WORKING GROUP REPORT ON LARGE SCALE COMPUTING FACILITIES AND RESEARCH NETWORK INFRASTRUCTURES

Final Version
6 January 2005

Report prepared by
Dany.Vandromme@renater.fr

Table of Content

Executive summary

The ESFRI (European Strategy Forum for Research Infrastructures) objective is to provide the EU with a prospective and strategic view of large scale research instruments, having a European significance, either in an EU funding approach, or more generally in Member States multilateral funding initiatives. Among the wide diversity of topics, deserving the ESFRI attention, the high performance computational facilities, together with the broadband research networks are a special case, since they are simultaneously a matter of research and at the same time support tools for all other science disciplines.

This report has been prepared by a dedicated working group, to provide the Forum with:

(i) a wide overview of the available resources available in Europe, and their relative position compared to similar instruments in other world regions

(ii) a set of guidelines and prospective views for the evolution of these instruments in the future, in order to keep Europe at a forefront position

(iii) an identification of science cases for which major progress may be critically dependent of the evolution of these instruments, either in a continuous streamwise mode, or eventually through a major breakthrough to achieve.

The report does not produce a complete and exhaustive overview of the computing and networking resources in Europe. However, most of these data are available through the updated work of the Arcade organisation (www.arcade-eu.org). The choice was taken to highlight a few examples, representing some of the main user communities. The next phase of the working group activity has been to bring focused information to the incomplete presentation of the European landscape in which the combination of computing resources, together with broadband network aims for the setup of a global Pan-European electronic infrastructure.

Because of the continuously evolving nature of the topic, this report does not intend to provide the Forum with a set of strategic decisions ready to take, but to trigger the discussion, for preparing in a near future, a clear roadmap and subsequent strategic guidelines for investments and policy for the European Union.

The facts coming out from the report are:

High Performance Computing

(i) There is an obvious relationship between the availability of large scale computing facilities and the support to a local hardware industry. The major HPC platforms are installed in US and Japan, where also companies are based, like IBM, NEC, TERA, INTEL, FUJITSU, HP, etc… Europe (except for the MareNostrum recently installed in Barcelona) does not belong to this, because the lack of major hardware industry.

(ii) The current funding model of hardware platform is based exclusively on national initiatives. There are no centrally funded HPC centres in Europe. This has, as a consequence, a very distributed (or even atomised) HPC landscape. The critical mass is therefore not there to challenge seriously US and Japan for centralized computing platforms.

(iii) For many scientific applications, there is not doubt that the GRID infrastructure will help improving the lack of resources, by making a better use

of distributed systems. However, GRID is unlikely to solve all scientific problems, and there are classes of scientific domains, which will always require very large centralised computing resources.

(iv) The most distributed nature of the GRID model will impact also the current peer-review model of projects for being granted computing resources. Except for project-based GRID infrastructure (like LHC GRID), there is a need to handle also the usage of GRID infrastructures at a policy level, without forgetting the existing established policy model in place for the use of the research network infrastructures. A better interaction between ESFRI and the eIRG is needed.

(v) The Earth Simulator (or even the various ASCI projects) is a one-of-a kind computer, which has the benefit to support the national HPC industry. Furthermore, it is attracting a lot of brain forces through the international scientific cooperation. The overall long term benefit of EU is questionable if no European resources are deployed for the research community.

(vi) Real impact on science of a large computing platform will be significant if associated with a scientific user community (like earth sciences in Japan, energy and space sciences in US) to promote it. Distributed or shared resources, like TeraGrid, or DEISA will benefit also to science, but in a much more dilute way, because not “owned” by a scientific user group. However, projects like MareNostrum (ranked 4th worldwide today) may also mark an inflexion point for HPC in Europe.

(vii) Storage is in many cases an integral part of HPC facilities. Today, there is a mismatch between application needs in high-performance online storage and what can be provided by existing facilities. Reaching the next level of capacity and performance, where data is kept online all the time, is central to advancing the state-of-the-art in existing applications as well as to enabling new applications.

Networking

(viii) The networking situation in Europe is relatively good, when compared to other world regions

(ix) The pan-European end-to-end communication issue still requires further improvement

(x) The collaborative organisational model for the NRENs and the European complement has a strong support of the EU. It has been copied worldwide and in so far “exports” the European model in other world regions.

(xi) The network infrastructure is still obtained from a lease mode with the telecommunication operators. The main issue for the future is to decide whether such infrastructures would be better owned by the research community. There are many issues related to this, in terms of capital investment and regulatory context as well, which need to be considered.

(xii) The use of network infrastructure is ruled today by the AUP compliance. If a per-project policy approach is to be needed in the future for allocating specific resources, the policy decision rules will have to be sorted out first, like for the GRID infrastructures.

Grid infrastructures

(xiii) GRID are about to provide users with interfaces to access any form of resources (computing, storage, processing, data transfer, etc…)

(xiv) GRID will support global progress of Science by allow virtual organisations to develop, whenever it will be required.

(xv) eIRG current activities are dedicated to develop a global approach and management for GRID access conditions (authentication, Authorisation and eventually Accounting).

(xvi) The current funding of the leading European GRID projects is not appropriate for providing a sustainable infrastructure. Recurring cost for the management of the service will need funding, which should be handled as the existing models (either by the users, or by its funding bodies). To minimize this, synergies between GRIDs and research networks and computing centres need to be harmonized rather than being duplicated.

As outcome of this preliminary work, it is shown clearly that several strategy and policy issues need a deeper discussion at the Forum level. Together with this long term guidance, the content of this report is the first milestone to produce a complete roadmap for the coming years, for the merging of the various components into a European e-infrastructure.

List of names symbols

AP: Asia-Pacific world region (see also www.apan.net)
BELNET: www.belnet.be: The organisation in charge of the Belgian NREN
CASPUR: www.caspur.it: Computing centre in Roma (Italy)
CEA: www.cea.fr: French Atomic Energy Research organisation
CEPBA : Centro Europeo de Paralelismo de Barcelona
CERN: www.cern.ch: Centre Européen de Recherche Nucléaire
CESCA : Centre de Supercomputació de Catalunya, Barcelona
CESGA : Centro de Supercomputación de Galicia, Santiago de Compostela
CIEMAT : Centro de Investigaciones Energéticas, Medioambientales y Technológicas, Madrid
CILEA: www.cilea.it: Computing centre in Milano (Italy)
CINECA: www.cineca.it: Computing centre in Bologna (Italy)
CINES: www.cines.fr: Computing centre in Montpellier (France)
CSIC : Consejo Superior de Investigaciones Científicas (National Council for Scientific Research)
CTTC : Centre Technologic de Transferencia de Calor, Terrassa
DANTE: www.dante.org.uk: Delivery of Advanced Networks To Europe, the coordinator of the European NREN consortium
DFN: www.dfn.de: The organisation in charge of the German NREN
eIRG: e-Infrastructure Reflection Group
EU: European Union or Europe region
FUNET: www.funet.fi: Finisch NREN
GARR: www.garr.it: Italian NREN
GEANT: Gigabit European Academic Network
GN1, GN2: FP5, FP6 funded projects to support pan-European backbone infrastructures
HEAnet: www.heanet.ie: Irish NREN
HEP: High Energy Physics
HPC2N: High Performance Computer Centre North, Umeå
HPC: High Performance Computing
ICT: Information and Communication Technologies
IDRIS: www.idris.fr: Computing Centre in Orsay (France)
IFAC : Instituto de Física de Cantabria, University of Cantabria, Santander
INEM : Instituto Nacional de Empleo, Madrid
INM : Instituto Nacional de Meteorología, Madrid
IN2P3: www.in2p3.fr: HEP research organisation in France
ITER : Instituto Tecnológico y de Energías Reno-vables, Santa Cruz de Tenerife
KTH: Kungliga Tekniska Högskolan, Stockholm
LHC: Large Hadron Collider
LUNARC: Centre for scientific and technical computing at Lund University, Lund
NORDUNET: www.nordu.net: Nordic NREN consortium
NREN: National Research and Education Network
NSC: National Supercomputer Centre, Linköping
PDC: Center for Parallel Computers at KTH
RENATER: www.renater.fr: French NREN
SNAC: Swedish National Allocations Committee
SNIC: Swedish National Infrastructure for Computing
SUNET: www.sunet.se: Swedisch NREN

Super SINET www.sinet.jp: Japanese Academic network
TOP500: www.top500.org: a public list of HPC facilities around the world, sorted by their theoretical computing power
UKERNA: www.ukerna.ac.uk: Organisation in charge of the UK NREN (SuperJANET)
UNICC: Unix based Numerically Intensive Calculations at Chalmers, Chalmers University, Göteborg.
UPPMAX: Uppsala Multidisciplinary Center for Advanced Computational Science
UPV : Universidad Politechnica de Valencia
VR: Swedish Research Council

List of technical symbols

CPU: Central Processing Unit (processor)
GIX: Global Inter-eXchange point
GNP: Gross National Product
IP: Internet Protocol
PoP: Point of Presence
SDH, PDH: Synchronous transport protocols for electrical high speed links
WDM: Optical transport protocol: Wave Division Multiplexing

1. Introduction

The working group operated mostly by using standard electronic tools for communicating and exchanging documents. There were 3 physical one-day meetings, organised in Munich (September 24th, 2003) and Paris (January 26th, and October 10th, 2004). The attendance was relatively weak (6-8 participants), and several countries, members of the Forum, did never participate.

At the second meeting, there were two invited speakers: Geerd Hoffmann, from the Deutsche Wetterdienst, and Olivier Martin, from the network division at CERN. A third meeting was organised in Paris, at which the discussion was started around the Barcelona HPC project, presented by Sergi Girona (CEPBA/UPC). Further discussions were pursued in the framework of the eIRG context, and the content of the report was related to the e-Infrastructure framework (www.e-irg.org).

Even if not all figures could be collected on time for this report, a more complete overview of the European landscape is available from the Arcade work at: http://www.arcade-eu.info/academicsupercomputing/.

The report was not intending to duplicate this work, but rather to analyse some issues related to the strategic or policy concerns for the future of research in Europe.

The main issues are, from the HPC side, the atomisation of resources, because of the current national level of investment decisions, which put Europe way behind US and Japan, in various aspects. From the Networking side, there is still a lack of strategic decision about the ownership of the infrastructure on a long-term perspective. This has obviously many consequences on the form of investment currently done as today. Even if the current model has certainly benefited to the European Telecom market, it is still questionable for a longer-term view. The prospective view about research networking in Europe was already worked out in the SERENATE study (www.serenate.org), which made suggestions and recommendations.

The report is organised in three main parts. The first one, dedicated to the High Performance Computing resources in Europe highlights the most visible ones, and shows how Europe is positioned in a global landscape. A second part will provide the same exercise as far as research networks are concerned, both at the national and European levels. It is recognized in this framework that GRIDs are a service infrastructure aiming to combine all these hardware and software components, including possibly widely distributed CPUs into an homogeneous resource accessible in a way transparent to users. A third part is dedicated to a few science cases, for which computing and networking will be critical to any significant scientific progress. The fourth part analyses some aspects of the situation in US and Japan and tries to elaborate some rationale and consequences for Europe research.

2. Computing

A first approach to the available computing power is not easy to present, since there are many different indicators, which could be used, giving different views of the situation. Among these indicators, we could consider:

(i) Total capital investment in HPC facilities
(ii) Total consolidated budgets of HPC facilities
(iii) Aggregated number of CPU available for public research activities
(iv) Aggregated computing power (extracted from the TOP500 world ranking[1])
(v) Massive storage capacity
(vi) Distribution and number of HPC facilities

There may be between one and two orders of magnitude in peak performance figures within a single investment cycle of 3 years for a given type of processor. Similarly, there is a wide range of discrepancies between performances of different processor architecture. The raw comparison of CPUs number is therefore a weak indicator in terms of computing power. However, it can be taken as an indicator of the national Policy regarding the importance of computational science.

Another missing indicator is the resource available for private research. The main players are primarily the manufacturing industries (car, aircraft, engine, drug, chemistry). However, as far as research is concerned, the industry activity is somehow reflected on the potential of the public research because of very frequent joint activities.

2-1. Global

The global situation of High Performance Computing is rather difficult to extract, because of a wide range of possible criteria to characterize it. A first approach would be to take the list of the largest facilities worldwide and analyze their geographical distribution. The basic source for this is obviously the TOP500 ranking, prepared jointly by Argonne Lab and Darmstadt University. From this, it is possible to show the number of systems in different world regions. Similarly, the aggregate computing power per region can also give useful indications. The main comparison between the two representations is certainly the gap between EU and AP/US, which is materialised by the performance of a single system located in Japan (Earth Simulator) ranked N°1 on the TOP500 list dated April 2004. In the November issue of the TOP500, the Earth Simulator has been pushed down to the third position, for the benefit of 2 new platforms in US, at NASA-Ames (SGI) and Los Alamos (IBM), while a European platform in Spain (Barcelona Supercomputing Centre) is now ranked at the fourth position (IBM).

[1] www.top500.org

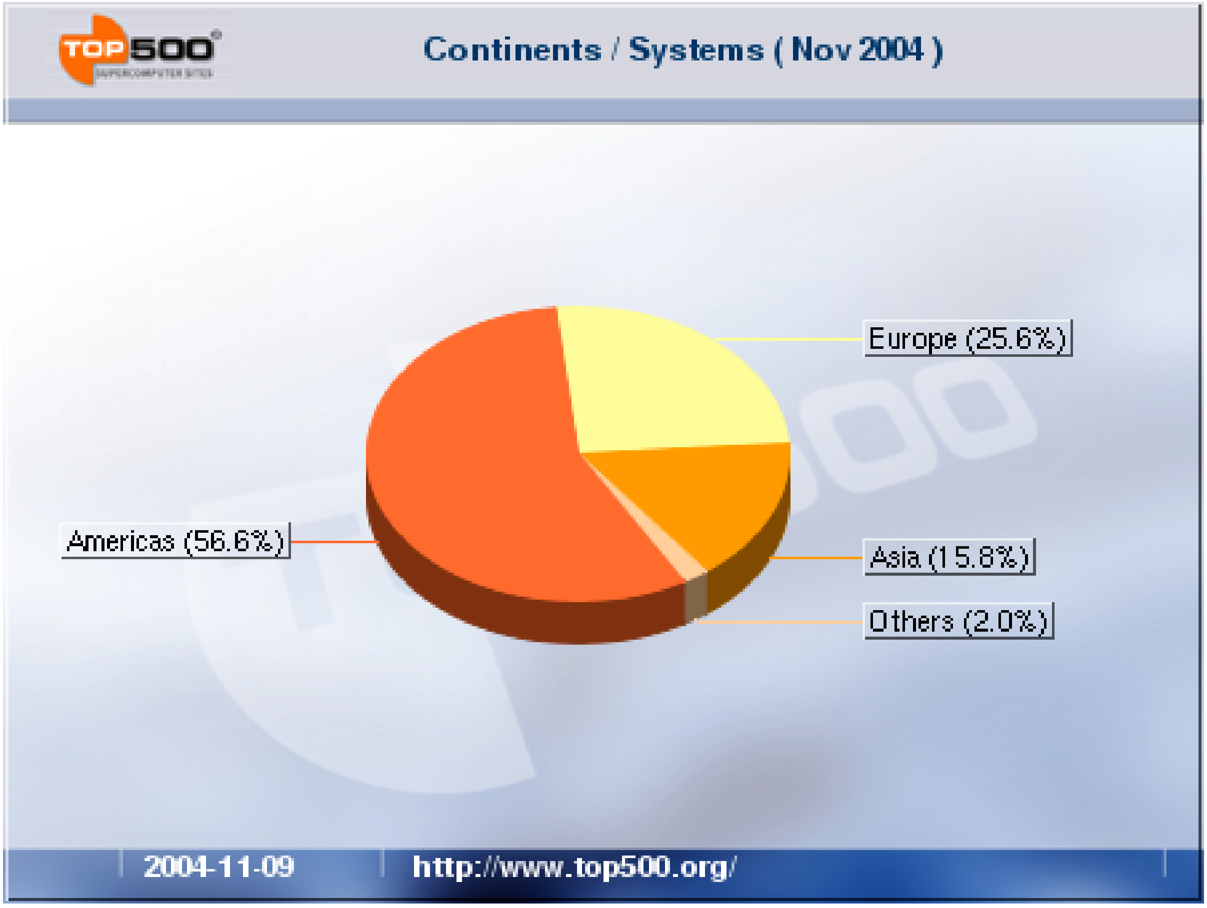
TOP500
Continents / Systems (Nov 2004)
Europe (25.6%)
Americas (56.6%)
Asia (15.8%)
Others (2.0%)
2004-11-09
http://www.top500.org/

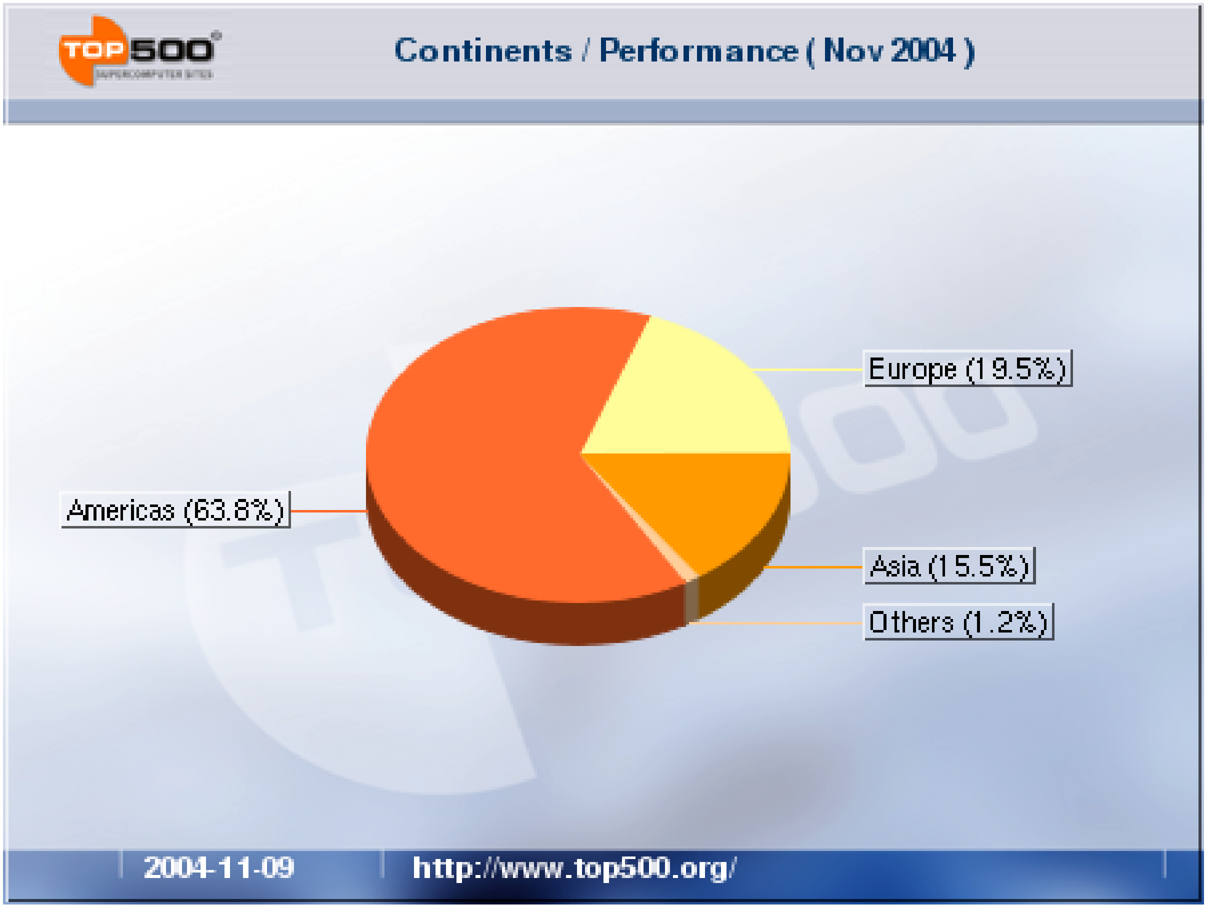
TOP500
Continents / Performance (Nov 2004)
Europe (19.5%)
Americas (63.8%)
Asia (15.5%)
Others (1.2%)
2004-11-09
http://www.top500.org/

Obviously these figures should be properly scaled with the user base figures, or any GNP-like figures, but it would go certainly beyond the scope of this report.

2-2. Per country basis

Looking at the same pictures with a country-sized focus shows similar trends:

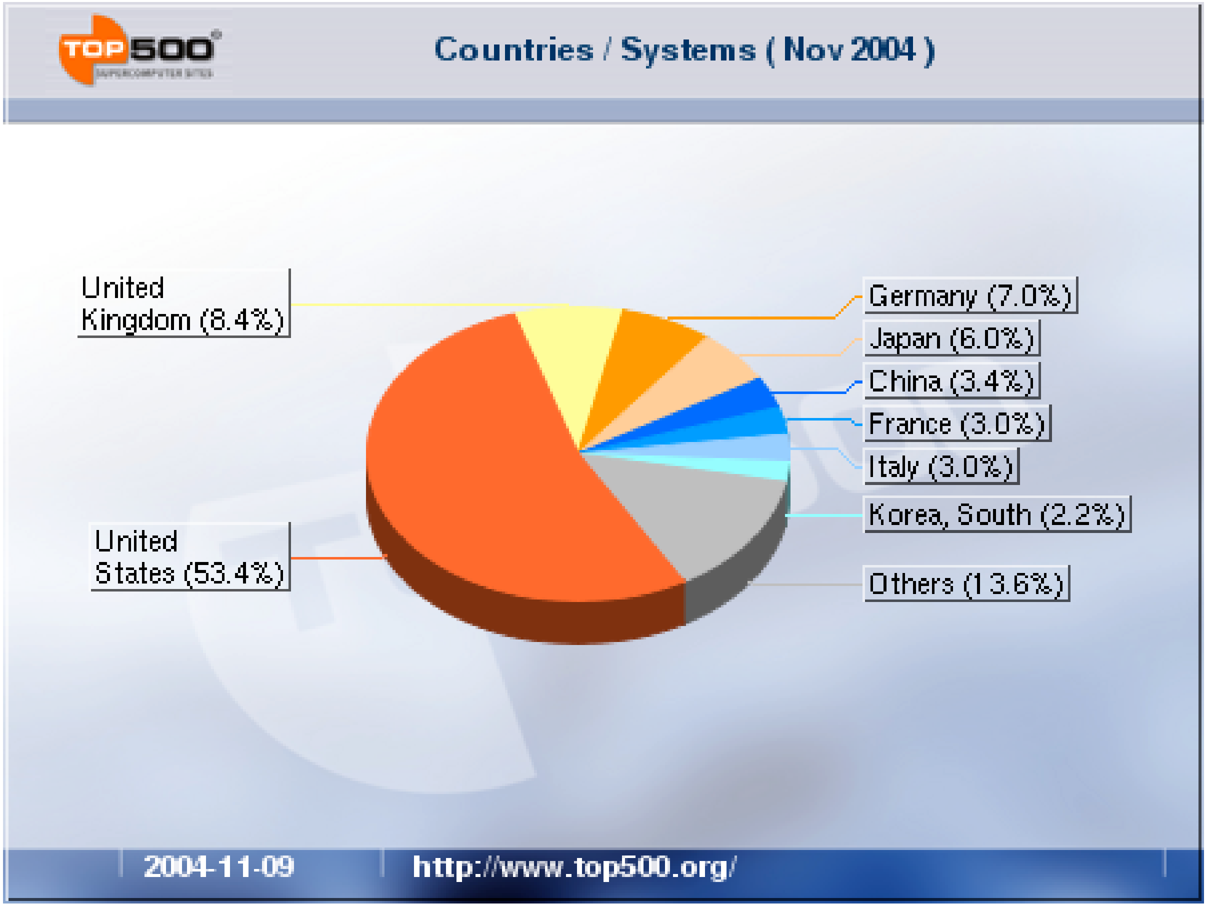

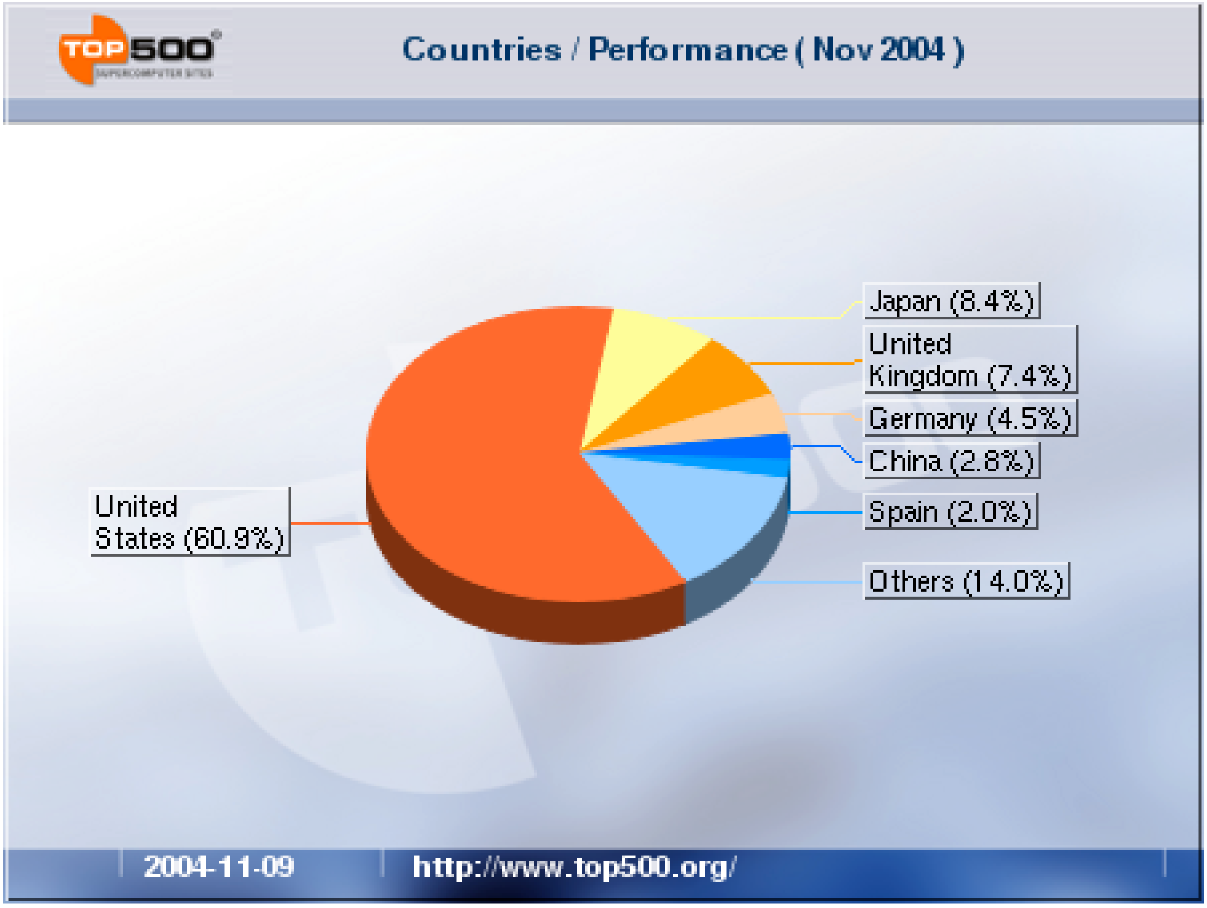

Belgium

General Description

In Belgium there are no centres for High Performance Computing that are available to the whole research community. Each university or research institute does provide computing power for its own users. The use of supercomputer in all institutes is limited or even non-existing. Computer centres now often decide to use clusters of ”PCs” to provide computing capacity.

Belgian GRID infrastructure

However BELNET started a BELNET Grid Initiative that initiated the start of BEgrid, a grid infrastructure that is open to all institutes that are connected to BELNET. Participants in BEgrid add their own equipment to the general shared infrastructure.

The Flemish Government has decided to invest 717 k€, spread over 4 years, in grid equipment that has to be integrated in BEgrid. Deployment of the first part of this equipment is foreseen before summer 2004.

Czech Republik

General Description

Building of the Czech Computing Grid started already in 1996, under the MetaCenter project, which has been part of the TEN-34 CZ programme and accompanied the TEN-34 CZ project of the first truly high speed academic network. The Grid was built over the three major computing centers in the country:

- Charles University in Prague,
- Masaryk University in Brno,

- West Bohemia University in Pilsen.

Since 1999 the Grid related activities are part of the CESNET research project. Under CESNET umbrella, both data resources – starting with the purchase of tape robot with 12TB online capacity in 1999 – and computational resources – starting with the purchase of the first PC based cluster with 64 CPUs in 2000 – gradually evolved. The Czech Grid (or MetaCenter) contemporary configuration contains almost 200 Intel IA-32 processors located in Prague, Brno and Plzeň, together with old SGI Origin systems (in Prague and Brno) with more than 70 MIPS processors and a mid-range Alpha server. The overall storage capacity is almost 5 TB disk space on-line and is expected to increase by another 5–10 TB disk space in 2004 (under an independent project for distributed data storage). The Grid resources are connected through Czech NREN CESNET2.

There is also participation in international Grid-related projects, most notably within the EU DataGrid project within the 5th Framework programme. The Masaryk University node participates also on the GridLab project (again within the 5th Framework programme). Currently CESNET becomes part of the EGEE consortium, continuing thus its activities into the EU 6th Framework programme. Thanks to rich external co-operation MetaCenter team has been a part of an international team which won most of the SC02 challenges, including the most bandwidth challenging Grid application.

Funding

Since 1999 the MetaCenter project as well as related Grid activities (AAI, PKI, Access Grid Points...) were supported by Ministry of Education, Youth and Sports within the framework of the research intent "High Speed National Research Network and Its New Applications". The research intent finished in 2003 and further development of computing Grid is ensured by new project "Optical National Research and Education Network and Its New Applications" accepted and supported by Ministry of Education, Youth and Sports of the Czech Republic for the period 2004-2010. The Grid development got no other official funding. At 2003, new grant programme for Information infrastructure for research and development was launched by Ministry of Education, Youth and Sports, but no Grid related project has been accepted in the first call.

At the 2003, the budget of "High Speed National Research Network and Its New Applications" research intend was 328 mil. CZK in total, consisting of governmental subsidy (208 mil. CZK) and CESNET's own contribution (120 mil. CZK). For the Metacenter project, there was dedicated 11 mil. CZK. Within the new research plan, the total budget for year 2004 is 351 mil. CZK (274 mil. CZK governmental subsidy). For this year about 10 mil. CZK is dedicated directly to Metacenter project.

Denmark

General Description

The Danish Research Computing Infrastructure is coordinated from the Danish Centre for Scientific Computing, DCSC. DCSC receives requests for two types of resources; computing time and hardware. There are four operational centres under DCSC. Scientists at national and public research institutions can apply for funds for equipment purchases, and a contract between DCSC and an operational centre specifies the funds available to the centre in return for purchasing and running a system for three years. A substantial number of the applications are sent out for international review and the applications are evaluated according to their scientific merits and originality. The funds go directly to the research groups and not to a centre. Technically, the centres receive the money from the research groups that obtained support from DCSC. This model for providing computing infrastructure is just under three years old and is up for international evaluation in the spring of 2004.

Equipment
The DCSC resources currently represents the accumulated resources for two years and thus are short the investment of one year to represent a full circle of investments.

- Three large clusters with a total of 1350 CPUs (Intel P4 2.0; 2.2 and 2.66 GHz)
- An IBM Regatta 80 CPUs (Power 4 1.4GHz)
- A SGI Altix of 64 CPUs (Itanium 2 1.4GHz)
- A 64 CPU SUN Enterprise 15000
- A 64 CPU SGI Origin 2000
- A 96 CPU SGI Origin 300
- A 40 CPU SGI Origin 3000

User-base
The DCSC user base is quite diverse and includes a large base of 'small scale users', the 'heavy' users includes

- Quantum chemistry
- Reaction dynamics
- Solid state physics
- Nano science
- Materials science
- Bioinformatics
- Bio-/Life-science
- Astrophysics
- Environmental science
- CFD – in particular windmill modelling

Funding
The overall budget comprises roughly 2M€ per year, is provided by government.

Finland
General Description
High performance computing facilities in Finland are mainly located at CSC in Helsinki. The current hadware resources are as follows:

- IBMSC: The IBM eServer Cluster 1600 (ibmsc.csc.fi) has the theoretical peak performance 2.2 Tflop/s, and is used for parallelized computing tasks in bioscience, physics, chemistry and engineering. The Finnish Meteorological Institute computes its weather forecasts on this computer.
- SGI Origin 2000: The SGI Origin 2000 (cedar.csc.fi) has the largest shared memory space in the Nordic countries, with a total of 160 gigabytes. The SGI Origin 2000 has been equipped with the largest variety of scientific software programs available in Finland.
- Compaq/HP Alpha Cluster: The cluster system, co-owned by CSC and the Helsinki University of Technology, is made up of 96 Compaq/HP AlphaServer DS10 servers. The hardware is intended for parallel applications. Some of the more important tasks that are run on the cluster are the simulations of material physics and molecular dynamics.
- Compaq/HP AlphaServer ES 45: Two 4 processor AlphaServer ES 45 servers are dedicated to the computation of serial calculations

- IBM RS/6000 H70: IBM server (laari.csc.fi) has several chemistry databases installed on it. Over 500 GB of disc space is available. The world's largest electronic spectrum database, SpecInfo, is on this server, with 700,000 chemical spectrums. It also includes the world's largest organic chemistry electronic database, Beilstein, with over 8 million chemical compounds. Yet another database, Gmelin, contains two million unorganic and organic metal compounds.
- Compaq/HP AlphaServer GS80: Two AlphaServer servers (mdserver1.csc.fi, mdserver2.csc.fi) and a three terabyte storage system make up the heart of the CSC metacomputer. The centralized disc server offers CSC’s customers a common file view and user environment regardless of which of the metacomputer servers is in use. Researchers can use the disc server to save their data temporarily for future large computing tasks or database retrievals.
- GeneMatcher2: Database search engine serving the bioinformation field (gepardi.csc.fi). The Paracel GeneMatcher2 hardware is used for analysis of microchip and proteomic data. The server is one of a kind in its technical capabilities in Finland. The system is composed of 16 processors from the Linux cluster and the GeneMatcher2 hardware with its 9216 special processors.
- Sun Fire V880 & ADIC Scalar1000: The CSC archive server (ds.csc.fi) is used for permanent archiving. The system is composed of a Sun V880 server and an ADIC Scalar 1000 tape robot. The storage capacity of the system equals 60 terabytes.
- Besides HPC environment, CSC offers expertise in high performance computing. The volume of the expert work exceeds 30 person years/year.
- CSC is the DEISA partner, and the main Grid center in Finland together with HIP, the Helsinki Institute of Physics.
- In the near future, CSC will actively invest on PC clusters.

Middle-scale HPC facilities can be found in Finnish universities (typically PC clusters 16-64 processors) for instance in Helsinki, Jyväskylä, and Turku.

France

General Description

In France, High performance computing resources are distributed on 4 major sites. These are:

- CINES (Montpellier), dedicated to support university research
- IDRIS (ORSAY), CNRS computing centre, located in Orsay. IDRIS is a leading partner of the FP6-funded GRID project DEISA.
- CEA Computing Centre, which is also located in the south of Paris region
- CC-IN2P3, located in Lyon, supporting mostly the HEP research activities, in close cooperation with CERN, FermiLab and SLAC.

For the first two centres, computing resources allocation is made on an annual basis, following open calls for proposals. The scientific projects are evaluated by a set of scientific panels (9 thematic committees), on a per discipline peer review, and the resources are allocated according to the scientific quality of the project and the ad-hoc matching with the computing architecture.

1) CINES (“Centre Informatique National pour l’Enseignement Supérieur”)

- IBM → 1 Tflops
    - IBM SP (P3) : 29 nodes HPC 16 proc summing up to 464 proc (with 304 GB)
    - IBM SP (P4): 2 nodes Regatta summing up to 64 proc (with 128 GB)
    - Switch COLONY with 6 TB Dsk
- SGI → 0.768 Tflops

  - O3800: 768 proc R14k (with 384 GB + 4 TB Dsk)

2) IDRIS (Institut pour le Développement des Ressources Informatiques Scientifiques – Development Institute for Scientific Computing)

- IBM → 1.3 Tflops
  - IBM SP (P4): 8 nodes Regatta 32 proc summing up to 256 proc (with 832 GB + 3.5 TB Dsk)
- NEC → 320 Gflops
  - 3 x SX5 (16 + 16 + 8 = 40) procs with 224 GB Dsk

3) CEA-CCRT (Commissariat à l'Energie Atomique – Atomic Energy Agency)

- COMPAQ/HP → 7 Tflops
  - 640 nodes of 4 procs + 2.5 TB RAM + 50 TB Dsk
  - x10 in 2005, x10 in 2009
- NEC → 350 Gflops

4) CC-IN2P3 (HEP focused)

- 900 proc cluster + 1 PB Storage

French GRID infrastructure

Several institutions are active in different EU-funded GRID projects like DATAGRID, EUROGRID and soon DEISA and EGEE. Besides these participations, the French ministry of research has supported a two year GRID initiative, e-Toile based on further adaptations of GLOBUS and using vTHD, a broadband test-bed infrastructure belonging to France Telecom. Since end of 2003, a major GRID project has started to deploy 5 000 CPUs in GRID architecture. This project, led jointly by CNRS, INRIA, Universities and the Research ministry will provide an operational infrastructure distributed over 8 laboratories connected to RENATER starting early 2005.

Germany

General Description

In Germany there are four supercomputer sites, which are open to most of the German scientific community, these are:

- Hochleistungsrechenzentrum Universität Stuttgart (HLRS)
- Leibniz Rechenzentrum München (LRZ)
- John-von-Neumann Computing Center Jülich (NIC)
- Konrad-Zuse Institut (ZIB) Berlin in close cooperation with Rechenzentrum Hannover (HLRN)

In addition there are several supercomputer sites, which are dedicated to specific communities Examples are:

- Rechenzentrum der Max-Planck-Gesellschaft (IPP Garching)
- Deutsches Klimarechenzentrum, Hamburg (DKRZ)
- Deutscher Wetterdienst (Offenbach).

The following statistics from LRZ is a snapshot of the past years about the usage profile from different disciplines. It is available in the LRZ web as well.

Usage (in % system time LRZ)

	Fluiddynamics	Solid State Physics	High-Energy Physics	Chemistry incl. Chem. Physics	Rest	Total
2000	48.5	5.2	37.0	0.1	9.2	100
2001	33.8	30.6	23.3	5.9	6.4	100
2002	38.2	26.7	20.4	9.1	5.6	100

German GRID infrastructure

D-Grid is a national initiative driven by research institutions, DFN, industry and the German Ministry for Education and Science. The purpose is to bundle activities for global, distributed and enhanced research collaboration based on internet services with the goal to build an e-science framework. The German Ministry for Education and Research expects common commitments and co-funding from research organisations and industry. A tender for e-science projects is expected to be carried out in Q3/04. The Ministry announced a funding profile of 5-10 million Euro per year starting in 2005 until 2008 (http://www.d-grid.de).

A prominent partner of this initiative is the Forschungszentrum Karlsruhe, which is running the large Grid Computing centre GridKa, currently serving 8 international particle physics experiments. GridKa is a Tier-1 centre for the LHC Computing Grid and will develop over the next years as is shown below (http://grid.fzk.de):

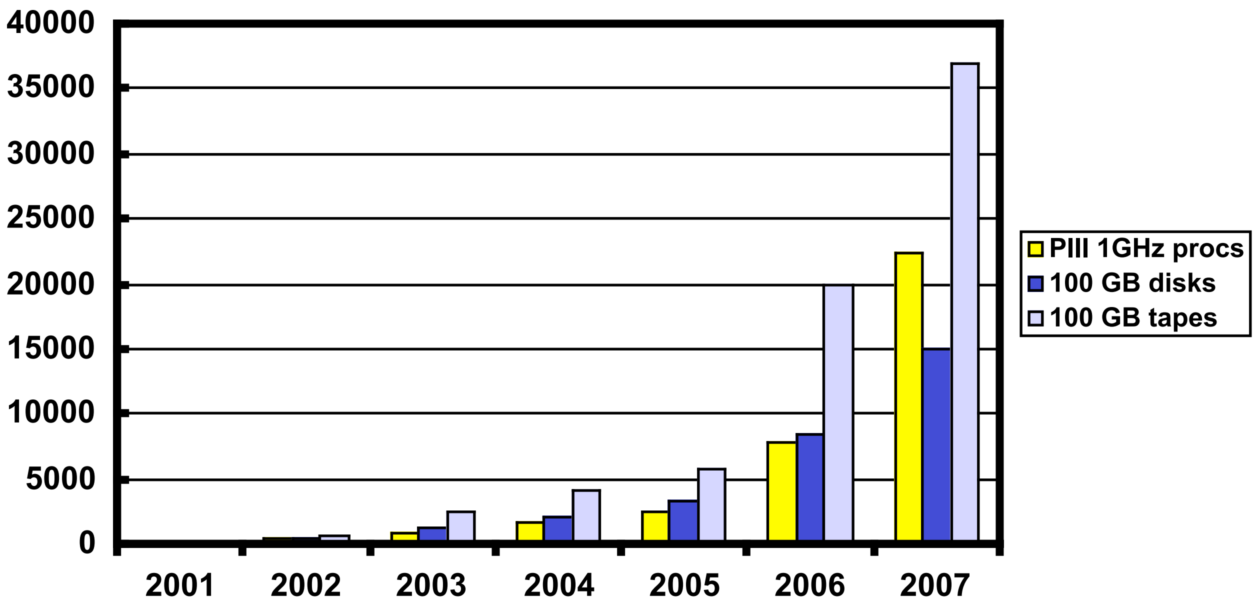

Greece

General Description

- Aristotle University of Thessaloniki: Beowulf cluster (15 PIV/2800, 8 2*PIII/1266), 4 TB Storage SATA RAID Array.
- Greek Research and Technology Network/HellasGrid - HG01 @ Demokritos: Beowulf cluster (68 2*Xeon/2800), 10 TB Fibre Channel SAN, 10TB Tape library. Under development: 6 Beowulf Clusters 128 Xeon/3000), 20 TB storage, 50 TB Tape Libraries, 4 Access Grid Nodes.
- National Technical University of Athens - Computing Systems Lab (CSLAB): SMPs: 2 x Sun Enterprise HPC 450 servers (quads, 4 CPUs each), 1 quad AMD Opteron/2400 (4 CPUs), Beowulf clusters: 8 x dual Xeon/2800 with Myrinet & Gigabit Ethernet, 16 x PIII/500 with SCI, 8 x dual PIII/800, 4 x dual PIII/1266 with

SCI & Myrinet & Gigabit Ethernet, 12 AMD Athlon/1600, Interconnects: Gigabit Ethernet, Scalable Coherent Interface (1GB/sec), Myrinet (2+2GB/sec), Storage: 2 TB storage SATA RAID Array

- Foundation for Research and Technology-Hellas / Institute of Computer Science (ICS-FORTH): Beowulf Cluster 64 AMD/1266
- Institute of Accelerating Systems and Applications (IASA): Beowulf cluster (14 2*Athlon/2000, 15 2*Xeon/3000), 3 TB Storage
- Research Academic Computer Technology Institute (CTI): 10 PIII/800)
- University of Athens: HP V2600/48
- University of Macedonia: Beowulf Clusters (65 PI-II/266, 6 PII/266, 4 PIV/2000)
- University of Patras: Beowulf Cluster (10 PIII/800)

Grid Infrastructure

The Local Office of the Information Society – Ministry of Economy and Finance established in December 2002 the HellasGrid task force. Its task is to define the basic strategy and guidelines for the Greek national, regional and international grid activities. In the framework of the Hellasgrid Task Force a 2M Euro proposal has been approved that will develop computing, storage and virtual collaboration infrastructures, complemented with operation, middleware, security, training and dissemination activities. For more information http://www.hellasgrid.gr/

Ireland

General Description

The funding agencies in Ireland are: the Higher Education Authority (HEA), Science Foundation Ireland (SFI), Enterprise Ireland (EI), the Irish Research Council for Science, Engineering and Technology (IRCSET) and the Health Research Board (HRB).

- Funding has just been granted for Phase 1 of a central large scale facility in Ireland.
- Phase 2 funding will follow in 2005.
- There is an Origin 3800/40 + Altix 3700/20 at NUIG and a 220-CPU cluster at UCD.
- There are three 100-CPU clusters in UCC and 130-CPU + 80-CPU clusters in TCD.
- There are several small-scale clusters.
- There is funding for several medium-scale clusters (100-500 CPUs) in 2004/5.
- There is funding for two medium-scale data farms in 2004/5.
- Most scientific computing is still done on local facilities.
- Wide range of application areas, but few truly parallel applications.
- Most users develop their codes but use commercial libraries and tools.

Irish GRID infrastructure

- Majority of academic institutions connected to national grid (Grid-Ireland).
- Grid-Ireland is in EGEE, with ROC at TCD.
- Three medium-size national VOs (CosmoGrid, WebCom-G, MarineGrid).

Italy

General Description

In terms of budget, only the figures of CINECA (www.cineca.it) are reported here, being the numbers of the other Inter-university Computing Centres in Italy, CILEA (www.cilea.it) in Milan and CASPUR (www.caspur.it) in Rome, much smaller than those describing CINECA.

CINECA has to do with ICT at large. The total budget of CINECA (consolidated year 2003) is 43 M€, of which 10.7 M€ have been allocated for the high performance computing system. The provisional for 2004 is: total budget 43.8 M€, of which 10.5 M€ for the high performance system.

The budget for high performance system does include the leasing of the facility, the maintenance, the electricity and the cooling system (leasing and maintenance) and also the cost of the cost of the staff (about 50 persons out of 240 employees)

Beside that, the development of CINECA computing system is supported by 5 national agencies, which in total cover half of the leasing cost of the facility, not included in the budget of CINECA.

In conclusion the actual capital investment of CINECA for the high performance system is of the order of 35 M€ and the yearly budget is in the order of 14 M€ (10-11 from CINECA and 3-4 from other national agencies investing on the facility of CINECA.

Spain

General Description

There is no specific national HPC policy. Some HPC and networking activities are embedded in the programmes of the National Plan for Research and Development which is managed by the Scientific Research Council (CSIC).

HPC facilities

The Spanish resources are distributed between various centres across the country as:

- CEPBA: IBM SP Power3/128, SGI Origin2000/64, Compaq/HP Alphaserver/16, Compaq/HP Alphaserver/12
- CESCA: Compaq/HP Alphaserver/32, HP GS1280/16, IBM SP2/44, HP V2500/16, HP N4000/8
- CESGA: HP SuperDome/128, HP HPC320/32, Beowulf cluster (34 PIII/(550-1000))
- CIEMAT: SGI Origin 3800/160, SGI Altix 3700/64
- CTTC, Beowulf cluster (48 Athlon/800)
- IFAC: Beowulf cluster (72 PIII/1260)
- INEM: SUN HPC 10000/52
- INM: Cray X1/40
- ITER: Beowulf cluster (288 Athlon/1650)
- University of Barcelona: Sun FireCluster/152
- University of Valencia: SGI Altix 3700/100, SGI Origin2000/64
- University of Sevilla: Beowulf cluster (31 P4/1700)

More recently the Ministry of Science and Technology has provided 70 M€ funding for a computing center, equiped initially with a 40 Tflop/s IBM system to be installed in Catalonia. Delivery is done during 2004. This system, known as the MareNostrum (http://www.top500.org/sublist/System.php?id=7119), is ranked today at the fourth position in the top500, which makes it, the first resource for high performance computing in Europe. MareNostrum will be hosted in the "Barcelona Supercomputing Center - Centro Nacional de Supercomputacion - which is being created.

National GRID Infrastructure

CESCA and CESGA have create a grid infrastructure by linking their HP HPC320 systems.

The PDGE project (Proyecto Data Grid España) is a co-operation of five Spanish HEP institutes (IFAE, IFCA, IFIC, CIEMAT and UAM) with the aim to establish a grid testbed as part of the EU-DataGrid project.

UPV is creating an interdepartmental grid that will integrate a SGI Altix 3000/48, a Beowulf cluster/128 and about 3000 computers at several UPV departments.

Sweden

General Description

The Swedish Research Council (VR) funds the Swedish National Infrastructure for Computing (SNIC), which is a meta-centre with six affiliated HPC centres: HPC2N, LUNARC, NSC, PDC, UNICC and UPPMAX. SNIC has an annual budget of 4.5 M€. Detailed information on SNIC is available at http://www.snic.vr.se/

HPC facilities

Supercomputing facilities for the academia

- HPC2N: Beowulf cluster (240 Athlon/1667), IBM SP P2SC/68
- KTH (including PDC): Beowulf cluster (180 Itanium2/900), IBM SP P2SC/146, IBM SP Power3/84, IBM SP Power3/40, Beowulf cluster (80 Athlon/1400+24 PIII/866+8 Athlon/900)
- LUNARC: Beowulf cluster (128 P4/2530 + 65 Athlon/1600), SGI Origin 2000/116
- NSC: SGI Origin 3800/128, Beowulf cluster (400 Xeon/2200), Beowulf cluster (33 Athlon/900), Beowulf cluster (16 Athlon/850), Beowulf cluster (16 Athlon/800), Beowulf cluster (25 Athlon/700)
- UNICC: SGI Origin 2000/110
- UPPMAX: Sun Fire15K/48, Beowulf cluster (17 Athlon/1000)
- SweGrid: 6 Beowulf clusters (600 Intel P4/2800).

Allocation of resources

The Swedish National Allocations Committee (SNAC), sorting under SNIC, regulates the use of national computational, grid, visualisation and data warehousing facilities through evaluation of applications for such resources. It also stimulates the innovative use of these resources. Further information is available at http://www.snic.vr.se/

Swedish National GRIDs

The Knut and Alice Wallenberg foundation (KAW) awarded 2.5 M€ to establish SweGrid – a Swedish computational and storage grid with several hundreds of processors and a few hundred TB of storage. The computational and storage resources of SweGrid are located at the six national HPC centres. VR has allocated through its IT-Research Committee additional funds to the SweGrid project for research on grid-related issues. SNICs Strategic Technical Advisory Committee (STAC) serves as steering committee for SweGrid. SweGrid will is planned to be fully operational in early 2004. For further information, see http://www.swegrid.se/

United Kingdom

General Description

In UK, the computing services are organised in two distinct sets of resources: A general purpose computing service for academic research, and an excellence site for high performance computing. Added to these, there are also significant GRID initiatives, which are also mentioned.

1. Computing Services for Academic Research (CSAR)

Location:

- University of Manchester

System Configuration and Performance

The service consists of 5 separate systems as follows:

(i). **Turing**, a Cray T3E-1200E with 816 DEC Alpha EV65 600 MHz CPUs and a 3D Torus CPU Interconnect. The peak performance is 979 Gflop/s with 204GB of memory.

(ii). **Fermat**, an SGI Origin2000 with 128 MIPS R12K 400 MHz CPUs and an SGI NUMAlink CPU Interconnect. The peak performance is 102 Gflop/s with 128 GB of memory.

(iii). **Green**, an SGI Origin3800 with 512 MIPS R12K 400 MHz CPUs and an SGI NUMAlink CPU Interconnect. The peak performance is 410 Gflop/s with 512 GB of memory.

(iv). **Wren**, an SGI Origin300 with 16 MIPS R14K 500 MHz CPUs and an SGI NUMAlink CPU Interconnect. The peak performance is 16 Gflop/s with 16 GB of memory.

(v). **Newton**, an SGI Altix 3700 with 256 INTEL Itanium 2 (Madeson) 1.3 GHz CPUs and an SGI NUMAlink CPU Interconnect. The peak performance is 1.331 Tflop/s with 384 GB of memory. {Note that this system is currently being installed – it is planned to start by October 2003}

Budget and Staffing

The cost will be about £31M for a 7.5. year service, which started at the end of 1998. The number of staff is 18 covering service delivery management, systems support, operations, applications support and helpdesk.

Science Communities

- Aerodynamics & Flight Mechanics
- Aerospace Engineering
- Animal & Plant Sciences
- Astronomy
- Atmospheric Modelling
- Biology
- Biomedical Sciences
- Chemistry
- Coastal & Marine Sciences
- Engineering
- Environmental Science
- Genetics & Biometry
- Geography
- Geological Sciences
- Materials Engineering
- Mathematics
- Mechanical Engineering
- Meteorology
- Ocean/Earth Sciences
- Particle Physics
- Physics

- Process Integration

2. High Performance Computing (HPC)

General Description

Location:

The system is located at the Daresbury Laboratory with user and applications support being provided by teams at the Edinburgh Parallel Computing Centre (EPCC) in the University of Edinburgh and at the Daresbury Laboratory.

System Configuration and Performance

The service consists of only one system. Phase 1, installed in Q4 2002, has 40 32-processor IBM p690 frames, using POWER4 technology, connected by IBM's Colony switch. The peak performance is 6.66 Tflop/s with 1.28 TB of memory. It is currently the 12^{th} most powerful supercomputer in the world (according to the 21^{st} Top500 list).

It is planned to upgrade the system in Q2 2004 to provide a peak performance of 12 Tflop/s, with a further upgrade in Q3 2006 to 24 Tflop/s.

Budget and Staffing

The total cost of a 6 year service, which started in December 2003, is £53M. Excluding staff supplied by IBM for system maintenance, there are 5.75 Full Time Equivalents (FTE) for Systems Support and front-line User Support services and 7 FTEs for computational science and engineering support, although this is planned to reduce over the lifetime of the service.

Science Communities

The principal communities making use of the facility reflect the relative contributions of the various UK Research Councils. The largest contributor among these, with 90% of the time, is the Engineering and Physical Sciences Research Council (EPSRC) and principally supports physics (excluding particle physics and astronomy), chemistry and engineering. Consequently, groups in these areas dominate usage though there are also active communities in the environmental sciences (9%) and, increasingly, in the life sciences (1%). Today, the largest user groups are in ab initio materials science.

GRID Infrastructure

The UK e-Science Programme has established a National e-Science Centre (NeSC) at Edinburgh linked to a network of 8 Regional Grid Centres at Belfast, Cambridge, Cardiff, London, Manchester, Newcastle, Oxford, and Southampton. A Globus GT2-based Grid infrastructure, with resources, middleware and applications from the e-Science Centres and the Central Laboratory of the UK Research Councils (CCLRC), has been deployed.

3. Networking

3-1. Global

GEANT

General Description

The European networks for research and education have been working together for many years (since early 90s) to have in place a shared infrastructure to interconnect them with broadband capacities. The successive networks, which were deployed in the past, were named as EUROPANET, TEN-34 and TEN-155. The current version of the pan-European backbone is named GEANT. This network is the main outcome of a EC funded FP5 project (GN1), carried by the consortium of NRENs. The coordination of the consortium (and of the GN1 project as well) is made by DANTE (Delivery of Advanced Networks to Europe, www.dante.org.uk), a not-for-profit company, based in Cambridge - UK, which is owned by the NRENs. The GEANT network is interconnecting more than 30 countries today, with gb/s capacity for most of them. Added to the intra-Europe interconnection service, DANTE is also providing access to the other NRENs in the world, like North America, Asia-Pacific, Latin America, Mediterranean Basin and Middle East, South East Europe, etc… Furthermore, GEANT is also currently working on new opportunities for gaining connectivity to other world regions, like Eastern Europe, Central Asia, India, South Africa, etc…

Technology available for GEANT

The overall topology for the pan-European infrastructure is made of a set of PDH/SDH/WDM links. The core of the network is made of 10 Gb/s segments, while the edge part of the network is using links with various capacities between 2.5 Gb/s et 34 Mb/s.

Worldwide connectivity is provide to Europe with the following links:

- The North American (and other world wide) research and education networks through 3*2.5 Gb/s links, while US has just started reciprocating the financial support of the trans-Atlantic link by putting in place a 10 Gb/s circuit between Chicago an the GEANT PoP in Amsterdam.
- The commercial Internet through a limited number of peerings with local/regional/national ISPs, while most of the Internet connectivity remains under the NREN responsibility.
- The access to the Japanese University network (Super SINET), which is directly connected to the GEANT PoP in New York
- The access to other world regions with various capacities, like 300 Mb/s for Russia, 200 Mb/s with the Mediterranean and 622 Mb/s with Latin America.

Organisation

Policy decisions are taken by the Consortium of NRENs (NRENPC), which is coordinated by DANTE (www.dante.org.uk). TERENA (www.terena.nl), which is the association of the NRENs in Europe, participates also to the NREN Policy Committee.

DANTE has been created in 1993 and is a private not-for-profit organisation. DANTE services are only offered to European NRENs. The shareholders members (presently roughly 13) elect the Board of Directors. The DANTE technical and administrative team comprises 30 people working in Cambridge (UK).

Budget per year

The overall budget comprises roughly 50 M€ per year, is provided by the European NRENs connected to GEANT plus the European Commission.

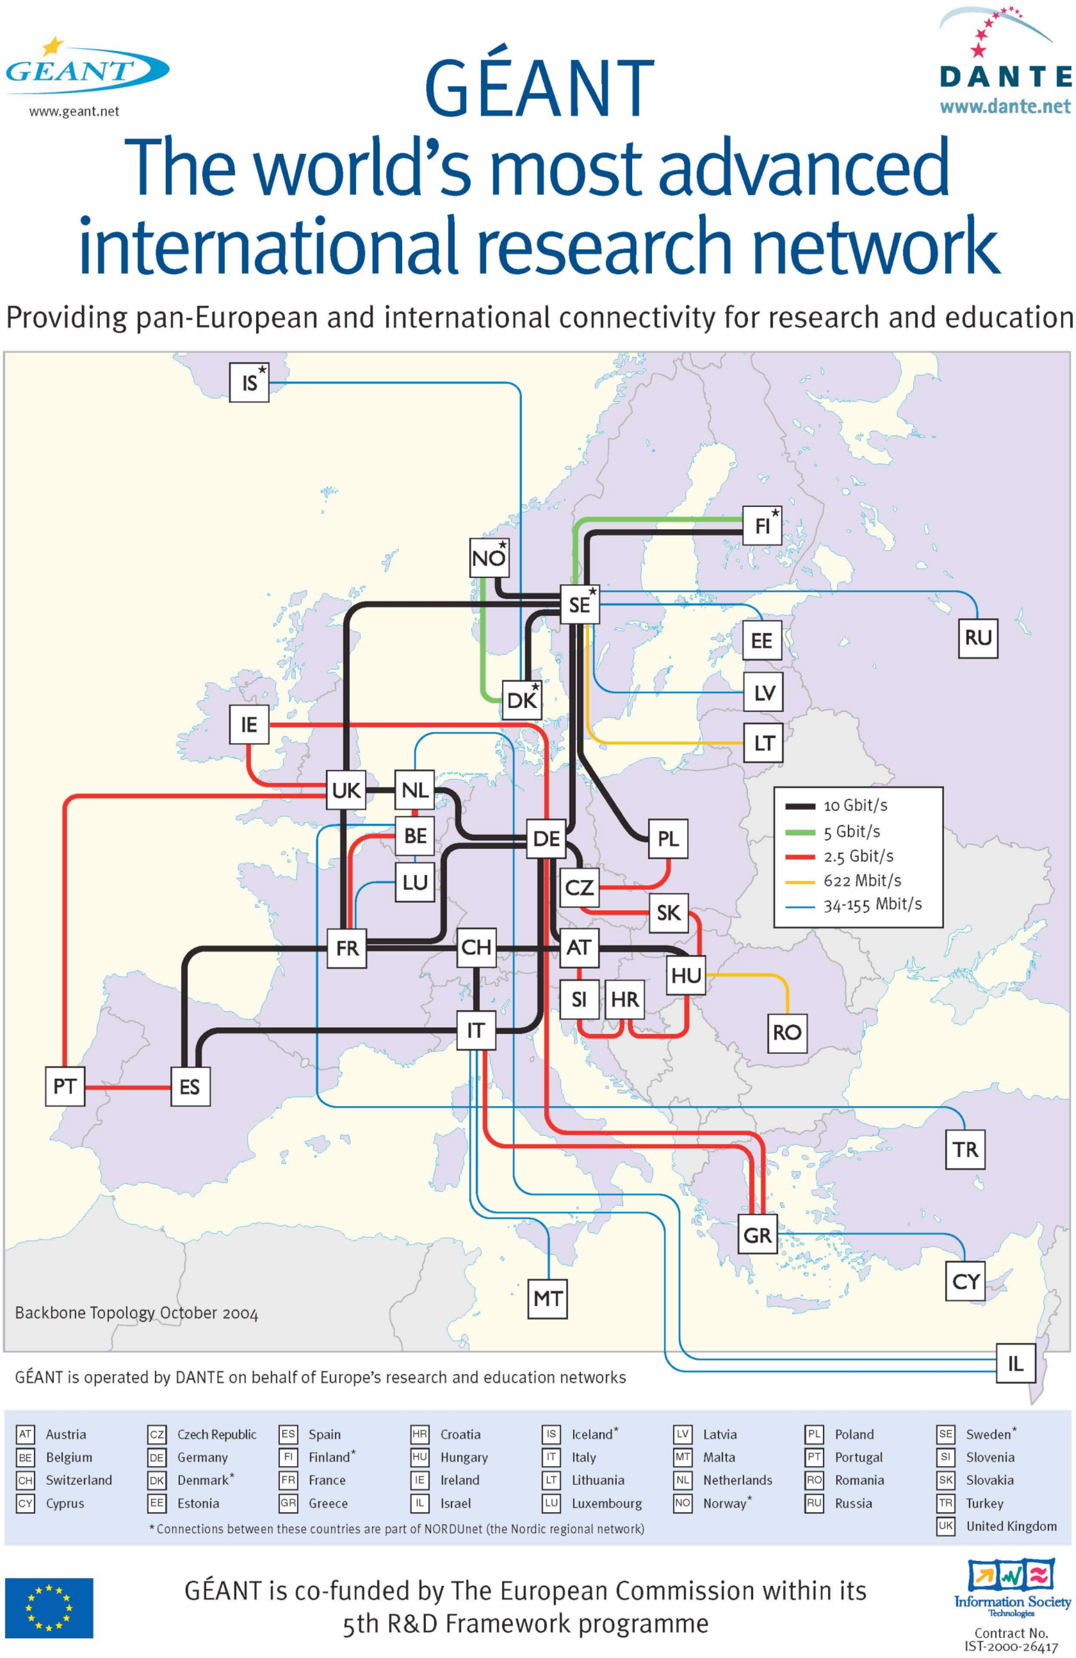

3-2. Per country basis

Belgium

General Description

BELNET is the Belgian national research network that started its activities in 1993 under the umbrella of the Federal Science Policy. BELNET, with its high performance Gigabit network, now offers the the best possibilities concerning internet access and the use of the research network. Since 1995 federal services and federal administrations are also allowed to make use of the BELNET services.

BELNET also manages BNIX, the Belgian National Internet Exchange. This exchange node is a central infrastructure by which Internet Service Providers (ISPs) that are active on the Belgian market can exchange traffic. This solution avoids that Internet traffic between Belgian Internet users has to go via other countries.

The BELNET network consists of a national and international infrastructure:

- The national network is built around two star-shaped structures, centralized in Brussels, from where datatramission lines at 2.5 Gb/s go to the 15 national POPs. This network is fully redundant to maximize availability.
- This national infrastructure is connected with European, North American (Internet2, CANARIE, Abilene, vBNS, …) and Asian research networks via a direct line with GEANT.
- It is also connected to the commercial Internet in the following ways:
    - The global Internet,
    - BNIX: direct connection to the Belgian Internet Exchange,
    - AMS-IX: direct connection to the Dutch Internet Exchange,
    - SFINX: direct connection to the French Internet Exchange.

The following types of customers are connected to BELNET:

- Administrations
- Education
- Research
- Universities

Today, about 500.000 end-users located in more than 150 Belgian institutes are getting broadband access to Internet via BELNET. Among these organizations are:

- All Belgian universities,
- Most of the Schools for higher, non-university education,
- Research centers (IMEC, KMI, VKI, Institute for Tropical Medicine, …),
- All federal scientific institutes,
- The FedMAN network that interconnects all federal administrations,
- Several federal institutes (the Belgian Senate, the Belgian Parliament, …),
- Several institutes belonging to the communities (Flemish or French),
- Several regional administrations.

Services

BELNET provides two basic services to its clients: unlimited access to the worldwide research network and fast access to the global commercial Internet. The access conditions and tariffs vary with the type and the activities of the institute that wants to connect:

- The official research and education institutes can have a free access to the worldwide research network but the access to the commercial internet has to be paid in function of the delivered bandwidth (tariffs at marginal costs),
- Administrations can connect to the worldwide research network and the commercial internet at a tariff in relation with the delivered bandwidth and at real cost,
- Private organizations can connect to the worldwide research networks but only for the strict use of research and development projects.

A BELNET connection also includes the following services:

- Registration of domain names,
- Access to native IP-Multicast,
- Permanent connection to an unlimited number of workstations, computers, …
- Free allocation if fixed and official IP addresses in relation with the needs of the customer,
- Use of FTP archives,
- Netnews, connection + full feed,
- 24x7 control of the access circuit,
- Helpdesk.

Czech Republik

General Description

Czech NREN is developed and operated by CESNET, z. s. p. o. Backbone core of Czech NREN (CESNET2) is composed by links with capacity of 2.5 Gb/s reaching all major entities within the national research and education network. Thirty-one universities and 229 hospitals, schools and libraries are connected to the CESNET, where the universities account for approximately 95% of all the generated traffic. Through the 31 university connections almost 200,000 students, professors and other staff members have connection to the network. The network topology has a character of multiple rings, which ensures redundancy for each PoP. Today 95% of the network provided by CESNET is fibre connections and the rest is microwave links.

The external connectivity consists of following links:

- 800 Mb/s (software-limited) commodity internet
- 2,5 Gb/s Géant
- 1 Gb/s Slovak NREN SANET
- 2,5 Gb/s Experimental connection CzechLight-NetherLight
- 2 Gb/s Czech Neutral Internet Exchange NIX.CZ

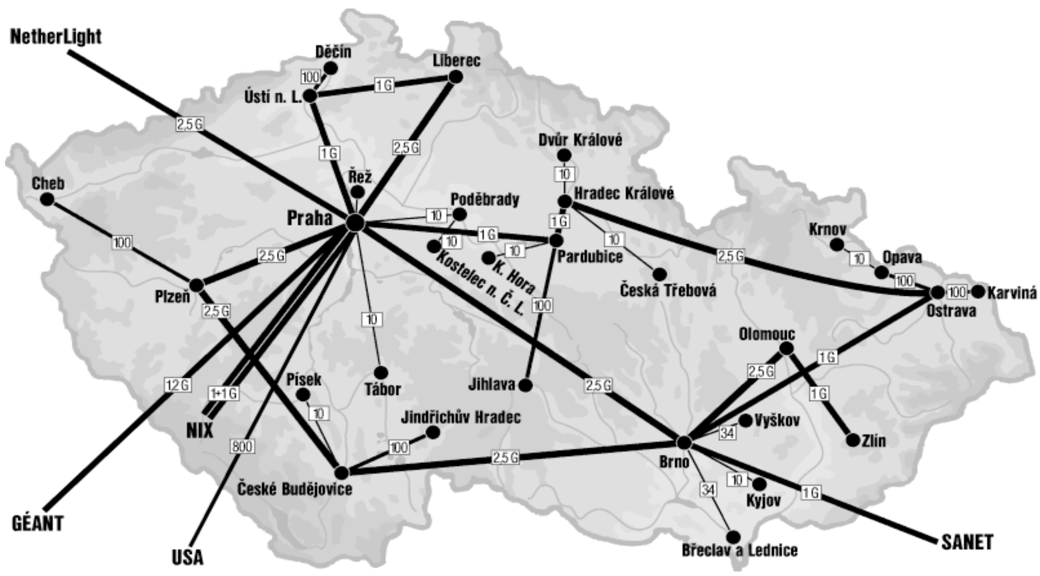

CESNET2 backbone topology

Denmark

General Description

The Danish Research Network, "Forskningsnettet", provides the network infrastructure for the Danish universities and other high level institutions of research and education. The Research-Network is based on a backbone ring with a capacity of 622 Mb/s + 2 x 100 Mb/s to which the four major universities are connected. From these universities regional connections are hosted to reach the remaining research institutions. Experiments with 10 Gb/s are currently being conducted.

- External links
  - The Geant network
  - Nordunet to Stockholm through a 2.5Gb/sec link
  - Nordunet to Oslo through a 2.5 Gb/s link

Budget per year

The overall budget comprises roughly 4.6 M€ per year, is provided as 50 % by government and 50 % by the institutions connected to the network.

Grid Infrastructure

The Danish Center for Grid Computing, DCGC, was established in August 2003 and seeks to set up a working Grid infrastructure in Denmark within the next two years. The Grid Center works closely with the Center for Scientific Computing and the Grid Centers primary responsibility is to provide the software and knowledge part of the infrastructure while the Center for Scientific Computing will still provide the hardware. In acknowledgement of the fact that Grid is still an active research area the majority of the 1 M€ granted for DCGC is spent on research scholarships which are matched 1:1 by the host institution of the researcher.

Finland

General Description

FUNET, the Finnish University and Research Network, serves the Finnish scientific community and connects over 80 research organizations and some 300 000 users in Finland. Funet also provides its user organizations with state-of-the-art data links to the global Internet network through the pan-Nordic NORDUnet network.

The backbone of FUNET is 2.5 Gb/s and FUNET has 2x2.5 Gb/s connections to NORDUnet network to Stockholm, Sweden.

France

General Description

All universities, research organisations and higher education institutions are interconnected through the RENATER network. RENATER is a public entity (“Groupement d’Intérêt Public”), of which the main stakeholders are:

- Ministries of Research and Higher Education
- CNRS (National Centre for Scientific Research)
- CEA (Atomic energy Agency)
- CNES (National Space Agency)
- INRIA (National Institute for Research in Computing Sciences)
- INRA (National Institute for Agronomical Research)
- CIRAD (Research Institute for Agronomical Research and Development)
- INSERM (Life Science and Health Research Institute).

RENATER provides its users with a broadband IP connectivity, together with access to other European and worldwide Research and Education networks through the GEANT backbone. RENATER provides also full Internet connectivity through commercial telecommunication operators backbone and also through a global Interexchange node in Paris, managed also by RENATER. Most users connect to RENATER through regional or metropolitan infrastructure. However the number of directly connected sites is increasing continuously, to allow a better end-to-end service to users.

Above the standard IP service, RENATER is also providing its users with advanced services support like Ipv6 connectivity, VPN capabilities, video-conferencing and VOIP, Quality of Service, CERT support, etc…

Technology available on RENATER

RENATER infrastructure is made of SDH/WDM links organized in closed loops, to allow optical redundancy in case of failure or maintenance. The standard link capacity is 2.5 Gb/s, with some paths with double connectivity, while there are still a few PoPs connected with lower capacity (622 Mb/s)

International connectivity of RENATER

- The GEANT network with a 10Gb/s access capacity,
- The North American (and other world wide) research and education networks through GEANT,
- The commercial Internet through a large number (80) of peerings with local/regional/national ISPs, at the SFINX (GIX) in Paris
- The commercial Internet through two global upstreams with 2.5 Gb/s each.

***Budget** per year*

The overall budget comprises roughly 27 M€ per year, is mostly provided by the Members of the “Groupement”, while 10% of this budget is coming from organisations connected to the network, which do not belong to the members.

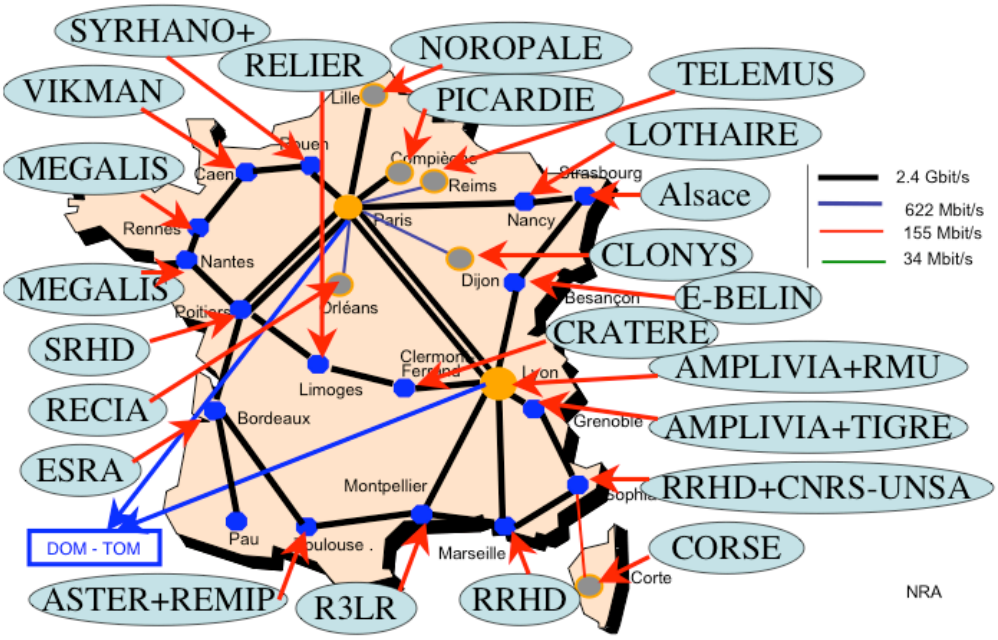

RENATER backbone topology

Germany

General Description

Almost all institutions of science, research and higher education in Germany are connected to the infrastructure of the Deutsches Forschungsnetz (DFN). The main part of the infrastructure is the Gigabit-Wissenschaftsnetz (G-WiN). The most important service available is DFN-Internet, providing IP-connectivity to research networks in Europe, North America and other parts of the world as well as global IP-upstream. Other services comprise Video-Conferencing (service: DFNVC), News delivery (service: DFNNews), a CERT for security issues (service: DFNCERT) etc. New services are being created based on user demands. Special access technologies are offered to users, a prominent example is a specialised solution for (national and international) GRID applications. Most of the services are configured (according to user's needs) with market offerings for components or partial systems which are all based on public tender procurements. G-WiN has been built with financial support by the federal ministry (BMBF), however the support had been restricted to the first three years only. Since 2003 financing comes from G-WiN users (i.e. institutions connected to G-WiN) only. DFN is integral part of European activities, which are organised in DANTE (where DFN together with other NRENs is owner) and TERENA (where DFN is member). More detailed information on DFN is presented at www.dfn.de.

Technology available for G-WiN

G-WiN is based on a SDH/WDM **platform** (provided by T-Systems, a subsidiary of Deutsche Telekom), which comprises a number of links. most of them at a capacity of 2.4Gb/s, but from 2004 on half of the links will be at 10Gb/s. This "toolbox" of links can be configured any time (nearly) on demand to form the network or specialised parts of the network.

The main service based on this **platform** is the IP Service (DFNInternet), which carries the overwhelming part of the traffic in the German research and education community. As of September 2003 the following distribution of **access links** to the G-WiN backbone (see below) is existing:

Less than 34Mbps	34Mbps	155Mbps	622Mbps	More than 622Mbps
213	198	54	5	3

In most cases an access links is for one research and education institution. However in some areas institutions are connected via a regional network.

External links *as part of the DFNInternet service are the link to :*
- the Geant network (presently a 2.4Gb/s link, from 2004 on a 10Gb/s link),
- the North American (and other world wide) research and education networks through GEANT,
- the commercial Internet through a huge number (85) of peerings with local/regional/national ISPs
- the commercial Internet through two global upstreams with 2.4 Gb/s each.

The **usage** of G-WiN (DFNInternet) is presently growing at a rate of a factor of 1.6 per year. The overall input volume transferred over the network is more than 1PByte per month. Several specialised solutions are being provided using the platform described above, the most prominent one is a lambda-connection between two supercomputer sites (Berlin and Hannover).

Organisation
The DFN association has been created in 1984 and is a private not-for-profit organisation with tax-exempt status. DFN services are only offered to institutions which are part of the research and education community in Germany. The members (presently roughly 300) elect a council of administration which elects the board. The DFN technical and administrative team comprises ca. 50 people working in to places in Germany (Berlin and Stuttgart). Most of the operation team is concentrated in Stuttgart.

Budget per year
The overall budget comprises roughly 40 M€ per year, is provided by the institutions connected to the network and it is managed by DFN.

Ireland
General Description
The academic network provider is HEAnet. HEA provides ~65% of the funding of the national research network HEAnet which in 2003 has a budget of 13 M€. The users pay ~30% of the cost.

All academic institutions connected by NREN (HEAnet) at 155 Mb/s - 2.5 Gb/s. HEAnet connections: 2.5 Gb/s to GEANT, 155 Mb/s to JANET, 1 Gb/s to QUB.

Italy
General Description
The Italian NREN is made, as most of others, of WDM link (2.5 Gb/s), running an IP layer to procure services to all users from universities and public research organisations. It is operated

by a consortium (GARR Consortium, www.garr.it), of which members are user organisations. Users sites are connected to the backbone nodes with PDH or SDH links.

GARR is connected to GEANT with a 10 Gb/s access port and has its own connectivity to the local and global Internet.

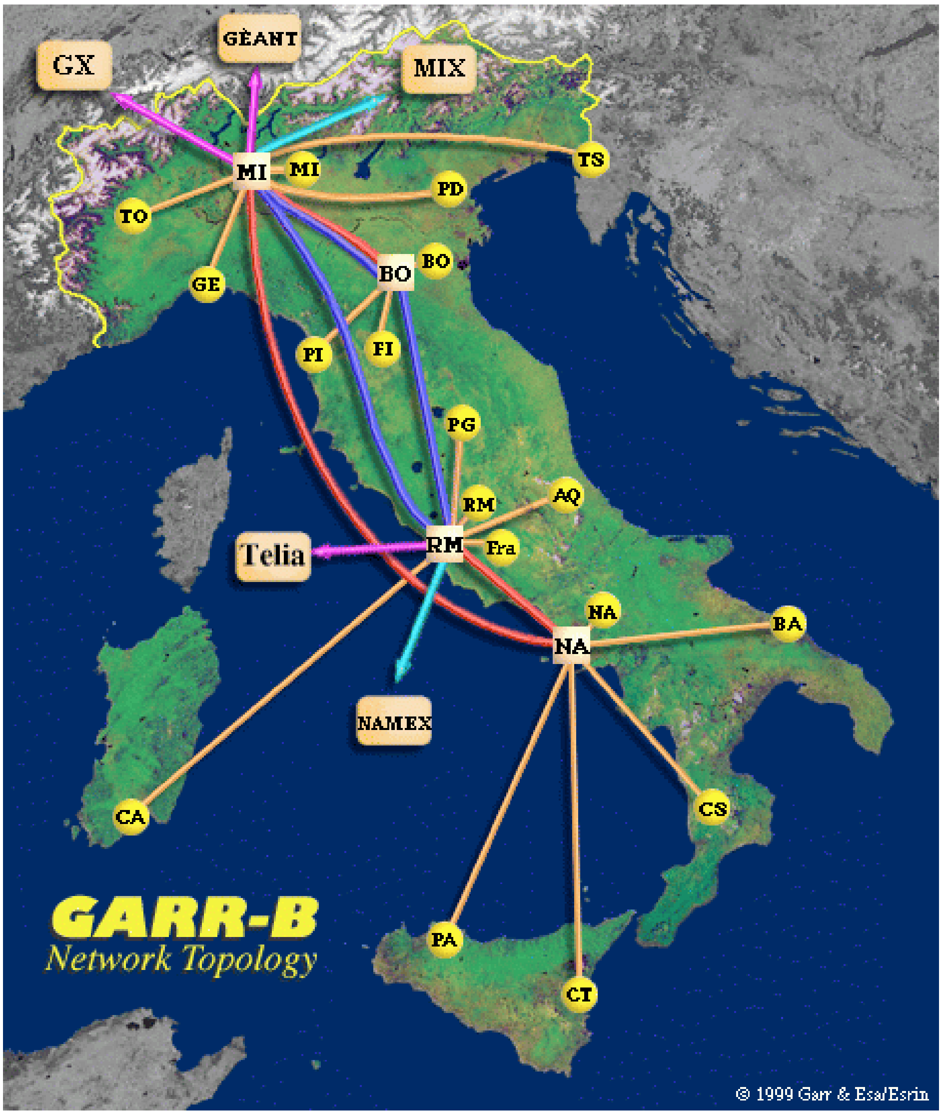

Spain

General Description

The backbone that supports communications services provided by RedIRIS is made up by a set of nodes distributed throughout the country. These nodes are interconnected by a meshed network with a 2.5 Gb/s core.

At present there are 18 nodes, one in each Autonomous Region. A node is a set of communications equipments that concentrate the backbone transmission media and access lines of the centres of each region. All these equipments are configured and managed by RedIRIS Network Operation Centre. At present RedIRIS gives access to more than 260 organizations, mainly universities and R&D centers.

Sweden

General Description

SUNET - the Swedish University Computer Network - interconnects local and regional networks at Swedish universities. The backbone capacity is 10 Gb/s. Major universities are typically connected at 100 Mb/s or 1 Gb/s. SUNET uses NORDUnet for international connectivity. The connection to NORDUnet is 2x2.4 Gb/s.

SUNET is coordinated by a board with representatives mainly from universities. The Swedish universities finance SUNET. The budget in 2003 is 17 M€. For further information see http://www.sunet.se/

United Kingdom

General Description

The UK NREN is SuperJanet. It is supported by JISC and operated by UKERNA. SuperJanet serves all Research and Education communities (www.ukerna.ac.uk) in UK and is also the background infrastructure for the e-Science programme. In terms of organisation and performances it is quite similar to DFN or RENATER descriptions, except the very centralized funding structure with JISC. Most of the advanced European NRENs use similar technologies and bandwidth capacities (IP/WDM with 2.5 or 10 Gb/s segments).

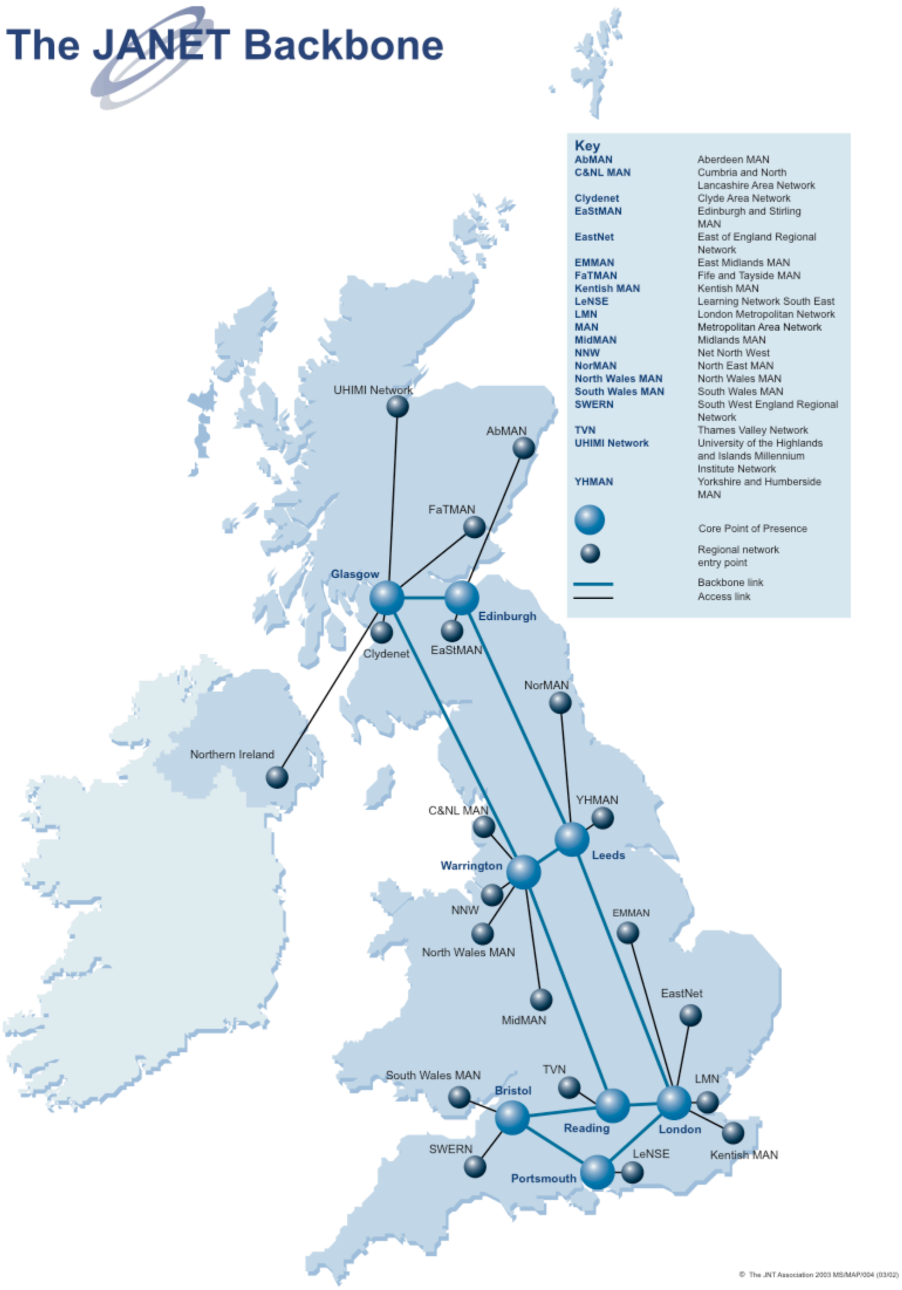
The JANET Backbone
Key
AbMAN Aberdeen MAN
C&NL MAN Cumbria and North Lancashire Area Network
Clydenet Clyde Area Network
EaStMAN Edinburgh and Stirling MAN
EastNet East of England Regional Network
EMMAN East Midlands MAN
FaTMAN Fife and Tayside MAN
Kentish MAN Kentish MAN
LeNSE Learning Network South East
LMN London Metropolitan Network
MAN Metropolitan Area Network
MidMAN Midlands MAN
NNW Net North West
NorMAN North East MAN
North Wales MAN North Wales MAN
South Wales MAN South Wales MAN
SWERN South West England Regional Network
TVN Thames Valley Network
UHIMI Network University of the Highlands and Islands Millennium Institute Network
YHMAN Yorkshire and Humberside MAN
Core Point of Presence
Regional network entry point
Backbone link
Access link
UHIMI Network
AbMAN
FaTMAN
Glasgow
Edinburgh
Clydenet
EaStMAN
NorMAN
Northern Ireland
YHMAN
C&NL MAN
Leeds
Warrington
NNW
EMMAN
North Wales MAN
EastNet
MidMAN
South Wales MAN
TVN
LMN
Bristol
Reading
London
SWERN
LeNSE
Kentish MAN
Portsmouth
© The JNT Association 2003 MS/MAP/004 (03/02)

4. Science cases

HEP

As a representative perspective for the future needs of the HEP community, the LHC experiment at CERN is chosen as an example. The following picture give some figures about the data transmission requirements:

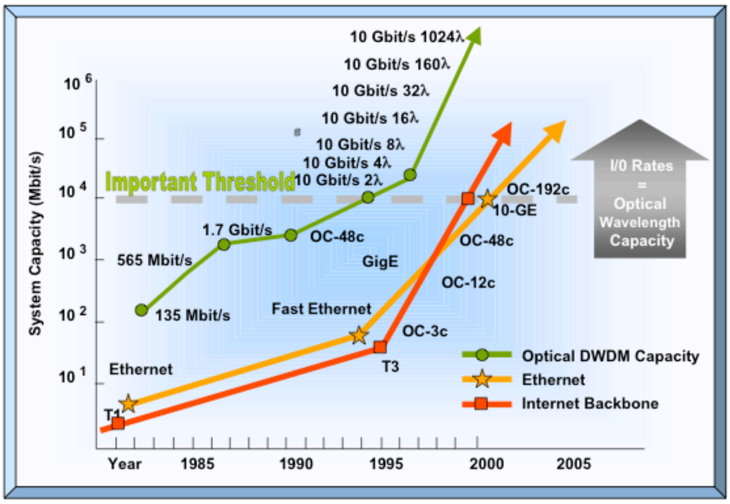

In terms of data processing, the LHC project expects to have Raw recording rate capacities of 0.1 – 1 GByte/sec, with the capability of accumulating data at 10-14 PetaBytes/year rithm, corresponding to about 20 million CDs each year, or in terms of storage to 10 PetaBytes of disk-space, with a processing capacity 100,000 of today's fastest PCs.

Another aspect of this project is to involve +5000 physicists belonging to around 250 organisations distributed in 60 countries. Beyond the scientific difficulties, a major challenges is therefore associated with supporting distance communication and collaboration around globally managed distributed computing and data resources

From an initial concept of hierarchical set of resources nodes (Tier model), the project management has increasingly involved the distributed GRID architecture concept as the appropriate tool for distributed data management.

This evolution, from large scale computing facilities, to distributed clusters and now towards fully distributed computing is properly represented in the figure below:

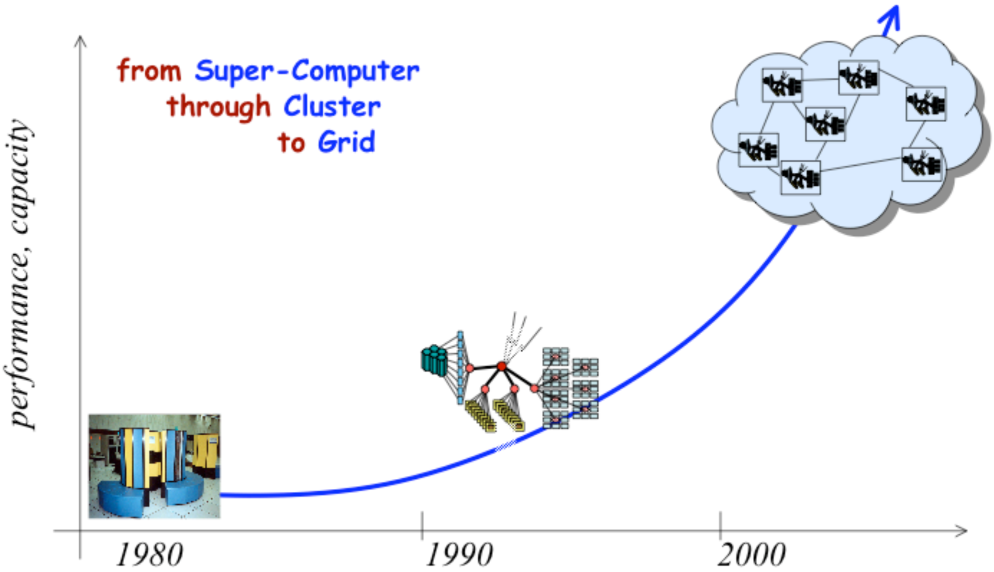

UK e-Science programme (http://www.rcuk.ac.uk/escience/)

In November 2000 the Director General of Research Councils, Dr John Taylor, announced £98M funding for a new UK e-Science programme. The allocations were £3M to the ESRC, £7M to the NERC, £8M each to the BBSRC and the MRC, £17M to EPSRC and £26M to PPARC. In addition, £5M was awarded to CLRC to 'Grid Enable' their experimental facilities and £9M was allocated towards the purchase of a new Teraflop scale HPC system. A sum of £15M was allocated to a Core e-Science Programme, a cross-Council activity to develop and broker generic technology solutions and generic middleware to enable e-Science and form the basis for new commercial e-business software. The £15M funding from the OST for the core e-Science Programme has been enhanced by an allocation of a further £20M from the CII Directorate of the DTI which will be matched by a further £15M from industry. The Core e-Science Programme will be managed by EPSRC on behalf of all the Research Councils.

What is meant by e-Science? In the future, e-Science will refer to the large scale science that will increasingly be carried out through distributed global collaborations enabled by the Internet. Typically, a feature of such collaborative scientific enterprises is that they will require access to very large data collections, very large scale computing resources and high performance visualisation back to the individual user scientists.

The World Wide Web gave access to information on Web pages written in html anywhere on the Internet. A much more powerful infrastructure is needed to support e-Science. Besides information stored in Web pages, scientists will need easy access to expensive remote facilities, to computing resources - either as dedicated Teraflop computers or cheap collections of PCs - and to information stored in dedicated databases.

The Grid is an architecture proposed to bring all these issues together and make a reality of such a vision for e-Science. Ian Foster and Carl Kesselman, inventors of the Globus approach to the Grid define the Grid as an enabler for Virtual Organisations: ‘An infrastructure that enables flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions and resources.’ It is important to recognize that resource in this context includes computational systems and data storage and specialized experimental facilities.

The Core Programme is structured around 6 key elements:

- A National e-Science Centre linked to a network of Regional Grid Centres
- Generic Grid Middleware and Demonstrator Projects
- Grid ‘IRC’ Research Projects
- Support for e-Science Pilot Projects
- Participation in International Grid Projects and Activities
- Establishment of a Grid Network Team

In its second phase, the Core programme has been modified to highlight 6 revised key activities:

- • A National e-Science Centre linked to a network of Regional Grid Centres
- • Support activities for the UK e-Science Community
- • An Open Middleware Infrastructure Institute (OMII)
- • A Digital Curation Centre (DCC)
- • New Exemplars for e-Science
- • Participation in International Grid Projects and Activities

In summary, the e-Science programme was begun in 2001-02 following the spending review 2000 and is being continued for the present spending review period, making a total OST investment in e-Science of £213M.

The programme has quickly demonstrated the utility of e-Science across disciplines in areas ranging from chemistry and biology to astronomy, atmospheric physics, engineering design, materials science and health and medicine. Scientists are engaging with each other in multidisciplinary teams to develop new Grid middleware across many pioneering projects and applications. The programme has seen significant industrial commitment, both major companies in the engineering and pharmaceutical sectors. The UK e-Science research community has also become significantly involved in major EU and other international programmes, including world standards bodies, as well as developing a focused long-term research agenda in computing, essential for maintaining the UK’s leading position in future Grid applications. The establishment of a national e-science Institute is acting as a focal point for activity engaging with international experts and ensuring that the highest academic excellence, standards and expertise in the UK are maintained.

Climate modelling and meteorology research

For many years, meteorology and climate modelling have been a highly demanding community in terms of computing and storage resources. Beyond the practical need to have performing tool to provide the civil society with ad-hoc weather forecast capabilities, the proper description of the global changes in climate are also of outmost importance for the understanding of large scale phenomena such as the greenhouse effets or the El Niño consequences.

The Meteoroly community is a rather well self-organized community in terms of HPC resources, with dedicated European computing centres like the ECMWF in Reading. Furthermore all the European Meteorology agencies are working routinely in close cooperation as:

- Austria
- Belgium
- Denmark
- Finland
- France
- Germany
- Greece
- Iceland
- Ireland
- Italy
- Luxemburg
- Netherlands
- Norway
- Portugal
- Spain
- Sweden
- Switzerland
- United Kingdom

Among the largest ones, the computing resources compared to the ECMWF are given below, and demonstrate the significant level of performances (in the Tflop/s arena) for a single community of users.

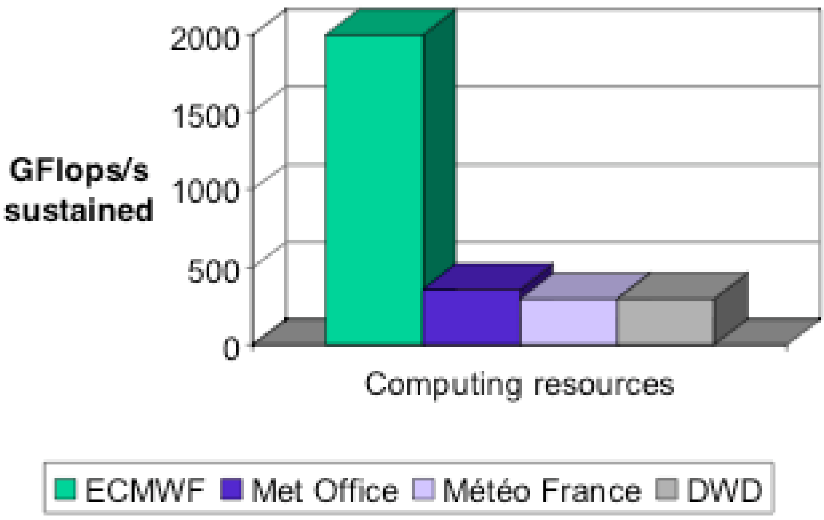

To prevent the saturation of the available resources, the community has also been very active in participating, as early adopter, to the deployment of GRID-type architectures. Based on the UNICORE software architecture, meteorology research was able to contribute significantly to the EUROGRID project and is now ready to jump onto the future ENES and EGEE projects. However, the computing applications of this community are rather heavy pieces of software with very long timescale for introducing changes or even slight improvements in the physical models or the numerical algorithms. In particular the move from vector machines (able to handle properly long vectors in structured matrices) to scalar type highly distributed is systems, is a real issue for this community. All these specific aspects make the community very proactive for requesting a European central computing facility, to complement the nationally funded existing resources.

To evaluate the weight of the climate modelling community in terms of computing needs, the following presents the main relevant computing centres, with their own ranking into the TOP500 list (dated June 2003) and the ratio of resources allocated specifically to climate modelling:

Country	**Centre**	**Top 500**	**Year**	**Performance Peak/Sustained**	**% climate**
Japon	Earth Simulator	1	2002	40960/*10000-20000*	50%
USA	NOAA NCAR	11 13	2002 2002	6758/*700* 6323/*650*	100% 100%
EU	ECMWF	15 16	2002	4992/ *500* 4992/ *500*	6%
UK	HPCX Hadley C.	12 161/180 474	2002 1997-99 2003 *2004* *2005*	6656/ *650* 788+763/*150* 256/60-*128*	8% 50-70%
DE	DKRZ	33	2003	1536/*500-800*	100 %
FR	IDRIS CEA	400 137 362? 10	2000 2002 2003 2001	320 1331 350 5120	30% 5% 35%

Another sensitive issuer for the Climate community is the processor architecture, i.e vector versus scalar processors. There are significant discrepancies between the peak performances (as displayed for instance on the TOP500 or many similar publications) and the real ones (so-called sustained performances). Recognizing for instance that the effective throughtput of a vector processor is like 30% of the peak value, while the ratio is down to less than 8% for a scalar architecture, the TOP500 may be reshuffled differently, in a more “honest” way. The picture below represents the TOP30 subset, according to this improved sorting, as presented recently in an ENES meeting (with pink color for climate studies):
(http://www.enes.org/Meetings/DeBilt2004/Presentations/index.html)

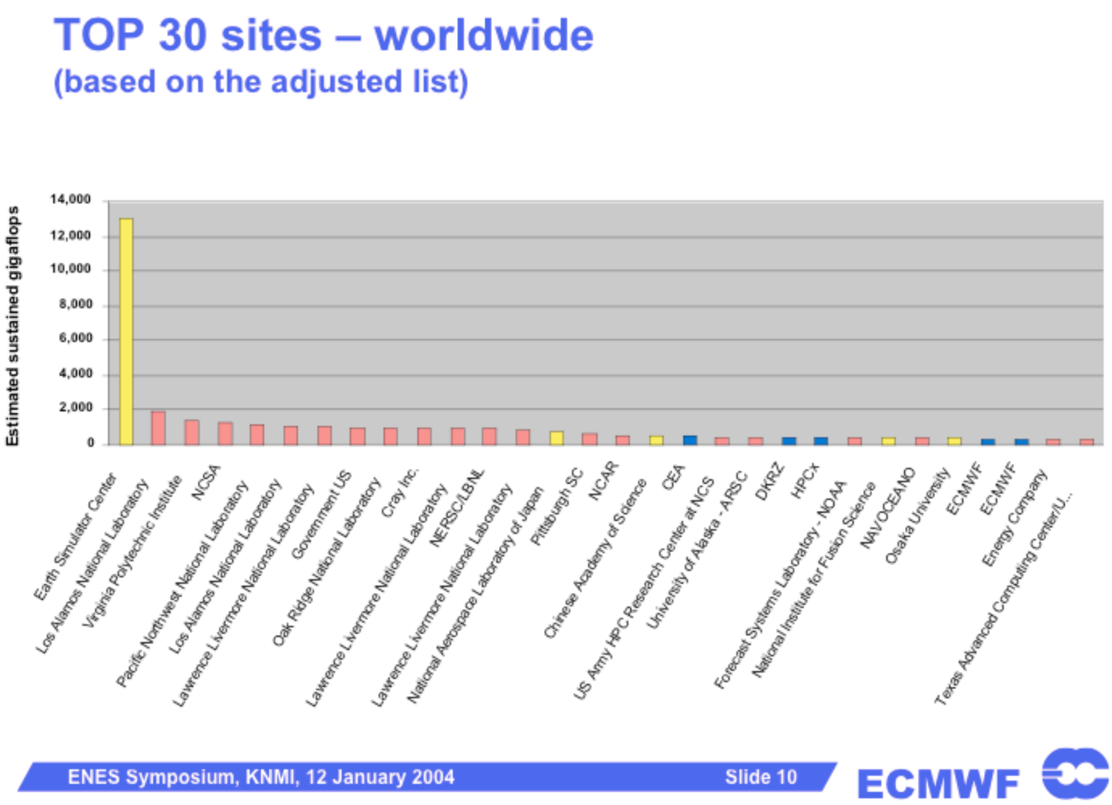

Now focusing on European resources, we end up with:

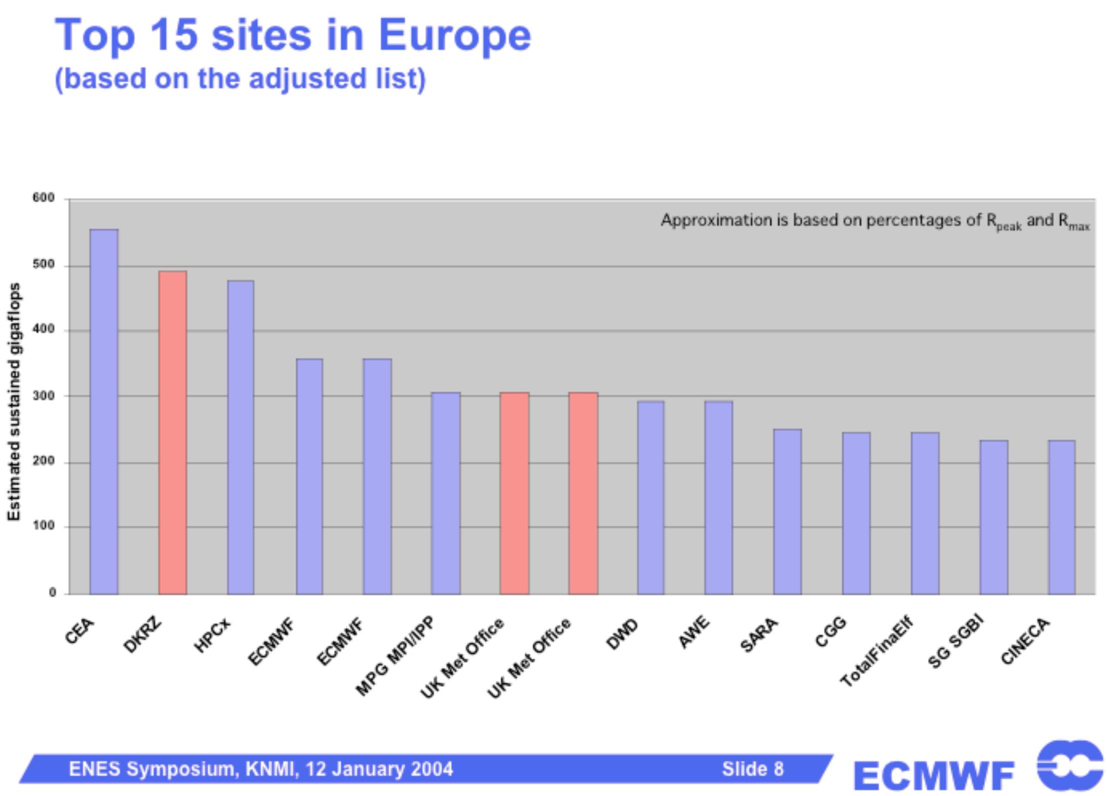

To complete this overview, there are significant disparities between countries for having dedicated HPC (like in Germany or UK) instead of a share in general purpose large scale facilities (like in France). There is certainly an issue to be discussed at the European level for the creation of multinational resource centre(s) on a model similar to the ECMWF facility, but dedicated to large scale climate modelling (including global changes) instead of mid term meteorology predictions.

Nano-sciences to nano-technologies

Nano-science regards the study of the physical, chemical and mechanical properties of objects whose size is in the order of few nano-meters. In this domain, thanks to the increasing power of computing facilities, numerical simulations are becoming a fundamental tool for theoretical studies. Nano-scale systems are 3d structures of atoms in the range from few hundreds up to few millions; only recently atomistic first principles simulations of such systems are becoming affordable.

Atomistic simulations can be classified in two main categories, classical simulations and quantum simulations. They differ in the degree of accuracy in which the inter-atomic forces are computed: in classical simulations the forces are calculated from parameterized functions, while in quantum simulations the forces are computed solving the Schrödinger equations of the coupled atomic and electronic system (to some degree of approximation). Some examples of this two class of simulations are Classical Molecular Dynamic, Ab-initio Molecular Dynamics (like Car-Parrinello methods), Classical and Quantum Monte Carlo simulations.

It is important to underline that, in order to understand and to study the physics of nanoscale processes, electrons have to be taken into account explicitly, and this does imply accurate quantum simulations. By mean of quantum simulations it is possible to investigate completely new phenomena, like those that arise from the coherent behaviour of electrons in nanoscale object (e.g. Quantum Hall Effect), but reaching this accuracy is highly demanding in computing resources. Thus, large computer facilities (computers with a power in the range of of 100TFlops-1Pflops) are fundamental for this class of computations.

Nevertheless quantum simulations (and computing power) are showing theirs limits in terms of scalability of the inherent physical system. By mean of quantum simulations just few hundreds of atoms for few picoseconds of simulated time are affordable. To bridge these limits new techniques are being developed, in example QMMM (Quantum Mechanics and Molecular Mechanics) is a technique that mix classical and quantum simulations, and this make the study of the interaction of the nano-scale objects with the environment, affordable. With the mixing of classical and quantum simulations, it is possible to use quantum accuracy when is needed (in the hot site) and classical simulations for the environment or the bulk material.

On the other hands, multi-scale simulations are being investigated to overtake the limit related to the difference between the time scale of electronic processes and relevant physical processes that involve nano-scale objects (e.g. Protein Folding). For multi-scale simulations classical and quantum atomistic techniques are integrated into a model of a continuum media. Both QMMM and Multi-scale approach require a huge computing power to afford systems whose size is relevant to study phenomena related to the behaviour of nano-scale objects.

Large computing facilities will also favour the technological transfer from the nano-science to the nano-technology, shortening the time from the discovery of some physical properties to the realization of a nano-scale device. In this cases the simulations will help uncover the fundamental principles governing the rational design of truly new materials.

Last but not least the basic building blocks of life (proteins or bio-molecules) are nano-scale objects and nano-scale simulations can help understanding the behaviour of these objects, moreover they can aid the building of devices that take advantage of bio-molecular properties, or devices that perform on bio-molecules some sort of manipulation.

Biotechnologies

From the genomic aspects…

The Bio-Informatics field is in rapid growth and the importance of Bio-IT is widely recognized in the industry. A typical application within Bio-IT currently is BLAST – Basic Local Alignment Search Tool, used for finding similarities in genomic data. The operations performed to do a BLAST comparison between two genome-sets is very simple and amongst the least demanding applications found in Bio-IT. This is interesting since even BLAST promises large computational problems in the future. The problem lies in the fact that speed of computers double approximately every 18 month while the size of the known genomic databases double every 10 months, this simply translates into the fact that running a BLAST search becomes increasingly slower even when keeping up-to-date with the hardware development – only continued investments in increasingly larger systems can keep the time of performing a BLAST constant. The rapid growth of the genomic databases could also strain the network as new database versions are available approximately every 6 hours, their size is measured in tens of gigabytes. Currently most groups seem to update their databases less frequently, e.g. weekly. Whether this represents their actual wish for update frequency or if it's a result of the available bandwidth is unknown. If the BLAST executions are moved into Grid computing the need for network bandwidth may raise.

Other, more challenging, applications are of high interest to the Bio-IT field including mapping gene-sequences to the production of proteins, simulating the processes that determine the absorption of molecules into cells and determining the sequence and timings of events that lead to protein-folding. The US is currently building a computer, which will, when fully finished, consist of more than 130.000 processors
http://www.research.ibm.com/resources/news/20031114_bluegene.shtml

To the associated functions

For most of the whole genome sequencing projects, the 3D structure and function of a large fraction of proteins remain unknown though their amino acid sequences are known[2]. This is still one of the key problems in bioinformatics and impressive progress has been made in the last years[3]. One of the most important contributions to the hardest version of this problem, the so called de novo protein structure prediction where no sequence similarity to proteins with known 3D structure could help the prediction, is the introduction of structure fragment assembly methods[4] using simulated annealing and genetic algorithms for optimization. We are currently investigating an alternative sampling strategy for this task.

On single processors we currently can optimize small proteins having less than 50 amino acids using a fragment library based on around 400 proteins having an average size of 200

---

[2] D. Baker et al., Protein Structure Prediction and Structural Genomics Science, v. 294, p. 93 (2001)

[3] C. Venclovas et al.,, "Assessment of Progress Over the CASP Experiments" PROTEINS:Structure,Function,and Genetics, v. 53, p. 585-595 (2003)

[4] P. Bradley et al.,, Rosetta Predictions in CASP5:Successes,Failures,and Prospects for Complete Automation, PROTEINS:Structure,Function,and Genetics, v. 53, p. 457-468 (2003)

amino acids and approximating each amino acid using a single reference atom. In this setting, the optimizations use 1GB of memory and take one day of CPU time on a 2.53 GHz Pentium4 PC.

The limitations on single processors and small clusters currently are the floating point CPU performance and the memory size. The communication bandwidth is in most algorithms, not critical, except if the cluster is used for large memory emulation and non-parallel prototypes of the algorithms.

Relevant problem sizes today, involve the optimization of proteins having a few hundred amino acids using a library based on a few thousand proteins with an average size of a few hundred amino acids and using a full atom representation of each amino acid (10 atoms on average, excluding hydrogen atoms). Such problems require systems with hundreds of nodes today. Larger problems, although relevant, are not considered yet due to their computational demands.

Astronomy

Among the various aspects of communications and computation, proper to the Astronomy community, two illustrations are given in this report: The first one address the networking interest for long base interferometers (JIVE), and the second one, the access to a global stellar data base.

For the European VLBI (Very Long Baseline Interferometry) Network, the radio astronomy central correlator is operated by JIVE[5] (the Joint Institute for VLBI in Europe) and is hosted at ASTRON[6] - the Netherlands Foundation for Research in Astronomy - which is located at Dwingeloo in The Netherlands.

eVLBI is a concept in which high capacity networks are used to transport the large amounts of data that VLBI measurements can produce. This has over the last few years become more appealing than the traditional approach taken to VLBI data transport, which has made use of high capacity magnetic tapes. The tape-based technology is ageing, becoming less well supported and proving expensive. On the other hand, telecommunications capacity in much of Europe has dramatically fallen in price - to the point where the eVLBI approach may prove more cost effective than the traditional VLBI techniques. Additional and significant advantages are also realised when adopting eVLBI; first and foremost amongst these is the notion of real-time eVLBI measurements. This is where observation and correlation can be done in near real-time thereby significantly reducing the time between observation and analysis of results from weeks to hours.

For the European eVLBI Proof of Concept (PoC), DANTE and the relevant NRENs have agreed to support up to 6 radio telescopes operating at speeds up to the full rate of 1Gb/s. These radio telescopes are listed in the table below along with the relevant NREN that will be providing the network access.

[5] http://www.jive.nl
[6] http://www.astron.nl

Country	Radio Telescope Location	NREN
DE	Effelsberg	DFN
IT	Medicina	GARR
NL	Westerbork	SURFnet
PL	Torun	PSNC
SE	Onsala	NORDUnet (SUNET)
UK	Jodrell Bank	UKERNA

Countries and NRENs participating in the eVLBI PoC

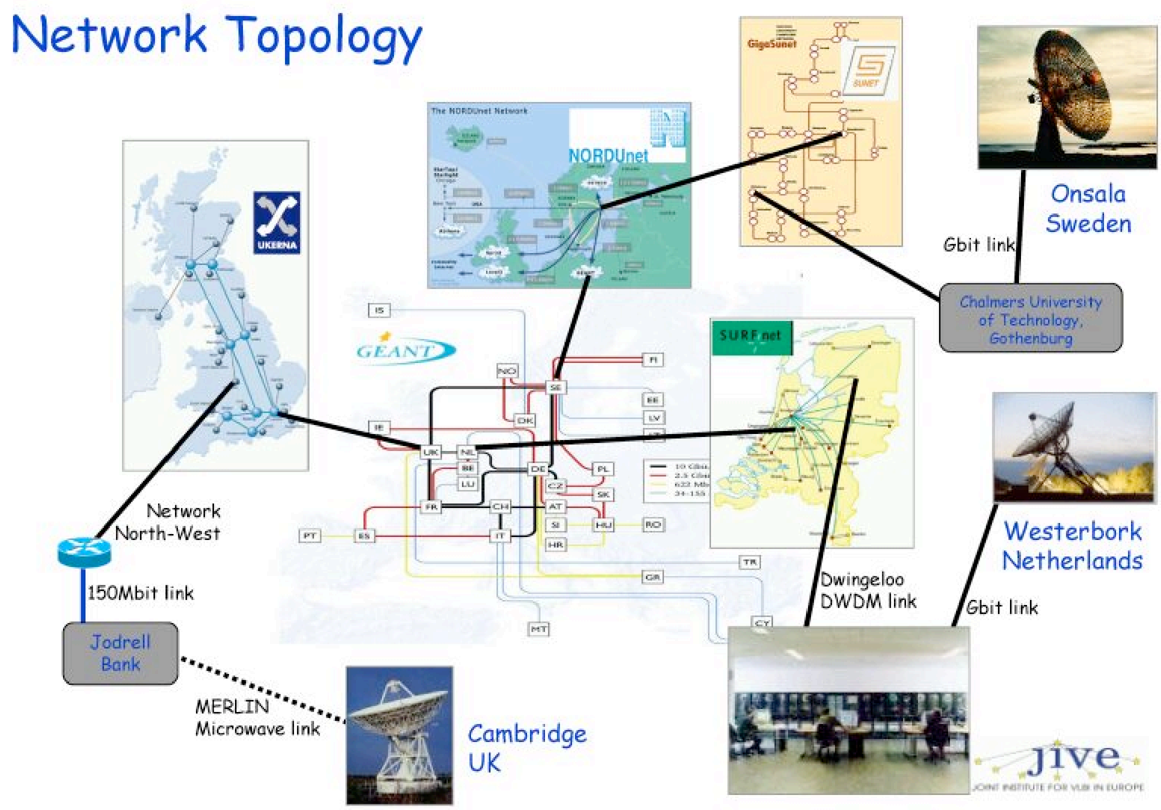

Beyond this Proof of Concept phase, which was demonstrated in early January 2004 with the first "network-produced picture"[7], there are now much stronger expectation for the community to have access to 10 Gb/s capacities and to be able to correlate pictures, not only with the first 6 instruments, but also with others located worldwide, especially in America and Asia.

The second aspect is a matter of data access. The CDS (Centre of Space Data), located in Strasbourg is a key element of the Virtual Observatory project. This centre is managing several databases of astronomical information and brings them at the disposal of users through a high performance web interface (http://cdsweb.u-strasbg.fr/). This portal gives access to various sources as:

[7] http://www.dante.net/upload/pdf/eVLBI_release_final_20040127171222.pdf

- **Simbad** reference database (Fr - US)
- **VizieR** catalogue service (Fr - Canada - US - Japan - India - UK - Russia - China)
- ftp access to catalogues: Astronomer's Bazaar - Submission guidelines
- **Aladin** sky atlas
- TOPbase database of the OPACITY project
- DENIS data release
- Dictionary of Nomenclature (Fr - Japan - Russia USA)
- INES Archive of IUE ultraviolet spectra

But also to focused bibliography sources and global projects links as :
- ADS* abstract service and scanned articles
- Astronomy & Astrophysics - CDS site*
- AJ* - ApJ* - PASP* mirror site at CDS
- A&A, A&AS and PASP abstracts
- A&A document map - ApJ document map
- Astrophysical Virtual Observatory - (AVO)
- Astrophysics Data Centers Executive Council - (ADEC)
- IDHA project
- MDA project
- Interoperability Standards and Tools for the Virtual Observatory
- GLU development site

Specific requirements for this type of global services are very demanding in terms of storage and transfer, with typical transfers of 20 TB size, which cannot be performed properly with standard protocols.

Astrophysics

There are various important problems to address in the field of Astrophysics, like Computational Fluid Dynamics, plasma physics or gravitational N-body to simulate the cosmological distribution of baryons (gas) and dark matter (n-body).

The type of problems which are currently feasible today are simulations with $N = 2000^3$ grid points, with the well known Reynolds number limitation ($N \sim Re^{-9/4}$), which prevents to have the required fine scale resolution. For particle-based simulations, the current limitation is in the range of 6 billions.

The major bottleneck is, in common with fluid dynamic research (including also meteorology), the size of the core memory. The significant progress to be expected will come from an increased resolution in solving problems (like more mesh points or a much larger number of particles). However, the challenge is not the “theoretical aggregate memory”, but the central core memory with very low latency, available to the computing units. As these problems are based on the handling of very long vectors and band-structured matrices, the large scale vector machines still represent a main issue, with regard to the various available architecture options.

As an objective, the goal of running $10,000^3$ nodes time-dependent computations with spectrum accuracy, is certainly a significant target to match the astronomical observed data sets.

LOFAR

The transnational project LOFAR (Low Frequency Array; www.lofar.nl) and its subproject LOIS (LOFAR Outrigger in Scandinavia; http://lois-space.org) comprise the building in Europe of a complex, wide-area, time-coherent digital sensor array for space and environmental studies. Funding of circa 70 M€ has been granted for the LOFAR project proper and the construction of this facility is now underway. Primarily intended for radio based astrophysics with a sensitivity that is 10-100 times higher than any existing radio telescope, LOFAR will operate in the 10-240 MHz frequency range thus providing, for the first time, a view on space at the long-wavelength limit at and immediately above the ionospheric radio frequency cut-off.

LOFAR's 13,500 digital electromagnetic sensors will be clustered in roughly 100 stations distributed over a region 400 km across Netherlands and North-Western Germany. Connected through purpose-built ultra high-speed fibre optical links capable of handling aggregated data streams of up to 25 Tb/s and with a 40 Tflops on-line compute capability, LOFAR will be at -- or even beyond -- the digital forefront when it goes on the air in the 2006-2007 time-frame. With its innovative, fully digital interferometric array design, LOFAR will allow multiple simultaneous beams to be synthesised in software, will possess extreme frequency agility, and will provide an exceptionally flexible system for distributed control, signal processing, monitoring, and remote operation. A LOFAR test station at Exloo, Netherlands, is already in operation.

The LOFAR sub-facility LOIS will extend LOFAR scientifically and technologically by providing a high-performance IT and telecom infrastructure. Distributed in southern Sweden with Växjö and Kronoberg County as hub, the LOIS software configurable network of approximately 2000 sensors and emitters, connected through a dedicated high-speed fibre optic network, will augment LOFAR with a deep-space radar capability. Thus the LOFAR-LOIS combination will provide ground-based atmospheric, ionospheric, magnetospheric, and solar physics (including particle storm forecasts and other space weather applications) capabilities which are truly outstanding.

In addition, LOIS will be used as a realistic, full-size proof-of-concept test-bed for telecom and IT network research and development. The circa 2 M€ funding obtained for LOIS to date, is being spent on hardware development, advanced numerical modelling, development of a GRID-enabled database manager for large amounts of streaming data capable of processing and storing combined LOFAR-LOIS data streams of up to 1 PB per day, and on the development of an adaptive GRID-service architecture system. A LOIS demonstrator sensor network, including radio units, a fibre optical network, and a prototype of the GRID-ified database manager has recently been installed around Växjö and proof-of-concept experiments are now being performed using this installation.

The successful utilisation of the LOFAR-LOIS combo for deep-space radar investigations of large portions of the inner heliosphere, will depend critically on the access to real-time transfer of data between the `local' LOFAR and LOIS network infrastructures at the 10-50 Gb/s level and on the pooling, with GRID technology, of the compute and data storage facilities at the two sites.

5. Concluding remarks

5-1. International situations

5-1-1. HPC in US

In the 70s, the supremacy of US in high performance computing was relying on several US vendors (CDC, IBM, etc…) but also on a one-of-a-kind parallel computer, known as the ILLIAC-4. In the early 80's, the US administration decided to base its high end computing resources on the standards of industry production. That led to the enhance development of CRAY computers for which the various federal agencies were among the best customers. However, in the 90s, the Administration realized that the unique force of industry and market was not sufficient to maintain the US at the cutting edge of the HPC performances. The competition with the Japanese manufacturers was certainly one element to consider at that time. Then the decision to stop real tests for nuclear weapons gave also a rationale for developing further the capacity of the numerical simulations. As a result, the US administration launched very ambitious initiatives named ASCI.

The ASCI vision (http://www.llnl.gov/asci-alliances/pp/002exsum.html#vision), as presented initially reads as:

To realize its vision, ASCI will create virtual testing and prototyping capabilities based on advanced weapon codes and high-performance computing. Virtual testing is the use of predictive simulations, based on experimental data, to assess and certify the safety, performance, and reliability of nuclear systems. Today, virtual testing and prototyping exist in rudimentary forms. Dramatic advances in computer technology have made virtual testing and prototyping viable alternatives to traditional nuclear and nonnuclear test-based methods. ASCI will provide computational and simulation capabilities that will help scientists understand aging weapons, predict when components will have to be replaced, and evaluate the implications of changes in materials and fabrication processes to the design life of the aging weapon systems. This science-based understanding is essential to ensure that changes brought about through aging or remanufacturing will not adversely affect the enduring stockpile. To meet the needs of stockpile stewardship in the year 2010, ASCI must solve progressively more difficult problems as we move away from nuclear testing. To do this, applications must achieve higher resolution, higher fidelity, three-dimensional, full-physics, and full-system modeling capabilities to reduce empiricism. This level of simulation requires high-performance computing far beyond our current level of performance. A powerful problem-solving environment must also be established to support application development and enable efficient and productive use of the new computing systems.

The ASCI program recognizes that the creation of simulation capabilities needed for virtual testing and prototyping is a significant challenge. This challenge is on par with many aspects of the original Manhattan Project and requires the science and technology resources available only at the national laboratories. This challenge will require close cooperation with the computer industry to accelerate their business plans to provide the computational platforms needed to support ASCI applications. Universities will also play a critical role in advancing the research and development needed for this unprecedented level of simulation.

This vision was traduced with various large scale initiatives, recognized with different colors (ASCI-Blue, ASCI-White, etc…) based on dedicated platforms, based on industry components (processors) put in configurations of unprecedented sizes. From this launch time, the computing platforms have evolved a lot. And the Teraflop capacity challenge has been largely over-passed today.

Although not detailed in this report, information about the various ASCI platform are available from http://www.llnl.gov/asci/platforms/platforms.html

Most figures corresponding to their performances are the ones supporting the information given in the TOP500 list.

Besides these ASCI initiatives, the NSF has also pushed for a significant increase in the efficiency of HPC, through the so-called TeraGrid initiative.

The TeraGrid is a multi-year effort to build and deploy the world's largest, most comprehensive, distributed infrastructure for open scientific research. By 2004, the TeraGrid will include 20 teraflops of computing power distributed at five sites, facilities capable of managing and storing nearly 1 PB of data, high-resolution visualization environments, and toolkits for grid computing. Four new TeraGrid sites, announced in September 2003, will add more scientific instruments, large datasets, and additional computing power and storage capacity to the system. All the components will be tightly integrated and connected through a network that operates at 40 Gb/s per second.

The ETF is being implemented through a series of coordinated NSF investments that began in FY 2000 with the Terascale Computing System(TCS) and continued in FY 2001 with support of the Distributed Terascale Facility (DTF). In FY 2002, TCS and DTF resources were integrated via an extensible, high-speed optical backbone, thereby creating the Extensible Terascale Facility. Under the NSF ETF activity, current partners include Argonne National Laboratory, the Centre for Advanced Computing Research (CACR) at the California Institute of Technology, the National Centre for Supercomputing Applications (NCSA) at the University of Illinois at Urbana Champaign, the Pittsburgh Supercomputing Centre (PSC) and the San Diego Supercomputing Centre (SDSC).

In October 2002, the Pittsburgh Supercomputing Centre (PSC) at Carnegie Mellon University and the University of Pittsburgh joined the TeraGrid as major new partners when NSF announced $35 million in supplementary funding.

Another $10 million in NSF awards in September 2003 adds four additional sites to the partnership: Oak Ridge National Laboratory (ORNL), Oak Ridge, TN; Purdue University, West Lafayette, IN; Indiana University, Bloomington; and the Texas Advanced Computing Center (TACC) at The University of Texas at Austin. Primary corporate partners are IBM, Intel Corporation, and Qwest Communications. Other partners are Myricom, Sun Microsystems, Hewlett-Packard Company, and Oracle Corporation (www.teragrid.org).

As a demonstration of the fast evolving feature of the ranking for supercomputers worldwide, the two main systems, as referred to in the latest edition of the www.top500.org are now again located in US. The first one is a DOE system built by IBM, providing 70 Tflops (http://www.top500.org/sublist/System.php?id=7101), while the second one, based on SGI products, is providing 52 Tflops, based at NASA Ames Research Centre (http://www.top500.org/sublist/System.php?id=7288).

5-1-2. HPC in Japan

There is a specific issue about the current situation of HPC in Japan, which is like in US, related to a strong commitment of the Japanese government to support its computer industry. In the past, Japan was successful to create a "Numerical Wind Tunnel", based at NAL, which was the biggest machine at that time, made of 156 vector processors, made by Fujitsu. Today, a more recent project, based on a partnership with NEC, correspond to the "Earth Simulator" facility, made of 640 nodes of 8 vector processors. This machine has been leading the TOP500 from 2002 to May 2004.

Beyond its intrinsic advantage as far as performances are concerned, the availability of this large computer is stressing the whole HPC world. It is certainly the only place where some computations can be performed today. As a result, there is now a real "affluence" of researchers to go there. This is very beneficial for the international cooperation, but on a longer term, presents a serious risk to produce a brain drain towards the Earth Simulator environment. Because of the human investment for developing numerical models for complex physics, and the long time scale of the research project, it would not make sense to develop models in Europe, which could be exploited in Asia (or in US). This concern is of increasing importance in several scientific communities, such as computational fluid dynamics or climate modelling.

5-1-3. Networks in US

Broadband networks have been obviously developed in US at the same time as the large HPC platforms were developed. Unlikely the European model, the US advanced networks were supported on a community based model. The Department of Energy (DoE) is supporting its own infrastructure (ESNET) dedicated to Energy Sciences, which includes the particle and high energy physics. NSF has supported its own backbone infrastructure (NSFnet) to interconnect HPC centres in universities, which was replaced later by the vBNS (and later the vBNS+ of MCI) and finally abandoned and replaced by a solicitation mechanism (similar to a call for proposals). The major US universities create a consortium (UCAID) to develop their own advanced network, known today as Abilene, which is one significant contribution to the Internet-2 project (http://www.internet2.edu/). To make it short, the Abilene network, which is the most advanced of the US academic networks is roughly comparable to the current European GEANT network.

Besides these production-quality infrastructures, the US is also promoting advanced infrastructures, both in terms of software developments and network infrastructure. Among them, there are the ETF (Extensive Terascale Facility), and the NLR (National Light Rail). The TeraGrid distributed testbed is a dedicated optical infrastructure used to interconnect the TeraGrid nodes at 40 Gb/s. Beyond this, another initiative took place to create a company, named FIBERCO, of which the role is to provide the academic research community with dark fibres across the US, to aim higher bandwidth than the currently existing ones.

5-1-4. Networks in Japan

The situation in Japan is rather simple compared to US. There is a basic optical infrastructure, available to the research communities, which is used, mostly by the Super SINET network.

Super SINET network (Science Information Network, www.sinet.jp) is operated by the NII (National Institute of Informatics).

Super SINET is an ultrahigh-speed network intended to develop and promote Japanese academic researches by strengthening collaboration among leading academic research institutes. The e-Japan Priority Policy Program announced by the IT Strategic Headquarters in March 2001 referred to this network, the world's fastest Internet for research, based on 10 Gbp/s optical communication technology. The National Institute of Informatics belonging to the Ministry of Education, Culture, Sports, Science and Technology has been operating the network since January 4, 2002.

For the time being, the network will be used as a basis for study information connecting universities and research institutes in the five fields of high energy and nuclear fusion; space and astronomical science; genome information analysis (bio-informatics); supercomputer-interlocking distributed computing (GRID); and nanotechnology. The Internet backbone connects research institutes at 10 Gb/s and the leading research facilities in the research institutes are directly connected at 1 Gb/s. The IT-Based Laboratory (ITBL) Project also uses Super SINET as its foundation.

5-2. Funding model issues

In the HPC world and for Research Networks as well, there is a widely accepted principle that the end user is not directly charged for the use of the resources. The most common principle is that these costs are supported by the research organisations themselves, either directly (a university or a research institute may buy a big computer or a broadband access to a NREN), or indirectly when the buyer is at a higher level (local or national government level, etc…).

For the HPC part, there is a clear borderline between the national funding and the European (or multinational) support. Besides the resources made available from international organisations (EMBL, ECMWF or CERN are certainly good examples), the general rule is that there is no funding in terms of hardware or capital investment at a high level than the national one. This is slightly modified by current FP6 projects like DEISA or EGEE, where the EC funds will contribute significantly to the development and the deployment of services on a pan-European scale. However, there is no incentive or intention from the Commission to provide a sustainable High Performance Computing service. Furthermore, these projects exclude any funding for HPC hardware. As some communities are expecting a continuous growth for the computing resources, there is certainly room for debating about the option for having massive computing equipment funded collectively by the EU member states. US (number of systems) and JP (size of the Earth Simulator) have clearly a significant advance over EU in this area. It is probably linked to the fact that both are directly or not, supporting their own hardware industry. However, there is a long term strategy issue for EU to loose the existing know-how and expertise in modelling activities, just because the most talented people will aggregate to the most powerful computing machines.

For the networking side, the situation is completely different. The NRENs are already organised into an European consortium, which interacts with the Commission to share the expenses of the pan-European backbone (GEANT today). This specific role of the Commission is very valuable as global cement to make a unique shared infrastructure, which serves simultaneously as multi-disciplinary communication backbone and as support for specific test-bed or discipline oriented usages. Furthermore, the core backbone is still a tiny part of the whole picture. In terms of funding, the global scenario to interconnect any combination of end users across Europe, follows the 1 : 10 : 100 figures. When 1 € is needed at the core level, 10 € are required at the national level and 100 € must be provided locally (campus or regional networks). In the current scenario, the EC contribution is therefore worth 0.5 €! while most efforts come from the countries (national and local funding). The capacity of the European networks is clearly at the same level has the ones in other regions, and the existing organisation model of the NRENs, with their long term partnership with the EC, is taken as a reference model for others places of the world, where the NRENs are not yet developed as they are in Europe.

5-3. Usage policy

The matter of usage policy is not very sensitive, because of the global funding model of these resources.
For the network side, because the funding is based on taxpayer's money, there is a strong requirement that the use to the Research network does not impact badly the commercial sector of the telecoms. Therefore all NRENs are enforcing the compliance of any user to an Acceptable Use Policy, which worded nationally but means that access to the infrastructure is restricted to professional activities in universities or research institutes. There is no a single European AUP, but all NRENs must have one.

For the HPC side, this is significantly different. The resources allocation is mostly done, based on the scientific quality of the activities, which can be reviewed globally for a research organisation, or specifically project per project, usually based on peer review model. However, in that case, the allowance is made for a given resource and it is not a general rule to have the same scientific evaluation accepted by different computing centres, especially in different countries. Any transborder exchange of resources will be limited by the upstream funding constraints of each HPC centre.

The GRID infrastructure, as it allows to access totally distributed resources will certainly have a strong impact on this local scientific review process, and will certainly be opposed, soon or later, to the very distributed funding of the hardware. Two computing centres of equivalent sizes, will certainly be happy to share resources and users, in a balanced scheme. However, if the centres are very disparate in size (and user base) then the balanced sharing will become much more difficult to achieve. It is expected that in the near future, the GRID infrastructure will develop its own policy rules to circumvent the existing model with a much more appropriate one. Current activities within the e-IRG forum are addressing this policy issue.

Working group participants

s.bassini@cineca.it	Sanzio Bassini	Italy	C
bilas@ics.forth.gr	Angelos Bilas	Greece	C
coghlan@cs.tcd.ie	Brian Coghlan	Ireland	C
francois.etienne@in2p3.fr	François Etienne	France	C
sergi.girona@bsc.es	Sergi Girona	Spain	C
jari.jarvinen@csc.fi	Jari Jarvinen	Finland	C
marcel.kunze@hik.fzk.de	Marcel Kunze	Germany	C
maglaris@mail.ntua.gr	Vasilis Maglaris	Greece	N
hugh.pilcher-clayton@epsrc.ac.uk	Hugh Pilcher-Clayton	United Kingdom	C
ullmann@dfn.de	Klaus Ullmann	Germany	N
vandenbroucke@helios.iihe.ac.be	Rosette Vandenbroucke	Belgium	N
vandrome@renater.fr	Dany Vandromme	France	N
vinter@imada.sdu.dk	Brian Vinter	Denmark	C
andyn@itn.liu.se	Anders Ynnerman	Sweden	C