SNIC Newsletter

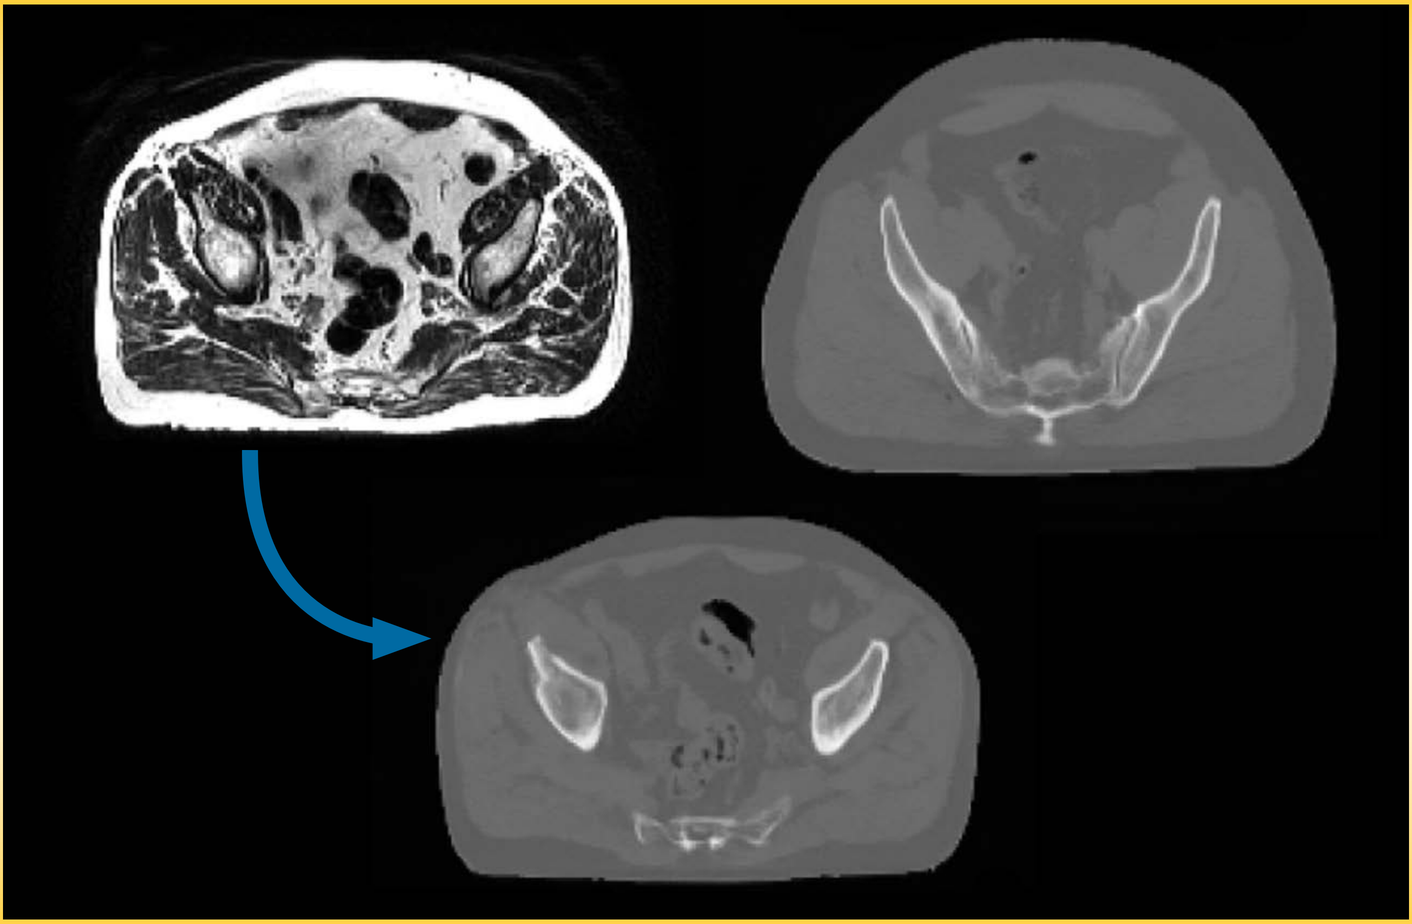

Winter 2022

Published by the Swedish National Infrastructure for Computing

Contents

View from the SNIC director's desk

Welcome to the last SNIC newsletter! By the time you read this, SNIC may already have ceased to exist, and, although tidings of joy are coming, this is not the occasion to be nostalgic. Instead, we would like to take the opportunity to reflect on what has happened over the last five years. During this period, which is referred to as SNIC 2.0, SNIC existed in the organisational form of a consortium comprising the ten largest research universities in Sweden and was hosted by Uppsala University. Initially the Swedish Research Council (VR) mandated that such infrastructures had to be set up in the form of consortia. Now, five years later, they have revised their point of view and decided that infrastructures should be established as organisations situated at single universities to simplify governance.

This change means that the next national infrastructure for providing high-performance computing (HPC) resources cannot be in the form of a consortium. A consortium is a somewhat unwieldy organisation with many different partners, each with their own agenda, and the infrastructures that were established as consortia had to be co-financed by the partners, partly with in-kind contributions. Thus, this alteration of the organisational form is advantageous in some senses. It will give the coming organisation – the National Academic Infrastructure for Supercomputing in Sweden (NAISS), which will be assuming SNIC's role from the 1st of January 2023 – more operational freedom. For example, a central theme of the application to VR for the new national infrastructure was the consolidation from several geographically distributed HPC centres to fewer, and preferably just one. Although that consolidation process had already commenced within SNIC, realising it was slowed down by the organisational structure and the requirements for consensus between the partners.

In hindsight, however, there is also a definite advantage with a consortium, namely that the partners are committed to collaborating. Over the last five years, a culture of national collaboration has matured within the SNIC consortium. For instance, there is strong agreement that the way resources are allocated to research projects is extremely important. The allocation of SNIC resources (such as a number of CPU hours on a particular HPC system or an amount of storage for data) is primarily handled by the Swedish National Allocations Committee (SNAC). The committee members have always worked together on the evaluation of applications requesting larger allocations on any SNIC resources. However, earlier on, members tended to only evaluate applications for smaller projects on resources at their local SNIC centre. The role of SNAC has evolved towards tighter collaboration with members now taking turns to make decisions on smaller allocations for several of the main SNIC systems. In addition, submission dates and assessment periods for applications are now coordinated so researchers can expect the same time frames across all the SNIC resources. This is advantageous, particularly if a project needs to request allocations on several SNIC resources for different parts of the research (such as running simulations on one system and analysing the resulting data on another). There are clear benefits for Swedish research with this kind of collaborative approach as the available SNIC resources can be distributed for optimal effectiveness such that the maximum number of eligible research groups are awarded allocations that meet the requirements for their research as closely as possible. NAISS is cognisant of this culture of collaboration and is keen to keep the cooperative spirit thriving in the new organisation.

Many voices in the SNIC community network also share an awareness of the importance of another area that benefits from national collaboration, namely the provision of user support. During SNIC 2.0, the collaboration between the centres has grown stronger: regular coordination meetings are held and it is now natural for questions from users to be directed to whichever SNIC centre can best assist, so an expert working at one centre might help a researcher using a system at another centre. The centres also jointly run monthly Zoom meetings where researchers can get assistance. Looking to the future though, it is important to be aware that, as a consequence of the preconditions associated with the VR call, the focus on providing support for the researchers using the various national resources in the Swedish HPC infrastructure has diminished significantly for NAISS. National user support has been an integral part of SNIC, supporting Swedish research just as much as the other pillars of the operation. Partners within the consortium emphasise the importance of such support (see, for example, https://www.pdc.kth.se/publications/pdc-newsletter-articles/2022-no-2/supporting-research-software-development-1.1208807). NAISS is naturally also very well aware of this shortcoming and of the need to identify and realise viable solutions, the major obstacle being the difficulties inherent in coordinating such activities without a pot to divide for funding such endeavours or at least a consortium agreement to rely upon.

At the time of writing, the signing of the NAISS agreement between the VR and Linköping University is expected to take place on the 12th of December this year. NAISS will be governed by a steering committee led by the venerable Jan-Eric Sundgren (https://en.wikipedia.org/wiki/Jan-Eric_Sundgren). SNIC guidelines and procedures will prevail until superseded by NAISS equivalents, with a new user agreement expected shortly.

As already mentioned in the previous newsletter, the swan song for SNIC and the first steps of the dance for NAISS are being devoted to ensuring a smooth transition at the change of the guard. The aim is that the researchers who are using the Swedish national resources for computation and data handling should not notice anything (apart from the change in the logo) as we enter the year of grace 2023, to which we look forward with confidence and great optimism. May it be a happy and peaceful new year for everyone!

Lars Nordström, SNIC Director

Note that applications for EuroHPC allocations on the large European HPC systems are submitted through EuroHPC calls, rather than via the earlier PRACE calls mentioned in the previous newsletter. Information about EuroHPC calls is available on the PRACE website: https://prace-ri.eu/hpc-access/calls-for-proposals.

NAISS FAQs

The computational and data storage resources currently under the umbrella of the Swedish National Infrastructure for Computing (SNIC) will instead be handled by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) from the start of 2023 onwards.

Here are answers to some common questions that researchers who have been using SNIC resources have asked about how things will work under NAISS.

What types of resources will NAISS provide to support Swedish research?

NAISS will be providing high-performance computing (HPC), storage, and data resources, including systems suitable for machine learning (ML) and artificial intelligence (AI) tasks. Initially NAISS will take over responsibility for all existing SNIC systems that were intended to run beyond 2022. From the perspective of researchers who use NAISS resources, 2023 will be extremely similar to 2022. NAISS will gradually set up new resources, replacing older systems as they are retired, to meet the growing needs of Swedish researchers using HPC resources for academic research. Quantum computing (QC) resources are not yet offered.

Does NAISS have a website? If not, will it have one soon?

The NAISS website, https://www.naiss.se, is under construction.

Our project was awarded a SNIC allocation that runs into 2023 – will anything about that change? Will we still be able to log in and use the systems/storage in the same way?

All SNIC allocations of time, storage or application expert assistance that run into 2023 will be continued under NAISS and will effectively stay the same. Login, software, hardware, the SNIC User and Project Repository (SUPR) portal and so forth will stay the same. Projects that have allocations will be advised of any minor changes, such as acknowledgement text to be used in publications or updates to user agreements in early 2023.

Our project has been awarded a new SNIC allocation that will start in 2023 - how will we log in and use the systems/storage?

In exactly the same way as previously under SNIC during 2022. The systems and storage will stay the same. The user and project portal, SUPR, will also stay the same.

How, when and where can our project apply for a NAISS allocation (that is, an allocation that will start in 2023 or later)?

From the 1st of January 2023, applications for NAISS allocations can be put in through the SUPR portal, just as under SNIC during 2022. Any applications that are submitted to SNIC very late in 2022 will be handed over to NAISS for consideration in 2023.

Where/how can our project get help to prepare an application for a NAISS allocation?

All applications for NAISS allocations will be done through the same SUPR portal as previously under SNIC. The current SNIC website (https://www.snic.se) gives details of all the documentation that needs to be provided, and the site will be kept online for a transition period until the NAISS web resources have been fully updated. If additional help is needed, contact info@naiss.se.

From the start of 2023 and onwards, who should we contact if our project has problems using NAISS resources?

Support for problems using NAISS resources should be sought in the same ways as under SNIC (for example, using the support form in SUPR). NAISS (including the local NAISS branches) will continue to run existing systems and provide support for them in terms of system support, as well as first-line and mid-level user support. Advanced user support, which has mainly been funded by Swedish universities as in-kind contributions to SNIC, will continue to be funded by the universities and will be coordinated by NAISS.

How will support services work under NAISS?

Support for users of NAISS resources in 2023 will largely be the same as with SNIC in 2022. The SUPR support form can be used, and, if help is needed when using a specific resource, the corresponding support website and e-mail address can be used. For example, with work on Tetralith, the https://nsc.liu.se/support site and support@nsc.liu.se email address can be used, and similarly for Dardel, Bianca and other systems. Information on support and ways to obtain help will continue to be present in the SUPR portal. The https://www.snic.se website and its support content will be kept online until the https://www.naiss.se site is complete.

Will NAISS be running training events to help researchers to improve their high-performance computing skills and/or to use NAISS systems?

Yes. NAISS will run training events for this purpose, building on the positive experience from the SNIC training. Information about these activities will be provided on the https://www.naiss.se website and through e-mail to researchers who are using NAISS resources.

Latest news from SNAC

The Swedish National Allocations Committee (SNAC) received 39 applications in the SNIC “Large Compute Fall 2022” call. Overall, the applications were of very high quality, and the committee would have liked to have been able to award more time. The available CPU resources were oversubscribed by 27%, with the Tetralith system in particular being almost 100% oversubscribed. The high pressure on resources was also apparent for the graphics processing unit (GPU) partition of Dardel, which was more than three times oversubscribed. As a result, the size of many allocations had to be reduced, and some allocations had to be shifted from Tetralith to the CPU partition of Dardel.

SNAC received 14 applications for the “LUMI Sweden Fall 2022” call, which requested more than twice the number of core hours that were initially available. Fortunately, the amount of available time was increased unexpectedly, and, as a result, it was possible to accommodate most applications with allocations of core hours quite close to the requested amounts. Note that, from this point onward, there will be only one regular LUMI call per year.

The new GPU partition of LUMI (known as LUMI-G) is presently in its pilot phase, and the plan is for it to be officially opened next year. SNAC has allocated time on the partition under the assumption that LUMI-G will become fully operational in the near future. However, due to the aforementioned uncertainties, access cannot be guaranteed at this stage.

During the transition from SNIC to NAISS, SNAC will continue to operate as usual to minimise the impact of the transition on researchers applying for allocations. After ten years as a member of SNAC, Philipp Schlatter will step down as the chair by the end of this year. Paul Erhart, who has been the vice-chair of SNAC, will take over as acting chair during the transition period.

Easily accessing European HPC via Puhuri

If you do research with the LUMI pre-exascale system, you are likely to use the work of the Puhuri project without being aware of it. Here is a behind-the-scenes look into the benefits of Puhuri and the future of accessing allocations on pan-European high-performance computing (HPC) systems and HPC-related services.

The Puhuri project was established in mid-2020 by the Nordic e-Infrastructure Collaboration (NeIC), the Swedish National Infrastructure for Computing (SNIC) and similar organisations from Denmark, Estonia, Finland, Iceland and Norway. The key driver for those countries in initiating the Puhuri project was the need for an authentication and allocation infrastructure (AAI) for the EuroHPC LUMI supercomputer that was being installed in Kajaani, Finland, at the time. Although LUMI is located in Finland, the system is hosted by a consortium of ten countries, and, as access to the system and management of allocations on it was to be handled both by the EuroHPC Joint Undertaking and by the consortium partners, it was clear that the partners needed a shared way to authenticate or identify researchers who accessed the system and to manage the allocations. It was decided that a suitable AAI would be developed under the auspices of NeIC in the form of the Puhuri project. The project built on experience from an earlier NeIC project, known as Dellingr, which looked into the feasibility of sharing and exchanging HPC resources (such as computing, storage and support services) between countries and across borders.

Puhuri is a bridge between researchers and HPC resources

The Puhuri project was set up to solve the problem of distributed resource allocation, that is, to find a way to enable researchers to have seamless cross-border access to the HPC resources they need. (“Seamless” access means that it is possible to get access to the resources in the same way, irrespective of where the resources or users are located.) The Puhuri system was therefore designed to create solutions for each step of the interaction between the users and the resources, from identifying or authenticating users to allocating specific resources and tracking resource usage. In principle, Puhuri constructs a bridge between the new LUMI supercomputer and the researchers who use it.

A key feature of the Puhuri system is the concept of portals where the allocation, access, and reporting processes can be handled. These portals can be provided by Puhuri with various levels of customisation for different countries or already existing national portals can be fully integrated using the Puhuri application programming interfaces (APIs). From the Swedish perspective, researchers who wish to use resources provided through Puhuri (currently the Swedish share of resources on LUMI) can request access through the SNIC User and Project Repository (SUPR). As more systems are added to the Puhuri service, there will be no need for researchers to request access to a resource provided via Puhuri through any other portal; it will be possible to handle all requests and interactions via SUPR (in a process that will utilise the Puhuri API). This will ensure that researchers will not have to learn to navigate through yet another allocation interface to get access to a resource in another country. In fact, Swedish researchers and people responsible for allocations for Sweden should not even notice that the resource management is being handled by Puhuri, rather than by SUPR, as the interactions are managed through the same interface, by virtue of a Puhuri API.

Simplifying remote access using identify federation

When providing researchers from across Europe with access to HPC resources, it is vital to be able to identify the person who is accessing a given resource. Within Sweden this is made possible through an infrastructure known as the Swedish Academic Identity federation, SWAMID. However, on the European level, no such all-encompassing identity federation has been available. So the Puhuri project has been working with the pan-European data network GÉANT and the FENIX infrastructure to develop the MyAccessID Identity and Access Management Service, which provides a common identity layer for Infrastructure Service Domains (ISDs) – in other words, MyAccessID enables researchers to use their home institution’s identity service to identify themselves with the appropriate level of assurance so they can easily access HPC resources located elsewhere.

Given the current situation in international relations, knowing exactly who is using particular HPC resources is of utmost importance. MyAccessID and the Puhuri system make it possible to accurately identify people who try to access European HPC resources, thus reducing the risks of unauthorised access.

Reporting usage and assigning allocations easily with Puhuri

The Puhuri system provides allocators (that is, the organisations that grant allocations for accessing particular HPC resources) with extensive reporting information about the utilisation of those HPC systems by individual research groups and researchers. Any HPC resource that is connected through Puhuri reports these statistics back to the Puhuri portal where the allocators, as well as the researchers using the resource, can view and make use of that information. In addition, when an allocator uses Puhuri to assign an allocation to a particular research group or project, the details about the allocation (such as the amounts of time and/or memory that were awarded) will be automatically propagated to the corresponding HPC resource and the allocation will be prepared for use. This is highly efficient as there is no need for the system staff from the organisation providing the HPC resource to take any action to establish the allocation. Automating both reporting and allocation handling in this way not only makes the process more robust, but also saves time and money for both the allocator and resource provider organisations.

Solving broader and future access with Puhuri

The Puhuri solution, which was initially developed for the LUMI system, was designed to be equally able to serve other HPC resources as well. For example, the Puhuri solution will be used in NeIC’s new project, NordIQuEst, which aims to build a Nordic-Estonian e-infrastructure for quantum computing, and also for the coming European quantum computer resource, LUMI-Q, which will be located at the IT4Innovations National Supercomputing Center in the Czech Republic. In principle, Puhuri can manage the AAI and reporting needs of any infrastructure that is sharing limited resources between researchers from multiple organisations or countries.

The trend in research and science is for there to be progressively more collaboration across national borders. At the same time, the cost of large-scale computing resources is approaching the limit of what can be financed with just the research budget of a single nation. Considering these facts, it is becoming clear that research in the future will rely more and more on multinational collaborations financing massive shared supercomputing and storage resources. The Puhuri system solves the multinational access and allocation management requirements that are arising from these trends, and the system is well placed to be the bridge between multinational resources and their users in the future.

Deep learning in medical image analysis

Cancer is one of the leading causes of death in Sweden. In recent years there have been about sixty to seventy thousand cancer tumours diagnosed each year and around twenty-three thousand deaths from cancer. Viewed from another perspective, these statistics indicate that one third of the Swedish population will suffer from cancer at some point in their lives. In addition, the world population is ageing, so the healthcare systems around the world will have to handle more cases of disease in the future, including more cancer cases, while there will be fewer people of working age. Consequently, the healthcare systems will need to make more efficient use of the resources available for dealing with cancer and other diseases, particularly in terms of how much time healthcare personnel can spend on each patient.

One way to achieve greater efficiency in healthcare is through the use of machine learning (ML), which can "learn" to perform specific tasks within the medical domain. For the training of large-scale ML models, high-performance computing (HPC) systems have become incredibly and increasingly important, especially when developing what are known as deep learning (DL) models. As an example, ML offers new perspectives in relation to cancer treatment by providing ways to automate routine and time-consuming manual parts of the radiotherapy pipeline. ML can free up significant amounts of time on the part of oncologists, radiologists, and radiation nurses, as well as reduce the amount of time the patients need to spend being assessed. For example, when planning for radiotherapy treatment, both a computed tomography (CT) scan and a magnetic resonance (MR) scan are conducted. A CT scan can take up to half an hour to perform, and an MR scan may take up to an hour. However, an almost-instant synthetic CT image may be generated from an MR image using ML, thereby saving the time spent performing the CT scan and reducing the total time by about a third. Identifying the exact location and extent of a tumour is another essential part of the radiotherapy pipeline that can be speeded up with deep learning. The process of segmenting – that is, manually marking the location of a tumour in anatomical medical images that are used to guide radiation treatment and marking the tissues that should not

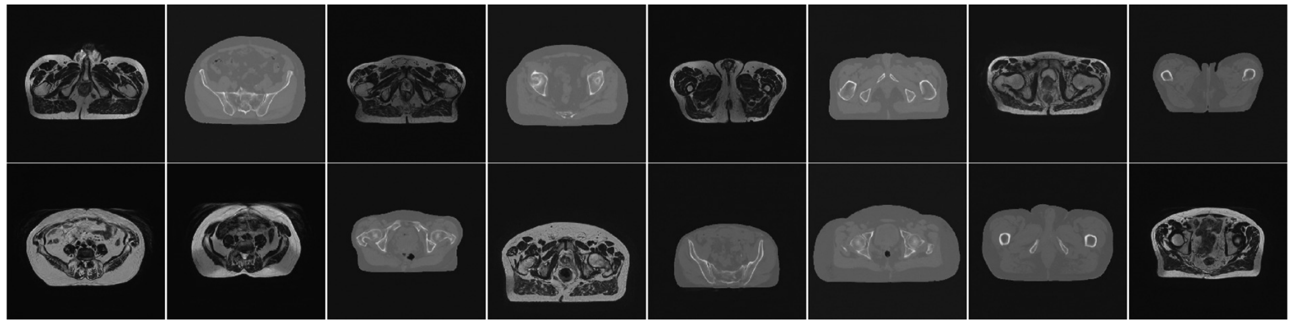

The sixteen random samples above are from an ML model that was trained on paired MR and CT images using images from one hundred patients. The patients had mainly been diagnosed with prostate, rectal, or gynaecological cancer and had undergone both MR and CT scanning. This means that there were both MR and CT images of the same tumours and organs available to help the model learn how to convert images from MR scans to CT images, and vice versa.

The illustration on the cover shows the result of converting a "source" MR scan (top left image) into an image in the style of a "target" CT scan (top right image) and thereby creating an almost instant synthetically produced CT scan image (bottom image).

Both illustrations are taken from material in a paper by Tommy and his colleagues (https://doi.org/10.1016/j.zemedi.2020.05.001) that is released under a Creative Commons licence (CC BY-NC-ND 4.0).

be exposed to radiation (called organs at risk) – can take up to several hours of work by a radiologist and other staff, whereas an automatic segmentation can be produced in less than a second with the help of ML.

Reducing the amount of time that radiologists and other healthcare personnel need to put into tasks like this is of great benefit as that makes it possible for more patients to be treated by a smaller number of staff in less time. However, it is a gargantuan task to manually create a program that explicitly takes into account all the possibilities that might be encountered while doing, say, a segmentation. This is where machine learning comes into its own as large-scale mathematical models are instead automatically trained to perform the desired tasks by using vast numbers of real medical images. The parallel computing capabilities of HPC systems enhanced with graphics processing units (GPUs) enable the models to be trained on tens of thousands of medical images in days instead of weeks or months. Where human medical staff undergo several years of training and then develop their skills further through the experience of examining and treating patients, these machine learning models are, in a sense, like medical staff who have a training and practice period of several years compressed into days, although, of course, each ML model is only able to do one particular task, namely the one it is specifically trained to do. The last few years have seen significant breakthroughs in DL development and usage, and it is encouraging and promising that current methods have also shown promise in medical applications.

Tommy Löfstedt and his colleagues in the Machine Learning group at Umeå University, in collaboration with researchers at the Umeå University hospital, are using the SNIC general-purpose system Kebnekaise to train deep learning models with the ultimate goal of automating the radiotherapy pipeline. They develop ML and DL methods to, for example, generate synthetic CT images from input MR images with improved quality and robustness. They also develop techniques to correct the artefacts in MR images arising from patient movements or scanner imperfections, which also, for instance, improves the quality of the synthetic CT images that are produced from those MR images. In addition, they develop resource-efficient solutions involving convolutional neural networks to automate other parts of the radiotherapy procedure, such as automatically detecting and/or segmenting tumours and organs at risk in MR and CT images. The researchers also work on data heterogeneity problems with medical imaging data, for example, where the distribution of different patient groups is non-uniform (for instance, with respect to age), which can lead to problems for the underrepresented patient group or groups. These problems are amplified by DL models in areas where there are small amounts of data, and thus it has become important to solve these problems so that, in the future, DL will be able to reliably assist the widest possible range of patients in terms of medical imaging.

Tommy's group and their research partners have been using about 150,000 core hours per month on the SNIC system, Kebnekaise, which is hosted at Umeå University. Their need for more GPU hours is increasing as their projects progress and as the group and their number of collaborations grow. While it may sound like they are using a lot of time on the SNIC supercomputer system, if the same work was being done on a desktop computer with a single GPU, it would take weeks, or even months, to train models that can be trained in a few days on Kebnekaise. More importantly though, once the machine learning models have been trained to perform appropriately, they only take an instant to use when, for instance, preparing radiation treatment for a patient. This saving of time is of great benefit to radiologists, oncologists, and patients alike.

For more information about this research group and their projects, see https://www.umu.se/en/research/groups/machine-learning.

SNIC Science Cloud gets extra GPUs

The SNIC Science Cloud (SSC) is a national cloud infrastructure based on three geographically distributed regions in Sweden: Umeå University in the north, Uppsala University in the east and Chalmers University of Technology in the west. The SSC is a complement to the SNIC supercomputing systems and provides researchers with more flexibility in how applications can be implemented compared to using traditional batch computing systems.

The SSC started in 2014 as an exploratory project and has since evolved into a fully fledged cloud infrastructure facilitating Swedish research. The allocation of resources in SSC differs from traditional batch computing systems where you are granted a number of core hours. Instead, projects are awarded "SNIC coins" for the SSC. You can then use the SNIC coins to pay for SSC resources in the same way you would in a commercial cloud.

Using a self-service portal, you and your research team will have full access to your research environment to set up compute instances with the Linux distribution of your choice, and to configure networks and storage to tailor it just right for your workflows. Also, just like in the commercial clouds, you can use an application programming interface (API) to programmatically provision your research environment on demand and make it elastic so that it scales up and scales down dynamically to the size that is just right for your research needs at any given time.

We are also happy to announce that additional graphics processing units (GPUs) have been added to the SSC this autumn. The east region now provides access to NVIDIA T4 and A2 GPUs, and in the north region you can access NVIDIA A100 GPUs. The T4 and A2 GPUs are designed primarily for fast machine learning. They will provide high single-precision performance and are ideal for your AI training or if you need high inference performance.

For information about when it is a good idea to use the SSC for your research, rather than other SNIC resources, see https://cloud.snic.se/about.

SNIC centre round-up

C3SE

The Chalmers e-commons infrastructure provides support in all aspects of digital research, including collecting and managing data. The infrastructure is now calling for pilot projects led by Chalmers researchers. The projects will be supported by Chalmers e-commons digital research engineers who will provide support as project partners in areas such as high-performance computing (HPC), artificial intelligence/machine learning (AI/ML) and visualisation. The visualisation projects will run within the national research infrastructure InfraVis (https://infravis.se).

Chalmers e-commons is now establishing the Swedish e-infrastructure node for the global Square Kilometer Array (SKA) radio astronomy infrastructure together with the Onsala Space Observatory. The data-driven life science (DDLS) national Data Science Node for Cell- and Molecular Biology will be established in close collaboration with the SciLifeLab Data Centre and the Chalmers Infrastructure for Systems Biology.

HPC2N

As of January 2023, the Kebnekaise system will not be available as a national resource. However, due to the high demand for its graphics processing unit (GPU) and large memory nodes amongst researchers in Sweden, HPC2N and the National Academic Infrastructure for Supercomputing in Sweden (NAISS) are holding ongoing discussions about making parts of Kebnekaise available nationally in 2023. For more information, see https://www.hpc2n.umu.se/kebnekaise-future-info. HPC2N is continuing to provide all its other current services.

LUNARC

The installation of COSMOS is progressing well, and most of the hardware has been delivered to Sweden. However, some minor components still need to be added due to issues with sourcing components. The data centre cooling infrastructure has been upgraded with new RITTAL in-row coolers, and the COSMOS hardware installation is planned to commence within a month. Deployment and acceptance testing is scheduled for January/February 2023.

To help new and existing LUNARC users at Lund University take advantage of COSMOS, LUNARC sent out a call, COSMOS Dreams, asking for proposals that the users would need assistance to implement. LUNARC will select some of these proposals and work with the users to implement them on COSMOS.

NSC

NSC's main systems Tetralith and Berzelius continue normal operations. The second phase of Berzelius is being scheduled for installation during the first quarter of 2023.

Linköping University is in the process of establishing NAISS, which will take over from SNIC on the 1st of January 2023 . NSC is assisting the university with this. The handover of present SNIC resources is in progress for a seamless transition from the perspective of users. At the end of November, SNIC and NAISS held an all-hands meeting to help facilitate the transition.

PDC

The GPU partition of the Dardel HPC system (which includes 56 GPU nodes) has been installed. That put the system in 5th place on the latest Green500 list (https://www.top500.org/lists/green500/2022/11) and 68th place on the TOP500 list (https://www.top500.org/lists/top500/2022/11) so it is now one of the most power-efficient systems in the world. HPE is building the final set of CPU-only compute nodes for Dardel, and the plan is for them to be installed in December.

In parallel, hardware from Dell for the new Dardel Cloud system has been brought up. PDC is currently working with the service providers IP Solutions and Engin IT to install the compute and storage middleware for the Dardel Cloud. The plan is for the system to be available for early user access at the beginning of 2023.

UPPMAX

UPPMAX has a new application expert, Matias Piqueras, with a focus on supporting the digital humanities user community. The storage system connected to Rackham, Crex, was reaching the limit of its previous capacity and has now been upgraded to a total of 9 PB with a corresponding upgrade to the number of files that can be handled. The graphics processing unit (GPU) capacity of the east region of the SNIC Science Cloud (SSC) has been extended from one dual NVIDIA A100 GPU node to 25 GPU nodes. These additional nodes have either one NVIDIA A2 GPU (20 in total) or one NVIDIA T4 GPU (4 in total). The SSC east region has also been equipped with new nodes with 284/768 GB RAM and 48 cores. The new capacity makes the SSC a suitable system for developing and evaluating distributed machine learning algorithms. Around 70 nodes with 256 GB RAM have been added to Bianca, which made it possible to increase the size of the login nodes, thereby supporting interactive usage. The number of dual A100 GPU nodes has increased from 4 to 10, and 12 new high-memory nodes with 512 GB RAM and 64 cores are coming soon.

HPC opportunities & organisation links

HPC in Sweden and Scandinavia

ENCCS: http://enccs.se

eSSENCE: http://essenceofescience.se

NAISS: https://naiss.se

NeIC: https://neic.no

SeRC: https://e-science.se

SeSE: https://sese.nu

SNIC: https://snic.se

SNIC centre project participation

BioExcel CoE: https://bioexcel.eu

EBRAINS: https://ebrains.eu

PerMedCoE: https://permedcoe.eu

EXCELLERAT: https://www.excellerat.eu

EOSC-Nordic: https://eosc-nordic.eu

DICE: https://www.dice-eosc.eu

HPC-Europa3: http://www.hpc-europa.eu

European HPC ecosystem

HPC in Europe: https://hpc-portal.eu

EuroHPC: https://eurohpc-ju.europa.eu

PRACE: https://www.prace-ri.eu

LUMI: https://www.lumi-supercomputer.eu

ETP4HPC: https://www.etp4hpc.eu

EOSC: https://eosc-portal.eu

General HPC news sources

HPCwire: http://www.hpcwire.com

insideHPC: https://insidehpc.com

Upcoming Swedish HPC training & events

Next year, a wide range of high-performance computing (HPC) training will continue to be provided for researchers in Sweden, even though SNIC is coming to an end. Various organisations, including NAISS, are committed to this important activity. Here are preliminary announcements for a variety of training events that will be held early in 2023 – please save the dates! More events are planned, and details will be made available next year.

- **Introduction to UPPMAX**: 9-11 & 13 January 2023, provided by UPPMAX
- **Introduction to HPC2N**: 12-13 January 2023, provided by HPC2N
- **AWK workshop**: 16-17 January 2023, provided by UPPMAX
- **Targeting chemical accuracy at the exascale with Quantum Monte Carlo**: 26-27 & 30-31 January 2023, provided by ENCCS and the TREX Centre of Excellence
- **Introduction to running R, Python and Julia in an HPC environment**: 8-10 February 2023, provided by HPC2N and UPPMAX
- **SciML GPU bootcamp**: 21-22 February 2023, provided by NVIDIA, OpenACC & ENCCS
- **Introduction to ALVIS**: February 2023, provided by C3SE
- **An introduction to the UNIX/LINUX command line**: 1 March 2023, provided by LUNARC
- **Dardel GPU hackathon**: 7 & 14-16 March 2023, provided by PDC, ENCCS, AMD & HPE
- **Cluster architecture and job submission**: 8 March 2023, provided by LUNARC
- **CodeRefinery workshop**: 21-23 & 28-30 March 2023, provided by CodeRefinery, UPPMAX, PDC & ENCCS