National Computational Infrastructure

Submission to the consultation on "Boosting the commercial returns from research"

Overview and Summary

This submission, from NCI, Australia's national, high-end, research computing service, considers the role of high-performance e-infrastructure[1] in advancing research and innovation, and the competitive advantage that it provides in generating commercial and economic outcomes. Recognition of this competitive advantage occurs throughout advanced economies and is manifest in the large investments in high-performance computing (HPC) being made in the USA, Japan, China, Europe, the UK and other nations, totalling about US$11B p.a. globally. Paralleling this is an increasing emphasis by governments on realising economic/commercial outcomes from publicly-funded research infrastructure, as exemplified in a directive from the office of President Obama to Heads of US executive departments and agencies regarding science/technology priorities for the FY2016 budget requiring that agencies "coordinate with one another and the private sector to promote innovation in high-performance computing to support national security, scientific discovery, and economic competitiveness". According to a recent IDC (industry analyst) study, the influence of HPC is such that now 97% of US companies adopting the technology believe that they could no longer compete or survive without it.

In Australia, investments in world-class e-infrastructure, established under its National Collaborative Research Infrastructure Strategy (NCRIS) and associated programs, have a strong track record in supporting[2] research of excellence, which have led to significant R&D outcomes that yield economic/commercial outcomes. In some cases, the benefits have been transformative, changing organisational practices, making possible system approaches to research and, in doing so, positioning HPC as core part of business-as-usual operations[3]. Such transformations are enabled not by any single component of the e-infrastructure but by the confluence of high-performance computing and data capabilities and expertise. This is certainly the experience of NCI over a number of years, with this holistic approach to the implementation of e-infrastructure now promulgated internationally as best practice and the way of the future[4,5].

Despite good returns on investment to date, and the growing portfolio of engagement with industry and supporting research that generates economic and commercial outcomes more generally, there exist impediments that limit the effectiveness of the "hard" infrastructure investments in delivering outcomes. These relate to the fragmentation of today's e-infrastructure (which is the subject of a separate review by the Department of Education) and to the existence of a "missing middle", caused by the paucity of the essential "soft" skills/infrastructure required to link end-users to the provision of e-infrastructure services. While these issues are not unique to Australia, the need to address them in the Australian setting is urgent to ensure competitiveness in a world economy increasingly driven by the translation of ideas to wealth.

The issues focused on in this submission relate to:

- *Skills, training and awareness:* Here, the relative paucity of high-level skills and expertise in advanced computational and data-intensive techniques, and insufficient investment to date in building these, is lim-

---

[1] Here, the term *high-performance e-infrastructure* refers to the ecosystem of high-end computing (including high-performance computers and clouds), data storage, networks, software, visualisation, and the associated skills and expertise in technical operations, and computational and data-intensive science.

[2] Exemplified in the (NCI) results of the NCRIS Research Infrastructure Survey undertaken by the Australian Government Department of Education in 2014.

[3] Dr Clinton Foster, Chief Scientist, Geoscience Australia.

[4] _, *Report of the Task Force on High-Performance Computing of the Secretary of Energy Advisory Board*, August 2014, http://energy.gov/sites/prod/files/2014/08/f18/SEAB%20HPC%20Task%20Force%20Final%20Report%2008-10-2014_1.pdf

[5] _, Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017–2020: Interim Report, http://www.nap.edu/catalog/18972/future-directions-for-nsf-advanced-computing-infrastructure-to-support-us-science-and-engineering-in-20172020

iting the translation of research into competitive outcomes, and the uptake of e-infrastructure technologies in industry. This is particularly true for the SME sector where the awareness and skills required to engage available services are rarely to hand. Accordingly, a more holistic view of future e-infrastructure investments will be required to complement the "hard" infrastructure with the necessary "soft" skills. Also curriculum developments to build capability in the researchers of the future (including training in high-performance computing and data methodologies for research students) are indicated. On-ramps, or centres of expertise, such as those that are being pursued in the UK, that build awareness in, and uptake by, industry, should be considered.

- *National research software capability:* Software, which is the "glue" of research e-infrastructure, and which transforms ideas into outcomes is a complex issue throughout the research and innovation spectrum. Software development, which is dependent on a strong base of relevant skills and expertise, takes place on a much longer timeframe than the corresponding developments in the hardware. Those countries with strong national (computational and data-intensive) software development capabilities will reap richer rewards from their investment in high-end e-infrastructure than those that do not. The national laboratories and major facilities in the US and the UK, for example, enjoy a concentration of expertise (e.g., 200+ staff) whereas Australia is yet to recognise the need. Such facilities/centres support the re-engineering of software of national importance (e.g., meteorological systems) to run on next-generation HPC systems, and are assisting independent software vendors to scale up the performance of simulation codes of importance to industry. Australia requires a similar national computational science capability, of an appropriate scale, to support its research and innovation sector, and to underpin the on-ramps to build industry's engagement with this technology. It is incongruous that Australia should be behind in this area, given its record of competitive advantage in software (rather than hardware) and the fact that international communities are increasingly looking to software innovation for higher performance in the future.
- *Shared infrastructure for shared costs and shared services:* The increasing cost of accessing world-class e-infrastructure is an issue for research institutions and industry alike. For SMEs, the costs are even more confronting, even at modest scale, when infrastructure and software may be needed for only a few weeks/months of every year. Pay-by-use models are required, with these now becoming available from some commercial cloud providers. However, the cloud does not provide necessary 'soft infrastructure', and cannot provide performance at the higher levels required for competitive advantage. The integration required for data-intensive applications in which "friction-free" computing leads to "workflows as a service" is another area not supported through cloud services, and an area that will reap the most transformative outcomes. It is here that the national e-infrastructure must be shaped to deliver not only shared costs, but also shared services (e.g., like those evolving at NCI) that benefit universities, government agencies and industry alike.
- *Long-term confidence, trust and security:* High-end e-infrastructure is now pervasive across the research and innovation spectrum, and the dependence on it has grown such that it is no longer an adjunct but a part of business-as-usual practice. The need to ensure long-term confidence in operational continuity and sustainability is important if industry is to be attracted to a national e-infrastructure, as is the need for a investment plan that is holistic in intent, and national in focus. Trust in the service, and the security and privacy of data (e.g., for medical data) and intellectual property, are critical in fostering an globally-competitive infrastructure for nationally significant outcomes in which industry, agencies, and universities can feel confident. Progress has already been made, but more is needed order to attain the "protected-level" status that will be necessary, with this having implications for the cost of recurrent operations.
- *A national e-infrastructure strategy:* A national governance framework will be required to (a) lead the development of a 5–10 year strategic roadmap for e-infrastructure and the integration of the component parts, (b) to coordinate the participation of universities, government agencies, and industry, and (c) to review progress towards the establishment of a national e-infrastructure to deliver shared services, shared outcomes, and shared costs.

Background

World-class research and innovation requires access to world-class research infrastructure, in which a world-class e-infrastructure is a core component, and a key determinant of international competitiveness. Over recent decades, there has been massive growth in computational power and storage capability, and correspondingly massive advances in our ability to harness these capabilities across the gamut of science and technology, simultaneously spawning the new disciplines and methodologies of computational and data-intensive science. The collective impact on knowledge discovery, and its commercial and industrial application, is such that now 97% of US companies that had adopted high-performance computing (HPC) believe that they could no longer compete or survive without it, according to a 2013 IDC presentation[6]. Its bearing on national economic competitiveness is driving the massive investments in HPC in the USA, China, Japan, Europe and most recently in the UK—in line with the now the common mantra[7] that "to outcompute is to outcompete", evidenced by the increasing number and performance of HPC systems in industry, government laboratories and universities in these countries/regions.

High-end computing is important to advanced economies, including Australia's, because it provides a 4–6 year advantage over that which is available from the commodity (cloud) marketplace. It enables vastly enhanced predictive power and reduces the time to market through model simulation, enables the processing, mining and analysis of large and complex data, and underpins research capabilities in areas where physical experimentation is impossible, too costly, or too dangerous.

Australia has made significant investments in world-class e-infrastructure and, in the context of national HPC, through the National Computational Infrastructure (NCI), is seeing benefits delivered, now and into the future, from research and innovation that have tangible economic, social and environmental outcomes in a range of areas including:

- Advanced extraction techniques for oil and gas—through (a) the determination of rock porosity using advanced pore modelling techniques, the business of an Australian start-up company (Lithicon) sold to a US multinational (FEI) for $76M in 2014, and (b) the use of the NCI HPC facility under contract by a leading mining company
- Simulation prototypes for renewable energy technologies, e.g., optimised combustion and combustion chamber design for biodiesel fuel
- Adaptation to regional climate change (sea level rise, flooding) through more accurate weather models (developed on the NCI research supercomputer) that will help reduce economic losses (e.g., avoid production shutdowns such as mine closures), and which will mitigate risks to human safety in extreme weather events (e.g., cyclone inundation modelling)
- Simulated design of advanced materials (e.g., nanotechnology toxicity testing)
- Economic gains in agriculture by modelling the recharging of aquifers in flood runoffs (made possible through use of land topology and vegetation data derived from satellite imagery), and advances in world food security through computational simulation that is enhancing the efficiency of the enzyme that controls photosynthesis in grains
- Underpinning a future transition to personalised medicine based on whole human genome data, processed with high-throughput/high-performance computing—with potentially large savings for future national health budgets

These examples span usage by industry directly, the science agencies of government (particularly, CSIRO, BoM, and GA), and university research sponsored by either by industry funding or the industry linkage pro-

[6] E. Joseph, *A Worldwide Overview of the HPC Market, Global Trends, Major Changes and Five Year Predictions*, IDC, October 13, retrieved from https://hpcuserforum.com/presentations/korea2013/IDCSeoulDayOnePresentation.pdf

[7] Matthew Faraci (Vice President (Communications), US Council on Competiveness), *American Business's Secret Competitive Weapon: HPC,* Forbes Magazine (May 2009) retrieved from http://www.forbes.com/2009/05/22/high-performance-computing-leadership-managing-hpc.html

grams of the national research funding councils. Coincidently, these also span a number of the priorities embodied in the Industry Growth Centres announced by the Prime Minister and the Minister for Industry in October 2014. In each case, the investment by the Australian Government in high-performance computational and data infrastructure through its national research infrastructure strategy has been successful and beneficial, and, in some cases, transformational to advancing research and innovation that yields economic and/or commercial outcomes.

Yet the potential transformative effect of the investments in e-infrastructure on commercial and industry outcomes is limited by some barriers that we now consider. Some of these may be addressed in the current reassessment of research infrastructure provision, and the development of the future research infrastructure roadmap, while others will require new bridges to be established between industry and the research sector, perhaps with assistance from programs of the Department of Industry.

Issues for consideration

Skills, Training and the "Missing Middle"

Wherever there exists a strong relationship between research and commercially/economically measurable outcomes that is advanced by access to high-end e-infrastructure, there exists a continuous bridge of skills between the core expertise of the industry, the research team that is driving the innovation (whether it be within the company or external to it), and the e-infrastructure provider. This is the case for large, mature companies that have the necessary critical mass of expertise in-house, or for start-up companies that have evolved in an expertise-rich environment of a research institution, supported by access to strong e-infrastructure and the associated expertise.

However, for many small and medium enterprises, the lack of core expertise in the use of the e-infrastructure is telling, and the skills gap between their internal capacity and that which is required to access major e-infrastructure, and particularly the HPC facilities (e.g., for engineering simulation and modelling that is required in advanced manufacturing), is a real barrier to uptake, and to the realisation of their vision or goals.

This was a principal concern noted in the influential Tildesley Report[8] ("A Strategic Vision for a UK E-infrastructure", 2011) commissioned by the UK Government in response to growing concerns that the UK was being "left behind" in the use of eScience/e-infrastructure technologies. This was further emphasised in the subsequent report[9] of the UK e-Infrastructure Leadership Council one year on, and the report of the eScience Leadership Council's Working Group on Engineering and Manufacturing[10] which identified the "missing middle" as the gap between scientific research (e.g., in academe) and its commercial exploitation and spin-out companies. The survey of UK manufacturing industry, undertaken in the preparation of reference [10], demonstrated not only a concerning lack of awareness of the availability of publicly-funded e-infrastructure in the 10% of firms that responded, but an even more concerning lack of awareness by the 90% of recipients which did not respond, and which was attributed to these firms not perceiving the survey was applicable to them, or not being able to engage with the survey because of a lack of knowledge.

In order to accelerate the translation of ideas to products (e.g., via prototype simulation), it was recognised in the UK that software and data "on-ramps" for industry, in the form of expertise centres to build awareness and provide a one-stop shop for access to the e-infrastructure, software and expertise (e.g., in computational manufacturing technologies) would be required. The UK can point to various activities (not all co-ordinated centrally by government) over the years, some in economically-depressed areas (e.g., in Wales, Scotland, and most recently Northern Ireland) that have been funded by the British Government and the European Commis-

[8] D. Tildesley, *A Strategic Vision for UK e-Infrastructure: A roadmap for the development and use of advanced computing, data and networks,* E-science Leadership Infrastructure Council, November 2011

[9] _,*E-infrastructure: The ecosystem for innovation — one year on*, E-Infrastructure Leadership Council, October 2013

[10] _, D*eveloping e-Infrastructure in the UK's Engineering and Manufacturing Industries*, eScience Leadership Council, June 2012 (a report commissioned for the Department of Business and Skills).

sion, together with its premier Hartree Centre (Science and Technology Facilities Council), which provides an internationally significant infrastructure and skills base that has a strong focus on industry and economic outcomes.

Almost certainly, comparable initiatives will be needed in Australia if we are to address corresponding problems, and support key themes, e.g., those associated with the Industry Growth Centres. The capacity to implement “awareness” programs, and scale up the expertise base, is quite limited within the current research infrastructure funding arrangements for the national HPC facilities, which are required to raise the bulk of their recurrent funding through partner (universities and science agencies) co-investment. These arrangements leave little latitude with which to boost industry uptake, particularly from the SME sector, since the funding base is both stretched and significantly prioritised by virtue of its source. Depending on the outcomes of this consultation, it is likely that “on-ramps” will need to engage additional expertise in areas of priority in industry, combined with business development capabilities in these centres through which to build bridges. The effective use of monies suggests that expertise needs to be concentrated around the HPC facilities in order to build critical mass and ensure skills diversity, and yet be mobile from this base(s) to drive the technical engagement with industry.

Compounding the problems in building engagement with industry is a general shortage of skills in eScience/eResearch/e-infrastructure capabilities, in Australia and internationally. Many of the next-generation of jobs in research and innovation, whether they be in industry, science agencies or academe, will require much heightened computational and data-intensive skills, and particularly in the context of the big data and data analytics that will drive a number of emergent industries. Present university curricula do not generate sufficient future staff for current organisations (e.g., science agencies, universities, facilities such as NCI, etc.), let alone for the likely future needs of Australian business, given international trends and the initiatives underway in advanced economies.

Accordingly, there is a side imperative to foster the development of essential skills through changes to the undergraduate and postgraduate science and technology curriculum, including a heighted emphasis on industry engagement. In the UK, the Tildesley Report recommended the embedding of eScience training in all relevant doctoral training, and activities are underway to expand the scope of the UK’s existing Doctoral Training Centres to include HPC and data-science training, as well as other activities to build an awareness and outreach program for industry, SMEs and universities.

Software

Software is a vital component in, and a complex issue for, all parts of the research and innovation spectrum. Some industries embody aspects of their IP in software codes that are written in-house, although this practice tends to be restricted largely to either substantial, mature corporations that have the full range of skills in-house, or to start-up companies that manifest their IP in a software product. For the majority of industry, however, the software that is used is either open source, or from a (commercial) independent software vendor (ISV).

Research organisations and industries that invest in engineering their own software are confronted by a software development lifecycle on the timeframe of a decade that is being outstripped by advances in hardware technology that occur on a 2–3 year timeframe. For such organisations, whether they be industries or science agencies of government or research universities, there is a growing legacy of software (open source or internal proprietary) that needs to be migrated to the new hardware technologies, and/or up-scaled in performance in order to maintain their competitiveness and take advantage of the capacity of current and next-generation hardware. A significant fraction of the ISV software is also waning in its performance competitiveness, relative to the capability of contemporary hardware. This is understandable since many ISVs are themselves SMEs and are driven by what the market can exploit, what the market is prepared to pay for, the cost of re-engineering codes, and the expertise that is available to them for the re-engineering of their code base.

Tildesley, in his report to the UK Government, conceptualised an ecosystem comprising industry partners, research institutions (including universities), and suppliers, including ISVs, that would create a "substantive and internationally-leading software base", built around a national infrastructure that would be accessed for mutual benefit. In such an environment, Tildesley envisaged that ISVs might provide "versions of their code for limited parties in exchange for its migration on to high-end systems". The UK has a number of expert resources on which to build, including the Hartree Centre, which has a substantial staff (150) and offers services a range of services including software development, applications optimisation, and HPC services provided on demand to industry. Their clients include the UK MetOffice (to which it is providing expert consultancy for the redevelopment of the UK's weather models for use on future generation high-petascale/exascale hardware), the UK motor (Jaguar, Land Rover, Bentley) and aerospace (BAE systems) industries, and other large and small/medium industries.

In the USA, the National Centre for Supercomputer Application (NCSA), the largest of the US NSF funded facilities, and the operator of the largest open access HPC system in the world (Blue Waters), operates a very successful Private Sector Program (PSP) that is spearheading the collaborative re-engineering by leading ISVs of engineering software to operate at high core counts on top tier HPC systems (including NCSA's Blue Waters system). This is driven by NCSA's partnership with approximately one-third of the Fortune50 companies, including Caterpillar, BHP Billiton, ExxonMobil, Rolls Royce, GE, and Procter and Gamble, amongst others, with the benefits to SMEs becoming available through more highly scaling tools with which to work. In addition, the NCSA PSP operates a dedicated HPC cluster for industry services and provides a high-tech mobile consulting team to industry, leveraging the substantial expertise (200+ staff) within NCSA for code performance optimisation and big data.

These two examples, which illustrate highly successful industry interaction, also demonstrate the critical role of a significant software development capability. It is this that Australia presently lacks, and which will be needed to foster deeper engagement with industry, and provide a strengthened skills/expertise base in computational/data sciences with which to support research and innovation. Without this, Australia will not benefit fully from its future investments in high-performance computing and data infrastructure (in research, in industry, and in operational production environments, such as meteorological facilities), and will lack the ability to track the competitive advantages that are, and will continue to be, reaped in the USA, Japan, China, Europe, the UK, and other advanced economies.

Shared Infrastructure for Shared Costs and Shared Services

The cost of accessing competitive e-infrastructure, both hardware and software, is a significant barrier in industry, for all but the largest of companies. It is a particularly significant issue in the SME sector for which access to computing resources and design software may be required for only a few weeks or months every year. Such concerns were identified in the Tildesley Report to the UK Government, in which it was commented, "significant economic growth will only be achieved through a shared infrastructure and a sharing of the costs of establishing and running it."

In Australia, there is significant progress to report on establishing shared infrastructure for shared services, particularly at the high-end where the costs exceed that which can be borne by any single institution. The NCI Collaboration (involving ANU, CSIRO, BoM, GA, research-intensive universities and consortia thereof supported by the Australian Research Council), which provides integrated and transformative high-performance services to support the work of industry, science agencies and universities, is a successful example of the benefits of shared costs and shared outcomes.

For industry, pay-on-usage models are required, in which software can be provided as a service—as is now the case in some commercial cloud service providers. The national research infrastructure will need to be correspondingly nimble and implement software licencing arrangements with ISVs that can meet this need. While NCI maintains a substantial software library, including software from commercial ISVs, the present licences are restricted to academic or research usage, and will need to be upgraded for commercial application.

The same nimbleness will also be required in hardware provision, with experience from NCSA's Private Sector Program suggesting that industry seeks access to the latest and most powerful computing hardware. To achieve this, NCSA leases a dedicated HPC cluster, and turns over the hardware on a regular basis. The Hartree Centre of the UK also provides commercial access for industry through a dedicated cluster. Similar arrangements could be implemented in Australia, within the major HPC installations that have a comprehensive and well integrated e-infrastructure.

While it is inevitable that commercial cloud service providers will be strongly involved in service provision to industry, the present hardware offerings are not yet competitive on price relative to the performance that is available, or required. While the situation will change, a differential must always remain since cloud providers operate at the commodity layer, which is a different performance and price specification to that of the HPC level of service that is sought by particular classes of industry. Further, the skills gap, referred to earlier, necessitates a higher, more expert and more application-oriented, level of service than is available from cloud providers. It is arguable, therefore, that the future advancement of industry's access to advanced computing will require a one-stop shop through national HPC centres, through which a mix of in-house and outsourced infrastructure provision will be integrated.

Long-term Confidence, Trust and Security

The Tildesley Report also identified impediments to industry uptake of e-infrastructure by a lack of confidence in its long-term future. It observed that "firms cannot interact easily with a fragmented national infrastructure without a clear forward commitment to a strategy." Tildesley went on to observe that "any move from a local, in-house provision of computing to the use of a shared infrastructure will only be achieved if a clear plan is developed and adhered to."

Australia similarly suffers from fragmentation of its national e-infrastructure, with previous planning having been conceived this in terms of disjoint infrastructures and services, rather than informed by an overarching goal of advancing research and innovation outcomes. While progress has been made by high-end facilities, such as NCI, in defragmenting the landscape by integrating the infrastructure and service component, a future national e-infrastructure needs to conceptualised from the outset as an integrated whole, in which the HPC facilities have a pivotal role with respect to national services and priorities.

Long-term certainty of service provision is also a significant issue for industry, a point also emphasised in the Tildesley report. Presently, the boom-bust cycles of infrastructure funding may be a disincentive to industry engagement with a national research infrastructure. A transition to a mainstream environment in which long-term certainty of operations, as well as the periodic refreshment of the technology will be required.

Compounding these concerns is the additional issue of the security, and the trust that can be placed in outsourcing arrangements. Progress towards a heighted security framework has been made over the past year at NCI, driven by its transition from a university supercomputing facility to a national, integrated e-infrastructure service—with expectations of a robust and reliable research production service being shaped by the requirements of the national science agency partners. While progress has been made, there remains more to do, including annual audits and the likely need to demonstrate certification at the "protected" level of security classification, in order to meet the security requirements of industry and government agencies, and associated also with data privacy requirements, particularly in the medical arena. These also will have implications for operating costs that will need to be factored into infrastructure planning and business models.

Other matters that require harmonisation include network provision, given the potentially fragmentary impact of the differences in business models of the providers of research network services to universities (AARNet), and to industry and the science agencies (i.e., commercial ISPs) — notably around data access and services. Authentication and authorisation services will also require further harmonisation.

National Focus for e-Infrastructure Planning

International benchmarks tell us that the establishment of an e-infrastructure of the requisite scale to provide internationally competitive facilities, the required level of operational robustness, the depth of expertise, and the diversity of capabilities, in which industry and other parts of the innovation sector can have confidence, requires a national, rather than a regional or institutional, focus. A suitable foundation from which to establish an industry-supporting e-infrastructure are the national high-performance facilities that provide a world-class capability, can demonstrate the requisite expertise and operational experience, and can provide the integration, performance and support required to accelerate innovation and productivity.

Only a national focus can establish an infrastructure capable of delivering the transformative outcomes that are needed in areas of national priority and significance. Almost certainly, one such priority will involve the use of whole human genome sequence data (now being generated by high-throughput sequencers, i.e., the $1,000 genome) to drive untold advances in medical research and clinical practice, and ultimately form the platform for personalised medicine that will deliver not only better health outcomes, but also bring under control the ballooning health care budget. A national computational genomics infrastructure will be needed to centralise the processing and analysis of the vast data streams and repositories, in order to optimise access, maximise its utility in research and clinical practice, and avoid fragmentation that would weaken system approaches to knowledge discovery, as well as overcoming concerns about the provenance and ethical handling of data.

The realisation of such an objective will necessitate a coherent roadmap for the national e-infrastructure—one that emphasises national priorities and research/innovation outcomes, and which overcomes the fragmentation caused by misaligned investments and insufficient focus on national goals.

Conclusion

The NCRIS investments in e-infrastructure, combined with those made by institutions, fulfil a critical role in supporting almost the entire spectrum of research and innovation in universities, public sector agencies, and industry. To achieve more in direct support of industry, and in the generation of commercial returns on research, national governance is required to lead the strategic development of a 5–10 year roadmap and the integration of the component parts, and to coordinate the participation of industry and commerce, government agencies, and universities.

A model for this may be the holistic pathway along which the UK is proceeding following the Tildesley review, involving the establishment of an *e-Infrastructure Leadership Council*, co-chaired by a Minister of the Crown (the Minister for Universities and Science in the BIS portfolio) and a senior representative of industry, and comprising representation from industry, government departments and agencies, the university sector and the research funding councils, and the charitable sector.

27 November 2014