VLAAMS SUPERCOMPUTER CENTRUM

COMPUTING POWER FOR WHAT MATTERS.

ANNUAL REPORT 2025

Supported by fwo

TABLE OF CONTENTS

It indicates a clickable link to an external webpage or online resource.

FOREWORD

2025 was a year of continued strengthening and expansion for the Flemish Supercomputer Center (VSC). In a research landscape that is becoming increasingly data- and compute-intensive, VSC continues to play a key role as a catalyst for excellent research and innovation in Flanders.

A major milestone was the inauguration of sofia, the new Tier-1 supercomputer within the VSC ecosystem. With sofia, Flanders' computing infrastructure has been strengthened for the future, both in terms of performance and energy efficiency. This investment reaffirms our ambition to continue providing researchers from academia, government, and industry with access to high-performance, reliable, and future-oriented computing resources. At the same time, sofia represents a natural next step within the VSC's tiered model, in which different systems complement and reinforce one another, from the local level to the national and international level.

VSC supports a broad user community from academia, government, and industry, combining computing power with expertise, training, and services. In addition to academic research, industrial use of VSC services continued to grow in 2025. An increasing number of companies and public organizations are making use of VSC services for research and development activities, often as a stepping stone toward larger-scale or more complex computational workloads. Through targeted support and accessible services, VSC contributes to innovation, knowledge valorization, and the strengthening of the Flemish research and innovation ecosystem.

At the same time, VSC remains an important gateway to European HPC infrastructure. Through active support for Flemish researchers in accessing and using EuroHPC systems, including LUMI, VSC helps ensure participation at the highest European level of high-performance computing. This connection between Flemish and European infrastructure is essential for maintaining competitiveness in a rapidly evolving international landscape.

This annual report provides an overview of the key developments, activities, and achievements of VSC in 2025. Thanks to the continued commitment of all partners, VSC is well positioned to maintain its central role in supporting research and innovation in Flanders in the years ahead..

> "VSC continues to play a key role as a catalyst for excellent research and innovation in Flanders."

VSC TEAM

KU Leuven
Leen Van Rentergem
Jan Ooghe
Ingrid Barcena

UAntwerp
Stefan Becuwe
Wim Vanroose

UGent
Johan Van Camp
Ewald Pauwels

VUB
Stefan Weckx
Ward Poelmans

UHasselt
Geert Jan Bex
Rafal Al-Tekreet

FWO
Caroline Volckaert

VSC AT A GLANCE

INTRODUCTION

The support and operation of High-Performance Computing (HPC) in Flanders are organized through the Flemish Supercomputer Center (VSC). VSC is a consortium in which the five Flemish university associations join forces to provide HPC infrastructure and support to the broader research community in Flanders, including universities, research institutions, industry, and government organizations.

As part of its service mission, VSC also offers a wide range of training programs designed to promote and facilitate the effective use of its infrastructure. The infrastructure is hosted across four hubs: the University of Antwerp, Vrije Universiteit Brussel, Ghent University, and KU Leuven. VSC is managed by the Research Foundation – Flanders (FWO).

Within the European HPC model, three levels are distinguished: the computing capacity available within research institutions (Tier-2), computing capacity whose scale and cost exceed the capabilities of a single institution and are therefore provided at the regional or national level (Tier-1), and the largest supercomputing infrastructures (Tier-0). VSC focuses on the Tier-2 and Tier-1 layers while facilitating access to European Tier-0 resources.

VSC serves as a key computing and data infrastructure hub for researchers and industry in Flanders and has demonstrated its value as a facilitator of research and innovation.

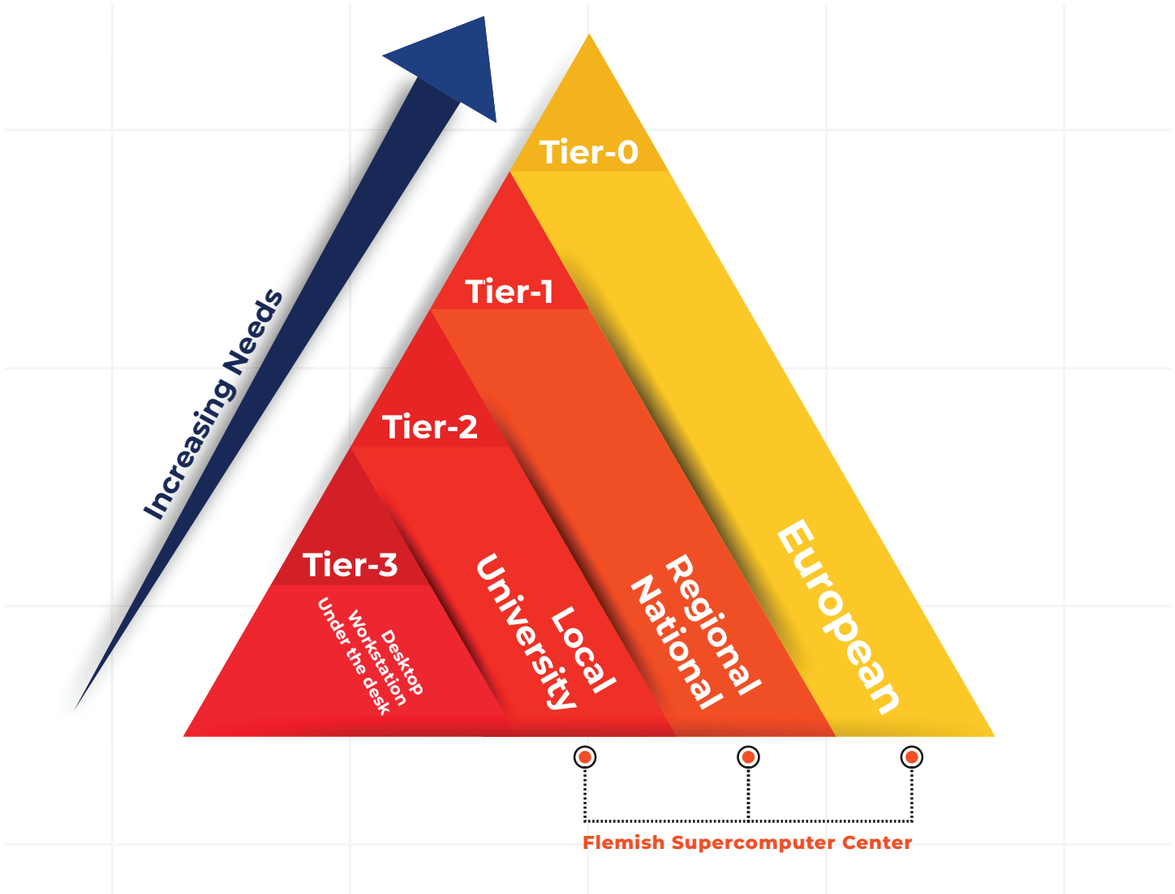

Figure 1. Tiering Model

VSC 2.0 STRUCTURE AND OPERATIONS

A strategic plan and funding framework were developed to ensure the future-oriented and sustainable growth of Tier-1 computing capacity. The objective is to continue supporting users' growing computational needs while keeping Flanders competitive in a rapidly evolving scientific and technological landscape driven by big data.

At the same time, under the impetus of the European Commission and the EuroHPC Joint Undertaking, substantial investments are being made in very large-scale computing infrastructure, particularly through European pre-exascale and exascale Tier-0 systems, as well as numerous initiatives related to Artificial Intelligence (AI). Significant attention is also being devoted to the development of European hardware technologies, driven in part by strategic geopolitical considerations.

In response to these developments, VSC is evolving into a renewed role within Flanders. While VSC originally had a primarily regional character and historically focused on users within the Flemish university associations, its user base is now being actively expanded across Flanders. This broader user community is encouraged and supported in taking steps toward participation in a European ecosystem of pre-exascale and exascale computing systems and regional competence centers, within which VSC aims to be a fully recognized partner.

The key principles of the VSC 2.0 strategic plan and its operational implementation can be summarized as follows:

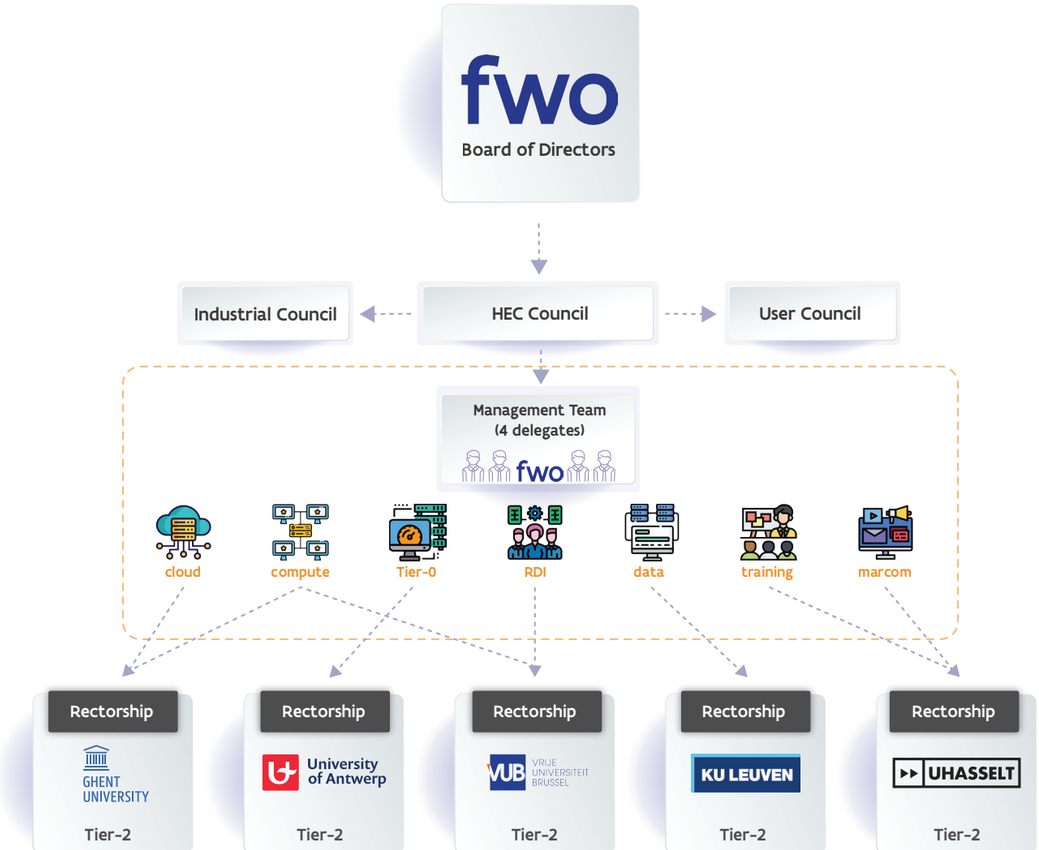

Figure 2. VSC 2.0 Structure and Operations

TIER-1 SUPERCOMPUTER PLATFORM

TIER-1 COMPUTE

TIER-1 HORTENSE

Throughout 2025, several infrastructure upgrades, maintenance activities, and service improvements were carried out on Hortense, the VSC Tier-1 Compute platform.

Infrastructure

POSTPONEMENT OF HORTENSE PHASE 1 DECOMMISSIONING

Tier-1 Hortense consists of two phases installed two years apart. Originally, the oldest part of the system, Phase 1, was scheduled to be fully decommissioned on November 1, 2025. This phase includes the oldest hardware in the Hortense installation: 384 CPU nodes based on the AMD Rome architecture and 20 GPU nodes. Due to the continued high demand for Tier-1 resources and the fact that Tier-1d (sofia) was not yet operational during 2025, an assessment was conducted throughout the year to determine whether Phase 1 could remain in service for a longer period. As a result, the decommissioning was postponed until mid-2026. The necessary support contract extensions were purchased at the end of 2025 to facilitate this decision.

EXPANSION OF NETWORK CAPACITY

The internal network infrastructure of the Hortense Tier-1 Compute cluster was further optimized. In February 2025, a major network upgrade was implemented within Ghent University's S10 data center to improve performance and resilience.

COOLING SYSTEM RETROFIT

During maintenance activities at the beginning of 2025, a significant amount of hardware was found to be unavailable due to corrosion and leaks in the direct liquid-cooling system serving the processing units. The issue was traced to an unexpected design flaw and had a substantial impact on system availability. At one point, nearly half of all GPU nodes were unavailable.

In May, VSC worked together with ATOS/Eviden and APAC to perform a complete retrofit of the direct-cooling infrastructure during an extended maintenance window. All direct cooling connections to the processing units were replaced. The operation was completed successfully, restoring the computing infrastructure to full operational capacity.

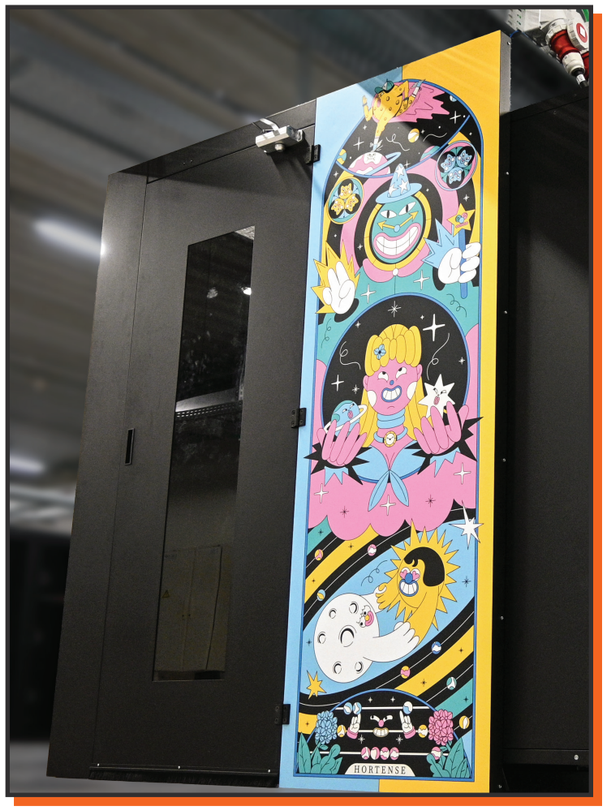

Figure 3. Tier-1 Hortense

User features en backend ontwikkeling

OPERATING SYSTEM UPGRADE

The operating system of the Hortense Tier-1 Compute cluster was upgraded to a newer version of Red Hat Enterprise Linux. A phased aproach was required due to the strategic decision to extend the lifetime of Phase 1. In May, the CPU nodes of Phase 2 were upgraded and two new login nodes running the updated operating system were deployed as the default access points. By the end of 2025, the upgrade had also been completed for the CPU nodes of Phase 1, which will remain operational until mid-2026. The GPU nodes in both phases are scheduled for upgrade in early 2026. Testing revealed significant issues related to PCI bus management on the GPU nodes, requiring additional investigation before deployment.

TIER-1 COMPUTE PROJECT PORTAL

Users of Tier-1 Compute Hortense can monitor their usage, available resources, and project duration through the Tier-1 Compute Project Portal (https://resapp.hpc.ugent.be). As VUB will use the Waldur platform for the new Tier-1 system, sofia, a major redevelopment cycle for the existing portal was considered unnecessary. Development efforts therefore focused on maintaining stability and resolving software issues.

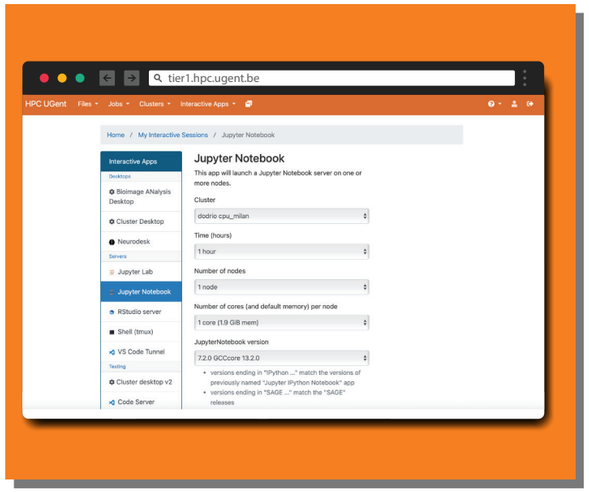

Figure 4. Screenshot of Hortense tier1.hpc.ugent.be

VSC ACCOUNT PAGE MAINTENANCE

All account and group management within the VSC is centralised on the VSC account page (https://account.vscentrum.be) and admin/API page (https://apivsc.ugent.be). This account page controls access to all Tier-1 and Tier-2 infrastructure of the VSC. The front- and backend of this service are maintained by UGent.

During 2025, this application was migrated to a Nomad platform within UGent, where maintenance is easier. The necessary implementations were also prepared and carried out so that UGent could address a new VSC range. At the end of September, the last user vsc49999 in the vsc4 range of UGent was assigned. Going forward, new UGent users are being assigned in the vsc5 range. This transition went smoothly thanks to thorough preparation.

In addition, several migrations of non-university research institutions were facilitated. In 2024, it was decided to treat these research institutions (VLIZ, INBO, ILVO, RBINS, RMCA) preferentially at the Tier-2 level, granting their researchers access to the Tier-2 infrastructure of KU Leuven and UAntwerp. The UGent team developed several automations to support this.

Scientific Projects

The table below provides an overview of the projects active on Tier-1 Hortense during 2025, categorized by call and project type.

Category	#projects	Allocated compute time	
		CPU Hours	GPU Hours
Ongoing projects from 2024	187	436,548,427	785,040
Cutoff 1 (3 February)	37	166,067,308	142,145
Cutoff 2 (2 June)	47	165,441,392	466,798
Cutoff 3 (6 October)	33	191,523,465	226,897
Starting grants	206	89,627,755	160,339
Collaborative grant (sub)projects	13	174,197,876	58,596
Industry projects	76	323,524,110	348,714

Table 1. VSC Tier-1 Compute Hortense – Overview of Scientific Projects

Usage and Availability

The figures below present a day-by-day overview of compute usage on Tier-1 Hortense. Both charts show the fraction of the maximum theoretical compute capacity utilized throughout the year.

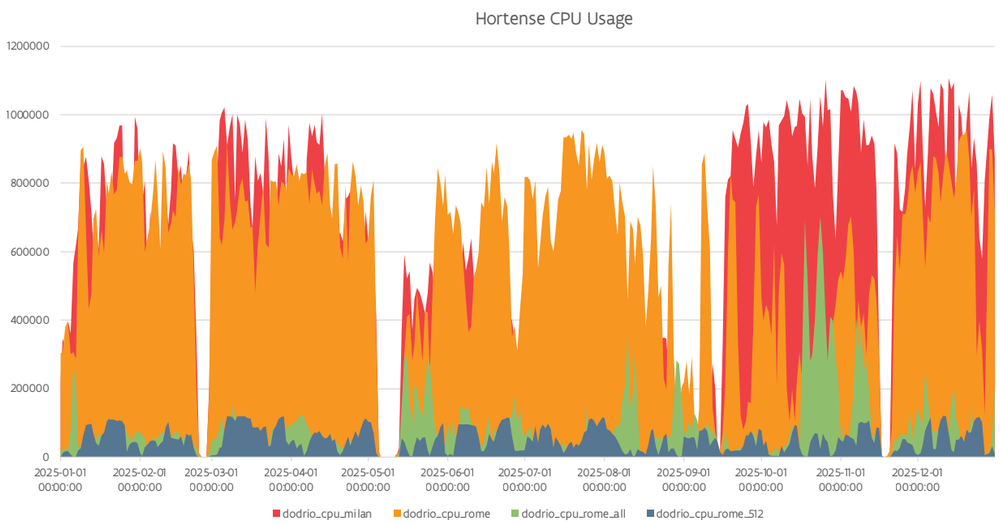

Figure 5. VSC Tier-1 Compute Hortense CPU Usage

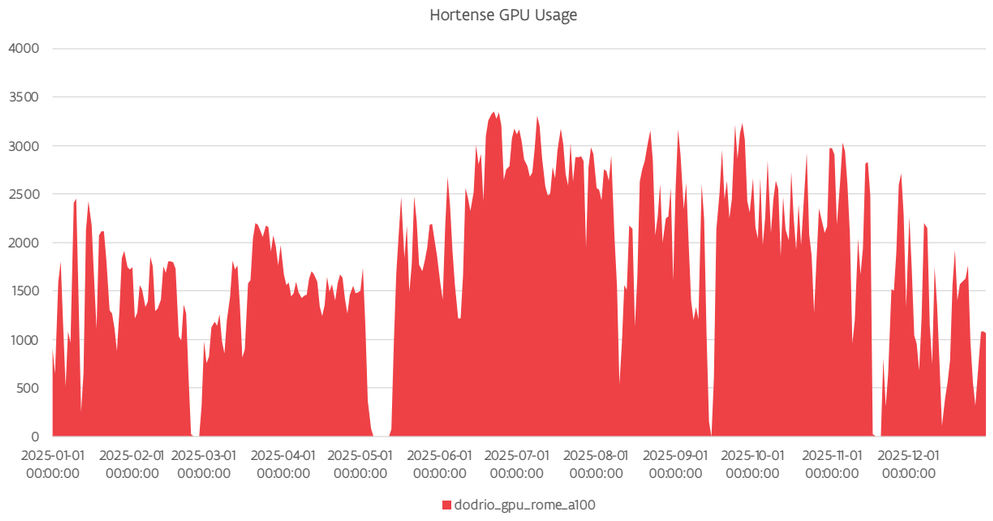

Figure 6. VSC Tier-1 Compute Hortense CPU Usage

During 2025, several scheduled maintenance activities temporarily affected system availability:

- February 24–28: planned maintenance to expand network capacity.
- May 5–14: extended maintenance for the cooling system retrofit.
- July 11: Disaster Recovery Procedure (DRP) test for the UGent S10 data center, resulting in temporary loss of network connectivity.
- September 12: urgent operating system update.
- November 17–21: scheduled maintenance including firmware, driver, operating system, InfiniBand, and Lustre file system updates, as well as migration of Phase 1 CPU nodes to a newer Red Hat Linux version.

Helpdesk

A centralized helpdesk supports the Tier-1 Compute platform. Support requests are handled collaboratively by multiple VSC institutions.

Queue	# Closed tickets 2025	Handled by
Tier-1 Compute (compute@vscentrum.be)	807	UGent (667) VUB (131) KU Leuven (6) UAntwerp (3)

Table 2. VSC Tier-1 Compute Hortense Helpdesk Activity

Through the dedicated Tier-1 Compute helpdesk (compute@vscentrum.be), users can submit problems or questions directly related to Tier-1 Compute Hortense. In 2025, a total of 807 cases were handled, a 21% increase compared to 2024. A breakdown of the time within which each submitted question was resolved is shown in the graph below. In 2025, over 90% of all incoming queries were resolved within one week.

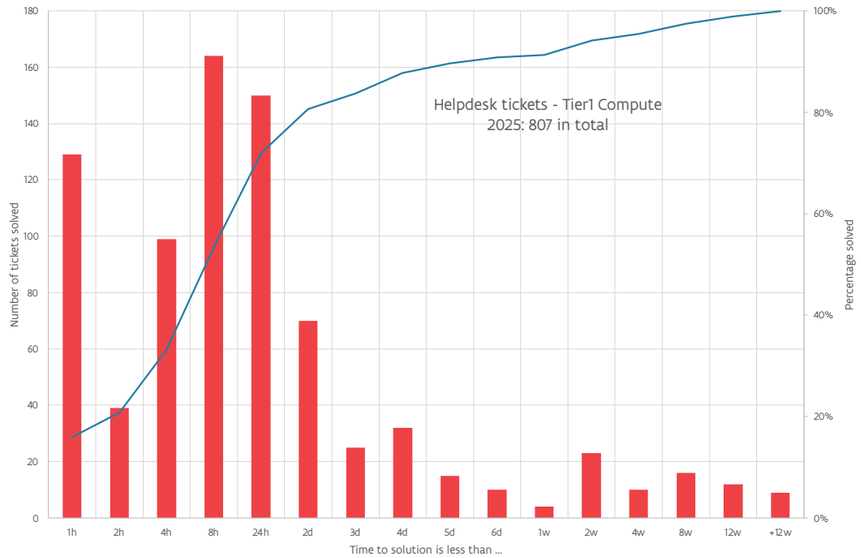

Figure 7. VSC Tier-1 Compute Hortense Ticket Resolution Time

In 2025, the HPC-UGent team also handled a number of software installation requests for the Tier-1 infrastructure. A more detailed analysis of these activities is provided in the section on the UGent Tier-2 infrastructure. A total of 70 requests for software installations on Tier-1 Hortense were submitted during the year: 26 from academic researchers and 44 from industrial users. In many cases, however, the requested software packages required the installation of additional software and libraries to satisfy dependency requirements.

As a result, more than 438 software installations were carried out on Tier-1 Hortense in 2025, including 42 GPU-related software installations.

User Survey

User satisfaction with the Tier-1 Compute infrastructure was also assessed through the 2025 VSC User Survey. A total of 90 respondents indicated that they make use of this VSC service to some extent. Of these respondents, 64% stated that Tier-1 Compute is critical to their research or professional activities.

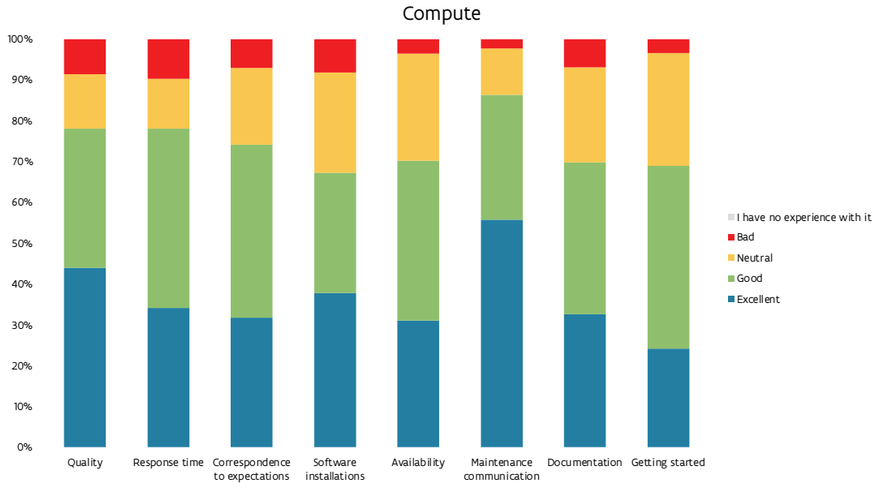

Figure 8. VSC Tier-1 Compute Hortense: User Survey 2025

The following aspects of the Tier-1 Compute infrastructure were rated as Good or Excellent by users:

- 78% – Quality
- 78% – Response Time
- 73% – Meeting Expectations
- 66% – Software Installations
- 70% – Availability
- 85% – Maintenance Communication
- 69% – Documentation
- 69% – Getting Started

Twenty-seven respondents provided additional feedback on the Tier-1 Compute platform through an open comment field. Most of the negative feedback came from a small number of respondents and focused on helpdesk response times (8 comments), the available documentation (8 comments), and the turnaround time for software installations (3 comments). Eight of these respondents also assigned several Bad ratings to the Tier-1 Compute service aspects listed above. This contributed to the lower scores for Response Time (78%), Software Installations (66%), Documentation (69%), and Getting Started (69%) compared to 2024.

The suggestions for improvement—although sometimes expressed in fairly general terms—can be grouped into the following categories:

- Documentation
  - Update the existing documentation.
  - Improve documentation for new and beginning users.
  - Provide practical, worked examples for as many scientific software packages as possible
- User Experience
  - Reduce helpdesk response times.
  - Provide additional helpdesk support for new users.
- Software
  - Shorten turnaround times for software installation requests.
- Infrastructure
  - Improve Tier-1 Compute I/O stability.
  - Increase the number of CPU and GPU nodes to reduce queue and turnaround times.
  - Provide a backup Tier-1 system that can be used during maintenance periods.

Some of these comments suggest that an increasing number of new users are quickly finding their way to Tier-1 Compute, even though they may not yet have acquired the necessary Tier-1 skills. This growing accessibility of Tier-1 Compute is a positive development; however, further strengthening user skills remains an important area of attention for all Tier-2 sites

One suggestion for future developments within the VSC ecosystem was:

- The introduction of a hosted LLM platform for end users.

What Our Users Say

"Very good collaboration!"

"Support is excellent and communication is very helpful and direct"

"Overall pleasant service"

"Keep up the good work!"

"Quality of the hardware is good considering that we are in a small region. Having our own infrastructure is very valuable"

TIER-1 SOFIA

Through the VSC 2.0 strategic plan and the associated investment funding, the Flemish Supercomputer Center (VSC) secured the resources needed to deploy a new Tier-1 supercomputer. This new infrastructure will succeed the current Tier-1 system, Hortense, and represents an important step in the continued development of Flanders' computing infrastructure for scientific and industrial research. In 2023, Vrije Universiteit Brussel (VUB) was appointed as the institution responsible for the procurement, installation, and operational management of the new Tier-1 system.

In 2024, the European public procurement procedure was launched. The process was completed in early 2025 with the award of the contract for the new Tier-1 system to NEC Deutschland GmbH. The evaluation of the submitted proposals considered not only the technical specifications of the system, but also its energy efficiency. The criterion of "cost per FLOP per watt" played a central role, with the objective of selecting a system that combines maximum computational performance with the most efficient possible energy consumption. The market consultation demonstrated that an architecture with a strong focus on GPU accelerators offered the best balance between performance and energy efficiency.

The new Tier-1 system will consist of four specialized partitions, each designed for a specific type of workload: a CPU partition for highly scalable applications, a high-memory partition for data-intensive workloads, a GPU partition optimized for AI applications, and a partition dedicated to visualization and interactive use.

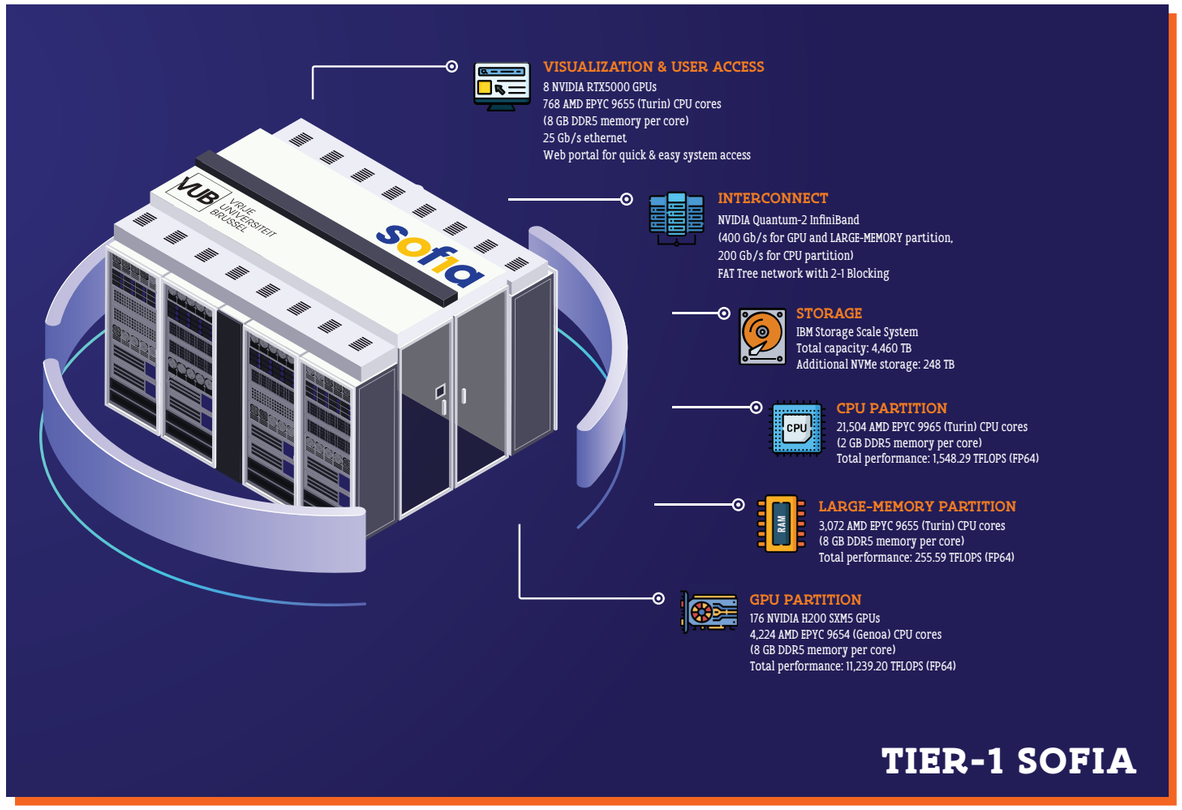

Figure 9. Overview of the Infrastructure and Technical Specifications of the Tier-1 Sofia

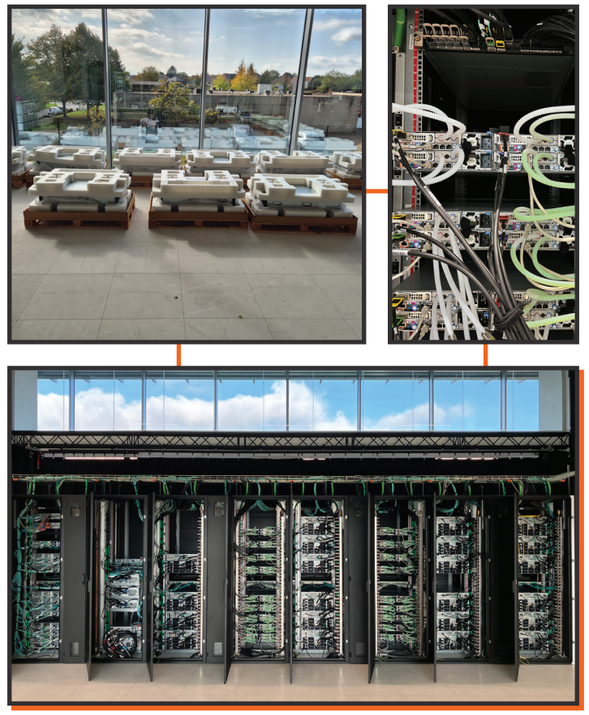

On February 20, 2025, the procurement contract was formally awarded, starting the countdown toward delivering the system into operation by November 20, 2025. In parallel, extensive preparations were underway at the Nexus data center in Research Park Zellik to accommodate the new supercomputer infrastructure.

By the end of August, most of the hardware had been delivered. One critical component, however, was still pending: the central heat exchanger. The new Tier-1 system uses direct liquid cooling to remove heat from the CPUs and GPUs as efficiently as possible. This technology enables a more compact system design while reducing energy losses associated with traditional air-cooling solutions.

The system operates with a closed cooling circuit, in which the coolant is cooled through the data center's water circuit via the central heat exchanger. In the longer term, the residual heat generated by the supercomputer is expected to be recovered and reused for heating other buildings within Research Park Zellik.

Because the central heat exchanger is an essential component of the cooling system, the supercomputer could not be brought online without it. To avoid further delays, it was decided to temporarily install two heat exchangers while awaiting delivery of the final unit. This solution allowed the installation and configuration of the system to proceed as planned.

In early November, the system was fully powered up and entered the testing phase. The original plan had been to launch a pilot phase in October, enabling researchers to submit projects for the first 2026 call. However, the delay in the delivery of the cooling infrastructure required this schedule to be revised. The current planning foresees a pilot phase starting in April 2026, allowing projects from the second 2026 call to run on the new system. The production environment is scheduled to become operational on July 1, 2026.

Despite the delay, an official inauguration ceremony was held on November 22. During the event, Minister-President Matthias Diependaele officially unveiled the name of the new VSC Tier-1 supercomputer: sofia.

Figure 10. (top) Architecture of the Sofia Tier-1 infrastructure; (middle) the Nexus data center in Research Park Zellik, where Sofia was installed; (bottom) server rack of the new VSC Tier-1 supercomputer featuring the Sofia logo

During the event, several users highlighted the importance of this type of infrastructure for research and innovation. Professor Wim Thiery explained how Tier-1 computing power is essential for advanced climate research. Ruben Sevenois of DUCO demonstrated how industrial users leverage such infrastructure to strengthen their innovation capacity. DUCO is a leading European provider of natural ventilation and solar shading solutions. Finally, Geert Van Minnebruggen of VIB outlined how access to large-scale computing infrastructure enables life sciences researchers to remain competitive at the highest international level.

The year concluded with the first benchmark results for sofia. These results are highly promising: all targeted performance objectives have been achieved, while energy consumption is lower than originally anticipated.

In 2026, the system will undergo the remaining testing and validation phases. A pilot phase will begin in April, allowing researchers to explore the capabilities of sofia, followed by the start of full production operations in July 2026.

TIER-1 DATA

KU Leuven is responsible for the infrastructure and operation of the Tier-1 Data platform. The platform is intended for active data processed on the VSC Compute and Cloud components. Publication and long-term storage take place on other platforms, although the availability of metadata facilitates the transition to this phase.

AVAILABLE INFRASTRUCTURE

The Tier-1 Data platform is built on four Lenovo Distributed Storage Solution for Spectrum Scale (DSS-G280) systems. It provides 27 PB of usable storage capacity (with dual copies), installed across two separate data centers with synchronous replication (mirroring), ensuring that data are protected against major incidents. In addition, the system supports file system-level snapshots for enhanced data protection, and data are encrypted at rest through software encryption. The maintenance contract for the storage infrastructure was extended. This ensures that the platform's capacity remains available for the next three years and provides sufficient transition time to carry out the migration to a future system. At the end of this period, the risk of disk failures becomes significant. The future system must therefore be prepared in a timely manner.

In 2025, the storage capacity for OpenSearch, which is used for metadata indexing, was also expanded. To support Globus usage, new Globus servers with enhanced connectivity were procured as well.

OPERATIONS AND USAGE

A minor and a major upgrade of iRODS (versions 4.3.4 and 5.0) were carried out, ensuring that VSC remains aligned with developments within the consortium. Membership in the iRODS Consortium and the Globus licenses were renewed for a further three years.

MANGO PORTAL

The ManGO Portal, the graphical user interface built on top of the iRODS infrastructure, continued to be developed by KU Leuven. In 2025, progress was made in expanding functionalities related to collection management, metadata processing, and access control within the platform. A series of targeted enhancements to both the frontend and backend made the system more robust and user-friendly.

These functionalities and improvements were released through five portal updates (v0.18.4–v0.23). The aim of these releases was to respond as effectively as possible to the evolving needs of researchers.

A first step involved making the Gallery view available on request, with the long-term goal of enabling users to activate this functionality themselves. In addition, a persistent and unique URL structure was introduced for collections and data objects, making it easier to reference and share collections and files. Existing access controls remain fully enforced.

This enhancement aims to make the user interface more intuitive and easier to navigate, while significantly expanding the underlying search capabilities. This includes support for metadata schema-based searches both within individual projects and across multiple projects, while always respecting user-specific access permissions. Preparations were also made for the further integration of OpenSearch indexes, enabling more advanced and efficient search functionality.

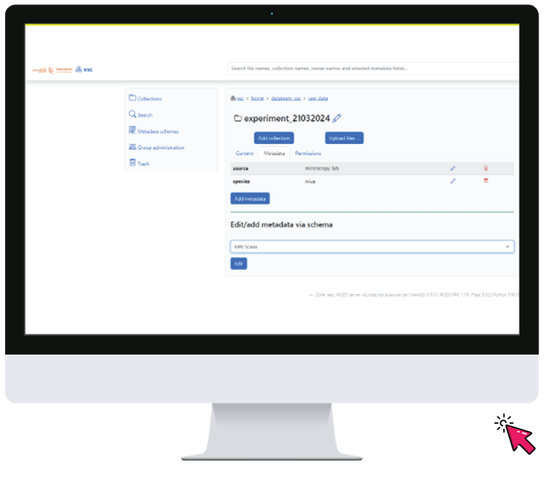

Figure 11. Screenshot of the VSC Mango Portal

In addition, development was initiated for folder upload functionality, which will simplify the management and addition of datasets through the portal. For uploading large datasets, however, it remains recommended to use alternative clients that are not subject to HTTPS limitations.

Finally, an important step was taken toward the automation of access management. With the introduction of YAML-based group and permission management, access rights can now be updated automatically based on predefined structures. This automation is reserved for accounts with group administration privileges and further enhances both the security and efficiency of access management within the platform.

The Python iRODS Client (PRC), which serves as the foundation of the ManGO Portal, was upgraded to version 3.1. In addition, all Python packages were updated to their latest versions.

Progress was also made in the area of data export as part of an archiving strategy for projects that are no longer active. If desired, users will be able to download entire collections, including their metadata, in a format that can subsequently be transferred to cold storage. This will be accompanied by a manifest file, ensuring that users can always identify which files are included in the archive package without having to download the archive itself. Once the required file has been located, the archive package can then be downloaded. An important aspect of this approach is that the contents of archived datasets remain discoverable without requiring the download of the entire archive.

TIER-1 DATA SERVICE DEVELOPMENT

In 2025, significant progress was made in the further development of the iRODS-based data infrastructure and its supporting workflows. A key focus was improving the accessibility and reliability of data transfers. Extensive testing was carried out using different transfer methods, including iCommands, iRODSFS mounts, and NFS mounts, across both Tier-1 and Tier-2 environments. These test campaigns provided valuable insights into system performance and formed the basis for further backend optimizations.

The iRODS NFS functionality gradually proved sufficiently stable for broader testing and was deployed on the VUB Tier-2 infrastructure, where it is currently being evaluated by research groups that use it to perform efficient data post-processing.

In addition, the potential of tools such as Xarray to support more efficient data transfers in data-intensive workflows was investigated.

In parallel, work was carried out to modernize the authentication framework. To eliminate the need for short-lived passwords and enhance security, a workflow was developed for PAM interactive authentication using tokens obtained through multi-factor authentication (MFA). Following successful validation in the test environment, this authentication method was also deployed in production. The existing login method remains available for the time being and is scheduled to be phased out gradually from 2026 onwards.

Finally, an important step was taken to improve the robustness of large-scale data transfers through the development of IRON, a new Go-based utility and command-line client for iRODS. This tool replaces the existing go-irodsclient developed by CyVerse and provides a more reliable and efficient solution for transferring large datasets.

TIER-1 DATA USAGE

The figure below shows the evolution of storage utilization throughout 2025.

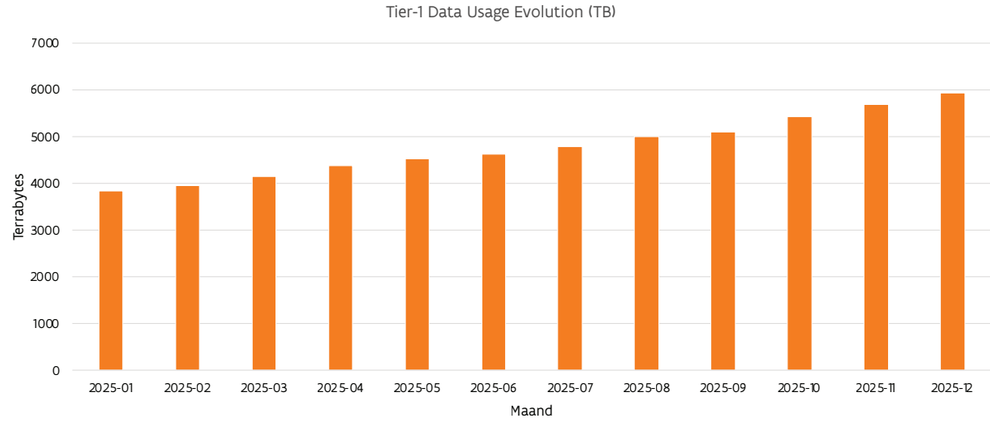

Figure 12. VSC Tier-1 Data Usage Evolution (TB)

GLOBUS DEPLOYMENT AND USAGE

Globus (www.globus.org) is a software solution that enables the secure, fast, and reliable transfer of data between different storage platforms. As shown in the figure below, all VSC sites use the platform to facilitate data transfers between various VSC infrastructure components and a large number of external endpoints.

In 2025, 847 users transferred more than 5.3 PB of data using Globus. Compared to the previous year, this represents an increase of more than 100%.

Figure 13. Screenshot of globus.org

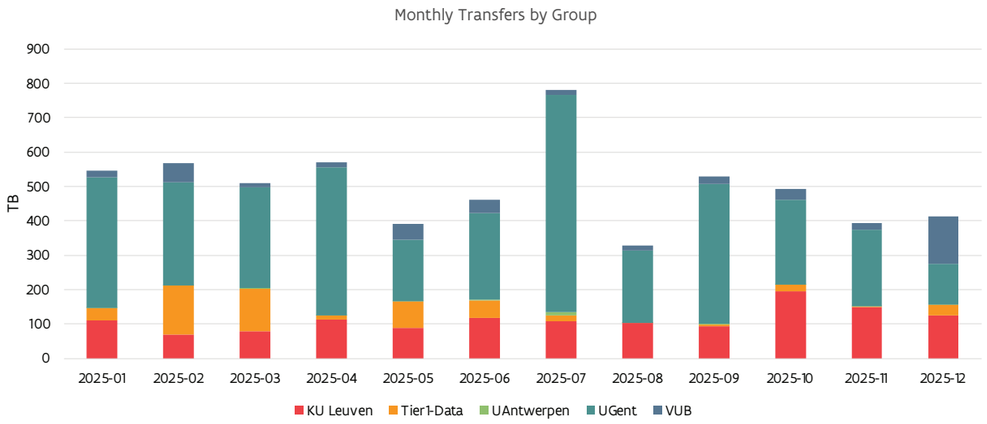

Figure 14. VSC Tier-1 Data – Monthly Transfers per Group

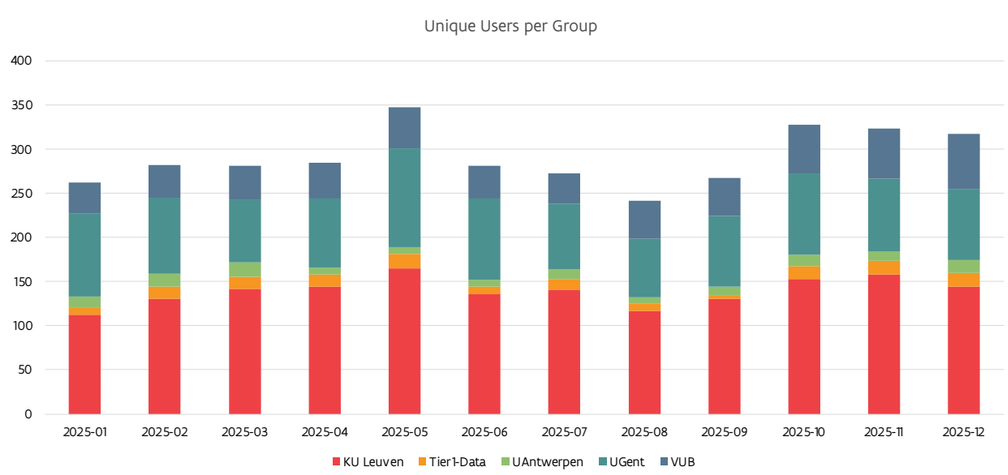

Figure 15. VSC Tier-1 Data – Number of Unique Users per Month per Group

User Survey

The 2024 user survey identified several action points that were addressed during 2025, including improvements to the documentation, particularly for Globus, increased attention to Globus transfer performance, and the investigation and testing of options for providing mount access to Tier-1 Data.

User satisfaction was again assessed through the 2025 VSC User Survey. For Tier-1 Data, 26 respondents indicated that they make use of this VSC service to some extent.

The following aspects of the Tier-1 Data infrastructure were rated as Good or Excellent by users:

- 85% – Quality
- 90% – Response Time
- 83% – Meeting Expectations
- 75% – Client Usability
- 92% – Availability
- 86% – Maintenance Communication
- 84% – Documentation
- 81% – Getting Started

Several categories received higher ratings than in the previous year, including Response Time, Availability, Maintenance Communication, Documentation, and Getting Started. Three other categories received slightly lower scores but still achieved strong overall ratings.

The limited number of comments did not highlight any specific concerns. Nevertheless, additional attention will be given to the various methods available for accessing and working with Tier-1 Data, as well as to maintaining clear and up-to-date documentation for these tools and services.

What Our Users Say

> "I am glad for the service from Tier-1. Very helpful!"

> "Mango is a very nice platform from what I've seen"

Allocation of Storage Capacity

Tier-1 Data projects are allocated through a project application process. In 2025, nine project applications were approved, representing a total storage requirement of 1,060 TB over a four-year period.

The three existing Collaboration Grants that make use of the Data component were renewed, accounting for a combined storage allocation of 7.015 PB.

User Support

At the beginning of 2025, training sessions were delivered to approximately 70 researchers, followed by a dedicated training session for Ghent University researchers in April. Later in the year, a specialized session was also organized for researchers from INBO. Several Globus training sessions were held throughout the year and attended by 38 researchers. During the KU Leuven HPC User Day on June 2, the topic "Tier-2 Data Management: From Ingest to Outflow" was also featured.

The Tier-1 Data documentation was updated with current information on login procedures using the Python iRODS Client, iCommands installations on VSC clusters, and new portal functionalities. In addition, the support team organized intake meetings with several research groups to prepare project applications and assist users in getting started with the platform.

The iRODS Consortium

During the annual iRODS User Group Meeting (iRODS UGM), KU Leuven delivered three presentations and two lightning talks:

- ManGO Platform Updates: ManGO Portal, ManGO Ingest, and ManGO Flow
- Metadata Schema Updates: JSON Schemas and Storage in iRODS
- FriGO: The KU Leuven Long-Term Archiving Solution with iRODS
- irods2dataverse: Python package for depositing an iRODS dataset into Dataverse (lightning talk)
- FriGO: Long-Term Archiving with iRODS in Action (lightning talk, including a live demonstration following the presentation)

Through its active involvement in the consortium, the support team not only remains closely informed about the latest developments but also benefits from opportunities to establish collaborations and jointly develop new ideas and solutions.

SHOWCASE

Bioinformatics for Biodiversity Genomics

Several projects granted access to Tier-1 Data use the platform to store genomic data. One project we would like to highlight aims to compare the evolutionary trajectories of two fish groups with contrasting evolutionary histories.

The first species, Clarias gariepinus (African sharptooth catfish), has a broad ecological and geographical distribution across Africa and a long evolutionary history. In contrast, species within the genus Enteromius are highly diverse and endemic to specific river systems. Despite considerable genetic differences, these species are often morphologically indistinguishable. The research investigates the genomic architecture underlying this species diversity and explores the genetic mechanisms that may contribute to reproductive isolation and speciation.

By comparing the evolutionary trajectories of these fish groups, the research team aims to gain deeper insights into the evolutionary history of African freshwater fish, the processes driving speciation, and the environmental history of African river systems. A key aspect of this project is the collaboration between researchers from multiple institutions, including Hasselt University, the Royal Belgian Institute of Natural Sciences (RBINS), and KU Leuven, who work with shared datasets stored on Tier-1 Data. Data processing can be carried out using local infrastructure as well as the various Tier-2 systems. Results are written back to the central Tier-1 Data platform, making them immediately available to all project partners..

TIER-1 CLOUD

Researchers using the VSC Tier-1 HPC services often require a flexible environment that allows them to use software or services beyond a large-scale HPC system, such as customized software packages, interactive data analysis, workflow portals, data visualization, and specific pre- and post-processing tasks.

The Tier-1 Cloud infrastructure provides these services through on-demand resources delivered in a flexible, cloud-like environment.

This Infrastructure-as-a-Service (IaaS) facility enables users to deploy and fully manage resources such as virtual machines (VMs), storage, and networking. A catalog of VM and orchestration templates is available, making it easy to deploy virtual machines with different configurations—such as web servers, basic clusters, and other environments—with just a few clicks.

STRATEGIC EVALUATION

In 2025, VSC conducted a thorough assessment of the future of the Tier-1 Cloud service. A reinvestment in the infrastructure will be required in 2026, making it necessary to evaluate strategic options and user requirements in light of the available funding.

Usage statistics and user feedback collected over more than seven years of Tier-1 Cloud operations were consolidated into an evaluation report that was presented to the VSC User Committee and the HEC Council. One of the key findings was that the Tier-1 Cloud service is used in a highly diverse manner by different project owners.

Based on the project descriptions submitted in funding applications, the following overview of commonly used tools was compiled.

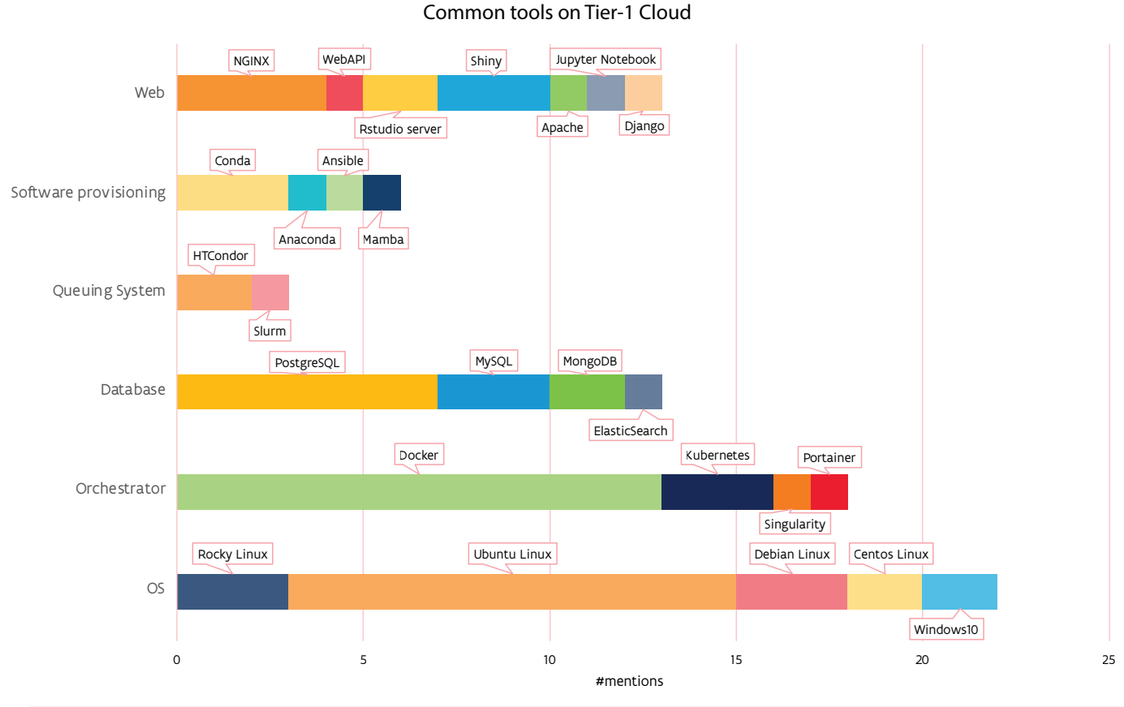

Figure 16. VSC Tier-1 Cloud – Most Used Tools

The VSC User Committee and the largest Tier-1 Cloud users were also consulted for feedback. The survey highlighted several clear conclusions:

- The VSC Tier-1 Cloud is a critical infrastructure component for a number of research infrastructures.
- Overall satisfaction with the service is very high.
- Users perceive the platform as highly stable, with excellent availability and a robust underlying storage platform.
- The VSC staff members responsible for supporting the Tier-1 Cloud infrastructure and helpdesk are considered highly valuable by the user community.

Figure 17. VSC Tier-1 Cloud – User Overview

In an additional survey, Tier-1 Cloud project owners were specifically asked about the use of the services they had developed on the Tier-1 Cloud infrastructure. Through these services, they reach a much broader community of (new) computational users, many of whom do not necessarily have a VSC ID and therefore do not appear directly in the VSC usage statistics. Based on the feedback received, the Cloud projects can be grouped into three categories of services:

- Sandbox-type virtual machines, used by project owners and power users for development activities, testing, and proof-of-concept purposes.
- Medium-scale services, developed by project owners for a limited user base (typically 10–15 users) within a clearly defined service offering.
- Large-scale services, through which project owners support a substantial user community (typically 100–1,000 users) and provide a broader range of applications within their research domain.

One example of such a large-scale service is the Belgian Galaxy Node, operated by VIB (Flanders Institute for Biotechnology) as part of the ELIXIR project (https://usegalaxy.be). This is one of the larger projects hosted on the VSC Tier-1 Cloud infrastructure.

A team of approximately five project owners from VIB Data Core manages an extensive collection of applications and datasets for the life sciences community. The service is used by more than 1,367 scientists from a wide range of Flemish research institutions. The figure below shows the distribution of users of this platform by institution (data provided by VIB Data Core).

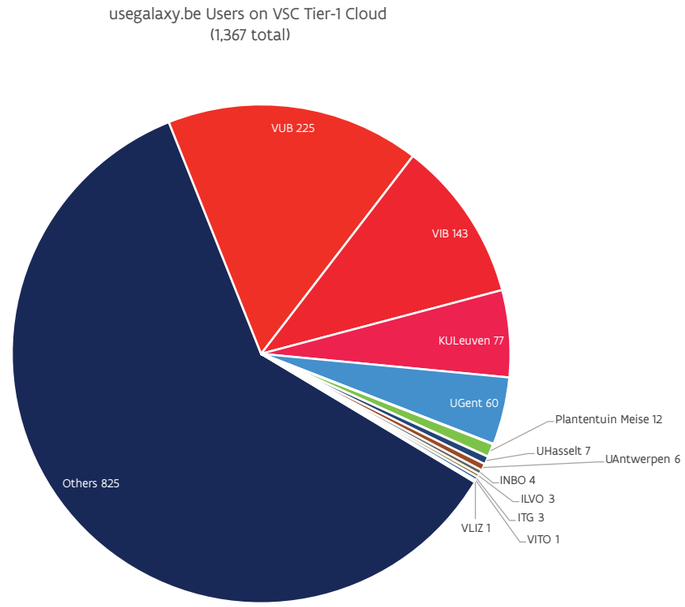

Figure 18. VSC Tier-1 Cloud – usegalaxy.be Users

Following this positive evaluation of the VSC Tier-1 Cloud service, the HEC Council approved a reinvestment plan. In 2026, investments will be made in new hardware, while the platform will simultaneously migrate to OpenNebula as a more cost-efficient technical framework. These activities, together with the gradual migration of all current users to the new platform, will occupy most of 2026.

OPERATIONAL ACTIVITIES AND USAGE IN 2025

In 2025, several functional updates were implemented to maintain and support the Tier-1 Cloud service:

- All servers within the Tier-1 Cloud infrastructure were upgraded to a newer version of the operating system.
- A new monitoring system based on Prometheus and Grafana was successfully deployed.

The graphs on the following page present several indicators of Tier-1 Cloud service usage in 2025.

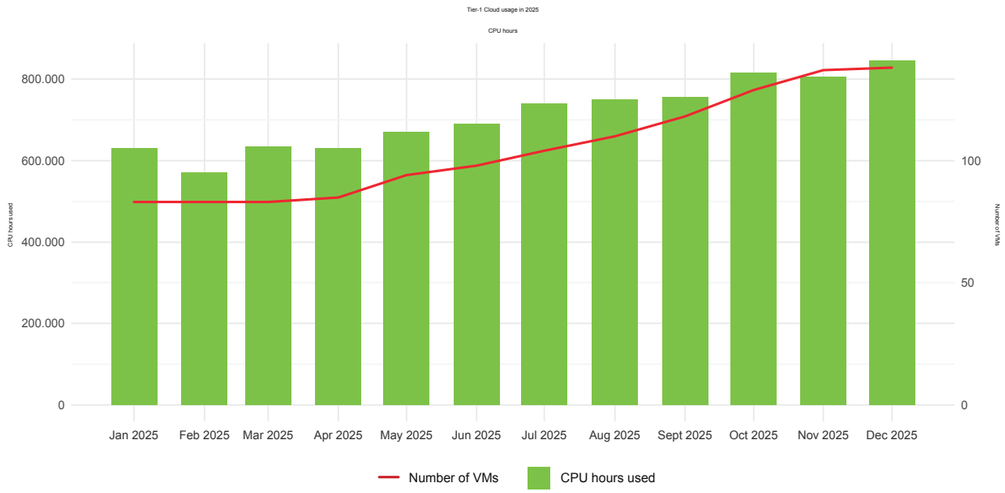

Figure 19. VSC Tier-1 Cloud – CPU Hour and VM Usage Statistics 2025

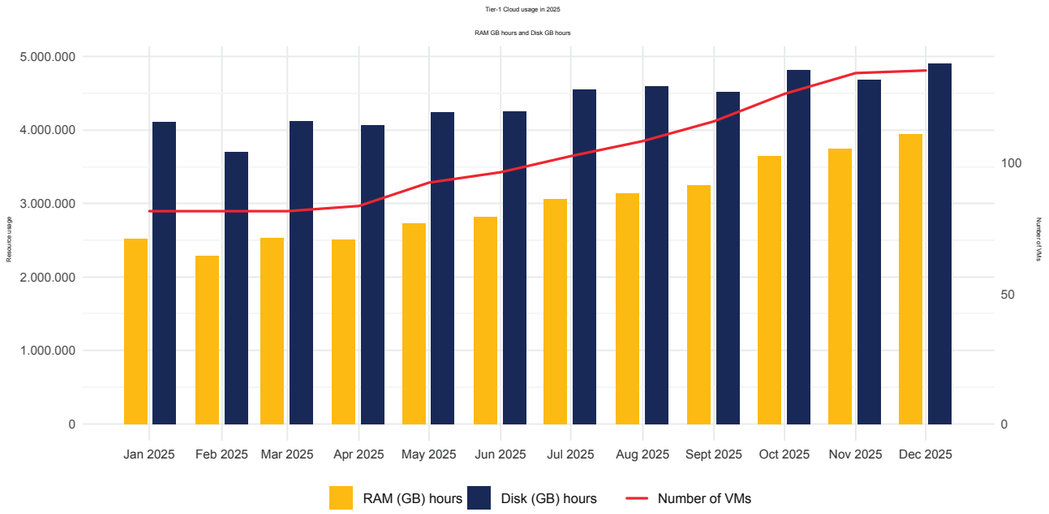

Figure 20. VSC Tier-1 Cloud – Memory, Storage, and VM Usage Statistics 2025

As in 2024, steady growth was observed in both the total number of virtual machines (VMs) in use and the amount of compute time consumed. In total, 33 different Tier-1 Cloud projects were active in 2025, consuming a combined 8,530,326 CPU hours of compute time.

Memory and storage usage also increased steadily throughout 2025, reaching a total consumption of 52,539,538 RAM GB-hours and 36,338,781 Disk GB-hours, respectively.

[1] Memory GB-hours are a unit used to measure HPC resource consumption. They are calculated for each job as the amount of allocated memory (in GB) multiplied by the job runtime (in hours).

A total of 11 Starting Grants and 3 Full Proposals were submitted, all of which were approved. The three Full Proposals included one new project and two projects building on an existing Starting Grant. In addition, several (sub)projects were created within the framework of the VIB Data Core Collaboration Grant, a consortium involving VIB, Ghent University, KU Leuven, the University of Antwerp, Vrije Universiteit Brussel, and Hasselt University.

The dedicated Tier-1 Cloud helpdesk (cloud@vscentrum.be) closed a total of 84 support tickets in 2025. More than half of all cases were resolved within 48 hours.

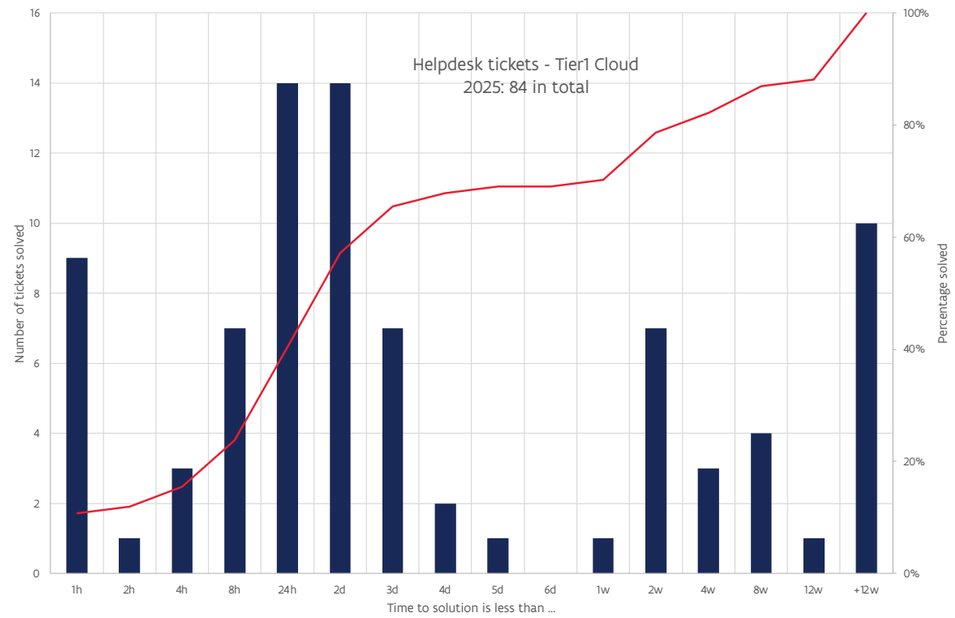

Figure 21. VSC Tier-1 Cloud – Helpdesk Tickets 2025

USER SURVEY

User satisfaction with the Tier-1 Cloud infrastructure was also assessed through the VSC User Survey conducted at the end of 2025. A limited number of respondents (16) indicated that they make use of the Tier-1 Cloud infrastructure to some extent. Of these respondents, 43% stated that Tier-1 Cloud is critical to their research or professional activities.

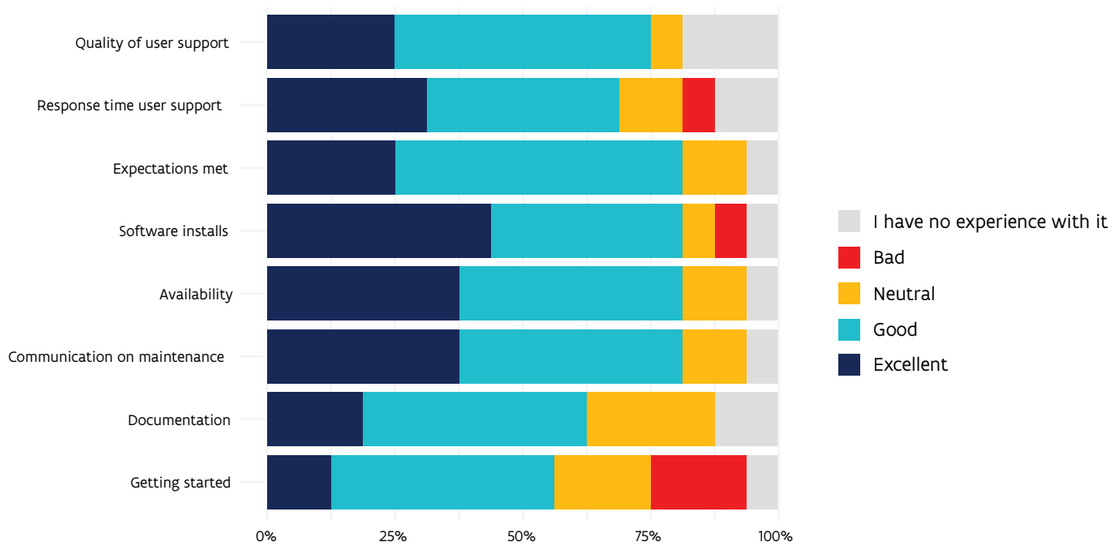

Figure 22. VSC Tier-1 Cloud – User Satisfaction 2025

The following aspects of the Tier-1 Cloud infrastructure were rated as Good or Excellent by users:

- 92% – Quality
- 79% – Response Time
- 87% – Meeting Expectations
- 87% – Software Installations
- 87% – Availability
- 87% – Maintenance Communication
- 71% – Documentation
- 60% – Getting Started

The following suggestions for improvement were provided:

- Infrastructure
  - Upgrade the infrastructure with additional and more recent GPUs (1 respondent).
- Documentation
  - Improve the documentation by adding FAQs (1 respondent).
  - Provide a step-by-step guide or demonstration for setting up a Cloud environment aimed at non-system administrators (2 respondents).
  - Provide more detailed documentation on the hardware available within the Tier-1 Cloud infrastructure (2 respondents).
- User Experience
  - Provide a simpler way to monitor service uptime than through the status page (https://status.vscentrum.be) (1 respondent)
- The following suggestions for new developments within the VSC ecosystem were also mentioned:
  - A container service.
  - A VSC-managed database service for non-system administrators.

RDI
CUSTOMER PORTFOLIO

The VSC infrastructure (Tier-1 Compute and Tier-1 Cloud) is also accessible to the Flemish private and public sectors for research and development purposes. Each new non-academic user is offered a free trial period through a so-called exploratory contract worth € 5,000. This contract allows users to request either 645,995 CPU core hours or 5,102 GPU hours, each accompanied by 1 TB of storage. This enables users to conduct a proof-of-concept experiment and assess whether the VSC infrastructure meets their specific requirements.

Since June 2023, and within the framework of the VSC 2.0 plan, these non-academic users have received both administrative and technical support from the VSC Research, Development & Innovation (RDI) Office. This central point of contact consists of one full-time equivalent (FTE) business developer and one FTE with a technical profile. In addition, the team is strengthened by two additional colleagues active within the European EuroCC project, who are responsible for promoting VSC services among external stakeholders.

In 2025, the RDI Office increased the visibility of VSC, and in particular its non-academic services, by actively participating in both domain-specific events and initiatives aimed at a broader audience. In this context, VSC explicitly positioned itself as an alternative provider of compute services for industry and the public sector by hosting booths at the Trefdag Digitaal Vlaanderen event in Ghent, the Knowledge for Growth conference in Antwerp, and the Digital in Action event in Brussels.

The Flemish AI community was also informed about the opportunities offered by the GPU partitions on both Hortense and the new supercomputer sofia for training AI models. This was achieved through participation in the FARI 2025 conference and the MateriNex event, as well as by establishing an official collaboration with sustAIn. brussels.

At the international level, the EuroCC colleagues were also actively involved. They participated in the EuroHPC Summit in Poland, the EuroHPC User Days in Denmark, and the annual general meeting with all EuroCC partners in Estonia.

Finally, the RDI team organized a webinar for the team of business advisors at VLAIO on the opportunities offered by HPC and the VSC service portfolio, with the aim of supporting SMEs and startups at an early stage. Within the same context, the Digital & Innovative Entrepreneurship ecosystem of the Federation of Belgian Enterprises (VBO) was also actively engaged.

These efforts resulted in a total of 24 new exploratory contracts being signed by non-academic users from both the private and public sectors during 2025, representing a year-on-year increase of 71%. In addition, 17 existing active (i.e., regular) contracts were extended by one year, and 8 completed exploratory projects were converted into regular contracts.

As a result, the VSC portfolio of active non-academic users reached a total of 39 contracts by the end of 2025, compared with 14 at the end of 2023. This represents an increase of almost 200% and highlights the growing importance of HPC infrastructure for industrial R&D activities in Flanders.

As shown in the figure on the following page, more than half of this portfolio (54%) consists of small and medium-sized enterprises (SMEs). In addition, VSC services are becoming increasingly established within the Flemish public sector, which now represents 10% of the portfolio. We expect this trend to continue in 2026.

In terms of sector diversity, 38% of our stakeholders are active in engineering or manufacturing companies. As in previous years, this group remains the largest representation within the portfolio. At the same time, we observed that in 2025, in line with global trends, the number of companies active in artificial intelligence development and making use of our GPU partitions increased significantly. These companies now represent 21% of the total portfolio.

The portfolio further includes users from the sectors of environmental services and utilities (10%), research and development (10%), healthcare and consumer services (8%), biotechnology (8%), and materials and chemistry (5%).

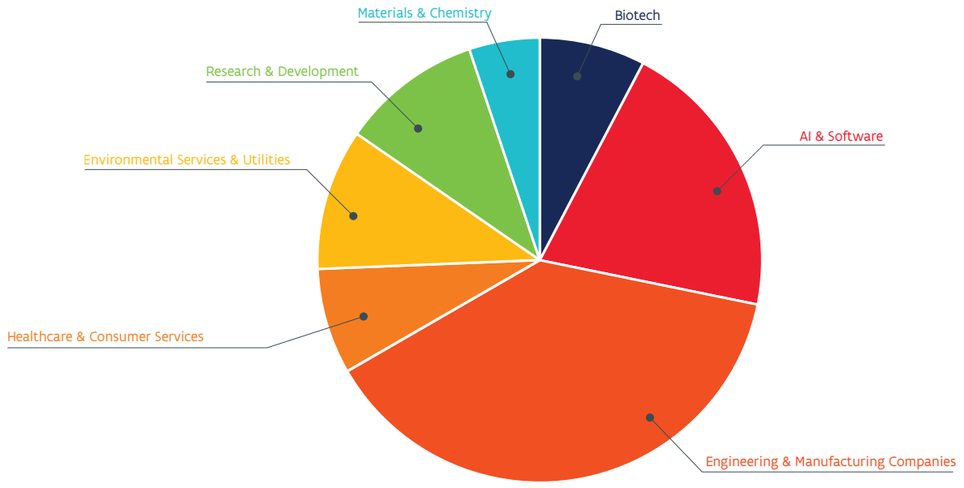

Figure 23. VSC RDI – Distribution by Sector

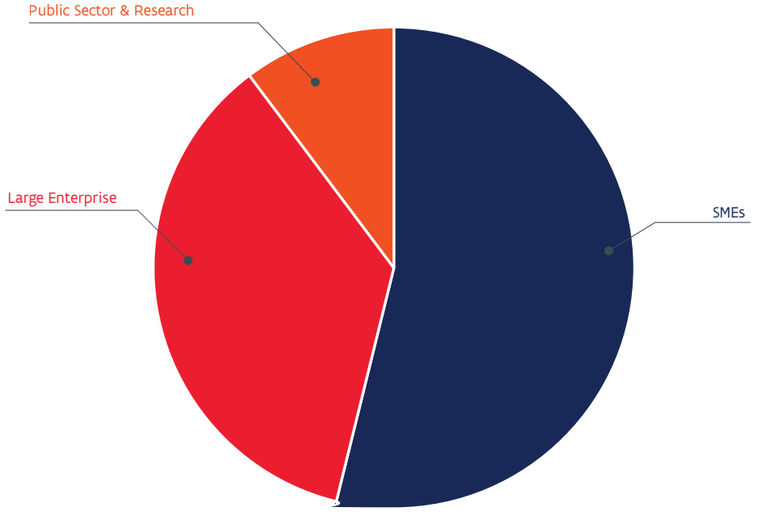

Figure 24. VSC RDI Users by Organization Type

In the 2025 financial year, a total amount of € 448,047 was invoiced, corresponding to an average monthly revenue of € 37,337. This revenue is entirely based on the effective use of compute capacity (CPU or GPU hours) and reserved storage capacity, according to the following pricing structure (including VAT): € 24 per CPU node per day, € 94 per GPU node per day, and € 16 per additional TB of storage per month.

This income is used, on the one hand, for further investments in the Tier-1 infrastructure and, on the other hand, enables the RDI Office to participate in or organize outreach activities.

Despite an increase in the number of users, revenue in 2025 was approximately 20% lower than in 2024. This is mainly explained by the fact that the Tier-1 infrastructure was unavailable to users for approximately one month in total (spread across several periods) due to maintenance activities.

SUPPORT

In addition to academic users, the VSC User Survey also assessed the satisfaction of non-academic users with the Tier-1 Compute infrastructure and its associated services. In total, 16 participants affiliated with private companies took part in the survey. Their evaluation, ranging from poor to excellent, is presented graphically below. The figure shows that VSC services continued to receive an overall positive evaluation from this target group in the past year. The quality of the service achieved the maximum score this year (16/16). The alignment with expectations (15/16) and communication regarding maintenance (15/16) were also rated as "Good" or "Excellent" by a very large majority of respondents (94%). Response time was positively evaluated by 88% of respondents (14/16), while 75% described the cluster as user-friendly.

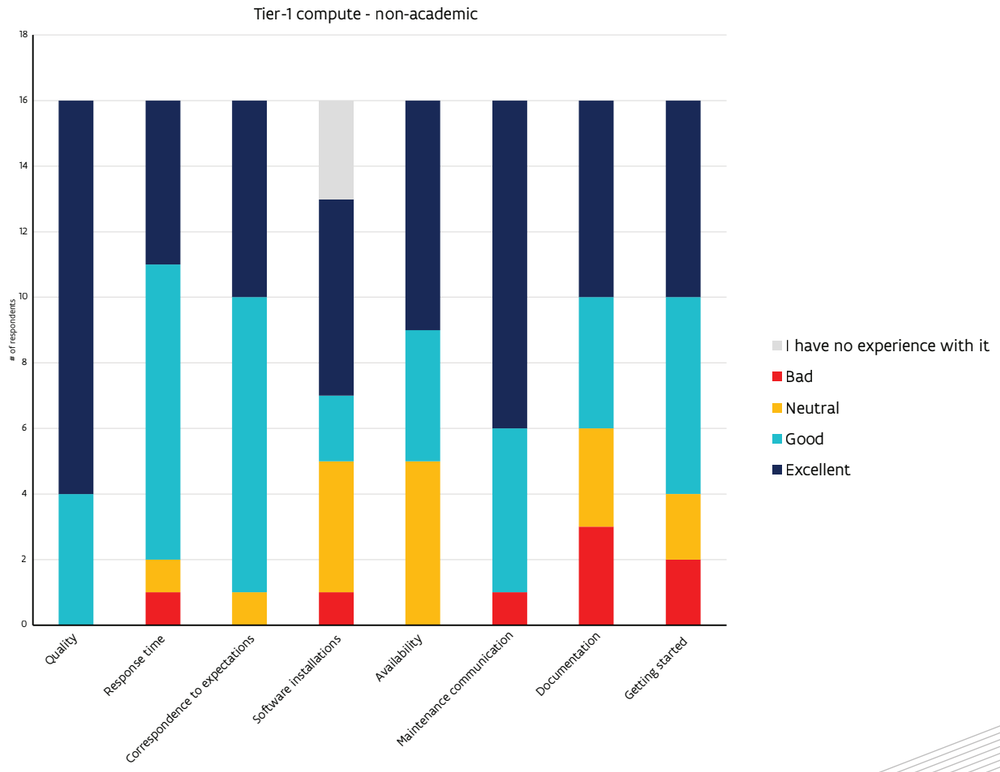

Figure 25. VSC RDI in VSC User Survey 2025

For Availability (69%), Documentation (63%), and Software Installations (62%), satisfaction levels were somewhat lower. These results indicate that, alongside the overall positive evaluation of the services, there remain areas requiring further improvement.

Regarding infrastructure availability, it should be noted that Tier-1 Compute was unavailable for approximately one month during the past year due to maintenance activities addressing serious hardware issues. This exceptional situation had a clear impact on the satisfaction score for this aspect. However, with the deployment of sofia, our new Tier-1 supercomputer, users will gain access to new and more powerful generations of GPU and CPU nodes. This modernization will not only result in more efficient computational performance but will also contribute to improved infrastructure stability and availability in the long term.

User feedback also indicates that software installations, and more specifically their turnaround time, are perceived less positively. In this context, it is important to explain the operational environment. The RDI team currently consists of two full-time staff members, with only one person responsible for all technical support for non-academic users, including software installations.

In 2025, a total of 693 support requests were registered, representing an increase of 70% compared to 2024. This significant growth inevitably impacts the processing time of individual requests. In addition, non-academic users often rely on commercial software. Such software is frequently more complex to install because the source code is not accessible, and the available documentation may be limited or restricted, making analysis and troubleshooting more challenging.

Although every request is handled with the necessary care, this workload requires more time per support ticket. If this growth continues, there is a risk of structural overload in the longer term. To proactively address this challenge, an adapted and more efficient support procedure is currently being developed. This procedure will come into effect once sofia becomes operational, with the aim of structurally improving both response times and the overall user experience.

TIER-2 SUPERCOMPUTING PLATFORM

UNIVERSITY OF ANTWERP

AVAILABLE INFRASTRUCTURE

In 2025, no changes were made to the existing Leibniz and Vaughan clusters. The configuration therefore remains as follows:

- 3 clusters, with CPU, GPU, and large-memory partitions
- 681 TF CPU, 106 TF GPU
- 17,216 CPU cores, 12 GPU devices, and 2 dedicated accelerators
- 620 TB scratch file system

In the autumn of 2025, a public procurement procedure was initiated to replace the Leibniz cluster.

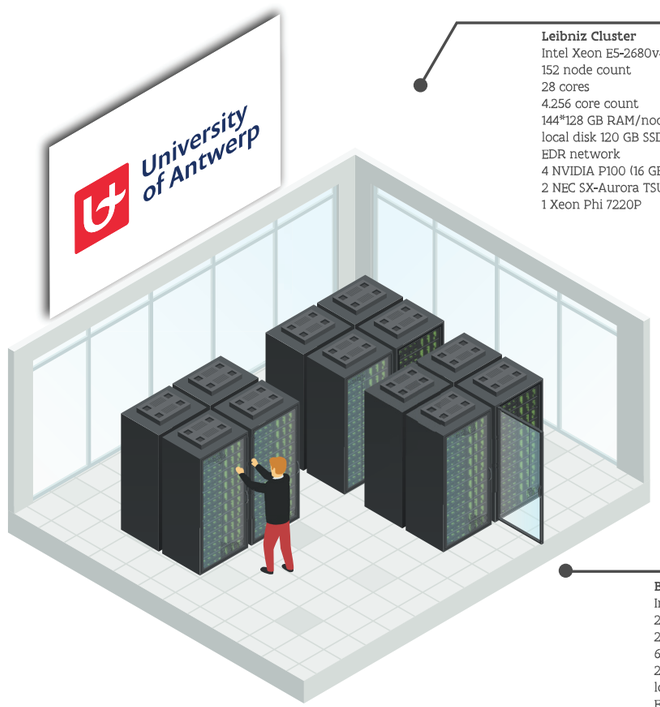

Leibniz Cluster
Intel Xeon E5-2680v4 (Broadwell)
152 node count
28 cores
4.256 core count
144*128 GB RAM/node, 8*256 GB RAM/node
local disk 120 GB SSD
EDR network
4 NVIDIA P100 (16 GB)
2 NEC SX-Aurora TSUBASA 10B
1 Xeon Phi 7220P

Vaughan Cluster
AMD Epyc 7452 (Rome) + AMD Epyc 7543 (Milan)
192 node count
64 cores
12.288 core count
176*256 GB RAM/node, 16*512 GB RAM/node
local disk 240/480 GB SSD
HDR100 network
4 NVIDIA A100 (40 GB)
4 AMD MI100 (32 GB)

Breniac Cluster
Intel Xeon Gold 6132 (Skylake)
23 node count
28 cores
644 core count
23*192 GB RAM/node
local disk 480 GB
EDR network

Figure 26. VSC UAntwerp Tier-2 Infrastructure

OPERATIONS AND USAGE

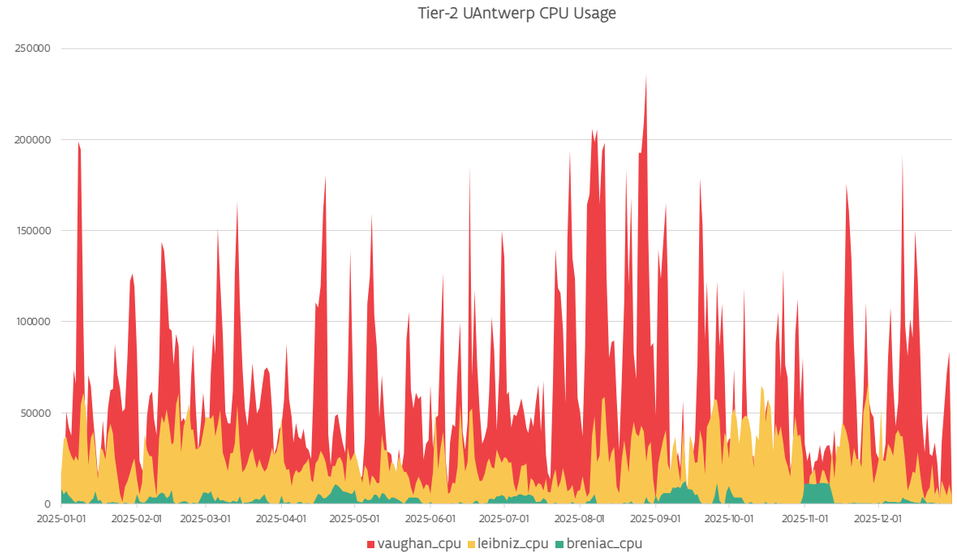

Figure 27. VSC UAntwerp Tier-2 CPU Usage

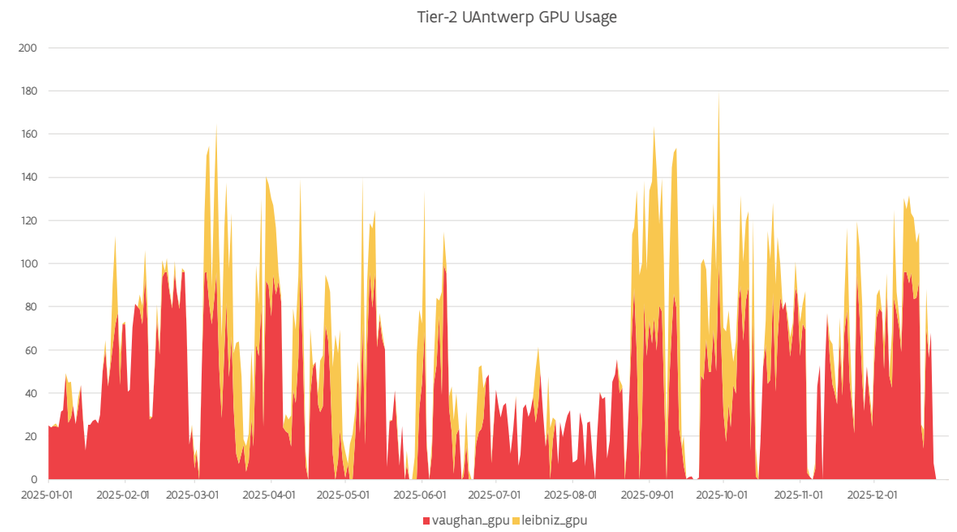

Figure 28. VSC UAntwerp Tier-2 GPU Usage

The number of interruptions in 2025, whether planned or unplanned, remained limited:

- January 16: short interruption due to issues with the scratch file system.
- April 28–29: login unavailable due to issues with the user and applications file systems.
- May 6–7: login unavailable due to issues with the user and applications file systems.
- June 9: short power outage, resulting in 97 nodes being unavailable.
- September 2: short power outage, resulting in all nodes being unavailable.

As shown in the overview above, issues occasionally occur with the storage systems. In particular, problems related to the scratch file system have proven difficult to resolve, as they are linked to the use of specific software applications. To address this, users are provided with guidance and best practices on how to use the different storage systems efficiently.

The upgrade to Rocky Linux 9 was carried out gradually between July and September 2025 following a rolling upgrade approach. Users were also given the opportunity to test the new version in advance and report any issues with their software.

Following the introduction of paid computing in 2024, a clear reduction in consumed compute time was observed that year. In 2025, a gradual increase in compute usage was observed again. At the same time, new research groups continued to find their way to the infrastructure, which is a positive development. In addition, during 2025, the migration of users from the Flanders Marine Institute (VLIZ), the Royal Belgian Institute of Natural Sciences (RBINS), and the Royal Museum for Central Africa (RMCA) from the Ghent University clusters to the University of Antwerp clusters was initiated.

To lower the barrier to entry, the Open OnDemand web portal became available in September 2025. This was also a specific request raised by users during the 2024 survey. The number of web applications currently offered through the portal is still limited but will be further expanded in 2026. The availability of the web portal has also led to requests to use it in courses within the regular academic curriculum..

ALLOCATION OF COMPUTE TIME

Since March 1, 2024, paid computing has been introduced for researchers within AUHA. Research groups can create one or more projects on which compute time can be consumed. No project proposal is required. An exception is made for students carrying out computations as part of their thesis work. Research groups can easily monitor their usage through a dedicated script.

USER SUPPORT

User support remains a central priority. This includes providing advice on optimizing workflows and computations, as well as ensuring efficient use of the infrastructure through active monitoring of jobs.
User support is primarily provided through the ticketing system. The number of submitted requests and their distribution across the different categories remained comparable to the previous year.

Category	Tickets
Accounts	96
Software	104
Tier-1	18
Other	118
Totaal	336

Table 3. Number of Helpdesk Tickets by Category in 2025

In addition, one-on-one meetings and discussions with research groups were organized to address specific questions or to (re)introduce the various VSC services. We also highlight the pitch presentations given for other core facilities within the university in January, the presentations during the Doctoral School Days, and the introduction given to the university's RDM team.

An important aspect of user support is the biannual introductory courses organized at the beginning of each semester: "Linux Introduction", "HPC@UAntwerp Introduction", and "Supercomputers for Starters", each consisting of two half-day sessions. To maintain interaction with users, the organization of lunchtime seminars on specific topics or as a supplement/refresher to the introductory courses is being considered for 2026.

In addition to the training courses offered within VSC, courses are also provided as part of the regular academic curriculum, namely "Scientific Computing Environments" and "(Parallel) Programming".

Furthermore, support is provided to master's students in using the infrastructure, as well as assistance with applications for Tier-1 and Tier-0 compute time and with the computational aspects of (inter-university) project proposals. In some cases, Tier-2 infrastructure proves to be sufficient, whereas users initially expected that a Tier-1 project would be required. It is important to identify the most suitable system for each type of project.

One of the topics covered during the introductory courses is the use of conda environments, a popular tool among Python users. Similar to the approach at Ghent University and Vrije Universiteit Brussel, the installation of conda environments directly on the storage systems is not encouraged at the University of Antwerp. Instead, users are guided to package these environments in containers.

At the beginning of November, the team also actively contributed to the VSC training course "Getting Scientific Software Installed" and promoted it among users. Following this session, a meeting was organized with users from the Department of Chemistry.

A special mention goes to a request from the Department of Philosophy in December concerning code optimization, which has since been successfully completed. The researchers ultimately found their way to the VSC team through an article discovered via a search engine. Continuing to strengthen our internal visibility therefore remains an important point of attention.

USER SURVEY

The 2024 survey identified several areas for improvement: documentation and keeping the barrier to using the infrastructure as low as possible. The documentation is kept as up to date as possible but will remain an ongoing point of attention. To further lower the barrier, the introductory courses were revised in the autumn of 2024 and, as mentioned above, Open OnDemand was also launched.

The survey conducted at the end of 2025 among all users provides, to some extent, an indication of how these efforts were evaluated. A total of 28 participants indicated that they use the Tier-2 University of Antwerp infrastructure to some extent.

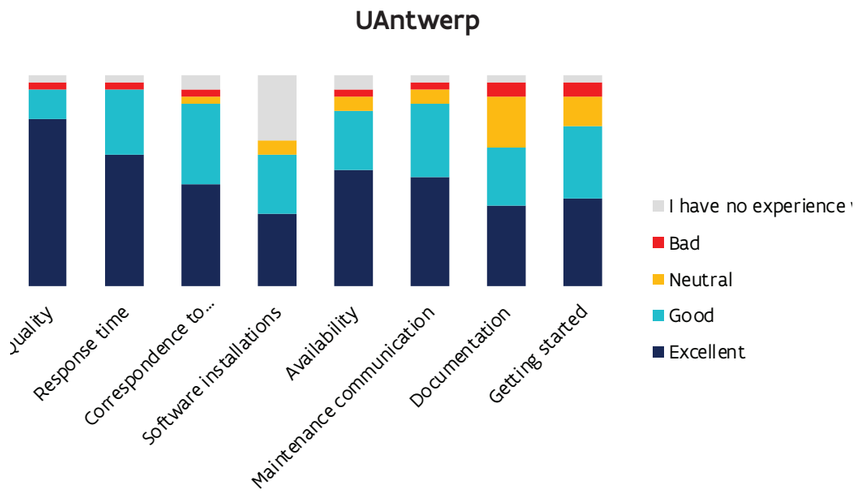

Figure 29. VSC UAntwerp Tier-2 User Survey 2025

The following aspects of the Tier-2 University of Antwerp infrastructure were rated as "Good" or "Excellent" by users:

- 96% (27/28) – Quality
- 96% (27/28) – Response Time
- 93% (25/27) – Meeting Expectations
- 90% (18/20) – Software Installations
- 89% (24/27) – Infrastructure Availability
- 90% (25/28) – Communication About Maintenance and Interruptions
- 68% (19/28) – Documentation
- 79% (22/28) – Getting Started

Compared to 2024, we observed a significant increase (10–20%) in satisfaction with software installations, infrastructure availability, and communication regarding maintenance and interruptions. There was also a slight improvement in the "Getting Started" score, providing additional motivation to engage users more actively through initiatives such as lunchtime seminars. Despite further updates and additions to the documentation, this aspect remains unchanged and therefore continues to be a permanent point of attention.

Looking at the comments from the 2025 survey, it is clear that documentation—particularly finding the required information within it—remains an area requiring attention, alongside the previously mentioned issues with the storage systems. A comment was also made regarding the fixed timing of the introductory courses (February and October). Therefore, the possibility of creating recordings is being considered so that the courses can be followed at any time, in addition to the two scheduled sessions.

There was also a request for more GPU nodes and additional local storage for bioinformatics and AI applications.

VRIJE UNIVERSITEIT BRUSSEL

AVAILABLE INFRASTRUCTURE

The Tier-2 infrastructure at Vrije Universiteit Brussel (VUB) is configured as follows:

- 3 clusters, with CPU, GPU, visualization/interactive, and large-memory partitions
- 293 TF CPU, 236 TF GPU
- 4,352 CPU cores, 32 GPU devices
- 28 TB memory

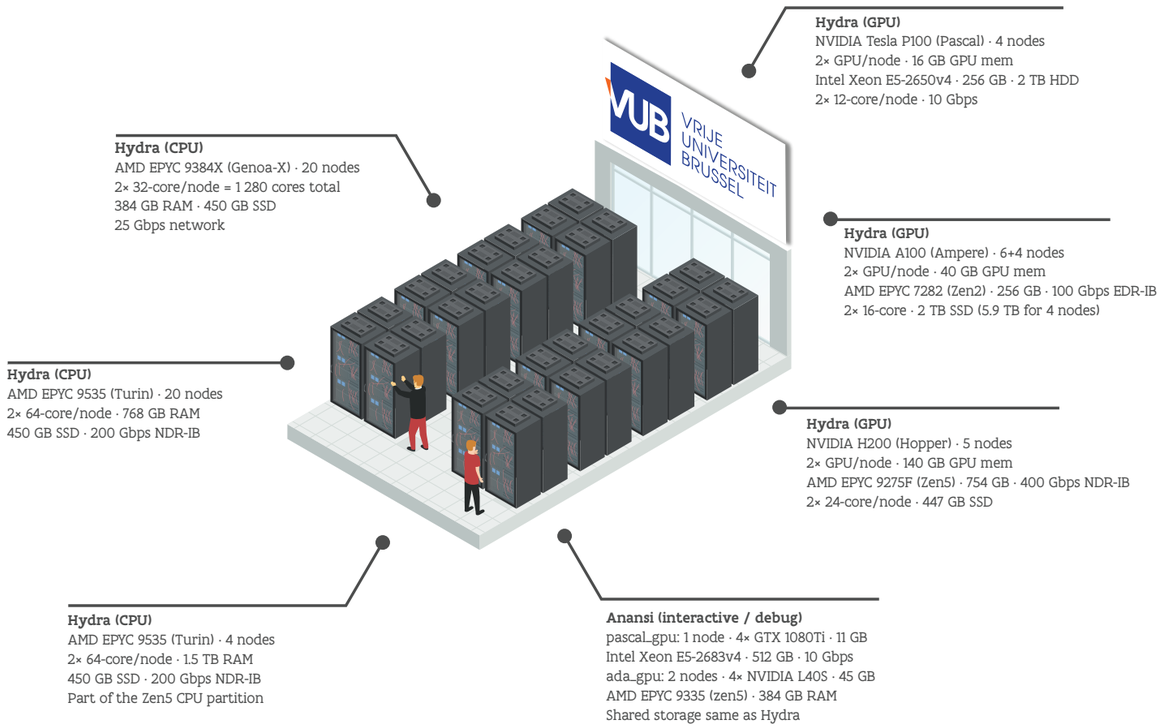

Figure 30. VSC VUB Tier-2 Infrastructure

At VUB, the decision was made to always implement expansions within the same Hydra environment, which is more efficient for both users and the management team. This does, however, result in a more heterogeneous cluster. Within the environment, efforts are made to hide this heterogeneity from users as much as possible by, among other things, providing identical software installations and assigning jobs to multiple partitions by default.

The following changes were implemented within the Hydra environment in 2025:

- New Open OnDemand web portal for HPC. This new portal has drastically lowered the entry barrier to the cluster and enables us to easily provide new (GUI) applications: https://hpc.vub.be/news/2025/vub-ondemand-release/
- New GPU nodes for interactive use on Anansi. These nodes are equipped with 4× NVIDIA L40S cards, which are ideally suited for interactive workloads: https://hpc.vub.be/news/2025/anansi-new-ada-gpu/
- New CPU nodes added to Hydra. Twenty-four new CPU nodes were added, including four CPU nodes with 1.5 TB of RAM to replace the older large-memory nodes. All these nodes are equipped with Turin CPUs (Zen 5) and an NDR InfiniBand (200 Gbps) network connection: https://hpc.vub.be/news/2025/hydra-new-turin-nodes/
- The cluster was upgraded to the next major version of the operating system, Rocky Linux 9. This operation was completed without downtime: https://hpc.vub.be/news/2025/update-rocky-9/
- Retirement of the Skylake partition. All Skylake nodes were decommissioned after serving the cluster reliably since 2018: https://hpc.vub.be/news/2025/bye-bye-skylakes/
- The login nodes were upgraded to Rocky Linux 9: https://hpc.vub.be/news/2025/hydra-login-rocky-9/
- A new command-line tool for recovering lost data was introduced. It is now easier than ever to restore accidentally deleted data: https://hpc.vub.be/news/2025/hpc-data-recovery/
- The notebook platform was decommissioned and replaced by the Open OnDemand web portal, as it provides a superior user experience: https://hpc.vub.be/news/2025/vub-ondemand-update-2024a/
- A closed pilot phase was launched for the newly acquired GPU nodes equipped with NVIDIA H200 cards, which have since entered production: https://hpc.vub.be/news/2026/hydra-new-hopper-nodes/
- Outreach efforts: throughout the year, introductory training sessions were provided to a wide range of research groups.

In addition to managing its own Tier-2 infrastructure, the IIHE team (VUB, together with ULB) also manages the grid infrastructure. This infrastructure is used, among other purposes, for processing data collected during experiments at the Large Hadron Collider (LHC) at CERN, as well as by the broader Flemish research community.

The grid infrastructure consists of:

- 1 cluster
- 218 TF
- 6,304 CPU cores
- 51 TB memory

In 2025, the capacity of the grid cluster remained unchanged.

OPERATIONS AND USAGE

In summary, 2025 can be characterized as follows:

- A total of 3.2 centuries of single-core CPU compute time were consumed.
- In addition, 21.4 years of GPU compute time were utilized.
- A total of 963,894 jobs were executed, the vast majority of which had a runtime of less than one hour.
- Approximately 30% of the jobs accounted for 99% of the consumed CPU compute time and 97% of the GPU compute time.

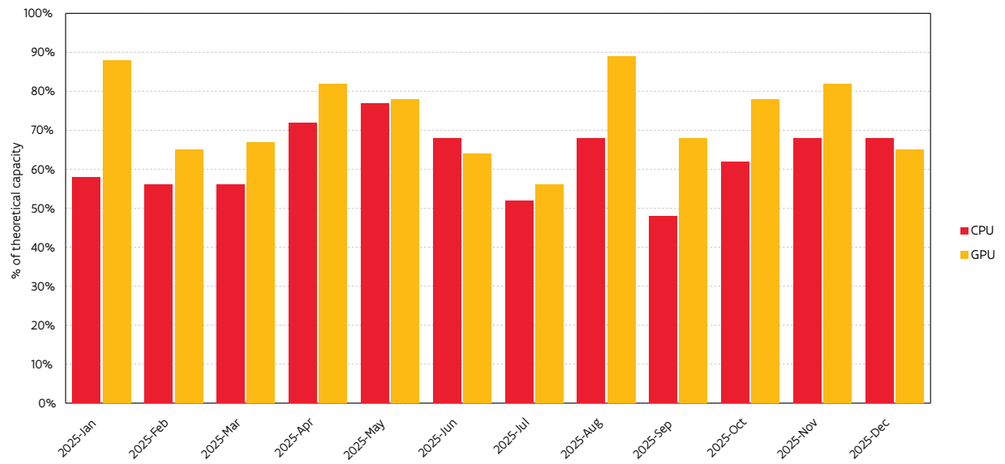

Figure 31. VSC VUB Tier-2 Usage

The average utilization in 2025 was 63% for the CPU partition and 74% for the GPU partition.

- In 2025, 369 new VUB VSC accounts were created.
- During 2025, there were 562 unique users of the VUB HPC clusters (i.e., users who submitted at least one job), compared to 383 in 2024.
- There were 513 unique users of the CPU nodes and 208 unique users of the GPU nodes.

The distribution of users by category is shown in the table below. For both CPU and GPU resources, the top 10% of users accounted for approximately 80% of the total usage.

	CPU Time	GPU Time	Share of Users
UZB	0.0%	0.0%	0.4%
Guest professor	2.9%	6.1%	1.2%
Professor	0.3%	0.1%	2.5%
Administrative/Technical Staff	2.3%	1.4%	2.1%
Non-VUB	0.9%	1.5%	8.2%
Student	13.6%	12.3%	30.6%
Researcher	80.1%	78.6%	55.0%

Table 4. VSC VUB Tier-2 Users by Category in 2025

The breakdown of the number of active users per month and by user type is shown in the figure below. The high number of students during the last two months is due to the use of the infrastructure for several courses.

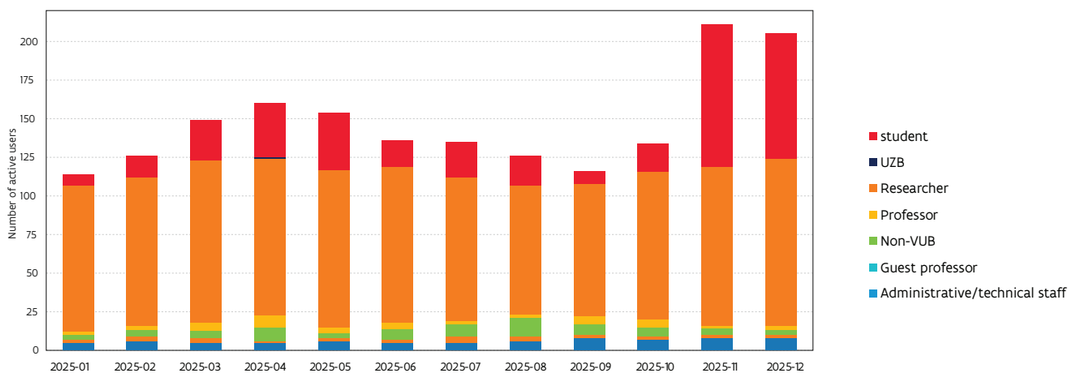

Figure 32. VSC VUB Tier-2 – Active Users on CPU Nodes

A breakdown by faculty is shown in the table below.

Faculty	CPU Time	GPU Time
Faculty of Social Sciences and Solvay Business School	0.4%	20.4%
Directorate ICT	2.3%	1.4%
Non-Vrije Universiteit Brussel	0.9%	1.5%
Faculty of Medicine and Pharmacy	2.1%	4.4%
Students	13.6%	12.3%
Faculty of Engineering Sciences	36.9%	24.5%
Faculty of Sciences and Bio-engineering Sciences	43.8%	35.3%

Table 5. VSC VUB Tier-2 CPU and GPU Usage by Faculty

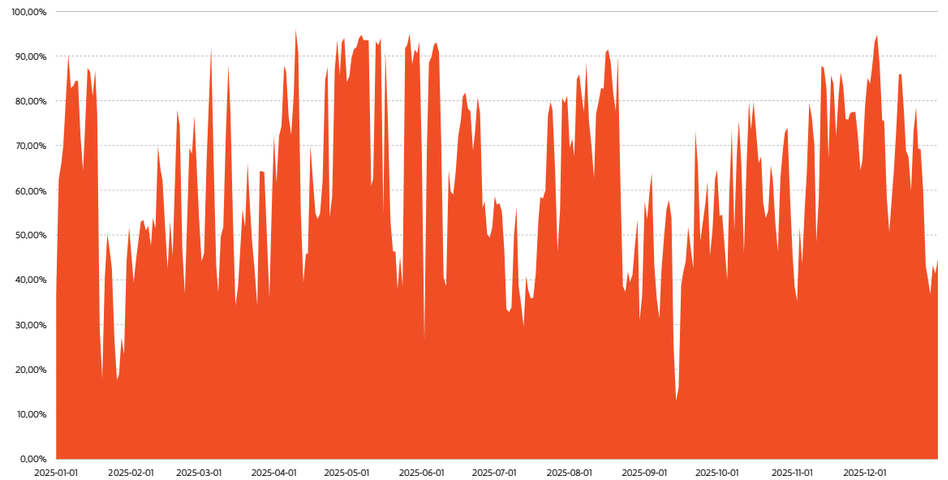

Figure 33. VSC Tier-2 VUB CPU Usage

For CPU usage, we observe the same trend as in previous years: a rather irregular usage pattern. This corresponds to a large number of relatively short jobs running on the system. Especially during weekends, system utilization is sometimes very low. This type of workload appears to be the dominant usage pattern at VUB.

Looking at the statistics (considering only jobs with a runtime of at least 30 minutes), 70% of the total compute time was consumed by jobs running on a single node, while only 25% was consumed by jobs using two or more nodes. Approximately 90% of the jobs ran for 20 hours or less, with 50% running for 5 hours or less. This corresponds to the irregular pattern shown in the graph.

Compared to 2024, jobs are running for longer periods. The number of single-node jobs has increased, while the size of a single node has also increased with the introduction of the Zen 5 nodes. Consequently, the waiting time for jobs to start is relatively low:

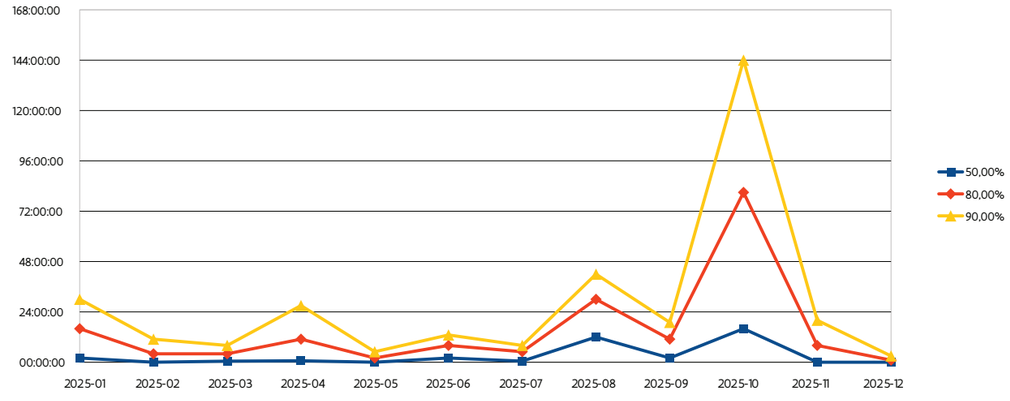

Figure 34. VSC VUB Tier-2 Monthly Wait Time for CPU Jobs

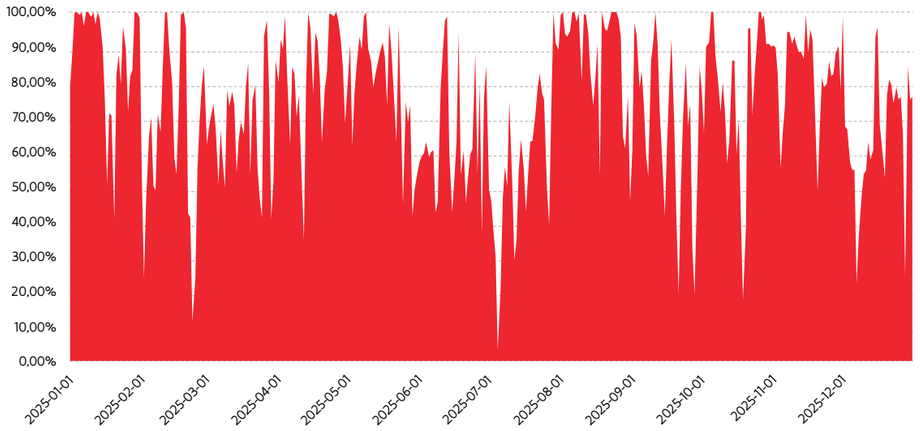

Figure 35. VSC VUB Tier-2 – Active Users on GPU Nodes

	75th percentile	90th percentile
Single CPU core jobs	00h17	05h31
Single node jobs	01h31	10h49
Multi node jobs	03h54	28h08

Table 6. Wait Time by Job Type (Percentiles)

The GPU partition, as in previous years, remained heavily utilized. There were several periods during which the utilization reached approximately 100%. Of all GPU jobs, 50% started within 2.5 hours and 90% within 27.5 hours. Approximately 80% of the jobs had a runtime of less than 8 hours. The vast majority of jobs (95%) used only one GPU, accounting for 78% of the total GPU compute time. Compared to 2024, we observe a shift toward larger GPU jobs.

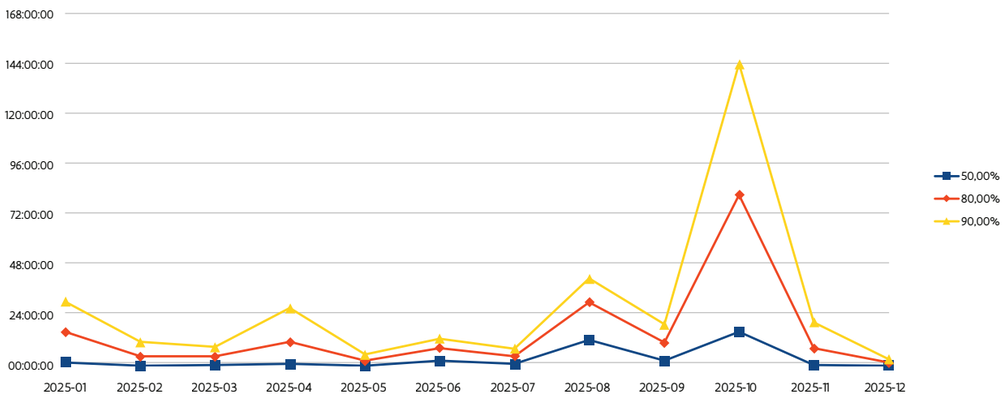

Figure 36. VSC VUB Tier-2 Monthly Wait Time for GPU Jobs

The GPU usage broken down by the number of active users is shown in the figure below.

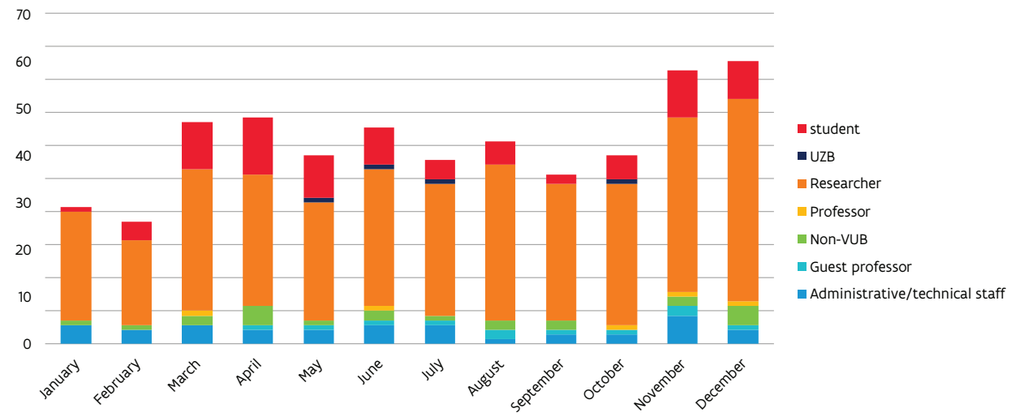

Figure 37. VSC VUB Tier-2 GPU Users per Month

In 2025, a total of 3,155 modules were installed across the various partitions, resulting in 1,317 unique software modules available to users, including 3,844 unique extensions.

More details about the usage of the Hydra environment can be found at: https://hpc.vub.be/news/2026/overview-of-2025/

There were no major unplanned interruptions in 2025. On April 14, maintenance was performed on the VUB firewall, temporarily making the cluster inaccessible, although all running jobs continued without interruption. Between September 12 and 16, access to the cluster was disabled for security reasons. A serious vulnerability had been discovered in the Linux kernel that could potentially be exploited. Until a patch had been applied to all components of the cluster, users were unable to log in, but all running jobs continued uninterrupted. No evidence was found afterward that the vulnerability had been exploited on the cluster.

In addition, several system upgrades were carried out. These upgrades were performed in phases to ensure that part of the cluster remained operational at all times. Such upgrades took place during the week of March 17, the week of July 17, and on November 24.

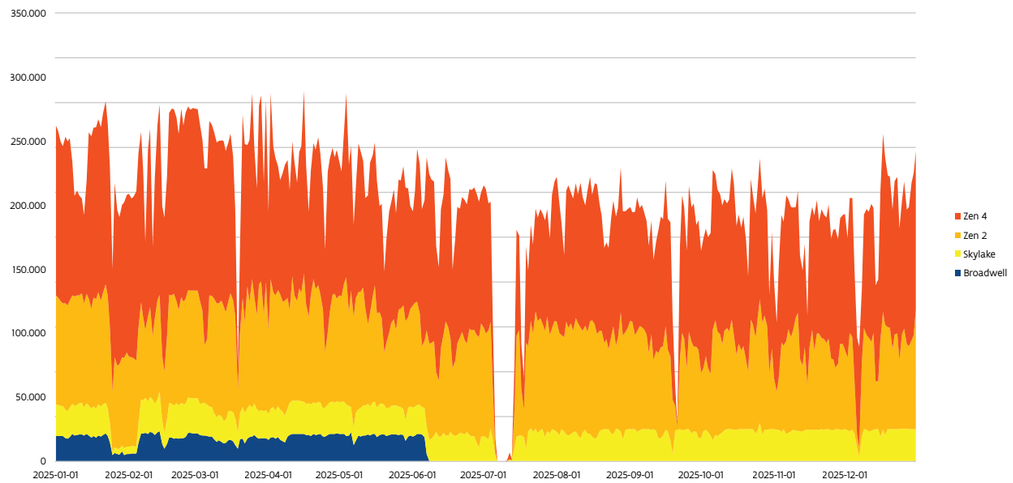

Figure 38. VSC VUB Tier-2 Grid Usage

The grid cluster is heterogeneous: it consists of a series of worker nodes equipped with different generations of Intel and AMD CPUs (with the most recent generation being Zen 4). The compute cluster is primarily used for data-intensive, single-core calculations (so-called High Throughput Computing), and its size is expressed in job slots, corresponding to one hyper-threaded CPU core.

The IIHE team also manages a Ceph storage system (938 TiB) in which users' home directories are stored and which also serves as the backend for the internal OpenNebula cloud. A replication level of 3 was chosen to ensure redundancy. In 2025, the capacity of the compute cluster remained unchanged at 218 TFLOPS, as did the capacity of the mass storage system (dCache), which remained at 11.2 PB.

One of the major achievements of 2025 was the relocation of the entire infrastructure from the Solbosch Campus to the brand-new Nexus data center in Zellik. The migration took place over one week, from July 6 to July 13, and involved the entire IIHE IT team (six staff members).

Most researchers use the so-called "glide-in" mechanism. In this approach, "pilot jobs" are submitted which, once active on a worker node, retrieve the actual "payload" from another location. For the calculation of the percentages mentioned above, only the compute time of the "pilot jobs" was considered, not that of the individual "payloads". It should also be noted that, within the grid environment, a user's workflow can be distributed across multiple sites in different countries, meaning that the percentages above represent only a part of the actual total compute time used.

Another major technical achievement was the implementation of a solution enabling CMS glide-in jobs to run on the VUB Tier-2 cluster (Hydra) through a dedicated compute element that routes the jobs and forwards them via a translation layer (BLAH) to the local Slurm batch systems.

In 2025, there were two planned maintenance periods on the grid infrastructure:

- On March 19, the dCache Storage Element (SE) was migrated to a new machine, and the entire dCache infrastructure was upgraded.
- During the week from July 6 to July 13, the complete cluster was relocated to the Nexus data center.

In addition, there were several unplanned interruptions:

- Multiple issues with the storage system (dCache pools failures), resulting in several worker nodes becoming unavailable: January 1, 2025 – January 31, 2025.
- Failure of two disks in a storage pool and an issue during the rebuild process: February 19, 2025 – February 21, 2025.
- HTCondor scheduler schedd02 (local jobs) overloaded: February 25, 2025.
- Network issues affecting the connection to Belnet: July 17, 2025.
- Restart of worker nodes in groups due to a kernel upgrade: September 19, 2025 – September 22, 2025.

ALLOCATION OF COMPUTE TIME

Researchers at VUB and its association have unrestricted access to their own Tier-2 infrastructure. The grid cluster is available upon request to the person responsible for this infrastructure. The use of the Tier-2 and grid infrastructure is free of charge.

USER SUPPORT

VUB has a dedicated contact point for all HPC-related questions via (hpc@vub.be). User requests are distributed across the different services offered.

Last year, a total of 999 tickets for Tier-2 (compared to 814 in 2024) and 249 tickets for the grid infrastructure (compared to 208 in 2024) were processed, representing an increase in both cases. For Tier-2, the distribution across the different services is shown in the figure below.

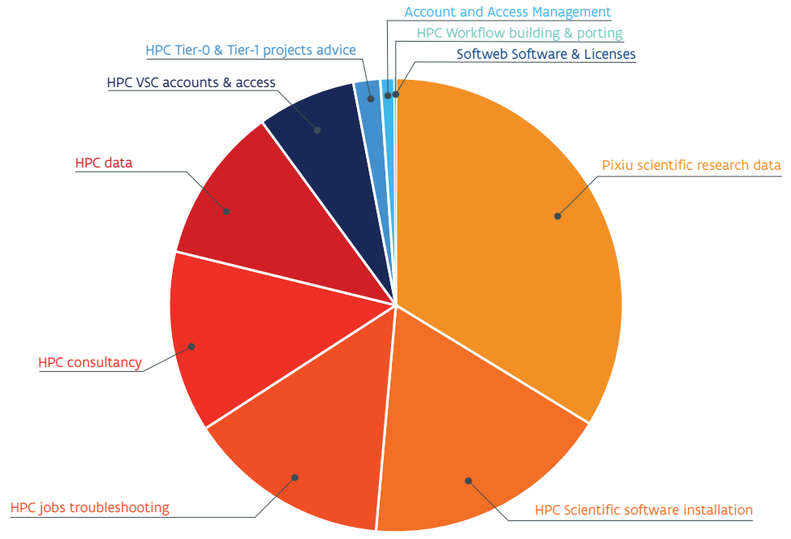

Figure 39. VSC VUB Tier-2 Tickets by Service

For the grid infrastructure, the distribution is as follows:

Accounts	86
Software	32
Other	131

Table 7. VSC VUB Tier-2 Grid Ticket Distribution by Service

For the Tier-2 tickets that were closed in 2025, the resolution time in working hours is shown below:

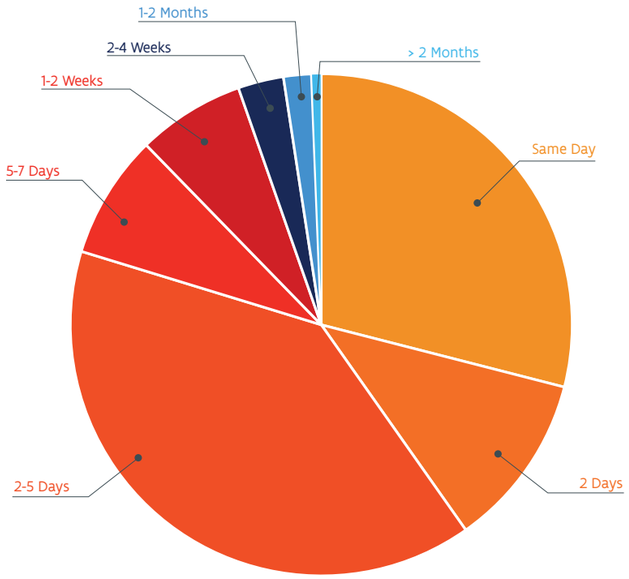

Figure 40. VSC VUB Tier-2 Ticket Resolution Time 2025

SPECIFIC SUPPORT

The courses “Introduction to Linux” and “Introduction to the Use of HPC at VUB” are organized twice per year.

In addition, several meetings were held with research groups to introduce them to HPC and VSC in general, and to encourage the transition from their own local machines to shared HPC infrastructure.

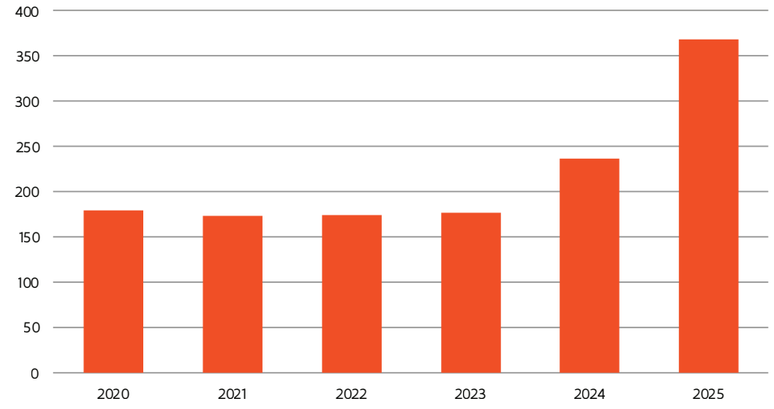

Figure 41. VSC VUB Tier-2 Users per Year

Different sessions with individual researchers were also organized. Particularly during the autumn, there were intensive discussions with several groups working on AI. The first results are now becoming visible: over the past two years, we have seen a strong increase in the number of requested accounts, and the number of active users is also steadily growing. In particular, during the last two months of 2025, a clear increase was observed in the number of users utilizing the GPU nodes. This has not yet translated into more applications for Tier-1 resources, and we continue to actively promote this.

USER SURVEY

In the annual VSC User Survey, 76 participants indicated that they use the VUB Tier-2 infrastructure. Overall, we observe results similar to those of the previous year.

Overall, we are very satisfied with these results, as they demonstrate that our services are highly valued by our users. Compared to 2024, we observe a slight decrease in most categories, while the score for “Meeting Expectations” has improved. We are especially pleased that the score for “Documentation” has improved.

We note that 42% of respondents selected “No Experience” for the Software Installations category. This suggests that users are satisfied with the software available on the VUB clusters or prefer to install their own software. The training on software installation has likely contributed to this.

“Getting Started” with Hydra still receives the lowest score, which is expected given the significant learning curve. In response, we organized additional training sessions for specific user groups and launched the VUB Open OnDemand web portal, which significantly lowers the barrier to getting started. However, these efforts have not yet resulted in an improved “Getting Started” score. We will therefore continue to invest in training opportunities and further explore ways to improve the user experience.

The score for Availability decreased by 6%. In 2025, there was only one unplanned downtime of Hydra to address a potential security issue. All other updates were carried out without requiring a complete downtime. The lower score in this category is therefore likely more related to waiting times and the availability of resources.

The satisfaction with response time decreased by 9%. The helpdesk experienced several periods with a higher number of users, resulting in a backlog in the handling of support tickets. Although this backlog has since been resolved, it led to longer response times for support requests.

As in 2024, the demand for GPUs remains high, and waiting times can sometimes exceed acceptable limits. We continue to actively encourage intensive GPU users to transition to Tier-1 resources. Investments were made in new GPU nodes, resulting in the recently announced hopper_GPU partition on Hydra. Additional investments were also made in the Anansi cluster for interactive use, and users are encouraged to select GPU shards in the ada_gpu partition.

Documentation remains a continuous work in progress. We regularly update and improve our documentation and continue to integrate VUB-specific material into the VSC documentation. Users are encouraged to inform us whenever documentation is unclear, incomplete, or contains errors.

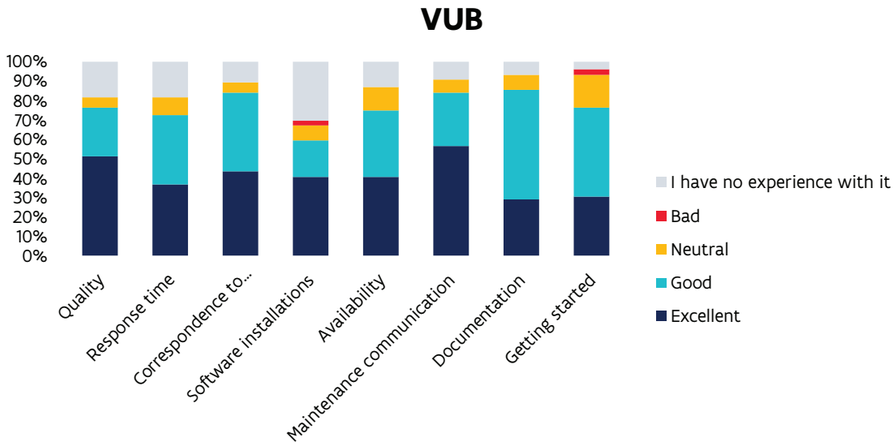

Figure 42. VSC VUB Tier-2 User Satisfaction

What Our Users Say

"The ICT helpdesk people are very responsive."

"Amazing work and support"

"I have only had excellent interactions with the HPC staff at the VUB! They have helped set up my software on the HPC, and even helped me debugging some specific problems."

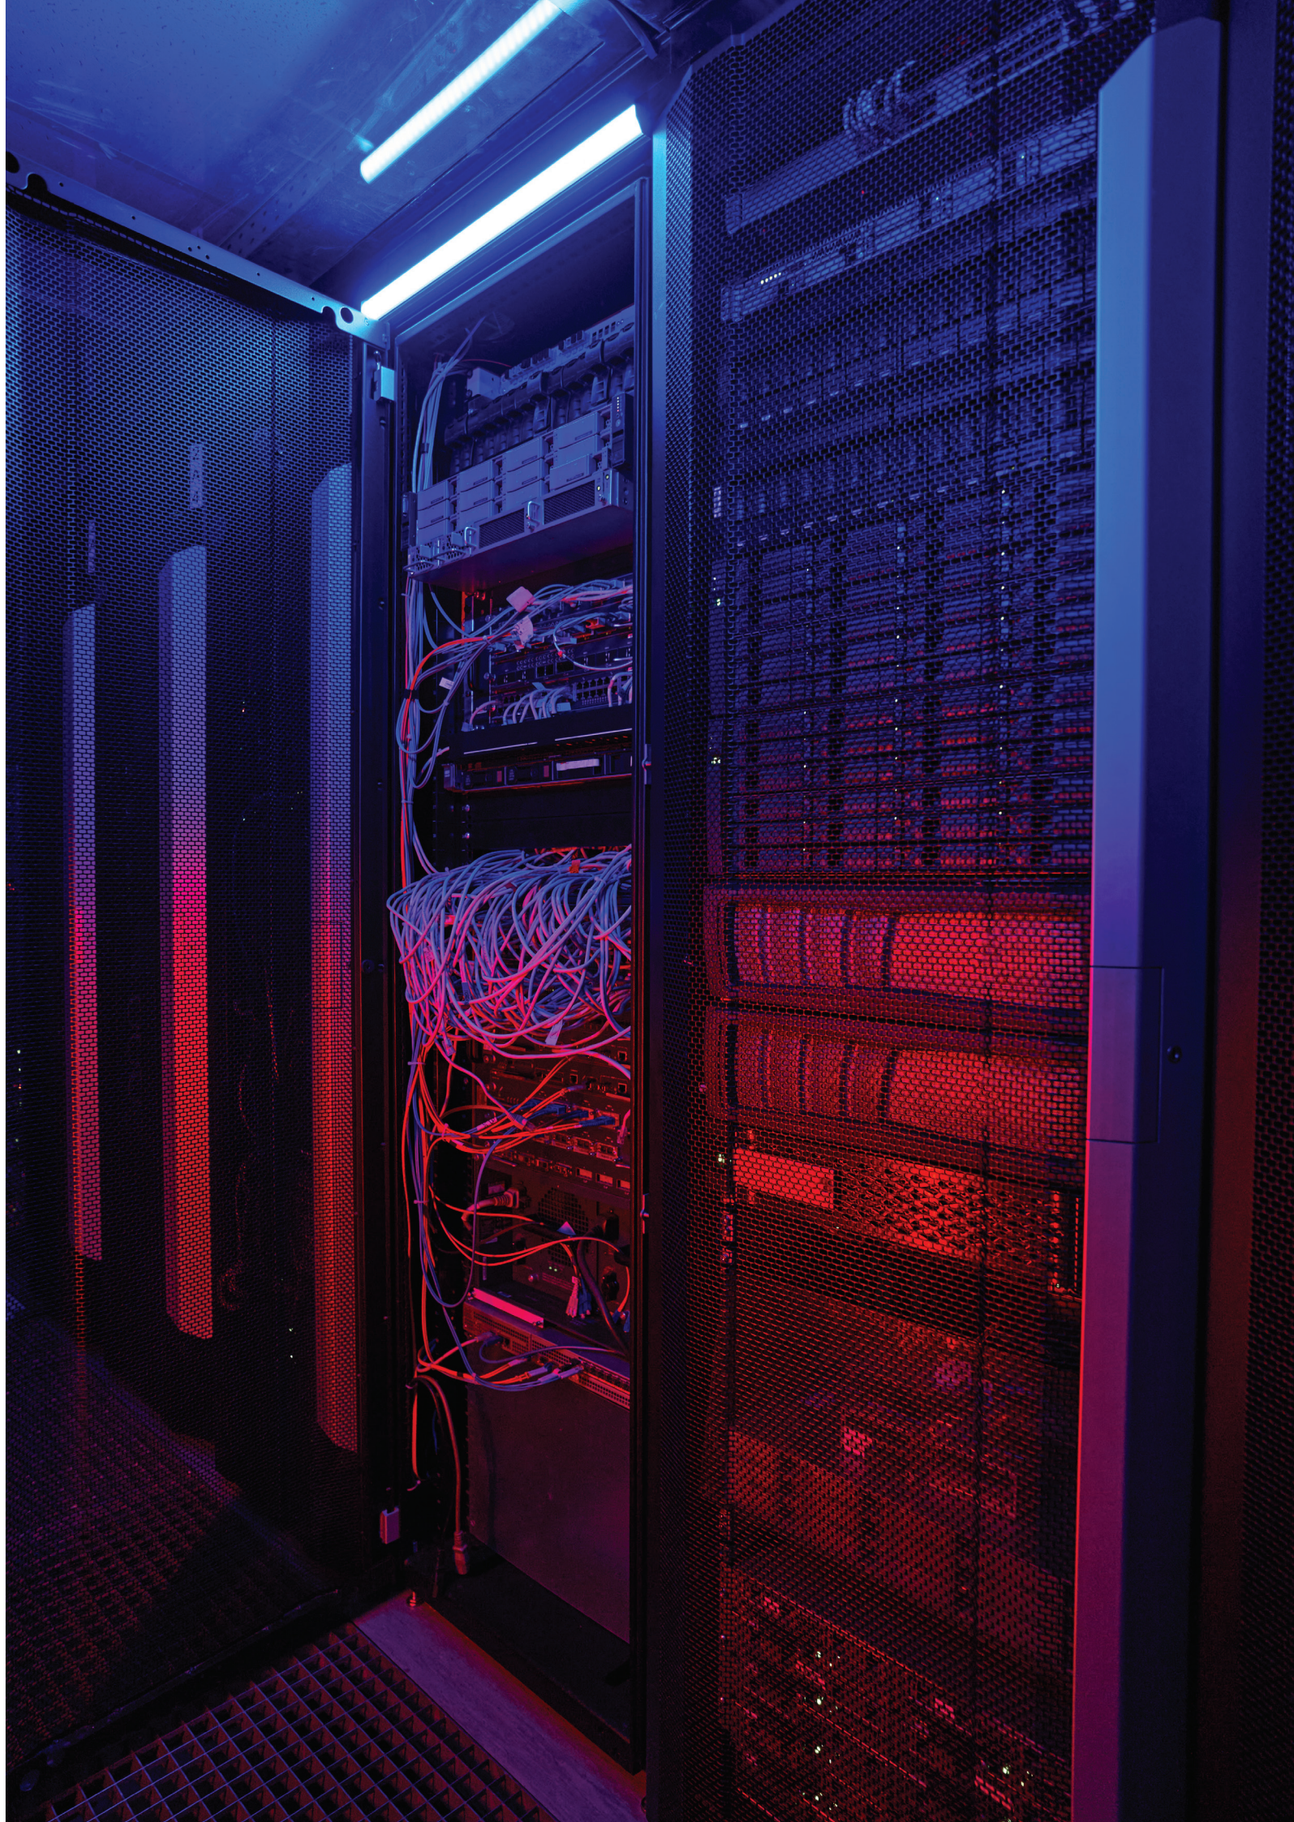

GHENT UNIVERSITY

AVAILABLE INFRASTRUCTURE

The Tier-2 infrastructure at Ghent University is composed of several clusters, each with specific characteristics:

- 8 clusters in production
- 4 CPU clusters: doduo, gallade, shinx, skitty
- 3 GPU clusters: joltik, accelgor, litleo
- 1 interactive debugging cluster: donphan
- A total of 25,264 CPU cores and 108 NVIDIA GPU accelerators.

During 2025, the following infrastructure and service changes were implemented:

- The skitty cluster was decommissioned.
- All clusters were migrated to the RHEL 9 operating system.
- Several network modifications were carried out in data center S10, impacting the entire Tier-2 infrastructure.
- The new GPU cluster litleo was introduced.

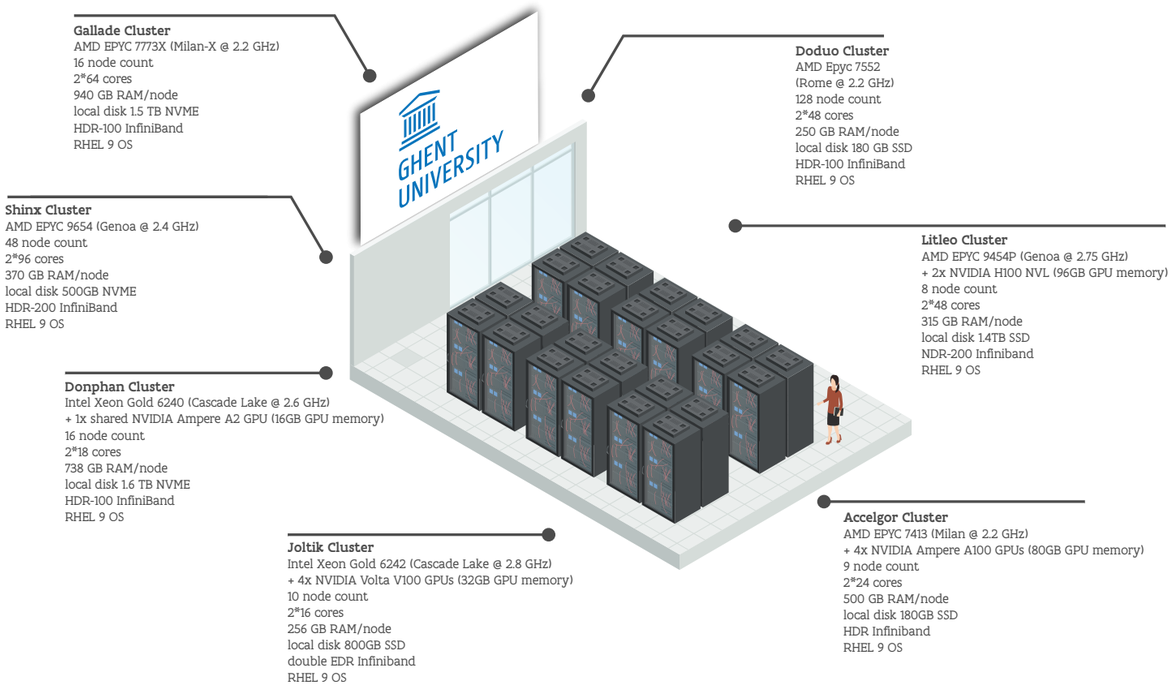

Figure 43. VSC UGent Tier-2 Infrastructure

USAGE AND AVAILABILITY

The figures below provide an overview of the day-to-day consumption of compute time on the different UGent Tier-2 clusters, separated into CPU usage and GPU usage. In total, 140,908,112 CPU hours of compute time were consumed in 2025. In addition, a total of 424,818 GPU hours were used.

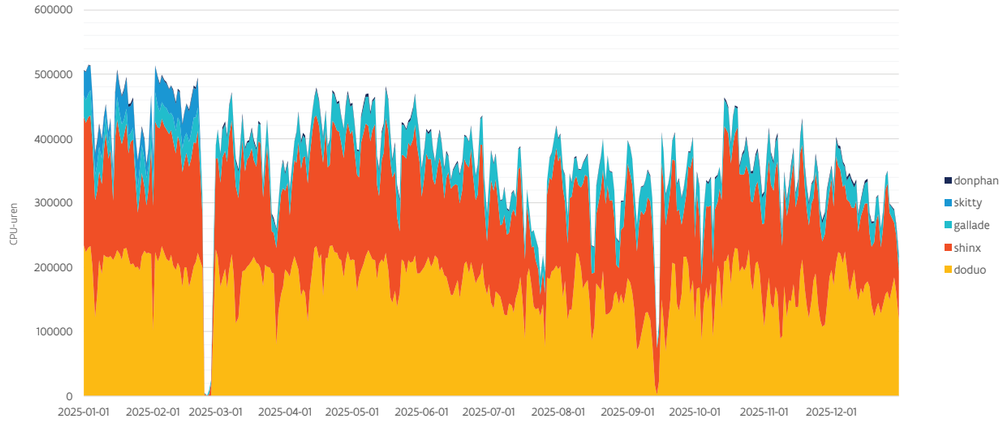

Figure 44. VSC UGent Tier-2 CPU Usage

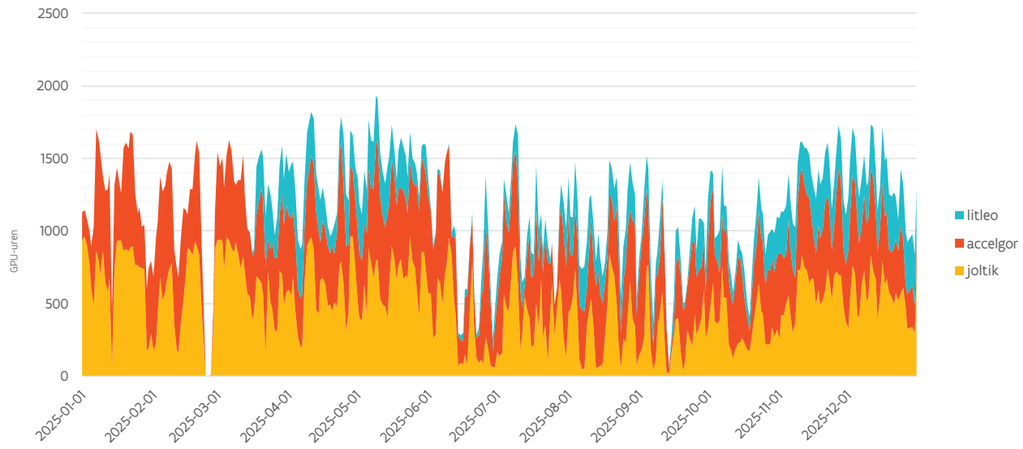

Figure 45. VSC UGent Tier-2 GPU Usage

There were several periods of unavailability of the Tier-2 infrastructure in 2025:

- February 24–28: scheduled maintenance for network modifications.
- July 11: HPC clusters were unavailable due to a scheduled data center test.
- September 12: urgent operating system update on all clusters.

A total of 2,887 users were active on the UGent Tier-2 infrastructure in 2025, representing an increase of more than 27% compared to the previous year.

As agreed upon during the establishment of the VSC, Tier-2 clusters can also be used by other institutions. The table below provides an overview of the number of users and their share of the total compute time consumed in 2025. The breakdown is made according to institutions that are consortium partners within VSC, other research institutions (e.g., VLIZ, RBINS, etc.), and companies.

Share of Tier-2 UGent Usage in 2025			
Institution	**# Users**	**CPU Usage**	**GPU Usage**
UAntwerp	34	2.91%	5.55%
Vrije Universiteit Brussel	16	0.00%	0.11%
UGent	2,806	96.25%	90.04%
KU Leuven/UHasselt	41	0.24%	3.72%
Other institutions and companies: RBINS, RMCA, VLIZ, INBO, UZGent, HOWest, KMI, HOGent, VIB, ILVO, VKI, Industry	81	0.60%	0.58%

Table 8. VSC UGent Tier-2 GPU Usage

The Tier-2 infrastructure is also used in the context of education. During the 2024–2025 academic year, the HPC infrastructure was structurally used within 35 courses. In 2025, 1,477 students were active, accounting for a total of 6% of all CPU compute time and 19% of all GPU compute time used. This includes all usage by Master's students, for example in preparation for their thesis work.

HELPDESK

User support for the UGent Tier-2 infrastructure is distributed across several helpdesk queues:

Queue	# Handled and Closed Tickets
Helpdesk Tier-2	3,004
Software installations	278

Table 9. VSC UGent Tier-2 User Support – Ticket Overview

Through the dedicated Tier-2 helpdesk (hpc@ugent.be), users can report issues or ask questions related to the infrastructure. The figure below shows the distribution of the time required to resolve each submitted request. Almost 90% of all requests were resolved within 24 hours.

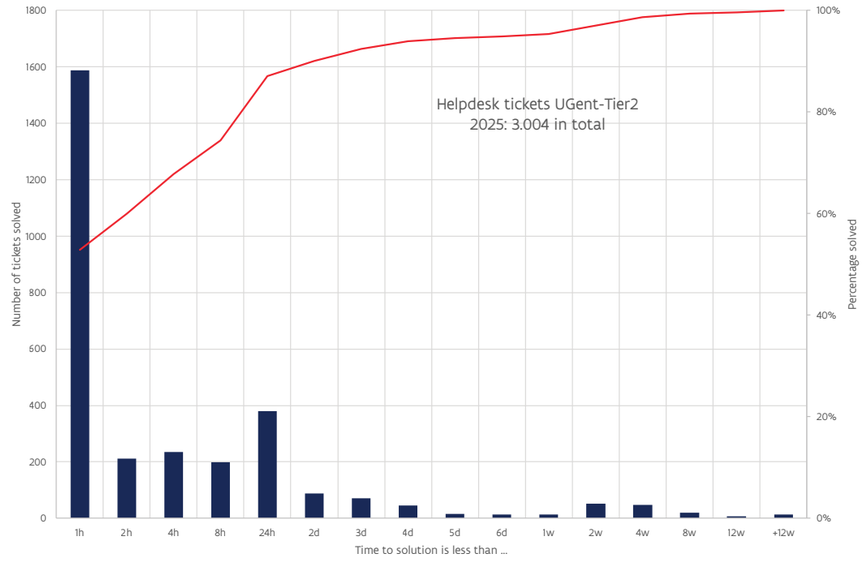

Figure 46. VSC UGent Tier-2 Helpdesk Tickets

Given the diversity of users of the Tier-2 infrastructure, it is interesting to observe that the support requests reflect a similar distribution. Support requests related to Tier-2 UGent are received from a wide range of institutions and companies.

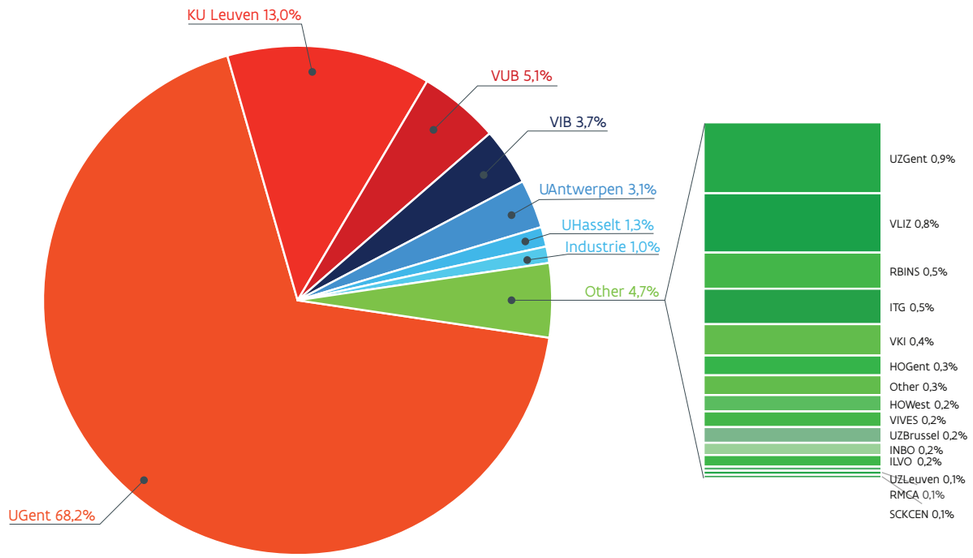

Figure 47. VSC UGent Tier-2 Helpdesk Requests by Institution

SOFTWARE INSTALLATIONS

For several years, a recurring concern raised in the user survey has been the need to reduce the turnaround time for software installations. To gain a deeper understanding of this issue, a comprehensive analysis was conducted in 2025.

In 2025, 278 support requests related to software installations were handled through the HPC-UGent software installation queue. Through this queue, as well as through collaboration with the VSC RDI Office, a total of 267 individual software installation requests were processed in 2025.

These requests included work for both the Tier-2 UGent infrastructure and the Tier-1 Compute platform. Compared to previous years, this represents a return to a more manageable number of approximately 200–250 software installation requests per year.

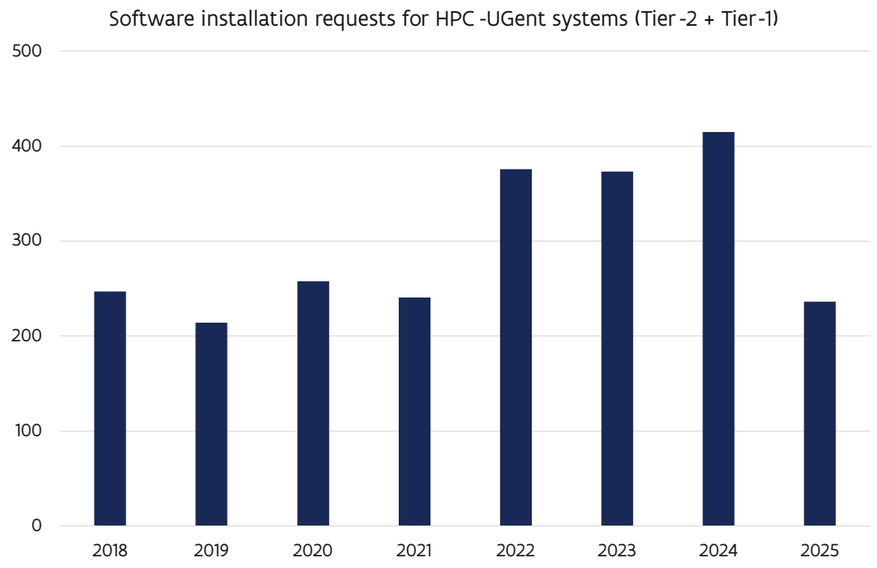

Figure 48. VSC UGent Software Installations

The main reason for this reduction is a more clearly defined policy regarding software installation requests, with several key points of attention:

- Requests are submitted through a form, which since 2025 requires more detailed information from the applicant. (https://www.ugent.be/hpc/en/support/software-installation-request).
- Each applicant is expected to write a report on the actual usage of the installed software.
- The high institutional cost of completed software installations is explicitly communicated to the applicant in euros. This cost is not charged to the applicant but is provided purely to raise awareness..
- Applicants are clearly informed in advance about the relatively long waiting time required to perform software installations.
- At the same time, users are informed about the possibility of installing simpler software themselves, for example through the vsc-venv tool, which enables users to set up a Python virtual environment (https://docs.hpc.ugent.be/setting_up_python_virtual_environments).

For the majority of received software installation requests (197), the software was installed on the Tier-2 systems of UGent. A total of 70 installations were performed on Tier-1 Hortense, either at the request of academic researchers (26) or industrial users (44).

Target System Software Installation	# Completed Requests	Share
Tier-2 UGent	197	74%
Tier-1 Academic Use	26	16%
Tier-1 Industry Use	44	10%
	267	

Table 10. VSC UGent Software Installations – Distribution

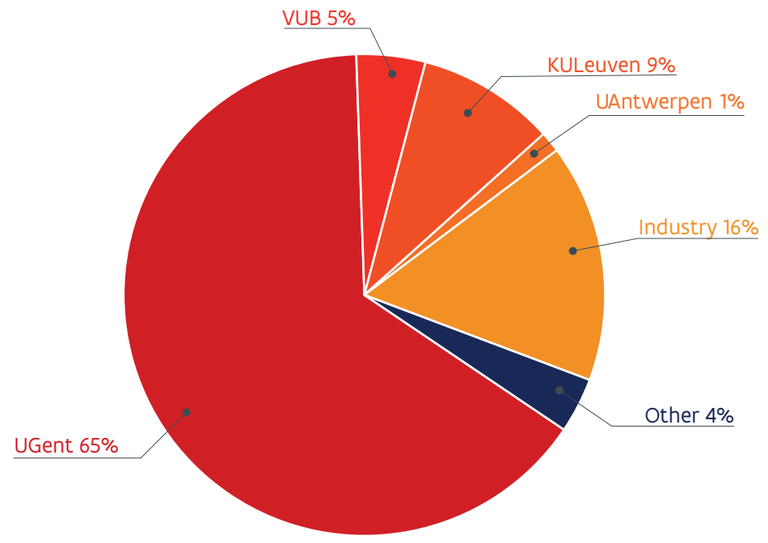

Figure 49. VSC UGent Software Installations by Institution

Somewhat surprisingly, only 23% of all software installation requests (Tier-1 Hortense and Tier-2 UGent combined) concern software that can run exclusively on GPU nodes. CPU-only software remains by far the most common, accounting for 77% of all installation requests. However, the dependencies associated with CPU-only software are considerably more extensive: installing a single CPU software package typically requires an additional seven software packages and libraries to satisfy its dependencies. This naturally has a significant impact on the complexity of the installation task and the required effort, which often increases sevenfold.

For GPU-related requests, this is much less of an issue, as the ecosystem of libraries and supporting software for GPUs appears to be less diverse. On average, only one software package is installed per GPU-related request.

It is therefore not surprising that the actual number of installed software packages is many times higher than the number of incoming installation requests. In 2025, an average of 1,800 software packages were installed for each CPU type within the UGent Tier-2 clusters, of which 50 were specific to each Tier-2 GPU cluster.

The amount of effort required to process and perform these software installations should not be underestimated. The objective of such installations is to make software centrally available to all users of the HPC systems. This centralized and broadly accessible approach generally requires more effort than, for example, a software installation performed by a user within their own account.

For nearly 70% of the requests, the work was completed within one working day. The remaining 30% of requests required more time. For 16% of the requests, more than one full working week was needed to complete the installation, and in complex cases, the effort exceeded one month of work.

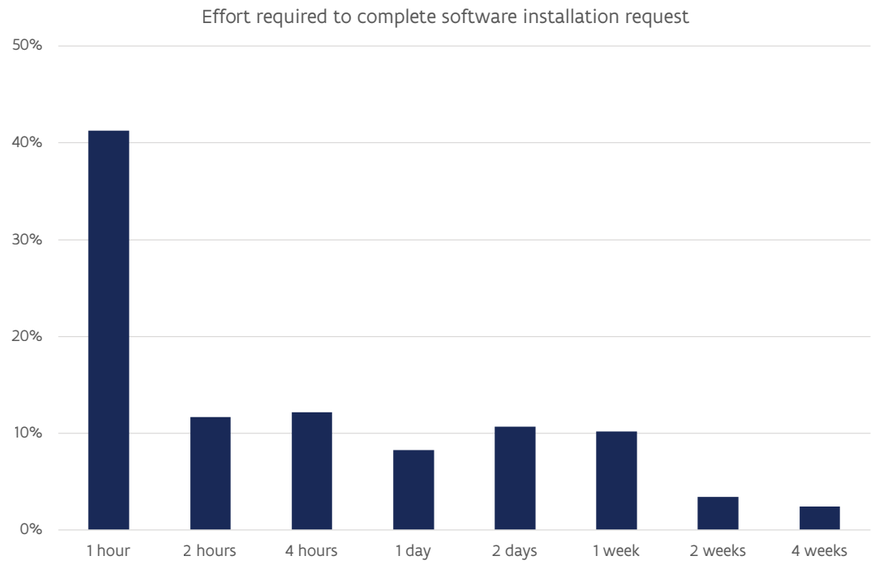

Figure 50. VSC UGent Software Installations – Resolution Time

These figures may seem relatively small at first glance, but due to:

- the high number of software installation requests (on average more than one per working day),
- the need to first install dependent libraries and tools,
- the significant time investment required for some installations, ranging from one week to one month of effort,
- the limited number of staff members responsible for performing software installations,
- a certain overhead per request due to back-and-forth communication,

this ultimately results in a waiting time for the applicant.

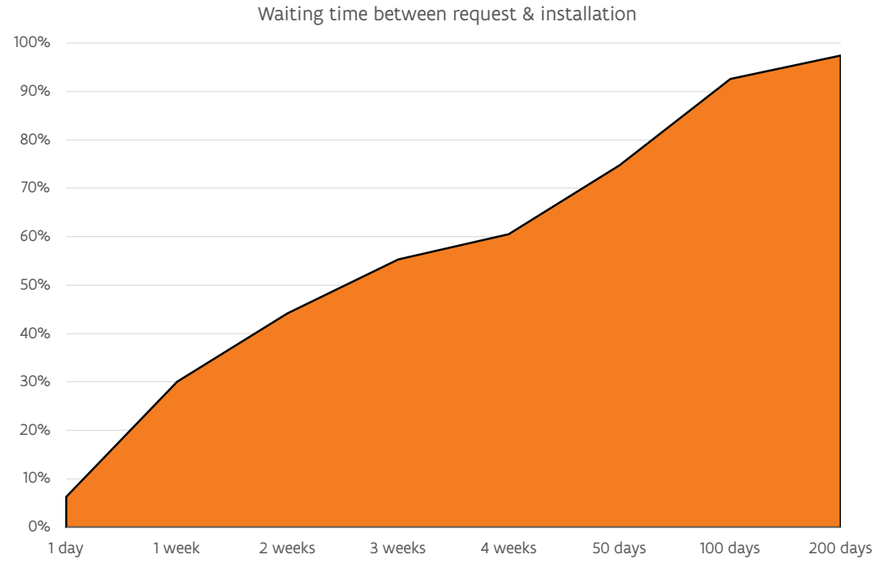

Figure 51. VSC UGent Software Installations – Wait Time

As shown in the figure above, the complete software installation process in 2025 was completed:

- within one week for 30% of the requests,
- within one month for 60% of the requests..

USER SURVEY

In the VSC User Survey conducted at the end of 2025, user satisfaction with the Tier-2 UGent infrastructure and services was also assessed. A total of 134 respondents indicated that they make use of the infrastructure to some extent.

33% of the respondents stated that the Tier-2 UGent infrastructure is crucial for their research or work, while more than 51% of respondents rated the infrastructure as "very helpful"..

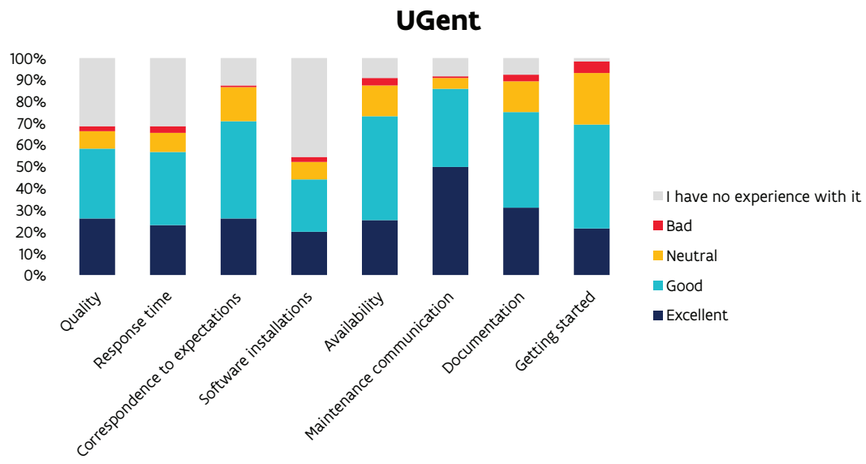

Figure 52. VSC UGent User Satisfaction

The following aspects of the Tier-2 UGent infrastructure were rated as "Good" or "Excellent" by users:

- 85% – Quality
- 83% – Response Time
- 81% – Meeting Expectations
- 81% – Software Installations
- 80% – Availability
- 94% – Communication Regarding Maintenance
- 81% – Documentation
- 70% – Getting Started

The complete list of suggestions for improvements, as provided by respondents in the 2025 User Survey, is as follows:

Infrastructure

- More CPU and GPU nodes to reduce waiting times (4)
- More memory per core (1)

Documentation

- More application-specific documentation (2)
- Updates to the existing documentation (1)

Software

- Faster response to software installation requests (2)
- The possibility to install software independently in a virtual environment (1)

User Experience

- Faster response to helpdesk requests (3)
- Higher maximum walltime, among others for AI workflows (2)
- Improved stability of interactive sessions (1)
- A more user-friendly and intuitive Whisper AI interface (1)
- Support for Windows or a sudo environment (1)

Data

- A dedicated platform for sharing and analyzing sensitive data (2)
- A shared central model directory for LLMs (1)
- Higher default storage quotas (1)

Training

- Training sessions on specific scientific software (1)

What Our Users Say

"You're offering an excellent service! I'm really grateful that I can use it."

"I greatly appreciate the HPC trainings that are offered by the University

"I am very happy with the service and support!"

KU LEUVEN/HASSELT UNIVERSITY

AVAILABLE INFRASTRUCTURE

KU Leuven and Hasselt University collaborate on the Tier-2 infrastructure. The infrastructure consists of:

- 2 clusters, 13 partitions
- 1,440 TF CPU, 1,290 TF GPU
- 26,000 CPU cores
- 131 TB memory
- 133 GPU devices / 513,792 CUDA cores

The KU Leuven Tier-2 cluster had a stable year in 2025 with few changes.

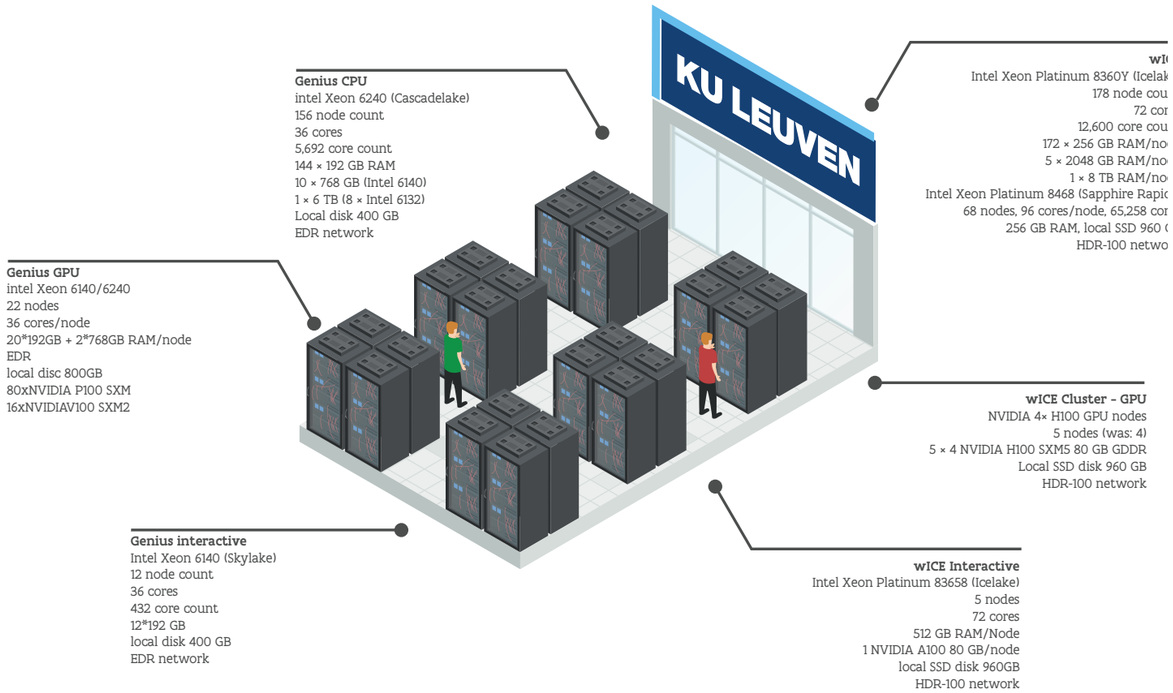

Figure 53. VSC KU Leuven Tier-2 Infrastructure

OPERATIONS AND USAGE

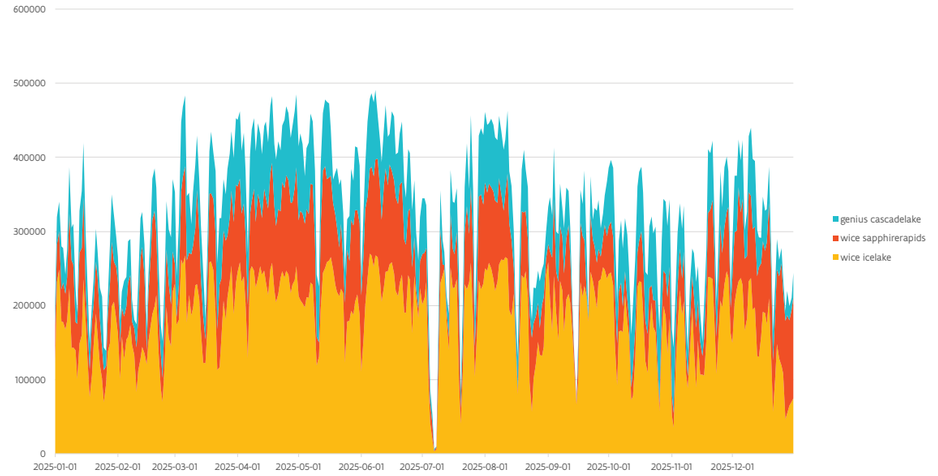

Figure 54. VSC KU Leuven Tier-2 – Thin Node System Usage

The figure above shows that the system was operational almost continuously throughout 2025. In July, maintenance was required on the Tier-2 cluster file systems, during which the system was made unavailable for a short period. A necessary update in September was gradually rolled out across all compute nodes, making downtime unnecessary.

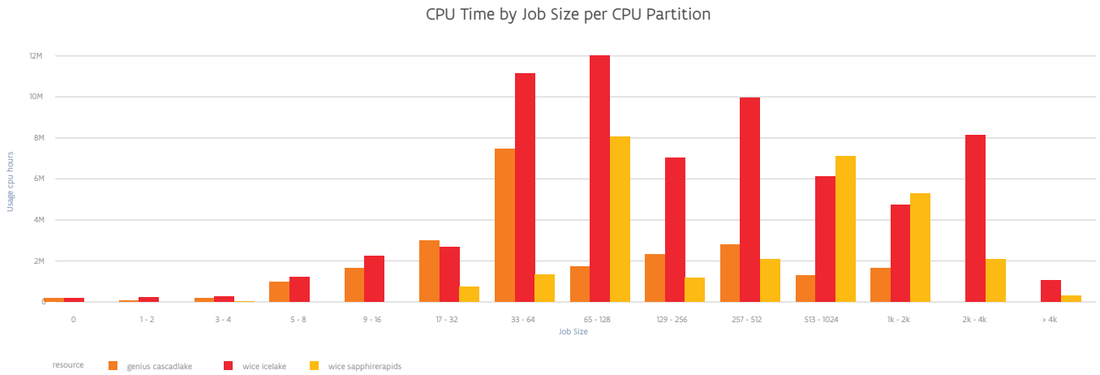

Figure 55. VSC KU Leuven Tier-2 – CPU Time by Job Size

The figure above provides a more detailed view of cluster usage. It shows that Genius is primarily used for smaller jobs, while jobs requiring a larger number of cores are executed on wICE. On wICE Sapphire Rapids, the partition with the highest number of cores per node, approximately 51% of the compute time is consumed by jobs using more than 512 cores. This represents a healthy distribution.

The most recent system is the most capable of handling large-scale parallel computing workloads. Since smaller jobs are mainly executed on the other systems, fragmentation of the Sapphire Rapids nodes remains limited. This enables more efficient scheduling of larger jobs. The most recent expansion of the wICE Tier-2 cluster was the first to use Direct Liquid Cooling. The experience gained over the past year with this technology is highly valuable, as the future system will also be entirely water-cooled.

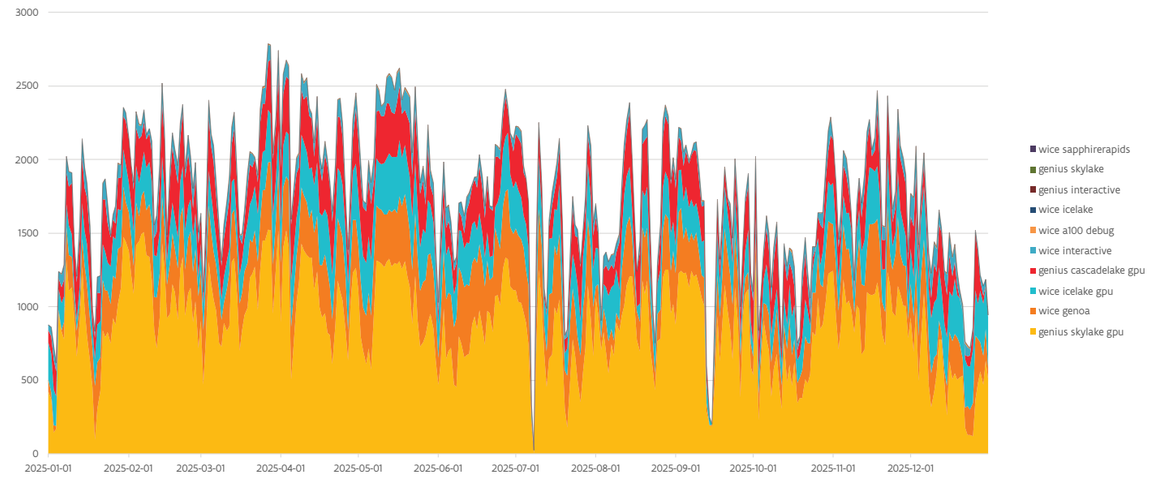

Figure 56. VSC KU Leuven Tier-2 – GPU Usage 2025

The GPUs in the KU Leuven Tier-2 cluster are an important component of the HPC environment, which is also reflected in their utilization. However, an important remark should be made regarding this figure. A significant portion of the GPU hours is consumed on the oldest partition, Genius Skylake GPU. This partition has been operational since 2018 and contains a large number of older NVIDIA P100 GPUs. As long as these GPUs remain in use and are in sufficiently good condition, they will continue to be available.

The wICE cluster consists of 16 A100 GPUs (4 nodes, each equipped with 4 GPUs), together with the more recent H100 generation (5 nodes, each equipped with 4 GPUs).

Depending on the workload, an A100 GPU can deliver four to ten times more computational performance than a P100 GPU. Therefore, the importance of the A100 and H100 GPUs is greater than what may appear at first glance from this figure. However, the rapid supply of specific components and repairs for the newer GPU nodes remains under pressure due to the high demand.

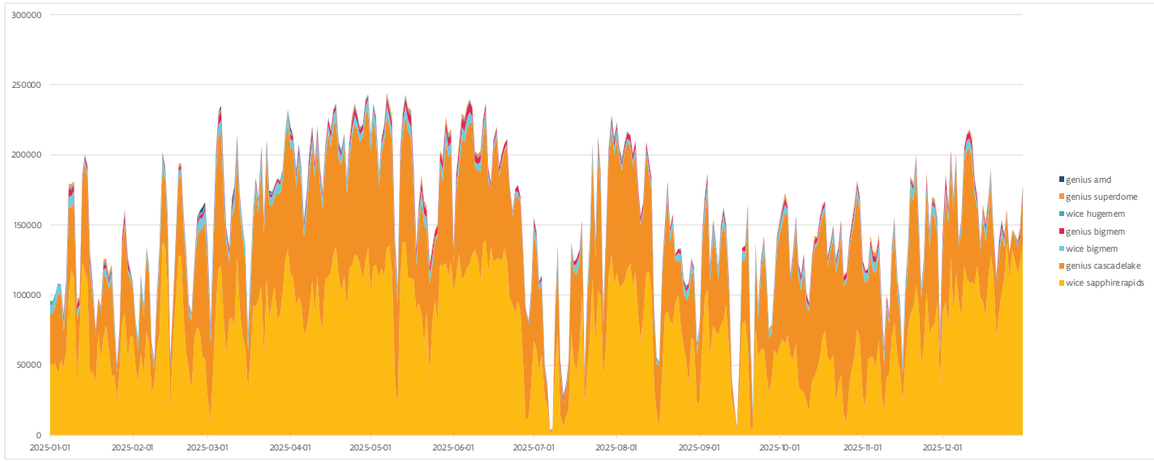

Figure 57. VSC KU Leuven Tier-2 – Specific System Usage

The KU Leuven cluster also includes nodes specifically designed for computations requiring larger amounts of RAM. These nodes are continuously in use (wICE bigmem and Genius bigmem). The nodes with very large memory capacities (the older Superdome and the wICE hugemem partitions) enable calculations that typically cannot be executed elsewhere. These nodes are regularly, but not continuously, occupied, which is in line with expectations for this type of workload. Since the Superdome and hugemem partitions each consist of only one node, the number of CPU hours shown in the figure remains limited. The usage of the interactive nodes is continuous. However, these partitions are not expected to maintain a constant full load. The Genius interactive partition was added this year and consists of the oldest Genius nodes, which have been kept in the system for a somewhat longer period.

For 2025, the parallel file system of the clusters should also be highlighted, as its capacity limit was reached. A healthy file system should maintain a maximum occupancy of approximately 80% of its total capacity. This provides sufficient margin to handle failing hard disks without significantly impacting performance. In 2025, the file system was continuously occupied between 80% and 90% of its capacity. Combined with occasional I/O-intensive computations, this regularly placed pressure on performance and made corrective interventions necessary at times.

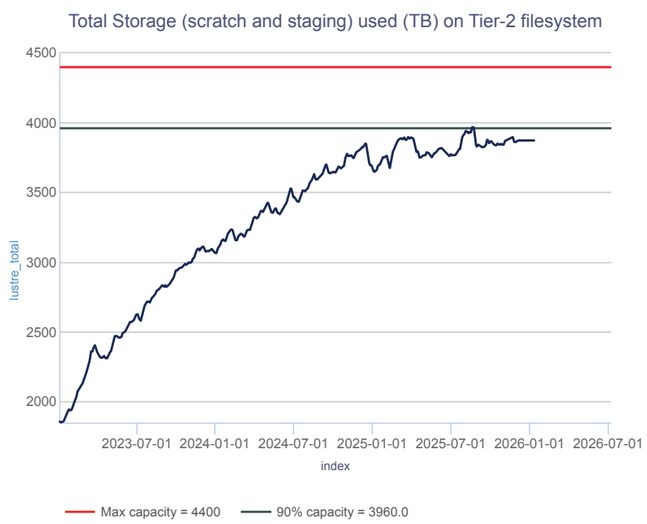

Figure 58. VSC KU Leuven Tier-2 – Storage Usage (Scratch and Staging)

The support team undertook many actions to raise awareness among researchers about this issue. First, researchers were asked to remove old data that was no longer immediately required for processing on the cluster as much as possible. The responsible users receive a monthly overview of the locations of outdated data, allowing them to take more targeted actions. These initiatives were supported by the local HPC Steering Committee.

It was decided not to invest additional resources in the old storage infrastructure but instead to allocate the available budget to provide sufficient storage capacity for the new Tier-2 cluster.

PROCUREMENT OF THE NEW TIER-2 CLUSTER

Following the preparation phase in 2024, a procurement procedure for a new Tier-2 cluster was carried out in 2025. The tender specifications were submitted to suppliers at the end of 2024. The first proposal was received on March 17. Following BAFO (Best and Final Offer) sessions with the candidates, the final proposals were submitted on May 15. In mid-June, the decision was announced, with NEC selected as the successful supplier. The new system consists of 50 compute nodes with a total of 9,600 cores, 24 NVIDIA B200 GPUs, and a parallel file system with a storage capacity of 5.6 PB.

Figure 59. VSC KU Leuven – New Tier-2 Cluster Overview

The installation of the system will be completed in 2026.

ALLOCATION OF COMPUTE TIME

In 2025, 2,090 users were active on the system, representing an increase of 16% compared to 2024. The KU Leuven/UHasselt clusters use a credit-based accounting system built on top of the scheduling software Slurm.

New users receive free compute time to become familiar with the system and to perform initial tests. For actual research activities, project credits are used. Credits can be requested through a simple procedure at minimal cost. This system encourages responsible use of the Tier-2 cluster. The principal investigator acts as the project administrator. They can grant researchers access to compute time and monitor the consumption of project credits. When submitting a computational job, users specify the project to which the credits will be charged. UHasselt researchers also work with credit projects, where the allocation of compute time is managed by the local support team.

Lecturers can also request compute time required for educational courses. A straightforward procedure is available to grant students access to the infrastructure. When students later conduct their own research projects, the transition to using HPC infrastructure for these activities becomes much easier.

USER SUPPORT

In 2025, approximately 4,555 requests submitted through the service desk were handled. 98.97% of these tickets received an initial response within the expected response time, and 93.08% of the tickets were subsequently further addressed within the expected resolution timeframe.

Type of Request	Tier-2 Count	Tier-1 Count
Accounts, access, project and storage management	3,263	42
User application support (incl. software installations)	582	13
System support	367	0
Other	119	21

Table 11. VSC KU Leuven Tier-2 – User Support

The diversity of the user community is also reflected in the variety of support requests. These requests originate from approximately 200 different research groups.

USER SURVEY

In the VSC User Survey conducted at the end of 2025, user satisfaction with the Tier-2 KU Leuven/UHasselt infrastructure and services was also assessed. A total of 155 respondents answered the specific questions regarding the KU Leuven/UHasselt Tier-2 environment.

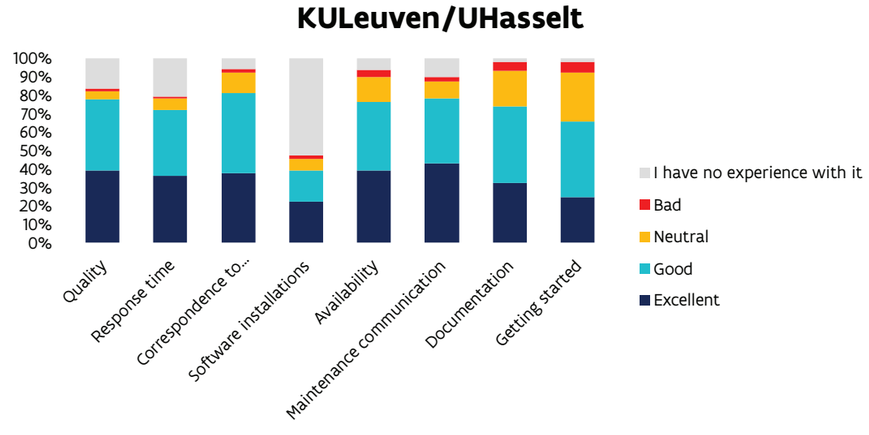

Figure 60. VSC KU Leuven Tier-2 – User Satisfaction 2025

The following aspects of the Tier-2 KU Leuven/Hasselt University infrastructure were rated as "Good" or "Excellent" by users:

- 93% – Quality
- 91% – Response Time
- 86% – Meeting Expectations
- 83% – Software Installations
- 82% – Availability
- 87% – Communication Regarding Maintenance
- 76% – Documentation
- 67% – Getting Started

SOFTWARE INSTALLATIONS

The results for software installations are better than last year. The efforts to introduce more automation into the software installation process have contributed to this improvement.

AVAILABILITY

System availability scored slightly lower than last year. The figures show that there was only one substantial downtime, but the transition to a new operating system and storage capacity issues may have contributed to this somewhat lower perception.

DOCUMENTATION

Documentation remains an area requiring continuous attention. Changes need to be incorporated in a timely manner, and more specific examples should be provided. The documentation was continuously updated, with particular attention given to Open OnDemand usage, data management, and system changes.

SUPPORT FOR NEW USERS

Getting started with the cluster is considered difficult by some users. In 2025, 10 introductory sessions were organized, with an average of around 20 new users attending each session. It is clear that supporting new users will remain an important point of attention.

GPU RESOURCES

As with other Tier-2 sites, there was a clear demand for additional GPU resources. This was taken into account during the procurement of the new Tier-2 cluster, within the limits of the available budget. The new GPUs are particularly well suited for AI workloads, but scaling up will continue to require the use of Tier-1 or Tier-0 resources.

PERFORMANCE OF LOGIN NODES

Comments regarding the performance of the login nodes can largely be attributed to the load on the file system. The use of the interactive partition is certainly part of the solution, but continued efforts to educate users about specific usage patterns that place unhealthy loads on the file system remain essential.

What Our Users Say

"Very helpfull and it has truly saved my research."

"Happy with what you are doing within my context"

"Keep up the good work."

"Really can't say where to improve."

"The team members,have always been very professional and helpful, the trainings are good for having to address such a broad group of users with very different use cases,the downtimes are minimal."

SPECIFIC SUPPORT

With Open OnDemand as a web portal, a new access point to the cluster has been introduced, making it easier for new users to get started. In 2025, more than 1,273 users used Open OnDemand to perform work on the cluster, representing approximately 61% of all users in 2025.

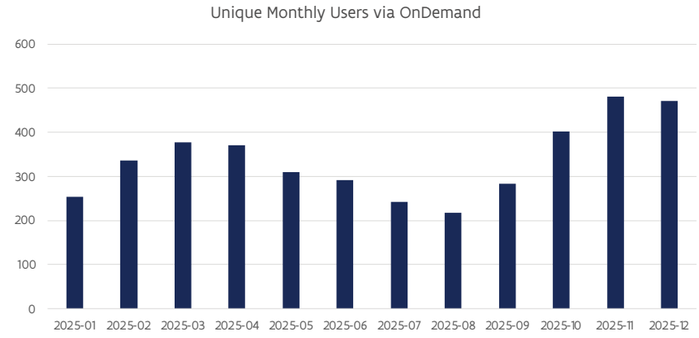

Figure 61. VSC KU Leuven Tier-2 – Unique Monthly Users via OnDemand 2025

The most frequently used applications are JupyterLab, Visual Studio Code Server, Interactive Shell, and RStudio. In 2025, several new applications were added, including Fluent, VS Code Tunnel, and a noVNC desktop. The latter will replace the NoMachine remote desktop, which will be phased out during 2026.

This year as well, several targeted initiatives were undertaken with research groups to further stimulate the use of HPC. For the Computational Biomechanics research group, the challenge was to transition from local Windows machines to the Linux-based cluster environment. This made it easier to run multiple simulations simultaneously. For the Brain and Cognition research group, the Slurm REST API was used to integrate data processing into their existing workflow.

On June 2, an HPC User Day was organized at KU Leuven, attracting around 60 participants. During this event, both workshops and presentations showcasing researchers' work on the cluster were featured. The workshops covered topics such as code profiling, high-throughput computing, and data management. In addition to a poster session, specific user stories were presented by researchers from Image and Speech Processing, Quantum Chemistry and Physical Chemistry, Neurophysiology, and Numerical Analysis and Applied Mathematics.

SHOWCASE

A compelling example of HPC support is the collaboration with the research group "Biomechanics of Human Movement". This group performs complex simulations that originally ran on a Windows workstation, based on MATLAB and linked to specific Windows libraries.

The HPC support team guided the group through the complete transition to Linux and helped them establish their simulation pipeline on the HPC infrastructure. Thanks to this migration, researchers can now not only perform the same analyses on the cluster but also achieve a much larger scale: significantly more models can be simulated simultaneously, and multiple researchers can run their computations in parallel without delaying each other.

In short, what was previously limited to a single workstation has evolved into a powerful, scalable HPC workflow that significantly accelerates the research.

TIER-0 SUPPORT

Since 2023, VSC has provided dedicated support for researchers who wish to use the EuroHPC infrastructure. This primarily concerns LUMI (Finland), the pre-exascale system in which Belgium/Flanders has invested, but users who wish to access other EuroHPC systems are also supported.

The following courses were organized:

- Getting Started with AI on LUMI (February 4–5)
- Advanced LUMI Training (March 3–7)
- Getting Started with AI on LUMI (May 27–28)
- Supercomputing with LUMI (June 2–3)
- Moving Your AI Training Jobs to LUMI: A Hands-On Workshop (October 8–9)
- Introduction to Supercomputing with LUMI (October 20–21)
- LUMI Performance Analysis and Optimization Workshop (October 22–24)

One or more researchers from Flanders participated in all courses.

In addition, the "Supercomputers for Starters" course was further developed and expanded with material specifically relevant to LUMI and other large-scale systems in general.

Work also continued on a dedicated version of an introductory LUMI course with greater focus on the specific situation in Flanders and Belgium. The online course is continuously kept up to date. (https://klust.github.io/intro-evolving).

Through the Tier-1 application form, researchers are offered the possibility to obtain GPU compute time on LUMI without having to submit an additional application. The only requirement is that their Tier-1 application has been approved by the TAB. In total, 25 researchers indicated their interest. Among them were several who had never previously used LUMI. Eventually, not all applicants received a LUMI project, either because their code was not compatible or because they did not respond to the request for additional information. Nevertheless, this provides a low-threshold way for researchers to gain access to LUMI.

A concrete example of more advanced support is the work with Julia on AMD GPUs. Users received guidance before and during the Hackathon in Oslo from an expert of the VSC Tier-0 team, in collaboration with the LUMI User Support Team. During these sessions, the efficiency of the submitted code was evaluated and further support was planned to optimally assist the continued development of the users' codes.

As a result, the VSC Tier-0 expert, together with the HPE LUMI Center of Excellence, developed a roadmap for the efficient development and execution of future code on LUMI, including recommendations for further scaling towards maximum capacity and for developing system-agnostic code.

In 2025, three calls for compute time on the Belgian share of LUMI were organized. Of the 18 submitted regular proposals, 15 were granted.

Date	Researcher(s)	Institution	Research Group / DepartmentV;
2025-09	Stijn D'Hondt	UAntwerp	Medicinal Chemistry
	Jesse Vos	KU Leuven	Center for mathematical Plasma Astrophysics
	Emanuele Sansone	KU Leuven	Processing Speech and Images
2025-10	Pranab Deka	KU Leuven	Center for mathematical Plasma Astrophysics
	Warre Dhondt	UGent	Center for Molecular Modeling
	Sun Wei	KU Leuven	Human-Computer Interaction
2025-11	Pranab Deka	KU Leuven	Center for mathematical Plasma Astrophysics
	Fabio Bacchini	KU Leuven	Center for mathematical Plasma Astrophysics
	Pieter Cnudde	UGent	Center for Molecular Modeling
	Arnout Maet	UGent	Center for Molecular Modeling
	Thomas Nicholas	UGent	Center for Molecular Modeling
	Stijn D'Hondt	UAntwerp	Medicinal Chemistry
	Fabio Bacchini	KU Leuven	Center for mathematical Plasma Astrophysics
	Fabio Bacchini et al.	KU Leuven	Center for mathematical Plasma Astrophysics
	Tinatin Baratashvili	KU Leuven	Center for mathematical Plasma Astrophysics

Table 12. LUMI Applications 2025 – Awarded to Flemish Researchers

In addition, there were 20 preparatory proposals ("benchmark" and "development"), of which 13 were approved. In total, 12 researchers submitted a LUMI-BE application for the first time.

The emails received through the ticketing system can be divided into two main categories: on the one hand, requests related to the LUMI-BE calls or assistance with applications for EuroHPC calls; on the other hand, questions related to LUMI accounts or specific issues encountered while working on LUMI.

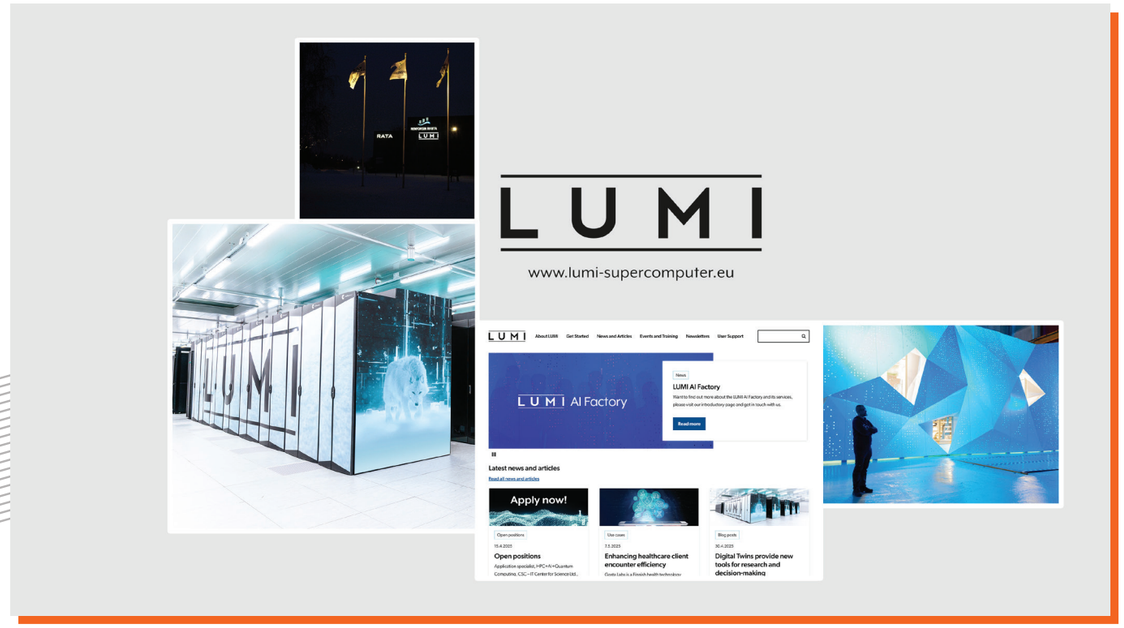

On December 17, the LUMI-BE User Day was organized at the Marie-Elisabeth Belpaire Building together with the VSC User Day. The LUMI-BE track featured five user presentations, in addition to a presentation on the EPICURE project and a plenary lecture:

- Michiel Van Ginderachter (RMI): Catching the Wind: Machine Learning Weather Prediction with Anemoi
- Cem Sevik (University of Antwerp): Quantum Simulations and Machine Learning for Future Materials Technologies
- Stefan Becuwe (University of Antwerp): Short Update on Getting Access to EuroHPC Infrastructure and Support
- Orian Louant (University of Liège): Implementation and Optimization of a Multi-GPU Discontinuous Galerkin Solver for Maxwell's Equations
- Fabio Bacchini (KU Leuven): Next-Generation HPC for First-Principles Simulations of Plasma in Astrophysics
- Laurent Bricteux (University of Mons): High-Fidelity Simulation of High-Speed Turbulent Flows in Low-Pressure Turbines for Advanced Jet Engines
- Michel Rasquin (Cenaero): Tiering Up

Presentations and recordings are available on the website. (https://www.enccb.be/LUMIBE2025).

The next User Day will be organized in the autumn of 2026.

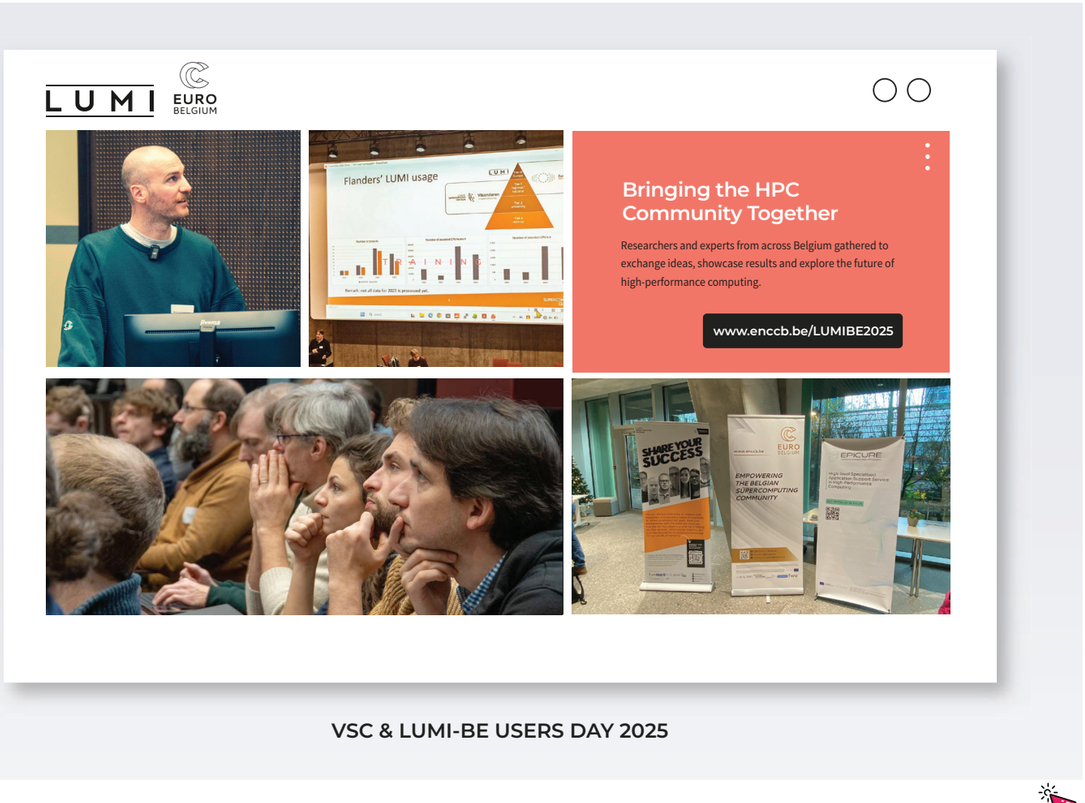

VSC & LUMI-BE USERS DAY 2025

Figure 62. Atmosphere of the Joint VSC & LUMI-BE Users Day 2025 in Brussels

TRAINING

VSC provides training activities primarily aimed at its current or potential users. These include employees from the private sector, public services, and, in terms of numbers, mainly researchers affiliated with the Flemish university associations and various research institutions.

We also observe a growing number of international participants in our training activities. This can partly be attributed to the increasing visibility of VSC abroad, particularly thanks to the close collaboration with EuroCC@Belgium in the areas of communication and training. This is undoubtedly a positive development, as it demonstrates that our activities provide added value.

Training is one of the core activities of VSC. Computations performed using supercomputing infrastructure are costly, both in terms of investment and operational expenses. Therefore, it is essential that this infrastructure is used efficiently. Another important consideration is that the use of this infrastructure and its services is a means to an end, rather than an objective in itself. It is therefore important that users can start working efficiently as quickly as possible and focus on the results of their computations rather than on the technical aspects of how those computations should be performed. VSC training activities make a significant contribution in this regard. While these short-term benefits are clearly important, professional use of this infrastructure will also, in the long term, increase the competitiveness of both our researchers and companies and contribute to the development of an ecosystem from which all parties can benefit.

Training courses and educational materials also contribute to the visibility and reputation of VSC. Analysis of the usage of the VSC website shows that the training page remains one of the most visited pages on the website. Training is also an area that lends itself well to collaboration with other HPC organizations, such as CÉCI, as well as internationally within EuroCC2.

The training activities can be divided into four categories, which either indicate the required prior knowledge or clarify that the topics are domain-specific::

Introductory courses are intended for all users of the infrastructure and are highly recommended for those who do not yet possess the necessary skills. These sessions are delivered by the local VSC staff members. They also provide researchers with the opportunity to become acquainted with the people who answer the questions they submit to the helpdesk. This removes the impersonal and anonymous nature of email communication and helps to lower the barrier for users.

To attend the intermediate-level sessions, the introductory courses are required as prior knowledge. These sessions therefore cover more specialized topics. The majority of these courses are intended for users who develop their own software, either for compute-intensive applications or for data pre- and post-processing.

Because these training sessions are more specialized and intensive than the introductory courses, they are not offered at every VSC site. Users are therefore encouraged to attend training sessions organized at other sites. The vast majority of these trainings are offered in a hybrid or online format, ensuring that geographical distance is not a barrier to participation.

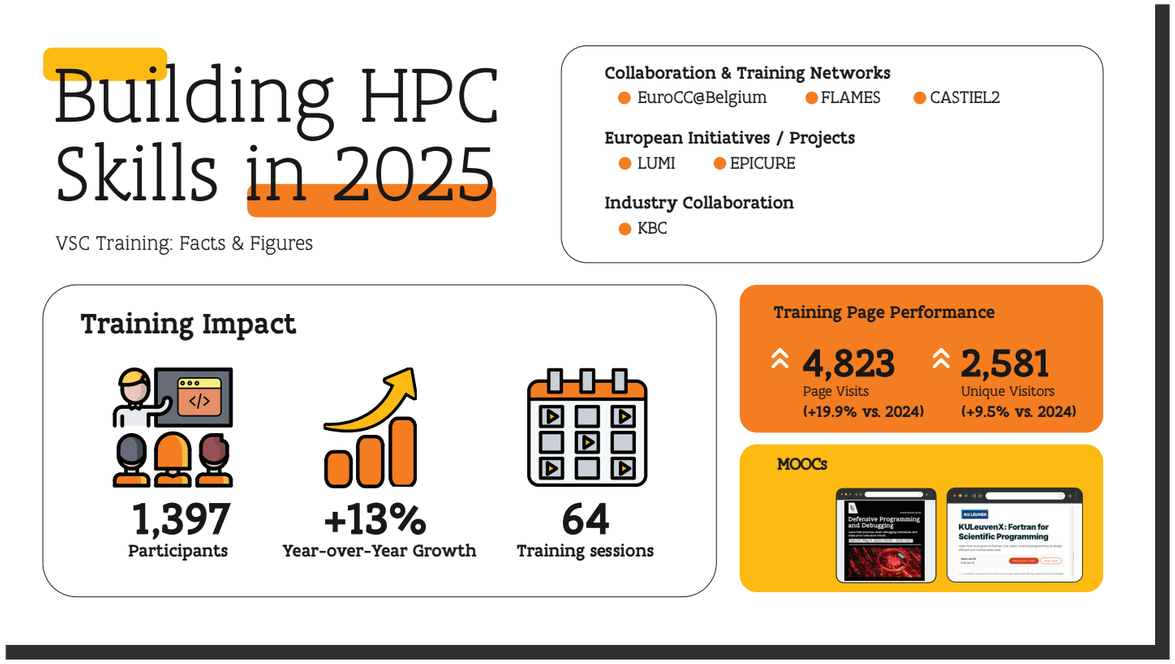

Figure 63. Overview of VSC Training Activities in 2025

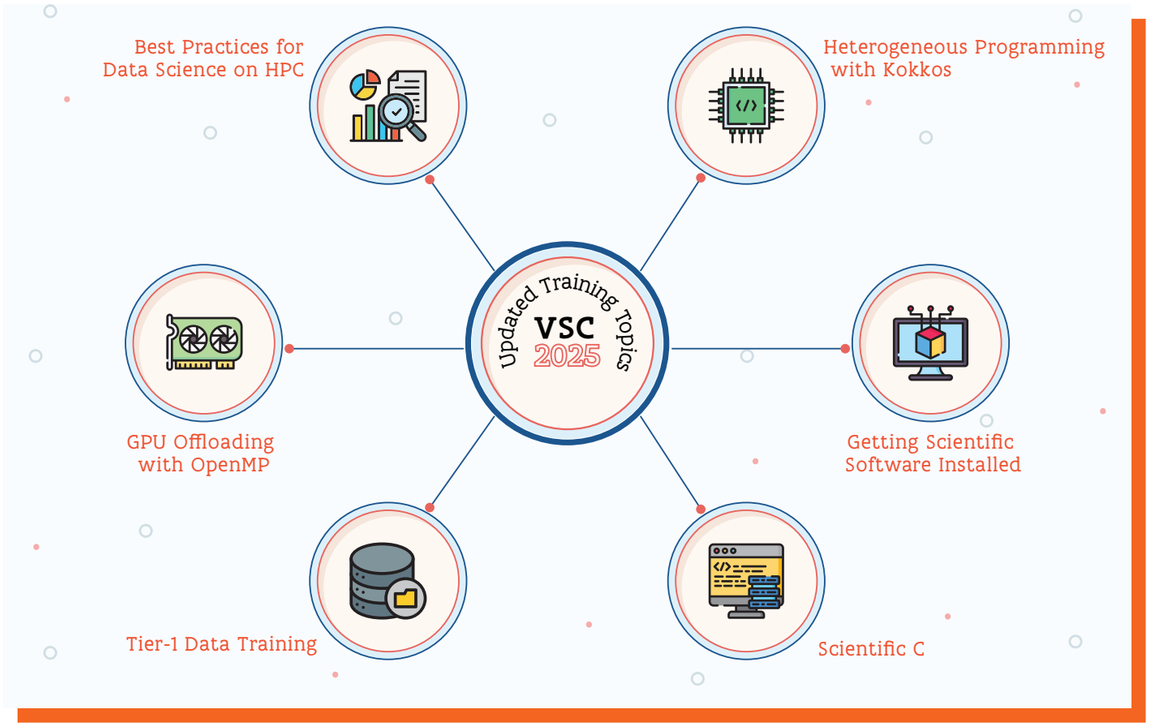

Figure 64. Selection of New and Updated Training Topics Offered by VSC in 2025

Advanced-level training requires even more experience and is more domain-specific than intermediate-level training, which often serves as a prerequisite. For these courses, VSC also relies on external instructors. They are often affiliated with EuroCC2 National Competence Centers or come from industry.

In total, 1,397 participants attended 64 training sessions, representing an increase of 13% compared to 2024.

VSC developed two MOOCs for PRACE: "Defensive Programming and Debugging" and "Fortran for Scientific Computing". These MOOCs were no longer available on FutureLearn, and therefore, work was carried out in 2025 to migrate the course material from FutureLearn to edX. In October 2025, "Fortran for Scientific Computing" became available on edX and had approximately 100 participants by the end of 2025. The KU Leuven team also continued working on "Defensive Programming and Debugging". This MOOC will become available in 2026.

Existing training courses are continuously updated to reflect recent developments, and new training topics are selected and developed each year. In 2025, training sessions were organized on "Getting Scientific Software Installed", "Scientific C", "GPU Offloading with OpenMP", "Heterogeneous Programming with Kokkos", and "Best Practices for Data Science on HPC". The VSC Tier-1 Data team also renewed its training courses on the use of these services and organized four training sessions.

VSC staff also participated in online training courses, often consisting of multiple days.

Training announcements are distributed to current and potential infrastructure users through the website, internal mailing lists, and social media channels. Targeted mailings are used to draw attention to specific courses when they are relevant to a particular audience or potential user group. The broad offering of EuroCC2 and LUMI training courses was regularly promoted through social media and the VSC website.

VSC also provided several tailor-made training courses in collaboration with other organizations, such as FLAMES (Flanders' Training Network for Methodology and Statistics), including "R Best Practices" (3 hours) and "Best Practices for Data Science on HPC" (4 hours).

For KBC, a training course was delivered on the use of Large Language Models (LLMs) on local infrastructure and how to apply them to proprietary data. This training consisted of two webinars attended by approximately 20 participants.

In collaboration with EuroCC@Belgium, a series of six webinars on the use of generative AI in software development and HPC was developed and delivered, attracting approximately 100 participants. In the same context, training material was also developed for a train-the-trainer event on "Multi-GPU Computing", organized by CASTIEL2 in February 2026.

As preparation for a hackathon organized by EPICURE, a webinar on "Python GPU Programming" was delivered, attended by approximately 40 participants.

Figure 65. Visual for the VSC Training "Deploying Large Language Models Locally"

USER SURVEY

The 2025 User Survey shows that users of our infrastructure and services rate the quality of the training offering very highly. This question was answered by 143 respondents.

- 1% – Low quality
- 18% – Average quality
- 62% – High quality
- 18% – Very high quality

The survey also assessed whether the number of training sessions across the different categories is sufficient. The offering of introductory and intermediate training courses appears to be adequate, but the results indicate a need for more advanced and domain-specific training.

- Introductory (121 respondents)
  - 17% – Too few
  - 79% – Sufficient
  - 4% – Too many
- Intermediate (90 respondents)
  - 20% – Too few
  - 79% – Sufficient
  - 1% – Too many
- Advanced (66 respondents)
  - 23% – Too few
  - 73% – Sufficient
  - 4% – Too many
- Domain/application-specific (57 respondents)
  - 37% – Too few
  - 62% – Sufficient
  - 1% – Too many

Out of 429 respondents, only 31% participated in any type of VSC training. It is interesting to relate this to the other survey results, which indicate that users often find it challenging to get started with the infrastructure.

When asked which additional training courses participants would find useful, there was a clear interest in AI and machine learning topics.

What Our Users Say

"I think the quality is already very high."

"Trainings are very informative."

"... the course material was very clear."

"I had a very positive experience... but I had to go through the video recording again... to get the full of it."

USERS DAY 2025

On December 17, 2025, the User Day brought together the Belgian and European HPC community in Brussels for an inspiring day dedicated to high-performance computing and artificial intelligence. The event took place at the Marie-Elisabeth Belpaire Building under the theme "Europe's Intelligence Runs on HPC". It highlighted how advanced computing infrastructures drive innovation across research, industry, and society.

With inspiring keynote presentations, parallel VSC and LUMI-BE sessions, and strong engagement from the community, the day once again provided a valuable platform for collaboration and knowledge exchange within the HPC ecosystem.

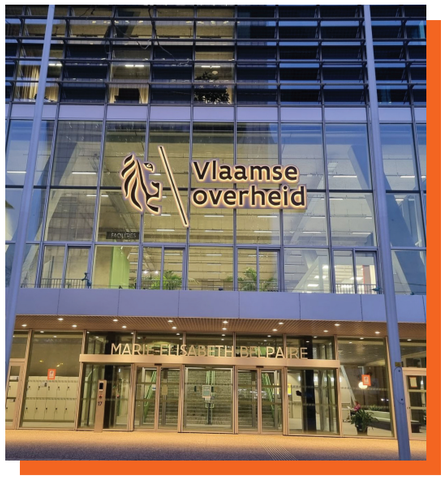

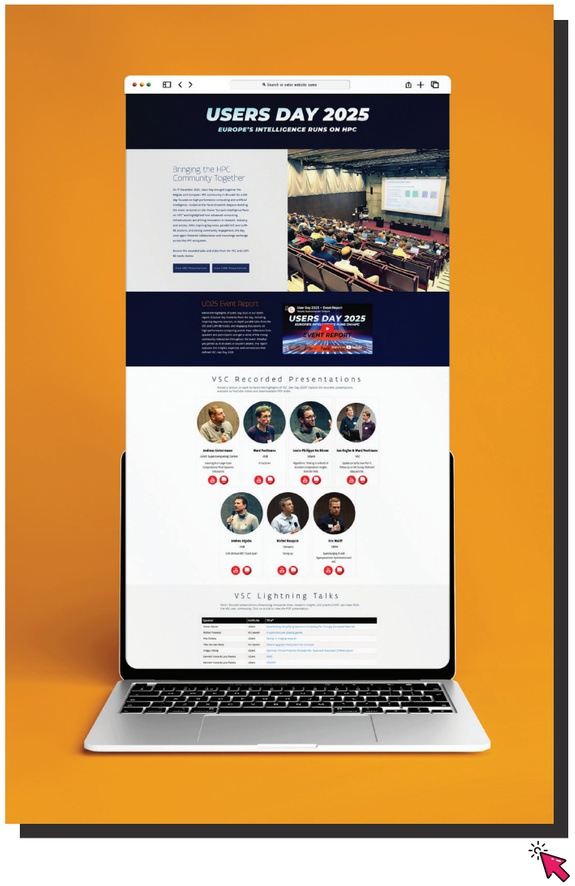

Figure 66. (top) Overview of the VSC & LUMI-BE Users Day 2025 webpage. (left) the Marie-Elisabeth Belpaire building in Brussels, venue of the event. (bottom) atmosphere of the day

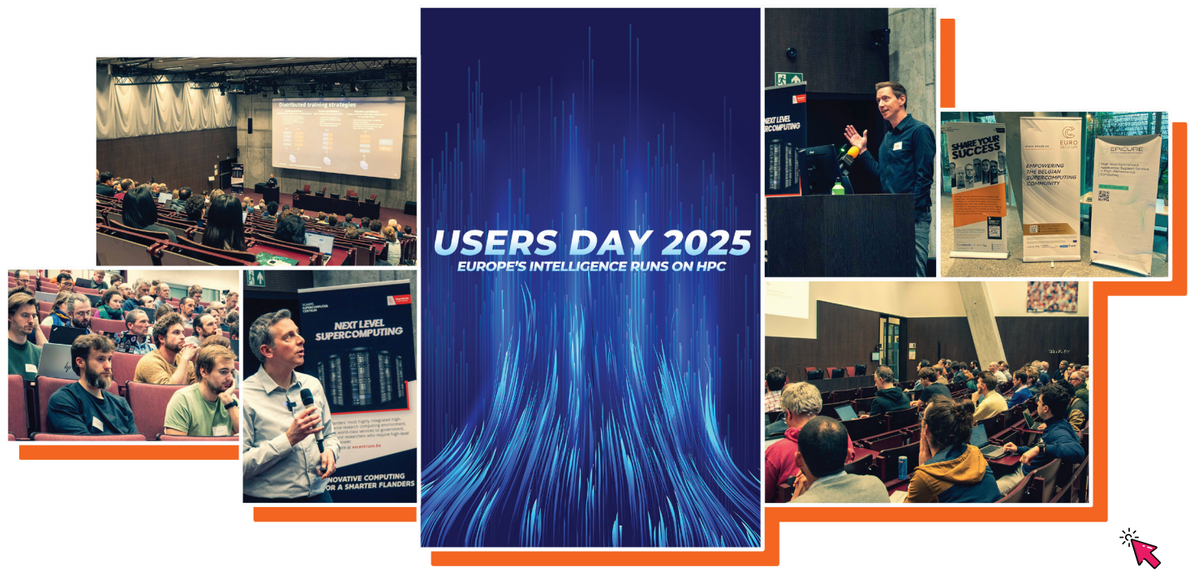

SPEAKERS

Andreas Lintermann

Jülich Supercomputing Center

Learning from Large-Scale Computational Fluid Dynamics Simulations

Ward Poelmans

VUB

AI Factories

Louis-Philippe Kerkhove

UGent

Algorithmic Thinking in a World of Abundant Computation: Insights from the Field

Jan Ooghe & Ward Poelmans

VSC

Update on sofia (new Tier-1), follow-up on VSC Survey 2024 and data policies.

Andres Algaba

VUB

LLMs Without HPC? Good Luck!

Michel Rasquin

Cenaero

Tiering up

Eric Wulff

CERN

Supercharging AI with Hyperparameter Optimization and HPC

To further extend the impact of the event, the presentations from VSC User Day 2025 were made available online as recorded videos and downloadable PDF versions of the original presentations. This allows participants to revisit the sessions and ensures that the shared expertise remains accessible to a broader audience beyond the event itself.

Discover the event report and relive the highlights of VSC User Day 2025 in the official aftermovie.

LIGHTNING TALKS

Speaker	Institute	Title
Simon Ganne	UGent	Downfolding: Simplifying Quantum Complexity for Strongly Correlated Materials
Michiel Provoost	KU Leuven	A supercomputer playing games
Pim Pullens	UGent	Reality in imaging research
Tibo Van den Eede	KU Leuven	Advancing graph theory with the computer
Xingyu Zhang	UGent	Optimize Infinite Projected Entangle-Pair State with Automatic Differentiation
Kenneth Hoste & Lara Peeters	UGent	EESSI
Kenneth Hoste & Lara Peeters	UGent	FOSDEM

Table 13. Overview of lightning talks presented during VSC User Day 2025. Presentations are accessible as PDF via the title.

Best Lightning Talk

During VSC User Day 2025, Tibo Van den Eede (KU Leuven – KULAK) received the award for Best Lightning Talk for his presentation "Advancing Graph Theory with the Computer". The presentation demonstrated how computational methods and HPC technologies can contribute to advanced mathematical research and applications in graph theory.

Figure 67. Tibo Van den Eede (KU Leuven Kulak) receiving the Best Lightning Talk Award at VSC User Day 2025.

LUMI Presentations

In addition to the VSC lightning talks, the event page also features presentations on research and applications supported by the LUMI infrastructure. These presentations are available as recorded videos and/or downloadable PDF slides.

OUTREACH

VSC WEBSITE

Introduction

The VSC website (http://www.vscentrum.be) serves as the central digital gateway to the services, infrastructure, and expertise of the Flemish Supercomputer Center. The website supports researchers, companies, students, policymakers, and the wider public by providing information on supercomputing infrastructure, training activities, research initiatives, and HPC services.

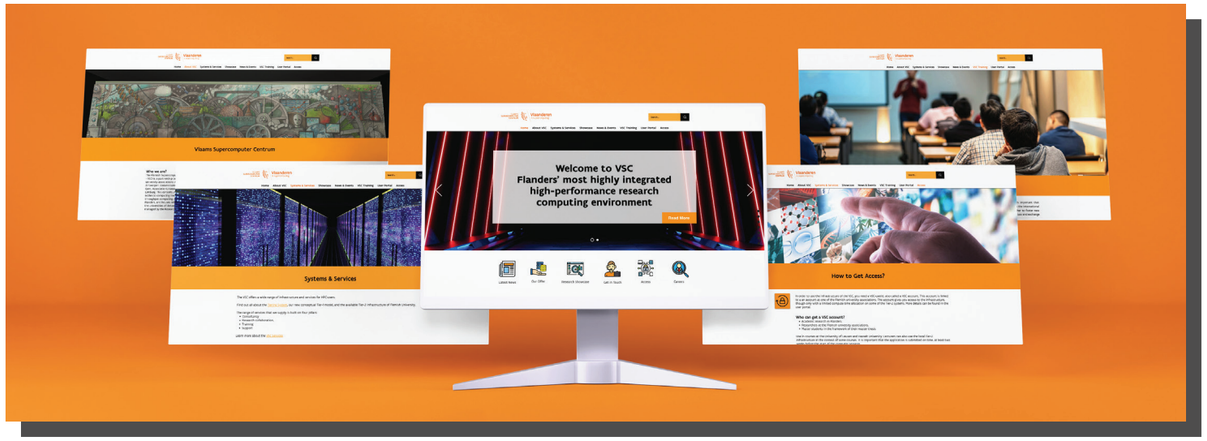

Website Objectives

The website aims to provide users with accessible access to VSC resources, support, and documentation. In addition, the platform strengthens the visibility of Flemish and European HPC initiatives and supports communication around training activities, events, research outcomes, and infrastructure developments.

Key Sections and Developments

- **Homepage** | A central overview of VSC activities, infrastructure, and the latest updates.
- **Services and Infrastructure** | Information on VSC supercomputers, cloud and data services, user support, and access opportunities.
- **Access and Accounts** | Guidelines for user accounts, project applications, and access to HPC resources.
- **Support and Documentation** | Access to user guides, technical documentation, training activities, and helpdesk support.
- **News and Events |** Updates on training courses, events, infrastructure developments, and community activities.
- **Research Showcase** | A growing collection of research stories demonstrating how researchers use VSC infrastructure across a wide range of scientific domains.
- **VSC Connect Blog** | A new blog section where VSC staff share insights on HPC, AI, data science, infrastructure, and user experiences.
- **Business Services** | A new webpage aimed at companies and innovation partners interested in exploring the potential of HPC and AI for industrial applications.
- **User Kit** | A download page containing the official VSC logo and guidelines for its use in posters and presentations.
- **Contact Information |** Information on how users can contact VSC and the different expertise centers.

User Engagement

In 2025, the website recorded more than 35,500 sessions and 16,500 unique visitors. Organic search traffic remained the primary source of website visits, while returning visitors represented a significant portion of the audience. This demonstrates strong user engagement and continued interest in VSC services and activities.

Technical Aspects

The website is supported by a stable technical infrastructure that ensures reliability, security, and performance. Throughout the year, several updates were implemented to further improve functionality, accessibility, and the overall user experience.

Accessibility and Usability

VSC continues to invest in an accessible and user-friendly website with clear navigation and well-structured content for diverse audiences, ranging from new users to experienced HPC researchers.

Future Plans

VSC will continue to develop the website as a central communication and support platform for the HPC community. Future developments will focus on additional research content, training activities, and further support for HPC, AI, and data-related initiatives.

Conclusion

The VSC website remains an important digital platform for the Flemish and European HPC community. With a growing range of research- and user-oriented content, the website supports the VSC mission of making high-performance computing accessible and visible to a broad audience.

NEW WEBPAGES AND INITIATIVES

In 2025, VSC further expanded its online presence with new webpages and digital initiatives targeting different audiences within the HPC community.

Business Services

The new Business Services page was developed to introduce companies and innovation partners to the opportunities offered by high-performance computing in industrial applications in an accessible way.

The page provides information on HPC infrastructure, access opportunities, training, and support for companies seeking to integrate artificial intelligence, simulations, and data-intensive applications into their digital transformation processes.

In addition, the page features real-world examples and testimonials from companies using VSC infrastructure across various sectors, including engineering, life sciences, energy, and industrial innovation.

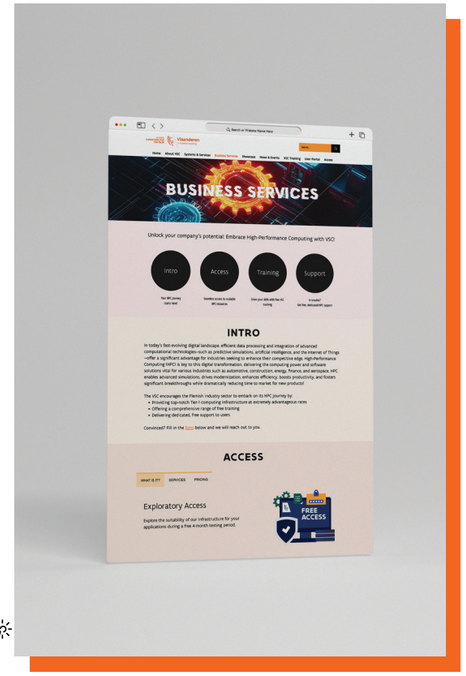

Figure 68. Overview of the Business Services Webpage

User Kit

To increase the visibility of VSC-related research, VSC launched the User Kit page. Through this page, researchers can easily download the official VSC logo in different formats and color variations, accompanied by guidelines for its correct use on posters and presentations. Researchers who include the VSC logo in their poster presentations can also submit their work for publication on the VSC website and social media channels. The inclusion of the VSC logo is entirely voluntary – never mandatory – but it helps highlight the role of VSC in the research and provides additional visibility to the work.

By making recognition as accessible as possible, VSC aims to build a growing collection of research outcomes that demonstrate the broad impact of high-performance computing across academic disciplines, government, and industry.

Figure 69. VSC User Kit Webpage Flyer with Practical Instructions

VSC Connect Blog

In 2025, VSC launched the VSC Connect Blog as a new communication platform for sharing insights, expertise, and discussions on high-performance computing, artificial intelligence, and digital research infrastructures. The blog provides VSC staff members with the opportunity to highlight technical, strategic, and community-oriented topics in an accessible way.

Throughout the year, three in-depth articles were published by VSC staff members and experts. These contributions generated strong engagement and visibility through both the website and social media channels, contributing to broader interaction within the HPC community.

The published interviews covered a range of topics within the HPC ecosystem:

- Geert Jan Bex on how VSC supports learning and professional development in HPC and scientific computing.
- Lander Willem, Chair of the VSC User Committee, on community building and the future of HPC.
- Carl Mensch on the opportunities and challenges of quantum computing.

Figure 70. VSC Connect Blog Banner for the Interview with Lander Willem

With the VSC Connect Blog, VSC strengthens its role as a knowledge platform and communication channel for HPC, AI, data science, and future computing technologies.

RESEARCH SHOWCASE

The Research Showcase section continued to grow in 2025 as a platform for highlighting scientific research supported by VSC infrastructure. The articles provide an overview of the broad range of innovation and research enabled by high-performance computing, data analysis, and artificial intelligence across diverse scientific disciplines.

Throughout the year, 33 new Research Showcase articles were published, featuring contributions from researchers and research groups from various Flemish knowledge institutions and research domains. The articles highlight both the scientific impact of the research and the role of VSC infrastructure and expertise within these projects.

The Research Showcase contributes to increasing the visibility of Flemish research initiatives and serves as an important communication channel for making complex scientific applications accessible to a broader audience.

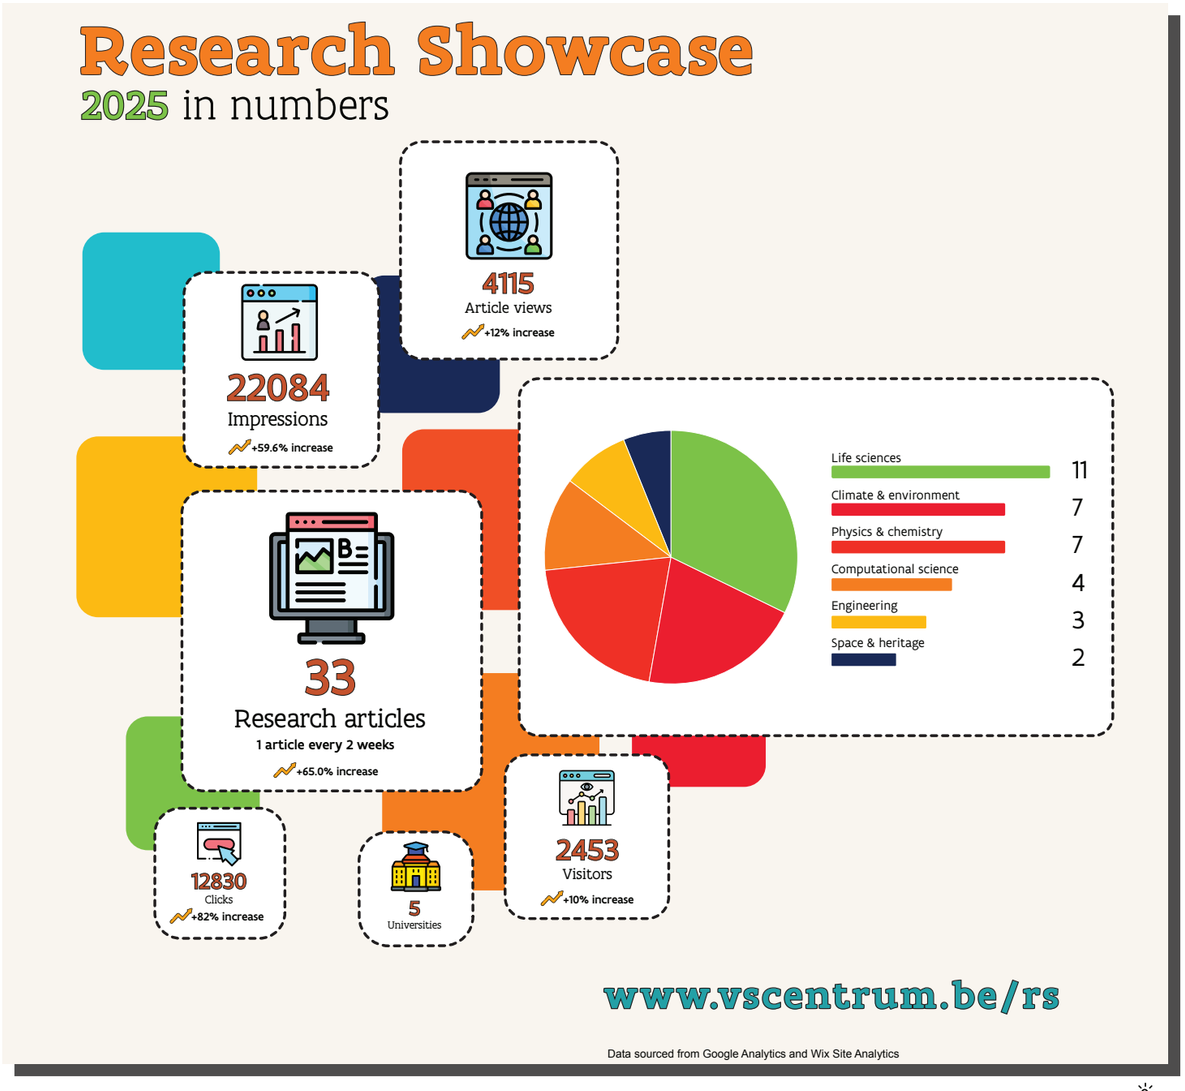

SOCIAL MEDIA AT VSC

Social media continue to play an important role in VSC's communication and outreach strategy. Through channels such as LinkedIn, YouTube, and Bluesky, VSC shares news, research outcomes, training activities, events, and infrastructure developments with a broad audience.

These platforms strengthen the visibility of high-performance computing and artificial intelligence, support community building, and contribute to the dissemination of knowledge within Flanders and beyond.

LINKEDIN: GROWTH OF THE VSC COMMUNITY

The VSC LinkedIn page (@vschpc) emains the most important social media channel for professional outreach and community engagement. The platform is used to inform researchers, users, companies, and partners about VSC services, research outcomes, training activities, events, and infrastructure developments.

In 2025, the number of followers increased from 1,957 to 2,860, representing a growth of more than 46%. This growth was supported by a consistent content strategy focusing on Research Showcase articles, success stories, training activities, VSC User Day, infrastructure updates, and news related to HPC and AI. A total of 234 posts were published on LinkedIn throughout the year.

The increasing engagement confirms LinkedIn as an important channel for strengthening the VSC community and increasing the visibility of HPC in Flanders.

Figure 71. New VSC LinkedIn Banner

YOUTUBE: KNOWLEDGE SHARING THROUGH VIDEO

The VSC YouTube channel (@vschpc) remains an important platform for education, knowledge sharing, and outreach. Through webinars, training sessions, recorded presentations, success stories, and thematic series, the channel provides users with continued access to expertise in HPC, AI, and data science.

In 2025, the channel was further expanded with recordings from VSC User Day, training materials, and research stories. As a result, YouTube continues to serve as a valuable resource for both new and experienced users and contributes to the visibility of research and innovation supported by VSC.

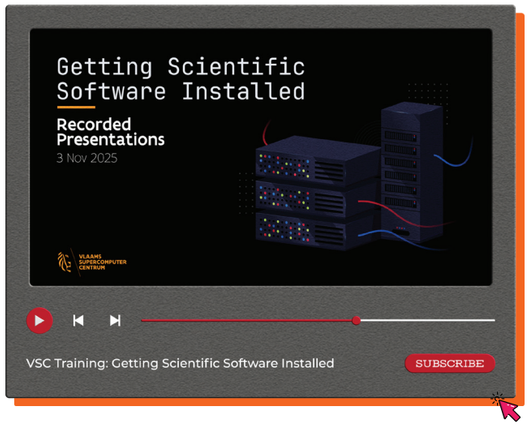

Figure 72. Series of training videos "Getting Scientific Software Installed", available via the VSC YouTube

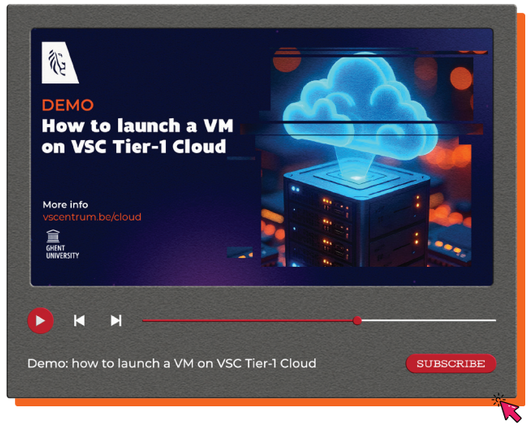

Figure 73. Instructional video on setting up a virtual machine within the VSC Tier-1 Cloud

BLUESKY

The VSC Bluesky account (@vschpc.bsky.social) continued to be used in 2025 as an additional communication channel for sharing news, events, training activities, and insights from HPC and AI research. Although the platform is still in a growth phase, it provides VSC with the opportunity to reach new audiences and maintain a presence on emerging social media networks relevant to the research and technology sectors.

Through an active presence across multiple social media platforms, VSC continues to invest in a broad and accessible communication approach that informs, connects, and inspires its users.

MAILING (e-NEWSLETTER)

The VSC e-Newsletter continues an important communication and outreach channel for informing researchers, users, partners, and other stakeholders about developments within the Flemish and European HPC ecosystem. Published four times a year, the newsletter reaches a growing community of more than 1,250 subscribers and active readers.

Each edition brings together relevant updates on infrastructure, training activities, research outcomes, events, and new initiatives. The newsletter includes:

- Announcements of training courses, workshops, and important deadlines
- Research stories and Research Showcase articles
- Updates on HPC, cloud, and data services
- News on infrastructure developments and strategic projects
- Reports on events and community activities
- Calls for project applications, collaboration, and participation in initiatives

In addition to the regular seasonal editions, targeted newsletters were also distributed to support specific campaigns and events, including VSC User Day, training series, infrastructure updates, and other outreach activities.

The newsletter is a valuable channel for maintaining contact with the VSC community and strengthening the visibility of HPC, AI, and data-intensive research. By combining technical updates, research stories, and practical information, it supports the dissemination of knowledge and stimulates engagement within the Flemish and European research community.

Figure 74. VSC E-Newsletter Autumn 2025 Edition

SUCCESS STORIES

Through the Success Stories series, researchers, companies, and organizations share their experiences using VSC infrastructure for research, innovation, and societal applications. The video interviews demonstrate the impact of high-performance computing across a wide range of scientific and industrial domains.

UNRAVELING A THOUSAND YEARS OF HISTORY THROUGH ANCIENT DNA

In 2025, VSC first published a Research Showcase article on the genetic history of the population buried beneath the historic city center of Sint-Truiden. The research received extensive attention from various Belgian media outlets and later formed the basis of an in-depth Success Story video featuring researchers Toomas Kivisild and Maarten Larmuseau.

By combining ancient DNA with archaeological and historical data, the researchers were able to reconstruct more than a thousand years of ancestry, migration, and community development. The study also revealed unexpected discoveries, including traces of plague infections from the fourteenth century that had not been documented in historical records.

Through large-scale genome sequencing and advanced data analysis, fragmented DNA samples were transformed into valuable historical information. These analyses were made possible by the computational power of VSC's Tier-1 and Tier-2 infrastructure.

This Success Story demonstrates how high-performance computing supports interdisciplinary research and enables new insights at the intersection of genetics, archaeology, history, and data science. At the same time, the example illustrates how research supported by VSC can generate not only scientific impact but also broad societal and media interest.

> "Without the Flemish Supercomputer Center, it would take us years to reconstruct genomes from the past. Thanks to this computing power, we can significantly accelerate our research."

MAARTEN LARMUSEAU
Professor at Human Genetics Department
KU Leuven

AI AND SUPERCOMPUTING FOR MORE ACCURATE LUNG CANCER DIAGNOSIS

The PIDGIN project at Hasselt University is an interdisciplinary collaboration between oncologists, pathologists, biomedical researchers, and data scientists. The project aims to develop new methods for more accurate lung cancer diagnosis by mapping the heterogeneity of tumors.

Through a two-part video series, the project demonstrated how researchers use artificial intelligence, machine learning, and VSC infrastructure to analyze digital microscopic tissue images. Using advanced algorithms, individual cells are detected, classified, and studied, revealing patterns that are difficult to identify using traditional methods.

The researchers made use of VSC Tier-1 Data and Compute resources for the storage, processing, and analysis of large volumes of imaging data. The infrastructure supported both the training of neural networks and the execution of large-scale analyses on high-resolution images.

This project illustrates how high-performance computing and artificial intelligence contribute to data-driven medical innovation. By providing researchers with access to powerful computing and storage infrastructure, VSC supports the development of more accurate, objective, and reproducible diagnostic methods, ultimately contributing to more personalized treatments for lung cancer patients.

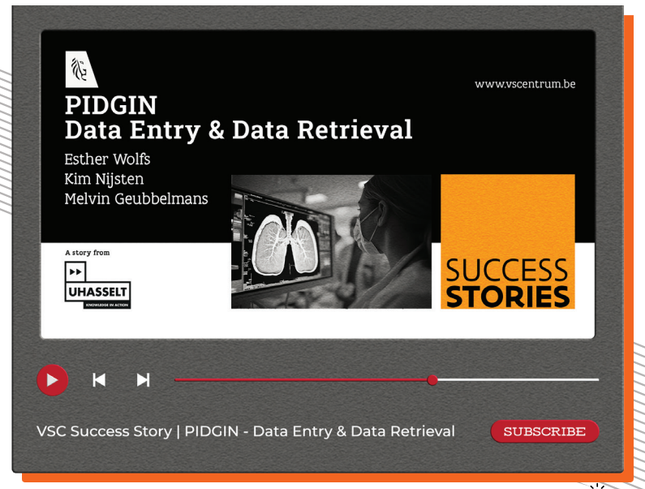

"Thanks to VSC, we have access to a powerful platform for storing, sharing, and analyzing large volumes of imaging data, enabling the development of advanced AI models for cancer diagnostics."

ESTHER WOLFS
Associate Professor
UHasselt

INTERNATIONAL COLLABORATION

EUROCC

Since 2023, VSC has participated in the EuroCC 2 project, the successor to the EuroCC 1 project launched in 2020. This initiative, supported by the EuroHPC Joint Undertaking and national authorities, maintains a network of 33 National Competence Centers (NCCs) across participating countries. The NCCs act as central access points in each country, connecting stakeholders with national and European HPC systems. They operate at regional and national levels to engage with local communities, map HPC competencies, and facilitate access to European HPC resources for users from the academic, private, and public sectors.

EuroCC Belgium brings together the expertise of VSC, Cenaero, and CECI and carries out a diverse range of activities. These include developing communication campaigns, providing a formal framework connecting HPC-related activities, creating and maintaining a comprehensive and transparent overview of HPC competencies and institutions, and serving as a gateway for industry and academia to providers with suitable expertise or relevant projects, both nationally and internationally. EuroCC Belgium also consolidates HPC training opportunities and presents them through a central platform, alongside training activities offered abroad and collected by other NCCs. Furthermore, it promotes the industrial adoption of HPC and facilitates access to both computing systems and expertise.

In 2025, VSC continued its strong commitment to supporting industry through EuroCC Belgium, with a particular focus on SMEs. This priority was reflected in the support provided to 170 SMEs. This result was achieved through intensified collaboration with ecosystem partners such as VLAIO, Agoria, Beltug, A6K, and FABI to strengthen lead generation and referral pathways, complemented by targeted outreach through LinkedIn.

More specifically, VSC provided SMEs with guidance on FFplus and AI Factory opportunities, practical support for industrial HPC users seeking access to computing infrastructures such as Hortense (Flanders), Lucia (Wallonia), and EuroHPC supercomputers, as well as quantum-focused webinars and introductory sessions. Renewed engagement with TTOs/KTOs also led to interactions with previously unknown university spin-offs and concrete follow-up actions regarding infrastructure access.

In parallel, collaboration with larger companies increasingly evolved towards a strategic dialogue on advanced computing, particularly quantum technologies. VSC led the Quantum Technologies working group within Quantum Circle, organized a webinar on quantum communication and a live event at The Beacon, and conducted interviews with experts to better understand sector-specific needs and generate new leads.

At the same time, VSC supported large enterprises in diversifying their access to HPC resources, including local Tier-1 systems and EuroHPC resources such as LUMI, as a response to geopolitical risks and to mitigate potential downtime associated with relying on specific supercomputing systems.

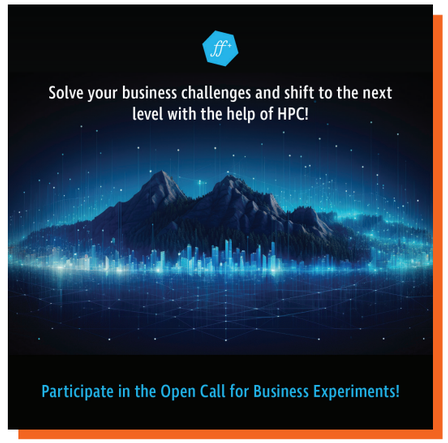

Figure 75. Campaign for the FFplus Open Call, targeting European SMEs and start-ups looking to leverage HPC for innovation

EPICURE

In February 2024, the EPICURE (https://epicure-hpc.eu) (European Pilot for Exascale Computing User Support) project, the "EuroHPC Application Support Project", was launched. The consortium is built around the hosting sites of current and future EuroHPC systems. VSC participates as part of the team supporting the LUMI supercomputer.

Within EPICURE, Level 2 and Level 3 support is provided to users of EuroHPC systems. Users can request this support directly as part of their compute time application. The support typically ranges from three to six months and includes services such as performance analysis, benchmarking, code refactoring, code optimization, and adapting applications to new architectures.

In addition, EPICURE develops so-called "Best Practice Guides" and organizes hackathons, training courses, and webinars.

VSC was actively involved in Work Package 2, "Code Porting, Enabling and Scaling", in 2024 and 2025. From 2026 onwards, VSC will also contribute to Work Package 3, "Code Optimization". In 2025, VSC was responsible for the "Best Practice Guide on Power Consumption Measurements in EuroHPC Systems". It also contributed to several projects, including the migration from ScaLAPACK to ELPA for solving large-scale eigenvalue problems.

Since access to EPICURE is not available for national or regional projects such as VSC Tier-1 projects, it is complementary to VSC's Tier-0 support. To increase awareness of EPICURE among Tier-1 users, a dedicated question will be added to the Tier-1 application form from 2026 onward.

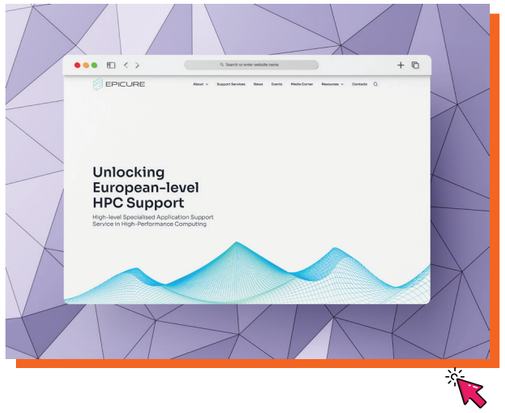

Figure 76. EPICURE Project Homepage

EUROSSQ-HPC

VSC is also involved in the EuroSSQ-HPC project, one of the EuroHPC initiatives focused on quantum computing. Within this project, VSC collaborates with SURF, the Netherlands eScience Center, Leiden University, Nikhef, and Delft University of Technology in the Netherlands, as well as GENCI in France.

The quantum computer will be hosted at the Amsterdam Science Park and integrated with SURF's Snellius supercomputer. At the end of October 2025, proposals from candidate suppliers were submitted. Following the evaluation of these proposals and synchronization with LuxProvide, which will operate a similar system, the project will formally start.

This will result in two semiconductor spin-qubit quantum computers being added to the eight other systems co-funded by EuroHPC, further strengthening the European quantum computing ecosystem.

MULTIXSCALE

Since 2023, the HPC team of Ghent University has been a partner in the MultiXscale EuroHPC Center of Excellence (https://multixscale.eu). This project, which runs until the end of 2026, is a collaboration between 16 European partners and aims to develop simulations that operate across multiple time and length scales.

The technical contribution of the project focuses on further developing and preparing the European Environment for Scientific Software Installations (EESSI, https://eessi.io). This is a shared collection of scientific software installations that allows a uniform environment to be offered on different systems, ranging from personal computers to virtual machines in the cloud and supercomputers. EESSI makes the work of scientists developing software for complex simulations easier and also allows results to be more readily exploited by others.

EUROHPC FEDERATION PLATFORM (EFP)

Since early 2025, the HPC team of Ghent University has been part of the consortium developing the EuroHPC Federation Platform (https://my-eurohpc.eu) , with a strong focus on the Federated Software Catalog component, which is based on EESSI, the European Environment for Scientific Software Installations (EESSI, https://eessi.io).

The goal of this platform is to simplify access to and use of the EuroHPC computing infrastructure, including the current EuroHPC supercomputers as well as the future AI Factories. A first version is expected to be operational by early April 2026.

COMMUNITY BUILDING

To maintain good relationships with existing user groups and establish contacts with potential new users from academic, industrial, or government circles, VSC contributes to building a number of communities. It does so within the context of several high-impact current themes such as AI and quantum computers. For example, VSC has previously organised various symposia on the significant influence of AI in structural biology and life sciences. Furthermore, VSC actively participates in a number of initiatives around quantum computers, and in recent years has organised a considerable number of webinars and events on quantum computers and their future impact across various fields. VSC is also involved as a core member of the Belgian Quantum Circle, which represents Belgian organisations in the field of quantum technologies.

Figure 77. Promotional visual for the webinar "Safeguarding Communications and Data in the Quantum Computing Age"

COLOPHON

The Flemish Supercomputer Center (VSC) is a virtual supercomputer center for both academics and industry. It is managed by the FWO, in collaboration with the five Flemish university associations..

Research Foundation Flanders (FWO)
Leuvenseweg 38
1000 Brussel
T 02 512 91 10

Published by
Dr. Hans Willems
Secretary General FWO

Text
Leen Van Rentergem, Jan Ooghe, Ingrid Barcena
Stefan Becuwe, Wim Vanroose
Johan Van Camp, Ewald Pauwels
Geert Jan Bex, Rafal Al-Tekreeti
Stefan Weckx, Ward Poelmans
Caroline Volckaert

Creation and Design
Rafal Al-Tekreeti

Published in **June 2026**

DISCOVER VSC

vscentrum.be