BIG GRID

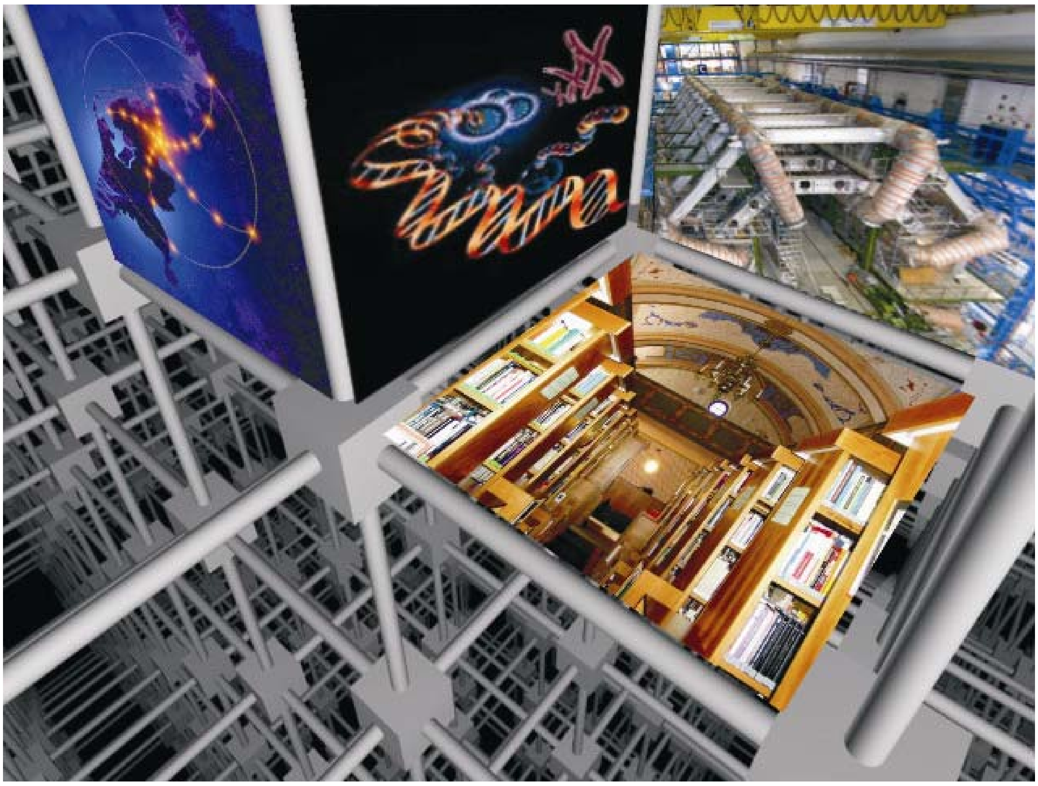

The Dutch e-Science Grid

Title	BIG GRID, the Dutch e-Science Grid
Date	13 October 2005
Proposed by	Stichting Nationale Computer Faciliteiten (corresponding author) *Laan van Nieuw Oost Indië 300, 2509 AN The Hague*
	Nederlands Bio-Informatica Centrum *Toernooiveld 1, 6525 ED Nijmegen*
	Nationaal Instituut voor Kernfysica en Hoge-Energie Fysica *Kruislaan 409, 1098 SJ Amsterdam*
Contact	Dr. P.H. Michielse, NCF, michielse@nwo.nl Laan van Nieuw Oost Indië 300, 2509 AN The Hague

In his 1999 book 'The Sun, the Genome, and the Internet', futurist Freeman Dyson wrote "The internet and the World Wide Web are permeating our society and changing the way we live... The networks of today are embryonic forms, destined to grow into mature structures whose shape and power we cannot yet imagine".
The infrastructure proposed here represents a giant leap in the maturation of the Dutch e-Science infrastructure.

Table of Contents

Executive Summary

Many scientific disciplines are presently undergoing technological revolutions that lead to a common challenge: managing a distributed data explosion. Detectors, medical imaging instruments, micro-arrays, and multi-sensor instruments are producing amounts of data that are rapidly exceeding the capacities of their current local data storage and computing environments. In many cases these 'exploding data' are distributed from the very start, being produced by different research groups or distributed sensor networks. Consider examples like genome and protein analysis data produced by many research labs in the world, biobanks containing patient data from a variety of hospitals, biodiversity data collected at the banks of the river Waal, historical archives and text corpora in many different places. Combining these datasets allows for completely new forms of research. Moreover, experiments generating petabytes of data per year, such as LOFAR in radio-astronomy and CERN in particle physics, need more data processing power than ever can be located in a single facility, with data utilized by researchers all over the world.

From an ICT perspective, these data have similar properties: all require reliable storage, comprehensive archiving, secure coupling and sharing. We propose to build and roll out a nation-wide grid-based e-Science infrastructure, BIG GRID, that strengthens the international position of the Netherlands in many scientific areas. BIG GRID encompasses data storage facilities and data processing services, enabled by grid services, for a requested budget of 30 M€ over a four-year period.

The science case for this proposal is the integral of many different science cases, reflecting the broad scientific community base. The realization of BIG GRID is crucial to the success and continuity of many Dutch research communities, covering important areas such as life sciences, astronomy, particle physics, meteorology, and climate research, water management, to name just a few. However, the very nature of the new infrastructure, a multidimensional collaboration enabler and accelerator, allows for direct participation of also social sciences, humanities, and even addresses communities in administrative domains, like digital academic repositories.

One basic ingredient for the proposed infrastructure is the network. The Netherlands are already in an excellent position, due to the world-class network services provided by SURFnet, the upgrade of which has been secured from the GigaPort-NG project. BIG GRID provides opportunities for enhanced international visibility. Dutch participation in international generic grid developments is already prominent (in flagship projects like EGEE and DEISA) and are on a national scale very well covered by the VL-e project. Coordinated by the Netherlands Genomics Initiative, NBIC is the key player for enabling informatics methodology for life sciences.

While the Netherlands is a leading player in the development of the grid, and has considerable expertise in bio-informatics, distributed sensors networks, and particle physics, the large-scale infrastructure to fully exploit this leading position is missing. The purpose of this proposal is to realize a science-wide national grid infrastructure. This puts the Netherlands at the forefront of grid developments, enabling many national ambitions. It enhances the excellent position of Dutch academic hospitals in patient data collections using the grid for biobanking. It enables major advances in drug discovery through combining data and through availability of massive compute resources for modelling. It allows industrial research labs, such as Philips, to both contribute to and profit from the available resources for engineering sciences. It positions LOFAR as the European centre for serving a variety of scientific communities using LOFAR data and the Netherlands as one of the Tier-1 sites for CERN's LHC experiments.

This proposal is a collaborative effort of NCF, NBIC and NIKHEF.

1. Introduction

Since the internet and the world-wide web have become a standard part of science, business, and society in general, the amount of information available in electronic form has steadily increased. We are currently experiencing a very steep increase in this amount, as data from a new generation of scientific instruments is appearing and being converted into even more information. These instruments include new particle accelerators, astronomical telescopes, 'omics' data, and other distributed sensors networks. The increase is furthermore due to a transition from traditional information archives, for example hospital (paper) files, to systems where all medical information for everyone will be kept electronically including successive photographs, X-ray pictures, and body scans. In addition, the publishing sector has recently started to digitally archive everything that is published. The Internet provides every PC with access to all public data, but in order to deal with, and effectively use, the very many and very large electronic collections, especially those including sensitive information, grid technologies are needed. This is a social revolution, a step towards information democracy that will significantly change the way we do science and will increase cooperation and decrease competition and duplication. The full consequence of this revolution is still hard to assess.

The information and data is made available in many different formats, units, access protocols, locations, and even different languages. Yet the user does not want to know about these differences when he accesses the information. The new research field of *semantic grids* addresses this problem. Personal privacy needs to be protected and some data may have strategic value and may be made available to only a selected group of users. Information may have commercial value and techniques need to be developed to profit from the usage of that data. Techniques to provide positive identification of grid users and their right to access certain pieces of information are part of security, another new field that grid developers have to address. Many of those and other necessary technologies are yet in their infancy and have to be developed further and scientific grids are the appropriate place to do so.

Several countries around the world are making large investments in developing and deploying grid infrastructures. China, the UK, and the US have each separately committed on the order of hundreds of millions of dollars to the development of grids. Most other developed countries are participating at more modest levels. Industrial interest is also apparent, IBM leading the pack with a billion dollar investment in this new technology. Investments in more traditional high-end computing infrastructures (supercomputers) are less visible but important to consider as well. For example, greater precision in meteorological observations allows in principle a correspondingly increased precision in results from climate studies. Computational sciences like climate modelling are best served by a tightly-coupled parallel architecture.

Figure 1 illustrates a grid infrastructure from the science community's point of view. The BIG GRID proposal concerns the bottom layer: the connections between the users, archives, and research facilities; the data archival facilities needed to house the data and information being produced; and the generic computational resources that the scientists need to perform their analyses and transformations.

The Netherlands was an early player in grid research – hosting the very first Global Grid Forum conference in 2001 – and a leader in some of the specific technologies needed to build a working grid. Dutch groups joined the very first international scientific grid projects and were able to educate top-level experts in grid technologies. The four-year project described in this proposal will build the Dutch e-Science infrastructure from its current status as an outstanding, internationally-known test bed into a world-class production infrastructure, thereby both maintaining our position at the forefront of grid development, and providing the resources needed for Dutch scientists to reap the full benefits of this new way of doing science.

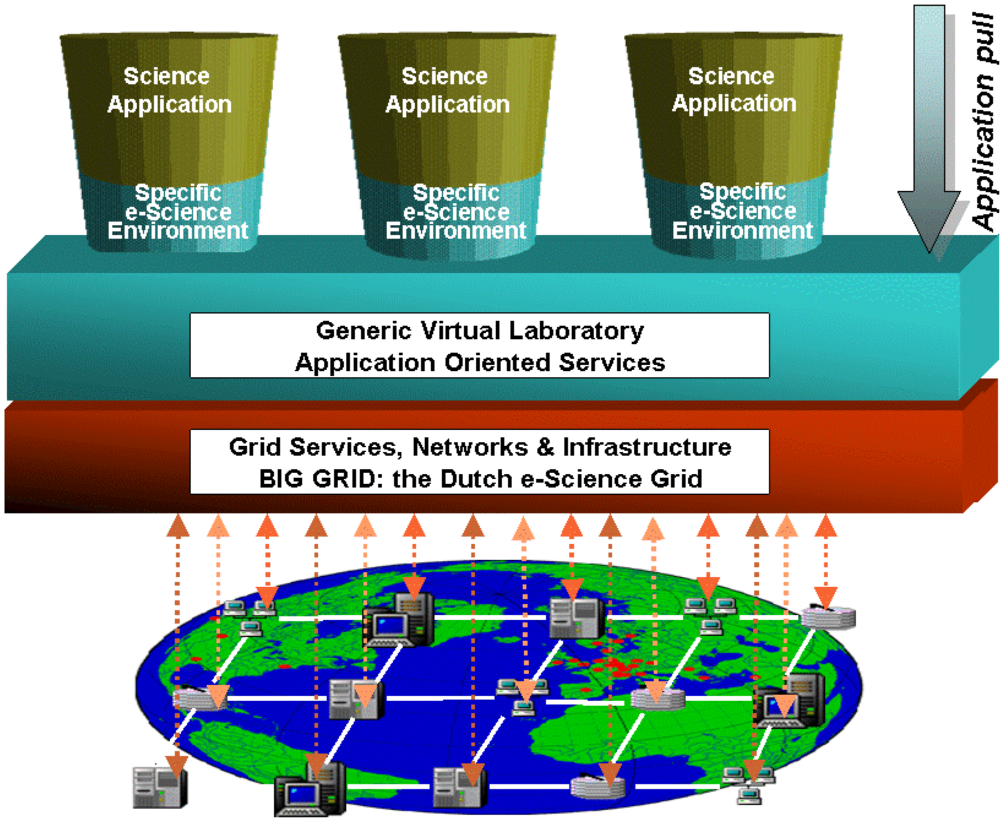

Figure 1: the Virtual Laboratory for e-Science approach to enhanced Science. The entire technology chain is studied, and generic solutions are developed and, as a result, re-use of components becomes possible. The BIG GRID will provide the compute power and storage capacity that will enable e-Science.

2. Problem Definition

In experimental sciences the tendency is to probe ever deeper, physics in matter, astronomy in the universe, medical research in life itself. In all research fields the experiments become more complex, and are driven by the increase in detector resolution, miniaturization, automation, and robotization. The data obtained per unit of time grows exponentially, while the intrinsic complexity of the data itself grows too. In many research fields even the most mundane computer tasks like, for example, making backups have recently already become cumbersome. This phenomenal increase in the amount of data has been called the *application data crisis* in data driven sciences. A problem complicating the data situation even further is the fact that in some cases the data are distributed over many physically distinct databases spread around the world. The situation is equally relevant for social sciences, while political organizations are dealing with the same problem in terms of digital information instead of experimental data. Computational Sciences deal with larger models and more simulation data. Some statistics that illustrate the crisis are presented in table 1.

Entity	**Approximate data volume**
The Bible	5 Megabytes (5 x 10^6 bytes)
Standard hospital X-ray image	5 Megabytes each
Medical imaging (fMRI)	1 Gigabyte per day (1 x 10^9 bytes)
Bioinformatics databases	500 Gigabytes each
Refereed journal articles	1 Terabyte per year (1 x 10^{12} bytes)
Satellite world imagery	5 Terabytes per year
Contents US Library of Congress	20 Terabytes
Internet archive 1996-2002	100 Terabytes
Current particle physics	1 petabyte per year (1 x 10^{15} bytes)
Particle physics in 2007 (LHC)	20 petabytes per year

Table 1: Examples of the information and application data crisis.

While this enormous increase (by more than a factor 1,000,000,000) in data and information presents practical and technical problems, the pill is sweetened by the potential for major advances in many scientific fields. Neutrino telescopes such as ANTARES or radio telescopes such as LOFAR now send massive amounts of individual signals to a remote powerful processing facility instead of following the traditional approach by using much slower local hardware to produce or process the images; the gain is the telescopes' ability to look in several directions at once, and to select a new direction in a fraction of a second (to catch a supernova in the act of exploding) instead of in hours.

The availability of large datasets from different fields also opens new possibilities when researchers can perform analyses on the combined data. A key activity in pharmaceutical research is the design of ligands, substances which bind specifically to other substances (those causing diseases for example). Several recent developments in bioinformatics have shown how ligand design can be significantly improved by using the results of studies of proteins present in a wide variety of species, especially when these species are properly categorized. This categorization information is found in biodiversity databases, which might seem at first glance to have essentially nothing to do with drug design. Presently, the combination of ligand design with biodiversity-based protein studies is performed by hand; time and money can be saved if rapid access to well-curated biodiversity data could be made available. The famous European Molecular Biology Laboratory (EMBL) spearheaded in the 90's the classical approach towards biological data integration which is to put all required databases and databanks on one computer. These approaches have worked fine in many disciplines for about a decade, but are now breaking down for two reasons. One, databases are growing too fast.

Two, the number of databases keeps growing, but the management time per local copy of a database (or databank) remains constant.

New possibilities due to this type of database federation are present for the alpha and gamma sciences as well, especially in combining the massive and now digitally-accessible datasets from various sources such as the Koninklijke Bibliotheek, the Centraal Bureau voor Statistiek, and the Belastingdienst. This example further emphasises the need for access control and privacy preservation already referred to for the medical and biological domain in the introduction. While these databases should be transparently accessible to authorized researchers, they must not be accessible to others, for a variety of legal and moral reasons.

Now that the Internet and the web are abundantly present in every day life, and networking and computing speed are improving at impressive rates (doubling performance in 9 and 18 months, respectively), federation and processing of these massive data sets is becoming possible via the developments that were sketched in the introduction.

An e-Science enterprise is driven by the researchers who use it and the data they generate. This enterprise requires a basic infrastructure in order to proceed:

- Data must be archived in such a way as to ensure its long-term preservation; convenient discovery (i.e. what data is present that would be of interest to the researcher?) and access is guaranteed to authorized researchers; leaks to unauthorized parties are prevented for sensitive data;
- It must be possible to communicate the data rapidly among the researchers working in the various collaborating universities and institutes;
- There must be sufficient computing power to enable analysis of these data.

In general, most applications, especially those driven by data, are perfectly well served by computing power in the form of clusters. However, one should not forget that there are very important fields of research (*e.g.* climate modelling) for which the traditional supercomputer architecture is absolutely necessary. Therefore it is also important to include supercomputer facilities in the national e-Science infrastructure.

As we will demonstrate below, much of the infrastructure needed by a wide variety of sciences is very similar. The most efficient solution to the national e-Science problem is to identify these common features and deploy a common e-Science infrastructure that takes care of these common problems. The engineering expertise, housing costs, and support infrastructure can be shared by all, and the common infrastructure will allow other, yet to be conceived synergies to arise between fields.

3. Science Case

In this section we present examples of the new scientific possibilities that are created by deploying the proposed national e-Science infrastructure. The communities represented in these examples are among those most advanced in pursuing the advantages of e-Science. An exhaustive list of science cases would be impossible to produce. For both academic and industrial research the infrastructure creates opportunities to position the Netherlands as a top class environment for distributed data handling and analysis. Moreover, all application areas served by the grid infrastructure, as outlined in the various science cases, have strong societal and economical relevance.

3.1 Biology

It is well known that the impact of new instrumentation (that could only be fully used thanks to an improved computational infrastructure) has been crucial for particle physics where it often led to Nobel Prizes. A similar technological revolution has recently taken place in biology where the study of single genes is being replaced by the mapping of complete genomes, and the study of single proteins by the study of interactions between all proteins within an organism. Figure 2 illustrates this so-called ‘omics’ influence in life sciences. The presence of the BIG GRID will enable major advances in biological research in several ways.

Figure 2: Essentially all aspects of health, from what we eat to the medicines we get are presently subject to intensified research. For example, all of the foods in the famous 'schijf van vijf' are presently researched with omics technology and subsequent computing.

First, it will enable the biological community to deal with its data explosion. As illustrated above, this community is now being faced with the need to make the transition from Gigabytes to Terabytes, while the rate of expansion of available data continues unabated. The BIG GRID infrastructure will supply these groups with local clusters and data caches, growing towards the aimed total national size of one petabyte in 2010. These caches will be remotely administered by experts at the large central facility, which will also have a petabyte-scale storage area for the biological sciences.

Second, the e-Science infrastructure enables a fundamental change in the research approach in the life sciences. Previous research was generally hypothesis driven: the researcher made a hypothesis and designed an experiment specifically to test this hypothesis. The data explosion allows analyses to be increasingly data driven, discovering surprising results by specialized searches through massive data sets. This is even further enhanced by the possibility to combine various datasets in order to increase the sensitivity of analyses, in much the same way that weather prediction is much more effective when input from several weather stations (instead of a single station) is included. Since the various datasets to be combined are located with various research groups around the country, their combination requires the infrastructure (network and storage resources) of the BIG GRID. This type of combined analyses also requires software similar to that being produced in related projects such as the BSIK VL-e project in order to proceed in an effective and transparent fashion.

The ultimate step in this process is to ‘virtualize’ the various heterogeneous data repositories produced by Dutch bioscientists. The result is that researchers can access all databases via the same protocols, and hence it will become easy to create data repositories and retrieve the relevant information whenever and wherever necessary. These developments will have a profound impact on life science experimentation and therefore also on healthcare. In addition it will revolutionize the way in which information is being used in the medical profession be it in medical hospitals or by other medical practitioners. Due to the medical, patient-related nature of the experiments, the data often are sensitive so that security must be guaranteed during the transfer and at the communicating sites. The life science measurement pipelines are presently subjected to international standardization to ensure that reproducible and comparable data and information sets are being produced that subsequently can be stored in (international) databases. The latter is important for the dissemination because publication in international journals increasingly demands the deposition of datasets such that other researchers can validate the quality of the experiment. This implies that the international position of this type of research becomes critically dependent on the quality of the e-Science infrastructure that is nationally available.

3.1.1 Drug design

G protein-coupled receptors (GPCRs) are the targets in the human body for 52% of today’s medicines. Almost all pharmaceutical companies market drugs that are GPCR agonists or antagonists aimed at diverse disease states such as hypertension, sleeplessness, inflammation, anxiety, or depression. The parallel introduction in the 1990’s of novel molecular biological, biotechnological, and computational techniques, has rationalized the process of drug discovery which previously was mainly guided by ‘trial and error’. The introduction of high-throughput and massively-parallel experimental techniques over the last five years additionally has caused the rapid population of data collection systems with experimental data. Databanks and databases provide access to all this data at a low level of sophistication, often without cross-references, visualisation, or inference facilities.

Pharmaceutical industries like Organon or Solvay are continuously working on improvements of their drug design pipeline. Cheminformatics and bioinformatics are playing an increasingly important role in the first five years of the drug-discovery timeline. Access to data of different origins has gained importance, and many databases have emerged that are queried routinely. The GPCR database (operated in part by Dutch researchers at the Centre for Molecular and Biomolecular Informatics, CMBI) aims at the integration of all existing GPCR-related data in a well-annotated information system that can be queried in an intelligent manner.

Available at http://www.gpcr.org/7tm/ the GPCR database system is the source of GPCR information for thousands of researchers around the world. The GPCR database is typically accessed nearly ten thousand times per day, and the list of users reads like the who-is-who of pharmaceutical research world-wide.
Further development of the GPCR database has slowed in recent years due to the continuous growth in data volume of the incorporated data types, and due to the continuous increase in the number of data types, databases, and databanks, that need to be incorporated. The infrastructure proposed here, and the collaborations that will be made possible by connections to similar centers in other countries, are needed to ensure continuation of the GPCR database’s position of world leadership in the drug-discovery world.

3.1.2 Biobanks: a key advance in transforming medicine

Biobanks supply the research community with biological samples and sample- or population-centric molecular, tissue, and clinical data. These samples, along with associated patient data, can be used to discover the genetic basis of diseases as well as determine the protein biomarkers that may predict disease onset, prognosis, or progression. These indicators of disease or biomarkers can then be developed into clinically administered diagnostic tools to better predict and monitor disease.

The ‘bank’ aspect is particularly important since the success of these analyses often relies on collecting and analyzing data from a large number of patients. While the importance of biobanks is widely recognized, their development is still faced with many ethical, legal, social, scientific, financial, intellectual-property, and information-technology challenges. To overcome these challenges and reach the promise of transforming healthcare, the public and private domains must engage in unprecedented levels of collaboration. The potential benefit to the national health and welfare more than justifies the effort needed to accomplish this.

Presently, most biobanks are patient oriented; i.e. all information about one patient can easily be extracted, but obtaining one type of information for all patients is difficult. Also, retrieving information for large patient groups may stress the biobank system to such extent that it may infer with daily healthcare. Therefore, duplication and reorganisation of hospital information systems into other more flexible (relational) databases appears to be necessary, and is enabled by the central storage facilities of BIG GRID. Patient-related data, clinical information and molecular data will be stored in different systems; using them in combined analyses is enabled by using the database virtualization techniques to be deployed on BIG GRID.

3.1.3 Electron tomography

An important goal of structural cell biology is to elucidate the mechanisms of the processes of life. One of the best techniques to understand these processes is to make a series of high-resolution three-dimensional images of functionally significant structures – for example organelles like the mitochondrion or Weibel Palade body – to understand how they change over time. Electron tomography is a powerful technique that provides such information about cells, organelles and macromolecular assemblies.

Electron tomography provides highly informative images of functionally significant cellular structures, at a resolution between forty and one hundred times better than that obtained with the most advanced optical microscopy. This three-dimensional, high-resolution, cellular-structure ‘movie’ comes with an associated data and computing cost; the images are constructed by combining a large number of two-dimensional ‘slices’ together to produce a three-dimensional result. The computing problem is trivially parallel and hence suited to the same cluster architecture needed for e.g. particle physics event processing.

On a national scale only two laboratories, at Utrecht University (UU) and at the Leiden University Medical Center (LUMC), were able to setup the required instrumental resources that are essential for the successful application of this technique. Both labs are embedded in national and international consortia aimed at maintaining and further developing high-resolution electron tomography. The consortia are the EU FP6 Network of Excellence on 3D electron microscopy[1], the Cyttron programme[2] to provide a window on the machinery of life, and the IOP genomics initiative[3]. These labs also collaborate with the VL-e project in anticipation of the e-Science challenges arising as a result of their research. The BIG GRID infrastructure will relieve these laboratories of the maintenance of an ever-increasing data archive; also it will allow other researchers across the Netherlands to profit from these facilities.

3.2 Ecology and bio-diversity

A variety of organisations have collected samples of species and made observations of their occurrence and abundance. These collections were made to study not only agriculture (pest organisms), water quality, bioprospecting, conservation, but also for research in systematics,

[1] See http://www.3dem-noe.org/
[2] See http://www.cyttron.org
[3] See http://www.senternovem.nl/iopgenomics/projects/ige03012_identificatie_van_de_oorzaken_van_aderve.asp

biogeography, population ecology, and other biodiversity related sciences. Recently these data are starting to appear electronically, scattered around various labs and institutes. At the same time, massive amounts of spatially explicit data on the collection sites (e.g. climate, soils, topography, features recorded by remote sensing) have become available. Currently, EcoGrid is being developed in the context of the VL-e project to be able to cope with this wealth of information; the goal is to enable data federation and collaboration within the ecological sciences to obtain new ecological insights. This is not only exciting but also socially very relevant, as our society becomes increasingly aware of the need to learn how to deal with the ecological challenges it is facing; among those relevant here are global change, emerging diseases, decreasing biodiversity, and waning resources. The project's ultimate goal is to accomplish a breakthrough in the effectiveness of decision and policy-making. EcoGrid has built the National Distributed Database of Flora & Fauna which contains all the data of private organisations that collect data, allied in the Association for Research of Flora & Fauna. It covers already 90% of all national spatio-temporal observations of species (10,000,000 records).
The hardware requirements for EcoGrid are not unique. Many other classes of the data nowadays contain spatial and/or temporal aspects, for example those from medical imaging, consumer behaviour, traffic, and security. Therefore EcoGrid is similar to the other data intensive sciences being deployed on the BIG GRID infrastructure, where the data is collected and stored in geographically distributed databases and is analyzed while using remote resources such as cluster computers or grid facilities. Analysis of distributed databases is enabled by the fast network connecting the infrastructure, and by the large data caches that will be available. The high-performance clusters will enable spatial data mining and inverse modelling tasks.

3.3 Food and Health

The technological advances in the fields of genomics, analytical chemistry and computer science have fundamentally changed the fields of food and nutritional research. Areas of research that are being reshaped by the omics techniques include:

- Functional genomics research is yielding new results on understanding the microbes involved in industrial food fermentations. In the last five years, more than 50 genomes of relevant microbes have been sequenced amounting to about 100 Megabytes of new gene data;
- A wide array of high-resolution and high-throughput genetics-research techniques is being used to study human intestinal microbiota. Given the fact that the intestinal microbiota consists of more than 1,000 species, their collective genome – termed microbiome – equals that of the human genome in size but has a much larger coding capacity;
- Studies of the impact of microorganisms on (intestinal) health, where data obtained with omics-based techniques reveal pathways involved in host microbe interactions, and will lead to biomarkers for gut health;
- Nutrigenomics, where the impact of food and food ingredients on health is studied by functional genomics technologies on both man and animal models. While presently 1,000 arrays are being analyzed per year, this amount is expected to grow exponentially as costs decrease.

In all these areas it is clear that rapidly increasing amounts of data are being generated. The estimates of the rate of increase range from a 10- to 100-fold increase within the next 5 years. Furthermore, combined (or fused) analyses, such as combining data from biobanks with those from clinical trials are increasingly common. In general, one can anticipate improved power in data analysis, and major scientific discoveries by combining many different data sets. Especially where links between microbes and man are being studied, different worlds meet, each with databases consisting of massive amounts of data that is generated at very different locations, and most likely with very different formats and attributes. There is strong need to ensure accessibility to these databases and organize their information in a structured way. Proper management of confidentiality of the data will be essential, both for intellectual-property protection and privacy reasons.

Within the Wageningen Center for Food Sciences (WCFS), data are being generated at five different locations, and there are many industrial partners with which data are being shared. Moreover, WCFS participates in many different consortia, among which the Centres of Excellence related to the National Genomics Initiative (Kluyver Centre for Genomics of Industrial Fermentations and the Innovative Cluster for Nutrigenomics). An infrastructure like BIG GRID that can support the storage, retrieval and integration of the different datasets at the different locations will be invaluable within the near future.

3.4 Astronomy

LOFAR started as a new and innovative effort to force a breakthrough in sensitivity for astronomical observations at radio-frequencies below 250 MHz. The basic technology of radio telescopes had not changed since the 1960's: large mechanical dish antennas collect signals before a receiver detects and analyses them. Half the cost of these telescopes lies in the steel and moving structure. A telescope one hundred times larger than existing instruments would therefore be unaffordable. New technology was required to make the next step in sensitivity needed to unravel the secrets of the early universe and the physical processes in the centres of active galactic nuclei.
LOFAR is the first telescope of this new sort, using an array of simple omni-directional antennas instead of mechanical signal processing with a dish antenna. The electronic signals from the antennas are digitised, transported to a central digital processor, and combined in software to emulate multiple conventional antennas, effectively creating a telescope looking simultaneously in eight different directions. The cost is dominated by the cost of electronics and will follow Moore's law, becoming cheaper with time and allowing increasingly larger telescopes to be built. So LOFAR is an IT-telescope. The antennas are simple enough but there are a lot of them – 25,000 in the full LOFAR design. To make radio pictures of the sky with adequate sharpness, these antennas are to be arranged in clusters that are spread out over an area of ultimately 350 km in diameter. (In phase 1 that is currently under construction 15,000 antennas and maximum baselines of 100 km will be employed). Data transport requirements are in the range of many Terabits/sec and the processing power needed is tens of Teraflops (1 Teraflop is 10^{15} floating-point operations per second).

Scientifically LOFAR will open up a new window on the Universe – exploring the low frequency (10-200 MHz) radio sky at high resolution for the very first time, focussing on these scientific goals:

- A very exciting application of LOFAR will be the study of the very early universe. It can be investigated whether – after the so-called Dark Ages when matter was neutral – the atoms were ionized again (‘re-ionisation’). The LOFAR frequency domain is ideally suited to this;
- By carrying out so-called large-sky surveys, LOFAR will make catalogues of (new) radio sources. These can be used to study exotic phenomena such as the formation of massive black holes, galaxies and clusters of galaxies. It is likely that this will lead to the discovery of new phenomena;
- LOFAR's large instantaneous beam will make it uniquely suited to observe objects that exhibit variable radio emission. Such objects include Gamma Ray Bursts (the largest explosions known to exist in the Universe), radio supernovae, intermediate black holes, flare stars and exo-planets;
- By observing the intense radio pulse that is produced when ultra-high energy cosmic rays hit the atmosphere, LOFAR has the possibility to study these phenomena in entirely novel approach. In astroparticle physics the origin of such extremely high-energy particles, and the corresponding accelerating mechanism represent central unresolved issues.

It should be realized that LOFAR can also be used as a more generic Wide Area Sensor Network. Sensors for geophysical research and studies in precision agriculture have been incorporated in LOFAR already. Several more applications are being considered, given the increasing interest in sensor networks that ‘bring the environment on-line’.

3.5 Particle Physics

Particle physics has made striking advances in the last fifty years in describing the intimate structure of matter and the forces that determine the architecture of the universe. Nevertheless, fundamental questions like: 'What is the origin of mass?' and 'What happened to anti-matter since the Big Bang?' remain. The answer to these, and other, questions are likely to come from the world's largest scientific instrument: the Large Hadron Collider (LHC), currently being built at CERN near Geneva. When the LHC begins operations in 2007, it will yield proton-proton collisions at an unprecedented centre-of-mass energy of 14,000 GeV thereby extending the discovery range for hitherto undiscovered particles (notably the elusive Higgs boson, assumed to be responsible for the masses of all elementary particles) and new phenomena by an order of magnitude. Nobody can predict the exact outcome of the LHC programme, but it is clear that the direction of future particle physics activities will be profoundly influenced by the results of the LHC.

In 1954, the Netherlands, together with 11 other European countries, founded CERN. Ever since Dutch researchers have been very visible at CERN as witnessed e.g. by the 1984 Nobel prize in Physics of Simon van der Meer for 'his decisive contributions to the large project which led to the discovery of the W and Z bosons' at CERN's SppS collider and the 1999 Nobel prize in Physics of Martinus Veltman and Gerardus 't Hooft for 'elucidating the quantum structure of electroweak interactions' confirmed to high accuracy by the experiments at CERN's LEP collider. At present, Dutch researchers play an important role in the design and construction of three of the four upcoming LHC experiments via NIKHEF (ALICE, ATLAS and LHCb). NIKHEF itself is a collaboration between FOM and four universities (RU, UU, UvA and VU) and coordinates all Dutch particle physics activities. The coming decade the Dutch LHC related activities will constitute about 2/3 of NIKHEF's research programme.

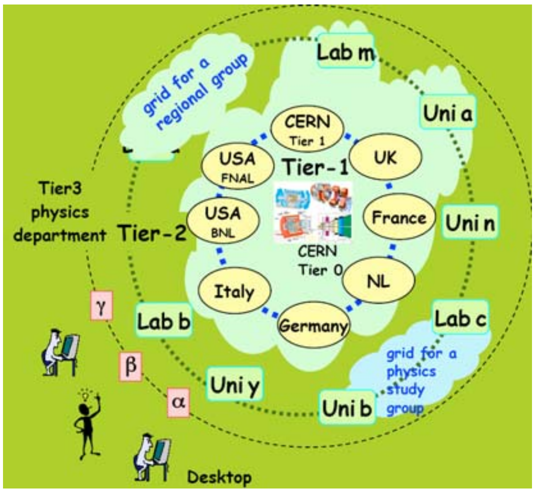

In operation, the LHC data volume will be astounding: 15 million Gigabytes (15 petabytes) annually. To handle this huge data volume and to allow the thousands of scientists scattered around the world to analyze the data, CERN envisages a multi-Tier hierarchical computing grid model: a Tier-0 centre at CERN connected to about ten Tier-1 centra scattered all over the world (see figure).

For several years, the Netherlands has been identified as a candidate LHC Tier-1 site. The Netherlands was on board very early in the LHC distributed-computing research effort, NIKHEF and FOM being one of the five major partners in the EU 5th framework 'EU DataGrid' project, and NIKHEF was one of the five 'core sites' of the European Data Grid testbed, which first functioned as a prototype of and later formed the nucleus for the current LHC Computing Grid. NIKHEF, together with SARA and the UvA, continued in the EU 6th framework 'Enabling Grids for E-science in Europe' (EGEE) project and by virtue of the excellent team built during these projects has gained international leadership in grid computing.

The presence of the BIG GRID infrastructure, supporting the LHC Tier-1, will ensure continued international prominence for the Dutch high-energy physics community during the LHC era. This leadership position has already started to show spin-offs in related areas, for example the Dutch leadership in the IT developments for the new field of neutrino telescopy in high-energy astrophysics.

3.6 Medical instrumentation design

New medical imaging techniques have enabled a revolution in clinical and preclinical research, and these techniques are rapidly finding their way into hospitals and research laboratories. The example described here, SPECT or Single-Photon Emission Computed Tomography, is certainly not the only beneficiary of this revolution.

SPECT has become an increasingly important preclinical imaging technique, due primarily to the recent development of new SPECT-labelled molecular imaging compounds. SPECT is a technique similar to PET (Positron Emission Tomography), the difference being that the radioactive substances used in SPECT have longer decay times than those used in PET, and that they emit single photons instead of two photons. SPECT can provide information about blood flow and the distribution of the radioactive tracer substances in the body. Despite the lower sensitivity and less detailed images, the SPECT technique has some key advantages: it is less expensive than PET, and SPECT centres are more suitable for a clinical setting as they do not have to be located near a particle accelerator.

To fully exploit the capabilities of SPECT instruments, the detector characteristics must be known to a high degree of precision, and the complex interactions of the photons on their way from the anatomic location of the tracer down to the detector must be well understood. Using complex simulation techniques (some of which were developed for particle physics) the resolution and discriminating power of the technique can be improved. Interactions of incident photons with the collimator and detector, including septal penetration, scatter and X-ray fluorescence, are significant sources of image degradation in applications of SPECT including dual isotope imaging and imaging using radioisotopes that emit high- or medium-energy photons.

Modelling these interactions using full Monte Carlo simulations is computationally very demanding. Even a single simulation of a full SPECT detector today would take a standard PC over 11,000 days (or over 30 years) to complete. The computational cluster component of the BIG GRID will enable significant reduction in the turnaround time for these instrument-design studies, and result in both more rapidly deployment of better instruments in clinical situations, and an advantage to Dutch industries performing these studies.

3.7 Cultural and Linguistic Studies

Determining how to provide researchers throughout the Netherlands with properly controlled, fast, and transparent access to scientific archives, and providing the basic infrastructure to contain the data, is part of the mandate of the DANS (Data Archiving and Networked Services), an initiative of the KNAW and NWO. The problems that archives generally are facing in the Social Sciences and the Humanities are challenging. Libraries and other data archives are increasingly turning to digital storage instead of paper or film; an example is the 'Het Geheugen van Nederland' archive operated by the Koninklijke Bibliotheek. As such archives increase in size and diversity of content, new and sometimes multi-disciplinary research becomes possible

Large databases with surveys on different subjects, collected in European countries over several years are hard to query, due to the volume of data and the lack of harmonization. In archaeology the digital data of excavations can vary from photos to geographic information sources, and from field reports to Computer Aided Design (CAD) drawings. Given the impact of the Malta treaty, that obliges construction projects to carry out archaeological research, more and more data will be available in the near future. Once combined and accessible, it could help to predict 'new' archaeological locations.

Linguists are already experimenting with analyzing large bodies of newspaper articles going back to early in the previous century. Such studies provide crucial insight into the development of the language over time, and one can imagine how much easier it is to process this information from a digital archive rather than a stack of newspapers. Other archives contain audio and visual data in addition to text. Given the Gigabyte typical size of a CD or DVD, one can quickly imagine how such archives can approach the petabyte sizes of particle physics as more and more historical information is digitized and added, and how large amounts of computing power could be used in analyzing trends in these registrations.

The BIG GRID infrastructure will provide the basic archival services needed to accomplish the goals of the DANS initiative.

3.8 Climate Research

'In 2007, a particularly severe storm causes the ocean to break through levees in the Netherlands making a few key coastal cities, such as The Hague, unliveable'. This scenario was described in an internal Pentagon report in 2003. Now, in reality, New Orleans has been flooded in September 2005. Hurricane Katrina has demonstrated that extreme climate events can affect the lives of thousands or even millions of people. In addition to the local destruction of homes and infrastructure, there was an immediate effect on oil prices, creating a worldwide impact.

The Dutch Centre of Climate Research (CKO), which involves researchers of the Royal Netherlands Meteorological Institute (KNMI) and of the Institute for Marine and Atmospheric Research Utrecht (IMAU), is very active in modelling the effects of global warming in the Western Europe. Changes and shifts of the climate system over the Atlantic sector will have a direct impact on the frequency and strength of storms in Western Europe, associated with storm surges and growing waves, threatening the Dutch coastal defences. Heavy rainfall events will pose a threat through the 'back-door', via the rivers. Major infrastructural measures to avert these threats have a lead time of several decades. For instance, after the February storm surge in 1953, with major flooding of the south-western part of the Netherlands, it took over 30 years to complete the well-known Delta Works (now used as an example of good practice).

Clearly, for a low-lying country like the Netherlands it is essential to have reliable data and forecasts of climate change and the frequency and intensity of extreme events over the next 10-20 years. This was illustrated in our Queen's 2005 'Troonrede' (a speech at the opening of the governmental year), saying: 'We are seriously confronted with the consequences of climate change[4]'. To produce such forecasts will require the best models of the climate system, an enormous network of observations of the state of the climate system (ocean, atmosphere, ice, land) and techniques to use past and present observations for prediction of future states. It also requires top-of-the-line computing facilities to perform the (ensemble) simulations with these climate models, for analyzing the data and for the production of long-term (10-20 year) forecasts.

Recent numerical simulations on the Dutch National Supercomputer (dedicated over a period of a few months) have given evidence that the probability of hot summers in Europe will increase dramatically under increasing greenhouse gases, and in further complicated by other indications that the temperature may actually decrease due to changes in the Gulf Stream resulting from polar ice fields reduction. For more accurate results, the resolution of the models will have to be increased several orders of magnitude, with associated smaller time steps, leading to huge computational requirements. The supercomputer facilities accessible via the BIG GRID infrastructure will enable the necessary computational advances, and provide transparent access to the observational data archives. Only in this way, the Netherlands will be able to take the right steps to protect its country and its people against the changing climate conditions.

[4] 'We worden indringend geconfronteerd met de gevolgen van klimaatveranderingen.'

4. Science Requirements Analysis

4.1 Life sciences

Data and information handling and processing are the most prominent e-Science activities for a wide range of life sciences-related fields. Experimental data come from many types of biological and medical detectors like sequencers for genomics, DNA micro-arrays for transcriptomics, mass spectroscopy for proteomics, and fMRI for brain imaging research. Another large volume of data used for research resides in databases like biobanks.

As explained earlier, the breakthrough potential mostly originates from the ability to combine these massive data sets, irrespective of their source. This could be either via on-the-fly fusion of data from different databases, or via the full integration of large parts of these databases in a centralized system. From an ICT point of view these data processing activities can be classified as database intensive. Most of these databases reside at the users' premises, however, ideas exist to let the data-fusion experiments produce new databases that can be shared by different users and could be localized centrally at a dedicated computer centre.

Transport of data must be secure. In certain cases (e.g. patient information) this is even demanded by legislation. Secure transport is possible either via encryption or via secure network connections that can be obtained through either (or both) Virtual Private Networks or Optical Private Networks. The emerging technologies for on-the-fly encryption and optically-transmitted data are a very promising development. The experiments on integration of biobank collections are just now beginning; hence the amount of data to be transported per action is limited to a few hundred Gigabytes. However, when biobanks will be combined with biological/biomedical experiments like genomics, proteomics, or fMRI, it can be expected that the amount of data to be transported will grow rapidly.

To deal with these data, large amounts of storage need to be present both at the experimental facilities – to hold their own experimental data, and for temporary buffering of combined databases during analysis – as well as at a larger data centre, in order to provide backup copies of these valuable data sets and more persistent storage for fused data sets being often requested by researchers.

4.1.1 Drug discovery data collection and handling

The GPCR data base project illustrates that data collection, when combined with validation, curation, annotation, and integration, allows scientists to think differently about their molecules. The sheer availability of all data regarding one molecular family has shown repeatedly to be an important inference engine for discoveries in many different directions. The GPCR database has been a driving force for many discoveries, both scientific and technical in nature. However, this progress has slowed in recent years due to a series of problems:

- There must be sufficient computing power to enable analysis of these data;
- The number of databases and databanks that needs to be inspected in order to get the GPCR data collection as complete as possible is slowly but steadily increasing while the amount of curator time per data collection remains constant;
- Access to remote data is not always equally easy. Access can be slow, restricted, or only intermittently possible;
- Most providers of data are database amateurs that store data in idiosyncratic ways. The computer-literate partners in the GPCR database project are willing to help in such cases, but access to remote computers is often hindered by technical or system manager generated problems.

4.1.2 Electron tomography

An electron tomography system can typically acquire about two hundred 2D projection images per hour. These images have dimensions of up to 4096x4096 pixels in size, and each pixel comes with a grayscale value (2 bytes). The resulting storage size for each 2D image is about 32 MB; the resulting 'tilt' series of slices that generate a 3D image measures about 6.5 GB. A collection rate of about 10 series per day yields 65 GB of data per day.

To generate a (3D) tomogram from a tilt series, two steps need to be taken. First, the 2D images need to be aligned to each other. Depending on the precise approach taken here, this alignment may be highly computing intensive. After alignment, the 3D images are generated. Again this is often a computationally-intensive task depending on the exact algorithm used. Each tomogram measures about 16 Gigabytes, giving a data production rate of 160 Gigabytes per day. A 32-CPU cluster can process a single dataset within a day.

To localize macromolecular structures within a tomogram, a specific implementation (Omnimatch) of the 3D cross correlation approach is used. Omnimatch is computing intensive and typically requires access to a 32-CPU cluster to process relatively small data sets within a day.

Databases are essential for recording a precise history and characterization of the samples, e.g. to record the exact procedure by which the sample was prepared. To be able to make use of the results obtained with electron tomography the results also have to be stored in a database that can be correlated with content in other databases for data-fusion based analysis. Within the field of electron tomography several activities are ongoing to define what metadata is indispensable for later queries. Prototype environments with these metadata capabilities are being developed in collaboration with Dutch researchers and have been deployed at SARA.

Electron microscopy, including electron tomography, requires a considerable expertise ranging from setting of the biological experiment and specimen preparation to the use of sophisticated computing tools. Therefore, facilities to sustain knowledge over a longer period of time are indispensable. In this context, access to database and visualization tools to access data for both educational and scientific purposes is important. Considering the relatively large data sets (>160 GB) access to high-performance networking as well as to better visualization resources (large screens, pre-processing data before displaying a few over the net on a display) are essential.

4.2 Astronomy

For the upcoming generation of data-intensive astronomical experiments, it is increasingly important that the data stream out of the central processor can be archived – at least for a period of about a month, preferably longer if possible – and further processed and analyzed in an efficient and transparent way. LOFAR has been planning to use grid resources for this phase of the computing task since an early stage of its development. Collaborations with SARA and the VL-e project have been set up to ensure requirements are sufficiently clear, and that interfaces are properly defined.

In contrast to particle physics, where the basic processing unit is a single event, the access profile for astronomical calibration and imaging requires processing of data in four-hour blocks. A typical LOFAR calibration application requires fitting a model containing about 100,000 parameter values to a 25-Terabyte dataset, performing several iterations. Image creation involves two-dimensional fast Fourier transforms on slices of the complete data cube. Both applications require substantial processing power, on the order of 2-5 Teraflops. This corresponds to 200 – 500 cluster nodes. These steps are carried out to a defined level on the Central Processor. There resulting intermediate data products are produced at 10 Terabytes per day and will be further processed by end-users in different ways, depending on the scientific goals.

In the LOFAR operational model this processing will take place at distributed science centres (comparable to LHC Tier 1 sites, with the LOFAR Central Processor forming the Tier 0). Given the varying requirements of the science applications, the astronomical community using LOFAR will need a large e-Science infrastructure, with the main constituting elements being:

- Broadband data links (10 to 40 Gigabytes per second) between the Central Processor and the Science Centers; these connections form part of the SurfNet6 blue-print being realized in the Gigaport Next Generation programme;
- An approximately 2 petabyte storage pool for the *Epoch of Re-ionisation* (EoR) application; given the access profile for calibration and imaging this storage is preferably located at or close to the EoR Science Centre;
- A flexible, order petabyte storage pool for (partly temporary) storage for the survey and transient imaging applications;
- Sufficient distributed computational power for deep calibration, imaging and analysis. The EoR application requires a 256-node computing cluster (using a 64-bit dual-core, dual-processor Xeon machine as a reference node), to be used in combination with a comparable amount of processing power in the Central Processor.

The other applications require comparable amounts of computational power. Transient and Cosmic Ray applications will have varying computational demands, which is one of the drivers for using a computational grid approach for them. These demands could be well-served initially with a 256 node cluster, to be extended to a 512 node cluster in 2008.

4.3 Particle physics

The particle physics community made a tentative choice to exploit grid technology to enable the LHC computing effort almost five years ago. Since then several projects around the world have been steadily advancing the state of the art in grid computing, as well as testing the LHC computing models in ever-greater scale and detail. For the last year these tests have been called 'Service Challenges' as they are to be regarded as progressive steps in testing that the LHC computing grid service can actually perform the task. Scientists constructing the LHC experiment software – a massive undertaking all by itself – have been participating in these challenges. These scientists have analyzed the outcome of the service challenges and hereby provided the particle physics community with a reasonably sound idea of the computing resources required in the medium-range future. These computing resource model analyses have been performed in great detail[5]. We make no attempt to explain these findings here, but instead present directly the results. The contribution of the BIG GRID to LHC computing will be a Tier-1 centre for the three experiments in which the Netherlands takes part: ALICE, ATLAS, and LHCb. Each of these experiments has expressed their total requirements for Tier-1 computing, which must be shared amongst the various contributing Tier-1s. The number of Tier-1s per experiment is about ten, setting the relevant scale of contribution at around 10 percent per centre.

The general characteristics of this facility are the following:

- High-speed (10 Gbit/s or more) connections to CERN and a few other centers in the world. This is necessary to transport the experimental data from Geneva to Amsterdam, to accept Monte-Carlo simulation results from Tier-2 centers, and to make all these data and analysis results available to researchers around the world;
- A large-scale data archive component, with disk and tape capacities of about one petabyte in 2007 (when the LHC turns on) ramping up to about six petabytes by 2010;
- A large high-performance computing cluster. Due to the loosely-coupled nature of the particle physics problem, high-end commodity PC rack systems are perfectly suitable to the task. This is equivalent to about 500 current (3.6 GHz dual Xeon) machines in 2007 increasing to 3500 current dual-CPU machines in 2010. The actual numbers of machines (and the corresponding price to performance ratio) will be lower given the steady increase of computer power per machine and per euro over time.

[5] See http://lcg.web.cern.ch/LCG/peb/LHCC/expt_reqts/

The following table gives the specific requirements for the BIG GRID Tier-1 components from 2006 until 2010, taking into account the considerations above. The contributions represent an average of 11% of the worldwide computing resources for the three experiments combined. A portion of the resources needed in 2006 will have already been funded from both the NIKHEF computing budget as well as money from the NCF proposal 'Framework voor versterking van de Nederlandse grid en e-Science infrastructuur'; this is taken into account in the financial case presented later in this proposal.

	2006	2007	2008	2009	2010
NL Tier-1 CPU (kSI2k†)	750	2150	5500	9500	12300
NL Tier-1 Disk (TB)	350	1100	2550	3800	6150
NL Tier-1 MSS (TB)	250	725	1825	3550	5750
† SpecInt2000, a standard unit for processor performance, see footnote 6.					

4.4 Medical instrumentation design

A modern *Single Photon Emission Computed Tomography* (SPECT) detector consists of high-definition solid-state detectors with a direction-defining collimator placed in front. To simulate such a detector setup and get accurate evaluation of the detector characteristics, about 10 million decays of the 'fake' tracer element, as contained in the patient's body, must be simulated. For each decay, a single simulated photon is generated, its path through the patient's body is simulated and tracked by the computer, and the photon ultimately winds up in a simulated SPECT detector.

One of the limiting factors in SPECT is that in real life, not only the desired photons from the tracer element, but also 'random' photons from naturally-occurring background radiation will be registered in the detector, as well as inherent detector noise. Actually, there will be about a factor 1,000 more of these 'noise' events than real events, and this 'noise' must be included in the simulation to conduct a realistic study of the detector's performance. So a single simulation requires 10^{12} events to be generated, tracked through the detector, and ultimately processed like the real detector data. A typical personal computer today can evaluate approximately 1,000 photon tracks per second, and thus a full simulation on the detector requires eleven thousand CPU days. Today, even a single simulation is stretching the limits of the available compute facilities, and it takes months to complete such a single run. Using 'in silico' evaluation of many instrument designs before construction seems utterly impossible.

Fortunately, like in the particle physics case, the simulated events are entirely independent. Only in the final analysis stage need the results to be combined to evaluate total instrument performance. This allows for 'conveniently parallel' simulation model, where a large cluster of loosely coupled commodity systems can complete the job in a reasonable amount of time. The computing needs stated above translate to 260,000 kSI2k-hours per full simulation (the SpecInt2000, SI2k, being the standard unit for processor integer performance[6]). With a cluster capacity of 2,000 kSI2k, a single run can be accomplished in five days; this capacity can be used immediately and is requested for 2006. A second cluster of the same size is requested for 2008, anticipating the use of this technique in other industrial-research areas.

4.5 Computational sciences

Computational Science can be viewed as the area of science which numerically models a wide range of phenomena. Typical examples are disciplines as theoretical chemistry and catalysis research, in which chemical reaction mechanisms are modelled over very small time scales, and fluid mechanics, in which the Navier-Stokes equations for numerous types of flows are modelled and approximately

[6] Standard Performance Evaluation Corporation, see http://www.spec.org/.

solved. Climate research typically combines atmosphere models (which are modelled by Navier-Stokes) with oceanography models, including heat exchange and turbulence models. The nature of the underlying phenomena (simulation of continuums) requires tightly-connected parallel systems, leading to the need of parallel, low-latency high performance systems.

As an example, many problems (but certainly not all) in computational science are numerically modelled by discretisation of the underlying differential equations. For a 3D problem, this leads to a 3D grid, in many cases time-dependent. The accuracy of the solution is determined by the resolution of the grid: the finer, the better. For numerical stability reasons, the time step is coupled to the dimensions of the grid: two times more grid points in each direction requires the time step to be reduced by a factor of two as well. So the problem size grows with a factor of 8, and the time step reduces with a factor of 2. Computational requirements overall grow with a factor of 16. Simulation output will grow with the same rate. And still, increasing the resolution by a factor of 2 in general is by far not enough. In fluid mechanics, typically the smallest turbulent eddies are extremely small compared to the grid sizes. And hence will not be modelled at all, except in very high resolution models. It is the effect of missing a small storm at the North Sea in meteorology. To avoid missing essential phenomena in simulations, a huge increase in computational power is absolutely needed.

4.6 Common requirements

The requirements cases presented above share several common themes, and this is the reason why a shared national e-Science infrastructure can serve them all.

- Data archival services, in *large* quantities are needed by all the disciplines, in order to deal with the information explosion. Most of the disciplines are well served with a model in which most of the storage is located centrally, but a significant (comparable) amount of storage is distributed amongst the science centres. This distributed storage is used both to house locally collected data as well as to provide local caching of frequently accessed datasets;
- A research network capable of dealing with traffic in these enormous datasets is required by all the disciplines. While we are convinced that the results of SURFnet and GigaPort-NG will be more than adequate for this part, it is so important that we mention it again here;
- Computing power is needed by all the disciplines to varying degrees. The performance-to-price ratio for PC-style clusters is much better than that of traditional supercomputer architectures; most of the groups represented here are served equally well by either type, hence the most benefit is gained by deploying a larger fraction of compute power as clusters;
- Support for the facilities is needed; in addition to the operations budget for the central facility, remote support is needed for the distributed component. The scientists are not computer operations experts, but the requirements above repeatedly ask for 'well-curated' data products; this requires operational expertise.

This generic structure has been 'discovered' by many in recent years, even by Microsoft[7]. In this technical report, they warn that the data explosion will be followed by a computational explosion, since a) more sophisticated analyses will be used, which require more computing power per byte of data, and b) many analysis algorithms are supra-linear, scaling as a power of the data size. In anticipation of this effect we have included a significant 'shared' computing capacity in the infrastructure to serve those sciences that have yet to experience the full impact of this effect.

[7] "Scientific Data Management in the Coming Decade", Jim Gray *et al.*, Microsoft TR-2005-10, January 2005.

5. Talent Case

The history of computing and data processing facilities in the Netherlands has turned out to be a key asset for planning research. It has enabled research groups to seek for both Dutch and foreign students and researchers with confidence. One of the assets they can offer is the transparent availability of state-of-the-art compute resources, with the administrative flexibility to make such resources available quickly. This has attracted both PhD students, so that new lines of research could be developed, as well as renowned visiting researchers, to exchange knowledge on a top level. The availability of these computing resources has helped Dutch research schools and institutes to attract many PhD students. Examples are the JM Burgers centre for fluid mechanics, with a current total of more than 200 PhD students, and the Debye Institute for chemistry and physics, with more than 100 PhD students. Also other institutes which rely on large compute resources have benefited from these stable and advanced facilities, for which continuity over the years was no question at all.

This situation can only be maintained when leading facilities can be provided in the future as well. All scientific areas described in this proposal, ranging from life sciences to astronomy, represent large communities of (junior and senior) researchers. A top class grid infrastructure will allow them to avail 'at the desktop' of a rich environment of data processing and storage facilities, brought to them by the already outstanding Dutch academic network provided by SURFnet. Such an excellent research infrastructure attracts the most talented researchers. For the data intensive scientist an infrastructure is a 'sine qua non', *i.e.*, he will not be able to do science unless such an infrastructure is available. Much effort must therefore be devoted to avoid any unnecessary threshold for accessing the infrastructure. The proposed grid expertise centre and help desk will fulfil an essential role here.

Not only using but certainly *building* the infrastructure will attract talented researchers from both the ICT and application areas. This currently involves NIKHEF, SARA and VL-e that play a role on the forefront of grid developments, both on an application level as well as on an implementation level. This has led to a pool of expertise, which has also been recognized in European collaborations, like LCG, EGEE and DEISA. In addition the ASCI computer science school has more than 300 PhD students that are trained among others via the Distributed ASCI Supercomputer of which the third version has just received its funding. Various grid based projects take place around this infrastructure that encompasses the two universities in Amsterdam, and the universities in Delft, Utrecht and Leiden.

The interest of young people for these grid developments is furthermore illustrated by, *e.g.*, the attendance at Master classes organized by the 'Vereniging Gridforum' and the success of the grid master course of the University of Amsterdam.

The next step is the linkage between this pool of expertise and the broad field of data-intensive sciences. Knowledge on the frontier of grid technology and the relevant scientific domain needs to be developed. Also for this challenge, the presence of an inspiring facility environment, and the philosophy to develop BIG GRID, are of great importance.

6. Innovation Case

A grid infrastructure enables academic and industrial research communities to communicate, collaborate and otherwise profit from the installed computing and storage resources. The proposed infrastructure enhances the excellent position of the Netherlands in networking by adding these essential resources. It thus strengthens the possibilities to implement the 'open innovation model' as advertised by Philips and brought into practice at their campus in Eindhoven. Here all type of technology companies are invited to one location for cross fertilization purposes. An e-Science based grid infrastructure as proposed here can further strengthen this model by letting people use the infrastructure to collaborate and share information, not necessarily constrained by the same physical location.

Both for academic and industrial research the infrastructure creates opportunities to position the Netherlands as a top class environment for e-Science varying from physics to astronomy to life sciences and engineering. Moreover, all application areas served by the grid infrastructure, as outlined in the various science cases, have strong societal and economical relevance and, in particular in case of the life sciences, direct impact on the society and the citizens. The following examples illustrate these opportunities.

- The pharmaceutical industry is continuously working on improvements of the *drug design pipeline*. Cheminformatics and bioinformatics play an increasingly important role in the first five years of the drug-discovery timeline. Access to data of different origins has gained importance, and many databases have emerged that are queried routinely. A grid infrastructure as described will greatly enhance the possibilities for combining and analyzing data, which will shorten the drug-discovery timeline;
- With the availability of a large grid infrastructure (and taking into accounts Moore's Law on increasing performance), in-silico design such as described in the '*medical instrumentation*' science case can become commonplace, completed in hours instead of months. This results in dramatically shorter turn-around time on instrument design, and thus a shorter time-to-market for improved medical equipment;
- The integration of patient-related, clinical and bio-molecular data for a specified collection of samples (*biobanking*) has huge economic value. Such databases would open the road for many new research programmes;
- Consider a system for planning by-pass operations. The patient is put under an MRI scanner to obtain the topology of the blood vessels where the operation has to take place. This information is sent to a grid infrastructure where a Lattice Boltzmann simulation of the blood through the vessels is realized, after which the medical practitioner can via a visualization station try out the best location for the by-pass operation on the computer.

Not only *using* the grid infrastructure will be innovative, *building* it is also a formidable challenge, which opens up many opportunities for innovative activities. The Netherlands are already very well positioned, through its participation in EGEE in fields such as grid security, grid authentication, authorisation, and accounting as well as in grid management. Through the national scale of the proposed science grid this will be the most encompassing grid yet known. This will undoubtedly stimulate regional and national economic activity. A parallel can be seen with the development of the Internet and the World Wide Web, which – in the Netherlands – started largely in and around scientific institutions (with a focus in the Science Park Amsterdam). This has lead to the emergence of a great variety of ICT related business activities in the Netherlands and the Amsterdam region in particular.

The current interest of industry in the Netherlands in grid developments in general is demonstrated by their participation in projects such as LOFAR and VL-e, and in the 'GridForum Nederland'. The industrial partners and members of these consortia represent not only users of grid technology but also prospective suppliers of this technology and value adding services. As was the case with the ISP industry (Internet Service Providers) in the slipstream of the development of the Internet it is very likely that a grid service provisioning industry will emerge and mature in the coming decade, which offers excellent opportunities for new economic activities.

7. Partnership Case

7.1 Dutch context

The science cases presented above illustrate the broad range of scientific activities in the Netherlands that will profit from the infrastructure proposed here. The extent of this range is demonstrated by the number and variety of organisations that have provided letters of support (listed in appendix A). In many cases their activities are only possible given the presence of the BIG GRID infrastructure. Indeed many of these activities have been approved (by e.g. BSIK grants) under the assumption that such an infrastructure will eventually be present. The current proposal is *explicitly* collaborative in nature; one of the prime reasons for constructing a grid infrastructure is to enable such collaborations. At the lowest levels of data archives, CPUs, high-speed networks, and grid techniques of task and data communications, the infrastructure is generic to a large degree and hence applicable to a wide range of activities independent of the actual scientific discipline involved. This is the reason for the convergence in requirements from the parties represented here, from cultural and linguistic studies to particle physics.

A prototype of the grid infrastructure is already present with a nucleus in Amsterdam at NIKHEF and SARA. Groningen has recently declared a distinct interest and is a partner (with NIKHEF, SARA, and the UvA) in the follow up to the EU 6th framework EGEE project; part of their participation in EGEE-II will be to begin building a prototype grid centre in Groningen similar to the one that now exists in Amsterdam, and to acquire the necessary expertise to expand this in the future. There are other grid ‘entry points’ at several locations in the Netherlands, mostly for partners in the bioinformatics community but also for KNMI and Philips Research.

7.2 International (European) context

Many other countries are proposing to build similar infrastructures, using exactly the same model: leveraging the lead of the particle physics community, and their drive to construct a working, large-scale system in short order, into a national facility that on the one hand makes important contributions to the international particle physics effort and on the other hand serves as a nucleus for the national e-Science infrastructure across all sciences. The UK e-Science programme is perhaps the most ambitious of these; other examples are GridKa (Karlsruhe, Germany) and Port d’Informació Científica (Barcelona, Spain). Particle physics is not the only source of such international convergences; the drug-design distributed database GPCR database as described here arose as a large-scale international collaboration from an EU-sponsored project that was headed by the EMBL in Heidelberg. Connections – both in the high-speed light path sense, but also people connections and grid-service connections – with these centres are already in place, enabling international collaboration for researchers in all fields. The Netherlands recently made a significant contribution to a worldwide malaria drug search (via computer simulation) in exactly this way: the international grid interoperation put in place for the LHC Computing Grid was used, with *no modifications,* by an international team of bioinformaticians to carry out the study. Enabling the collaboration consisted simply of adding the team members as authorized grid users, installing the relevant bioinformatics software, and requesting various institutes to support the search by supplying compute cycles. The Dutch compute cycles were supplied from the allocation of the VL-e project.

The security models being developed from grid-computing initiatives will be very important for communities dealing with sensitive data, such as biobanks. Not only is protection of patient data for privacy reasons quite important, but the international exchange of data will have to conform to legal policies in all participating countries.

The BIG GRID will enable Dutch scientists to compete effectively in the international arena and place the Netherlands in the top position to achieve the next wave of scientific breakthroughs.

7.3 The ESFRI list of opportunities

The European Strategy Forum on Research Infrastructures (ESFRI) has in its report to the European Commission (March 2005) drawn a ‘List of Opportunities’, containing concrete examples of new, large-scale Research Infrastructures which the scientific community in Europe will need in the coming decade. Scrutiny of this list demonstrates that many of the listed opportunities need the infrastructure we propose. This is most striking for the majority of the opportunities listed under ‘Biological and Medial Sciences’, such as ‘Bioinformatics infrastructure for Europe’, ‘European network of advanced clinical research centres’, ‘European network of bio-banks and genomic resources’ and ‘Model testing facilities for biomedical research.’ From the environmental sciences, the ‘European infrastructure for research in, and protection of biodiversity’ must be mentioned and from the social sciences the ‘European Social Survey (ESS)’, because this last one reflects the kind of research that on a national scale will be enabled by the DANS project.

The listed physics and astronomy research infrastructures (FAIR, SPIRAL II, KM3NET, ELT) are natural candidates to fully exploit a grid infrastructure, as is already the case for LOFAR and the CERN experiments. Last but not least it needs mentioning that one of the projects listed as ‘Global’, ITER (fusion physics), has been officially introduced to the European grid arena as an application area in the second phase of the EGEE project.

7.4 e-Infrastructure development in Europe

In sync with the work of ESFRI its sister organisation e-IRG (e-Infrastructures Reflection Group) has been working on a long term European strategy for e-Infrastructure. ESFRI and e-IRG operate side by side; e-IRG focuses only on genuine general IT infrastructures irrespective of disciplines or specific sciences – encompassing networking, highly advanced computing and grids and data storage. e-IRG has produced an European roadmap for e-Infrastructures in July 2005 during the Luxembourg EU chairmanship. In this roadmap the overarching vision is that significant strategic investments in e-infrastructure are to be made by the Commission and the member states together if Europe wishes to remain competitive. Many of the investments that the Netherlands will have to make to comply with these demands are by and large covered in the BIG GRID proposal, as they include, among others, large data storage facilities, high performance computing facilities and an authentication and authorisation infrastructure. The opportunity for BIG GRID is huge in this respect, as other countries have to follow 'normal procedures' in order to create the kind of funding necessary. BIG GRID has the potential to make the Netherlands one of the 'first movers', which is interesting from both a scientific and an economic point of view.

8. Technical Case

All science and innovation cases described above require computer networks, computing, and storage technology on an unprecedented scale. Linking these together in a coherent way, using web and grid technology and employing the generic e-Science paradigm, is a *conditio sine qua non* for progress in these fields, as well as for maintaining international competitiveness. But also the actual construction and operation of such a large infrastructure involves new challenges in *scaling* and *validation* of the web and grid middleware, the bridge to the applications, and last but not least the sheer number of physical machines. Yet, within the timescale and budget proposed, it will be possible to build and operate such an infrastructure, and there is enough skill in the Netherlands to realise and exploit this. This is supported by the performance of the team in earlier projects like EU DataGrid, EGEE, DEISA, LCG in Europe, and in national programmes such as VL-e, GigaPort-NG, and the NCF NL-grid project, and by initiatives such as GridForum Nederland.

To be successful, the project has to follow the philosophy that the basic e-Science infrastructure has to be standardized as far as possible, and that the application dependent generic part should follow international standardization where or whenever applicable.

8.1 Web and grid technology

Since the emergence of large operational infrastructures a few years ago, significant progress has been made in building software that enables access to geographically dispersed resources that are managed by independent organizations. From its humble beginnings of linking up just a few high-end computer systems, grid middleware has matured such that it can now service distributed infrastructures of over 20,000 computers spread over 150 organisations, and this infrastructure continues to grow. This has become possible by significant advances in both software design, the software engineering methodology, and by growing coordination of the policies and operations that govern international grid operations.

By and large, the maturing of the web and grid service technology has resulted in international consensus on the software architecture that will be the basis for generic resource access. This concept of Service-Oriented Architectures, and the Open Grid Services Architecture OGSA that leverages this model, allows flexible mixing and interconnection of both web and grid services, and of higher-level services that are application-oriented. It is supported both by industry and by open source projects like the Globus Toolkit and Apache. Standards bodies like the Global Grid Forum (GGF) are consolidating, and projects that were previously fervently independent are now considering interoperability their prime goal, not only on the software level, but also extending to operational coordination. This means that building the BIG GRID today is both timely and feasible with respect to the grid middleware.

Key parts of this software and operational infrastructure have been designed, implemented and coordinated by a team of experts that will also be building the BIG GRID. The international reputation of the team in this field of work is high, and key team members hold prominent positions in the international community, covering the full spectrum from global trust coordination, to deployment management, and down to resource information services.

To reach the new scales foreseen for the BIG GRID, validation and testing of new middleware will be often required. Based on our experience in the various international grid test beds at NIKHEF, from the building and operation of a grid fabric research cluster, in the challenges faced daily on the European grid infrastructure of EGEE, and in the scaling and validation programme of VL-e, we know that a dedicated ‘certification’ and testing environment is needed. On this certification environment, experimental services and new technologies can be evaluated, without wreaking havoc on the actual production facility. Since this certification environment can be operated with a lower level of assurance, it can be cost-effectively combined with the hardware scaling tests facility – which also ensures that a sufficiently diverse test bed is available to pilot various different services.

8.2 The application-dependent generic layer

Such a nationwide distributed infrastructure has to be developed along the lines that those parts that are application dependent, but generic for the application itself, have to be developed near the application groups that develop the applications. For this development to be effective, the application groups should avail over a local infrastructure that is at the same time part of the BIG GRID, as explained in more detail in section 8.4. Implementing such a distributed system, allowing local control whilst maintaining integration with and accessibility of the grid at large, is an essential component of the BIG GRID that we envisage.

The application-dependent generic layer is essential for the uptake by science domain experts. Without it, the grid would fail to reach most of the scientists, as they are rightfully interested in science, and not in manipulating computers. This key issue is being addressed in the Netherlands today in groundbreaking projects like VL-e and the NBIC BioRange and BioAssist programmes, and on the European and global scale by projects such as HealthGrid, CrossGrid, MyGrid and many others. The prominent role of the domain scientists in projects like VL-e and bioinformatics programme BioRange will ensure adequate measures to bridge this gap. The BIG GRID grid is essential to ensure that their commitment is matched by actual benefits in terms of capacity.

At the same time, there will also be a large body of central services and resources that are forming the core of the NL-science grid. Generic services such as bulk storage and large databases, but also the software used to access such services, like brokers for compute and storage resources, or shared databases for sensitive information, have to be realized and maintained centrally.

8.3 Building the physical infrastructure

This proposal is aiming at a facility with a raw computing power of over 35,000 of today's personal computers and over 30 million Gigabytes of storage in its final stage. This is of an unprecedented scale, but certainly feasible.

First of all, the continuing improvement in the price to performance ratio, usually referred to as Moore's Law, is essential to reach this target. Moore's Law implies that the CPU capacity of 9,200 of today's computers could be realised in the year 2009 with only 1600 physical boxes. These boxes will be far more advanced than the current systems, probably featuring quad-core CPUs running at extremely high frequencies. Also, more and more processors can be accommodated in a single box.
The technology to build clusters is today a commodity that can be bought from many hardware vendors. Clusters of up to 500 processors exist in the Netherlands today. Increasing the scale to 2,000 machines is certainly feasible.

For storage on disk and tape, the annual improvement in the price-performance ratio is even higher, with the capacity per Euro spent doubling every year. This has a considerable advantage: the bulk of the storage will be required in the later phases of the project, when among others LOFAR and the LHC will start actual data taking, but the supporting network and storage infrastructure (i.e. the systems needed to manage the tape storage, and to move data between tape and disk) will need to be tested well prior to data taking. A new aspect to mass storage will be the high rates at which data is to be stored, where the mass storage system is almost directly linked to the data producing centres of LOFAR and the LHC experiments via dedicated optical networks. Moreover for (bio)medical applications secure communications and storage are of utmost importance

Current exploratory work via 'service challenges', and experience from CERN, the only multi-petabyte data store in Europe today, indicates that the construction of such a facility at the projected scale is entirely feasible. Periodic tests in the context of the LCG service challenges and of LOFAR will be used to monitor bottlenecks in the system and address the operational challenges before mass data collection starts, and before the other science cases will experience such problems.

Also, by virtue of the science use cases, this entire effort is firmly embedded in international collaborations that are building the LHC Computing Grid project. Via this collaboration we will have access to the expertise of those few tens of sites in the world that will be building similar large-scale systems.

Like for the web and grid software, also the physical infrastructure needs a certification environment to do tests on scaling. Experience in building the current (small scale) compute facility at NIKHEF, and the knowledge obtained from the implementation of the NL-Grid project, has shown that with every sizeable increase, system components (from computer networks up to high-level schedulers and the storage broker software) that were functioning correctly before will now fail in unexpected ways. In such cases, it is essential to have a testing environment were such problems can be diagnosed and fixed, but it is also essential to experiment with new technologies that will prevent scaling problems from occurring on the production system. Such a 'certification' facility will by nature have a lower guaranteed service level, but could still be used in many of the science use cases above, specifically if application-level middleware can gracefully handle system failures. Thus, the capacity of the certification facility is not 'lost' to science.

For realistic tests, the size of the certification facility need only be approximately 15% of the full size, since artificial loads and special test scenarios can be used for stress testing. As stated before, it is this same facility that can also be used to evaluate and certify the generic software needed to enable the science grid to operate.

8.4 Distributed facilities

The NBIC BioAssist programme aims at supporting the various medical, pharmaceutical and biological application domains with the adequate ICT infrastructure. For that reason a centralized-decentralized model for support was being developed. This model in a natural way links up with the role bio- and neuro-informatics plays in the support of life science applications such as genomics and neurosciences and ecology. The model was recognized by the NWO midterm NGI evaluation committee as the most adequate one for bioinformatics support.

To be able to realize such a situation the necessary infrastructure has to be established. In this process the Dutch academic hospitals play a pivotal role. BioAsssist has started from their own funding with a small project to connect these hospitals with SARA and intents to further up-scale this infrastructure under this proposal. Therefore an amount of money is asked to further realize grid nodes in the various academic hospitals as well as in some of the major institutes like the cancer institute NKI and the brain research institute as well as the TNO institute for food and medical research. This centralized/decentralized approach will use secure communication (on the fly encrypted and via dedicated optical links) with the SARA data storage to be able to store and retrieve the large amounts of data that will be produced in the future in the biodiversity, biobanking, neuro-informatics and drug design cases.

BioAssist has decided to base itself on service oriented architecture more in particular on the combination of web and grid services. This fits in well with the infrastructure and middleware being asked for in this proposal.

8.5 National and international embedding and synergies

The BIG GRID will greatly strengthen the national science position in the international arena, because it allows easy adaptation towards European infrastructures based among others on Géant and further overseas international e-Science infrastructures like for instance the BIRN collaboration in the USA. It will have to do so based on sound technology such as Service Oriented Architectures more in particular preferably based on OGSA and implementations of the Web Service Resource Framework such as the Globus Toolkit version 4.

To realize such an infrastructure demands close collaborate with the various grid infrastructure and network activities that take place in Europe such as the EU projects EGEE and DEISA.

9. Business Case: Building the BIG GRID

To enable the science cases, their requirements will have to be met by the infrastructure in terms of capacity. This infrastructure includes a central core location, with tape, high-speed disks and mass cluster computers, as well as geographically distributed facilities with compute and disk servers. At the same time, 'people ware' is needed to turn the `bare metal' in to an excellent scientific instrument.

9.1 Infrastructure

To account for Moore's Law, implying the doubling of compute power to be bought per Euro spent every 18 month, and the improvements in storage which halves the cost per Terabyte every 12 months, the science requirements have been expressed in performance units throughout.

For compute power, the unit in this case is in the *SpecInt2000 (SI2k)*, which mainly takes into account the performance of processors in doing calculations with integer numbers. For the class of systems relevant to our scientific applications, the expected floating point performance improvements are similar to those of the integer performance. One kSI2k corresponds roughly to a single 2.5 GHz Intel Xeon processor.

For storage, the unit of expression is Terabytes. Like with compute power, the price-performance ratio for storage is improving exponentially, and at an even faster rate: every 12 month. Approximately half of this improvement is obtained by increasing actual storage density (which helps to bring down not only investment costs, but also the housing cost component). The other half is price reductions as a result of improvements in the production process.

To clarify the effects of the performance improvements over the project life time, the proposed infrastructure has been defined in terms of the capacity required. For computing, for example, 594 kSI2k new capacity required in 2006 can be satisfied with 143 'boxes' containing two Intel Xeon CPUs, each with a dual core, running at 2.5 GHz. In 2009, only 36 boxes would be needed, that would then cost only 25% of the 2006 investment.

This clearly indicates that a 'just in time' acquisition model is the most cost-effective way to construction. Structurally leaving capacity unused – *i.e.*, buying the hardware too early – implies that one cannot profit from the performance increase over that period, and thus one ultimately gets less capacity than would have been possible otherwise. It also allows for some remaining uncertainty in the acquisition time line, especially in the life sciences domain, to be resolved. On the other hand, this model needs an agile acquisition process, able to respond almost instantly to the application needs once they arise. The organisational structure outlined in section 9.2 ensures that the turn-around time from scientific need to actual deployment meets this agility requirement.

9.1.1 Infrastructure costs

The total cost of creating the infrastructure contains not only the direct investments in hardware, but also the cost of housing the equipment in a suitable environment, supplying the electricity, cooling the generated heat, and in general ensuring the correct operation of the facility. The costs shown in the tables below are based on a detailed cost estimate calculation, presented in Appendix B, which also contains the breakdown in local versus remote facilities. Only totals for investment and exploitation over the entire duration of the project are shown here. Also consult appendix B for details on the capacity ramp-up scenario for computing and storage.

Investments

Where possible, acquisition costs are based on extrapolations of actual purchases in 2004 and 2005. The cost of a compute system with a capacity of 1 kSI2k (typically a 2.5 GHz Intel Pentium based system) in 2005 was €810. This includes sufficient system memory and disk to satisfy the science applications mentioned in this proposal. These costs have been extrapolated using Moore's Law. A similar estimate has been made for storage prices, using actual quotes from manufacturers in 2005.

Hardware maintenance

After acquisition, there needs to be on-going hardware maintenance of the equipment. In the exploitation budget we have assumed a constant overhead factor of 5% of the equipment investment to cover maintenance over a 4-year exploitation period.

Housing

Modern computer systems consume a significant amount of power (approximately 200W per box), which at current energy prices translates into €105 per machine per year. To this, cooling costs (approximately an additional 50% is needed to cool away the generated heat) and housing costs (the costs of the controlled environment, cabinets and special flooring) should be added. For a rack containing 25 systems, this amounts in total to 9 k€ per year. The same type of exploitation costs has been taken into account for tape and disk storage, assuming that the power consumption per m^2 floor space for disk and computing is approximately similar. Based on the currently available disk storage systems, we estimate that in 2006 one rack will contain 44 Terabytes of disk space, and assume equivalent floor space needs for 44 Terabytes in a tape robot.

For the central facility, the housing quality is an important issue. The ability to handle a large amount of equipment and the necessary cooling and power infrastructure is not to be taken lightly. Therefore, organizations that have proven expertise in hosting large computing and storage facilities, will also act as hosts for the central components of the BIG GRID.

For the distributed facilities, the power and housing costs will be borne by the host organization, thus reducing the housing costs for the equipment located there. For those systems located at the laboratories of the scientists, remote support and maintenance must be provided, which must be added to the exploitation costs. To account for this, an additional surcharge of 100% has been added to the maintenance cost (the nominal 5% overhead on the investment cost) for those facilities. Besides, a significantly sized compute facility (1024 CPUs in 2006 and an equivalent amount of compute power added in 2008) will be housed by Philips Research in their Eindhoven facilities. The housing cost of that part of the distributed facilities will be taken up by Philips Research, thus significantly lowering the overall housing cost for the BIG GRID.

Economic life time

For computing equipment, an economic life time of four years is feasible. This implies that also after the BIG GRID project itself ends in 2009, the equipment bought in 2007 and later will remain available, which (due to Moore's Law) constitutes the bulk of the capacity. Of course, this capacity can and will be used by scientists in 2010 and beyond.

9.1.2 Compute facilities

Based on the science use case requirements in section 4, an informed estimate of the processing capacity needs can be derived. This includes both capacity available at the central facility as well as that available in the distributed clusters. The most detailed capacity requirements are available for the particle physics case, whereas for the life science community these requirements may be more susceptible to modifications. The total compute capacity available in the BIG GRID for each year is shown below.

Year	**2006**	**2007**	**2008**	**2009**
Capacity central facility (kSI2k)	1250	4150	8500	14000
Distributed capacity (kSI2k)	2900	2900	7300	13600
Total computing capacity (kSI2k)	3294	7050	15800	27600

Location	**Investment (k€)**	**Exploitation (k€)**	**Total (k€)**
Central Facility	3088	2413	5501
Distributed	3201	848	4049
Total			9550

The costs shown in the table above include both the central and distributed clusters. The investment also amalgamates the cost for connecting the machines to a local network, and allows for sufficient system disk and memory to run all scientific applications (in 2005, this corresponds to 3 GB of internal memory and a 120 GB local disk per node). Part of the 2006 acquisition need has been covered already by existing funds from NCF. The budget for 2007 and 2008 includes scheduled replacement of current test bed facilities.

9.1.3 Mass storage

As is the case of compute clusters, the particle physics area has estimated the required amount of data storage in detail. Also for storing life science data and imagery, substantiated estimates are available: in 2009 an aggregated capacity of 500 TB should be available, with a growth path towards this goal over the 2006-2008 period.

For storing archive data in the social and cultural science domains, the required reliability level of the persistent storage is higher than in the physics and astronomy domains, where copies of the data will be available in other locations as well. Thus, data should be stored redundantly on tape, requiring twice the amount of tape storage space. This has been made explicit by doubling the tape requirements of the shared pool in 2007, as shown in appendix B. The data for the life sciences domain will be present both in the mass store at the central facility and in the distributed facility, thus also providing redundancy.

Year	**2006**	**2007**	**2008**	**2009**
Mass store capacity (TB)	1020	3560	9600	14100

Storage type	**Investment (k€)**	**Exploitation (k€)**	**Total (k€)**
Tape storage	2960	1552	4512
Fast Disk	3939	1638	5576
Total			10088

Also in this table part of the acquisition need in 2006 (50 TB) has already been covered by funds from the NCF ‘Framework versterking Nederlandse grid en e-Science Infrastructuur’ project.

Cache storage

The storage foreseen in the BIG GRID is made up of two classes: reliable mass storage systems backed by magnetic tape and 'cache' disks for use during data manipulation and data acquisition. The former (mass storage) can by its nature only be provided at the central facility. The latter (cache and scratch disk) should be located close to the actual processing centres, and thus be located both at the central facility and in the distributed clusters. Equipment costs for cache storage is also cheaper than reliable disk mass storage by about 40%.

The actual cache capacity needed at the central facility (i.e. where reliable mass storage is also available) is related solely to the number of processing nodes. Taking current best practice of allocating a non-node bound cache of 20 GB per processor at the current CPU performance (which satisfies all data intensive applications today), this translates into a disk to compute ratio of 54 GB/SI2k. This ratio has been applied for calculating the cache storage needs both at the central facility and for the cache storage co-located with the distributed compute facility at Philips Research in Eindhoven.

The needs for data storage at the distributed facilities must be added to this. That local storage is used largely for data acquired locally or accessed frequently from the local compute clusters. Each local cluster will avail over 10 TB of disk space in 2006, and this amount will grow with 10 TB in the years 2007 and 2008. In 2009 and 2010, the increase will be 20TB and 30 TB per location, respectively.

Year	**2006**	**2007**	**2008**	**2009**
Capacity at Central Facility (TB)	70	190	460	870
Distributed Capacity (TB)	220	340	560	800
Total cache capacity (TB)	290	530	1020	1670

Location	**Investment (k€)**	**Exploitation (k€)**	**Total (k€)**
Central Facility	311	177	488
Distributed	456	150	606
Total			1094

9.1.4 Specialized low-latency computing service

For 15 years, NCF has taken care of high performance computing, and will continue to do so. For this historical reason, this proposal does not apply for a high performance computer, but does ask for a small share, to be able to bring the high-end computer on a level which holds pace with the other investment requirements in this proposal for cluster capacity and for data storage and processing. The ultimate advantage here is to have a science grid, with all components in balance, with access for numerous scientific disciplines, and with optimal opportunities for using the right compute facility for the right application. Looking at European developments, the actual balanced system in the context of this proposal is estimated to have a peak performance of 40-45 Teraflops. This leads to a high-end computer share of 5 M€. NCF will apply this amount, together with its own funding, for the purchase of such a system for computational scientists from various fields. To give an idea on exploitation costs, for a period of 4 years, the exploitation of a high performance system (power, cooling, administration, floor space) amounts to about 70% of the purchase price.

Location	**Investment share (k€)**	**Exploitation (k€)**	**Total (k€)**
Central Facility	3000	2000	5000
Total			5000

9.1.5 Networking equipment for the central infrastructure

The central facility must be extremely well connected to the national backbone of SURFnet 6, and be able to interface directly to the lambda infrastructure. Also the approximately 1800 machines must be connected to GigaPort-NG and be able to communicate at full speed with the data stores. Given the port density of current routers, two high-performance routers will be needed to sustain the networking needs of the central facility. Such routers will cost 200 k€ each. The operational costs of the network equipment are significantly higher than the 5% used for the maintenance of compute and storage systems since, by virtue of there being few, these routers need high-availability maintenance contracts, incurring a 20% yearly exploitation cost of the initial investment, in addition to the regular housing cost.

Location	**Investment (k€)**	**Exploitation (k€)**	**Total (k€)**
Central Facility	400	295	695
Total			695

9.2 Organizational aspects

By its nature a grid infrastructure requires a good balance between central and distributed activities. For most activities, essential for grid middleware development and roll out, we can rely on the existing national and international collaborations in which Dutch groups participate on a very high level. Nationally the VL-e project is more than adequate to fulfil this role, while in the international context the EU-project EGEE plays an important role. The interaction of the key Dutch parties, which propose to build the Dutch science grid, with these international developments is excellent.

With scaling up the grid infrastructure to serve an ever larger number and variety of user communities we need to install on a national level a well balanced structure. Key components discussed in the sequel, are:

- A steering group;
- Operations management;
- Grid expertise centre and help desk.

9.2.1 Steering group and management

For the modest investments in grid infrastructure in recent years an organizational structure (steering group) has been put in place. This group, with representatives from a.o. VL-e, ASTRON, SARA, NIKHEF and NCF, has decided on purchase of equipment (clusters and storage, both centrally and locally) with special attention to the right timing of purchases. This model has worked well, but needs a broader base to realize BIG GRID.

The composition of this steering group will therefore be renewed to achieve good representation of the various scientific user communities (physics, astronomy, life sciences, humanities, industrial research labs), that will use the grid infrastructure. Key persons, with a clear knowledge of the needs of their community, will be appointed. A full time operations manager (see below) is q.q. member of the steering group. Main responsibility of the task force is to decide on priorities, budget issues and actual purchases.

Steering group and management	**Cost (k€)**
Operations Manager (4 years, 125k€/year full cost)	500
Total	500

9.2.2 Operations management

Though the grid is distributed in nature, it is also true that housing and operating computing and storage equipment is a ‘craft’. This is even more so for software installation, distribution, production, etc. Few centres qualify for such a task. SARA, as the long standing Dutch national HPCN centre and since many years deeply involved in grid projects is a likely candidate to provide a large part of grid operations. This however does not imply exclusivity. It is expected that the ‘Rekencentrum Groningen’ (RC-RUG), especially in connection with LOFAR, will also play a role as housing centre. Collaboration has already taken shape by the recent involvement of RC-RUG in the EGEE project in the ‘Operations’ and ‘User support’ activities. The grid infrastructure furthermore includes smaller distributed facilities. These will be (remotely) operated from the national grid operations centre. Part of the infrastructure (in the order of 15%) will serve as testing and certification environment. NIKHEF is prepared to take up its role as housing and operations partner here, which is a continuation of its current role in both EGEE and VL-e. Because grid operations management needs to be coordinated on a national scale, we opt for a ‘single’ operations manager who oversees the national context.

The national operations management should be part of (and can largely be derived from) the international operations management structure as is being developed in the EGEE project and will encompass the Tier-1 organisational structure that will be set up for CERN’s LHC experiments. To facilitate the operations management a yearly (manpower) budget of 2 fte is requested (800 k€).

Grid systems operations	**Cost (k€)**
Grid systems managers (4 years, 2 people, 100 k€/year full cost)	800
Total	800

9.2.3 Grid expertise centre and help desk

BIG GRID requires the formation of a grid expertise centre, including a helpdesk. The national helpdesk can be viewed as one entry point for telephone or email support to users. Questions like ‘why does my job give an error’ might not be easy to solve for a user of a distributed infrastructure and should therefore be tackled by the helpdesk. Good helpdesk working models developed in international grid projects will be implemented on the national scale. This also allows the Dutch research community to profit from the support structure that will take shape on an international scale.

A crucial part in the adoption of e-Science and the use of grid technology is the link between the underlying scientific discipline and the ‘grid experts’. In order to efficiently use the ICT possibilities, there is a growing need for knowledge, which combines the scientific knowledge with practical skills in the area of grid technology. We feel that building this knowledge eventually will be very important for the success of the overall hardware investment. We therefore envisage the need for ‘Domain Application Analysts’ to form the core of the grid expertise centre. The expertise centre will build and maintain close contacts with the current Dutch grid expert groups, with the activities of the ‘Vereniging Gridforum Nederland’ and to the dissemination activities of VL-e and EGEE. The centre will also provide courses, contributions to symposia and conferences, tutorials, business days, booklets, etc.

Grid expertise centre and help desk	**Cost (k€)**
Help desk (4 years, 2 people, 75k€/year full cost)	600
Domain application analysts (4 years, 4 people, 100k€/year full cost)	1600
Total	**2200**

9.3 Capacity summary

By presenting the capacity and capabilities of the BIG GRID in abstract units like SpecInt2000 and Terabytes, the actual scale of the facilities may not be immediately obvious. Making the comparison with today's technology, once more emphasises the unique scale of the proposed infrastructure, and makes it clear how it will enable key scientific breakthroughs in all sciences.

	BIG GRID capacity in 2009	**In today's (2005) technology this would correspond to**
Total compute cluster capacity	27 600 kSI2k	10 530 processors
Mass storage on tape	7 910 Terabyte	1 700 000 DVD's
Fast disks	6 190 Terabyte	1 317 000 DVD's
Disk caches and data acquisition space	1 670 Terabyte	20 875 home PC disks

The 45-Teraflop supercomputer in the BIG GRID qualifies as the 4th fastest system in the world.

9.4 Financial summary

Summarizing all equipment investments, maintenance, operations management, and the user support and expertise centre, the total cost to realize the BIG GRID is € 29,926,000.

Cost Item		**Cost (k€)**
Central facility:	Equipment investment	10 698
	Exploitation (operations, maintenance, housing)	6 073
Distributed facilities:	Equipment investment	3 657
	Exploitation (operations, maintenance)	998
Specialized low-latency compute service		5 000
Personnel costs		
Operations manager		500
Grid systems management		800
Helpdesk and support		2 200
Total		**29 926**

Glossary

ASCI	Advanced School for Computing and Imaging
ASTRON	ASTRronomisch Onderzoek Nederland
BIG	Budget voor Investeringen in Grootschalige onderzoeksfaciliteiten
BSIK	Besluit Subsidies Investeringen Kennisinfrastructuur
CAD	Computer Aided Design
CAVE	Cave Automated Virtual Environment
CERN	Centre Europeènne pour Recherche Nucleaire
CKO	Centrum voor Klimaat Onderzoek
CWI	Centrum voor Wiskunde en Informatica
DANS	Data Archiving and Network Services
DEISA	Distributed European Infrastructure for Supercomputing Applications
e-IRG	e-Infrastructures Reflection Group
e-Science	enhanced Science
EGEE	Enabling Grids for E-sciencE
ESFRI	European Strategy Forum on Research Infrastructures
EZ	(Ministerie van) Economische Zaken
fMRI	Functional Magnetic Resonance Imaging
FOM	Fundamenteel Onderzoek der Materie, de NWO-stichting voor Natuurkunde
Gigaport-NG	Gigaport Next Generation (BSIK project)
HEP	High-Energy Physics (particle physics)
HPC	High Performance Computing
ICT	Informatie en Communicatie Technologie
IMAU	Instituut voor Marien en Atmosferisch onderzoek Utrecht
KNAW	Koninklijke Nederlandse Academie van Wetenschappen
KNMI	Koninklijk Nederlands Meteorologisch Instituut
LCG	Large hadron collider Computational Grid
LEP	Large Electron Positron collider
LHC	Large Hadron Collider (with the Alice, ATLAS and LHCb experiments)
LOFAR	LOw Frequency ARray
NBIC	Nederlands Bio-Informatica Centrum
NCF	Stichting Nationale Computer Faciliteiten
NGI	Netherlands Genomics Initiative
NKI	Nederlands Kankerinstituut
NIKHEF	Nederlands Instituut voor Kernfysica en Hoge-Energie Fysica
NREN	Provider of National Research Network (like SURFnet)
NWO	Nederlandse Organisatie voor Wetenschappelijk Onderzoek
OCW	(Ministerie van) Onderwijs, Cultuur en Wetenschap
RC-RUG	Rekencentrum Rijksuniversiteit Groningen
RU	Radboud Universiteit
SARA	Stichting Academisch Rekencentrum Amsterdam
SI2k	SpecInt2000, performance unit for integer computations
SURF	Landelijke Stichting ter bevordering van het gebruik van ICT in het hoger onderwijs
SURFnet	Dutch high-quality Internet for higher education and research
TNO	Nederlandse Organisatie voor toegepast-natuurwetenschappelijk onderzoek
UU	Universiteit Utrecht
UvA	Universiteit van Amsterdam
VL-e	Virtual Laboratory for e-Science (BSIK-project)
VU	Vrije Universiteit (Amsterdam)

Appendix A. List of Supporting Organisations

Letters of support are available from:

Organisation	**Signatory**
Amsterdam Centre for Computational Science	Prof. dr. D. Frenkel, scientific director
Amsterdam Medisch Centrum	Prof. dr. L.J. Gunning-Schepers, chair executive board
ASTRON	Prof. dr. Harvey Butcher, director
Centrum voor Klimaatonderzoek (CKO)	Prof. dr. G.J. Komen, scientific director
Centrum voor Wiskunde en Informatica (CWI)	Prof. dr. J.K. Lenstra, director
DANS (KNAW and NWO)	dr. Peter Doorn, director
Debye Institute	Prof. dr. L.W. Jenneskens, director
Dutch Institute for Catalysis Research (NIOK)	Prof. dr. ir. B.M. Weckhuysen, scientific director
European Organization for Nuclear Research (CERN)	dr. Robert Aymar, director general
FEI Company	dr. A.F. de Jong, manager research and technology
IBM Nederland	Ir. A.A.J. Reuver, GP executive
J.M. Burgerscentrum	Prof. dr. ir. G. Ooms, scientific director
Koninklijke DSM N.V.	dr. Ir. J.G.H. Joosten, director corporate technology
LogicaCMG Netherlands B.V.	dr. Geleyn Meijer, innovation director
Materials Science Centre	Prof. dr. J. Knoester, scientific director
Nationaal Regie-Orgaan Genomics	Prof. dr. P. Folstar, director
Nederlands Instituut voor Hersenonderzoek	dr. Jaap van Pelt (steering group neuro-informatics)
Organon	dr. Rene van Schaik, head molecular informatics technologies Dept. of Molecular Design & Informatics
Philips Research	Prof. dr. Fred Boekhorst, senior Vice President
Stichting SURF	Prof. dr. Wim B.G. Liebrand, director
Universiteit van Amsterdam, Faculty of Science	Prof. dr. K.J.F. Gaemers, dean
Vereniging GridForum Nederland	Mw. M. van Dusseldorp, chair
Vereniging Onderzoek Flora en Fauna	Mw. drs. A.J.J. Lemaire, coordinator
Vrije Universiteit Amsterdam	dr. G.W. Noomen, chairman of the Board, and dr. A. Kaldewaij, Registrar
Wageningen Centre for Food Sciences	Prof. dr. W.M. de Vos, programme director, microbial functionality and safety
WTCW N.V.	drs. Johan Vos, director

Appendix B. Detailed Cost Breakdown

Distributed facilities for life sciences and the R&D cluster at Philips Research

Computing								
year	Compute kSI2lk	nonBIG	BIGNeed kSI2k	Increment kSI2k	Investment k€	Maintenance k€	Remote k€	Total k€
2006	2900	200	2700	2700	1378	69	69	1515
2007	2900		2900	400	129	75	75	279
2008	7300		7300	4400	891	120	120	1131
2009	13600		13600	6300	804	160	160	1124
Total					3201		848	4049

Cache and Local Data Disks								
year	Storage TB	nonBIG	BIGNeed TB	Increment TB	Investment k€	Maintenance k€	Remote k€	Total k€
2006	220		220	220	275	14	14	303
2007	340		340	120	75	18	18	110
2008	560		560	220	69	21	21	111
2009	800		800	240	38	23	23	83
Total					456		150	606

Central facilities

Computing													
year	Compute capacity requirements kSI2k					nonBIG kSI2k	BIGNeed kSI2k	Increment kSI2k	Investment k€	Maintenance k€	Housing k€	Σexplt k€	Total k€
	Tier-1	Astro	Bio	Shared	Total								
								184					
2006	750	200		300	1250	656	594	594	303	15	182	197	500
2007	2150	1400		600	4150		4150	3740	1202	75	403	478	1680
2008	5500	2100		900	8500		8500	4350	881	119	638	758	1639
2009	9500	3300		1200	14000		14000	5500	702	154	826	981	1682
Total									3088	364	2049	2413	5501

Mass Store, tape component													
	Storage TB					nonBIG TB	BIGNeed TB	Increment TB	Investment k€	Maintenance k€ (1)	Housing k€	Σexplt k€	Total k€
year	Tier-1	Astro	Bio	Shared	Total								
2006	250	200	60	20	530	50	480	480	600	30	61	91	691
2007	725	1000	125	30	1880		1880	1400	875	74	186	260	1135
2008	1825	3000	250	280	5355		5355	3475	1086	128	406	534	1620
2009	3550	3300	500	560	7910		7910	2555	399	148	520	668	1067
Total									2960	380	1172	1552	4512

Mass Store, fast disk component													
	Storage TB					nonBIG TB	BIGNeed TB	Increment TB	Investment k€	Maintenance k€	Housing k€	Σexplt k€	Total k€
year	Tier-1	Astro	Bio	Shared	Total								
2006	350	100	30	10	490		490	490	980	49	70	119	1099
2007	1100	500	60	20	1680		1680	1190	1190	109	191	300	1490
2008	2550	1500	125	70	4245		4245	2565	1283	173	375	548	1830
2009	3800	2000	250	140	6190		6190	1945	486	197	474	671	1157
Total									3939	527	1111	1638	5576

Cache and Local Data Disks (NAS)													
	Storage TB					nonBIG TB	BIGNeed TB	Increment TB	Investment k€	Maintenance k€	Housing k€	Σexplt k€	Total k€
year	Tier-1	Astro	Bio	Shared	Total								
2006				70	70		70	70	88	4	10	14	102
2007				190	190		190	120	75	8	22	30	105
2008				460	460		460	270	84	12	42	54	138
2009				870	870		870	410	64	16	62	78	142
Total									311	40	136	177	488