CSCS STRATEGY

2004 – 2010

ACCELERATING SCIENTIFIC DISCOVERY

October 30th 2003

Taskforce CSCS of the ETH Board

Table of contents

1. Summary

As a world leader in scientific research and education, the Swiss academic community must ensure that its research groups working at the frontiers of numerical sciences have access to the necessary resources in the best possible way. In this view, the mission of the national centre for scientific computing, CSCS, is to set and manage an appropriate framework for the support, the planning, the coordination and the promotion of scientific computing for numerical sciences; this framework will also serve as a stimulant and as the integrative house for interdisciplinary collaborations of the field. As such, it contributes significantly to the effort of acquiring the necessary momentum for sustaining growth for the benefit of Switzerland, targeting at the acceleration of scientific production.

This document describes the strategy proposed for the years 2004-2010. The objectives can be summarized as

- support of excellence in research, in particular for the production and development of knowledge at the frontiers of science
- resources coordination of the perimeter: how to guarantee flexible investment for new projects at a sustainable cost
- implementation of the roadmap evolution of the HPCN systems
- diffusion among the scientific communities (corporations, academic, individuals, foundations) the information and communication programmes
- excellent communication platform between the CSCS and the rest of the Swiss research community.

The objectives lead to the three strategic goals of CSCS until 2010:

- Goal 1: development of thrust areas of scientific computing: High Performance Computing and Networking (HPCN) methods and software

- Goal 2: operations of complementary HPCN computer architectures, at least one of them being an internationally leading system

- Goal 3: positioning the centre both nationally and internationally by partnering with key players

2. Introduction

2.1. Trends in Scientific Computing

Several trends can be observed in scientific high-end computing:

1. the trend of supercomputing centres becoming additionally "*superdata centres*"
2. the increasing necessity to visualize massive amounts of data
3. the trend towards networked computing

On the application side, we observe:

1. new research fields becoming HPCN customers like biology and medical sciences. In contrast to classical application fields such as physics, engineering, and chemistry, these fields do not have a tradition in mathematical modelling and programming, and therefore, these customers will benefit of intensive dissemination of knowledge and education in scientific computing
2. the growing demand for new algorithms in modelling and statistics for these new application fields
3. some research fields like e.g. climatology develop so-called community models as the common software framework for modelling and simulation.

2.2 Present situation at CSCS

The value of a scientific computing centre is made up by two elements: firstly, the available infrastructure or technology which is oriented towards capability computing instead of capacity computing like in other computing centres, and secondly, the unique competencies of the people at the centre. The departure of the group of the former director Prof. Parrinello, deprived CSCS of a big share of these competencies, namely the application field or computational sciences field.

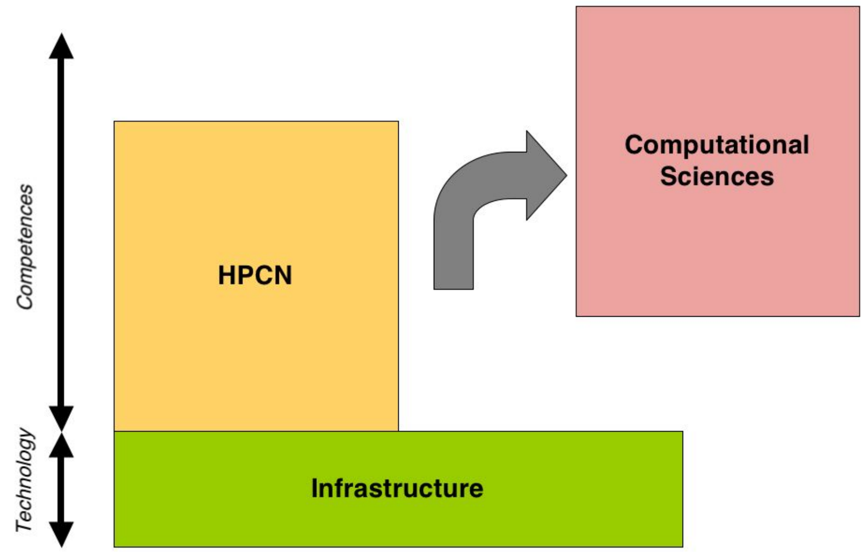

Figure 1: Situation of CSCS after the departure of Prof. Parrinello and his group

CSCS is not viable on the basis of HPCN operations alone. Whereas it was still sufficient in the 1980s to operate powerful calculators only, high-performance computing centres of the 21st century have to offer scientific competencies and flexible services in order to be attractive.

Two scientific disciplines are deeply connected with CSCS' core business: computational sciences and scientific computing. For computational sciences, i.e. research by numerical applications in natural sciences and engineering, one can assume that it would be extremely difficult to persuade a research group to leave its academic environment and to settle at CSCS. Such an attempt would require very special circumstances beyond scientific reasoning. Furthermore, from the view of other users, CSCS would become biased towards a specific discipline or user group. It will therefore concentrate on the sector of *scientific computing*, i.e. on the development of methods, algorithms, and software tools.

3. Business model

We propose a new business model for CSCS which enhances its scientific standing, fosters a mutual enrichment of HPCN operations and science sections and allows CSCS to grow by itself. This will enable CSCS to fulfil its mission statement as formulated by the ETH Board:

> *CSCS is the Swiss National Supercomputing Centre, providing, developing and promoting technical and scientific services for the Swiss research community on the fields of high-performance and high-throughput computing. It is a centre of competence that pioneers new information technologies; collaborates with domestic and foreign researchers, and carries out its own research and development in computational sciences and scientific computing. CSCS is a scientific and service centre and quality must be the overriding criteria guiding all its choices.*

The future business model for CSCS allows the centre to act as an enabler and a leading house rather than being a service unit that accepts orders from outside customers. The model rests on three pillars, the three strategic goals of CSCS until 2010:

Goal 1: development of thrust areas of scientific computing: High Performance Computing and Networking (HPCN) methods and software

Model: CSCS will develop a very strong branch in scientific computing. It will carry out its own research in HPC methods and software development. CSCS will devise these competences in specific thrust areas which it deems to be promising in the near and medium-term future (see chapter 3.1). The development of these thrust areas will be based on applications in strategically important research fields (see chapter 3.2). Establishing the necessary human resources and know-how in scientific computing will let CSCS bridge between applications in computational sciences and technology. This way, CSCS will gain attractiveness for national and international top researchers and it will strengthen the interconnection between HPCN and scientific services.

Goal 2: operations of complementary HPCN computer architectures, at least one of them being an internationally leading system

Model: CSCS will continue to operate in Manno different HPC computer architectures, at least one of them being an internationally leading system in the upper quarter of a widely accepted global ranking such as the TOP500 list. The selection of these machines will take place on the basis of user requirements in existing and to-be-developed application fields, and on the basis of technology forecasting and proactive benchmarking. CSCS will not be dedicated to specific HPC architectures or manufacturers. The resources on the top systems are allocated on a competitive basis to scientists with research projects of the highest quality.

CSCS will reinforce its position as the national centre coordinating the development of HPC architectures and solutions in the ETH Domain, and will enable shared computing as Load Manager.

Goal 3: positioning the centre both nationally and internationally by partnering with key players

Model: CSCS will actively network in the national and international context. It will offer scientific computing services for collaborations with and between academia, government, and industry. CSCS will lead partnerships and will enable incubators for the exploitation of HPCN.

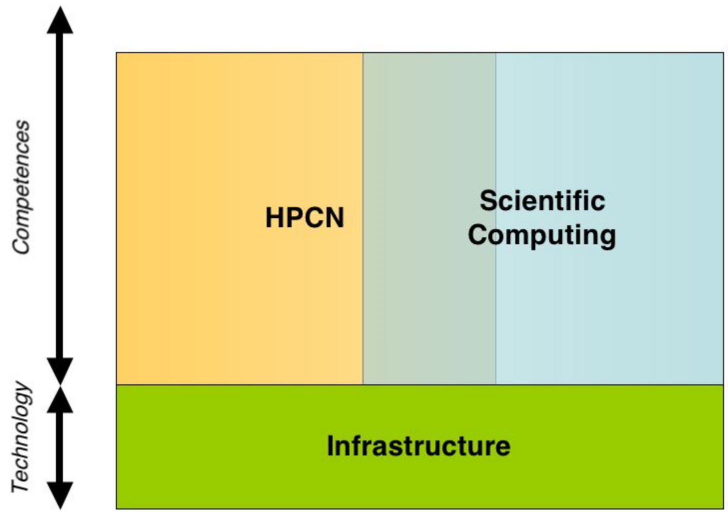

Figure 2: The future building blocks of CSCS

3.1. Thrust areas

The pace, at which the different areas of scientific computing (Goal 1) will develop, is controlled by two parameters:

- Firstly, it depends on the needs and requirements of the user community, not only the current, but also new customer groups to be acquired in the near future.

- Secondly, CSCS will initially focus on domains of high strategic value, either by supporting one of the strategic directions of the ETH domain, by accelerating the transfer of knowledge and methods to other application fields, or by giving CSCS an instant visibility. Early wins in 2004 and 2005 are crucial for CSCS.

We therefore propose that CSCS delves into these topics and covers these following disciplines of scientific computing in the future:

Discipline	Description
Visualization	Imaging, virtual reality,
Optimisation, Numerics & Benchmarking	Optimisation, selection and development of solvers and other mathematical methods
Data Intensive Computing	Data mining, knowledge management, information management and retrieval, machine learning, geographical information systems
Distributed Computing	Grid computing, networking, communications, agents, distributed services
Education & Training	Organisation of courses and summer schools, visiting scientists
Modelling Framework Support	Hosting and fostering a software framework that is being used as the community model of a research discipline

Table 1: The six thrust areas in scientific computing for CSCS

The different areas of scientific computing and the different application fields of computational sciences form the "texture of high-performance computing" (Figure 3). This texture is the basic architecture on which we can build the new CSCS. The different rows of scientific computing are building blocks placed one on the other with the HPCN hard- and software infrastructure forming the basis of the centre. On this basis, user services are developed. These services themselves are the building blocks of the toolboxes needed by the different scientific computing fields as listed in Table 1. The different application fields vertically integrate these toolboxes: the same set or different toolboxes may be used by different fields, thus cross-fertilising the various scientific computing disciplines.

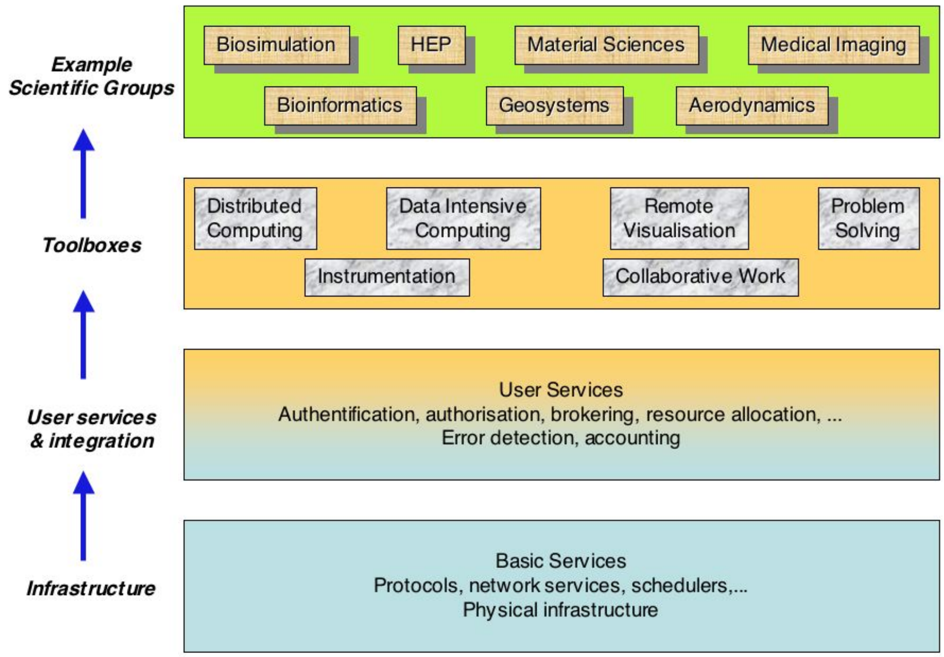

Figure 3: The texture of high-performance scientific computing

According to the proposed plan, as a rough estimate, 30 people will be involved in research and development of scientific computing and will offer scientific services in these fields. About 15 FTE will be allocated to HPCN operations. 4 to 5 positions will be assigned to management and administrative work, thus keeping the overhead of CSCS as low as 8 to 10%.

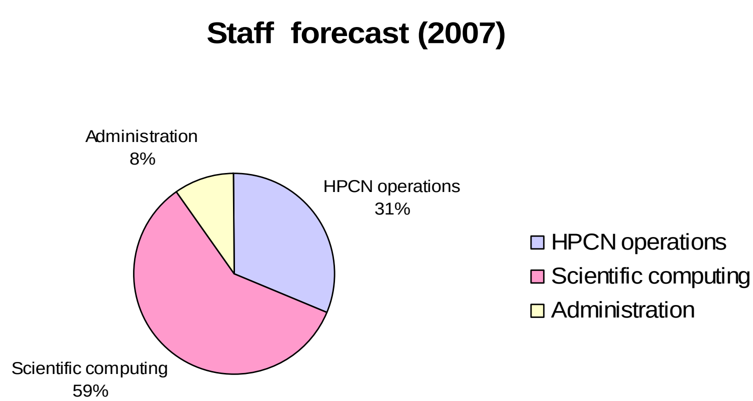

Figure 4: Future staff composition of CSCS

3.2. Focal applications

Until 2000, the application –and expertise- portfolio of CSCS has been developed along four main domains, traditional of high performance computing from its inception: **engineering (CFD, energy, automotive, aerodynamics), computational chemistry, computational physics, meteorology and climate modelling**. Relying on its existing strengths, CSCS will keep on developing these competences and collaborations, that will also benefit of the thrust areas as described under Table 1.

A more recent thrust of the ETH domain is the sector of molecular and material sciences. By continuing the support for Prof. Parrinello's research group, CSCS's customer portfolio comprises one of the top players of this discipline.

In this development phase, it is very important to further enrich the portfolio of applications towards new communities, like **biology and biomedical sciences, particle physics, and Grid computing and services.**

The terms "life sciences" and "bio-informatics" cover a vast area of different research directions ranging from genomics to simulations of cell processes. Besides, life sciences are one of the thrust areas of the ETH domain. Although a few centres of gravity already exist for bio-informatics and bio-simluation in Switzerland (the Swiss Institute for Bio-Informatics, the Arc Lémanique, the Zurich *fcgz*, the Basel Bio Zentrum,), we are confident that CSCS can play an important national role in developing this major field of computational sciences. In this range of scientific explorations, many aspects are not yet covered by the aforementioned institutions, and even in their current focal directions, CSCS is a welcome partner for developing complex algorithms and software systems as well as for adapting existing systems to high-end capacity resources. Furthermore, the wealth of important pharmaceutical industries in Switzerland makes bio-informatics a perfect field for joint-effort projects between academia and industry. As a first move, CSCS will enter the field of visualization in bio-informatics. Visualization has been a traditional stronghold of CSCS, and the value of the existing work was highly appreciated by the evaluation team of the 2002 peer review. First contacts with the Swiss bio-informatics scene have shown high interest in a contribution of CSCS in structural visualization. Thus, CSCS will contact the group of life scientists in Western Switzerland (Arc Lémanique) and the Swiss Institute of Bio-Informatics.

Another important development in HPCN is the concept of grid computing. In a first step, CSCS will participate as a tier-2 centre of the CERN LCG (LHC computing grid). Although particle physics is one of the disciplines to be reduced or to be restructured according to the strategic plan of the ETH domain, LHC computing offers a natural and easy way to have first experiences with concepts and software in distributed computing. The memorandum of understanding with the Swiss labs and the members of LCG has very recently been signed. As CSCS may also participate in the DEISA initiative – presented in the sixth framework programme of the European Union-, knowledge transfer to other application fields will be easily attained.

Another significant development will be the combination of grid technology and bio-informatics. First talks with bio-informaticians in academia and industry have been very encouraging. In conjunction with the visualization project mentioned above, the Swiss Bio-Grid could evolve over the next couple of years with CSCS as the leading house. At the time of writing, the memorandum of understanding for creating the Swiss Bio-Grid is being negotiated.

The lines of action above described represent a continuation of directions of the recent past and have already been identified during the 2002 peer review. In addition to these, three new strategic thrusts will be initiated at CSCS to give it new academic impact.

Application Field	CSCS Contribution	Status
Engineering, Computational Chemistry, Physics	software framework and expertise in application optimisation	ongoing work
Meteorology and Climate Modelling	contract work for Meteo Swiss, Prism collaboration	ongoing work
Material Sciences	hosting CPMD-software & support to the Parrinello group	ongoing work
Grid Computing	towards e-science • Computational Grid • Information Grid • Knowledge Grid	Tier-2 Centre of LHC Computing, CoreGrid and Swiss Bio-Grid
Earth Science & Global Modelling	software framework for supporting community	first contacts planned for 4Q 2003
Life Sciences	development of visualization/imaging techniques	first talks with Swiss bio-scientists completed
Information Science	hosting of data and development of filtering and retrieval methods	customer acquisition will start in 2004 (libraries, archives, hospitals)

Table 2: Focal applications of CSCS

3.2.1. Community Modelling Framework

At the 18th International Supercomputing Conference ISC2003, Jim Gray of Microsoft Research outlined what he called "The Evolution of Science": the development from an observational activity to an analytical discipline to computational science and finally to data exploration science. The data used during the last phase is either instrumental (particle colliders, gene arrays, etc.) or simulation data. In the latter two phases, the scientific community of a discipline often develop a common software infrastructure to be used for modelling and exploration tasks. The best-known example for such a community model is the NCAR climate model. Here, HPCN centres can offer an essential added value as community models must be developed and fostered over a much longer time period than the three to four years typical Ph.D./post-doc cycle. HPCN centres can guarantee the necessary continuity for these modelling frameworks. Furthermore, a community model typically covers a multitude of domains of scientific computing, from data modelling to simulation to archiving and visualization. Becoming the host of a community model would give CSCS the unique chance to develop many of the competencies listed in Table 1 and to gain instant visibility in the academic world.

The possibility to become the hosting centre of a community model depends largely on talking to the right people at the right time. The time window of opportunity is typically short and credibility is crucial. In addition, the application must fit into the strategic directions of the ETH domain. ETH Board has outlined five research domains to intensify in the 2004-2007 period: information sciences, life sciences, microscopic material sciences, micro-/nanotechnologies, and natural hazard and risk management. Each field offers CSCS a variety of opportunities. So far only the field of natural hazards has actively begun to develop new community models. Whitepapers have been published, describing the necessity to create such models for earth surface processes, geodynamics, and earth system modelling. CSCS will explore the possibility of becoming the hosting centre of one of these community models.

Several international developments back the analysis that solid earth simulation and natural hazards is an up-coming major trend in computational science: In 2002, Japan opened the Earth Simulator Centre. In 2003, Australia will launch its Australian Computational Earth Systems Simulator (ACcESS). The European Union is discussing the opening of a European earth observation data centre, hosting the increasing wealth of data provided by satellites. Earth modelling is one of the most important number-crunching applications (5 of the current top 15 centres in the TOP500 are completely or mostly dedicated to earth modelling). It has become a data-intensive application due to remote sensors and instrumentation. One could say, that the field of modelling solid earth and natural hazards is currently in the state in which bio-informatics was about five years ago. Furthermore, natural hazards are a national and scientific priority in Switzerland. It is therefore deemed to be the perfect application for CSCS' purpose.

3.2.2. Benchmarking

With the development of beowulfs, clusters with high performance communication systems, clusters of fat nodes, and wide-area distributed computing system such as grids, there is a growing need in the exchange of experience and information on benchmarking, testing methods and the assessment of total ownership cost. The reinforcement of this existing activity at CSCS on the present architectures towards new communities (earth modelling, molecular science and biology, for example) will give the necessary support to leading groups in Switzerland.

3.2.3. Data-Intensive Computing and Knowledge Management

High-throughput computing has become one of the most important application fields in HPCN. The growing volume of data can be simply seen from the fact that the amount of data in the computing centres worldwide grows at a higher rate than Moore's Law, which means that data volume grows faster than processor speed. The processing of large amounts of data, from archiving to filtering to the retrieval and management of knowledge, will therefore represent another focal point of CSCS' activities.

Data-intensive computing will bring new groups of customers to CSCS. High-energy physics, medical research and hospitals, libraries, archives, and large-instrumental sites are facing challenges in handling and visualizing their data. As it is more efficient to transfer a condensed result found by filtering or machine-learning instead of the original data, or to transfer a visualization, data-intensive computing will fertilise other areas of scientific computing as well. Data-intensive computing will not be restricted to numerical data but will also extend to textual data. Thus, information retrieval techniques will also be of importance. CSCS will start to actively establish new collaborations in 2004 by contacting groups that produce and use large amounts of data. One small first step towards this direction is the operation of a knowledge database on research projects in Switzerland.

3.3. Education and Training

One of the main shortfalls of CSCS is still its insufficient integration into the academic scene of Switzerland. When comparing CSCS to other national HPCN centres, it is striking how low the activities in training and education are. For example, the Finnish CSC estimates that 1,500 students and researchers participate at its courses and summer school annually. Although CSCS serves an academia of similar size than CSC, it reaches not one tenth of this number. None the less, offering courses and summer schools, is an extremely efficient way to acquire new and future customers. The creation of a course-offer to be launched in 2005 is therefore one of the strategically important activities for CSCS (Goal 3).

3.4. Technology Strategy

CSCS currently operates two different architectures: a parallel-vector system with a single shared memory and a cluster of fat SMP-nodes with scalar processors and distributed memory. As different HPCN applications require different HPCN architectures, the centre will most likely continue to maintain two distinct systems. However, the emphasis will be shifted away from the schism between vector and scalar processor types to different system architectures.

Although technology prediction is always partially a look into the crystal ball, some basic developments can be observed:

- Firstly, scalar processor architectures are converging. As a result of the first step of convergence, observed in the last two years, the border between so-called RISC (reduced instruction set) and CISC (complete instruction set) processors was blurred. The current HPCN market is dominated largely by the IBM Power processor family and the Intel Itanium processor type. The Power processor, a typical RISC architecture, has grown its instruction set so extensively during its evolution in the 1990s to 2002 when the last type was put on the market, that one can hardly speak of a "reduced" instruction set anymore. On the other hand, the Intel processor family, a typical CISC system, started to incorporate features of RISC processors like speculative instruction execution.

- Secondly, vector and scalar processor architectures are converging, too. The newest incarnation of the IBM Power-4 technology, the 970 released in mid-2003, has a small vector processing unit added to the scalar processor core. Other developments try to add vector processing capabilities to scalar computers on the system level: The "Blue Planet" project of IBM and NERSC wants to add software-vector processing capabilities to the IBM Power-5 architecture to be released in the next years. Hitachi has already a massive-parallel system with pseudo-vector processing capabilities in its portfolio.

- Finally, on the system level, even pure vector systems like the latest NEC technology are designed as massive-parallel systems, meaning that they have to be programmed with the same programming model based on MPI as scalar massive-parallel systems.

A better discriminator for HPCN architectures than processor architecture is the communication between the different processing units because it defines the suitability for different applications. Systems with extremely large numbers of processors, high nearest-neighbour communication bandwidth but low bisection bandwidth, are very well suited for bio-informatics, for instance. Other applications like computational fluid dynamics rely heavily on very high bisection bandwidths. The key is the balance between processor performance, memory bandwidth and communication bandwidth on different levels of the system. CSCS will therefore concentrate on different overall designs that will enable it to offer best performance for different sets of key applications (Goal 2).

4. Organisation and governance

For CSCS achieving the scientific and technological goals described above, the external and internal organisation of the centre must be adapted to attain the necessary flexibility and autonomy. In a joint effort, ETH Board, ETH Zurich and CSCS developed a basic organisational concept which on the one hand gives CSCS a certain independence and on the other hand establishes reporting and steering mechanisms. It is based on the FLAG[a] concept of the federal administration. Furthermore, the new orientation of CSCS is reflected in its new internal organisation.

The FLAG concept foresees a strategic steering entity placed between the institution and its umbrella organisation. The steering board negotiates a performance agreement and a global budget for the institution with the umbrella organisation. The board surveys the performance of the institution using reporting indicators and itself reports to the umbrella organisation. Although the institution remains an administrative part of the umbrella organisation, it is not connected to it in a hierarchical sense with top-down instructions.

Conforming to best practices, five key functions have to be defined in the new organisational scheme:

a) strategic steering
b) scientific advice
c) customer representation
d) resource allocation
e) technology selection

For an entity of 50 FTE, it would be inadequate to create a permanent board as a part of CSCS' organisation for each of these functions. Instead, we propose customers to be represented in a self-organised, grass root fashion, similar to many user groups in indus-

---

[a] FLAG: „Führung mit Leistungsauftrag und Globalbudget" – Management by performance mandate and global budget

try. Furthermore, we think that technology selection does not require a permanent panel but can be organised as a project when required. A pragmatic, lean, and efficient organisation is achieved through the new managing structure.

Governance

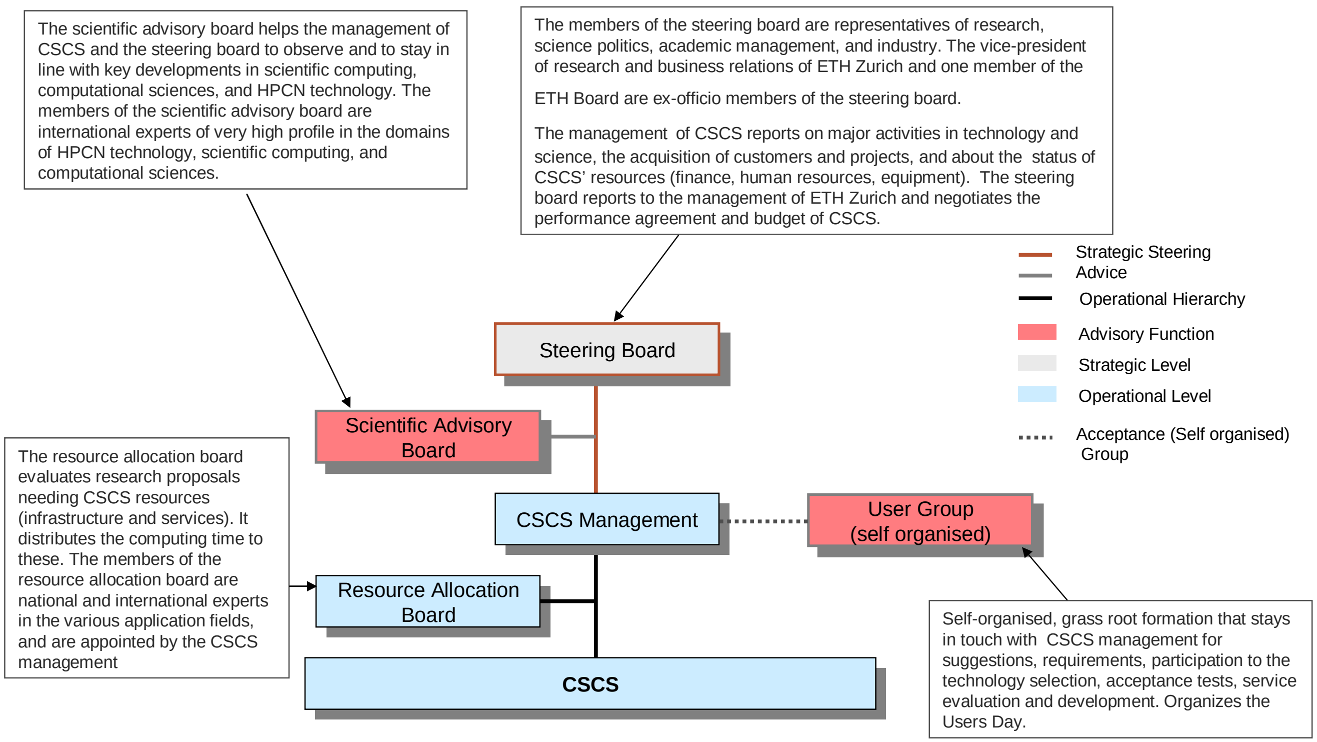

15