NAISS

National Academic Infrastructure
for Supercomputing in Sweden

Annual Report 2025

Executive Summary

In early 2025, NAISS underwent an interim review that resulted in a very positive report strongly supporting the new national organisation, including suggestions about how to build stronger national collaborations with partner universities.

These were incorporated in the proposal for renewed long-term funding for the period 2027-2030 that NAISS submitted to The Research Council autumn 2025, in particular focusing new funding on a significantly expanded user support organisation.

During the year, the infrastructure provided some 82 million core-hours to 3,352 unique approved research projects. At least 1,447 scientific papers acknowledging NAISS resources were published, and numerous PhD theses defended. This emphasises what a critical resource NAISS is across all fields of science in Sweden, both for high-performance computing, artificial intelligence and sensitive data processing.

A new national data centre facility has been completed in Linköping. The procurement for the EuroHPC mid-range system Arrhenius was concluded in summer 2025, which resulted in a contract for a system with GPU partition using 1,528 Grace-Hopper GPU superchips, a separate CPU partition with 54,272 AMD x86 cores, and almost 30PB of storage. The installation of the system is currently being finalised, and it is expected to be available to research users in late spring 2026.

The year also represented a major focus on artificial intelligence, where NAISS established Sweden's new AI Factory, Mimer, in collaboration with RISE AB. This is a rapidly expanding NAISS unit that provides extensive support services to both academia and industry.

As part of the AI Factory investments, NAISS will also host a new dedicated computer resource entirely focused on AI usage; the contracts are currently being finalised together with EuroHPC, targeting an installation towards the end of 2026.

Professor Erik Lindahl started as new NAISS Director in January 2025, and the staffing of the organisation has expanded significantly, in particular related to new EuroHPC projects.

The infrastructure leadership has been reorganised with dedicated associate directors for each module, including the recruitment of Dr. Rossen Apostolov as AI Factory Director, and Dr. Joel Hedlund as NAISS's Data Science Director.

Professor Matts Karlsson has been appointed NAISS Director of External Relations, Andreas Johansson as Director of Technical Operations, Associate Professor Peter Münger has the role as Director of Allocations, and Dr. Henric Zazzi is Acting Director of User Support.

To create a clear brand for the joint national computational infrastructure, Linköping University has also decided to no longer operate its previous NSC centre and instead contract NAISS to provide both its local resources, operate services for collaborations with external actors, and host the special AI systems funded by the Knut and Alice Wallenberg Foundation. This gives NAISS a unique mandate not only as the Research Council-funded infrastructure but representing all major public investments and strategic efforts related to HPC, AI and data in Sweden.

There are numerous future challenges, particularly related to aiding users through the sometimes-difficult transition of moving workloads to accelerators. This is not unique to NAISS, but a necessary change occurring all over the world in response to hardware roadmaps, energy efficiency and requirements for green computing, which in turn are subject to the laws of physics.

The researcher-directed advanced user support at the partner universities is a critical resource for this, and NAISS will help organise a series of activities and training events to aid this type of work in collaboration with the Swedish e-Science research initiatives SeRC and eSSENCE.

Cover photo: US participants from the EU-US Emerging Leaders Visitors Programme visiting NAISS to discuss strategies in HPC and AI with the NAISS leadership, March 2026.

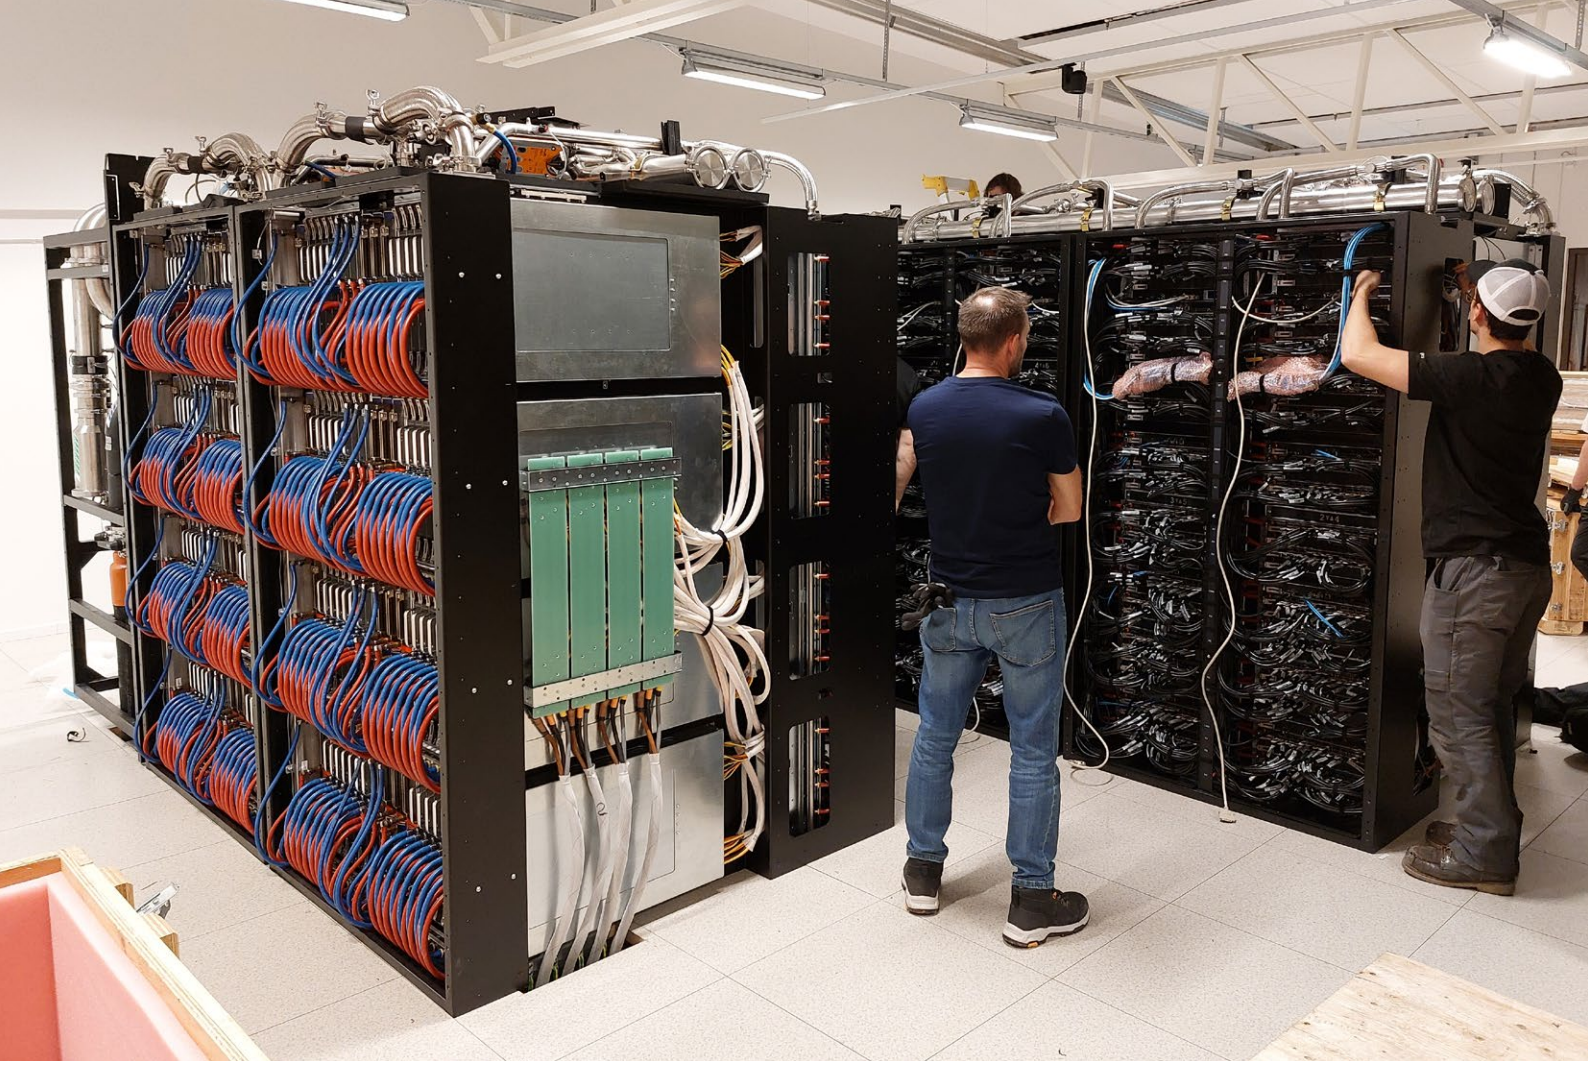

Installation of the Arrhenius system in the new computer hall in Kärnhuset, December 2025.

Contents

Photographers: Jonathan Pakvis, Andreas Johansson, Anna Jänis, Magnus Johansson, Thor Balkhed, Gabrielle Beans (Tommy Löfstedt)
Layout: Jonathan Pakvis
Print: LiU-Tryck 2026

Director's Overview

The year 2025 marked a period of consolidation, unprecedented expansion, and a strategic pivot for NAISS. After two intensive years of establishing contractual agreements, collaborations, changes in staffing and securing additional funding, we have been able to focus on expansion and the future.

While we easily take it for granted, the most important task of NAISS is to deliver stable services to our users, and just as in previous years our staff did an outstanding job achieving this for the Dardel, Tetralith, Alvis and Bianca systems. We are thrilled to have more users than ever and a similar record in the number of projects.

We are doing our best to justify to funding agencies how important this computational infrastructure is for the research and that the government's investments need to keep up with the usage, while being careful not to argue resources should be moved from other research funding - because that would mean the project grants paying the students doing the research.

Demand outstrips supply

A defining characteristic was the sheer scale of researcher demand and NAISS's robust operational response. The current systems represent some 30 petaflops of computational capacity, which NAISS used to deliver approximately 82 million core-hours to 3,352 approved research projects.

NAISS is unique internationally in that our infrastructure is directly responsible both for the largest HPC allocations and thousands of small users, e.g. in bioinformatics. While this leads to high demands, e.g. for user support, training, and requirements for different types of services, it also provides strong synergy effects and close collaborations with stakeholders across society.

One challenge is that demand significantly outstrips available resources; some proposals are rejected each year, and in particular for the largest requests it is common with substantial cuts.

While I hope the next-generation resources currently being installed will alleviate some of this, we know from experience that users quickly adapt and find new ways of turning computational power into scientific breakthroughs.

1,447 papers acknowledging NAISS

Sweden in general, and NAISS in particular, have a strong tradition of handling this competition through a fair, transparent, and scientifically prioritised tiered allocation process that we will continue to build upon.

It will remain trivial for new users to get a small allocation to try out NAISS resources, while the largest users will need to provide more extensive technical and scientific justifications used in peer review.

In particular, with the rapid expansion of computing in Europe, and Sweden's new role in EuroHPC, it is also important to remember that users can apply in the EuroHPC calls too. Arrhenius is one of the systems that will be available in those calls, together with a dozen other resources.

While the requirements of these proposals can seem daunting at first, my experience is that they are merely different rather than more difficult, and both NAISS and EuroHPC staff can help with benchmarks and recommendations.

The ultimate metric of our success remains sci-

entific impact, with at least 1,447 scientific papers acknowledging NAISS resources published in 2025, and I want to take this opportunity to thank our users for remembering to acknowledge the infrastructure in their work - it is crucial to show the value we are delivering together.

To meet this rapidly escalating demand, NAISS successfully secured and finalised key infrastructure investments that will increase Sweden's national capacity substantially. Recognising that AI is fundamentally disrupting science as much as industry – and that fields like life science are becoming inherently AI and HPC-driven – we executed a strategic pivot toward artificial intelligence.

Strong support from the Research Council

The new general-purpose Arrhenius resources will be a great resource both for HPC and AI workloads, and during the year we also deployed Sweden's new AI Factory (Mimer) together with RISE AB.

The AI Factory will not only provide dedicated AI resources, but arguably the most important investment is a broad set of user support competence shared between academia and SMEs (small and medium enterprises). To house these massive investments, a new, state-of-the-art national data centre facility was completed by rebuilding a large hall at Linköping University.

The strategic direction of NAISS was validated early in the year through a comprehensive interim review by the Research Council. The resulting highly positive report supported our new national organisation and provided valuable recommendations for collaborations with partner universities.

These insights were directly incorporated into the proposal for renewed long-term funding (2027-2030) that NAISS submitted to the Research Council in autumn 2025, which prioritises an expanded, researcher-focused support organisation.

Restructuring leadership to focus governance

Internally, the organisation underwent some changes following my appointment as Director in January 2025. We restructured our leadership to feature dedicated module directors, ensuring focused governance across our AI, Data Science, External Relations, Technical Operations, Allocations, and User Support divisions.

We also professionalised our internal workflows through the national rollout of standardised project management and code-handling tools like OpenProject and GitLab.

Externally, NAISS leadership aggressively championed Swedish research interests on the global stage, participating in the EU AI Summit in Paris, the EuroHPC Summit Week in Krakow, and our staff was present at numerous conferences such as ISC and SC25.

We are also facing new challenges: While accelerators provide orders-of-magnitude performance and efficiency gains, they require substantial effort from researchers to adapt codes, or we might have to accept that it will no longer be possible to use traditional in-house codes in the future.

Transition to AI-driven architectures

The AI revolution will exacerbate this; in a few years it will not just be a matter of moving traditional floating-point codes to GPUs, but we will have to accept that traditional floating-point performance will go down, and future advances will instead rely on massive amounts of low-precision tensor operations.

This is by no means easy, but the computational industry has seen several of these paradigm shifts. I am confident our users will be able to adapt to this one too, and we will do our best to help.

As we close 2025, NAISS is not only delivering world-class computational power but is structurally equipped to guide the Swedish research community through the complex, inevitable transition to modern accelerator- and AI-driven computing architectures.

NAISS 2025 Highlights

- **Two new resources on the way:** The Arrhenius system is close to completion, we secured a new EuroHPC AI resource, while the Research Council, Vinnova and the government came through and provided all the national co-funding necessary. Thank you!
- **The Mimer AI Innovation Factory was launched:** This marks a new phase where NAISS increasingly provides support also to industry, which will emphasise the importance of AI and HPC in Sweden and hopefully lead to more resources overall.
- **Our users keep delivering outstanding impact:** Nothing makes us prouder than the record-breaking 1,477 scientific studies published last year with the help of NAISS resources, and the numerous PhD theses defended. This is the value Sweden gets for the investments, and what keeps our researchers internationally leading.

Erik Lindahl
Professor and NAISS Director

Deputy Director: Björn Alling

Throughout 2025, NAISS successfully navigated the challenge of operating across three critical time horizons. First, we maintained our steadfast commitment to delivering highly valued services and user support on our existing systems.

Second, we dedicated substantial effort to the preparation and deployment of our upcoming national systems, Arrhenius and the Mimer AI Factory, along with their respective support structures.

Finally, we looked to the future, actively planning our funding proposals for the 2027–2030 operational period while embracing our expanding responsibilities in a rapidly evolving computational landscape, particularly within artificial intelligence.

Specialised skills driving innovation

As our ambitions expand, we have encountered an increasingly competitive global market for expert staff in HPC, AI, and data science. We are actively recruiting to support our new initiatives and to seamlessly replace highly skilled personnel who move on to other key roles in the sector.

While this high turnover presents operational challenges, we are immensely proud of NAISS's role as a national incubator for advanced technical expertise.

By cultivating this critical mass of competence, we are effectively supplying the wider Swedish society and industry with the highly specialised skills needed to drive future innovation.

During the year, NAISS actively contributed to the national e-infrastructure investigation led by Professor Pär Weihed (LTU), which examined the future organisational structure of Sweden's digital research infrastructure.

We were very pleased to see the investigation firmly validate NAISS's current role in the landscape and underscore our strategic importance for the future of Swedish research.

Sweden has a long and fruitful tradition of Nordic collaboration within the HPC domain. Historically, the Nordic e-Infrastructure Collaboration (NeIC) has served us well by enabling joint projects, shared services, and valuable exchanges among e-infrastructure professionals.

Commitment to Nordic collaboration

However, as computational ambition levels in our respective countries dramatically increase, traditional frameworks are shifting. Following strategic changes within NeIC's previous host, NordForsk, and the separation of the Nordic Tier-1 for the CERN collaboration, the future role of NeIC has come under review.

NAISS remains strongly committed to our Nordic collaborations as a special case of the extended EuroHPC infrastructure, and together with our neighbours we are discussing how we can best take this to the next level, in particular to provide more joint training events.

Finally, the preparatory work for Arrhenius, our new joint national and EuroHPC JU system, has progressed according to plan. Introducing the first collaborative Swedish-EuroHPC system in Sweden presented unique administrative hurdles, which we successfully navigated.

Unwavering national support

This smooth transition would not have been possible without the unwavering support of key national actors, including our partner universities, the Research Council, and Kammarkollegiet (the Legal, Financial and Administrative Services Agency).

We are especially grateful to our host, Linköping University, for their vital leadership, as well as their comprehensive legal and administrative support throughout this complex process, and all our partners for having a bit of patience as we have navigated two large procurements, built a new data centre, negotiated several contracts, written a new proposal for long-term funding, and started an AI Factory in 2025!

Administration, Staff, and Organisation Development

Anna Jänis
Head of Administration

NAISS is in a state of rapid expansion. As our user base and mandates have grown, so too has the dedicated personnel required to sustain our world-class services. Functionally, the NAISS workforce is divided into three pillars: Administration and Management; Operation of Computational Resources and Storage; and User Support and Training.

Linköping University (LiU) serves as the host for NAISS, and this is also where our administration is located. Conversely, to maintain proximity to the research community, our user support and training staff are distributed nationwide across six branches at partner universities: Chalmers, KTH, Linköping, Lund, Umeå, and Uppsala University.

Relocation to modern premises

To ensure that local expertise translates into national impact, NAISS asks that branch staff dedicate a minimum of 50 percent of their employment to the national infrastructure. To support this distributed workforce, NAISS prioritises administrative efficiency.

We have streamlined the bilateral agreements with our partners so that only a single document has to be signed each year, and updates can be handled within 24h over email to reduce administrative workload.

To accommodate our expanding administration and technical teams, the NAISS headquarters relocated to new, modernised office premises in November 2025. Beyond providing workspace, these upgraded facilities feature purpose-built rooms specifically designed to host in-person classes and professionally broadcast remote training events, directly supporting our expanding educational mandate. We warmly welcome you to visit our new premises to see these capabilities first-hand and share a coffee with us!

The scale of our operations is directly reflected in our workforce metrics. Over the course of the year, approximately 90 individuals were actively engaged within NAISS, corresponding to a total of 48 Full-Time Equivalents (FTEs), with 20 FTEs allocated entirely to user support, 17 out of which were located at partner sites.

To achieve national cohesion across these distributed teams, we deployed a new national internal NAISS Staff Meeting for the first time in May 2025, and a separate NAISS User Forum in October 2025. These events proved instrumental in aligning our support teams, sharing best practices, and gathering direct, actionable feedback from our users.

Major recruitment efforts

The mandate from the EuroHPC Joint Undertaking to establish the Mimer AI Factory in Sweden triggered a major recruitment drive throughout 2025. During the summer, NAISS successfully appointed a dedicated Director for the AI Factory, as well as a Data Director.

The remainder of our recruitment efforts focused heavily on securing specialised Application Experts within the AI and Machine Learning domains. While we successfully onboarded several exceptional experts, the process underscored the intense global competition for AI competence, with highly qualified international candidates occasionally declining offers due to relocation hurdles or lucrative competing offers.

Welcoming high-level visitors

2025 saw a record-breaking number of high-level visits and meetings hosted by the infrastructure. Highlights included hosting His Majesty the King in conjunction with Linköping University's 50th-anniversary celebrations, the Minister for Higher Education and Research Lotta Edholm, and the Minister for Public Administration Erik Slottner.

We also received the Board of the Knut and Alice Wallenberg Foundation, and a delegation from the Ministry of Finance. Furthermore, as a demonstration of our integration into European strategies, we were honoured to have the European Commissioner, Henna Virkkunen, visit the AI Factory for discussions at the RISE premises in Stockholm.

Technical Operations and Infrastructure Management

Andreas Johansson
Technical Director

Gert Svensson
Deputy Technical Director

Throughout 2025, the operation of our existing national systems proceeded with high reliability. Dardel at KTH, Tetralith at Linköping University, Alvis at Chalmers, and the Swedish partition of the LUMI supercomputer continued to deliver critical AI and HPC compute resources to Swedish researchers. Furthermore, researchers working with sensitive data were continuously supported by the Bianca resource at Uppsala University.

The national storage service, Swestore, remained fully operational throughout the year; however, a strategic decision was finalised to decommission it during 2026 in favour of deploying a modernised, cost-efficient object storage service.

Major procurements for systems and storage

Similarly, the Swedish Science Cloud (SSC) maintained steady operations across its Uppsala and Umeå regions, with comprehensive plans established to transition its capabilities to the Persistent Compute and Data partition of the upcoming Arrhenius system in 2026.

To secure NAISS's operational capacity for the coming years, several major procurements were aggressively advanced during 2025. Notably, the procurement for the Mimer AI Factory entered the competitive dialogue phase with qualified vendors during the autumn.

NAISS, in its role as the hosting entity, led the technical specification and proposal evaluation, while the EuroHPC JU managed the administrative processes.

Following three intensive dialogues, specifications were finalised, and vendors were invited to submit their final offers. This system, expected to be fully operational by the end of 2026, will represent a massive expansion of Sweden's national AI capabilities.

Concurrently, at the end of the year, NAISS finalised the procurement of storage servers to establish the new national object storage service. With initial hardware deliveries scheduled for March 2026, this infrastructure will launch with a targeted 4 PB of disk capacity, with architectural preparations already in place for further expansion in the latter half of 2026.

Proactive investments in cooling

NAISS operates two cluster rooms at the Linköping data centre, with the majority of our systems housed in the Kärnhuset facility. During 2025, significant structural adaptations were completed within Kärnhuset to accommodate the Arrhenius system, which utilises approximately half of the newly expanded hosting capacity.

The remaining space has been strategically reserved to host the upcoming Mimer AI Factory, with facility adaptations scheduled to begin as soon as the final hardware specifications and contracts are signed.

Proactive investments were also made to upgrade cooling redundancy, significantly mitigating operational risks associated with district cooling fluctuations.

Emphasising our commitment to sustainable operations, the solar panels installed on the roof of Kärnhuset successfully generated 235 MWh of electricity to help power the systems throughout 2025.

Rigorous evaluation of tenders

The procurement of the EuroHPC mid-range system, Arrhenius – which was initiated prior to 2025 – reached its successful conclusion this year. Following initial expressions of interest from ten companies, eight were formally qualified.

In February 2025, four comprehensive tenders were received and rigorously evaluated by an expert group drawn from across the NAISS branches. Following vendor clarifications and oral hearings with the top two candidates, final bids were submitted in early May.

The evaluation clearly identified Hewlett Packard Enterprise (HPE) as the winning bidder. Hardware deliveries officially commenced at the end of the year, with final system approval and user availability slated for 2026.

User Support and Community Engagement

Dr. Henrik Zazzi
Acting Director of User Support

The mission of the NAISS user support is to guide and empower scientists in utilising Sweden's advanced computational infrastructure. Providing this support presents complex challenges, as our user base requires a broad spectrum of assistance, ranging from system administration to highly specialised interventions of research software engineers. To meet these demands, our support organisation is predominantly composed of PhD-level scientists with extensive HPC software experience.

During 2025, our agreements with partner universities matured, allowing us to transition from a network of local providers into a truly unified national organisation. This consolidation enabled the formation of discipline-specific expert teams across scientific fields. Each of these teams is guided by a leader and is designed to seamlessly integrate external collaborators from other national research infrastructures, resulting in a highly cohesive and efficient support ecosystem.

Translating user feedback into improvements

The User Support Advisory Committee (USAC) serves as a vital link between the research community and NAISS, ensuring that services are aligned with scientific needs. Over the past year, USAC strengthened its engagement with a series of roadshows that brought researchers and NAISS staff together. These sessions provided invaluable opportunities to address technical challenges, gather user insights, and promote deeper collaborations. The USAC members, representing institutions across the country, work in close coordination with NAISS support to translate these insights into operational improvements. We strongly encourage all users to volunteer for this important committee!

To handle the scale of user demand, our ticketing system ensures that all requests are prioritised efficiently, and researchers are guaranteed that their inquiries are tracked and addressed within predictable timeframes.

Engaging new users with targeted outreach

As detailed in the table below, our performance in resolving tickets highlights our commitment to timely and reliable support. Crucially, our resolution speeds and handling are now continuously evaluated against the rigorous Key Performance Indicators (KPIs) established by the EuroHPC Joint Undertaking.

Looking ahead, we will strengthen collaborations with regional sites and expand our targeted outreach to engage new user communities. Establishing strong partnerships with national organisations remains a cornerstone of our efforts to streamline support.

We are deploying new collaboration tools, strengthening our internal technical competencies, and advancing our documentation practices, all with the goal of better helping our users to quickly find the information they need without even asking – but you are always warmly welcome to get in touch with me or my team!

Queue	Resolved	Within 24h	Within 2d	Within 7d	Within 30d
Alvis	446	229 (51%)	270 (61%)	360 (81%)	403 (90%)
Dardel	1,149	535 (47%)	673 (59%)	910 (80%)	1,087 (95%)
Tetralith	374	144 (39%)	176 (47%)	244 (65%)	318 (85%)
Bianca	312	64 (21%)	95 (30%)	190 (61%)	270 (87%)

Ticket resolving performance per NAISS system in 2025.

Training and Skills Development

Dr. Joachim Hein
Training Manager

User training remains an indispensable pillar of the NAISS mission. Providing world-class hardware is only half the equation; ensuring that all researchers across Sweden possess the skills to effectively and efficiently harness these advanced computational infrastructures is critical to accelerating scientific discovery.

In 2025, NAISS significantly expanded its educational outreach, delivering an unprecedented 52 dedicated training events to the national user community - a substantial increase from the 37 events held in the previous year.

Consistent with our mandate to lower barriers to entry and democratise access to high-performance computing, all NAISS training programs are provided free of charge to participants.

Several new courses introduced

The dramatic increase in our user base and projects during 2025 required a proportional scaling of our training. To accommodate this escalating demand, our key introductory courses – including the comprehensive NAISS Introduction Training Week, Introduction to Linux, and system-specific onboarding for PDC, NSC, and Alvis – were offered up to four times throughout the year.

We also placed heavy emphasis on fundamental programming and environment management, running recurring sessions on Introduction to Git, Cluster Architecture and Job Submission, and Command Line 201.

To prepare the Swedish research community for our next-generation hardware in Arrhenius, we also delivered targeted training on the efficient use of the handful of Grace Hopper nodes on Dardel, to make sure that users can exploit the potential of our new accelerator architectures.

To remain responsive to the evolving technological landscape and the complex needs of the scientific community, several new courses were introduced to the 2025 curriculum.

Due to the explosive growth in artificial intelligence, we launched a dedicated Large Language Model (LLM) workshop. We also introduced training on ParaView and Trame for Scientific Visualisation on HPC resources and array computing in C++ using Eigen.

In direct response to user feedback and requests, a new module focused on debugging HPC applications was developed and deployed. Our curriculum also continued to support in-depth courses such as Introduction to parallel programming using MPI, running and building Singularity and Apptainer containers, VASP best practices, and comprehensive multi-day workshops on Python, Pandas, R, Julia, and MATLAB in HPC environments.

Prioritising sensitive data training

As medical research increasingly relies on NAISS, training users in secure handling of sensitive data has become a top priority.

In 2025, we hosted specific educational tracks for these environments, including Introductions to sensitive research data, intermediate and in-depth courses, which were highly appreciated. We co-hosted several events with the National Bioinformatics Infrastructure Sweden (NBIS), including a five-day Programming Formalisms course and a dedicated Awk workshop.

We also partnered internationally to deliver the CodeRefinery tools workshop, ensuring Swedish researchers are trained in modern, reproducible software engineering practices. Recent training events have been recorded and made available on YouTube.

Nine “Zoom-in” sessions delivered

Complementing our structured courses, we continued the popular “Zoom-in” interactive support. The training team organised nine of these open sessions in 2025.

The events provide researchers with direct access to NAISS experts, allowing them to troubleshoot technical barriers, ask complex architectural questions, and receive tailored, one-on-one guidance for optimising their computational workflows.

Finally, to expand our community leadership, NAISS staff proudly led the national Swedish effort for the collaborative “Mondays with MATLAB” initiative, successfully hosting three distinct events over the year.

Allocations and Resource Management

Associate Professor Peter Münger
Allocations Director

The National Allocations Committee (NAC) bears the overall responsibility for distributing NAISS computational and storage resources among researchers across Swedish academia. While the primary criterion for allocation is scientific excellence, demonstrated need and the efficient utilisation of resources are also critical evaluating factors. Recognising the diverse computational requirements across different scientific domains, NAISS employs a tiered allocation structure classified into large, medium, and small requests.

Swift access to entry-level resources

Large proposals, which represent significant economic value and address the scarcity of our most expensive, high-end resources, undergo rigorous formal peer review and technical evaluation. Medium-scale proposals undergo a targeted technical evaluation to ensure the requested computational resources are used efficiently. Conversely, the process for securing Small-scale allocations is designed to be highly accessible and swift, with approvals typically granted within a single working day.

For these entry-level allocations, applicants essentially need to provide only a project title and pass a simple eligibility check (though subsequent renewals do require a brief progress report and a very short new proposal).

Providing these streamlined resources to groups with limited computational needs is highly cost-efficient for Sweden, ensuring broad national impact while consuming only a small fraction of our total capacity.

A unified proposal for all resources

To manage this tiered system, the NAC relies heavily on the dedicated National Allocations Committee Working Group (NAC-WG). Decisions regarding Large-scale proposals are based on comprehensive assessments comprising external scientific reviews, an internal scientific review by a NAC member, and detailed technical reviews conducted by the NAC-WG.

Medium-scale proposals are evaluated by NAC-WG members in close collaboration with NAISS application experts, while Small-scale proposals are efficiently processed directly by the NAC-WG. I would like to take this opportunity to extend our deepest gratitude to the many scientists across the country who volunteer their time to provide external peer reviews; our national allocation system simply would not function without their invaluable contributions.

In 2025, we implemented a major simplification to the resource application process. Previously, researchers were required to submit separate applications for computational power, storage, cloud resources, and environments for processing sensitive personal data, which could get tedious when there was significant overlap.

To streamline the user experience, we consolidated this structure, allowing and encouraging researchers to apply for all required resources within a single, unified proposal. This new framework was launched during the autumn 2025 allocation round, which also marked a milestone as the first round where researchers could apply for Large-scale resources specifically dedicated to handling sensitive personal data.

Steady growth of active users

A defining metric of the infrastructure's success and the escalating demand for computational power is the continuous, uninterrupted growth of our user base over the past decade. From 3,327 users in 2016, the national infrastructure has experienced steady annual expansion. This growth trajectory accelerated in recent years, rising to a record of 8,611 active users in 2025. NAISS provisions resources to this expanding user base across a vast spectrum of scientific fields and Swedish academic institutions.

While traditional, compute-heavy research fields at major universities remain prominent users, a defining trend of 2025 was the explosive growth in AI-driven research, alongside a rapidly increasing demand for large-scale storage, particularly within life sciences.

Overall, the demand for NAISS resources continues to significantly outpace available capacity. To maximise our scientific impact and support as many researchers as possible, NAC generally addresses this oversubscription by strategically reducing the size of granted allocations rather than outright rejecting scientifically sound proposals.

Statistics

Requested/allocated compute resources

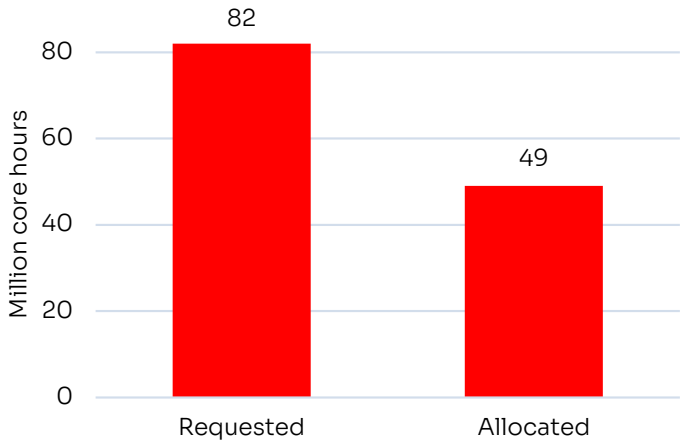

Requested and allocated compute resources in total on all NAISS resources, normalised to the Dardel CPU performance per core during the calendar year 2025.

Approved/denied compute proposals

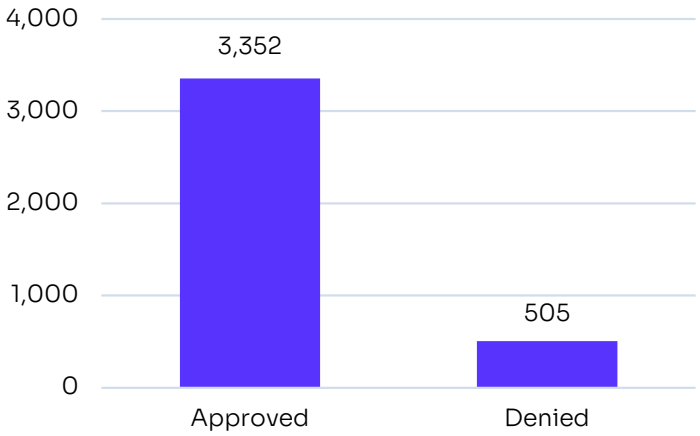

Number of proposals requesting NAISS resources during some part of 2025.

Active users

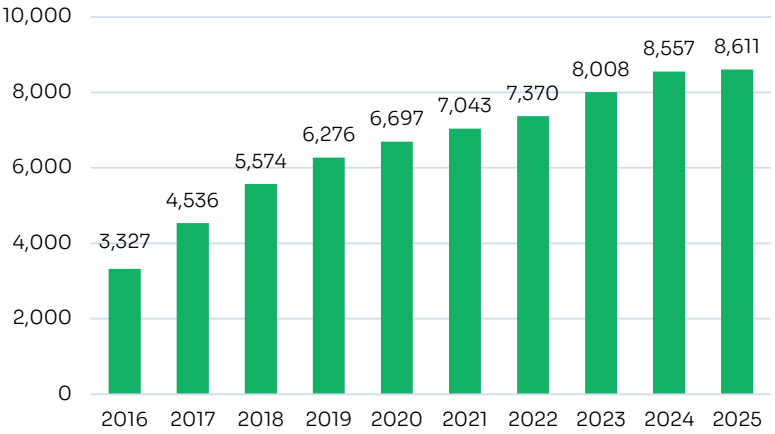

Number of users associated with any active project during each year. Over the last decade, the NAISS user base almost tripled; from 3,327 to a record 8,611 users in 2025.

Compute resources per field

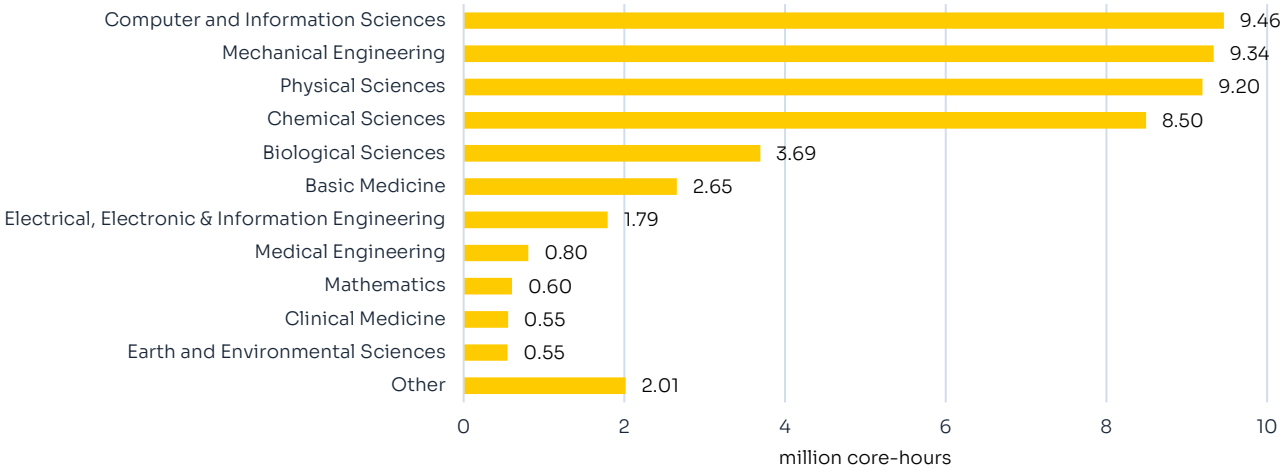

Amount of NAISS compute resources used per field, based on the classification codes selected at proposal submission. Many of the applied AI projects across fields are classified as Computer and Information Sciences, which is the reason for its new dominance.

Storage resources per field

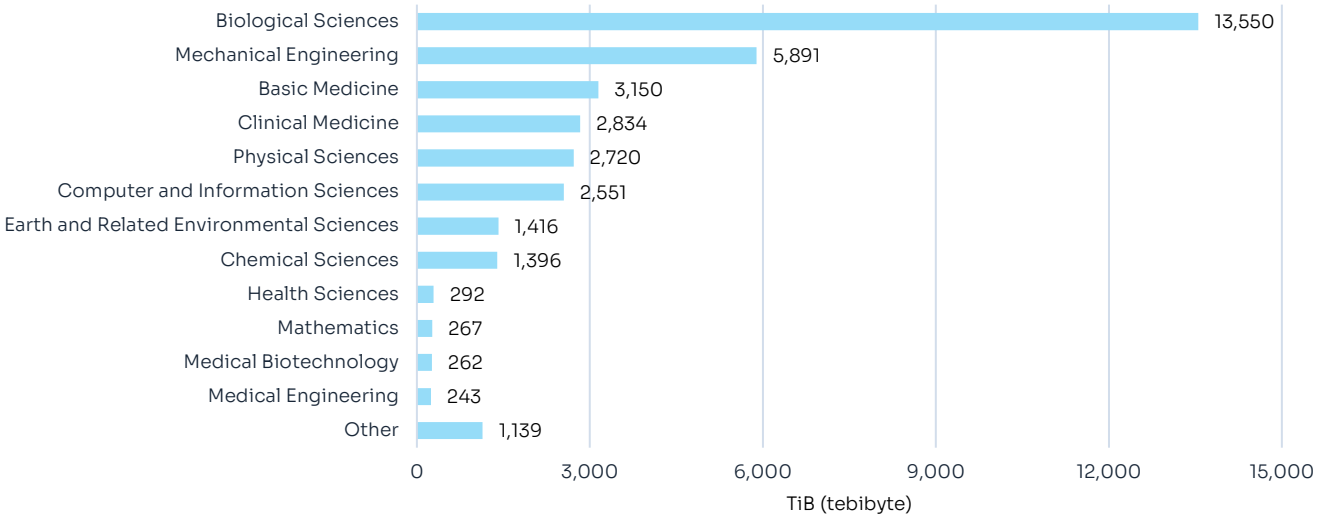

Total amount of NAISS storage resources used per field, based on the classification codes selected at proposal submission. The dominance of Biological Sciences is due to sequencing data in general, and sensitive data in particular.

Mimer: Establishing Sweden's AI Innovation Factory

Dr. Rossen Apostolov
AI Factory Director

During its inaugural year, the Mimer AI Factory successfully established the core organisational, technical, and outreach foundations required to serve as Sweden's premier national platform for AI support and infrastructure services.

Moving rapidly from initial setup to active operation, we successfully filled approximately 90 percent of our targeted technical, scientific, and operational roles.

This robust recruitment enabled the formation of dedicated expert teams covering AI technologies, scientific domains, service delivery, training, trustworthy AI, outreach, and strategic collaborations. With internal governance and user interaction processes now firmly established, we are strongly positioned to efficiently drive hundreds of projects annually.

Massive demand for practical AI support

National accessibility and visibility expanded rapidly throughout 2025. Following the launch of the comprehensive Mimer web platform as a central entry point, we cultivated a strong digital presence, reaching 300 dedicated newsletter subscribers and over 1,000 followers on LinkedIn.

This visibility translated directly into operational impact: our service pipeline is now fully active, successfully managing 196 active or completed projects for researchers, industry, and public-sector stakeholders.

A major outreach milestone was our prominent participation at Stockholm Tech Arena; one of Sweden's largest technology and innovation events. Through a dedicated booth, we achieved strong exposure toward industrial actors, startups, investors, and public-sector stakeholders, generating over 150 strategic follow-up contacts and confirming a massive national demand for practical AI adoption support.

To lower the barriers to AI adoption, capacity building and skills development remain core components of our service portfolio. During the year, we developed and deployed over 20 targeted training activities – including specialised modules, workshops, and webinars – designed to strengthen national competence across disciplinary and sectoral boundaries.

Procurement of dedicated system finalised

Furthermore, to ensure strong national anchoring, we continuously rotate our plenary meetings across different Swedish locations, engaging directly with local research communities.

In parallel, intensive discussions are ongoing with science parks, research organisations, innovation support actors, and Swedish resource providers to develop collaboration models that strengthen the broader national AI landscape. Contact us if you want to become part of this extended network!

On the infrastructure front, the finalised procurement of the dedicated Mimer AI supercomputer represents a crucial strategic milestone, laying the technical foundation for massive future service expansion and enabling the long-term integration of Swedish and European AI infrastructures. We look forward to sharing details about it as soon as we can.

Focus on scaling and long-term impact

In the meantime, our team actively facilitated academic and industrial access to existing EuroHPC resources, successfully securing 300,000 GPU hours for Swedish projects to date. Additionally, our data team is actively prototyping Trusted Execution Environments and sensitive data sharing to drive the establishment of a future European Health and Life Science data lab.

Ultimately, our first operational year demonstrates that Mimer has swiftly progressed from its establishment phase into active national relevance. With core structures in place, demand increasing, and technical readiness advancing, our next phase will focus entirely on scaling service delivery, broadening national access, and translating this immense initial interest into measurable, long-term impact for Sweden.

Financial Statement

	Result 2023	Result 2024	Result 2025	Budget 2026
INCOME				
Research Council base funding 2023-2026	115,000,000	115,000,000	115,000,000	115,000,000
Partner contributions to user support	34,000,000	35,500,000	37,000,000	37,000,000
LiU administrative support to NAISS	1,000,000	1,000,000	1,000,000	1,000,000
RC dedicated NeIC funding 2024-2026		3,589,000	3,589,000	3,589,000
Research Council Arrhenius co-funding	0		48,870,000	50,000,000
EuroHPC Arrhenius operations (their 35%)			24,550,969	16,050,000
Research Council AI Factory co-funding				180,000,000
EuroHPC AIF HW operations (their 50%)				10,649,000
EuroHPC AIF Service Hub			43,640,927	
EU ENCCS Competence centre		4,415,000	4,255,095	
EU Virtual training academy			1,416,249	
Sum	150,000,000	159,504,000	279,322,240	413,288,000
EXPENSES				
Management	4,795,938	6,414,943	6,364,323	8,600,000
New system operations (Bianca), staff		829,487		
Communication	395,871	1,055,845	1,109,026	1,200,000
Old systems operations	71,472,553	70,180,933	64,298,621	44,000,000
Allocation SUPR/SAMS	2,141,813	2,277,483	2,709,011	2,700,000
User support, including management	22,395,981	26,790,327	29,940,471	33,700,000
Training; personnel in user support	1,116,505	663,415	427,000	4,100,000
LiU center storage		4,616,324	247,973	
Bianca, UPPMAX			10,000,000	
Arrhenius startup costs		766,234	6,424,616	17,000,000
Arrhenius investment				254,780,000
Arrhenius operations				25,000,000
Arrhenius insurance				4,000,000
AIF hardware investment				159,216,000
AI Factory hardware startup				2,000,000
AI Factory hardware insurance				1,100,000
Object storage investment				4,000,000
Object storage operations				2,500,000
Other hardware				2,500,000
National data centre investment costs			17,868,057	
AI Factory Service Hub personnel			1,398,681	39,800,000
ENCCS		4,490,000	3,527,435	
EU Evita			462,682	
NeIC		3,603,462	4,208,923	3,574,538
NAISS other items				5,000,000
NAISS other operations	127,699		2,014,544	
Other costs	86,108	674,425	522,629	
Sum	102,532,468	122,362,878	151,523,992	614,770,538
Cost increases (salaries and operations)				8,257,301
Result	47,467,532	37,141,122	127,798,248	-209,739,839
Prognosis for full funding period				2,667,062

User Story: Tommy Löfstedt

Tommy Löfstedt, an associate professor and docent in computing science at Umeå University, is leading a group of eight people doing research in machine learning and computer vision.

The application areas are wide-ranging but focusing mainly on projects related to medical imaging, life sciences, and software security.

One current project funded by the Swedish Childhood Cancer Fund is about improving automatic segmenting of medical images for children. There are already segmentation models that perform well on adults, but much less so on children.

"There is a scarcity of quality medical images of children, partly and luckily because a low number of children have cancer. Children's images are different in their noise distribution and anatomy, but children also may not lie still during the scanning, which produces motion artefacts. And so the resulting images are more difficult to segment," Tommy Löfstedt explains.

"We are making progress"

His group is trying to address the problem by developing AI models which handle these differences.

"The goal of the project is to improve the segmenting results for children without sacrificing the quality for adults. And we can see that we are making progress."

Another project he is involved in which is launching soon is about analysing Android software to find harmful code, with funding from WASP (Wallenberg AI, Autonomous Systems and Software Program).

Tommy Löfstedt has a Medium allocation on the GPU-based NAISS system Alvis, but admits they have occasionally been hitting the ceiling.

"It hasn't been a huge problem so far – we have been able to keep running anyway. But I've been thinking that maybe it's time for us to apply for a Large allocation."

The process his team is using involves training the same model over and over, with different settings, to determine which settings work best.

"We may end up training 10,000 models or more. I would say that these models today are medium-sized. They for sure don't require hundreds of GPUs. Generally, one node per model is sufficient," Tommy Löfstedt says.

Tricky to estimate future needs

He and his team rarely encounter problems – they have a set-up that works and usually manage the processing all on their own.

Should they ever need a special version of a particular software, NAISS's support is quick to assist. But he wishes that it would be easier to scale usage on existing allocations.

"You apply for Medium and Large resources for one year at a time, and it is sort of a guessing game. You never know how many of the project grants you apply for will be approved, and so it is pretty hard to know in advance what kind of resources you will be needing in six months' time."

"If there was a way to adjust it continually, it would be great."

User Story: Rocío Mercado Oropeza

Rocío Mercado Oropeza, originally from California, came to Gothenburg in 2018 for a postdoc. Today she is an assistant professor at the Data Science and AI Division at Chalmers University of Technology where she is heading AIME, the AI Laboratory for Molecular Engineering.

Her team is building generative AI tools that can be used to design molecules with specific properties and speed up the process of finding promising candidates for new materials or drugs. The applications are wide-ranging.

"We have 10^{33} possible drug-like compounds that could exist out there. How do we find the best ones for different kinds of diseases and so on, the best treatments?"

"And that is where our tools come in. They are helping us design molecules that meet all of our desired constraints, or as many as possible, while also searching through this chemical space," Rocío Mercado explains.

All tools offered as open source

All AIME tools are open source and shared on GitHub, free for anyone to use.

"That is very important to me and the collaborations I choose. Many of the use cases we do see of our tools are other academic groups. If people use them for business purposes then that's great, but I'm not making any additional money off of this."

Training the models which her tools are based on requires large compute resources, both CPU and GPU, across different stages of the development process. Her team has allocations on several NAISS systems.

With deep molecular generative models, the molecules are treated as text, and there is an architecture with different weights. Training them efficiently requires at least GPUs, as they usually run into millions of data points.

"Sometimes we do a bit of molecular simulation to generate training data. For that we use Dardel which is dedicated for this, as opposed to Alvis which we use more for training of our machine-learning models."

"Lots of opportunities"

Since Dardel is quite crowded and is soon also being retired, AIME has increasingly started using the Wallenberg-funded Berzelius system. Rocío Mercado says it is a lot faster, but they have yet to figure out how to run all of their simulations on it.

AIME receives funding from many different sources, among them WASP (Wallenberg Autonomous Systems Program), WISE (Wallenberg Initiative Materials for Sustainability), and the Research Council. The group is also collaborating with Intel and Merck who are funding a PhD student and a postdoc on her team.

And in September 2025, Rocío Mercado was awarded the European Research Council's prestigious EUR1.5 million Starting Grant, for a project about generative models for polymers.

"There are a lot of opportunities, and I was lucky to get many of them and recruit people to work on these various topics early on. Now I feel quite satisfied with the size of the team."

User Story: Mats Granath

The buzz around quantum computing is getting increasingly louder. But will such computers ever become commercially viable, and can they be built large enough to set off the next big revolution?

Things are definitely brewing also in Sweden, with Gothenburg being one of the main hubs. At University of Gothenburg, Physics Professor Mats Granath is developing quantum error correction algorithms in a project which is part of the Wallenberg Centre for Quantum Technology (WACQT), funded by the Knut and Alice Wallenberg Foundation (KAW).

Using quantum data to build realistic models

Famously, a quantum computer bit, or qubit, is represented by a superposition and can feature multiple possible states simultaneously, not only 0 or 1. This generates errors, so to become useful a quantum computer needs powerful error correction – something that is not easily done.

"A quantum computer is exceptionally sensitive to environmental noise. It is impossible to directly measure the information without affecting it. So you perform measurements by sort of probing around, but not disturbing, the logical information," Mats Granath explains.

Decoding these measurements with a classical algorithm can reveal what type of errors that are the most probable, and those decoders are the ones that he is researching and developing.

Running a quantum computer is costly, and it is difficult to generate sufficient amounts of data. The solution is to use actual quantum computer data to build a realistic model.

"We are training a neural network, and we have a simulation model running in the background which is also building experimental data. The elegant thing about our method is that we are sort of generating data while we are training."

Mats Granath uses the NAISS Alvis system to train the models, which is sufficient for the time being. He tried using Berzelius (the KAW-funded AI-optimised supercomputer) too, but the algorithms did not run as smoothly on that system.

"We are not yet at the stage where we can do real-time decoding. This is being developed, but it will require a dedicated hardware with more hard-coded instructions."

Mats Granath has been into quantum computing since 2019, but it is only recently that he has been able to use real quantum hardware. He has access to IBM quantum computers and is hoping to soon get his hands on the 25-bit quantum computer that Chalmers is constructing.

Excellence cluster for quantum computing

The WACQT funding runs until 2030, but there are plans to establish a national excellence cluster in quantum technology. Mats Granath has contributed to an application to the Swedish Research Council and Vinnova which opened in April.

"This will be hugely important to advance quantum research in Sweden. We will likely focus on specific areas, like developing and optimising parts of the hardware so that it can be used in larger systems," he says.

So when will we have a stable, fully functional, large quantum system?

"Something might be available ten years from now. But that is a very optimistic estimate."

Appendix – NAISS Committee Members

The NAISS Steering Committee

Jan-Eric Sundgren, Chair
Senior Advisor

Frida Bender
Associate Professor,
Stockholm University

Lars Kloo
Professor,
KTH Royal School of Engineering

Bengt Persson
Professor,
Uppsala University

Katrine Riklund
Professor and Dean,
Umeå University

Julia Wiktor
Associate Professor,
Chalmers University of Technology

Anders Ynnerman
Professor,
Linköping University

National Allocations Committee – Spring 2025

- **Chair:** Paul Erhart, Professor, Chalmers University of Technology
- Xue-Song Bai, Professor, Lund University
- Leif Eriksson, Professor, University of Gothenburg
- Evan Patrick O'Connor, Associate Professor, Stockholm University
- Outi Tammisola, Professor, KTH Royal Institute of Technology
- Olav Vahtras, Professor, KTH Royal Institute of Technology

National Allocations Committee – Autumn 2025

- **Chair:** Paul Erhart, Professor, Chalmers University of Technology
- Shervin Bagheri, Professor, KTH Royal Institute of Technology
- Xue-Song Bai, Professor, Lund University
- Leif Eriksson, Professor, University of Gothenburg
- Evan Patrick O'Connor, Associate Professor, Stockholm University
- Rocío Mercado Oropeza, Assistant Professor, Chalmers University of Technology
- Olav Vahtras, Professor, KTH Royal Institute of Technology

User Support Advisory Committee (USAC)

- **Chair:** Marie Skepö, Professor, Lund University
- Mikael Carp, System Manager, Karolinska Institutet
- Nicolas Delhomme, Senior Researcher, Umeå University
- Nandan Haloi, Postdoctoral Researcher, KTH Royal Institute of Technology
- Michael Holmboe, Associate Professor, Umeå University

Appendix – NAISS training events in 2025

User-facing NAISS training events

Date(s)	Title	Duration	Registered
15 January	Introduction to Alvis	2 h	23 registered; 21 participated
16-17 January	Awk workshop Joint event with NBIS	2 days	27 registered; 17 participated
21 January	Introduction to Linux	3 h	31 registered; 13 participated
4 February	Cluster architecture and job submission	2 to 3 h	50 registered; 29 participated
13 February	Basic Singularity: Running and building Singularity containers	1 day	58 registered; 35 participated
26 February	Using ParaView and Trame for Scientific Visualisation on HPC resources	4 hours	36 registered, 39 participated
4 March	Introduction to Python	1 day	67 registered; 16 participated
5 March	Array computing in C++ using Eigen	6 h	15 registered; 10 participated
5 March	Introduction to Alvis	2 h	51 registered; 47 participated
20 March	Introduction to PDC Systems	1 day	40 registered; 32 participated
7 March	Transferring files to/from HPC clusters	3 h	37 registered; 9 participated
19 March	Introduction to Bianca: Handling sensitive research data	1 day	24 registered; 11 participated
24-26 March	Introduction to running R, Julia, and Matlab in HPC	3 days	45 registered; 16 participated
3 April	Introduction to Linux	3 hours	20 registered; 10 participated
8-9 April	Matplotlib of Publications	2 days	25 registered; 21 participated day 1 and 18 participated day 2
10 April	Introduction to Alvis	2 hours	40 registered; 34 participated
11 April	Introduction workshop to Alvis	2 hours	9 registered; 5 participate
15 April	NSC Introduction to Tetralith/Sigma	2 hours	9 participated
24-25 & 28-29 April	Introduction to Python and Using Python in an HPC environment	4 days	65 registered; 20, 21, 20 resp. 17 participated
5 – 9 May	Programming Formalisms	5 days	13 participated
16 May	Transferring files to/from HPC Clusters	3 hours	15 registered, 4 participated
20 May	Selecting Modules	2 hours	14 registered; 7 participated
28 May	Introduction to Alvis	2 hours	20 registered; 16 participated
23 May	Bianca in-Depth: Improve your handling of sensitive research data	1 day	2 participated
27 May	Cluster architecture and job submission	2.5 hours	40 registered; 18 participated
2-3 June	Intermediate Bash and Linux	2 x 3 hours	40 registered; 15 participated
10-11 June	Introduction to Pandas	2 x 4 hours	43 registered;

Date(s)	Title	Duration	Registered
13 August	Introduction to Alvis	2 hours	27 registered; 14 participated
28-29 August	Awk workshop Joint event with NBIS	2 days	8 registered; 4 participated
2 Sept	Cluster architecture and job submission	2.5 hours	43 registered; 28 participated
5 Sept	Log in and Transfer Files to/from HPC Clusters	3 hours	29 registered; 15 participated
10 Sept	Selecting software modules	2 hours	16 registered, 6 participated Event suffered from Zoom outage
15 Sept	Introduction to Bianca: Handling Sensitive Research Data	1 day	23 registered, 12 participated
16, 17, 23, 24 Sept	Introduction to parallel programming using Message Passing	4 mornings	14 registered; 8, 6, 4 resp. 4 participated
22 Sept	Introduction to Linux	3 hours	31 registered; 7 participated
1 Oct	Introduction to Alvis	2 hours	25 registered; 18 participated
3 Oct	VASP Best Practices	3 hours	11 participated
6-7 Oct	Introduction to PDC	2 mornings	23 registered; 17 participated
6-8 + 10 Oct	Introduction to running R, Julia and Matlab	4 days	14 registered, 6 participants for R, 2 each for Julia and Matlab
7 Oct	Debugging of HPC applications	2 hours	13 registrations, 7 participants
13 Oct	Running and building Singularity/ Apptainer containers	1 day	31 registered, 16 participated
14 Oct	Intro to Python for HPC	1 day	30 registered, 12 participated
20 Oct	Linux Command line 101	3 hours	21 registered; 13 participated
23-29 Oct	Programming Formalisms Joint event with NBIS	5 days	6 participated
3-7 Nov	Introduction to Git	5 days	29 registered; 16 participated
14 Nov	Log in and Transfer Files to/from HPC Clusters	3 hours	22 registered, 6 participated
18 Nov	Bianca In-Depth: Improve your handling of sensitive research data	1 day	2 participated
19-21 Nov	Large Language Model Workshop	3 days	80 registered; 56 participated
25 Nov	Running jobs on clusters		34 registered; 18 participated
27-28 Nov 1-2 Dec	Introduction to Python and Using Python in an HPC environment	2 + 2 days	37 registered; 13 participated
3 Dec	Introduction to Alvis	2 hours	37 registered; 26 participated
4-5 Dec	Linux command line 201	2 x 3 h	74 registered; 32 attended

Mondays with Matlab

Date	Topic	Participants
12 May	Parallel Computing with MATLAB	25 from Sweden
9 Sept	MATLAB usage in HPC Computing Environment	15 from Sweden
23 Sept	Building Scalable AI Workflows with MATLAB	11 from Sweden
Total	Attending at least one workshop	39 from Sweden

Zoom-in virtual user support roadshows

Date	Time	Event
30 January	14:00	Zoom-in
13 March	14:00	Zoom-in
10 April	14:00	Zoom-in
8 May	14:00	Zoom-in
12 June	14:00	Zoom-in
11 Sept	14:00	Zoom-in
9 Oct	14:00	Zoom-in
13 Nov	14:00	Zoom-in
11 Dec	14:00	Zoom-in

CodeRefinery Joint Nordic Events

Dates	Title	Duration	Participation
25-27 March and 1-3 April	CodeRefinery workshop (spring)	6 x 3.5 h	25 registrations from Sweden, 378 in total
9-11 Sept + 6 following Wednesdays	CodeRefinery workshop (autumn)	3 x 3.5 h 6 x 2.5 h	66 registrations from Sweden, 400 in total
25-27 Nov	Python for Scientific Computing	3 x 3 hours	494 registrations in total, statistics by country not available

NAISS
Linköping University
581 83 LINKÖPING
SWEDEN

naiss.se
linkedin.com/company/naiss-se