The GRID: Towards a European Metacentre

ENACTS, an EC-funded Cooperation Network in High Performance Computing, presents a series of in-depth scientific and technological studies on the various aspects of the Grid. The six studies were performed in 2002-2004 and form a basis for a Grid Metacentre Demonstrator Application. The aim of the Demonstrator is to draw together the results from all the studies and evaluate their practical consequences for operating a pan-European metacentre and constructing a best-practice model for collaborative working amongst individual facilities. The Demonstrator feasibility report is available at http://www.enacts.org/demonstrator.pdf

Grid Service Requirements

HPC Technology Roadmap

Grid Enabling Technologies

Data Management

Software Efficiency & Reusability

Distance Learning and Support

ENACTS Network Partners:

- EPCC, United Kingdom
- ICCC Ltd, Czech Republic
- UNI-C, Denmark
- CSC, Finland
- ENSL-Lyon, France
- TCD, Ireland
- FORTH, Greece
- CINECA, Italy
- CNR-ISTM, Italy
- Parallab, Norway
- PSNC, Poland
- UPC, Spain
- University of Linköping, Sweden
- ETH-Zurich, Switzerland

Grid Service Requirements

The Authors:
EPCC, University of Edinburgh, UK, www.epcc.ed.ac.uk
PSNC, Poland, www.psnc.pl
January 2002

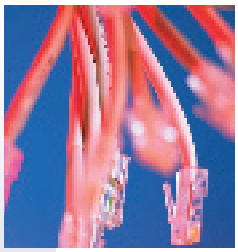

Computational grids can be complex and difficult to use, but their usability can be enhanced by the provision of simple transparent user services. The objective of this study is to specify the level and quality of services users require from a computational Grid.

One of the most promising solutions to the problem of opening up access to Large Scale Facilities is the formation of Computational Grids. These are distributed, high-performance computing and data infrastructures that manage geographically and organisationally dispersed, heterogeneous resources. There is considerable consensus that "The Grid" provides a blueprint for harnessing and exploiting truly distributed computing. Grids must provide common interfaces and access mechanisms, standardised management and uniform user authentication and authorisation for distributed resources. It is important that they offer a comprehensive and consistent set of tools and services for accessing and managing dynamic collections of widely distributed resources: CPUs, storage systems, communications systems, real-time data sources. The term "middle-ware" is used to refer to the glue which makes these distributed resources transparently available to the user.

Crucial to achieving this vision is the need to understand:

- what services are needed to support the user's activities;
- parameters by which the quality of these service can be measured;
- for each group of users an indication as to the minimum quality of service they require for a grid to be useful.

For example consider data storage services with the storage capacity available a quality parameter. Different groups of users will have different storage capacity needs, a gigabyte may be sufficient for some, where others may need petabytes. The report was written in consultation with potential users from a representative selection of groups.

Service categories covered by the report include:

- Job submission, monitoring and runtime control -- what sort of services need to be provided for monitoring and control of executing jobs (e.g. computational steering, intermediate result access, interrupt and restart.)
- Visualisation and post processing -- including local and remote access to specialist facilities.
- Data management and storage -- including access and movement of data between the compute resource, storage facilities, post processing resource and user.
- Authentication and security -- preventing non-users accessing the resources, preventing rival groups accessing each others results etc.
- User Support -- what level of support is needed, can it be provided remotely?
- Accounting -- including all the services required to attach a cost to any usage of the grid.
- Collaboration environments and communications -- services to support user collaboration and interaction (e.g. web servers, collaborative programming environments, voice and video conference facilities)
- Remote access services -- services for coupling large-scale computing and data systems to scientific and engineering instruments.

The report is available to download:
http://www.enacts.org/gridservice.pdf

HPC Technology Roadmap

The Authors:
NSC, University of Linköping, Sweden, www.nsc.liu.se
CNR-ISTM, University of Perugia, Italy, www.istm.cnr.it

April 2002

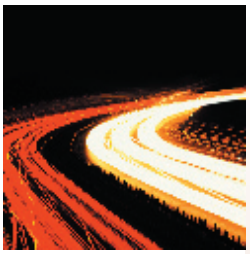

The objective of the study was to determine the likely impact of current technological and economic trends on the hardware architectures of future HPC systems and applications software.

Massively Parallel Processors (MPP) use standard RISC-based microprocessors connected together with proprietary networks and have been the mainstay of high-end computing for the last 5 - 10 years. However, despite advances in software standards to improve code portability, they provide a complex programming model and consequently a entry-level barrier for new users. The small marketplace for such products, coupled with the, still, significant research and development costs probably means that the end is in sight for MPP.

The key question is what will the next generation of HPC systems look like? Will they be distributed memory PC/workstation clusters? Will they be clusters of shared-memory servers - or will the bottom drop out of the high-end market and academics will have to look to custom-built machines again? These are some of the questions which this ENACTS study panel will ask. The topic is very timely due to the very different approaches being taken in the US and Japan - the world's largest producers of computer hardware:

- US: in 1998, the President's IT Advisory Committee (PITAC) found that Federal investment in IT was inadequate (even though it is probably the biggest driver of the economy) and recommended a longer-term strategy to build, and make available for academic use, high-end systems - culminating in 1 Pflops by 2010, though specialised machines such as IBM's BlueGene may attain this performance in niche areas as early as 2005. Their approach to meeting this goal is to use commodity processors and communications to build distributed memory systems. The rationale is that economies of scale and rapid improvement of commodity parts, in contrast to the smallness of the high-end market, make a customised approach unprofitable. Projects such as ASCI (Nuclear Weapons Stewardship), are pumping research and development money into US companies to underpin this approach, but meeting their goals will require aggregating components on an unprecedented scale. The result is that not only will the US have the most powerful computers, but its hardware industry will have developed state-of-the-art interconnect technology.

- Japan: while the Japanese agree that commodity-based distributed memory systems are economically attractive, useful for some applications, and may become more attractive in future as software and tools are developed, the Japanese vendors are sticking with customised vector processors. Japanese vendors believe these systems still have the edge and offers superior performance for HPC applications. This will be resolved by market forces and the rate at which commodity components develop.

In Europe we are a consumer of this technology and are therefore free to make the most appropriate choice for the application. But in making that choice we must not head down blind alleys. In this study, ENACTS attempts to investigate the critical technologies and trends and write a Sectoral Report on the HPC technology roadmap over the next 5 - 10 years and indicate the basis for a future technology scoping project.

The report is available to download:
http://www.enacts.org/hpcroaadmap.pdf

Grid Enabling Technologies

The Authors:
IESL-FORTH, Greece, www.forth.gr
ETH-Zurich, Switzerland, www.ethz.ch
December 2002

The objective of the study was to evaluate the available technologies for implementing a computational Grid. This study uses the service definitions from the ENACTS Grid Service Requirements Study to evaluate different technologies which could be used to implement these services. This study also attempts to ascertain which technologies show the greatest promise and what important services cannot currently be supported.

The development of applications and tools for high performance computational Grids is complicated by the heterogeneity and frequently dynamic behaviour of the underlying resources; by the complexity of the applications themselves which often combine aspects of supercomputing and distributed computing and by the need for frequent node communication.

A low level tool kit can provide basic mechanisms such as communication, authentication, network information and data access. Mechanisms are used to construct higher-level metacomputing services such as parallel programming tools and schedulers. However, users also need an integrated set of higher-level services that enable applications to adapt to heterogeneous and dynamically changing metacomputing environments.

Chapter 2 of the Report gives **an overview of the most commonly used middleware for implementing a Grid: GLOBUS, UNICORE and Legion,** with their main characteristics and their differences and possible future developments**.**

Chapter 3 presents tables of **testbeds of computational and data grids**, and **applications,** together with a short description of their history and functions. The analysis of application types and categories of scientists who use them is given, along with an attempt to point out future needs.

Chapter 4 deals with the **current status of Grids** in Molecular Sciences: Physics, Chemistry and Biology. It turns out, that applications which can run on the existing computational Grids are at a very preliminary stage. New algorithms should be developed which do not require extensive communication between the nodes. Biophysics and high energy physics are the most promising fields for developing such algorithms.

Chapters 5 describes the authors' **personal experience of installing GLOBUS** and **UNICORE** on local network, thus constructing small local computational Grids, which allowed for a few applications to be run. Although the significance of the computational Grids is based in their large size, the present experience showed that the existing middleware such as GLOBUS provides tools for unifyinga large number of local area computers.

The report is available to download:
http://www.enacts.org/gridenabling.pdf

Data Management and Assimilation

The Authors:
Trinity College Dublin, Ireland, www.tcd.ie
CINECA, Italy, www.cineca.it

October 2003

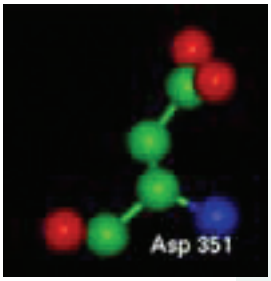

The objective of this project is to gain an understanding of the problems associated with storing, managing and extracting information from the large datasets increasingly being generated by computational scientists, and the technologies, which could address these problems. The report also addresses the integration of data management activity in a GRID environment.

The archiving, preservation, retrieval and exploration of data is of growing importance as computer systems enable the rapid generation of extremely large and complex data sets. Users typically have access to three levels of compute services - a national HPC facility, a dedicated front-end or local departmental server and a personal workstation - and require to move data between these distributed systems to perform post-processing and analysis. Increasing the ease and efficiency of data transfer and access would greatly enhance the amount of science that can be done on HPC facilities and allow for inter-disciplinary sharing of data. Providing distributed data services requires the ability to access conceptually "local" file systems in a uniform fashion, regardless of the physical location of a computation. Moreover, in an increasing number of scientific disciplines, large data collections are emerging as important resources for the wider scientific community. In domains as diverse as global climate modelling, high-energy physics and computational biosciences, the volume of interesting data will soon be measured in Pbyte.

The communities of users who need to access and analyse this data are often large and are almost always geographically distributed, as are the computing and storage resources that these communities rely on to store and analyse their data. This combination of large datasets, geographic distribution of users and resources and computationally intensive analysis results in complex and stringent performance demands that are not satisfied by any existing data management infrastructure. A large scientific collaboration may generate many queries, each involving access to – or supercomputing-class computations on – Gbyte or Tbyte of data. Efficient and reliable execution of these queries may require careful management of Tbyte caches, gigabit/s data transfer over wide area networks, co-scheduling of data transfers and supercomputing computation. It is also essential to provide accurate performance estimations to guide the selection of dataset replicas and other advanced techniques that collectively maximise use of scarce storage, networking and computing resources.

The ENACTS Network attempts to promote and provide advice on good practice in data management. We aim to investigate and evaluate currently available software and hardware, explore new standards and support their development, develop good practice for users and centres and investigate platform-independent and distributed storage solutions. The aim is to identify requirements and components common to different systems, to enable us to apply different technologies in a co-ordinated fashion to a range of data-intensive application domains, for the benefit of users of HPC centres.

The Data Management Report will be of interest to:

- Researchers who will gain knowledge and guidance further enabling their research;
- Research centres that will be better positioned to give advice and leadership;
- European research programmes in developing more international collaborations;
- Industry, especially bio-pharmaceutical companies.

The report is available to download:
http://www.enacts.org/datamanagement.pdf

Distance Learning and Support

The Authors:
ICCC Group, Czech Republic, www.iccc.cz
UNI-C, Denmark, www.uni-c.dk
May 2004

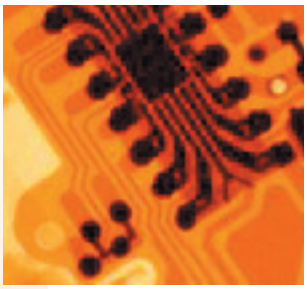

With few notable exceptions, the actual beneficiaries and users of tomorrow's Grid Technology have not yet established a dialogue regarding standardisation of technologies and tools used in education. Remarkably, formal remote training and distance learning are not rated as high priorities by the users of the Grid. The Study attempted to address the issue of the most relevant and effective distance learning method in the context of European and International Grid Communities.

Users accessing facilities through a pan-European metacentre will, typically, require training and support from remote centres in order to make best use of the available facilities. The WWW-based technologies are emerging to accomplish this, but are largely untried outside intranets. Here, we aim to determine the most appropriate support and training methods and the enabling technologies.

One of the key factors in increasing the uptake of HPC, is the ready availability of appropriate training for users and prospective users. Although many of the participants in ENACTS offer training to users, the majority of the courses they provide are instructor-led, language dependent and tailored to users of particular machines or tools. While there have been some instances of LSFs taking courses "on tour" to other areas of Europe, it has been difficult to share or re-use training material in a systematic way, due to difficulties in customising courses with a high practical content. There has been even less progress on sharing course development for similar reasons. WWW-based training offers opportunities to provide a flexible, self-paced approach to training, which is available to users of different platforms, via the Internet.

Recent developments in high-performance networks, computers, information servers and display technologies are of particular interest to the participants. These make it feasible to design network-enabled tools that incorporate remote compute and information resources unto local computational environments and collaborative environments that link people, computers and databases into collaborative sessions. The Grid takes this concept one step further, by enabling teleimmersion, the use of immersive virtual reality systems over a network, where the resources, simulations and data are remote from the user's display environment. Distance learning is one of the applications areas keenly watching developments in these areas.

Through the Sectoral Report, ENACTS aims to recommend a flexible suite of software tools to support the development and administration of interactive courses in HPC for scientific researchers and to provide an interactive teaching environment which allows users to concentrate on course material. Although technology-based training tends to be cheaper in the long run, the principal costs are all up front in course development. It is estimated that it takes between 50 and 300 hours to develop one hour of technology-based training, hence the importance of developing course material which can be widely used and easily customised.

Through the adoption of open standards, scalable design, flexible authoring tools, adaptability to different IT infrastructures, content and user management capabilities and customisation (languages, environments, implementation services and support), ENACTS aims to establish a collaborative framework of best practice for the development of teaching materials. These will be of benefit to all participants and infrastructure operators in Europe.

The report is available to download:
http://www.enacts.org/distancelearning.pdf

Software Efficiency and Reusability

The Author:
Parallab, University of Bergen, Norway, www.parallab.no
July 2004

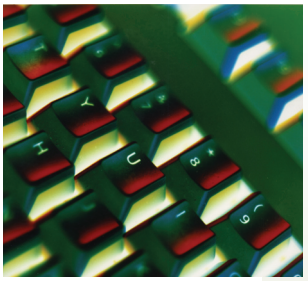

The uptake of computer simulation by emerging computational science disciplines and the constant quest for larger and more efficient computer resources by traditional disciplines leads to frequent evolution and revolution in software technologies that aim to support these developments. The Software Efficiency and Reusability Study as carried out by the ENACTS project focuses on recent developments and emerging software technologies as these affect the efficiency and reusability of new and existing applications. In addition, the Study looks at portability and interoperability of software and data between facilities and how these impact program efficiency.

The Study is performed by ENACTS partner Parallab, Bergen Center for Computational Science at the University of Bergen (Norway) with input from the other partners in the ENACTS consortium.

Increasing software reuse leads to higher human productivity and increased software quality. However, it is not always pursued in practice. The report produced by the Study reviews some general considerations related to software reuse, including its place in the software development life cycle but also the influence of non-technical aspects (like economics and management). The importance of standards and standardisation in modern programming and scripting languages, protocols, mark-up languages, and meta languages is discussed as well as the impact of Open Source.

The efficiency of sequential, parallel, and distributed applications depends on the infrastructure on which it is executed. Application performance is influenced by the details of the underlying hardware topology, network topology, software topologies, and not in the least the data topology. Measuring and tuning the performance of such applications is a complex task since it requires a good understanding of potential bottlenecks that can occur at the various levels in the computational infrastructure.

From a user's perspective, the development cycle and turnaround time of an application in a grid infrastructure are important. Successful reuse of existing applications in an efficient manner is a major acceptance criterion for emerging grid environments by the computational science community. Heterogeneous grid environments however pose daunting reusability and efficiency problems. The Study addresses all of the above issues in detail.

The report is available to download:
http://www.enacts.org/softwareefficiency.pdf

ENACTS is a Co-operation Network in the EC's "Improving Human Potential - Access to Research Infrastructures" Programme.

In HPC the key developments are in the area of Grid computing, and ENACTS will facilitate the move towards European Grid, a "virtual infrastructure" - where each researcher, regardless of nationality or geographical location, has access to the best resources and can conduct collaborative research with top quality scientific and technological support.

In ENACTS, our strategy involves close co-operation at a pan-European level - to review service provision and distill best practice, to monitor users' changing requirements for value-added services, and to track technological advances, in the following areas:

- Grid Service Requirements
- HPC Technology Roadmap
- Grid Enabling Technologies
- Software Efficiency and Re-usability
- Distance Learning and Support
- Data Management

For more information please see:
http://www.enacts.org

ENACTS Network is funded by the EC
under the contract HPRI-CT-2000-40015