

HIGH END COMPUTING MANAGEMENT PLAN

Issue

1. NERC invests at least £3M per year in the provision of high end computing resources to the environmental science community. This document provides a framework for investment decisions and the management of the HEC programme.

Mission

2. To ensure that NERC's provision of high end computing supports the delivery of science priorities set out in the NERC strategy, '*Next generation science for planet Earth*'.

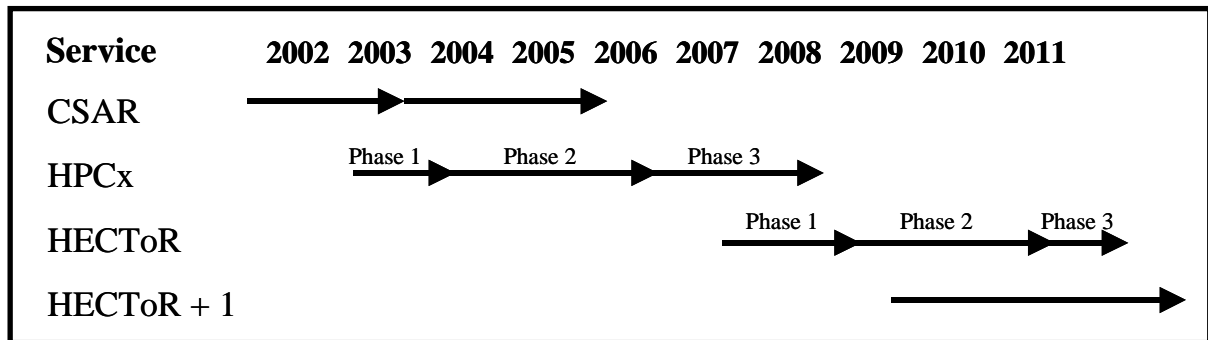
Fate of this Management Plan

3. The plan will be considered by NERC's Science and Innovation Strategy Board and Executive Board. Ultimately, it will be owned and used by the NERC HEC Manager to guide decision-making.

Introduction

4. In this document, High End Computing is defined as computing facilities above and beyond those generally provided as desktop PCs. This includes national services (also referred to as High Performance Computing, HPC, or supercomputers), mid-range machines (clusters), access to Computational Science and Engineering (CSE) support, development of model codes, data, visualisation and informatics. Technology in all areas of computing is developing very rapidly. The HEC community is aware of the emerging grid infrastructure, including the European Grid Initiative, but currently this does not satisfy their requirements.
5. HEC is a vital tool for environmental science and its use supports world-leading research carried out by NERC-funded scientists. A requirement for High Performance Computing was highlighted in the *Science and Innovation Investment Framework 2004-2014*, which states '*The UK is a world leader in Earth System Science modelling. Maintaining this leadership edge will require: high performance computing and data-handling facilities far in excess of those currently available*'
6. Use of HEC underpins research in a number of areas of NERC science. NERC users of national HPC services are members of one of several consortia, based on their science area and the computer models they use. Currently, these consortia are:
 - Atmospheric and Polar Science
 - Oceanography
 - Mineral Physics
 - Continental Shelf Modelling
7. EPSRC are managing agents on behalf of the Research Councils for the national services. They lead in the procurement and management of services. NERC secures access to national HPC services for its users by arrangement with EPSRC.
8. At the beginning of 2007, there is currently one national service, HPCx, run by the University of Edinburgh and CCLRC (Daresbury). The next national service, called HECToR, is being procured and is due to be operational in October 2007. The current strategy is that each service runs for six years, with a new service being procured every three years, as shown below. The cost (HPCx £55M, HECToR ~£120M) and long-term nature of these procurements makes it difficult for NERC to switch quickly from one type of provision to another, eg we have committed funds to the HECToR

service for six years. There is a tension between the long-term procurement of national HPC services and the demand for HEC resource from the community. Through a rolling planning process, NERC must be sensitive to the science strategy needs of the community and ensure the whole community is made aware of the opportunities that are available. The environmental science community have an obligation to engage in this process to ensure the provision of HEC is fit for purpose.



Operational timetable of national HPC services

Context

9. This plan is related to a number of other strategies.

- NERC’s new science strategy, *Next generation science for planet Earth*. The cross-cutting science theme ‘technology’ and the non-science theme ‘science infrastructure’ are closely related to HEC
- A Strategic Framework for High End Computing, written by the High End Computing Strategy Committee, a body convened by EPSRC with representatives from the UK academic community
- NERC technology strategy. Aims to provide guidance to NERC on the management of technology and technology development investments and to inform internal and external stakeholders of NERC’s technology requirements and priorities.
- NERC eScience strategy. Proposes a number of key principles to support integrating e-Science in to NERC activities.
- NERC science information strategy. Aims to ensure that internal and external demands for structured data and information can be met.
- NERC information strategy. Ensures NERC meets legal requirements for data and information availability and accessibility, encourages cross-disciplinary working, seeks to improve quality assurance and control of data and protect NERC’s intellectual property rights.
- NERC information technology strategy.
- NERC business information strategy. Defines NERC’s business information needs and provides a common framework for future developments.

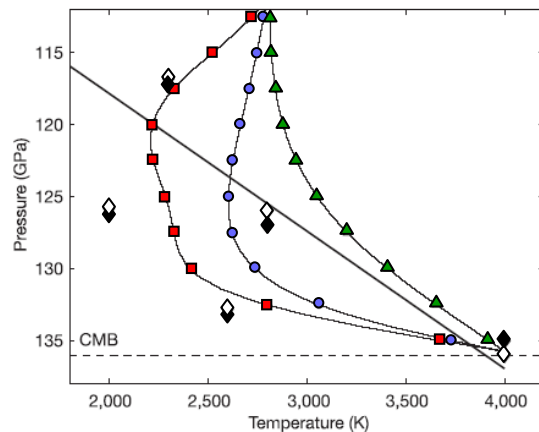
Key Principles

10. **NERC will provide a diversity of HEC resources for the needs of its community**
 The NERC community has a range of requirements for HEC and consequently we need to offer a range of solutions. The technology develops very quickly and consequently the portfolio of technology offered may change quickly to reflect this.

11. Model development can often be achieved using mid-range machines, whereas production runs need access to the most powerful computers available. Some models never use more than 100 processors while others can use thousands simultaneously. We will ensure these contrasting requirements are met.
12. It is recognised that flexibility of access is essential, and any solution will be considered. We will promote shared access to mid-range machines for the whole community. This may be provided by ensuring NERC-funded mid-range systems are set up to provide access for other users, or by partnering with others in a centrally provided facility.
13. Ideally, applicants for resources will be able to choose from a continuum of hardware provision. Offering a flexible service presents its own challenges, not least the long time periods needed for procuring and operating a national HPC service. It is essential the community engage in discussions about requirements and inform the HEC Manager of their future needs. There may be an emerging need to provide for more data rich sectors like bioinformatics or genomics, for example.
14. NERC programmes that use HEC include BAS, NCAS, Oceans 2025, QUEST, RAPID and SOLAS. A number of existing and new consortium grants are also users, for instance HiGEM and CASCADE. Programmes that anticipate using a large proportion of NERC's capacity must signal their need in the early stages of planning to ensure adequate resources can be made available. New programmes that have signalled their need for national HPC resources include the National Centre for Earth Observation and RAPID-WATCH.
15. **NERC is committed to nationally provided HPC services.** NERC scientists, across all disciplines that currently use HPC, need access to the largest available machines at some point in their research. The cost of the largest supercomputers makes it prohibitively expensive for NERC to procure a service for its community without entering in to partnership.
16. Access to national HPC services will be provided through consortia, based on science area and models used. Consortium leaders provide a valuable service to their communities – they allocate resources, act as a point of contact for issues relating to the service and are a conduit for exchange of information with the HEC Manager. To facilitate carrying out this role, each consortium will be given a management allocation of resource on the current national service, determined on a case-by-case basis.
17. **We will encourage new science areas that may be able to benefit from the use of HEC**
The following science areas have been identified as potentially benefiting from a greater use of the national services: Earth Observation, seismology, physical volcanology, bioinformatics, genomics. We will encourage new users of HEC and ensure the whole community is aware of the opportunities provided by HEC.

Mineral Physics Consortium.

This consortium uses about 20% of NERC's capacity on the HPCx service. The following is an example of some of the science made possible by modelling on HPCx.



Phase boundary and calculation points for the perovskite and post-perovskite polymorphs of MgSiO_3 .

We show the variation of temperature with pressure for three vertical profiles in a mantle convection model. These come from 'hot' (triangles), 'cold' (squares) and 'average' (circles) parts of the mantle. Also plotted is the phase boundary from previous work; MgSiO_3 has a perovskite structure above this line, and a post-perovskite form below it. The 'cold' and 'average' profiles cross this boundary twice, thus predicting a thick post-perovskite layer underlain by a thin perovskite layer. The 'hot' profile does not cross the phase boundary at all. It should be noted that this linear phase boundary is approximate; deviation from this will change the transition pressure. The dashed line indicates the CMB. The open and solid diamonds show the positions for which we have *ab initio* calculations of the seismic properties of the perovskite and post-perovskite, respectively.

Ab initio calculations performed on HPCx were used to show that the transition of the lower mantle mineral perovskite to the newly discovered post-perovskite phase is able to explain many of the properties of the lower most mantle (D''). Combining the *ab initio* results with seismic modelling showed that the post-perovskite model is able to explain the topography and location of the D'' discontinuity, the observed shear-wave anisotropy, the observation that the D'' discontinuity is stronger in S-wave models, and the observation of a deeper, weaker discontinuity. (Wookey et al, Nature 2005).

Lidunka Vocadlo (Earth and Planetary Science Letters 2007) carried out *ab initio* molecular dynamics simulations on HPCx and calculated the elasticity of hcp-Fe, bcc-Fe, FeS and FeSi at core conditions. The calculated shear wave velocities of all phases studied were significantly higher than those inferred from seismology. This discrepancy can only be explained if the inner core is partially molten.

Steve Parker and colleagues at Bath have been using HPCx to study the structure of periclase/brucite and calcite nanoparticles in aqueous conditions. They investigated the effect of mineral composition and particle shape on the mineral-water interface and on the transport and adsorption of ions and molecules at room temperature. Very fine-grained materials (nanoparticles) comprise only a small fraction of the mass in the environment but they represent a large fraction of the number of particles in aqueous environments and hence they account for most of the surface area of geological materials. Thus they control a large proportion of the adsorption, reactivity and transport in the environment.

18. **NERC supports EPSRC's role as managing agents of HEC for the research councils**

The EPSRC community is the largest user of national HPC services (approximately 75% of the capacity of HPCx compared with 20% for NERC). EPSRC staff in Swindon organise the procurement and management of the national services, and lead reviews of HEC provision in the UK. NERC does not have the relevant experience in-house, and supports EPSRC in this role.

19. We will commit to the procurement of national HPC services as early as possible in the procurement. This provides the opportunity for NERC to have maximum influence over the service that is finally delivered and ensure it meets the needs of our science community.

20. To realise the full benefits of EPSRC's role, NERC must be actively engaged in the decision-making processes related to HPC provision. We will ensure both NERC Swindon Office and academic representation on appropriate boards and committees. The roles of NERC representatives will be clearly defined, with key members of the academic community plugged in to the decision-making processes. The HEC Manager will act as a conduit for information, providing input to EPSRC from NERC management and the academic community where required. The HEC Manager will ensure NERC representatives are adequately briefed, and in turn they will feed back information to the HEC Manager. Members of the community have a responsibility to provide relevant information to the HEC Manager.

21. **The NERC HEC Steering Committee will provide guidance and advice to the HEC Manager**

The Steering Committee is selected from the NERC community. It carries out peer review of applications for resources, advising on the most appropriate platform the work to be undertaken, and moderating resource requests where appropriate. The current membership includes members of the NERC Cluster Computing Forum, and a representative of the HEC Strategy Committee. The membership will remain flexible in anticipation of changes in priority.

22. **NERC will support software development to exploit new technology**


HPC is evolving rapidly, sustained by exponential speed-ups in technology. Software must be flexible so that it can exploit new architectures as they arise. Without investment in software development, UK researchers will not remain internationally competitive in the areas of science that require the power of future generations of machines.

23. Attracting the best scientists to work on software development is a challenge due to the lack of recognition and a career structure. We will work with other Research Councils to provide recognition and a career structure for computational scientists who develop software. We support the proposed cross-council software programme.

24. Many existing codes do not maximise the benefit of the new and emerging hardware architectures. This issue is not unique to NERC-funded environmental scientists – if possible we will seek to form partnerships for software development programmes where there are benefits to doing so.

Oceanography Consortium

The oceanography consortium is based at the National Oceanography Centre, Southampton. It uses about 40% of NERC's capacity on the HPCx service.



A snapshot from the OCCAM model of ocean salinity in the Southern Ocean. Saltier waters (browns and reds) from north of the Antarctic Circumpolar Current are exchanging with fresher waters (blues) from south of the current. These exchanges also transport heat south across the world's most powerful circulation.

The Oceanography Consortium has used the National HPC resources to study the role of the ocean in climate. The climate models, used in policy decision-making, are run at low resolution and therefore fail to capture oceanic features such as major currents and eddies. These are important for the global transport of heat and fresh water. In order to address this issue we have run a $1/4^\circ$ eddy 'permitting' model and, more recently, a $1/12^\circ$ eddy-resolving model of the global ocean covering the past 20 years. The $1/12^\circ$ model, run on **HPCx**, is the first such experiment internationally - the output is available as a resource for researchers world-wide.

Researchers at NOCS have used the models to study the role of ocean heat transport variations in recent warming of the North Atlantic sea surface temperature (SST). The model SST variability compares favourably with satellite and *in situ* observations and reveals a similar pattern of surface warming over the last 20 years, by over 0.25°C , across much of the North Atlantic. These long-term rises in SST coincide with increasing ocean heat content in the upper few hundred metres, consistent with recent observations. In particular, two episodes of strong heat gain in the mid-latitudes region, in the early and late 1990s, are due to periods of increased heat transport across 36°N into the region, coincident with decreased heat transport out of the region across the Spain-Greenland section. Averaged within the two regions, surface heat fluxes in the model do not show any major trends over 1985-2003 and are within the range of uncertainty in observational estimates over 1985-97. This confirms that the pattern of recent warming in the North Atlantic, particularly in mid-latitudes, is strongly influenced by variations in ocean heat transport.

The transport of properties by baroclinic eddies in the ocean has been widely regarded as potentially important for the time-mean budget of properties. The $1/12^\circ$ model has been used to examine two eddy transport mechanisms (advective and diffusive) for temperature and salinity in the Southern Ocean. The key finding is that the direction and magnitude of eddy transport of any tracer can be explained by the large-scale distribution of the tracer. This may provide a way to estimate eddy transport of tracers (e.g. oxygen, nutrient, etc) which are difficult to observe in the ocean.

25. **NERC will support the training of the next generation modellers**
Inexperienced researchers will be provided with basic training via the HPC consortia. This training, including basic software instruction and programming best practice, may be provided within the consortium or by CSE support at one of the national services.
26. Students will be provided access to national services, provided the HEC Steering Committee approves their applications.
27. The 'People' section of *Next generation science for planet Earth* states that skills gaps will be addressed. The community has identified a skills gap in trained computational scientists. We will develop offering PhD students the opportunity to attend one of the national HEC Training Centres at Warwick and Edinburgh, to take modules in computational science and engineering or complete an additional MSc.
28. **International collaboration**
NERC has international partnerships with the Earth Simulator in Japan, and we are discussing software collaboration with the National Centre for Atmospheric Research in the USA. We will continue to support the emerging European consortium for HPC infrastructure that offers the potential of a machine to match the capability of that provided in the USA and Japan. The increased capability of a joint service must be tensioned against the inevitable loss of control entailed. NERC supports the inclusion of HPC on the ESFRI Roadmap.
29. **Data Analysis and Data Management**
Section 10.8 of *Next generation science for planet Earth* states that data will be managed effectively for the benefit of society and the science community. Models require and generate large volumes of data that require transmission and storage. Communicating these data (input/output, or I/O) is now a limiting factor for some models rather than processor power. NERC scientists will be provided with opportunities to address I/O issues, through collaboration with the national services or utilising the new software development programme. Long-term storage and access to appropriate data will be a challenge – we will seek appropriate solutions.
30. Some users have a requirement for data analysis at the same location as data production. Currently, this need is not well catered for at the national services, and should be addressed.
31. **Application, audit trail, peer review**
We will provide clear information and support for users to apply for HPC resources, through a clear and informative web site and adequate support for HPC in Swindon Office.
32. An audit trail will be maintained for all awards of HPC resources, from application and peer review through to utilisation. Appropriate output and performance measures will be collected annually to ensure an appropriate level of audit for resource use.
33. The HEC Steering Committee will be responsible for the peer review of applications to use HPC resources.

Atmospheric and Polar Science Consortium

This consortium is managed from the National Centre for Atmospheric Science, and includes NCAS and BAS scientists. It uses about 35% of NERC's capacity on the main HPCx service. The following is a case study of the Maritime Continent and the key processes that govern the local hydrological cycle, made possible by modelling on HPCx.

Objective: To understand the source of model biases in the Maritime Continent region.

The climate of the Maritime Continent region in the UK Met Office HadGAM1 atmospheric model has been evaluated, and a simple parameterization that crudely represents the effects of small-scale circulations around the islands has been developed and tested. In common with many climate models, HadGAM1 underestimates precipitation in the Maritime Continent region and is excessively wet in the western Indian Ocean and through much of the Pacific Intertropical Convergence Zone (ITCZ).

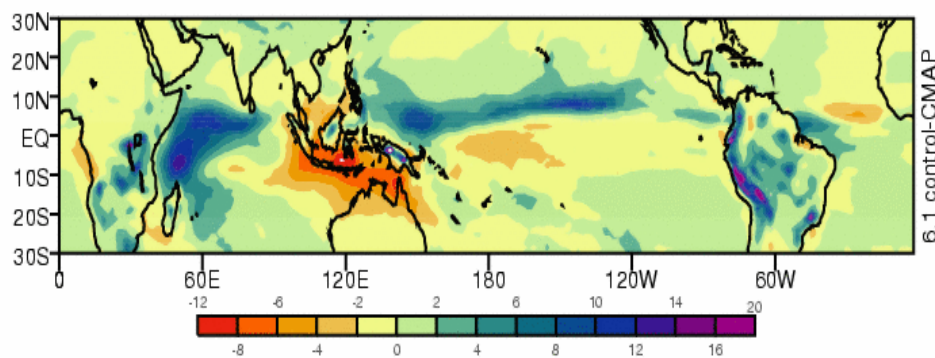


Figure 1. Difference in precipitation in the December-February season between the HadGEM1 model and the CMAP precipitation climatology. Negative differences indicate a dry bias in the model

HadGAM1 incorporates a surface tiling scheme whereby coastal grid-points have a representation of both land and ocean surfaces in varying proportions, and surface fluxes are computed as a weighted average over the tiles. A simple parameterization of sea breeze effects has been introduced, which increases the surface latent heat flux based on the difference between the ocean and land tile surface temperatures. In fact the latent heat flux on coastally tiled points around the Maritime Continent has been found to be too low, so the sea breeze parameterization has been scaled to correct for this underestimate. Precipitation increases around the Maritime Continent and reduces in both the Pacific ITCZ and the western Indian Ocean, in almost the negative pattern of the original HadGAM1 error.




Figure 2. The change in precipitation in HadGAM1 in the Dec - Feb season due to the inclusion of the sea breeze parameterization. The colour scale is the same as for Fig. 1

The model experiments, carried out on HPCx, have shown that the Maritime Continent dry bias is probably due to a number of different factors, both local and regional. There appears to be compensation between regional biases of opposite sign, in a pattern that is systematically reduced by correcting each of the contributory factors.

34. **Funding**

The procurement cycles of the national service require the funding for HPC to be uneven between years. A capital spike occurs at the start of a service, approximately every three years. Capital funding may be secured directly from the Office of Science and Innovation, from NERC's capital line or from the European Union Framework programme.
35. Funding for access to national services will be free at the point of use for NERC funded science, with funds allocated internally to pay for access.
36. **NERC will ensure it gets best value from investments in mid-range machines.**

When NERC invests in a mid-range machine, either as part of a grant or a programme, we will put processes in place to ensure adequate review of the procurement. If necessary, we may expect the operators to allow access to the machine by users external to the institution to which it is awarded. We will maintain an inventory of machines purchased using NERC funds.
37. **The HEC Community**

The International Review of Research using HPC in the UK reported in December 2005. One of the recommendations of the review was to build a computational science community by increasing interactions and collaborations between disciplinary groups. We will encourage interdisciplinary activities such as participation of NERC consortia in the Computational Collaborative Projects Steering Panel.
38. **Annual Report**

A Report of NERC High Performance Computing will be submitted annually to the Services Review Group or equivalent. This report will contain output and performance measures, science highlights, hardware/software developments, a forward look at new services, new programmes and their likely impact on provision. It will draw on information from sources such as HSC and be NERC specific.
39. **Risk Management**

The risks associated with HEC are mainly financial, due to relatively large sums of money involved, and scientific, when the demands of the science are not met by the provision of resources. These risks will be managed through STAR (System for Targets and Risks), NERC's risk management system. The NERC HEC Manager will be the owner of the risks.
40. **Review**

This framework provides guidance for future investment in HEC. It should be an evolving document, providing a template for implementation and will be reviewed at least annually by the HEC Steering Committee.