

UK Participation in HPC at the European Level

This report analyses the vision, policies and support, which currently exist for HPC at the European level, and recommends the strategic approach that the UK should take towards participation in it.

Executive Summary

- 1. Strategic Importance.** High Performance Computing (HPC) and e-infrastructure in general are drivers of economic growth and societal well-being. They are also vital for maintaining international competitiveness in the generation of knowledge and its application. A European Commission (EC) Communication¹, issued in February 2012, highlights the strategic nature of HPC and *“calls on Member States, industry and the scientific communities, in cooperation with the Commission, to step up joint efforts to ensure European leadership in the supply and use of HPC systems and services by 2020”*.
- 2. Need for Action at the European Level.** The scale and cost of the infrastructure needed to support the most demanding computationally and data intensive applications requires some organisation and cost sharing at the European level, although maximum benefits can only be realised through integration of this with balanced investment at national and regional levels. The EC Communication recognises that *“the increasing importance of HPC for researchers and industry, as well as the exponential rise in the investments required to stay competitive at world level, have led to a common understanding that ‘Europeanisation’ of this domain would benefit everyone”*.
- 3. UK Engagement.** The UK has a long history in Europe of leading the debate on HPC, providing innovative thinking on how to move forward, and leading some key projects. The UK has an equally long history of failing to back this up with sustained engagement with the EC due, in the main, to the lack of coherent ownership by any UK organization. With its focus on excellence, support for SMEs and societal challenges, and streamlined funding mechanisms, the next EC programme of support for research and innovation, Horizon2020, is a major opportunity for the UK. The UK would derive greater benefits from investments in HPC by the EC and other Member States if it employed a more joined-up approach to its engagement in planning, implementing and exploiting HPC at the Europe-wide level, recognising the breadth of the interface with European activities and exploiting both bottom-up and top-down channels of influence. **Significant opportunities exist to enhance and share the associated cost of e-infrastructure and its exploitation for UK research and innovation through the EC Action Plan in its Communication on HPC and the Horizon2020 programme, and the UK should make it a priority to engage coherently with these. (Recommendation 1).**

¹ http://ec.europa.eu/information_society/newsroom/cf/document.cfm?action=display&doc_id=891

4. **Academic Research.** Across an expanding range of disciplines, internationally-competitive academic research is dependent on access to leading-edge e-infrastructure, including high-end HPC, approaching exascale capabilities by 2020, that can only be afforded at the European level (tier 0). The Partnership for Advanced Computing in Europe (PRACE)² is providing access to tier-0 resources through peer review based solely on excellence. The UK position in PRACE is unsustainable, in that its researchers are winning a disproportionate share of the resource compared with the UK contribution. **To ensure that the most computationally demanding UK research has access to the next generation exascale platforms, it is essential that the UK engages with PRACE and sets a timetable for achieving full partnership by mid-2013 (Recommendation 2).**
5. **Sustainability of PRACE.** The EC Communication promises continued support for PRACE and encourages it to expand its role *“to (i) pool national and EU funds, (ii) set the specifications and carry out joint (pre-commercial) procurement for leadership-class systems, (iii) support Member States in their preparation of procurement exercises, (iv) provide research and innovation services to industry, and (v) provide a platform for the exchange of resources and contributions necessary for the operation of high-performance computing infrastructure”*, with the aim of becoming *“a globally leading e-Infrastructure”*. However, Member States’ commitments to PRACE expire at the end of 2014. **The UK should play an active role over the next year in transforming PRACE so that it can sustainably support, via an equitable contribution model, world-leading open research by academia and industry through the provision of a level of computing infrastructure and access which is beyond the capabilities of any one European nation (Recommendation 3).**
6. **Data Deluge.** To a greater or lesser extent, in all areas of computational science, the strategic importance, complexity and size of data sets are growing rapidly, so that most researchers regard data management to be as big a problem in the approach to exascale as the scalability of their codes. Furthermore, at all levels in the e-infrastructure, from microprocessor architectures to data archives, the cost of data storage and movement has become non-negligible. More than any other aspect of HPC, this demands coordinated action at local, regional, national and international levels, and between researchers and providers. Efforts in this direction at European level are comparatively recent, involving several uncoordinated initiatives, although the pace of activity is accelerating. **The UK should take an integrated approach to e-infrastructure for data-intensive and compute-intensive applications and engage immediately with work to develop this at the European level, to ensure that the UK achieves full integration with, and maximum added value from, the imminent establishment of European data services (Recommendation 4).**
7. **Support for Industrial Use of HPC.** Few SMEs have the capability to exploit directly HPC infrastructure or funding at the European level. Nevertheless, increasing business opportunities and competitive advantage exist from doing so. SMEs require support from universities or large companies with experience in European projects. Support is also needed for Independent Software Vendors (ISVs) to scale their codes to make effective use of tier-0 systems, so that applications and tools are available to enable businesses to exploit higher levels of the e-infrastructure effectively. **The UK should establish support structures on a regional basis to provide bridges for SMEs to engage effectively with the opportunities provided through EC**

² <http://www.prace-ri.eu/>

- funding and shared European infrastructures, and should encourage the EC to establish mechanisms for sharing best practice across Europe (Recommendation 5).**
8. **Share of the HPC Technology Supply Chain.** Relevant UK companies should be encouraged to join the European Technology Platform ETP4HPC³, which will develop a strategic research agenda for HPC technology, although many SMEs will not have the capability to do so directly, or be sufficiently knowledgeable about mechanisms such as Pre-Commercial Procurement (PCP) that the EC is promoting to support HPC technology R&D. **To ensure that the UK is positioned to derive increased value from its domestic HPC supply chain, the UK should establish a support structure for companies that could contribute to the HPC technology supply chain, to pool knowledge and expertise and broaden UK engagement with ETP4HPC (Recommendation 6).**
 9. **Pre-Commercial Procurement.** PCP provides an opportunity for UK companies to benefit from substantial Horizon2020 funding. **The UK should clarify its position on Pre-Commercial Procurement, identify technology developments where UK companies have, or could acquire the capabilities to contribute and, through active engagement with ETP4HPC, encourage the EC and Member States to direct their PCP spend towards them (Recommendation 7).**
 10. **Coordination of UK HPC Support.** The UK needs an authoritative voice for its HPC interests that is both outward-facing to influence the EC and inward-facing to ensure effective coordination of UK HPC activities. The breadth and complexity of HPC activities at both European and national levels, coupled with the growing reach into business, research and development, policy-making and emergency response, demand a focus of expertise within government and raised awareness across all government departments. This must be the centre of a network that spans all user communities and extends vertically through government, acting for and reporting to the e-Infrastructure Leadership Council. **The UK should establish an executive arm of the e-Infrastructure Leadership Council, an "e-Infrastructure Directorate", that is user-centric and tasked with achieving more coherent engagement between HPC at the European level and within the UK, connecting both top-down and bottom-up channels of communication with the EC (Recommendation 8).**
 11. **Training.** To maintain its competitiveness, the UK must build business and research communities capable of exploiting all levels of the HPC ecosystem. The advanced skills must be an integral part of postgraduate training, through Centres for Doctoral Training (CDTs) and MSc programmes, but the basic foundation of mathematics and computing competence must be laid at undergraduate level. Opportunities exist to leverage training through expertise that exists across Europe, particularly for those technologies that only exist at the European (tier-0) level, e.g., building on the PRACE Advanced Training Centres⁴. **The UK should introduce training in the use of HPC technologies employed at all levels in the ecosystem for all postgraduate research students in relevant fields and, particularly, it should participate in the development and provision of training on tier-0 facilities on a Europe-wide basis, e.g., through the PRACE Advance Training Centre programmes (Recommendation 9).**

EC Support for HPC and e-Infrastructures

12. The EC has a long history of funding HPC and e-Infrastructures, dating back to the Europort and the High Performance Computing and Networking (HPCN) Technology Transfer Node network

³ <http://www.etp4hpc.eu/index.html>

⁴ <http://www.prace-ri.eu/PRACE-Advanced-Training-Centres>

programmes of the mid-1990s, and before that, in the 1980s, GPMIMD. These programmes were managed by the HPCN Unit of the Information Society Directorate General (DG INFSO) (then called ESPRIT). Today, the EC provides support for research and innovation through three programmes: the Framework Programmes for Research and Technological Development (currently we are at Framework Programme 7 – FP7), the innovation related activities of the Competitiveness and Innovation Framework Programme (CIP) and the European Institute of Innovation and Technology (EIT).

13. UK universities generally do well out of the Framework Programmes; UK industry generally does less well. While some of the big companies (e.g., BAE Systems and Rolls Royce) are successful in exploiting EC funding, many companies view the programmes as overly bureaucratic, daunting and lacking value, and most SMEs are unable to engage at all. It is difficult to get a complete picture, because EC funding has been distributed across many areas, but, in FP7 and the preceding programmes, the UK has, over time, roughly recovered in total cash terms what it has put in.
14. From 2014, Horizon2020, with a requested budget of EUR 80 billion (at 2011 prices, for the seven years 2014-2020), will bring together all EC research and innovation funding. However, it is as yet unclear if Member States will agree to this increase. Horizon2020 will prioritise spending on projects with an immediate impact on growth in jobs through, amongst others, large-scale pilots, demonstrators for key technologies and developments closer to market. Simplified rules and auditing procedures will be introduced to reduce the average time to grant to 100 days.
15. Horizon2020 will comprise three main pillars: **excellent science** (EUR 24.6 billion) to support the best researchers in a bottom-up approach, strengthening of the role of the European Research Council (ERC), which includes e-infrastructures; **industrial leadership in innovation** (EUR 18 billion), specifically aimed at investments in key technologies, providing greater access to capital and support for SMEs; and tackling **societal grand challenges** (EUR 32 billion) to address political priorities. Horizon2020 funding will be allocated on criteria of excellence, with no geographical consideration.
16. The focus on excellence, support for SMEs and societal challenges, together with streamlined administrative processes, match UK priorities and address some of the barriers to wider UK participation by both academia and industry.

EC Vision for HPC in Horizon2020 and PRACE

17. The EC vision for HPC is spelled out in its 2012 Communication¹. The transition from today's petascale computing to exascale, during the rest of this decade, will entail radical innovation in technologies and applications, which *“offers opportunities to industrial and academic players in the EU to reposition themselves in the field”* and, in particular, get *“the EU back on the world scene as a leading-edge technology supplier”*.
18. On current technology trends, it is widely felt in Europe that exascale HPC will not be affordable by any individual nation. The several interconnected technology challenges that must be overcome has led to the US DOE extending the original highly ambitious timeframe of building the first exascale machines by 2018 to 2022-24. In a parallel initiative to the US International

Exascale Software Project (IESP)⁵, the European Exascale Software Initiative (EESI)⁶ has mapped out the system and applications software roadmap. Innovation is required at all levels, creating opportunities for UK academia and industry. One likely avenue is that the ultra-low power processors and memory systems, required to meet acceptable power ceilings for exascale machines, will drive the mass market for mobile devices.

19. The EC Communication calls for investment in HPC by Member States, the EU and Industry to double to around EUR 1.2 billion pa, with roughly half of the additional EUR 600M for hardware procurements, a quarter for training and a quarter for software development. It would like to see 10% of the public sector hardware budget used for PCP to *“develop and maintain native EU supply capabilities that cover the whole technology spectrum from processor architectures to applications”*. Such PCP actions could be eligible for EU co-funding if they are directed towards this EU-level mission. The total EC expenditure on HPC during the seven years of Horizon2020 is expected to be around EUR 600M.
20. The EC’s proposed Action Plan for EU HPC renewal should be governed through an industry-led technology platform (industry has responded with the ETP4HPC proposal (see 43)) and, for science, through PRACE, linked together by centres of excellence for the deployment and application of HPC software and services.
21. PRACE AISBL was established as Belgian company in April 2010, following a UK-led initiative over six years, beginning with a meeting with the German Wissenschaftsrat in March 2004, following which the *“HPC in Europe Taskforce”* organised a community-wide effort to write the Scientific Case⁷ and HPC was included on the 2006 ESFRI Roadmap⁸. The original rationale behind PRACE, that there should exist in Europe a small number of leadership HPC systems of a scale that no single Member State could reasonably afford and that these systems should be made available to computational scientists from around the EU through a peer-review process that focuses on the quality of the science delivered through the capability that the HPC systems provide, has become diluted with the adoption of an in-kind contribution model dominated by Germany and France. So the current PRACE mission statement says it is *“to enable high impact European scientific discovery and engineering research and development across all disciplines to enhance European competitiveness for the benefit of society”*.
22. PRACE and ETP4HPC are proposing to play complementary and coordinated roles: PRACE will continue to provide HPC services to support the needs of scientific and industrial applications, while ETP4HPC will identify research priorities for the development of HPC technologies. PRACE will have complete responsibility for Europe’s high-end, tier-0, computing infrastructure and its use, identifying priorities of the societal and scientific grand challenges that could benefit from future HPC technologies. ETP4HPC will provide PRACE with the point of view of industrial technology providers on topics relevant to this infrastructure and identify priorities for development of HPC technologies leading to a competitive and sustainable HPC vendor industry in Europe. The two organisations will work together in a coordinated manner, with clearly separated roles, to deliver the vision in the EC Communication. Although PRACE will pilot PCP for HPC (see 35), the ETP4HPC will play the leading role in its implementation.

⁵ http://www.exascale.org/iesp/Main_Page

⁶ <http://www.eesi-project.eu/pages/menu/homepage.php>

⁷ <http://www.hpcineuropetaskforce.eu/>

⁸ http://ec.europa.eu/research/infrastructures/index_en.cfm?pg=esfri-roadmap§ion=roadmap-2006

23. The EC Communication, while praising the efforts of PRACE, states that, “*Member States are encouraged to set up HPC competence centres that facilitate access of industry and specifically SMEs to HPC services, and should support supercomputing centres to transfer expertise to them*”, clearly understanding that PRACE is not the only solution to all HPC challenges and aligning to the emergent UK model for HPC Centres of Excellence. While PRACE can provide the high-end research capabilities, which are out of reach of individual nations, broad industrial engagement at this international level can only be achieved by building a “bridge” of industry support structures at both national and regional levels.
24. The UK position as a General Partner within PRACE is widely regarded as anomalous. UK HPC infrastructure is comparable to that of the Hosting Partners (France, Germany, Italy and Spain, who have each committed to provide EUR 100M of petascale computing cycles in the period 2010-14) and both our HECTOR⁹ and DiRAC¹⁰ BlueGene/Q machines could easily qualify as tier-0 systems (they are, respectively, 32nd and 20th fastest in the world²⁶, and the STFC Blue Joule system is 13th). In fact, the UK only contributes 5% of HECTOR to the shared Distributed European Computing Initiative (DECI, formerly DEISA), which makes up the tier-1 level of the PRACE infrastructure (there are plans to contribute a similar fraction of Blue Joule). In contrast, UK scientists have won 8.6% of the tier-0 resources, although the UK success rate of 39% is lower than the average of 56%. Of the proposals which meet PRACE’s quality standard, not all of which are awarded resources because the demand exceeds what is available, the UK share is 12.3%. Notably, the largest single allocation, 144M core hours, was to a UK climatology project. The imbalance between the UK contribution to PRACE and resources won is the main reason the Hosting Partners are calling for a juste-retour mechanism to be introduced. PRACE is currently engaged in a strategic review to establish a sustainable financial and governance structure beyond 2014. It aims to agree a plan for the next phase of PRACE by mid-2013.
25. Responsibility for the activities chosen to implement the EC Communication within the EC lies with the Research Infrastructures Unit, which is aligned with the “Excellence in Science Base” pillar of Horizon2020. It is not clear if this is the correct place for this responsibility to lie, as the majority of industrial funding for computing systems research lies within the “Industrial Leadership and Competitive Frameworks” pillar.

Recommendation 1: Significant opportunities exist to enhance and share the associated cost of e-infrastructure and its exploitation for UK research and innovation through the EC Action Plan in its Communication on HPC and the Horizon2020 programme, and the UK should make it a priority to engage coherently with these.

Recommendation 2: To ensure that the most computationally demanding UK research has access to the next generation exascale platforms, it is essential that the UK engages with PRACE and sets a timetable for achieving full partnership by mid-2013.

Recommendation 3: The UK should play an active role over the next year in transforming PRACE so that it can sustainably support, via an equitable contribution model, world-leading open research

⁹ <http://www.hector.ac.uk/>

¹⁰ <http://www.stfc.ac.uk/Our+Research/24711.aspx>

by academia and industry through the provision of a level of computing infrastructure and access which is beyond the capabilities of any one European nation.

The Data Deluge

26. The importance, complexity and size of the data sets used in academic and industrial research, and commerce, are growing at an enormous rate, due to ubiquitous sensors in our environment, large data-producing instruments in scientific research and large-scale computer simulations. Our e-infrastructure capability to support data storage and movement has not been developing as rapidly as the capabilities of the instruments and computers creating data. Addressing this gap is urgent and will be expensive. It requires coordinated action at local, regional, national and international levels, because researchers typically need end-to-end data management capabilities, extending from international facilities to their desktop. Also, the cost of data movement and storage has become non-negligible at all levels in the infrastructure, extending from the power cost of moving data on a chip to the cost of multi-petabyte archive systems.
27. The “one size fits all” approach to the provision of high-end computers that has satisfied researchers in many disciplines to a large extent, and permitted providers to adopt a “technology push” strategy, will not work cost-effectively for the provision of data services, because of the diversity of user requirements. Instead, we need to develop an “applications pull” strategy, and this will require partnership between users and providers, and community data solutions. Some research communities are more organised than others in this respect, but where a community doesn’t have a strategy for storing its large data, it will be necessary for communities with common requirements to form, develop and adopt data management strategies.
28. This vertical organisation must be embedded in horizontal organisation of e-infrastructure provision, spanning computers, networks and data storage. This coordination of the e-infrastructure also spans all of the tiers in the HPC pyramid, since large-scale simulation and modelling is often needed to interpret data, or to generate it in the first place (tier 0), pre- and post-processing of data often can be performed most efficiently on large cluster systems (tier 1 or 2), and visualisation must be performed locally (tier 2). Where there is an immediate need for data interpretation, e.g., for medical diagnosis and emergency response, large data sets must be transmitted rapidly to an appropriate HPC resource, processed and results shipped back in a guaranteed turnaround time. This will require radical changes in operating modes for HPC and these changes are essential if the full benefit of e-infrastructure is to be realised by society.
29. Several initiatives at European level have been driven by applications needs. Most prominent has been the LHC Grid, required by CERN to post-process LHC data internationally, that sparked grid and e-science initiatives world-wide. The resulting grid infrastructure in Europe is organised through the European Grid Infrastructure (EGI)¹¹, which spans more than 30 countries. More recently, following prioritisation by ESFRI, the EC has funded a preparatory phase for the European Life sciences Infrastructure for Biological Information (ELIXIR)¹² to construct and operate a sustainable infrastructure for biological information in Europe to support life science research and its translation to medicine and the environment, the bio-industries and society. PRACE is undertaking an update of its Scientific Case and the need for integrated data services is

¹¹ <http://www.egi.eu/>

¹² <http://www.elixir-europe.org/>

emerging as a top priority, particularly from medicine and climate science. In a top-down response, the EC has funded EUDAT¹³ to provide a pan-European solution to the challenge of data proliferation in Europe's scientific and research communities, which is making progress towards an established pan-European data infrastructure, as well as working to build ties with PRACE and EGI.

30. Coordination and an overall strategy for data services across Europe are lacking at this stage. It is not clear whether PRACE, EUDAT, or some other organisation will take responsibility. The UK must engage with these developments to ensure that they complement our national data services, and that mechanisms are put in place to allow UK data services to interoperate with community based solutions and pan-European projects, such as EUDAT. The diversity of the requirements, the complexity of the issues, technically, operationally and organisationally, and the urgency to respond to growing research needs demands that the UK strategy for participation in HPC at European level includes support for data-intensive applications.

Recommendation 4: The UK should take an integrated approach to e-infrastructure for data-intensive and compute-intensive applications and engage immediately with work to develop this at the European level, to ensure that the UK achieves full integration with, and maximum added value from, the imminent establishment of European data services.

Industry Use of HPC

31. While universities and academic researchers are generally experienced in drawing down EC funding and exploiting e-infrastructure at all levels, UK industry generally is not. Support structures need to be created to facilitate industry engagement. SMEs in particular require either regional support structures, such as clusters around centres of expertise (computer centre or university), or to be part of the supply chain of a large company which defines their HPC strategies.
32. While many publicly funded HPC centres talk about industry usage of their HPC systems, there is much less of this across Europe in reality than might be expected. Germany and the UK have HPC centres that excel at this and the French have been gradually growing their activities. The Italians have limited industrial use of systems at their main centre, as do the Spanish. The UK regions have specific access support programmes for SMEs (e.g., HPC Wales¹⁴, Supercomputing Scotland¹⁵), as do the Germans in the Stuttgart area through HLRS. The French have recently established the HPC-PME¹⁶ initiative through their involvement in PRACE. There is an understanding at the EC that this is an under-developed area across Europe.
33. The Germans see HPC as a talisman of German excellence in technological design and manufacturing. The German government spends more than three times the UK expenditure on HPC equipment for its major universities and sees the provision of such equipment as a key enabler of German industrial strength through collaboration with its HPC centres and universities. This comparison is for HPC equipment alone, not the entire e-infrastructure including, e.g. networking, for which UK expenditure is around £200M pa.

¹³ <http://www.eudat.eu/>

¹⁴ <http://www.hpcwales.co.uk/>

¹⁵ <http://www.supercomputingscotland.org/>

¹⁶ <http://www.initiative-hpc-pme.org/>

34. It is not clear, at this point, how well PRACE can engage with industry and particularly with SMEs. PRACE today is essentially a creation of the academic sector and this is reflected in the Calls for Access, wherein the focus is on academic or large-scale projects using highly scalable codes. It is highly optimistic to believe that an SME, which is attempting to transition beyond high-powered workstations, or attempting to understand how HPC will be useful to its business, will be able to exploit the most powerful supercomputers in the world in a single step, or even speak the same language. However, PRACE is currently going through an attempt to re-invent itself to achieve sustainability beyond 2014 and respond to the EC Communication. Although it is not clear how successful this will be, it is clear that a significant re-invention will be required to reach down to the start-ups and SMEs.
35. A major barrier to industry use of the more powerful systems in the e-infrastructure is the need for their applications codes to exploit high degrees of parallelism. This is increasingly true also at the multicore chip level. Encouragement and support is needed for ISVs to scale their codes to make effective use of tier-0 systems, so that applications and tools are available to enable businesses to exploit higher levels of the e-infrastructure effectively.
36. The recognition across Europe that HPC drives industrial competitiveness and hence economic growth justifies increased efforts by BIS, TSB and regional development agencies to grow the number of SMEs and heavy industry (e.g., chemicals, manufacturing) starting to use HPC for simulation and modelling. This is a leadership and driving role. It requires a communications process to explain and demonstrate to companies what simulation and modelling can do for their businesses, and hand-holding from entry to working knowledge of simulation and modelling, enhancing their ability to apply and communicate their new product solutions. In the short term, this should build on initiatives such as Supercomputing Scotland and HPC Wales; in the long term, sustainability will depend on increasing the supply of suitably trained graduates.
37. Support for HPC as a tool for UK larger industries might be provided by a central group within BIS with the mandate to enable UK large industry to benefit from product simulation and modelling. Their approach would be, say, to list the top 25 UK companies (e.g. in manufacturing, drug creation) as driven by their volume of export and/or IP creation, in order to define a set of industry verticals. The representatives for each chosen industry vertical within BIS would meet regularly and each would also be in regular contact with the nominated companies in their vertical. They would act as the point of communication to raise awareness of each vertical's ability to use simulation and modelling, including matchmaking to regional support establishments that can provide advice, training, and an appropriate level of machine access, and to communicate opportunities for strategic new or improved products.
38. The same BIS group should periodically conduct a gap analysis to identify which industry verticals, or volume products within verticals, are not represented and which skill sets need to be added to the regional establishments. It is expected that this will require investment by industry, as well as start-up investment by the regional establishments. The support structure should provide UK companies with a path towards, and an on-ramp for the ETP4HPC.

Recommendation 5: The UK should establish support structures on a regional basis to provide bridges for SMEs to engage effectively with the opportunities provided through EC funding and shared European infrastructures, and should encourage the EC to establish mechanisms for sharing best practice across Europe.

The HPC Technology Supply Chain and ETP4HPC

39. The worldwide HPC market is forecast¹⁷ to be EUR 17.3 billion in 2013, with a CAGR of 7.5%, the European HPC market representing 25% of the worldwide total or EUR 4.6 billion.
40. Within the HPC supply sector, the UK is recognised as having a relatively strong track record for innovation. Previously, companies like Inmos, Clearspeed and Quadrics provided world-class components for HPC at the highest levels. More recently, ARM, Xyratex and Gnodal are developing products that have the opportunity for innovation around exascale, e.g., the Mont Blanc Project¹⁸ co-ordinated by the Barcelona Supercomputing Centre (BSC) in partnership with ARM, Gnodal and Bull. This EU funded project has the objective of designing and developing a next-generation energy-efficient HPC prototype system. While initiatives in other member states have sought to coordinate the HPC supply sector, this has not occurred in the UK. These technologies, algorithms and applications, which are developed for exascale over the coming decade, will inevitably see more general adoption in the ICT sector, particularly in other large scale e-commerce systems such as those that drive Paypal, eBay, Google, Amazon etc.
41. France has been lobbying the EC, via the highest levels of Government, for the past few years to use EC funding to establish and grow a European HPC industry. This would obviously benefit their indigenous HPC company Bull, but their rationale for this is deeper than simple naked self-interest. The French Defence Agency believes that it is a matter of national security that France (and by implication Europe) should have an indigenous ability to develop and build computer systems using components designed and manufactured in Europe. They look to the USA, China and Russia and see their governments supporting HPC companies and growing indigenous computing industries, and believe that Europe must do the same.
42. The PCP process designed by the EC is very similar to the UK SBRI process (indeed their documentation references SBRI). The overall PCP model has been created in order to tackle a difference between how Europe and the USA benefit from their basic research expenditure. Although the EU and US have similar levels of spend on basic and applied R&D (2% of GDP in Europe versus 2.5% of GDP in the USA), the USA has, over the past 30 years, followed a public procurement policy that encourages early procurement by the public sector of new innovative products prior to their commercial release. This has led to a great difference between public procurement of R&D results by Europe and the USA (EUR 2.5 billion in Europe versus EUR 50 billion in the USA annually). This twenty times difference is felt by the EC to be a major contributing factor in Europe's inability to build large innovative companies, particularly in the IT sector. The French, in particular, have been pushing for PCP to be used in Horizon2020 to strengthen the EU HPC supply base. A EUR 10M PCP pilot is proposed in PRACE-3IP, the 3rd PRACE Implementation Project which starts in Jul 2012, such that a larger activity can be implemented in Horizon2020, probably at the level of around EUR 60-80M per year for the seven year duration. The UK will be engaged in this pilot activity and funding for the UK contribution of EUR 700k is being sought from TSB with the support of EPSRC. In the broader PCP context, the UK must assess if this is a sensible use of public funding and, if the EC decides to invest heavily in PCP in Horizon2020, how UK companies can benefit from what may be significant tranches of funding over the next seven years.

¹⁷ A strategic Agenda for European leadership in Supercomputing HPC2020, IDC#SR03S

¹⁸ <http://montblanc-project.eu/>

43. The proposed European Technology Platform for HPC, ETP4HPC³, is an industry led initiative formed to define European research priorities in all segments of the HPC value chain to increase the value created in Europe from future HPC systems and solutions. In addition to the industrial partners, Allinea, ARM, Bull, Caps Entreprise, Eurotech, IBM, Intel, ParTec, STMicroelectronics and Xyratex, six supercomputing centres from Spain (BSC), Germany (Fraunhofer, FZJ, LRZ), Italy (CINECA) and France (CEA), are also founding members. EPCC in the UK plans to join. The ETP4HPC vision document¹⁹ recommends prioritising research in a number of areas of particular relevance for existing UK expertise (e.g., ARM, Allinea and Xyratex):

- programming paradigms;
- energy driven HPC;
- architecture evolution;
- scaling I/O and storage with processing;
- new storage solutions for big data.

Without strong engagement by the UK, there is a risk that the ETP4HPC is dominated by interests that are not aligned well with those of UK companies. The TSB could play a significant support role to encourage UK company engagement.

44. Within the EC, this industry agenda is being driven inappropriately through Research Infrastructures and, hence, into the “Excellence in the Science Base” pillar of Horizon2020, rather than the more appropriate ICT area. This risks jeopardising Europe’s international excellence in the application of HPC, which is dependent on access to the most powerful machines wherever they are developed, by seeking to employ some of the funding for these to support European hardware providers, which might result in the procurement of less competitive systems in the short term.

45. There is an opportunity for the UK to position its companies to maximise the benefit from EC actions, such as PCP, to grow the European HPC supply chain. Achieving a truly competitive HPC industry in Europe will require billions of euros consistently invested over many years. So the entry cost may be prohibitive. The alternative, particularly for the UK, is being a designer of components or IP in that space, following the example of companies like ARM.

Recommendation 6: To ensure that the UK is positioned to derive increased value from its domestic HPC supply chain, the UK should establish a support structure for companies that could contribute to the HPC technology supply chain, to pool knowledge and expertise and broaden UK engagement with ETP4HPC.

Recommendation 7: The UK should clarify its position on Pre-Commercial Procurement, identify technology developments where UK companies have, or could acquire the capabilities to contribute and, through active engagement with ETP4HPC, encourage the EC and Member States to direct their PCP spend towards them.

¹⁹ <http://www.etp4hpc.eu/documents/ETP-key-documents.html>

Other EC e-Infrastructure-Related Activities

46. In the EC ICT programme, the UK has played a much greater role in setting strategy than PRACE through the PlanetHPC²⁰ Support Action led by EPCC and engagement by many UK companies and universities in the HiPEAC²¹ Coordination Action. However, these activities have been largely bottom-up and do not operate at governmental level.
47. Several FP7 initiatives relate to HPC and the emerging e-infrastructure, but are not well connected even in Brussels. Future Emerging Technology (FET) Flagships are at the proposal development stage. The two major research data-management initiatives, EUDAT and ELIXIR, currently sit in Research Infrastructures. These will be brought together under “Excellence in the Science Base” in Horizon2020, creating the opportunity for a much more coherent development. This demands that the UK also brings its strategies for compute and data together.
48. GÉANT²² is the high-speed European communication network dedicated to research and education. In combination with its National Research and Education Network (NREN) partners, GÉANT creates a secure, high-speed research infrastructure that serves 40 million researchers in over 8,000 institutions across 40 European countries. Co-funded by FP7, GÉANT is the e-infrastructure at the heart of the EU’s European Research Area. The project partners are 32 European NRENs, TERENA and DANTE. GÉANT is operated by DANTE on behalf of Europe’s NRENs. The UK is a key member of GÉANT and, in fact, hosts DANTE. GÉANT recently announced migration to the latest transmission and switching technology designed to support up to 2Tbps (terabits per second) capacity across the core network, effectively future-proofing Europe’s critical network up until 2020. 500Gbps capacity will be available across the core network from first implementation, delivering circuits across Europe that will allow individual users to transfer data at speeds of up to 100Gbps, thereby enabling faster collaboration on critical projects and meeting the rapidly increasing demand for data transfer. In the UK, companies who are engaging with UK universities can transfer data to and from their collaborators on JANET. This is not always the case on GÉANT, or other NRENs (some of which don't allow company data on at all).

UK Leadership

49. With the advent of Horizon 2020 in two years’ time, there is an opportunity for the UK to decide what its European HPC policy should be. Given the UK strengths are in applications, the main objective driver should be to maximise our effective use of HPC technology by academia and industry. This will only be achieved through a proper dialogue with industry and academia, and to be successful in Europe a top-down and a bottom-up approach will be crucial for success.
50. The rapid evolution of the enabling technologies and the expanding range of applications demand a user-centric approach to strategic governance, as recommended in the 2011 e-IRG White Paper²³. This will involve a shift away from provider-driven technology push towards user-driven applications pull. It will be a challenge to engage the user communities and develop a balanced partnership between them and the providers.
51. The process of annual work programme development by the EC is remarkably transparent for such a complex funding programme. The easiest way for a company or university research group

²⁰ <http://www.planethpc.eu/>

²¹ <http://www.hipeac.net/>

²² <http://www.geant.net/pages/home.aspx>

²³ <http://www.e-irg.eu/publications/white-papers.html>

to become engaged in a particular area of programme activity is to attend the regular “Information Days” which happen before each call for proposals. The EC, generally through projects called Coordination & Support Actions, also holds regular collaboration meetings for each programme activity. Attending these meetings, taking part in discussions and helping to write documents and reports gives those who engage the chance to influence the text of the next work programme and its corresponding call for proposals. Project Officers are generally very approachable and keen to discuss ideas face-to-face and provide general guidance on what they are looking for in their particular area. It is also common for those involved in funded projects to be invited onto an “Expert Panel” which advises the particular Unit or Directorate on future directions for their funding. These activities are very bottom-up, but those who engage in them find the process both intellectually and financially rewarding (via subsequent project proposals). A number of companies and institutions in the UK are good at taking this approach.

52. The top-down approach is more problematic and something the UK has been poor at until recently, while some other Member States have been much better organised. Over the years, as Framework Programmes and their various funding rules and research objectives have been established, Member States are given the opportunity to comment on the plans. Until recently this task was under-resourced and not seen as a priority by BIS. It is clear that recently this process has changed and much more work is being done in this area by the BIS International Knowledge and Innovation Unit. However, commenting on plans is a secondary position – it is much better to promote your own country’s position so that it shapes others’ plans. Many Member States form clear, agreed, country-wide positions on issues, and lobby the EC and other Member States to see their point of view. Germany is a master of this approach (France is to a lesser extent). The UK is not. This has led to too much high-level policy being set by countries such as Germany and France, particularly in the area of HPC, with the UK’s voice not being heard at all.
53. One of the key problems in European HPC is that PRACE was supposed to bring together the funding bodies that support national HPC. The UK therefore engaged in PRACE through EPSRC. While this was the right thing to do, most countries have delegated their role in PRACE at all levels (from the PRACE Council downwards) to their HPC Centres. This means there has been no proper dialogue at an intergovernmental level with regard to PRACE (hence some of the problems it faces today). This has left the door open for countries such as France and Germany to lobby the EC with regard to HPC with very little discussion (either support or opposition) from the UK.
54. UK leadership should sit within an organisation whose core responsibility is for completing the symbiotic circle of research, industry and government funding streams for the benefit of research competitiveness, national GDP, and job maintenance and creation. Currently, the obvious choice is BIS. Leadership should be by an executive team, an “e-Infrastructure Directorate (e-ID)” or the recently-formed “Programme Board”, with the core function of being able to reach into domestic and international business and research via co-workers, ideally within the same department. While a major target must be to influence Horizon2020 funding streams to the mutual benefit of UK industry and academia, the top-down approach, it must also connect effectively with the UK stakeholders who are engaged in EC projects and who play an equally important role in influencing the EC from the bottom up. It should be specifically tasked with creating a joined-up approach to Europe, sharing information gathered by those at the coal-face and presenting a consistent UK strategy.

55. The mandate of the e-ID should recognise that industry and academia are fellow travellers, i.e., have overlapping, though not always identical, objectives and mutually stand to benefit. Its objective should be to position the UK as one of the leading European countries in terms of its HPC strategy, as measured by political, technical and fund-creating influence. Specifically, the UK should to be seen by the EC DG CONNECT²⁴ staff as one of the first to run new ideas by for inputs, and as an example of how to utilise programmes and affect growth. UK staff should be seconded to DG CONNECT at a strategic level.

Recommendation 8: The UK should establish an executive arm of the e-Infrastructure Leadership Council, an "e-Infrastructure Directorate", that is user-centric and tasked with achieving more coherent engagement between HPC at the European level and within the UK, connecting both top-down and bottom-up channels of communication with the EC.

The Wider International Position

56. An EESI report has recently surveyed HPC initiatives world-wide²⁵. Following significant investment over the last few years, the US is still dominant, with seven systems performing at the petascale and two 20 PFlop/s systems anticipated by the end 2012, one of which, Sequoia at Lawrence Livermore National Laboratory, has just reached No. 1 on the TOP500 list²⁶ with 16.32 PFlop/s performance. To date, the US has led much of the discussion around exascale computing, investing in the IESP and funding programmes such as DARPA's Ubiquitous High Performance Computing program. President Obama's administration listed exascale computing as one of the Grand Challenges of the 21st Century in its "Strategy for American Innovation". His 2012 budget included funding for the development of an exascale system, providing the DoE with \$126M.
57. Asia has been investing heavily in supercomputing technology and now hosts two of the five fastest computers. Dominant countries in Asia are Japan and China. In 2011, Japan reclaimed top ranking with its Next Generation Supercomputer (NGS) or "K Computer", installed at the RIKEN Advanced Institute for Computational Science (AICS) in Kobe, that is now No. 2 on TOP500 list with 10.51 PFlop/s performance. Five strategic fields have been selected to exploit the system: life sciences and drug design; new materials and energy; global climate change prediction for disaster prevention/mitigation; manufacturing technologies; and the origins of matter and the universe. Japan has set up a PRACE-like entity, the "Innovative HPC Infrastructure (HPCI)" including the NGS and other supercomputing sites, to establish the hierarchical organization of the NGS with other supercomputers, to set up a large-scale storage system, and to establish an HPC consortium which will run the HPCI. Japan has funded a two-year project to develop plans for exascale. It has also increased efforts to encourage industry use of HPC.
58. China, whose Tianhe-1A supercomputer was ranked No. 1 in the TOP500 Nov 2010 edition, remains the fastest rising country in the region and even in the world, with an increasing number of systems listed in the TOP500. China has invested heavily in supercomputing technology and Tianhe-1A now holds fifth position in the TOP500. The Chinese government has actively promoted independent innovation to avoid reliance on foreign technology. Although China's

²⁴ From 1 July 2012 DG-INFSo will be renamed DG-CONNECT

²⁵ <http://www.eesi-project.eu/pages/menu/publications/investigation-of-hpc-initiatives.php>

²⁶ <http://www.top500.org/list/2012/06/100>

recent top-end systems are impressive, it is notable that the peak performance/sustained performance ratio is high, leading to questions over how applicable these systems are.

59. The recent G8 Research Councils' Initiative on exascale applications for global issues provides funding to enable collaboration between US and European partners and focuses on the applications of exascale, rather than hardware. Most exascale initiatives across the continents identify a similar set of global issues as scientific drivers for exascale computing. These include climate change, sustainable energy sources and nanotechnology.
60. Russia has recently increased its spending on HPC and has a desire to become a major player in the HPC market. The "Lomonosov" supercomputer located in Moscow State University, at 0.90 PFlop/s, is the fastest supercomputer in Russia and the 7th in Europe, ranking 22nd in the TOP500 list of Jun 2012. It was manufactured by the leading Russian HPC vendor T-Platforms, a rapidly growing private company established in 2002 to provide supercomputing systems, software and services in Russia. T-Platforms has installed almost 200 supercomputer systems and is currently expanding its activities into Europe and Asia. Collaboration between the EU and Russia was strengthened by the launch of the first joint EU-Russia ICT call in 2009, focusing on programming models and runtime support, performance analysis tools for HPC, and optimisation, scalability and porting of codes. The call includes three stages of two years supported for a total of EUR 6M (EUR 4M EU + EUR 2M Russia). Two projects have been funded under this call, the Application Performance Optimisation and Scalability (APOS) project to develop optimised versions of scientific and industrial codes, which are scalable and portable across heterogeneous and homogeneous architectures, and the HOlistic Performance System Analysis (HOPSA) project, looking specifically at the interplay between system and application performance.

Training

61. In the long term, the ability of the UK to exploit e-infrastructure at all levels to deliver world-class research, and economic and societal benefits will depend on the availability of graduates with appropriate mathematical and computing skills. The strategy presented here for UK participation in HPC at the European level is dependent on the most advanced of these skills and the relevant training will typically be provided at the postgraduate level, although this should be seen as one part of a much broader skills requirement for the UK to remain competitive.
62. PRACE has established six Advanced Training Centres across Europe⁴, including EPCC in the UK. These are developing a standard curriculum for advanced computing. PRACE also runs a series of summer schools. The UK can lever these training opportunities and the material developed for them. It should develop a strategy for postgraduate training across all disciplines that provides relevant skills in using e-infrastructure.

Recommendation 9: The UK should introduce training in the use of HPC technologies employed at all levels in the ecosystem for all postgraduate research students in relevant fields and, particularly, it should participate in the development and provision of training on tier-0 facilities on a Europe-wide basis, e.g., through the PRACE Advance Training Centre programmes

Richard Kenway, Peter Coveney, Sean McGuire, Oz Parchment, Mark Parsons