

## EPSRC strategy for the developing landscape of Tier-2 HPC in the UK

### Introduction

#### High Performance Computing (HPC) landscape in the UK

The HPC ecosystem in the UK consists of Tier-1, Tier-2 and Tier-3 (local university resources). The Tier-1 resources, which provide the current highest capability research HPC machines, are ARCHER and DiRAC. These machines provide the capacity to deliver leading edge science and the capability to develop next generation codes. UK researchers can also access Tier-0 resources through PRACE. Along with the Tier-0 and Tier-1 resources, the UK computational science community requires training machines, entry-level capacity machines such as those in the Universities and at Tier-2 level, and both experimental and customised machines. The diagram below shows the HPC landscape in the UK.



Figure 1: The UK HPC ecosystem

### Tier-2 HPC

Computer-supported modelling and simulation is now widely recognised as the third `leg` of Scientific method, alongside theory and experimentation and the EPSRC and partner Research Councils have been committed to supporting this through our investment in HPC. A globally competitive supercomputer and local university clusters have been available to researchers in the UK for many years. However, the gulf in capability from a local university system to the National Service (ARCHER) has acted as a barrier for some researchers and

communities to access and realise the full benefits of HPC. This capability gap was recognised and in 2011-12, EPSRC invested a total of £8 million capital in the establishment of 5 regional Tier-2 Computing Centres of Excellence with matching funding from host institutions and collaborators. The established Tier-2 centres are the level of provision above the local HPC cluster and below the UK National Facility.

The scientific impact of the Tier-2 facilities has already been significant, with each facility featuring in an average of more than 100 grant applications totalling nearly £50m and its use acknowledged in at least 280 publications. The remit covered by these facilities span a wide range of scientific disciplines from life sciences to earth sciences and engineering.

The Tier-2(T2) level of HPC provision is fundamental because it provides a diversity of computing architectures, which are driven by science needs and are not met by the national facilities or universities. This is because the National HPC Service must meet the needs of the whole UK community and so cannot specialise in specific novel architectures or novel requirements. The different types of computing requirements provided by Tier-2 include high-throughput and GPU computing and are supported by local expertise in these areas.

The T2 layer also provides easy access through a light touch review process for users, a rapid turnaround in experiments and allows small compute runs to be carried out which are not suitable on a national machine. It also enables researchers to carry out proof of concept studies which can in turn help users to migrate to the National Facility ARCHER.

The Tier-2 centres are available to multidisciplinary researchers from across the UK. An important feature of the T2 level is that it provides local access, training and support for users. Therefore, the centres have the ability to broaden access to researchers new to HPC for example certain science communities or universities that do not have access to local resources.

After the EPSRC's initial investment in this layer of HPC in 2012, this infrastructure now needs to be refreshed and reinvigorated to ensure that it continues to be fit-for-purpose in the rapidly changing e-infrastructure landscape over the next 4-5 years. Tier-2 HPC forms a vital part of an integrated e-infrastructure landscape. The continuation of a sustainable and cutting edge e-infrastructure ecosystem is vital in allowing EPSRC to deliver its Strategic Goals and support excellent and innovative science and engineering research.

A coherent strategy for developing and delivering the UK's next Tier-2 HPC ecosystem is essential in driving forward advancements in computational science in the UK because it provides the diversity of architectures and easy access routes needed to do a diverse range of science. Tier-2 HPC is an essential component of an integrated e-infrastructure landscape, and is vital for maximising the benefits obtained from both

national and local facilities. It enables integration both vertically from Tier-3 to Tier-1 and horizontally, between other Tier-2 centres.

## **Vision**

Our vision is to work towards a diverse and flexible UK Tier-2 HPC provision with the capability and capacity to meet the scientific needs of the academic and industrial computational science community. The diverse computational science needs and user requirements, including the range of compute job shapes and sizes, needs to be met. In addition, we have a goal to enable researchers to future-proof their research in relation to new and emerging hardware architectures for the future.

Our goal is for T2 to encourage and widen participation from new users and those from different disciplines to access the service. We recognise that this cannot be achieved without building up support and expertise in software engineering skills and training for computational science.

We will work towards a Tier-2 layer that is integrated with the HPC ecosystem across the UK, both vertically (into tier 1 and tier 3) and horizontally across to other tier 2 hubs.

Furthermore, the T2 layer should provide comprehensive national access for researchers within the UK which is driven by scientific need and location of expertise and is not defined by region.

The evolving Tier-2 ecosystem will be enhanced by the sharing of expertise and collaboration between centres.

## **Key Features to achieve the vision**

- A landscape of diverse hardware architectures throughout the UK. This is in recognition of the need for a diverse range of hardware architectures to support the different requirements of the computational science community. These different architectures will also help to shape future Tier-1 systems for the UK.
- Encouraging skills and expertise in software engineering to support the range of T2 architectures which will be available in the UK.
- Investing in a number of individual Tier-2 hubs, which are located in the most appropriate area for a particular technical expertise or scientific need; this could be hosted by a consortium e.g. technical or a single institution.
- Encouraging easy access mechanisms to the whole Tier-2 layer, which allow access to researchers, both local and UK wide, and which encourages access by new users and disciplines.
- Encouraging industrial collaboration and use of T2 HPC where appropriate.

- Investing in Tier-2 Hubs that enable multidisciplinary computational science but with the centre of gravity in the engineering and physical sciences remit.
- Encouraging a community within the Tier-2 ecosystem that shares expertise.
- Encouraging future proofing and sustainability of codes by enabling access to new hardware architectures.
- Encouraging a Tier-2 layer that is integrated with the HPC ecosystem across the UK, both vertically (into Tier 1 and Tier 3) and horizontally across to other Tier 2 hubs.
- Encouraging sustainability of an agile, flexible and responsive Tier-2 provision which is able to meet the evolving needs of the computational science community.