WE ENHANCE EXCELLENT RESEARCH FOR A BETTER WORLD

UNINETT sigma2

Content

3	Managing Director's Corner
4 – 15	Use Cases and Research Activities
16 – 21	Sigma2 Services
22 – 25	International Services
26 – 28	Presenting the Norwegian e-infrastructure

Published by: UNINETT Sigma2 AS
Address: Postboks 4769 Torgard, 7465 Trondheim

Editor in Chief: Gunnar Bøe
Editors: Vigdis Guldseth & Kjersti Strømme

Layout: HK Reklamebyrå

Photo:
Page 3 Andreas Bach/Sigma2
Page 6 Photo of Frigessi by Oslo University Hospital
Page 7 Photo of Bentsen, Private
Page 12 Photo of Carlsson by Ola Sæther/Uniforum
Page 16 Espen Ali Johansen/Uninett
Page 23 Photo of Reuter by University of Bergen
All other photos by Shutterstock

THE MANAGING DIRECTOR'S CORNER

This is the first annual report we create in this format. The intention is to give a good insight into our activities and more important, what great science has been achieved by using the national e-infrastructure we provide. The list of the scientific disciplines using the Sigma2 national e-infrastructure is getting too long to list here, but there are some great examples of research projects later in this report. At the end of the report you can also find more information about the development of services from Sigma2 and our international involvement.

We have with great determination been supporting the government strategies for open data and FAIR[1] principles. The Archive has been increasing exponentially, in line with the governmental guidelines and policies for access and reuse of scientific research data. The Resource Allocation Committee has been leading the way in Norway with their requirement that all projects using storage resources need to have a Data Management Plan. This has been facilitated by our internationally recognized tool easyDMP. This tool is now being used by almost 180 organizations both inside and outside Norway.

Another highlight in 2020 was the inauguration of the supercomputer Betzy. For researchers in Norway this means a tremendous increase in capacity. Many projects can now do science on a scale not possible before.

Advanced user support was one of the objectives specifically mentioned in the government Digitalisation strategy for the higher education sector 2017-2021. This is an important service for Sigma2 and we regularly see excellent results from this work.

Sigma2 is also mentioned in The National Strategy for Artificial Intelligence, published in 2020. During 2020, we have seen the first projects starting to use our e-infrastructure services related to AI/ML. We are looking forward to 2021 with great expectations when we get access to the Norwegian share of the new pre-exascale machine LUMI. LUMI will be one of EU's flagships for HPC and AI/ML infrastructure and our preparations for the use of this machine is well underway.

In 2020, we reached the fulfilment of the goal set in 2015 to create a national e-infrastructure as the transformation from four local sites to two national sites for e-infrastructure. Before Sigma2 was established in 2015, four universities were responsible for a locally owned and operated supercomputer. One of the major goals of Sigma2 was to consolidate this e-infrastructure and take the overall responsibility, as well as reduce the number of sites down to two. The plan for doing this was to gradually phase out the local machines and replacing them with new national e-infrastructure.

We would like to emphasize the good collaboration we have with our partner universities in achieving all this. We see that the work we do together in our Metacenter is very much appreciated in the feedback from many of our users. We have worked to strengthen the Metacenter and hope to deliver even better services in 2021.

Finally, a thank you to our nearly 2000 users who inspire the work we are doing. We are honoured to provide e-infrastructure services for the Norwegian research community to enhance excellent research for a better world.

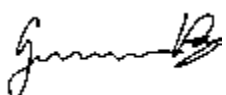

Gunnar Bøe
Managing Director

"It is of utmost importance for Norwegian researchers to have access to e-infrastructure at a high international level in order to be competitive."

[1] Findability, Accessibility, Interoperability, and Reuse of digital assets

Use Cases and Research Activities

WHO IS USING OUR E-INFRASTRUCTURE?

At the end of 2020, we had more than 2000 active e-infrastructure users on our resources. 321 computing projects used the HPC facilities and there were 140 storage projects on NIRD. The number of allocations requested for both HPC and storage indicates that the demand for resources continues to increase. So does the number of users, which increased by around 20% last year. We also have a high utilization on our systems, around 87% in the last year.

Researchers, particularly in climate modelling, life sciences, materials technology, physics, chemistry, geosciences and language technology have a wide-ranging need for reliable computing power and data storage. The capacity also includes 10.7 MCPUh (million CPU hours) reserved to the Norwegian Institute of Public Health for Covid-19 simulations since April 2020, analyzing, forecasting, and reporting virus spread scenarios to the government. Hence our infrastructure makes an important societal contribution.

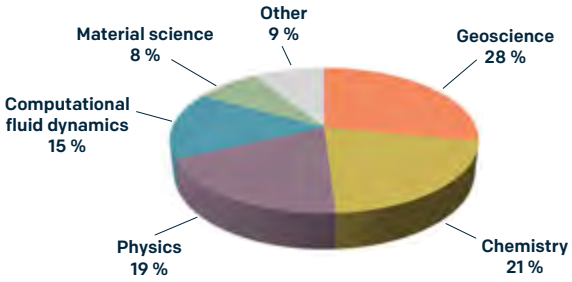

HPC usage by field of science

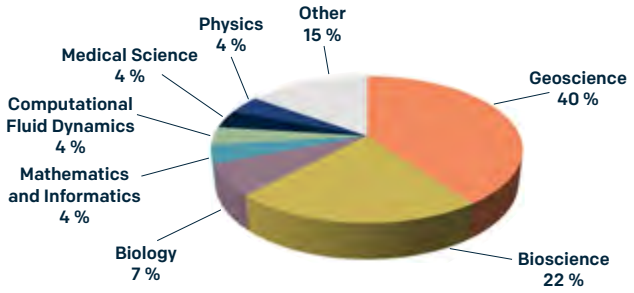

Storage usage by field of science

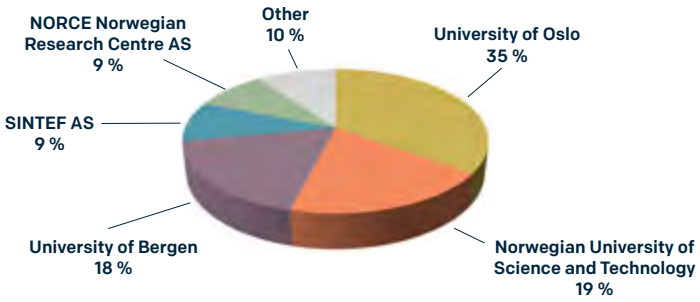

HPC usage by organization

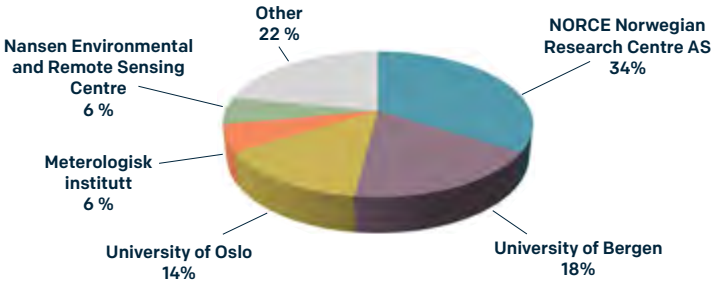

Storage usage by organization

Resource allocation – Who decides?

Applications for computing, storage facilities and advanced user support resources for projects are evaluated by a Resource Allocation Committee (RFK). The RFK is composed of leading Norwegian scientists.

The RFK evaluates the proposals and awards access to Sigma2 computing and storage resources twice per year. The allocation process follows a strictly regulated and defined process. Proposals for access must demonstrate scientific excellence, scientific need and that the resources requested will be used efficiently.

FIGHTING **COVID-19**

The Norwegian Institute of Public Health (FHI) has been completely dependent on the national e-infrastructure in their simulations during the Covid-19 pandemic. The Institute does not have their own HPC resources and are thus reliant on Sigma2 to prioritise and carry out the calculations on short notice.

Arnoldo Frigessi
Professor, Faculty of Medicine, University of Oslo

Early in the pandemic, when we first started to hear about the now-familiar terms "the reproduction number R" and "infection pressure", FHI's Infectious Diseases Epidemiology and Modelling Team, led by Department Director Birgitte Freiesleben de Blasio, was already running simulations on virus spread scenarios on the Sigma2 supercomputers.

FHI has developed different models to simulate people moving between geographical areas and individuals interacting in virtual representations of society. Data from various sources are added to adjust the models, and then a vast amount of computational power is needed to simulate different scenarios for virus spread or expected vaccine efficacy. Based on this statistical evidence the team and FHI report and advise the authorities during the pandemic.

"

The support we have received from Sigma2 along the way has been excellent. Without their resources and unbelievably efficient support, Norway would not have had in time the evidence available to build its containment and prevention strategies, which, as we know, turned out to be the most drastic ones in recent history, but also the most effective ones."

Professor Frigessi, on FHI's Modelling Team.

SIMULATING CLIMATE CHANGE

Mats Bentsen
Researcher at NORCE

The research community agrees that climate and climate calculations are of great social importance in the future. Therefore, it is important to ensure that the research environments have access to substantial compute resources.

NorESM is an earth system model for simulating climate change. Researchers all over Norway and Europe use this model, which is maintained through the project "Infrastructure for Norwegian Earth System Modeling" (INES). The model is recognised by the Intergovernmental Panel on Climate Change (IPCC) who has used research based on NorESM simulations in its climate report.

A climate model is complex, and a typical simulation uses thousands of processor cores simultaneously, preferably over several weeks in a row. NorESM has run up to 25,000 cores on Betzy and GPU porting of the code is underway. This makes the application suitable for testing on the new supercomputer LUMI.

"

We are almost 'insatiable' when it comes to the desire for increased computing power. Many of the projects that depend on carrying out climate simulations would not be conducted on this level of ambition without access to Sigma2's resources. Our simulations create large amounts of data that must be stored securely, processed efficiently, found by, and shared with other researchers. At present, only Sigma2 can offer suitable HPC and storage resources that the national climate research community jointly has access to."

Mats Bentsen

LEARNING AND EXPERIENCING WITHOUT HUMAN INVOLVEMENT

What is commonly called "deep learning" has become the dominant paradigm in much contemporary machine learning research and development. Deep learning emphasizes on training the computer to learn by experience to perform tasks such as self-driving cars, facial recognition, or natural language processing.

In Natural Language Processing (NPL) break-through advances in recent years have been fuelled by massive neural language models, the best-known instance is BERT (Bidirectional Encoder Representation from Transformers), introduced by Google, LLC in 2018. These models are computationally expensive to train, refine, and user training can take up to several GPU months. Fine-tuning a pertained model for a specific application typically requires at least several GPU days. The models are currently only available for English and a few additional languages.

Mature and open-source deep learning frameworks like TensorFlow and PyTorch (by Google and Facebook respectively) in principle allow researchers without in-depth specialist training to conduct large-scale deep learning experiments that effectively parallelize across multiple GPUs or even multiple multi-GPU nodes.

At present, there is only very limited GPU capacity available for research usage in Norway. Professor Stephan Oepen and his colleagues at the Department for Informatics at UiO should be able to make good use of the powerful LUMI supercomputer.

COMPUTATIONAL ANALYSIS OF SEVERE MENTAL ILLNESS

NORMENT – Norwegian Centre for Mental Disorders Research is a Centre of Excellence (CoE) using the national e-infrastructure resources. Their main research goal is to find answers to why some people develop severe mental illness.

Psychiatric disorders are recognized as leading causes of morbidity globally and are among the most expensive disorders to affect humans. Identifying the underlying pathophysiology is imperative and can lead to major health benefits, through better treatment and prevention strategies. The heritability is high, but most of the genetic factors and the interplay with environment are unknown, as well as the underlying brain mechanisms.

It is becoming increasingly clear that multiple factors influence most mental disorders. In such a scenario with a very high number of risk genes, each gene has a tiny effect. This makes it difficult to determine an individual's risk, and to identify disease mechanisms that can be used for development of new effective treatments. Further, studies of the complex interplay with environmental stressors further complicates the analyses. It is increasingly clear that many small brain abnormalities seems to contribute to disease characteristics, not a single brain structure.

With help from the national e-infrastructure resources provided by Sigma2, the researchers from NORMENT aims to extend these approaches and develop new tools that can leverage the rapidly growing amount of genetic and brain imaging data and rich collections of various clinical measurements which they have access to.

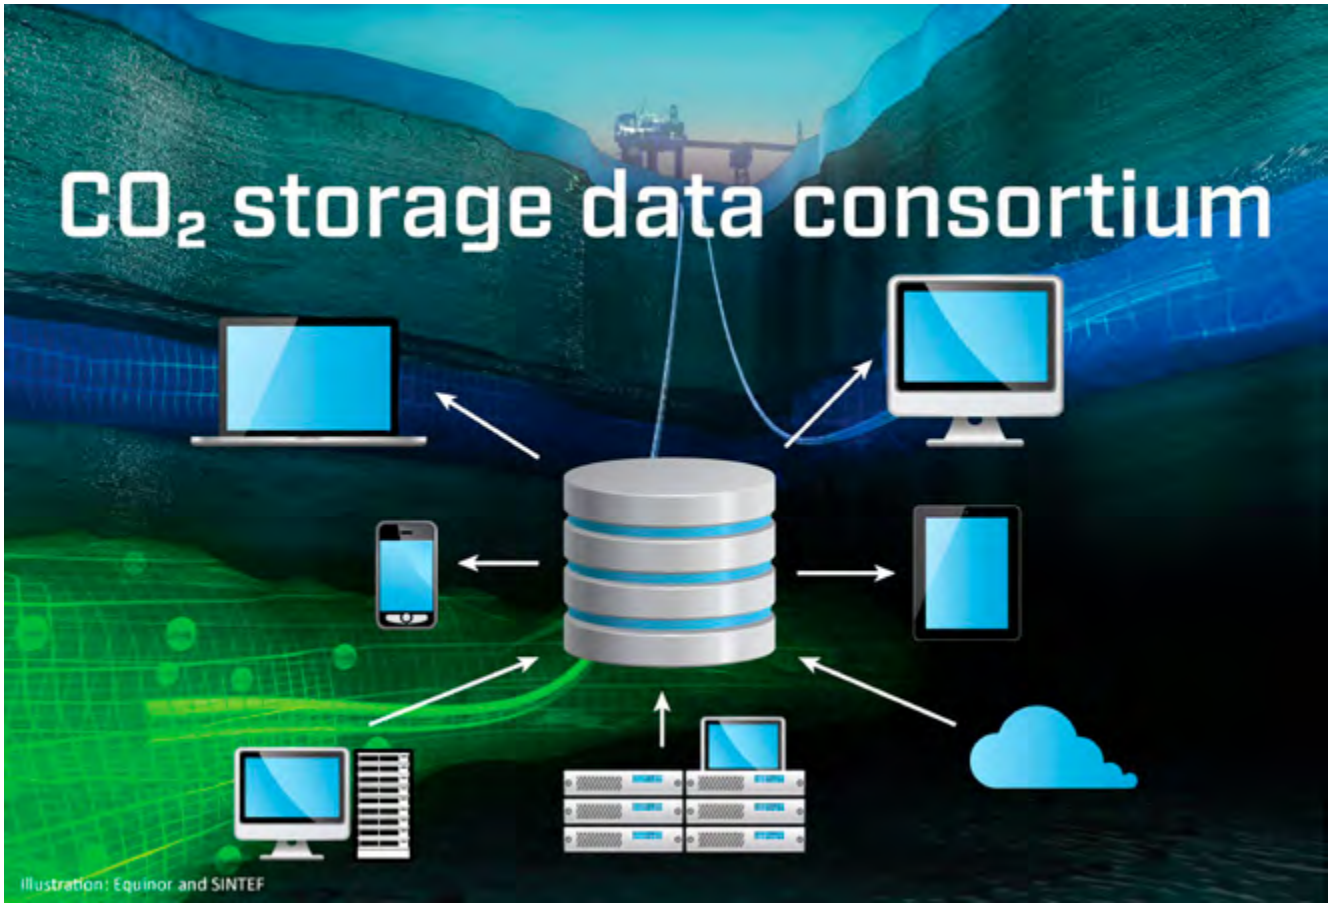

LAUNCHING THE CO2 DATASHARE PORTAL

In February 2020, the CO2 Datashare launched a web-based digital portal for sharing reference datasets from pioneering CO2 storage projects. The new portal enables researchers and engineers to improve their understanding, reduce costs and minimize uncertainties associated with CO2 storage.

The CO2 Datashare Portal was granted Advanced User Support from Sigma2. The digital platform of the CO2 Datashare Portal builds on Sigma2s infrastructure for data storage combined with a portal developed using the open-source software CKAN.

CO2 capture, transport, and storage (CCS) is a process where CO2 emissions are captured from large industrial plants, transported in pipelines or ships, and stored in rocks deep beneath the surface so it will not enter the atmosphere. More than 3400 industrial CCS plants will be needed by 2050 if we are to mitigate the impact of climate change by meeting the critical target of no more than a two-degree rise above pre-industrial. Lack of relevant data has been a barrier for CCS research and deployment, a barrier that can be overcome by access to open data.

Therefore, the CO2 DataShare project was established to accelerate the deployment of CCS by providing open access to CO2 storage data.

> **"The interest in CO2 storage just continues to grow. Now, a new reference dataset from Norway's proposed Smeaheia storage field, developed by Equinor and Gassnova, is available on an open access basis for researchers and engineers around the world to use and learn from."**
>
> Grethe Tangen, Senior Research Scientist, SINTEF.

CO2DataShare partners are SINTEF, Equinor, University of Illinois and IEAGHG.
The project is supported by the CLIMIT programme and US Department of Energy.

AUS
ENABLING CODE FOR HPC THROUGH ADVANCED USER SUPPORT

In one example of from 2020, an AUS project was used to adapt and optimize a much-used pharmaceutical application allowing for larger and more complex pharmacokinetic models to be run on the HPC-infrastructure.

The AUS has provided the end users with improved infrastructure, shifting the workload from personal workstations and laptops to an HPC-cluster, in which the following user experiences confirm.

> **"The AUS has been of great value to our research group and has allowed us to overcome obstacles that we previously couldn't. We are now able to run multiple, more complex models at an HPC, which is a drastic upgrade to our infrastructure. Overall, we are very grateful for the results from this AUS, which has drastically increased our research capabilities due to improved infrastructure."**

Markus Herberg Hovd (MH) M.Sc. Pharm,
Department of Pharmacy, University of Oslo.

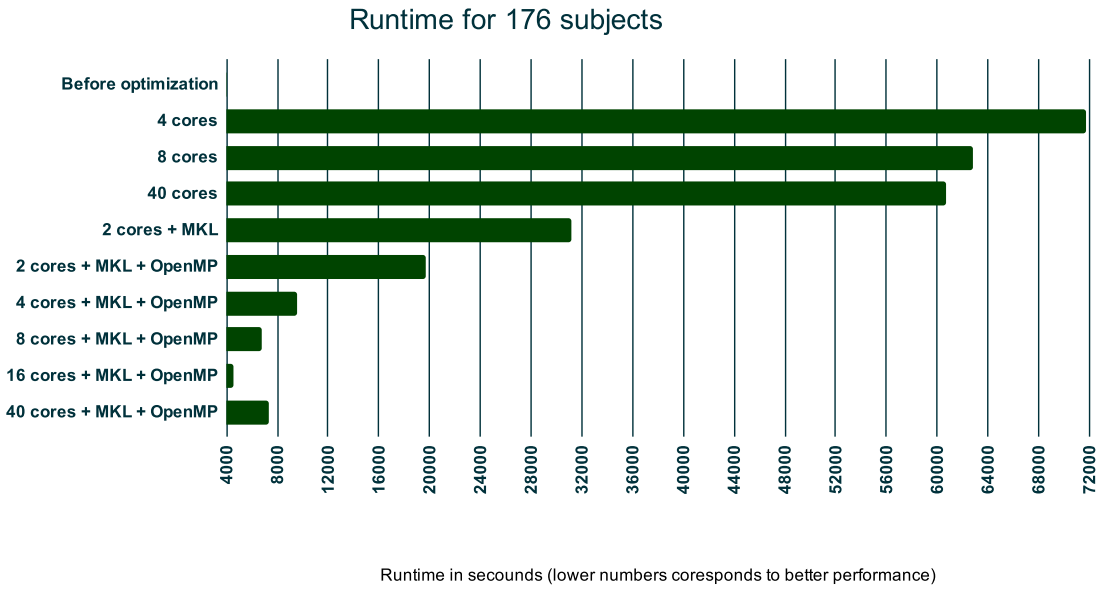

Plot showing the performance increase (lower numbers are better), with different levels of optimization.

LEARNING MORE ABOUT THE SUN

Mats Carlsson,
Professor at Rosseland Centre for Solar Physics, UiO

The Sun is of great importance for life on earth. Increased knowledge about the Sun's atmosphere and space weather is important as the Sun conditions can impact human health, the performance of technology we use on Earth and can even cause damage to critical infrastructures, such as the electric grid.

The HPC project "Solar Atmospheric Modelling" uses the applications Bifrost and Dispatch to simulate the Sun's atmosphere, the outer layers from which sunlight comes. The simulations are compared with observations from solar telescopes on the ground and space, to learn more about how the Sun, and thus other stars, works.

The applications are highly parallel, which mean that they use numerous processor cores to perform calculations faster. On Betzy, Bifrost has been tested on more than 131,000 cores, and application developers are collaborating to adapt the code to the GPU. This makes the applications particularly suitable for piloting the new supercomputer LUMI.

"**In the long run, we hope the increased knowledge can provide a better understanding of solar storms and better forecasts of space weather. To achieve this, we must simulate much larger areas of the Sun than we have done so far. This requires even larger computing resources as well as the development of new algorithms and methods. We hope the new LUMI computer can help with this."**

Mats Carlsson

ENVIRON-MENTALLY FRIENDLY AVIATION

Aviation is one of the rapidly growing sources of greenhouse gas emissions. Unlike other liquid fuels, such as gasoline or diesel with established alternatives (battery or electric power), the most promising alternative to current aviation fuel is bio-derived aviation fuel, which is still in the earliest stages of development.

The biomass to aviation fuel (B2A) project run by Professor De Chen at the Department of Chemical Engineering, NTNU and his colleagues, aims to develop a new competitive technology platform to produce the next generation of aviation fuels from lignocellulosic biomass. The costs of biomass, catalysts, carbon yield and catalyst stability have been identified as the most sensitive parameters for the cost of biofuels.

To meet the challenges, the project proposes three main technical solutions to identify low-cost lignocellulosic biomass and biomass waste, develop low-cost stable catalysts and achieve high carbon yields of jet fuels, where catalysts play in critical role. Development of low cost and highly active and selective catalysts relies on a better understanding of surface reactions of the complex reactions.

The project aims to elucidate the catalyst structure, properties, and performance relationship by density functional theory (DFT) calculations and to provide principles for catalysts rational design and maximize the catalyst performance and yield of aviation fuels.

MACHINE LEARNING FOR PREDICTING "CHATTERING ALARMS"

Alarms systems play a vital role to ensure safely and reliability in the process industry. In a process plant environment, operators are notified by alarms if a process diverges from normal operating conditions.

Alarms are triggered to warn the operator about irregular incidents. When several alarms go off simultaneously it is known as "alarm floods", and the operator may miss the crucial alarms. Ideally, an alarm should therefore inform the operator about critical conditions only and provide guidance to a set of corrective actions associated with each alarm.

Three scientific articles were recently published based on research conducted via the NIRD Service Platform and the Deep Learning Toolkit. An Advanced User Support (AUS) project contributed when a new method for assessing dynamic chattering alarms was developed. The results were used to train and evaluate a Deep Neural Network. The model was then tested against the ability to predict alarm chatter.

A modified approach based on run lengths distribution was developed to evaluate the likelihood of future alarm chatter. The method has allowed categorizing historical alarm events as alarms that will (or will not) show chattering in the future.

Finally, categorized alarms have been used to train a Deep Neural Network by using TensorFlow, whose performance has been evaluated against the ability to predict alarm chatter.

Associate Professor Nicola Paltrinieri at the Department of Mechanical and Industrial Engineering at NTNU and his colleagues recently benefited from an AUS project.

HIGH LATITUDE COASTAL CIRCULATION MODELLING

Knowledge of coastal ocean transport processes is vital for predicting human impact on the coastal marine environment. Coastal industry discharges pollutants and nutrients into the ocean, and this has an impact on the coastal marine ecosystem. To understand and predict such impacts require a high level of understanding of near-shore transport processes and we need coastal ocean circulation models to calculate concentrations and pathways of spreading of discharges.

The researchers in Akvaplan-niva collects and analyse data from the coastal environment and we develop and use high-resolution coastal circulation models as a tool to understand the human impact. Our main area of interest is the Norwegian coast, but we also use our ocean models to study the coastal ocean in the Arctic and Antarctic. High-resolution atmospheric data is added to force the ocean circulations models. Using MPI parallel applications, it is safe to say that the models require lots of calculating resources and the project used 9 million CPU hours in 2020 alone. The large amounts of model data and simulations also need to be safely stored. During a typical run they produce 100 files of 100 GB each.

"We need to store data that we use for forcing ocean circulation models. This data is continuously updated by development of large-scale circulation models. The amount of data storage will increase because models develop towards higher resolution and datasets generally becomes larger."

Ole Anders Nøst, Head of Section Oceanography and Modelling, Akvaplan-niva AS.

Sigma2 Services

HIGH PERFORMANCE COMPUTING

At Sigma2, we provide the classical HPC service, but with a varying degree of performance along the computation to data access capability axis. In addition, we also provide a common software environment for scientific computing. While not being pure HPC, this is important to many researchers for conducting their computer driven science, as they might neither have the capability nor the resources to provision this internally within their research groups. Our HPC service is fuelled by three major systems, Betzy, Fram and Saga.

The supercomputer Betzy is named after Mary Ann Elizabeth (Betzy) Stephansen, the first Norwegian woman with a PhD in mathematics.

Betzy, the newest and most powerful supercomputer in Norway, was inaugurated in December 2020. She went straight into and was ranked at 56th place on her introduction on the TOP500 list, a global ranking of supercomputers.

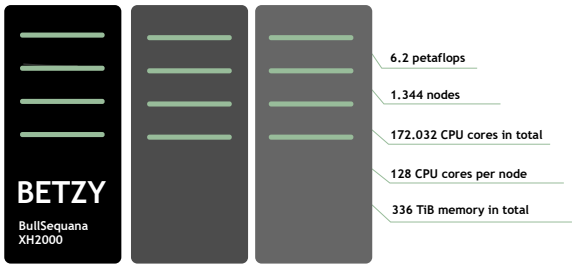

Betzy, with her 172,000 cores (computational units) is by far the largest system, providing the highest FLOP performance and aggregate I/O capability. However, per computing unit (core) Betzy delivers half the bandwidth of the two sister systems. Betzy is targeted at massively parallel applications, solving large (manycore) computational problems with large memory and disk requirements.

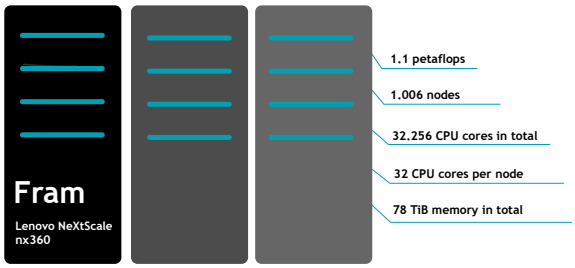

Fram offers medium computational and disk capability to computational science not needing or able to exploit capabilities of the two other systems.

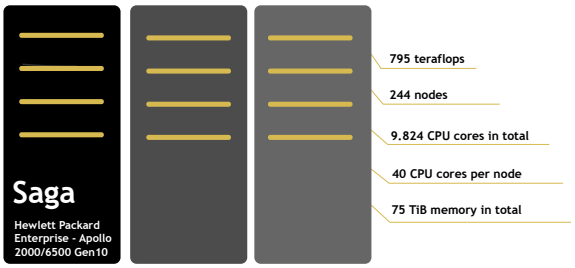

The **Saga** machine, with her 10,000 cores, is hardly an HPC system, but its disk subsystem offers high-performance random access to bioinformatics research and similar fields of science.

Decommissioned systems

After a long and faithful service, the two supercomputers Stallo and Vilje were decommissioned in 2020.

Stallo was installed in 2012 and upgraded in 2013 and has since run 173 projects with national allocation, and thus contributed to a vast number of scientific publications.

By the time of installation in 2012, **Vilje** went straight into 44th place on the TOP500 list. This is the best Norwegian listing ever. Vilje has since contributed to numerous research projects, 176 to be specific. In the later years "Meteorologisk institutt" has used about 1/3 of Vilje's power to calculate weather forecasts.

Vilje was involved when the Estonia case was resumed in 2020. Simulations run on Vilje to calculate the force needed to make the four-meters-high and about 1-meter-wide hole in the hull were conducted by Professor Jørgen Amdahl at the Department of Marine Technology, NTNU.

DATA STORAGE SERVICES

NIRD – The National Infrastructure for Research Data is an eco-system of storage services designed to support the scientific research in every steps of the Research Data Life Cycle, from data planning to data collection and storage, to data analysis and finally data sharing, archiving and re-use. The guiding principles are Open Science and Open Data.

The NIRD ecosystem consists of easyDMP, NIRD Project Storage, NIRD Service Platform and NIRD Research Data Archive.

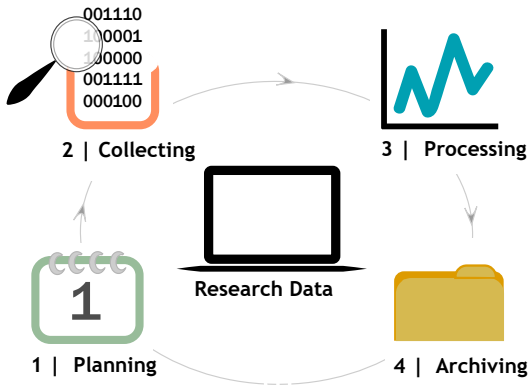

The Research Data Life Cycle

easyDMP – Data Planning

The easyDMP is a tool designed to facilitate the creation and maintenance of Data Management Plans (DMPs). easyDMP supports templates that have either a rigid structure or consist of questions with canned answers so that the users are guided through the different aspects of data management step-by-step.

In 2020, the portfolio of templates offered in easyDMP has been enriched, including templates following Horizon 2020 recommendations and the Science Europe recommendations, as recommended by the Norwegian Council of Research. Furthermore, easyDMP supports institutional templates, i.e. templates that have questions customised to a specific institution and prefilled answers, to greatly enhance the user experience in filling up the DMP.

NIRD Project Storage

NIRD Project Storage provides massive storage resources to projects which are actively running a scientific investigation. These are scientific projects which need large storage resources (above 10 TB) or special requirements, such as needs for a data intensive computing platform. The NIRD Project Storage is connected to the NIRD Service Platform to enable data intensive computing and enhance accessibility and discoverability of the data through services.

In 2020, NIRD storage consisted of circa 11 PB of storage geo-replicated on two different locations, to ensure data integrity and facilitate the data staging to the HPC systems. Analysis of the usage patterns of the NIRD Project Storage shows that the need for long term data preservation and the need for controlled data sharing has become increasingly more relevant among the user communities, thus requiring a re-design of the NIRD infrastructure. This is currently done by the NIRD2020 project and the new infrastructure will be procured in 2021.

NIRD Service Platform

The NIRD Service Platform consists of two clusters with computing resources orchestrated by a Kubernetes engine with a CPU/GPU architecture to enable services and computing environments to run data intensive computing workflows such as pre-/post- processing, visualization and AI/ML analysis.

The Service Platform is facing increasing popularity since it allows computing over a large volume of data without data movement. Furthermore, the possibility to spin-off community specific portals to expose large datasets to external users allows the sharing and reuse of the data in line with the current policies and best practices in research. Services can run permanently or on-demand through the NIRD Toolkit.

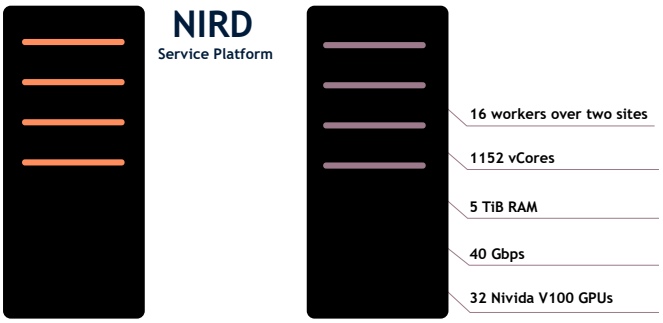

The NIRD Service Platform is a computing platform located in Tromsø and Trondheim.

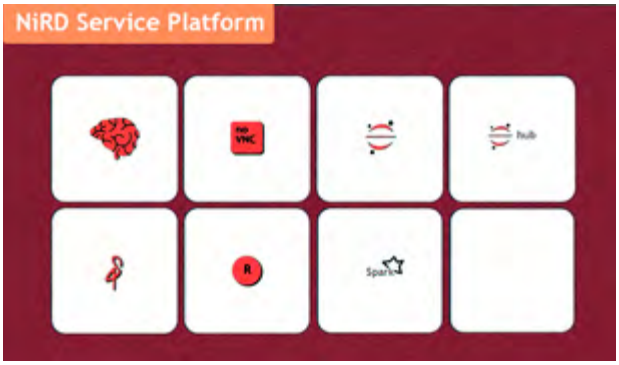

NIRD Toolkit

The NIRD Toolkit is a Research Platform to easily spin-off software on-demand on the NIRD Service Platform, thus avoiding cumbersome IT-operations.

With the NIRD Toolkit, computations on large amount of data with software such as Jupyter Notebook/Jupyter Hub, Spark, R-Studio, and several popular artificial intelligence algorithms are possible in a few clicks.

In 2020, the mechanism which allows end-user software to be added to the platform was set in pilot production, including security measures.

NIRD Research Data Archive

The NIRD Research Data Archive is an archive solution for long preservation. The archive came into service in June 2014 and has ever since stored 244 datasets, corresponding to circa 500 TB in total (more than 1 PB including replicas).

The increasing awareness of researchers of data management and FAIR data is shown in the sharp rise in the number of datasets deposited during the last years. While the NIRD Data Storage ensures mechanisms to administer access and preserve business confidentiality, the datasets in the Research Data Archive are discoverable and openly accessible.

Dataset deposited during the last 5 years in the NIRD Research Data Archive with Open Access policy.

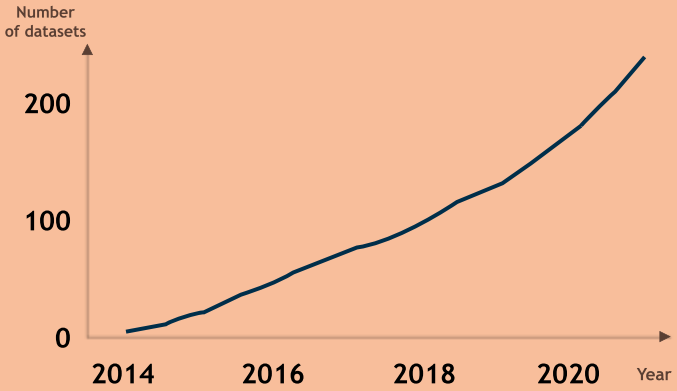

Dataset deposited during the last 5 years in the NIRD Research Data Archive with Open Access policy.

Sensitive Data Services

The TSD - Services for Sensitive Data offers a remote desktop solution for secure storage and high-performance computing on sensitive personal data. The service has been designed and set up to comply with the Norwegian regulations regarding individuals' privacy. TSD is developed and operated by the University of Oslo in partnership with Sigma2. Since 2018, TSD has been part of the national e-infrastructure.

A new agreement was signed in 2020 and Sigma2 now owns 1600 cores in the TSD HPC cluster, corresponding of the 80% of the total HPC capacity available to the end users (not privately owned). Sigma2 owns in addition 1370 TB of the TSD disk, delivering storage for general use or HPC use.

BASIC USER SUPPORT

Every year Sigma2 carries out user surveys among the research projects that make use of the national e-infrastructure resources. The surveys allow the user community to communicate their satisfaction with the resources and services and point to areas that need increased attention.

Basic user support provides a national helpdesk that deals with user requests such as issues relating to access, systems usage or software. Support tickets are handled by highly competent technical staff at the Metacenter universities. Our user surveys consistently return high scores on both the technical and administrative support channels, with 5.17 of 6 and 5.30 of 6 points respectively in last year's survey. The survey also shows that the users are satisfied with the other services, with a score from 4.7 to 4.9

ADVANCED USER SUPPORT

Advanced User Support (AUS) is a service at the apex of what a national e-infrastructure can offer, combining competence from the experts in the Metacenter with the advanced technical resources. AUS offer research projects the extra high-level support that can really make the difference, and additionally benefits all users by way of implementing new services or improving efficiency and utilization of resources.

The structure of the support can be visualized as in the following figure, where on one side the 1st-line support handles a lot of tickets, some of which require more work by experts, and ultimately some lead to AUS projects on the other side.

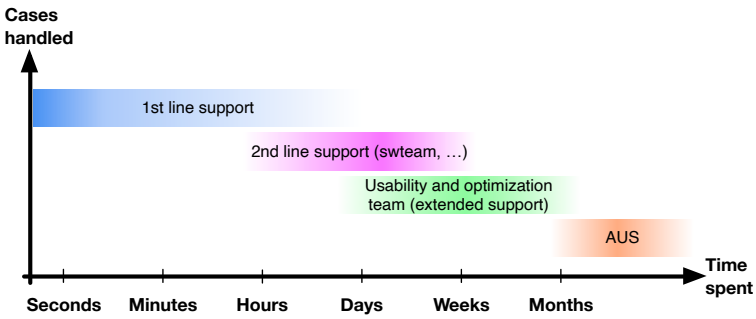

Time spent and cases handled by our support services.

CRAAS – COURSE RESOURCES AS A SERVICE

Launched in 2020, CraaS is a service for researchers who require e-infrastructure resources to be used in a course or workshop for research purposes. The service intends to help researchers who are interested in providing courses or workshops with national e-infrastructure resources as a necessary tool. Nearly 300 users from almost 10 projects have so far utilized the CRaaS-service.

"With your solution, students always have something to practice on and more resources than they can have at the University. You can't imagine how valuable it is for a teacher like me who needs a Unix computer and for whom students are not experts. I've been talking to some teacher colleagues in my home country, France. It is a pain for them not to have a solution like this. Thank you again for your precious time, and Sigma2 for offering this solution."

Thibault Michel Joseph Tubiana,
Postdoctoral fellow at Department of Chemistry, University of Bergen.

The partnership within the NCC and the international collaboration will enable the use of the partners' experience and strengths in supporting industry, public sector and academia.

EUROHPC NATIONAL COMPETENCE CENTRE

Making industry and SMEs HPC-ready

The EuroHPC National Competence Centre (NCC) in Norway is a partnership between NORCE, SINTEF and Sigma2. The partnership offers a novel blend of technical, administrative, domain and business development competence, available infrastructure and collaborative networks with the industry. This was formally established in 2020.

Building on existing successful competences and business models, the Norwegian Competence Centre aims to support the Norwegian industry, the public sector and academia in using HPC-services. The centre intends to help small and medium-sized companies, which can benefit greatly from HPC, High Performance Data Analytic, Machine Learning, Artificial Intelligence, and related technologies, but rarely have the capacity and expertise within their own business. By offering support and competence sharing, the centre will also drive innovation and knowledge building.

"

NORCE with its longstanding collaboration with industry on various levels sees a great potential in the NCC. The flexibility of the NCC to adjust the collaboration with companies, especially SMEs, to the needs of the industry partners will allow for a targeted help and foster developments in the industry to the maximum possible level."

Klaus Johannsen, NORCE.

Widespread collaboration

The Norwegian Competence Center will collaborate with the various national initiatives around AI, DIH (Digital Innovation Hubs), and the Research Council of Norway, Cluster and Centre of Excellence programs, in order to provide a larger network of services being offered to SME´s, as well as providing services to innovation and SME-related work connected to these initiatives.

"

HPC and associated technologies can be a significant enabling factor for Norwegian industries and public sector. However, the cost of maintaining these technologies, also at the competence level leads to a barrier of adoption. The NCC will try to fill this gap and operate as a vital asset for increased adoption and enablement of HPC and related technologies for Norwegian industry and public sector."

Josip Zoric, SINTEF.

International Services

COLLABORATIONS ACCROSS THE BORDER

Norwegian researchers can get access to services, including compute cycles on very large systems, through Sigma2's partnership in international collaborations. This is especially interesting for researchers that otherwise are not covered at the national level, for example due to the size of the system, or a technology-specific solution. Sigma2 has also in 2020 participated in several international initiatives and collaborations in the field of high-performance computing and data infrastructure.

PRACE

Enhancing European competitiveness for the benefit of society

The aim of PRACE (Partnership for Advanced Computing in Europe) is to cover the top level (Tier-0) of the European HPC ecosystem, by establishing a permanent infrastructure of Tier-0 compute systems. The experts at the Metacenter and our national e-infrastructure also enable us to give valuable contribution to PRACE projects, e.g., Norway is leading the task on containerized services on HPC within PRACE 6IP. The PRACE infrastructure is open to Norwegian researchers and Norway has had several successful projects running on PRACE Tier-0 systems.

Well over 230 million CPU hours have been awarded Norwegian-led projects awarded in the various PRACE calls. In 2020, Norwegian tsunami researchers were awarded 70 million computing hours from PRACE, the project is led by Dr. Finn Løvholt from the Norwegian Geotechnical Institute. The computing resources were granted form one of Europe's fastest supercomputers, the Italian MACRONI-100.

In 2020, Professor Nathalie Reuter (UiB) was elected to the PRACE Scientific Steering Committee, a very prestigious and competitive appointment.

LUMI

State-of-the-art systems for world-class research

The LUMI (Large Unified Modern Infrastructure) HPC system is one of the three pre-exascale supercomputers co-funded by the European High-Performance Computing Joint Undertaking (EuroHPC JU). As a member of the LUMI consortium, Sigma2 invested 20 MNOK in 2020 and Norway owns a fixed share of this infrastructure, with further opportunities to trade shares between consortia members. In addition, Norwegian projects are eligible to apply for the EuroHPC Joint Undertaking (JU) 50% share of this system.

LUMI will provide predictable access to Norwegian researchers, to a high-end system approximately 80 times the capacity of the currently largest Norwegian system. Being equipped with the latest GPU technology, LUMI will not only provide a next level of capability to classical HPC sciences, but also to AI/ML research in Norway.

In 2021, selected projects will start pilot testing.

NeIC

Nordic role model for cross-border collaborations

Sigma2 is a partner in NeIC – the Nordic e-Infrastructure Collaboration – bringing the national e-infrastructure providers in the Nordics and Estonia together. Hosted by NordForsk, which facilitates cooperation on research and research infrastructure across the Nordic region, NeIC is based on strategic collaboration between the partner organizations, CSC (Finland), SNIC (Sweden), UNINETT Sigma2 (Norway), DeIC (Denmark), RH Net (Iceland) and ETAIS (Estonia).

NeIC projects with Sigma2 and Metacenter participation with activities in 2020

Glenna2
Glenna2 concluded in 2020 after having continued the collaboration of the first Glenna project phase. The project provided added value to the Nordic national cloud and data-intensive computing initiatives by supporting national cloud initiatives to sustain affordable IaaS and PaaS cloud resources, and then using such national resources to establish international collaboration in collaboration with user communities. The Glenna2 project leveraged the pooled competency of the Nordic e-infrastructure providers and performed an assessment of future hybrid cloud technology that was communicated to the national initiatives, additionally supporting use of resources and creating a Nordic support channel for cloud and big data.

NICEST2
The second phase of the Nordic Collaboration on e-infrastructures for Earth System Modeling, NICEST2, focuses on strengthening the Nordic position within climate modeling by leveraging, reinforcing and complementing ongoing initiatives. It builds on previous efforts within NICEST and NordicESM, and planned activities includes enhancing the performance of and optimizing workflows used in climate models, so that these can be run in an efficient way on future computing resources (like the EuroHPC LUMI system); widen the usage of and expertise on evaluating Earth System Models and developing new diagnostic modules for the Nordic region within the ESMValTool; creating a roadmap for FAIRification of Nordic climate model data.

Puhuri
Puhuri aims for developing seamless access to the EuroHPC LUMI supercomputer, as well as to other e-infrastructure resources used by researchers, public administration and innovating industry. The project will develop and deploy trans-national services for resource allocation and tracking, as well as federated access, authorization and group management.

PaRI
When the Covid-19 pandemic hit, the importance of cross-border collaboration and time-critical information exchange became evident.

> **"Already being a distributed organisation, NeIC was able to respond to the pandemic quickly, and drawing on suddenly un-spent travel budgets, PaRI was set up to facilitate exchange of information with an exceptional number of individual researchers."**
>
> Hans A. Eide Sigma2's representative in the NeIC Board.

PaRI also addresses the demand for services based on the real-time comparison of different epidemiological analyses of the same data, as well as a major need for exchange of more person-sensitive information between authorities, researchers and different countries.

CodeRefinery2
The goal of the CodeRefinery project is to provide researchers with training in the necessary tools and techniques to create sustainable, modular, reusable, and reproducible software. The result of this project are training events, training materials and training frameworks, as well as a set of software development e-infrastructure solutions, coupled with necessary technical expertise, on-boarding activities and best practices guides, which together form a Nordic platform.

EOSC

A gateway to information

The European Open Science Cloud (EOSC) is an environment for hosting and processing research data to support EU science Sigma2 participates in three EOSC projects: EOSC-Hub, EOSC-Nordic and DICE.

EOSC-Hub (2018-2020) brings together multiple service providers to create the Hub: a single contact point for European researchers and innovators to discover, access, use and reuse a broad spectrum of resources for advanced data-driven research. For researchers, this means broader access to services supporting their scientific discovery and collaboration across disciplinary and geographical boundaries. Sigma2 has contributed with work on data management and service for sensitive data.

EOSC Nordic is focusing on the implementation of the EOSC on a regional level. The activities include promoting open science policies for cross-border research, provisioning of services and support for the implementation of FAIR principles. Sigma2 is responsible for one of the work packages, whose goal is to pilot services in the Nordic region that in the future can scale-out in the EOSC. Proof of concept has been demonstrated for cross-border processing for cloud and portal use cases and is in progress for sensitive data use cases. Sigma2 is also actively contributing to mapping policy for open data in the country and the Nordic region.

DICE is aiming at powering up the EOSC with services easily available to the end users, possibly free at the point of use. In DICE, services are offered through EOSC portals and marketplaces through the virtual access mechanism, while the EU covers the costs of the services to the service providers. The aim is to facilitate the cross borders research while avoiding difficult mechanisms for national allocation of resources to foreign research collaborators. Sigma2 has pledged storage and sensitive data resources to the EOSC through DICE.

SUPPORT TO RESEARCH INFRASTRUCTURES

ESFRI, the European Strategy Forum on Research Infrastructures, is a strategic instrument for research infrastructures in Europe, with a mission to facilitate multilateral initiatives leading to the better use and development of research infrastructures, at EU and international level.

Sigma2 experts support the work og the ESFRI forum. ESFRI has established a Strategy Working Group on Data, Computing and Digital Research Infrastructures (SWG DIGIT). The group consists of 15 members, including one from Sigma2, who are all e-infrastructure experts proposed by the ESFRI Forum and European Commission representatives. Additionally, the Norwegian ministry og Knowledge has appointed an expert from Sigma2 to the Forum itself.

In **EuroHPC**, the Infrastructure Advisory Group (INFRAG) is one of the advisory groups making up the Industrial and Scientific Advisory Board. It is composed of 12 members appointed by the EuroHPC Governing Board, one of which is from Sigma2. The INFRAG provides advice to the Governing Board on the acquisition and operation of the supercomputers, drawing up and regularly updating the multi-annual strategic agenda for such acquisitions.

SIGMA2 EMPLOYEES

Gunnar Bøe
Managing Director

Jenny Andrea Amundsen
Special Advisor

Andreas Bach
Senior Systems Developer

Hans A. Eide
Special Advisor

Vigdis Guldseth
Senior Advisor

Steinar Gundersen
Coordinator

Maria Francesca Iozzi
Senior Advisor

Stein Inge Knarbakk
Senior Project Manager

Roger Kvam
Senior Project Manager

Marius Linge
Coordinator

Carl Thomas Stene
Project Controller

Helge Stranden
Senior Advisor

Kjersti Strømme
Head of Communications

Lorand Janos Szentannai
Senior Advisor

PRESENTING THE NORWEGIAN E-INFRA-STRUCTURE

This is UNINETT Sigma2

UNINETT Sigma2 AS (Sigma2) has a strategic responsibility for and manages the national e-infrastructure for large-scale data and computational science in Norway. We provide services for high-performance computing and data storage to individuals and groups involved in research and education at all Norwegian universities and colleges, and other publicly funded organizations and projects. In addition, we also coordinate Norway's participation in Nordic and European e-infrastructure organisations and projects. Our activities are jointly financed by the Research Council of Norway (RCN) and the Sigma2 consortium partners, which are the universities in Oslo, Bergen, Trondheim and Tromsø.

The board is chaired by the managing director of Uninett and consists of four members from the consortium partners, a legal expert from a national research institute and one external representative from outside Norway. In 2020, the board consisted of:

- Tom Røtting, (Chairman) CEO, Uninett AS
- Morten Dæhlen, formerly Dean at the Faculty of Mathematics and Natural Science, UiO
- Nathalie Reuter, Professor at the Department of Chemistry, UiB
- Ingela Nyström, Professor, Uppsala University
- Kenneth Ruud, Pro-Rector Research, UiT
- Therese Løvås, Head of Department of Energy and Process Engineering, NTNU
- Øivind Hennestad, Corporate Lawyer, SINTEF

Solveig Kristensen, Dean, Faculty of Mathematics and Natural Science at UiO and Helge Dahle, Dean, Faculty of Mathematics and Natural Science at UiB entered the board in 2021.

[Information about financial results]

ABOUT THE METACENTER

An important foundation for the successful collaboration on e-infrastructure in Norway is the Metacenter, a collaboration of five organisations to pool competencies, resources and services.

The Metacenter consists of highly qualified staff located at the partner universities and in Sigma2, counting nearly 50 people.

The core activities in the Metacenter are:

- Joint operation of the national HPC and storage systems.
- Unified user support, including a common national help desk.
- Application management.
- High Level Support/Advanced User Support.
- Shared training and dissemination resources and activities.
- Resource allocation (including common access management across all national systems).
- Participation in international partnerships and projects related to e-infrastructure.
- Projects' participation and development.

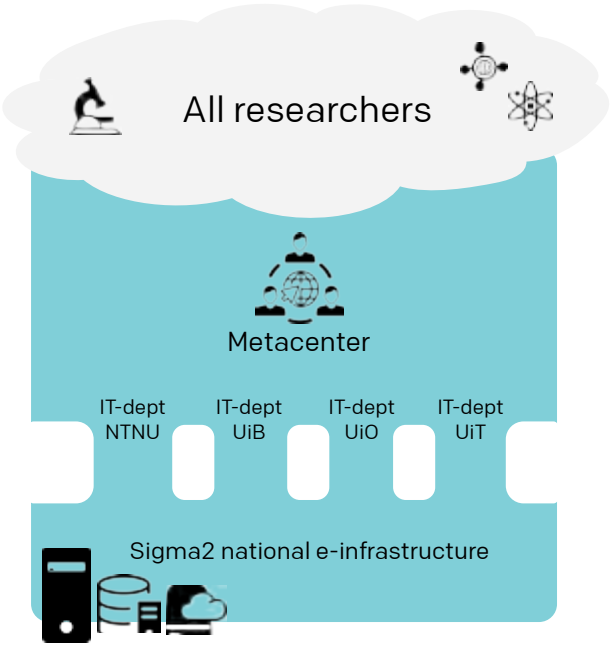

During 2020, we have worked actively with the identity of the Metacenter to increase the organisation's strength. We have created a new platform, which includes a clear vision for our work: "We enhance excellent research for a better world". This work will continue in 2021.

Vision and key strategic areas

Sigma2's vision is to provide a permanent, predictable and cost-efficient e-infrastructure with access based on scientific merit in order to maximize the impact and return of scientific research.

Our values are:

- customer focus
- corporate social responsibility
- openness

Our key strategic areas are to:

- provide advanced compute- and data services addressing user's needs.
- facilitate for international research services, cloud-based resources and common components.
- combine Sigma2 infrastructure with the capabilities and competences of partner institutions to achieve excellent research services.

UN sustainability goals

Through its broader impacts, the national e-infrastructure will help address societal challenges on a global level. Referring to the UN Sustainable Development Goals, nearly all of the 17 goals are addressed through research enabled and supported by e-infrastructure, either directly as in environment and resource modelling, or by extension through support of research infrastructures and enabling collaboration.

Sigma2's core value of social corporate responsibility is important in this context as well. We have focus on diversity and equal opportunities in our recruitments and our work environment is monitored closely. For the external environment, Sigma2 puts great emphasis on re-use of ICT equipment and for procurement and operations re-use of heat is an important area. For our latest HPC system, Betzy, the re-use of the energy is 100%.