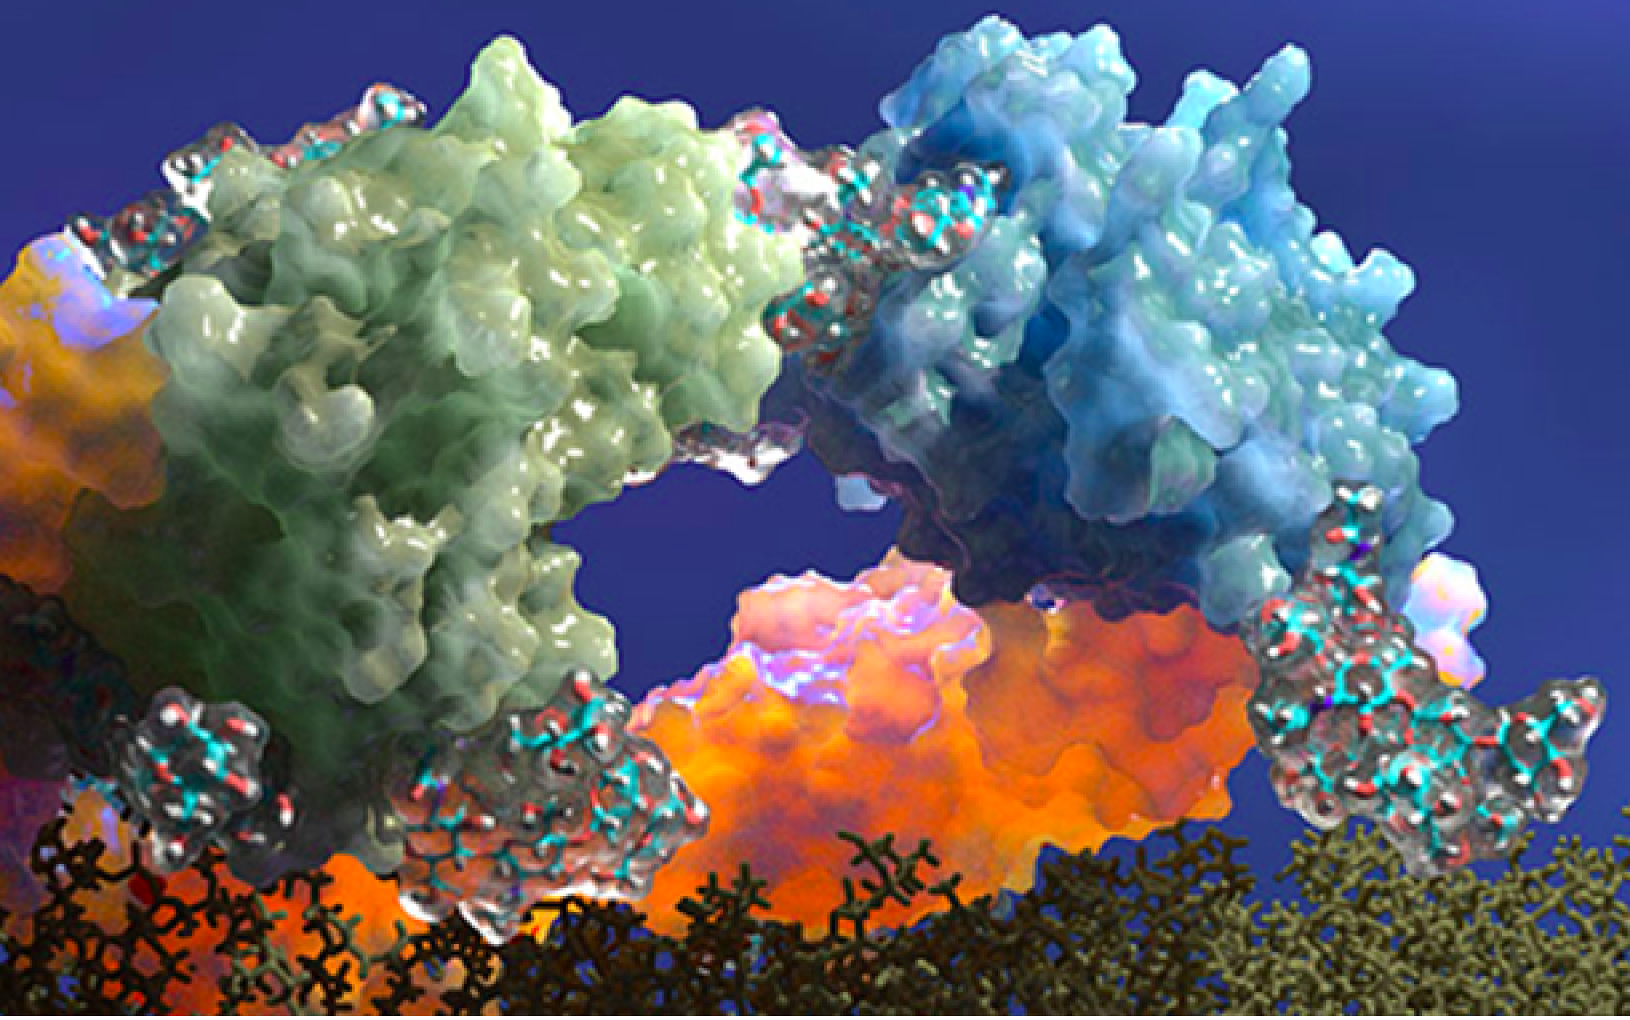

CSC

The Scientific Case and User Requirements for High-Performance and Data-Intensive Computing in Finland 2017-2021

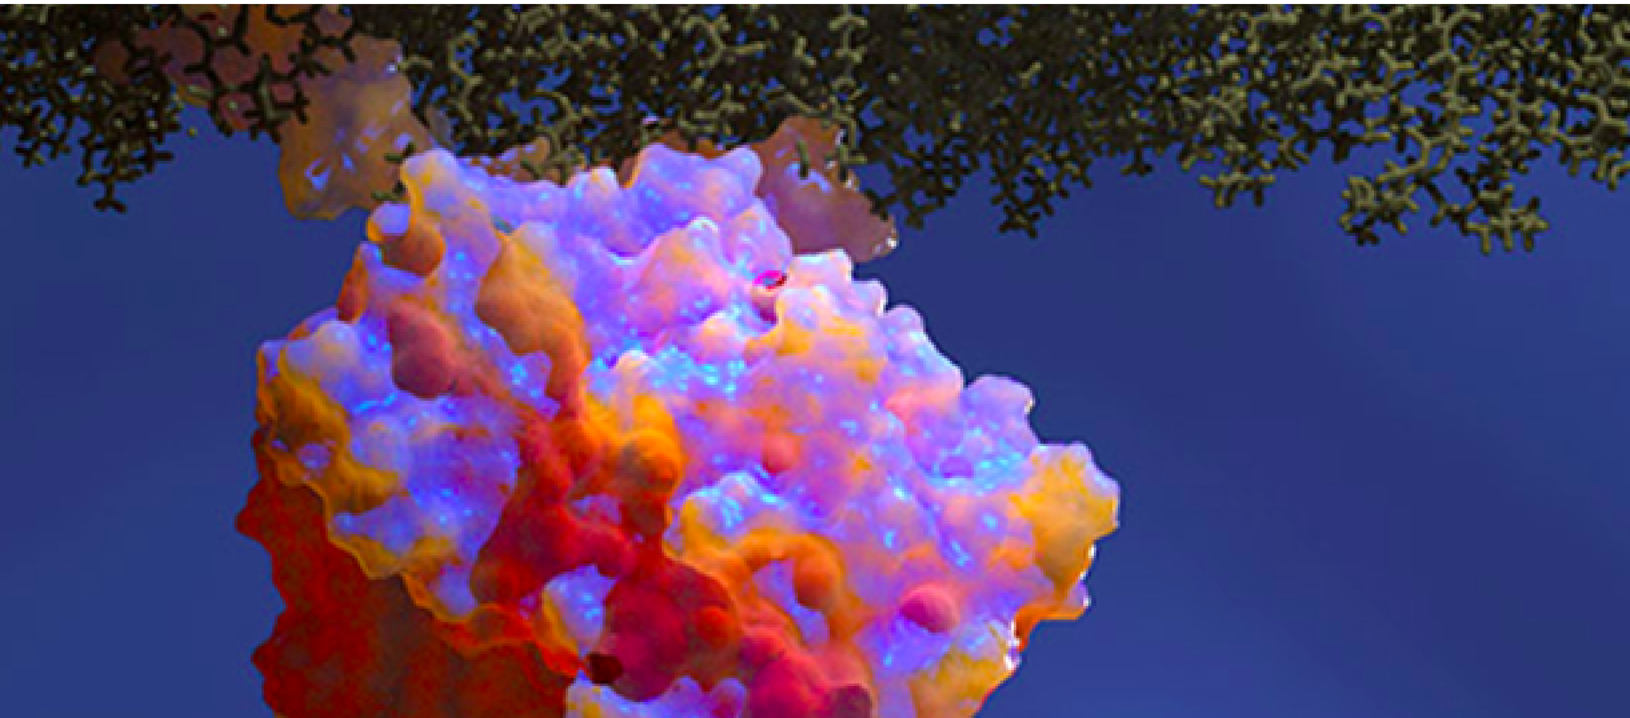

The Scientific Case and User Requirements for High-Performance and Data-Intensive Computing in Finland 2017-2021

Table of Contents

Cover image: N-Glycosylation as determinant of EGFR conformation in membranes. Kaszuba, Grzybek, Orlowski, Danne, Rog, Simons, Coskun, Vattulainen. PNAS 112, 4334 (2015) Copyright: Jyrki Hokkanen, CSC

Executive summary

This report details the case for renewing the centralized research infrastructure at CSC - IT Center for Science (CSC). Traditionally the level of computational science in Finland has been on a high level, supported by regular investment to maintain an internationally competitive infrastructure. Without a regular investment there is a risk that this position will be lost, as the international competitiveness of research in Finland will decrease. This would not only harm the scientific community, but would also negatively impact the Finnish society and economy, as the ideas and minds of science often help companies to innovate and to utilize new technologies.

The future prospects of computational scientific research in Finland are highlighted by the following examples:

- Finding new tools for diagnosis, and treatment of diseases by understanding tumor genomics and cellular signaling, and via advanced microscopy.
- Novel materials obtained by computational materials modeling combined with big data methodology.
- Towards clean and efficient energy production by modeling nuclear fusion and other power plants
- Understanding climate change through improving the coupled climate models, finding feedback mechanisms and modeling sea-level change due to glacier ice loss.
- Improving human worldview by studying fundamental topics in astrophysics and high-energy physics: galaxy formation, nature of dark energy and physics beyond the standard model.
- Emerging fields such as machine learning and its applications in areas such as natural language processing.

The cases are established by discussing with scientists and through a series of six seminars arranged at universities throughout Finland.

From these discussions and cases, supplemented by a large user survey, we analyze the requirements placed on the future infrastructure, and recommendations are drawn up to meet these needs. These include:

- Ability to efficiently run computational jobs ranging from single-core and single-node to massively parallel jobs employing tens of thousands of CPU cores simultaneously. To be in line with CSC’s previous procurements since 2000, and to maintain the relative competitiveness level globally, the new infrastructure should be able to produce total computational capability of around 7 Pflops in 2017, with an additional upgrade in 2019.
- The requirements of data-intensive computing need to be addressed. This in practice entails larger amount of main memory, large and versatile storage systems, and also support for cloud environments.
- To meet these requirements funding for the infrastructure investments should be 30-35 million euros in 2017-2019.

Introduction

Computing and data-driven approaches are ubiquitous in science today; they are utilized at the advancing forefront in almost all fields of science. In addition to traditional users of computational methods, the humanities and social sciences are introducing computational methods as a consequence of digitization and accelerating growth in data volumes.

The combination of modelling and simulation is the third pillar of scientific research, together with theory and experimentation: with them one can compare complex models against measured data. With large-scale computing, we are able to understand phenomena that are too slow or fast, too small or too large for a human to observe: through simulation we can reproduce these processes step by step. Simulation can complement and replace experimentation should the experiments prove too difficult, expensive, dangerous or unethical to realize. The ever larger simulations produce enormous amounts of data that need to be stored and processed, and the computationally heavy simulations are often followed by a data analysis stage requiring high performing data storage systems, as well as computing architectures suited for data analysis.

Not only simulations produce tremendous amounts of data. During the last few years, the amount of experimental data has increased dramatically, and the trend seems to continue at an accelerating pace. DNA-microarrays and next-generation sequencing approaches reveal more details about our genome, satellites provide us with measurements of both the Earth and the universe over a wide spectrum of frequencies, and experiments in particle physics will overwhelm all available disk space with data. The scientific community must be ready not only to save these data once they are created, but above all, to analyze them efficiently in order to find and establish new knowledge.

The needs of the scientific communities are very diverse. Top-tier supercomputers are needed for running large simulation models utilizing tens of thousands of compute cores in parallel. Equally important is high-throughput compute capacity for executing many related smaller computational jobs in order to gather statistics and to investigate a wider number of parameter values. Advanced data storage systems are needed: fast disks for the computing systems and flexible provision of long-term storage for archival purposes as well as for distributing large datasets among research communities. Especially in data-intensive computing, cloud technologies are also gaining a foothold since they allow custom environments to be developed for large communities and enable sensitive data to be handled. Also, in data-intensive life science the computing is mostly limited by the speed of storage and main memory, and not the central processing unit (CPU) performance.

Specific examples of these computational challenges studied in Finland include global warming and our future climate; modeling fusion power plant plasma and reactor materials; understanding and inventing complex novel materials; complex chemistry and nanoscience; large-scale protein dynamics, protein association and aggregation; bioinformatics; elementary particle physics; space physics; as well as intense fluid-dynamics simulations for e.g. modeling biomedical flows, gas turbines, combustion engines, or green aircraft. In addition to the established users of computational methods, it is possible to foresee many new fields emerging and start employing high-end computing.

By looking at this range of examples, it becomes clear that platforms for scientific computing are the only class of top-end research infrastructure that is applicable to numerous scientific problems across the disciplines - experimental equipment is most often aimed for serving one community of researchers, even for solving just one particular problem. The high standard of Finnish computational science is partly due to the internationally competitive national scientific computing infrastructure at CSC, combined with decentralized solutions and international resources.

This report has been written during the autumn of 2015. It is based on data that has been gathered from CSC's customers, as well as researchers at research institutes, to get an up-to-date view of the visions and user requirements of the scientific community in Finland. The data has been collected through personal interviews with researchers and through a series of six birds-of-feather seminars arranged in Espoo, Jyväskylä, Oulu, Kuopio, Lappeenranta, Tampere and Turku. A user survey was sent to 472 researchers in Finland representing the top users of compute, data and cloud services, researchers at research institutes, and a sample of all CSC's customers. In addition, the batch job statistics of CSC's current set of machines were carefully analyzed to establish further insight on the usage trends and future infrastructure needs.

Infrastructure for high-performance and data-intensive computing in Finland

Historical development

The computational capability of high-performance computing systems have been gathered into the Top500 list, which twice a year lists world's 500 fastest supercomputers at that time. Ever since the establishment of the list in 1993 there has been a clear trend: exponential growth. Throughout the history of the Top500 list the computational capability of all spots on the list have grown by up to 1.8 each year.

In Figure 1 the ranking of CSC's systems are shown on the Top500 list, as well as a graph showing the aggregate performance of CSC's computing environment over time. CSC's flagship system has typically been among the world's 100 fastest supercomputers for a few iterations of the list, losing its position at a relatively rapid pace. In order to remain competitive with the peer centers, a regular investment is needed to upgrade the infrastructure.

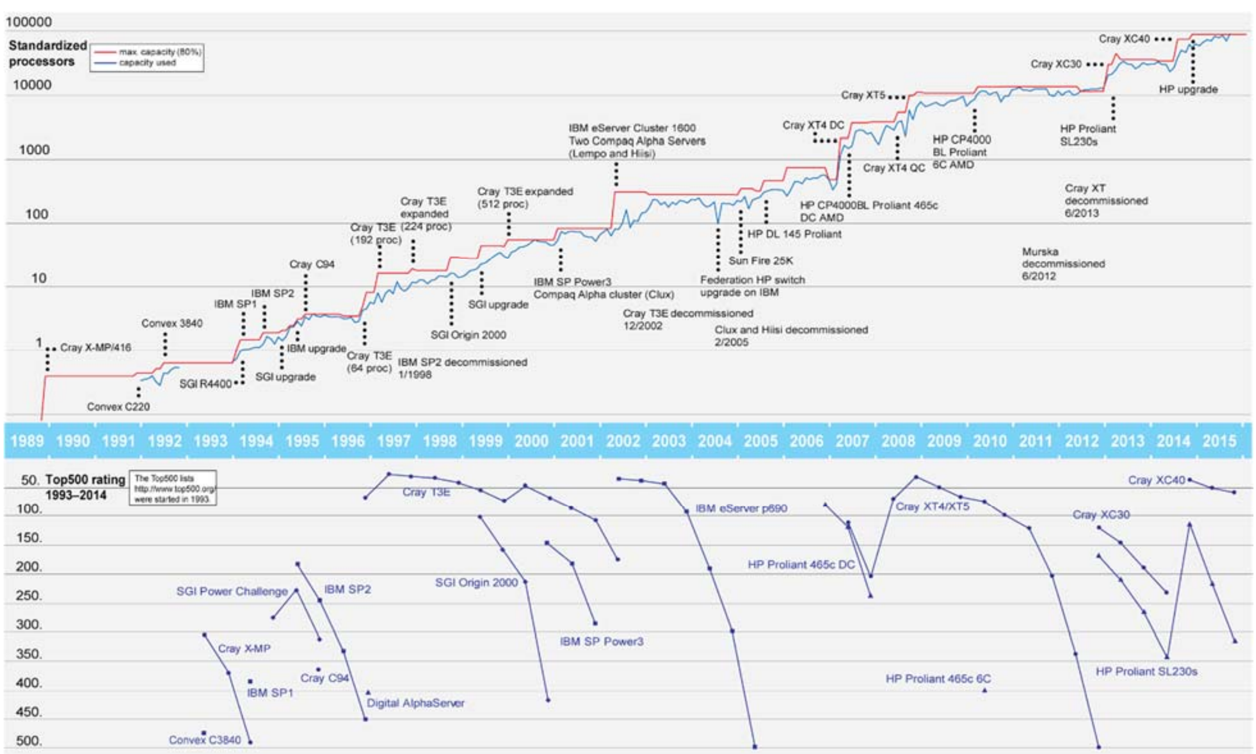

Figure 1: Performance development of CSC's computing environment over time.

Current infrastructure at CSC

Capability computing	A Cray XC40 supercomputer Sisu is meant for large-scale parallel computation which requires supercomputing resources, from several hundreds to up to thousands of cores. It comprises of 1688 compute nodes, each containing two 12-core Haswell processors with 64 GB of memory, interconnected via the high-performance custom Cray Aries network. The first phase of Sisu was installed in 2012, followed by a major expansion and a CPU upgrade into its current shape during summer 2014.
Capacity computing	Taito provides a cluster environment for a wide range of scientific software that uses anything from one core to a few hundreds of cores per job. Taito is a heterogenous supercluster manufactured by Hewlett Packard, with 985 nodes based on Intel CPUs connected with an Infiniband network. Approximately one half of the machine has 64 GB of memory per node, while the other half has 128 GB of memory per node. For jobs requiring large amounts of main memory there are 26 nodes with 256 GB of memory and 2 nodes with 1.5 TB.
Cloud computing	CSC's cloud computing platforms, cPouta and ePouta, offer flexible high performance computing resources with via Infrastructure as a Service (IaaS). cPouta is intended for general-purpose use and ePouta for organizations needing a secure private cloud. cPouta is provisioned from the Taito supercluster whereas ePouta is at CSC’s Espoo data center. The users connect to ePouta via private network solution.
Accelerated computing	Specialized Taito accelerator partition consists of 38 nodes accelerated with Nvidia based GPUs, and 45 nodes accelerated with Intel Xeon Phi.
Storage	HPC Systems Sisu and Taito use a Lustre parallel file system hosted on a DDN storage system, with 4 PB of storage capacity. For long-term storage CSC provides an HPC Archive, complemented by the Ministry of Education’s IDA service for open data. cPouta is currently also using the DDN storage system, but both ePouta and cPouta are taking in use storage servers based on the Ceph system. ePouta provides also a private NFS service.

The case for the next-generation e-infrastructure in 2017

Current infrastructure will not provide sufficient capabilities after 2017

An efficient and easy-to-use centralized infrastructure for scientific computing is a competitive edge for universities throughout Finland. This enables cutting-edge computational research by established groups, while also enabling new upcoming researchers to get access to competitive resources. At the same time the infrastructure helps to attract skilled researchers to Finland and in retaining talent in Finland. Computational sciences, e.g. computational chemistry and physics, have been rated highly in status reviews by Academy of Finland during the past couple of years.

Today, the infrastructure provided by CSC is in general sufficient for Finnish researchers, but it is clear that the situation will rapidly change. For researchers using computationally intensive methods, such as computer simulations, novel breakthroughs require larger and more accurate simulations. Alternatively, more simulations are needed to improve statistics and to fully investigate the processes in question.

In the case of data-intensive research, both experimental and simulation data is exploding in size. Already now researchers have ambitious plans that are unfeasible with the current infrastructure. Several disciplines also require compute capacity with large main memory, a resource which is already now turning into a bottleneck.

Cost-efficiency

The power efficiency of computing platforms is improving at a rapid pace. This means that a newer machine can provide much higher scientific throughput using the same amount of electricity. Assuming a four year lifetime, the investment can be estimated to be approximately half of the total cost of ownership of the infrastructure. Prolonging the lifetime of the infrastructure beyond its useful lifetime still incurs a significant costs (electricity, administration, servicing, etc.), while in terms of science per euro the performance is very poor compared with international competitors. Thus no money is saved by delaying the investment.

Computational science is also comparatively affordable, with a sufficient investment one can support hundreds of users from almost all fields of science during the next four years, whereas an experimental research infrastructure (possibly of comparable investment) usually serves only single community or discipline, even a single research group.

The socioeconomic impact of computational science

Breakthroughs in scientific problems of high societal impact, such as understanding climate change, designing new medical treatments or finding new materials, require often carrying out so called grand challenges i.e. computational tasks of unprecedented complexity. In practice, being able to facilitate these demands ever-growing capability and capacity for computation as well as storage and data analysis. It should be noted that solving these grand challenges will not only bring answers to the academic research questions but produce potential being seeds for commercial innovation.

An up-to-date infrastructure for scientific computing is also for the good of the competitiveness of industrial R&D, as modeling and data analysis play an important role therein. A tight collaboration between higher education institutions, academic entrepreneurship, and start-up companies emerging from academic research have great potential in creating new jobs and tax revenue. For example, data about genes and healthcare data of entire populations can be applied for instance in diagnostics of complex diseases and making more safe, precise and cheaper medicines, and new biotechnological applications. CSC's resources can act as a stepping stone for small companies in Finland, for instance present their scientific software solutions in the international setting on scalable platform.

Access to international research infrastructures

PRACE, the European research infrastructure for high-performance computing, and the INCITE program in the United States provide top-end (Tier-0) resources on some of the world's largest supercomputers. To gain access to these highly competed resources, the research groups have to prove technical readiness for Tier-0 as well as an ambitious scientific case. Finnish researchers have been very successful in gaining access to PRACE resources, and an important enabler has been the fact that the step from national to international resources is not too large and the technical and scientific readiness has matured on the national resources.

The advanced infrastructure has also enabled CSC to attract and host international infrastructures. For example, the life science data infrastructure ELIXIR connects Finland with global life science data resources. CSC is the ELIXIR Finland node, and already contains petabytes of curated high-quality life science data from all over the world to Finland.

Datacenter infrastructure know-how and ecosystem

Finland has been investing for quite some time in establishing expertise in data centers and data center infrastructure. We have attracted large global players to build their data centers in Finland by utilizing the competitive edges we have in this field: energy and cost efficiency, societal stability and high standards for security. Thus it is rational to host the data centers needed by scientific research in Finland; and short-sighted to let the established expertise drift out to other countries.

Selected scientific cases in Finland

In the following, a number of scientific cases are presented, selected to represent the computational research in Finland. These have been obtained by discussing with scientists at different stages of their career. They are listed in an alphabetical order. Furthermore, the related needs for the ICT infrastructure are being discussed among the science cases.

Astrophysics and high-energy physics

Elementary particle physics: CMS and ALICE experiments

Elementary particles are particles which do not consist of other particles at a finer scale. There are both elementary particles which are the building blocks of all ordinary matter (sc. fermions), as well as particles (bosons) which mediate forces between fermions. The Standard Model of particle physics is a very successful theory collection that explains all known properties of elementary particles and their interactions, with the exception of gravity. Recently, the neutral Higgs boson was found by the Large Hadron Collider (LHC) at CERN. This particle was the last part of the Standard Model which existence had not been experimentally verified.

Helsinki Institute of Physics (HIP) is involved in the analyses of the CERN CMS experiment, and long term preservation of the CMS data. Additionally HIP is involved in efforts to provide open access to it, and leads a pilot project to bring the data to public use at schools. The HIP group analyses a particular decay channel of the J/psi particle, interesting because it can indicate contributions from new physics beyond the Standard Model. Researchers of HIP are also involved in the searches for charged Higgs bosons. The ALICE experiment is a dedicated heavy-ion collision experiment to exploit the unique physics potential of nucleus-nucleus interactions at LHC energies. The goal of the ALICE experiment is to study the physics of strongly interacting matter at high temperatures and extreme energy densities, where the formation of a peculiar phase of matter, the quark-gluon plasma, is expected. Related work at HIP is focusing on the data analyses.

Monte Carlo simulations of the CMS and ALICE detectors and the physical processes, are a crucial tool for analyzing the data. The work require a very large amount (up to 2 PB) of storage and cloud-like computational resource for data-intensive computing.

Formation and evolution of galaxies

Theoretical Extragalactic Research Group at Helsinki University, led by Prof. Peter Johansson studies the formation and evolution of galaxies. Their computational research answers fundamental questions on the assembly of the stellar component in galaxies and the relationship between stars and dark matter, formation mechanisms of elliptical galaxies, the formation of supermassive black holes in the early Universe and their subsequent impact on the evolution of massive galaxies, as well as the detailed dynamics of supermassive black holes and the stars around them. All of these also directly link to the observational community by providing clues for analyzing data; observational signatures that allow distinguishing between different galaxy formation mechanisms, accurate predictions for the merger rates of supermassive black holes, etc.

It is clear that the computational demands will increase in the future due to more complex simulations as well as the need to improve statistics through a larger number of simulations, each employing thousands of compute cores. The work described above is critically dependent on a competitive national computing infrastructure to remain relevant in the community. This enables faster development of new models and makes it possible to solve scientific grand challenges on national supercomputers. Physically accurate galaxy simulations are typically run on up to 2000 cores, and also require a decent amount of memory per node, at least 128 GB but ideally up to 256 GB.

Nature of dark energy

Euclid is a cosmology satellite mission of the European Space Agency (ESA), to be launched in March 2020. The primary goal of Euclid is to solve the central mystery of modern cosmology: the nature of dark energy - i.e. why does the Universe expand in an accelerating pace. Whereas ESA's previous satellite mission, Planck, concentrated on the first 500 000 years of the history of the Universe, Euclid concentrates on later stages, the last 10 billion years of the 13.8 billion year age of the Universe. Euclid and Planck complement each other ideally in improving on our understanding of the Universe. Euclid will be the leading project in cosmology in the coming decade.

Euclid produces a very large amount of observational data: a three-dimensional map of the Universe that will have a great impact in both cosmology and astrophysics as well as human worldview in general, and form a basis for future research in these fields. The data from the satellite will be analyzed by the Euclid Science Ground Segment (SGS) consisting of 9 national Euclid Science Data Centers (SDC). The Finnish contribution to Euclid is to participate in the Euclid SGS, and in particular to provide one of the SDCs, the SDC-FI. SDC-FI will be set up in the CSC Kajaani datacenter and will eventually contain petabytes of disk space.

Physics beyond the Standard Model

It has become clear that the highly successful Standard Model of particle physics is not the ultimate "theory of everything" but rather a phenomenological theory valid at low energies. Cosmological observations have also shown that the mass of the Universe mostly comprises of dark matter and dark energy, neither of which the Standard Model is able to explain. Professor Kari Rummukainen's Computational Field Theory group at University of Helsinki is a key group in Finland in developing and validating theoretical models that extend or correct the Standard Model, in particular supersymmetry and the technicolor theories. This is carried out using computational methods, mainly the lattice Quantum Chromodynamics (QCD) method, which is a lattice-based method for investigating technicolor-like theories. Key scientific questions include gravitational wave production in phase transitions in the early Universe, and their possible observation. To this end, the group is involved in the upcoming eLISA mission by ESA, which is the first space based mission for measuring gravitational waves.

The group has developed their own code (SU2), which is tuned for modern supercomputers and can utilize tens of thousands of cores simultaneously. It is also able to utilize not only normal CPUs but also GPUs. The requirements for this science case is a large number of cores connected with a high performance interconnect, while neither memory nor I/O are major bottlenecks. Even if the code supports GPUs the group still prefers supercomputers based on normal CPUs as they are easier to use and develop for, and because the GPU version typically does not scale to a sufficiently large number of nodes.

Engineering and industrial applications

Computational fluid dynamics in energy technology

Energy technology studies technologies and systems for generating, distributing and using energy. Key issues are improving efficiency and environmental friendliness of the technologies. Many of the problems studied in Finland are solved using computationally demanding methods, e.g., three dimensional simulations of fluidized bed furnaces, finite-element simulations of electromagnetic engines, and optimizing wind power installations using computational fluid dynamics (CFD) simulations.

A potential wind farm site has to be thoroughly surveyed and the wind climatology analyzed before installing any wind turbines. Modeling atmospheric boundary layer (ABL) flows over complex terrains containing, e.g. hills, forest, and lakes can help in optimizing the location and design of wind farms. Numerical modeling of wind flows using CFD has become a popular technique during the last few decades. This kind of modelling has been done for example at the Lappeenranta University of Technology. The CFD simulations for energy technology rely and benefit from the available high-throughput (mid-range) computational resources.

Inverse problem: Mapping groundwater resources

In an inverse problem, one attempts to reconstruct a structure from a set of low-dimensional observations, a problem typically much harder than simulating the observation given a known structure. Inverse problems can be applied in several fields, including medical imaging, image processing, mathematical finance, astronomy, geophysics, nondestructive material testing and subsurface prospecting. One of the scientific cases solved in the Centre of Excellence in Inverse Problems Research (2006-2011 and 2012-2017) is mapping groundwater resources based on seismic signals generated by earthquakes. In Finland the groundwater resources are plentiful, but in many parts of the world the resources are much scarcer and the wellbeing in the region is dependent on accurate estimates of their capacity. Groundwater reserves have traditionally been mapped using boreholes but this process is slow and expensive to implement on a large scale, and the results often lack accuracy. In the University of Eastern Finland this problem has been extensively studied and custom software and methods have been developed. The groundwater inversion problem is not a heavy user of data storage, but does require significant computational resources. With the current resources the methods have successfully provided high-quality two-dimensional estimates for groundwater resources, but the ultimate target is realistic three dimensional maps of the resources. Current two-dimensional reconstructions use thousands of cores, while the realistic three dimensional reconstructions with hundreds of millions of elements will require tens of thousands of cores.

Simulation of prototype fusion reactors

A mix of sustainable energy sources with low carbon emissions are urgently needed. Nuclear fusion is the main candidate for baseload electricity production in the long term, because fusion reactors can operate independently of environmental conditions and they also have a very high power yield. In the ITER experiment, a tokamak-based, power-plant scale fusion reactor is currently under construction in France as a world-wide collaboration. ITER is an important step towards DEMO, which will be the demonstration reactor for electrical power production.

Numerical simulations are essential for the design and operation of these research reactors. Fusion experiments are very expensive and computer simulations are therefore used to plan the tests in order to minimize the risk of damaging the structure of the reactor. Advances in high-performance computing together with technical challenges of tokamaks has also revitalized the interest in other designs, such as stellarators. The world's largest stellarator Wendelstein 7-X will start operation at Max-Planck-Institut für Plasmaphysik later this year. Unlike tokamaks, stellarators are inherently steady-state devices and thus readily suited for continuous operation.

Several Finnish Universities together with the Technical Research Centre of Finland (VTT) form the FinnFusion consortium which participates in these important international projects. Prof. Mathias Groth’s group at Aalto University studies the challenging plasma physics problems in ITER and W7-X using a number of simulation codes, among them ASCOT and Elmfire, which have been developed in Finland. The stability of the fusion process has continued to be one of the main technological challenges, and improved computer simulations are together with experiments vital to overcome it. The larger size of the upcoming tokamak based experiments (ITER, DEMO), and the complex geometry of the stellarators (W7-X) also poses new plasma physics computational challenges, especially so when whole-device simulations are needed. Elmfire and ASCOT both benefit from large amount of memory per node, but large simulations also require a fast interconnect.

Fusion reactors will produce large quantities of highly energetic particles, which can potentially harm the walls of the reactor. Understanding the full effect of these particles pose an inherently multiscale computational grand challenge at the extreme of high-performance computing that involves a number of complex, interdependent processes occurring over a wide range of length and time scales. At the University of Helsinki, Prof. Kai Nordlund’s research focuses on using molecular dynamics and accelerated dynamics techniques on supercomputers to discover the kinetic mechanisms governing surface and microstructural evolution. Recently the group has recognized that to be able to predict reliable operation of DEMO and commercial power plants, the materials simulation methods need to be extended to a more fundamental, quantum mechanical level. This deeper level of simulations will require a massive increase in supercomputer capacity compared to what is currently available.

Life sciences and medicine

Bioimaging

Bioimaging has recently emerged as one of the most important research methodologies in biomedical research globally, with 70% of high impact publications relying on advanced microscopy. Bioimage informatics is an emerging field that deals with the analysis and storage of bioimaging data. It has been speculated to revolutionize biomedical research, and its potential is well illustrated by a recent pan-European survey, which showed that 95% of life scientists consider image analysis to be very important for their work, but most consider it to be also the most difficult part of their work and the part lacking most in infrastructure and personnel.

A European-wide open access network of imaging centers, Euro-BioImaging, is currently in the process of being formed, and Finland has been elected to lead it, hosting its ERIC seat in Turku. Turku BioImaging also coordinates national and Nordic imaging networks and the upcoming Finnish Euro-BioImaging Node, where other partners are the universities of Helsinki and Oulu. Bioimaging is also becoming a major strategic focus area of both the universities in Turku, where it is increasingly attracting the attention of the local industry as well. Finland is one of the European spearheads also in bioimage informatics, and for instance one of the largest software tools available in the field, BioImageXD, has been developed in Finland.

It is evident that in the next few years the need for data storage and cloud computing services in imaging will increase substantially in Finland, due to for instance the Euro-BioImaging infrastructure and its requirements, but also due to local needs. It has been recently decided that in Finland bioimaging services will focus on full 3D services that cover both image acquisition and analysis. This means that amounts of data can easily become large. Functional nation-wide image data repositories and other bioimage informatics services will most likely be essential not only for providing the necessary bioimaging services for the Finnish life science community.

Cancer genetics

Cancer is a disease involving two unique genomes: germline, and that of the respective tumor. The rapid advances in genomic technologies are now enabling whole genome analysis of individuals and cancers. This will finally allow thorough dissection of germline and somatic genetic variation contributing to uncontrolled growth of cells. The Finnish Centre of Excellence in Cancer Genetics Research, lead by Acad. Prof. Lauri Aaltonen, aims to unravel the genetic components of human cancer susceptibility, and to develop computational methods to fully benefit from the massive influx of data from high-throughput, whole-genome scale experiments. The group and the whole biomedical research field are heavily investing in these emerging technologies. For instance single-cell approaches are currently transforming biomedical research and clinical practice. In these, noisy data is produced by multiple acquisition methods in high-throughput manner, requiring large and fast storage capacity and often special handling of sensitive (patient) data. Ultimately, the goal is to translate these findings into clinical benefits, such as novel approaches in cancer risk prediction, prevention, diagnosis, and treatment.

Many of these high-throughput techniques can be foreseen being put into practical clinical use, impacting healthcare on the population level. Computational science provides the methodology and tools that are essential in dealing with terabyte-scale data produced by modern high-throughput experiments in molecular biology. In addition to the massive quantity of data, pronounced complexity of data due to multiple heterogeneous sources must be considered. As the whole-genome sequencing and other high-throughput approaches such as imaging become mainstream, the bottleneck in cancer genetics is shifting from data generation to data analysis and interpretation. It is clear that the availability of high-throughput data from thousands of samples in the near future will necessitate the development of computational methodologies and present high demands for the available computational infrastructure.

Cellular signaling

All major diseases are related in one way or another to receptors and the dysfunction of cellular signaling. Carbohydrates and glycolipids and their effect on receptors and cellular signaling are a good example of the role computational studies play. Binding of ligands and the availability of active pockets are key to the biological function of receptors. Glycolipids may shield or expose these active pockets and have a significant effect on the binding of ligands, which may have direct link to the treatment of diseases such as cancer. Unfortunately, glycoconjugates attached to proteins are not visible to experimental techniques, even though they may have such a profound biological and medical relevance. Since they are not available in the known protein structures that are based on experimental data, computational methods are the only way to probe their relevance and to design further experimental tests to verify the hypothesis.

Prof. Ilpo Vattulainen (Tampere University of Technology) and his team studies these topics employing state-of-the-art molecular dynamics simulations. To push the understanding of the biological function and chemical reactions involved, the simulations need ever increasing accuracy and timescales. With the currently available computational resources in Finland, 10 µs simulations of drug-protein interactions are possible; to get accurate assessment of the interactions which can be directly compared with laboratory experiments up to 1 ms (hundredfold) timescales are needed. In addition to timescales, it is also important to perform a large number of simulations to get a representative sample of the phenomena. In order to be at the forefront of computational biological research, the computational resources available to Finnish researchers should be, if not comparable, then at least not lacking far behind those of the leading groups in the world. In practice this means supercomputing resources with fast compute cores and a high performance interconnect. The size of main memory is not a limitation.

Sequencing Initiative Suomi

The Sequencing Initiative Suomi (SISu project) lead by Prof. Aarno Palotie, is one of the largest human sequencing initiatives in Europe. Whole-genome and whole-exome sequence data from Finnish samples can be combined with decades of data gathered in the Finnish health and welfare sector, enabling breakthroughs in personalized health care. The researchers have already unraveled genetic components affecting growth and health of newborns, leading to new health recommendations, as well as repurposing of drug molecules to offer new hope for cancer patients. New genomic services are currently piloted in Finnish health care, e.g., the Kardiokompassi project, GeneRISK project and Tekes Digital Health Revolution project.

Biomedical research which depends/uses on large scale genomic datasets presents some distinct challenges for storage infrastructures and supercomputing. First, the basic genome analyses and processing steps read mapping, variant calling and recalling, require huge storage capacity (order of 100s of gigabytes per one human genome) and intensive and challenging I/O profile consisting of lots of random I/O and large write operations suggesting that fast local disks should be available. Additionally, analysis of new samples often rely on the use of large shared reference datasets and thus the storage and compute infrastructure should support the use of these within the cloud and/or federated between clouds. Finally, the ICT infrastructure for management and analysis of biomedical data must also be certified for handling sensitive information.

Materials science, chemistry and nanoscience

Efficient solar cells

Efficient and cheap solar cells can be constructed by using light absorbing molecules called chromophores or dyes. The efficiency of dye sensitized solar cells depends on the absorption spectrum of the chromophore and the ability of the chromophore to transfer the absorbed light energy and the corresponding electron to the electrode. The best chromophores absorb light in the whole range of the solar spectrum, from near-infrared radiation to ultraviolet. Prof. Dage Sundholm's group at the University of Helsinki studies such molecular light absorption processes computationally at quantum level. Molecular simulations of the light absorption processes require quantum chemical calculations of the interaction between light and the chromophore including the energy and electron transfer. Simultaneous calculations on all molecules taking part in these reaction steps are computationally very demanding since they involve excited states, dynamic energy and electron transfer processes, and interactions with molecules in the liquid phase or with e.g. a surrounding protein molecule.

The predictive power of quantum chemical studies depends on the employed level of theory. Although density-functional theory (DFT) is nowadays the "workhorse" of quantum chemistry, there are still many problems, such as some of the above mentioned, which require wavefunction-based high-level theories. Such methods scale only moderately and typically require huge amounts of memory (10-100GB per core) and their performance rely on efficient, low-latency run-time storage, typically ~TB's per job.

First-principles design of materials

Materials are at the heart of our society. Basically every new commercial product - be it related to health and environment, clean energy, heavy industry, information and communication technology and more - largely depends on novel, new, or improved materials. The value and range of applications of computational materials science is illustrated in the following brief examples:

- Artificial hips need metal alloys that are hard, biocompatible and have the same elastic properties and thermal expansion as natural bones.
- About 40% of the energy presently used by humankind is wasted as heat. In order to turn some of this waste heat into useful electricity we need thermoelectric materials with an efficiency about a factor of 2-4 higher than what we have today.
- Novel nanomaterials, e.g., carbon nanotubes in transparent and flexible electronics and nanoparticles and nanostructured microparticles as drug delivery systems.
- Novel batteries and novel catalyst materials that transform electricity into useful chemicals.

Identifying the best materials for a particular application is a significant challenge. Computational materials science uses atomistic simulations and multi-scale modeling to investigate materials at the atomic level, providing insight into material properties and enabling the design of entirely new materials to meet specific requirements. By calculating the key characteristics of materials, including candidate materials that are not even known to exist so far, one can screen a huge chemical-compound space by grouping materials on the basis of particular characteristics. This, in turn, facilitates the identification and selection of optimal materials for targeted applications. Researchers at Aalto University - the groups of Academican Risto Nieminen and of several other senior scientists - will put this approach into practice within an EC-funded effort by developing big data analytics for materials science.

The research towards new materials via the route described above will pose heavy demands on the available e-infrastructure and heavily utilize the full range: throughput computing, supercomputers and storage. Key materials science codes used in the research, such as GPAW and CP2K, require supercomputer resources that support runs with up to ten thousand core each. Other key applications, such as VASP and Quantum Espresso, need hundreds of cores per run with a lot of memory per node and very fast local storage. In addition, it will require cutting-edge storage capabilities: vast amounts of fast storage as well as long-term storage with a capability for data mining and global sharing of data.

Metal nanoparticles in biological processes

Research at the nanometer-scale in physics, chemistry and biology has brought these core disciplines together and has opened fascinating views to future applications of nanotechnology in electronics, material sciences, medical diagnostics and treatments, as well as clean technology. Technological motivations spur the drive to miniaturize working devices and components, and to optimize energy consumption and operation speed. Nano-structured materials are expected to be lighter but stronger than conventional ones, be more energy efficient, and to have programmable functional properties. Nanoscience research is expected to have a great impact on biological and medical research as well.

From a theoretical and computational point of view, understanding phenomena and processes in the nanoscale matter and molecular systems brings challenges to incorporate the quantum nature of atomic and molecular interactions to large-scale temporal and spatial dynamics in the proper environment. Solving these fundamental problems and achieving breakthroughs requires substantial knowledge database of dynamical material properties at the interfaces between physics, chemistry, and structural biology. In order to make inorganic nanostructures useful for and even conceptually compatible with biology, dynamical processes at the inorganic/organic interface must be addressed and understood. This is a significant current challenge in nanoscience frontier research.

Acad. Prof. Hannu Häkkinen is one the biggest users of CSC's computing resources. One of the recent research interests of his group is to study how nanometer-sized metal clusters are interacting in biological environment (viruses, proteins). Such an example is the joint experimental and computational study of the interaction of thiol-stabilized gold nanoclusters with capsid proteins of enteroviruses and cell membranes. They have managed to demonstrate a unique procedure for achieving robust, direct, site-specific labeling of enteroviruses for future structural studies of virus uncoating, for electron microscopy investigations of virus entry and pathways into cells in vitro and in vivo.

This field is computationally very intensive, and has typically been able to fully utilize the largest available resources at a given time. The researchers use massively parallel programs requiring a tightly-coupled supercomputer. The success of Finnish researchers in this internationally highly competitive field has been partly due to available supercomputing resources. To secure the success of this research area in the future it is essential that the national supercomputing resources are kept up to date to maintain the competitive edge, and that the service has sufficient resources for a top-tier service layer, including support for the development of the materials science code GPAW, which is used and developed in all computational nanoscience groups in Finland.

Earth sciences

Atmospheric feedback mechanisms

One of the greatest challenges facing humankind within the next hundreds of years is the climate that is warming at a quick pace. The cause for the warming is thought to be the emission of greenhouse gases, mainly carbon dioxide and methane, but research is also being carried out to investigate other causes and feedback mechanisms. In a feedback process in general, changing one quantity changes a second quantity, and the change in the second quantity in turn changes the first. Positive feedback amplifies the change in the first quantity while negative feedback reduces it. The term "forcing" means a change which may drive the climate system in the direction of warming or cooling.

Forcings, feedbacks and the dynamics of the climate system determine how much and how fast the climate changes. In study of climate change, understanding feedback mechanisms is crucial for understanding global warming, because they may either amplify or diminish the effect. Therefore these are key for determining the climate sensitivity and future climate. Many both positive and negative feedback mechanisms have been identified, but their relative strengths are not known well enough and many of them are still missing from numerical climate models. Even the net outcome of the mechanisms has not been established, i.e. whether the global warming will be just amplifying itself and accelerating, or if there exists a thermostat in Earth's climate.

This case is being addressed in the Academy of Finland Center of Excellence in Atmospheric Science, lead by Acad. Prof. Markku Kulmala. This research group is one of the leading groups in climatology in the world, and has been on top of citation ratings for several years. The research within this CoE lies on three pillars: various environmental measurements, satellite data and multiscale modelling. The multiscale models research start from nanoscale (quantum chemistry), reaching out towards global atmospheric models. Assessing the significance of various feedback mechanisms and finding new ones requires following improvements to the multiscale climate models:

- more physics incorporated into the models
- higher resolution; denser spatial grids and shorter time steps
- longer simulations, i.e. more time steps taken.

Achieving these require orders of magnitude higher computational capability than currently available. This research would directly benefit from installing a larger supercomputer, and respectively this research would be badly harmed by a non-sufficient or competitive computational infrastructure. They estimate their data-processing needs (fast storage, data-intensive computing, capacity for storing and sharing the experimental data) to increase by a factor of 10,000 during the 2017-2022 timeframe.

Earth-system model development

Active work is being done within the Intergovernmental Panel on Climate Change to assess scenarios for climate change based on different emission levels to aid decision makers. In Finland the Finnish Meteorological Institute (FMI), the University of Helsinki, and the University of Eastern Finland, have been active in providing modelling support for the fifth IPCC report. This work is being continued within the Coupled Model Intercomparison Project (CMIP), an international collaboration for deciding what and how coupled models will be used for the sixth IPCC report. FMI will participate in the modelling work and CSC will provide one of the data nodes. The CMIP project, as well as the feedback from the science community in general is pushing researchers to move to full Earth system models (ESMs) with coupled air, ocean and land models. This is driving a large increase in the computational demands and also requires new models. EC-Earth model is developed by a large European consortium including 27 partners from 10 countries, with FMI being the Finnish representative. EC-Earth couples the IFS atmosphere model, Nemo ocean model, LPJ-GUESS vegetation model and TM5 chemistry model. The Finnish contribution is closely linked to improving aerosol modeling in TM5. The computational and data demands are strongly increasing for global models:

- Running multiple coupled models instead of individual models

- Simulations moving to higher resolutions
- Longer time-periods are simulated
- Ensemble simulations to gather statistics.

Another research topic are regional simulations with much finer resolution and physics. Here the community in Finland is moving to Harmonie which is also used for numerical weather prediction in Finland.

Single full-Earth system simulations require up to 2000 cores, while ensemble simulations comprising multiple simulations run in parallel with long simulated times multiply the required computational resources. The future ensemble simulations will require higher computational capability than currently available, storage space and I/O performance are also a major bottlenecks. Another challenge is the need for long term storage to fulfill community requirements for storing simulations for up to 5 years.

Changing climate and sea-level rise

Understanding and predicting the behaviour of the Antarctic ice sheet on decadal to centennial time scales is of considerable importance for global policy making, mainly due high uncertainty in the AIS contribution to sea level change. Significant regions of the Antarctic ice rest on bedrock below sea level and may be vulnerable to marine ice sheet instability, a mechanism for rapid and potentially irreversible discharge of ice. Due to large uncertainty as to which basins are vulnerable to the ice sheet instability, and how rapid their unstable retreat might be, even the sign of the contribution of the Antarctic ice to sea level rise remains uncertain according to the fifth assessment report of the Intergovernmental Panel on Climate Change. The group of Prof. John C. Moore, University of Lapland, is going to conduct long term, high-resolution enhanced model simulations of the Antarctic ice sheet. Simulations will be carried out using the Elmer/Ice package and certainly exceed in size the current ice sheet simulations, as marine ice sheet dynamics demand extremely high resolutions, and therefore require mid/large-scale supercomputer resources, typically runs with a few thousand cores. These simulations will produce state-of-the-art results in numerical glaciology.

Unraveling fundamental space weather processes

The most familiar manifestation of space plasma is the aurora. They are caused by space weather, the complicated chain of plasma physical events starting from solar eruptions and ending in energetic particle precipitation into the Earth's atmosphere. Space weather can harm technological systems or human life through, e.g., geomagnetically induced currents that affects ground-based power grids, and deterioration of radio and radar signals affects positioning and aviation in the Arctic and Antarctic regions.

Large-scale global models are a key in modern space research tying together the local in situ phenomena measured by spacecraft missions, and their effects and causes far away from the observing position. The spatial domain controlled by Earth's magnetic field is vast (millions of kilometers at a direction), while the physical phenomena are determined by effects occurring in 10-100 kilometer scale. Therefore, the global models need to cover enormous volumes with extremely good resolution. With the aid of a European Research Council Starting grant awarded in 2007, Prof. Minna Palmroth's group Finnish Meteorological Institute (FMI) has developed the world's only hybrid-Vlasov based global space physics model called Vlasiator. Compared to other approaches Vlasiator offers global space physics in unprecedented detail with very low noise, allowing new insight in the grand challenges of space physics.

The simulations are highly parallel, and have been performed on tens of thousands of cores on CSC's machines. Even so, to go from two dimensional simulations modelling the Earth systems on a plane, to realistic three dimensional simulations an order of magnitude larger resources are needed. For Vlasiator 64 GB of memory per node is sufficient, but 128 GB would be better. The fast storage system needs to provide capacity on the order of hundreds of TB with very good bandwidth, to enable restarts to be written in reasonable time.

Emerging fields

Machine learning

The modern society is being radically transformed by the global megatrend of digitalization of all fields of science, industry and the rest of the society. In the core of this revolution is learning from data, alternatively called machine learning, studied e.g. in Acad. Prof. Samuel Kaski's group in Aalto University. Finland is particularly strong in this field which combines computer science and advanced statistics. The field is extremely computationally intensive, and in particular requires coupling massive data sets to large-scale computational resources. Utilization of this significant asset Finland has, and staying ahead in the development of the extremely rapidly advancing field, requires sufficient computational resources of which a supercomputer is a key part.

Natural language processing

Natural language processing (NLP) is a cross-disciplinary field that is concerned with enabling computers to process and generate speech and text. This multidisciplinary field intersects with linguistics, computer science and machine learning. NLP research in Finland is done, for example, at the University of Helsinki (Department of Computer Science, Department of Modern Languages) the University of Turku (BioNLP Group). The research in Turku focuses on the various aspects of natural language processing, ranging from corpus annotation to machine learning theory and applications. The main application areas are centered on biological, biomedical, and clinical texts. This technology has been used at the University of Tampere to identify and track pain levels in cancer patients based on existing electronic medical records.

NLP is able to efficiently utilize capacity systems. The main tools do not scale to multiple nodes but are able to utilize reasonable compute resources by running multiple jobs in parallel. I/O can be a bottleneck since the methods typically need to work on large amounts of small files which is not performing well on parallel file systems. GPUs are starting to show up in NLP, because many machine learning methods used are moving into deep learning (neural networks) for which GPUs are very suitable. This field may be able to utilize them very well in a cluster environment. Another key algorithm in NLP is Hadoop map/reduce-style processing which could be very useful if supported by the next generation systems. Staging local copies of large public data sets is extremely useful for NLP applications. There are significant advantages in storing big data sets centrally and near adequate computing resources as opposed to researchers having to deal with copying, storage and backup and public availability of such large data sets themselves. Language data sets are relevant not only to NLP but also to social science in general. Central storage furthermore facilitates repeatability of research.

A summary of user requirements

CSC serves a wide range of scientific fields and groups, and this is reflected in the requirements for compute infrastructure. Based on the user survey (see Appendix 1), the birds-of-a-feather sessions and the interviews of customers it is clear that the needs vary from databases and cloud environments for data intensive computing to demanding parallel simulations requiring tens of thousands of processor cores.

Support for compute-intensive research

Users from almost all fields of science required a competitive supercomputer with a high performance interconnect able to maximize the scalability of their applications. In the user survey, approximately 40% of the respondents were planning to use thousands of cores, or even tens of thousands of cores for their production jobs, which is a scale dedicated supercomputers are required. These users typically did not need local disks, but did need access to a fast parallel file system for I/O. Many would also have benefited from more memory per node than currently available. Approximately 47% users requested hundreds of cores, and many could also easily utilize the supercomputer, except where very high amounts of main memory or local disks is needed.

On the other hand, a large user segment (51% in survey) needed at most a few hundred cores for their computational jobs. Additionally many users incorporate pre- or post-processing steps in their workflow, which do not support large scale parallelism. For these needs, a more scale-out-style resource may be a more cost-efficient solution as compared to a tightly-connected supercomputer.

Support for data-intensive research

The amount of data researchers need to process is increasing massively, due to digitization of data, new experimental methods and due to ever larger simulations. To support this more emphasis needs to be put also on supporting data intensive computing. For data-intensive research performance is typically not limited by numeric floating point operations as in traditional simulations, but by how efficiently data can be loaded from storage and random-accessed in memory. This in practice entails larger amounts of main memory (see memory recommendation), large and performant storage (see storage recommendations) and also support for cloud environments.

Cloud resources

CSC's cPouta compute cloud, and ePouta compute cloud for sensitive data, have gained hundreds of users in last two years. It is expected that the usage will continue to grow, of the questionnaire responders some 40% are considering using cPouta in the future. Cloud resources enable one to support customized environments to ease the installation of complex applications and workflows.

Docker containers providing ready-to-run packages of software may gain momentum in the near future and should also be supported. These have less overhead than virtual machines, and can be run in a normal compute cluster or supercomputer.

Completely new ways of using computational systems can also be developed, e.g. providing web services to researchers tightly integrating to large data sets or setting up customized development environments for software development. Many of its current users comes from bioscience, and the ePouta HPC cluster will be a very important component in enabling the community to handle sensitive genome data.

The cPouta capacity has already attracted several paying customers. We anticipate that research institutes will buy capacity from ePouta where they can integrate their existing resources via lightpath or MPLS network technology.

Increased amount of large-memory nodes

A common request from users is the need for more main memory. In the survey (question 6) 66% of the respondents reported that they need 128 GB or more main memory per node, while 16% of the respondents require more than 1 TB or more. This need exists for highly parallel simulations, e.g. galaxy simulations and space weather simulations, as well as for small and medium sized problems in e.g. material science. The biggest user of cloud resources is the bioinformatics community, and they have a great need for large main memory for many data intensive analysis methods and for gene sequencing.

GPU resources

The majority of users did not consider GPUs to be currently that useful for their research, but at the same time there is a small but active community in Finland developing GPU based applications. In the survey 53% of the respondents reported that they are ready to utilize GPUs for computing, and 33% reported that they are ready to modify their application to utilize GPUs. For some methods, e.g. deep learning, GPUs provide a superior platform in terms of performance and total cost.

On the other hand, in the free text comments many users specifically warned against investing too heavily into GPUs. The most cited reasons were: 1) No GPU version of their main application exists, 2) only a subset of features are supported in the application, 3) it is time-consuming and technically challenging to program or modify applications, and 4) GPU based applications have limited scalability and/or performance.

It is clear that investing into a significant size GPU environment needs to be matched by a significant investment in training and porting assistance at CSC.

Storage

Larger computational resources need to be complemented by larger and faster storage systems. Many questionnaire responses and BoF seminar discussions highlighted large storage needs in terms of capacity and performance, both for compute intensive simulations as well as for new data intensive computing methodologies. The HPC systems will need a fast parallel file system, most likely based on Lustre. Using burst buffers based on NVRAM or SSD disks the effective bandwidth may be increased dramatically. The storage for cloud resources will most likely be based on the Ceph storage technology, which also supports object storage.

Many fields of science, e.g. quantum chemistry and bioinformatics, need fast local disks for large sets of temporary data that do not fit into the main memory. Fast SSD or NVRAM based local disks are needed for a large enough number of nodes.

Archival storage is needed for storing large data sets. In many fields of science the raw simulation data of publications should be stored for many years. Services for sharing large data sets publicly is also needed and IDA covers this need mostly.

Advanced data management by e.g. using data identifiers and ontologies will allow long-term data management according to the standards. The practices and the services are driven by various open science and data initiatives (e.g. ATT, EUDAT).

Object Storage service enables for advanced users storing and sharing large data sets and to add metadata. The data owner can control data access and the service interfaces are open to Internet.

Especially in life sciences, sensitive data is common because of human medical and health information which require tools, storage and compute platform designed for raised security level. Sequencing and various imaging devices produce a lot of data and require large computable accessible storage as well as long term storage capacity. Also some reference datasets can be accessed only with permission of data owner. This may lead multiple copies of the same read only data unless there is an advanced data management solution to enable access to common data for the authorized users.

Digital preservation solutions for research data (PAS) is out of the scope of this procurement.

Recommendations

Level of funding

Internationally and nationally, the installed compute capacity has historically increased in performance by a factor of 1.8 each year. During the last two years, this progress has clearly slowed down. This slowdown is at least partly due to the economically challenging times worldwide, and may be only temporary.

Using a conservative estimate of performance increasing by a factor of only 1.4 each year, means that to maintain the international ranking the total installed compute capacity should be 7 PFlops in 2017 and 13.5 PFlops in 2019. Based on public processor roadmaps one can estimate that the peak performance per node is approximately 3 TFlops in 2017, assuming each node is configured with regular CPUs. Over the years the price of each node has been fairly stable, and one can thus estimate that reaching the desired peak performance in 2017 will require an investment of at least 15 million euros. The pricing and technology is less clear in 2019, but one can estimate that an additional investment of at least 9 million euros will be needed in 2019.

The compute capacity needs to be coupled to a performant and large storage system. The size of this storage should be 10 - 20 PB in the first phase, based on the historical increase in storage capacity at CSC, the storage requirements of the customers, and storage capacity installed in other systems. To reach sufficient performance significant investment needs to be done into solid state disks (SSD or NVRAM based), that are able to buffer the slower large file systems. The cost of this storage is at least 5 million euros in the first phase, with an additional upgrade in 2019 to match the larger compute capacity.

These estimates give a total cost of 30 - 35 million euros to be invested in 2017 - 2019, giving an infrastructure that is relevant until 2021. Naturally this estimate is very rough, and the exact price depend on the dollar exchange rate, technological development and the actual configuration of the infrastructure. This number is slightly larger than last time (25 million euros), and is based on the need to maintain the same level of international competitiveness, as motivated in the section “The case for the next-generation e-infrastructure”.

The ideal situation would be that the infrastructure is funded using a total cost of ownership model, where all costs would be covered: investment, electricity, service contracts, and work of CSC system administrators. This funding would cover the true cost of the system over the expected life-time of four years. In that scenario funding level should approximately be double the investment as argued above i.e. 60-70 million euros.

Recommended compute infrastructure

Heterogeneous infrastructure

To meet the user requirements detailed in the previous chapter, we recommend that the infrastructure comprises two systems, or two specialized partitions. The infrastructure should support the high productivity of the end user across their complete workflow. This involves, among other things: efficient data management, rich set of developer tools, ease-of-use and high reliability. This should be taken into consideration in both individual hardware and software investments as well as the overall architecture of the HPC ecosystem.

Capability system	A supercomputer that supports jobs that utilize hundreds, thousands, or even tens of thousands of cores. The user community for these kind of machine can adapt to more specialized hardware, but it is recommended that this machine should be programmable with an MPI + OpenMP approach. The storage performance should be sufficient to support the computational performance, and burst buffers should be incorporated to cache disk traffic to the fast parallel file system. The amount of memory per node should be increased from Sisu up to, e.g., 128 GB. Docker containers should also be supported, if possible.
Capacity system	This system caters to the users who need at most a few hundred cores per job, thus the interconnect does not need to be as efficient. These users often are not able to decrease memory consumption through distributed parallelism, thus the main memory per node should be greater than for the capability system. It should also provide a wider range of specialized nodes for different use cases: high-memory nodes (more than 1 TB), nodes with fast local disks and a sizeable portion of nodes with GPUs. Docker containers should also be supported via the regular queueing system. Finally, this system acts as a platform for cloud provision.
Cloud resources	The cPouta and ePouta cloud resources should be further developed, and the hardware of the capacity system should be specified in view of the needs of the cloud users. These include large main memory and efficient work storage.
Storage	The amount of storage should be significantly greater than in the current infrastructure. Growth estimates are exponential and thus computable storage needs are likely to exceed 10 PB in 2017. Not only the capacity and performance but also usability and the data management features are important.

Number of installation phases

Traditionally CSC has installed systems in two phases. The benefit of two phases is that it takes some time for customers to fully utilize a significantly larger resource. By installing the machine in two phases the unused resources may be smaller, and the final configuration can be more performant due to technological development. This can also allow some flexibility in tuning the system during the user requirements. At the same time it may lead to worse bids due to uncertainties in extrapolating benchmark results to the second phase architecture. Costs may also increase if the first phase hardware have to be discarded. The decision of the number of phases needs to be carefully considered by the procurement project, basing on vendor roadmaps and pricing indications.

Part of the funding should be reserved for incremental upgrades in response to capacity needs that emerge during the lifetime of the system (for example, nodes with specialized CPU or a large amount of memory), independently from the number of main phases.

Installation time

The installations of both the supercomputer and cluster have traditionally been done simultaneously. The drawback is that the installation time is disruptive for customers, and also puts a strain on CSC's human resources. One could decrease disruptions and utilization gap by installing them at separate times. The schedule is dependent on the technology roadmaps, as well as the utilization rate of the various systems. It may be beneficial to install the cluster replacement first; Taito is significantly smaller than Sisu and the replacement for Sisu will most likely be installed in the same datacenter as Sisu leading to a long service break.

Next generation infrastructure should be available to the entire academic research community in Finland

The computing resources supplied by CSC have been offered to researchers in Finnish universities and polytechnics, financed by the Ministry of Education and Culture. However, there is a considerable need of high-performance and data-intensive computing capacity in state-owned research institutes of other ministries.

At the moment, their use of CSC's resources is possible only when agreed and paid separately by the research groups themselves. Among such institutes are, e.g., Finnish Meteorological Institute (FMI), National Institute of Welfare and Health (THL), Natural Resources Institute Finland (LUKE), Environmental Administration (SYKE), and VTT Technical Research Centre of Finland. The recent document Enterprise Architecture of Scientific Computing Resources and Services (*Tieteellisten laskentaresurssien ja niiden oheispalvelujen kokonaisarkkitehtuuri*) recommends to eliminate such barriers, which are artificial in the scientific community itself and hinder collaboration between computational research groups in Finland.

CSC's computational resources should be a national resource used by all researchers in the public sector on the same conditions. The centralization of heavy computing capacity needed by all parties is cost-effective and saves public money.

A proper operational model should be made up to enable the use by researchers in the state research institutes. The availability of sufficient computing capacity would strengthen the position of Finnish research groups in international consortia.

Conclusions

1. **The infrastructure should be upgraded in 2017**
   To ensure the international competitiveness of Finnish research and to enable the scientific grand challenges to be solved.

2. **The funding for infrastructure investments should be 30-35 million euros in 2017-2019 or 60-70 million euros when total life-cycle costs of the investments are included.**

   This ensures the systems provide the same level of international competitiveness as in history in the years 2017-2021. The high quality of Finnish computational sciences would easily justify even larger investments.

3. **The infrastructure should address the diverse needs of the scientific community in Finland**

   A common theme for all fields of science is the major increase in the data requirements, at the same time significant computational capabilities are needed for simulations and data-processing. The user community is also expanding as ever-more fields of science have a need for computational and data services. Cloud computing and application distribution platforms enable users to manage custom complex software stacks. Specially designed platform and isolated environments increase the security to process sensitive data. The computational needs for simulations range from specialized capability systems enabling large parallel runs, to cost-effective capacity systems enabling on large numbers of smaller runs.

4. **The next generation infrastructure should be available on same terms for all academic research in Finland**

   A substantial share of scientific research is being carried out in state research institutes and in university hospitals, which contrary to universities are not able to utilize CSC resources free of charge. This barrier is artificial, and it is detrimental to the competitiveness of the Finnish scientific community.

Author details and acknowledgements

This report is commissioned and published by CSC – IT Center for Science Ltd, 2015.

Corresponding author	Dr. Pekka Lehtovuori, Director, Services for Research CSC – IT Center for Science pekka.lehtovuori@csc.fi, mobile +358 50 3819 723 www.csc.fi
Editors	Dr. Sebastian von Alfthan, Senior Applications Specialist Dr. Pekka Manninen, Senior Applications Specialist CSC – IT Center for Science
Working group at CSC	Dr. Juha Fagerholm, Dr. Jussi Heikonen, Dr. Sami Ilvonen, Jarno Laitinen, Maria Lehtivaara, Olli-Pekka Lehto, Dr. Martti Louhivuori, Dr. Kimmo Mattila, Dr. Nino Runeberg, Dr. Peter Råback, Samuli Saarinen, Dr. Seija Sirkiä, Olli Tourunen, Dr. Thomas Zwinger, Dr. Jan Åström

Acknowledgments	Prof. Sampsa Hautaniemi Prof. Hannu Häkkinen Assoc. Prof. Peter H. Johansson Prof. Olli A. Jänne Dr. Pasi Kankaanpää Acad. Prof. Samuel Kaski Acad. Prof. Markku Kulmala Dr. Taina Kurki-Suonio Adj. Prof. Perttu Lantto Dr. Susan Leerink Dr. Tomas Linden Dr. Timo Miettinen Academician Risto Nieminen Prof. Kai Nordlund Res. Prof. Minna Palmroth Dr. Joni Pietikäinen Prof. Taina Pihlajaniemi Dr. Esa Pitkänen Prof. Kari Rummukainen Dr. Joonas Sorvari Prof. Dage Sundholm Prof. Ilpo Vattulainen Prof. Jan Westerholm

Appendix 1 - User survey results

The survey was sent to 472 persons in total, out of which 76 completed the questionnaire. These persons comprise the top users of CSC's compute, data and cloud resources, as well as a sample of all CSC users in an attempt to also get feedback from regular small-scale users. In addition to this, it was also sent to persons representing research institutes who currently cannot use CSC's resources, at e.g., Finnish Meteorological Institute (FMI), State Technical Research Centre (VTT), National Institute for Health and Welfare (THL).

1. Field of science (detailed list)

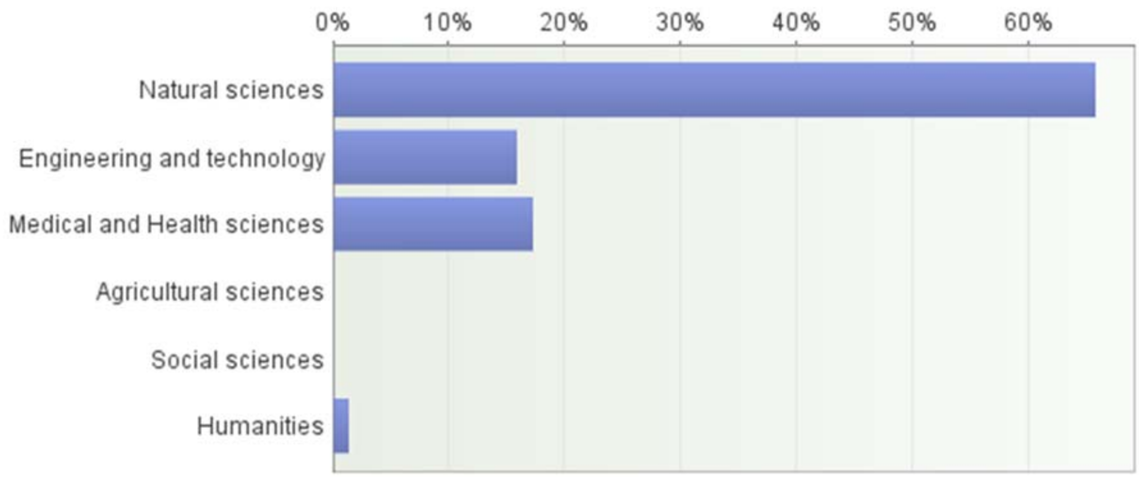

2. When I use CSC's resources I use my main application(s) mainly through the following methodologies (select all that apply)

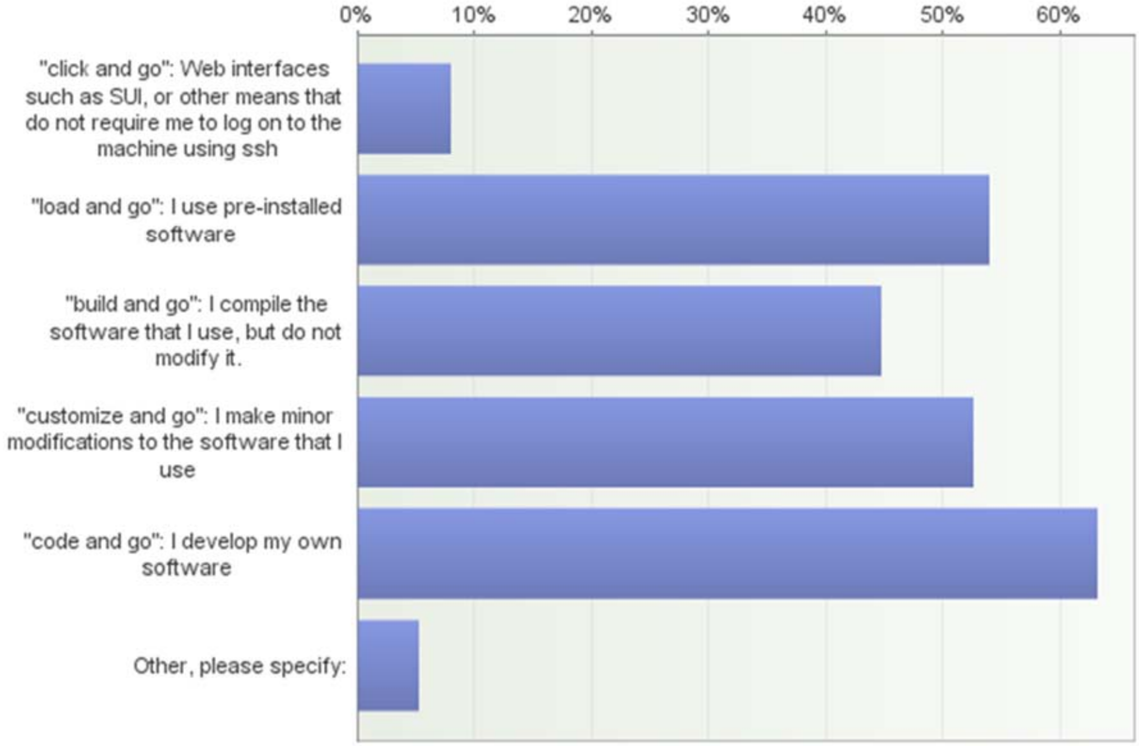

Avoimet vastaukset: Other, please specify:

- cPouta usage
- teach and go
- I use CSC's cloud compute service
- Prefer the Pouta-model

3. Which of these modes of using the system does your research mainly depend on (select all that apply)?

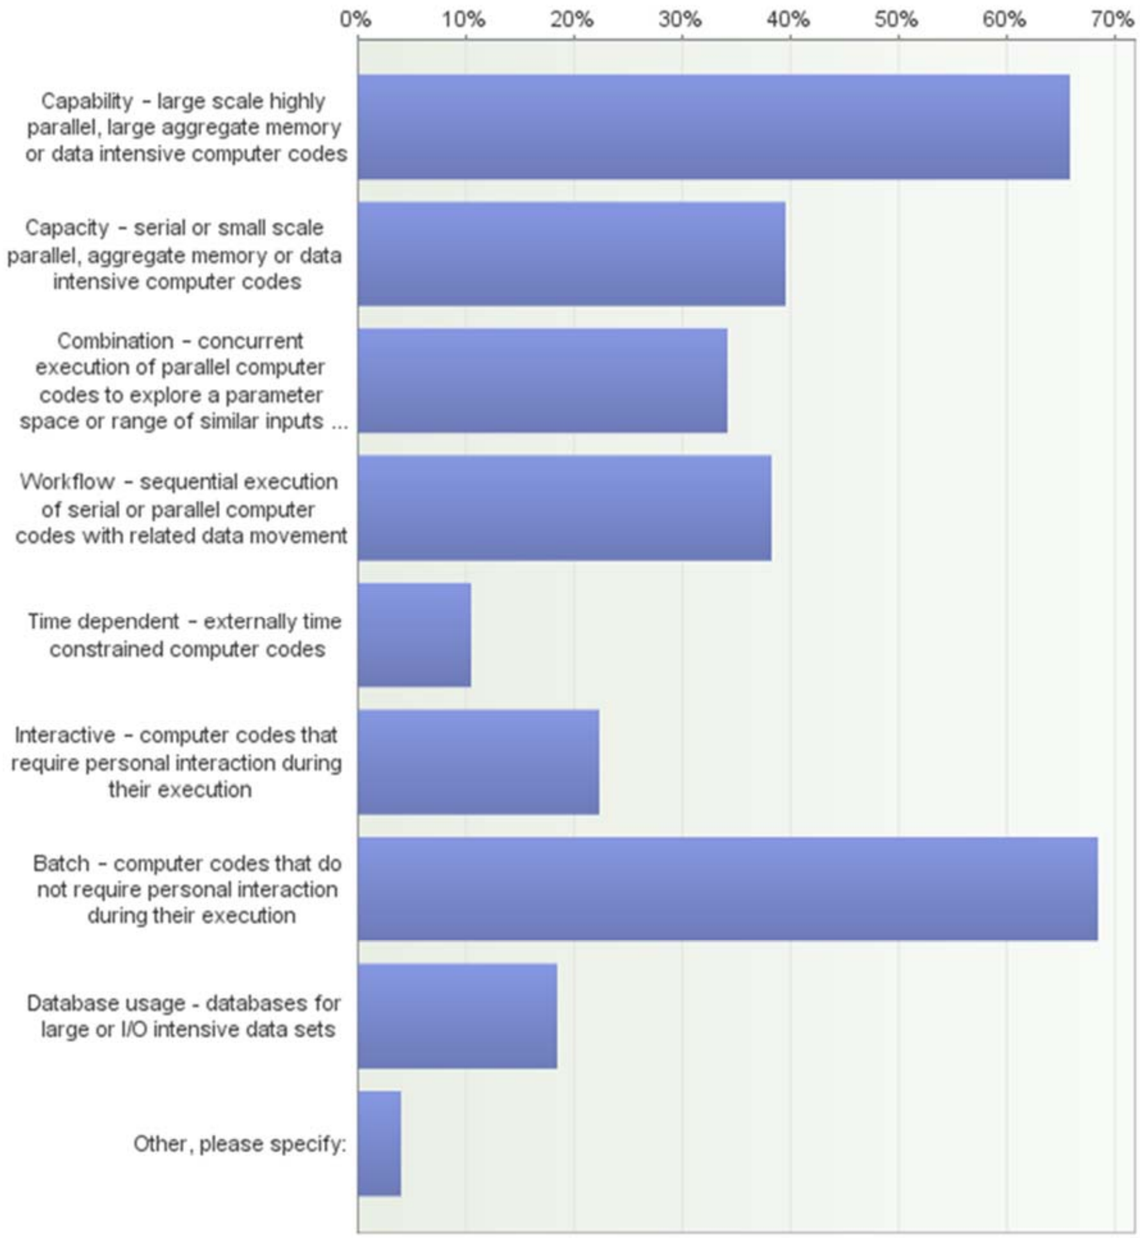

Avoimet vastaukset: Other, please specify:

- large memory, short time (<1 day)
- Embarrassingly parallel computation with very large storage and RAM requirements

4. My estimated CPU hour consumption per year in the next few years is

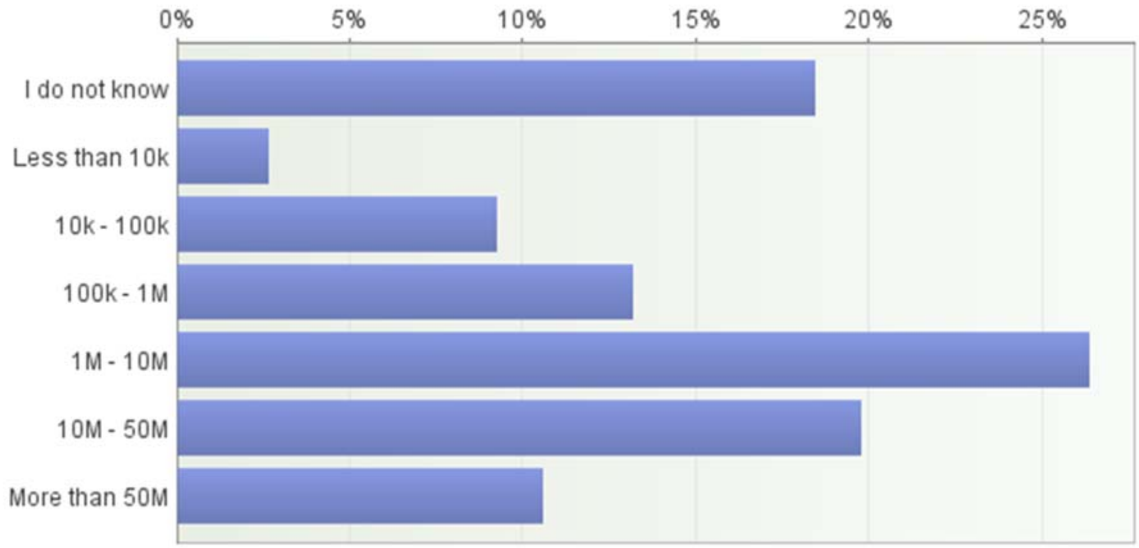

5. The largest runs I intend to perform in the next few years require

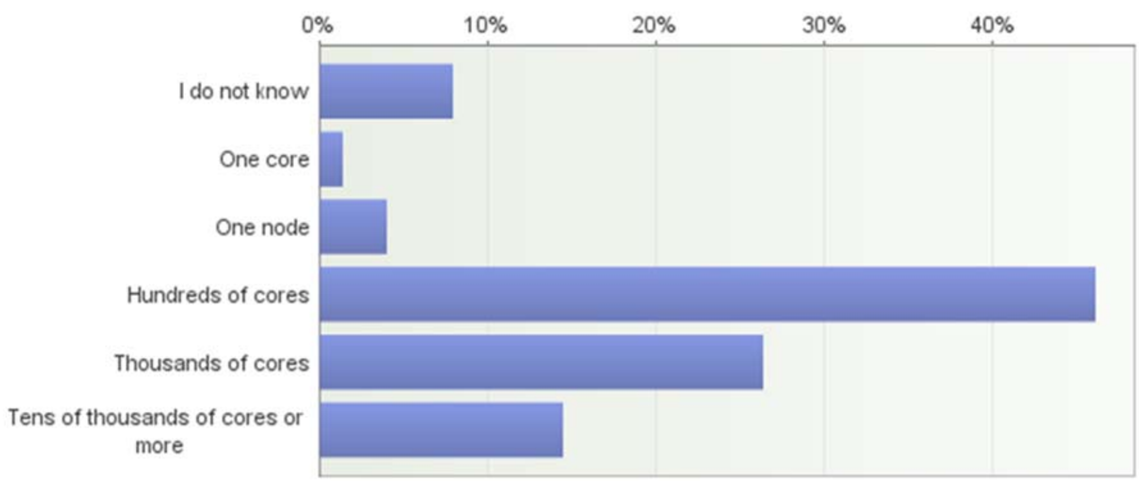

6. The largest amount of memory I will need per node (node comprising, e.g., 2 x 20 compute cores) in the next few years is:

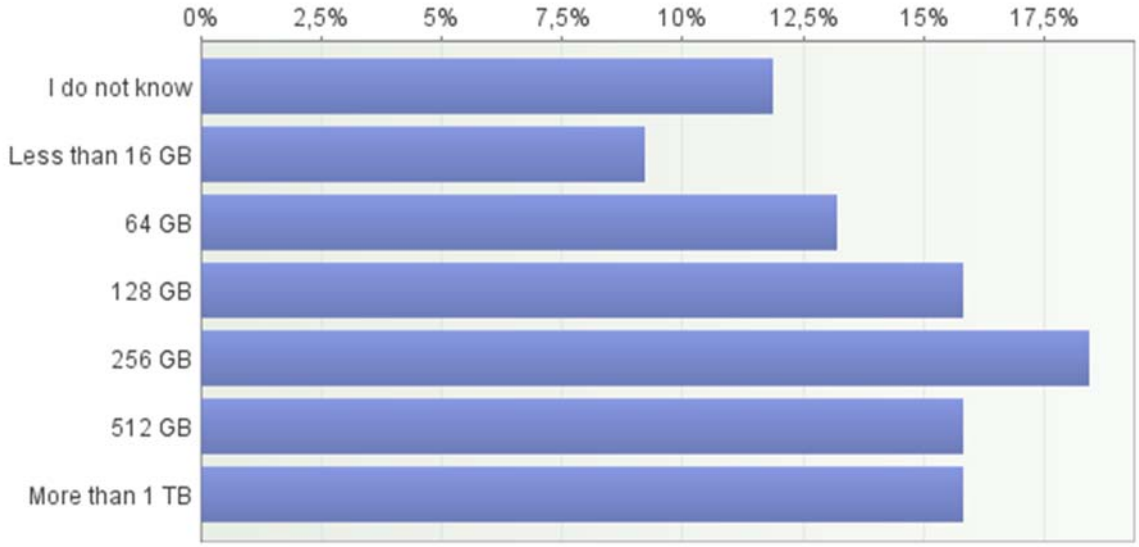

7. Please estimate how much capacity you will need in the next few years for data that is stored on a fast parallel file system which does not guarantee long term storage. This corresponds to the current work disk.

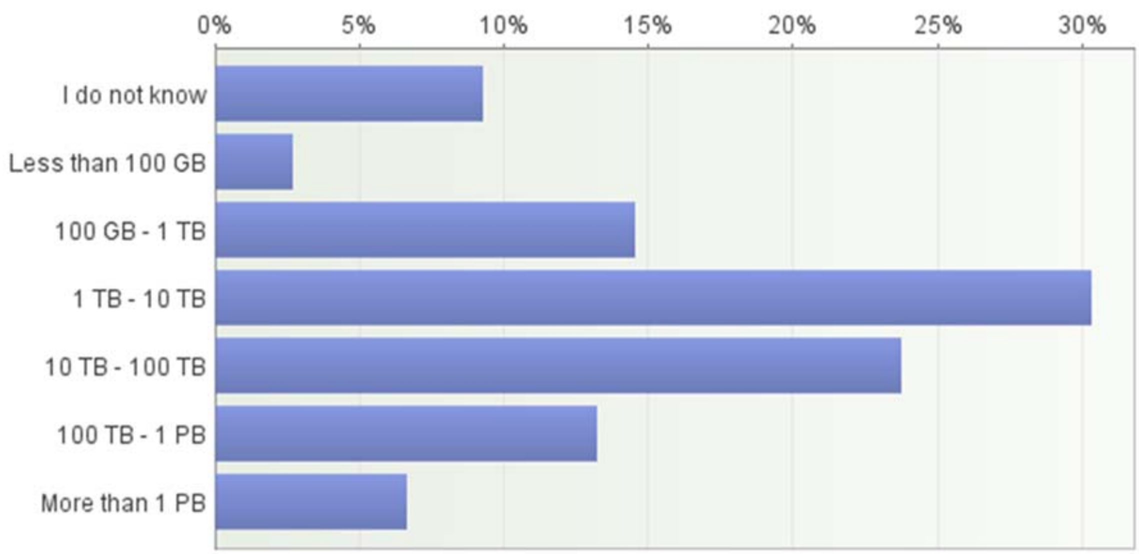

8. Please estimate how much capacity you will need in the next few years for data that is stored for the duration of a typical computational project (up to a few years). This corresponds to the current project disk area.

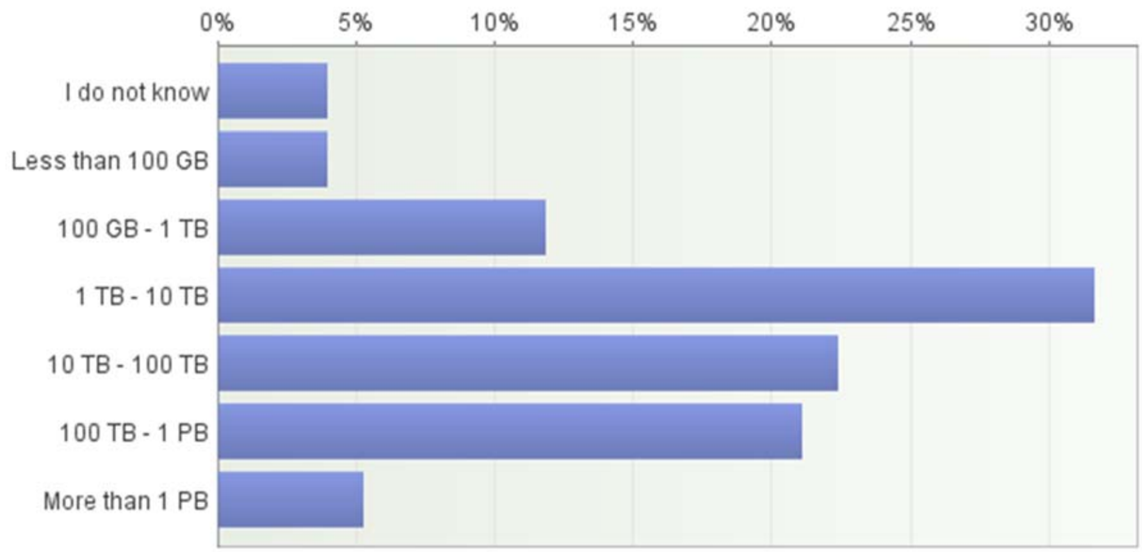

9. My estimated need for permanent data storage (archive) in the next few years is

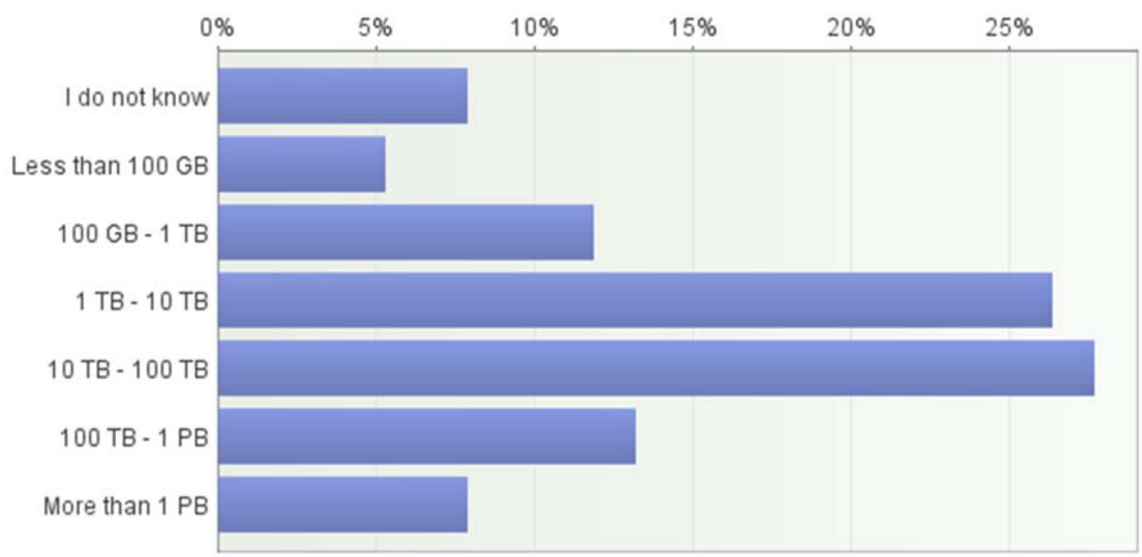

10. The main bottleneck(s) in my applications and workflow are (select all that apply)

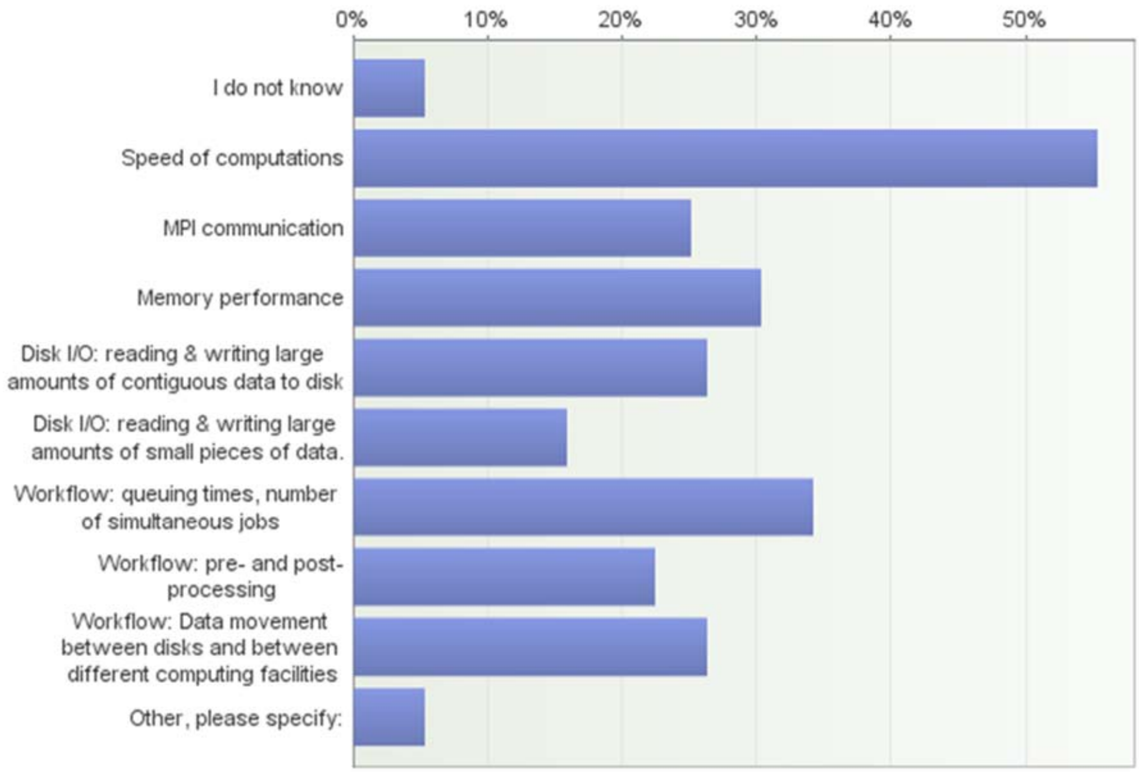

Avoimet vastaukset: Other, please specify:

- Queuing times due to low number of huge mem nodes
- Memory size
- limited cloud capacity
- Graphics processing

11. To improve the throughput of my research I would be prepared to do the following changes to my applications (select all that apply)

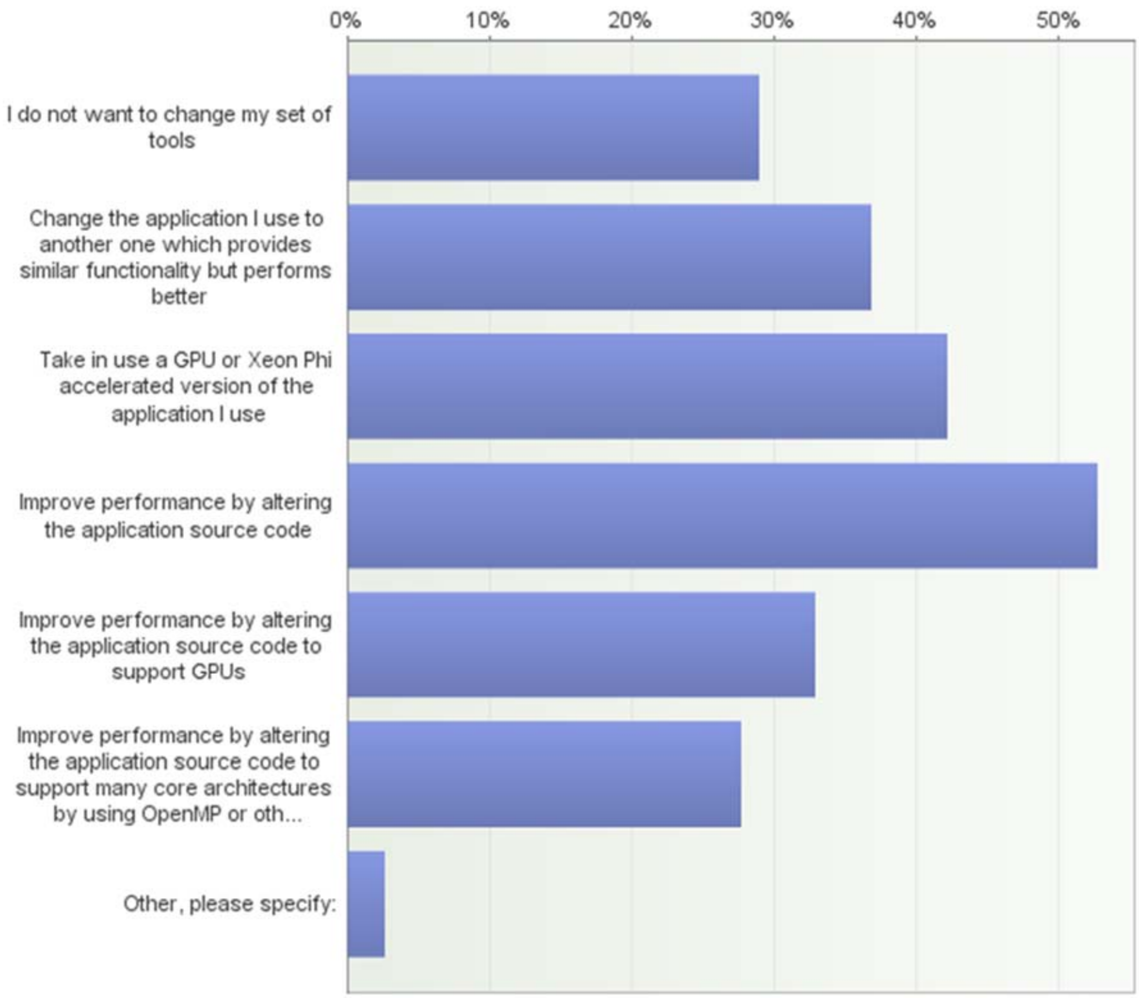

Avoimet vastaukset: Other, please specify:

- The application codes are constantly being developed and optimized.
- I do can not want to change my set of tools

12. New technologies; If available I am ready to utilize these (select all that apply)

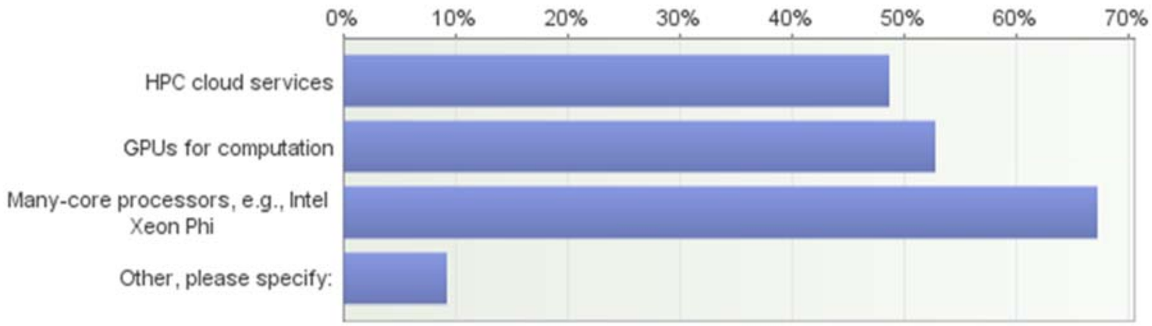

Avoimet vastaukset: Other, please specify:

- Data intensive clound nodes (lots of ssd or similar)
- If the toolset would work well with GPU/Xeon Phi we would be ready to use them. I am not yet convinced that this is the case.
- hadoop type commodity clusters
- I may use GPGPUs in the future, but not ready to utilize them yet.
- ARM resources if they become competitive with x86-64 resources.

13. I have plans to use CSC's Pouta cloud for (select all that apply)

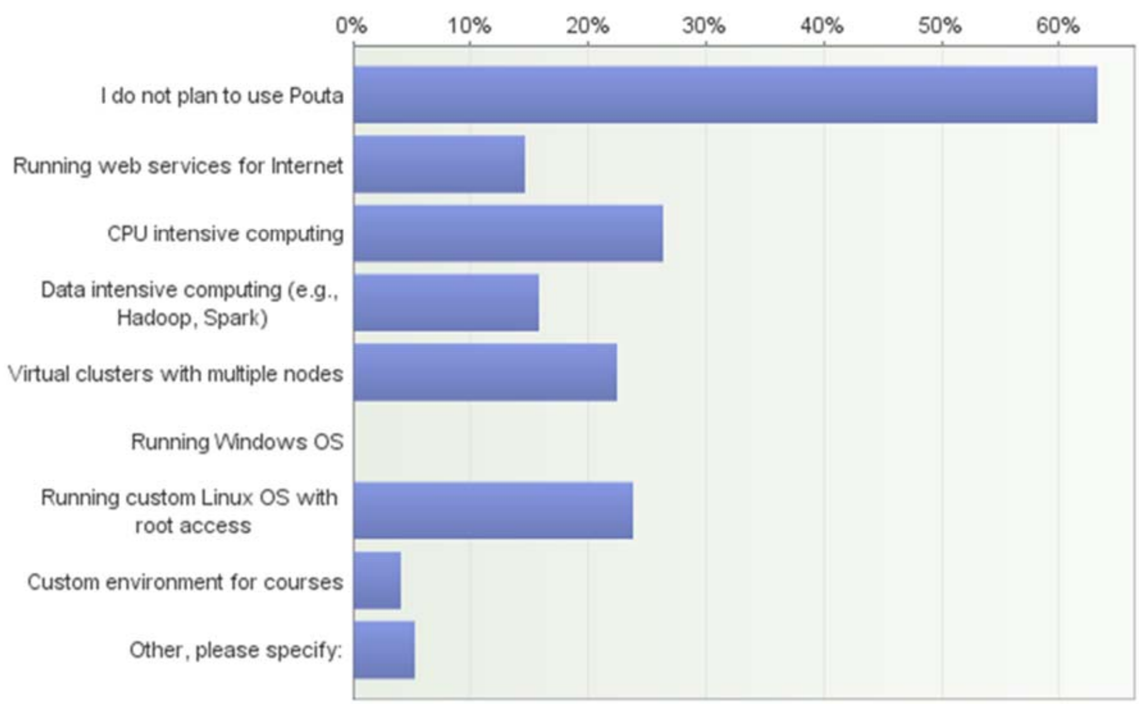

Avoimet vastaukset: Other, please specify:

- Do not know yet
- Making large simulation datasets available to other researchers
- I don't know
- Run clustered database serving web server

14. My CSC Pouta cloud use requires (select all that apply)

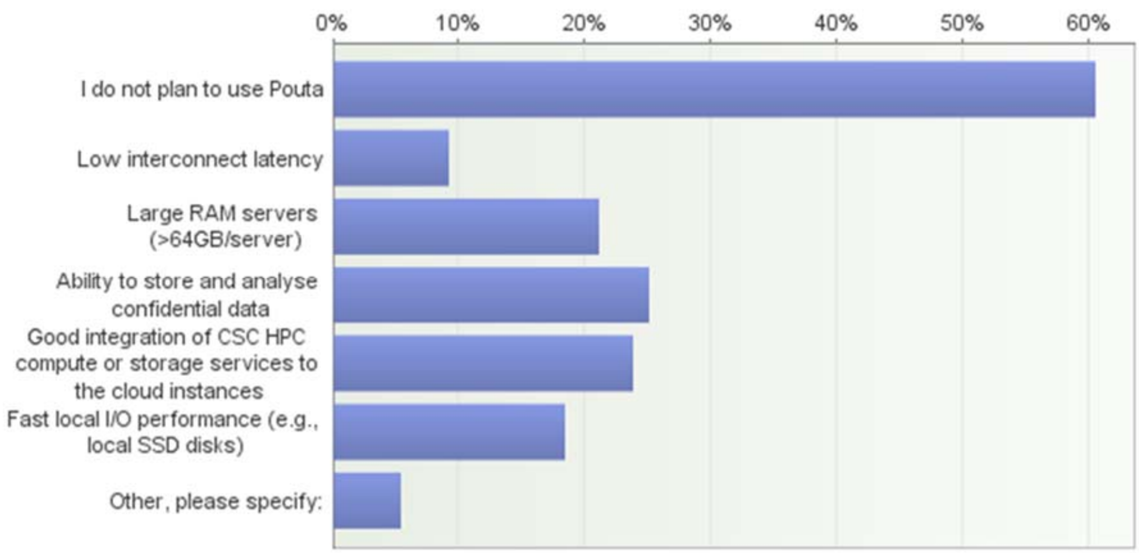

Avoimet vastaukset: Other, please specify:

- Do not know yet
- Having a certified storage location for sensitive data
- Big capacity, large data amounts, good external network for data transfers
- Anything which makes database I/O fast.

Appendix 2 - Usage statistics

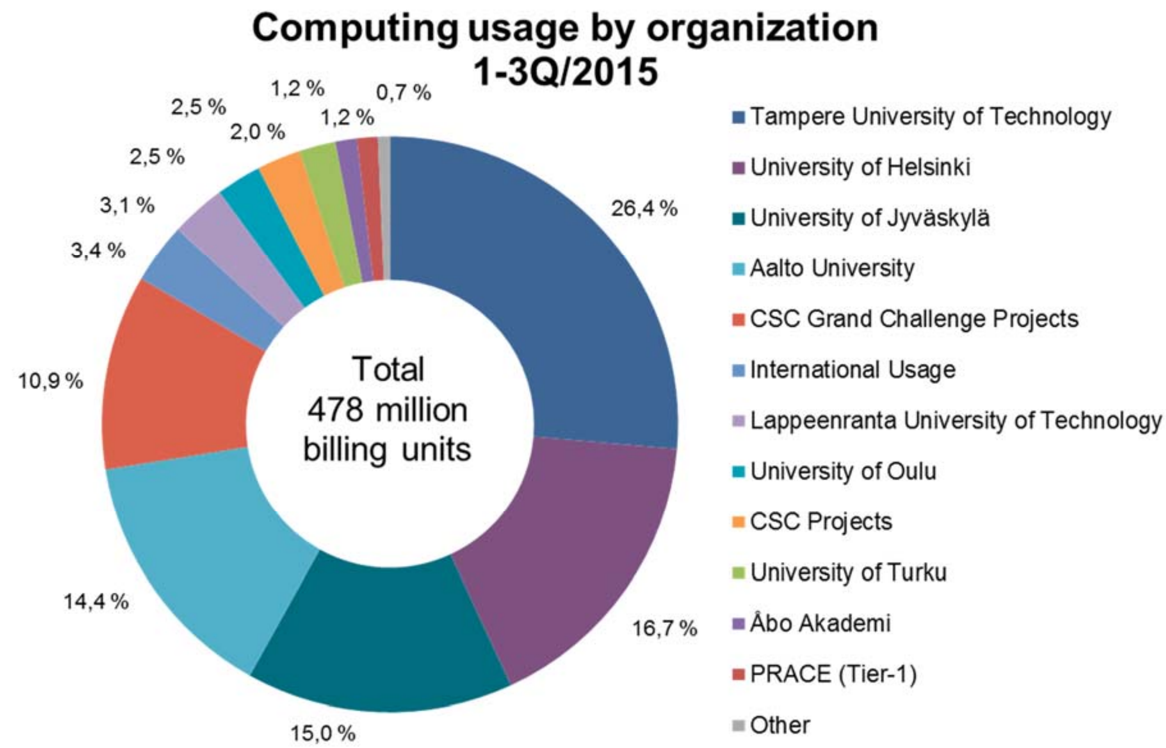

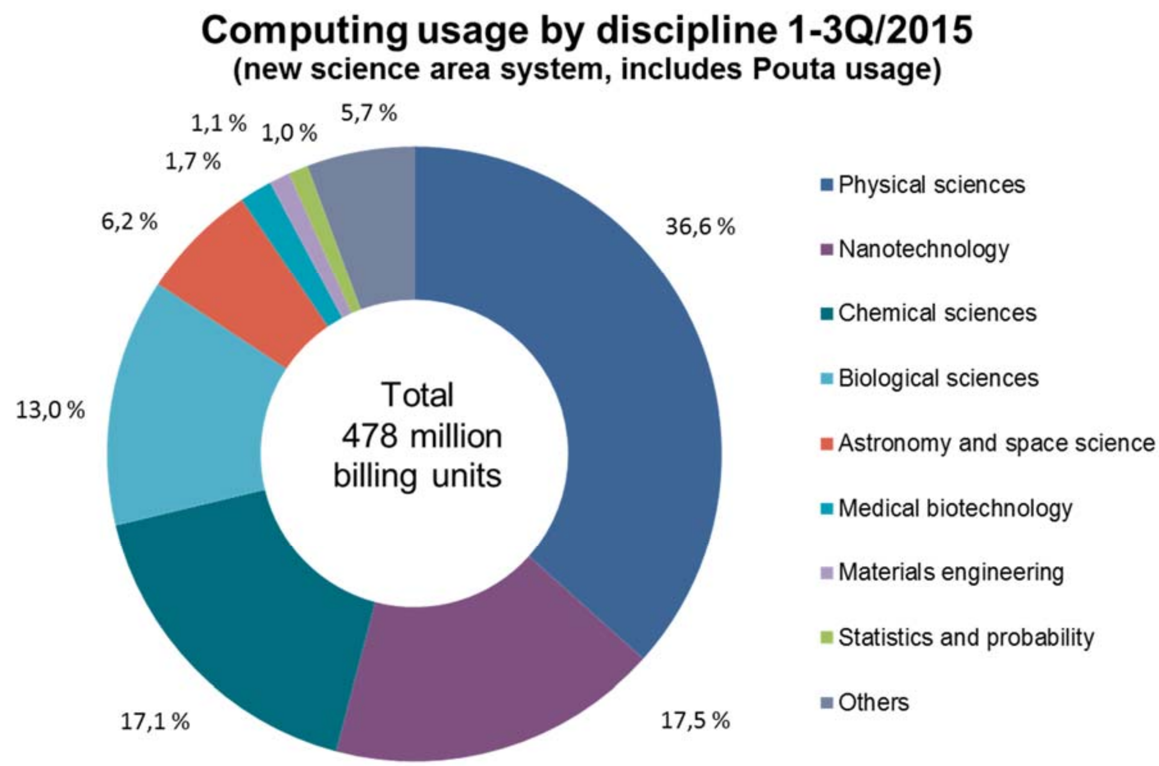
Computing usage by discipline 1-3Q/2015
(new science area system, includes Pouta usage)
Total
478 million
billing units
36,6 %
17,5 %
17,1 %
13,0 %
6,2 %
1,7 %
1,1 %
1,0 %
5,7 %
Physical sciences
Nanotechnology
Chemical sciences
Biological sciences
Astronomy and space science
Medical biotechnology
Materials engineering
Statistics and probability
Others

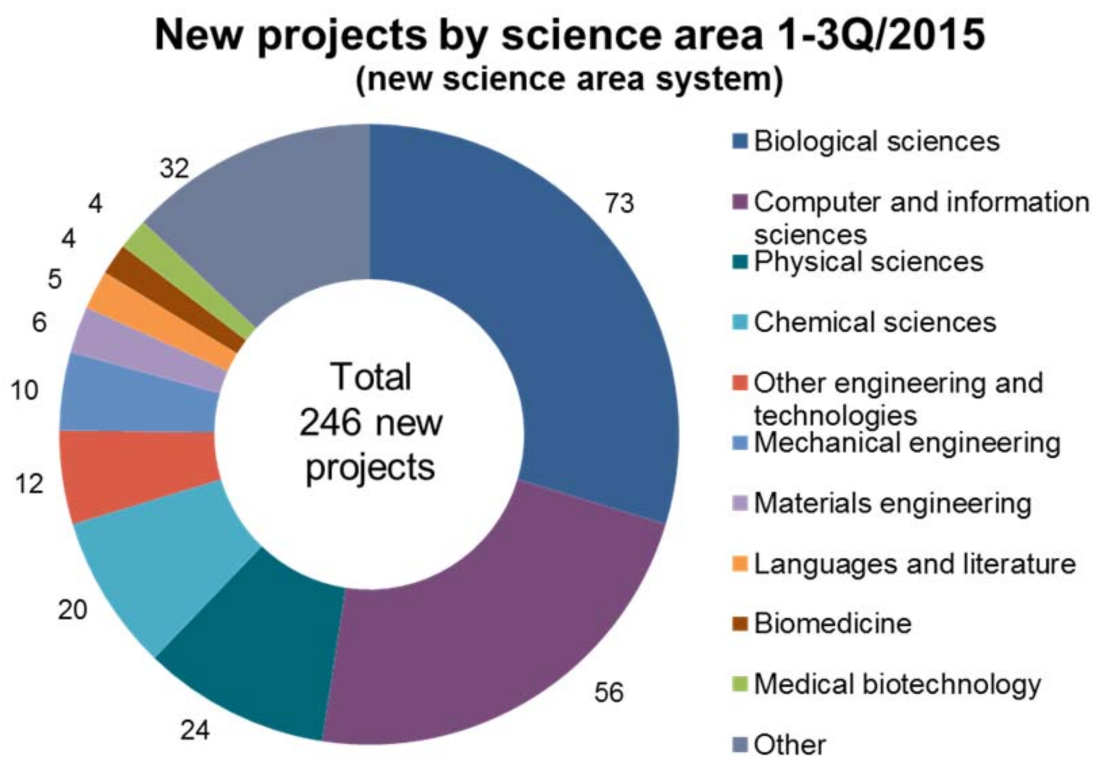
New projects by science area 1-3Q/2015
(new science area system)
Total
246 new
projects
73
56
24
20
12
10
6
5
4
4
32
Biological sciences
Computer and information sciences
Physical sciences
Chemical sciences
Other engineering and technologies
Mechanical engineering
Materials engineering
Languages and literature
Biomedicine
Medical biotechnology
Other

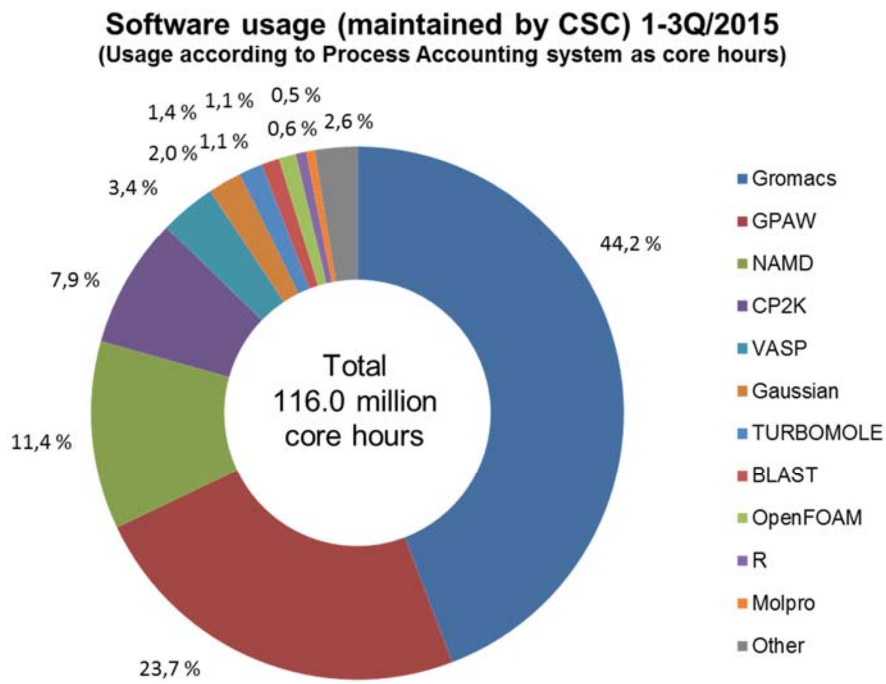
Software usage (maintained by CSC) 1-3Q/2015
(Usage according to Process Accounting system as core hours)
Total
116.0 million
core hours
44,2 %
23,7 %
11,4 %
7,9 %
3,4 %
2,0 %
1,4 %
1,1 %
1,1 %
0,5 %
0,6 %
2,6 %
Gromacs
GPAW
NAMD
CP2K
VASP
Gaussian
TURBOMOLE
BLAST
OpenFOAM
R
Molpro
Other

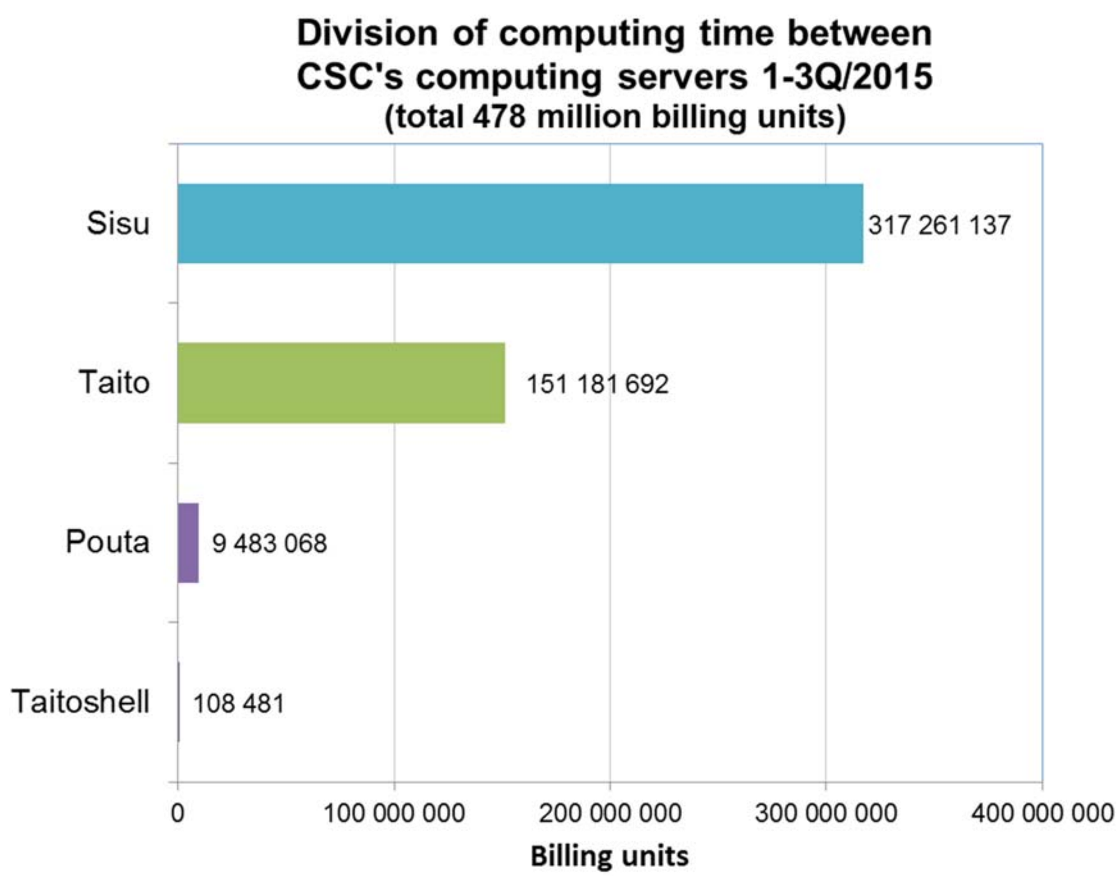
Division of computing time between
CSC's computing servers 1-3Q/2015
(total 478 million billing units)
Sisu
317 261 137
Taito
151 181 692
Pouta
9 483 068
Taitoshell
108 481
0
100 000 000
200 000 000
300 000 000
400 000 000
Billing units