The University of Manchester

Hadley Centre Review 2006: Supercomputing

Final Report

W T Hewitt, A M Jones, Z Chaplin

12 March 2007

NB:
This report was prepared for Defra and MoD as part of the 2006 Review of the Hadley Centre and should not be further distributed without specific permission from the authors or Defra/MoD.

Contact details:
A M Jones
andrew.m.jones@manchester.ac.uk
+44 (0) 161 275 0635

Manchester Computing
The University of Manchester
Manchester M13 9PL
United Kingdom

http://www.manchester.ac.uk/manchestercomputing

Issue 1 (Draft), W T Hewitt, A M Jones, Z Chaplin, 2 February 2007

Draft for comment by Defra and MoD

Issue 2 (Draft), W T Hewitt, A M Jones, Z Chaplin, 21 February 2007

Revised, incorporating feedback from Defra, MoD, Met Office & Hadley Centre, and additional investigation requested by Defra & MoD.

Issue 3 (Final), W T Hewitt, A M Jones, Z Chaplin, 27 February 2007

Report issued as final.

Issue 4 (Final), W T Hewitt, A M Jones, Z Chaplin, 12 March 2007

Minor correction.

Executive Summary

Introduction

The 2006 Hadley Centre review

Risk Solutions and the University of Manchester have been contracted by the Department for the Environment, Food and Rural Affairs (Defra) and the Ministry of Defence (MoD) to undertake a strategic review of the Met Office Hadley Centre.

The status of the Hadley Centre as a high profile world class research centre is a key part of the UK's international policy on climate change, underpinning the UK's (leading) position in the global climate change agenda.

In the field of climate science, computational modelling is an essential (and arguably even the primary) means of investigation. The use of supercomputing to run long and computationally expensive simulations of the Earth's climate underpins much of the science output and resulting policy advice of the Hadley Centre. Thus, a significant part of the Hadley Centre's capability and operation is the supercomputing provision, which accounts for over half of the total Defra/MoD climate research funding at the Hadley Centre.

Reflecting the critical role of supercomputing, a distinct element of the overall review was focused on the supercomputing provision and exploitation. This report presents the findings of that supercomputing review.

Scope of the review

"Supercomputing provision" potentially means all elements of the supercomputing service (e.g., supercomputer systems, software, management, support, staff, etc); related systems and services such as the mass storage system, the post-processing and visualization facilities, etc; plus expertise and support in supercomputer programming, model development, training, etc.

Whilst this provides a very broad potential scope for a supercomputing review, the issues raised during the review process together with our understanding of the concerns of DEFRA/MoD and the strategic supercomputing issues facing the Hadley Centre have led the review to focus on the following key questions:

- Has the supercomputing service provided quality and value for money in support of the climate research mission?
- Is the Hadley Centre "slipping behind" in terms of supercomputing capacity – and does this matter?
- What are the supercomputing provision options to provide the optimum service to support the climate research mission going forward?

The Impact of Supercomputing

In all disciplines, world class science (and in this case, resulting policy advice) is fundamentally enabled and sustained by access to world class resources – both facilities and people. Indeed, it is clear from history that step changes in science are sparked by step changes in instruments of discovery. This is especially true in the realm of computational science, for which the enabling facility is supercomputing or high performance computing (HPC). This critical enabling resource must also include sufficient depth of HPC expertise to optimally exploit the supercomputers.

Supercomputing is unique among major scientific instruments in two ways:

- No other major facility can be applied with real benefit at the leading edge of every field of scientific research (and many non-scientific fields).
- No other major facility combines such large capital costs with such a short lifecycle.

Supercomputing also shares many characteristics with other major facilities – notably that the ongoing costs of operation are significant; and that the research teams require access to sufficient depth of expertise *in the facility as well as in the science* to get the best results.

Key Findings

The Met Office supercomputing service delivers good value for money to the Hadley Centre

The Hadley Centre currently fulfils its supercomputing requirements through a service provided by the Met Office under a customer supplier type arrangement, consuming around 60% of the overall Defra/MoD funding for the climate programme at the Hadley Centre. This review finds that the supercomputing service provided to the Hadley Centre by the Met Office is of a high professional standard and represents good value for money.

Sustainable reputation and science leadership is inherently dependent on supercomputing

"The accuracy of climate predictions is limited by computing power." [Stern Review, 2006]

The underpinning contribution of the Hadley Centre to the UK government's leading role in the global climate change agenda critically depends on its ability to deliver policy advice backed by leading climate science subjected to the highest standards of international scientific peer review. This science is developed through using supercomputers and observational data to continually test and enhance the climate models. The available supercomputing power limits both this science development and the application of this science (encapsulated in the models) to answer policy questions.

The 2006 Hadley Centre Review International Expert Panel report concludes:

"*The Panel agreed ... that Hadley, even with the proposed Met Office investment in a new supercomputer in 2008, was falling behind and would continue to fall further behind other climate research organisations in computing resources. While in the short term Hadley would be able to retain their position and reputation because of the excellence of their staff, their leadership position could not be sustained longer term without computing capability and capacity comparable with those of their peers overseas.*"

Climate science is readily able to exploit a major step change in supercomputing capacity

Foreseeable policy drivers on the Hadley Centre outputs (in particular the need to reduce uncertainty and the need to provide more localised predictions) require increased supercomputing capacity, through better representation of the science, through increased model resolutions and through increased ensemble set sizes. A conservative estimate of the required increase in supercomputing capacity, endorsed by the 2006 Hadley Centre Review International Expert Panel, shows a need for a step-change increase in supercomputing capacity. Specifically, the science is readily able now to deliver customer benefits from a supercomputer up to 1600x bigger than the current one.

Key Recommendations

The Hadley Centre should continue to meet its supercomputing needs through the Met Office

This review concludes that the current business model, under which a single supercomputing service (which may deliver multiple technologies/architectures and a mix of capability and capacity services) is provided to both climate research and weather forecasting users is highly likely to remain the best service delivery option for the foreseeable future. This core service should be supplemented by a range of external support/consultancy contracts where a critical mass of expertise is already established in external organisations and cannot be justifiably replicated internally.

There is a clear scientific case, and the basis of a strong outputs-led business case, for a multiple order of magnitude increase in supercomputing capacity

DEFRA, MoD and the Hadley Centre should urgently form and promote the business case for a sustainable step change increase in supercomputing capacity

The fundamental role of available supercomputing capacity in creating science, underpinning policy advice and sustaining international reputation, together with the economics of climate change as presented by the Stern Review form the basis for a clear outputs based business case for a multiple order of magnitude increase in supercomputing capacity for the Hadley Centre. **Further, the UK's adoption of a high profile and leading role in the global climate change agenda cannot be undermined by inadequate (supercomputer) support for the science underpinning this role.**

The UK's climate research supercomputer (also used jointly for operational forecasts) is now around 20x slower than the worlds fastest climate supercomputer (the Earth Simulator). Further, the current upgrade target (in 2009) will not deliver a system among the world's leading climate supercomputers.

We recommend that DEFRA, MoD and the Hadley Centre *urgently* form and promote an outputs based business case for a multiple order of magnitude increase in supercomputing capacity (i.e. capital investment and sustainable operational costs).

In support of this, we note that the 2006 Hadley Centre Review International Expert Panel report concludes "*The costs to the UK of adapting its basic infrastructure ... to climate change could run to many hundreds of billions. Without better information on future climate and climatic extremes, the UK is likely either to lose or to waste many billions per year through under- or over-investment in adaptation.*" In addition, the 2006 Stern Review states "*The accuracy of climate predictions is limited by computing power. ... It is important to continue the active research and development of more powerful climate models to reduce the remaining uncertainties in climate projections.*"

The Met Office should establish sustained Computational Science and Engineering (CSE) support

We recommend that a sustained and dedicated CSE effort – i.e., a team or function with a team whose primary skill base is in developing and supporting the optimal exploitation of the supercomputing through parallelisation, optimisation and numerical algorithm work should be established. Sustained CSE effort will become essential to exploit the next generation of supercomputers, independent of which product is chosen, as all currently viable roadmaps show dramatic increases in parallelism and evolving changes in memory hierarchy and bandwidth.

This CSE support will almost certainly be best provided through established external providers, as this is a highly specialised activity that is not easily created quickly, and for which a critical mass of expertise (beyond what is needed for any one customer) is needed to provide the best service. However, it is essential that the degree of CSE expertise continues to grow in the climate applications support team to provide an appropriate interface between the end users and the CSE experts.

The Met Office should expand its efforts to proactively evaluate new supercomputing technologies for the service

The Met Office has commissioned a study into the total cost of ownership aspects and procurement strategy implications with regard to commodity clusters forming part of the future supercomputing service portfolio. We have engaged with that study and whilst it has not yet reached firm conclusions, it is likely that the study will conclude that the case for deploying commodity cluster systems as a significant fraction of the service is not yet evident.

However, we recommend that the Met Office expand its proactive evaluation of new technologies as they become relevant, including commodity clusters, hybrid architectures and new generations of massively parallel processing (MPP) systems. This is especially true when considering significant changes in the size of the supercomputing capacity as recommended above.

The Met Office should enhance support for visualization and data exploration

Visualization is an essential part of the process of understanding (and subsequently communicating to policymakers) the results of the climate simulations. The Met Office should enhance facilities (including software) for both individual (i.e., desktop) and for presentations to policymakers and the public (e.g., powerwalls); and most importantly put in place training and support to ensure effective use of visualisation and data exploration.

The Met Office should continue to evaluate service partnerships

The Met Office is undertaking an evaluation of the potential for supercomputing service partnership with EPSRC and NERC in more depth than is practical under this review. This review has engaged with the Met Office at multiple levels on this issue, and we are satisfied that the Met Office is conducting a full and proper analysis. Whilst the Met Office evaluation is still too early to provide conclusions itself, this review has so far seen no evidence of compelling benefits to the Met Office, the Hadley Centre or their customers arising from the these partnership proposals. In particular, it appears initially that partnership will struggle to successfully mix a capability usage dominated requirement (EPSRC/NERC) with that of a capacity usage dominated requirement (Met Office).

We further recommend that the Met Office continually evaluate other potential partnership options as their context arises, including academia and other organisations, both nationally and internationally.

Summary of Key Review Outputs

Findings

- The Met Office supercomputing service delivers good value for money to the Hadley Centre
- Sustainable reputation and science leadership is inherently dependent on supercomputing
- Climate science is readily able to exploit a major step change in supercomputing capacity

Recommendations

- The Hadley Centre should continue to meet its supercomputing needs through the Met Office
- There is a clear scientific case, and the basis of a strong outputs-led business case, for a multiple order of magnitude increase in supercomputing capacity
- DEFRA, MoD and the Hadley Centre should urgently form and promote the business case for a sustainable step change increase in supercomputing capacity
- The Met Office should establish sustained Computational Science and Engineering support
- The Met Office should expand its efforts to proactively evaluate new supercomputing technologies for the service
- The Met Office should enhance support for visualization and data exploration
- The Met Office should continue to evaluate service partnerships

Contents

Executive Summary ... 3

Summary of Key Review Outputs ... 7

A: Introduction and Background ... 11

1 Context ... 13
1.1 The 2006 Hadley Centre review ... 13
1.2 Scope of the supercomputing review ... 13
2 Science Drivers for Supercomputing ... 14
3 Document structure ... 15

B: Supercomputing ... 17

4 The Impact of Supercomputing ... 19
5 Supercomputing Technology ... 20
5.1 Terminology ... 20
5.2 Architecture ... 20
5.3 Scalability ... 22
5.4 Memory performance ... 22
6 Performance Measurement ... 23
7 Computational Science & Engineering ... 24

C: Value for Money ... 25

8 Findings & Recommendations ... 27
9 Overview of the current service ... 28
9.1 Supercomputers ... 29
9.2 Mass storage ... 29
10 Usage ... 29
11 Quality ... 30
12 Cost and availability ... 30
12.1 The science need for high availability ... 31
12.2 The user need for high availability ... 32
12.3 Summary ... 32
13 Computational science and engineering support ... 33
14 Data exploration and visualization ... 33

D: Size Matters ... 35

15 Findings & Recommendations ... 37
16 International Importance ... 38
17 The world's leading supercomputers ... 38
18 Does size matter? ... 39
19 Earth Simulator Centre – Case Study of a World Leader ... 40
20 How big? ... 40
20.1 Not just big ... 41
20.2 Clever too ... 41

E: Service Delivery Options ... 43

21 Findings & Recommendations....45
21.1 Choosing change 45
22 Current planning process 46
22.1 Evaluating the requirement 46
22.2 Investigating the options 46
23 Current plans for the future service 47
24 Broader service....47
25 Options for satisfying the future need 47
25.1 Option 1 – separate Hadley Centre supercomputing service....48
25.2 Option 2 – supercomputing service provided by the Met Office....48
25.3 Option 3 – partnership, e.g., with collaborators such as the Research Councils 48
25.4 Option 4 - outsourcing....49
25.5 Summary 49

F: Appendices....51

A The Supercomputing Review Team 53
A.1 Manchester Computing at the University of Manchester....53
A.2 W T Hewitt 54
A.3 A M Jones.... 54
A.4 Z Chaplin.... 54
B Hadley Centre Initial Information Request.... 55
B.1 The current arrangements....55
B.2 Future requirements....55
B.3 Supporting information.... 56
C Benchmark Organisation Initial Information Request 58
C.1 Providing a Supercomputing Service 58
C.2 Future requirements.... 58
C.3 Usage Information 58
D Community & stakeholder consultation questions.... 60
D.1 Consultation question 1: Does size matter? 60
D.2 Consultation question 2: The need for high availability....60
D.3 Consultation question 3: The supply options 60
D.4 What next?....60
E Current "Top500" supercomputers.... 61
E.1 Climate supercomputers in Top500....61
E.2 UK supercomputers in Top500.... 62

A:
Introduction and Background

1 Context

1.1 The 2006 Hadley Centre review

Risk Solutions and the University of Manchester have been contracted by the Department for the Environment, Food and Rural Affairs (DEFRA) and the Ministry of Defence (MoD) to undertake a strategic review of the Met Office Hadley Centre.

The status of the Hadley Centre as a high profile world class research centre is a key part of the UK's international policy on climate change, underpinning the UK's (leading) position in the global climate change agenda.

In the field of climate science, computational modelling is an essential (and arguably even the primary) primary means of investigation. The use of supercomputing to run long and computational expensive simulations of the earth's climate underpins much of the science output and resulting policy advice of the Hadley Centre. Thus, a significant part of the Hadley Centre's capability and operation is the supercomputing provision, which accounts for over half of the total Defra/MoD climate research funding at the Hadley Centre.

Reflecting the critical role of supercomputing, a distinct element of the overall review was focused on the supercomputing provision and exploitation. This report presents the findings of that supercomputing review.

1.2 Scope of the supercomputing review

"Supercomputing provision" potentially means all elements of the supercomputing service (i.e., supercomputer systems, software, management, support, staff, etc); related systems and services such as the mass storage system, the post-processing and visualization facilities, etc; plus expertise and support in HPC programming, model development, training, etc.

Whilst this provides a very broad potential scope for a supercomputing review, the issues raised during the review process together with our understanding of the concerns of DEFRA/MoD and the strategic supercomputing issues facing the Hadley Centre have led the review to focus on the following key questions:

- Has the supercomputing service provided quality and value for money in support of the climate research mission?
- Is the Hadley Centre "slipping behind" in terms of supercomputing capacity – and does this matter?
- What are the supercomputing provision options to provide the optimum service to support the climate research mission going forward?

The third question naturally contains a number of elements including: the role of the relationship with the Met Office operational supercomputing service; considering a separate Hadley Centre supercomputing service; the potential role of external HPC service providers; architectural diversity and balance in the supercomputing service portfolio; balancing investment in expertise with compute power; balancing investment in core supercomputing with supporting storage and post-processing capability, etc.

2 Science Drivers for Supercomputing

The science underpinning the requirement for supercomputing is discussed in the overall review's final report prepared by Risk Solutions and additionally in the report of the 2006 Hadley Centre Review International Expert Panel. It is also well covered in the Hadley Centre's own documentation, including for example, their 5 year strategy paper. Aspects of these are summarised below.

Climate change modellers face particular challenges balancing the different demands on their resources. These requirements, and developments in the underlying climate science, have led to increasing demands for complexity and model resolution, including for example:

- Increased vertical and horizontal resolution, over both land and oceans
- Improved representation of the dynamics of the atmosphere, and of many processes in the atmosphere and oceans, and inclusion of more physical processes
- Inclusion of interactions and feed back between constituent systems.

This and the need to quantify uncertainty through repeated (ensemble) model runs mean that climate centres around the world require increasing amounts of computer power, post processing and mass data storage. The Hadley Centre has illustrated the nature of this increased demand through a simple example from their 5 year strategy, which estimates that even a modest increase in model capability arising from science and policy drivers would result in an increased requirement for supercomputer power as follows:

- to double the resolution in all dimensions [a factor of 8 increase in supercomputer capacity required]
- improved science and increased complexity (such as chemistry and earth system feedbacks) [a factor of 8 increase]
- for ensemble predictions with a version of this model which has reduced resolution and complexity [a factor of 25 increase]
- leading to a total increase of a factor of 1600
- or, if reduced resolution/complexity are used for ensemble modelling, to a factor of ~80x.

That is, the science is readily able now to deliver customer benefits from a supercomputer up to 1600x bigger than the current one.

While modelling strategies enable best use to be made of the available computer resources, the demand has increased much more rapidly than was previously envisaged. At the time of the last comprehensive review of the Hadley Centre in 2000 it was anticipated that the computing power available to the centre would increase by a factor of 30 to 70 in 2003. In fact it increased by a factor of 12 in 2004. A recent upgrade provided an additional increase of 20%; however the current resources clearly fall short of what was anticipated or what could be used.

The 2006 Hadley Centre Review International Expert Panel report concludes:

"*The Panel agreed … that Hadley, even with the proposed Met Office investment in a new supercomputer in 2008, was falling behind and would continue to fall further behind other climate research organisations in computing resources. While in the short term Hadley would be able to retain their position and reputation because of the excellence of their staff, their leadership position could not be sustained longer term without computing capability and capacity comparable with those of their peers overseas.*"

3 Document structure

This report is a part of the overall 2006 Hadley Centre review and so, although largely self contained, should be read in that context, i.e., alongside the Final Report of the overall review (prepared by Risk Solutions) and the report from the International Expert Panel (included within the Risk Solution report).

Following the context in this section, Section B presents a discussion on computational science and supercomputing.

Section C describes the current supercomputing service at the Hadley Centre and includes an assessment of its quality and value for money.

Section D addresses the key question of sizing the capability and capacity of the supercomputing, recognising issues of reputation, international competitiveness, and scientific need.

Section E explores the future of supercomputing at the Hadley centre, including service provision options and opportunities.

A number of appendices (Section F) provide supporting information to the report.

B: Supercomputing

4 The Impact of Supercomputing

Computational science (the use of simulation, modelling and data analysis to undertake research, to develop and maintain products, and to process and learn from information) is a crucial and fundamental science that has already become a core part of all modern science, engineering and innovation – including climate research and impact assessments. In the USA, and in Germany, a number of well documented studies have shown the economic benefit associated with high performance computing (HPC), e.g., the 2005 President's IT Advisory Committee (PITAC) report [http://www.nitrd.gov/pitac/reports/] stated that "*Computational science is now indispensable to the solution of problems in every sector, from traditional science and engineering domains to such key areas as national security, public health, and economic innovation ... [that were] previously deemed intractable or beyond imagination.*"

In all disciplines, world class science (and in this case, resulting policy advice) is fundamentally enabled and sustained by access to world class resources – both facilities and people. Indeed, it is clear from history that step changes in science are sparked by step changes in instruments of discovery. This is especially true in the realm of computational science – for which the enabling facility is supercomputing or high performance computing (HPC). This critical enabling resource must also include sufficient depth of HPC expertise required to optimally exploit the supercomputers.

Supercomputing is unique among major scientific instruments in two ways:

- No other major facility can be applied with real benefit at the leading edge of every field of scientific research (and many non-scientific fields).
- No other major facility combines such large capital costs with such a short lifecycle.

Supercomputing also shares many characteristics with other major facilities – notably that the ongoing costs of operation are significant; and that the research teams requires access to sufficient depth of expertise *in the facility as well as the science* to get the best results.

It is an additional defining characteristic of supercomputing, although perhaps not unique, that the technology of supercomputing itself evolves at a pace comparable to, or even faster than, the research undertaken using it. This is one of the key reasons underpinning the need for significant HPC-specific expertise supporting the procurement, operation and use of supercomputer services.

This pace is also one of the drivers of the short lifecycle, as the compute power for a given budget doubles roughly every 14 months – i.e., faster than the Moore's Law derivation (doubling in speed every 18 months) that is commonly applied to processors or personal computers. It means that over a period of 2-5 years (depending on technology and operational/service issues) a given supercomputer service will become obsolete in terms of both cost (i.e., the same capability/capacity can be provided more cheaply by a newer systems, even taking the capital replacement into account) and scientific capability (i.e., the ability to undertake leading science, as many competitors will have significantly advanced their supercomputing capability).

Given these competing aspects – high scientific impact and high depreciation costs (large capital, short lifecycle) determine the appropriate investment is both critical to optimal scientific (or business) success and, equally, hard to accurately and reliably establish. This is discussed with reference to the Hadley Centre, as one of the key review questions, in Section C.

5 Supercomputing Technology

5.1 Terminology

Supercomputing is also referred to as high end computing (HEC) or high performance computing (HPC) and, less frequently, high performance technical computing (HPTC) or high end scientific computing. Although some people make distinctions (usually of emphasis or scale) between these various terms, in general use they are used interchangeably. In addition to these, HPC has over recent years become almost synonymous with the term parallel computing since essentially all supercomputers employ parallel processing (i.e., multiple processors). Note, however, that not all parallel computers (systems with multiple processors) would be considered supercomputers – for example, a coupled network of a dozen PC or servers is unlikely to be regarded as a supercomputer.

5.2 Architecture

There are many papers and conventions that address the categorisations of supercomputers. However, a practical "consensus" of these many conventions (albeit qualitative and with "soft" boundaries) that folds in issues such as architecture, cost, typical usage, vendor positioning, operational aspects and performance is used here.

Supercomputers architectures today fall into roughly three categories:

- "MPP" systems – systems that are primarily designed and targeted to deliver the needs of capability supercomputing, i.e., the application of a substantial proportion of the machine to a single simulation. MPP (the acronym stands for massively parallel processing, although as noted above effectively all supercomputers systems are massively parallel today and so has extended beyond its original usage) implies that the system and technology has been inherently architected as a single massively parallel system encompassing processors, memory, global interconnect, operating systems, physical packaging etc. RAS (reliability availability and serviceability) features are usually directed foremost towards capability workloads. MPP supercomputers are usually only cost effective at large scale. Current examples include Cray's XT4 and IBM's BlueGene.
- "SMP" clusters – systems that consist of a collection of compute nodes (each with multiple processors) integrated together into a single system. These SMP nodes (the acronym stands for symmetric multi-processor, and again has grown beyond its original usage, as few "SMP" nodes are truly symmetric these days) are often designed with a focus on delivering capacity and throughput workloads, which is reflected in the RAS features. The nodes are usually primarily designed for commercial and commodity markets (databases, web servers etc) and are then applied to the HPC market. This class of supercomputer includes "commodity clusters" (2 or 4 socket x86 computers, often connected by Gigabit Ethernet) and systems such as IBM's p-Series with IBM's HPS interconnect or Bull's Novascale servers with Quadrics QsNet interconnect. As the attention to software stack (operating system, cluster management, performance and productivity tools etc) increases, along with the hardware integration and performance levels, these SMP clusters blur the distinction between themselves and MPP class supercomputers.

- "Novel" architectures – again a misnomer in many ways, as all architectures are "novel" at some point in their development and adoption lifecycle. However, at the time of this report, there is a particular community activity in several novel architectures, including the STI Cell Broadband Engine (CBE or just Cell), reconfigurable computing (especially Field Programmable Gate Arrays, FPGAs), and floating point accelerators (e.g., ClearSpeed). Perhaps the most general trend arising from all of these is the increased interest in "multi-paradigm" computing or the use of hybrid processing elements encompassing "traditional" processors together with one or more of Cell, FPGA, ClearSpeed, etc. Individually, these architectures are highly unlikely to be of sufficiently maturity to impact the Hadley Centre's computing needs for the reasonable future; however the issues associated with programming multi-paradigm computing are equally relevant to other more mainstream developments such as multi-core processors, and the increasing complexity of local and global memory hierarchies.

5.2.1 Vector computers

Whilst in the past, and even today in some taxonomies, vector supercomputers were treated as a separate class, this is arguably no longer true. Vector supercomputers employ a particular type of processor (the vector processor, discussed below) but at a system level they fall into either the MPP or SMP cluster classes above. Cray's vector supercomputer line (currently X1E, and the future product known as BlackWidow) increasingly shares characteristics with its scalar line (XT4 and successors) and can be regarded in most aspects as a MPP system, albeit with "special" processors. NEC's SX series of vector supercomputers straddles the boundary between MPP and SMP cluster. In many architectural respects they are SMP class systems as they are designed at the node level. However, they represent the highest levels of system wide integration of nodes to form a supercomputer and thus have many of the characteristics of an MPP system.

A vector processor is able to run mathematical operations on multiple data elements simultaneously. This is in contrast to a scalar processor which handles one element at a time. The vast majority of processors are scalar. Vector processors were common in the scientific computing area and formed the basis of most supercomputers through the 1980s and into the 1990s. Today almost all commodity scalar processor designs include some vector processing instructions, typically known as SIMD (Single Instruction, Multiple Data). Modern video game consoles and consumer computer-graphics hardware in particular rely heavily on vector processors in their architecture. In particular, vector systems provide much higher memory bandwidth than most other processors/systems. The importance of memory bandwidth is discussed below.

5.2.2 Multicore

Of current high profile in the media – both specialist and non-specialist – at the time of this report is the arrival of multi-core processing. This applies to each of the classes above and means that within a given processor there two or more processing cores – i.e., the bits that do the work, such as mathematical calculations. Arguably multi-core processors are not new – they have been available for a few years. What has changed over the last year or so is the widespread acceptance that all future (scalar) processors are likely to be multi-core. This has significant effects on memory bandwidth and resulting programming and performance (see below).

5.3 Scalability

With all of the architectures above there is a rapid trend to increase the total compute power of a machine by increasing the number of processors rather than just increasing the performance of those processors (indeed multi-core means that the performance of an individual core is unlikely to increase rapidly). Some efficiency is lost when running codes on many processors (as clearly some effort has to be devoted to the coordination of those many processors in a single job) and eventually there are no further gains which can be made as the cost of communicating between processors outweighs the benefit of additional processors. This code scalability becomes increasingly important and is dependent on system, code, and data set among other things. Addressing scalability is, however, essential for achieving increased performance.

5.4 Memory performance

Whilst the speed of a processor in performing mathematical operations (especially floating point operation per second or FLOPS) has always been the substantial factor in the real performance of a supercomputer, the performance of the memory hierarchy has equally been critical to utilising this processing power. With increasing complexities of memory hierarchy through the advent of multi-core processors and the growth of systems sizes in terms of numbers of processor or nodes, memory performance is increasingly recognised in the HPC world as the critical aspect of a supercomputer. It has been noted (by Bill Dally of Stanford) that the cost of memory bandwidth is an order of magnitude greater than that of processing power.

There are two aspects to memory performance – latency and bandwidth. Latency most notably affects the memory system performance for small data sizes – a few bytes. Bandwidth most notably affects the memory system performance for larger data sizes – few Kbytes and upwards.

Latency is usually shortest (best) and bandwidth greatest (best) at the most local cache levels such as cache specific to a core, and the systems exhibit progressively longer latency and less bandwidth i.e., worse performance, moving through the memory system hierarchy to processor cache that may be shared between a number of cores, to node memory that may be (asymmetrically) shared between a number of processors, to global memory that may be part of another node across the interconnect.

In the case of most scientific simulations, the local node memory bandwidth is of most importance, followed by the bandwidth and latency of the global memory (i.e., across the interconnect). Many scientific codes are also tuned (by HPC experts and even by compiler software) to optimise their usage of local cache memory, where data is able to be fitted (or blocked) into cache.

Different numerical algorithms and scientific codes place different requirements on memory performance. In the case of finite difference algorithms, such as that underpinning the Hadley Centre climate models, where the computational density compared to the volume of data needed for that computation is low, the most stringent demand is on getting enough data into the processor to extract its performance - i.e., on node memory bandwidth. Data sizes also tend to be large enough to make cache less effective (noting that the NEC SX6/8 vector systems used at the Met Office do not employ cache memory as their node memory subsystem is optimised for this kind of large data bandwidth use). Thus, for this class of code, systems with optimal memory bandwidth per [dollar, system, square foot, watt, processor, etc] rather than optimal FLOPS per [dollar, … metrics as above] are normally required. This has usually meant vector supercomputers but this is not always true, and for some data sizes and problem partitioning, systems such as IBM's Power based p-Series or HP's Itanium 2 based Integrity Superdomes (which are both high memory bandwidth scalar systems) offer better solutions.

A note on multi-core processors – these usually share the bandwidth to a given element of the memory systems – so that a dual core processor has only half of the local memory bandwidth compared to a single core processor and so on. This places memory bandwidth at the heart of the programming change to multi-core paradigms.

It is fair to say that any foreseeable supercomputer will employ parallelism significantly beyond that currently used and the resulting impact on local and global memory bandwidth will need to be addressed by programmers. Increasingly, algorithmic development is focused on recognising and working with the challenges of the memory hierarchy. The days of relying on vector computers delivering the bandwidth will soon fade, and although vectors will continue to offer the best solutions for many aspects of the node memory bandwidth problem they too will need active programming consideration.

6 Performance Measurement

It is useful to measure and compare against some standard the speed of supercomputers. Each supercomputer has a theoretical maximum performance for floating point computations, known as the peak performance. In practice, no real application achieves this theoretical performance level, often missing it by an order of magnitude or more, depending on the application and on the architecture of the supercomputer. The performance (speed) that a real application achieves is known as the sustained performance and this also varies between applications. It is worth noting that climate modelling codes are usually subject to sufficient optimisation and tuning effort to enable them to attain reasonable sustained performance on deployed architectures.

The international Top500 list is the most prevalent performance ranking systems of the world's supercomputers. The Top500 uses the High Performance Linpack (HPL) benchmark application as a measure of sustained performance. It should always be borne in mind that the Top500 is really a list of the top 500 scores on HPL, not truly a list of the top 500 supercomputers, but over the last 20 years these two things have correlated well. On most architectures, HPL is able to achieve between 30-80% of the peak performance. Most real user applications (e.g., fluid dynamics, molecular modelling, etc) vary between 5% and 50% of the peak performance.

HPL in particular measures a system's floating point performance, with processor FLOPS and cache size/speed being critical. HPL does not really measure either node or global memory performance, which we have noted above is critical for many climate codes.

Whilst HPL (and thus the Top500) has its detractors and limits as a benchmark representative of the broad range of scientific codes, the 20 years history of the Top500 means it has a useful role to play in tracking the industry trends over time. The difference between the top and bottom of this list is nearly two orders of magnitude and so when considering the list as a whole, the difference between a system where HPL performs much better than a real code (e.g., a climate code) and another system where the performance is more comparable can be neglected for most analyses.

7 Computational Science & Engineering

In common with other major scientific instruments of discovery or research facilities, effective use of supercomputers requires that (climate) research teams have access to sufficient depth of expertise *in the facility as well as in the science* to get the best results.

The complexity of the supercomputing ecosystem and its nature as a rapidly developing technology (evolving at a pace comparable to, or even faster than, the research undertaken using it) are the key reasons underpinning the need for significant HPC-specific expertise supporting the procurement, operation and use of supercomputer services.

The underlying algorithmic and software engineering innovations required to continue to exploit this rapid evolution of supercomputers requires a sustained and dedicated CSE effort – i.e., a team or function with a team whose primary skill base is in developing and supporting the optimal exploitation of the supercomputing through parallelisation, optimisation and numerical algorithm work.

Such a sustained CSE effort will become essential to exploit the next generation of supercomputers, independent of which product is chosen, as all currently viable roadmaps show dramatic increases in parallelism and evolving changes in memory hierarchy and bandwidth.

It is important to note that this effort is required now, to enable models and algorithms that can be optimally used on the next generation of supercomputers in a few years time, which will in turn result in science and papers a few years after. That is, the impact of this CSE effort will not be fully seen for several years. But the converse is also true – failing to provide the appropriate CSE effort now will not have its full detrimental impact on the science output for some years – and by then it will be far too late to recover the situation without a further period of sustained CSE effort.

This CSE support is almost certainly best undertaken in partnership with established external providers, as this is a highly specialised activity that is not easily created quickly, and for which a critical mass of expertise (beyond what is needed for any one customer) is needed to provide the best service. However, it is essential that the degree of CSE expertise continues to grow in the climate applications support team to provide an appropriate interface between the end users and the CSE experts.

C:
Value for Money

8 Findings & Recommendations

> The Met Office supercomputing service delivers good value for money to the Hadley Centre

The Hadley Centre currently fulfils its supercomputing requirements through a service provided by the Met Office under a customer supplier type arrangement, consuming around 60% of the overall Defra/MoD funding for the climate programme at the Hadley Centre. This review finds that the supercomputing service provided to the Hadley Centre by the Met Office is of a high professional standard and represents good value for money.

> The Met Office should establish sustained Computational Science and Engineering (CSE) support

We recommend that a sustained and dedicated CSE effort – i.e., a team or function with a team whose primary skill base is in developing and supporting the optimal exploitation of the supercomputing through parallelisation, optimisation and numerical algorithm work should be established. Sustained CSE effort will become essential to exploit the next generation of supercomputers, independent of which product is chosen, as all currently viable roadmaps show dramatic increases in parallelism and evolving changes in memory hierarchy and bandwidth.

This CSE support will almost certainly be best provided through established external providers, as this is a highly specialised activity that is not easily created quickly, and for which a critical mass of expertise (beyond what is needed for any one customer) is needed to provide the best service. However, it is essential that the degree of CSE expertise continues to grow in the climate applications support team to provide an appropriate interface between the end users and the CSE experts.

> The Met Office should expand its efforts to proactively evaluate new supercomputing technologies for the service

The Met Office has commissioned a study into the total cost of ownership aspects and procurement strategy implications with regard to commodity clusters forming part of the future supercomputing service portfolio. This review has engaged with that study and whilst it has not yet reached firm conclusions, it is likely that the study will conclude that the case for deploying commodity cluster systems as a significant fraction of the service is not yet evident.

However, we recommend that the Met Office expand its proactive evaluation of new technologies as they become relevant, including commodity clusters, hybrid architectures and new generations of massively parallel processing (MPP) systems. This is especially true when considering significant changes in the size of the supercomputing capacity as recommended above.

The Met Office should enhance support for visualization and data exploration

Visualization is an essential part of the process of understanding (and subsequently communicating to policymakers) the results of the climate simulations. The Met Office should enhance facilities (including software) for both individual (i.e., desktop) and for presentations to policymakers and the public (e.g., powerwalls); and most importantly put in place training and support to ensure effective use of visualisation and data exploration.

9 Overview of the current service

The current supercomputing service of the Hadley Centre is provided by the Met Office IT division.

The supercomputing service provided by the Met Office includes:

- Approximately 3.7 TF of NEC SX6 and SX8 vector supercomputers
- Approximately 800TB of mass storage (Hadley Centre share)
- Management and operation of the supercomputer service
- 5 on-site dedicated support staff provided by NEC (3 for hardware, 2 for software)
- Met Office Supercomputer Support team (5 staff)
- Mass storage support team
- Service desk

The service is provided on a customer supplier basis. Approximately 60% of the Hadley Centre's DEFRA funded climate programme (i.e., around £7M per year) is paid to the Met Office as a fixed annual fee for this service. This buys the Hadley Centre just over a 50% share of the overall supercomputing service.

Within its total allocation, the Hadley Centre manages the usage of supercomputing primarily through the prioritisation of scientific management.

The Met Office plans the future operation of its supercomputers on a rolling 5-year programme, in which the coming year's requirements are specified as closely as is reasonable on the basis of known requirements. Hadley Centre staff are consulted throughout the procurement cycle.

Whilst the Hadley Centre's initial disclosure report and strategy documents did not directly address supercomputing in any detail, they did make some clear points regarding the climate models and the process of modelling which are relevant. In particular, Chapter 3 (Building Capacity and Reducing Uncertainty) of the "Hadley Centre Achievements 2000-2006" document submitted to the review by the Hadley Centre discusses the different models (HadCM3, HadGEM1a and successors) and the collaborations in place to develop them.

9.1 Supercomputers

In early 2005 the Met Office accepted the first production NEC SX-8 supercomputer in the world. This system provides additional capacity to the NEC SX-6 computer systems used by the Met Office in Exeter. The SX-6 and SX-8 machines are divided into 'nodes' with each node containing eight processors. Currently the SX-8 has 16 nodes, and there are two SX-6 clusters; one with 19 nodes, and the other with 15 nodes. Each SX-8 node is twice as powerful as an SX-6 node, so together the combined systems deliver over thirteen times the sustained power of the previous Cray T3E computers, which the SX-6 replaced in 2004.

For operational resilience, the systems are divided between two halls, with the SX-8 and the 15 node SX-6 cluster in one hall, and the 19 node SX-6 cluster in the other. This enables the operational forecast to be maintained even if one of the halls is unavailable.

The nodes are accessed through front-end machines. There are six of these, three in each hall. The front end machines are scalar machines and perform tasks such as compiling the computer code, which is better performed on a scalar machine. The resulting program is then run on the vector nodes.

As was the case on the previous supercomputer, a Cray T3E, the operational model is run on multiple processors but as each individual processor is that much more powerful than the T3E ones, far fewer are required. For instance, the operational global model is typically run on 24 SX-8 processors compared to 144 T3E processors.

9.2 Mass storage

The output from the simulations is stored in the tape library. The current capacity is around 1.4 PBytes, with the Hadley Centre share being 800TB (around 60%). There was some criticism from the Hadley Centre staff about the access time of the tape store but we note that the Met Office plans to address this are well advanced. **It is important in future procurements to maintain the balance between the supercomputer and mass storage**, but in particular noting that latency/bandwidth from the tape-store to supercomputer if inadequate will stall productivity.

10 Usage

The service supports around 300 users, of whom around 100 are Hadley Centre accounts. Around 20 of these Hadley Centre users account for the bulk (75%) of the Hadley Centre consumption of the supercomputer resource. Some 200,000 Hadley Centre compute jobs are processed by the system per year. All of these numbers are typical for a supercomputer service of this size and nature. Approximately ¼ of the supercomputer related support calls to the helpdesk arise from Hadley Centre users.

The vast majority of the individual jobs run on the supercomputer by Hadley Centre users are fairly small by supercomputer centre standards (14 cpu-hours and 8GB of memory or less and using only a few cpus at once). However it must be made clear that these individual jobs must be run in sequence and together form the climate simulation of interest. The simulation is split into this sequence of jobs to optimise scheduling and usage of the supercomputer and to help guard against system failure before job completion. Together with the overall simulation length, it is the seer quantity of these jobs (to create the statistics required for meaningful climate research) that defines the need for the capacity of supercomputer.

However, there are also a number of jobs requiring large individual compute time (over 200 cpu-hours) and these drive a different type of supercomputing requirement - that of capability.

Balancing capacity and capability usage is one of the big issues facing most large supercomputer services, especially those aimed at delivering a prestige science mission. The "biggest" science (highest profile, leading edge, or other metric) tends to come from the capability workload whereas the volume of science required to sustain a leading position overall in a scientific field comes tends to come from the capacity workload. In addition, capability computing requires leading edge supercomputing architectures that are often not the most cost-effective architecture for capacity computing requirements.

11 Quality

It is difficult to measure "quality" of a supercomputing service and many supercomputer centres and service contract procurements struggle with this issue. Hard metrics such as uptime or utilisation do not properly measure the service quality in the same way that the stress co-efficient of a spring does not measure the comfort of a mattress.

In undertaking this review, we have noted the highly professional manner of the Met Office supercomputing team and this may be one of the best (albeit subjective) indicators of a high quality service.

One common solution to this issue is to rely on user satisfaction measures, e.g., annual user surveys. In this respect, the staff workshop conducted early in this review provides the most relevant material. One theme arising in particular from the workshop is that the overall computing ecosystem, including elements such as visualization, post-processing and model development tools appears to require investment to better support the scientific process.

It is also important to consider the processes that are in place to monitor and act on user feedback and other quality data. The primary mechanism for this is the Supercomputer Users Working Group (SUWG). Meetings of the SUWG are held monthly and include service reports from the Supercomputer Support manager. The main interface between the end users and the supercomputer service comes via the Climate Research IT support team. This represents a good bridge between user domain experience and supercomputing service expertise.

12 Cost and availability

The Met Office supercomputer service is designed to include a significant degree of high availability as needed by the operational forecasting. This usually adds significant cost to a service, as seen by compute work that does not need this availability. Note that it is the period over which availability is measured that is the critical factor. All types of sustained research compute work require high availability over practical research timescales (e.g., 3 months) and the investment in availability over this period is usually paid back in full though increased research throughput. However, most research work does not require high availability over short timescales (e.g., 3 hours) and it is the cost of providing this level of availability that is usually disproportionately expensive.

This should not be confused with a valid requirement for high availability over the life of the service (years). Reducing the overall availability over the service life represents a reduction in the effectiveness of the investment, unless by reducing the availability requirement, the initial (capital) investment is greatly reduced. This may be best explained by way of example. Assuming all other factors were equal (they rarely are!) a £30M supercomputer providing 100TF at 90% lifetime availability is better value for money than a £30M supercomputer providing 75TF at 99% lifetime availability. An accountant might note that the first case is "losing" 10% (£3M) of its investment, but it is still providing more output (TF) for its £30M than the latter case.

The concern is that the cost of the Met Office operational requirement for high availability is being shared with the climate research requirement leading to higher costs for the Hadley Centre than otherwise. Against this, it must be noted that the Met Office service delivers the high level of availability in part through deploying a service sized significantly bigger than the base requirement. Thus, it may equally be argued that the Met Office investment in high availability through additional capacity is actually subsidising the Hadley Centre workload via the use of "spare capacity".

It must be noted here that this is not a case of the "cluster vs. supercomputer" argument that often appears in less informed comment, as either cluster or supercomputer installations may be architected to be of high or low availability (although it is true that many cluster installations are procured at the lower end of availability). Clusters may often have much lower initial capital costs but over the course of a multi-year service to deliver a given research mission the total cost of ownership (TCO) can be comparable to a supercomputer.

The Met Office has commissioned a study into the total cost of ownership aspects and procurement strategy implications with regard to commodity clusters forming part of the future supercomputing service portfolio. This review has engaged with that study and whilst it has not yet reached firm conclusions, it is likely that the study will conclude that the case for deploying commodity cluster systems as a significant fraction of the service is not yet evident.

The climate research computing of the Hadley Centre may be able to be provided more cost-effectively through a service that is designed to a different balance of availability and cost. This applies to both capacity and capability workloads. This question was extensively investigated during the review, including through the community & stakeholder consultation exercise, addressing the question "is it generally true that the Hadley Centre use of supercomputing does not have a requirement for very high availability over short periods (3 hours) and reliability over long periods?" with the subsequent potential result that a significant proportion of the supercomputing investment should be devoted to high end capacity computing with "medium" availability – i.e., supercomputing that has high availability over a period of a few weeks (overall computing cycles delivered over the course of a large ensemble experiment), capable of cost-effectively running very large numbers of compute jobs, where each compute job still requires significant computer power in itself.

12.1 The science need for high availability

It is clear that availability levels for climate research do not need to be as high as for short-term weather forecasting applications. Thus, an outage of (say) 3 hours *occasionally* would not pose a problem, though outages of 3 hours *frequently*, would.

However, it is equally true that supercomputer services delivering research missions around the world continue to put reliability and availability high on the list of required system features.

In particular, relatively high availability is important for climate research applications most notably because major "production" runs (which have model horizons of 100 or more years) take a long time to run (several months). It is important that they are started as close to customer delivery deadlines as possible, so as to maximise the amount of science that can be incorporated in the model. If only a low-availability machine was available, runs would need to be started earlier with a larger time contingency (to allow for machine outages and recovery of the runs) and therefore less up-to-date science could be incorporated in the model, thus yielding poorer quality results.

In the case of the Hadley Centre, where the production output drives either policy advice or major outputs such as IPCC assessments there is a strong link between the quality of the output and the availability of the supercomputer that distorts the normal tension between availability and cost for research use of supercomputers.

12.2 The user need for high availability

Availability (and other related characteristics such as reliability and quality) have a strong interaction with the scientific users/researchers:

- Long or repeated outages are frustrating and distracting for users. High availability contributes to good staff morale and hence to their creative and intellectual output;
- High availability affords greater predictability and thus permits better planning and efficient use of scientists' time, plus increased confidence and ambition in submitting computational workloads, thus yielding more rapid scientific advances;
- High availability allows scientists to devote more time to the science and interpretation of climate runs and less time to recovering from run failures and developing software to mitigate against the impacts of frequent run failures;
- High availability speeds problem solving, as trial and error de-bugging processes (of either code or science), and "what-if" experiments can be turned around more quickly.

12.3 Summary

Without a full cost benefit analysis and a clear definition of the approach and usage scenarios etc, which is beyond the scope of this review, we are cautious in drawing any firm conclusions regarding availability solutions. We recommend that the Hadley Centre and Met Office continue to investigate this issue for the future. It is worth making three notes on this:

- Where a major change to critical business processes is being considered (such as moving to multiple architectures for availability and cost balancing) the general rule should be that if the move is not compellingly beneficial it should not be done – i.e., a marginal benefit is not sufficient to balance out the risk and impact of transition, nor the unknowns in future requirements and usage;
- In this case, we must also ask – it may be that the cost/availability balance set by the operational forecast is different to that of climate research (high cost), but does the climate research share of that cost, given that the operational forecast would need the provision in any case, still provide good value to climate research, even if that same availability/cost balance would not provide best value if delivering the climate research mission alone?

- It is often quoted in the HPC industry that "*a capability computer can do capacity work, but a capacity computer cannot do capability work*". This is true, and for a policy driven programme such as the Hadley Centre, retaining flexibility for both capability and capacity workloads against unknown future evolutions of customer requirements is important. Maximum flexibility is normally achieved through deploying "traditional" supercomputers configurations with high availability and capability focused architectures.

13 Computational science and engineering support

One key service element that is missing at the Hadley Centre is any significant computational science and engineering (CSE) expertise – i.e., a team whose primary skill base is in developing and supporting the optimal exploitation of the supercomputing through parallelisation, optimisation and numerical algorithm work.

We refer to the discussions in Sections 5.3, 5.4, and 20.2 that highlight some of the areas where CSE expertise underpins scientific discovery using modelling, and to Section 7 which addresses the critical role of CSE and software engineering in sustained the science leadership of the model and is use.

It is worth noting that the current common consensus among the computational science community (including end users) is that typical supercomputing investment needs to be rebalanced in favour of CSE skills and other software support as opposed to the supercomputer hardware – without diminishing the required increases in supercomputing power!

14 Data exploration and visualization

Just running the simulation is clearly not sufficient, the scientist needs to assimilate the results to develop his or her understanding of the science. Visualization of the results is a major component of this process and consists of two phases:

- **Exploration**: what the scientist does (possibly with colleagues) to explore the wealth of results, and gain an understanding of those results. This could be envisaged as looking at large numbers of graphs, charts, pictures, movies, and animations, for a variety of situations with the result of some scientific insight by the climate scientist. This also includes using data-mining and other software assisted exploration techniques.
- **Presentation:** where the climate scientist is presenting results to others whether that be in academic papers, presentations at academic conferences or papers and presentations to policymakers (DEFRA etc.,) or the public. Typically a small number of charts etc., are used to summarise the exploration phase.

The software and hardware used in both of these phases is essentially identical, although used in different ways.

The importance of both data exploration in large science and of presentation in high profile and policy-informing output leads this review to conclude that the provision of enhanced visualization facilities across the scientific process from individual (desktop) to group/collaborative usage (e.g., powerwall) be pro-actively considered. We note in particular that the data volumes available to explore are very significant (800TB) which will place strong demands on both hardware and software for these facilities. These facilities can also be used to enormous effect with policy makers and the public during the presentation phase.

In the material provided to the review team excellent use is made of visualization, but it is notable that most of the graphs and charts are 2 dimensional, with little use made of 3 dimensional charts. Further oral presentations to the international expert review panel made little use of animation. Indeed the animations that were shown were animations of 2D charts.

Most of the graphics and charts are produced using PV-Wave, and the Hadley Centre staff indicated that they were changing to IDL in the near future. This software change in itself does not fully address the concerns of the review team.

There is a wealth of techniques from computer graphics and visualization that appear not to be used by Hadley Centre staff, which is down mainly to inertia – a) the staff are not aware of these techniques, and b) the software available does not provide these techniques.

It must be stressed that just putting in place the additional facilities will not be sufficient, a concerted training effort must be undertaken to ensure take up and use of these facilities.

D:
Size Matters

15 Findings & Recommendations

Sustainable reputation and science leadership is inherently dependent on supercomputing

"The accuracy of climate predictions is limited by computing power." [Stern Review, 2006]

The underpinning contribution of the Hadley Centre to the UK government's leading role in the global climate change agenda critically depends on its ability to deliver policy advice backed by leading climate science subjected to the highest standards of international scientific peer review. This science is developed through using supercomputers and observational data to continually test and enhance the climate models. The available supercomputing power limits both this science development and the application of this science (encapsulated in the models) to answer policy questions.

The 2006 Hadley Centre Review International Expert Panel report concludes:

"*The Panel agreed … that Hadley, even with the proposed Met Office investment in a new supercomputer in 2008, was falling behind and would continue to fall further behind other climate research organisations in computing resources. While in the short term Hadley would be able to retain their position and reputation because of the excellence of their staff, their leadership position could not be sustained longer term without computing capability and capacity comparable with those of their peers overseas.*"

Climate science is readily able to exploit a major step change in supercomputing capacity

Foreseeable policy drivers on the Hadley Centre outputs (in particular the need to reduce uncertainty and the need to provide more localised predictions) require increased supercomputing capacity, through better representation of the science, through increased model resolutions and through increased ensemble set sizes. A *conservative* estimate of the required increase in supercomputing capacity, endorsed by the 2006 Hadley Centre Review International Expert Panel, shows a need for a step change increase in supercomputing capacity. That is, the science is readily able now to deliver customer benefits from a supercomputer up to 1600x bigger than the current one.

There is a clear scientific case, and the basis of a strong outputs-led business case, for a multiple order of magnitude increase in supercomputing capacity

DEFRA, MoD and the Hadley Centre should urgently form and promote the business case for a sustainable step change increase in supercomputing capacity

The fundamental role of available supercomputing capacity in creating science, underpinning policy advice and sustaining international reputation, together with the economics of climate change as presented by the Stern Review form the basis for a clear output based business case for a multiple order of magnitude increase in supercomputing capacity for the Hadley Centre. Further, the UK's adoption of a high profile and leading role in the global climate change agenda cannot be undermined by inadequate (supercomputer) support for the science underpinning this role.

We recommend that DEFRA, MoD and the Hadley Centre urgently form and promote an outputs based business case for a significant increase in supercomputing capacity.

In support of this, we note that the 2006 Hadley Centre Review International Expert Panel report concludes "*The costs to the UK of adapting its basic infrastructure ... to climate change could run to many hundreds of billions. Without better information on future climate and climatic extremes, the UK is likely either to lose or to waste many billions per year through under- or over-investment in adaptation*." In addition, the 2006 Stern Review states "*The accuracy of climate predictions is limited by computing power. ... It is important to continue the active research and development of more powerful climate models to reduce the remaining uncertainties in climate projections*."

16 International Importance

It is a matter of policy that the Hadley Centre be one of the world's leading climate research centres. This fosters an element of national competition in the global cooperation on climate research. The UK's credibility as a leading player on the climate issue directly affects its ability to influence and drive the global agenda in the best direction for both the world as a whole and the UK specifically.

It is often raised as a concern that supercomputing becomes a "technology race", in which ever larger machines are justified (subconsciously or otherwise) by the fact that competitor organisations have them. However, this is an artefact of the nature of supercomputing – no other instrument of scientific discovery applies so effectively across so many disciplines, nor increases performance so dramatically each generation. Thus, regular orders of magnitude increase in scale are inherent to the value of supercomputing, with a multitude of relevant climate research ready to exploit supercomputers much bigger than currently available or planned.

It is clear that the potential impact of climate change is so significant (as shown by the Stern Review) that investment by the UK of several million pounds annually on understanding its mechanisms and potential impacts seems justifiable. Supercomputing, as the fundamental enabling resource of modelling, is clearly a major part of this investment need.

17 The world's leading supercomputers

Whilst not in itself a key factor in determining supercomputing investment, it is nevertheless useful to compare the supercomputing capacity available to the Hadley Centre with that of other climate centres worldwide, and also with other supercomputers within the UK.

The two tables in Appendix E show the rankings of climate research supercomputers (E1) and UK based supercomputers (E2) and in the November 2006 list of the world's Top500 supercomputers.

This list is ranked based on the supercomputer's recorded speed (measured in Teraflops, TF) against the HPL performance benchmark and is updated twice-yearly. It should be noted of course that speed of different codes (models) vary across the different types of supercomputer, and that a supercomputer that is good at running HPL may or may not be good at running climate models. Nevertheless the 20 years of historical data together with the breadth of systems (geography, type, etc) covered means that the Top500 is one of the most useful indicators of the supercomputing industry.

The world's current fastest supercomputer [the 360TF BlueGene/L at Lawrence Livermore National Lab (LLNL)] is also shown in both tables, along with the UK's fastest [the 40TF Cray XT3 at the Atomic Weapons Establishment (AWE), shaded blue], Europe's two fastest systems [94TF JS21 cluster at Barcelona and 64TF Bull cluster at Commissariat à l'Energie Atomique (CEA –the French equivalent of AWE), both shaded in yellow], and the UK Met Office system (2TF NEC SX8, shaded in green) which is not among the Top500 fastest systems in the world.

Looking at the climate table (E1) it is noticeable the range of different countries present – Japan, South Korea, USA, EU (ECMWF), China, and Japan. Clearly many governments are able to justify the significant investment required to procure – and run – a Top500 supercomputer for climate research. It is also noticeable that the USA has 9 supercomputers in the Top500 for climate research – more than any other country.

Looking at the UK table (E2) it is noticeable the range of different sites present – from government (defence) to academic research to industry (mostly finance). A variety of different organisations are readily able to justify the significant investment required to procure – and run – a Top500 supercomputer.

Finally, it can be seen that the UK's climate research supercomputer (also used jointly for operational forecasts) is now around 20x slower than the worlds fastest climate supercomputer (the Earth Simulator). Looking forward to 2009 when the next UK Met Office upgrade is due (currently intended to be a 8x increase) it is reasonable to expect (by extrapolation from the historical data) that the fastest climate research supercomputer by then will be some 4x faster than the newly upgraded system. Or to put it another way, the current upgrade target will not deliver a system among the world's leading climate supercomputers.

18 Does size matter?

It is often argued that a world leading supercomputer is not needed for world leading climate research. In our opinion, this can be true with appropriate investment in staff and other resources over a medium timescale, but world leading climate research is not sustainable over any reasonable period of time (years) without continued access to supercomputers amongst the world's fastest. This applies equally to any other key resource – excellence in one area (e.g., scientific staff) can cover the gaps of weaker areas – but only for so long. Excellence can be achieved with limited facilities but sustained excellence requires leading edge facilities.

As addressed in the context of CSE support in Section 7, it is important to note that inadequate supercomputer resources now will restrict the development of the science and the application of the models over the next few years, and thus the full detrimental impact on the science output of inadequate supercomputing power will not become evident for some time.

19 Earth Simulator Centre – Case Study of a World Leader

Review team members were able to visit the Earth Simulator Centre in Japan to discuss the Earth Simulator project including the current supercomputing service and future plans. This was through discussions with Shigemune Kitawaki of the Japan Agency for Marine-Earth Science and Technology (JAMSTEC) which hosts and runs the Earth Simulator Centre.

The project started in 1997 by developing a world leading supercomputer to enable a comprehensive understanding of the global changes such as global warming. It was funded by the former STA (Science and Technology Agency of Japan, now MEXT: Ministry of Education, Culture, Sports, Science and Technology). The Earth Simulator computer is made up of 640 nodes; each node is a custom version of the NEC SX-6 supercomputer. The system came into service in February 2002 with a peak performance of 40TFlops and was the world's fastest supercomputer for over 2 years. It made a major impact on the international community with the step change in capability it offered the Japanese scientists. In the context of this review there are two important aspects pertinent to the Hadley Centre:

- The substantial increase in capability enabled breakthroughs in their climate research.
- The system has attracted world class collaborators to work with JAMSTEC. Indeed the Hadley Centre has an allocation of resources on the facility. This is enabling the Hadley Centre to undertake work that cannot be accommodated on the existing Met Office facility, but more interestingly will still not be possible on the 2009 planned upgrade to the Met Office supercomputer (almost 7 years after the Earth Simulator was deployed).

In 2008 the Earth Simulator expect to upgrade from 40TFlop peak performance to 150TF-200TF and a further upgrade to around 340TFlop in 2010. This upgrade will be funded from within their current operating budget – i.e. they are not expecting another major capital investment from the government.

20 How big?

In section 2 we refer to the **clear scientific case for a significant increase in supercomputing capacity** – we repeat here the Hadley Centre's conservative estimate (endorsed by the 2006 Hadley Centre Review International Expert Panel) of the computing needs of future science and policy driven modelling requirements, leading to an estimate of 1600x performance required in the near future. This can be compared with following: "*simple extrapolation of climate science and model requirements is easily able to consume a 10,000,000x increase in computer power – without considering extending the ensemble set sizes*" [Getting Up to Speed: The Future of Supercomputing, 2005, National Research Council, National Academies, USA].

Clearly then, the answer to how big will never be big enough for the foreseeable future. It is equally clear that a target of 8x is inadequate (based on approximately level spend, although arguably even approximately level spend should target 15-20x). A 1600x increase is most likely outside the budget of the UK government, however this still leave two orders of magnitude to address. With arguments placed purely on future national economic grounds (drawing from the 2006 Stern Review for example) a significant budget increase should be considered. When combined with the international political impact of the UK's role in the global climate agenda the arguments for a significant budget increase gain considerable strength. Whilst an Earth Simulator scale project for the UK may not be currently economically feasible, it should at least be actively considered and, if necessary, decided against rather than simply being ruled out without proper consideration.

In addition to the clear scientific case, the fundamental role of available supercomputing capacity in creating science, underpinning policy advice and sustaining international reputation, together with the economics of climate change as presented by the Stern Review form **a strong basis for a clear outputs based business case for a multiple order of magnitude increase in supercomputing capacity for the Hadley Centre**. Further, the UK's adoption of a high profile and leading role in the global climate change agenda cannot be undermined by inadequate (supercomputer) support for the science underpinning this role.

We recommend that DEFRA, MoD and the Hadley Centre urgently form and promote an outputs based business case for a significant (order of magnitude) increase in supercomputing capacity (i.e. capital investment and sustainable operational costs).

In support of this, we note that the 2006 Hadley Centre Review International Expert Panel report concludes "*The costs to the UK of adapting its basic infrastructure ... to climate change could run to many hundreds of billions. Without better information on future climate and climatic extremes, the UK is likely either to lose or to waste many billions per year through under- or over-investment in adaptation.*" In addition, the 2006 Stern Review states "*The accuracy of climate predictions is limited by computing power. ... It is important to continue the active research and development of more powerful climate models to reduce the remaining uncertainties in climate projections.*"

20.1 Not just big

In overall terms what is most important to the Hadley Centre is the delivery of their science. The size of the system therefore should be appropriate to their goals in this area and sized such that it can deliver the science now and over the lifetime of the system. There is little to be gained by obtaining a very theoretically fast computer if it doesn't deliver the science. The determining factor is the limitation of the application software and the characteristics of the system that will deliver the better research rather than the overall size of the system. It is worth noting that the performance of an individual compute job is only one contributing factor to the volume and quality of science output. Application work-flow, usability, ease of management and system flexibility, the system's ability to be able to handle mixed workflows and applications, and its ability to manage jobs and work flows through its memory and its storage systems are also important elements.

20.2 Clever too

It is also prudent to note here that resource limitations have always been the force behind innovations. In computational science, a healthy tension between raw supercomputer power and continued algorithmic innovation is essential for scientific progress. Limits in the foreseen available computing power, and in the nature of this power (e.g., multi-core, memory bandwidth, increasing parallelisation), direct the effort in algorithm development. In turn, the regular deployment of more powerful supercomputers is essential to exploit the limits of these algorithmic advances in delivering new science.

Whilst this will never be an argument for not deploying larger supercomputers when there is easy evidence of at least a 1600x climate modelling performance requirement, it does mean that the drive to innovate with this larger facility must be maintained.

E:
Service Delivery Options

21 Findings & Recommendations

> The Hadley Centre should continue to meet its supercomputing needs through the Met Office

This review concludes that the current business model, under which a single supercomputing service (which may deliver multiple technologies/architectures and a mix of capability and capacity services) is provided to both climate research and weather forecasting users is highly likely to remain the best service delivery option for the foreseeable future. This core service should be supplemented by a range of external support/consultancy contracts where a critical mass of expertise is already established in external organisations and cannot be justifiably replicated internally.

> The Met Office should continue to evaluate service partnerships

The Met Office is undertaking an evaluation of the potential for supercomputing service partnership with EPSRC and NERC in more depth than is practical under this review. This review has engaged with the Met Office at multiple levels on this issue, and we are satisfied that the Met Office is conducting a full and proper analysis. Whilst the Met Office evaluation is still too early to provide conclusions itself, we have so far seen no evidence of compelling benefits to the Met Office, the Hadley Centre or their customers arising from these partnership proposals. In particular, it appears initially that partnership will struggle to successfully mix a capability usage dominated requirement (EPSRC/NERC) with that of a capacity usage dominated requirement (Met Office). A supplementary recommendation is that, should clear benefits emerge from the evaluation, a proactive process of stakeholder engagement should be pursued to ensure that these benefits are understood by funders, customers and operators of the supercomputing service(s).

We further recommend that the Met Office continually evaluate other potential partnership options as their context arises, including academia and other organisations, both nationally and internationally.

21.1 Choosing change

Where a major change to a critical business process (such as the supercomputing service) is being considered the general rule should be that if the move is not compellingly beneficial it should not be done – i.e., a marginal benefit is not sufficient to balance out the risk and impact of transition, nor the unknowns in future requirements and usage. This recommendation of principle applies particularly to the previous two issues – i.e. moving to multiple architectures for availability and cost balancing, or major partnerships for service provision.

22 Current planning process

22.1 Evaluating the requirement

The Hadley Centre has senior and technical representation on the future supercomputer project board and in the project team for the replacement supercomputer to ensure the users requirements are met. In defining the requirement, the benchmarks, which are agreed with the key users and are based around the Unified Model, play a key role. Benchmark configurations (including capability and capacity data sets) are chosen based on predicted usage patterns of the supercomputer. The Hadley Centre's workload typically represents 50% of the benchmark weighting. It is also worth noting that over time the model resolutions used by the Hadley Centre increase towards those used by the operational forecasting near the start of the service – and thus the Hadley Centre can have confidence in the supercomputers ability to deliver on reasonable growth in resolution from their own initial baseline. Additionally, detailed functional requirements are developed, for the Hadley Centre by members of the project team who have a very detailed knowledge of the needs of the application (the Unified Model).

In sizing and balancing the facility/funding across compute resource, data storage, staffing and other operational costs, the biggest challenge is reported to be relative sizing of the super-computer and MASS storage. The Met Office has a model for estimating MASS storage costs for a given amount of super-computer power, which has proven to be reasonably accurate in the past. The model has fudge-factors that are used to adjust rates of data production based on changing user needs. Human resources are based on the effort to port, optimize and support a number of versions of the Unified Model. Large code changes, such as conversion to MPP, are planned in advance (note: see comments elsewhere in this report on CSE support). As appropriate, network infrastructure is upgraded in line with the supercomputer needs and is estimated by the supercomputer support team with user input as to changing requirements. Operator and systems support requirements on Met Office staff are estimated before final solution selection based on system complexity. The costs of these support elements are included in the cost of ownership (which also includes factors such as risk functionality, etc) of the system and in the evaluation process to ensure best value for money.

22.2 Investigating the options

The Met Office has a scrutinised business case process that will analyse the evidence before making a decision. All realistic options available to them will be considered by the process. The process includes an analysis of direct costs, risks and benefits. Typically, they try to quantify as many of these as possible. The current HPC project has a team of technical, business and financial experts. This team draws together a case which includes the facts and a recommendation. The case is scrutinised by members of the project team (to provide project assurance) including members of the project board. The case is summarised for executive approval as defined below. The case is further scrutinised by the Bid and Investment Appraisal Committee, an internal group of Directors and other senior staff of all professional disciplines including business and financial. The case is submitted to the Met Office Board, which is comprised of both Executive and non-Executive Directors. Cases of this size also need to be approved by MoD (Ministry of Defence) for both investment and capital loan approval.

Collaboration plays a strong role in the eventual exploitation of the supercomputers, models, and science – including links with other supercomputer sites such as the Earth Simulator – and this is covered in more detail in the overall review report from Risk Solutions.

The Met Office staff keep aware of the supercomputer industry and user community, including technology awareness and computational science developments. Where appropriate topics merit it (e.g., evaluation of clusters or of new technologies such as Cell processors) the Met Office should (and does) employ expert external consultants to support and advise the business processes.

23 Current plans for the future service

As referred to elsewhere in this report, the current Met Office procurement plan targets a least an 8x performance improvement to be delivered in early 2009, with a budgeting assumption of approximately level spending with previous procurements. Section D (chapters 15-20), and especially chapter 20 “How big?” discusses this sizing in more detail.

The Met Office is undertaking an evaluation of the potential for supercomputing service partnership with EPSRC and NERC in more depth than is practical under this review. This review has engaged with the Met Office at multiple levels on this issue, and we are satisfied that the Met Office is conducting a full and proper analysis. Whilst the Met Office evaluation is still too early to provide conclusions itself, this review has so far seen no evidence of compelling benefits to the Met Office, the Hadley Centre or their customers arising from the these partnership proposals. In particular, it appears initially that partnership will struggle to successfully mix a capability usage dominated requirement (EPSRC/NERC) with that of a capacity usage dominated requirement (Met Office). A supplementary recommendation is that, should clear benefits emerge from the evaluation, a proactive process of stakeholder engagement should be pursued to ensure that these benefits are understood by funders, customers and operators of the supercomputing service(s).

Where a major change to critical business processes are being considered (such as major restructuring of service delivery partnerships) the general rule should be that if the move is not compellingly beneficial it should not be done – i.e., a marginal benefit is not sufficient to balance out the risk and impact of transition, nor the unknowns in future requirements and usage.

24 Broader service

The supercomputing service needs to be broadened to encompass a significant degree of computational science support and appropriate post-processing and data exploration/visualization facilities. This is discussed more fully in chapters 7, 13 & 14.

25 Options for satisfying the future need

There are essentially 4 options for providing the supercomputing service:

- In-house provision by the Hadley Centre (i.e., separate to the Met Office operational service)
- Current arrangements (i.e., provided by the Met Office jointly with the operational service)
- Shared resource with other major collaborator(s) – e.g., the Research Councils, EU partners, etc
- Outsourcing and other supply mechanisms

25.1 Option 1 – separate Hadley Centre supercomputing service

This option is regarded, both by the review team, and the majority of respondents to the consultation exercise as highly unlikely to deliver any significant benefits over the current arrangements (c.f. option 2 below). In particular, there is currently no compelling evidence (see section 3.6) that the Hadley Centre needs a different *type* of supercomputer from that required for other Met Office types of modelling (*i.e.,* short-range weather forecasts). However, as mentioned previously, as part of the development of the Business Case for replacement of the existing Met Office supercomputer a study has been commissioned to consider whether a cluster-type solution could offer advantages over traditional supercomputer architecture, for climate research. Even aside from any possible benefit of different technology solutions (e.g., clusters) providing optimal benefit for the Hadley Centre, the overall service provision duplication cannot be justified, and is readily ruled out.

25.2 Option 2 – supercomputing service provided by the Met Office

This is the current business model, under which a single supercomputing service (which may deliver multiple technologies/architectures as necessary) is provided to both climate research and weather forecasting users. This is discussed in depth in Section C, but here we note that it possesses several significant advantages over other options, including:

- Economies of scale are being secured with respect to system and software management;
- Maximum synergies between weather forecasting science and climate change science are being obtained since essentially the same code, running on different time-scales, is being used on the same machine. This key benefit is discussed more fully in the Risk Solutions report.

This review concludes that this option is highly likely to remain the best delivery option for the foreseeable future, noting however that the evaluation of the partnership options (below) have yet to be completed by the Met Office.

25.3 Option 3 – partnership, e.g., with collaborators such as the Research Councils

The Met Office has considered partnerships for supercomputing services with a range of different organisations (including UK government and European organisations) in the past. These past discussions did not result in eventual partnerships. However, partnership is now being considered again by the Met Office, with the Research Councils (EPSRC and NERC) covering procurement and service operation. This proposal is based on the possibilities of economies of scale and the opportunity to advance science more rapidly through a more joined-up approach than would be possible using separate facilities. We refer to chapter 23 for more details.

It is fair to say, that the academic community has not overwhelmingly endorsed this potential partnership, nor does the benefit appear to be clear within the Hadley Centre. We note again, however, that the Met Office board are undertaking a full evaluation of the partnership options in more depth than is practical under this review, and that we have engaged with that review and are confident it is working properly and towards verifiable conclusions.

25.4 Option 4 - outsourcing

The Met Office's Business Case will review all options for supercomputing service provision to the Met Office (including the Hadley Centre). However, it is reported that the Met Office feels there is significant risk in using outsourcing as a major part of the supercomputing service delivery. This is in contrast to the reason that many companies pursue the outsourcing option – i.e., to lower both overall cost of ownership and service delivery risks. Note that outsourcing does not, necessarily, have to mean moving the supercomputers to another location. Many outsourcing customers retain the supercomputers within their own facilities and simply outsource the operation of the facility and/or the service.

Large scale distributed computing using "free" resource from the Internet as typified by climateprediction.net are simply another form of outsourcing (albeit with low apparent cost, complexity of data management, and unpredictable delivery of the overall resource) and thus needs to be considered along with other outsourcing solutions.

Given the data constraints imposed by the modelling, and the existence of a high quality and professional supercomputing service team with the Met Office, including modern hosting facilities, it would seem unlikely that a significantly outsourced service (even if located at the Exeter site) would deliver practical benefits.

In addition, the UK does not appear to have any significant number of sustained success stories of outsourcing scientific supercomputing services (in contrast to other countries – the USA for example uses outsourcing widely in government supercomputing centres). This lack of clear track record in turn limits the market interest in providing such services at affordable costs – as experienced by EPSRC under the HECToR procurement process.

An exception to this lack of UK track record is EPSRC's outsourcing of their HPC services, to a consortium of CSC, the University of Manchester and SGI for the CSAR contract (£30M, 1998-2006) and to a consortium of the University of Edinburgh, Daresbury Laboratory and IBM for the HPCx contract (£53M, 2002-2008). Both of these outsourcing programmes appear to have worked well. In the case of CSAR, the contract successfully combined industry and academic solution providers. This CSAR contract also successfully used the PFI mechanism.

25.5 Summary

Clearly a combination of the above options is possible and different options may be most appropriate for the different elements of the service

In particular, the computational science support will almost certainly be best provided through established external providers, as this is a highly specialised activity that is not easily created quickly, and for which a critical mass of expertise (beyond what is needed for any one customer) is needed to provide the best service. However, it is essential that the degree of CSE expertise continues to grow in the climate applications support team to provide an appropriate interface between the end users and the CSE experts. (This interface model appears to work well for the current supercomputing usage.)

This review concludes that the best model will be a core service provided by the Met Office (probably with a mix of capability and capacity services) supplemented by a range of external support and consultancy contracts where the critical mass of expertise is already established in external organisations.

F: Appendices

A The Supercomputing Review Team

A.1 Manchester Computing at the University of Manchester

Manchester Computing (MC) is the University of Manchester's centre for Research Computing services, support and related research. MC has provided a number of regional, national and international services, notably the CSAR national HPC service in partnership with CSC and SGI, and leadership of the UK's National Grid Service. Manchester Computing also undertakes a wide range of research in supercomputing, Grid/e-Science, high performance visualization and collaborative working technologies.

Individually and collectively MC staff provide consultancy services within the University and externally to academia, industry and government at both technical and strategic levels. MC has run over 10 supercomputing procurements in the last 8 years for internal and external customers.

Key combined expertise and experience of the review team members:

- Strategic planning of supercomputing requirements, resources and services
- Supercomputing service definition and innovation
- Management of supercomputing services
- Costing and financial management of supercomputing systems and services
- Evaluation of supercomputing technologies and architectures including benchmarking
- Procurement of supercomputing systems and services, including EU/OJEU procurement
- IT Governance
- Application of supercomputing to real world problems
- Large scale scientific data processing & analysis
- Supercomputing practice in the MoD and academia
- Building inter-site collaborations and relationships, including international development
- Project leadership
- Presentation, communication and major proposal architecting/writing
- Climate modelling, including algorithms and science of the Unified Model
- Scientific applications on supercomputers
- Supercomputing user support
- End user of supercomputers

A.2 W T Hewitt

Terry Hewitt is an internationally known figure in supercomputing and is currently Director of Research Computing at the University, leading the research support activity of Manchester Computing and provides strategic leadership of the research computing services at the University. Terry has been active in scientific computing since the early 1970s and has extensive experience in programming and supporting researchers particularly in supercomputing and visualization. In the late eighties he established a parallel processing service at the University and since that time he has had increasing responsibility for the development of high end computing and high performance visualization services both locally and nationally. More recently he has been active in the development of e-Science again both locally and nationally. This service role has been complemented by a diverse research and development programme that he has led, receiving funding from EPSRC, NERC, ESRC, DTI, EU and international companies. Terry is also an honorary lecturer in the School of Computer Science where he lectures on the Advanced Computer Science and Computational Science MSc courses and supervises a number of MSc and PhD students. Terry is a Fellow of the Institute of Mathematics and its Applications, Chartered Mathematician and Chartered Scientist. He has been Chairman of an International Standards Working Group, Chairman of the European Association for Computer Graphics, and has lectured extensively throughout the world.

A.3 A M Jones

Andrew has over 10 years of experience of scientific research using supercomputing, with expertise in the challenges and application of supercomputing to scientific research, including technical, cultural and business issues associated with applying this successfully to real world problems. Andrew's career has spanned being an end-user, technical/CSE support, buyer and manager of HPC facilities, and has gained and applied this experience through work at academic, industry and government organisations.

In his current role at Manchester Computing Andrew is responsible for strategic developments in supercomputing and related activities. This includes developing the strategy for supporting computational based research (including supercomputing); developing relationships with other national and international centres using supercomputing, and with major supercomputing vendors; ensuring the HPC services and research programmes evolve and innovate in response to changing user needs and external technology/business changes; and acting as the lead for the outreach and promotion activities associated with the HPC services and research.

A.4 Z Chaplin

Zoe Chaplin has over 8 years experience of the UM on a variety of (supercomputing) platforms and for a number of different versions. She also has experience of performance optimisation and parallelisation of other scientific codes as well as application support for a variety of codes.

Following her degree in Mathematics from the University of Bristol and a Masters in Fluid Dynamics from Cranfield University, Zoe worked in the New Dynamics and later UM Systems groups at the UK Met Office, developing a new dynamical core for the Unified Model (UM), with specific responsibility for the mesoscale model covering the UK and the implementation of the boundary conditions required to run this configuration and later implementing changes to the UM for new releases of the model.

B Hadley Centre Initial Information Request

B.1 The current arrangements

Contractual and funding:

- Details of the statement of requirements defining this service
- Details of the contractual structure under which this service provided
- Details of the SLA/SLD under which this service is provided
- Details of the funding/costing of this service, including any performance related elements, service credit regimes
- Details of the funding process with DEFRA (e.g., how is the funding acquired – source/process – and how is it matched to the requirements?)

Operational matters:

- Details of in-service/ongoing monitoring and review processes
- Details of the UKMO organisation providing the service, including at least the following: staff numbers, staff experience and training, staff career progression and retention issues, overview of financial operations, impact of the Hadley supercomputing provision to other supercomputing services provided from UKMO, etc
- Details of representation of the customer (DEFRA/Hadley) in operational, procurement and other strategic matters affecting the service
- Operational details of the relationship between the Hadley service and the UKMO forecasting needs
- Details of the security procedures (required for UKMO Operational Forecasting) with respect to the impact on the Hadley service (e.g., ability to collaborate with external groups)

Other:

- Recommendations for representatives at various levels of both Hadley and UKMO to contact for detailed follow-up in later parts of the study (e.g., investigating experience and perceptions of the various service provision options) - management, support, end users

B.2 Future requirements

Evaluating the requirement:

- How are requirements from the end user community obtained?
- How is the facility/funding sized (balanced) across compute resource, data storage, staffing and other operational costs?
- How is the diversity of requirements reconciled?
- How are competing requirements prioritized and balanced?

- How are user workflow practices & environment [culture, security, tolerance of change etc] taken into account?
- How is the need for leadership and innovation taken into account - i.e., stepping beyond the users often conservative statements of nature/scale of future needs?

Investigating the options:

- How are the pros and cons of the various service provision options evaluated and balanced? (e.g., partnerships with other organisations such as the Research Councils)
- Are there any limits (e.g., export or other trade) that limit the international source of either technology or services specific to the Hadley requirement?
- What is the role of external and internal collaboration in satisfying requirements?
- How are the evolving technology and architectures matched with the evolving applications and algorithms?
- How are the risks, costs and benefits of new/disruptive technologies (such as Cell, FPGA) evaluated?
- How are the end user and applications base engaged over these topics?
- What conflicts of interest exist between the Hadley, UKMO and other potential service providers?

B.3 Supporting information

Usage of the service:

- Number of Hadley user accounts on the supercomputer
- Total number of user accounts on the supercomputer
- Number of "projects" into which the Hadley accounts are organised
- Number of Hadley user accounts together that consume >75% of the total Hadley usage (i.e., how many of the users are "big" users?)
- Number of UKMO user accounts together that consume >75% of the total UKMO usage (i.e., how many of the users are "big" users?)
- Average number of jobs run per year by Hadley users
- Average number of jobs run per year by all users
- Breakdown (within reasonable effort of data collection) of job sizes (cpu, memory, disk, time, etc) of the Hadley users
- Breakdown (within reasonable effort of data collection) of job sizes (cpu, memory, disk, time, etc) of the UKMO users
- Average number of calls per year to the helpdesk (specific to supercomputing) from Hadley users
- Average number of calls per year to the helpdesk (specific to supercomputing) from all users
- Average number of calls per year to the helpdesk (all queries) from all users

- Any comments on the variability (e.g., seasonal) or other characteristics of the statistics requested above

Profile of the service:

- Example *Hadley* disseminations (to public, collaborators, internal, or management) such as presentations/papers/reports describing the Hadley usage of the supercomputer service and results arising from that usage
- Example *UKMO* disseminations (to public, collaborators, internal, or management) such as presentations/papers/reports describing the Hadley usage of the supercomputer service and results arising from that usage
- Other high profile (e.g., public mass media, or very senior internal/stakeholder) disseminations discussing the Hadley supercomputing

C Benchmark Organisation Initial Information Request

C.1 Providing a Supercomputing Service

Overview:

- Please provide an overview of the supercomputing service provided to the climate research centre/users, including main hardware & infrastructure etc.

Contractual and funding:

- Please describe the Statement of Requirements (or equivalent Service Level Agreement etc) that defines the supercomputing service provided to the climate research centre/users. Where possible please provide a copy of the SoR/SLA etc.
- Please describe the contractual structure under which this service provided.
- Please provide an overview of the funding/costing of this service, including any performance related elements, service credit regimes, etc.

Operational matters:

- Please describe the main in-service/ongoing monitoring and review mechanisms for the supercomputing service.
- Please describe the organisation providing the service, including elements such as staff numbers and type/experience, staff training, career progression and retention issues, etc.
- Please describe how the climate research centre users (customers?) are represented in operational and strategic matters affecting the service (including procurements).

C.2 Future requirements

Evaluating the requirement:

- Please describe how the requirements from the climate research user community obtained and prioritized/balanced across the diversity of compute resource, data storage, staffing and other requirements.
- Please comment on the role of service leadership and innovation (e.g., stepping beyond the users statements of nature/scale of future needs) in defining the future service.

C.3 Usage Information

Usage of the service:

- Please state the number of climate user accounts on the supercomputer
- Please state the number of "projects" into which these climate accounts are organised
- Please estimate the number of climate user accounts together that consume >75% of the total usage (i.e., how many of the users are "big" users?)
- Please estimate the average number of jobs run per year by the climate users

- Please comment on the typical breakdown of job sizes (cpu, memory, disk, time, etc) of the climate users

Profile of the service:

- Please provide copies of example disseminations (to public, collaborators, internal, or management) such as presentations/papers/reports describing the climate usage of the supercomputer service and results arising from that usage

D Community & stakeholder consultation questions

D.1 Consultation question 1: Does size matter?

Current and planned supercomputers of the Met Office will continue to lag significantly in size (or speed) behind other climate centres world wide.

Can a world leading climate science programme – and thus the policy advice to support the UK's international leadership role in the global climate change/impact agenda – be sustained without a supercomputer capability among the world's fastest for climate research?

D.2 Consultation question 2: The need for high availability

Is it generally true that the Hadley Centre use of supercomputing does not have a requirement for very high availability over short periods (3 hours) and reliability over long periods?

A result of this would be that a significant proportion of the supercomputing investment should be devoted to high end capacity computing with "medium" availability – i.e., supercomputing that has high availability over a period of a few weeks (overall computing cycles delivered over the course of a large ensemble experiment), capable of cost-effectively running very large numbers of compute jobs, where each compute job still requires significant computer power in itself.

D.3 Consultation question 3: The supply options

There are essentially 4 options for providing the supercomputing service:

- In-house provision by the Hadley Centre (i.e., separate to the Met Office operational service)
- Current arrangements (i.e., provided by the Met Office jointly with the operational service)
- Shared resource with other major collaborator(s) – e.g., the Research Councils, EU partners, etc
- Outsourcing and other supply mechanisms

What are the pros and cons of each of these options?

In particular we are interested in views on an arrangement that has the capability computing provided through collaboration with UK academia and/or the capacity computing provided by a specialist external capacity computing provider.

D.4 What next?

We welcome views on other topics that the supercomputing review should investigate that are not covered above.

E Current "Top500" supercomputers

E.1 Climate supercomputers in Top500

Climate Rank	World Rank	Site	Country	Use	System	Year	Speed (HPL Rmax)	%peak	GF per CPU
	1	DOE/NNSA/LLNL	USA	Research	IBM BG/L	2005	280.60	76	2.1
	5	Barcelona	Spain	Medicine	IBM JS21 PPC	2006	62.63	66	6.1
	7	CEA	France	Defense	Bull IA64 Quadrics	2006	52.84	83	5.3
1	14	Earth Simulator	Japan	Research	NEC vector	2002	35.82	87	7.0
	15	AWE	UK	Defense	Cray XT3	2006	32.50	80	4.2
2	29	Korea Meteorological Administration	S Korea	Weather/Climate Research	Cray X1E	2005	15.71	85	15.4
3	35	Naval Oceanographic Office (NAVOCEANO)	USA	Weather/Climate Research	IBM p5	2006	13.99	60	4.6
4	38	ECMWF	UK	Weather/Climate Research	IBM p5	2006	13.99	82	6.2
5	39	ECMWF	UK	Weather/Climate Research	IBM p5	2006	13.99	82	6.2
6	48	Naval Oceanographic Office (NAVOCEANO)	USA	Weather/Climate Research	IBM p5	2006	11.99	82	6.2
7	57	China Meteorological Administration	China	Weather/Climate Research	IBM p4+	2005	10.31	47	3.2
8	67	ECMWF	UK	Weather/Climate Research	IBM p4+	2004	9.24	56	4.2
9	69	Japan Meteorological Agency	Japan	Weather/Climate Research	Hitachi SR11000-K1/80	2005	9.04	84	113.0
10	70	Japan Meteorological Agency	Japan	Weather/Climate Research	Hitachi SR11000-K1/80	2005	9.04	84	113.0
11	129	Japan Meteorological Agency	Japan	Weather/Climate Research	Hitachi SR11000-J1/50	2005	4.99	82	99.9
12	167	National Centers for Environmental Prediction	USA	Weather/Climate Research	IBM p4+	2004	4.38	56	3.8
13	193	NCAR (National Center for Atmospheric Research)	USA	Weather/Climate Research	IBM p4	2003	4.18	50	2.6
14	277	Naval Oceanographic Office (NAVOCEANO)	USA	Weather/Climate Research	IBM p4	2004	3.72	51	2.6
15	346	Forecast Systems Laboratory - NOAA	USA	Weather/Climate Research	HPTi Aspen Xeon	2002	3.34	49	2.2
16	437	NOAA/Geophysical Fluid Dynamics Laboratory (GFDL)	USA	Weather/Climate Research	SGI Altix	2005	2.87	93	5.6
17	438	NOAA/Geophysical Fluid Dynamics Laboratory (GFDL)	USA	Weather/Climate Research	SGI Altix	2005	2.87	93	5.6
18	439	NOAA/Geophysical Fluid Dynamics Laboratory (GFDL)	USA	Weather/Climate Research	SGI Altix	2005	2.87	93	5.6
N/A	N/A	UK Met Office	UK	Weather/Climate Research	NEC SX8/128M16	2005	2.02	98	15.8

E.2 UK supercomputers in Top500

UK Rank	World Rank	Site	Country	Use	System	Year	Speed (HPL Rmax)	%peak	GF per CPU
	1	DOE/NNSA/LLNL	USA	Research	IBM BG/L	2005	280.60	76	2.1
	5	Barcelona	Spain	Medicine	IBM JS21 PPC	2006	62.63	66	6.1
	7	CEA	France	Defense	Bull IA64 Quadrics	2006	52.84	83	5.3
	14	Earth Simulator	Japan	Research	NEC vector	2002	35.82	87	7.0
1	15	AWE	UK	Defense	Cray XT3	2006	32.50	80	4.2
2	20	Cambridge University	UK	Research	ClusterVision/Dell	2006	18.27	65	7.8
3	38	ECMWF	UK	Weather/Climate Research	IBM p5	2006	13.99	82	6.2
4	39	ECMWF	UK	Weather/Climate Research	IBM p5	2006	13.99	82	6.2
5	43	HPCx	UK	Not Specified	IBM p5	2006	12.94	84	5.1
7	67	ECMWF	UK	Weather/Climate Research	IBM p4+	2004	9.24	56	4.2
7	127	Financial Institution (J)	UK	Not Specified	HP Opteron	2006	5.04	56	2.7
8	132	Financial Services (M)	UK	Finance	IBM LS20 Opteron	2006	4.92	56	2.5
9	133	Financial Services (M)	UK	Not Specified	IBM LS20 Opteron	2006	4.92	56	2.5
10	148	University of Edinburgh	UK	Not Specified	IBM BG/L	2005	4.71	82	2.3
11	166	Imperial College ICT HPC	UK	Research	Dell Xeon IB	2006	4.40	25	2.6
12	196	University of Reading	UK	Not Specified	IBM JS20 PPC	2006	4.15	64	5.6
13	243	Bank	UK	Not Specified	IBM HS20 Xeon	2004	3.76	19	1.2
14	248	Credit Suisse/First Boston	UK	Finance	IBM HS20 Xeon	2004	3.76	41	2.5
15	264	Geoscience (C)	UK	Geophysics	IBM x Xeon	2004	3.76	54	3.3
16	265	Geoscience (C)	UK	Geophysics	IBM x Xeon	2004	3.76	54	3.3
17	300	Financial Services (O)	UK	Not Specified	IBM HS20 Xeon	2006	3.59	52	3.8
18	307	Fujitsu Services	UK	Information Processing	HP Superdome	2006	3.53	53	2.1
19	330	Government	UK	Not Specified	HP Integrity Superdome	2006	3.43	75	4.5
20	375	University of Nottingham	UK	Not Specified	Sun Opteron	2005	3.15	72	3.1
21	440	Bank (J)	UK	Finance	HP Xeon	2005	2.86	56	3.6
22	441	Bank (J)	UK	Finance	HP Xeon	2005	2.86	56	3.6
23	442	Bank (J)	UK	Finance	HP Xeon	2005	2.86	56	3.6
24	443	Bank (J)	UK	Finance	HP Xeon	2005	2.86	56	3.6
25	444	Bank (J)	UK	Finance	HP Xeon	2005	2.86	56	3.6
26	445	Bank (J)	UK	Finance	HP Xeon	2005	2.86	56	3.6
27	446	Bank (J)	UK	Finance	HP Xeon	2005	2.86	56	3.6
28	447	Bank (J)	UK	Finance	HP Xeon	2005	2.86	56	3.6
29	448	Bank (J)	UK	Finance	HP Xeon	2005	2.86	56	3.6
30	467	EDS	UK	Not Specified	HP Superdome	2006	2.85	53	1.9
N/A	N/A	UK Met Office	UK	Weather/Climate Research	NEC SX8/128M16	2005	2.02	98	15.8

Hadley Centre Review 2006: Supercomputing

Final Report

Manchester Computing
The University of Manchester
Manchester M13 9PL
United Kingdom

http://www.manchester.ac.uk/manchestercomputing