The Importance of Regional e-Infrastructure

Within the National Landscape

A submission compiled by the Regional e-Infrastructure Centres Network (ReICN)

ReICN is a network comprising the five Regional Centres that were established by EPSRC in 2012, and that have been operating ever since. The Centres, and their member universities, are:

ARCHIE-WeSt	Glasgow Caledonian University, University of Glasgow, University of Stirling, University of Strathclyde, University of the West of Scotland
HPC Midlands	University of Leicester, Loughborough University
MidPlus	Queen Mary College London, University of Birmingham, University of Nottingham, University of Warwick
N8 HPC	Durham University, Lancaster University, University of Liverpool, University of Leeds, University of Manchester, Newcastle University, University of Sheffield, University of York
SES	Imperial College London, University of Cambridge, University College London, University of Oxford, University of Southampton*

All five Centres have been running Regional e-Infrastructure Centres for at least four years, and in several cases were running multi-institutional services before that.

This report Summarises our Collective experience of the advantages and disadvantages of regional provision within the Nation's e-Infrastructure landscape, and suggests several mechanisms by which the nation would benefit from including a vibrant regional tier within the plans for the post-ARCHER landscape. The report also includes a substantial evidence base to support the recommendations we present herein.

The ReICN Executive consists of: P.M. Rodger (Chair, Warwick), S.J. Cox (Southampton), M.T. Dove (QMUL) D.C. Hogg (Leeds), S.D. Kenny (Loughborough), R. Martin (Strathclyde), P.A. Mulheran (Strathclyde), A. Richards (Oxford) and C. Taylor (Manchester)

Additional input to this submission was provided by: T. Metcalf (UCL), G. Sinclair (Manchester), A.N. Real (Leeds), W. Groenewald (Warwick) and A. Fakhrkonandeh (Warwick)

* Also includes King's College London since 1st January 2016

Executive Summary

Regional e-Infrastructure Centres are an essential component of an integrated e-Infrastructure landscape, and are essential for maximising the benefits obtained from both national and local facilities. They contribute directly to the development of high value digital skills, research leadership, economic growth and productivity in the UK.

Continued benefit from Regional Centres within a national ecosystem requires a cohesive long-term plan for UK investment in e-infrastructure and associated skills. Such a plan is essential if the digital capabilities are to be available for the demands we currently anticipate, while also delivering the flexibility to exploit new opportunities. The on-going culture shift toward collaboration and sharing, embodied in Regional Centres, is also integral to expanding the use of digital technologies within research and industry, thereby stimulating innovation and underpinning competitive advantage.

Key Selling Points for Regional Centres within an integrated UK ecosystem

1. Increased capability and capacity: readily accessible to research groups, enabling them to achieve more rapid and effective research impact.
2. More effective two-way communication of priorities and initiatives: creating a 3-level network (National ↔ Regional ↔ Universities) to focus and filter discussions
3. Successful pathways for broadening access to HPC
4. Increased diversity of e-infrastructure, skills and networks: readily accessible to individual researchers, and enabling and de-risking rapid experimentation with new ideas and architectures.
5. Effective vehicles for regional investment and resource sharing: leveraging University investment, and enabling regional research strengths to become gateways for SMEs into high-end computing.
6. Personal networking: face-to-face discussions within 2 hours travel (‘*Regionality*’) is a powerful enabler for research collaboration, technical support, training and knowledge exchange.
7. Effective platforms to link with other key investments, such as the Henry Royce Institute, Farr Institute, Alan Turing Institute or the National Automotive Innovation Centre.

Recommendations

1. That a strong and vibrant Regional (Tier 2) layer be embedded within the next realisation of the UK e-Infrastructure ecosystem.
2. That recognition of the ‘on-the-ground’ benefits of the regional centres that is backed up with a commitment to a sustainable, long-term strategy that includes capital and operational funding.
3. That funding for the Regional layer be delivered through a research-driven grant scheme that is strongly informed by Research Council experience with both CDT and Programme Grant initiatives so as to ensure:
    i. Appropriate diversity of provision across the UK;
    ii. Comprehensive coverage of researchers within the UK;
    iii. Flexibility to adapt spending plans in the light of both regional research initiatives and technological developments;
    iv. Continuing quality of provision through mid-term review.
4. That consideration be given to an immediate short-term tranche of refresh funding to ensure the expertise built up within the Regional Centres since 2012 is not completely lost before the post-ARCHER plans can be implemented in 2018-19.

Contents

Introduction

Until the e-infrastructure initiative that created the five EPSRC regional Tier-2 centres in 2012, high-performance capability for scientific research in England and Scotland largely consisted of a single central Tier-1 HPC resource and whatever local resources an institute was prepared to establish. A few discipline-based collaborative resources were available (Dirac, UKMHD), but were not available to the vast majority of UK researchers.

The creation of the regional Tier-2 Centres in 2012 demonstrably created new research opportunities, and produced a major advance in any strategy to develop a new generation of scientists who are literate in simulations and high-performance computing.

At the level of meeting demand, the provision of Tier-2 regional centres has been extraordinarily successful. The Tier-2 centres have provided different types of computing that are not being met by national facilities, including high-throughput computing, high-memory installations, smaller-scale facilities and GPU computing. The added capacity from this initiative was quickly met by demand, and subsequently by high quality research outputs.

In many cases this demand was not appropriate for national facilities, since national facilities need to be optimised for users with very large demands that challenge the boundaries of the technology, and so it is strategically unwise for the national facilities to be diluted to meet needs that are best served by regional Tier-2 centres.

This experiment in Regional Computing was not universal. Only about half the research-active Universities in England and Scotland were supported. However, their experience provides a wealth of information for designing the next generation e-Infrastructure landscape to optimise the UK's future in a highly competitive global environment.

Within this report we identify a number of crucial and unique features that Regional Centres add to the National e-Infrastructure landscape; we identify features that must be included in—and others that must be excluded from—design of the next generation landscape; and we suggest funding models that could be used to implement this. We also provide an evidence base to support our proposals.

Advantages of a strong regional layer

Regional Centres should play an essential role in any strategy to develop wider high-end computing literacy within the scientific community. The balance they strike between *regionality* and critical mass enables them to do far more than simply bridge between local and national facilities: it enables them to be a powerful tool for shaping high-end e-Infrastructure access, training and support within a national strategy.

Regionality brings with it a sense of local ownership that cannot be captured by national provision. This extends to access, training and support. Whilst ARCHER has excellence in these areas, it is essentially only available to the large scale users and thus is not appropriate for developing large numbers of new users who do not (yet) satisfy the requirements for—or, indeed, do not yet recognise the potential of—using a Tier-1 centre.

Critical mass brings with it a level of expertise and innovation that cannot be matched by most Universities. This can lead to much more powerful training and support than can be provided locally, without losing the ability to adapt to individual needs—as often must happen when working on a national level.

Based on these foundations, and derived from our experience of the last four years, we have identified a number of highly desirable features of the national e-Infrastructure ecosystem that would be delivered most effectively by a strong and vibrant regional tier.

1 Ready access to increased capability

Regional centres extend researcher horizons significantly, while the scale or resources still enables them to employ only light-touch review processes to ensure effective usage. This makes the centres very effective at enabling researchers to explore new ideas and the possibilities for scale-up with minimum inertia, and so enables both greater impact for the research and optimal efficiency on national facilities when they are subsequently employed.

The importance of bureaucracy-free access in enabling major new research avenues was underlined in a number of interviews (page 12), while surveys by the Regional Centres suggest that up to 10% of researchers have used the regional facilities for proof-of-concept research that has subsequently been migrated on to ARCHER.

2 Communication of priorities and initiatives

The one-to-very-many mapping of researchers onto national facilities leads to very inefficient communication: either communications are flooded by too much attention to individuals, or they are starved by the need to be relevant to everyone. The presence of a universal tiered structure, enhancing two-way communication between each tier, can enable highly effective sifting, sorting and propagation of information. Bidirectional networks with about three layers (local, national and regional in this case) are generally found to provide excellent targeting of information flow, with little degradation of information from filtering between levels.

3 Broadening access to HPC

The Regional Centres have been very successful in allowing many researchers to broaden their computational horizons. In some cases this has involved making HPC available to institutions that, thus far, have not had adequate access. Examples include the involvement of Glasgow Caledonian University, University of Stirling, and the University of the West of Scotland in ARCHIE-WeSt, and the use of MidPlus by a research group at Aston. In other cases, the wider access has involved helping non-traditional disciplines discover the power of high-end computing. In particular, easy access to Regional Centres has empowered research groups in the Social Sciences and Economics to discover just how much they can now gain through the analysis of very large social and media databases using high-end computing facilities.

4 *Diversity of e-infrastructure, skills and networks*

The Tier-2 centres have provided different types of computing that were not being met by national facilities, including high-throughput computing, high-memory installations, remote hardware accelerated visualization, smaller-scale facilities and GPU computing. Most Universities do not have the resources to support several significant installations and so are forced into a one-size-fits-all strategy. National facilities must usually be justified by a proven high level of need within the national community, and so are not well placed to accommodate new technologies or research communities with novel requirements. Regional Tiers provide an optimal balance, pooling resources to enable several architectures to be supported, exhibiting a natural diversity through their accountability to different regional research strengths, and facilitating low-risk experimentation with new systems.

An excellent example is Emerald (SES), which was the largest GPU machine in Europe when first deployed (374 GPUs). At the time, GPU technology was considered novel and to have a narrow application and skills base. Most university GPU machines were small (<10 GPUs) and operated in silos within departments. GPUs are now accepted technology. The second most powerful computer[*] and the ten most energy efficient computers[†] in the world are based on GPUs. The SES investment at an early stage has enabled the UK to be at the forefront of these developments, and has allowed a great deal of exciting research to be conducted ahead of

[*] http://www.top500.org/

[†] http://www.green500.org/

international competitors. It is also important to stress that access to SES GPU infrastructure and expertise has always been open to all UK researchers, and there are long-time users from the universities of Bath and Manchester. Thus by responding to regional clusters of expertise, the Regional Centres can enable the whole UK to benefit from early uptake and experimentation with new technology.

5 Effective and efficient regional investment and resource sharing

Regional centres provide a showcase through which to develop and project the research strengths of a region, and should therefore facilitate collaboration and knowledge exchange between Universities, Industry and Commerce within the region. This is likely to be particularly important for SME engagement, which typically does not have the level of expertise and resources that enables large industries to exploit national centres.

Opportunities also abound for University-Government co-investment. The funding initiative that created the ReICN centres has already demonstrated just how powerful a tool Regional Centres can be for regional investment. The initial investment from Government was more than doubled by additional investment from the member Universities—in the first year! Very significantly, the matched University investment went well beyond equipment sharing to co-ownership. At present, this has been achieved as a one-off experiment. If some level of continuity of Government funding for regional centres can be provided, then it will help embed a culture of strategic co-investment by University groupings in large research infrastructure.

6 Personal networking: diversity with critical mass through "regionality"

Face-to-face interaction is an essential element of networking, just as time spent in discussion is an essential component of developing new collaborations. The barrier created by requiring more than 2 hours travel cannot be overestimated. Regional Centres are ideally placed to exploit this, being close enough to develop powerful collaborations and promote knowledge-exchange between experts who would otherwise be isolated within their University, and in the process creating new centres of excellence in aspects of computational science and modelling that others in the region can engage with. This point was emphasised by users during the interviews (page 12), with comments such as "*…it not only provides computers, but also the community. Other facilities only offer CPU time*" (Eberhard, Aston)

7 Effective platforms to link with other key investments

Major new research initiatives—such as the Alan Turing Centre, the Henry Royce Institute, the Farr institute or the National Automotive Innovation Centre—raise demands for cutting-edge e-infrastructure that are too large for individual Universities to meet, but too specialised for National Centres to address without a major funding and procurement exercise. However, such research initiatives usually create regional priorities that leave Regional e-Infrastructure Centres well placed to meet, having sufficient resources but small enough size to be able to adapt to meet at least the initial demands such initiatives create, and therefore provide space for more considered responses to be crafted as the initiative develops.

The research enabled by a strong regional layer

The scientific impact of the Regional Centres has already been significant, with *each* facility featuring in an average of more than 100 grant applications totalling nearly £50M, and its use acknowledged in at least 280 publications. Projects using the facility cover a wide range of applications in chemistry, mathematics, physics, engineering, computing, life sciences, earth sciences and economics. A set of case studies illustrating the range of applications is attached.

The expected impact of the Regional Centres over the next 5–10 years was explored in a series of Town Meetings held at each of the Centres at the start of 2016. Some of the main priority areas identified during these meetings are set out below, organised using the EPSRC Outcomes Framework. In all cases, the researchers felt strongly that much of the work would be ideally suited to Regional Centres, exploiting the scale and diversity of hardware, the regionality of support, and the ability to build a community of experts around research initiatives and meet with them frequently.

Productive Nation

Materials: New and improved materials are fundamental to achieving breakthroughs in many aspects of product design and manufacturing. Materials innovation is critical in underpinning manufacturing industry and driving economic growth. Particularly important examples are 2D nanomaterials, quantum technologies, biomaterials and materials for energy efficient ICT, energy storage, nuclear power systems, and uses in extreme environments. Modelling, at multiple scales and levels of theory, will play an essential role in the discovery, design and characterization of new materials, and then optimizing and scaling up the synthesis and manufacturing processes.

Virtual Engineering: Virtual engineering has the potential to transform industrial practice by developing and testing complex products *in silico*, improving product quality and reducing time to market. It involves the integration of modelling, simulation and visualisation technologies to explore and optimise all aspects of product performance and usability, in a virtual innovation environment. It draws on and integrates a very broad range of computationally intensive modelling and simulation methods, from chemical dynamics, through computational mechanics and fluid dynamics to autonomous multi-agent systems to derive holistic simulations of whole products and ecosystems. It has been applied extensively in the automotive, aerospace, oil and gas industries, but has the potential to be applied far more extensively across all sectors.

Intelligent and Adaptive Quality Control: Developments in real time data analytics and artificial intelligence will make it possible to develop control systems for complex multi-stage processes that are able to identify unexpected variations in one process within the system, diagnose the likely consequences, and then design adaptations in other processes that will correct for the initial variations. This research is at the heart of the "Industry 4.0" vision, but will have impact well beyond that project.

Healthy Nation

Digital Human: Healthcare could be transformed through truly predictive models. The challenge is to develop computational models from the bottom up – from cells and pathways, through complex organs, to individuals and populations – that are able to predict the outcomes of healthcare interventions (or none). This requires models of normal function and of disease and it requires that these models can be personalised using individual data (*e.g.* imaging, physiological measurements, medical records, genome, self-reporting), so that diagnoses and interventions can also be personalised.

Pattern Discovery: As a complement to the model-driven *Digital Human* approach, there is also huge potential in a data-driven approach to healthcare. The explosion of health-related digital data has created an unprecedented opportunity is to find patterns in this complex multivariate data that can support stratified or precision medicine, *e.g.* identifying subgroups of patients who are likely to respond to a particular treatment, or detecting unanticipated side-effects of treatments.

Drug Development: Computer aided drug design continues to evolve as a crucial, cost-effective tool for identifying new lead compounds, validating targets, and predicting toxicity. New methods that combine deep mining of gene expression and proteomics data with fundamental calculations of molecular and electronic structure will lead to an ever greater prominence for computational screening in drug discovery.

Resilient Nation

Low-carbon/Clean Energy: Computational fluid dynamics (CFD) simulations involving turbulence, vortex flow, fluid-structure interactions *etc.* will be central to developing efficient, low environmental impact wind and tidal turbines. Simulations of turbulent reacting flows, including large eddy simulations (LES), will be important in designing high fuel-economy, low pollution combustion systems. Computationally intensive modelling and simulation will also be critical for nuclear power (plant design, reprocessing and waste storage), while modified CFD methods for high temperature plasmas will underpin longer-term work on nuclear fusion. Other important areas that depend fundamentally on modelling and simulation are hydrogen fuel cell design and carbon capture.

Sustainable Infrastructure: Increasingly complex and interconnected infrastructure is an essential feature of modern society. System-wide modelling will be crucial in understanding, designing and controlling sustainable infrastructure, such as that governing water, power, communications, employment, transport and health. Computationally intensive modelling involving many of these elements will be needed to understand the emergent properties of this complex network, to inform planning decisions, and ultimately to design and implementing optimal control mechanisms.

Natural World: Understanding and predicting the behaviour of the natural world is fundamental to developing a resilient nation. On-going research is needed to develop better computational models, and to integrate more and more extensive data sets, into packages that will provide more reliable and longer term predictions of climate change, weather, flood risk and erosion, and that will underpin strategies for mitigating the effects of natural disasters.

Connected Nation

Data Science: Developing new, efficient algorithms for data analytics is fundamental to many of the challenges outlined above. Health and city data are extremely high-dimensional, with complex temporal structure, of variable reliability, and subject to missing data and artefacts. New methods will be needed to make predictions reliable and specific. It is also a common requirement not only to make predictions from data, but also to provide estimates of their reliability.

Internet of Things: The internet of things (IoT) is again is an important underpinning technology for productivity, health and resilience. Sense-making, given the self-organising nature of the IoT, the local interactions and information flows, and the large, heterogeneous, distributed data output will present a significant, computationally complex challenge that, in practice will have to be solved using distributed computational resources. Large-scale modelling and simulation of the IoT will be an important tool in developing methods that can be deployed in the field.

Cognitive Computing: The interpretation of images, video, text, and speech are examples of intelligent technologies that will be important in realising both the Connected Nation ambition and in underpinning goals and opportunities in the other three themes. State-of-the-art methods (*e.g.* CNNs) rely on computationally intensive training that is often only practical using significant computational resources.

Important Principles and Potential Models

Guiding Principles

For the reasons enumerated above, we believe that there are a number of key features that must be built into the national e-Infrastructure landscape.

- It must include a vibrant and cohesive regional (Tier 2) layer.
- It must provide both diversity and capacity within the regional layer.

- It must develop a healthy synergy between Regional and local University e-Infrastructure, that recognises the greater strength of some Universities while still enabling all Universities to access appropriate parts of the Regional layer.
- It must recognise and accommodate the need for continuity and long-term financial planning in successful Regional Centres; in the process, it should recognise the value of Government funding in leveraging further investment from Universities.
- It must recognise that e-Infrastructure requires premises, support staff and software in addition to the hardware; in the case of Tier 2, all these elements must be located to exploit the benefits of regionality.
- It should foster the development of local communities of experts, enabling critical mass to form across institutions within a Regional Centre.
- It should facilitate the exploitation of high-end e-Infrastructure by industry, particularly (in the case of Regional Centres) by SMEs.

Mechanism for Funding

We consider here three potential scenarios relating to a Regional layer within the next generation of UK e-Infrastructure landscape. All three scenarios presuppose the existence of a national centre:

1. No Government / Research Council funding for a regional tier;
2. EPSRC regional centres, funded through an appropriate research-driven grant scheme;
3. A model similar to that used for national facilities, where service contracts are awarded through a tender response.

Scenario 1 would negate all the advantages of regional centres that we have already catalogued in this submission. It would also step backwards from the outcomes of the Tildesley report and would undo the progress the Regional Centres have made over the last four years in establishing a more connected e-infrastructure that engages better with a wider user base. Some inter-University cooperation is likely to continue, but this will fall far short of providing e-Infrastructure for the nation. We conclude that scenario 1 is completely untenable.

Scenarios 2 and 3 both allow for regional definition of the e-Infrastructure provision, but differ in the extent to which the facility is research- (scenario 2) or service- (scenario 3) driven.

Scenario 3 requires a detailed service definition to be constructed centrally so that it can be tendered against. As a result, it will tend to be blind to regional clusters of expertise, and will need to impose diversity rather than allowing it to emerge in response to local strengths. It will, however, make it easier to deliver universal coverage within the Regional layer.

Scenario 2 will be responsive to University research needs, will allow partnerships to form naturally, and will naturally generate diversity through regional variations. Some innovation will be needed in constructing a competitive grant scheme that can include diversity and universal coverage across the cohort as criteria for selecting individual winning grants. We note, however, that similar issues arose with the national call for CDTs, and suggest that similar processes would again be successful for a Regional e-Infrastructure call.

Any scheme must also allow for both on-going technology refreshes (every 2–3 years) and flexibility to incorporate new technology as it emerges. While both scenarios 2 and 3 readily allow for technology refreshes, responsiveness to new technology is far better driven at a regional level where it can emerge naturally in response to the goals and expertise of the research community. We note also that Research Councils have experience with this, since similar flexibility is built into programme grants.

We therefore conclude that a research-driven grant scheme (Scenario 2) would be the best way to deliver the key features identified at the beginning of this section. The grant scheme should be informed by experience with both CDT and Programme Grant schemes, particularly in relation to balancing provision across the cohort of Centres funded, enabling flexibility to adapt plans during mid-term technology refreshes, and being subject to mid-term review.

Evidence Base

Training and Support

Information on training and support has been obtained by surveying the ReICN Regional Centres, conducting interviews with users in our member Universities and potential users from other Universities, discussions with organisations such as the HPC Short Course Consortium, and consultation with a range of CDTs.

The ReICN Regional Centres were not funded to provide training, and operational costs were only funded for one year. As a result, there is a wide diversity of provision, and considerable reliance on existing mechanisms within the member institutions. The number of staff FTEs for support ranged from 2–8 across the Regional Centres, though much of this was exploiting local provision, and was not always able to be coordinated across the whole regional centre. Where coordination was possible, it was greatly valued by the users.

All centres provided a significant level of support for loading and maintaining major research software packages and libraries. This was seen as essential by all users and potential users. Lack of coordination of commercial software licenses, and lack of flexibility on the part of some software vendors, was seen as a significant barrier to more effective use of the Regional Centres in some areas of research, and is clearly an issue that needs to be addressed as part of the whole research computing ecosystem.

The need for reliable access to high quality training in advanced research computing techniques was a consistent theme in all the interviews and consultations. It is noted that a number of national initiatives are also relevant in this context (*e.g.* SSI, HPC Short Courses and a number of excellent domain-specific summer schools), but when consulted, users consistently noted that the timing and capacity of these training mechanisms did not accommodate the needs of a substantial percentage of computationally-oriented researchers. There remains a need for regional coordination, and potentially provision, of training to ensure all researchers are sufficiently well skilled in the computational methods they require.

A number of different strategies were adopted amongst the ReICN Regional Centres. SES made use of summer schools run at Oxford, Southampton and Cambridge; MidPlus made MSc modules delivered at one member available to the other institutions *via* access grid, with supporting practical exercises run locally; N8 devolved training to local institutions. All these training experiments proved to have their merits, but it is clear that more deliberate investment—exploiting and coordinating the advantages of Regional Centres, CDTs and national training networks—is needed to ensure that a consistently high level of training becomes available across the UK. Thus must not be left as an afterthought in the next major funding round.

Outreach

ReICN studied the activities that had been undertaken by the five EPSRC Regional Centres to engage users from Universities beyond their main membership.* The study involved a mixture of gathering quantitative data on actual use of the five facilities, and conducting a survey and interview process with a selection of people at non-member institutions. The main findings of this study are given below.

---

* The list of twenty-four universities that are members of one of the five EPSRC Regional Centres is given in the description of ReICN on page 1

- ARCHIE-WeSt and HPC Midlands currently host projects from external University groups, N8 HPC has several projects that involve users from outside N8 institutions, MidPlus has fostered use through collaboration between member and non-member Universities, and there is legacy external activity through SES.[*]
- In general the external projects arise through academic contacts, and might involve visiting researchers and Fellows, as does the external activity at MidPlus.
- Successful outreach requires dedicated staff and sustained effort, which is not a priority of the centres at present, which instead focus on supporting consortia members. For example, N8 requires its project leads to be a member of its already large consortium, whilst collaborators can come from beyond the consortium.
- The services utilised by these projects include access to the HPC resource, plus the training, support and consultancy that each centre provides.
- Access arrangements vary, but at present the centres usually require project leaders from non-member institutions to cover the operating costs of their projects at an appropriate academic rate. Some Regional Centres do provide some free exploratory access to non-member Universities, but this is necessarily limited.
- The most obvious barriers, beyond raising awareness, to uptake by external University groups are the access costs and software licencing; software is usually licensed to individual Universities rather than to a Regional Centre itself, and vendors are often reluctant to extend academic licence fees to Regional Centres.
- Future regional provision could be expanded in scope, but requires more operational staff as well as hardware to support a wider user community.
- Future access arrangements and operating costs for academics out-with the regional consortia could be aligned to those of the National centres.

Industrial Engagement

One of the key drivers of the 2012 e-Infrastructure funding was the realisation, clearly expounded in the Tildesley Report,[†] that UK industry would require easy access to a nationally coherent e-Infrastructure in order to remain globally competitive and to exploit fully the advantages that e-Science has to offer. It was thought that this would best be achieved by the public and private sectors sharing a common infrastructure, with Industry having access to expertise and skills in High Performance Computing through knowledge exchange with the Academic community.

Consequently, a major driver for the ReICN Regional Centres was to reach out and promote HPC to Industry. This has proven to be very successful in a number of sectors: across all five regional centres there have been 135 joint Academic-Industry projects utilising HPC, 18 direct Industry usage projects and a further 16 Knowledge Exchange activities involving the Regional HPC Centres in the space of three years. It is also illustrated through a number of the attached case studies.

ReICN has catalogued and analysed Industry Engagement activities across the Regional Centres in order to determine best practice and develop an effective Industry Engagement strategy. Its

[*] ARCHIE-WeSt has projects from Heriot-Watt and the Glasgow School of Art; SES has projects from Bath and Manchester; HPC Midlands has one current project and is in negotiations with four other external academic groups; N8 HPC has several projects which involve users from outside the N8 institutions such as the University of Bradford and Salford University; MidPlus has a project with Aston.

[†] "A Strategic Vision for UK e-Infrastructure", https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/32499/12-517-strategic-vision-for-uk-e-infrastructure.pdf

members have also examined the barriers to greater engagement. * The key conclusions/outcomes are:

- Almost without exception, Industry Engagement has grown out of established relationships between Universities and companies;
- This is more difficult to achieve for national facilities, which do not have their own academic-industry community to explore;
- National facilities are well placed to deliver HPC-on-demand to companies that already understand their HPC needs, while Regional Centres—with their close links to member Universities, and hence to domain-specific expertise—are well placed to work with companies that do not (yet) understand their HPC needs, and so are key enablers in exporting HPC to Industry;
- Regional Centres are working to build upon this strong foundation and seeking to establish a nationally coherent Industry Engagement HPC strategy; succeeding in this will be difficult without sustained investment in the Regional layer, and the current hiatus in funding for Regional Centres risks undermining the progress made so far;
- Best practice is still emerging on how to balance competition and cooperation between Universities in Regional Centre – Industry engagement;
- Effective industrial engagement does require Regional Centres to have significant, dedicated human resources.

Based on these conclusions, we recommend that any Industry Engagement Strategy built into the next generation of the UK e-Infrastructure landscape should:

- Recognise that different Regional Centres have different research strengths with different skills profiles and should enable those characteristics to be mapped to Industry requirements;
- Recognise that relationships are fundamental and will protect the integrity of any Academic – Industry relationship;
- Exploit the advantages of local industry/university interaction where ever feasible;
- Provide mechanisms for easy project migration to other centres and platforms when the industry's requirements are not well matched to the regions expertise and facilities;
- Enhance collaboration between national centres, such as Hartree, and Regional Centres that exploits the synergy between scale of resources and domain expertise;
- Ensure that the strategy is adequately resourced with the requisite support staff and facilities.

Delivering such a strategy requires a coherent infrastructure across the Regional Centres, and may at times require elements of common authentication, access, support procedures, operating environments and a common approach to security. This is only achievable by establishing a *coherent* Regional layer within the e-Infrastructure landscape.

* "High Performance Computing for Industry", http://www.ses.ac.uk/wp-content/uploads/sites/79/2015/11/HPC-for-Industry-blue-Report.pdf

Personal Interviews

Telephone interviews were conducted with ten research leaders, selected from both inside and outside the Regional Centre member universities. Three key attributes were consistently cited as substantive benefits of working with Regional Centres.

- *Versatility:* researchers who used both national and regional facilities found the Regional Centres were more closely linked with their users, and were able to offer a more flexible and versatile service which was able to accommodate different types of users.
- *Good support and infrastructure:* the size of the regional centres was found to generate a good balance between support staff and users, providing a high level of expertise coupled with good responsiveness. Interviewees also valued the research community they found associated with the Regional Centres.
- *Appropriate scale for the task:* many of the researchers interviewed noted that the regional facilities were well sized to provide most of their research computing capacity, while allowing scale-up, optimisation and targeting of their large jobs that really did require the resources of a national (tier 1/0) facility.

A more complete description is included in the supporting documents.

Case Studies

A selection of case studies, sampling the variety and quality of research enabled by the Regional Centres, is listed on the next page. The selection includes world-leading science in fields such as modelling DNA and rotary aircraft, and searching for extra-terrestrial life. Real-world problems with economic impact are addressed in fields of health, renewable energy, advanced manufacturing and public and environmental safety. In some projects smaller, local, companies have partnered the work, while others have involved larger companies such as Johnson & Johnson and Airbus. More details and links to the original studies are provided in the supporting documents.

Case Study title	**Scientific impact**	**Economic impact**	**Renewable energy**	**Health**	**Industrial engagement**	**Advanced manufacturing**	**Public & environ. safety**	**Engineering design**
SgurrEnergy: Using High Performance Computing for Large Scale Renewable Energy Projects		X	X		X			
Developing State-of-the-art Desalination Processes using Molecular Dynamics		X		X				
Ship Propeller Boss Cap Fin Optimisation using ARCHIE-WeSt		X						
Stopping Alzheimer’s in its Tracks: Using HPC to Understand and Prevent Beta-amyloid Aggregation		X		X				
Molecular Dynamic Simulation Study of Nanometric Cutting of Single Crystal Silicon at Elevated Temperatures		X				X		
Utilising HPC for the Prediction of Allosteric Binding Sites for Drug Discovery		X		X				
Structural Analysis and Design Optimisation using ANSYS Remote Solver		X			X			X
Modelling mutations for cancer cure	X			X				
Environmentally Friendly Heterogeneous Catalysis	X							
Slowing the aging process	X			X				
HPC Adds Extra Dimension To The Search For Extraterrestrial Life	X							
Modelling the coastal effects of a Tsunami		X					X	
Imaging Software Brings the Brain into Fuller Focus	X			X				
Computing Power Helps Researchers Unlock DNA’s Mysteries	X			X				
From biodiesel to detergents		X	X					
University of Manchester and Johnson & Johnson Protein Conformational Change using Molecular Simulation				X	X			
N8 HPC - An HPC Stepping Stone: Wind and Tidal Turbine Design		X	X					
Improving safety by predicting the lifespan of composite materials		X			X	X	X	
Increasing turbine life through novel blade coating		X			X			X
Simulations find answers for cheaper, more efficient thin-film photovoltaics		X	X		X			
High Fidelity CFD Simulations of Rotary Wing Aircraft	X	X						
DFT calculations in applications of NMR crystallography to organic molecules of importance to the pharmaceutical industry and in supramolecular self assembly		X		X	X			
Regularity of classical three-dimensional fluids	X						X	X

Statistical Indicators

The following data is the average* for an individual ReICN Regional Centre.

Average regional centre metrics calculated with available values		
Objectives	**2015**	**Lifetime (2012-2015)**
World class and world leading scientific output		
# journal papers for period	168.5	278.5
Impact case studies or press releases in period	8.3	29
Graduate and post doctorate training and support		
total number of users	454.3	404.5
Increased impact and collaboration with industry		
% use by industry	19.20%	19.20%
#industrial partners	12	56
#groups with industrial collaboration	11.7	29
#joint papers	3.5	27
Strengthening of UK's international position		
% use by o/seas	0%	2%
#o/seas collaborators	17.25	33
Operational metrics		
# active accounts per month	79.5	55.31
% utilisation of system	91.28%	83.04%
Other metrics currently recorded		
# grants mentioning regional HPC centre	31.3	108.7
total value of grants mentioning regional HPC centre	£14.9M	£48.5M
#industrial studentships	9.5	32
# KTPs / KTNs	2	7
#CDT interaction	4	11
#progressions to Tier 1	9.5	20.5

* The different Centres have different reporting requirements, and so not all indicators are available for all Centres. Averages have therefore been calculated only over those Centres for which data is available.