EPSRC
Engineering and Physical Sciences Research Council

SOFTWARE AS AN INFRASTRUCTURE

1. Introduction

A coherent strategy for developing and delivering the UK's future e-infrastructure[1] needs is essential in driving forward the continued development of a globally competitive research base within the UK. Since 2009 there have been a number of BIS, Research Council and community led reports and reviews on e-infrastructure, outlining our future requirements and how e-infrastructure should be funded and supported.[2]

The most recent report is the UK e-Infrastructure Strategy for Science and Business commissioned by BIS and chaired by Dominic Tildesley. Within this report a 10 year roadmap for investment will be proposed for networks, data and storage, compute, software and algorithms, people and skills and security and authentication. The additional £145M of capital investment in e-infrastructure recently announced by BIS builds on the initial outputs from this report.

The impact of this additional investment is still to be determined; however support for software will be important in maximising the long-term value that the community gains from this.

Computer-supported modelling and simulation is now widely recognised as the third `leg` of scientific method, alongside theory and experimentation. Many phenomena can be studied only by using computational processes such as complex simulations or analysis of experimental data, including that produced by large facilities such as Diamond and ISIS. Data analysis software is equally as important, particularly with the recognition of datasets being research outputs in their own right. Software is where much intellectual property, knowledge and understanding resides and this is why software has such longevity: people replace their hardware, but don't dispose of their codes.

The large suite of codes used in research therefore needs to be regarded as a research infrastructure in its own right, requiring support and maintenance along the innovation chain, and throughout its lifecycle. As such it is important for EPSRC to be able to articulate our strategy for investing in software, ensuring that our funding (current and future) adds value to the complex and evolving e-infrastructure eco-system and supports the needs and requirements of the EPS communities.

2. EPSRC Future Strategic Framework and Plans

2.1 STRATEGIC FRAMEWORK

Over the last five years EPSRC has invested approximately **£9 million per annum** in software. This has covered the spectrum from new algorithm development at the leading edge of research

[1] E-Infrastructure refers to a combination and interworking of digitally-based technology (hardware and software), resources (data, services, digital libraries), communications (protocols, access rights and networks), and the people and organisational structures needed to support modern, internationally leading collaborative research be it in the arts and humanities or the sciences. This definition reflects a broader understanding of e-Infrastructure as defined in the report "Delivering the UK's e-Infrastructure for Research and Innovation."

[2] Report of the e-Infrastructure Advisory Group 2011; Delivering the UK's e-Infrastructure for Research and Innovation report 2010; Review of e-Infrastructure 2009

applications through software development to code maintenance. It also includes training, and community support activities, such as networking. The current portfolio of activities has supported a thriving community of computational scientists who are recognised internationally. A summary of these activities is shown in **Annex 1**.

We now want to produce a strategic framework for engineering and physical sciences research that will allow us to provide the appropriate level and type of support across the software landscape.

The strategic framework below has been formulated with input from members of the community *via* a workshop held in October 2011. A draft strategy document was published at the end of 2011 and comments from the community were invited. The comments received have led to some minor updates to this document and have also fed into the associated action plan.

It is important to note that it is our intention that the scope of this strategic framework expands beyond users of the national HPC system, and covers the pyramid of software usage including users of local clusters, and desk-top systems.

2.2 ADDITIONAL INFORMATION

2.2.1 Strategic Goals:

The framework is structured around the five strategic goals, the three corporate EPSRC strategic plan goals of **Shaping Capability**, **Developing Leaders** and **Delivering Impact** and two additional goals that we have developed with the community, that are more specific to software related research and development: **Ensuring Trust** and **Planning for the Future**. These strategic goals describe what we would like to achieve through our support of software as an infrastructure.

2.2.2 Objectives:

Against each strategic goal there are a number of key objectives that were identified through engagement with the community. These objectives outline streams of work that will enable us to meet our strategic goals for software.

2.2.3 High Level Interventions:

Against each objective a number of high level interventions have been developed with the help of the community. These interventions describe, in broad terms, activities that EPSRC and the community could pursue. Interventions have been mapped to show whether they would have a short, medium or long-term impact on the software landscape. The actions arising out of these interventions, following feedback and input from the community, are articulated in the associated action plan. Some of the interventions planned over the longer term may need further development before clear actions can be set out.

2.2.4 Mechanisms for implementation: What we will do in the future?

During the community workshop, EPSRC worked with attendees to build more detailed plans on what the high level interventions for each objective could mean in practise. This information has been supplemented by discussions within EPSRC and planned activities that have come out of these discussions have been incorporated into the action plan associated with this strategy.

Strategic Goals	Objectives	Interventions with Short-term impacts	Interventions with Medium-term impacts	Interventions with Long-term impacts
Shaping Capability: Building the capability to deliver high quality and important research through the development of sustainable and robust software.	*Identification of new areas and grand challenges*	• Work with EPSRC themes and the community to understand existing research needs and requirements, linking with the EPSRC Shaping Capability activity. • Gather and disseminate exemplars of computationally enabled grand challenges.	• Identify new research areas and user communities, working with Research Councils, Institutions and the community (academic and industrial). • Broker relationships with nascent communities.	• With the community and other research councils develop a number of cross-cutting grand challenges that are relevant to both industry and academia. Provide support for collaborative and cross-disciplinary team working in these areas.
	Enabling and promoting collaboration	• Support more effective networking across communities, e.g. Maths, computer science, HPC, computing and industrial and academic users. • Support more effective collaboration between existing projects supported by EPSRC.	• Support more effective collaboration across communities and change "roll-your-own" culture. • Promote more effective working with the large facilities such as ISIS, Diamond and CLF on software development.	• Invest in tools and frameworks that help software and hardware developers to work together to co-design hardware and software. • Develop international links and opportunities that are beneficial to the UK community.
	Research and Development	• Work with EPSRC themes and the community to map existing research needs and requirements, linking with the EPSRC's Shaping Capability activity. • Develop and re-engineer existing code for existing architectures in key areas and encourage code consolidation where appropriate.	• Develop and re-engineer key existing code for new architectures.	• Develop novel code in key application areas. • New coding platforms for current and future architectures.

Developing Leaders and Skilled People: Building the software correctly by building on current expertise, supporting the careers of researchers (including developers) in this domain, providing training, and actively encouraging collaborative working and inter-disciplinarity.	*Training*	• Provide training and skills development to students and researchers that is suited to their requirements and level of expertise.	• Work with industry to develop internship models, increasing researcher understanding of the industrial environment and the issues and challenges faced by industry.	• Support the development of a clear career path for developers and researchers in this domain, from PhD onwards.
	Career Path Support	• Work with institutions that are successfully supporting and developing career paths for software developers. Develop best practise that can be disseminated to institutions. • Acknowledge contribution of software specialists in research proposals. • Raise awareness in institutions of the research potential, importance and value that software engineering can bring, highlighting available investment in the UK.	• Build careers at the interface, by supporting researchers to work in different research environments (e.g. maths, computer sciences and application areas) through discipline hopping type activities and networking. • Develop measures and success features that Universities and Research Councils can use to recognise innovation in software provision. • Encourage and support international collaboration with centres of excellence through the exchange of people and expertise.	• Support the development of a clear career path for developers and researchers in this domain, from PhD onwards. • Provide appropriate long-term support mechanisms for potential and established leaders e.g. through fellowships and underpinning platform type funding.
Delivering Impact: Maximising the academic, societal and economic impact of current and future investments in software by providing appropriate user	*Joint funding models*	• Actively seek to broker greater industrial involvement in current investments.	• Develop joint funding models with Industry and International partners to sustainably support software through "Follow-on funding" type activities.	• Develop longer-term, strategic engagement with potential user communities, to understand the cross-cutting issues that are affecting industry. Determine whether an industry club model would be appropriate.

support, developing business models and promoting collaboration with industrial and international partners.	*Supporting Innovation*	• Support innovation in software through the short, medium and long-term. Providing opportunities for innovation wherever it might occur. • Work with existing centres to identify and disseminate best practise for software development and exploitation.	• Support innovation in software through the short, medium and long-term. Providing opportunities for innovation wherever it might occur. • Develop future business models for software exploitation	• Support innovation in software through the short, medium and long-term. Providing opportunities for innovation wherever it might occur
	User support	• Invest in optimising the use of highly regarded current codes. • Provide user support including support for software updates and maintenance, recognising different levels of user community knowledge.	• Provide support for the development of sustainable and adaptable code. • Build on existing work in the CCPs and SSI to support a `code depository` and MOT centre.	
Ensuring Trust: Ensure that developers and users can trust in a robust and reliable software	*Quality of code*	• Develop testing and verification suites.	• Develop peer review models that allow the separate assessment of the science case *vs* the software development component.	• Encourage culture of software publishing and peer review scrutiny, with acknowledgement of developer contribution.

infrastructure by providing support for verification, testing and use.	*Sustainability of code*	• Incentivise researchers to share software outputs through OpenSource. • Build communities of developers and users around software codes.	• Develop joint funding models with Industry and International partners, enabling early involvement in projects from the start, and building trust between developers, researchers and industry partners. • Provide support for professional software developers to be involved in projects from the ground up, ensuring that reliability and robustness is incorporated from the start.	• Understand and characterise trends in hardware development to allow the community and EPSRC to plan future investments required to provide robust and reliable software infrastructure.
Planning for the future: Developing a longer-term strategy and sustainable funding streams for software support.	• Deliver and communicate a clear, shared strategy developed in dialogue with the community. • Develop interconnectivity with other funding agencies to provide a clear statement of strategic intent on investment in e-infrastructure and hardware. • Understand and characterise trends in hardware development, allowing the community and EPSRC to develop future software requirements and challenges that require investment. • Work to establish suitable mechanisms for funding software, and work with the community to plan for long-term sustainability of funding. • Support innovation in software through the short, medium and long-term. Providing opportunities for innovation wherever it might occur.			

Annex 1: EPSRC activities in software development over the last 5 years.

Activity	Value
Computational Science and Engineering (CSE) support to users of national service (NAG Ltd): • Help Desk. • The distributed Computation Science and Engineering (dCSE) team provides longer term, tailored support (3 – 12 months) to individual research groups in specific code development activities. 50 person years of support are provided through the contract with EPSRC. • Training courses: NAG provide ~70 person days of trainer time to HECToR users per year. To date, roughly 1000 people have attended these courses.	£3 million per year
Service level agreement with STFC, Daresbury: • Computational science support to the UK academic community in research areas of interest to EPSRC, in order to maximise exploitation of local and national supercomputing facilities.	£2.4 million per year
Collaborative Computational Projects (www.ccp.ac.uk): • CCPs bring together a community of major UK groups in a given field of computational research to tackle large-scale scientific software development projects, maintain and distribute code, and provide training and user support. They are community led. Eight new CCPs have just been funded from 10/2011-4/2015.	£2.7 million
HPC Software development projects: • Funded through 3 calls for proposals.	£13 million
Science and Innovation Award: • Centre for numerical algorithms and intelligent software.	£4.5 million
Software Sustainability Institute: • This provides training, best practice and software engineering support to increase the long term sustainability of research software, including HPC codes. Funded 2010-2015.	£4.4 million
Extreme Computing sandpit: • This activity brought together experts from across the HPC, computer science, mathematics and scientific application areas to identify novel approaches to the development of software to exploit the next generation of HPC hardware. Three projects were funded.	£2.4 million
HPC Short Course Centre: • This virtual centre is a collaboration between 13 UK HEIs to provide an advanced training programme in HPC methods for researcher and students in the UK. The Centre has funding for three academic years 2010-2013.	£0.3 million
NSF/EPSRC Software Development Call: • Bringing together researchers from US and UK to develop research projects in computational chemistry. Initial funding for networking and travel, followed by funding for research projects. EPSRC funding is matched by NSF funding.	£3 million
Crossing the Chasm: • This activity aims to promote the wider uptake and exploitation of new and existing UK e-Infrastructures by promoting the two-way flow of knowledge between users of e-infrastructure (particularly academic researchers) and providers. Funded 2011-2013.	£250,000
e-Science Platforms: • E-research South (Oxford, Southampton, Reading)	£2.9 million

(http://www.eresearchsouth.ac.uk/) • White Rose Grid e-Science centre (York, Leeds, Sheffield) (http://www.wrgrid.org.uk/) • National e-Science Centre Research Platform (Edinburgh, Glasgow) (http://research.nesc.ac.uk/node/259) • Belfast e-Science Centre (http://www2.besc.ac.uk/)	
Numerical Algorithms and High Performance Computing network: • Providing a focus for a new collaboration between numerical analysts, computer scientists and developers and users of software and HPC within the nodes, supported by the necessary administrative organisation.	£130,000