

Research Computing Strategy and Objectives 2018

This Research Computing Strategy document updates and replaces the version produced in 2016.

John Owen
Carol Sandys
September 2018

Table of Contents

Research Computing Strategy and Objectives 2018.....	1
Introduction	4
Structure of the Document.....	4
Formulating this Strategy	5
Strategy Statement	6
Chapter 1: University Research Strategy: Implications for Research Computing	8
Dubai	9
Chapter 2: The Current Context.....	10
Research Computing Needs of Birmingham Researchers	10
IT Services Support for Researchers	10
Chapter 3: Advanced Research Computing (ARC) and BEAR (Birmingham Environment for Academic Research) Overview.....	12
Governance.....	15
BEAR and Life Sciences.....	15
BEAR and AI.....	16
ARC Partnerships and Posts	17
Investment in BEAR and ARC	17
Chapter 4: Current Team Activities.....	19
Engagement	19
Data Management	19
Training	20
Building Communities of Practice/Common Interest	21
ARC/BEAR Events	21
Research Software Engineering	22
Software Applications	23
Training for Research Software Engineering: Software Carpentry.....	23
Data Science.....	24
AI and Deep Learning.....	24
Visualisation	24
Levels of RSE Support.....	24
Collaboration Across Borders	25
Architecture Infrastructure and Systems.....	25

Chapter 5: Technology Review	26
Introduction	26
Rationale	26
BEAR Infrastructure.....	26
Technology Outlook.....	27
Future Planned Technological Developments (2018-2020)	29
Chapter 6 : Research Support in Partnership with Other Areas of IT Services.....	31
Chapter 7: The National Scene.....	32
Notes from National eInfrastructure Survey 2017	32
Other Topics from the National e-Infrastructure Survey.....	33
The National Infrastructure Roadmap	36
Chapter 8: Aspects of a Successful Advanced Research Computing Team: People	37
Chapter 9 : Summary of Development and Implementation Objectives	42
Postface to the Second Edition	44
Glossary.....	45
Appendix 1 : Progress Report on the Objectives set out in the Research Computing Strategy 2016 (Published May 2018)	47
Appendix 2 : IT Needs of Birmingham Researchers.....	60
Appendix 3: ‘You Said, We Did’ – a Summary Report of the 2017 Research Computing Survey	63
Appendix 4 : The Consultation Process.....	69
Appendix 5 : The Advanced Research Computing Funnel	70
Appendix 6 : Research Computing Management Group Remit.....	71
Appendix 7 : Membership of the Research Computing Management Committee (August 2018)..	74
Appendix 8 : Technical Details of BEAR Infrastructure	75
Compute.....	75
Storage	76
Ancillary Services	76

Introduction

The first Research Computing Strategy was published in July 2016¹. It provided a snapshot of the University's research computing environment at that time set in both a local and national context. 27 ambitious Objectives were set aimed at consolidating and extending research computing services, strengthening our links with the research community and playing a key role in supporting the University's research ambitions. Good progress has been made against these objectives, which is summarised in Appendix 1. Other innovations have seen significant improvements to the services and solutions available to researchers.

The research community's views on research computing support from IT Services were assessed in April / May 2017. Whilst some core issues were highlighted, feedback on the BEAR services (Birmingham Environment for Academic Research) was generally good. These findings and the action taken in response are summarised in Appendix 2. A more detailed report was also produced² as was a widely available 'You Said, We Did' article³; this article is included as Appendix 3: 'You Said, We Did' – a Summary Report of the 2017 Research Computing Survey³.

The advanced research computing resources available through BEAR have developed and grown significantly in the last two years and Birmingham now has a reputation as a centre of excellence within the Sector for advanced services. There is much still to be done however as researchers' needs develop and expand.

Structure of the Document

A concise statement of the Strategy for Research Computing appears at the start of the document. A number of Development and Implementation Objectives which outline the forward looking ambitions of Advanced Research Computing are set within the text, and these are summarised in Chapter 9 at the end of the document.

The remainder of the document sets out in Chapters the context within which IT support for research computing operates and how the strategy and objectives have been developed. Research computing support is a rapidly changing and developing area and this additional background information is viewed as useful in positioning the current service offering and vision within a University, national and international context.

1

<https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/documents/public/website-revised-2015/survey/Research-Computing-Strategy-2016-v1m1.pdf?>

2

<https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/documents/public/website-revised-2015/survey/Survey-Overview-Report-v1m1.pdf?>

3

<https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/documents/public/website-revised-2015/survey/IT-Survey-YouSaid-WeDid-v1.3-Public.pdf>

Formulating this Strategy

This Strategy follows extensive consultation and meetings conducted by the Head of Advanced Research Computing and members of the team with colleagues throughout the University and at all levels, from PVC for Research and Knowledge Transfer, through College Directors of Research and Principal Investigators to research post-graduates.

The team has always enjoyed a good relationship with the Birmingham research community, and during the two years since publication of the last Strategy this has expanded and matured. Exploiting this strong team relationship with our user base has shaped the user-focussed development of this Strategy.

Inevitably this document concentrates on the services built around advanced architectures and specialist software engineering. Other associated services which are part of the general support provided by IT Services, such as desktop support, are described in Chapter 6 of this document.

The consultation process is summarised in Appendix 4.

Strategy Statement

The University has an ambitious yet achievable plan to raise our academic performance by growing research activity and increasing staff and student numbers over the next decade. This will involve:

- Increasing research awards by 10% per annum to £350m
- Securing a top 100 position in all major international rankings
- Recruiting 1,400 additional academic and research-focused staff
- Being among the UK's top 5 research intensive universities

BEAR services are the solid foundation on which many University research projects already rely. The Advanced Research Computing (ARC) team will maintain and develop these core services, applying the latest technologies to deliver reliably, securely and cost-effectively.

Working in collaboration with and guided by the needs of research groups from every College, cost-efficient solutions will be built primarily on infrastructure in our own data centres and capitalising on our significant in-house expertise both for service innovation and accessible support. With investment, BEAR will continue to evolve to provide the quality environment which a top 5 research intensive University demands.

We will do this by:

Growing our Partnerships with Researchers

- We will listen to our research base to understand the advanced computing tools they require to become leaders in their field. IT will be an enabler of these advances.
- We will work in partnership with our researchers to develop their specialist IT-related skills to accelerate and raise the quality of their research. In many areas IT competence will be an essential pre-requisite to success. This is perhaps no better illustrated than in developments in AI which already span disciplines, from Computer Science to Law, from the Arts to Genetics.

Further Developing and Growing ARC as a High Performing Team

- We value our people, and will deliver leadership to maintain and grow a flourishing and strong team with a user focussed ethos. Staff retention and careful selection are essential to delivering this.
- We will support and invest in our people, trusting in their commitment and expertise, and provide a flexible environment where they can respond effectively to the rapid pace of research-led demands.

Developing the BEAR infrastructure

- We will maintain and develop the quality and reliability of BEAR services that we have already established.
- We will deliver the best and appropriate leading edge facilities to keep Birmingham amongst the UK leaders in advanced computing facilities to support research.

- We will identify and engage with other areas of IT Services to deliver enhancements which rely on their underpinning services

Ensuring the Investment in BEAR and ARC are Sufficient to Deliver the University's Ambition

- We will maintain the current annual investment cycle so that BEAR continues to evolve, building on the success we have already achieved.
- We will lobby for increased investment in BEAR services which addresses inflationary pressures and, at a minimum, guarantees current service levels for the growing number of researchers attracted to Birmingham through the University's ambitious strategy.
- More ambitiously we will bring forward carefully analysed and costed proposals for service development which will support Birmingham's drive to be among the UK's top 5 research intensive universities.

Growing our Partnerships and Reputation beyond the Campus

- We will engage at a national and international level to stay on top of IT innovation and Sector developments, and select those which will provide the most benefit to the Birmingham community
- We will work to build effective partnerships with the NHS to facilitate delivery of the University's life sciences strategy.
- We will build meaningful and high level partnerships with our industrial and IT partners so Birmingham benefits from informed ambition and can influence the development of their products to our advantage.
- Selectively we will engage with our peers at other leading institutions and laboratories to ensure our researchers' collaborations receive the support they require.

Our success will be judged by the number of Birmingham researchers benefitting from and investing in BEAR services, by positive feedback through diverse channels, and by surveys and a selection of case studies demonstrating the contribution BEAR and ARC have made to quality research at Birmingham.

Chapter 1: University Research Strategy: Implications for Research Computing

The University has an ‘ambitious yet achievable’ plan to raise our academic performance by growing research activity and increasing staff and student numbers over the next decade. This is set out in ‘Birmingham 2026’⁴ which is the next stage of the process begun with the publication of the Strategic Framework 2015 – 2020⁵.

Focusing on research, the document sets out a vision for the University of Birmingham in 2026 which will be:

A more research intensive environment with at least 40% of our income from research, enabling a significant growth in research output and strengthening the translation of research into economic and societal benefit.

The document continues:

We will do this by:

- Investing in people and the intellectual environment
- Strengthening our research position

We aim to:

- Increase research awards by 10% per annum to £350m
- Secure a top 100 position in all major international rankings
- Recruit 1,400 additional academic and research-focused staff
- Be among the UK’s top 5 research intensive universities

Clearly IT, and Advanced Research Computing in particular, will continue to have a key role to play in making these goals a reality.

Building on the foundations of traditional High Performance Computing domains in science and engineering, demand from new areas, principally in the life sciences, has grown significantly within recent years. The exploitation of data and data storage are now major drivers as well as enablers across all research areas. The University is also now a partner of the Alan Turing Institute and, based on our experience within the last 12 month, Artificial Intelligence (AI) with its extensive computational demands will become the next major challenge for fit-for-purpose, centralised IT services. Through our recently announced strategic relationship with IBM (November 2018) BEAR will be well-equipped to deliver a globally-significant AI platform in support of the University’s ambitions, whilst continuing to meet our commitments to the existing user base.

The BEAR facilities provided by ARC are already used extensively with the rate of user registration increasing as new disciplines discover the value of advanced computing resources. 1,400 additional

⁴ <https://intranet.birmingham.ac.uk/strategic-framework/birmingham-2026.aspx>

⁵ [https://intranet.birmingham.ac.uk/planning/documents/public/Shaping-our-Future-\(PDF---5,61MB\).pdf](https://intranet.birmingham.ac.uk/planning/documents/public/Shaping-our-Future-(PDF---5,61MB).pdf)

academic and research-focused staff will inevitably place further demands on the resource, representing as it does a greater than 50% increase on the current potential user base. Careful planning and additional investment will be required if the University is to fulfil its promise to provide “outstanding facilities” and develop “the necessary underpinning infrastructure and technology to be world leading in these research areas”⁶

Objective 1. We will carefully monitor the usage of BEAR services, ensuring efficiency and matching demand against available resource. Where projections indicate demand cannot be met from current provision we will bring forward business cases to argue for additional resource.

Dubai

The University’s objective of developing a campus in Dubai opens up new and exciting opportunities for research but also poses many support challenges. The University’s IT infrastructure to support research is hosted almost entirely on the Birmingham campus. The strategy for delivering advanced computing services for research in Dubai has yet to be determined. Three principal approaches need to be considered:

- Deliver **remotely from Birmingham**. This has the benefit of drawing on the existing investment but depending on the pattern and topic of the research could require high bandwidth (and therefore expensive) network connections. Some local support will be required and licence conditions may impose constraints or additional costs.
- Deliver services **locally in Dubai**. Removes the networking constraint but requires investment in infrastructure, facilities and both wider and deeper skills in Dubai. Some system management could be done from Birmingham. Licence conditions may impose constraints or additional costs.
- Use **public cloud**. Whilst our current strategy is primarily to exploit our own campus infrastructure and expertise, including private cloud developments, our roadmap does anticipate the selective and managed use of public cloud. Public cloud offers advantages but also has limitations and risks, a cost-premium and needs to be managed effectively. Again licences conditions for software may have a significant impact in some cases as could the new support structures required.

It is likely that the eventual service model will be a combination of all three options to varying degrees.

Objective 2. When research requirements in Dubai are identified, we will contribute our expertise to identify options and potential solutions.

⁶ Ibid

Chapter 2: The Current Context

Research Computing Needs of Birmingham Researchers

Birmingham currently has approximately 3,000 research active staff who between 2015 and 2017 produced 15,100 papers. Their research support needs are diverse and requirements and opportunities may develop and change rapidly, driven by the success (or otherwise) of research grant applications, and University strategic initiatives such as the developments in Life Sciences. IT Services needs the skills, resources and flexibility to respond to these needs and to facilitate and not encumber research outcomes. The University needs to recognise and plan for the IT resource implications of strategic research initiatives.

Increasingly the University is involved in national and international collaborations, with both academic and industry partners. The IT resource to support research therefore needs to stretch beyond the confines of the campus to lower barriers and facilitate effective collaboration.

There are areas where IT Services does not yet have a sufficiently close relationship with research teams to fully understand and represent their needs. Exploration and development of needs and solutions will grow year on year as this Strategy is realized. The recent investment in engagement, training and RSE resource is already making a difference to many more research groups in an ever-widening spectrum of research disciplines.

IT Services Support for Researchers

Just as research is collaborative, so is the support for research. We work closely with various of the University's Professional Services; partnering with the Library Digital Assets Team who lead on data management and publications support, with Research Facilitators in Finance to support grant funding applications and with the Planning Office to support the REF process. Other collaborations stretch beyond the campus, and the Advanced Research Computing (ARC) team has established successful partnerships with our principal suppliers. Often these partnerships are at a very high level within these major companies, giving access at director level and to development laboratories to influence product development to Birmingham's advantage

Support for researchers from IT Services is principally delivered through two Sections – Advanced Research Computing, and End User Services. Additionally, the Application Services Section supports the associated administrative systems, including all Library systems. The division of responsibility, and other routes to support, are shown diagrammatically below. To ensure homogeneity of service all initial service calls are directed through the IT Service Desk, but inevitably a number of these calls develop into much more detailed consultations or even long-term partnerships.

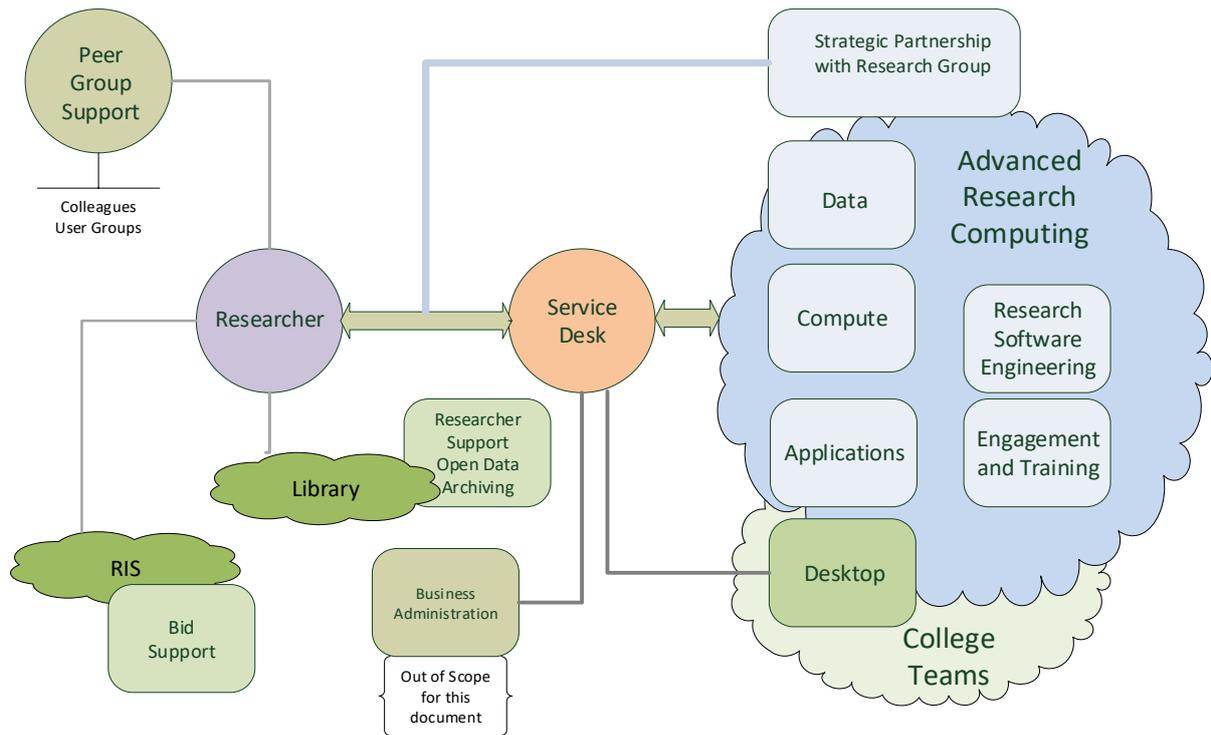


Figure 1: University Services for Researchers

Chapter 3: Advanced Research Computing (ARC) and BEAR (Birmingham Environment for Academic Research) Overview

Previously known as the Research Support Section, ARC works in partnership with Research Groups and Institutes to design, build, deliver and support a set of advanced and specialist services for researchers.

Available to all Birmingham researchers and their collaborators, collectively these services are known as BEAR (the Birmingham Environment for Academic Research). ARC's services are designed to meet the compute or data intensive needs of a broad and growing spectrum of research disciplines where desktop or local solutions are neither practical nor cost effective. As of June 2018, BEAR serves over 650 compute projects with over 3500 users registered for at least one BEAR service.

The past two years has been a period of significant investment in ARC both to support the award winning infrastructure⁷ and to expand our engagement with users. This expansion resulted from the clear demand for more of the 'excellent service' commended in the VCs Review of IT Services in 2015.

The team has added six new members in the last year. ARC now (September 2018) consists of 20 team members (6.5 funded by research groups⁸) organised in three Groups:

The **Engagement Group's** remit ranges from developing the Section's communications through all available media, developing and delivering a fully-featured training programme, fostering and facilitating the BEAR User Group and Special Interest Groups, identifying the requirements of the user base and supporting the delivery of services.

Architecture Infrastructure and Systems has responsibility for designing, building and operating the infrastructure and systems to deliver the specialist services needed by the University's researchers. They employ some of the latest technology to deliver what is now one of the best and most fully featured set of services for research available in a UK HEI.

The **Research Software Group** was formed in October 2017 with the goal of building a core team of Research Software Engineers (RSEs) (in line with one of the goals from the 2016 Strategy). The process of recruiting will continue throughout 2018 as the new Group takes shape and begins delivering services. The RSEs' primary focus is on supporting users to exploit the potential of computing in their research and to share best practice in software development in order to ensure reproducible and referenceable research. They have adopted the slogan of the national Software Sustainability Institute: "Better Software, Better Research."

⁷ HPC Wire Awards 2017

<https://www.hpcwire.com/off-the-wire/hpcwire-reveals-winners-2017-readers-editors-choice-awards-sc17-conference-denver/>

HPC Wire Awards 2018

<https://blog.bham.ac.uk/bear/2018/11/12/hpcwire-award-winners/>

⁸ By the Centre for Computational Biology (CCB) / Life Sciences (CaStLeS) initiative – 3 posts

<https://www.birmingham.ac.uk/research/activity/mds/centres/computational-biology/index.aspx> and

<https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/bear-cloud/index.aspx>

By the PRISM2/ CASIM research groups in Engineering and Physical Sciences (EPS)

<https://www.prism2.org/> and Centre for Advanced Simulation and Modelling (CASIM2).

By the CLIMB project in Life and Environmental Sciences (LES)

<https://www.climb.ac.uk/>

By the College of Arts and Law (CAL)

By Aston University to support the HPC-Midlands-Plus Tier-2 service

Leadership for ARC is provided by the Head of Section and Deputy. Note that the Head of ARC has a brief which extends beyond BEAR to wider areas of concern to the University’s researchers.

Although ARC has seen a healthy growth in funding for posts over the last two years, funding for the underpinning technology has remained relatively static in real terms. It is imperative that funding is adequate to meet the high expectations and growing demands of the University’s researchers, including the ability to develop further technologically advanced solutions. Over recent years, BEAR services have developed rapidly and have built a national reputation for both innovation and performance. Failure to invest means a real terms threat to service.

In order to minimise barriers to access and to limit the cost of administrative overheads, a meaningful default level of service is typically free at the point of use to all. Some users with particularly demanding workloads augment the standard service by funding additional resources.

BEAR exists within a wider eco-system of provision so some groups will have access to additional research community specific resources or, in very specific cases, access to national compute and subject specific repositories or resources such as HPC-Midlands-Plus⁹, ARCHER¹⁰, the European Bioinformatics Institute (EBI)¹¹, Elixir¹², and the 100k Genome project¹³. In the case of High Performance Computing (HPC), our on-campus BlueBEAR system is often the proving ground, without which these groups would be unable to gain access to these facilities. A summary of BEAR Services is provided in Figure 2.

The core BEAR services are therefore essential to new researchers and those with no or only modest funding and for those needing to demonstrate the value of their research to attract funding. In other words, BEAR services are vital tools to support the University in its aim to “lift the middle”.

The extent and breadth of work undertaken by the team is best highlighted via the ‘Advanced Research Computing Funnel’ which shows processes and services from conception, through feasibility and development and into ‘business as usual’ service. The funnel for June 2018 is presented in Appendix 5.

⁹ <http://www.hpc-midlands.ac.uk/>

¹⁰ <http://www.archer.ac.uk/>

¹¹ <https://www.ebi.ac.uk/>

¹² <https://www.elixir-europe.org/>

¹³ <https://www.genomicsengland.co.uk/the-100000-genomes-project/>

Figure 3: BEAR Services



BEAR Services in Summary

- [BlueBEAR](#), our Linux-based, batch processing High Performance Computing (HPC) cluster
- [BEAR Cloud](#), the next generation of powerful interactive virtual machines for compute or data intensive work - built for flexibility, accessibility & efficiency.
- [CaStLeS \(powered by BEAR\)](#), - powerful compute and storage to meet the challenges of life sciences research in all Colleges.
- [BEAR DataShare](#), offering dropbox-type facilities for collaboration both on and off campus
- [BEAR Research Data Store](#), providing storage for ‘work-in-progress’ data, including a facility for groups with exceptionally large requirements to purchase extra capacity. All data are kept securely on campus, replicated between two data centres and regularly backed-up. Each research project is allocated up to 3 TB of storage on request.
- [BEAR Archive](#), offering secure, tape-based archiving for the long-term retention of all valuable data associated with a completed research project (that is not required to support publication). Up to 20 TB is provided per project for up to 10 years; additional storage can be purchased if required. The previous [Research Data Archive](#) now provides the underlying storage for the UBIRA eData Repository, designed for the long-term retention of open data associated with publications. The UBIRA service is managed by the Library Digital Assets Team. BEAR funds and provides up to 1 TB of storage for each publication, with a default and guaranteed retention period of 10 years. An extended retention period may be requested, and additional storage can be purchased if required.
- [BEAR Research Data Network](#), catering for the fast transmission of exceptional data volumes, generated by specific scientific equipment on campus.
- [BEAR Gitlab](#), an on-campus software repository supporting academic programmers through the whole software lifecycle.
- [BEAR Software](#): advice, coaching and coding services from Research Software Engineers.
- [BEAR View](#), offering visualisation, collaboration and large scale video-conferencing facilities including a 4metre by 2metre display screen, 2 and 3-D imaging and movement tracking.
- [BEAR Plus](#), additional computational resources available to Birmingham researchers through [the Tier2 regional centres](#) and other national resources. (Note, some of these national resources may be paid-for services).
- [Exclusive BEAR](#), enabling research groups with demanding or time-sensitive workloads to buy centrally managed and dedicated compute capacity. Costs for management and all supporting infrastructure (racks, switches and of course electricity) are usually met by IT Services.
- [BEAR Database Service](#), runs in a high-availability cluster and is available both from the BlueBEAR cluster and from the campus network.
- [BEAR Galaxy](#) (for computational biologists) - a platform for data-intensive biomedical research (analysis and process sharing).

Governance

The development and operation of all BEAR services are overseen and governed by the Research Computing Management Committee. The Committee is chaired by a research active academic with members drawn from each College as well as the Team and reports to the Academic IT Reference Group, chaired by the PVC for Research and Knowledge Transfer. The Committee Remit and current Committee membership can be found on the BEAR web site¹⁴ and are reproduced in Appendix 6 and Appendix 7.

Reflecting the rapidly growing demand for data storage and priority to address research data management, the Committee also serves as the University's Research Data Storage Management Board, overseeing policy and provision of the centralised archive and 'working' data stores.¹⁵

This 'Research Computing Strategy for the University of Birmingham' is proposed by IT Services with the full support of the Research Computing Management Committee.

There is also a BEAR User Group chaired by one of the academic members of the Management Group.

Governance of the new Life Sciences investment (CaStLeS)¹⁶ is managed by a Life Sciences Executive, established to allocate processor and storage resources purchased by this award and take forward the strategy for the Life Sciences. The Executive is advised by a Strategic Oversight Group drawn from the Colleges of Medicine and Dentistry, Life & Environmental Sciences, and Engineering and Physical Science. It will provide regular reports to the Research Computing Management Committee and via them to the Academic IT Reference Group.

Objective 3. We will continue to provide Service Reports and Metrics to each meeting of the Research Computing Management Committee

Objective 4. We will provide Service Reports and Metrics to each meeting of the CaStLeS (Life Sciences) Strategic Oversight Group

BEAR and Life Sciences

Life Sciences are now placing heavy demands on BEAR services. The University's significant investment in the CaStLeS infrastructure has markedly strengthened and accelerated computational intensive research in this area. That initial tranche of kit however is approaching end of life and so to maintain momentum further investment in the next 12 – 24 months will be necessary.

ARC has also seen a growing demand in recent months to strengthen and develop our interaction with the NHS, notably University Hospitals Birmingham¹⁷. Patient confidentiality in this area is paramount.

¹⁴ <https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/bear-governance.aspx>

¹⁵ Research data storage policies can be found at <https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/research-management-policy.aspx>

¹⁶ CaStLeS – Compute and Storage for the Life Sciences <https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/bear-cloud/index.aspx>

¹⁷ <https://www.uhb.nhs.uk/home.htm>

Currently BEAR can only suitable for the analysis and storage of anonymised or pseudo-anonymised data, although this is not without risk as confidentiality breaches cannot be ruled out entirely. This may happen for instance by mistake, or when the patient pool is small (as is the case for certain rare diseases) and other data sources could help identify an individual. If potentially patient identifiable data is to be processed then significant additional security measures will be required.

Given the risks it would be wise to handle all patient data (anonymised, pseudo-anonymised or identifiable) within a separate security domain on campus. This concept is known as a Data Safe Haven (DSH), and will require significant investment to implement. UCL¹⁸ and Manchester¹⁹ have already implemented DSH's as have other consortium members within the Health Data Research UK initiative, (HDR-UK)²⁰. Having gained a foothold in HDR-UK²¹, Birmingham must take full advantage of the leading position it now has with respect to large, diverse patient datasets and longitudinal studies available through University Hospitals Birmingham NHS Foundation Trust. Implementation of a DSH is therefore an imperative. This will be a complex and costly development relying on input from many areas of IT Services. A DSH is also highly likely to be a pre-requisite for the successful development of the Birmingham Life Sciences Park (BLSP)²²

Objective 5. We will work to build effective partnerships with the NHS to facilitate delivery of the University's life sciences strategy

Objective 6. We will work with others to ensure the essential technology requirements and costs for advanced research computing at the Birmingham Life Sciences Park are understood

BEAR and AI

During 2018 ARC purchased and installed two of IBM's flagship PowerAI computing nodes²³. The success of this implementation coupled with ARC's growing international reputation and the University's ambitious research ambitions led to the establishment of a strategic relationship between IBM, ARC and the University. Initially, and by early 2019, this will deliver:

- A further 9 PowerAI water-cooled nodes making Birmingham the largest such site in the UK and amongst the leading sites in Europe;
- IBM-sourced training for novice and expert AI users, and the RSE community
- An IBM Client Technical Specialist (CTS) on site for one day a week for a year to engage with researchers and the ARC technical teams to facilitate optimum use of the AI resource.

A longer term collaboration (initially over 3 years) between IBM UK Research, IBM UK Development Teams and the University will be energetically explored.

¹⁸ <https://www.ucl.ac.uk/isd/services/file-storage-sharing/data-safe-haven-dsh>

¹⁹ <http://www.itservices.manchester.ac.uk/cybersecurity/programme/datasafe/>

²⁰ <https://www.hdruk.ac.uk/>

²¹ <https://www.birmingham.ac.uk/news/latest/2018/02/54-million-funding-to-transform-health-through-data-science.aspx>

²² <https://www.birmingham.ac.uk/research/activity/birmingham-life-sciences-park/index.aspx>

²³ <https://www.ibm.com/uk-en/marketplace/deep-learning-platform>

Initially, but not exclusively, the focus of these AI developments will be in the life sciences and computational biology area. ARC will engage with the newly established University Turing Fellows to ensure a wide audience enjoys the benefits of this partnership.

Objective 7. The IBM partnership marks a significant opportunity for AI development in Birmingham. ARC's programme will be an inclusive one to make these resources available to the widest possible audience. These developments will be aligned to and integrated with the University's AI and Data Science Strategy.

ARC Partnerships and Posts

Much of the work relating to Life Sciences is focussed through a strong and vibrant partnership with the Centre for Computational Biology²⁴ and Professors Jean-Baptiste Cazier, George Gkoutos and Ben Brown. Three posts within the team are funded through this link.

One post is part funded through the CLIMB²⁵ initiative, led in Birmingham by Professor Nick Loman.

Funding for one FTE comes from the Prism2 / Casim²⁶ group in EPS, led by Professor Jeff Brooks.

A further post is financed by the College of Arts and Law.

The Tier-2 national network²⁷ has led to a partnership with the University of Aston and a shared Research Software Engineer post.

Beyond Birmingham ARC has established senior level partnerships with our main suppliers, principally Lenovo and IBM.

Investment in BEAR and ARC

It is imperative that adequate funding for BEAR continues. The advanced computing facilities that BEAR provides are now a fundamental component of the well-founded laboratory expected by research councils.

A default allocation of resource, be it compute or storage or software support, adds competitive advantage to research bids. Whilst the budget for this comes through the institutional allocation made to IT Services, the funding is provided by the 'indirect costs' levied on all research grants. So whilst we may colloquially refer to these services as 'free at the point of use', research activity has already paid for them.

²⁴ <https://www.birmingham.ac.uk/research/activity/mds/centres/computational-biology/index.aspx>

²⁵ <https://www.climb.ac.uk/>

²⁶ By the PRISM2/ CASiM research groups in Engineering and Physical Sciences (EPS) <https://www.prism2.org/> and Centre for Advanced Simulation and Modelling (CASiM2).

²⁷ <https://epsrc.ukri.org/research/facilities/hpc/tier2/>

Centralisation, coupled with dialogue with the research community, provides an efficient and effective way of delivering services. ARC can bring both a high degree of expertise to system design and economies of scale which ensure Birmingham's researchers benefit from the best possible facilities.

Objective 8. We will maintain the current annual investment cycle so that BEAR continues to evolve, building on the success we have already achieved.

Objective 9. We will lobby for increased investment in BEAR services which addresses inflationary pressures and, at a minimum, guarantees current service levels for the growing number of researchers attracted to Birmingham through the University's ambitious strategy.

Objective 10. More ambitiously we will bring forward carefully analysed and costed proposals for service development which will support Birmingham's drive to be among the UK's top 5 research intensive universities.

Chapter 4: Current Team Activities

Engagement

The VC's Review in 2015 highlighted the need for greater engagement with the research community by ARC on the basis of the potential for significant benefits to research. In response to this recommendation and the analysis of the Cubane Benchmarking exercise in 2016, the Research Engagement Group was formed to spearhead the corresponding mission.

There are clear challenges for effective engagement not only around messaging and media but also around making BEAR services accessible to a broad cross-section of University users. Alongside these immediate and very practical hurdles, the Group is also tasked with finding routes to raise awareness of the potential of these computational tools for many more disciplines which previously would not have considered employing them. The growing publicity and awareness around Artificial Intelligence (AI) is clear and acting to open doors for initial discussions, as is the drive for better data security.

As the BEAR service offering expands and adds functionality, the Group is increasingly busy and welcomed as they seek to introduce BEAR to potential users and listen to the needs and ideas of researchers. We will maintain and develop our communication channels building on the effective strong based of talks, special events, drop-in sessions, the web site, blogs and tweets.

The results of their labours are reflected in the increasing user base for BEAR.

Objective 11. We will continue to promote BEAR across all Colleges and grow the user base in line with requirements using a variety of methods aimed at reaching the widest possible audience.

Equally important to researchers are other IT Services developments which affect their working environment. These may impact on a personal level (for example characteristics of a VPN service which enable off-site access to resources) or at a group level. Some research groups run complex IT systems in their own right and changes to core components such as firewall policies or the availability of advanced features such as IPv6 can have significant impact on their activities. Early engagement will improve satisfaction levels and service quality.

Objective 12. We will work with our colleagues in IT Services to promote the wider portfolio of services which are of benefit to researchers. This will range from appropriate and early, collaborative consultation through to widespread dissemination of service characteristics and benefits.

Data Management

The Group is the focus for support and outreach activity relating to data management, drawing on skills and resources from other members of ARC and other Professional Services, most notably the Library. They subscribe to and promote the FAIR data management principles – that data should be Findable, Accessible, Interoperable and Reusable²⁸.

Data storage is provided through the RDS or BEAR Archive resources.

²⁸ <https://www.force11.org/group/fairgroup/fairprinciples>

Training

One core and supporting deliverable of the Engagement Group is to co-ordinate a BEAR training programme and ensure it co-exists and complements the training provided by other units e.g. the Library, People and Organisational Development (POD) or the Graduate School (and, in the future, the possible Data Science and Artificial Intelligence Institute). At the same time a goal is to work with these other units to ensure all relevant training from whatever source is made easily identifiable and accessible to researchers. Where appropriate, the Group contributes to training sessions run by others e.g. the Library's Raising Your Research Profile training.

The Team continues to co-ordinate delivery of the long-standing foundation sessions 'Introduction to Linux' and 'BEAR Necessities' but at a much increased frequency in response to clear demand. The appointment of RSEs now means a set of Software Carpentries²⁹ courses are scheduled, using both certificated ARC instructors and trainers from Colleges. These Carpentries courses have proven to be extremely popular – the first one was full just 1 hour after opening registration, whilst the second reach capacity even more quickly – in just 13 minutes. Due to the popularity, ARC will fund another year of membership of the Carpentries, allowing for a further 6 instructors to be trained. Data Carpentry³⁰ and HPC Carpentry³¹ will also be offered.

Informal 'drop-in' sessions are held weekly at locations around campus. These are open to all research staff and students. Whilst questions on any BEAR related topic are welcomed, to-date the bulk of enquiries have concerned data, data management plans, and secure storage

Following soon will be Deep Learning courses from our newly qualified Nvidia Deep Learning³² Ambassador, with developing plans to add both the HPC and Data Management course during 2018/19. Training makes significant demands on ARC time and necessarily must be balanced against other activities including building new solutions and general user support. However, training is viewed as a generally good investment both to empower researchers and to reduce the number of the most basic 'how to' support calls.

A significant challenge for the Group is to locate suitable training venues. Because of course timing and duration student PC clusters are usually unavailable. Activities have relied heavily on the generosity and flexibility of the CCB in making their training room available.

ARC receives no specific funding for researcher training. All current activity is therefore funded from the infrastructure budget.

Objective 13. We will continue to promote and grow the BEAR-related training courses, publishing courses on the BEAR web site, and maintaining attendance and demand statistics

Objective 14. We will expand the curriculum of training courses, responding to demand from the Colleges. A current priority is for Data Science and AI training and requests in these areas are expected to grow significantly.

²⁹ Software Carpentry teaches researchers the computing skills they need to get more done in less time and with less pain. <https://software-carpentry.org/about/>

³⁰ <https://datacarpentry.org/>

³¹ <https://hpc-carpentry.github.io/>

³² <https://developer.nvidia.com/deep-learning>

Building Communities of Practice/Common Interest

By building communities and facilitating networking events, the Group supports fora (both physical and virtual) to connect researchers with similar needs or interests. These may focus around specific software applications through Special Interest Groups (SIGs) or around a community such as the Academic Programmers SIG or arranging bespoke events for particular groups e.g. bringing a research group together with the experts in ARC or introducing ARC to the Digital Humanities Programme in CAL (CAL Digital).

Through harnessing the collective experience and knowledge from diverse areas, understanding and good practice are built and may be documented for wider dissemination and benefit. The learning too completes the feedback loop to influence the design and implementation of the next iteration of the service.

Currently active SIG's are:

- Matlab
- Finite Element Analysis
- Computational Fluid Dynamics
- Stata
- Academic Programmers
- Bioinformatics

- A Data Science / GPU SIG is being planned

'The Hacker Within³³', is a world-wide peer learning group for sharing skills and best practices for research computing and data science. The Birmingham Chapter is jointly supported by ARC and members of the research community and meets monthly to share topics useful in our data analysis and software development workflows.

The Engagement Group works not only with research groups and researchers but increasingly with other Professional Services units engaged in supporting research. The relationship with the Library is well established and productive. Areas targeted for investment of more effort include:

- College Research Support Services – supporting grant applications and the creation of Data Management Plans
- The Graduate School
- CAL Digital - developing awareness and exploring the application of technologies, including BEAR services, for the Digital Humanities
- IT Services' College-based teams
- New staff and post-graduate induction programmes.

ARC/BEAR Events

The Engagement Group is now routinely involved in events such as

³³ <http://www.thehackerwithin.org/UoB/>

- Post Graduate Open Days
- New staff briefings
- New Post Graduate Research (PGR) induction/briefings (by invitation)
- School / Department or research group events
- The University Research Conference
- Birmingham Fellows induction days

As ARC's influence grows beyond campus, the Engagement Team is also called on to organize events to showcase BEAR and its underpinning technology, all in the context of its purpose to provide more and better tools for research. Not only does this provide a showcase for innovation and service excellence at Birmingham but also attracts influencers and decision makers, as well as other institutions keen to learn from ARC's experience, for example UCL, Oxford, and Cardiff.

Research Software Engineering

Recent years have seen a growth in the recognition of the importance of software and software engineering in the research process – 'better software, better research'³⁴ Whilst for a long time coding has been a constant of most research processes, recent growth has been driven on two fronts: the growing availability of (big) data; and a recognition that 'bad software' means 'bad research'. There are numerous examples of mistakes in the analysis leading to incorrect conclusions, and where results cannot be reproduced^{35, 36}. Despite its importance, evidence suggests that many people developing software in research environments have no formal training, are self-taught and struggle with the conflicting time pressures of completing a research project and testing software for accuracy and reproducibility. Code is written for a 'one-off' job and so never contributed to the wider community; many people are likely to have developed code for the same task over the years, wasting effort. Of course if a researcher is adopting someone else's code they need to be sure of its accuracy.

The role was well described in the preface to the 2017 'Research Software Engineers: State of the Nation Report'³⁷ published by the Software Sustainability Institute:

Most research would be impossible without software, and this reliance is forcing a rethink of the skills needed in a traditional research group. With the emergence of software as the pre-eminent research tool used across all disciplines, comes the realisation that a significant majority of results are based, ultimately, on the skill of the experts who design and build that software.

Research software engineering provides a framework for accurate and reproducible data analysis. This has led to the creation of the Research Software Group (RSG) populated by Research Software Engineers.

The Research Software Group's goals are:

³⁴ Software Sustainability Institute www.software.ac.uk

³⁵ "Over half of psychology studies fail reproducibility test" M.Baker, Nature <https://www.nature.com/news/over-half-of-psychology-studies-fail-reproducibility-test-1.18248>

³⁶ [10.31219/osf.io/z48cm](https://doi.org/10.31219/osf.io/z48cm)

³⁷ https://zenodo.org/record/495360#.W2QT_WyWzF9

- Enable the University of Birmingham's research community to get the best from their research software.
- Provide specialist software engineering advice and support to researchers and research software engineers (RSEs)
- Help to enhance the University's reputation for high quality research
- Help researchers get the most from BEAR services, maximising the return on the University's investment in BEAR.

The Group does this by providing advice (generally or specifically on BEAR services) and coaching/coding sessions.³⁸

Software Applications

In the background the Group build and maintain software applications on BEAR (BlueBEAR, BEAR Cloud and CaStLeS) – more than 300 by July 2018, and rising. The Group responds to many and ongoing requests to mount software on these platforms.

Much of the software base is open source or provided with similar terms and conditions and so can be deployed without incurring additional licencing costs. Other software however must be purchased. ARC has a modest software budget which is used to maintain a core portfolio of software which has benefit to a broad user base. This core changes little from year to year, with little opportunity to retire programmes and consequently little budget to add new software. On occasion a research group or groups will find funds to purchase software licence, and it is sometimes possible to make this available to a wider audience, for example if a site licence is cost effective. A major institutional initiative in recent years saw the Colleges come together to fund a Matlab site licence. This is a very popular and comprehensive software suite and has delivered significant benefit to students (undergraduate as well as postgraduate) and research staff.

Maintaining the core BEAR software portfolio while dealing with increased licencing costs (almost invariably tied to the number of direct users or members of the University) will be a growing challenge going forward.

Objective 15. We will maintain our dialogue with the user base to ensure the software portfolio remains relevant and delivers broad value. We will also lobby for increased funding where the developing ambitions of the University's research strategy leads to increased software licencing costs.

Training for Research Software Engineering: Software Carpentry

RSG makes a significant contribution to the ARC training programme, in particular by supporting and delivering the Software Carpentry initiative.

³⁸ <https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/rsg/research-software-group.aspx>

Data Science

The concept of “Data Science” is changing from something that a small number of particular researchers would do, to being something in which virtually all researchers across the University will need to be proficient. This encompasses using computers to extract knowledge from both structured and unstructured data. It is no longer enough to employ people to do this - all researchers will soon need to be able to do it themselves. As a result, our RSEs need to be able to train and coach researchers in these core skills. Likewise the RSEs themselves will require ongoing training in this rapidly developing area.

AI and Deep Learning

One member of the RSG has qualified as an NVIDIA³⁹ University Ambassador, and Deep Learning instructor. We will continue to invest in training our RSEs in this highly important growth area. ARC has recently invested in IBM Power AI systems. The RSG will help the research community to maximise the use of these powerful machines for their research.

Visualisation

This RSG plans to help researchers to find the best way to visualise their research data, thus providing them with the best opportunities to make breakthroughs in their research. This visualisation may involve simple graphing of data, large screen 3D displays, or the use of Virtual Reality or Augmented Reality hardware - particularly as such hardware is expected to dramatically reduce in price over the coming months and years.

Levels of RSE Support

Funding for all the current members of RSG is guaranteed. This core group allows us to offer support up to a maximum of 20 half-day sessions for coding and coaching.

We envisage as the next phase of the service that we will also work with PIs who need significant RSE support in preparing their bids to include RSE time, for example 40% FTE for 3 years. We would then combine such bids into RSE posts within the RSG. If possible we will create a suitable amount of permanent RSE posts and then allow for buy-out of their time on research grants. This will allow us to retain a team of skilled and experienced RSEs from project to project, without having to fund them centrally - alongside the core of centrally funded RSEs. This model is clearly scalable and will allow the group to grow to whatever size is needed to support such work. In the medium term we may need to offer fixed term roles to match the initial grant funding, with the intention of turning these into permanent posts as the pipeline of such funding becomes established.

Objective 16. We will develop the RSE service with the aim of creating a sustainable pool of expertise available to commit to medium to long term software engineering engagements that are essential to research projects

Objective 17. Subsequently, we will work with PIs and Research Support Services to include funding for RSE time within research funding bids.

³⁹ www.nvidia.co.uk

Research Software Engineering is a rapidly growing field where good staff are scarce and therefore in high demand. RSEs require breadth and depth in their knowledge and an ability to relate to a variety of researchers with differing skill levels to determine and deliver what is required.

Collaboration Across Borders

In many respects the development of the RSE community mirrors that of the Open Source community. Developments are not for profit and collaboration benefits everybody by improving quality and reducing effort. A 'research software engineering' community has existed since computers began to be used in research, but not under that name and not initially recognised as a community. Formalisation of the RSE role began in the UK but is now a global initiative. We were extremely pleased to host the International RSE Conference in Birmingham in September 2018⁴⁰. Even with 300+ delegate places available the conference sold out by 1st August 2018.

Objective 18. Through the RSG we will work with researchers from across the University to deliver 'better software', and expand our skill sets to support the burgeoning growth of AI and Data Science.

Objective 19. We will seek appropriate funding for the RSG through both central allocations and contributions from research grants.

Objective 20. We will provide annual reports on activity to the Academic IT Reference Group

Architecture Infrastructure and Systems

Architecture Infrastructure and Systems take the lead on technology developments for BEAR. This is highlighted in the following Chapter.

⁴⁰ <https://rse.ac.uk/conf2018/>

Chapter 5: Technology Review

Introduction

Advanced Research Computing provides a wide range of services, from high performance computing and storage to software engineering coaching and training. These services are only possible due to the underlying infrastructure designed and implemented by the team in association with our partners OCF⁴¹, Lenovo⁴², Mellanox⁴³, IBM⁴⁴ and other key partners.

BEAR services have already been described. Appendix 8 gives more detail of the technology which underpins the services.

Rationale

To enable the services to be developed and scaled over time, we use framework agreements. These give us stability and predictability when purchasing and, in turn, enables the development of technical partnerships with vendors which would not be possible via other means. For example, the team have direct access to the development teams at Lenovo and are closely engaged in the planning for development of new platforms, enabling the University to influence the design of next generation systems.

The ability to provide and sustain the infrastructure is essential, and as such we do not aim to have the largest UK deployments but, instead, to have a diverse and responsive system that caters for the needs of the broad community of researchers at Birmingham. However, the continuing growth in demand for capacity and new applications of technology (e.g. deep learning and cognitive computing) mean that constant growth and development are required.

Part of the BEAR hardware strategy is to invest in new developments to enable researchers to have access to cutting edge technology and thus the tools to support the delivery of new research. This will always be a small component of the overall infrastructure investment. Aligned with this, we expect our teams to be continuously learning in order to be able to adequately support these new technologies.

BEAR Infrastructure

BEAR Services are designed to be capable of expansion, rather than taking the approach of some sites to buy large HPC clusters as a one-off exercise and renew them every 5 years. Our strategy is to design and implement incremental services and systems which can be grown over time. This enables spending to be profiled over multi-year purchases and our services to be highly tailored to evolve in response to the needs of the research community.

⁴¹ <http://www.ocf.co.uk/>

⁴² www.lenovo.com

⁴³ <http://www.mellanox.com/>

⁴⁴ www.ibm.com

In addition to core funded activity, researchers are able to invest in resource as part of the BEAR services: this could be access to compute resources, specialised resource or additional capacity. Resources purchased by research groups are specified, procured and managed by Advanced Research Computing as part of the service. By specifying the hardware platform, this enables integration into the systems and provides the team with a standard management platform. A resource rather than hardware is provided, and guaranteed for a full five-year lifetime (regardless of any failure that might impact individual pieces of equipment). This is beneficial to the research group as warranties sold with new hardware components are typically limited to 1 or 3 years.

High speed interconnects and networking are enabled across the BEAR hardware estate, this is to help ensure:

- Network capacity is not a barrier to resources
- Fast access to data sets
- Flexible use of systems for different classes of workload
- Ability to move systems between resources

The overall architectural intent is to provide fast integrated services wherever possible, for example research data storage is accessible across all BEAR compute and private cloud platforms as well as from client desktops and devices. This closely integrated approach underpins the strategy of a campus-based infrastructure. Many research areas generate large and growing datasets and moving these volumes of data to a distant compute resource is inefficient and can be costly.

Technology Outlook

Over the past few years, we have seen CPU core density increase dramatically which has led to a corresponding increase in Thermal Design Power (TDP)⁴⁵ for the CPU sockets. To gain a competitive advantage, we have elected to implement a water-cooled strategy for all systems. This takes the form of both direct-to-node (where water is taken into the systems and across the CPU etc) water cooling for high-density compute systems and rear door heat exchanger units for additional heat-recovery and where direct-to-node is not possible. This has seen the University of Birmingham be the first site in the UK to deploy modern direct-to-node water cooled HPC systems with Lenovo as well as the first site in Europe to deploy IBM's water-cooled Power9 AI systems. Over the next few years, there are significant increases in density expected to high-bin CPUs as used for HPC compute and we believe we are well placed with our research-focussed data centre to accommodate these future systems and technologies. There is a growing trend from other vendors to deliver these types of high-density systems, and our water-cooled data centre has been designed with the flexibility to accept a variety of platforms from different vendors.

In addition to Intel's processor technology, there are systems from IBM and AMD starting to ship and significant likelihood of ARM⁴⁶ based systems becoming an HPC contender over the next few years. It is planned to keep abreast of these technologies and systems and small-scale evaluation used strategically to evaluate the performance and cost implications (both hardware and management overhead) given the workload presented by Birmingham's researchers.

⁴⁵ The TDP is the maximum amount of heat generated by a computer chip or component

⁴⁶ <https://www.arm.com/>

Traditionally GPU⁴⁷ usage at Birmingham has been low, however there are strong trends towards additional requirements from the emerging deep-learning and cognitive areas (e.g. Alan Turing institute). Investment has traditionally been limited in this area, however we have recently invested in 2 IBM Power9 AI systems (as used in the SUMMIT HPC system⁴⁸ - which as of June 8, 2018 was the fastest supercomputer in the world) to better support these workloads. We anticipate additional GPU applications and demands over the coming years and are working with both IBM and Lenovo to investigate better water-cooled solutions to enable efficient heat recovery. The implementation of GPU systems is likely to have demands on the storage infrastructure and the use of NVMe⁴⁹ (Non-Volatile Memory Express) to support the workloads is expected to be required – this is to ensure the Return on Investment on GPU systems can be maximised.

Spinning disk technology is likely to continue to play a significant role in supporting large volumes of research data. Whilst NVMe and flash technologies are becoming main-stream, the cost per TB is still significantly higher than spinning disk. In addition to this, given the volumes of data being generated, changes to the way data is stored and tiered are likely to be required. The current use of replication has both performance and cost implications and the level of availability for research data is different to business systems requirements. Over the past few years, there has been much discussion around the use of object storage, however few scientific workloads are well developed to use object and so the traditional POSIX⁵⁰ access methods are likely to remain for some time. This is even recognised by public cloud providers who are implementing such front-ends to their storage systems.

Aligned to this is the growing use of public cloud widely in IT infrastructures. As part of the Research Councils' (RCUK) Working Group for Cloud^{51,52} Birmingham has been well connected with understanding the challenges of public cloud. In many cases where the systems are busy, the cost advantage of public cloud is diminished - cloud providers' business model is to ensure their systems are fully utilised and so running an already full system in the cloud is not cost effective. There are also numerous well documented scaling issues with HPC workloads in public cloud. We do however plan to investigate and develop the capability to make strategic use of public cloud for HPC workloads to run in addition to the on-site capacity. This is anticipated to take the form of bursting specific workloads, in addition to the utilisation of specific technologies (FPGA⁵³, GPU). In general, the use of FPGA technology is not viewed as a high demand for Birmingham, in part due to the steep learning curve required to utilise the technology.

Objective 21. We will continue to monitor and evaluate the role public cloud can play in supporting researchers at Birmingham and adopt our approach so that the right solution is delivered to maintain competitive advantage.

⁴⁷ Graphics processing unit

⁴⁸ [https://en.wikipedia.org/wiki/Summit_\(supercomputer\)](https://en.wikipedia.org/wiki/Summit_(supercomputer))

⁴⁹ https://en.wikipedia.org/wiki/NVM_Express

⁵⁰ <https://en.wikipedia.org/wiki/POSIX>

⁵¹ <https://cloud.jiscinvolve.org/wp/>

⁵² <https://cloud.ac.uk/>

⁵³ Field Programmable Gate Array https://en.wikipedia.org/wiki/Field-programmable_gate_array

To support cloud services and to enable researchers needs for access to data, it is anticipated that a Science DMZ⁵⁴ needs to be developed on campus. This is likely to be progressive work based on the initial Research Data Network design but is expected to be managed by the Network Team in IT Services. This is an area where Birmingham is significantly lacking compared to peer institutions and the requirements of this are not currently well understood. Based on discussions with peer sites, implementing a 100Gbs Ethernet ring to enable high speed interconnect internally with appropriate firewall-bypass technology is likely to be required. As Advanced Research Computing are regularly engaged with researchers, it is essential they are involved in the development and design of such a facility to ensure it is appropriate and fit for purpose. It is also likely that a less corporate approach to technology can be implemented, for example the use of other brand and white-box technology may be able to deliver value for the system. The implementation of SDN⁵⁵ Layer 2 tunnels and IPv6⁵⁶ technology are likely to be required to enable connections to partner institutions.

Future Planned Technological Developments (2018-2020)

Research Data Storage – the current hardware is almost five years old and is filling rapidly. During autumn 2018, the infrastructure is planned to be replaced and as part of this, the architecture for data placement and replication will be reviewed in conjunction with discussion with researchers. Initial discussions have already indicated that there is appetite within the research community to allow non-replicated storage (but higher capacity) to be implemented, provided there is adequate backup and Disaster Recovery capability in place. As part of this, we are considering the use of erasure coded object storage as a backing store to traditional POSIX access. An object layer will also enable support for direct access to those requiring it who are starting to cloud-enable their applications.

CaStLeS – this storage and compute solution is key to supporting the growing life-sciences workloads and investments at Birmingham. The hardware will require refresh over the period and careful review of changing requirements is necessary to ensure systems that are purchased are appropriate for the changing workloads. For example: additional HPC compute or GPU resource may be required. We have found the use of BEAR Cloud Virtual Machine (VMs) effective at on-boarding researchers to the facility and in many cases have been able to assist them with moving from VMs to HPC batch compute.

BlueBEAR – as part of the standard strategy, BlueBEAR will continue to be invested in over-time and will see older generations of hardware drop out. There is likely to be regular skips over processor technology (for example missing Intel Skylake⁵⁷ and moving to Cascade Lake⁵⁸). Small scale evaluation of other chip technologies is likely as these enter and mature in the market.

BEAR Cloud – development of technologies to better support high throughput computing and accelerator access (e.g. GPUs) into VMs is anticipated. We plan to implement self-service provisioning of VMs to end users and expect to provide a small standard allocation to any researcher as “sand-pit”

⁵⁴ DMZ: In computer security, a **DMZ** or **demilitarized zone** (sometimes referred to as a **perimeter network**) is a physical or logical subnetwork that contains and exposes an organization's external-facing services to an untrusted network, usually a larger network such as the Internet
[https://en.wikipedia.org/wiki/DMZ_\(computing\)](https://en.wikipedia.org/wiki/DMZ_(computing))

https://en.wikipedia.org/wiki/Science_DMZ_Network_Architecture

⁵⁵ Software Defined Network https://en.wikipedia.org/wiki/Software-defined_networking

⁵⁶ <https://en.wikipedia.org/wiki/IPv6>

⁵⁷ [https://en.wikipedia.org/wiki/Skylake_\(microarchitecture\)](https://en.wikipedia.org/wiki/Skylake_(microarchitecture))

⁵⁸ https://en.wikichip.org/wiki/intel/microarchitectures/cascade_lake

systems, with larger requests requiring approval. In addition to this, we are reviewing the use of container technology to support easier workload deployment and the ability to move towards better “reproducible science”

Objective 22. We will maintain and develop the quality and reliability of BEAR services that we have already established.

Objective 23. We will deliver the best and appropriate leading edge facilities to keep Birmingham amongst the UK leaders in advanced computing facilities to support research.

Objective 24. We will maintain close relationships in the IT industry, keep abreast of technology developments to ensure we are in a good position to take advantage of developments. To support this we will allow time and invest in opportunities for the team to expand their knowledge and skills.

Chapter 6 : Research Support in Partnership with Other Areas of IT Services

ARC relies on partnerships and services delivered by other areas of IT Services:

Park Grange (Research) Data Centre

The bulk of the BEAR infrastructure now operates from the new Park Grange Data Centre. Bringing this facility on-line was the culmination of many year's effort and design work by members of ARC, and a significant financial demonstration of trust and commitment by the University in ARC's work to supporting research activities.

Operational responsibility lies with colleagues in Core Infrastructure.

Objective 25. The water-cooled Park Grange data centre is highly energy efficient. We will measure and monitor this efficiency to support Birmingham's commitment to sustainability.

Core Infrastructure

The Core Infrastructure Section includes within its remit responsibility for the on-campus networks (excluding the Research Data Network), off-campus gateways and the Wifi network. All are of course of importance to the research community.

Basic web site infrastructure for research groups is also provisioned through Core Infrastructure.

End User Computing

As previously mentioned End User Computing incorporates the College-based teams and concentrates principally on desktop-level support.

Customer Experience and Communications

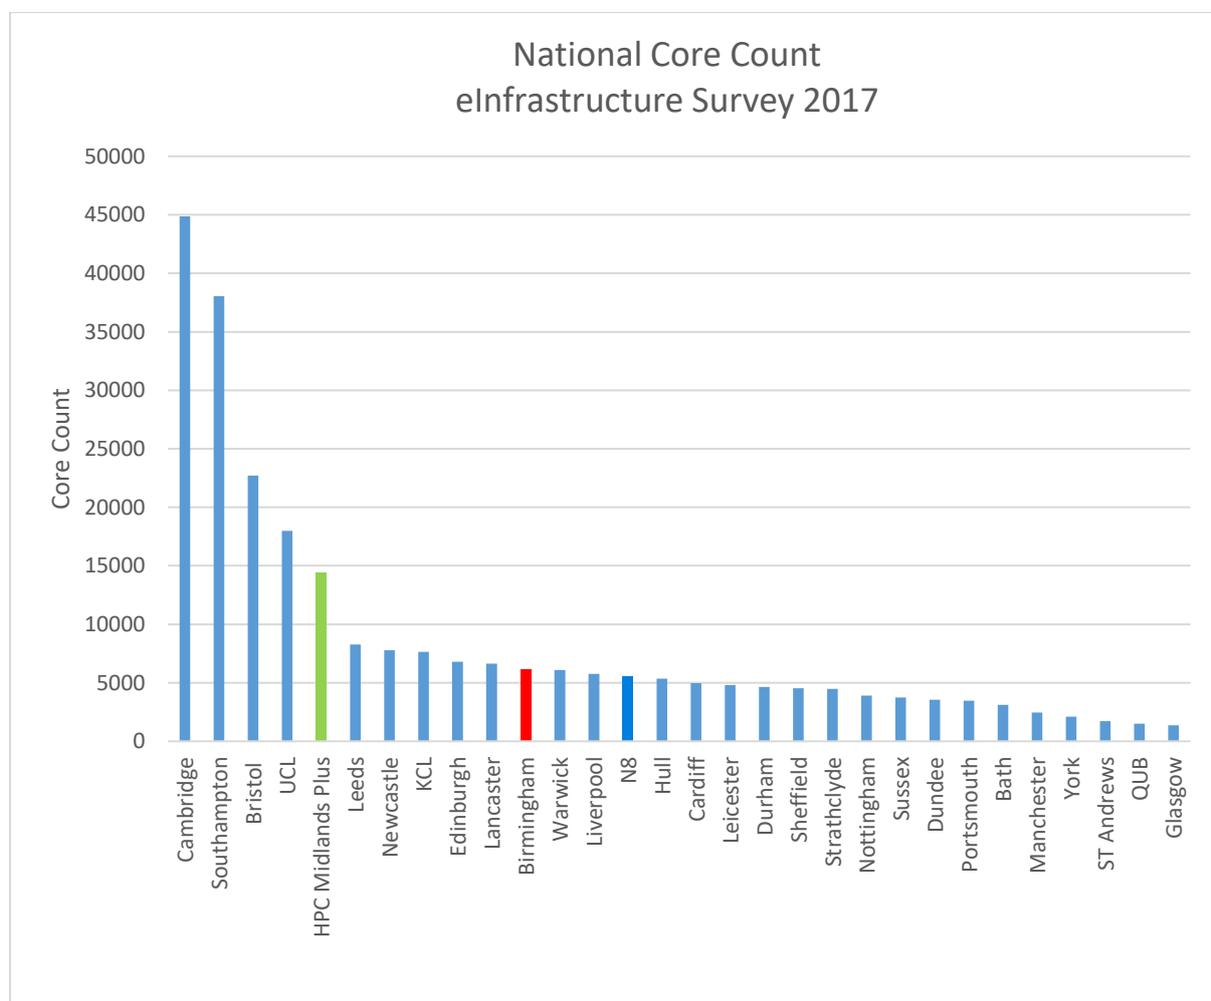
Enhanced communication of the broader range of IT services of benefit to researchers will be channelled through this group.

Chapter 7: The National Scene

Notes from National eInfrastructure Survey 2017

Most research active HEIs in the UK who run High Performance Computing systems are member of the national HPC Special Interest Group⁵⁹. Each year this body conducts a wide-ranging survey of its members, looking at installed infrastructure, site priorities and issues, and sector trends.

An indication of the size and (by association) investment in an Institutes' advanced computing resources is the number of cores⁶⁰ in its HPC system. There are of course limitations in such a metric – a large number of very old cores is likely to be less performant than fewer, newer cores - but it does provide a basis for comparison. The 2017 figures for core counts are illustrated below:

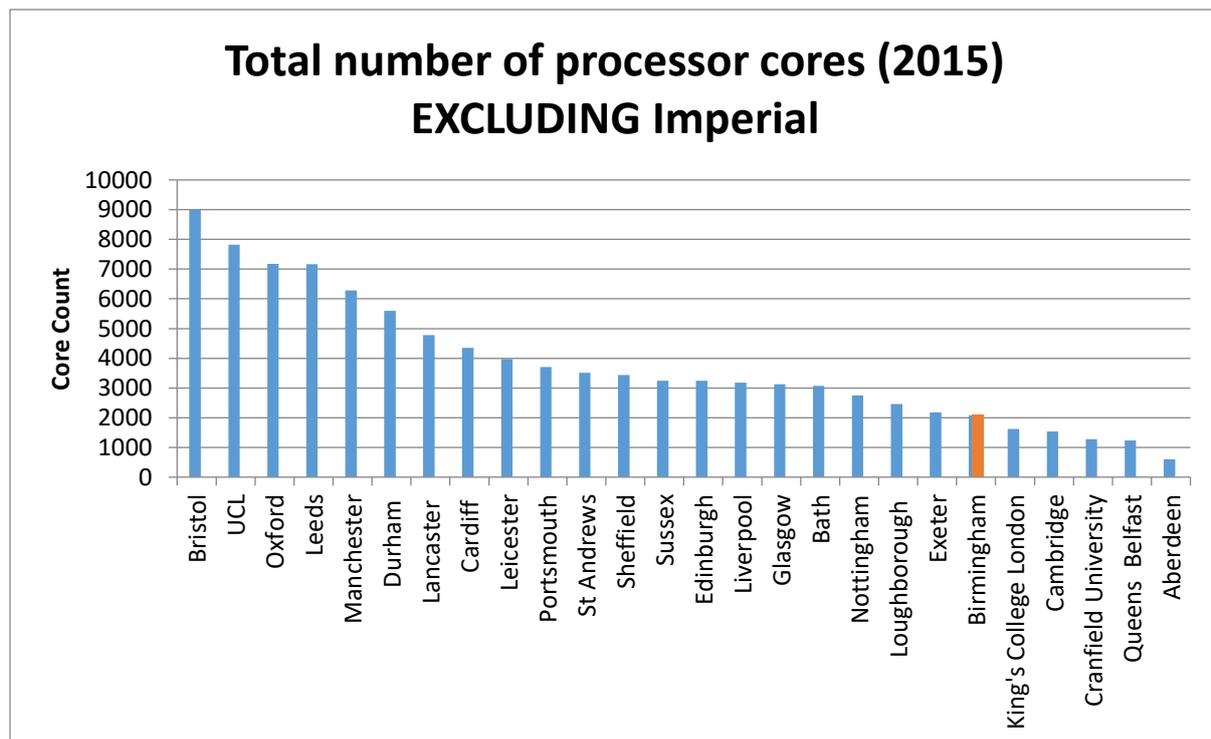


Not all sites consistently respond to the survey, and there may be variation in the way cores are counted (for example, are departmental systems included in the overall core count?) but from these

⁵⁹ <https://hpc-sig.org.uk/>

⁶⁰ <https://www.computerhope.com/jargon/c/core.htm>

data Birmingham lies within a fairly close knit pack centred around the middle of the table – in actual ranking we are 11th out of 30 returns. (In the 2015 survey Birmingham was 22nd out of 27 returns with the distribution shown below). Some sites list Tier 2 facilities separately (HPC Midlands Plus and the older, previous generation N8 facility) whilst others (Cambridge) group them in with their campus resource (this in part explains why Cambridge jumped to number one position in the current survey, from 24th in 2015.) Imperial, who were by far and away the leaders in 2015, did not reply to the 2017 survey.



The data show that most sites have enjoyed a large increase in core count over the two year period, but these data are of course transitory and depend where a site is in its replacement cycle. Birmingham has a rolling programme of infrastructure upgrades (rather than a periodic 'big-bang' total replacement) but our 2017 figures were boosted by having significant numbers of older generation kit still operational which was about to be retired.

Other Topics from the National e-Infrastructure Survey

Security and Identity

Security and Identity featured largely in this year's survey with the recognition that data needs to be properly secured and managed under regulatory frameworks such as the Data Protection Act and the General Data Protection Regulation (GDPR). The investments in health informatics and medical bioinformatics (MRC) and administrative data (ESRC) has emphasised the need for secure data storage

and trusted research environments. Systems now need to meet ISO27001⁶¹, NHS IG Toolkit⁶² and Cyber Security Essentials⁶³ standards, adding considerable additional costs for equipment, staff and certification.

The Report recommended: It needs to be recognized that secure systems certification should be funded as part of developing capital investments in health and social data research.

RCUK-sponsored JISC projects to integrate two key existing **Authentication, Authorisation and Accounting Infrastructure (AAAI)** technologies - namely Assent⁶⁴ and SAFE⁶⁵ - are approaching completion. The Authentication Service Assent and Safe Share are already in service at a number of institutions. The associated Authorisation and Accounting service SAFE is currently used by Archer⁶⁶, DiRAC⁶⁷, the Hartree Centre⁶⁸ and the EPSRC⁶⁹-sponsored Tier 2⁷⁰ HPC facilities. Successful integration of these services will be a significant step towards implementing a coherent National Authentication, Authorisation and Accounting Infrastructure that serves the needs of UK and international research collaborations.

These access and resource management infrastructures work in several common research settings, especially in areas where data must be securely accessed, stored and transported.

Birmingham's researchers need access to these resources

Objective 26. We will work with our IT Services colleagues to meet security requirements and deliver Assent, Safe Share, and Data Safe Haven (subject to a clear identification of need and a funding source.)

Networking

Nationally JISC are looking to the next generation of networking and the upgrade of the national backbone to speeds of 40 Gps or greater. Success will depend on better designed campus network

⁶¹ https://www.british-assessment.co.uk/services/iso-certification/iso-27001-certification/?gclid=EAIaIQobChMI5cmBt6eo3AIVq5XtCh0ZxQMzEAAAYAAAEgLn3fD_BwE

⁶² <https://www.igt.hscic.gov.uk/>

⁶³ <https://www.gov.uk/government/publications/cyber-essentials-scheme-overview>

⁶⁴ <https://www.jisc.ac.uk/assent>

Assent provides access to a wide range of services including cloud infrastructures, high performance computing, grid computing and commonly deployed services such as email, file store, remote access and instant messaging. It enable single-sign on capability within, across and between organisations, where secure communication is ensured.

⁶⁵ <https://www.jisc.ac.uk/safe-share>

Safe Share enables authorised researchers to access and share data safely and securely, from your own project site – while still meeting strict information governance rules, which may vary according to the sensitivity of the data

⁶⁶ <http://www.archer.ac.uk/>

⁶⁷ <https://dirac.ac.uk/>

⁶⁸ <https://www.hartree.stfc.ac.uk/Pages/home.aspx>

⁶⁹ <https://epsrc.ukri.org/>

⁷⁰ <https://epsrc.ukri.org/research/facilities/hpc/tier2/>

interfaces (being explored through the Jisc End-to-End⁷¹ initiative) to enable the transportation of large datasets and flows between facilities.

Objective 27. We will work with our IT Services colleagues to design the next generation of the campus network to ensure it meets the needs of the University's research community.

Virtual clusters.

The emergence of virtual clusters as a complementary service to traditional HPC batch processing was noted, highlighting that these reduce “the barrier of entry to those new to cluster computing and quickly improves their productivity and time to result.” CLIMB and BEAR Cloud were two of the first such virtual clusters to be provisioned as a campus and national resource.

RSE Activity

The quick action of the Research Councils to allow grant monies to be used for RSE activities was also highlighted as being key to this rapid expansion of software engineering support.

Aspects of a Successful Research Computing Team: People

The report went into detail about teamwork and the characteristics of a successful research computing team. This is of significant importance and is discussed in more detail in a later section of this report.

The Report concluded:

RECOMMENDATIONS AND NEXT STEPS

There were no new messages emerging from the discussion but the strong agreement amongst a diverse set of attendees representing the breadth of the science base was considered a strong statement in itself and their definite endorsement as to where further investment was needed.

Three emerging areas of activity will now be explored:

- 1. Additional funding to support methodological and skills development in big data analytics*
- 2. Formalising a policy network to take forward emerging policy issues that need a sector-wide response such as harmonisation, incentives, interoperability and standards.*
- 3. Networking opportunities to bring together researchers from different disciplines to share best practice*

A Holistic Approach

Looking nationally the report recognises that a holistic approach is key if we are to deliver more productive research and research systems. The Government's Industrial Strategy sets four grand challenges –

- growing the AI and data-driven economy,
- STEM and digital skills and improving the digital infrastructure,
- clean growth and future of mobility, and

⁷¹ <https://community.jisc.ac.uk/groups/janet-end-end-performance-initiative/document/janet-end-end-performance-initiative-project>

- the aging society

and this brings further requirements for continued investment in e-Infrastructure resources. Computational and/or data science approaches underpin almost all research across all industrial sectors and this requires both resources and the skilled people to drive innovation. This can only be delivered through a programme of co-design via strong partnerships between industry and academia

The National Infrastructure Roadmap

UK Research and Innovation (UKRI)⁷² is developing a research and innovation infrastructure roadmap⁷³ with the aim of delivering a long-term (until approximately 2030) research and innovation infrastructure view based on an understanding of existing UK infrastructure (and key international facilities in which the UK participates), future needs (research, economic and social), and resulting investment priorities. The intention is to publish the final report by April 2019.

The roadmap will:

- identify future research and innovation capability priorities
- identify opportunities for increasing inter-connectivity
- support development of UKRI's overall long-term investment plan
- promote the UK as a global leader in research and innovation
- set out the major steps needed to reach the long-term vision.

The scope includes facilities, resources and services that are used by the research and innovation communities to conduct research and foster innovation in their fields. They can include:

- major research equipment (or sets of instruments)
- knowledge-based resources such as collections, archives and data
- e-infrastructures such as data and computing systems and communication networks

The roadmap will be structured in the following sectors:

- Biological sciences, health and food
- Environment
- Energy
- Physical sciences and engineering
- Social sciences, arts and humanities
- Computational and e-infrastructures

It is recognised that computational and e-infrastructures cut across all sectors.

Within this context research computing at Birmingham could make a significant contribution and open up possibilities for future research income.

⁷² <https://www.ukri.org/>

⁷³ <https://www.ukri.org/research/infrastructure/>

Chapter 8: Aspects of a Successful Advanced Research Computing Team: People

ARC has always placed great emphasis on teamwork and the ability to relate to our users. This ethos begins at recruitment – we recruit as much for personal qualities and commitment as we do for technical competence. These themes have been explored in recent national reports. Extracts are included here as our touchstone and describe the characteristics we reflect as a high performing team in research support.

Objective 28. We will provide leadership to maintain and grow a flourishing strong team with a user focused ethos. We will learn from and apply proven best practice to facilitate this.

The National e-Infrastructure Survey⁷⁴: Extracting directly from their document the authors reported:

A successful and sustainable Research Computing team brings together people with a wide range of skills, interests and experience, including, in no implied order of importance, some or all of:

- *traditional system administrators, including, but by no means limited to, Linux, storage and Virtual Machines. This has recently expanded into Cloud in some form, whether that's building local clouds and/or exploiting commercial clouds. Technical staff are now as much system designers rather than system administrators.*
- *outreach, with the skill to talk to both the traditional (STEM, for example) users and non-traditional users (arts and humanities for example). This should also include knowledge of other resources that may be available, such as the Tier2 centres, the national service (ARCHER) and international resources such as PRACE*
- *applications specialists, especially to maintain a wide range of applications using a build system such as EasyBuild⁷⁵ (other systems are available). This would include handling the wide range of licensing methods for commercial and non-commercial software*
- *storage/data specialists including storage management, based on a solid understanding of data management*
- *trainers, especially for getting started training. Knowing where to find other training, for example that offered by ARCHER and suppliers such as NAG⁷⁶, and running those on-site is essential. Applications training on specific applications, and/or in specific research domains, may or may not be seen as part of the remit of Research Computing support. This includes support for train the trainers approaches advocated by the Software Sustainability Institute⁷⁷ and Software Carpentry⁷⁸.*
- *Research Software Engineers or similar to work more in-depth with researchers to advise on coding and how to make optimal use of the facilities available. Such RSEs must be aware*

⁷⁴ Ibid

⁷⁵ ARC has already adopted EasyBuild for application deployment

http://hpc-sig.org.uk/wp-content/blogs.dir/sites/63/2017/06/2017-05-11_EasyBuild_and_HPC.pdf

⁷⁶ <https://www.nag.co.uk/>

⁷⁷ <https://software.ac.uk/>

⁷⁸ Ibid

of other initiatives, such as the Performance Optimisation and Productivity Centre of Excellence in Computing Applications

- *Internal ambassador/champion to get the collaboration from local IT Teams and Jisc to enable operations and developments*
- *expertise in associated areas such as local and remote visualisation and support for high-end collaborative conferencing*
- *expertise in running procurements, and persuading a procurement department that cheapest often isn't best in this specialist area. Managing long-term relationship(s) with supplier(s) is also essential.*
- *management able to make a case for sustained and predictable investment to the funders, for example in a University this would typically be the Pro-Vice-Chancellor for Research or equivalent, and also to motivate existing and potential users of the Research Computing service to make the case from the researchers' perspective*
- *a futurist, to keep abreast of longer-term developments that may (or may not) become relevant to the Research Computing team*
- *The majority of the members of the team need to be outward focussed, enjoy networking with users and peers at other institutes and be well connected both inside and outside of their HEI*
- *It is not generally the case that any team member is only involved in a narrow area - for example, the outreach team needs to work closely with the applications specialists to make introductions where appropriate and to have an understanding of the challenges that others in the team face to avoid over-promising on initial contact*

The above maps closely to the ARC team's and philosophy. Each Group within ARC is critically dependent upon the others and proximity of working delivers far-reaching benefits.

In 2015 a HEFCE Report looked at 'Characteristics of high-performing research units'⁷⁹ – an extract from this report is included below. Here too there are many similarities between the working practices and ethos of the ARC team work and research groups, and these characteristics provide an excellent template for best practice. Indeed after one lengthy, introductory discussion with a Birmingham professor he concluded "Your team [ARC] works just like a research group. You have a problem, you have no idea how long it is going to take you to solve that problem [as often this has never been done in Birmingham – or sometimes elsewhere - before], you research it, you develop a hypothesis, you test it, and you implement it."

The report's conclusions and defining characteristics are detailed below (italicised). Comments on how these characteristics chime with the work of ARC are interspaced in the text.

79

http://www.hefce.ac.uk/media/HEFCE,2014/Content/Pubs/Independentresearch/2015/Characteristics,of,high-performing,research,units/2015_highperform.pdf

From our analysis we identified eight observations that are associated with high research performance and warrant further investigation. They are:

- *In high-performing research units more of the staff have PhDs, professorial positions, international experience and externally funded salaries*

Whilst not a pre-requisite for recruitment, staff who understand and have experience of the research process help the wider team to build strong and trusted relationships with research groups and are able to share that outlook with their colleagues.

Of the 20 staff in ARC as of 1 September 2018, 6 have or will soon receive a PhD. Others have many years experience working alongside research groups.

Some staff have national and international experience, serving on strategic working groups and forging collaborations across academic and the commercial sectors.

7 staff are funded from research groups.

- *High-performing research units are focused on recruiting the best and retaining them*

Finding staff of the right calibre and a very competitive and limited pool is a challenge. Nevertheless ARC's growing reputation has helped us appoint and retain high-performing staff.

ARC is a new Section within IT Services, formally established in 2015, although previously it existed as a team within another section. 17 of the 20 team members have been directly appointed to the research role.

- *High-performing research units provide training and mentorship programmes to develop staff, while offering rewards for strong performance*

ARC has a successful training programme in place, although formal training courses in this advanced area are few and far between. More relevant has been attendance at national and international conferences and working groups, with senior staff mentoring more junior colleagues.

The commitment to mentoring has also been demonstrated by employing 5 computer science undergraduates on extended placements over the last two years. We believe our commitment to train these students in a working environment has enhanced their CVs whilst our investment in developing these promising individuals has paid dividends in the work they have been able to deliver for the team and the service.

- *Staff within high-performing research units display a distinct ethos of social and ethical values.*

The ability to appoint the majority of staff to the research-focussed team has enabled us to build an ethos within the team, where an outward facing and engaging attitude is just as important as technical ability.

Additionally, one member of the team is a STEM Ambassador⁸⁰, who offers their time and enthusiasm to help bring STEM subjects to life and demonstrate their value and opportunities.

The formation of a chapter of the 'Women in HPC'⁸¹ initiative is also under consideration.

- *The leaders of high-performing research units have earned 'accountable autonomy' within their higher education institution*

ARC has earned the trust of the research groups through our track record, and the expertise, effectiveness and responsiveness which characterises our work. We work in a dynamic environment in an orderly but not over-bureaucratic way, mirroring in many respects the approach of the research groups themselves.

- *High-performing research units have strategies that are real, living and owned, and more than merely a written document*

Whilst this is, of course, a written document it is developed in liaison with academic colleagues. It provides a direction of travel but its relevance is under constant review. The rapidly changing landscape of research means the team must be flexible and able to respond to change rather than be constrained by a rigid rules-based approach.

- *High-performing research units receive more income per researcher than the average research unit*

Whilst ARC is not directly eligible to apply for research grants, support costs may at times be added to academic cases. Several research groups have invested in additional computing or storage resources (above the default) to meet more demanding needs. The time of the RSEs within ARC is not currently charged back as this free-at-the-point-of-use approach is seen as necessary to meet the broadest of needs with a significant but limited investment of time. However in future we foresee the possibility of RSE time being added to research grant applications where a large but not full-time resource would be beneficial. ARC would then amalgamate these fractional investments to employ and retain full-time RSEs.

- *High-performing research units enable and encourage researchers to initiate collaborations organically as opposed to using a top down approach*

ARC is a meeting place for researcher from across campus. Our wide-ranging view of research activity allows us to make connections and facilitate collaboration where links might not otherwise be identified. As just one example of this a close research link now exists between genetics and theology (Edmondson)⁸².

⁸⁰ <https://www.stem.org.uk/stem-ambassadors>

⁸¹ <https://womeninhpc.org/>

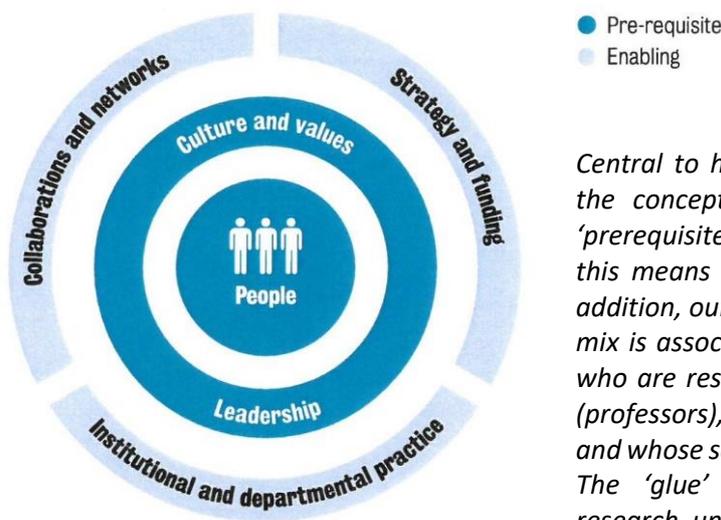
⁸² <https://www.birmingham.ac.uk/research/heroes/ancient-texts.aspx>

Looking at these observations, it is apparent that they can be mutually reinforcing and interact in positive ways. We have developed a conceptual model that helps explain how these observations may interact.

Key themes associated with high-performing research units

- *People*
- *Leadership, culture and values*
- *Strategy and funding*
- *Collaboration and networks*
- *Institutional and departmental practice*

Figure 4 Concept model for describing characteristics of high-performing research units



Central to high-performing research units – and the conceptual model in Figure 4 – are two ‘prerequisite’ characteristics. The first is *People*: this means recruiting and retaining the best. In addition, our analysis suggests that a certain staff mix is associated with high performance, ie staff who are research trained (PhDs), who are senior (professors), who have international experience and whose salaries are funded by external sources. The ‘glue’ that creates the high-performing research unit is its research culture, underlying values and leadership. All the high-performance

research units we spoke to had a degree of earned or accountable autonomy – that is they were allowed to get on with what they were doing, partly as it was recognised that they were successful due to their strong leadership and the research culture of the unit.

Three ‘enabling’ characteristics allow people and leadership to thrive and they are depicted in the outer circle in Figure 4. They are: collaboration and networks, a coherent strategy and diverse funding sources, and supporting institutional and departmental practices. From our research it is not clear whether these three criteria are prerequisites for high research performance, or whether they simply enhance such performance.

Chapter 9 : Summary of Development and Implementation Objectives

Throughout the Chapters of this document which follow, a number of objectives have been set where IT can support the University's ambition of delivering a more research intensive environment. The background to these Objectives is described later and they are collected here to inform the Development Plan for BEAR and ARC. They are set in the context of the Strategy Statement outlined earlier.

Objective 1. We will carefully monitor the usage of BEAR services, ensuring efficiency and matching demand against available resource. Where projections indicate demand cannot be met from current provision we will bring forward business cases to argue for additional resource.

Objective 2. When research requirements in Dubai are identified, we will contribute our expertise to identify options and potential solutions.

Objective 3. We will continue to provide Service Reports and Metrics to each meeting of the Research Computing Management Committee

Objective 4. We will provide Service Reports and Metrics to each meeting of the CaStLeS (Life Sciences) Strategic Oversight Group

Objective 5. We will work to build effective partnerships with the NHS to facilitate delivery of the University's life sciences strategy

Objective 6. We will work with others to ensure the essential technology requirements and costs for advanced research computing at the Birmingham Life Sciences Park are understood

Objective 7. The IBM partnership marks a significant opportunity for AI development in Birmingham. ARC's programme will be an inclusive one to make these resources available to the widest possible audience. These developments will be aligned to and integrated with the University's AI and Data Science Strategy.

Objective 8. We will maintain the current annual investment cycle so that BEAR continues to evolve, building on the success we have already achieved.

Objective 9. We will lobby for increased investment in BEAR services which addresses inflationary pressures and, at a minimum, guarantees current service levels for the growing number of researchers attracted to Birmingham through the University's ambitious strategy.

Objective 10. More ambitiously we will bring forward carefully analysed and costed proposals for service development which will support Birmingham's drive to be among the UK's top 5 research intensive universities.

Objective 11. We will continue to promote BEAR across all Colleges and grow the user base in line with requirements using a variety of methods aimed at reaching the widest possible audience.

Objective 12. We will work with our colleagues in IT Services to promote the wider portfolio of services which are of benefit to researchers. This will range from appropriate and early, collaborative consultation through to widespread dissemination of service characteristics and benefits.

Objective 13. We will continue to promote and grow the BEAR-related training courses, publishing courses on the BEAR web site, and maintaining attendance and demand statistics

Objective 14. We will expand the curriculum of training courses, responding to demand from the Colleges. A current priority is for Data Science and AI training and requests in these areas are expected to grow significantly.

Objective 15. We will maintain our dialogue with the user base to ensure the software portfolio remains relevant and delivers broad value. We will also lobby for increased funding where the developing ambitions of the University's research strategy leads to increased software licencing costs.

Objective 16. We will develop the RSE service with the aim of creating a sustainable pool of expertise available to commit to medium to long term software engineering engagements that are essential to research projects

Objective 17. Subsequently, we will work with PIs and Research Support Services to include funding for RSE time within research funding bids.

Objective 18. Through the RSG we will work with researchers from across the University to deliver 'better software', and expand our skill sets to support the burgeoning growth of AI and Data Science.

Objective 19. We will seek appropriate funding for the RSG through both central allocations and contributions from research grants.

Objective 20. We will provide annual reports on activity to the Academic IT Reference Group

Objective 21. We will continue to monitor and evaluate the role public cloud can play in supporting researchers at Birmingham and adopt our approach so that the right solution is delivered to maintain competitive advantage.

Objective 22. We will maintain and develop the quality and reliability of BEAR services that we have already established.

Objective 23. We will deliver the best and appropriate leading edge facilities to keep Birmingham amongst the UK leaders in advanced computing facilities to support research.

Objective 24. We will maintain close relationships in the IT industry, keep abreast of technology developments to ensure we are in a good position to take advantage of developments. To support this we will allow time and invest in opportunities for the team to expand their knowledge and skills.

Objective 25. The water-cooled Park Grange data centre is highly energy efficient. We will measure and monitor this efficiency to support Birmingham's commitment to sustainability.

Objective 26. We will work with our IT Services colleagues to meet security requirements and deliver Assent, Safe Share, and Data Safe Haven (subject to a clear identification of need and a funding source.)

Objective 27. We will work with our IT Services colleagues to design the next generation of the campus network to ensure it meets the needs of the University's research community.

Objective 28. We will provide leadership to maintain and grow a flourishing strong team with a user focused ethos. We will learn from and apply proven best practice to facilitate this.

Postface to the Second Edition

This second Research Computing Strategy was endorsed by the Academic IT Reference Group, chaired by PVC Prof Tim Softley, at its meeting on 17th January 2019.

There was discussion at this meeting about the need to explore “SMART” metrics and quantitative performance indicators. Future funding decisions will more closely link investment to research outputs. The breadth and diverse application of advanced research computing services as well as its indirect impact on outputs make such objective performance metrics difficult to obtain. The pace of technology change may well introduce the need for new metrics within the lifetime of a Strategy; for example AI received only sparse mention in the 2016 Strategy; in 2018 it is of significant prominence. There is no consistent measure of provision, effectiveness or reporting between HEIs and so comparisons with our competitors are near impossible.

It was agreed to proceed with the 28 largely qualitative Objectives defined in this document but that further thought would be given to improved metrics with a view to introducing these with the Third Edition of the Strategy, due in 2021.

Dr John Owen
Head of Advanced Research Computing
January 2019

Glossary

AI	Artificial Intelligence
ARC	Advanced Research Computing – the team that supports the BEAR infrastructure
ATI	Alan Turing Institute
BEAR	Birmingham Environment for Academic Research
BLSP	Birmingham Life Sciences Park
CAL	College of Arts and Law
CASIM2	Centre for Advanced Simulation and Modelling
CaStLeS	Compute and Storage for the Life Sciences
CCB	Centre for Computational Biology
CLIMB	Cloud Infrastructure or Microbiology
CPU	Central Processing Unit
DMP	Data Management Plan
DMZ	Demilitarized Zone (used in computing to mean part of a campus network)
DSH	Data Safe Haven
EBI	European Bioinformatics Institute
EPSRC	Engineering and Physical Sciences Research Council
ESRC	Economic and Social Research Council
GDPR	General Data Protection Regulation
GPU	Graphics Processing Unit
HDR-UK	Health Data Research (UK)
HEI	Higher Education Institute
HPC	High Performance Computing
IPv6	Internet Protocol (version 6)
MRC	Medical Research Council
NVMe	Non-Volatile Memory Express
PGR	Post Graduate (Research) students
PI	Principal Investigator (on a research project)
POD	People and Organisational Development

POSIX	Portable Operating Systems Interface
PRISM2	Partnership for Research in Simulation of Manufacturing and Materials
RCUK	Research Councils (UK)
RSE	Research Software Engineer
RSG	Research Software Group
SDN	Software Defined Network
SIG	Special Interest Group
STEM	Science, Technology, Engineering and Mathematics
TDP	Thermal Design Power
UKRI	UK Research and Innovation
VM	Virtual Machine

Appendix 1 : Progress Report on the Objectives set out in the Research Computing Strategy 2016 (Published May 2018)

John Owen
Carol Sandys
Advanced Research Computing.

Summary

The Research Computing Strategy (July 2016) set 27 objectives to grow and improve the services provided to researchers. A review of these objectives in May 2018 conclude that significant progress had been achieved for 20 of these; some progress had been made on a further 5; whilst 2 still needed significant action.

Document Format

Colour coding of text:

Some of the original objectives were action points; some aspirational. These are highlighted in the text below in blue and green



Action



Aspiration

Colour coding of response boxes:

The table below is colour coded as green, amber and red

- **Green** means significant progress has been made
- For **Amber** some progress has been made but there's more to do
- **Red** shading means there's still a lot more work to be done

Nomenclature

BEAR – the Birmingham Environment for Academic Research – is the umbrella term for the advanced research support services offered by IT Services.

CaStLeS – Compute and Storage for the Life Sciences; resource made available to this research domain by an investment for IT infrastructure.

In the 2016 Strategy the team that supports and develops BEAR was known as the Research Support Section. In 2018 the name was changed to **Advanced Research Computing (ARC)** to better reflect the services offered and the activity of the team. The new name (abbreviated to ARC in places) is used throughout this report.

Objective 1. Through growth of our outreach activities we will develop a better understanding of the research needs, and make the case for adequate resourcing to develop solutions to address these.

Advanced Research Computing (ARC) has developed significantly over the past two years, including with the addition of a number of new posts and a real focus on understanding and working with research groups. ARC now comprises three groups with one entirely focused on Engagement and Training. Team members from the other two groups also have regular contacts with researchers at different levels, from early planning through operational support and the provision of coaching and support for software engineering.

ARC already offers regular starter training for use of BlueBEAR and is developing a broader training programme, dovetailed with training offers from other key units on campus such as the Library and the Graduate School. The new posts mean that 'contact hours' with research staff and students has increased radically and are reflected in the positivity of feedback received and the evolution of BEAR services. Although we now offer weekly Drop-In sessions at different locations on campus, carry out One to One sessions as well as Group sessions, the 'IT Needs of Birmingham Researchers' Survey conducted in April/May 2017, showed that there is still work to do to raise the profile of the service.

Objective 2. We will work with the research community to identify, resource and make available software applications to meet their research need. Our aim is to keep this under regular review with an ambition of providing a 'research toolbox'. This extends beyond BEAR and would include, for example, support for electronic laboratory notebooks.

The BEAR application landscape has changed radically since this objective was set. Most application builds are now automated and are made available both on BlueBEAR and to BEAR Cloud/CaStLeS VMs for maximum flexibility and convenience. The number of applications has approximately doubled to over 280 with new applications or versions added almost on a weekly basis, ranging from the most basic to complex sets of dependent code and even being asked to resurrect archaic code. There are still barriers for some disciplines e.g. where the cost of commercial licences is prohibitive. The site licence for Matlab has been popular as well as significantly reducing the administrative overheads.

The provision of electronic lab notebooks has proven challenging with no one solution applicable across all disciplines. Notebooks continue to be purchased and supported on a local basis and there seems little prospect of this changing.

The ARC team has however implemented Jupyter Notebooks within the BEAR Cloud environment. This allows the creation and sharing of documents that contain live code, equations, visualizations and narrative text.

Objective 3. We will continue to deliver a broad-based research support service, and to demonstrate the power of analysis and data analytics to established and new users

The trend identified in the 2016 Strategy has continued and accelerated with take-up of BEAR services by an increasingly broad set of users with diverse research interests. Our storage offerings and sync and share collaboration tool have proved particularly popular as

basic productivity tools. The Research Data Store (RDS) has developed and provides safe and secure storage for research data, with a free allocation of 3 TB per project.

Although exploited fully by a subset of users, the high speed and highly reliable storage which underpins all BEAR services is valuable for those with demanding compute requirements. Data may be processed on the different platforms and shared appropriately by users on the owning project without duplication or migration.

The CaStLeS infrastructure, dedicated to life sciences research, provides an efficient and easy to access resource which is increasingly being exploited within this research domain.

Both CaStLeS and BEAR Cloud provide a Virtual Machine (VM) service which is popular with users who are less familiar with the traditional High Performance Computing (HPC) environments; some of whom go on to develop the skills and confidence to use the more advanced compute services. The increased frequency of introductory training and additional support that we now offer means that we are able to assist this process and enable more users to exploit the power of BlueBEAR HPC

Objective 4. We will seek funding for the necessary staff resource to deliver an appropriate programme of IT Services provided technical training, with initiatives and courses targeted to specific disciplines where appropriate. The programme will be approved by the Research Computing Management Committee, published on the BEAR web site and reviewed annually.

As a result of a successful bid through the Compact process the team grew by six additional members of staff during 2017/18. One of these posts was for a Research Training and Engagement Officer, and this new post, together with three other existing positions (one of which is fixed term), have been brought together to form an Engagement Group.

Our programme of drop-in sessions, covering research data storage and management, and latterly broader aspects of the BEAR service, have grown and now occur weekly in different locations across campus. The frequency of our 'Introduction to Linux' and 'BEAR Necessities' courses will increase in 2018. At the start of the 2017/18 session the Matlab online training academy resource was purchased and made available to all across campus. Discipline specific courses occur on an ad-hoc basis as requested and where we can provide support. The network of Special Interest Groups has been re-invigorated offering peer-to-peer training and support opportunities.

Our Research Software Group provides coaching sessions on code development, and one of the team is certified as an NVIDIA Deep Learning Ambassador with a programme of courses on artificial intelligence which he is now empowered to deliver to members of the University.

These opportunities are promoted widely and regular reports on outreach activities are provided to the Research Computing Management Committee.

Objective 5. We will initiate a dialogue with the Graduate School and the Colleges with a view to developing a complimentary programme of data science to provide all postgrads and research staff with a foundation set of skills. This will be funded and delivered centrally.

Our links with the Graduate School have developed, with the Head of ARC now sitting on the School Board.

'Data Science' has many interpretations. In the context of this Objective the term was in part used to define what has now become known as Research Data Management (RDM). RDM is a core component of our frequent drop-in sessions and the need for data management plans and (open) data archiving has been promoted by the cross-unit and evolving Open Data programme involving IT Services, Library, Planning Office and other areas of Corporate services and the academy.

More broadly the ARC team have been involved in developing research software coding skills, good practice in software engineering (through the Research Software Engineer – RSE – programme), through the Academic Programmers' (RSE) Special Interest Group, and in planning and preparation for the University's membership of the Alan Turing Institute and the proposed Data Science and Artificial Intelligence Institute.

Objective 6. Recognising that the Team are our most valuable resource, we will take appropriate steps to retain their skills through maintaining the stimulating and rewarding environment in which they work.

Birmingham's reputation within the international research computing community has grown and we are recognised as delivering quality and innovative work. It is our ability to continue with this programme that has drawn new staff to join us and encouraged existing staff to stay. We are constantly aware however that this is a competitive environment with similar jobs within UK HE appearing weekly. Our staff are valuable targets.

Team working is at the heart of our philosophy. The team is strong and supportive – to others within the team as well as our researcher colleagues – and works hard to remain as such. Increasing the team size by nearly 50% during 2017/18 has been a challenge but has been achieved successfully. This reflects in our recruitment process where we select as much on an ability to work as part of the team as on technical ability.

There have been challenges to our ability to remain as one team throughout the period, not least by the lack of office accommodation within Elms Road. The expanded team is now dispersed across the building whereas previously they were (mostly) located in adjacent offices along one corridor. This co-location promoted a culture of informal collaboration and support, the team functioning very much along the lines of a research group, bouncing ideas off each other and stimulating development. This degree to which this has been lost is to be lamented and efforts will continue with an aim of co-locating the entire team.

Objective 7. We commit to providing adequate and appropriate training for the team, and also to a continued presence on national and international bodies to share and learn best practice, and promote Birmingham's reputation as a technology leader in the sector.

All training requests have been met and the team's representation on national and international bodies has remained and grown.

Research Software Engineering (RSE) – the recognition that 'better software means better research' - is an area which has grown rapidly since the strategy was produced in 2016. Through the team, Birmingham is now a member of the national RSE Special Interest Group, on the Leadership Committee of that SIG, and their national / international conference will come to Birmingham in September 2018.

Objective 8. We propose the creation and promotion of an advanced computing graduate trainee scheme, similar to the one already operated by the University for management trainees. In relation to research computing, we propose offering such opportunities to Computer Science (or similar) graduates and growing skills that would benefit the research environment in Birmingham, and nationally.

Although we have not instituted a formal trainee scheme, we have employed three Computer Science students during the period. Initially these were summer job appointments but all three remained with us throughout the year working on a part-time basis. This was a highly successful initiative and the quality of students excellent. ARC benefitted from their presence and as did they.

Our two year-round students left us in Spring 2018 to take up year-in-industry appointments, and we intend to expand the programme by offering an industrial year appointment to another student. The opportunity was advertised across EPS although the expectation is that this will be offered to another Computer Science student.

Objective 9. Birmingham must pay close attention to national networking developments and be in a position to take full advantage of new and ground-breaking services offered by the JISC and others.

We have been unable to take advantage of these national networking developments due to a lack of resourcing from our Network colleagues and a lack of prioritisation given to developments which will support research activities. This is within a context where data transfer is increasing and two projects in particular – COMAPRE and Prof Ben Brown's collaboration with US National Laboratories – are stressing our networking capabilities. This is an area of major concern within ARC and one that must be addressed in the coming period.

Objective 10. We will work with Colleges to identify and connect locations which would benefit from a Research Data Network connection.

The Biosciences building, IBR and ITM are connected to the RDN. Connections are in place to Physics although not yet activated. Developments within the Particle Physics area mean that this location is likely to go live in the near future.

Other sites being investigated include SportEx and Haworth

Objective 11. Through our highly skilled staff we will ensure appropriate security measures are in place to safeguard our valuable data. Inadequate security will see us being excluded from some joint activities (eg with the NHS) or lose grant income where we cannot safeguard commercial and other data.

BEAR data is secure and maintained to the highest standards. We recognise that we hold valuable and sometimes sensitive research data

BEAR has a sufficiently high standard of security to satisfy most requirements. For instance, Rolls Royce have confidence in the levels of security we have put in place to safeguard their data. Other research projects are now encrypting their data at source before storing it on the research data store (RDS).

We have not yet implemented on-disc encryption although all the necessary components are in place and all that remains is the staff effort to implement these. Encryption forms part of a broader roadmap for storage developments which will see a major infrastructure refresh in 1H2019.

However we still fall short of being able to satisfy requirements particularly for NHS data and joint NHS projects. The need to implement new levels of security such as the NHS Toolkit, SafeShare and Data Safe Haven, are likely necessary components for emerging NHS interactions such as the Health Data Research consortium (HDR), Birmingham Health Partners and the Life Sciences Park. These initiatives will place significant demands on many areas of IT Services, including and beyond the BEAR team, and there is currently significant uncertainty as to how these will be resourced, both by personnel and financially.

Objective 12. We must be in a position to take advantage of nationally developing services to ensure IT facilitates and does not hold back collaborative research.

The objective in this context related to a streamlining of authentication methods, removing the need for multiple log-ons to different and geographically scattered systems. Although the technology (JISC Assent - <https://www.jisc.ac.uk/assent>) has been available for several years the Security and Authentication teams within IT Services have not been able to support its deployment on campus.

This remains a requirement and one that is the bedrock for other essential security and authentication developments and the gateway to future collaborative research areas.

Objective 13. We will continue to develop our skills and experience to be in a position to fully contribute to the design and operation of the regional infrastructure if and when funding becomes available. Our proposal will be based on developments of the BEAR architecture planned during 2016.

The EPSRC call to replace the MidPlus infrastructure came much sooner than anticipated on 17th June 2016. Birmingham was involved in two bids – one from the Midlands (MI) consortium, and one based around the OpenStack Cloud environment. Birmingham led on the latter, with Queen Mary University London and Cardiff, but on this one was unsuccessful. Disappointingly EPSRC thought the bid was not innovative enough!

The Midlands bid did however succeed and led to the release of the Tier-2 HPC-Midlands-Plus service which opened to pilot users in April 2017.

<https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/HPC-MidPlus/index.aspx>

Objective 14. We will seek to expand our services to include advice and support for researchers wishing to apply for time on national and international resources, enabling Birmingham to take full advantage of the hierarchy of services.

Having been part of the EPSRC funded 'MidPlus' regional tier-2 service, Birmingham joined the consortium of six other Universities to successfully bid for the replacement 'HPC-Midlands-Plus' Tier-2 compute service, funded by EPSRC and made available in 2017.

<https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/HPC-MidPlus/index.aspx>

Problems at Loughborough however (the lead site) mean this service is only just becoming stable. (May 2018)

Tier-2 support through ARC is also the gateway to other national Tier-2 facilities.

A member of the team became an 'Archer champion' in 2016 with a remit to promote and support work on the national Tier-1 service. Notwithstanding some successes more work remains to be done to fully support Birmingham's research community in using Archer, which will be picked up by the new Engagement and Research Software groups within the BEAR team.

Objective 15. We will continue to invest in the BEAR infrastructure, steadily increasing our investment year on year to ensure we are competitive.

Our investment in BEAR has continued, albeit starting from a low base, and we are now competitive in the sector although not sector leading. Recurrent and capital funding streams support BEAR, as well as income resulting from one-off initiatives. Facilities also received a major boost in 2016 when £2M was awarded to pump-prime IT developments for life sciences research. This established the CaStLeS infrastructure.

Recurrent, maintenance funding has remained relatively static over the period. Capital investment is used to provision compute, storage and research networking. Increasingly it

is also being used to renew aging systems. As demand for services has grown we have maintained a steady growth in investment. Year on year growth is reported below relative to the 2015/16 benchmark:

	2016 / 17 %	2017 / 18 %	2018 / 19 (requested) %
Compute	25%	14.8%	13.5%
Storage	25%	29.6 % (a)	9.3%
Research Networking	Networking was not funded separately in 2015/16	15%	54% (b)

(a) Third tier of storage – to tape – introduced with bulk of new equipment cost met through a separate funding bid.

(NB This provides an active and cost-effective storage layer, not a backup, although backup is also done on the same device)

(b) Major equipment refresh required

Objective 16. We will maintain our technical lead in the sector by developing and nurturing our staff. We will build a ‘private cloud’ on campus for the benefit of Birmingham researchers and as a technology demonstrator of what can be delivered regionally and nationally.

Our BEAR Cloud service launched in October 2016 and to date has been mainly exploited by life sciences users through the CaStLeS infrastructure. This in turn built on our innovative experience gained from building the national CLIMB infrastructure.

These developments have built on exceptional technical expertise and a strong commitment to team working and peer support from the individuals who constitute Advanced Research Computing.

Objective 17. We are in a position of national leadership. We will exploit this further by working with national bodies to deliver ground breaking service developments. We will seek to fund this by winning grant income for IT projects from bodies such as the JISC

We have maintained and developed our position of national leadership and have facilitated research groups in winning grant income. ARC has so far not been successful in winning grant income in its own right however. Conflicting priorities have meant that so far only one grant proposal has been submitted – to EPSRC to develop one of the national Tier-2 services. This was unsuccessful, perhaps because the infrastructure proposed laid beyond EPSRC’s model for Tier-2 services, but the process did help us develop better relationships with Queen Mary and particularly Cardiff which, as the main host site for HPC Wales, has long been seen as holding a prominent position within the high performance computing community in the United Kingdom. Links developed here remain strong to this day, most

recently in shaping the governance of the national High Performance Computing Special Interest Group.

We still believe there are opportunities in this space which a stronger team will enable us to revisit.

Objective 18. We will maintain and develop an appropriate and effective storage environment, including encryption where relevant.

Storage provision for researchers has increased significantly since 2016 with XXX PB now provisioned. Most of this is replicated and all is backed up.

The default free allocation per project is now 3 TB of working space, and 1 TB of archive space. Additional storage space can be purchased and whilst a few users have done so, for most this allocation meets their needs

There has been high demand for a dropbox-like facility which was addressed by the introduction of BEAR DataShare. This provides many of the features of Dropbox at a very cost-effective price and has the advantage that data remains on campus and under the University's control. Access however is also available to external partners. As of May 2018 more than 3000 users were registered to use DataShare.

Encryption has not yet been implemented as a services but are the components are in place to facilitate roll-out during 2018.

Objective 19. Along with partners from the Library and RIS, we will support the Planning Office in delivering an effective research data management system by October 2017.

The decision was taken during 2016 to base the user interface and storage service on a new installation of ePrints backed by the on-campus Research Data Archive storage. A spin-out company from the University of Southampton (authors of ePrints) were engaged to develop the solution. However the project suffered a number of delays with service release not occurring until May 2018.

Objective 20. We will take possession of the new research data centre during the second half of 2017. This facility will then be equipped with new BEAR infrastructure and enable Birmingham to respond to campus and national development opportunities. We will design all future services to take full advantage of the resilience potential offered by the later delivery of a secondary research data centre in a refurbished Elms Road facility.

A number of issues, not least the presence of badgers on the Park Grange site, delayed the start of the build until August 2017. The building was handed over to the University on 14th May 2018 and the migration of kit to the new facility began just a few days later.

Some BEAR kit will remain in Elms Road for resilience but the future strategy for delivering a secondary data centre, be it in a refurbished Elms Road or in another location, is still not clear.

Objective 21. We will propose a programme for supporting and managing a collaborative software development environment and repository.

Funding for new members of staff during 2017 has facilitated a major expansion of the BEAR team and enabled the creation of a Research Software Group, currently with 5.5 members of staff. This group is now engaging with individuals and groups to facilitate software development and improve the quality of software. They offer advice, coaching and coding. (*"Better software, better research"*⁸³)
<https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/rsg/research-software-group.aspx>

Our on-campus software repository BEAR Gitlab is now operational. This is based on the de facto standard Github but keeps all data securely on campus.
<https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/bear-gitlab/gitlab.aspx>

An on-campus special interest group ('Academic Programmers' otherwise known as RSE) has been established by Professors David Parker and Jean-Baptiste Cazier, in association with the BEAR Research Software Group.

Objective 22. Galaxy will be the first subject specific portal to be supported on the new BEAR Cloud infrastructure. Working alongside researchers we will add further portals, always with the aim of increasing ease of access and effectiveness of research.

Galaxy is being used extensively within bioinformatics and the national Phenome Centre.

In other areas the use of subject-specific portals has not been progressed although virtual machines tailored to a researchers requirements are made available through BEAR Cloud. This seems adequate for the user base at the moment. With the growth of the Research Software Group within the BEAR team more detailed and specific interactions with individual researchers are now a possibility. This will guide the development of more user-friendly interfaces to advanced computing resources.

Objective 23. By working with our academic and industrial colleagues we will maintain and develop Birmingham's reputation as a sector leader in the provision of centralised research computing services.

⁸³ Software Sustainability Institute www.software.ac.uk

Birmingham's research profile has benefitted significantly by the availability of the BEAR resource. Notable areas include the whole field of life sciences research, spearheaded through the Centre for Computational Biology and the collaborative COMPARE project. Professor Nick Loman's work supported by the CLIMB infrastructure won two international awards which were presented at Super Computing 2017 in Denver:

- Best Use of HPC in Life Sciences, and
- Best HPC Collaboration (Academia/Government/Industry)

The BEAR team now have a significant profile, both nationally and internationally. Members of the team participate in the following strategic groups and activities:

- RCUK working group on Cloud
- National Project Directors Working Group
- Spectrum Scale user group
- High Performance Computing national special interest group
- HPC Midlands Plus Management Board and Technical Committee
- National e-Infrastructure Survey (2017) co-author
- Academic Applications Forum
- National Research Software Engineer Leaders' Group
- Archer Champions' group

The development of the innovative BEAR Cloud environment and the commitment to environmentally friendly all-water cooled research data centre have received wide spread attention. We retain and have developed very high level and strategic relationship with our industrial partners, notable Lenovo and IBM.

Objective 24. Through the use of our Research Computing Funnel we will publish our portfolio of existing and planned services. This will include horizon scanning for developing technologies and, where possible and appropriate, planning for their timely introduction.

The funnel has been maintained and updated regularly, with the portfolio of services being built in consultation with the user base. New, but not bleeding edge, technologies figure highly in the work of the team with for example the adoption of a wholly water cooled approach for kit in the new research data centre (commissioned May 2018 but in planning since 2016), the launch of the BEAR Cloud virtual machine architecture, and a critical examination of object storage for part of our storage infrastructure refresh planned for early 2019.

Objective 25. We propose employing two staff within IT Services (Application Services) to work alongside researchers to develop and maintain research-focused web-sites. Over time these staff will develop expertise to address many of the underlying and common requirements of research projects and hence speed development in a cost effective way whilst maintaining Birmingham's intellectual property.

It took longer than expected but by May 2018 two Research Software Engineers (RSE's) with a focus on web developments were in post. One is funded by CAL, the other by IT Services, and both sit within the Research Software group within ARC. This move away from locating the posts in Application Services was a result of a growing recognition over the period that, although web remains important, the needs of the Colleges are much more diverse and demanding than such a narrow focus would suggest. The two, relatively junior, new posts will also benefit from peer group support of other RSE's within the group.

Objective 26. We will build links on and off campus facilitating peer to peer support and cross-disciplinary research.

During 2017/18 Advanced Research Computing was significantly enhanced by the addition of 6 new members of staff. Amongst other important developments this enabled us to create an Engagement Team within the Section and increase the frequency of our training and drop-in sessions.

The existing 'special interest group' network on campus has been revitalised – although much more remains to be done – and a new SIG for Application Programmers has emerged. Interest in application programming – or Research Software Engineering (RSE) -has emerged strongly over the period since the last version of the Strategy, recognising that 'Better Software means Better Research' and the BEAR team has responded by creating a Research Software group. This team is already extremely busy working with researchers from across the University and facilitating cross-discipline links.

Our representation on external bodies has grown and our relationship with our main technology partners is excellent.

2016 saw the launch of our BEAR DataShare service, providing dropbox-like facilities including links to external bodies but with the data kept on campus for security. This has proved to be our most popular service with more than 3000 users registered users at May 2018.

Objective 27. We will survey researchers to gain a better understanding of the IT needs to support research, and build these findings into future strategy updates and implementation plans.

A survey of the IT Needs of Birmingham Researchers was begun in April 2017; analysis took almost a year. A brief summary (You Said, We Did) is available at

<https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/documents/public/website-revised-2015/survey/Research-IT-Survey.pdf>

More detailed reports will form an Appendix to the 2018 revision of the Research Strategy.

In summary, BEAR facilities were well received and the three main concerns, common across all Colleges, related to broader issues of IT support:

1. IT Services needs to be more responsive to user requests.

2. There was significant frustration at being denied administration rights to desktop, personal computers. This, coupled with perceived long wait times for a response from the Service Desk (as highlighted in [1] above), was highlighted as adversely impacting research activity.
3. Many researchers commented that desktop computers provided by the University/College are inadequate

IT Needs of Birmingham Researchers

Results of a Survey conducted in
April / May 2017

Executive Summary

IT Services

Executive Summary

This survey was commissioned by PVC Tim Softley to better understand the IT needs of the University's research community. It ran from 12th April until 12th May 2017. This report analyses the survey responses; an action plan for addressing the issues raised will be developed separately. To date there has not been significant analysis of the responses by College, other than overall satisfaction levels.

563 people viewed the survey and 481 contributed answers to all or part of it. EPS, LES and MDS had similar response rates with 133, 117 and 128 responses respectively from each College, whilst CAL and SS contributed 56 and 49 responses.

81% of respondents were staff, 17% students, and 2% did not declare a status.

Several questions invited textual responses about priority developments in the overall IT service provided for researchers. Although the intention was to encourage respondents to approach the issues from different perspectives, similar comments, mostly highlighting operational issues, were made in response to all questions. The three most frequently mentioned items to improve the service were:

1. IT Services needs to be more responsive to user requests. Long wait times for what are often perceived as non-challenging tasks causes a high degree of frustration. The response times quoted for request submitted through the Service Desk (a default value of 10 working days) were seen as unacceptably long.
2. Related to this there was significant frustration at being denied administration rights to desktop, personal computers. Wait times of several days for what were often perceived as straightforward tasks were frequently highlighted, with comments that this impacted adversely on research activity.
3. Many researchers commented that the desktop computer provided to them by the University was inadequate for their need.

Other requests for improvements included:

4. Further investment in BEAR and particularly BlueBEAR (HPC) facilities
5. Better WiFi provision and coverage. (There were also comments about the suitability of the general wired network which need further investigation but are most likely linked to the growing need to transport large quantities of data)
6. Better facilities for remote access (virtual private networks, VPN)
7. Better support for Linux and MAC desktops, and permission to purchase these.

Where request for new services were highlighted these were principally around:

8. 'Big data' – capacity for storage and guidance and training for analysis.
9. Support with programming, which is increasingly becoming a fundamental part of the research process. This mirrors developments nationally and the rise in importance of the 'Research Software Engineer' role.

BEAR services, both compute and data, are valued and appreciated by those who use them. Disappointingly though only 25% of respondents are BEAR users, and many have not even heard of BEAR.

Many of the people who are not using BEAR but have heard of it apparently have misconceptions about the facilities. Respondents quoted lack of privacy of data, overly bureaucratic process, and high cost of use as barriers. Many did not seem to realise that the standard provision is free at the point of use.

On the data front there is a general level of satisfaction with the centrally provided Research Data Store (RDS) and Research Data Archive (RDA), although when asked where research data is stored the RDS came in third place behind 'On my own or Research Group Storage' and 'On Department/School or College storage'. Worryingly 18 people said they did not have enough storage space so they have to throw data away, whilst at least one person admitted their primary storage module was a USB stick.

Demand for storage space for both working data and archived data was high, as expected, with a predicted requirement growing at greater than 1 PB per year.

There were some negative comments about the stability and functionality of BEAR DataShare in comparison to Dropbox.

Most researchers are aware of the University's and Research Councils' requirements around Open Data, although 21% (62 respondents) said they were unaware.

The need for more training came through loud and clear. The scope covered existing facilities (BEAR, existing applications eg Matlab, and new areas such as big data, data management and data analytics.)

Overall there was a positive impression of the support provided by the wider IT Services' and by the BEAR / Research Computing team in particular. The average satisfaction level for IT Services (on a scale of 1 – 5) was 3.26, whilst for BEAR it was 3.79. There were however some hotspots within Colleges, particularly LES, which will be investigated.

Appendix 3: 'You Said, We Did' – a Summary Report of the 2017 Research Computing Survey

Background

The first University wide survey of the IT needs of researchers was commissioned by PVC Tim Softley and ran from 12th April until 12th May 2017. This brief document looks at the major themes and summarises IT Services' response.

563 people viewed the survey and 481 contributed answers to all or part of it. EPS, LES and MDS had similar response rates with 133, 117 and 128 responses respectively from each College, whilst CAL and SS contributed 56 and 49 responses.

81% of respondents were staff, 17% students, and 2% did not declare a status.

Your Concerns

The table below lists your concerns and what we have done about them. We've also colour coded the responses as green, amber and red

Green	we believe significant progress has been made
Amber	some progress has been made but there's more to do
Red	there's still a lot more work to be done

Of course you may disagree with our assessments; if you do please let us know. (Good as well as bad feedback welcome!)

You Said	We Did
Principal Concerns	
Across the Colleges three principal concerns came through loud and clear	
IT Services needs to be more responsive to user requests. Long wait times for what are often perceived as non-challenging tasks causes a high degree of frustration. The response times quoted for request submitted through the Service Desk (a default value of 10 working days) were seen as unacceptably long.	We appreciate that waiting for a response to a service desk call can be frustrating and 10 days is a long time to wait for a fix. Analysis suggests calls are resolved more quickly than this with the average being between 4 and 5 days depending on the College. A number of College IT teams now have dedicated resource aligned to research support and this will help improve responsiveness in key areas.

<p>There was significant frustration at being denied administration rights to desktop, personal computers. Wait times of several days for what were often perceived as straightforward tasks were frequently highlighted, with comments that this impacted adversely on research activity.</p>	<p>Developing a policy on granting administration rights on desktop and laptop machines straddles an uneasy boundary between convenience for you as the user and security risks which could potentially affect many more people. We have a good record in Birmingham in protecting the campus from malicious attacks and do not want to jeopardise this.</p> <p>Current policy is that all justifiable requests to the Service Desk for admin rights will be granted (subject to review by your College IT Manager). Depending on the College this may be a temporary elevation of rights, or a more permanent arrangement.</p> <p>In the longer term we are looking at automating this procedure so that rights may be granted in a speedier fashion and revoked when no longer required.</p>
<p>Many researchers commented that desktop computers provided by the University are inadequate</p>	<p>Desktop computers are provided from College budgets and independently of IT Services. These concerns have been fed back through your College IT Manager</p> <p>If specialist IT hardware is required to perform research activity, where possible, a request for funding should be written into the research grant. Assistance with specifying and costing the hardware is available from IT Services.</p>
<p>Other Comments</p>	
<p>Calls for further investment in BEAR and particularly BlueBEAR (HPC) facilities</p>	<p>BEAR infrastructure (compute, storage and research networking) receives significant funding each year through IT Services budgets. During 2016 we also received substantive investment to create the CaStLeS environment, providing dedicated infrastructure for research into life sciences and available to all Colleges working in this area.</p> <p>For many years this investment enabled us to grow the BEAR environment. However some of</p>

	<p>the kit is now in need of replacement meaning the budget is now more concentrated on renewal, leaving limited opportunity for further growth.</p> <p>During May this year our new multi-million pound, research-focussed data centre will open providing capacity for future growth for several years to come.</p>
Further investment in BEAR (Part 2)	<p>BEAR services are not just about the kit. Infrastructure needs skilled people behind it to deliver the best for the community.</p> <p>The Advanced Research Computing (BEAR) Team was already populated by highly skilled individuals but in the last 12 months we have been able to expand the team. 6.5 new posts have been created in the team in the last year and all have been filled by further skilled individuals.</p> <p>There's no shortage of work and these new people make a real difference to the ability of researchers to exploit the technology.</p>
Better WiFi provision and coverage	<p>We are moving away from our current wifi vendor and replacing the infrastructure with Hewlett Packard's Aruba equipment, widely acknowledged as the joint best network solution in the world. The project is benefiting from significant investment and well underway with a planned completion date in 2019. Several areas have already been converted including Computer Science, the Main Library, Alan Walters building and Poynting Physics</p>
Better facilities for remote access (virtual private networks, VPN)	<p>A new remote access and VPN service was launched in February 2018. It is available to all staff and research postgrads, and can be requested through the Service Desk.</p>
Better support for Linux and MAC desktops , and permission to purchase these	<p>All justifiable requests to purchase Linux and Mac computers will be granted (subject to funds being available and review by your College IT Manager).</p> <p>There are a number of staff across IT Services that are trained to support Linux and Mac Operating Systems, with a number of College IT teams having</p>

	dedicated Linux and Mac specialists aligned to research support
Big data – capacity for storage and guidance and training for analysis.	<p>Demand for data and storage has grown significantly. The centralised Research Data Store (RDS) is available to all researchers with a free allocation of 3 TB per project (larger volumes may be purchased at a cost effective price). Storing your data on the RDS has several advantages:</p> <ul style="list-style-type: none"> • High speed connectivity to BEAR services • Data are backed-up regularly • Data are replicated between two data centres delivering security against equipment failure and resilience. • Data are managed by a professional team of technical experts, further safeguarding integrity. <p>Drop-in sessions on data management are held regularly in various locations across campus. Further support in data analysis is available through the Research Software Engineers and the Researcher Training and Engagement Officer (see below for further details)</p> <p>This is a very dynamic area and more support will become available during the coming months.</p>
Support with programming , which is increasingly becoming a fundamental part of the research process. This mirrors developments nationally and the rise in importance of the ‘Research Software Engineer’ role.	<p>Within the Advanced Computing (BEAR) team we have now created a Research Software Group with 5.5 Research Software Engineers (RSEs).</p> <p>https://intranet.birmingham.ac.uk/it/teams/infrastucture/research/bear/rsg/research-software-group.aspx</p> <p>This team provides programming Advice, Support for BEAR users, and Coding support which can involve embedding an RSE within a research group for up to 20 half-day sessions.</p>
Other Issues	
Misconceptions about BEAR facilities	Several ‘urban myths’ seem to have developed access to BEAR facilities. Full details of BEAR can be found at www.bear.bham.ac.uk but in summary:

	<ol style="list-style-type: none"> 1. The service is provided free of charge to the vast majority of users. Only when exceptionally large resources are required (for example, more than 3 TB storage for a project) is a contribution to costs necessary. 2. Requesting access to BEAR is straightforward. A simple request made through the Service Desk is all that is required. 3. We already have a Dropbox-like facility, known as BEAR DataShare which allows the sharing of data with collaborators both within the University and externally.
<p>Stability of BEAR DataShare</p>	<p>At the time of the survey (April/May 2017) we had experienced some issues with BEAR DataShare. These were worked on as a priority and the service is now very reliable.</p> <p>We have added functionality over the year and the service is now widely used across the campus (approaching 3000 users/collaborators).</p>
<p>A need for more training came through loud and clear. The scope covered existing facilities (BEAR, existing applications eg Matlab, and new areas such as big data, data management and data analytics.)</p>	<p>Within the Advanced Research Computing (BEAR) Team we have now created a Research Engagement Group with a dedicated Researcher Training and Engagement Officer. Our training programme and supporting documentation is under active review as we begin to expand our offering.</p> <p>Full details of training courses can be found at https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/bear-training/index.aspx</p> <p>Amongst many initiatives are:</p> <ul style="list-style-type: none"> • As a University we have now signed up to the Matlab Academy Online Training Suite – a series of online courses which significantly increases the training material available to us. • Birmingham is a member of Software Sustainability Institute (SSI) and through them have access to Software Carpentry material (https://software-carpentry.org/) . This year, six Carpentries Instructors have been trained who will then be empowered to deliver training using material from the SSI

How Did We Do? Satisfaction Level with IT Services and BEAR Services (out of 5)	
Satisfaction level for IT Services	3.26 out of 5
Satisfaction level with BEAR	3,79 out of 5

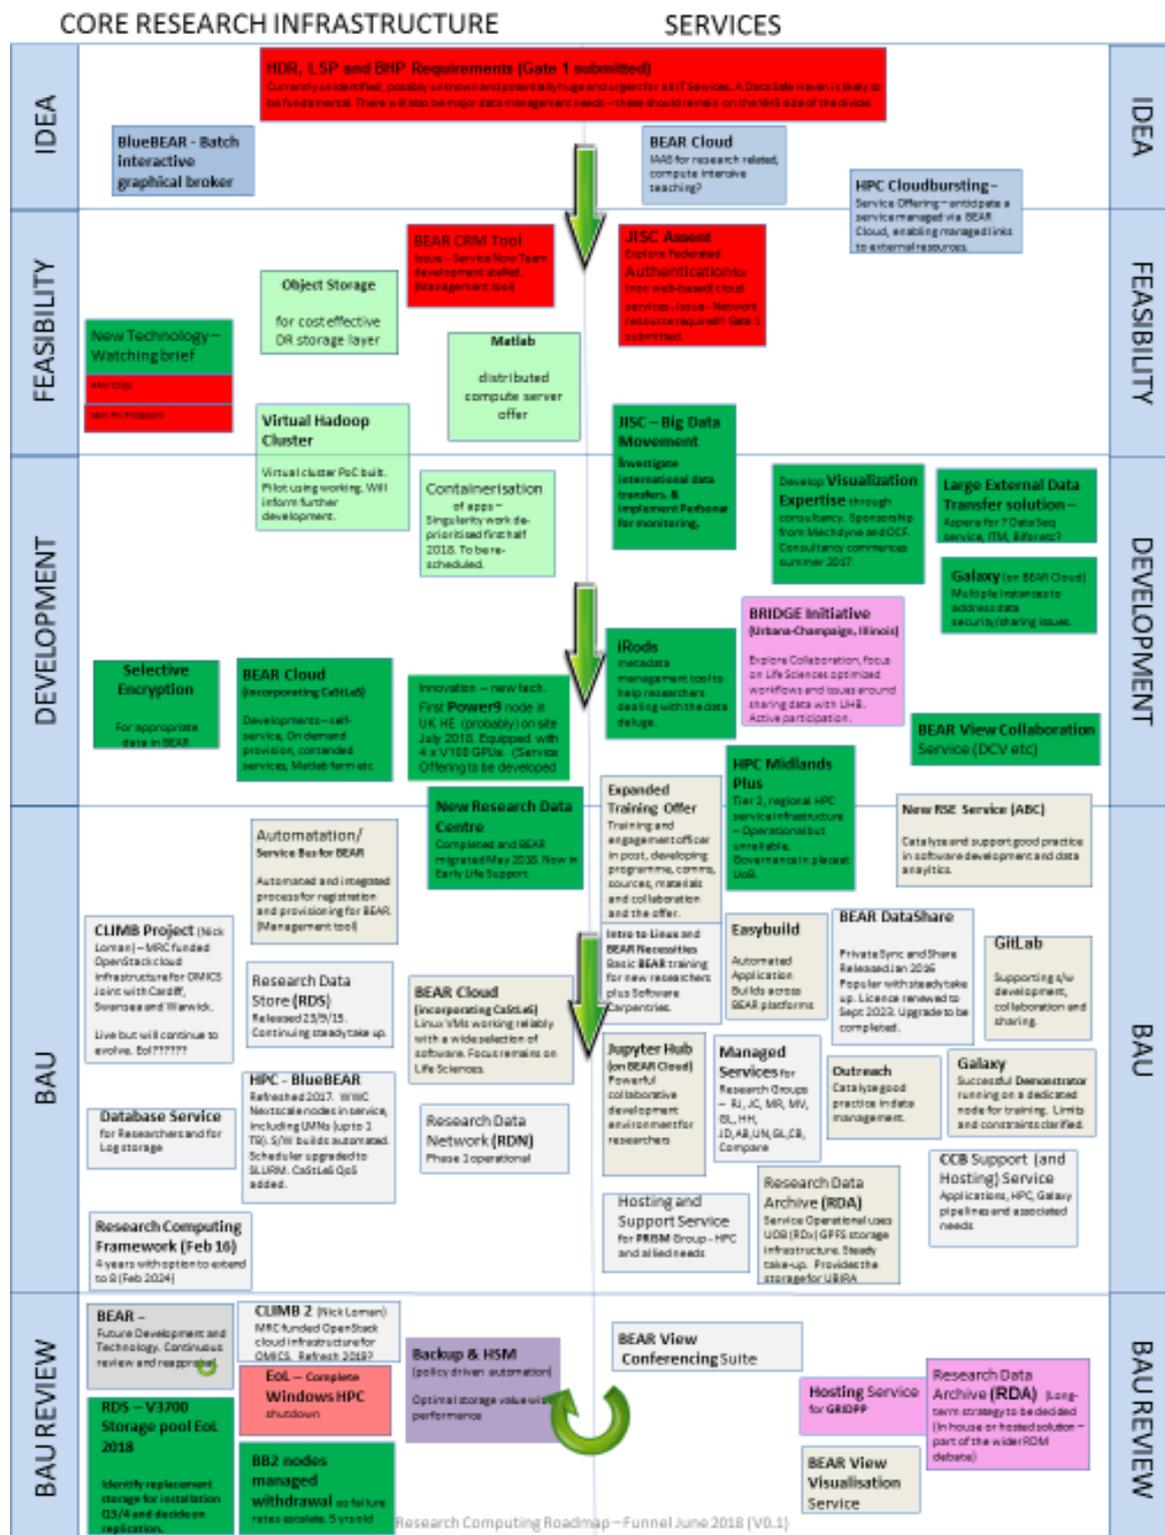
Appendix 4 : The Consultation Process

Members of ARC are in daily contact with members of the research community and much of this strategy derives from these interactions.

Discussions have occurred at both the Research Computing Management Group and the Academic IT Reference Group. Several members of the Reference Group have been consulted individually.

Meetings have been held with a wide range of users, from post-graduate students to Heads of School, College Directors of Research and the PVC for Research and Knowledge Transfer.

Appendix 5 : The Advanced Research Computing Funnel



(This proposed remit with some amendments to the original was proposed in July 2018 but is not yet ratified. This will be on the agenda for the next meeting in Autumn 2018)

Research Computing Management Committee

Terms of Reference and Remit

Membership of the **RESEARCH COMPUTING MANAGEMENT COMMITTEE** is drawn from academic researchers and Advanced Research Computing (ARC) staff in IT Services. It exists to ensure the academic needs of researchers are being addressed by BEAR (the Birmingham Environment for Academic Research), and to oversee and advise on its operation and development.

ORGANISATION AND MEMBERSHIP OF THE MANAGEMENT COMMITTEE

1. The Management Committee will have a maximum of FIFTEEN members: Ten academic representatives of the Colleges who have an active interest in using central, advanced computing facilities in their research; three from ARC; a senior representative from Computer Science and the Committee Secretary.
2. The Committee shall be quorate if at least three members of academic staff are present, representing at least 2 colleges.
3. The Pro-Vice-Chancellor for Research and Knowledge Transfer has a standing invitation to the Committee.
4. Two representative from each College will be sought; Colleges will be asked to nominate members.
5. The Chair of the Committee will be a member of academic staff and agreed by the members of the Committee.
6. The Chair of the BEAR User Forum shall be a member of the Committee, with the expectation that they will be an existing academic member of the Committee. If no such member is willing to serve then exceptionally another member of the University's academic staff may be co-opted and attend meetings.

7. If a member of the Committee is unable to be present at a meeting a deputy may attend in their place.
8. Other members may be co-opted onto the Committee or invited to attend specific meetings, at the discretion of the Chair.
9. Meetings will be held at least THREE times a year; additional meetings may be called at the discretion of the Chair.
10. Membership of the Committee and Committee positions will be reviewed every two years.

REMIT OF THE MANAGEMENT COMMITTEE

The Committee exists to:

1. advise the Academic IT Reference Group, and through it the University Research Committee and PVC Research and Knowledge Transfer, on the research computing strategy for the University, and guide its implementation.
2. advise ARC on the academic need for services to be delivered by BEAR.
3. recommend new services and developments to BEAR which would benefit researchers. Where funding exists, the Committee may authorise such developments subject to their compatibility with ARC plans and capacity. Where no funding exists the Committee will make representation to the University.
4. act as the University's Research Data Storage Management Board (RDSMB) and undertake the duties as defined in the Research Data Storage Management Policy
5. consider operational policies for BEAR, based on recommendations made by the ARC (after consultation). Such operational policies would include but are not be limited to:
 - a. default filestore allocations,
 - b. run-time allocations,
 - c. configuration of the scheduling system to maintain an equitable share of the resources for all users.
 - d. data management and security.

6. consider and adjudicate on any special requests for resources outside of the norm (for example exclusive use of a large number of nodes over an extended period of time, or requests for unusually large amounts of filestore) which have been referred by ARC. Users may also make a direct approach to the Management Committee where they wish to appeal a decision on such a matter by ARC.
7. consider suggestions from Users for further development of BEAR, and hear any concerns.
8. guide and oversee implementation of the strategy for BEAR.
9. make recommendations on University strategy as it affects the operation of BEAR.
10. collaborate with IT Services to create and agree the research computing strategy as it relates to BEAR.
11. appraise proposals from research groups to fund major enhancements to BEAR which have been referred to the Committee by ARC.
12. through a standing sub-committee, consisting of at least 4 academic members of the Committee, manage applications for allocation of the University's share of processing capacity in HPC Midlands Plus and subsequent Tier-2 services where Birmingham is a consortium member, and other national services.
13. oversee and advise on any costs or charging model which may be applied for use of BEAR.
14. guide the development and advise on the Communication and Engagement Plan for BEAR.
15. consider and adjudicate on requests for access to BEAR for external organisations, or individuals who are not members of the University, and any charges which may apply.

REPORTING LINE

The Management Committee reports to the Academic IT Reference Group and then to University Research Committee.

Last updated:
July 18

Appendix 7 : Membership of the Research Computing Management Committee (August 2018)

BEAR Management Group

Dr Colin Rowat	CoSS, Economics (Chair)
Dr Alessandro Mottura	EPS, Metallurgy and Materials
Dr Hassan Hemida	EPS, Civil Eng
Prof David Smith	EPS, Maths
Dr Xiaoming Cai	LES, GEES
Prof George Gkoutos	MDS, Cancer Sciences
Dr Henry Chapman	CAL, Archaeology
Dr Liza Jabbour	CoSS, Economics
Prof Jean-Baptiste Cazier	MDS, Cancer Sciences
Dr Louise Van Zoest	LES, Psychology
Prof Andrew Howes	EPS, Computer Science
Dr John Owen	ARC, IT Services
Carol Sandys	ARC, IT Services
Debbie Carter	ARC, IT Services (Secretary)
Aslam Ghumra	ARC, IT Services

Appendix 8 : Technical Details of BEAR Infrastructure

As at June 2018, the infrastructure components of BEAR are:

Compute

BlueBEAR, our Linux-based batch processing High Performance Computing (HPC) cluster

2017-18 saw a significant expansion in BlueBEAR, including launching the new water-cooled compute nodes. However, this will be tempered by the need to decommission old BlueBEAR 2 (SandyBridge) compute nodes. The shared compute node types accessible to all researchers in the University are currently:

- 75 older SandyBridge-based compute nodes (16 cores, 32GB RAM)
- 5 older SandyBridge-based large memory nodes (16 cores, 128 or 256GB RAM)
- 3 older SandyBridge-based GPU nodes (16 cores, 32GB RAM, NVIDIA K20 or Q5000 GPUs)
- 24 Haswell-based compute nodes (24 cores, 128GB RAM)
- 92 Broadwell-based compute nodes (20 cores, 128GB RAM)
- 7 Broadwell-based large memory nodes (20 cores, 256GB or 512GB RAM)
- 2 Broadwell-based GPU nodes (20 cores, 128GB RAM, P100 GPU)
- 1 IBM Power9 node with 1TB RAM and 4 NVIDIA V100 GPUs

In addition to these shared resources, several research groups have invested in BEAR, paying for dedicated compute resources.

BEAR Cloud: the next generation of powerful compute resources built for flexibility, accessibility and efficiency. BEAR Cloud capacity includes a number of systems purchased with CaStLeS funding.

- 121 hypervisor compute nodes (Broadwell 20 cores, 128GB RAM)

BEAR Plus, additional computational resources available to Birmingham researchers through the regional HPC Midlands Plus consortium and other national resources (e.g. ARCHER). Some of these national resources may be paid-for services.

- **CaStLeS Compute:** CaStLeS (Compute and Storage for Life Sciences) provides powerful compute and storage to meet the challenges of life sciences research in all Colleges.

CaStLeS includes two compute components, virtual machines as part of BEAR Cloud, and a set of compute nodes associated with BlueBEAR:

- 1 large memory node with 1TB RAM and a pair of NVIDIA P100 GPUs
- 1 large memory node with 1TB RAM
- 1 IBM Power9 node with 1TB RAM and 4 NVIDIA V100 GPUs
- 12 standard compute nodes (Broadwell 20 cores, 128GB RAM)

BEAR Galaxy (for computational biologists) - a platform for data-intensive biomedical research (analysis and process sharing)

Storage

BEAR Research Data Store (RDS), providing storage for 'work-in-progress' data. Data storage is intended to be accessible across BEAR platforms as well as to client end-points and specialised equipment across the campus. The research data storage platform provides high-availability, reliable and backed-up storage. The RDS currently operates with 2 synchronous copies. The architecture for long term growth and sustainability of storage services is currently under review to ensure storage can be provided at a cost-competitive price point. The research data store includes capacity funded from CaStLeS also.

BEAR Research Data Archive (RDA), delivering resilient storage exclusively for the retention of data associated with research projects and publications.

BEAR database service. The service runs on MariaDB which is compatible with MySQL. The service runs in a high-availability cluster and is available both from the BlueBEAR cluster and from the campus network.

BEAR DataShare, offering dropbox-type facilities for collaboration both on and off campus

Ancillary Services

BEAR Gitlab, providing code control, revision management and issue tracking to enable researchers to produce and manage better software.

BEAR View, offering visualisation, collaboration and large-scale video-conferencing facilities including a 4 metre by 2 metre display screen, 2D and 3-D imaging and movement tracking.

BEAR Research Data Network (RDN) to cater for the fast transmission of high volume data, generated by experimental equipment on campus.