Enabling the UK knowledge economy

A new STFC e-infrastructure strategy designed to support and strengthen UK research and UK industrial productivity

"As well as physical and digital infrastructure, we need to make sure that we also have in place an effective data infrastructure. This means the right elements for an economy in which open data drives growth, efficiency and innovation."

UK Industrial Strategy Green Paper 2017

Science & Technology Facilities Council

V2.7-final (25/01/18)

Contents

1. Introduction

1.1 Over 100,000 UK scientists, from both academia and industry, rely upon digital or 'e' infrastructure to extract value and knowledge from experimental data[1]. This knowledge underpins the UK economy, facilitating growth, job creation and innovation. For the purposes of this strategy e-infrastructure is defined as the combination of:

- Computers and software (for modelling, simulation, calculation and data analysis);
- Data storage facilities, data curation and data archiving;
- Software infrastructure for managing platforms, virtual organisations, workloads, data and security;
- The network connecting it all together (enabling data to be accessed and moved);
- The people with the necessary expertise in scientific software and data intensive science.

1.2 This strategy identifies the e-infrastructure challenges facing the Science and Technology Facilities Council (STFC), its research communities and the UK. It also highlights where, in tackling these challenges, additional strategic investment is considered essential. STFC is also playing a leading role in working with Research Councils UK (RCUK) to develop the national e-infrastructure (NeI) that the UK requires for the future and in doing so it can draw upon its breadth and depth of experience in this area. This in turn will help achieve STFC's strategic goals of delivering world-class research, innovation and skills and deliver our ambition to be a world leader in the analysis and application of data.

1.3 We know that putting in place the e-infrastructure which is needed by UK researchers cannot be achieved by STFC acting in isolation. Increasingly there is a need to share both expertise and resources across organisational boundaries and, in many cases, across international boundaries. Against this background, STFC's e-infrastructure strategy will contribute to a larger debate with the other Research Councils and partner organisations such as Joint Information Systems Committee (Jisc), the UK Met Office and the Culham Centre for Fusion Energy (CCFE). With input from industry from the E-infrastructure Leadership Council (ELC), these deliberations will ultimately lead to the development of an agreed NeI strategy for implementation by UK Research and Innovation and BEIS.

1.4 The *Industrial Strategy Green Paper* published in January 2017 recognised the need to put in place the "right elements for an economy in which open data drives growth, efficiency and innovation". STFC supports the drive for scientific data to be open to all, recognising that holding data in such a way that it can be interrogated and analysed easily can help maximise the value drawn from the UK's investment in research. Supporting open and reproducible science in turn accelerates the transformation of research into commercial advantage. These important principles run through many of the actions proposed in this document.

1.5 This strategy is written at a time of considerable pressure on public expenditure and the case for any additional investment needs to be especially compelling. Sometimes investment in e-infrastructure can be seen as an overhead cost, but this strategy shows that it is, in fact, an essential enabler of growth. Experience overseas echoes this, with a recent review of Australian research infrastructure highlighting the need for a strategic,

[1] See for example, the website Statista.com for numbers of physical, chemical, biological and biochemical scientists working in the UK over the period 2011-2017.

whole-of-Government approach to the funding of critical national research infrastructure. It is interesting to note the parallels here with the UK's own House of Lords Inquiry in 2013 which also identified the need for a long-term strategy and investment plan for all national scientific infrastructure.

1.6 Whilst some incremental improvement may be possible within the restrictions of existing resources, the absence of a broadly agreed NeI strategy, together with an underpinning national investment plan, potentially puts the UK economy at a disadvantage when compared to our competitors. We believe swift action is needed to address this, but there is a firm foundation on which to build and STFC welcomes this opportunity to set out its own e-infrastructure vision. Comments on the direction of travel we have outlined here would be very welcome.

2. Executive summary

2.1 <u>*Helping develop a sustainable UK wide e-infrastructure strategy and investment plan*</u>

The UK's national research infrastructure deserves the same recognition as any other productivity enhancing part of the UK's national infrastructure. STFC is well placed to work with partners to help develop a UK wide e-infrastructure strategy and investment plan which is fit for the 21st Century.

Action 1 – Helping develop a sustainable UK wide e-infrastructure strategy and investment plan

STFC will continue to play a major role in the development of a UK wide e-infrastructure strategy. In doing this, STFC is collaborating with the Engineering and Physical Sciences Research Council (EPSRC) and other Research Council partners through the Research Councils UK (RCUK) NeI group. This will help inform the future work of UK Research and Innovation as it develops a new UK research infrastructure roadmap and capital investment plan and help guide the strategic discussions taking place within the industry focused UK E-infrastructure Leadership Council (ELC).

2.2 <u>*Developing the skills the UK requires*</u>

The Government's *Industrial Strategy Green Paper* identified that the UK has an urgent need for talent development both in research computing, in high performance data analytics, and in data and information management. Most notably, scientific software development and advanced data analytics are both areas where skills are in very short supply. STFC contributes directly to tackling this deficit and developing UK talent through a range of different initiatives and its research activities. But there is a need for more investment to develop the next generation of data-intensive science and data-intensive engineering experts. By keeping the skills of its team fresh, and attracting the best talent, STFC will help to provide new and innovative solutions to the various e-infrastructure challenges the UK faces.

Action 2 – Investing in skill development

STFC has a key role to play in developing the very specialist scientific computing, research software engineering and data analytic skills that the UK needs. It will seek to expand this role as opportunities occur and will ensure that skills development is central to existing initiatives as well proposed new initiatives.

2.3 <u>*Co-ordinating STFC's scientific computing efforts*</u>

Recognising that resources are finite, it is vital that STFC's e-infrastructure work is fully co-ordinated both across the organisation and with partners. To this end, it is proposed to establish a new 'Scientific e-Infrastructure Governance and Oversight' group.

Action 3 – Co-ordinating STFC's scientific computing efforts.

STFC will establish a 'Scientific e-Infrastructure Governance and Oversight' group with cross directorate and facility representation. This group will be accountable to STFC Executive Board for the delivery of the strategy and for ensuring co-ordinated action across the scientific computing requirements of all STFC's programmes, and with partners such as the other Research Councils. This will require the development of a prioritised e-infrastructure roadmap and investment plan closely dovetailed with the development of the UK wide e-infrastructure strategy and investment plan. It will also include appraisal of emerging opportunities such as cloud computing.

2.4 <u>*Meeting the need for access to High Throughput Computing*</u>

The data processing, analysis and event simulation requirements of particle physics, astronomy and nuclear physics (PPAN) researchers, facilities users and others require advanced High Throughput Computing (HTC) facilities. By co-operating across organisational and project boundaries to make UKT0 a reality it will be possible to ensure the consolidation of computing resources and provide HTC access in a cost effective manner.

Action 4 – Meeting the need for access to High Throughput Computing

Recognising the need for HTC access for UK researchers, and the need for cost-effective provision, STFC will continue to support the UKT0 consortium and work with partners to secure the investment needed to provide UK researchers with access to a UKT0 facility.

2.5 <u>*High Performance Computing capability for UK researchers*</u>

The needs of academic and industrial researchers are changing. The need to develop increasingly complex models and simulations and to manipulate ever larger sets of data brings in its wake demands for access to High Performance Computing (HPC) and data intensive science capabilities.

Action 5 – Helping provide up-to-date High Performance Computing capability for UK researchers

Recognising the critical necessity of HPC capability for its programme as part of an integrated UK landscape, STFC supports: (i) The proposed replacement for the current HPC provision provided through the Distributed Research Utilising Advanced Computing (DiRAC) Facility; (ii) ARCHER 2; and (iii) Continued access through Partnership for Advanced Computing in Europe (PRACE). These assets, together with those at the Hartree Centre, will ensure that STFC researchers have access to the world-class facilities they need for the medium term. STFC is committed to working with key partners, including the other Research Councils, UK Research and Innovation, and others, to ensure that HPC provision is co-ordinated as part of the NeI capability.

2.6 *Building UK High Performance Data Analytics capability*

High Performance Data Analytics (HPDA) is an increasingly important and emerging field whereby extremely large datasets can be analysed to uncover useful information and meaningful patterns. It can be considered to be the specialist intersection between high performance computing and increasingly large datasets. It is a resource needed by both academic and industrial researchers.

Action 6 – Building UK High Performance Data Analytics capability

Through the members of the DiRAC community, and through the Hartree Centre and the Ada Lovelace Centre, STFC will collaborate with partners in projects such as JASMIN (a NERC funded project) to further develop its existing HPDA skills base. This will enable STFC to provide a data science service to both industry and academia which will complement the research and services provided by EPSRC's Alan Turing Institute.

2.7 *Improving the UK research data network*

The UK research community has an increasing need to rapidly transfer research data between organisations, institutes and research facilities, both within the UK and internationally. High bandwidth and good end-to-end network performance is crucial in meeting this need.

Action 7 – Improving the UK research data network

STFC supports initiatives to improve the UK research data network. As part of this it will seek investment to improve network capability between STFC sites, between STFC-supported research facilities and with major university user communities. It will do this in liaison with Jisc's Joint Academic Network (Janet) and other stakeholders to obtain the end-to-end performance improvements needed by UK scientists.

2.8 *Improving productivity by establishing the Ada Lovelace Centre*

Large scientific facilities are becoming generators of ever increasing amounts of data. The quantity of data and complexity of experiments requires both advanced computational modelling and data science expertise. The Ada Lovelace Centre is a new approach that concentrates scientific computing resources and skills alongside the Harwell-based research facilities. This has the potential to generate a step change in productivity, accelerating the translation of experimental data into research outputs.

Action 8 – Improving productivity by establishing the Ada Lovelace Centre *The rapidly changing nature of modern research requires a step change in the way researchers access the scientific computing resources they need. STFC, and CCFE, will seek investment to establish the Ada Lovelace Centre for computational modelling and data science alongside the UK large research facilities based at the Harwell Campus and at Culham. This will provide an important catalyst that has the potential to raise UK scientific productivity, improve the competitive edge of UK science and extract the most value from the multi-billion pound investment in research facilities.*

The diagram below shows the relationships between the RCUK NeI, UKT0 and the Ada Lovelace Centre[2].

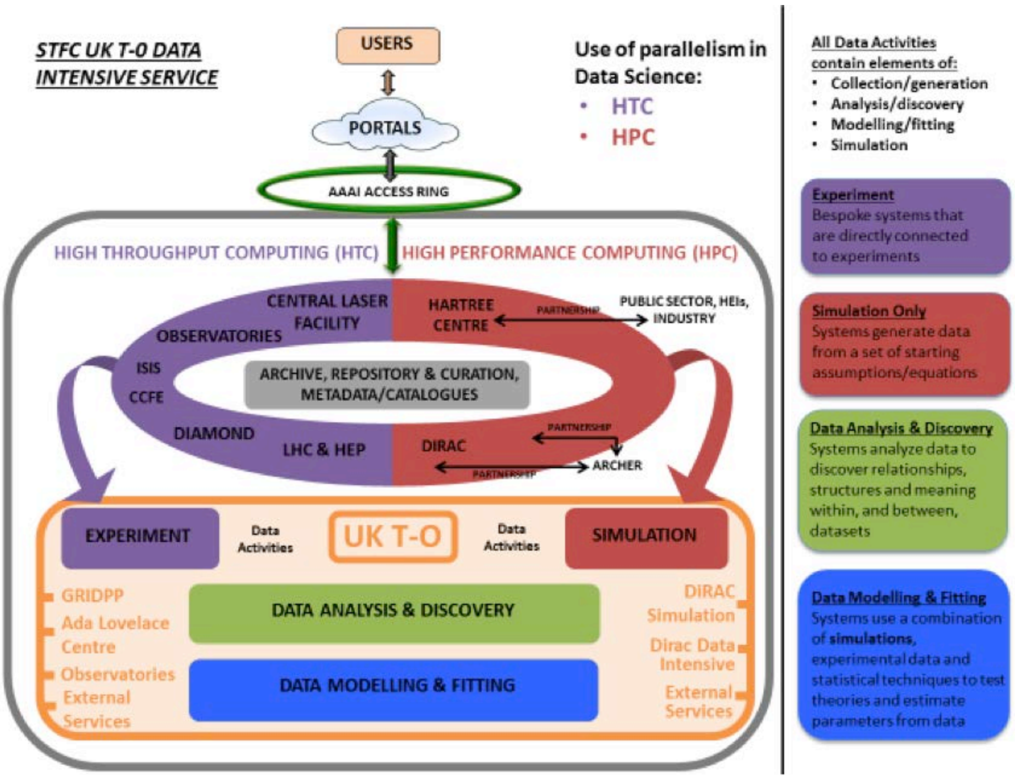

2.9 *Building competitive advantage*

[2] Diagram courtesy of Jeremy Yates and Clare Jenner.

The Hartree Centre seeks to transform UK industrial competitiveness by accelerating the adoption of HPC, HPDA and cognitive technologies. It has the promise to deliver transformative gains for UK industry through collaborative industrial research projects.

> ***Action 9 – Building competitive advantage***
>
> *The Hartree Centre is now an increasingly important resource for UK industry and a centre of excellence in terms of how to apply HPC, cognitive computing and Big Data expertise to a wide variety of industrial challenges. Continuing to develop the Hartree Centre's collaborative approach with industry will boost the UK's competitive edge and help deliver economic growth and job creation. The success to date shows that the Hartree approach can deliver competitive advantage to UK industry, but investment needs to be maintained over the long term.*

2.10 *A leadership role*

UK e-infrastructure must be developed with international collaboration in mind. STFC is uniquely placed to provide the input needed to ensure that the UK's view is taken into account in the development of any collaborative infrastructure across international borders.

> ***Action 10 – A leadership role***
>
> *The uniquely global nature of STFC's science and research requires us to plan, establish and operate infrastructures not only across organisational boundaries, but across international boundaries. STFC is in a key position to continue to provide leadership to those international collaborations relating to e-infrastructure and to link this to the strategic discussions taking place in the UK as it develops its own e-infrastructure strategy. It will therefore build on its existing position to play an active role on the international stage as an advocate for the UK's e-infrastructure interests.*

3. The STFC e-infrastructure vision

3.1 STFC's e-infrastructure vision is straightforward:

"To work with our partners to develop and deliver cutting-edge e-infrastructure solutions for UK academia and industry in order to advance computational and data-intensive science and innovation".

3.2 Successfully delivering this vision will involve developing the world-class e-infrastructure components that UK researchers need to maximise their scientific, technological and economic output. These researchers include all those who use STFC-supported research facilities either in the UK or overseas – whether they are industrial users from large companies or SME's, or those from the academic research community. In this context it should be noted that although STFC directly supports research in particle physics, particle astrophysics, astronomy, accelerators and nuclear physics (PPAN)[3], its role in providing access to large-scale research facilities and e-infrastructure extends to a much broader research community.

3.3 In 2015 an STFC Computing Strategic Review was undertaken which assessed the organisation's computing (e-infrastructure) capability and emerging requirements. In summary, the review's key recommendations included meeting the growing needs of data-intensive science, sharing resources across different communities, building skills and taking a leadership role in all aspects of e-infrastructure, both at home and overseas. In 2016, STFC undertook an exercise examining the balance of its research programme and this identified that computing support is an area where demands are growing rapidly. The findings of both these reviews remain valid and have been used to help inform and guide this strategy.

4. The key e-infrastructure challenges facing STFC and the UK

4.1 In recent years a number of key e-infrastructure challenges have emerged which provide stimulus for change. These include:

a) There is a need to develop a UK wide sustainable NeI strategy and investment roadmap – *with adequate resourcing this could meet the present and future computing, data and networking needs of academia and industry;*

b) UK researchers require STFC to be at the forefront of scientific computing resources, HPDA, research data networking and scientific software systems if we are to remain competitive on the world stage – *this means that the UK needs to consider each of these elements and how best to provide a world-class research 'eco-system' which meets the needs of UK researchers;*

c) Research experiments carried out at the UK research facilities supported by STFC or elsewhere at international facilities such as CERN, European Synchrotron Radiation Facility (ESRF) and European Southern Observatory (ESO) are generating ever greater volumes of data – *the opportunity exists to transform the efficiency of research in the UK;*

[3] The PPAN research community.

d) There is a need for UK companies to be supported in extracting maximum competitive advantage from HPC and data analytics – *the Hartree Centre has the potential to be one of the world's foremost centres for collaborative industrial research.*

e) Ensuring that UK science sustains its global impact – *for UK researchers to maintain their leading role in global science programmes, UK e-infrastructure and international e-infrastructure strategies need to dovetail together.*

5. Building on a successful base

5.1 STFC has three distinct but interrelated functions:

- Frontier science: developing and exploiting a world leading programme of particle physics, astronomy, nuclear physics and space science, to resolve some of the most ambitious and challenging questions facing us today and support innovation and skills development in our partner universities, research institutes and facilities;

- Scientific facilities: planning, designing, constructing and providing stewardship for world-leading, large-scale research facilities used by a broad range of academic and industrial researchers, enabling and driving research, innovation and skills training to deliver national multidisciplinary research priorities;

- Research and Innovation Campuses: building national research and innovation campuses around our facilities and laboratories to promote hubs for academic and industrial collaboration on research, skills training and smoothing the pathway to market thereby stimulating jobs and growth.

5.2 These three programmes are underpinned and linked by five world-leading capabilities in technology that cannot be cost-effectively and reliably procured on the open market: detectors and instrumentation, accelerators, specialist engineering, optics and e-infrastructure.

5.3 In carrying out these three functions, STFC provides a range of e-infrastructure support activities, some of which are operational in nature and others which are more developmental. Delivery is through the STFC Scientific Computing Department (SCD)[4], the STFC Programmes Directorate, the Hartree Centre, and Digital Solutions (formerly Corporate ICT). Collectively these activities can be considered to fall within the following broad areas:

a) Supporting the diverse e-infrastructure needs of the PPAN community:

- The Programmes Directorate directly supports PPAN research and skills development through a broad programme of activity, both in universities and at research facilities. It also provides support to the DiRAC community to provide access to dedicated supercomputer capability across five complementary installations. The DiRAC Facility is essential to astronomy, particle astrophysics, cosmology, nuclear physics and particle physics researchers. Further details of the DiRAC Facility are attached at Annex A.

[4] SCD comprises around 150 full-time equivalent staff and has a total budget of around £12m. Around 80% of this budget comes from sources outside STFC.

- The operation of the Large Hadron Collider (LHC) brings with it significant data and compute challenges in manipulating and analysing the experimental data generated. With funding from the Programmes Directorate, SCD has for many years provided computing support for GridPP[5], directly operating the UK's 'Tier 1'[6] LHC computer facility and supporting university 'Tier 2' centres around the UK.

b) Supporting the effective operation of the UK's large-scale research facilities

- The Harwell Campus in Oxfordshire is home to many of the UK's large-scale research facilities including the Diamond Light Source (the UK's national synchrotron science facility), the ISIS Neutron and Muon Source and the Central Laser Facility. These are world-class facilities and keeping them at the forefront of science requires continual upgrades and improvements. STFC supports this ongoing process with its e-infrastructure expertise, including specialist scientific software development, data analysis and data storage services. It is currently engaged in a new pilot program to enable the better exploitation of research by closer coupling of software, data and compute resources. This is a close collaboration between STFC facilities and the CCFE and is a precursor to an ambition to establish a new multi-disciplinary centre (the proposed Ada Lovelace Centre). This concept is discussed further in Section 12.

c) Working in partnership to assist other research communities:

- STFC's SCD has worked with RAL Space and the Natural Environment Research Council (NERC) to design, build and support an innovative SuperData Cluster for High Performance Data Analytics (HPDA). This allows researchers such as those investigating climate change to utilise and access ever increasing environmental data sets. This was a genuine world-first system – known as JASMIN – and similar HPDA systems are now being built in the USA.

- STFC works in collaboration with the EPSRC and other Research Councils to provide scientific software support for the UK research community. The Collaborative Computational Projects (CCP) initiative is now over 30 years old and SCD staff support the key scientific application programs for nine CCP collaborations and five High End Computing Consortia for implementation on leading edge HPC systems. In addition, the Software Engineering Support Centre (SESC) is another EPSRC-funded project hosted by SCD. SESC aims to promote software engineering in the research community by providing advice and training on practical software engineering tools and techniques.

d) Data networking for the UK research community

- Digital Solutions supports the existing STFC network. However, the existing network arrangements are not sufficient to allow the fast transfer of experimental data between facilities and external institutions. Current and future data transfer demands from the UK research community generally, from the PPAN community and from the

[5] GridPP manages the UK's involvement in CERN's Large Hadron Collider Computing Grid project and is a collaboration of Particle Physicists and Computing Scientists from 19 UK universities, STFC and CERN.
[6] There are two definitions of 'tiers'. These are: the 'Branscomb' Supercomputing/HPC tiers, namely 0 (leadership), 1 (national), 2 (regional), and the CERN High Throughput Computing (HTC) tiers, namely 0 (CERN), 1 (national e.g. RAL), and 2 (regional/university) centres.

large research facilities are such that a research data network is required to enable the rapid transfer of large quantities of data between organisations.

- To achieve this, high bandwidth connections supporting good end-to-end performance are crucial. Jisc's Joint Academic Network (Janet) is capable of allowing the fast transfer of data, but its capability can only be fully realised if all parts of the network are fit for purpose and the end points do not constitute a bottleneck restricting performance.

e) Providing support for industrial usage of HPC, Big Data and cognitive technologies:

- The Hartree Centre, located at the Sci-Tech Daresbury campus in Cheshire, is supported by significant direct investment by the UK Government and IBM. Its aim is to help transform the competitiveness of UK industry by providing industrial access to HPC resources for simulation and modelling as well as cognitive technologies for data analytics, together with the associated skills and knowledge needed. Key strengths of the Hartree Centre are access to competitive HPC as well as to HPDA resources, and the ability to bring together cross-disciplinary teams to bridge the gap between academic research and provide solutions which meet particular industry needs.

f) Providing e-infrastructure leadership within the UK and overseas:

- STFC is one of the few organisations with the depth of e-infrastructure knowledge and expertise necessary to allow it to provide strategic leadership. STFC does this by providing input at both national and international levels, working in partnership with many organisations. One example is collaborative working within the EU-T0 consortium whose vision is to "Create a hub of knowledge and expertise to co-ordinate technological development and optimise the investment of the funding agencies in their existing data centres by broadening, simplifying, and harmonising, driven by well-defined user requirements".

- STFC is also an active participant in the Worldwide Large Hadron Collider (LHC) Computing Grid (WLCG) project, a collaboration of more than 170 computing centres in 42 countries which provides global computing resources to store, distribute and analyse the data generated by the LHC. This expertise in Big Data is now being brought to bear in the design of the European Square Kilometre Array (SKA) data centre. This type of strategic input is increasingly important as computational and data intensive analytic approaches have become ubiquitous research tools.

- STFC is also a major player in many other European e-infrastructure initiatives. Amongst others, these include the European Open Science Cloud (EOSC)[7], EGi[8] and EUDAT[9], as well as the Research Data Alliance (RDA)[10] and many other EU Horizon 2020 projects.

[7] EOSC (European Open Science Cloud) is a collaborative venture to support Open Science in Europe.
[8] EGi is a federation of 20 cloud providers and over 300 data centres, spread across Europe.
[9] EUDAT's mission is to develop pan-European collaborative data infrastructure for science.
[10] RDA seeks to develop and adopt international infrastructure that promotes data-sharing and data-driven research.

g) Maintaining and building specialist e-infrastructure skills

- The skill and expertise of those working within STFC are vital to the UK's ability to extract value and knowledge from HPC and from Big Data. Maintaining and building the skill base is therefore essential and STFC is engaged in a broad spectrum of training-related activities from computing apprenticeships to post-doctoral training in computational science and data science; STFC has recently funded eight university-based Centres for Doctoral Training (CDT) in Data Intensive Science.

- It is important that STFC keeps abreast of new developments (such as new computer architectures) in order that the opportunities such advances offer can be fully evaluated.

5.4 The STFC track record in scientific computing is impressive and the impact of its work is felt in many areas of science. For example, the DiRAC Facility was established to provide world-leading distributed HPC services to the PPAN research community. HPC-based modelling is an essential tool for the exploitation of observational and experimental facilities in astronomy and particle physics, as this technology allows scientists to test their theories and run simulations from the data gathered in experiments. A second example is the work of SCD within the CCPs that bring together leading UK expertise in key fields of computational research to tackle large-scale scientific software development. Each project represents many years of intellectual and financial investment and helps achieve optimum use of the whole range of hardware available to the scientific community, from the desktop to the most powerful national supercomputing facilities.

5.5 STFC is therefore well placed to respond positively to the challenges identified in section 4 above, and has the ambition and skills needed to do so. But meeting the ever-increasing demands of its 'customers' within the restrictions of existing levels of investment is increasingly difficult.

6. Partnership, collaboration and strategy

6.1 As can be seen from the description of activities above, the STFC approach has been to work in collaboration with others wherever this helps to achieve STFC organisational objectives and realise the e-infrastructure vision. This provides a good way of sharing both costs and risks. In some cases, this may involve supporting formal partnerships or consortia arrangements, or working directly to forge new partnerships with industry or universities. In others, it is using influence and advocacy to steer the strategic direction of international collaborations such that the maximum benefit can be derived for UK science into the future.

6.2 Co-ordination of HPC and other e-infrastructure activities for academic research in the UK is the joint responsibility of RCUK. STFC has been working with partners within RCUK to develop a roadmap that seeks to optimise the provision of scientific computing facilities and data e-infrastructure across the UK. The RCUK NeI roadmap is being developed jointly by all the Research Councils together with input from Jisc, the Met Office and industry. The roadmap is based upon an intention that multiple locations can work together to offer complementary scientific computing and data services. This work will be an important input for the development of a comprehensive national plan that also takes account of the needs of industry.

6.3 The impending creation of UK Research and Innovation in 2018 will bring together the seven Research Councils, Innovate UK and the research funding and knowledge exchange elements of the Higher Education Funding Council for England (HEFCE) into a single strategic body. This provides a great opportunity for even closer collaboration and a unique opportunity to promote new ideas and new approaches. Notably, the suggestion within the recent *Industrial Strategy Green Paper* that UK Research and Innovation could develop a new research infrastructure roadmap to provide the modern infrastructure needed to support research is especially welcome. STFC looks forward to playing an active part in these important discussions.

6.4 At Governmental level, the ELC provides a high profile forum for strategic ministerial level discussion of UK wide e-infrastructure issues with representatives from industry. In the future, working within UK Research and Innovation, STFC would wish to continue to provide input to ELC's deliberations[11] and assist in the development of both a NeI strategy and roadmap and a capital investment plan.

6.5 The UK's national research e-infrastructure deserves the same recognition as any other productivity-enhancing part of the UK's national infrastructure, or arguably higher recognition, since this is a major growth area for the future. Delivering internationally competitive e-infrastructure on a long-term and sustainable footing requires collaborative action by a wide range of organisations delivering on a shared agenda.

Action 1 – Helping develop a sustainable UK wide e-infrastructure strategy and investment plan

STFC will continue to play a major role in the development of a sustainable UK wide e-infrastructure strategy. In doing this, STFC is collaborating with EPSRC and other Research Council partners through the RCUK NeI group. This will help inform the future work of UK Research and Innovation as it develops a new UK research infrastructure roadmap and capital investment plan, and help guide the strategic discussions taking place within the industry focused UK ELC.

7. Developing the critical skills the UK needs

7.1 The Government's *Industrial Strategy Green Paper* identified that the UK has an urgent need for talent development both in research computing, in HPDA, and in data and information management. Most notably, scientific software development and advanced data analytics are both areas where skills are in very short supply. STFC contributes directly to tackling this deficit and developing UK talent through a range of different initiatives:

- Provision of computer apprenticeship placements;
- Graduate industrial work placements;
- Specialist training such as provision of specialist summer schools;
- PhD placements[12];

[11] Professor Tony Hey, STFC Chief Data Scientist, is the present Co-Chair of the ELC.

[12] Between 2012 and 2015, 941 PhD students were funded by STFC in the areas of astronomy, nuclear physics and particle physics. Of those, 28% of the graduates took up positions in the private sector, with over 70% of those working in software development, data analysis, engineering and finance.

- Supporting eight CDTs in data intensive science[13] and the development of data science as a distinct specialism for scientists;
- Joint appointments with industry, such as those under the auspices of the DiRAC consortium;
- The promotion of research software engineering as a distinct specialism, recognising the importance of nurturing individuals who combine expertise in scientific programming with an intricate understanding of research.

7.2 It is through application of these specialist skills that the research problems faced by industrial scientists and their academic colleagues can be tackled. Whilst STFC maintains a core of highly skilled computing professionals, many of those who have participated in training within STFC have eventually taken their specialist skills to work in industry or to set up new high-tech companies. In this way STFC helps contribute to the STEM skills pipeline, helping to generate the wealth and jobs needed by the UK economy.

7.3 The recent decision by Government to increase the number of PhD studentships in STEM subjects is a very positive step and will, in time, help address the existing skills shortage in this area. There is a need for more investment to develop the next generation of data-intensive science and data-intensive engineering experts and STFC will continue to play its part in this. By keeping the skills of its team fresh, and attracting the best talent, STFC will help to provide new and innovative solutions to the various e-infrastructure challenges the UK faces.

Action 2 – Investing in skill development

STFC has a key role to play in developing the very specialist scientific computing, research software engineering and data analytic skills that the UK needs. It will seek to expand this role as opportunities occur and will ensure that skills development is central to existing initiatives as well proposed new initiatives.

8. Tackling the challenges – Keeping the UK's competitive edge in research

8.1 A changing environment

8.1.1 The UK punches above its weight as a research nation. A 2013 report by Elsevier examining the comparative performance of the UK research base[14] found that whilst the UK represents just 0.9% of global population, 3.2% of R&D expenditure, and 4.1% of researchers, it accounts for 9.5% of downloads, 11.6% of citations and 15.9% of the world's most highly-cited articles. This position can only be maintained in the future if the UK is at the forefront of scientific computing, data analytics, research data networking and scientific software systems. This requires leading edge research e-infrastructure.

8.1.2 The research landscape is however changing fast. Large-scale research facilities both in the UK and overseas are generating very large – and very rapidly increasing – volumes of experimental data. With the ever increasing size of datasets, the common access model for industry and academia of "search, access, download and use" will no longer be viable.

[13] Through the STFC Programmes Directorate, 88 students are supported across the eight CDTs, with a further 40 or so funded from other sources. Each of the CDTs brings together a range of academic and industrial partners.

[14] 2013 report prepared by Elsevier for the former Department of Business, Innovation and Skills (BIS).

8.1.3 Connecting data to users, and to appropriate data intensive computing analytics, requires high bandwidth and end-to-end network performance not currently in place across the UK. What is needed is the development of national data infrastructure and expertise that can facilitate the inter-linking of nationally important datasets and research facilities[15]. Achieving this will require significant re-design of existing local network structures and needs to be a central part of the UK strategy referred to earlier, addressing the changing needs of today's researchers. The emergence of UK Research and Innovation provides an ideal opportunity to move in this direction in a co-ordinated fashion.

8.1.4 From an STFC perspective, this overall direction of travel requires:

- Localised investment in improving the research data network associated with the UK's large-scale experimental facilities (mainly based at STFC's Harwell and Daresbury campuses). This will help support not only the research undertaken by the PPAN community supported by STFC, but also those researchers from other disciplines utilising large research facilities;
- Investment in computing hardware and data storage facilities within STFC to cope with the increasingly large and complex scientific data sets now being generated. This will help maximise the efficiency of research and its translation into innovation;
- Maintaining support for the DiRAC Facility and for access to PRACE, utilised by many PPAN researchers and others;
- Continued investment in the Hartree Centre's supercomputing and cognitive computing resources to maintain its unique role in assisting UK industry in exploiting these technologies;
- Continuing with specialist training initiatives which will help provide the expertise needed by the UK.

8.1.5 Whilst recognising that resources are finite, there is little doubt that without investment of this nature the competitive advantage of UK research will be rapidly eroded.

8.1.6 Against this background, there is a need to ensure that action is co-ordinated across all areas of STFC activity and it is proposed to establish an STFC 'Scientific Computing Governance and Oversight' group which crosses facility and directorate boundaries. This group will be led by the Director of SCD, and will be tasked with delivering this strategy as well as the co-ordination, planning and organising of STFC's scientific computing efforts. It will complement the work of STFC's Computing Advisory Panel and Science Board. In this way it will be possible to extract the most cost effective return from STFC investment and deliver improved research outcomes. It will also be consistent with the findings of the 2015 Computing Strategic Review which calls for STFC areas to work more closely together to find ways of sharing computing infrastructure, with new projects encouraged to make use of existing expertise and infrastructure.

[15] This is an approach that both Germany and the United States are adopting.

Action 3 – Co-ordinating STFC's scientific computing efforts.

STFC will establish a 'Scientific e-Infrastructure Governance and Oversight' group with cross directorate and facility representation. This group will be accountable to STFC Executive Board for the delivery of the strategy and for ensuring co-ordinated action across the scientific computing requirements of all STFC's programmes, and with partners such as the other Research Councils. This will require the development of a prioritised e-infrastructure roadmap and investment plan closely dovetailed with the development of the UK wide e-infrastructure strategy and investment plan. It will also include appraisal of emerging opportunities such as cloud computing.

8.2 Providing access to world-class compute resources

8.2.1 In general the required computing infrastructure for world-class research may be broadly characterised by three types of computer architectures:

- HTC – needed for very large-scale data applications;
- Supercomputing/HPC – for applications requiring low latency, high-bandwidth communications between large numbers of compute cores;
- HPDA – for data intensive applications requiring close coupling of vary large datasets with HPC resources.

8.2.2 Commercial cloud providers such as Amazon, Microsoft and Google are now offering a range of virtualised hardware and services that in future may become cost-effective alternatives for some types of scientific computing provision. The RCUK Cloud Working Group has recently investigated the current commercial offerings and the major conclusions of the group are attached at Annex B. At present although cloud computing is used for some scientific computing purposes, it does not yet provide a cost-effective alternative to 'direct' provision for most of STFC's leading edge HTC/HPC/HPDA applications. Cloud computing is also currently unsuitable for very high end HPC applications but this is a rapidly changing area and needs to be kept under review.

8.2.3 As we approach the end of Moore's Law[16], STFC needs to be aware of newly emerging technologies. Microsoft have deployed field-programmable gate array co-processors in its large-scale data centres and other IT companies are producing special purpose chips. For example, IBM have produced a very low power 'TrueNorth' chip to simulate complex neural networks and Google have developed a custom 'Tensor Processing Unit' designed specifically to speed up machine learning tasks. Universities are also exploring novel computer architectures. At the University of Manchester, the Advanced Processor Technologies Research Group have designed and built the SpiNNaker massively parallel computer that is inspired by the structure and function of the human brain. At the University of Oxford, the Networked Quantum Information Technologies Hub (NQIT) aims to deliver a prototype quantum computer on a very aggressive time-scale of only a few years. Such small scale quantum computers may well be of value for simulating real quantum systems, as first suggested by Nobel Prize winner, Richard Feynman. It is

[16] Moore's law is the observation that the number of transistors in a dense integrated circuit doubles approximately every two years. The miniaturisation of silicon technology has now reached the point where it is now no longer possible to shrink the feature size on the chip and increase the speed of the chip. There are now many discussions about the 'end of Moore's Law'. For example: https://www.technologyreview.com/s/601441/moores-law-is-dead-now-what/

important that STFC is aware of all these developments and able to advise on their possible exploitation for STFC science.

HTC

8.2.4 STFC is working with a range of partners in the UKT0 consortium[17] which seeks to provide access to HTC facilities through collaboration. By co-operating across organisational and project boundaries it will be possible to ensure the consolidation of computing resources and provide HTC access in a cost effective manner. This is a particular consideration for STFC given the data processing, analysis and event simulation requirements of PPAN research projects. The recently published *STFC PPAN Balance of Programme report* also acknowledges the need for additional funding and efficiency gains from the UKT0 initiative.

8.2.5 UKT0 is an initiative which has broad support from STFC supported scientists and facilities. It is a strong foundation on which to build and, as it becomes formalised, it will be a strong and effective body for co-ordinating compute resources across STFC projects. The UK's involvement in the SKA project will certainly place a large load on STFC's provision of HTC systems.

Action 4 – Meeting the need for access to High Throughput Computing

Recognising the need for HTC access for UK researchers, and the need for cost-effective provision, STFC will continue to support the UKT0 consortium and work with partners to secure the investment needed to provide UK researchers with access to a UKT0 facility.

Supercomputing/HPC

8.2.6 The current main HPC providers to scientists are:

- The DiRAC Facility. DiRAC provides the PPAN community with dedicated Tier 1 access;
- The Hartree Centre at Daresbury. The Hartree Centre is the major provider of supercomputing resources and expertise to industry;
- The Advanced Research Computing High End Resource (ARCHER) Tier 1 supercomputer based at the University of Edinburgh. This is the primary academic supercomputer for the non-PPAN UK research community. SCD staff provide HPC support for major ARCHER user communities;
- The EU PRACE consortium. Participation in PRACE allows STFC users access to European Tier 0 facilities.

8.2.7 The DiRAC community is examining options for upgrading their supercomputer capability whilst the Hartree Centre has just completed the procurement of its new supercomputing platform, Scafell Pike. In addition, EPSRC, together with input from Hartree and other organisations, is working to define the requirements for an ARCHER 2 machine as a replacement for ARCHER, the current national supercomputer for the non-PPAN research community.

[17] See Annex C for details of the UKT0 consortium.

Action 5 – Helping provide up-to-date High Performance Computing capability for UK researchers

Recognising the critical necessity of HPC capability for its programme as part of an integrated UK landscape, STFC supports: (i) The proposed replacement for the current HPC provision provided through DiRAC; (ii) ARCHER 2; and (iii) Continued access through PRACE. These assets, together with those at the Hartree Centre, will ensure that STFC researchers have access to the world-class facilities they need for the medium term. STFC is committed to working with key partners including the other Research Councils, UK Research and Innovation, and others to ensure that HPC provision is co-ordinated as part of the NeI capability.

HPDA

8.2.8 Data Analytics is an increasingly important and emerging field whereby extremely large datasets can be analysed to uncover useful information and meaningful patterns. It can be considered to be the specialist intersection between HPC and the increasingly large datasets becoming available to academic and industrial researchers. The need to develop ever larger complex models and simulations and to manipulate ever larger sets of data brings in its wake demands for access to HPC and data intensive science capabilities.

8.2.9 The Hartree Centre is collaborating with IBM on a major R&D project on Big Data and cognitive computing to investigate how cognitive computing platforms might be of benefit to UK industry. Most major manufacturers are now building systems that combine HPC and Big Data capabilities and in the future STFC will work with NERC and others to expand the capabilities of the JASMIN HPDA system and to deliver computing support for users of the large-scale facilities by the establishment of the Ada Lovelace Centre. There is also potential for the JASMIN and DiRAC teams to collaborate on the provision of common HPDA infrastructure for their communities, possibly utilising cloud technologies. This approach is supported by STFC and NERC.

8.2.10 The emphasis of the Hartree Centre on practical, collaborative Big Data projects with industry complements the more research-oriented focus of the Alan Turing Institute (ATI). Similarly, the data science activities of the Ada Lovelace Centre will be focused on the 'Big Scientific Data' produced by users of the Harwell site facilities. The ATI is a collaboration between the five founding universities – Cambridge, Edinburgh, Oxford, UCL and Warwick. Hartree and Ada Lovelace Centre staff will establish working collaborative relationships with both the London ATI centre and with the five university partners.

Action 6 – Building UK High Performance Data Analytics capability

Through the members of the DiRAC community, and through the Hartree Centre and the Ada Lovelace Centre, STFC will collaborate with partners in projects such as JASMIN (a NERC funded project) to further develop its existing HPDA skills base. This will enable STFC to provide a data science service to both industry and academia which will complement the research and services provided by EPSRC's ATI.

8.3. Research data networking

8.3.1 The UK research community has an increasing need to rapidly transfer research data between organisations, institutes and research facilities, both within the UK and further afield. High bandwidth and good end-to-end network performance is crucial in meeting this need but university network architectures typically lack the necessary 'fast paths' to support data intensive flows of research data.

8.3.2 Jisc's Joint Academic Network (Janet) is currently the key mechanism for allowing the fast transfer of experimental data. It provides the bandwidth and reliability required of the UK network backbone to enable the fast transfer of prodigious amounts of data. But this is insufficient if there are bottlenecks between university networks and the research facilities that generate the data. These bottlenecks can arise because university firewalls and policies associated with teaching and other needs of the university may conflict with the needs of data intensive research.

8.3.3 Whilst a number of universities have begun to address these issues, a coherent community wide approach is required. Without investment in appropriate local network engineering within university campuses, and between the research facilities themselves, UK researchers will be hindered by these limitations.

8.3.4 Implementing appropriate 'Research Data Transfer Zones' at the end points is an approach that will provide the end-to-end connectivity needed between the generators of high volume datasets at the national facilities and remote users at universities[18]. This needs to be coupled with providing users with easy-to-use software tools for initiating large data transfers.

Action 7 – Improving the UK research data network

STFC supports initiatives to improve the UK research data network. As part of this it will seek investment to improve network capability between STFC sites, between STFC supported research facilities and with major university user communities. It will do this in liaison with Jisc's Joint Academic Network (Janet) and other stakeholders to obtain the end-to-end performance improvements needed by UK scientists.

8.4 Scientific software systems

8.4.1 Well-engineered software is now critical for helping extract and managing the experimental data produced by large research facilities. It is also a vital component in helping to analyse that data. SCD's scientific Software Engineering Group has three main areas of activity:

- Software engineering and support (using expertise to assist computational scientists)[19];
- Data management tools and data pipelines for researchers;
- Research output recording systems.

8.4.2 SCD's Software Engineering Support Centre is funded by EPSRC to provide the tools, training, specialist expertise and software support that UK researchers require. This skillset is vital to

[18] In the US, the National Science Foundation has invested nearly $100m in implementing secure 'Research Data Transfer Zones' at over 80 universities.
[19] The Hartree Centre's Research Software Engineering Group provides a similar function for HPC.

get the best from UK research facilities, with engineers from SCD working closely with scientist colleagues who are responsible for the operation of those facilities.

8.4.3 Another example is work done by the STFC Research Data Group to develop a catalogue (ICAT[20]) to capture experimental meta-data. This enables researchers to find and combine datasets from various sources to get the best outcomes from the experiments carried out at the large facilities. This system was developed by STFC and is now in use both within the UK research facilities (ISIS and Diamond) and as well as further afield (at ESRF in France and the Oak Ridge Spallation Neutron Source in the USA). Keeping this core of software expertise enables researchers to extract the maximum benefit from experimental data. It is therefore a cost effective investment for the UK.

8.4.4 As the quantities of data become larger the need for software and data analysis expertise becomes ever more critical. Whilst SCD at the Harwell site provides support for the UK's large research facilities, and for many of the researchers that use them, the Hartree Centre is the main provider of specialist HPC software and HPDA expertise to industry. Again this is an area of strength for STFC.

8.4.5 The DiRAC Facility also has a dedicated software engineering team that provides support for a number of the research community's software systems. In addition, the DiRAC Facility runs a training programme in scientific software engineering.

9. Increasing the productivity of the UK's large research facilities

9.1 Large scientific facilities are becoming generators of ever-increasing amounts of data. For example, in 2017 alone the LHC is likely to generate some 50 petabytes[21] of data. Based on current projections, the SKA Observatory, once operational, may produce up to 300 petabytes per year.

9.2 Unsurprisingly, this type of exponential growth is also being seen at the UK's own large research facilities. For example, at Diamond more data was generated in the 2015/16 year than in its entire period of operation since 2007. This reflects the rapid development of detector technology and the increasing utilisation of Diamond by UK scientists. All of the STFC-supported UK research facilities at Harwell (Diamond, ISIS and CLF) are experiencing rapid rates of data growth and this situation is replicated at CCFE and elsewhere.

9.3 This revolution in experimental and observational science is not just about coping with large amounts of data. The increasing complexity of experiments, sometimes utilising more than one of the large facilities, often requires advanced computational modelling to interpret the results. There is also an increasing requirement for the facilities to provide near real-time feedback on the progress of an experiment as the data is being collected.

9.4 These trends require a close coupling between software, data and compute resources. The handling, analysis, visualisation, integration, modelling and interpretation of experimental data is now such that the degree of scientific computing expertise and resources needed can be extreme and is sometimes not easily available to researchers through their home institutions. Key issues here include the complexity of data (and ability to combine data sets) and variation in the researchers' knowledge of HPC, the potential of scientific software and advanced analysis techniques, as well as access to sufficient computing resources.

[20] ICAT provides an interface to large-facility experimental data.
[21] One petabyte is 1,000 terabytes; one terabyte is 1,000 gigabytes.

9.5 This evolving situation has led STFC and the CCFE – who share many of the same problems at their own research facility – to conclude that there is a need for a new approach that concentrates the scientific computing resources and expertise needed alongside the UK's large-scale research facilities. The Ada Lovelace Centre will be a way of generating a step change in the quality and quantity of the science delivered through the large facilities, accelerating the translation of experimental data into research outputs. It would also deliver a new generation of data scientists and software developers who will be working at the frontiers of their disciplines and can go on to careers either in science or in industry.

9.6 The Ada Lovelace Centre is therefore an exciting concept that seeks to tackle the productivity challenges facing the UK in the emerging field of data science by establishing a cross-disciplinary data science centre. It is consistent with the ambitions of the *Industrial Strategy Green Paper* which points out that investing in science, research and innovation is "not just a few people in labs making breakthroughs, but about adopting new and more productive ways of working". This philosophy applies equally to maximising the productivity of research facilities as it does to maximising the outputs of industry.

9.7 The Ada Lovelace Centre's mission is unique and it will work in tandem with other initiatives such as those underway at the Hartree Centre and the Edinburgh Parallel Computing Centre. It will also work to establish close links with the ATI. Further details of the Ada Lovelace Centre and the opportunities it offers are attached at Annex D. The diagram below shows the relationships between the RCUK NeI, UKT0 and the Ada Lovelace Centre[22].

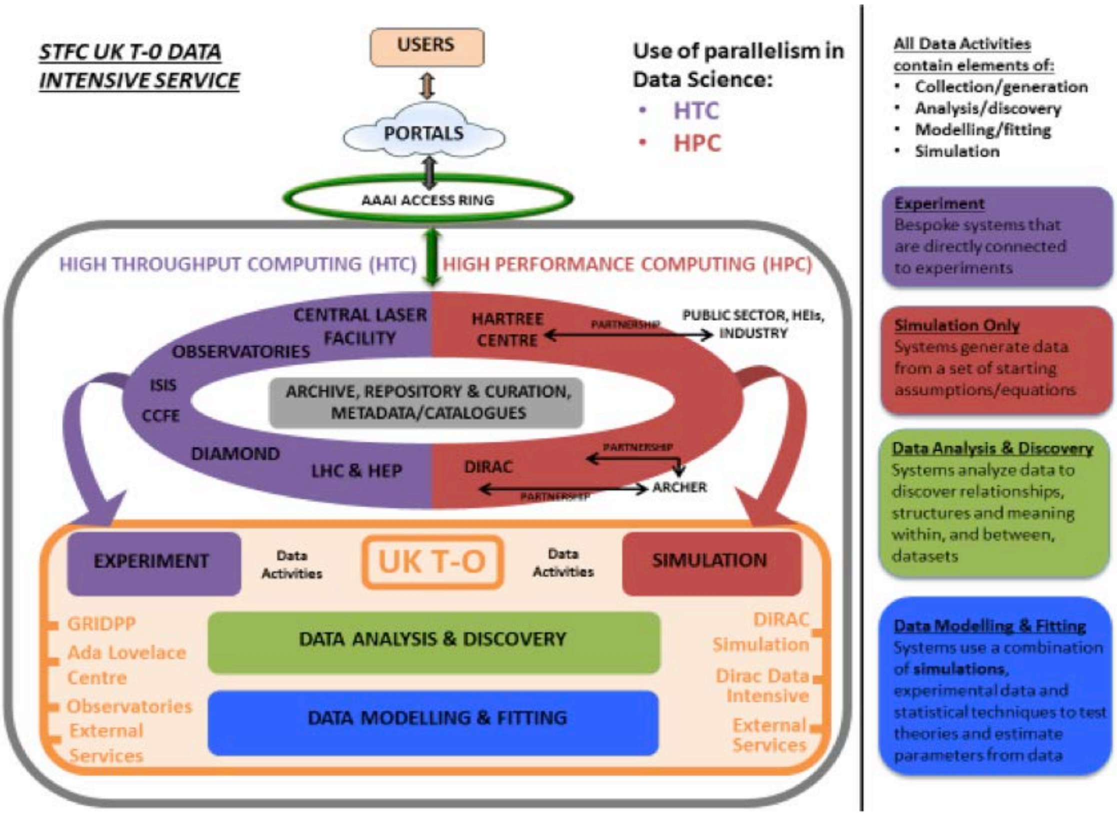

[22] Diagram courtesy of Jeremy Yates and Clare Jenner.

Action 8 – Improving productivity by establishing the Ada Lovelace Centre

The rapidly changing nature of modern research requires a step change in the way researchers access the scientific computing resources they need. STFC, and the CCFE, will seek investment to establish the Ada Lovelace Centre for computational modelling and data science alongside the UK large research facilities based at the Harwell Campus and at Culham. This will provide an important catalyst that has the potential to raise UK scientific productivity, improve the competitive edge of UK science and extract the most value from the multi-billion pound investment in research facilities.

10. Extracting competitive advantage from High Performance Computing

10.1 STFC supports industrial engagement with scientific computing in a number of ways. In some cases it is through the work carried out within partnerships such as DiRAC, where mechanisms such as joint appointments with industry are utilised. STFC also supports the Hartree Centre which seeks to transform UK industrial competitiveness by accelerating the adoption of HPC, HPDA and cognitive technologies. It does this through collaborative research projects that require the Centre's specialist scientific computing skills and technologies. In this way it can deliver transformative gains for UK industry.

10.2 A key strength of the Hartree Centre is the ability to bring together cross-disciplinary teams to bridge the gap between academic research and solutions that meet particular industry needs. It therefore helps address the gaps in HPC skills and knowledge which may hamper economic growth.

10.3 Recognising the need for support for Big Data and cognitive technologies for UK industry, in 2015 Government invested £115.5m over five years in the Hartree Centre. This has been augmented by a package of technology and onsite expertise of up to £200m from research collaborator IBM. Examples of some of the variety of the Hartree Centre's work include:

- Partnering in a four year £20m+ collaborative project known as ADDoPT (Advanced Digital Design of Pharmaceutical Therapeutics), which involves working with major pharmaceutical companies to apply advanced data analytics and modelling capabilities. This is helping optimise manufacturing processes and so strengthen an industry whose turnover is some £60bn;
- A collaboration with Alder Hey Children's NHS Foundation Trust is developing cognitive technology to improve the customer experience and free up valuable consultants' time;
- Working with a small start-up business to develop new software security systems which are then incorporated into their products.

10.4 The continuing development of the Hartree Centre as a centre of expertise for industry is vital. It is becoming clear that emerging needs are likely to be focused on how to solve data-centric challenges such as those in the *Industrial Strategy Green Paper*. Tackling these challenges is beyond the reach of many individual companies (particularly SME's), so it is important that the services on offer through the Hartree Centre are made widely known.

10.5 Looking ahead, to support its future work the Hartree Centre needs to continue developing the menu of services it offers to industry. Experience to date has already shown the value that developing a national centre of excellence can bring. However, it is likely that some continuing investment from UK Research and Innovation will be required to maximise its effectiveness. Further details can be found at Annex E.

Action 9 – Building competitive advantage

The Hartree Centre is now an increasingly important resource for UK industry and a centre of excellence in terms of how to apply HPC, cognitive computing and Big Data expertise to a wide variety of industrial challenges. Continuing to develop the Hartree Centre's collaborative approach with industry will boost the UK's competitive edge and help deliver economic growth and job creation. The success to date shows that the Hartree approach can deliver competitive advantage to UK industry, but investment needs to be maintained over the long term.

11 Helping ensure that UK science sustains its global impact

11.1 Collaborative action is a theme running throughout this strategy. This includes collaboration with UK-based institutions such as universities, collaboration with other Research Councils or collaboration with Government, industry and international organisations.

11.2 Scientific research is inherently international and STFC collaborates with partners in a wide range of countries to develop and support the large-scale science programmes and facilities required by the UK's science communities. Examples include CERN, SKA, ESO, the European Spallation Source etc. For these collaborations to be effective they need to utilise a common e-infrastructure which supports the global research community's ability to access, transfer and process data. It is also important that any decisions enable the UK to extract the maximum value from its international subscriptions to these multi-national research projects.

11.3 What this means for the UK is that our own e-infrastructure must be developed with international collaboration in mind and cannot be seen as independent from these international discussions. STFC is uniquely placed to provide the input needed to ensure that the UK's view is put forward and is taken into account in the development of any collaborative infrastructure which spans international borders. It therefore has a major role to play in providing input to discussions at European Strategy Forum on Research Infrastructures (ESFRI)[23] and participating in critical European e-infrastructure collaborations such as GÉANT[24], PRACE[25], EGi[26] and EUDAT that provide the fundamental networking, HPC and data services for European research.

[23] The ESFRI has a key role in policy-making on research infrastructures in Europe. In particular ESFRI contributes to the development of a strategic research infrastructure roadmap.
[24] The GÉANT network is the largest and most advanced research network in the world, connecting over 50 million users at 10,000 institutions across Europe.
[25] PRACE offers leadership scale computing and data management resources and services through a peer review process.
[26] EGi e-infrastructure is publicly-funded and comprises over 300 data centres and cloud providers spread across Europe and worldwide.

11.4 STFC also has a role in representing the UK in discussions further afield with, for example, research laboratories in USA and elsewhere that have a similar mission. It is also involved in forming partnerships with institutes elsewhere in the world, including via programmes such as the Global Challenges Research Fund and Newton Fund.

> ***Action 10 – A leadership role***
>
> *The uniquely global nature of STFC's science and research requires us to plan, establish and operate infrastructures not only across organisational boundaries, but across international boundaries. STFC is in a key position to continue to provide leadership to those international collaborations relating to e-infrastructure and to link this to the strategic discussions taking place in the UK as it develops its own e-infrastructure strategy. It will therefore build on its existing position to play an active role on the international stage as an advocate for the UK's e-infrastructure interests.*

12. Resources

12.1 The direction suggested by this strategy requires new investment. It is recognised that this is against a background of resource constraints, but holding public investment at 'flat cash' in such a fast changing area will lead inevitably to the rapid erosion of the UK's competitive edge as a leading scientific nation. The warning signs are already evident that this is taking place – the 2013 Elsevier report referred to earlier identified that the UK research base shows areas of potential vulnerability and states that maintaining high-quality research outputs whilst R&D expenditure is broadly stable or decreasing cannot be achieved indefinitely.

12.2 A 'flat cash' environment, imposed over many years[27], also erodes the UK's ability to maintain leadership in instrumentation, facilities and exploitation and has an adverse impact on the UK's scientific output. This adverse impact is amplified if UK researchers are denied the e-infrastructure tools they require. Of course some minor incremental improvement through consolidation may be possible if extra resources are not provided, or if investment is held at 'flat cash' for a further prolonged period. But failure to invest adequately in the long term ignores the essential enabling role of e-infrastructure and will have a serious negative impact on the UK's long-term future as a research-driven economy.

12.3 Annex F provides a broad estimation of the extra investment which would be necessary to make this strategy a reality and enable UK scientists to achieve more, and more quickly.

[27] The *2013 STFC Programmatic Review* noted that 'flat cash' would result in a cumulative 37% reduction in volume if extended over four years.

13. Conclusions

13.1 The rapidly changing environment described within this strategy demands the development of a coherent UK e-infrastructure strategy and investment plan. STFC is ready to play a full role but believes it needs to be based upon a step change in investment if UK research is to remain internationally competitive.

13.2 The actions outlined in this strategy, taken together, will provide a sound basis for moving forward. We have identified high level solutions to the challenges facing STFC and UK research, but recognise that co-ordination and integration will be required. We believe that the actions are suitably ambitious and will not only increase research productivity but also ensure that UK researchers (both academic and industrial) have access to the tools and expertise they need. They will also make a significant contribution to building the UK's skills in the critical areas of HPC and data science.

13.3 This document has been circulated widely within STFC and has been discussed at the Computing Advisory Panel and Science Board. We would welcome any further comments on the proposals outlined above.

Annex A – The DIRAC Facility

Overview

- The DiRAC Facility – Distributed Research utilising Advanced Computing – has provided integrated computational resources for the STFC-funded theoretical astronomy, particle-astrophysics, nuclear physics and particle physics communities.
- Since 2009, the DiRAC Facility has provided two generations of a dedicated Supercomputing and HPC service. It has also been a pioneer of scientific software and computing hardware co-design with industry, and has worked closely with technology companies such as IBM and Intel.
- The science case for the full remit of STFC PPAN science has demonstrated the need for a 10-fold increase in computing resources available to the DiRAC research community to deliver the new insights required to advance our knowledge in theoretical astronomy, particle-astrophysics, nuclear physics and particle physics.

Description of the proposed DiRAC services and the communities supported

- The DiRAC Facility has identified a combination of computing services needed to achieve its desired science outcomes:
    - <u>HPDA Systems:</u> Novel inter-operability and petabyte-scale low latency memory/disk systems are needed in data-intensive scientific areas to confront observations with simulations of increasingly complex physical systems, and thereby drive scientific discovery. The DiRAC Facility intends to work closely with the JASMIN teams in SCD and RAL Space to provide a common architecture and environment.
    - <u>Supercomputing capability:</u> The DiRAC community will continue its development programme of cutting-edge scientific codes that can fully exploit extreme-scaling multi-petaflop systems to deliver world-leading science cost-effectively. In addition, physically realistic large-scale Computational Fluid Dynamic problems require access to HPC systems with very large memory footprints (>300TB) so that codes can 'follow' structures as they form to understand the interplay of different processes.
    - <u>Data Curation Service:</u> This is required to store the 100 petabytes of scientifically valuable DiRAC simulation data outputs, tagged with metadata. This allows the datasets to be re-used by other researchers and further increases the impact of DiRAC simulations.

- The UKT0 facility will provide the DiRAC Facility users with the data modelling capability to interpret measurements from Advanced LIGO, ALICE, ALMA, ATLAS, CMS, DES, DESI, eBOSS, LHCb, and Planck experiments.

Annex B – The potential of cloud computing

The RCUK Cloud Working Group has recently investigated the current commercial offerings and has made the following observations:

Public cloud, on premise hosting and storage costs

- Computing for the research community is underpinned by the ability to host large amounts of persistent storage collocated with resources for processing and analysis together with a high level of utilisation.
- Current pricing models for persistent storage on public cloud are consistently higher than the total cost of ownership equivalent for on-premise hosting. A recent comparison with Amazon storage options – Glacier, S3 standard, Elastic Block Magnetic and Elastic Block SSD – shows that these options range from 2 to 14 times more costly than an on-premise solution.
- These figures also do not take into account the additional costs that would be incurred for input/output (i/o) operations which are part of the overall tariff that providers charge for storage. With the high-level of utilisation of data and therefore i/o operations by the research community, the costs become even less competitive.

Hybrid cloud model

- Given the above issues with the long-term hosting of scientific data in public cloud, a hybrid public cloud/on-premise strategy provides an alternative possibility. In this scenario, the research infrastructure provider retains the hosting of its data but opportunistically bursts compute to third party providers when demand requires more compute performance.
- This takes advantage of the fact that prices for cloud compute are more competitive than cloud storage. It is predicated on the availability of good WAN (Wide Area Network) connectivity between the public cloud and on premise systems.
- Using the AWS Spot market mechanism, for example, to bid for spare capacity in the cloud, one can start to approach the equivalent price of an on-premise compute solution.
- However, larger instance types requiring more memory or cores are much more expensive in the commercial cloud and for many workloads cannot deliver suitable performance at any cost. These higher specifications are important for many research workloads.
- Finally, to take advantage of this burst model, data must be moved in and out of the public cloud. For large datasets used in the research community this is likely to become very expensive because of public cloud provider’s egress charges.
- Recently both AWS and Azure have provided a waiver for academic use. Taking the AWS waiver, for example, this is only favourable if the cloud tenant spends a large amount of money on other AWS services, because the egress waiver is limited to 15% of the total spend.

The RCUK Cloud Working Group concludes:

- In summary, infrastructure providers in the research community are monitoring pricing of public providers as part of an ongoing process of reappraisal to look for cost effective means to use public cloud.
- Current commercial pricing models mean that storage of data is much more expensive than on premise taking into account total cost of ownership.
- Some scenarios for bursting to take advantage of compute on public cloud are becoming more competitive but egress charges are likely to remain a barrier to cost effectiveness for the research community.

Annex C – The UKT0 consortium

Overview

- The UKT0 consortium is science-led and aims to co-locate, share and aggregate HTC computing resources and expertise, where appropriate, to allow different projects with similar high throughput computing needs to access the necessary computing resources required by the projects as and when needed. Such consolidation of computing resources will be more cost-effective than providing each project with its own, siloed resources and will also enable faster time to science.
- The UKT0 HTC facility is unfunded at present although it has been supporting some minor components of the STFC science-funded programme on the GridPP infrastructure to demonstrate the principle. When funded it will build on existing infrastructure on the STFC National Research and Innovation campuses at Harwell and Daresbury, at the EPSRC UK Research Data Facility (RDF) in Edinburgh, and several other major centres in UK universities.
- A major investment is required to provide the data processing, analysis, and event simulation requirements of the new generation of PPAN experiments.

Description of the UKT0 HTC collaboration and the communities supported

The UKT0 HTC facility will provide efficient computing support for a consortium that includes present and planned research projects across the entire STFC science portfolio and that of CCFE. The activities currently involved are:

- The LHC Experiments (ATLAS, CMS, LHCb and ALICE)
- Other current and future High Energy Physics Experiments (ILC, NA62, MICE, T2K, SNO+, HyperK, DUNE)
- The Square Kilometer Array (SKA)
- The Large Synoptic Survey Telescope (LSST)
- Dark Universe space mission (EUCLID)
- Low Frequency Radio Astronomy (LOFAR)
- Advanced-LIGO
- Cerenkov Telescope Array (CTA)
- Dark Matter (Lux-Zeplin)
- Diamond Light Source (Diamond)
- ISIS Neutron and Muon Source
- Central Laser Facility (CLF)
- RAL Space
- Culham Centre for Fusion Energy (CCFE)

- Other future PPAN projects may join the UKT0 collaboration as they are approved by STFC's peer review system. The computing requirements of the international SKA project will be very demanding and require careful scrutiny. They may require dedicated HTC analysis of streaming SKA data which has operational implications for the UKT0 facility.

- The Ada Lovelace Centre will rely on the UKT0 HTC system for its computing resources for the facilities and their users. Without such computing resources, the user communities of the facilities will struggle to analyse the data being produced on a timescale comparable to that of our international collaborators and the UK will rapidly lose its competitive edge and its leadership role.
- The UKT0 HTC facility will also support the computational plasma and materials simulation needs of the CCFE. The Culham Centre is near to the Harwell campus and CCFE requires access to significant computing for its research and development programme based around the JET and MAST Tokamak confinement systems in the UK, and for the international ITER project in France.
- The UKT0 facility will be managed by an appropriate governance structure and also provide the HTC capability and data archiving resources required by the DiRAC consortium.

The diagram below shows the relationships between the RCUK NeI, UKT0 and the Ada Lovelace Centre[28].

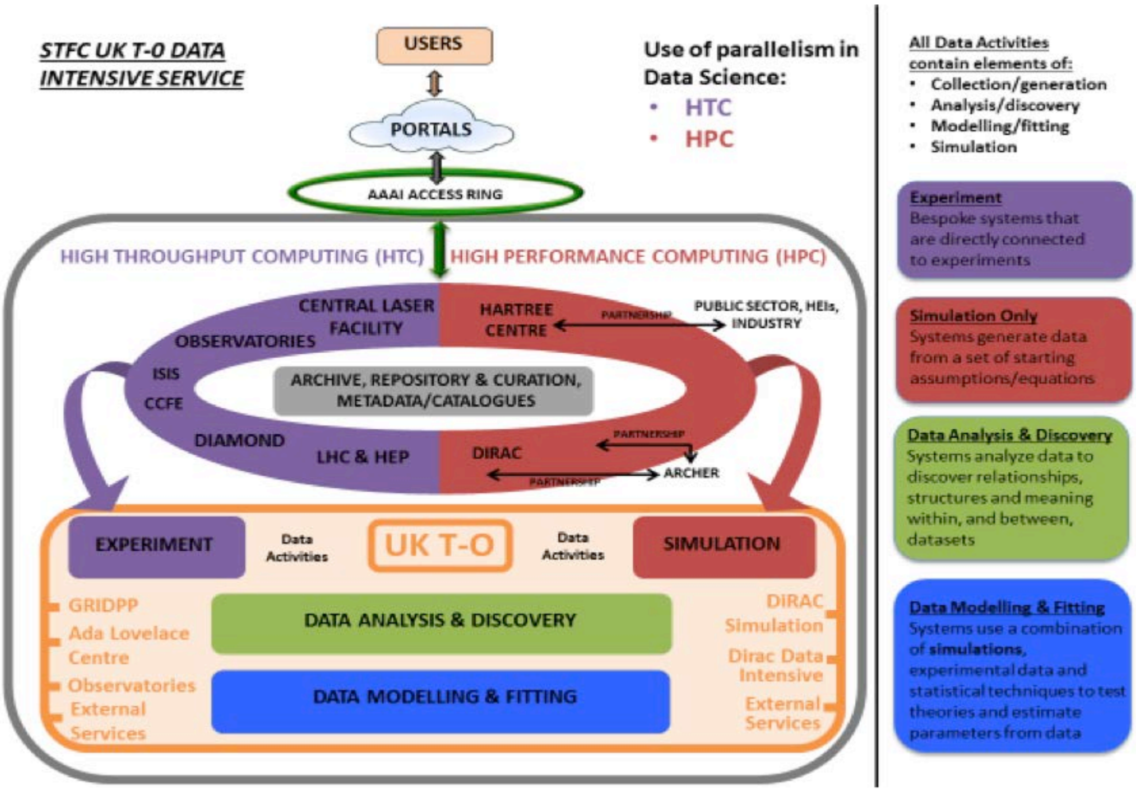

[28] Diagram courtesy of Jeremy Yates and Clare Jenner.

Annex D – The Ada Lovelace Centre

Overview

- Significant investment is required to establish the Ada Lovelace Centre, an integrated, cross-disciplinary data intensive science centre, for better exploitation of research carried out at our large-scale national facilities including Diamond, ISIS Neutron and Muon Source, CLF and CCFE.
- The Ada Lovelace Centre has the potential to transform research at the facilities through a multidisciplinary approach to data processing, computer simulation and data analytics. The impact will be felt across the many science disciplines and communities these facilities support, including industry and academia.
- In addition, skills in the UK will be furthered through the engagement with some of the most advanced scientific data and compute challenges.

Description of the proposed centre and the communities supported

- The UK's large-scale national facilities – Diamond, ISIS Neutron and Muon Source, CLF and CCFE – enable and deliver world-leading science across the physical and life sciences. Their science impact is wide-ranging, with examples ranging from the environment, pharmaceuticals, healthcare, and nanotechnology, to advanced materials, energy, robotics and engineering.
- Experimental science is undergoing a revolution driven by next generation instruments, the dramatic advances in detector technology and the increasing complexity of experiments. These factors result in ever greater volumes of data that require advanced computational modelling and data analytics to interpret and find signals in the data. There is also an increasing requirement for the facilities to provide near real-time feedback on the progress of an experiment as the data is being collected.
- Further complexity comes from research that combines data from different experiments. In addition, there is a need for intelligent text mining of research papers and exploitation and re-use of data sets via databases. All of these trends require a closer coupling between software, data and compute resources.
- The Ada Lovelace Centre will integrate the unique capability of the UK's large-scale experimental facilities with the decades-long track record of STFC's SCD in supporting scientists with state-of-the-art software, large-scale computing services and data infrastructures. The Ada Lovelace Centre will provide computing hardware, build software, and provide computational and data analytics expertise that will spark a paradigm shift in the capability of scientists to design, analyse and interpret experiments.

- The Ada Lovelace Centre investment would be used to:

    - Build capacity in advanced software infrastructure for the handling, analysis, visualisation, integration, modelling and interpretation of experimental data;
    - Build and curate scientific databases of analysed and published data to enable long-term exploitation and re-use of the data;

    - Develop the skills in data science and scientific software of new generations of scientists educated at the frontier of their disciplines in the fast-expanding field of data intensive science;
    - Supply computational hardware to enable the data intensive analysis, simulation and modelling to be performed. The hardware will be provided by UKT0 HTC facility (see below).

- In the first five years of the Ada Lovelace Centre, a common software foundation will be laid across the many experimental research techniques available at these facilities. This will act as a catalyst for the creation of global research communities in academia and industry.
- While the mission of the Ada Lovelace Centre is unique, this centre will be an essential component of a broad-based ecosystem of UK initiatives, and will seek partnerships with the Hartree Centre and the ATI.

Annex E – The Hartree Centre

Overview

- The Hartree Centre has a mission to transform UK industrial competitiveness by accelerating the adoption of HPC, HPDA and cognitive technologies. This will be achieved through collaborative research projects that leverage relevant knowledge, research, skills and technologies from academia and industry to deliver transformative gains for UK industry and Government.
- Key strengths of the Hartree Centre are access to competitive e-infrastructure, HPC and HPDA resources and the ability to bring together cross-disciplinary teams to bridge the gap between academic research and solutions which meet particular industry needs.
- Early Hartree Centre projects addressed problems in chemistry, engineering and the life sciences, but emerging industry needs are increasingly focused on more complex, data-centric challenges such as those identified in the Government's new Industrial Strategy. These include applications in advanced manufacturing (Industry 4.0) and smart cities.
- More than four out of five technology start-ups never make it to the commercial world, due to their inability to cross the 'Valley of Death' – the virtual chasm that separates research technology demonstrations from credible industrial prototypes. The UK Government's Industrial Strategy will provide further opportunities for the Hartree Centre to demonstrate the effectiveness of its novel technology adoption activities in increasing UK productivity and bridging the gap between research and industry.

Description of Hartree Centre plans for increased industry collaboration

- The Hartree Centre currently has a partnership with IBM Research to demonstrate the benefits of cognitive computing technologies – such as Watson – through a series of proof-of-concept and proof-of-technology projects that tackle a wide range of industry challenges.
- The Hartree Centre brings together teams of Industrial end-users with STFC domain scientists working with both IBM and STFC staff with transversal skills capable of producing digital assets which have the potential to be exploited jointly for commercial and economic gain.
- The added-value activities in the Hartree's new Snowdonia procurement – an ATOS Bull system that will deliver around 3.4 Petaflop/s of computational power – will provide further opportunities for STFC, vendors and industry partners to develop partnerships where infrastructure and skills are brought together for industry's benefit.

- To support its future work the Hartree Centre needs to provide:

    - Access to state-of-the-art heterogeneous HPC systems which are flexible enough to cope with changing workloads and large enough to allow industry to work at a scale that was previously impossible;
    - Secure data areas where customer data can be kept in a confidential way and accessible to research data networks which enable fast and secure data transfer;

- Authorisation, authentication and access models appropriate for non-expert users (e.g. via tablets and smart phones);
- Early access to a range of emerging technologies and tools to evaluate their suitability for academia, industry and Government;
- Partnerships with vendor organisations to transfer expertise to STFC staff to enable them to support industry uptake and to develop exploitation plans to exploit the digital assets created for mutual gain and industry benefit;
- Training and staff development programmes to develop the skills which RCUK/UK Research and Innovation and UK industry will need to successfully profit from large-scale data analytics and cognitive computing;
- Migration paths to in-house, regional and national HPC systems, and to commercial cloud systems, on the completion of proof of technology and concept projects;
- Partnerships with academic and industrial research organisations, and with the E-infrastructure Leadership Council (ELC), to share expertise on novel systems and emerging technologies so that future industry requirements can be more effectively fed into a NeI roadmap.

A unique role for the Hartree Centre in international collaborations

- The UK is unique in having a national HPC and HPDA centre dedicated to working with industry. The Hartree Centre and its focus on facilitating return on research innovation for the benefit of industrial competitiveness is now a topic of increasing global concern. This has opened up opportunities for collaboration with some of the US Department of Energy laboratories, particularly those involved in the HPC CORAL programme (Laurence Livermore National Laboratory, Oak Ridge National Laboratory and Argonne National Laboratory).

- The Hartree Centre’s additional focus on emerging technologies and energy efficient computing and the strong links with vendors has also positioned STFC better for engaging in the European Commission’s forthcoming Extreme Scale Demonstrators programme, designed to produce prototype pre-exascale and exascale systems in Europe by 2022.
- HPC, data analytics and machine learning are becoming increasingly intertwined and inter-related – the Hartree Centre’s involvement with all three of these technologies equips the Centre well to play a bigger role in Europe and internationally, particularly in articulating the challenges and opportunities for industry to maximise the benefits of early adoption of machine learning.

Annex F - Resources

1. The following tables show the additional resource requirement for the next five years in the following key areas:

   - DiRAC
   - UKT0
   - The Ada Lovelace Centre
   - The Hartree Centre
   - Training and skill development
   - Scientific Data Networking improvements

2. The intention is to provide a broad picture of the level of investment necessary, including both capital and resource, whilst recognising that each component will still be required to undergo the necessary project approval and oversight processes.

3. It should be noted that the electricity costs for running computing centres are increasing and is likely to increase further in the future. For example, the current electricity bill for GridPP Tier 1 is around £1m per annum and rising.

Table 1: DiRAC

DiRAC indicative costs £m[29]	**2018/19**	**2019/20**	**2020/21**	**2021/22**	**2022/23**	
Compute	12.00	8.00	12.00	10.00	11.00	
Disk storage	3.50	3.00	4.00	4.00	4.00	
Tape storage	0.20	0.20	0.20	0.20	0.20	
Networking	0.25	0.00	0.25	0.00	0.00	
Operations staff	0.60	0.60	0.60	0.60	0.60	
Software engineering staff	1.40	0.50	0.50	0.50	0.50	
Licencing, maintenance	0.10	0.10	0.10	0.10	0.10	
Power and cooling	2.20	2.20	2.20	2.20	2.20	
Capital works in data centres req'd to house equipment	2.50	0	0	0	0	
Total	**22.75**	**14.60**	**19.85**	**17.60**	**18.60**	**93.40**

[29] Table includes contribution to the DiRAC Data Curation Service.

Table 2: UKT0

UKT0 indicative costs £m	**2018/19**	**2019/20**	**2020/21**	**2021/22**	**2022/23**	
Unfunded CPU	2.30	3.30	3.40	3.70	5.00	
Unfunded Disk	.51	.82	.86	1.06	1.50	
Unfunded Tape	.16	.29	.35	.38	.46	
Other infrastructure	.40	.60	.70	.80	.90	
Operations and middleware staff	.70	.70	.70	.70	.70	
Software engineering	0	0	0	0	0	
Licencing, maintenance	.10	.10	.10	.10	.10	
Power and cooling	1.00	1.50	1.50	2.00	2.00	
Total	**5.17**	**7.31**	**7.61**	**8.74**	**10.66**	**39.49**

Table 3: The Ada Lovelace Centre

Ada Lovelace Centre indicative costs £m	**2018/19**	**2019/20**	**2020/21**	**2021/22**	**2022/23**	
Unfunded CPU	2.70	2.70	2.70	2.70	2.70	
Unfunded Disk	0.50	0.70	0.70	0.90	0.30	
Unfunded Tape	0.14	0.10	0.10	0.20	0.10	
Other infrastructure	0.40	0.60	0.70	0.80	0.90	
Operations and middleware staff	0.80	0.90	1.20	1.20	1.25	
Software engineering	0.85	2.40	2.75	3.15	3.40	
Licencing, maintenance	0.50	0.50	0.50	0.50	0.50	
Power and cooling	0.30	1.00	1.20	1.50	1.50	
Total	**6.19**	**8.90**	**9.85**	**10.95**	**10.65**	**46.54**

Table 4: The Hartree Centre

Hartree Centre indicative costs £m	**2018/19**	**2019/20**	**2020/21**	**2021/22**	**2022/23**	
Compute	0	3.00	0	15.00	3.00	
Disk storage	0	1.00	0	3.00	1.00	
Tape storage	0	0.50	0	0.50	0	
Networking	0.50	0.50	0.50	0.50	0.50	
Operations staff	1.20	1.20	1.30	1.30	1.40	
Software engineering staff	3.00	3.20	3.50	3.70	4.00	
Licencing, maintenance	2.50	2.50	2.50	2.50	2.50	
Power and cooling	0.90	0.90	0.90	0.90	0.90	
Capital works in data centres req'd to house equipment	0	0	15.00	0	0	
Total	**8.10**	**12.80**	**23.70**	**27.40**	**13.30**	**85.30**

Table 5: Training and skill Development (allied to Ada Lovelace Centre and Hartree Centre)

Training and skill development indicative costs £m	2018/19	2019/20	2020/21	2021/22	2022/23	
Ada Lovelace Centre	0.50	0.70	0.70	0.75	0.75	
Hartree Centre	0.30	0.35	0.40	0.45	0.55	
Total	**0.80**	**1.05**	**1.10**	**1.20**	**1.30**	**5.45**

Table 6: Scientific Data Networking improvements

Networking indicative costs £m	2018/19	2019/20	2020/21	2021/22	2022/23	
To provide a scientific data network infrastructure at STFC sites and STFC funded groups at Universities	7.00	3.00	2.00	2.00	0	
Total	**7.00**	**3.00**	**2.00**	**2.00**	**0**	**14.00**