ICTRegie

Netherlands ICT Research and Innovation Authority

Towards a competitive ICT infrastructure for scientific research in the Netherlands

December 2008

Distance separates only bodies, not the mind

Desiderius Erasmus

ICTRegie
Netherlands ICT Research and Innovation Authority

Title	Towards a competitive ICT infrastructure for scientific research in the Netherlands
Number	08 NROI 258
Author	Prof.dr.ir. L.J.M. Nieuwenhuis
Version	1.7
Responsibility	Prof.dr.ir. G. van Oortmerssen, Director ICTRegie
Date	11 December 2008

Table of Contents

1. The challenge

1.1. The Netherlands: a knowledge economy in a globalized world

The Netherlands aims to be one of the most competitive knowledge economies in the world. For many years, the Netherlands has invested in building a world-leading national research and education network. This advanced network has accelerated the development of public broadband connectivity. In the Global Competitiveness Index 2008, the Netherlands moved up two spots to the 8th place[1], supported by a number one position for technological readiness, including factors like pc and internet usage and broadband internet subscribers. Information and communication technology (ICT) is the *general purpose technology*[2] of our time given the critical spillovers to the other economic sectors.
To be competitive with a knowledge based economy in a globalized world requires innovative services and products with high added value. The Netherlands is considered to be one of the most globalized countries in the world, and therefore knowledge in general and scientific knowledge in particular is of paramount importance. Innovation driven growth requires well-educated people, a strong research environment and an innovation stimulating environment.

1.2. The development and importance of e-science

The application of ICT in scientific research has generated the development of e-science. E-science is collaborative science performed using large-scale distributed computing and data resources. E-science is changing the way scientific research is conducted in an increasing number of scientific fields and opening up new research methods. In its report "Goed gereedschap is het halve werk" (February 25, 2008), ICTRegie presented an overview of the relevant infrastructures, organizations and domains. The report shows that maintaining the recognized high level of scientific output in the Netherlands requires an excellent ICT infrastructure. Even more, the Netherlands has great opportunities in the field of the development of e-science. These opportunities are originating from directed investments in the network (SURFnet6/GigaPort), in extension and renewal of supercomputers (Huygens, Blue Gene), and in projects like LOFAR, VL-e and BiG Grid, next to NWO research programs like STARE, GLANCE and VIEW.

1.3. Ministerial request

Apart from the conclusions mentioned above, the study presented in the mentioned ICTRegie report informed the minister of OCW that substantial extra money might be needed to realize the stated ambitions, and that a number of bottlenecks needs to be resolved. ICTRegie proposed to develop a long term (ten year) vision in consultation with the main stakeholders, in which the key points would be addressed. That proposal was accepted by the Minister in a letter on April 22, 2008. The minister asked for a more elaborate investigation resulting in an advice on the entire ICT infrastructure for research with conclusions and recommendations that are supported by all parties involved.

---

[1] Global Competitiveness Report 2008, World Economic Forum

[2] Innovation Policy for Development: an Overview, Manuel Trajtenberg. Working Paper No. 6-2006, July, 2006, Tel Aviv University, NBER and CEPR

2. Summary of conclusions and recommendations

2.1. Conclusions

The conclusions are summarized below and elaborated further in Section 6 of this document.

- The investigations confirm that SURFnet6 in its concept and applied technologies is internationally recognized as world leading[3]. NetherLight in Amsterdam is a major hub in the worldwide network of lightpaths. The Netherlands has very good computing facilities with Huygens and LISA at SARA, the Blue Gene at CIT in Groningen (LOFAR) and the DAS-3 clusters. The VL-e project is an excellent Bsik-project that has developed a number of prototype services that need to be developed further. The BiG Grid project has pushed the e-science infrastructure and its applications to a much higher level. SARA is hosting, operating and servicing many of the above mentioned systems and supports the scientific community. The growing volume and increasing sensitivity (life science) of data drive the need for more reliable and secure data storage functionality.
- The number of scientific fields using the current ICT infrastructure is rapidly expanding. NCF and SARA have recently published 42 interviews with professors at Dutch universities in a large variety of scientific fields using the Huygens computer[4]. Each of them points out the importance of high performance computing for his or her field, the societal relevance and impact and their petascale challenges and opportunities when ICT advances. The use of the ICT infrastructure in the Social Sciences and Humanities is emerging, e.g., the Virtual Knowledge Studio supports researchers in these fields in the creation of new scholarly practices and in their reflection on e-research in relation to their fields[5]. The digitization of the world and the creation and dissemination of (scientific) knowledge is subject of a study conducted by the AWT (Advisory Council for Science and Technology Policy).
- The committee Dutch Roadmap Large-scale Research Facilities (Commissie Van Velzen), recently presented its recommendations, which have been adopted by the government[6]. The recommendations are based on a well-defined set of criteria to evaluate and rank investment proposals[7]. These criteria are tailored to domain specific facilities rather than the generic ICT infrastructure. The Committee advised to develop an accommodating policy for the ICT infrastructure. The Committee asserts that the Netherlands has an excellent ICT infrastructure and its continuation and excellence need to be secured as it is a necessary precondition for the effective use of all large-scale research facilities.
- The current funding and organizational structure complicates the provision of ICT infrastructure by limiting the autonomy and flexibility of the collective of involved organizations.
    - The infrastructure requires long term policies, because there are elements that go well beyond the lifespan of a typical project, including the investment in knowledge and experience to

[3] For instance, SURFnet contributed to the number one ranking of The Netherlands in the Earnest Foresight study. Refer to http://www.terena.org/activities/earnest/ws4/gfx/20070925-JS-geographic.pdf
[4] Het belang van High Performance Computing voor Nederland, December 2008
[5] http://www.virtualknowledgestudio.nl/
[6] Letter to the Tweede Kamer, December 4th 2008
[7] Committee Dutch Roadmap Large-scale Research Facilities, October 2008

build the infrastructure and such aspects as housing and (increasingly important) energy supply.
- There is no clear locus of general responsibility, due to partially overlapping distributed activities and responsibilities across the many involved permanent and temporary organizations.
- Multiple programs complicate alignment and synchronization due to the variety of objectives and evaluation criteria used by distinct programs granting investment proposals for the ICT infrastructure for scientific research.

2.2. Recommendations

Achieving the desired top position as a knowledge economy asks for competing on scientific and economic markets. That can only be successful when the scientific community in the Netherlands has the best ICT infrastructure available at its disposal. Maintaining the Netherlands in the top of science requires permanent investments and developments in computers, databases and tools for the ICT infrastructure. ICT is significantly impacting the way scientific research is conducted: we see rapidly increasing use of large-scale and complex modeling by a growing number of scientific fields and enormous amounts of data from high tech instruments like the LHC or DNA sequencers that need to be stored, analyzed and presented using advanced visualization methods combined with worldwide, interdisciplinary collaboration.

The Netherlands can leverage its current strong position by increasing the coherence between organizations, activities and the components of the ICT infrastructure. Therefore ICTRegie proposes four actions, which are supported by the key parties involved. They will be elaborated in Section 7.

- **Position all ICT infrastructure for science developments and operations under the SURF umbrella**. This means creation of an entity SURF Computing and Data Facilities as part of the SURF organization, next to SURFnet. By doing so, the responsibility for the ICT infrastructure is put at the general board of SURF, including all universities and research organizations, being the key stakeholders. SURF as a single entity managing the entire ICT infrastructure, will result in coherent long term user driven strategies, policies and technology roadmaps, which are adjusted on a yearly basis.
- **Provide to SURF an annual budget from a dedicated fund to cover expenses for permanent innovation, development and operation of the required ICT infrastructure.** This approach removes competition with domain specific research projects and investments. It enables long term strategic planning and long term infrastructural investments and increases the flexibility and transparency in budget allocation.
- **Establish an e-science Research Center (e-SRC) in cooperation with universities, NWO and NWO institutes, conducting e-science research supporting innovations of the ICT infrastructure and applications for science.** The e-science research center secures the continuity of the very valuable VL-e research activities and develops new research activities required for the ICT infrastructure. It facilitates both multidisciplinary research and cooperation across universities, research organizations and industry. A substantial part of its funding should be provided by SURF, based on a long term research plan.
- **Acquire a key position in the innovation, development and operation of a pan-European ICT infrastructure for collaborative scientific research.** The key position is to be achieved by building a European data center node providing large-scale data, computing, networking, and visualization services. The successor of Huygens and expansion of capacity computing power has to be aligned with this ambition. The data center node is connected to the rest of the world through a grid using excellent network and exchange facilities. The proposed key position contributes to the Dutch scientific output and attracts talented researchers, foreign R&D investments and international companies.

2.3. Next steps

After a basic decision has been taken on the recommendations, it is advised to charge SURF with the implementation, in cooperation with the other organizations involved. To monitor and supervise the implementation, a special steering group with an independent chairman should be set up. Special attention is needed for the restructuring of SARA in a public and a private part, for the positioning of NCF, for the creation of the proposed e-science Research Center (e-SRC) and for the reallocation of existing funds.
For the financing of research programs of the e-SRC beyond the proposed funding through SURF, a proposal needs to be included in the ICT Theme for the 2009 FES round. ICTRegie takes that into account.
As for the timing, the aim should be to have the new structure in place before the end of 2009. That means a principle decision should be made early 2009, which is also needed to safeguard the Principal Partnership position of the Netherlands in PRACE.

3. To be the best, the Netherlands needs the best

To become one of the most innovative knowledge economies in Europe and the world, the scientific community in the Netherlands needs to have the best ICT infrastructure possible. This is a necessary condition to create scientific and economic competitiveness on global markets. The main forces driving the need for further improvements are the rapidly growing amounts of data, the increasing number of scientific fields using e-science, and the fast expanding global collaboration.

3.1. Growing amount of scientific data

Experiments using detectors, medical imaging instruments or multi-sensor arrays, but also library projects building digital archives all lead to one common challenge: managing the volume and complexity of distributed data to be processed. Often, the data are already distributed from the very start, being produced by different research groups or distributed sensor networks. Experiments and simulations generating Petabytes of data per year need more data storage and processing power than ever can be located in a single facility, with data utilized by researchers all over the world. At the same time, handling the complexity and combining the data allows for completely new forms of research. Complex problems are understood in the context of the overall system they belong to.

3.2. Increasing number of scientific fields use e-science

Many scientific disciplines are undergoing technological revolutions now: computers and computing techniques increasingly penetrate science. This approach is becoming important for a wide range of scientific fields. The positive effects of using large scale ICT infrastructures are well-documented and well-known in areas such as physics and chemistry. These front runners on using e-science and computational science both benefited from and contributed to the development of excellent infrastructures. As is already true for computational science, the use of e-science is not confined to a small subset of large-scale research. In the Netherlands, e-science and computational science has successfully reached out to a wide range of scientific fields, e.g., life sciences, pharmaceutical science, cognitive sciences, high energy physics and astronomy. Similar developments in other countries are extensively described in policy documents around the world.

3.3. Expanding globalization of research activities

The development of international grid-based infrastructures, applications, and collaboration environments goes hand in hand with the internationalization of research. Technology is pushing more globalization through improved connectivity and network performance and availability of global infrastructures. However, the main drivers of globalization are growing research needs, triggering the move towards large-scale virtual research environments. An example is Bioinformatics. At the European scale it is almost as tightly organized as on the Dutch scale. At the European scale the EBI (European Bioinformatics Institute) is the natural leader for large projects. One of these projects is the ESFRI project Elixir, which aims to set up the next generation European infrastructure for bioinformatics. This infrastructure will consist of a data warehouse at the EBI that is connected to many niche databases, but it will also include software that addresses interoperability issues such as grid applications and web-services.

4. Social and economic impact

"It is excellence by innovation that we want, and for that you need to stick out your neck. Now is the time for us to really fight for that." - Jan Peter Balkenende, Rotterdam December 3, 2008

Large-scale research facilities are of strategic importance

Large-scale research facilities are of invaluable strategic importance for the Dutch knowledge economy and for a flourishing innovation climate[8]. These facilities are essential for scientific progress and for conducting top research and have large social and economical impact. Large scale research facilities attract top scientists leading to a concentration of human capital and act as node.

Key innovation areas building on excellent scientific research

The economic opportunities build on top of the scientific opportunities. The Netherlands has defined six key innovative areas: High-tech Systems and Materials, Flowers & Food, Water, Creative Industry, Chemistry, Life-Sciences and Health[9]. These innovation areas are driven by scientific results from the above mentioned scientific fields that are demanding the proposed ICT infrastructure. An excellent ICT infrastructure for scientific research accelerates scientific fields that drive innovations in the key innovative areas of the Netherlands. Clear examples of such are life sciences and chemistry. But also climatology, which is of importance for the Netherlands in particular, and energy.

Computational science enhances progress in many areas

Computational e-science enhances theoretical and experimental progress in many areas of science critical to the scientific and societal needs of the 21st century. Successes have been documented in areas such as biotechnology (e.g., genomics, cellular dynamics), advanced energy systems (e.g., fuel cells, fusion, alternative energy), nanotechnology (e.g., sensors, storage devices), and environmental modeling (e.g., climate prediction, pollution remediation, flood prediction). Computational science offers the best near-term hope for progress in answering many scientific questions in such disparate areas as the fundamental structure of matter, the functions of enzymes, the production of heavy elements in supernovae, the climate change and the spread of infectious diseases.

Hub function in Amsterdam

Through a number of advanced ICT infrastructure projects, the Netherlands has become an attractive country for scientific research. The hub function of the Amsterdam Internet Exchange (AMS-IX), the global connectivity achieved through SURFnet/GLIF and the active participation by Dutch groups in various international projects all have contributed in attracting high-tech companies and research centers. The increased visibility as center of expertise in the field of e-science combined with excellent ICT infrastructure for scientific research will further increase the attractiveness of the Netherlands for academic and industrial R&D.

---

[8] Dutch Roadmap Large-scale Research Facilities and Innovation Platform
[9] Ministry of Economic Affairs

Education of e-scientists

A recent study by the Council of Competitiveness in the USA shows that internationally best-of-class industries agree that high performance computing contributes to their innovativeness[10]. These companies indicate that scarcity of talent is one of the important barriers for broader use. The Netherlands can adapt to these needs and offer educational programs building on knowledge and practice gained from using the scientific ICT infrastructure.

Large-scale Research Facilities depend on the ICT infrastructure

The Dutch Roadmap Large-scale Research Facilities evaluated investment proposals for large scale research facilities using a set of well-defined criteria and peer review processes. The importance of ICT infrastructures for science is illustrated by the heavy reliance of initiatives that are supported by the Roadmap Large-scale Research Facilities. All initiatives in which the committee suggests active participation (grades A and B) rely heavily on collection, aggregating, processing and distributing data, for which ICT infrastructures are needed[11]. Their project descriptions address the issues of grids and ontologies, data collecting, remote sensor systems, distributed sensor systems, pan-European experiments, data acquisition and processing, data handling, virtual laboratories, and data quality of very large collections.

Advanced ICT infrastructures for science stimulate the general ICT infrastructure development

The SURFnet case demonstrates that constantly pushing the envelope on networkfacilities for education and research has a stimulating effect on the development of the general national network. Many optical fibre projects have been started from the seeds laid by SURFnet, and by applying the technical and organizational knowledge and experience developed. The same can be said on the development of high performance computers and large storage systems. Multicore computers are now becoming the standard, even for desktop applications, and the exploitation of the available power by smart software relies on the experience gained in the scientific domain. The use of grids and grid software is developing as well and gradually changing the architecture of the general available services. Fast and advanced data access, data mining and search algorithms find their way in general applications outside the field of research.

---

[10] Benchmarking industrial use of high performance for innovation, Council of Competitiveness, USA 2008

[11] Clarin, ESS, XFEL, E-ELT, ESS-PN, KM3NET, ICOS, LIFE WATCH, European Biobanking and Biomolecular

5. Collaboration requires matching our peers

The effort that other countries and regions put into developing and ensuring continued presence of high level ICT infrastructures for science underlines its significance. The ambitions of other countries serve as an indication of perceived relevance, of the level of competence that must be matched, of what is needed to effectively collaborate and what must be done in order to be world-leading.

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 in the order of hundreds of millions of dollars to the development of grids. Most other developed countries are also investing at significant levels. Industrial interest is also apparent, with IBM leading the pack with a billion dollar investment in this new technology.

5.1. United States

The United States stresses the importance of their leading position in networking and information technologies and they put in a lot of effort to fend off what they see as a strong competition from the Asian Pacific region, China and the European Union. A recent report highlights in particular the need for expertise, large-scale, long-term, multidisciplinary activities and visionary, high-payoff goals. This is intended to complement the goals and vision for leadership in the physical sciences articulated in the American Competitiveness Initiative.[12]

5.2. United Kingdom

In the UK it is stated that significant investment is needed in the development of relevant skills, and researchers must be provided with access to a balanced and flexible high end computing infrastructure that is amongst the best in the world, with regular upgrades to keep pace with the accelerating rate of technological advance. It is stated that continuing funding is necessary because simulation is becoming a bigger part of science and engineering[13].

5.3. Nordic Countries

The Nordic countries acknowledge that their combined strength is bigger than their individual strengths. They aim to reinforce the competitiveness and visibility for Nordic research worldwide further by offering a well worked-out virtual system for sharing e-science resources and competences. Their ambition is to be a world leading actor in e-science infrastructure and technology development, with resulting cutting edge research in many areas of science. Their goal for pursuing this is to ensure that the Nordic industry can have access to the

[12] Leadership Under Challenge: Information Technology R&D in a Competitive World, An Assessment of the Federal Networking and Information Technology R&D Program, Presidents Council of Advisors on Science and Technology, august 2007
[13] Strategic Framework for High End Computing, High End Computing Strategic Framework Working Group, 2008

technology experts of the 21st century, skilled with the advanced computing technologies that are the foundation for service oriented business models and a knowledge-based industry.[14]

5.4. European cooperation

The Netherlands has a very strong international reputation in networking organizations as GEANT, GLIF and TERENA, in High Performance Computing (HPC) oriented communities such as DEISA and PRACE, and in grid-oriented communities such as EGEE and EGI. The Netherlands is considered an attractive partner for participation in world leading collaborative projects.

The e-Infrastructure Reflection Group aims to support the creation of a political, technological and administrative framework for an easy and cost-effective shared use of distributed electronic resources across Europe. Where the European Strategy Forum for Research Infrastructures (ESFRI) is concerned with the creation of initiatives for building physical resources, amongst which are e-infrastructural resources, the e-IRG is concerned with the entire lifecycle of e-infrastructures, what also includes training, policy, coordination and efforts to increase synergy effects between Europe and other regions.[15]

At present, the Netherlands participates in PRACE, the Partnership for Advanced Computing in Europe preparing the creation of a persistent pan-European HPC service, consisting of several tier-0 centres providing European researchers with access to very large capability computers and dataservices and forming the top level of the European HPC ecosystem. ESFRI has identified HPC as a strategic priority for Europe. The pan-European HPC service will be a part of the European Research Area, in which the Seventh Framework Program is prepared to invest hundreds of millions of Euros[16].

The Netherlands is Principal Partner of PRACE, together with Germany, UK, France and Spain, and with nine regular partners. The Principal Partners have expressed their interest in hosting one of the tier-0 systems along the course of the cooperation, in our case endorsed by the Ministry of OCW. In fact, only a limited number of countries (including the Netherlands) have the expertise and the critical mass to significantly contribute to this infrastructure and by doing so, enhancing their position and attractiveness for advanced research activities, as well by academia as by enterprises.

14 Nordic eScience - Research, Education, and Sustainable Infrastructure Services - A strategy document for the Nordic Council of Ministers, 2007-07-17
15 http://www.e-irg.eu/
16 http://www.prace-project.eu/

6. Evaluating the ICT infrastructure

In this section, we briefly evaluate the current situation of the ICT infrastructure and how to anticipate on the changing needs originating in the many scientific fields involved and in the fast development speed of information and communications technologies. We have tried to avoid duplication with the overview in our earlier report, "Goed gereedschap is het halve werk" 08-NROI-031.

6.1. ICT infrastructure

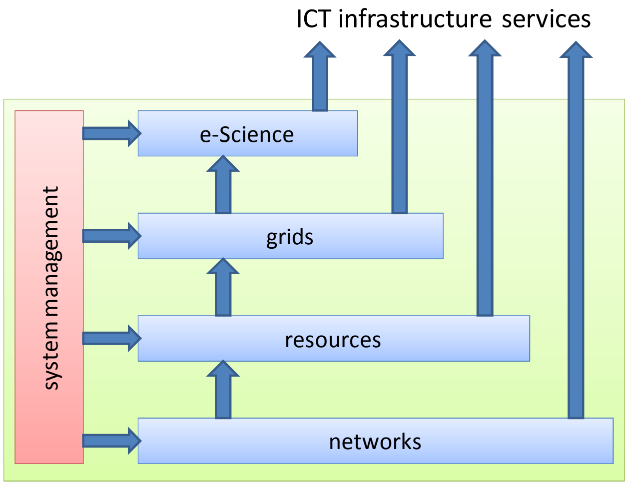

Figure 1 Architecture of the ICT infrastructure for scientific research

Figure 1 shows the architecture of the services that the ICT infrastructure provides to the scientific community, revealing building blocks and their relationships. Key ingredients for the ICT infrastructure are people and their expertise to develop and operate the ICT infrastructure efficiently and effectively and support members of the scientific community in this. Without such people, the value of the ICT infrastructure will significantly reduce. The systems of the infrastructure comprise communication networks, various resources (like computers, data storage, visualization, and scientific instruments), grid software and e-science software. System management refers to stable and reliable housing, operations, maintenance, and support activities of all building block systems. The architectural view shows a hierarchy of services (top level services make use of the underlying services) and points out that scientific users can use services of distinct levels directly. The architecture underlines the increased mutual dependency of the building blocks. This also reflects the need to develop an integral strategy that adapts to the changing needs of the scientific community and technological developments.

6.2. Networks

The SURF organization together with the institutions for research and education, the government, and the industry partners have built SURFnet6, which is the current generation of the Dutch national research network, offering a state-of-the-art hybrid (IP and lambda-based) network. This network has been built through a joint effort in GigaPort and GigaPort NG Network projects. Figure 2 shows the network of optical lambdas reaching

Europe as of May 2008. SURFnet6 is widely recognized as a world leading network[17]. These developments have attracted top scientists and international companies and need to be continued at the current level.

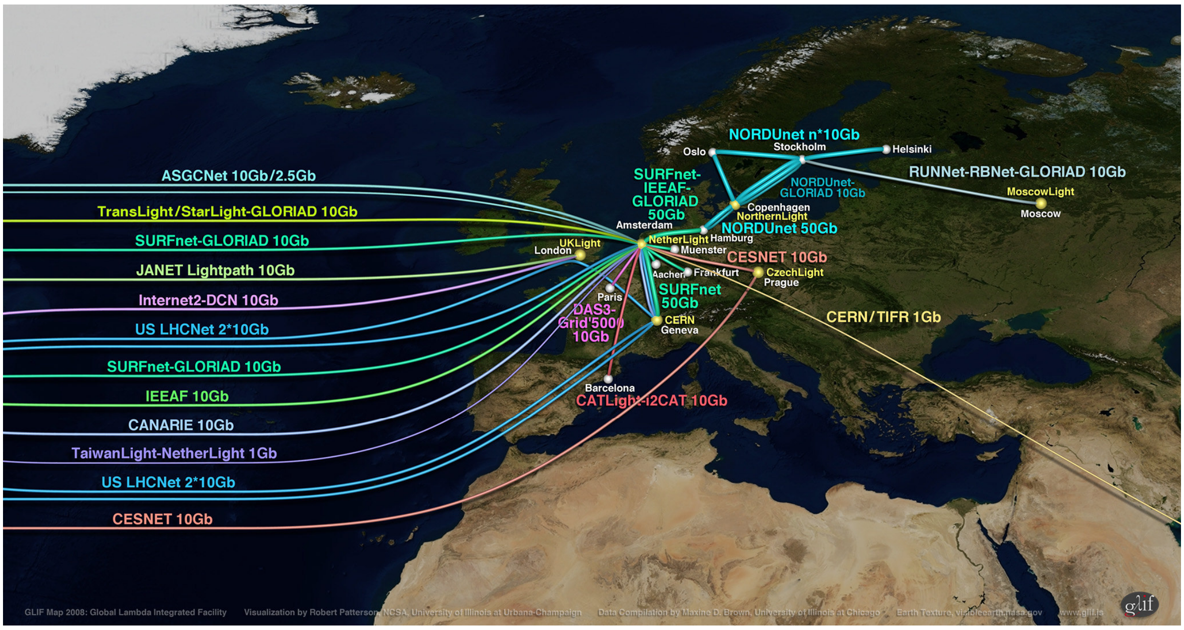

Figure 2 SURFnet's lightpath hub in Amsterdam

6.3. Resources

Data resources

Currently, both NCF and the BiG Grid project implement storage facilities and data access and processing facilities. LOFAR and the Large Hadron Collider (LHC) are examples of projects involving Dutch scientists in which vast amounts of data are (or will be) created. However, the trend of increased volumes of data to be stored and accessed increasingly applies to other disciplines. In life sciences for example, modern DNA sequencing equipment generates several TeraBytes per run, a data rate comparable to the LHC experiments. In humanities and social sciences, preservation of data and making them available for shared use is the very reason for the existence of DANS (Data Archiving and Networking Services). During the coming years, BiG Grid heavily invests in mass storage and cache heading for over 10 million GigaBytes (or 10 PetaBytes) of storage in its final year. The research facilities selected by the Committee *Roadmap Large-scale Research Facilities* (grades A and B) require collection, aggregating, processing and distributing of huge amounts of data. It may be expected that the data volume (from a wide range of scientific fields) continues to grow (exponentially), requiring continuous investments in storage capacity and resolving issues of accessibility and security. Grids and e-services are needed to retrieve, analyze and visualize data from distributed resources.

17 For instance, SURFnet contributed to the number one ranking of The Netherlands in the Earnest Foresight study. Refer to http://www.terena.org/activities/earnest/ws4/gfx/20070925-JS-geographic.pdf

Computing resources

Today, the national *capability supercomputer* is the Huygens system at SARA with a 60 Teraflop/s Power6-based IBM low-latency computer with a large amount of directly addressable memory. Latest in 2011, the Huygens needs to be replaced by a new Petaflop-range world-class supercomputer to serve the needs of the scientific community in the Netherlands alone[18]. The availability of the national supercomputer has strongly contributed to the development and international standing of whole disciplines in computational science, such as theoretical chemistry, computational fluid dynamics, climatology, catalysis and medical simulation. The Stella (IBM Blue Gene/P) at CIT of the University of Groningen is needed to do the real-time processing of the LOFAR sensor data. NCF and SARA also supported universities in providing *capacity* computing through the national compute cluster LISA. Other systems include the DAS-3 (Distributed ASCI Supercomputer) cluster, which is an initiative of the ASCI research school with participation of the universities in Amsterdam (UvA and VU), Leiden, Utrecht and Delft. The national capacity computing resources are further expanded upon by the BiG Grid project. The need for capacity computing cycles is exponentially growing and requires further continuous investments in new systems.

6.4. Grids

A grid refers to an infrastructure that enables the integrated, collaborative use of high-end computers, networks, databases, visualization facilities and scientific instruments owned and managed by multiple organizations. Grid applications often involve large amounts of data and/or computing and often require secure resource sharing across organizational boundaries, and are thus not easily handled by today's Internet and Web infrastructures. Since 2006, the BiG Grid project is heading for the development of a science-wide national grid infrastructure. The grid infrastructure is used by the Nikhef high energy physics community (linked as a Tier-1 site to the LHC at CERN[19]), by the astronomy community (LOFAR), by the life sciences community as represented by NBIC, by alpha science communities like DANS, and by computational sciences communities represented through NCF.

To a large extent, the maturing of the web and grid service technology has resulted in an international consensus on the software architecture which will be the basis for generic resource access. This concept of Service Oriented Architectures allows flexible mixing and interconnection of both web and grid services, and of higher-level services that are application-oriented. Standards bodies like the Open Grid Forum (OGF) are consolidating, and interoperability between various middleware ‘solutions’ (i.e. resource management software) is now high on the agenda. New approaches such as ‘cloud computing’ promise to lower the threshold for resource use, but need further study and development for the application in the scientific research domain.

The Netherlands is a leading player in the development of the grid (a.o. through the engineering and development activities of Nikhef), and has considerable expertise in bio-informatics, distributed sensors networks, and particle physics applications. The BiG Grid project[20] provides the Netherlands with the large-scale infrastructure to fully exploit this leading position. This puts the Netherlands at the vanguard of grid developments. The investment in both building knowledge and expertise and expending facilities needs to continue and possibly increased matched to the needs of the scientific community.

---

18 Het belang van High Performance Computing voor Nederland, December 2008
19 http://lcg.web.cern.ch/LCG/lhcgridfest/
20 NCF, BIG GRID, the Dutch e-Science Grid, October 13th, 2005

6.5. E-science

E-science refers to a worldwide development to bridge the gap between scientists in application domains and the developments of ICT. E-science comes in different flavours, but the common goal is to make the most efficient use of the very fast developing ICT infrastructure in all fields of science and research. The following three lines of action can be distinguished.

- The Virtual Laboratory line is developing software (middleware) providing generic e-science services and generic application-domain oriented services. Examples of generic e-services are workflow management, metadata management and knowledge extraction, content access and browsing, reasoning technologies, and security services. Examples of generic application-domain oriented services are found in areas like bioinformatics, medical informatics, and food informatics. The project involves researchers from public organizations (universities and other research organizations) and private organizations (Philips, Unilever, IBM, LogicaCMG, FEI).
- The computational science line is the field of study concerned with constructing mathematical models and numerical techniques and using powerful computers to analyze and solve complex scientific and engineering problems. The vision is to advance innovative, interdisciplinary research where complex multi-scale, multi-domain problems in science and engineering are solved on distributed systems, integrating sophisticated numerical methods, computation, data, networks, and novel devices. Computational science is very widely used in the Netherlands[21], in Europe[22] and the rest of the world. The Netherlands scores well in this field, with most cited scientists in chemistry and mathematics among the key users. Also the very development of methodologies in chemistry, water management and fluid dynamics takes place in The Netherlands.
- The third line of action is the human and digital interface to end users, including various facilities for visualization. Researchers with clear or partly articulated computational problems in their field need highly skilled personnel to understand and address their issues and provide real and practical solutions.

The Netherlands are in the forefront of the e-science development benefitting from the excellent grid and network in the Netherlands and supporting a wide range of scientific communities. The prototype services need to be further developed, both to provide better support for the research areas for which they were developed and to broaden their scope to include other research areas. In order to provide the scientific community with stable and reliable products, research activities need to be complemented by (industrial) software engineering development and operational support.

The Netherlands can build on top of a very strong position in all three lines of action. Further, scientific fields in social sciences and humanities expose a growing interest in building large-scale datasets and powerful computational analysis. The economic and scientific interest of for example the application-domain oriented e-services developed together with NBIC (genomics, ontologies for biomedical information, medical imaging)

---

[21] Computational e-Science: Studying complex systems in silico. A National Coordinated Initiative, P.M.A. Sloot; D. Frenkel; H.A. Van der Vorst; A. van Kampen; H.E. Bal; P. Klint; R.M.M. Mattheij; J. van Wijk; J. Schaye; H.-J. Langevelde; R.H. Bisseling; B. Smit; E. Valenteyn; H.J. Sips; J.B.T.M. Roerdink and K.G. Langedoen, Feb 2007

[22] Scientific Case for European Petascale Computing, European High Performance Computing Initiative, The Scientific Case for a European Super Computing Infrastructure, Petascale Computing in Europe, 2006

requires continuation of the e-science research activities. The Bsik VL-e consortium has expressed interest in doing so by setting up a center for e-science.

Finally, the growing attention for e-science and its importance has led the AWT to write a report, that will be published in January 2009. From communication with the editors, it seems that there is consistency between the recommendations of the AWT and ICTRegie, the latter focusing on the infrastructure needed for e-science development and applications.

6.6. System management

System management refers to housing the equipment (hardware and software) in a suitable safe and secure environment, supplying the electricity, cooling the generated heat, and in general ensuring the correct operation of the facilities. Moreover it includes the necessary support and development activities for users in the scientific community. Today, SARA hosts among others the national supercomputer Huygens, the national compute cluster LISA, parts of the national academic network SURFnet6, the GLIF European node NetherLight, Kennisnet, AMS-IX, BiG Grid, and the LHC Tier-1. SARA hosts grid based data storage for various projects, including LOFAR, LHC Tier1 (NIKHEF), VL-e, NBIC activities and for Cognitive Sciences. SARA provides the mass storage and archiving facilities for VL-e and BiG Grid.

SARA is not only providing specific services for the scientific community, but is also conducting commercial activities, including co-location and hosting, ASP and database services, on-demand services, and network design and operations. Customers include regional education and medical networks, VNU, Royal Library (KB), and companies like KPN, BT, and Google. SARA has decided to bring its commercial activities in a separate limited company called Vancis. SARA will focus on the advanced ICT infrastructures management, support of computational science and e-science, and application development activities. This separation provides transparency to the customers and diminishes difficulties due to differences in activity types (task organization versus commercial organization).

Today, SARA was selected to manage SURFnet6 after a European tendering process setup by SURFnet. Other contributions of SARA to the ICT infrastructure are based on participation in projects of limited duration. The inherent uncertainty hinders SARA in doing long term investments in the system management, support and development for the research infrastructure. However, investments are needed to create a reliable and stable taskforce with talent and expertise and a high-end data center to provide state-of-the-art computing infrastructure services as well as to store the growing amount of scientific data (originating from unique and expensive experiments) and handle sensitive data (e.g., originating from medical files). A centralized data center for advanced system management functions that are not available on the commercial market requires long term investments in housing, experienced and qualified personnel, and research and development.

6.7. Current organization and funding

Dutch Roadmap Large-scale Research Facilities

Since many years, funding for the ICT infrastructure is mainly coming from dedicated programs for large-scale research infrastructures (NWO, BiG Grid) and from scientific impulse programs (Bsik). In programs for research infrastructures, i.e., the Dutch Roadmap Large-scale Research Facilities, proposals are evaluated using a set of well-defined criteria[23]. Parties involved with providing ICT infrastructures are capable of fulfilling the criteria

[23] See the committee Dutch Roadmap Large-scale Research Facilities, October 2008

for project proposals, and they are indeed large-scale research facilities. However, the relations between the criteria and the (expected) results of creating excellent ICT infrastructures are more complicated than with other facilities. The Committee *Dutch Roadmap Large-scale Research Facilities* concludes that the ICT infrastructure requires accommodating policy[24]. The Committee also rightly states that the ICT infrastructure is a necessary precondition for all large-scale research facilities. The committee therefore concludes that continuation of SURFnet6 and further developments in the field of e-science are essential.

Current organization structures complicate infrastructure provision

The ICT infrastructure for scientific research is currently provided by the organizations SURF, NCF and SARA in cooperation with projects conducted within the Bsik program (GigaPort NG, VL-e, LOFAR) and NWO programs (BiG Grid, STARE, GLANCE, VIEW). This situation is the result of a historical process where visionary and ambitious people every time saw opportunities to develop and build the components of the excellent ICT infrastructure that the Netherlands is provided with today. Despite these good results, the current fragmented organizational situation complicates the coordinated development of the ICT infrastructure, in which more than ever the relation between the constituting elements defines the overall capabilities. The autonomy of the collective organizations results today in separate coordination activities, which impair flexibility and extend the timelines.

Infrastructure requires long term policies

Although the economic lifetime of ICT systems is not much longer than a few years, the ICT infrastructure for scientific research includes many elements that go well beyond the lifespan of a typical project. This includes the knowledge and experience to build the infrastructure and such aspects as housing and (increasingly important) energy supply. The relatively short lifetime of systems force the use of just-in-time procurement mechanisms, building another tension between the period of enjoying the profits of investments and bearing responsibility for them. Project funding complicates the strategic decision making process that is needed here.

There is not one locus of responsibility

The different organizations and projects all strive for the best possible ICT infrastructure, but they also have a partially overlapping set of responsibilities for the ICT infrastructure. There is a separation of concerns, and many of the services that are required by scientists require activities to be distributed over organizations. Furthermore, these activities are carried out by organizations with a different structure, financing model and lifespan. Even when the organizations themselves are able to decide who should do what, funding parties that are interested in the aggregate results have difficulties in locating and pinpointing an actual locus of responsibility.

Multiple programs complicate alignment and synchronization

Finally, the projects that are granted can only offer minimal assurance on to what extent they use the same criteria, adhere to the same goals and judge the same about the need for a project proposal. While the organizations synchronize their proposals to a reasonable extent, this is even harder to do for granting organizations.

[24] Dutch Roadmap Large-scale Research Facilities, October 2008, pg 24

6.8. Users of the ICT Infrastructure

SARA and NCF have recently published 42 interviews with professors at Dutch universities in a large variety of scientific fields using the Huygens supercomputer. Each of them points out the importance of high performance computing for his or her field, the societal relevance and impact and their actual Petascale challenge.[25] More than 100 research groups rely on the use of the national supercomputer Huygens and the national compute cluster LISA, and publish high quality papers in high-ranked scientific journals. Every Dutch university and research center has an excellent connectivity through SURFnet6. Many important scientific disciplines like high energy physics, astronomy and life sciences (NBIC) have early adapted grid technologies and are now using the BiG Grid infrastructure. Other disciplines like cognitive sciences, medical sciences and media are growing.
The use of the ICT infrastructure in the social sciences and humanities is emerging. The Virtual Knowledge Studio KNAW aims to support researchers from these disciplines in the Netherlands in the creation of new scholarly practices, and to stimulate reflection on e-research. A core feature of the VKS is the integration of analysis and design in a close cooperation between social scientists, humanities researchers, information technology experts and information scientists. This integrated approach should provide insights in the e-research that can contribute to new research questions and methods in the humanities and social sciences. The VKS collaborates with the Erasmus University Rotterdam in the Erasmus Virtual Knowledge Studio KNAW (in short: Erasmus Studio) based in Rotterdam and with the Maastricht University in the Maastricht Virtual Knowledge Studio KNAW based in Maastricht.[26]

6.9. Relevance for industry

Apart from the direct interest of a few high-end and large industrial research entities (like Philips and DSM), industry is most interested in the expertise and knowledge generated by the availability of an advanced ICT infrastructure. It allows them to apply and use this expertise in their own research labs and develop their own infrastructures (e.g. Shell). The ICT industry (equipment, software and services) has a similar interest in the growing base of expertise, but also in the opportunity to contribute to innovative developments and the participation in pilots and test beds.

Considering this industry position, it seems that there are no realistic opportunities for Public Private Partnerships in operating the proposed advanced infrastructure. For the development activities, the proven model of public tenders, as profitably and successfully used by GigaPort/SURFnet should be continued, as well as the industry cooperation in research consortia like in VL-e.

[25] Het belang van High Performance Computing voor Nederland, December 2008
[26] http://www.virtualknowledgestudio.nl/

7. Recommendations

As motivated in the earlier sections, achieving the desired top position as a knowledge economy asks for competing on scientific and economic markets. Maintaining the position of the Netherlands in the top of science requires permanent investments and developments in computers, databases and tools of the ICT infrastructure.

Today, the Netherlands has a very good starting point: a world leading position in the field of communication networks, a competitive position in supercomputing (Huygens) and leading edge knowledge development in de field of grid technologies and e-science. The Netherlands can use this position to further strengthen it in order to attract talented scientists pushing the limits even further and to attract international companies for locating their research and development in the Netherlands.

To make this possible, it is required to increase the coherence between organizations, activities and the components of the ICT infrastructure. The ICT infrastructure must be used by a wide range of scientific fields, including alpha and gamma sciences. The activities of the VL-e project need to be continued and extended. The funding mechanisms must be changed, and become independent of existing mechanisms for investment and research projects. The ambition is to bring all components of the full ICT infrastructure at the level of the current world leading position for communication services. To achieve this goal, ICTRegie has formulated the following recommendations.

7.1. Governance of the ICT infrastructure

Position all ICT infrastructure development and operations under the SURF umbrella.

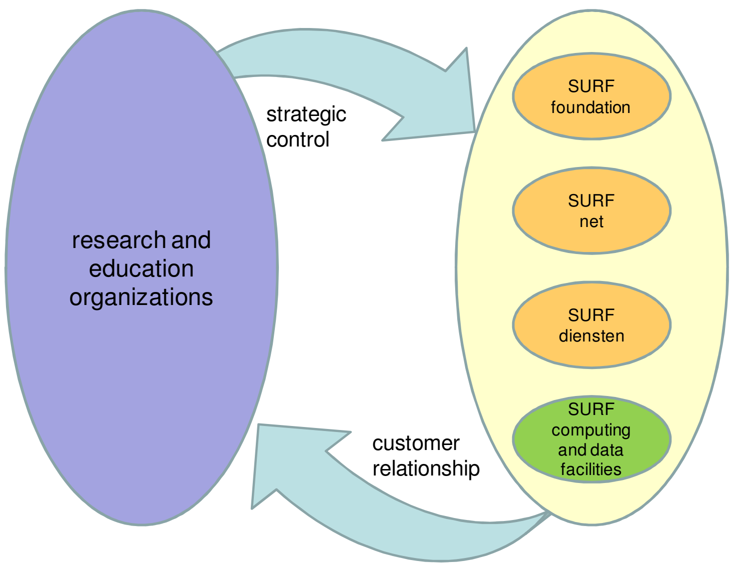

Figure 3 Proposed SURF organization

In SURF a new entity is envisaged. This entity could be named Computing and Data Facilities (in Dutch *Rekenfaciliteiten*, referring to the original meaning of the name of SURF…). This entity is responsible for the development, operations and support of the e-science services, grids, and resources (computing resources, data and visualization resources). This relates to activities currently executed by NCF, SARA and part of the activities conducted by the projects VL-e and BiG Grid. The research activities related to e-science and grids are positioned within a new organization for e-science research (e-science Research Center), which will be discussed in Section 7.3.

- **NCF**: The foresight, planning, standardization and international representation activities of NCF will be located as a center of expertise in SURF Computing and Data Facilities. The activities related to granting computing time on the supercomputer and computer clusters will remain with NWO. SURF Computing and Data Facilities and NWO establish an agreement that NWO continues to allocate computing time for supercomputing, based on submitted proposals and a peer review evaluation process. The current financing of the national computing facilities by NWO is transferred to SURF. A proposal for using supercomputing and cluster computing facilities may be part or is accompanied by a research proposal submitted to one of the NWO research programs.
- **SARA**: The research support activities of SARA will become part of SURF Computing and Data Facilities. This includes the system management of (part of the) infrastructure building blocks, engineering and development. This part includes the management of pre-competitive building blocks, not available on the commercial market. This includes also advanced support to members of the scientific community to take the best possible advantage of the ICT infrastructure. SURF Computing and Data Facilities is the national e-science support center (public organization) financed by the ICT infrastructure fund.
  The new structure creates improved conditions for long term investments in reliable and stable housing facilities, training and education of qualified personnel and participation in (national and international) research programs. This organization will also act as center of expertise on management of advanced systems and provide consultancy and advanced support to research organization in the Netherlands.
  The merger of SARA parts and SURF requires a decision in principle followed by a due diligence, consultations and consent of all involved parties.

Governance structure

The key driver for this recommendation is to create more coherence, which is achieved by positioning the relevant existing organizations within the SURF organization. The general board of SURF includes representatives of all organizations using the national research and education network in the Netherlands. It is recommended to form a *strategic board for policy decisions on the ICT research infrastructure* within the SURF organization. These board members are the key stakeholders for the scientific research conducted in their organizations. This board advises on centralization or decentralization of ICT infrastructure facilities, investment plans, developments etc. The board members are also responsible for aligning ICT investments at the national and institutional level.

The effectiveness of the governance structure and innovation model applied by SURF and SURFnet has been proven in practice. Characteristics of the SURF's governance process are a formalized periodical process of strategic planning with wide participation of users and their local organizations leading to long term support at the board level. This provides stability and continuity of the governance at all levels with sufficient distance between the user role and strategic control. Typical is also the clear financial distinction between innovation and exploitation and the absence of other organizations, at board level e.g., the ICT industry[27]. The ICT industry is involved through innovative tendering processes.

---

[27] Innovatie bekeken, bekeken innovatie, SURF innovatie in de praktijk bij SURFnet, K. Neggers, F. Van Iersel, Maart 2006.

Integrating the NCF and SARA functions within SURF means a significant simplification of the governance.

Development and Operations

The main responsibility of SURF Computing and Data Facilities is the development and operations of the computational and storage infrastructure. SURF employs high qualified engineers capable of designing and developing production systems in close co-operation with suppliers on the commercial market. Besides networks, also software systems for e-science and grid functions need to be developed in close operation with industrial parties. SURF Computing and Data Facilities employs qualified personnel that support members of the scientific community in using the infrastructure. SURF as an organization can implement a one stop shop supporting users for all services provided by the entire ICT infrastructure.

Contract research for external parties

The organization of innovative research activities are implemented as it is organized for SURFnet. SURF is responsible for defining the innovative research programs needed to innovate on the ICT infrastructure. This means that the present practice of contracting scientific organizations to do relevant applied research will be continued. Part of that will be done by the e-science Research Center, based on a long term program that is funded by SURF, in a similar model as the government is funding the long term research at TNO.

7.2. Funding of the ICT infrastructure

Provide an annual budget from a dedicated fund to cover expenses for permanent innovation, development and operation of the required ICT infrastructure for scientific research.

It is strongly recommended to finance the continuous development and innovation of the infrastructure (mainly driven by very fast advancement of information and communication technologies) through a dedicated fund. Then, ICT Infrastructure investment proposals do not have to compete with research proposals that are for example submitted for R&D subsidies from FES. This is already the case with the supercomputing facilities, which are financed through a dedicated facilities budget of NWO. However, the volume of this fund is too small to accommodate the growing needs for both supercomputing and cluster computing as well as the growing needs for data storage and data retrieval. An annual budget takes away the uncertainty created by impulse money for infrastructure, like in the BiG Grid project. When such impulse money is spent, it is very uncertain to what extent the built infrastructure and expertise can be further developed and maintained.

A dedicated fund enables the new organizational entity to develop a long term strategy providing the means to balance investments across the various components of the infrastructure and the expert team. The long term strategy and policies are developed by SURF and are subjected to decision making processes by universities and research organizations, which are all represented in the board.

In addition to the proposed central funding for network innovation, connected organizations contribute for € 19 million annually for network usage covering the operational costs.

Because of the very different needs and use patterns for computing and storage facilities over time, as compared with network provisioning, it is not considered realistically possible to have the operational costs for these facilities billed in a similar way as for the network. Therefore a lump sum central financing is proposed for innovation and operations, similar as is in existence today with NWO/NCF.

7.3. E-science Research Center

Establish an e-science Research Center in cooperation with universities, NWO and NWO institutes, conducting e-science research supporting innovations of the ICT infrastructure and applications for science.

ICTRegie recommends founding an e-science Research Center (e-SRC). This is a center that bundles e-science research in the three lines of actions described in 6.5, i.e., virtual laboratories, scientific computing and human and digital interface to end users, including visualization. This center creates a platform for multidisciplinary research by research groups from various universities. The e-SRC should be founded through a cooperation agreement between these parties and NWO and should be located closely to the data center hosting systems of the ICT Infrastructure.

Development through push and pull

Figure 4 clarifies the relationship between the e-SRC and SURF Computing and Data Facilities. The e-SRC is a Research Center conducting scientific research in close cooperation with research groups, including computer science, from various universities. The e-SRC is the multidisciplinary platform for further co-operation on e-science research. E-science research within e-SRC includes research and development of prototype tools, similar to VL-e. Also, e-SRC itself is an advanced user of the infrastructure. Hence, there is direct coupling of services provided by the infrastructure and needs from the scientific community.
When sufficiently mature, prototype tools of e-SRC are transferred to the SURF Computing and Data Facilities organization and integrated in the (software) development of reliable (production) systems that can be used by the scientific community. The e-SRC bridges the gap between scientific research and development and operation of e-science tools.

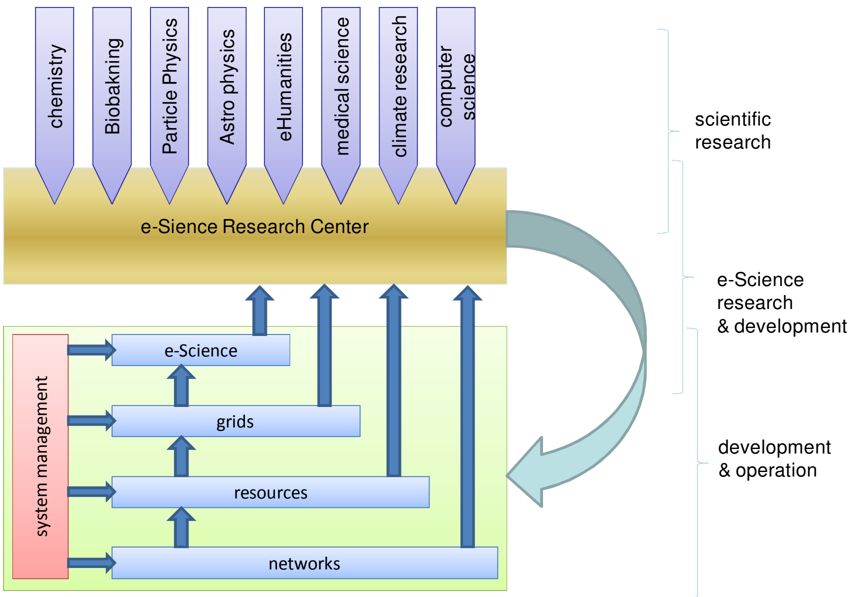

Figure 4 The E-science Research Center and ICT infrastructure for scientific research

Multidisciplinary expertise

The development of multidisciplinary expertise is extremely important. The cooperation of computer science and life sciences, with groups working on biodiversity, iobanking, metabolomics, biomolecular informatics and cognitive sciences need to be continued and intensified. Other relevant fields include climate induced weather extremes and services for the humanities. Several data repositories are currently in use in the humanities via initiatives of KNAW's Virtual Knowledge Studio and DANS.

Funding

Fundamental research for e-SRC is partly funded through SURF from the ICT infrastructure fund, through a long term strategic research program (similar to the government funding of TNO). Additional funding will be needed by NWO, universities and government agencies (competitive, peer reviewed programs). Strategic as well as applied research can be financed. Other funding will have to be found in the usual national and international sources, as well as in joint projects with application disciplines (like in life sciences or physics).

7.4. The European ecosystem

Acquire a key position in the innovation, development and operation of a pan-European ICT infrastructure for collaborative scientific research.

ICTRegie recommends combining the above mentioned actions with the ambition to acquire a key position in the innovation, development and operation of a pan-European ICT infrastructure for collaborative scientific research.

There are unprecedented opportunities for the Netherlands. The Netherlands can play an important role in hosting the EGI.org: the head office of the European Grid Initiative. Even more important, the Netherlands is Principal Partner of PRACE, together with the leading countries Germany, UK, France and Spain. PRACE wants to build a world-class pan-European high performance computing centre. The pan-European HPC service will be a part of the European Research Area, in which the Seventh Framework Program is prepared to invest hundreds of millions of Euros[28].

The key position is to be achieved by building a European data center node providing large-scale data, computing, networking, and visualization services. The successor of Huygens and expansion of capacity computing power has to be aligned with this ambition. The data center node is connected to the rest of the world through a grid using excellent network and exchange facilities. The proposed key position contributes to the Dutch scientific output and attracts talented researchers, foreign R&D investments and international companies. The profile of the contribution of the Netherlands to the pan-European ecosystems further depends on the consultation and negotiation with the Prime Partners in PRACE. The decision making in the Netherlands can be evaluated using the criteria of the Dutch Roadmap Large-scale Research Facilities.

7.5. Next steps

After a basic decision has been taken on the recommendations, it is advised to charge SURF with the implementation, in cooperation with the other organizations involved. To monitor and supervise the implementation, a special steering group with an independent chairman should be set up. Special attention is needed for the restructuring of SARA in a public and a private part, for the positioning of NCF, for the creation of the proposed e-science Research Center (e-SRC) and for the reallocation of existing funds.

For the financing of research programs of the e-SRC beyond the proposed funding through SURF, a proposal needs to be included in the ICT Theme for the 2009 FES round. ICTRegie takes that into account.

[28] http://www.prace-project.eu/

As for the timing, the aim should be to have the new structure in place before the end of 2009. That means a principle decision should be made early 2009, which is also needed to safeguard the Principal Partnership position of the Netherlands in PRACE.

8. Financial consequences

The financial consequences of the recommendations are based on an analysis of the average investments made during the last five years and investment proposals submitted during 2008. The current innovation of the ICT infrastructure is financed by NCF (through direct funds from NWO for capability computing, capacity computing, storage) and through funds from dedicated projects like GigaPort Next Generation (network innovation), VL-e (e-science), BiG Grid (storage, capacity computing) and university funding (capacity computing). SARA receives funding from the aforementioned organizations and projects. SARA invests profit from private activities in the ICT infrastructure for scientific research.

The budget to implement the recommendations for the coming 10 years requires a fund of € 62.5 million annually (2008 price level). A carefully orchestrated replacement excludes the need for further impulse financing. This is a significant improvement on the mentioned € 100 million in the earlier report “Goed gereedschap is het halve werk”. Table 1 identifies the immediate needs starting in 2009, except for the budget related to the European key position, which is needed earliest in 2011. Part of the needs for the next few years are covered by reservations currently made, e.g., for Gigaport/SURF, NCF, VL-e and BiG Grid.

	R&D	innovation	operation
e-science	1.5	2.0	0.5
grid	1.5	1.5	0.5
storage		3.0	1.5
capacity		2.5	2.5
capability		6.0	4.0
network	1.0	9.0	
system management	0.5	0.5	2.5
support			2.0
	4.5	24.5	13.5

	R&D	innovation	operation
European key position		12.0	8.0
		12.0	8.0

Table 1 Annual budget for the ICT infrastructure (in € million)

The budget for the ICT infrastructure will be governed by the SURF organization. The budget covers costs for R&D, innovation, and operation. SURF will use the R&D budget for contract and program research that will be conducted by the e-SRC and by universities and other research organizations. The innovation budget is spent on hardware investments (annual depreciation) and software purchases (licenses) and costs for engineering by SURF engineers, together with industry where possible. This budget includes also activities such as software hardening and production readiness for services. The operation budget is used to cover expenses like hosting, power and maintenance, operations, operational support and maintenance.

The budget is subdivided into amounts for each of the previously introduced building blocks of the ICT infrastructure and budget for scientific support.

E-science and grid

The e-science and grid R&D budget of € 3.0 million (€ 1.5 + € 1.5 million) is intended for research that supports the innovation of the ICT infrastructure. This R&D is directed towards Virtual Laboratories, computational science and interfaces (virtualization). SURF will commission this research to the e-science Research Center, based on a joint long term research program. The e-science and grid innovation budget of € 3.5 million (€ 2.0 + € 1.5 million) is intended for conversion of prototypes into reliable services, together with the software industry. The e-science and grid operational budget of € 1.0 (€ 0.5 + € 0.5 million) is intended for activities including software maintenance, software operations and operational support.

Storage

The storage innovation budget is € 3 million covering mass storage (tape storage and fast disk) and cache storage. The budget is based on careful analysis of the growing need and based on experienced gained from the BiG Grid project. The storage operation budget is € 1.5 million. In practice, the ratio between investment and operation turns out to be 2:1.

Capacity and capability computing

The capacity computing innovation budget is € 2.5 million per year. The budget is based on careful analysis of the growing need and based on experienced gained from the BiG Grid project. The operation budget equals € 2.5 million. In practice the ratio investments and operational costs turns out to be 1:1.

The capability computing innovation budget is € 6 million and is a continuation of the current policy to provide the community with a supercomputer at the level of Huygens. It is expected that the operational costs will increase due to growing power requirements. The required operational budget for the successors of Huygens equals € 4 million.

Network

The network R&D budget is € 1.0 million and includes research contributing to the innovation of the network. The network innovation budget equals € 9.0 million. The R&D and innovation budgets are a continuation of the current policy. During the last five years, € 8 million originated from Bsik projects and € 2 million was funded by the Ministry of OCW (Surfworks)[29]. The network innovation includes engineering, innovative procurement and investment in hardware. Innovative procurement enables collaboration with industry. The operation of the network is being paid for by the organizations connected to SURF, including organizations for education and research. These operational costs amount to € 19 million annually and are not included in Table 1.

System management

The management R&D budget is € 0.5 million. This budget covers costs for developing knowledge and experience for management of pre-competitive systems for scientific research. The management innovation budget is € 0.5 million. This budget covers collaborative engineering for system management. Today, these costs are covered by profits originating in private activities of SARA, which will not be the case in the proposed new organization. The management operation budget is € 2.5 million annually, needed to cover the depreciation and generic running costs for a data center (reliable and secure housing and facilities related to energy and cooling). Other operational costs are covered by operational budgets mentioned elsewhere in the table.

---

[29] SURFnet also receives € 2.0 million annually for innovations related to Kennisnet (SURFnet/Kennisnet). These costs are beyond the scope of the ICT infrastructure for scientific research.

Support
An operational budget of € 2.0 million is needed to finance a team of 15 to 20 experts in various application fields supporting the members of the scientific community.

European key position
The Netherlands has the opportunity to take a key position in Europe. The key position will strengthen the position of the Dutch scientific community in Europe and in the world by sharing and exchanging research facilities all over the globe and create a hub of attraction for e-science related R&D. This collaboration is based on a balanced “bringing and taking” with Principal Partners in the PRACE consortium. It will bring the Netherlands at the top-5 in Europe. An investment of € 20 million annually from the Netherlands is required to acquire a top-5 position in Europe. Other countries as well as the EU will provide a multiple of this amount from which the Netherlands directly benefits.

9. Appendix A

Throughout the process, desk research has been performed, covering policy documents from the Ministry of Economic Affairs (OCW, EZ), CBS, CPB, het Innovatieplatform and ICTRegie. Desk research also covered Dutch and European policy documents from various organizations. Further, policy documents from the United States and a range of European countries have been studied. In addition, workshops and many interviews with involved organizations and persons have been conducted, both at executive and board level. ICTRegie has achieved consensus on the content of this advice. Although ICTRegie remains responsible, it would like to express its gratitude for information, feedback and help received from many persons, including

- dr. P.J.C. Aerts
- prof. dr. H.M. Amman
- drs. J.P. Barthel
- drs. H.W. te Beest
- dr. ir. A. Berg
- dr. ir. H. Dijkman
- drs. P.W. Doop
- dr. G.J. Heimeriks
- prof. dr. L.O. Hertzberger
- ir. P.P. 't Hoen
- prof. dr. F.M.G. de Jong
- dr. M.J. Koornstra
- dr. ir. C.Th.A.M. de Laat
- prof. dr. W.B.G. Liebrand
- dr. G.R. Meijer
- dr. ir. P.H. Michielse
- ir. C.A.M. Neggers
- dr. S.J. Noorda
- prof. dr. ir. G. van Oortmerssen
- dr. ir. A. Osseyran
- drs. A.J. van Rijn
- prof. dr. C.G.M. Sterks
- prof. dr. J. de Vlieg
- prof. dr. W.B.G. Liebrand

10. Appendix B

An overview of the abbreviations in this rapport is presented below.

AMS-IX	Amsterdam Internet Exchange
ASCI	Advanced School for Computing and Imaging
AWT	Advisory Council for Science and Technology Policy
BiG Grid	the Dutch e-Science Grid
Bsik	Besluit Subsidies Investeringen Kennisinfrastructuur
CERN	European Organization for Nuclear Research
CIT	Donald Smits Centre for Information Technology at the University of Groningen
CWI	Centrum voor Wiskunde & Informatica
DANS	Data Archiving and Networked Services
DAS	Distributed ASCI Supercomputer
DEISA	Distributed European Infrastructure for Supercomputing Applications
EBI	European Bioinformatics Institute
EGEE	Enabling Grids for E-sciencE
EGI	European Grid Initiative
e-IRG	e-Infrastructure Reflection Group
ESFRI	European Strategy Forum on Research Infrastructures
FES	Economic Structure Enhancing Fund
GLANCE	Global Computer Science
GLIF	Global Lambda Integrated Facility
HPC	High Performance Computing
KNAW	Royal Netherlands Academy of Arts and Sciences
LHC	Large Hadron Collider
LOFAR	Low Frequency Array
NBIC	Netherlands Bioinformatics Centre
NCF	Netherlands National Computing Facilities foundation
NIKHEF	National Institute for Subatomic Physics
NWO	The Netherlands Organization for Scientific Research
OCW	Ministry of Education, Culture and Science
OGF	Open Grid Forum
PRACE	Partnership for Advanced Computing in Europe
SARA	Stichting Academisch Rekencentrum Amsterdam
STARE	STAR E-science
TERENA	Trans-European Research and Education Networking Association
VIEW	Visual Interactive Effective Worlds
VKS	Virtual Knowledge Studio
VL-e	Virtual Laboratory for e-science