

**The vision for networks, data storage systems
and compute capability**

January 2006

Preface

In July 2004 the Treasury, Department of Trade and Industry (DTI) and the Department for Education and Skills (DfES) published the “Science and Innovation Investment Framework 2004-014”, which set out the government’s ambitions for UK science and innovation over that period, in particular their contribution to economic growth and public services.

A section of the Framework addressed the need for an e-infrastructure for research. It proposed an Office for Science and Innovation (OSI) lead a steering group to focus discussion and assess requirements for its development.

To implement this recommendation, the OSI steering group was formed and commissioned a study to help inform the process of achieving the objectives set out in the Framework document. The study was tasked with establishing a high-level “road map” of the current provision of the UK’s “e-Infrastructure” to support research, and in doing so help define the development this infrastructure.

The steering group subsequently formed six working groups to develop this road map of the e-Infrastructure in greater detail in specific areas. These working groups were tasked with producing the following reports:

1. Information creation and data creation
2. Preservation and curation
3. Search and navigation
4. Virtual research communities
5. Networks, compute power and storage systems
6. Middleware, AAA (authentication, authorization, accounting) and digital rights management

The individual reports are intended to represent the informed opinion of the working groups and their contributors and to guide discussion and future development of the e-infrastructure. The working groups have worked closely together. Although each report is free-standing, a synthesis of all the reports and major issues has also been produced which will provide a framework for any potential departmental bids in the next Comprehensive Spending Review in 2007 and for future planning and development of the e-infrastructure for research.

Prue Backway
Office for Science and Innovation
Department of Trade and Industry

Contents

Preface	1
Contents	2
Executive Summary	3
Introduction	6
Networking.....	7
Introduction	7
The present situation.....	7
The aim	8
Key messages.....	9
Computing Capability.....	12
Introduction	12
The present situation.....	12
The aim	13
Key messages.....	13
Data Storage Systems	16
Introduction	16
The present situation.....	17
The aim	18
Key messages.....	19
Concluding remarks.....	21
Appendix 1: Working group contributors	22
Appendix 2: Drivers for Networking.....	23

Executive Summary

This report brings together the work and recommendations of the Networking, Computing Capacity and Data Storage Systems sub-groups formed by the DTI steering group to address the issues and challenges related to the development of a UK e-Infrastructure for research.

Key findings and messages are:

On Networking:

- Diversity: The network infrastructure of the future should not be driven by a conservative approach based upon "one size fits all".
- Services - not just bandwidth: Networks of the future should not be seen by the end user as low level technical offerings (i.e. long lived "bandwidth" contracts), but instead should appear as schedulable and dynamic "services" presented to the end user.
- Integrated component: The network should not be the "afterthought" assumed to be provided somehow by someone else - it should be a first class component of the future e-Infrastructure planning.
- Business model: The current business model of traditional telcos is inhibitory - and needs to be restructured for the good of the UK. *This requires government policy intervention.*
- Mobility: Mobile access is crucial - researchers are often not "at their desk". When on the move pervasive wireless access is demanded along with an international roaming authentication system.
- Research Network: A research network is required on which not only applications can be investigated but also research conducted into the network infrastructure itself.
- International comparison: The UK, USA and Europe lag behind other countries in networking which have leapfrogged us in certain respects. For example several Eastern European countries have much more penetration of fibre infrastructure (e.g. CZ and PL) and the Asia pacific region is ahead on mobile access.

On Compute capacity:

- Capability vs. capacity: At the national level, leading-edge facilities are essential to support both capability computing (targeting substantial fractions of a tightly coupled resource at a single problem) and capacity computing (running many smaller, largely independent jobs simultaneously).
- System Architectures: For both capability and capacity systems, there needs to be a diversity of architectures as one size does not fit all.
- Traversing the Computational Pyramid: The provision of intermediate scale computational resources, between the local university resources and the UK national facilities, and the appropriate policies to allow co-ordinated use of the resource is essential.
- Visualisation: A national UK e-Infrastructure should include visualisation with a set of rendering generators capable of delivering sufficiently high resolution imagery, on demand, to the user.

e-Infrastructure for research in the next 10 years

- **Parallel Software:** Continued speed-up of simulations will be dependent on exploiting parallelism. An initiative in software for massively parallel systems is essential.
- **Going Multicore:** A trend that we will experience over the next decade is the move to multicore systems that is multiple CPUs on a chip. This will have a large impact on software, as new algorithms and programming methodology will need to be developed and supported to enable researchers to use these architectures effectively.

On Data Storage:

- **Data Volumes:** In many research fields data production is doubling annually and the research communities are not prepared, nor are the facilities in place for management of these data volumes.
- **Open Access:** The funding source of the generation of the data is also a driver in terms of storage and long term curation. Data generated through publicly funded activities may soon require long-term access to be mandatory.
- **Legal Requirements:** Both legality issues and co-analysis issues will lead to data maintenance periods becoming mandated by funding providers.
- **Data Storage Technologies:** New technologies such as holographic storage systems may be developed sufficiently for use within the next decade, but at this time we are reliant on disk and tape technology to meet the increasing need.
- **Data management and retrieval:** At the present time it is often the data storage systems that are constraining large scale data movement (rather than the networks).
- **Data Systems Software:** There is a need for a software development and support for effective use of data storage systems – this is largely covered in other reports in this series but noted here for completeness.
- **Stable support:** As communities will have become dependent on massive openly accessible databases, stable financial support will become vital. Both the UK and the EU need to develop funding mechanisms that reflect the real costs and importance and of strategic data repositories (to at least match efforts in the U.S. which have almost twenty years of heritage). Such mechanisms will reflect the real staffing costs associated with the data repositories.

On Cross cutting issues

- **Resilience and self management:** Faults and attacks on networks and other resources will inevitably occur, and so they must be built to deal with this itself.
- **Security:** Work is already underway on security, for example to harmonise authentication and authorization across the world (based on Shibboleth and associated policy extensions). However, security needs to be treated as a system-wide concern that includes co-ordination of policies as well as technologies.
- **Personnel:** There is a need to provide support and training in computational science and data storage technologies and capabilities. This needs to be

e-Infrastructure for research in the next 10 years

enhanced and rolled out across all disciplines as part of any e-Infrastructure strategy.

Key recommendations are:

In creating future network infrastructures:

- I. Fibre infrastructure reaching to most research sites, and giving international access through SuperJANET. This will preferably be accompanied by changes in the provisioning structure that encourages fibre infrastructure providers.
- II. Wireless connectivity and mobility for all, provided seamlessly, is an increasingly important area for new services required by researchers on the move and who will be using a range of access devices of widely different capabilities.
- III. Robust, resilient, self managing networks are basic necessities in order to provide the Quality of Service that users will increasingly expect, even in research environments, and at the same time the networking research needed to achieve this will inform the provision of networks as a crucial and critical part of e-Infrastructures.
- IV. Research Networks need to be provided in order to develop the necessary know-how, and the UK urgently needs to make the necessary investment.

On future Compute Capacity:

- I. Maintain the support at national and University levels to enable the continuing development of the hierarchy of computational resources for capacity and capability computing.
- II. Ensure better co-ordination between funding agencies, service providers and policy developers to ensure effective use of co-ordinated computational resources.
- III. Ensure scientific software development and support is given a priority to enable algorithm development and deployment across new architectures and parallel systems.

On Data Storage Systems:

- I. Develop tools, and policies to support the appropriate use of managed data storage
- II. Provide national storage facilities for reliable storage for very high-volumes of data (greater than tens of PB each).
- III. Provide discipline specific data storage facilities which have optimised their storage to support the access paradigms required by their communities, and exploit secure backups at institutional, regional or national storage facilities.

On cross cutting themes:

- I. Provide effective training for researchers on these technologies. Provide support for appropriate post doctoral training.;
- II. Ensure security infrastructure is co-ordinated across networked resources.
- III. There needs to be coordinated effort of stakeholders in the development of the e-Infrastructure to ensure most effective investment.

Introduction

This report is one of six developed as part of the OSI considerations on the requirements of a UK e-Infrastructure for research. It brings together the work of three sub-groups (listed in appendix 1) who addressed the issues of Networking, Computational Capacity and Data Storage Systems separately. Although each of these issues is sufficiently well defined that they can be considered in isolation, there are clearly some issues of integration of the three technologies in supporting research that need to be considered. There are a number of areas particularly in data storage where the material covered in other reports in this series, on the creation to curation of data, are eminently relevant as are findings within the report on middleware and security. There are also a number of cross-cutting themes that are not considered here or in the individual reports, but which are covered in the overview report.

For the UK to maintain, enhance and promote its international position in computational science and related research disciplines it is essential and necessary that UK researchers are provided with an e-Infrastructure which enables an integrated approach to computing, data resources and research facilities across the UK and internationally. These resources should include the hierarchy of computational resources from desktop PCs to high-performance computers (HPC), data storage, visualization, underpinned high-performance network services. Research practice has changed from a predominantly local environment to one where leading edge research comes only through a global virtual collaboration and hence the provision of computational, data and network resources needs to be considered in this context. Problems with the present situation have been identified in the recent International review of HPC¹:

“Perhaps due in some measure to present funding practices, many research groups in the UK appeared to be deficient in adapting to modern programming practices, anticipating the need for data management and visualization, and developing interoperable software and algorithm environments. If these deficiencies were remedied, we should witness the accelerated emergence of dramatically improved capabilities for making critical discoveries in many scientific domains ...”

In this report we attempt to pick up on these already identified problems and attempt to provide a view of requirements for the next decade. This report is meant as a starting point for discussion and further consideration.

The next sections provide the findings for each area, followed by concluding remarks.

¹<http://www.epsrc.ac.uk/ResearchFunding/FacilitiesAndServices/HighPerformanceComputing/InternationalReview/default.htm>

Networking

Introduction

Research practice has changed from a predominantly local environment to one where leading edge research comes only through a global virtual collaboration. To enable researchers in the UK to participate in such global collaborations the UK must have a world class, pervasive, dynamic and resilient network capable of meeting the needs of diverse and demanding research use patterns. Virtual researchers can operate in many places concurrently, yet network provision is still primarily based on the assumption that researchers are statically located in one place. The vision must be to remove the barriers that separate one researcher from another, that distance the researcher from a remote instrument, and that also inhibit researcher mobility. To achieve this vision the network needs to be based on fibre connections to research sites combined with pervasive wireless access. It should provide services ranging from normal (internet) IP to dynamic path provision; it should shed the outmoded telco-like business model, and provide ubiquitous and mobile access services.

This quote below captures succinctly the challenge that this vision addresses.

"We are impressed by the evidence of growing network requirements from all areas of research. The needs will grow dramatically over the next 5-10 years, in all disciplines and in all countries. Examples are given where research would become much more efficient if network speeds were increased by one or two orders of magnitude. Examples were also given where there is the possibility of starting completely new research activities that were prohibited until now by lack of very high performance networking facilities..."

-[The SERENATE Report for the European Union]

The present situation

The UK national education and research network, SuperJANET4 (SJ4) is provided by the JISC through UKERNA. SJ4 provides a first class IP production service with ample capacity for IP traffic. This is connected to Universities and other sites through Regional Networks (RN), which is adequate for the normal internet traffic. Some Institutions (e.g. Royal Holloway, Brunel or Stirling University) still suffer from low capacity connections due to prohibitive leasing rates - in turn inhibiting network intensive research from being developed. Some research sites (e.g. Jodrell Bank) have until recently suffered from the same problem and have only been solved on a local basis.

The JISC and UKERNA have worked very closely with the UK research community in the last few years, building the foundations for a much more dynamic network structure, and resulting in the UK's first research network (UKLight) which provides connectivity into the global optical network. This has been an excellent precursor to SuperJANET5 (SJ5).

e-Infrastructure for research in the next 10 years

SJ5 is about to be rolled out in the UK. UKERNA are to be commended on SJ5 which was designed to meet the evolving needs of its entire community, ranging from education and teaching to high demand research. SJ5 provides a managed fibre infrastructure which reaches all RNs and gives the capability of dedicated bandwidth services in a flexible and cost effective way.

The SJ5 rollout has not solved the problem of some institutes connected via low capacity links to their RNs, and at present this problem is not likely to go away due to continuing inhibitory "premium" costs and lack of competition on non metropolitan routes.

Thus at present the UK stands on the threshold of an excellent physical research infrastructure, but it has incomplete reach.

Other services described below are in some cases just being put in place (e.g. roaming network access has just been piloted, and Shibboleth authentication and authorization is being deployed). There is very little by way of higher level services under development.

The aim

The Network e-Infrastructure should be driven by the idea of ensuring excellence, quality, capability and competitiveness in UK research. Science and research should not be artificially inhibited through lack of network capability. In particular the above should not be the case purely due to outdated provisioning models.

The network must be viewed not just at technical capability at the infrastructure level but also as services expressed in the terms encapsulates what users actually want to do. This will be ever more true as distributed computing evolves to treating interconnectivity services in the same way as job scheduling services.

At the technical level the target should clearly be to achieve as close as possible to 100% fibre infrastructure to all major research sites². This will ensure that the fundamental infrastructure will not limit UK research and collaboration. To achieve this aim there should be a policy to encourage the emergence of more "utility-like" fibre providers and the provision of open access "city fibre exchanges" where all fibre providers, leasers and renters in a region can interconnect.

UK researchers should benefit from truly pervasive roaming mobile access following from the first steps which are being taken today. The ad hoc and

² The UK is part way there with the SuperJANET5 infrastructure, but this does not reach all research end points and several institutions still have only 155 Mbit/s connections.

e-Infrastructure for research in the next 10 years

disparate ways of providing on-the-move access should be consigned to history as soon as possible in recognition that research is not done in one place. This redoubles the need to achieve a global authentication and authorization system to enable sign-on and access to resources equally from anywhere

Using this infrastructure, services can be built which are fit for purpose, rather than employing fix-ups to make up for the deficiencies of the underlying structure and business model. These range from easy set up and tear down of end-to-end quality and bandwidth-assured services tailored to different classes of application, through reliable high speed integrated data movement and management services, to voice-over-IP based telephone conferencing and multimedia conferencing that are easy to set up and use.

The network should provide a flexible crash and burn facility for network researchers, which in turn can be used to both develop new ways of networking, as well as undertake realistic tests with real users.

Much (or all) of this could be provided and managed by UKERNA who already engage fully with the research community and appreciate these issues. UKERNA can only do this, however, if policy is set at government level to allow the vision to be enacted.

Key messages

- **Diversity:** The network infrastructure of the future should not be driven by a conservative approach based upon "one size fits all". This completely inhibits development of new ways of working which could potentially cause a sea change in efficiency and capability. Instead the network should permit many and diverse use patterns to co-exist, ranging from real time instrument control through to the traditional packet based many to many services.
- **Integrated component:** The network should not be the "afterthought" assumed to be provided somehow by someone else - it should be a first class component of the future e-Infrastructure planning.
- **Business model:** The current business model of traditional telcos is inhibitory - and needs to be restructured for the good of the UK. Existing telcos are vertical structures where the telco owns the entire stack. As a result they feel the need to (by default) charge for high yield high level services appropriate to business, but wholly inappropriate to research infrastructure. They fail to recognise that research projects will never need or buy the expensive high level services. As example research rarely needs the 99.9999% availability of a premium SDH circuit, but would really like a 98% available wavelength.

e-Infrastructure for research in the next 10 years

This requires government policy intervention.

- **Services - not just bandwidth:** Networks of the future should not be seen by the end user as low level technical offerings (i.e. long lived "bandwidth" contracts), but instead should appear as schedulable and dynamic "services" presented to the end user. These services should be presented in the terms most natural to users needs, e.g. *Virtual Visualisation service* to a set of remote researchers so that they can all see and manipulate MRI images and simulations as if they were in the same office, or *Guaranteed high speed file transfer service* to allow a computational chemist to guarantee that a simulation checkpoint data file is moved to an expensive HPC machine in time to use a scheduled processing slot.
- **Mobility:** Mobile access is crucial - researchers are often not "at their desk". When on the move pervasive wireless access is demanded along with an international roaming authentication system. Wireless networking needs to be better integrated into the whole, with solutions to its attendant problems of intermittent connectivity and variable-quality bandwidth as well as heterogeneous technologies.
- **Security:** Much work is already underway, e.g. to harmonise authentication and authorization across the world (for example based on Shibboleth and associated policy extensions). However, security needs to be treated as a system-wide concern and work undertaken to minimize the possibility of intrusions and their consequences.
- **Resilience and self management:** Faults and attacks will inevitably occur, and so the network must be built to deal with this itself. To support critical applications, the network will have to be more self-managing than it is now - this implies a degree of autonomic operation. Ultimately, work in this area aims to improve the user-level Quality of Service (QoS) through the provision of a resilient network infrastructure.
- **Research Network:** The technologies of overlays and programmable technologies (and others as appropriate) need to be harnessed into producing a research network on which not only applications can be investigated but also research conducted into the network infrastructure itself.
- **International comparison:** The UK, USA and Europe lag behind other countries which have leapfrogged us in certain respects. For example several Eastern European countries have much more penetration of fibre infrastructure (e.g. CZ and PL) and the Asia pacific region is ahead on mobile access.

e-Infrastructure for research in the next 10 years

The European Union has recognised this in its SERENATE report - a strategic study aiming at providing input to the European Commission on initiatives targeted to keep the evolution of the European research networking at the forefront of world-wide development. The European network provider is now actively offering wavelength services and is engaged in developing user to network interaction services.

<http://www.serenate.org/>, <http://www.geant2.net/>

Very recently, a new project has been approved by the European Union, called "OneLab" (An Open Networking Laboratory Supporting Communication Network Research Across Heterogeneous Environments), under the Research networking testbeds initiative (activity code IST-2005-2.5.6). OneLab will build on the successful PlanetLab testbed, incorporating into that platform both passive and active measurements across multiple link technologies, and so enabling the next generation of network experiments and prototyping.

The US research community has recognised this, and as example has for some time had several consortia run groups providing multiple wavelength access to institutions (e.g. StarLight, NLR, CENIC). The NSF has set in place the GENI programme which will support research, design, and development of new networking and distributed systems capabilities by creating new core functionality, developing enhanced capabilities, deploying and validating new architectures, building higher-level service abstractions, building new services and applications, and developing new network architecture theories.

<http://www.nsf.gov/cise/geni/>

The world community has recognised this, as can be seen from the emergence of the Global Lambda Infrastructure Facility group (GLIF). It brings together some of the world's premier research networking providers and users who are working together to develop an international infrastructure by identifying equipment, connection requirements, and necessary engineering functions and services

<http://www.glif.is>

Computing Capability

Introduction

Computational science, the scientific investigation of physical processes through modeling and simulation on computers, has become generally accepted as the third pillar of science, complementing and extending theory and experimentation.

The introduction and utilisation of computational resources is now indispensable to the solution of research challenges and problems in every sector, from traditional science and engineering domains to such key areas as national security, public health, and economic innovation. This trend is set to continue over the next decade. For the UK to compete in the global economy and international research arena, it must take advantage of its excellent track record in scientific discovery and technological innovation and provide an e-Infrastructure for this to continue and expand.

High performance computing (HPC) satisfies the most demanding scientific goals pushing the boundaries of researchers' ambitions and stretching the development of hardware and software technologies in dimensions that often prove beneficial outside the HPC research arena. This enables computational science to tackle a class of problems and grasp opportunities that cannot be approached in any other way. These opportunities include genome sequencing, biomedicine, and drug design; molecular engineering of new materials and devices; engineering efficient and environmentally sound technologies and energy sources; forecasting global climate change and finding strategies to deal with it; and understanding the fundamental interactions of elementary particles and the origin and fate of the universe.

Through the emergence of "*Grid technologies*", which provide a framework for integration across computational resources, it is becoming increasingly important to consider issues that impact co-ordinated use of the resources. This includes security, usage policies, resource booking practice, and software compatibility and licensing. Since the development of a UK e-Infrastructure will build on these technologies then the related co-ordinating issues will be key to effective use of computational resources in the future. The reports on AAA and middleware cover these issues in depth but they should also be kept in mind here.

The present situation

The hierarchy of computational resources is often captured in the notion of a pyramid with a broad base (many smaller systems) and a narrow top (a few leading-edge systems) usually classified as follows:

- Leadership computing: Capability Computing
- National supercomputers: High-end Capability-Capacity Computing

e-Infrastructure for research in the next 10 years

- University-Department Group Compute Clusters: Local Capacity Computing
- Personal computers and workstations: Personal Computing

The UK to date has been well served at the leading edge by the HPCx and CSAR computational systems and by the special-purpose QCDOC machine. and The HECToR program promises to build on this excellent HPC experience. SRIF and other similar funding mechanism have to date enabled Universities to invest in Tier 2 type systems for local use.

In the last two years the National Grid Service (NGS: www.ngs.ac.uk) has been developed that provides a Core capability of what might be considered individual Tier 3 systems, but that are connected together to provide a larger computational "Grid" resource. The NGS is expanding by the addition of University and other resources. This trend is one which will shape the requirements of the next decade and beyond.

The aim

The aim over the next decade must be to maintain the support at national and University levels to enable the continuing development of the hierarchy of computational resources. The co-ordinated use and development of those systems will require increased co-ordination at the funding level, the service provision level and at the policy level. It is likely that the international co-ordination requirements will also increase.

In the future exploitation of computational systems, software and the support of software systems will be as important as the hardware; ranging from new algorithms (to maximise parallelism exploited) to community codes that are portable efficiently across several architectures. To ensure accessible and scaleable software across computational resources will require new approaches.

Similarly the integration of computational resources with data storage systems, together with analysis and visualization systems is key to science outputs. Any vision of compute power for the future must be accompanied by a matching vision of power to interpret the results obtained. Indeed the greater the compute power, the greater (potentially) the volume and complexity of the resulting output.

Key messages

Capability vs capacity: At the national level, leading-edge facilities are essential to support both capability computing (targeting substantial fractions of a tightly coupled resource at a single problem) and capacity computing (running many

e-Infrastructure for research in the next 10 years

smaller, largely independent jobs simultaneously). A balance between these two types is essential because both are fruitful scientifically and are often needed on the same scientific project, for example; generation of simulations may require capability computing, while the analysis of these simulations may best be done on capacity machines.

System Architectures: For both capability and capacity systems, there needs to be a diversity of architectures as one size does not fit all. The UK needs at least tightly coupled distributed-memory machines and large shared memory systems with special allowance for special-purpose machines (e.g., QCD machine). While all of these machines operate primarily in standard batch mode, they should be scheduled to allow interactive access and even co-scheduling with other facilities (e.g., for visualisation).

Traversing the Computational Pyramid: The provision of intermediate scale computational resources, between the local university resources and the UK national facilities, is essential. Today, there are several barriers that stand in the way of researchers making effective use of high end computing. One is that the transition from local/departmental resources to high end national resources is far from seamless (the potential energy barrier is high).

Hence a strategy of providing lots of capability machines across the UK, with the proviso that a percentage of their activity must be given up to provide UK-capacity, is an appropriate component for a model of a UK e-Infrastructure. The economies of these systems is also changing, for example, with the increasing market share of cluster systems built from commodity, air-cooled components, the floor space requirements of systems are getting bigger again. Heat output is going up, cooling is becoming a problem, and power bills can no longer be ignored in an energy concerned world.

Visualisation: A national UK e-Infrastructure should include visualisation with a set of rendering generators capable of delivering sufficiently high resolution imagery, on demand, to the user. (At present, high resolution (4k x 2k) is only available on desktop but the expected resolution of laptops and mobile devices will steadily increase over coming years.) The e-Infrastructure will therefore require visualisation power – a rule of thumb might be that for every ten units of compute power one unit for visualisation is required.

Security: As noted above security needs to be treated as a system-wide concern and as more computational resources become part of the e-Infrastructure it is important to ensure that a comprehensive security policy framework exists.

Parallel Software: While there is a good chance that processor speed and bandwidth will continue to follow Moore's law, network latency is more or less stuck. Developing scalable codes is getting harder. As a result continued speed-up will be dependent on exploiting parallelism.

e-Infrastructure for research in the next 10 years

Going Multicore: A trend that we will experience over the next decade is the move to multicore systems, that is multiple CPUs on a chip. Desktop machines are already available with dual core processors and manufacturers have produced roadmaps indicating up to 16 core processors in the relatively near future. This will have a large impact on software, as new algorithms and programming methodology will need to be developed and supported to enable researchers to use these architectures effectively.

Personnel: The most successful practitioners of computational science today are multidisciplinary teams that include, for example, mathematicians and computer scientists. These teams have set up a software infrastructure, including a support infrastructure, for large codes that are well maintained and extensible beyond the set of original developers. This needs to be enhanced and rolled out across all disciplines as part of any e-Infrastructure strategy.

Data Storage Systems

Introduction

In the next decade we will see new experimental facilities coming online that will generate data sets ranging in size from 100's of Terabytes to 10's of petabytes per year. Such enormous volumes of data exceed the largest commercial databases currently available by one or two orders of magnitude. There are many instances where sensors are generating petabytes of data from engine pressure, temperature and vibration sensors checking the health of an aircraft, to high-throughput sensor experiments in fields such as environmental and earth observation, and of course human health-care monitoring.

Data generated from sensors, satellites, high-performance computer simulations, high-throughput devices, scientific images and so on will soon dwarf all of the scientific data collected in the whole history of scientific exploration. Consider the production of X-ray data by the present generation of electron synchrotron accelerators. At 3 seconds per image and 1,200 images per hour, each experimental station generates about 1 Terabyte of X-ray data per day. At the next generation 'DIAMOND' synchrotron currently under construction the planned beamlines will generate many Petabytes of data per year, As of today, the largest commercial databases range from 10's of Tbytes up to 100 Tbytes. In the coming years, we expect that this situation will change dramatically in that the volume of data in scientific data archives will vastly exceed that of commercial systems

We should not forget either the data sets generated and analysed within the social sciences or the digital repositories being built in the humanities. In the UK, the total storage requirement for the social sciences has grown from around 400 Gigabytes in 1995 to more than a Terabyte in 2001. Growth is predicted to grow in the next decade but the total volume is not thought likely to exceed 10 Terabytes by 2010.

There are many issues that need to be considered including the generation of appropriate metadata and curation technologies. Here we concern ourselves only with the technology to store the data.

In all communities, data volumes of PB/year are expected within the next decade:

- In the environmental sciences, new satellite instruments and more use of ensembles of high resolution models will lead to a number of multi-PB archives in the next decade (within the U.K. the Met Office and European Centre for Medium Range Weather Forecasting already have archives which exceed a PB each. Such archives will need to be connected to the research community with high-speed networks.

e-Infrastructure for research in the next 10 years

- In the astronomy community, there are now of the order of 100 experiments delivering 1 or more TB/year with the largest at about 20 TB/year, but in the near future the largest will be providing 100 TB/yr and by early in the next decade PB/year instruments will be deployed.
- In the biological sciences, new microscopic and imaging techniques, along with new sensor arrays and the exploitation of new instruments (including the Diamond Light Source) are leading to, and will continue to, an explosion in data production.
- Within the particle physics community, the need to exploit new instruments at CERN and elsewhere is leading to the development of new storage paradigms, but they are continually on the bleeding edge, both in terms of software archival tools, and the hardware which is exploited.
- In general the scientific research that uses the computational resources described above will create TB of data through simulation.

Legal requirements to keep research data as evidential backup will become more prevalent. Communities are recognising the benefits of meta-analyses and interdisciplinary analyses across data repositories.

With the plethora of instruments and simulation codes within each community, each capable of producing different forms of data, heterogeneity of data types coupled with differing indexing strategies will become a significant problem for cross-platform analyses (and the concomittant data retrieval). The problem will be exacerbated for interdisciplinary analyses.

The present situation

In many fields data production is doubling every two years or so, but there are a number of fields where in the near future, new generations of instruments are likely to introduce major step changes in data production. Such instruments (e.g. the Large Hadron Collider and the Diamond Light Source) will produce data on a scale never before managed by their communities.

Thus far disk storage capacity increases have met the storage demands, and both tape and disk storage capacities are likely to continue to increase although there are always concerns about being on the edge of technological limits. Most present technologies are designed to a specific bit error rate e.g. 1^{e-15} , for modern petabyte disk farms and multi-petabyte archives; this error rate is too high. A further challenge is that errors are bursty - that is a failing system tends to have multiple errors. For large storage systems, the time to recover from a failure is becoming too short for effective management.

New holographic storage systems are beginning to hit the market, but thus far are not yet scalable nor fast enough to compare with older technologies. In addition much of the demand for new and increased capacity storage is driven by mass market consumer applications where the performance and reliability

e-Infrastructure for research in the next 10 years

demands may be much lower than required for large scale storage of valuable data.

Consequently, while storage capacities are likely to meet demand, major problems are anticipated in both the ability to retrieve and manage the data stored. Although storage capacities and network bandwidth have been roughly following Moore's Law (doubling every two years), neither the speed of I/O subsystems nor the storage capacities of major repositories have kept pace. Early tests of data movement within the academic community have found in many cases that their storage systems have been the limiting factor in moving data (not the network). While many groups are relying on commodity storage solutions (e.g. massive arrays of cheap disks), as the volumes of data stored have gone up, random bit errors are beginning to accumulate, causing reliability problems. Such problems are compounded by many communities relying on ad hoc storage management software, and few expect their solutions to scale with the oncoming demand. As volumes go up, finding and locating specific data depends more and more on sophisticated indexing and cataloguing. In addition to experimental facilities, existing High Performance Computing facilities do not adequately provide for users whose main codes produced huge volumes of data. There is little thought given to the cost of creating large datasets, or recreating them if they are lost. E.g. There is little provision of backup for large research datasets, and widespread international access is typically limited to links through major funded international projects.

The aim

In ten years, the UK community should expect to have available both the tools, policies and the infrastructure to allow scientists to exploit the deluge of data expected. Creation, analysis, storage and archival of data will be possible across the full range of scientific endeavour, with appropriate actions being determined by conscious tactical, strategic or policy decisions and not on ad hoc technical or financial opportunities. The infrastructure is likely to consist of

- I. Tools, and policies to support the appropriate use of managed data storage
- II. One or more national storage facilities that provide reliable storage for very high-volumes of data (greater than tens of PB each).
- III. A number of discipline specific data storage facilities which have optimised their storage to support the access paradigms required by their communities, and exploit secure backups at institutional, regional or national storage facilities.

All these storage facilities will all have links with international peers that enable international collaborations to exploit distributed storage paradigms without explicitly bidding for linkage funding. Such links will include both adequate

e-Infrastructure for research in the next 10 years

network provision and bilateral policies on the storage and access of data for international collaborators.

The UK data analysis facilities will have large and efficient I/O subsystems that can support data intensive high performance computing, and will be linked by high performance networks (possibly including dedicated bandwidth networks) to national and international data facilities.

Researcher will use sophisticated software tools to efficiently manage large archives on commodity hardware, and on the development of methodologies to improve the bandwidth to storage subsystems, making cost effective large scale data storage a reality. At the same time both distributed searching (across archives) and server-side data selection (to minimise volumes of data transfer) will available as the demand for access to storage will have continued to outstrip the bandwidth.

Many communities will have developed automatic techniques to capture metadata, and both the data and metadata will be automatically mirrored to national databases and potentially available to a wider community even as they are exploited by individual research groups.

As communities will have become dependent on massive openly accessible databases, stable financial support will become vital. Both the UK and the EU will have developed funding mechanisms that reflect the real costs and importance and of strategic data repositories (to at least match efforts in the U.S. which have almost twenty years of heritage). Such mechanisms will reflect the real staffing costs associated with the data repositories.

Key messages

Data Volumes: In many research fields data production is doubling annually and the research communities are not prepared, nor are the facilities in place for management of these data volumes.

Legal Requirements: Both legality issues and co-analysis issues will lead to data maintenance periods becoming mandated by funding providers.

Technologies: In the near term we see the existing technologies of disk and tape drives as needing to expand to meet the demand. It is possible that new technologies such as holographic storage systems will be developed sufficiently for use within the next decade.

National Storage facilities: There is a need for national facilities to support research community data storage of greater than tens of PetaBytes each. These will need support for staffing as well as the required hardware.

e-Infrastructure for research in the next 10 years

Data management and retrieval: This topic is covered in depth in other reports in this series, but it should be noted here that at the present time it is often the data storage systems that are constraining large scale data movement (rather than the networks).

Personnel: There is a need to provide support and training in data storage technologies and capabilities to communities who are creating or on the brink of creating large amounts of research data.

Software: There is a need for a software development and support for effective use of these data storage systems – this is largely covered in other reports in this series but noted here for completeness.

International relations:

Stable support: As communities will have become dependent on massive openly accessible databases, stable financial support will become vital. Both the UK and the EU will have developed funding mechanisms that reflect the real costs and importance and of strategic data repositories (to at least match efforts in the U.S. which have almost twenty years of heritage). Such mechanisms will reflect the real staffing costs associated with the data repositories.

Open Access: The funding source of the generation of the data is also a driver in terms of storage and long term curation. Data generated through publicly funded activities may soon require long-term access to be mandatory.

Concluding remarks

Several cross-cutting issues have come up throughout the document. These include security, co-ordination across resource providers, interoperability of software systems, and support for research communities in using these resources. Education and training of computational scientists is a necessary part of any e-Infrastructure strategy. In addition this requires that there is an academic career path for the interdisciplinarians who are involved in computational science.

There needs to be much more co-ordination of funding and policy development of the Research Councils, JISC, HEFCE and other stakeholder agencies. to ensure an effective and efficient development of an e-Infrastructure that will truly support research. There needs to be more cohesion between the way e-Infrastructure, HPC, network and other resource provisions are funded, and the mechanisms for granting access to these resources.

Such an e-Infrastructure should be sustainable, that is, receive an investment over a long period to ensure stability for UK researchers. It takes a lot of research investment in building up the capabilities to effectively use these resources and that investment will be wasted if the e-Infrastructure is not successfully maintained.

In summary the UK needs a strategy where the e-Infrastructure provides the integration fabric for a wide range of computational resources including computers, visualization, data storage and the underpinning network capacity and services. The strategy must enable transitions across, and allow co-ordinated use of the whole spectrum of resources, from desktop and departmental systems through campus and regional or functional resource pools/grids and in co-ordination with international infrastructure developments.

Appendix 1: Working group contributors

Report Editor: Anne Trefethen

Network:

- Co-ordinator: Peter Clarke
- Contributors:
 - o I Andonovic
 - o J.Crowcroft
 - o D.Hutchison
 - o R.Tasker
 - o D.Williams.

Computational Systems Contributors:

- Co-ordinator: Ron Perrott
- Contributors:
 - o Richard Kenway
 - o Peter Lee
 - o Peter Coveney,
 - o Stephen Pickles
 - o Simon Cox
 - o Ken Brodlie
 - o Neil Geddes

Data Storage Systems:

- Co-ordinator: Bryan Lawrence
- Contributors:
 -

The individual reports have also been circulated to several groups including the e-Science Technical Advisory Group, the e-Science Directors forum, the HPC SIG and comments from individuals in those groups have been incorporated into this report.

Appendix 2: Drivers for Networking

An excellent cross section of applications which are drivers for change can be seen through the IGrid series of conferences. Examples from the IGrid 2005 meeting in San Diego can be seen in the brochure which can be found at

<http://www.igrid2005.org/program/>

Drivers for change in applied research include:

- Very high bandwidth point-to-point connections empower a range of compute intensive research:
 - Compute and Data centres (HPC centres, Research Data centres) need guaranteed co-scheduling of expensive distributed resources - including the network interlink.
 - Large research facilities & instruments (CERN, Telescopes, Microscopes, Simulators,
- Virtual communities. Leading edge research is more and more achieved through massively distributed collaborations on a global scale - "virtual communities" (VCs). Often these VCs access instruments or facilities which are remote from the people. There is no need or desirability for people and instruments/data centres to be co-located. This requires:
 - Ubiquitous roaming wireless access - more and more researchers are "on the move"
 - Truly interactive virtual presence
 - Very high definition video
- Dependability and predictable connections are needed for
 - Medical research where researchers are remote from an expensive imaging instrument.
 - Remote Instrument control
 - Supercomputer operations (links simulations, interactive visualisation to reduce time to insight).
- Secure services are demanded by:
 - Medical applications
 - Industrial collaborative projects

Drivers for change in communications research include:

- Many elements of the internet are past sell by date (like TCP, routing protocols)
- Optical technologies are advancing - requiring advances in dynamic service provisioning protocols and realistic testbeds.
- The application of overlay and P2P techniques, and in general the adoption of programmable networking, offer promise for future networks.

Drivers for change in industry ↔ university research

e-Infrastructure for research in the next 10 years

- Electronic distribution represents an enormous opportunity for games/broadcast industries in the UK
- No extensive broadband infrastructure in the UK that facilitates the development, test and demonstration of new on-line concepts
 - to develop the know-how, capability and protocols to support electronic distribution.
 - quality of service, particularly issues associated with latency in situations where the aggregated data levels are high.
 - video on demand and an extended mix of video services
 - studying the impact of storage - mirrored or more distributed
 - resilience and how this may make services unattractive
 - the development of tool kits to support a new distributed enterprise culture, enabling new workflows and businesses to develop

Drivers for change in society

- Pre-emption of massive computing resources to cope with disasters (storms, eruptions, building fires)
- Computer aided medical diagnostics and surgery.