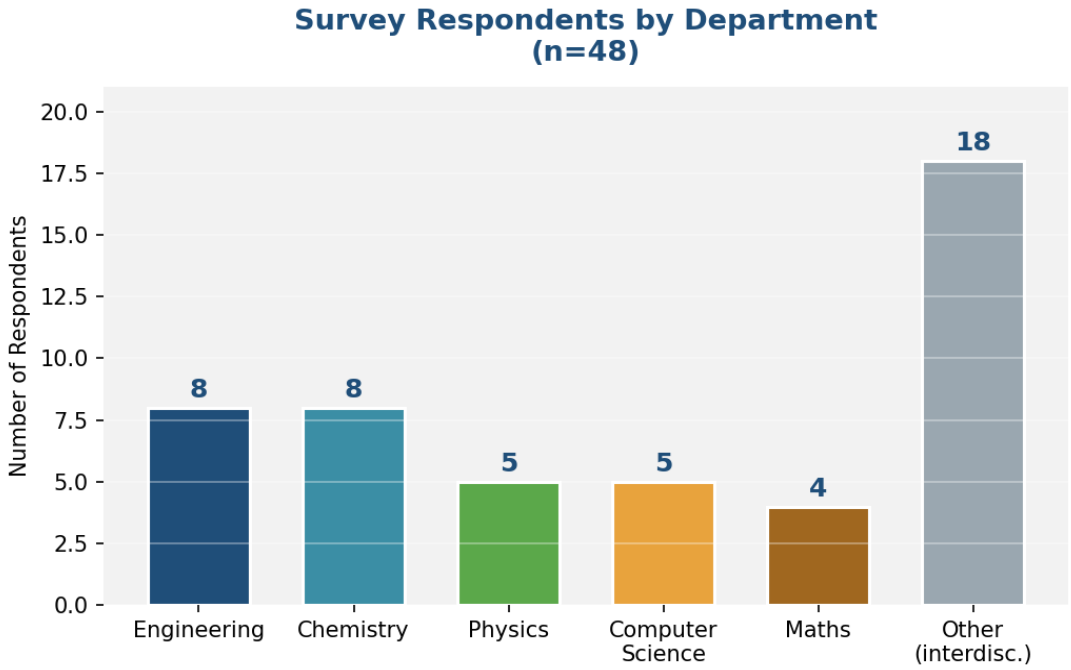


Figure 1: Survey respondents by department (n=48; one of the 49 respondents did not answer this question)

Academic staff formed the largest respondent group (26 respondents, 54.2%), followed by postgraduate researchers (10, 20.8%), postdoctoral researchers (9, 18.8%), and undergraduates (3, 6.2%). The spread demonstrates that the infrastructure serves users at every stage of the academic career pipeline, from undergraduates encountering research computing for the first time to professors whose entire research programmes depend upon it.

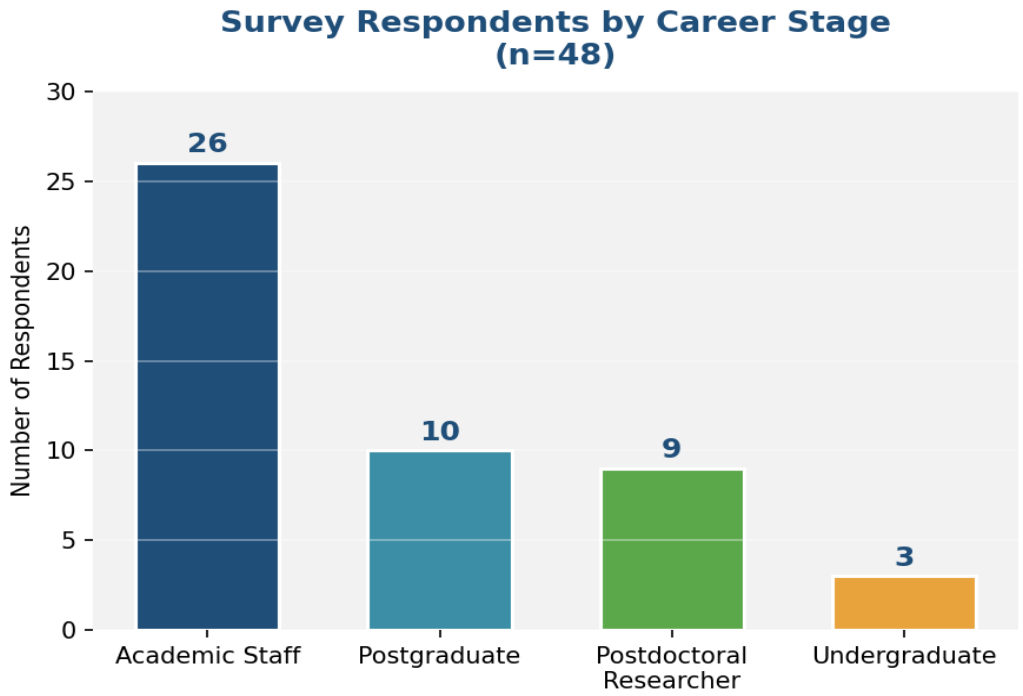


Figure 2: Survey respondents by career stage (n=48; one of the 49 respondents did not answer this question)

Platform usage was dominated by HPC: all 47 respondents who specified a platform used HPC, and 48 out of 49 used Barkla specifically, confirming it as the central DRI asset. A smaller number also accessed HTC (5 respondents), ARCHER2 (6), and Bede (3), typically alongside rather than instead of Barkla.

4.2 Computational Tasks

Respondents used the infrastructure for a range of computational tasks. Modelling and simulation was the most commonly reported activity, covering engineering (CFD, structural mechanics, finite element analysis), physics, materials science, nuclear engineering, and earth sciences. Machine learning and AI was the second most reported category, used in computer science, health data science, medical imaging, pharmacology, and, more recently, in music and communications. Data analysis was the third, covering biochemistry, public health, psychology, and environmental science. Smaller numbers reported rendering and visualisation, or specialist work such as Monte Carlo sampling, crystal-structure prediction, and computer algebra.

This range is itself a finding. Liverpool's DRI is used not only by the disciplines that have always relied on HPC, such as physics, chemistry, and engineering, but by a widening group of departments for which computational methods are becoming a normal part of the work.

5. Impact: The Evidence

5.1 Essentialness: Establishing the Baseline

Before looking at the separate effects of the infrastructure, it helps to set out how essential the researchers who responded considered it to be.

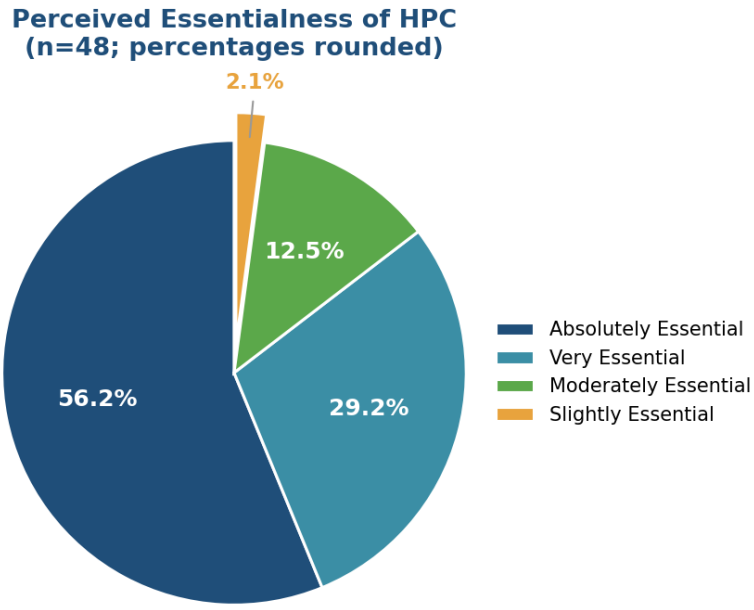


Figure 3: Perceived essentialness of HPC to research (n=48; one of the 49 respondents did not answer this question)

85.4%	of survey respondents rate HPC as 'Absolutely Essential' or 'Very Essential' to their research.

93.8%	report that their results would be impossible or significantly compromised without HPC access.

Two survey questions bear on this directly. When asked how essential HPC was to their research, 85.4% of the 48 respondents who answered chose 'Absolutely Essential' or 'Very Essential'. When asked whether their results would be possible without HPC, 93.8% said they would be impossible or seriously compromised. The two questions approach the same issue from different angles and produce figures within eight points of each other, which suggests the responses are measuring a real and shared dependence rather than a passing opinion. For most respondents, HPC is a condition of the research being possible at all, not a means of doing it more comfortably.

The interview participants described the same dependence, in stronger terms, when asked what HPC made possible in their own research.

> *"I couldn't do the research which I do without the HPC. I couldn't get it to nearly the quality of research which I can achieve with HPC. It just wouldn't be feasible computationally."*
> **— Dr Samuel Ball, Lecturer in Health Data Science**

> *"If I didn't have access to HPC, all I could simulate would be aerofoils or small, unrepresentative wings on my local computer. But anything more realistic, which makes the work competitive on an international scale, would not be possible."*
>
> **— Dr Sebastian Timme, Senior Lecturer in Computational Aerodynamics**

5.2 Research Impact Statements: Survey Evidence

The survey asked respondents to rate their agreement with six statements about the impact of HPC on their research. The results are presented in Figure 4 below.

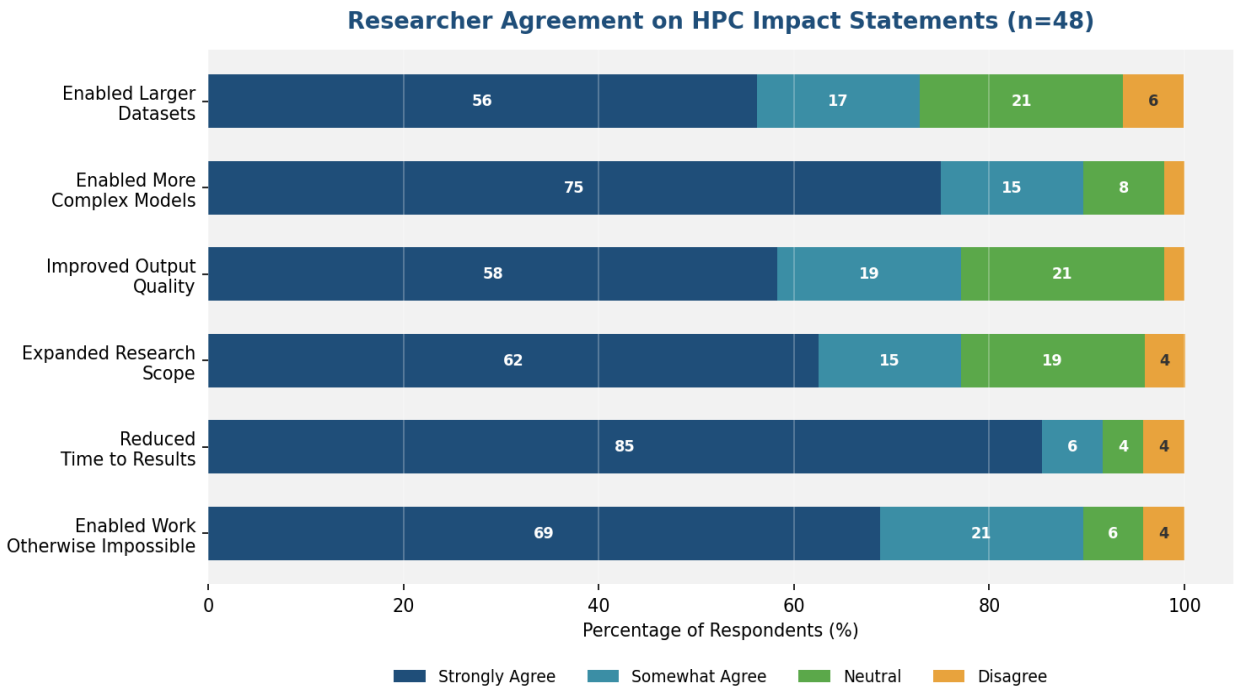


Figure 4: Researcher agreement on HPC impact statements (n=48; one of the 49 respondents did not answer this question)

The strongest agreement concerned the time taken to produce results: 85.4% strongly agreed and a further 6.2% somewhat agreed, a combined 91.6%. The agreement was also high that HPC had let them run more complex models (89.6% agreeing in total), do work that would otherwise be impossible (89.6%), and widen the scope of their research (77.1%). The lowest agreement, on enabling larger datasets, still reached 72.9%. No statement fell below 70% agreement, and the two statements about what HPC makes newly possible, complex models and otherwise-impossible work, scored almost identically, which is consistent with the interview accounts in the sections that follow. Figure 4 sets out the full distribution.

5.3 Speed and Scale: The Research Acceleration Effect

The benefit reported most often, in both the survey and the interviews, was a reduction in the time taken to complete demanding computational work. In the survey, this was the single highest scoring impact statement, with 91.6% agreeing. The interviews show what that figure represents in practice, with two participants giving concrete measures of the speed-up they see.

> *"It's about scale. With our own GPUs, we can handle smaller tasks, but with big data we cannot. Recently we have been dealing with over half a million images. Our computers cannot deal with that, so we have to use Barkla."*
> **— Professor Yalin Zheng, Professor of AI in Healthcare**

Zhiheng Wang, an MSc student modelling blood flow in the human eye, gave a measured example. Fluid-dynamics simulations that took days on a CPU completed in seconds on the GPU resources of Barkla 2. The gain was not only in the single run: because each run finished quickly, he could correct errors and repeat the analysis many times within his project, a cycle that would not have been possible on a personal machine. Dr Samuel Ball described a similar order-of-magnitude change in his own work, estimating that tasks ran roughly ten times faster on Barkla, so that a model left to train overnight was ready by morning rather than part-finished after several days.

> *"If I didn't have access to it, it would significantly delay the progress. Because each simulation takes days, and if I get something wrong, I need to rerun the whole programme again and wait another day."*
> **— Zhiheng Wang, MSc Computer Science Student**

Professor Yalin Zheng's group in AI and Healthcare works with datasets of more than half a million medical images. The group's own workstations, each with a dedicated GPU, cannot hold or process data on that scale, so the work moves to Barkla. Without it, the group could not train and test new AI models on the volume of data that clinically useful results require.

5.4 Quality and Scope: Doing Better Science

Speed is not the only gain the interviews describe. Several participants explained that HPC lets them work with full resolution data instead of downscaled approximations, run larger models, and repeat analyses enough times to test how stable their results are. Each of these affects the quality of the findings, not just the speed at which they are produced, and the survey reflects this: 77.1% of respondents agreed that HPC had improved the quality of their output, and 89.6% that it had let them run more complex models.

Dr Samuel Ball's work on uncertainty in medical image AI shows the point. His method resamples a neural network many times on the same image, varying the inputs slightly to see whether the prediction stays stable or shifts. Doing this well needs considerable compute. On a desktop machine he would have to cut the image resolution or the number of samples; on Barkla he can run the full method, which produces firmer statistical results and models that reach the standard of reliability clinicians require before they will use them, for example within the NHS.

> *"With the HPC, because I could use the full images, the models trained better, they were more accurate. They had more clinical utility. They were just generally a lot better. And because of the method that I'm using, you resample the AI multiple times. Because it's orders of magnitude faster, I can resample it a lot more, which gives more robust statistical analysis."*
> **— Dr Samuel Ball, Lecturer in Health Data Science**

In Materials Chemistry, Professor George Darling's group uses HPC to run materials discovery simulations at a scale that is not achievable on desktop hardware. The group's in-house methods run many simulations in parallel to identify optimal structural configurations across large compositional spaces. This capability has produced multiple high-profile publications in journals including Science and Nature.

> *"Now, we can quite reliably predict where experimentalists should try to make a material: what elements and what ratios you could mix to make a new material. In the past, you simply couldn't do the calculations well enough to get reliable predictions."*
> **— Professor George Darling, Senior Lecturer in Materials Chemistry**

5.5 Research Outputs: Publications, Datasets, and Software

The survey collected data on research outputs produced over the previous five years and attributed at least in part to HPC access. Among the 35 respondents who reported publication counts, the totals came to 264 publications, ranging from 1 to 30 per individual. Respondents additionally reported 79 datasets and 26 software tools produced with HPC support.

These figures should be read as a floor rather than a complete count, since they reflect only the 49 respondents to this survey and only the outputs each chose to attribute to HPC. Even so, 264 publications from a single survey sample, alongside 79 datasets and 26 software tools, indicates that the infrastructure feeds into a wide body of work. The outputs span journals, conference proceedings, preprints, and technical reports across physics, engineering, chemistry, the health sciences, mathematics, and environmental science, which matches the spread of departments seen in the respondent profile in Section 4.

> *"Local HPC is required in addition to national HPC; they don't fulfil the same role. There is a huge amount of development work that requires a local HPC, as well as work for undergraduate project students, where you can't anticipate what they are going to do, so you can't apply for time in advance. Barkla gives you the chance to try things immediately."*
> **— Professor George Darling, Senior Lecturer in Materials Chemistry**

5.6 Grant Competitiveness and External Funding

The relationship between DRI access and grant competitiveness is a recurring theme in both the interview and survey data. Several interviewees described the availability of Barkla as a tangible asset in research grant applications, something that can be listed as an in-kind contribution from the university and that signals to funders the operational credibility of the proposed work.

Dr Dzanis Litskevich, whose nuclear-engineering simulation work runs across several EU and EPSRC projects, made this point directly. Being able to offer HPC-based computation as part of a bid has a clear monetary value, because the CPU hours can be costed and entered as an in-kind contribution from the university. For a team that has to show a funding panel, it has the computing capacity to deliver access to Barkla is a concrete and credible answer.

Professor Yalin Zheng's group shows how access to the infrastructure can bring in outside investment. When the COVID-19 pandemic began in 2020, the HPC team set up a link with Amazon Web Services that secured US$50,000 of GPU time for the group to work on AI-assisted COVID-19 diagnosis. The work led to several peer-reviewed papers and a presentation at the SC20 Supercomputing Conference. At its peak the allocation was being used at about $10,000 per week, which gives a clear measure of the value that good HPC infrastructure and support can unlock.

> *"Our HPC colleagues have helped to facilitate collaboration with industry. A good example is that when COVID started, the HPC team set up a link for us with Amazon. AWS provided us with $50,000 worth of GPU time to work on AI models. We have presented*

> *our work at a supercomputing conference and published several papers from that collaboration. It has been very successful."*
> **— Professor Yalin Zheng, Professor of AI in Healthcare**

Dr Sebastian Timme noted that, although Liverpool's facility is at the smaller end for a Russell Group university, it has supported national and European projects in computational aerodynamics and made possible the development of open-source CFD codes now used in aircraft simulations across the field.

5.7 Societal, Environmental, and Economic Impact

The impact of Liverpool's DRI extends beyond the university's own publications and grants. Several interviewees described research enabled by HPC that is contributing directly to real world challenges of global significance.

Healthcare and Medicine

Professor Yalin Zheng's group is developing AI tools for the screening of diabetic retinopathy, a leading cause of blindness and a complication affecting an estimated 110 million people in China and around four to five million in the UK. The group's EPSRC Global Challenges Research Fund award resulted in AI models that have been patented by the university and formed the basis of a spin-out company developing commercial screening tools. Barkla is integral to training and evaluating these models at the scale required for clinical validity.

Dr Samuel Ball's work on uncertainty quantification for medical image AI addresses one of the key barriers to clinical adoption: the inability of AI models to communicate when they are uncertain about a prediction. His research, directly dependent on Barkla for its computational feasibility, aims to produce systems that recognise when to defer to a clinician. NHS consultants have identified this capability as one of the most important missing features in current AI tools. Ball's earlier work on brain imaging for children with cerebral malaria in Malawi represents a further example of global health impact made possible by HPC.

> *"We developed a new AI model that has been patented by the university, and the university also made a spin-out from one of the patents. So you can see we do basic research through to translational industry impact."*
> **— Professor Yalin Zheng, Professor of AI in Healthcare**

Sustainable Aviation and Net Zero

Dr Sebastian Timme's computational aerodynamics work feeds into the aerospace sector's effort to reduce fuel burn and carbon emissions. His group develops CFD codes and runs them on representative aircraft-wing models to predict aerodynamic behaviour, and these predictions inform the design of more efficient aircraft. Because the simulations involved are large and require many runs, the work depends on access to HPC; on local hardware Timme reports being limited to small, unrepresentative cases that would not be competitive internationally.

> *"Aerodynamics used for aircraft design is very important, so developments in my field can help with aircraft design. What I've done in the past is develop some tooling, free codes, and then run them on representative aircraft models, which require very many simulations."*
> **— Dr Sebastian Timme, Senior Lecturer in Computational Aerodynamics**

Clean Energy and Nuclear

Dr Dzanis Litskevich's work on molten salt reactors, carried out within a consortium of universities, national laboratories, and industrial partners, addresses the longer-term problem of sustainable nuclear energy. The technology could use spent nuclear fuel, currently treated as waste, as a fuel source. Barkla has carried much of the modelling and simulation behind this work, and the group's ability to offer HPC based analysis as an in-kind university contribution has strengthened multi-partner funding bids.

Materials Discovery

Professor George Darling's materials-chemistry group works at the leading edge of computational materials discovery, using HPC to predict which combinations of elements might form new materials for batteries and fuel cells. The work is basic research, but it has clear long-term implications for clean energy storage and the wider move to low carbon energy.

The survey supports the interview accounts of wider effects, with some variation between categories. Asked about societal effects, 48.9% of respondents agreed (somewhat or strongly) that their HPC enabled research had contributed to them. The figures for environmental (37.8%) and economic (29.5%) effects were lower, and more respondents were unsure about these. This fits the fact that many researchers in more basic disciplines see their work as adding to knowledge rather than producing direct applied outcomes, which is an accurate account of how the research contributes.

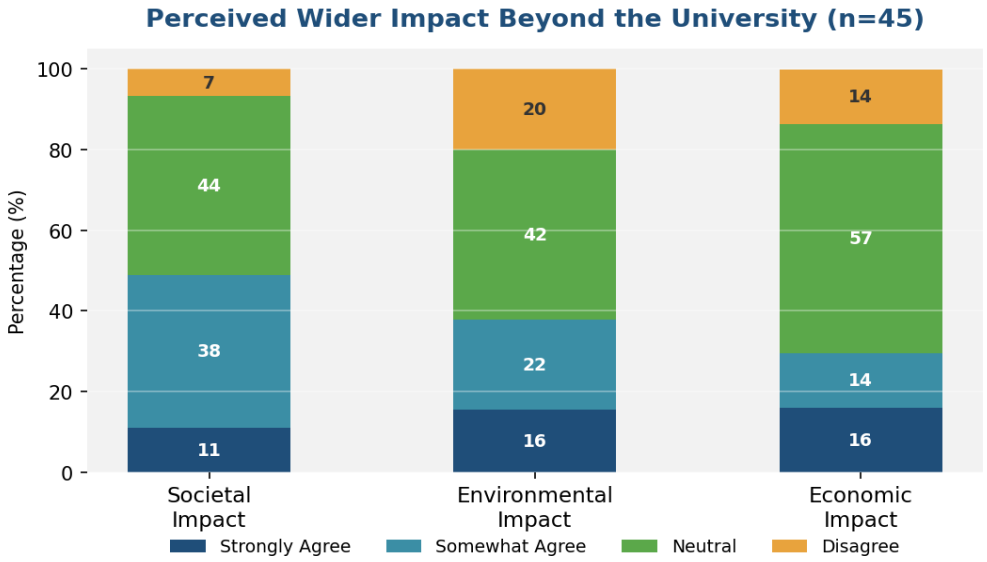


Figure 5: Perceived wider impact of HPC-enabled research beyond the university (n=45)

6. Support, Training, and the Human Infrastructure

6.1 The Role of Technical Support

A capable system still has to be usable. Both the survey and the interviews point to technical support as the factor that turns access to HPC into usable research, and the survey provides a direct measure of how that support is rated.

Of the 49 respondents, 73.5% (36 people) had used the HPC support services. Of those, 29 of 35 who rated the service, or 83%, judged it 'Very Effective' or 'Extremely Effective'; the remaining six rated it 'Moderately' or 'Slightly Effective', and none rated it poorly. The interviews give the same picture in more detail and explain what that support consists of in practice.

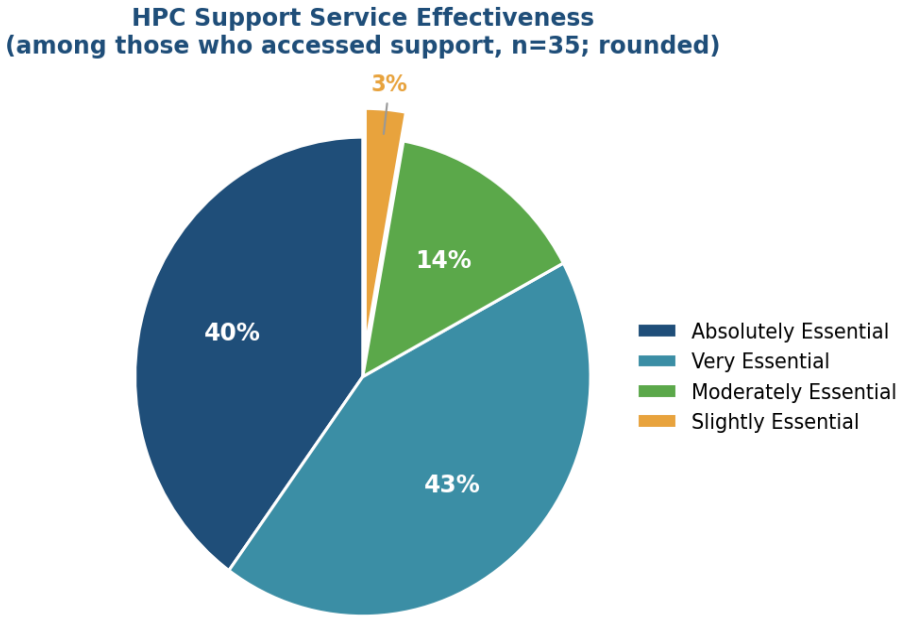


Figure 6: Perceived effectiveness of HPC support services (among those who accessed support, n=35)

The four interviewees who were asked about support all described it in positive terms, consistent with the survey rating above. Dr Samuel Ball said the team's replies, through the Teams channel and directly, were quick and came from people who knew the system, and he credited a departmental talk by the HPC team with drawing in colleagues who had not used HPC before. Professor Zheng said email support was answered within about a day and credited the team with arranging the AWS partnership behind the COVID-19 work. Dr Litskevich described the Barkla 2 team as quick to respond and willing to install requested software without delay. Zhiheng Wang, as a newer user, made the same point from the other end of the experience scale, which is set out in the quotation below.

> *"I registered on the mailing list, so I got the email about the drop-in sessions and I went to those. When I wanted to install a package, they responded within a day and got it installed for me. It is very supportive."*
>
> **— Zhiheng Wang, MSc Computer Science Student**

> *"I think it's the best experience I've had with high performance computers and supercomputers, especially with the updated Barkla 2. It's getting better and better with each update, and I am really happy with what we've got now."*
> **— Dr Dzanis Litskevich, Lecturer in Nuclear Engineering**

6.2 Training and the Research Computing Clinic

Of the 49 respondents, 38.8% had received formal HPC training, and of those, 82% rated it 'Very Effective' or 'Extremely Effective'. The Research Computing Clinic, a drop-in service, was known to many respondents: most of those aware of it had either attended and found it helpful or were aware but had not yet been. A small number did not know it existed, which points to a need to publicise it more widely, particularly to newer users.

Several interviewees and survey respondents offered consistent suggestions for improvement. The onboarding experience for new and less technically experienced users needs to start at a more foundational level. Dr Ball noted that students coming to HPC training sessions often needed Linux basics before they could engage meaningfully with cluster concepts. Dr Litskevich suggested short introductory video series and more regular structured onboarding events, particularly at the start of the academic year.

> *"For PhD students, that would be a great opportunity. For those who are involved in modelling and simulations, very often they come in and they don't know anything about HPC. So it's going to be a good idea to have this at least once a year. One or two days, or one full day would probably be enough, to give an overview."*
> **— Dr Dzanis Litskevich, Lecturer in Nuclear Engineering**

7. Challenges and Limitations

The evidence of impact set out above sits alongside a set of clear limits. Left unaddressed, these limits will hold back what the infrastructure can contribute to the university's research.

7.1 Capacity and Queue Times

The limit raised most often, in both the survey and the interviews, is over-subscription: more demand for compute than the clusters can meet at once, which lengthens job queues. Dr Timme noted that a request for two compute nodes can mean a wait of close to a week, which slows the work considerably. Professor Zheng's group, which may need to run 20 experiments at once before a conference deadline, often cannot get the GPU time needed at the moment it is needed. Several free text survey answers asked directly for more nodes, more GPUs, shorter queues, and more compute overall.

> *"For a Russell Group university, it's probably one of the smaller facilities. If you think about places like Bristol, who have a massive national investment in a new supercomputer, it's quite a different scale that they can operate on."*
> **— Dr Sebastian Timme, Senior Lecturer in Computational Aerodynamics**

7.2 GPU Demand and Accelerated Computing

A specific and growing pressure sits within the broader capacity problem: demand for GPUs. AI and machine learning work accounts for part of this, but a large and rising share comes from research codes that use GPUs purely as accelerators, with no AI component, and this accelerated computing workload is now a main driver of GPU use. The H100 cards are heavily queued, with waits of several days reported. The scale of the speed-up explains the demand: the HPC team notes that a job taking 24 hours on a CPU can finish in under an hour on a GPU, so as more software is rewritten to use GPUs, the pressure will grow rather than ease. One practical point for planning is that the current generation of GPU hardware is the last expected to use air cooling; any future data centre expansion will need to provide direct liquid cooling.

7.3 Storage Quotas

File and storage quotas were raised as a practical constraint, particularly by researchers in simulation-heavy disciplines who generate large numbers of intermediate files. One survey respondent noted that in symbolic computation work, quota restrictions caused jobs to fail before completion. Another described repeatedly reaching file number limits whilst using only a fraction of their total storage allocation, owing to the volume of small files their simulations produce. In both cases, the quotas resulted in lost time and reduced productivity for the researchers affected. It should be noted that quotas serve a necessary purpose in maintaining overall system performance; the HPC team provides monitoring tools and documentation to help, and additional storage can be purchased by research groups or addressed through use of local scratch space. Where further investment in general storage is made, quota limits could be adjusted accordingly.

7.4 Onboarding and Accessibility

How much technical fluency HPC requires varies widely across the user community. Researchers with a computing background manage the learning curve without much difficulty, helped by the documentation and the support team. For those from less technical disciplines, the terminal, SSH, Linux commands, SLURM job scripts, and containerisation form a real barrier to getting started. One survey respondent put it plainly in a free text answer:

> *"Onboarding for postgraduate students and researchers: complete terror when you ask them to use HPC."*
> **— Survey free text response**

This gap has consequences for the university's strategy. A large and growing body of research in the social sciences, health, environmental science, and the arts depends on computational methods. If Liverpool's DRI stays in practice within reach only of the traditional computing disciplines, it will not serve as a sizeable and growing part of the university's research community.

7.5 Physical Capacity Constraints

Dr Timme reported that the university's current data centre is full and cannot take new rack installations. If that is so, it sets a hard ceiling on near-term growth in capacity and needs to be dealt with as part of strategic infrastructure planning before the other capacity investments recommended in this report can proceed.

8. Implications and Recommendations

The recommendations below follow from the evidence in this report and are addressed to the University of Liverpool's senior leadership, Research IT, and the N8 CIR partnership. They reflect the points raised most often across the survey and the interviews.

Recommendation 1: Treat Expanded Compute Capacity as a Strategic Priority

Demand for HPC and GPU resources at Liverpool is greater than the current supply. Over-subscription, extended queue times, and the inability to respond to time-sensitive research needs all represent direct opportunity costs to the university's research productivity. The university should measure its HPC capacity against comparable Russell Group institutions and set out a multi-year plan to close the gap. Competing for a place among the top 100 universities globally requires computing capacity at a matching level.

Recommendation 2: Prioritise GPU Expansion to Meet AI and Machine Learning Demand

The growth of machine learning and AI methodologies across all research disciplines means that GPU availability is now a critical bottleneck. The existing H100 nodes are heavily oversubscribed. The university should invest in additional high memory GPU nodes as a matter of urgency. GPU computing is no longer a specialist requirement; it is becoming a general research need across the sciences, social sciences, and increasingly the humanities.

Recommendation 3: Address the Data Centre Capacity Constraint

If the current data centre cannot accommodate new hardware, this represents an existential constraint on the university's computational research capacity. The university should urgently investigate expansion options, including partnerships with the Science and Technology Facilities Council at Daresbury Laboratory, cloud burst capabilities, and potential new data centre investment. Without resolving this physical constraint, all other investment recommendations become moot.

Recommendation 4: Review Storage Provision and Quota Flexibility

Storage quotas serve a necessary role in protecting overall system performance, but they are a recurring source of friction for researchers generating large numbers of files. The university should review storage provision so that quota limits can be raised where further investment in general storage is made, communicate the existing option for research groups to purchase additional storage, and promote effective use of existing resources such as local scratch space. Clearer guidance and monitoring tools would help researchers manage allocations before jobs are affected.

Recommendation 5: Expand Structured Onboarding and Training, Supported by Adequate Staffing

Onboarding is the most consistently cited obstacle to broader HPC adoption. Research IT already runs an Introduction to HPC workshop tailored to groups outside the traditional HPC disciplines, and this provision should be expanded and more widely publicised, alongside short workflow specific video tutorials and a regular onboarding schedule tied to the academic year. Crucially, training capacity depends on staffing: Research IT has been reduced from nine to four staff in recent years, and

restoring capacity would directly enable the expanded training and support the user community needs. The Research Computing Clinic should also be promoted more visibly to newer users.

Recommendation 6: Invest in Research IT Staff and Retain Institutional Expertise

The human infrastructure behind Liverpool's DRI is as important as the hardware. The quality of support, evidenced by the 83% 'very' or 'extremely' effective rating in the survey, reflects a capable and dedicated team. Dr Litskevich's observation that the university should invest in people alongside hardware is well founded. Staffing levels should be commensurate with the growing complexity and breadth of the user community, and the expertise of key individuals should be retained through adequate recognition and succession planning.

Recommendation 7: Improve Documentation, Visibility, and Awareness of the Infrastructure

Multiple interviewees and survey respondents noted that many potential users across the university are unaware of what Barkla can do for them or are discouraged by a perception that it is only for technical specialists. The university should invest in a dedicated, up to date web presence for HPC, including dynamic status pages, user guides, FAQs, and case studies illustrating impact. Research success stories enabled by HPC should be actively highlighted through the university's communications channels, helping to normalise HPC use across disciplines.

9. Conclusion

The evidence in this report points in one direction. Liverpool's digital research infrastructure, built around the Barkla HPC cluster and the HTCondor HTC platform, is not a background utility. It is a working part of the research itself, used by scientists, engineers, clinicians, mathematicians, and a growing number of social scientists and humanities researchers to pursue questions they could not otherwise address.

The figures are consistent. Of those who answered, 93.8% said their results would be impossible without HPC, 85.4% rated it absolutely or very essential, and 91.6% said it had cut the time taken to get results. These are not small changes to how the work is done. They mark the difference between research that happens and research that does not; between AI models accurate enough to help clinicians and models that fall short; and between aircraft designs tested in simulation and designs left to physical prototyping alone.

The accounts behind the figures carry the same weight: AI models that can tell a clinician when their prediction should not be trusted; aerodynamic simulations that feed into lower emission aircraft; reactor models that may one day turn nuclear waste into usable energy; materials predicted on the cluster for experimentalists to try to make; a pandemic response that produced peer reviewed AI tools within months through a partnership the HPC team arranged; and a master's student modelling blood flow inside the human eye to improve the diagnosis of diabetes.

The report also records clear and pressing limits. The clusters are oversubscribed. The data centre is full. Demand for GPUs is outrunning supply. Storage limits are interrupting active work. And researchers from less technical backgrounds face barriers that, left in place, will keep a growing group of potential users away from the systems.

The message for the University of Liverpool's leadership is straightforward. The return on investment in digital research infrastructure is not only financial. It shows up in papers that would not have been written, grants that would not have been won, diagnoses that would not have been reached, and materials that would not have been found. The question is not whether to invest, but whether the university can afford not to.

> *“I think they should look into the investment in HPC and in general computational infrastructure, in the same way they look into physical facilities. With the emergence of AI technology and the growing need for computational power, this kind of facility becomes critical for any research organisation which aims for leadership in the world and in the UK. And one more important thing: it's not only the hardware we should invest in. Also, we should invest in people. We should invest in training and invest in the staff who are supporting this infrastructure. They're doing an amazing job in my view.”*
>
> **— Dr Dzanis Litskevich, Lecturer in Nuclear Engineering**

Appendix: Survey Data Summary Tables

A.1 Respondent Profile

Category	n	%
Academic Staff	26	54.2%
Postgraduate Researcher	10	20.8%
Postdoctoral Researcher	9	18.8%
Undergraduate	3	6.2%
Engineering	8	16.7%
Chemistry	8	16.7%
Physics	5	10.4%
Computer Science	5	10.4%
Mathematics	4	8.3%
Other (interdisciplinary)	18	37.5%

A.2 Key Impact Statistics

Impact Statement	Strongly Agree	Agree (total)
HPC has reduced time to produce research results	**85.4%**	**91.6%**
HPC has enabled work that would otherwise be impossible	**68.8%**	**89.6%**
HPC has enabled more complex models	**75.0%**	**89.6%**
HPC has expanded the scope of research	**62.5%**	**77.1%**
HPC has improved output quality	**58.3%**	**77.1%**
HPC has enabled larger datasets	**56.2%**	**72.9%**

Note on response numbers. Sixty-seven survey responses were received. Forty-nine were completed in full and form the basis of the quantitative analysis; the 18 partial responses contributed data on platform use and career stage where available. One of the 49 respondents did not answer the profile and impact questions, so the figures and percentages for those questions are based on 48 responses and are labelled n=48; questions answered by all 49 are labelled n=49. Percentages are calculated on the number who answered each question and are rounded, so they may not total exactly 100%.