VLAAMS SUPERCOMPUTER CENTRUM

Flanders is Supercomputing

ANNUAL REPORT 2023

With support from

FWO Research Foundation Flanders
Opening new horizons

Contents

Foreword

The importance of computing power has increased significantly, and today it is even indispensable for a region to remain competitive within the many domains of expertise in science and technology. VSC wants to contribute to strengthening the innovation potential in Flanders by offering services and infrastructure at regional (Tier-1) and local (Tier-2) level for all RDI activities in academia, industry and government. The VSC also monitors international developments and ensures alignment with the European HPC objectives. Given these recent developments, the VSC must grow into a renewed role within Flanders.

A new strategic plan and financing for a future-oriented and sustainable expansion of the Tier-1 level has been implemented from 2023, so that users continue to be supported in their need for computing power, and Flanders can remain competitive in a rapidly evolving, big data-based scientific environment. and technological world.

In the newly designed structure in which the VSC remains a virtual, decentralized organization under the supervision of and in collaboration with the FWO, the operational tasks are carried out by the five university associations, four of which are also an infrastructure hub. For the benefit of the growing group of users, contact with (potential) users takes place via service points that function as a single point of contact, and where each association takes responsibility for organizing a specific service VSC-wide. For example, KU Leuven will focus on data storage, UAntwerp will focus on the European Tier-0 component, UGent will focus on cloud computing and compute, UHasselt will focus on marketing and communication and coordination of training, and VUB will organize a service point for users outside the knowledge institutions.

During 2023, the second phase of Tier-1 Hortense was also installed with additional CPU and GPU partitions and, after a thorough evaluation, the location of the next Tier-1 was decided, with VUB being responsible for the purchase, installation and exploitation thereof. The procurement procedure will start from January 2024 with the aim of awarding the purchase by January 2025. The new Tier-1 should then become operational by November 2025.

The VSC is now prepared for a challenging future.

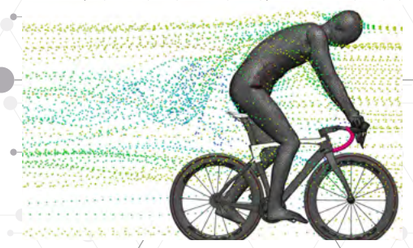

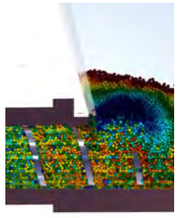

KU LEUVEN

Leen Van Rentergem
Jan Ooghe
Ingrid Barcena

University of Antwerp

Annie Cuyt
Stefan Becuwe

GHENT UNIVERSITY

Johan Van Camp
Ewald Pauwels

VUB VRIJE UNIVERSITEIT BRUSSEL

Stefan Weckx
Ward Poelmans

UHASSELT

Geert Jan Bex
Rafal Al-Tekreeti

fwo

Caroline Volckaert

In the Spotlight

Below you will find a summary of the most important developments and facts from the VSC Annual Report 2023.

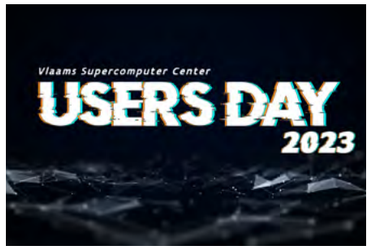

The VSC Highlighted

Introduction

The support and operation of High Performance Computing (HPC) in Flanders is organized within the Flemish supercomputing center, the VSC. The VSC is a consortium in which the five Flemish associations join forces to offer HPC infrastructure and support to the broad research community in Flanders, consisting of universities, knowledge institutions, companies and government. In addition, within its mission as a service provider, the VSC offers a range of training courses that should promote the use of the infrastructure. It houses infrastructure in four hubs: UAntwerp, Vrije Universiteit Brussel, UGent and KU Leuven.

The VSC is managed by the FWO, Research Foundation - Flanders.

In the European model for HPC, a distinction is made between three levels: the computing capacity that research institutions have (Tier-2), the computing capacity whose needs and costs exceed an institution and which is provided at the level of a region or a country (Tier-1) and the very large scale computing infrastructure (Tier-0). The VSC focuses mainly on the Tier-2 and Tier-1 layers, and thus provides stepping stones between Tier-3 and Tier-0. For the Tier-0 layer, the VSC offers support to its users.

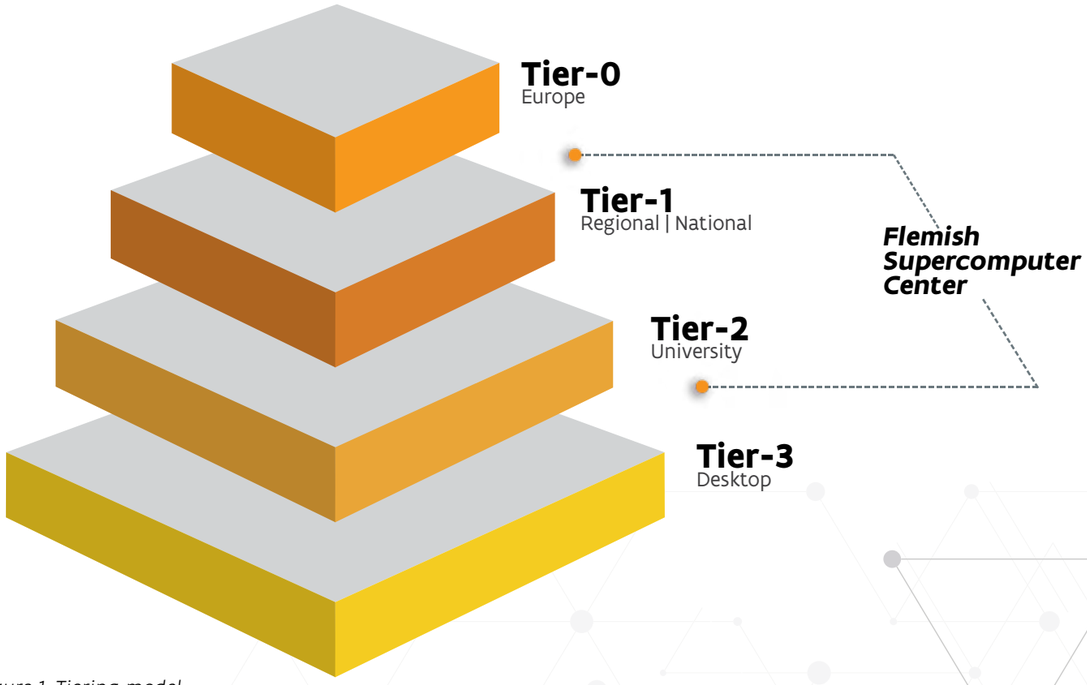

Figure 1. Tiering model

Structure and Operation of VSC 2.0

A new strategic plan and financing for a future-oriented and sustainable expansion of the Tier-1 level was approved at the end of 2022, so that users can continue to be supported in their need for computing power, and Flanders remains competitive in a rapidly evolving, big data-based scientific and technological world. At the same time, under the impetus of the European Commission through the EuroHPC Joint Undertaking, there is a strong focus on very high computing power, mainly in the form of a European (pre-)exascale Tier-0 level, but also on our own development of hardware technology, not in the least from geopolitical strategic considerations.

Given these recent developments, the VSC must grow into a redefined role within Flanders. Where the VSC has so far had a purely regional character, with a focus on users within the university associations, especially from a historical perspective, the user base within Flanders must be expanded more actively, and that expanded user base must also be encouraged and supported to take steps towards a European ecosystem of (pre-)exascale clusters and regional expertise centers, within which the VSC wishes to be a full partner.

In 2022, the VSC 2.0 future plan and its operational elaboration for VSC were approved by the board of directors. The guidelines can be summarized as follows:

Figure 2. Structure and operation of VSC 2.0

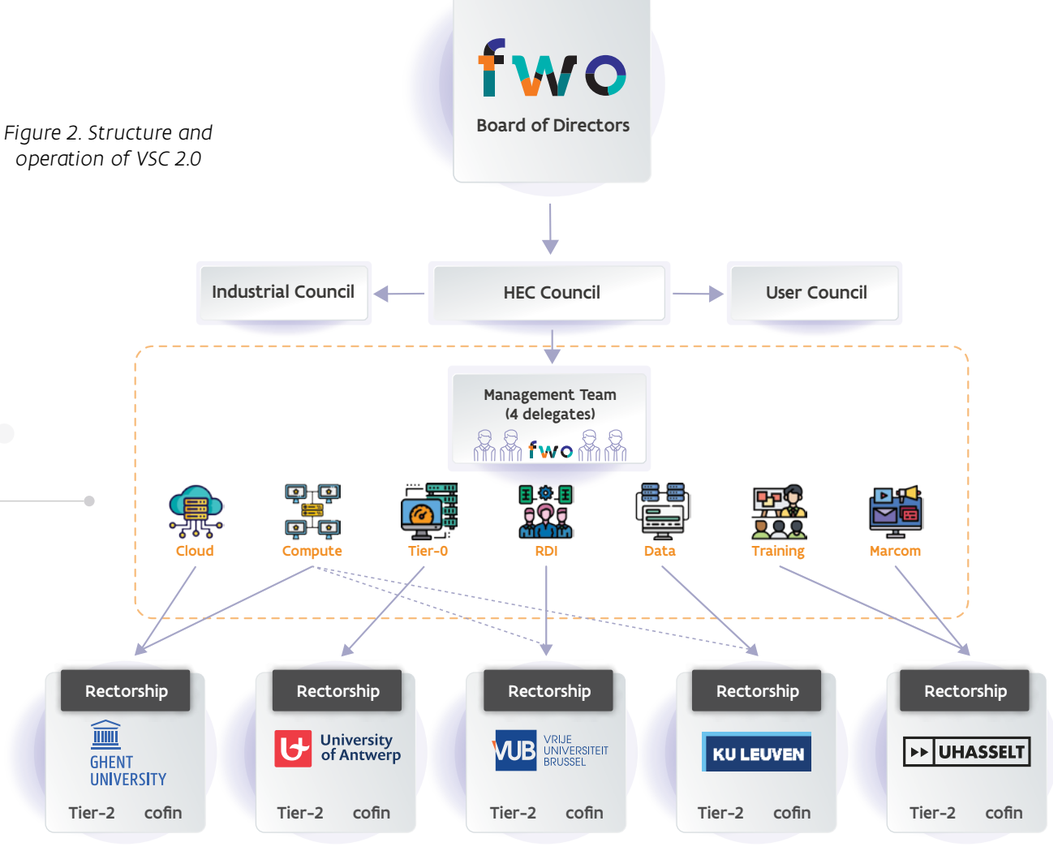

VSC 2.0 as a Regional Service Center

The VSC will remain a virtual, decentralized organization, under the supervision of and in collaboration with the FWO. The operational tasks are carried out by the five university associations, four of which are also infrastructure hubs. For the benefit of the growing group of users, VSC 2.0 will position its operational tasks more clearly, by streamlining contact with (potential) users through service points that function as a single point of contact, and with each association taking responsibility for the internal organization of a specific service VSC-wide. A service point will analyze a user's question and, if necessary, call help from more specialized experts within the VSC to survey: distribution of respondents offer a solution to the user. For example, KU Leuven will focus on data storage, UAntwerp on the European Tier-0 component, UGent on cloud computing and compute, UHasselt on marketing and communication and coordination of training, and the VUB will organize the service point for users outside the knowledge institutions. Organizations outside the broad group of knowledge institutions, such as companies and government agencies, contact a specific service point.

A Management Team for Daily Management

The existing governance with a High End Computing Council (HEC Council), an Industrial Council and a User Council will be retained, as will the positioning of the HEC Council within the FWO. A four-member Management Team (MT) is being established for daily management, with employees who are involved in the daily activities. In concrete terms, each hub will delegate one employee to the MT.

Vsc User Survey

To measure the satisfaction of the various VSC services, all VSC users were invited for the first time to participate in a survey between November 7 and December 18, 2023. The Tier-1 platform (compute, data and cloud), the Tier-2 platform on the four hubs, the training offering and communication were discussed. This survey will be repeated annually so that we see a (hopefully positive) evolution in the results.

Of the 467 responses, 372 were complete. As can be seen in the graph below, there was a nice spread between the various institutions, research institutions and companies. About a hundred respondents left their email addresses so that we can contact them.

Judging from some responses, we infer that services may not have been described clearly enough and/or we are unsure whether respondents actually used a particular service. In future surveys we will no longer ask for suggestions for improvements in general terms, but per service.

Later in this annual report, the various hubs discuss the results for the specific service(s) for which they are responsible. The management team examines how to act on the specific feedback.

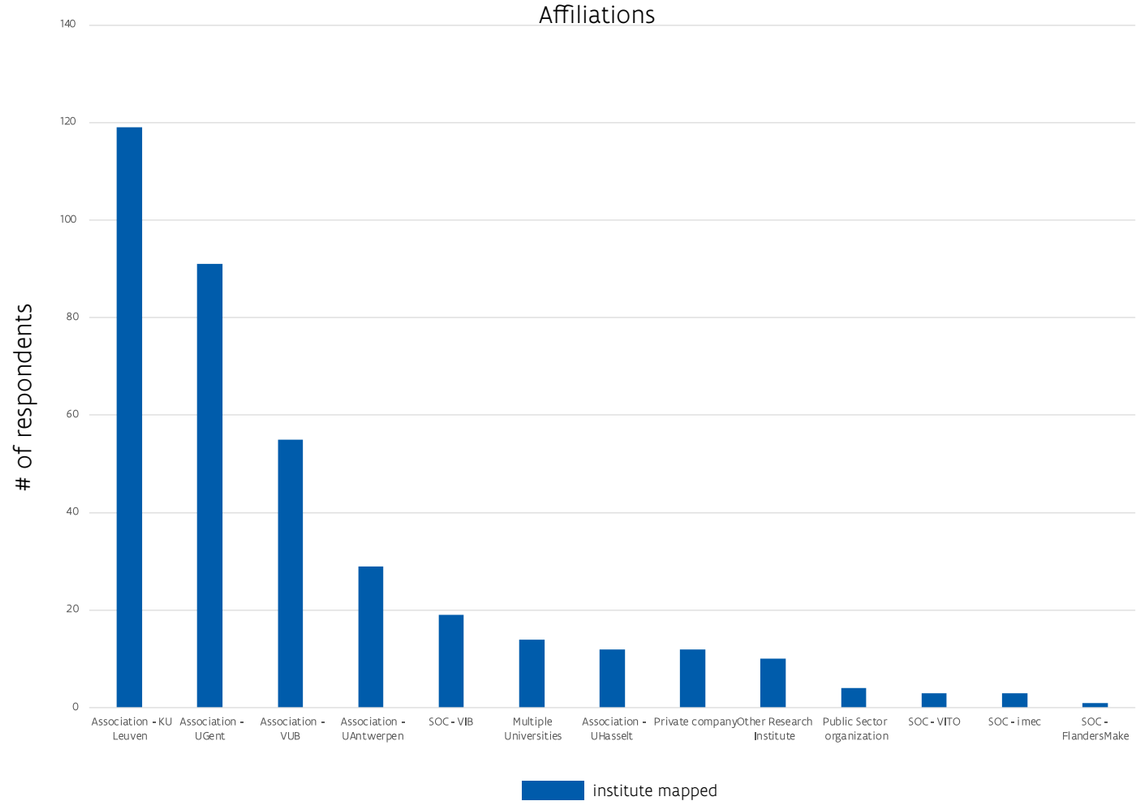

Figure 3. VSC User Survey

Tier-1 Supercomputer Platform

Tier-1 Compute

Tier-1c

During 2023, several improvements were carried out in and around the VSC's Tier-1 Compute setup Hortense.

Infrastructure

Adjustment of UGent data center for Tier-1 Compute Hortense

In order to optimize energy consumption, an additional cooling installation was installed in the Ghent University data center, where Tier-1 Compute Hortense is housed.

Additional dry coolers were placed on the roof of the S10 data center of Ghent University. Conduits through the roof and pump units were installed in the HPC room, and electrical works were finalized. The new cooling was brought online during the maintenance window of June 26-30, 2023.

After this, some cooling stability issues were noted, compounded by the higher outside temperature over the summer. For this reason, the cooling control was further tweaked, eventually making it possible to bring all Milan nodes reliably online in mid-August 2023. During Q4, another minor leak in the cooling circuit was identified and repaired. This improves the stability of the cooling.

Phase 2 Tier-1 Compute Hortense

The second phase of Tier-1 Hortense was also installed during 2023. A second partition of 384 CPU compute nodes was added, based on the AMD 7763 Milan CPUs. An additional GPU partition was installed, with 20 accelerator compute nodes, each equipped with four A100-80 accelerators (80 in total). The parallel file system capacity was approximately doubled, from 3 PB to 5.4 PB. The network infrastructure (incl. InfiniBand) was modified to accommodate these changes.

In addition to the hardware, these new partitions were tested, installed and configured so that they are usable by the end user. On May 15, the login nodes of Tier-1 and the various partitions (including new partitions gpu_rome_a100_80 and cpu_milan) were gradually brought online.

All ongoing Tier-1 projects were given 'free' access to cpu_milan until the next cut-off round (early July), in order to properly stress test this new partition under load.

At the kick-off meeting of May 26, 2023, phase 2 of Tier-1 Hortense was finally announced in production.

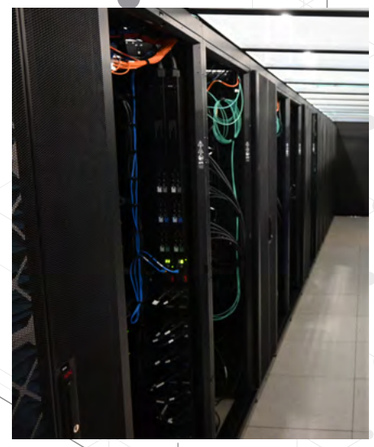

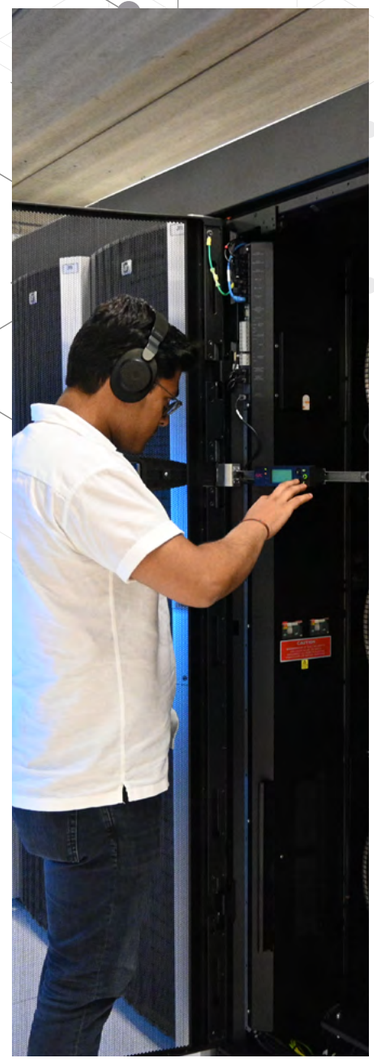

User features and backend developments

Debug queue

A specific 'debug' queue was added to the Tier-1 Compute cluster Hortense. The jobs sent to this queue start immediately. here is no such guarantee with other queues, where the (un)availability of resources can result in waiting times before jobs start running. This queue is ideal for end users to debug and test jobs, workflows, etc.

Further Developments to the Tier-1 Compute Project Portal

Tier-1 Compute Hortense users can easily consult their usage, available resources and runtime on the Tier-1 Compute project portal (https://resapp.hpc.ugent.be). For new features and bug fixes, a cycle of adjustments was carried out on this resource application: a total of 25 pull requests were completed, including bugs and smaller features. A notable implementation in the backend is that VSC staff can now easily enter and create new Tier-1 Compute projects via the resapp web portal as a single point of information. Automation then ensures that these projects are created on Tier-1 Hortense, that access is set for the indicated vsc-ids, and that allocated computing time and storage space is provided.

New Developments on the VSC Account Page

All account and group management within the VSC is centralized within 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 is maintained by Ghent University.

In addition to a whole series of bug fixes and feature implementations, steps were also taken to give more research institutions easier access to VSC infrastructure. The identity provider was added as a source for the account page for several research institutions: Artevelde Hogeschool, HOGent, Hogeschool West-Vlaanderen, ILVO, INBO, RBINS, VIB. For some of these research institutions, authentication runs smoothly this way. For others it is less straightforward, requiring a highly technical dialogue with the IT department of the research institution involved. But the feasibility of this workflow was well mapped out, with a view to adding even more research institutions.

Scientific Projects

Below is an overview of the number of academic projects that were running at Tier-1 Hortense in 2023, divided by category and cut-off.

Category	#Projects	Allocated Computing Time	
		CPU Hours	GPU Hours
Cut-off 1 (6 February)	31	112 529 414	254 899
Cut-off 2 (5 June)	25	70 182 476	278 101
Cut-off 3 (2 October)	34	97 542 520	414 541
Starting grants	116	46 500 000	52 100

Table 4. Overview of academic projects at Tier-1 Hortense in 2023

Usage and Availability

The graphs below provide a day-by-day overview of the calculation time used on Tier-1 Hortense. The first plot shows the CPU consumption (in CPU hours), with the arrival of the second phase (dodrio_cpu_milan_cpu_hours) very visible. The use of the twenty high-memory Rome nodes (with 512 GB of memory per node) is shown in orange, and is clearly proportionally lower. The plot below shows the GPU consumption (in GPU hours).

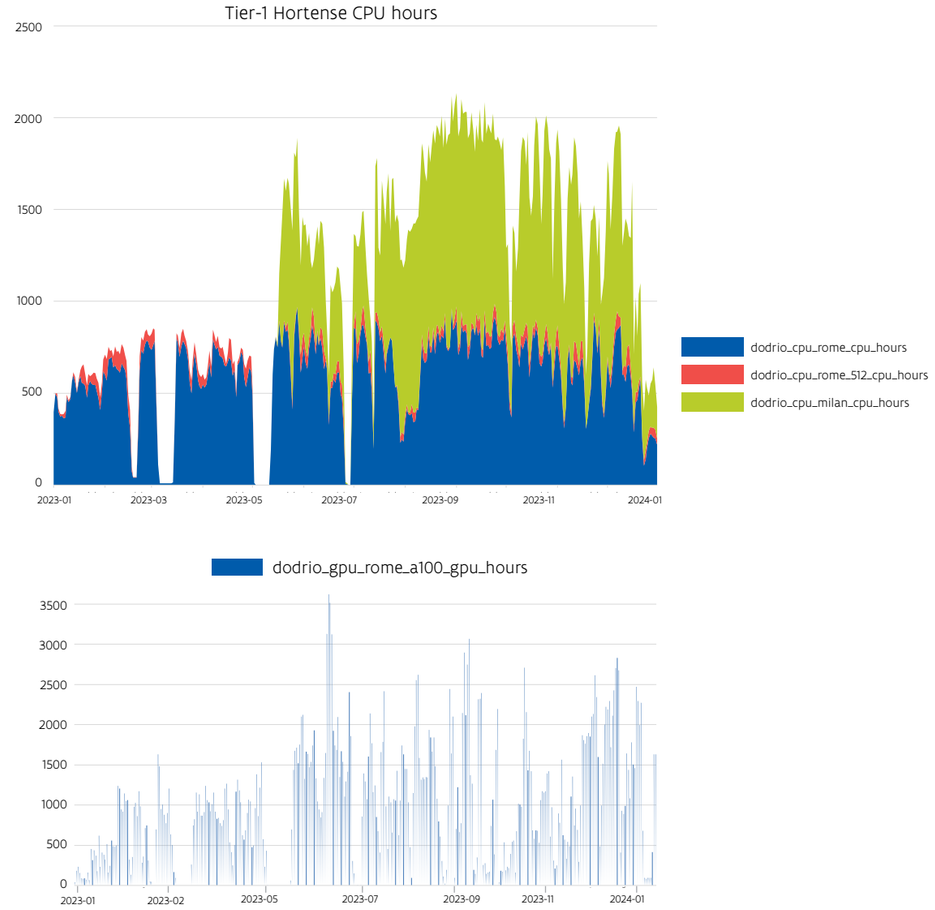

Figure 4. Overview of CPU and GPU use of Hortense.

The introduction of phase 2 and the new cooling system unfortunately led to several unavailabilities, especially in the first half of the year. These are also clearly visible in the graphs above.

- 17-21/2: cooling system instability (unplanned)
- 6-15/3: planned maintenance
- 2-15/5: planned maintenance
- 6-7/6: Cooling system instability (unplanned) – all Milan nodes unavailable
- 8-26/6: cooling system instability (unplanned) – part of Milan nodes unavailable
- 26-30/6: planned maintenance
- 12-13/7: cooling problems (unplanned)

After tweaking the cooling over the summer, the system operates properly again in the second half of the year.

Helpdesk

Several helpdesk queues are an indication of the efforts of Ghent University for the Tier-1 Compute platform.

Queue	#processed and closed tickets
Tier-1 Compute compute@vscentrum.be	616
Software installations	293
Industry (UGent partim)	121

Table 5. Tickets Processed at Tier-1 Compute

Through the dedicated Tier-1 Compute helpdesk (compute@vscentrum.be), users can report problems or questions regarding. Tier-1 Compute Hortense. A breakdown of the time it took for each logged question to be resolved is shown in the following graph.

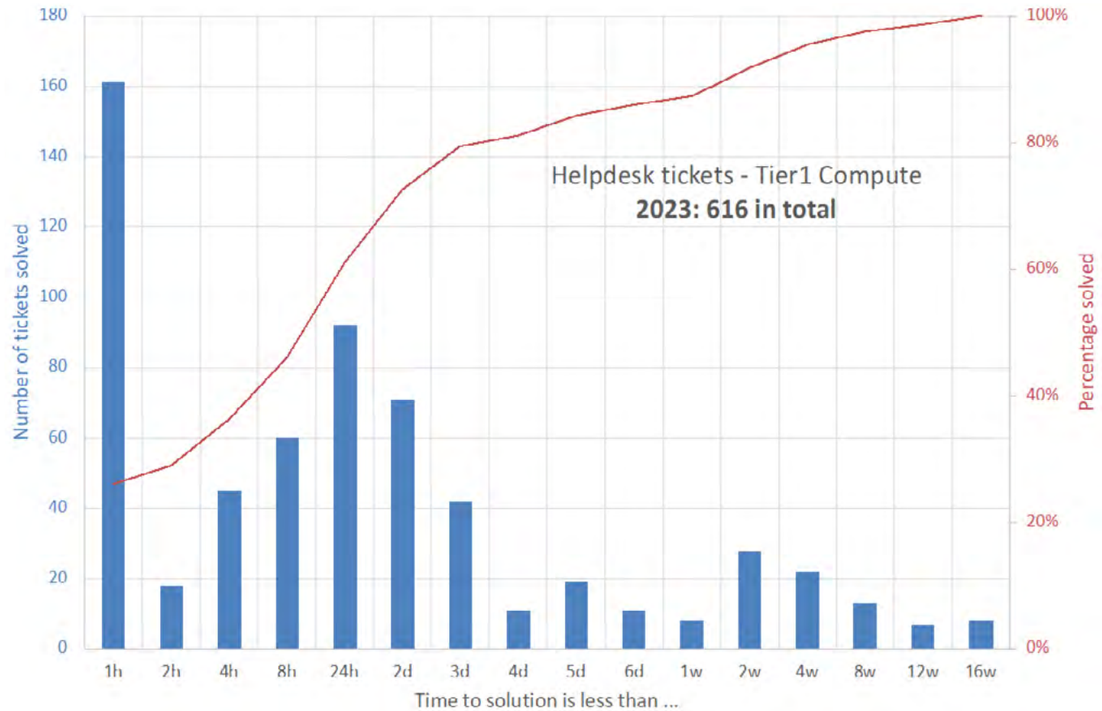

Figure 5. Resolution times for questions submitted to the Tier-1 Compute helpdesk

Furthermore, in 2023, a total of 293 software installation requests were handled by the Ghent University team. A strict division between questions for Tier-1 Compute and the Tier-2 UGent setup is not easy, which is why all tickets are also reported here. As a result of these questions, 3,514 software packages (including dependencies) were installed and/or updated, of which 2,201 for Tier-1 and 1,313 for Tier-2. The response time for software installations is noticeably longer than for other tickets, and is discussed in more detail with the Tier-2 Ghent University setup.

Finally, the Ghent University team also took on several tickets for industry users of Tier-1 Compute, in order to facilitate an optimal start-up of the RDI team (VUB). This mainly concerns the management of 116 vsc-ids for companies, the account management of which is still largely at Ghent University.

User survey

User satisfaction for the Tier-1 Compute infrastructure was part of the VSC user survey at the end of 2023. 153 different respondents indicated that they used this VSC service to a greater or lesser extent.

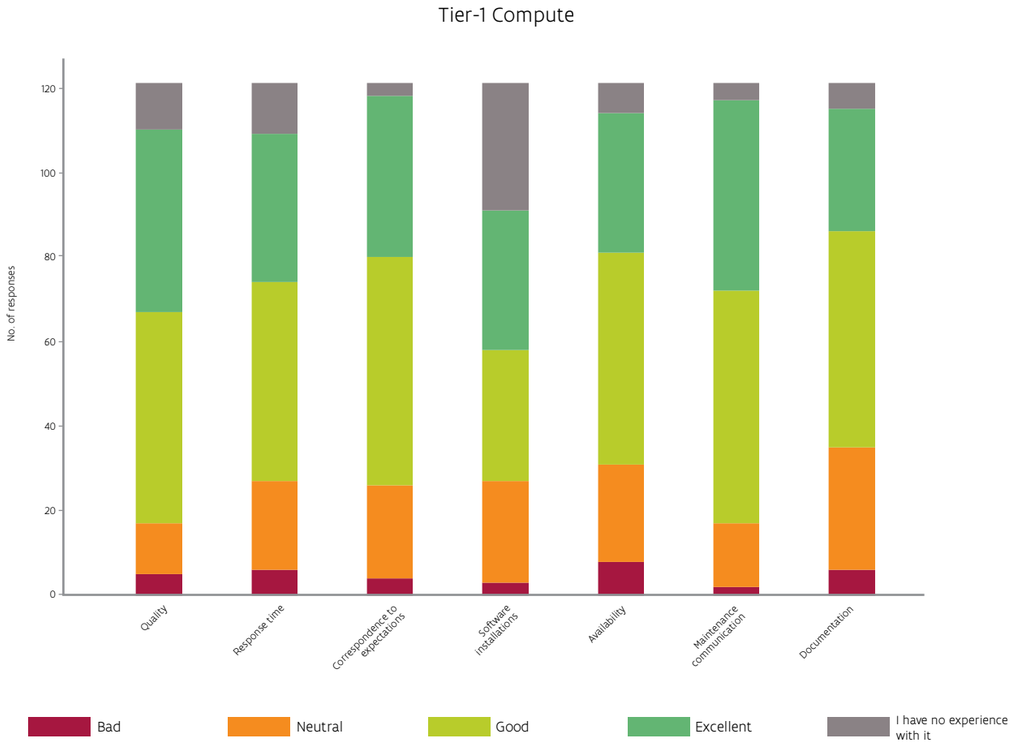

Figure 6. Tier-1 Compute User Survey

The following components of the Tier-1 Compute infrastructure were rated as 'good' or 'excellent' by users:

- 85% (93/110) - quality
- 75% (82/109) - response time
- 78% (92/118) - agreement with expectations
- 70% (64/91) - software installations
- 73% (83/114) - availability
- 85% (100/117) - communication regarding maintenance
- 70% (80/115) - documentation

Possible suggestions for improvements were listed:

Software

- faster response time for software installations (2 respondents)
- documentation about miniconda (1 respondent)
- info about toolchain compatibility (1 respondent)

Infrastructure

- no alternative (for industry) when Tier-1 unavailable (1 respondent)
- more nodes with more memory (1 respondent)
- increased stability of the web portal (1 respondent)

Documentation

- more support for Slurm (3 respondents)

As a possible suggestion for new developments it was mentioned:

- consider developing two redundant Tier-1 Compute setups (2 respondents)

Tier-1d

With the VSC 2.0 plan and the associated stimulus funding, we have obtained the resources to purchase a new Tier-1 supercomputer. This new machine will serve as the successor to the current system, Hortense. The first step in purchasing the new Tier-1 is determining the location.

Several universities showed interest in being responsible for housing this new supercomputer. After careful evaluation, a closed call has been launched within the VSC, based on the criteria also used for the LUMI supercomputer. Four evaluation criteria were used, each with a maximum of 25 points:

- Experience in purchasing, installation and management. The quality and experience of the universities in the field of comparable systems was examined. It is essential that the chosen party has sufficient expertise to effectively manage the next Tier-1 supercomputer.
- Quality of housing facilities. The infrastructure, security and connectivity to Belnet have been assessed. In addition, the plans to prepare the location for Tier-1 were examined.
- Estimated Total Cost of Ownership (TCO). The clarity and correctness of the estimated TCO and the methodology used to calculate it have been evaluated.
- Quality of service to users. The quality of service was examined to effectively operate a Tier-1 machine. The support procedures and tools used were also important.

Each criterion required a minimum of 13 points, and a total of at least 60 points had to be achieved. An independent panel of three experts conducted the assessment. They assessed each file individually and then reached a consensus. Their advice was submitted to the HEC Council on June 23, 2023 for approval.

KU Leuven and the Vrije Universiteit Brussel have each submitted a file. The panel awarded the Vrije Universiteit Brussel one more point. The HEC council has then decided that the Vrije Universiteit Brussel will purchase, install and operate the next Tier-1.

In the autumn, informative discussions with suppliers were started to find out what the possibilities are within the available budget. A delegation attended the Supercomputing fair in Denver, CO, USA. The intention is to start the procurement procedure from January 2024 with the aim of committing the purchase by January 2025. The machine should then become operational by November 2025.

Tier-1 Data

KU Leuven is responsible for the infrastructure and operation of the Tier-1 Data platform. The platform is positioned for active data processed on the VSC compute or cloud components. Publication and long-term storage take place on other platforms, but the availability of the metadata will facilitate the transition to this phase.

The platform went into production in April 2023.

Available Infrastructure

The Tier-1 Data platform is made up of four Distributed Storage Solutions for Spectrum Scale (DSS-G280) from Lenovo. It has a usable storage capacity of 27 PB (double copy), installed in two different data centers with synchronous replication (mirroring), so that the data is protected against major accidents. In addition, the system can use snapshots at file system level for additional data protection and data is encrypted at rest via software encryption.

There is also a Quality and Test environment around the infrastructure. This is used in the development of new features for the portal and new iRODS releases.

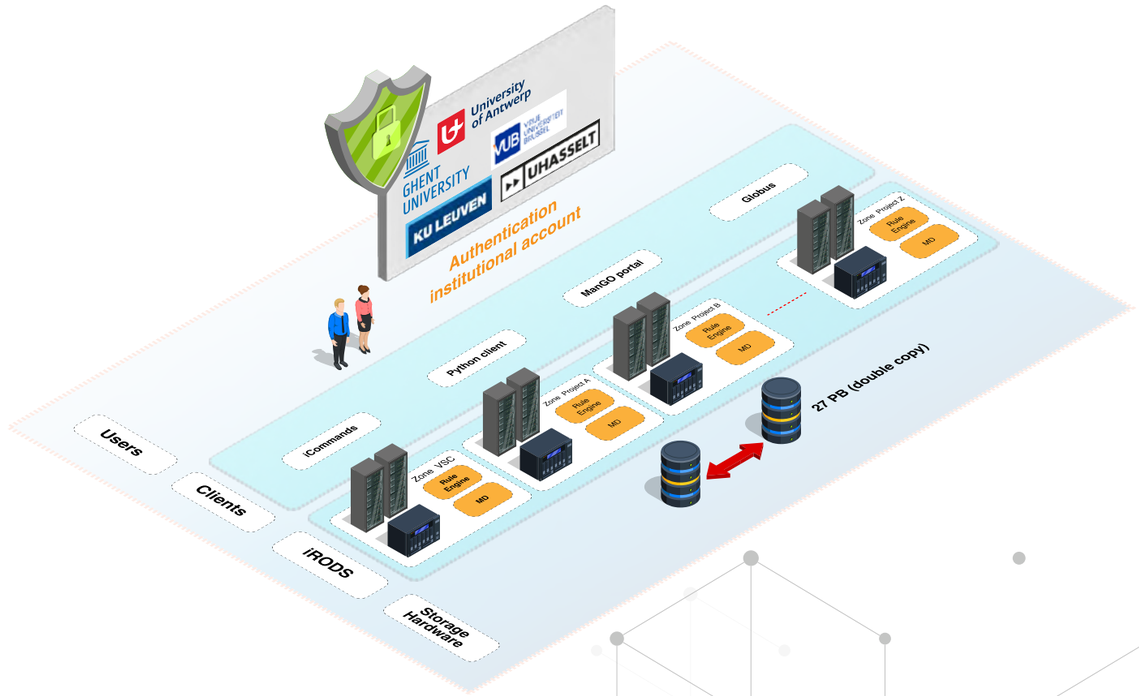

Figure 7. Tier-1 Data platform

Exploitation and Use

iRODS 4.3.1 was released in 2023. This technical upgrade included some features that were needed within the context of the VSC. For example, there is now support for i-commands on Red Hat, which is required for installation of the client on Tier-1 Hortense. The upgrade also includes support for OpenID authentication. After testing, the upgrade was rolled out in the production environment. The new clients were also rolled out on the various Tier-2s.

However, Tier-1 Data is not only accessible via command line clients. ManGO portal, the GUI environment on top of the iRODs infrastructure, was further developed by KU Leuven. Important new features are: basic management of delegated access rights, management of metadata with metadata schemas with version control, a schema editor, import/export of metadata schemas in JSON format, inspect & extract metadata from datasets with, among others, Apache Tika ... . In 2023, six main versions of the portal were released (v0.9, v0.10, v0.11, v0.12, v0.13 and v0.14) plus a number of additional minor and bug fix releases. This shows the rapid evolution of the ManGO portal to best meet the needs of researchers.

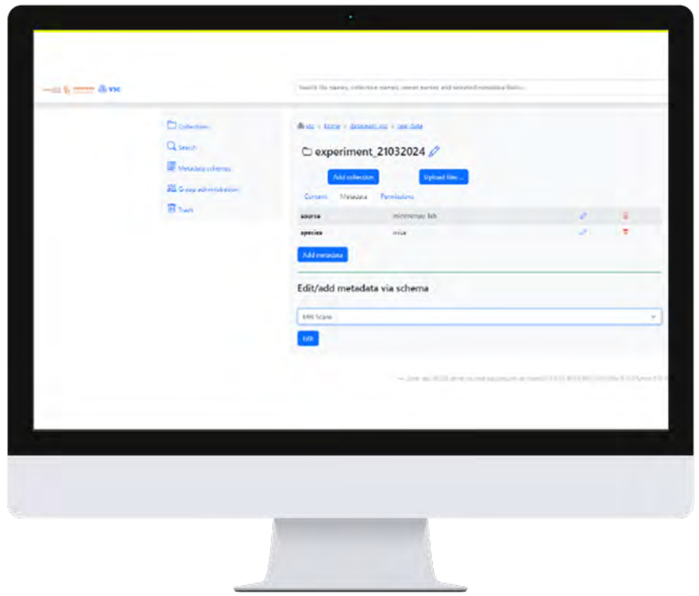

Figure 8. Screenshot of VSC Mango portal

Not only the portal, but also the entire service is constantly being developed. An SFTP ingestor tool was created so that data can be automatically uploaded to an ingress zone. This is a step in enabling workflows where, for example, data sets from experiments can be automatically uploaded to the Tier-1 Data platform. This new tool also helps to securely upload data from non-trusted devices such as mobile phones or autonomous camera devices directly to the Tier-1 Data platform.

Tier-1 Data Usage

The graph below shows the evolution of the data since production in the production environment. For the climate group, the migration of data to the production zone has not yet been completed. A total of 1.3 PB of data is stored.

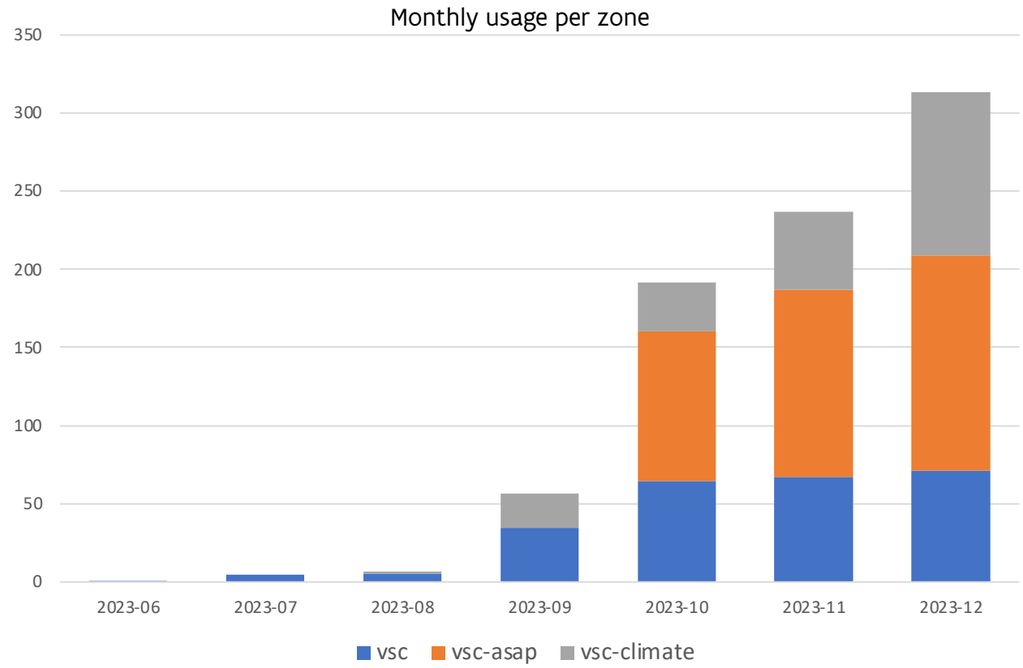

Figure 9. Tier-1 Data usage

Globus Rollout and Use

Globus online (www.globus.org) is a software solution to exchange data in a secure, fast and reliable way between different storage solutions. In 2022, Globus was fully rolled out to all VSC sites. As the graph below shows, the solution is used to perform transfers between different VSC components and a large number of endpoints. About 100 users moved a total of more than 2 PB of data with Globus. Use is expected to continue to increase in 2024.

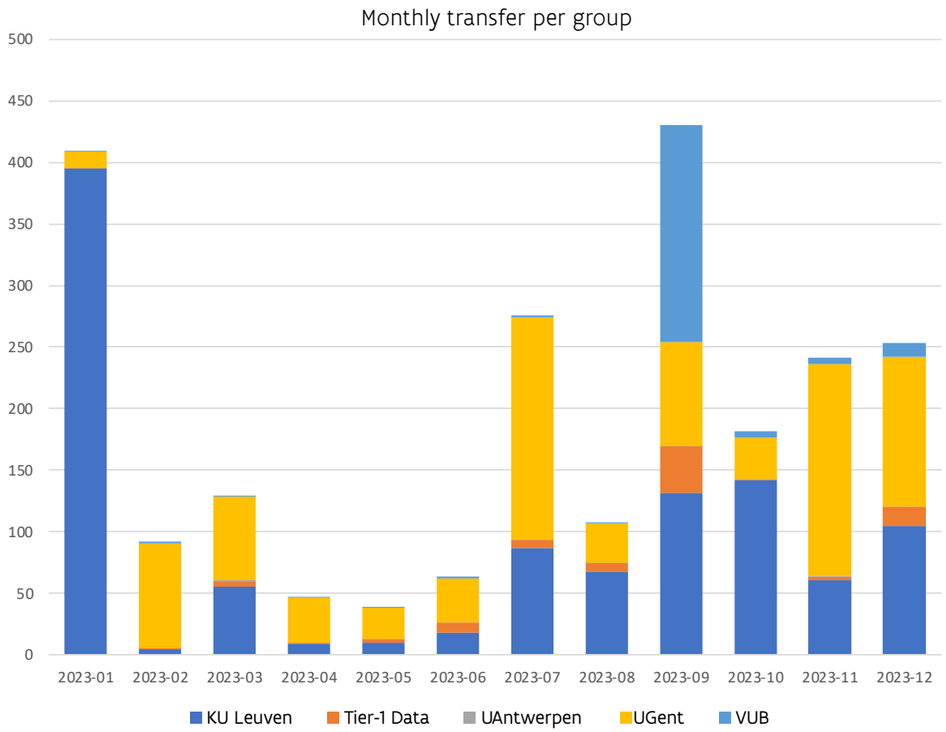

Figure 10. Globus deployment and use

User satisfaction was asked in the VSC user survey at the end of 2023. For Tier-1 Data, 79 different respondents indicated that they used this VSC service to a greater or lesser extent.

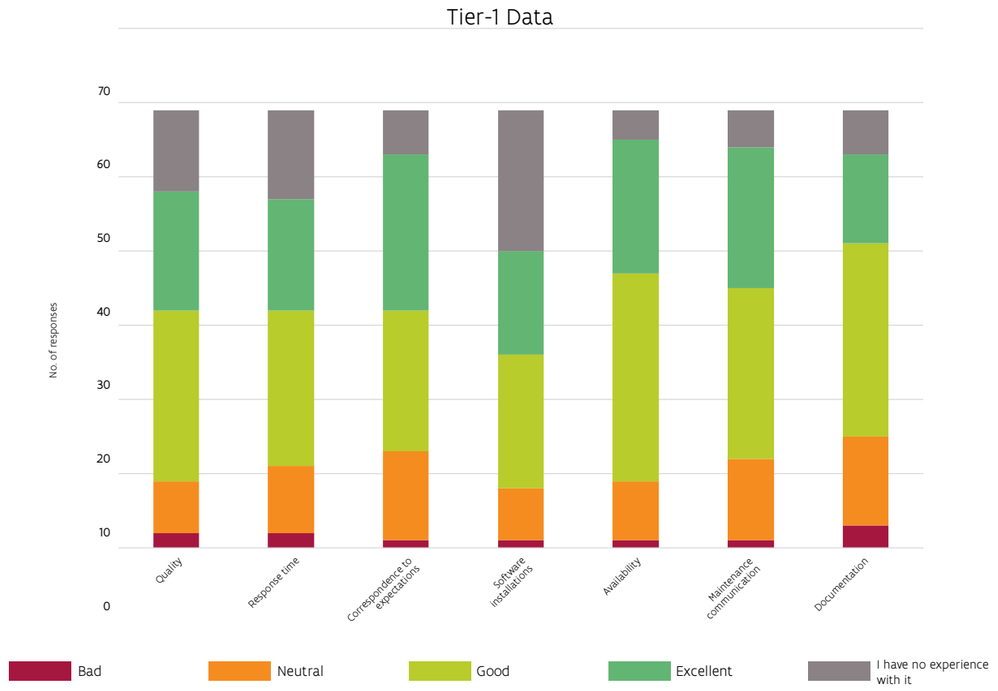

Figure 11. User Satisfaction Tier-1 Data

The components of the Tier-1 Data infrastructure below were rated as 'good' or 'excellent' by users:

- 83% quality
- 78% response time
- 76% agreement with expectations
- 79% software installations
- 83% availability
- 79% communication regarding maintenance
- 72% documentation

There are certainly some lessons to be learned for next year. The results are good, but attention to documentation with examples is an area for improvement. Next year, more specific comments will be requested for the data component. The question regarding software installations will be better formulated based on the usability of the various clients.

Allocating Storage Capacity

Tier-1 Data projects are awarded via project applications. In 2023, there were four project applications that will use a total of 730 TB of storage over a four-year period.

In addition, a new collaboration grant was requested by the storage4climate group, which had already built up 1 PB of data during the pilot phase. In this collaboration grant, it is expected that the required storage capacity will grow to 6 PB in the next four years.

User Support

In 2023, the focus was on the start of the production phase of Tier-1 Data. The pilot users were encouraged to submit a request for regular access and to migrate data to the production environment. The climate group is the largest user here.

With the launch, the documentation on the VSC website was also refined. An exploratory discussion was held with various research groups. This has led to a number of starting grants and complete project applications.

Specific Support

On May 9, the production environment was officially launched via an online launch session. About 90 participants followed the session. The technical infrastructure was explained, followed by a demonstration of using the online portal and via the Python clients.

Tier-1 Data was used in two hackathons, including one from Digital Humanities. The participants were able to retrieve the data needed for the hackathon from Tier-1 Data.

The iRODS Consortium

During the annual iRODS UserGroup Meeting (iRODS), three presentations were given by KU Leuven and one Lightning Talk.

- "ManGO: A web portal and framework built on top of iRODS for active research data management"
- "Towards rich and standardized metadata in iRODS"
- "rirods: An R client for iRODS"
- "Dataverse integration dashboard: pulling data from iRODS"

Due to their involvement in the consortium, the supporters are not only well informed about the new developments, but collaborations can also be found to develop new ideas.

Showcase

Flanders BioImaging (FBI) is a consortium that provides infrastructure and services for imaging facilities spread across the Flemish universities. They produce a large volume and diverse set of imaging data linked to a large number of research projects. Research data management, sharing data and collaborating around this data are important aspects of the service. Tier-1 Data can play an important role in the workflow here. On the one hand, in addition to Tier-1 Data, VSC also offers the resources via the Tier-1 Compute and Tier-1 Cloud, and on the other hand it can also integrate with external compute facilities. At the same time, it provides a path to publish data to external repositories. The consortium applied for a starting grant in 2023 to test this workflow. Following this, a collaboration grant was also applied for.

A success story over Tier-1 Data was published on the VSC website and will be followed up with a report on the further progress of the collaboration grant.

Watch the video

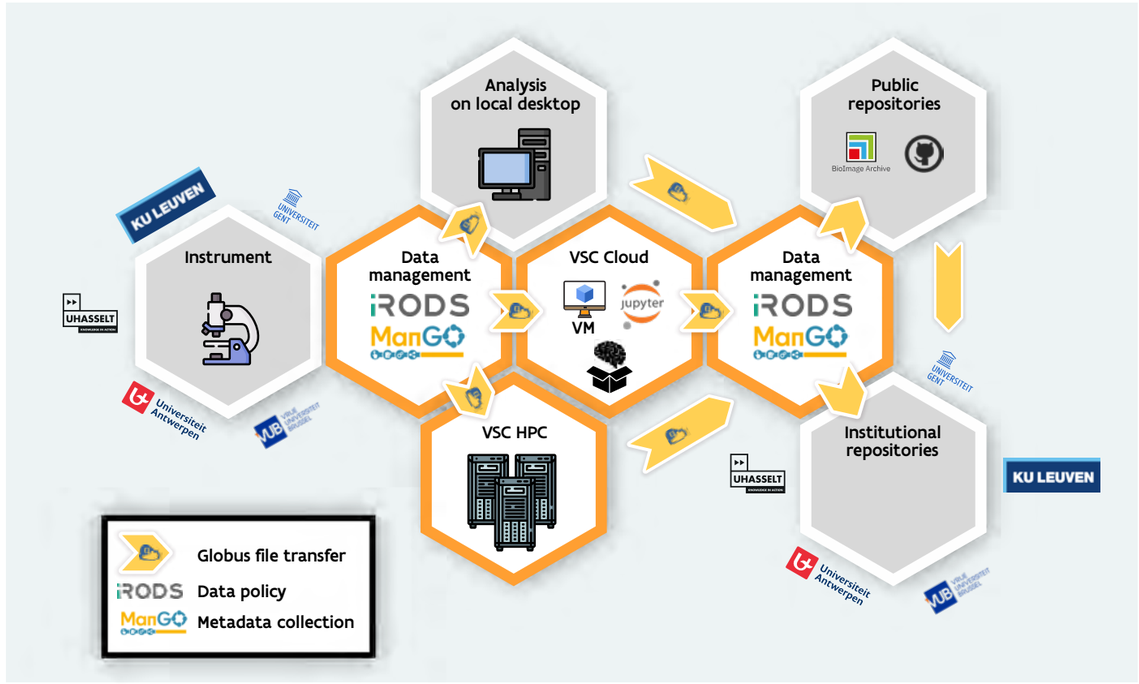

Figure 12. Flanders BioImaging workflow, Reference (Tatiana Woller)

Tier-1 Cloud

Researchers using the VSC Tier-1 HPC services also often need a more flexible environment in which they can use software or services in addition to a large-scale HPC system, such as custom software packages, interactive data analyses, workflow portals, data visualization and specific pre- and post-processing tasks.

The Tier-1 Cloud infrastructure addresses this need by implementing a cloud service based on OpenStack that offers 'on demand' resources in a more flexible and cloud-like manner. This is an Infrastructure as a Service (IaaS) facility that provides the capacity to deploy resources such as virtual machines (VMs), storage and networking with full control by the users of these resources. A catalog of VM and orchestration templates is available, making it easy to set up virtual machines with different flavors, a web server, a basic cluster, etc. with just a few mouse clicks.

The Tier-1 Cloud setup for the VSC has been running for some time now and has reached a specific target audience. In 2023, this VSC service was critically examined and it was investigated where improvements and simplifications are possible. The results of this review indicated that the underlying software solution, OpenStack, is a powerful but technically complex cloud platform. This often requires specific knowledge and regular (expensive) assistance from the vendor to perform basic tasks. At the same time, a review of the current Tier-1 Cloud projects shows that no special cloud facilities are required and functionalities are sufficient. Other platforms (e.g., OpenNebula) offer similar functionalities as OpenStack that are certainly sufficient for a full-fledged Tier-1 Cloud platform. The support costs of OpenNebula, for example, are lower and the configuration is also simpler. In the long term, a switch to another technical framework will therefore be appropriate.

There is currently a support contract with Red Hat for Openstack, until mid-2026, and the current Tier-1 Cloud infrastructure functions well and properly. It is not desirable to modify this setup before the end date is nearer. If a switch to OpenNebula (or another platform) is in order, preparations and efforts can be planned for this by 2025.

An important first step towards a possible platform switch has already been taken by the introduction of HashiCorp Terraform orchestration. This is an infrastructure as code tool that allows users to easily build complex cloud infrastructures using templates. This tool is also compatible with multiple cloud providers (e.g., AWS). The introduction of this orchestration provides a more platform-agnostic experience for users of the Tier-1 Cloud platform.

The Tier-1 Cloud setup was also examined in terms of hardware. This setup is made up of a mix of new hardware and older (ex Tier-1 Compute) hardware. With the phasing out of BrENIAC in 2023, newer Skylake nodes became available within the VSC and it was investigated whether they could replace/upgrade older Broadwell nodes in the Tier-1 Cloud setup. Based on a cost-benefit analysis, it was decided not to carry out this replacement. After all, the Tier-1 Cloud CPU hypervisors are used to a limited extent, where the limited added value of (slightly) more recent CPUs does not outweigh the downtime that this operation would entail.

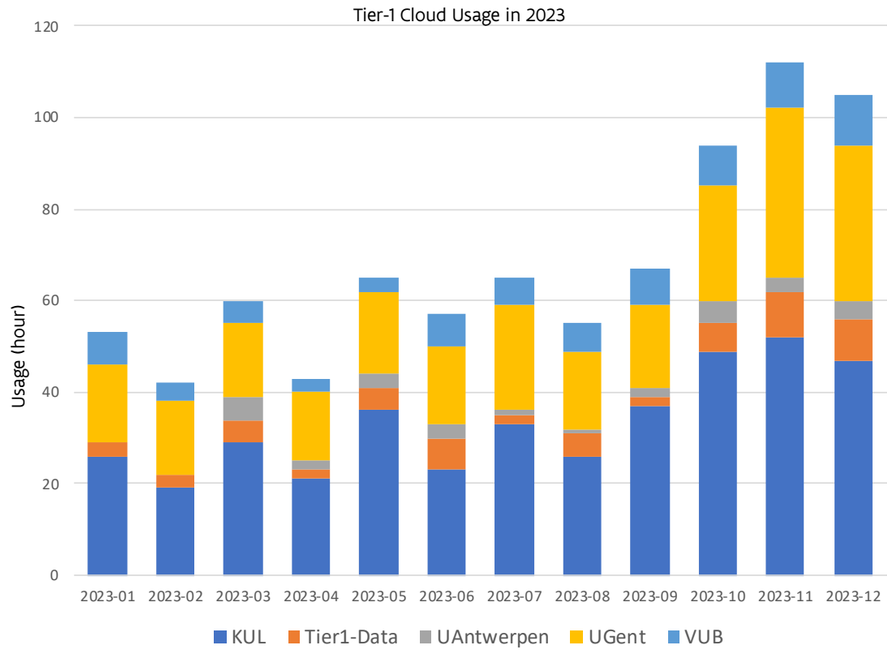

Figure 13. Overview of Tier-1 Cloud in 2023

Fairly steady activity is noticeable for all Tier-1 Cloud projects in 2023. A sharp increase in the number of active virtual machines in November and December can be linked to one larger cloud project that had a burst in activity. Ten new projects started in 2023, including seven starting grant projects and three full proposals. A total of 4,593,852 CPU hours were consumed on the Tier-1 Cloud platform in 2023.

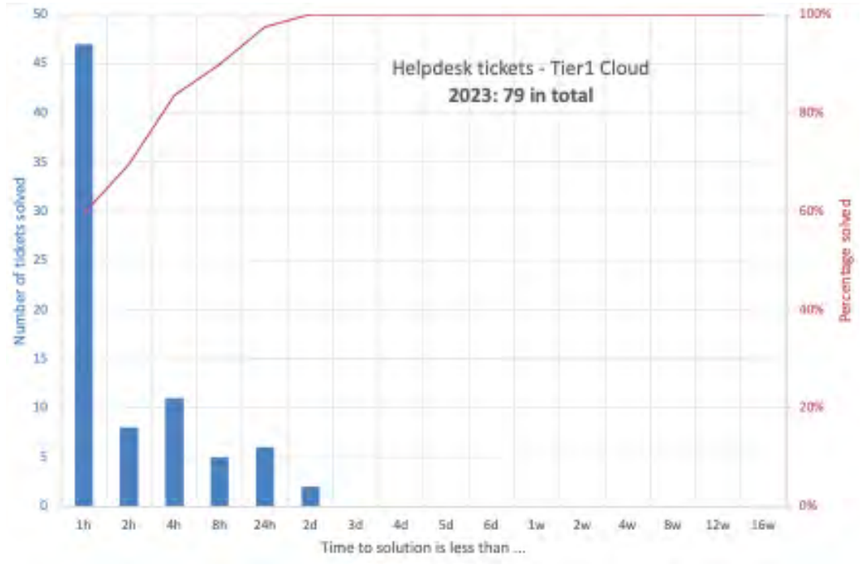

Figure 14. Activity and usage of Tier-1 Cloud projects in 2023

The dedicated Tier-1 Cloud helpdesk closed a total of 79 tickets in 2023. The time to resolve a ticket was very low, with the majority of tickets already resolved within one hour.

User satisfaction – as surveyed in the VSC user survey at the end of 2023 – for the Tier-1 Cloud infrastructure is quite high. 82 different respondents indicated that they use this VSC service to a greater or lesser extent. The following components of the Tier-1 Cloud infrastructure were rated 'good' or 'excellent' by users:

- 92% (36/39) - quality
- 82% (32/39) - response time
- 88% (38/43) - agreement with expectations
- 83% (24/29) - software installations
- 86% (38/44) - availability
- 83% (35/42) - communication regarding maintenance
- 76% (32/42) - documentation

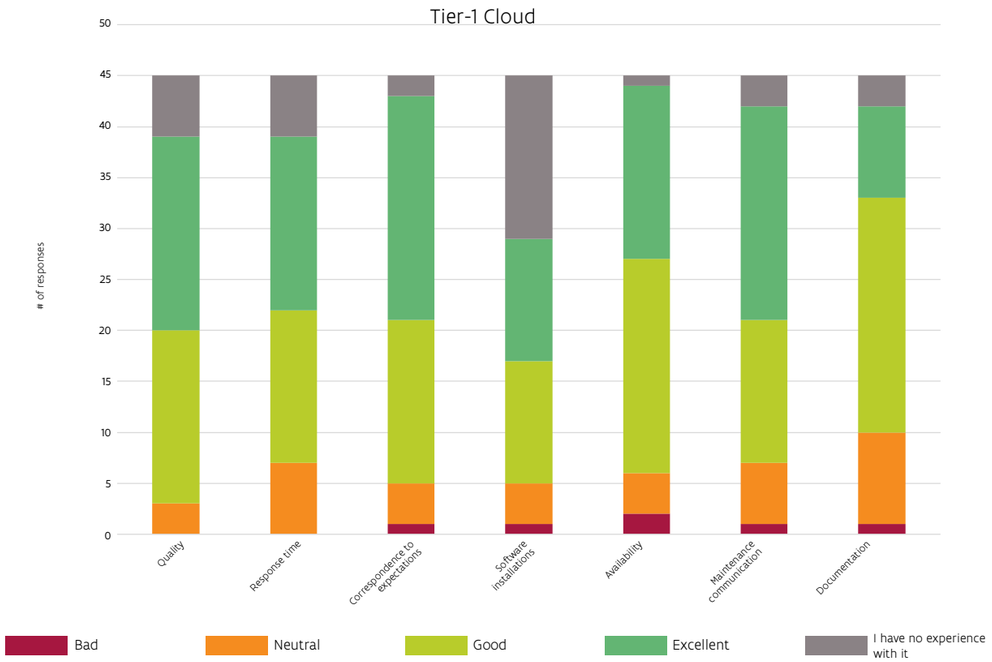

Figure 15. Tier-1 Cloud User Satisfaction

Possible suggestions for new developments were:

- Cloud services for teaching (1 respondent)
- Sandbox servers for research groups (1 respondent)

RDI

With the VSC 2.0 plan, the Vrije Universiteit Brussel was designated as the Research, Development & Innovation office of the VSC. This office was started in 2023. A budget for 2 FTE was provided in the VSC budget. The most important achievement of 2023 is therefore the successful recruitment of 2 people in June. Someone with the profile of a business developer and someone to support the RDI users was recruited. These people are settling in.

The tasks that were previously performed by various VSC hubs (mainly the Tier-1 hosting institution) were transferred to the VUB. The support for RDI users was split into a first and second line: the first line is handled by the VUB, the second line by the Tier-1 hosting institution (UGent). Agreements were made regarding the follow-up of questions and there is a monthly meeting to discuss the state of affairs. For RDI users, everything is transparent: their only point of contact for support is still (compute@vscentrum.be) and the Ghent University ticket system processes all questions. In addition, the intention is that (info@vscentrum.be) will be used as the single point of contact for new RDI users.

Most of the work in 2023 focused on the legal side of the matter: a new agreement has been concluded between the FWO and the various university associations that appoints the VUB as administrative coordinator for RDI. A lot of time has also been spent drawing up new contract templates. Because the hosting of the Tier-1 compute cluster has now been split from the administrative coordinator, a rework was required. The new contracts are now being signed by the company, the FWO, the VUB and the hosting institution of the Tier-1 (UGent).

Due to the expansion of Hortense with the Milan partition, an adjustment in the price for use was necessary. Based on the cost of the hardware, the reserved power, the electricity consumption and the human resources, the new cost price amounts to:

	Price
CPU	€ 24 / node / day
GPU	€ 94 / node / day
Storage	€ 16 TB / month (from > 1 TB)

Table 6 Cost for Tier-1 Compute

All prices include VAT. The actual billing is done at the level of CPU or GPU hours. Storage capacity is based on requested capacity, not what is in use.

There is a free exploratory grant worth € 5,000 to give new customers a taste of our services. One can divide this amount between CPU and GPU resources as desired.

Even though the focus in 2023 was on smoothing the administrative flow, we took various opportunities to market our services. For example, a poster was presented on the Elixir national day. Based on the results of the Stretch Innovation study, which was ordered by the industrial council of the VSC to explore the HPC potential in Flanders, a large number of leads were followed up. In addition, a partnership was concluded with Make-IT-Fit to allow Limburg companies to also taste the VSC services. This already gave a number of leads.

Now that the foundation of the RDI office is in place, the focus will shift to active outreach. For example, work will be done on a service catalog and customized training for RDI users in 2024.

Customer Base

Eight exploratory grants were started in 2023. Five exploratory grants have also been converted into paid contracts in the last two years. The companies mainly use Tier-1 to perform fluid dynamics calculations.

In addition, ten paid contracts were also concluded. Invoices worth € 718,304.33 were sent out. The income in 2023 was abnormally high due to the one-off very extensive use by one company.

	Number of Exploratory Grants
2020	7
2021	8
2022	13
2023	8

Table 7. Tier-2 Overview of Exploratory Grants and Paid Contracts for VSC Tier-1 Compute in 2023

Support

Included in the use of Tier-1 Hortense by RDI users is also user support. This is organized into a first and second line via (compute@vscentrum.be). In 2023, a total of 273 tickets were handled and resolved for RDI users.

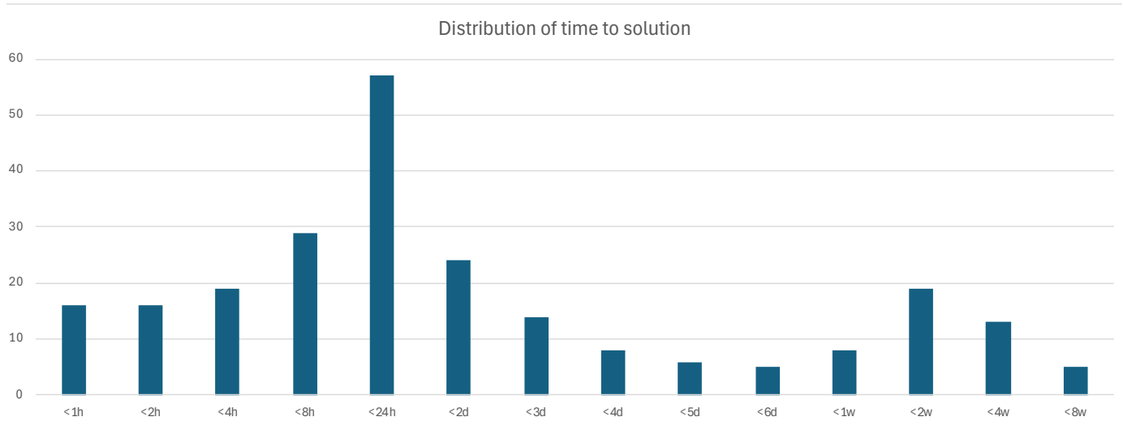

Figure 16. Resolution time for support tickets for Tier-1 Hortense RDI users in 2023

Tier-2 Infrastructure

This section provides an overview of the Tier-2 infrastructure that is available within the various Flemish universities. Their use is also illustrated.

University of Antwerp

Available Infrastructure

In 2023, the Vaughan cluster was expanded with 40 additional nodes and the capacity of the scratch file system was expanded to 610 TB. In addition, the old "Hopper" nodes were removed from the system and replaced by 24 nodes from the former Tier-1 BrENIAC.

The Tier-2 infrastructure consists of the Vaughan, Leibniz and BrENIAC clusters:

- 681 TF CPU, 106 TF GPU
- 17.188 CPU cores
- 78 TB memory
- 12 GPU devices / 72704 GPU cores

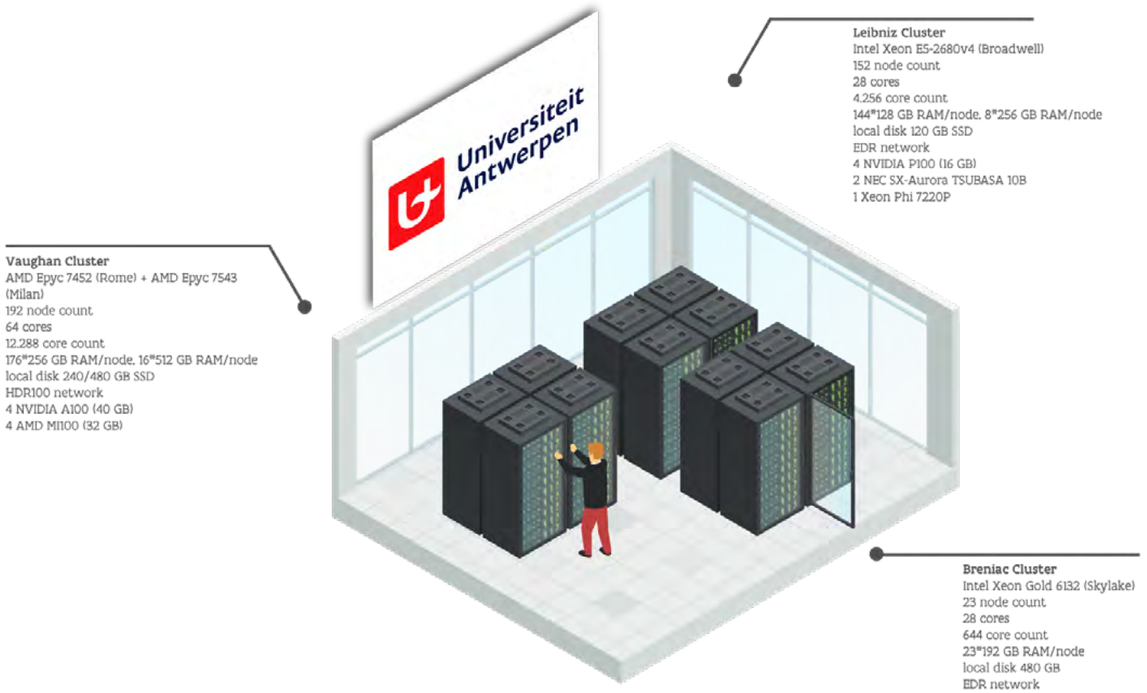

Figure 17. Tier-2 infrastructure at UAntwerp

Exploitation and Use

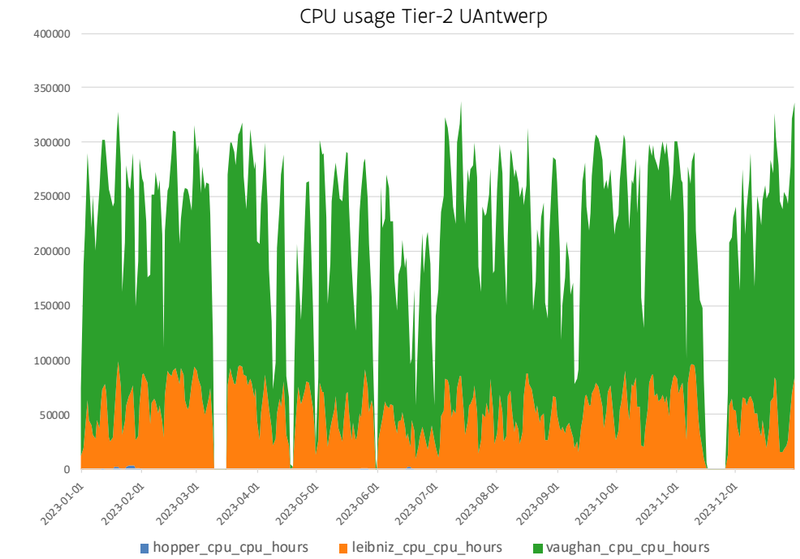

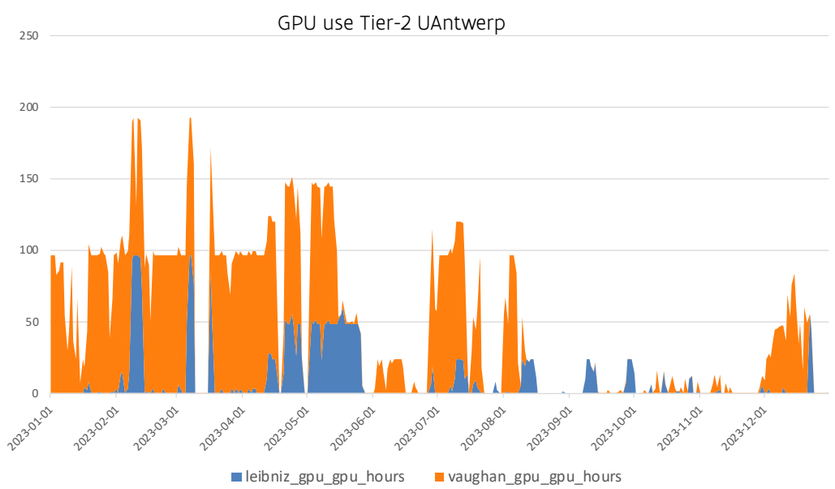

Figure 18. Overview of CPU and GPU usage at Tier-2 UAntwerp

Early March: unforeseen major power outage, complete system down. When restarting there were problems with the scratch file system. Downtime: 1 week.

Mid-April: planned interruption to expand the cooling system in preparation for the installation of new nodes. Downtime: 2 days.

End of April: unforeseen interruption due to problems with the Infiniband network. Downtime: 4 days

End of June: unforeseen interruption due to a leak in the cooling circuit. Downtime: 1 day.

Mid-November: planned interruption that we will coincide with work on the high-voltage cabin on campus. Due to problems with the upgrade of the scratch file system, the interruption lasted longer than expected (due to copying files and reinstalling the file system). Downtime: 2 weeks.

Allocating Calculation Time

In 2023, researchers from UAntwerp and its association had completely free access to the Tier-2 infrastructure. At the request of the Department of Research, Innovation & Valorisation at the University of Antwerp, paid compute time is being introduced. This will happen from spring 2024. Various information sessions were organized on this subject.

> "The VSC supercomputers allow us to study the aerodynamics of airborne wind energy systems in detail. Using 392 cores, the simulation takes approximately 8 hours to complete one crosswind flight loop."
>
> **Niels Pynaert**
> **Faculty of Engineering and Architecture**
> **UGent**

User Support

Category	Tickets
Accounts	89
Software	157
Tier-1	11
Other	100
Total	357

Table 8. Tier-2 handled tickets

User support is central. We are thinking of tips for optimizing workflows or calculations and the efficient use of the infrastructure through active monitoring of the jobs.

In 2023, a news section was introduced on our website, where smaller announcements are posted in a quick and accessible way.

We are also looking further into integration with Globus, with possible extensions to OneDrive, and iRODS to give users more options in terms of data transfer.

Initially, user support is provided via the ticketing system.

In addition, there were also 1-on-1 conversations or conversations with research groups to discuss questions.

Also important in the context of support are the semi-annual introduction courses that are organized (February/March and October): "Linux introduction" (2 half days), "Supercomputers for starters" (1 half day) and "HPC@UAntwerp introduction" (2 half days).

In addition, there is guidance for master students in using the infrastructure, and support with requests for computing time at Tier-1 and Tier-0 and with computational aspects of (inter-university) project applications.

In 2023, the software stack was significantly expanded and brought more up to date.

We switched from our own custom build scripts in the 2020 Intel environment to shared (EasyBuild) build scripts (which are also in use on the other VSC hubs). In addition to more recent versions of many packages that were previously available, this change has also allowed us to install new software more quickly and allow us to work more closely with the other VSC hubs. These more recent versions (and new software) are spread across two environments (Intel and Foss) for the various CPU types in our clusters, including GPU-supporting packages.

In addition to the courses within the VSC, courses are also provided in the regular program, namely "Scientific computing environments" and "(Parallel) programming".

VSC Survey

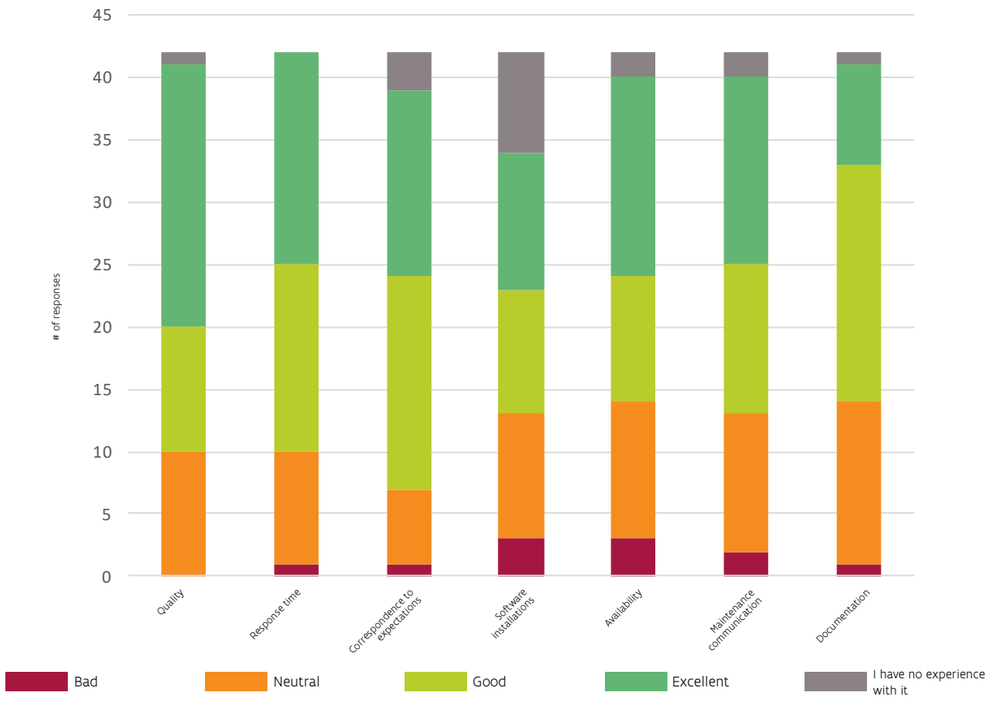

Figure 19. Tier-2 survey UAntwerp

The survey identified four work points: software installations, availability, communication regarding interruptions and documentation.

As described above, over the course of 2023 we have changed the way we install software. We hope to see a positive evolution in the 2024 survey.

Some unforeseen interruptions have led to a certain dissatisfaction regarding availability led. Of course, we always hope to reduce this to a minimum. Related to this is the communication surrounding the unavailability. Unforeseen interruptions cannot be announced in advance, and after an interruption we prefer to perform the necessary testing, where users naturally hope to be able to work again as quickly as possible.

The VSC documentation, which is shared for all VSC users, has been significantly updated in 2023. However, users from Antwerp apparently view it a little more critically.

Vrije Universiteit Brussel

Available Infrastructure

The Tier-2 infrastructure at the Vrije Universiteit Brussel (VUB) looks as follows:

- 1 cluster / 6 partitions
- 235 TF CPU, 236 TF GPU
- 3.624 CPU cores / 166.912 GPU cores
- 25 TB memory

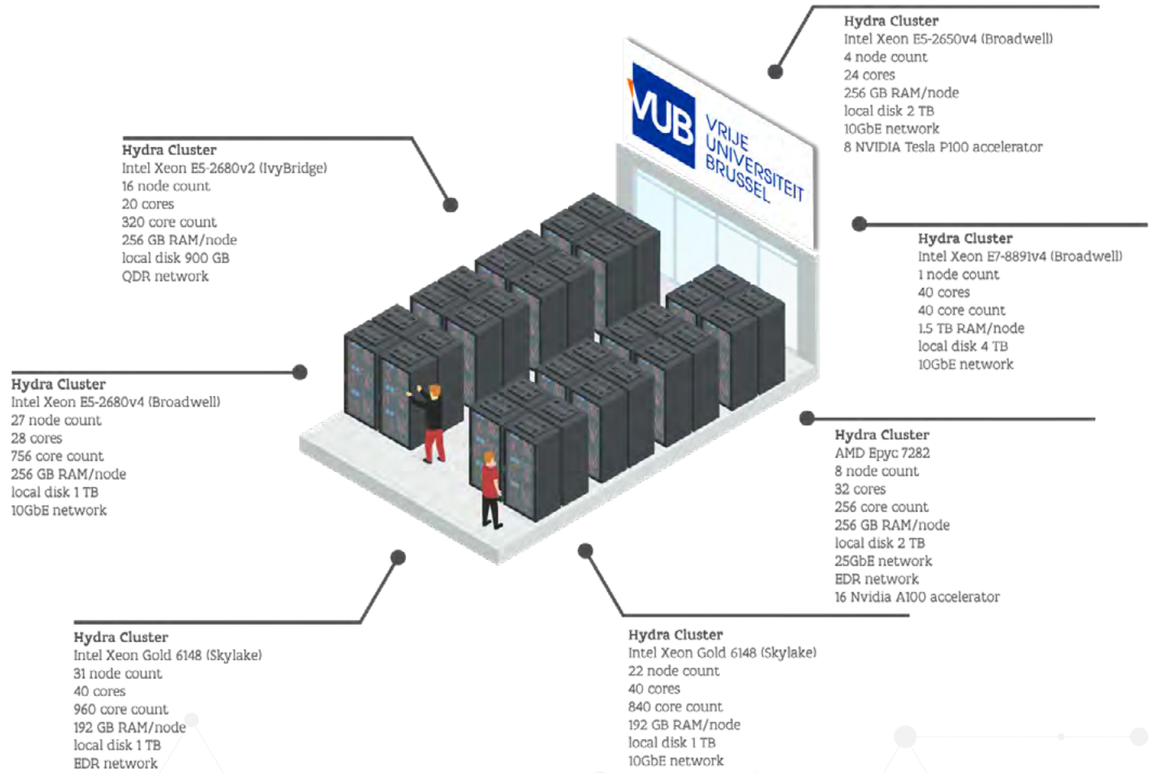

Figure 20. Tier-2 infrastructure at VUB

At the VUB it was decided to always do extensions within the same Hydra environment, which is more efficient for both the users and the management team. This results in a more heterogeneous cluster. Within the environment, an attempt is made to hide this heterogeneity from users as much as possible by, among other things, offering identical software installations and assigning jobs to multiple partitions by default.

The following changes were made to the Hydra environment in 2023:

- The last nodes of the Ivy Bridge generation have been taken out of service.
- Two additional GPU nodes were added, each with two NVIDIA A100 GPUs on board.
- 16 old Skylake BrENIAC nodes were recovered and added to Hydra.
- All systems in the cluster were updated to Rocky Linux 8.8. This is the biggest change to the software stack in years. In preparing for this operation, almost every piece of our stack was examined.
- The scratch file system was replaced by a new hybrid system with an NVMe and hard disk drive disk pool. We now have 65 TB fast NVMe storage and 590 TB hard disk drive storage. The migration between both pools is invisible to the user.
- A platform based on JupyterHub was launched on (https://notebooks.hpc.vub.be).
- The application procedure for a VSC account was further automated. For staff members, applications are now approved within 5 minutes, for students there is a form that professors can use to request an account.

In addition to its own Tier-2 infrastructure, the VUB - together with the ULB - also manages the grid infrastructure, which is used, among other things, to process the data collected from experiments with the Large Hydron Collider (LHC). at CERN, but also used within the Flemish research community. The grid infrastructure consists of:

- 1 cluster
- 205 TF
- 6.104 CPU cores
- 51 TB memory

At the end of 2023, the capacity of the cluster was expanded to 205 TF by adding 16 new nodes with 2x AMD Epyc 9334 each. In addition, the dCache system was expanded to 11.2 PB.

The biggest change in 2023 was the introduction of a Ceph storage system for the users' home directories and as a backend for the local OpenNebula cluster. A lot of time was also spent migrating and upgrading all existing systems to this new storage cluster.

Exploitation and Use

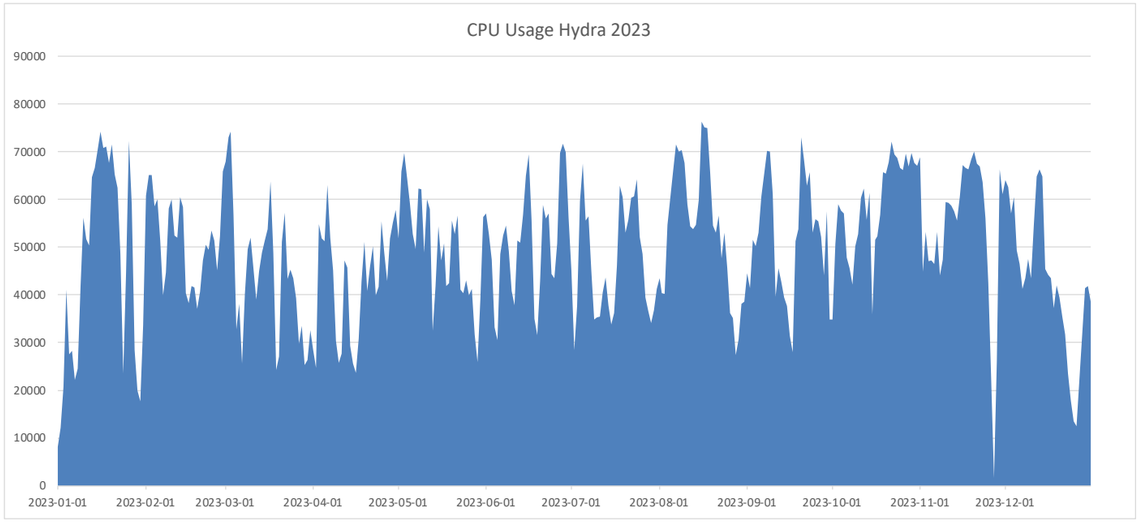

Figure 21. Overview of CPU and GPU usage on Tier-2 VUB.

The average usage in 2023 was 61% for the CPU part and 66% for the GPU part. A total of 19,807,543 CPU hours (2,261 years on 1 core) were used and 155,615 GPU hours (17 years on 1 GPU). The bulk of the computation time (98%) was consumed by 20% of the jobs. There were a total of 349 unique users active on the cluster, 327 unique on the CPU portion and 110 unique on the GPU portion (so some on both). Of these users, 209 were staff and 140 were students. On both CPU and GPU, 10% of users were responsible for 75% of usage.

For the CPU part we see the same trend as last year: a very erratic pattern. This corresponds to many but relatively short jobs on the system. Especially during the weekend, the load on the system is sometimes very low. One explanation is that several large users have successfully submitted Tier-1 projects. If we look at the statistics (where we only count jobs that took at least 1 hour), 72% of the computation time was used by jobs that ran on 1 node, and even 48% by jobs that used 1 CPU core. About 80% of jobs ran for less than 28 hours, 50% even in less than 6 hours. This corresponds to the erratic pattern in the graph.

About 90% of jobs start within 32 hours, half even immediately. This is broken down:

	75% percentile	90% percentile
Single CPU core jobs	2,75 hour	16 hour
Single Node jobs	5,50 hour	32 hour
Multi node jobs	7,50 hour	45 hour

Table 9. Allocation of compute time

The GPU part is a lot more popular, just like last year. On several occasions we have approximately 100% load. Of the GPU jobs, 50% started within a few minutes and 90% within 31 hours. About 80% of jobs lasted less than 20 hours. The vast majority (92%) of jobs use only 1 GPU which also represents 87% of the GPU time.

There have been no major unplanned outages this year. On November 27 and 28 there was a scheduled maintenance to deploy the new scratch and upgrade all systems to Rocky Linux 8.8. On October 10 there was a configuration issue with the job scheduler that had a major impact and on July 14 there was a general network outage that also affected the Tier-2 cluster.

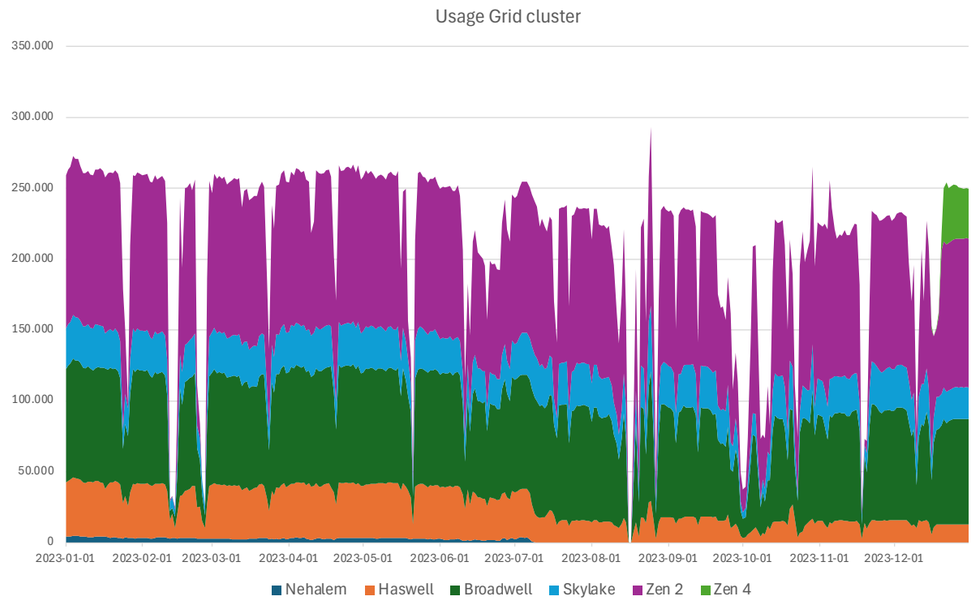

Figure 22. Usage statistics of the grid infrastructure (2023)

The Begrid cluster is heterogeneous: the cluster consists of different series of worker nodes with different generations of mainly Intel CPUs. The purpose of the cluster is mainly to perform data intensive, single-core calculations (so-called High Throughput Computing) and the size of the cluster is expressed in job slots which corresponds to 1 hyper-threaded CPU core. The average utilization of the cluster is approximately 80%.

On July 17 and 18 there was a reboot of all worker nodes after a planned power outage. Between August 16 and 23, the dCache system was upgraded and there were problems with the xrootdlocal daemon on the worker nodes. There was a planned power outage in the data center between November 17 and 20.

The grid infrastructure, managed by VUB/ULB, is also intensively used by researchers from various institutions: VUB, UAntwerp and UGent. The remaining computing time on the grid infrastructure is used by researchers from ULB and UCL.

Most researchers use the so-called “glide in” mechanism. This involves submitting “pilot jobs” that, once active on a worker node, will collect the “payload” elsewhere. For the calculation of the percentages mentioned above, only the calculation time of the “pilot jobs” was taken into account, not of the individual “payloads”. Moreover, it should be noted that with Grid, a user’s workflow can be spread across different sites even in different countries and the above percentages therefore only reflect a part of the actual computing time used.

Toekennen rekentijd

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

User Support

The VUB has a specific contact point for all HPC questions (hpc@vub.be). User questions are divided among the various services offered. Last year, a total of 684 tickets for Tier-2 and 189 for grid were processed, which is in line with previous years. For Tier-2, you can see the breakdown across services in the graph below.

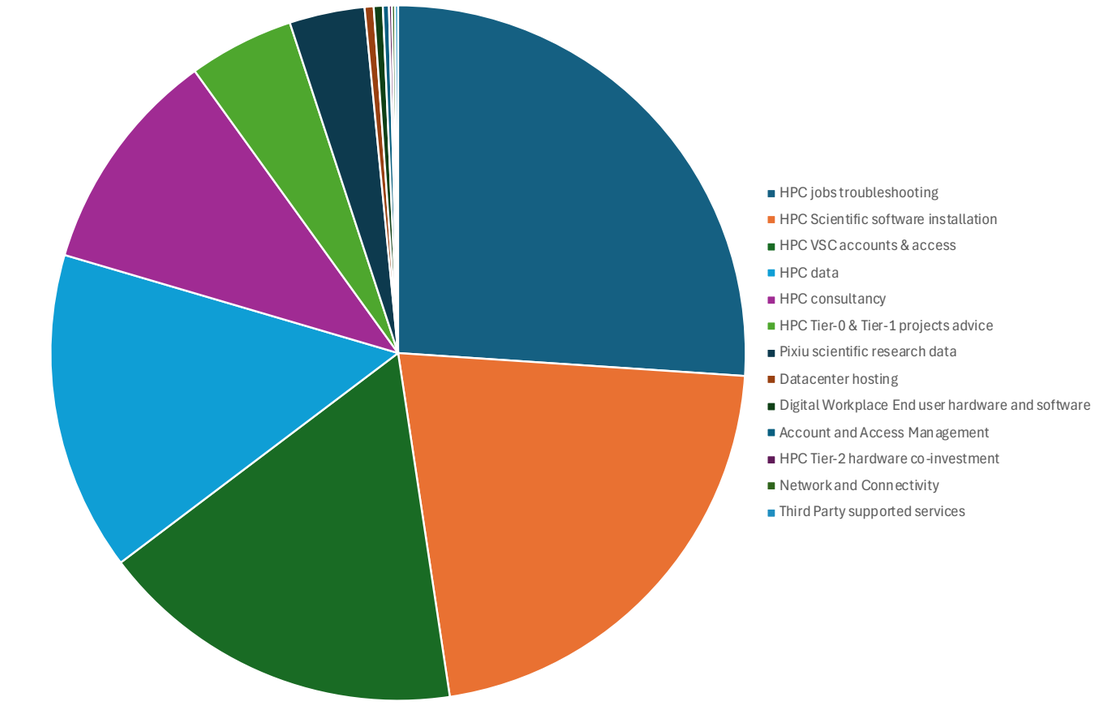

Figure 23. Distribution of tickets across the various services offered to VUB

For grid the distribution is:

Accounts	73
Software	14
Other	102

Table 10. Overview of grid tickets handled at VUB.

The time to resolution expressed in working hours for Tier-2 is below:

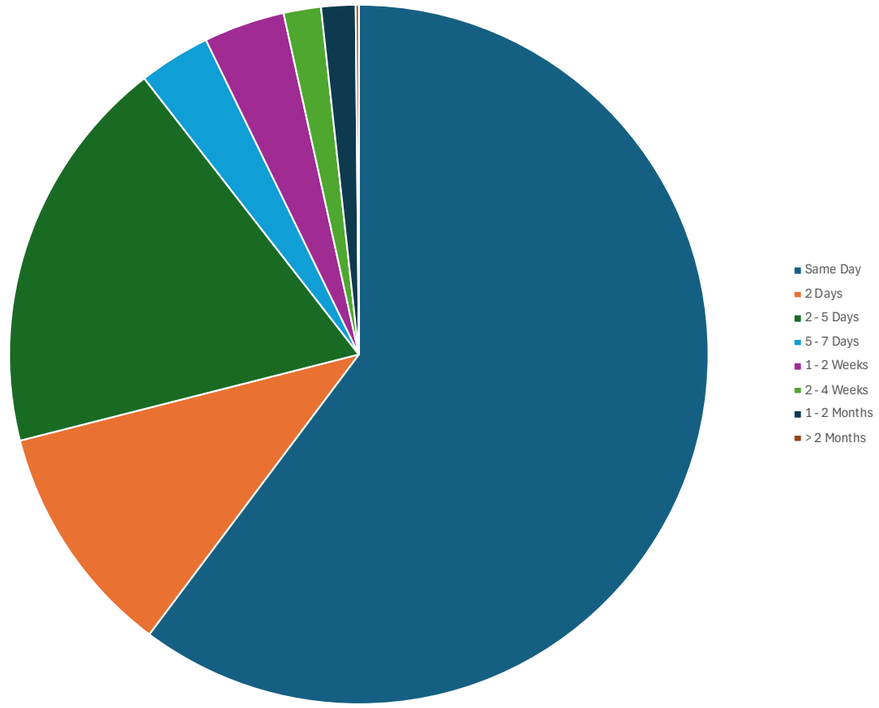

Figure 24. Resolution time for Tier-2 tickets at VUB.

The vast majority of questions are resolved within a few (working) days.

Specific Support

The courses "Introduction to Linux" and "Introduction to the use of HPC at the VUB" are organized twice a year. In addition, there have been several meetings with research groups to introduce them to HPC and VSC in general (and move away from their own machines). A number of groups were also encouraged to try the Tier-1 Cloud platform in addition to the usual push towards Tier-1 Compute from major users.

The VUB climate group has also received specific support to migrate from BrENIAC to Hortense.

The annual user survey has been replaced by a VSC-wide version. There were 55 respondents from the VUB. Overall we see great satisfaction. That is only slightly lower for documentation. The exact cause of this is not immediately clear.

User feedback was mainly about the need for more GPU computing power and the wait time to start jobs. In the coming year, we want to focus even more on migrating users towards Tier-1 and Tier-0 to respond to this complaint.

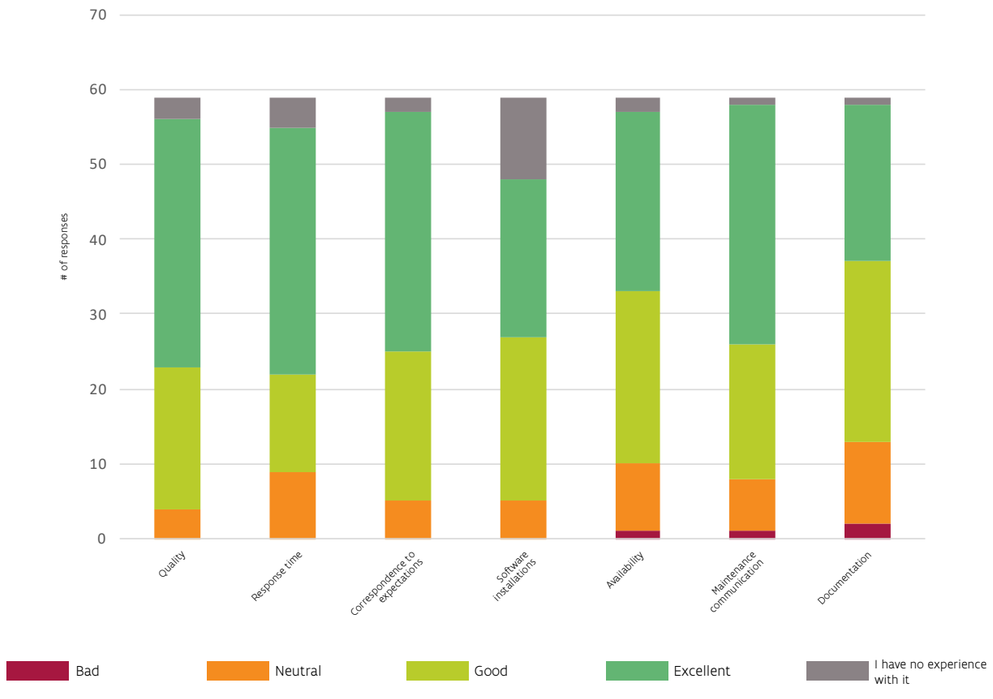

Figure 25. Tier-2 survey at VUB.

Showcase

Adrián Díaz and David Bickel from the Bio2Byte research group teach a course on Structural Biology to master's students. In the computational part of this course they use AlphaFold and talk about molecular dynamics simulations. For this part, students need access to an HPC cluster, but every year this has proven to be a difficult task. The launch of the notebook platform (https://notebooks.hpc.vub.be) and the automatic flow to create VSC accounts makes this go more smoothly. Only via the web browser and with minimal effort, students can now use HPC power to carry out for the hands-on sessions. They make extensive use of the possibilities of Jupyter notebooks to combine code and text.

"Our research relied heavily on numerical simulations that required a significant amount of computing power. Specifically, simulations of prominences and their formation were performed on more than 300 cores. These calculations were performed using the VSC's advanced Tier-2 (Genius) resources."

Veronika Jerčić
Plasma Astrophysics
KU Leuven

Ghent University

Available Infrastructure

The Tier-2 infrastructure of Ghent University consists of of seven clusters, each with specific characteristics:

- 4 CPU clusters
- 2 GPU clusters
- 1 interactive debug cluster

A total of 21 712 CPU cores, 92 GPU accelerators (NVIDIA)

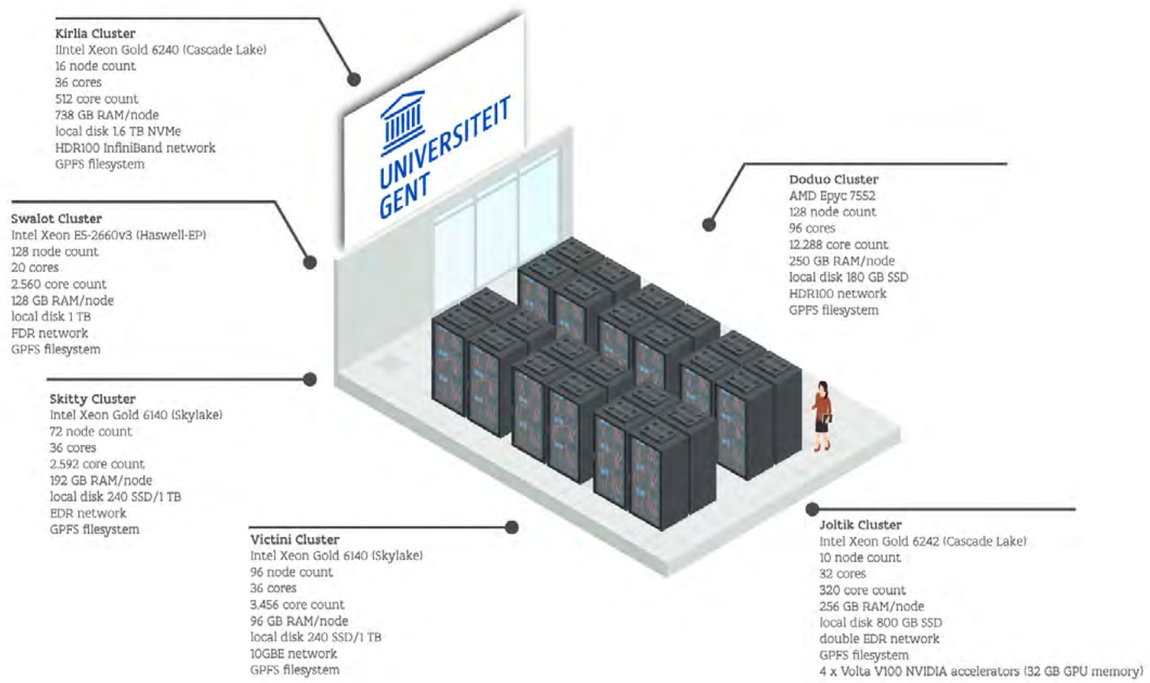

Figure 26. Tier-2 infrastructure at UGent

The following infrastructure and service changes were implemented during 2023:

- New cluster donphan was introduced. This is a debug/interactive cluster to replace cluster slaking, which was decommissioned.
- New cluster gallade was introduced. This is a new large-memory cluster to replace cluster Kirlia, which was decommissioned.
- A new service was rolled out to support teachers who use the HPC infrastructure of Ghent University. A kick-off was held on June 22 and documentation was recorded (https://docs.hpc.ugent.be/teaching_training/).

Usage and Availability

The figures below provide an overview of the day-by-day consumption of computing time on the various Tier-2 UGent clusters, broken down by CPU usage and GPU usage. A total of 101,969,556 CPU hours of computing time were used in 2023, which would correspond to 11,640 years of computing on one CPU core. Furthermore, a total of 395,358 GPU hours were used. This would correspond to 45 years of computing on one GPU unit.

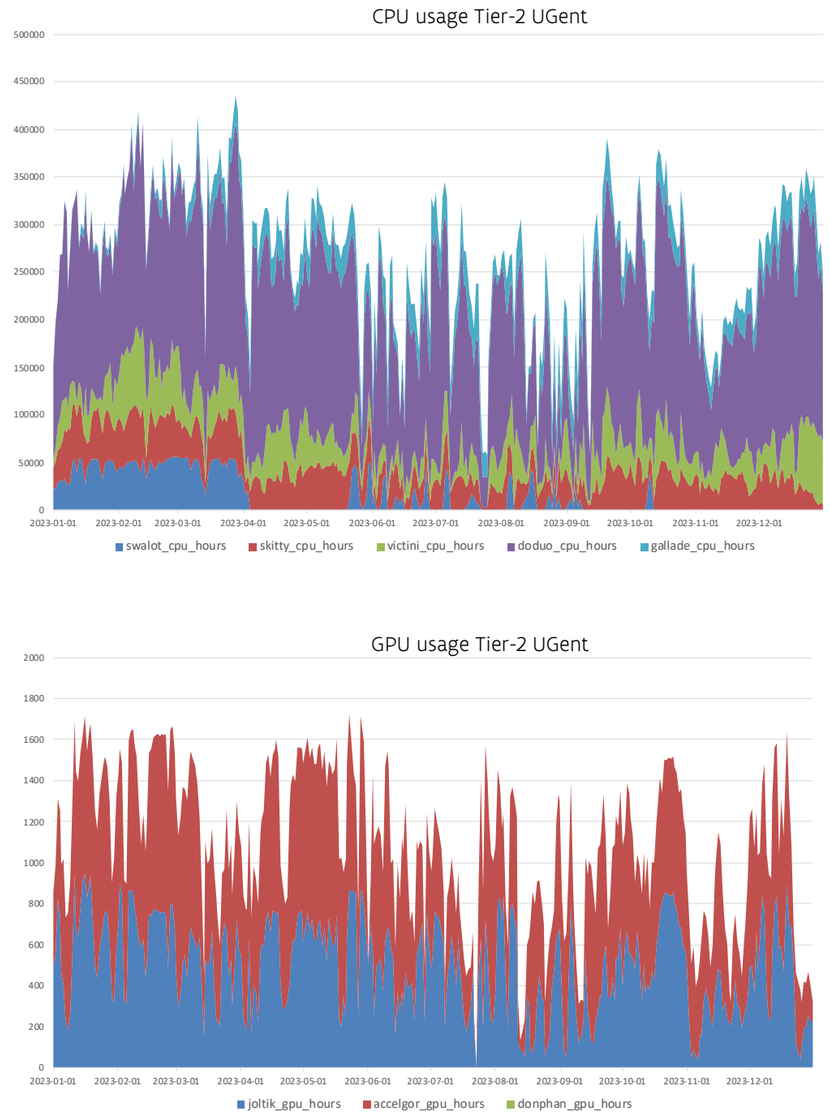

.Figure 27. Overview of CPU and GPU usage at Tier-2 UGent

There was a limited number of unavailabilities of the Tier-2 infrastructure in 2023:

- 2023/06/16: all clusters unavailable due to power outage (unplanned)
- 2023/09/11-12: Tier-2 shared file system crash (unplanned)

A total of 1,700 users were active in 2023 on the Tier-2 UGent infrastructure. A growing number of students used it in the context of lessons and theses: 725 students. The distribution according to the use of computing resources is:

	Percentage share in total of all resources used	
	Researchers	Students
CPU Resources	93.44%	6.56%
GPU Resources	89.67%	10.33%

Table 10. Share of researchers and students at Tier-2 UGent.

As always, the Tier-2 clusters were very well used by several power users, who were responsible for a large part of the usage. These were encouraged to submit a Tier-1 Compute application. A growing tail of very diverse users is also clearly visible, both for CPU (left) and GPU (right). There is great potential here for HPC users, in very diverse scientific domains.

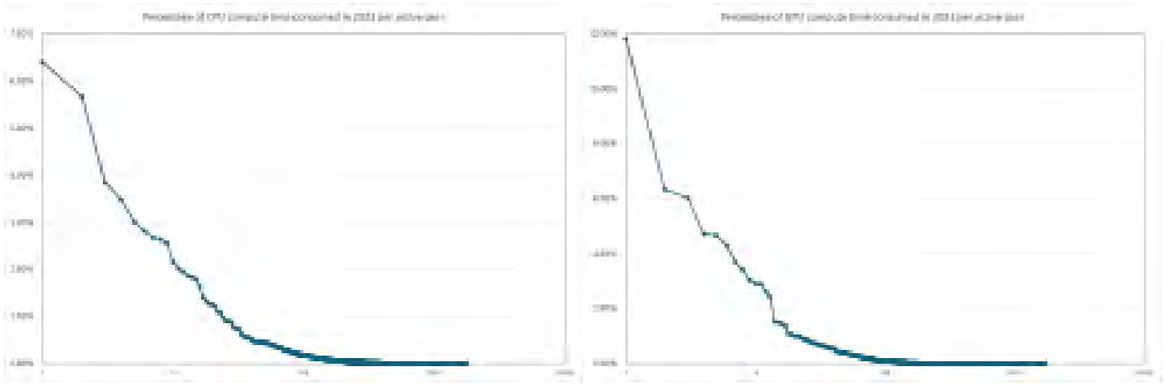

Figure 28. Percent or resp. CPU/GPU calculation time used in 2023 per active user

As agreed when the VSC was established, Tier-2 clusters can also be used by other institutions. The table below provides an overview of the number of users and their share in the total computing time consumed in 2023. A breakdown has been made according to the institutions that are consortium partners in the VSC and other research institutions (e.g. VLIZ, RBINS, etc.) and businesses.

Share in use of Tier-2 UGent in 2023			
Institution	#users	CPU usage	GPU usage
UAntwerp	32	3.26%	6.75%
VUB	20	0.92%	4.10%
UGent	1 564	95.10%	86.89%
KU Leuven/UHasselt	35	0.08%	1.70%
Other institutions and companies	49	0.65%	0.56%

Table 11. Share of the different users at Tier-2 UGent.

Helpdesk

The user management for Tier-2 UGent is divided into a number of helpdesk queues:

Queue	#processed and closed tickets
Helpdesk Tier-2	2 312
External account management	82
Software installations	293

Table 12. Overview of tickets handled at UGent.

Via the dedicated Tier-2 helpdesk (hpc@ugent.be) users can report problems or ask questions related to this infrastructure. The graph below shows a breakdown of the time it took for each logged question to be resolved. More than 80% of all questions were resolved within 24 hours.

Figure 29. Resolution time for Tier-2 tickets at UGent.

The Tier-2 of Ghent University was also used in 2023 by several companies and research institutions that are not inherently part of the VSC consortium, e.g., INBO, ILVO, VLIZ, etc. For the external account management for these institutions the helpdesk resolved 82 technical questions.

As already mentioned in the text about Tier-1 Compute, in 2023 a total of 293 software installation requests were handled by the UGent team, divided between Tier-1 Compute and Tier-2 UGent. As a result of these questions, 1,313 software packages (including dependencies) were installed and/or updated for Tier-2. The response time for software installations is noticeably longer than for other tickets, as shown in the graph below. Typically, each centralized installation requires significant R&D and much longer development time.

Figure 30. Resolution time for software installation requests to UGent.

User Survey

In the VSC user survey at the end of 2023, user satisfaction for the Tier-2 UGent infrastructure and services was also questioned. 154 respondents indicated that they use it to a greater or lesser extent.

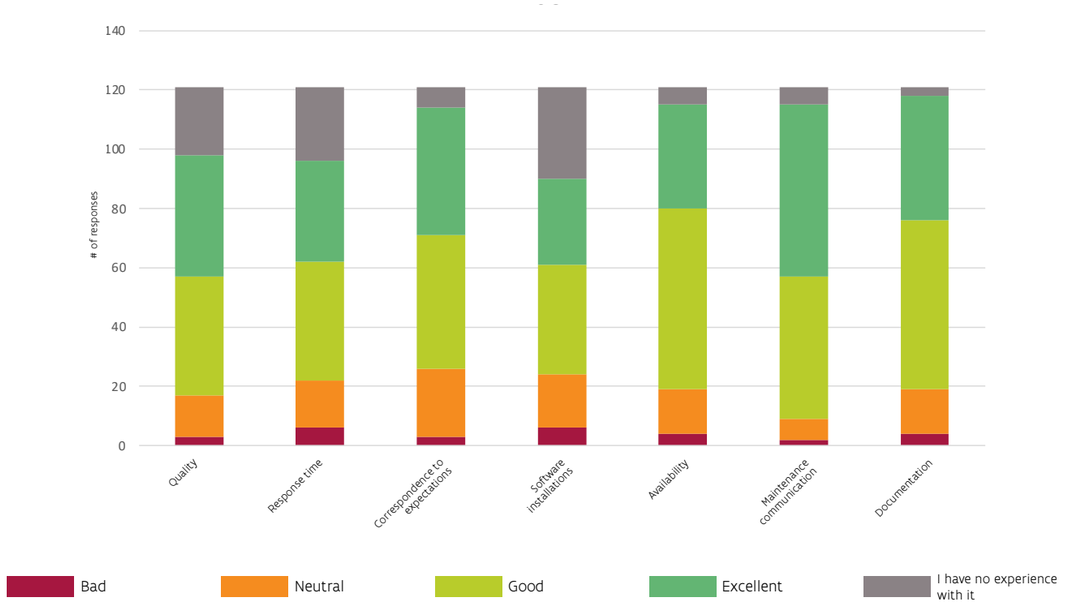

Figure 31. Tier-2 survey at UGent.

The following components of the Tier-2 UGent infrastructure were rated as 'good' or 'excellent' by users:

- 83% (81/98) - quality
- 77% (74/96) - response time
- 77% (88/114) - agreement with expectations
- 73% (66/90) - software installations
- 83% (96/115) - availability
- 92% (106/115) - communication regarding maintenance
- 84% (99/118) - documentation

Possible suggestions include::

Infrastructure

- More stable mounting of shares via kinit (2 respondents)
- More clusters (1 respondent)
- More RAM per node (1 respondent)

Documentation

- Clearer indication of relative/absolute paths on different clusters (1 respondent)
- Examples of job files (1 respondent)
- How can you install software/module yourself? (1 respondent)
- How can you work with conda environments? (1 respondent)
- How can you work with VS Code? (1 respondent)

Software

- Jupyter notebooks for R (2 respondents)
- Faster response to software installation requests (4 respondents)
- Easier sharing of software installations in a VO (2 respondents)
- Easier deployment of Python deep learning applications (1 respondent)
- An Abaqus desktop app (1 respondent)

User Experience

- Higher job walltime (2 respondents)
- Info/impact on job/queue status across the different clusters (2 respondents)
- Larger $ VSC_DATA folder (1 respondent)

Possible suggestions for new developments were:

- Specific data storage for sensitive data (1 respondent)
- Nextcloud-like interface with data on HPC (1 respondent)

▶ KU Leuven/Hasselt University

Available Infrastructure

KU Leuven and UHasselt work together for the Tier-2 infrastructure.

The infrastructure consists of:

- 2 clusters, 6 partitions
- 1.790 CPU TF, GPU 752 TF
- 23.194 CPU cores
- 120 TB memory
- 117 GPU devices / 368.640 GPU cores

There were no major changes in 2023. The new wICE cluster was added in 2022 and went into production in 2023 after a long pilot period. The scratch storage capacity was expanded in 2023.

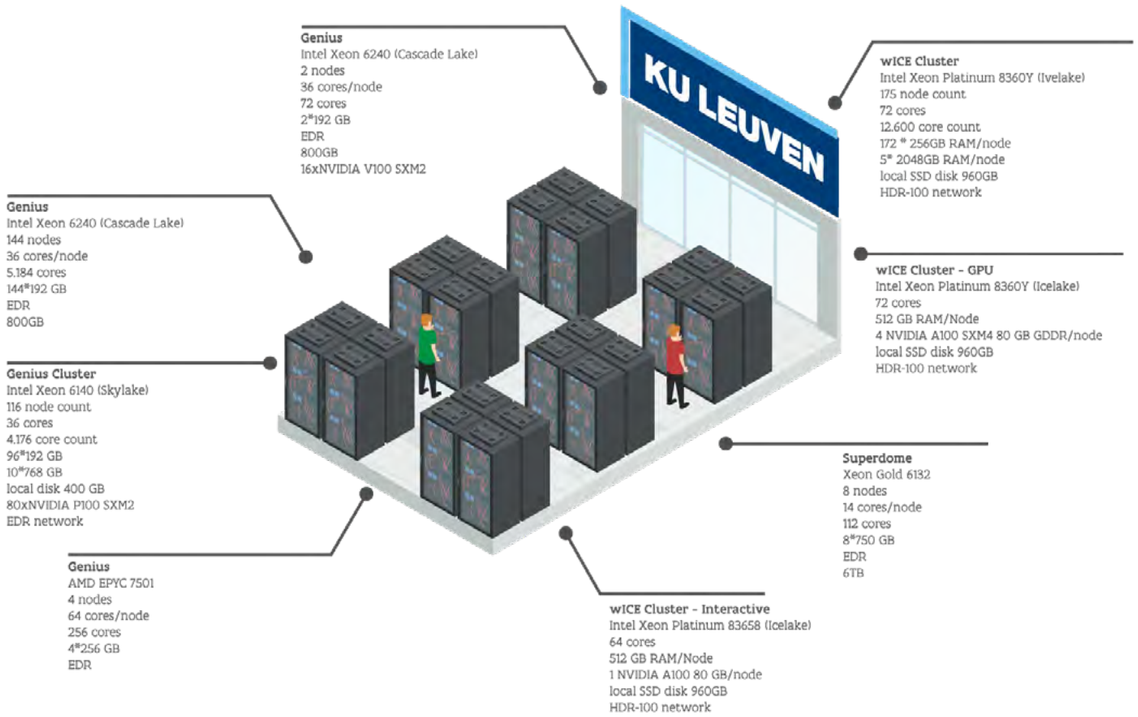

Figure 32. Tier-2 infrastructure hour KU Leuven – UHasselt

Exploitation and Use

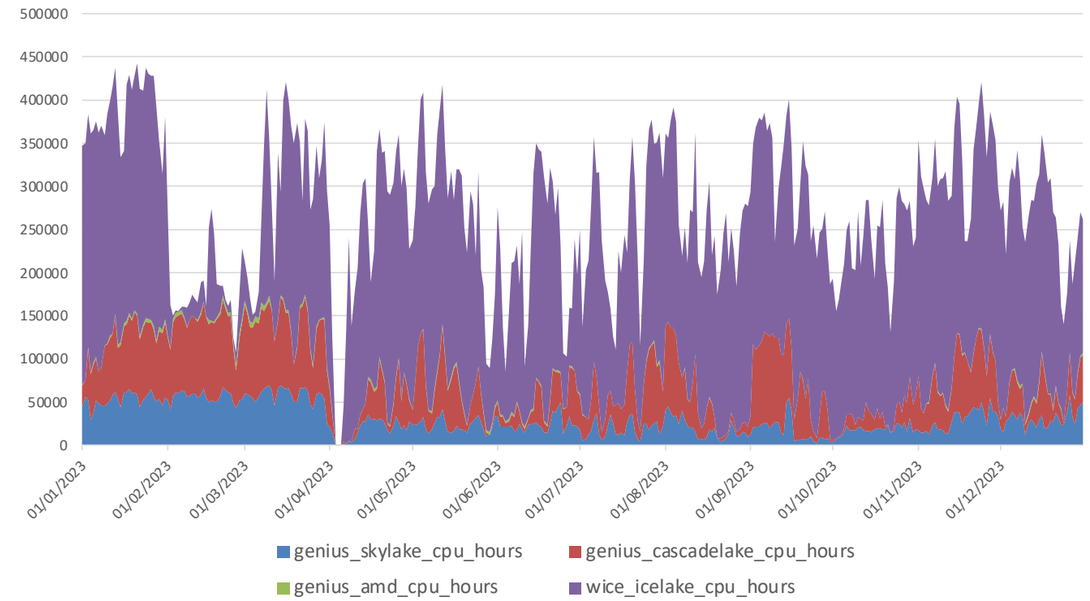

Figure 33. Overview of CPU usage at Tier-2 KU Leuven/UHasselt

The Genius/wICE CPU utilization graph shows that the systems are in continuous use. The uptime of the machine has been very good. The Tier-2 machines were only unavailable during the upgrade of the scratch file system in April, combined with the complete transition to Slurm. wICE was already set up with Slurm. So the researchers already had several months to convert their workflows to Slurm and run them on wICE. After the update, it was no problem for users to get started with Genius.

A number of important events can also be deduced from the graph. A decrease in the use of wICE can be observed in February/March. This is the transition from pilot phase to production phase of the new Tier-2 cluster wICE. As soon as the news of the production phase was well known to the users, wICE was quickly used again by all researchers. Further down the chart it is clear that Genius' share is declining. This is now the oldest cluster. The researchers clearly prefer the new cluster that has more cores per node and also provides better performance with optimized software installations.

On September 25, there is also a decrease in the use of Genius. On this day, all Genius nodes were upgraded from CentOS 7 to Rocky Linux 8. This upgrade was necessary because CentOS has changed its release strategy and version 7 will no longer receive updates in 2024. The OS migration also means that all user software on the cluster must be reinstalled. With the different combinations of software version, compiler version and CPU version, this quickly amounted to around a thousand installations that also had to be tested. The entire process went smoothly and since the researchers were already working with Rocky 8 on wICE, the transition on Genius went smoothly and the system could be put back into production immediately.

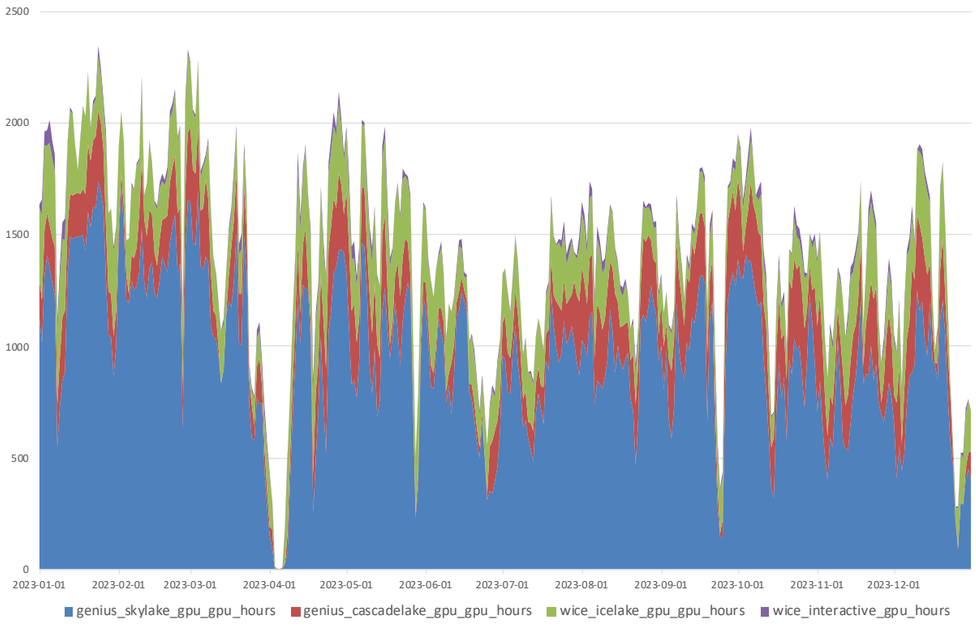

Figure 34. Overview of GPU usage at Tier-2 KU Leuven/UHasselt

The GPUs in the KU Leuven Tier-2 cluster are an essential part of the HPC environment and this is also reflected in the occupancy. The dips in the graph are related to the same updates that were already described in the CPU usage. A few additional comments can be made about the GPU graph. The computing capacity of the different GPU generations has increased very quickly. But the AI revolution/hype has also caused prices for GPU devices to rise sharply. The Genius Skylake partition contains the largest number of GPUs in our Tier-2 environment (80). These are from the NVIDIA P100 generation. wICE contains only 16 GPUs (4 nodes with 4 GPUs each) of the more recent A100 generation. Depending on the workload, an A100 can do four to ten times more computations per unit of time than a P100. The importance of the A100 GPUs is therefore greater than what is perceived at first glance on this graph. However, in the first year of production there were problems with two of the A100 GPU nodes. On one node it took over a month to fix the hardware problem. Since one node represents 25% of the available A100 resources, this has a significant impact on usage figures.

Allocating Computing Time

In 2023, 1,400 users were active on the system. The KU Leuven/UHasselt clusters use credit accounting system that is included in the scheduling software. In 2023, this was switched from the old scheduler (Moab Accounting Manager – MAM) to Slurm. To provide a similar user experience, the user commands that were available in MAM have been converted to a Slurm equivalent.

New users are given computing time to familiarize temselves with the system and to perform initial tests. Project credits are used for the actual work. This system promotes responsible use of the Tier-2 cluster. The principal investigator is the manager of the project. He can give researchers access to the computing time and also monitor the computing time used. When submitting a job, the project to which the credits are charged is specified. The UHasselt researchers also work with the projects. Computing time is allocated by local support.

Teachers can also request computing time that is required for courses. This keeps the entry barrier for researchers to switch to the Tier-2 infrastructure as low as possible. Credits can then be requested via a simple procedure and at minimal cost. In 2023, the process for teachers was simplified so that accounts are automatically provisioned for students in a course.

User support

Questions from users are received via email (hpcinfo@kuleuven.be).

In 2023, approximately 3,500 questions were handled. 87% of these tickets were answered within response time and subsequently 91% of tickets were resolved within the expected working time.

The diversity of the user community is also reflected in the questions. They originate from about 200 different research groups.

Type of Question	Quantity
Accounts, access, project and storage management	2 109
User application support (including software installations)	681
System support	414
Other	346

Table 13. Overview of tickets handled at KU Leuven/UHasselt.

User Survey

In the VSC user survey at the end of 2023, user satisfaction for the Tier-2 KU Leuven/UHasselt infrastructure and services was also questioned. 130 respondents from KU Leuven and 12 from UHasselt indicated that they use the Tier-2 Cluster to a greater or lesser extent.

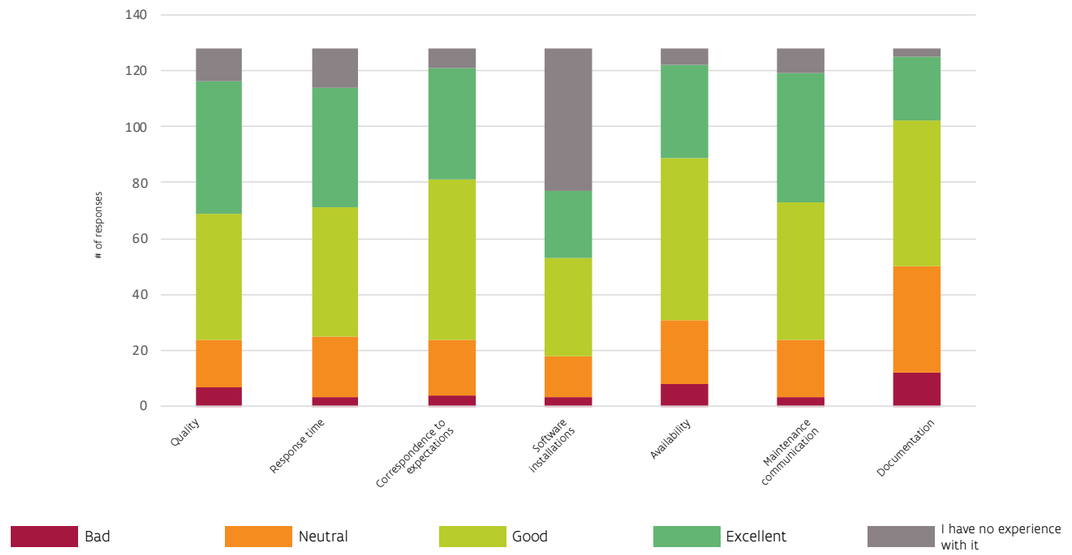

Figure 35. Tier-2 survey KU Leuven/UHasselt.

Percentage of users who answered the questions well or excellently:

- 79% quality
- 78% response time
- 80% agreement with expectations
- 76% software installations
- 74% availability
- 81% communication regarding maintenance
- 60% documentation

Important to note here are the comments about the documentation. In a year with two important changes that had an impact on users (new scheduler and new operating system), more documentation is being searched. There is a clear demand for more and specific examples.

Specific Support

With the Open OnDemand as a web interface, there is a new access to the cluster that can make it easier for novice users. A specific initiative was taken to bring this to groups in the domain of Humanities.

Showcase

The increase in computing power often also means that a domain under investigation can be simulated at a higher resolution. A good example of this is research into the microclimate in forests. Previously, average temperatures recorded at weather stations for a 1 square km domain at a standard height of 2 m were used to model species' ranges. This is an important method to project the impact of a changing environment on the distribution of a species. But in forests under a canopy, the temperature can differ greatly from the temperature measured in a weather station. The sGlobe lab of KU Leuven Department of Forest, Nature and Landscape has developed models with microclimate data (resolution 25 square meters) and shows that distribution models are more accurate than the models based on macroclimate data. They can also detect small environments in the landscape where a species can survive, even when there are changes in macroclimate. The new models can help formulate policy decisions (Microclimate reveals the true thermal niche of forest plant species).

The research group has published the 'ForestClim' dataset, which contains high-resolution bioclimate variables, so that other researchers can also use it further (ForestClim—Bioclimatic variables for microclimate temperatures of European forests).

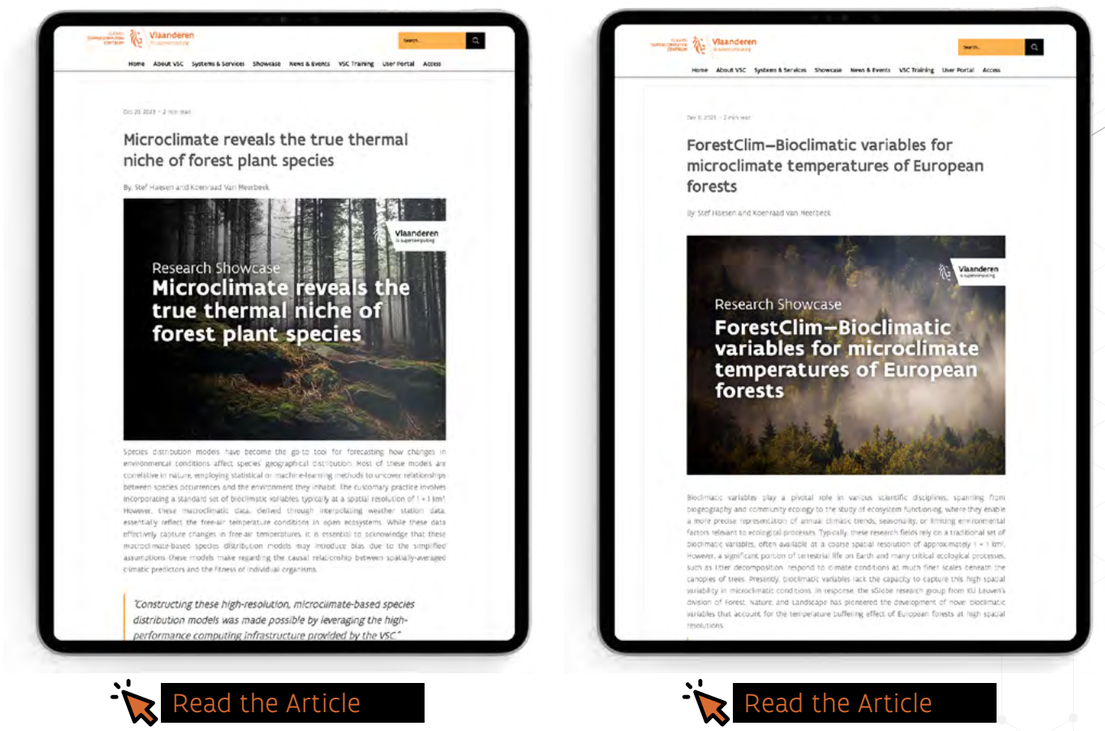

Figure 36. Both articles published on the "VSC Research Showcase" page.

Tier-0 Support

From 2023, there will be specific support within the VSC for researchers who want to use the EuroHPC infrastructure. Initially this concerns LUMI (Finland), the pre-exascale machine in which Belgium/Flanders has invested, but users who want other EuroHPC machines are supported as well. One FTE was active in 2023. The search for additional staff continues.

In addition to LUMI-specific courses, the available staff also participated in the "EuroHPC Virtual Workshop – AI on HPC", "Best Practices in HPC training", "Performance aware C++ programming" and training courses on Julia and OpenFOAM.

The following courses were announced to all universities:

- LUMI – developers (February 14-17, May 30-June 2 and October 3-6),
- LUMI – users (May 9 and 16 and September 21),
- LUMI – Hackathon (April 17-21 @CSC and November 27-December 1 in Krakow).

There were no participants from Flanders at the hackathons, but there were often people on the waiting list for the other courses.

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

In 2023 there were three calls for computing time on the Belgian share on LUMI. The following applications were granted:

2023-03	Jelle Vekeman	UGent	Preparatory	Center for Molecular Modelling
	Ahmadreza Mehdipour	UGent	Preparatory	Center for Molecular Modelling
2023-04	Cem Sevik Jovana Vlahovic	UAntwerp	Preparatory	Condensed Matter Theory
	Pieter Cnudde Massimo Bocus Jenna Mancuso	UGent	Regular	Center for Molecular Modelling
	Kenneth Hoste	UGent	Development	VSC
	Ahmadreza Mehdipour Alen Thykkoottahil Mathew Robin De Baker	UGent	Regular	Center for Molecular Modelling
	Olivier Beyens	UAntwerp	Preparatory	Medicinal Chemistry
2023-05	Olivier Beyens	UAntwerp	Preparatory	Medicinal Chemistry
	Olivier Beyens	UAntwerp	Regular	Medicinal Chemistry
	Hans De Winter	UAntwerp	Preparatory	Medicinal Chemistry

Table 14. Approved LUMI-BE applications in 2023.

In addition, (at least) one researcher has also submitted an application via the EuroHPC calls.

Researchers were given the necessary feedback on their applications, both before and after submission.

To attract more users to Tier-0 machines, the question was added to the request form for Tier-1 whether there was interest in GPU computing time on LUMI. All those users were contacted, and their responses varied: "We are already working on it", "We are still considering it, but don't have time for it right away", "No interest for the time being". In any case, it remains a point of attention to reach more users. In total, there were eight requests from new users (who have not used LUMI before) in 2023, a number that will hopefully increase in the coming years.

The emails that arrive via the ticketing system belong to two major categories: on the one hand, requests for the LUMI-BE calls, or assistance with requests for EuroHPC calls, and on the other hand, questions regarding. accounts on LUMI or specific problems with working on LUMI.

On November 6, the LUMI-BE user day was organized at FWO, together with CÉCI. In addition to presentations from EuroHPC, PRACE, the Belgian National Competence Center and LUMI, there were seven presentations by users:

- Josephine Wood (EuroHPC JU): EuroHPC JU: Leading the way in European Supercomputing
- Serge Bogaerts (PRACE): The Third Phase of PRACE - A User-centric Association
- Benoît Dompierre (NCC Belgium): EuroCC Belgium: Empowering the Belgian supercomputing community
- Kurt Lust (LUST): LUMI introduced: Opportunities and limitations
- Wouter Ryssens (ULB): Dense matter in the cosmos: nuclei, paste and explosions
- Christophe Geuzaine (ULiège): GmshDDM on LUMI: first runs of a new solver for large scale time-harmonic flow acoustics problems
- Denis Haumont (RMI): Destination Earth: using LUMI to improve the prediction of extreme weather events
- Olivier Beyens (UAntwerp): Design of novel DPP8 and DPP9 inhibitors using cosolvent molecular dynamics simulations
- Michel Rasquin (Cenaero): High fidelity flow simulations of the boundary layer transition on a high-pressure turbine vane in view of accurate predictions of the heat flux distribution
- Pierre Beaujean (UNamur): Pushing the boundaries of molecular simulations with LUMI: multi-million atoms simulations on a realistic dyed cell membrane
- Thomas Gillis (UCLouvain): Faster MPI: GPU-to-GPU Communication on Slingshot-11

There were 53 participants. The next user day will be organized in the autumn of 2024.

Following this LUMI-BE user day, the MOCCA code of Wouter Ryssens (ULB) was examined. It suffers from load balancing problems that limit large-scale computations. It is being investigated how this can be remedied to gain experience with code modernization and parallelization on LUMI. A performance analysis is planned for the spring of 2024.

Training

The VSC provides training that is primarily aimed at its users or potential users. These are employees in the private sector, public services and, in terms of numbers, mainly researchers affiliated to collaborate the Flemish university associations and the various knowledge institutions.

We are also gradually noticing a small, but steadily growing, number of foreign participants in our training activities. This is partly due to greater brand awareness of the VSC abroad, but also to the network of National Competence Centers (NCCs) that work together in the context of EuroCC. This is without a doubt a positive trend.

Training is important for the VSC. After all, computations using supercomputing infrastructure are expensive, both in terms of investment and operation. It is therefore expected that this infrastructure is used efficiently. The training courses contribute greatly to this. Although the financial aspect is of course important in the short term, professional use of this infrastructure will also ensure greater competitiveness of both our researchers and our companies in the longer term.

Education and training materials also contribute to the image and reputation of the VSC. The analytics of the VSC website show that the training page is one of the most visited on the site. Training is also a subject that lends itself well to collaboration with other HPC organizations such as CÉCI, but also internationally within PRACE and EuroCC.

The courses can be divided into four categories that either indicate the required prerequisits knowledge or clarify that they are domain-specific topics:

- Introductory
- Intermediate
- Advanced
- Specialist courses & workshops

Introductory courses are intended for all infrastructure users and are highly recommended for those who do not yet have the necessary skills. Local VSC staff lead these sessions. This also offers researchers the opportunity to get to know the people who answer the questions they ask the helpdesk. This removes the impersonal and anonymous character of e-mail traffic and therefore lowers the threshold.

To follow the sessions at the intermediate level, the introductory courses are required as prior knowledge. These sessions therefore cover more specific topics. The majority of these courses are intended for users who develop software themselves, either for computationally intensive applications or for pre- and post-processing of data. Because these courses are more specialized and intensive than the introductory courses, they are not given at every VSC site. Users are therefore encouraged to attend the training courses on another site. Many of these training courses are offered online, so that geographical distance no longer plays a role, which should lower the barriers.

Advanced level courses require even more experience and are more domain-specific than intermediate courses, which often require prior knowledge. The VSC also invites external instructors for these courses. They are often affiliated with EuroCC National Competence Centers or come from industry. In 2023, the VSC organized a “Performance-aware C++” training in collaboration with EuroCC@Belgium, the Belgian Competence Center. EuroCC and LUMI training courses were regularly promoted via social media and the website.

In total we welcomed 1,057 participants in 58 training sessions.

Some courses do not fit into any of these three levels: either they are too domain-specific, or they include the entire introductory to advanced level. This explains the term “specialist”, as external specialists are often called upon for this. These events are often part of community building. More information about this can be found in that section.

The VSC developed two MOOCs for PRACE, “Defensive programming and debugging” and “Fortran for scientific computing”. The first is no longer available and a replacement training is in preparation. “Fortran for scientific computing” is still offered and has already been followed by more than 1,500 participants from 98 countries.

Naturally, the existing training courses are always kept up to date with recent developments, but new training topics are also selected and developed every year. For example, in 2023 a training on the use of AI in software engineering was given.

VSC staff also followed training courses online (half or full day) and at ISC 2023 in Hamburg, Germany.

The announcements are distributed to (potential) users of the infrastructure via the website, internal mailing lists and social media. Targeted mailings draw attention to specific training courses if they may be useful for a limited target group or for potential users.

Figure 37. VSC & EuroCC Belgium team at ISC 2023 Hamburg

VSC training 2023

1057
participants

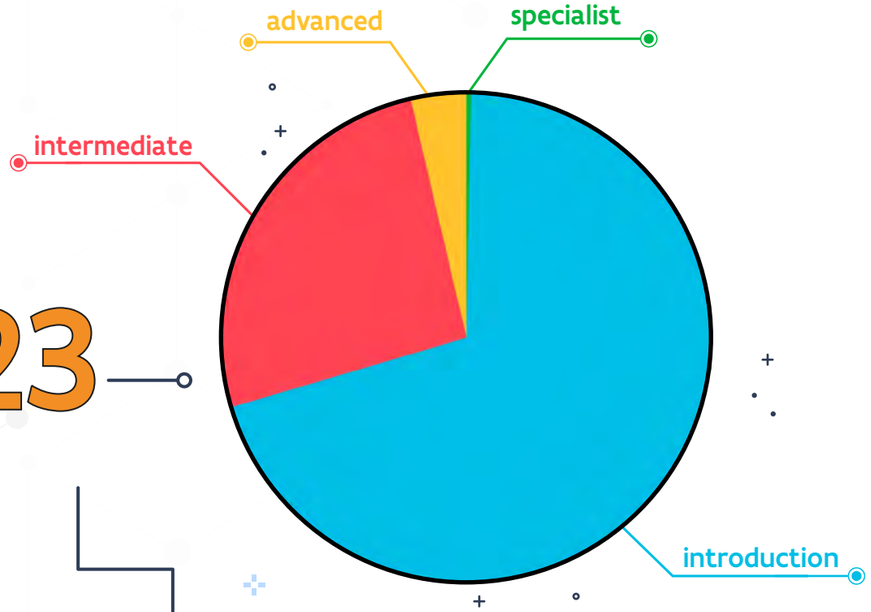

57
training events

1 external specialist

15 topics

Machine learning
MPI
OpenMP
C++
C
HPC intro
Linux
Python
Debugging
Fortran
Optimization
Data science
Containers
Software engineering
Interactive computing

PRACE MOOCs

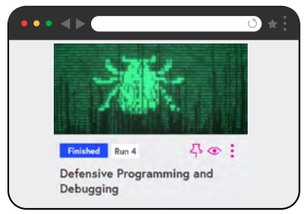

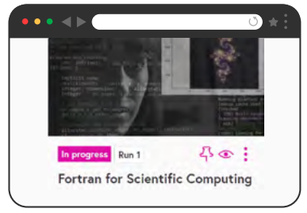

Events

VSC Users Day

The event, which took place on October 24, 2023 at FWO, Hoek 38, Brussels, was a resounding success. The theme of the day revolved around optimizing workflows for maximum HPC efficiency. We planned this event down to the last detail with those attending in mind, providing a diverse selection of thought-provoking lectures and engaging lightning talks. The day provided an unparalleled opportunity to connect and engage with fellow VSC HPC users, reaffirming our commitment to fostering a strong community.

Speakers

Frederik De Ceuster
KU Leuven
Scientific Programming in The Age of AI

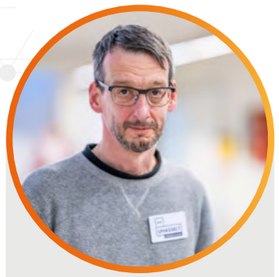

Geert Jan Bex
VSC
ChatGPT for HPC

Frédéric Wautelet
University of Namur
Ansible Lightspeed

Ewald Pauwels, Jan Ooghe
VSC
Beyond traditional HPC: Tier-1 Cloud and Tier-1 Data, Open OnDemand, containers

Sebastian Munck
VIB Bio Imaging Core
Flanders BioImaging – Towards efficient centralized research data management and analysis of bioimaging data

Alex Domingo, Kenneth Hoste
VSC
Best Practices for HPC or How to get your work done faster?

Discover the web page of VSC Users Day 2023

Lightning Talks

Below is a list of participants in the lightning talks. Click on the title of the presentation to view it.

Speaker	Institute	Title
Olivier Beyens	UAntwerp	DPP9 inhibition design using cosolvent molecular dynamics simulations
Alexander Botzaki	VIB	VSC infrastructure as cornerstone for VIB technology training operations
Kurt De Grave	FlandersMake	RAISE: a challenge to hone your AI-on-HPC skills
Adrián Díaz	VUB	Interactive teaching in computational biology with Jupyter notebooks
Nicholas Jannsen	KU Leuven	End-to-end simulations as an indispensible tool for the ESA PLATO space mission
Cécile Kremer	UHasselt	Evaluating COVID-19 testing policies in primary schools
Selma Mayda	UAntwerp	From paintings to batteries: our DFT journey through CdS and Li batteries on Tier-1 and Tier-2
Derrick Muneki	VUB	Hydroclimatic data digitization campaign of the INERA archives in Yangambi, DRC & the state archives of Belgium
Jérôme Neirynck	KU Leuven	Freewind project
Jochen Schütz	UHasselt	High-fidelity parallel time-stepping schemes
Astrid Sierens	UHasselt & VUB	The role of venue-based superspreading in the transmission of infectious diseases
Lander Willem	UAntwerp	Modeling COVID-19 in Belgium using the VSC HPC cluster
Max Yudayev	KU Leuven	How to train better time-series AI models, on bigger data, in shorter time, on smaller GPUs?
Jure Oder	VKI	Validation of liquid metal pool-type nuclear reactor CFD simulations

Table 15. List of Lightning Talks attendees (click on the title to view the presentation).

Figure 38. Atmospheric images of the event

In 2023, VSC organized a number of networking events, such as the Quantum Computing session and Data-driven Approaches in Life Sciences in collaboration with the University of Antwerp.

VSC Lunch Session on Quantum Computing

In April 2023, VSC organized a two-hour online Lunch Session with a focus on Quantum Computing, a rapidly emerging technology that solves problems that are too complex for classical computers. The session, led by leading experts in the field, included presentations from Dr. Ariana Torres, Dr. Mikael Johansson and Dr. Koen Groenland. Dr. Torres discussed the fundamental concepts and current state of the art in quantum computing, Dr. Johansson explored the integration of quantum computers with high-performance computing via LUMI-Q, and Dr. Greenland investigated the potential future applications of large-scale, error-free quantum computers. The session ended with a fascinating Q&A session. Among the participants were researchers, enthusiasts and government officials. The event highlighted the shared journey towards achieving quantum advantage and is available to watch on VSC's YouTube channel.

Figure 39. Illustration of the sessions

Dr. Ariana Torres
SURF

Dr. Mikael Johansson
LUMI-Q / CSC - IT

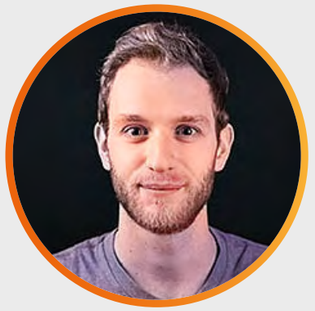

Dr. Koen Groenland
QuSoft / Quantum Amsterdam / UvA

Day of Science

Last November, we embarked on an extraordinary journey of innovation, presenting three groundbreaking use cases made possible by our state-of-the-art supercomputers. At Hasselt University, the Innoptus Solar Team fascinated attendees with their solar car, emphasizing the future of green energy and supercomputing. The Museum of Ghent University delved into advanced climate models and revealed how our supercomputers decipher the secrets of the atmosphere at lightning speed. In the UFO building of Ghent University we demonstrated our commitment to sustainability with innovative recycling processes and biodegradable plastics.

Figure 40. Illustration of the publication of the event on social media (left) and atmospheric images of the DVW in UHasselt (right)

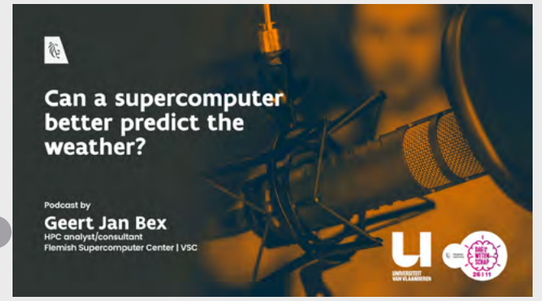

In addition, our VSC colleague Geert Jan Bex shared insights about supercomputers and weather forecasting in a live podcast during the University of Flanders activity. The event was a testament to our commitment to discovery and excitement in the world of supercomputing.

Listen to the podcast here

Outreach

VSC in the Press

The VSC enables significant scientific advancements in Flanders and beyond. As a result of its prominent position, we have garnered media coverage in various national and international publications in 2023. This coverage primarily highlights the hosting of the new Tier-1 at the VUB, which will be advantageous for both academia and industry.

- On November 9, 2023, the website **Techzine Europe** published that Belgium will receive a new Tier-1 supercomputer. This will be placed in Brussels, but will be accessible to all scientists in Flanders.

The supercomputer will be in production by Fall 2025. This will be located in the Green Energy Park of the Vrije Universiteit Brussel. The project requires a total investment of 12 million euros, with the university sharing the costs.

Link to the article

- The article on **Computable.be**, published on November 8, reports that the Vrije Universiteit Brussel (VUB) will start up the Flemish Tier-1 supercomputer in the Nexus data center in Zellik in the autumn of 2025. The Flemish Supercomputer Center (VSC) has decided that the VUB will be responsible for housing and maintaining this supercomputer, which represents an investment of twelve million euros.

Link to the article

- **Engineeringnet.be** has reported that the Flemish Supercomputer Center (VSC) has decided that the next Flemish Tier-1 supercomputer will be housed and managed by the Vrije Universiteit Brussel (VUB). The supercomputer will be located in the Nexus data center at the Green Energy Park site in Zellik. This investment of 12 million euros is expected to be operational in the autumn of 2025.

The Nexus data center, located at the Zellik research park, will serve, among other things, as a data center for the VUB and the UZ Brussels. It is positioned as the greenest and most efficient data center in Belgium.

Link to the article

- The article by **DataNews** of November 8 discusses the preparations of the Flemish Supercomputer Center (VSC) for their next Tier-1 supercomputer, which will be housed at the Vrije Universiteit Brussel (VUB). This supercomputer will be available to all scientists and is planned to be operational around October-November 2025, located at the Green Energy Park site in Zellik. The total investment for the supercomputer is 12 million euros. The specifications of the machine and the builder are not yet known, as the European tender for construction has yet to take place. Ward Poelmans, department head of Scientific Data and Compute at the VUB, states that the tender will be completed by the end of 2024 or early 2025.

Although there is a lot of attention to artificial intelligence and language models, the new supercomputer remains a general-purpose machine, since the costs are high and the machine will be used by all Flemish universities.

Link to the article

VSC Website

Introduction

The VSC website (www.vscentrum.be) is a crucial platform that supports the research community by providing comprehensive information about our supercomputing resources, services and performance.

Website Objectives

Our website aims to facilitate easy access to VSC resources, support researchers with detailed documentation and keep the community informed about the latest news and events.

Key Features and Sections

- **Homepage**: Get a dynamic overview of VSC activities.
- **About Us**: Learn more about our mission, vision and history.
- **Services**: Explore our supercomputing services and resources.
- **Access and Accounts**: Step-by-step guide to gaining access.
- **Support and Documentation**: Access user manuals, guides and helpdesk services.
- **News and Events**: Stay up to date with our latest news and event announcements.
- **Success Stories and "Research Showcase"**: Read or watch all about successful research projects facilitated by VSC.
- **Contact information**: Find out how to contact us.

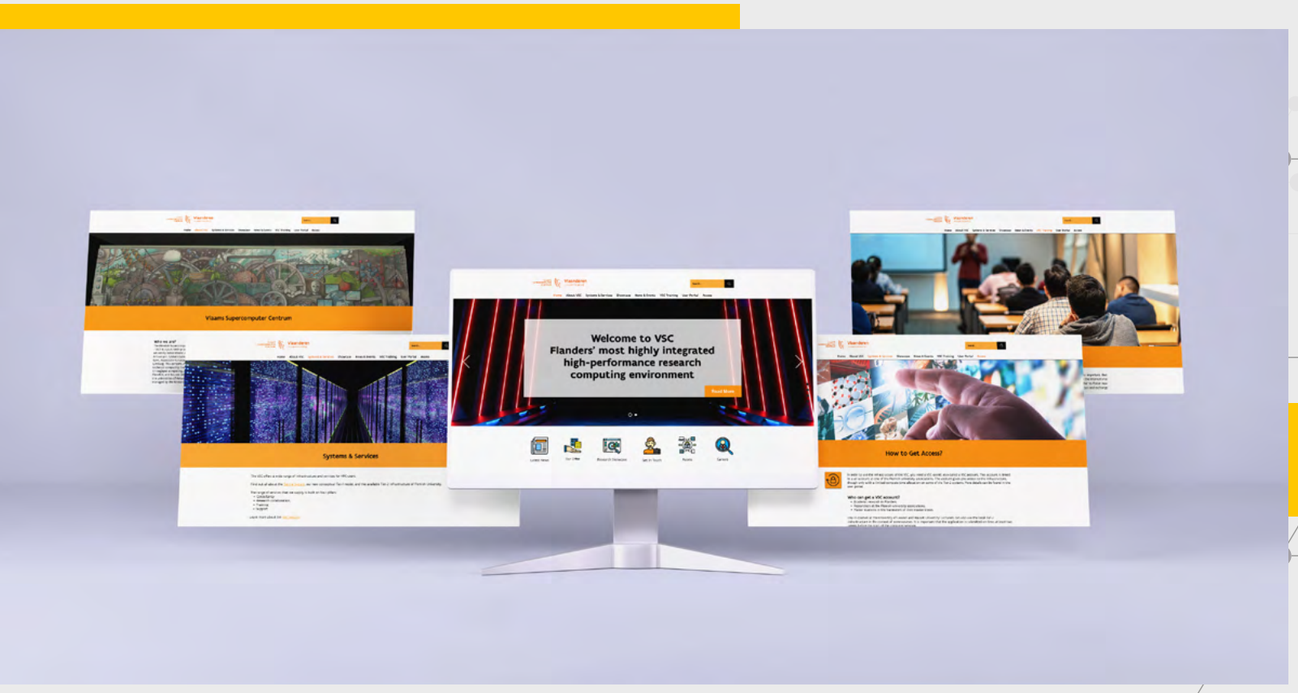

User Engagement

Last year our website saw significant user engagement, with over 26,129 visitors and 57,277 page views.

Technical Aspects

Our website is supported by a robust technical infrastructure, which ensures reliability and performance. Over the past year we have implemented several updates to improve functionality and security..

Accessibility and Usability

We are committed to making our website accessible to all users, with a user-friendly design and easy navigation..

Future Plans

Moving forward, we are planning to introduce new features and enhance the user experience. Our objective is to consistently address the needs of our users and offer the best support possible.

▶ Conclusion

The VSC website is a valuable resource for our research community, providing essential information and resources. We encourage everyone to visit (www.vscentrum.be) for the latest updates and support.

www.vscentrum.be

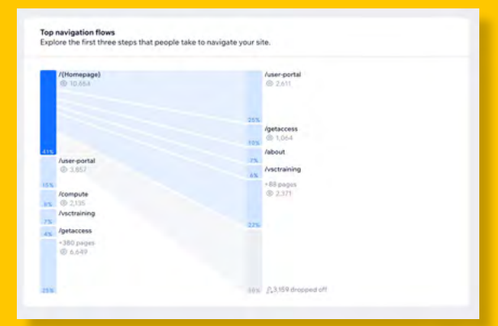

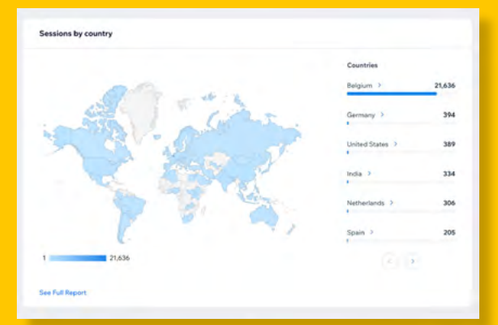

Social Media Engagement and Community Involvement

Expand our presence through social media

In today's digital age, social media plays a crucial role in connecting our community and promoting innovations in high-performance computing (HPC). The Flemish Supercomputer Center has significantly increased its social media presence and uses platforms such as LinkedIn and YouTube to create engagement with HPC enthusiasts, researchers and industry professionals.

Figure 41. VSC LinkedIn rollup banner

LinkedIn's Growth and Impact

Our LinkedIn Page has become a central platform for sharing the latest updates, research breakthroughs and event announcements. Over the past year, we have seen substantial growth in our following, which has increased our ability to disseminate important information and foster professional connections within the HPC community. Regular posts about progress, user stories and collaboration opportunities have not only increased visibility, but also attracted a diverse audience interested in developments within HPC.

YouTube Channel: a visual gateway to HPC

The VSC YouTube Channel (VSC YouTube) serves as a dynamic platform for sharing educational content, webinars and event recordings. Our video content provides an accessible way for users to understand complex HPC concepts, see applications of supercomputing, and stay up to date on VSC activities. The channel's growing subscriber base is a testament to the value of visual learning and the engaging nature of our content.

Engaging the community via the VSC newsletter.

The VSC newsletter, sent quarterly, has become an essential tool for maintaining a continuous connection with our subscribers. Each edition is packed with insights into recent HPC projects, upcoming events and success stories from our users. The newsletter not only keeps our audience informed, but also encourages active participation and feedback, fostering a sense of connection and community among VSC users.

Growing Number of Subscribers

Our newsletter subscriber list is steadily growing, reflecting the increasing interest and engagement within the HPC community. The reach of the newsletter ensures that our latest developments and opportunities are communicated effectively, which stimulates collaboration and innovation.

Research Showcase

Introduced in 2022, the "Research Showcase" page on VSC's website serves as a vibrant showcase of the diverse and cutting-edge research facilitated by VSC's cutting-edge infrastructure. This dynamic blogging platform features tailored versions of original research articles, carefully edited by the research author to highlight the crucial role of VSC in these innovative projects.

The proposed research projects span all disciplines and demonstrate the wide range of scientific research and technological advances supported by VSC. Each featured study not only showcases the researchers' innovative methods and key findings, but also illustrates how VSC's resources were essential in overcoming complex computational challenges and achieving research goals.

"The long-range coupling makes the numerical integration of these equations very challenging, and the large number of simulations required for this project would not have been possible without the HPC infrastructure of the Flemish Supercomputer Center."

Robin Msiska
UGent

"The resources provided by the VSC infrastructure play a major role in reducing the computation time when running our simulation model for many different parameter variables. This makes it possible to investigate the impact of a wide range of scenarios, such as different screening intervals, in a repeated testing policy."

SIMID Research Group
UHasselt, UAntwerp

Through this page, VSC underlines its commitment to advancing cutting-edge research and providing unparalleled support to the academic and scientific community. This page not only celebrates the achievements of researchers, but also inspires future collaborations and discoveries, thus strengthening VSC's position as a cornerstone of scientific progress in Flanders and beyond.

Discover the Research Showcase

"VSC played a crucial role in this research by providing the necessary computational resources for the CFD modeling of the Swirling Flow Reactor."

Thomas Holemans & Zhu Yang
KU Leuven

Success Stories

The following are some testimonials about the use of supercomputers for research, industry, innovation and social benefits. We show examples of how universities and companies, private and public organizations use supercomputers to realize breakthroughs of great scientific or economic importance. These achievements, many of which were achieved thanks to access to very powerful supercomputers and HPC experts at the Flemish Supercomputer Center, have long provided researchers and companies with immense efficiencies, both in working hours and in euros.

Mpacts | Leveraging Expertise

The Mpacts simulation software, designed to model the behavior of granular materials in various industrial processes, has achieved significant performance improvements thanks to the collaboration with the Flemish Supercomputing Center (VSC). Granular dynamics, inherently more complex than molecular dynamics, involves complicated contact detection and rotational degrees of freedom, leading to computationally intensive simulations. One of the biggest challenges for Mpacts has been inefficient memory access since interacting particles are not stored contiguously in memory. Experts from VSC proposed a solution: sorting the particles so that particles that are close to each other in the simulated space are also close to each other in the computer memory. This technique reduced cache misses and significantly improved software performance. Thanks to VSC's support and expertise, Mpacts now works more efficiently and provides faster and more reliable simulations for industrial applications. This collaboration highlights the value of VSC's services in optimizing computational software for high-performance computing environments.

"The opportunity to work with the highly qualified HPC experts at VSC was a very interesting and rewarding experience. I would recommend anyone involved in supercomputing to contact the VSC and take advantage of their available expertise to improve your calculations."

Simon Vanmaercke
Co-founder of MPacts

PUXANO | Matchmaking

Puxano, an innovative biotech company offering structure-based protein research services, distinguishes itself by developing proprietary technologies such as the Procondor platform for automated design of protein constructs. This platform optimizes protein sequences, crucial for effective expression and purification of proteins.

By participating in the 'Call for extended HPC support' of the Flemish Supercomputer Center (VSC), Puxano was introduced to the Data Science Institute of UHasselt (DSI). DSI, with more than 150 data scientists, redefined the scope of the Procondor project and designed an optimal database structure. This collaboration between Puxano, VSC and DSI demonstrates how strategic support and expertise lead to significant advances in biotech research and developmen.

"We are very positive about the collaboration with the academic world. The approach was very professional. Stijn took the time to analyze the problem and all results were delivered on time. He managed to understand our needs by asking the right critical questions. That was very valuable for the success of the project.

This is also what the VSC strives for, as Mia Vanstraelen, chair of the VSC Industrial Council, notes: "VSC is committed to stimulating innovation in Flanders. In addition to offering high-quality HPC infrastructure and consulting services, we focus on bridging knowledge and experiences between academia and industry and vice versa."

Wouter Van Putte
Co-founder & CEO Puxano

Tier-1 Cloud | Eliza Depoorter & Charlotte Peeters

In this success story we dive into the transformative impact of the VSC Tier-1 Cloud platform, specifically through the experiences of Eliza Depoorter and Charlotte Peeters from LM-UGent. Faced with the challenge of managing extensive omics datasets, they turned to our cloud services for a solution tailored to their bioinformatics needs.

At LM-UGent, where leading research in bioinformatics is flourishing, the demand for a robust and flexible environment was essential. VSC Tier-1 Cloud, built on OpenStack, provided them with the perfect ecosystem. They used custom software packages, interactive data analysis tools and workflow portals to streamline their operations. Critical to this was the VSC Tier-1 Cloud's Infrastructure as a Service capability, which allowed them to deploy virtual machines and storage systems that seamlessly fit their scientific workflows.

The ease of setting up databases and web servers using our ready-to-use catalog of templates further improved their efficiency. This enabled Charlotte Peeters and Eliza Depoorter to focus more on their research, resulting in significant time savings and maximizing the impact of their findings. Their success story highlights not only the versatility of our cloud solutions, but also the transformative potential when cutting-edge technology meets cutting-edge research.

"We experience that the VSC infrastructure is very stable and well organized. The available documentation is extensive, and the support is truly excellent."

Eliza Depoorter
LM-UGent

Tier-1 Data | Sebastian Munck & Ingrid Barcena Roig

In a world where performing compute-intensive tasks on large data sets poses significant challenges, finding efficient ways to manage and share research data is essential. The VSC Tier-1 Data platform offers a transformative solution for researchers such as Sebastian Munck, assistant professor at KU Leuven and VIB BioImaging Core. Sebastian's bioimaging projects benefit greatly from the platform's robust metadata and data management capabilities, streamlining workflows and ensuring data integrity and accessibility. Furthermore, Ingrid Barcena Roig, our colleague at KU Leuven, highlights the platform's advanced functionalities, which improve collaboration and accelerate research results. The VSC Tier-1 Data platform stands as a cornerstone in enabling researchers to securely store, manage and share research data, driving innovation within scientific disciplines.

"We benefit greatly from the use of Tier-1 Data by having access to optimized and efficient research data management. We can rely on a great VSC support team and promote fair and open data along with efficient and fast data transfer. We recommend Tier-1 Data because we want to focus on science, acquire Figures and provide the best service to our users. With the VSC, our data is in good hands. We can rely on their expertise and benefit from a large support team that provides the service we need."

Sebastian Munck
Assistant Professor | KU Leuven
VIB BioImaging Core

Diamcad

In the heart of Antwerp's famous diamond district, Diamcad stands as a pioneering force in the art of diamond cutting, driven by cutting-edge technology and the immense computing power of VSC supercomputers. Under the leadership of Bart De Hantsetters, the visionary general manager, and supported by Jonas Tollenaere, a dedicated doctoral student from KU Leuven, Diamcad has embarked on a revolutionary journey.

Armed with the ability to explore millions of potential cut options, Diamcad meticulously creates custom cut patterns tailored to the unique characteristics of each diamond. From size and weight to complex internal flaws, every facet of the diamond is optimized, ensuring unparalleled brilliance and value. This innovative approach not only improves the precision of diamond cutting, but also sets new standards of excellence in the industry. Diamcad's fusion of traditional craftsmanship with advanced computational techniques not only guarantees superior quality, but also ushers in a new era of diamond perfection.

"The collaboration with the VSC went very smoothly from the start. We quickly accessed the cluster and found the website documentation and software libraries sufficient to start our application. In addition, communication about cluster downtimes and maintenance was clear and concise."

Jonas Tollenaere
PhD student | KU Leuven

Atlas Copco | Support

Atlas Copco has had an extremely positive experience working with the VSC, which provides fast and professional support in setting up simulations and software tools. With an extensive library of simulation and calculation tools readily available and new tools quickly installed on request. Tom Saenen, technology developer at Atlas Copco, emphasizes the invaluable importance of the help they have received. Ewald Pauwels, scientific coordinator of the Ghent University High Performance Computing infrastructure at Ghent University & VSC, has played a crucial role in the success of this collaboration.

"Our experience with the VSC team has been very positive whenever we have had a technical question. They were really interested in finding a solution. It really felt like they were working with us to achieve the results that Atlas Copco wanted to achieve. They were ready to share their insights and expertise in solving challenges."

Tom Saenen
Technology Developer
Atlas Copco

International Collaboration

EuroCC

VSC has been participating in the EuroCC 2 project since 2023 as a successor to the EuroCC 1 project that started in 2020. This is the initiative supported by EuroHPC JU and local authorities to maintain a network of more than 30 National Competence Centers (NCCs) in the participating countries. The NCCs act as central access points in each country between stakeholders and national EuroHPC systems. They operate at regional and national levels to engage with local communities, map HPC competencies and facilitate access to European HPC resources for private and public sector users.

EuroCC Belgium joins the forces of VSC, Cenaero and CÉCI, and carries out a diverse range of tasks that includes setting up communication campaigns, providing a formal framework between HPC-related activities, developing and presenting a comprehensive and transparent map of HPC competencies and institutions, and acting as a gateway for industry and academia to providers with suitable expertise or relevant projects, both nationally and internationally. EuroCC Belgium also monitors the HPC training offer and combined it in a central location, together with training courses abroad, collected by other NCCs. Furthermore, it stimulates industrial adoption of HPC and facilitates access to systems and expertise.

In 2023, VSC supported industry-oriented training via EuroCC Belgium, such as an OpenFold training at Puxano Talks, and webinars on Quantum Computing, with 140 participants from companies.

Through outreach campaigns and participation in events such as Tech Meets Chem, Smart ventilation in mid-sized buildings, Data-driven Approaches in Life Sciences, Energy Innovations, ELIXIR, and Advanced Engineering, VSC built contacts with more than 100 companies and recruited 30 companies.

In addition to a direct connection with industry, we also maintain contacts with other organizations and local authorities, such as POMs, VLAIO, BioTope, TTOs, City of Antwerp, VAIA, AGORIA, MCA, Flemish Aerospace Group, imec and Agentschap Nathour & Bos, to identify the need for computer capacity and expertise in different types of organizations.

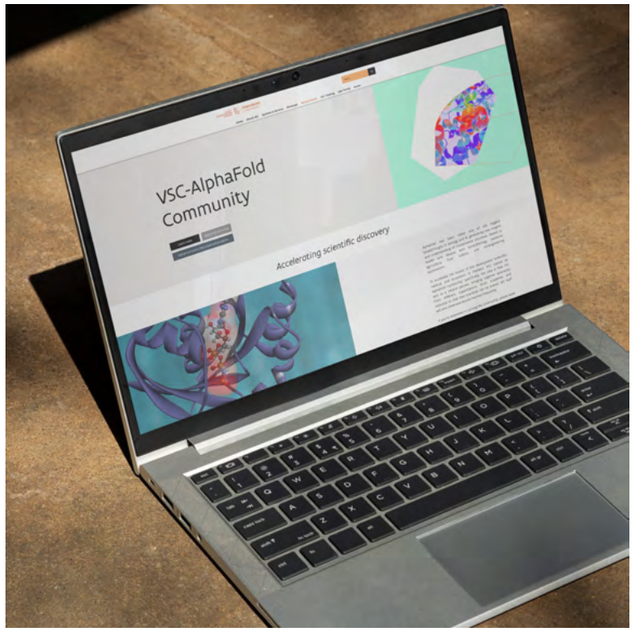

Figure 42. VSC AlphaFold community webpage

Community Building

To maintain good relationships with current user groups and establish contacts with potential new users from academic, industrial or government circles, VSC helps build a number of communities. It does this in the context of a number of themes with a high impact, such as AI. and quantum computing. For example, VSC has already organized a number of events around the significant influence of AI in structural biology and life sciences, given the major impact of tools such as AlphaFold2 on these domains. Furthermore, VSC actively participates in a number of initiatives around quantum computing, and has organized a significant number of webinars in recent years on quantum computing and its future impact on various fields.

PRACE

The PRACE 2 program ended on March 31, 2023, after several extensions. National "resources" are no longer distributed via PRACE calls. The PRACE-IP projects have also come to an end. Calls for actions are now made by EuroHPC.

Given the role that EuroHPC plays in the European landscape, PRACE wants to transform into an association of users and HPC centers in Europe from 2024. It wants to put the user communities at the forefront and therefore put the user at the center. In addition, it wants to provide support for the best possible use of the EuroHPC infrastructure and further promote international cooperation. Serge Bogaerts explained all this during the LUMI-BE user day.

Appendix

Personnel List

Institution	Name	% employment within the framework of HPC
UGent	Wouter Depypere	100%
UGent	Stijn De Weirdt	100%
UGent	Alvaro Simon Garcia	100%
UGent	Andy Georges	100%
UGent	Kenneth Hoste	100%
UGent	Ewald Pauwels	100%
UGent	Balazs Hajgato	100%
UGent	Kenneth Waegeman	100%
UGent	Bart Verheyde	100%
UGent	Danny Schellemans	10%
UGent	Johan Van Camp	30%
UGent	Fréderic De Leersnijder	20%
UGent	Wim Waeyaert	20%
UGent	Bruno Cardon	20%
UGent	Koen Van Hecke	5%
UGent	Tim De Roeck	15%
UGent	Debby Elbers	10%
UGent	Ruth Degroote	10%
UGent	Michel Raes	15%
UGent	Myriam Mertens	10%
UGent	Dieter Roefs	60%
UGent	Denis Krišťák	100%
UGent	Filip Kružík	100%
UGent	Petr Kral	25%
UGent	Hafsa Naeem	25%
UGent	Martin Sakin	100%
UGent	Peter Hardon	100%
KU Leuven	Herman Moons	20%
KU Leuven	Leen Van Rentergem	30%
KU Leuven	Jan Ooghe	100%
KU Leuven	Jan De Laet	100%
KU Leuven	Maxime Van den Bossche	100%
KU Leuven	Mustafa Dikmen	100%
KU Leuven	Mag Selwa	100%
KU Leuven	Alexander Vapirev	100%
KU Leuven	Ingrid Barcena	100%
KU Leuven	Ehsan Moravveji	100%
KU Leuven	Steven Vandenbrande	100%
KU Leuven	Jef Scheepers	100%
KU Leuven	Wouter Van Assche	100%
KU Leuven	Wouter Lampaert	100%
KU Leuven	Louis Roeben	100%
KU Leuven	Jo Vandeginste	50%
KU Leuven	Yorick Poels	100%
KU Leuven	Peter veraedt	100%
KU Leuven	Guy Van Den Bergh	100%

Personnel List

Institution	Name	% employment within the framework of HPC
KU Leuven	Rudy Rys	40%
KU Leuven	Tom van Mierlo	100%
KU Leuven	Tom Vanhout	20%
KU Leuven	Koen Vanherck	30%
KU Leuven	Lennart Peeter	100%
KU Leuven	Bart Vanneste	100%
KU Leuven	Sofie Pieraerd	10%
KU Leuven	Philip Brusten	10%
KU Leuven	Filip Severants	100%
KU Leuven	Jo Wijnant	100%
UAntwerp	Franky Backeljauw	100%
UAntwerp	Stefan Becuwe	100%
UAntwerp	Kurt Lust	50%
UAntwerp	Carl Mensch	20%
UAntwerp	Michele Pugno	60%
UAntwerp	Engelbert Tijskens	100%
UAntwerp	Robin Verschoren	100%
UAntwerp	Annie Cuyt	5%
UAntwerp	Wim Vanroose	5%
UAntwerp	Koen Decauwsemaecker	10%
UAntwerp	Herwig Kersschot	10%
UAntwerp	Wim Holemans	5%
UAntwerp	Bertin Peeters	5%
VUB	Alex Domingo Toro	100,00%
VUB	Samuel Moors	100,00%
VUB	Stéphane Gérard	100,00%
VUB	Ward Poelmans	100,00%
VUB	Jan Turek	58,33%
VUB	Cintia Willemyns	58,33%
VUB	Olivier Devroede	30,00%
VUB	Joachim Verschelden	15,00%
VUB	Johan D'Hondt	10,00%
VUB	Micky Mattens	10,00%
VUB	Serge Morabito	10,00%
VUB	Stefan Weckx	10,00%
VUB	Alain Graulus	5,00%
VUB	Henri Vanroelen	5,00%
VUB	Luc Eulaerts	5,00%
VUB	Marc Pierloot	5,00%
VUB	Bernadette Van Gansbeke	2,00%
VUB	Jan Paredis	2,00%
UHasselt	Geert Jan Bex	100%
UHasselt	Rafal Al-Takreeti	50%

Tier-1 Compute

Title	Applicant	Institution	Department
Growth of Nitrogen doped diamond: dynamics, bonding and charge transfer.	Vanpoucke, Danny Eric Paul	UHasselt	Institute for Materials Research (IMO-IMOMEC)
Discovery of selective DPP8/9 inhibitors: Identification of hydrophobic affinity sites	Beyens, Olivier	UAntwerp	UAMC / Department of Medicinal Chemistry
Free energy perturbations for the optimization of DPP9 selective inhibitors	Beyens, Olivier	UAntwerp	UAMC / Department of Medicinal Chemistry
Studies of fluid-scale nonlinear surface waves in collisionless plasmas	Magyar, Norbert; Bacchini, Fabio	KU Leuven	Centre for mathematical Plasma Astrophysics, Department of Mathematics
Extracting Meaning from Multilingual Text with Deep Neural Networks	Vanroy, Bram	KU Leuven (& UGent)	Computational and Formal Linguistics (ComForT) / Center for Computational Linguistics (CCL)
Characterizing the adsorbate-induced response of metal-organic frameworks for applications in carbon capture and water harvesting from desert air.	Goeminne, Ruben	UGent	Center for Molecular Modeling / EA17
Hydrogen diffusion pathways in a clathrate	Lamaire, Aran	UGent	Center for Molecular Modeling
Evaluating the role of acidity on alkene protonation in zeolites	Cnudde, Pieter	UGent	Center for Molecular Modeling
Potential Energy surfaces of Hydrogen on Iron surfaces for the Exploration of the Hydrogen Embrittlement Effect in Pipeline Steel	Meier, Lukas	UGent	Center for Molecular Modelling
Selective catalytic conversion of CO2 to ethanol over high entropy alloys	Yan, Dengxin	UGent	Prof. Mark Saeys's group at the Laboratory for Chemical Technology (LCT) in the Department of Materials, Textiles and Chemical Engineering
Computational Fluid Dynamics Based Process Intensification of CO2 Capture: Solvent Regeneration in a Gas-Liquid Vortex Reactor	Chen, Siyuan	UGent	Laboratory for Chemical Technology Department of Materials, Textiles and Chemical Engineering (EA11)
Effect of phase mixing on collisionless plasma turbulence	Bacchini, Fabio; Pucci, Francesco	KU Leuven	CmPA, Department of Mathematics
High-throughput screening of covalent organic frameworks for carbon capture	De Vos, Juul	UGent	Center for Molecular Modeling (EA17)
FineSegment	Picron, Cédric	KU Leuven	ESAT-PSI, Visics
High-troughput simulations of antisolvent addition to solvated polymer solutions - continued	Denayer, Mats	VUB	DSCH - ALGC
The first general relativistic kinetic simulations of black hole jet formation in the center of the Milky Way	Küchler, Lorenzo; Ripperda, Bart; Hertog, Thomas; Vercnocke, Bert; Mayerson, Daniel;	KU Leuven	theoretical physics group, physics department
Operational MHD modelling of the solar corona for space-weather applications	Brchnelova, Michaela	KU Leuven	Centre for mathematical Plasma Astrophysics (CmPA), Department of Mathematics
Conformational dynamics of membrane proteins	Mathew, Alen T	UGent	Centre for Molecular Modelling
Electron scattering simulation with deep learning molecular dynamics	Zhang, Zezhong	UAntwerp	EMAT / Department of Physics
Genotype imputation from low-pass cell-free DNA sequence data	Becelaere, Sara	KU Leuven	Laboratory for Human Evolutionary Genetics / Department of Human Genetics
First-principles investigation of superconductivity in transition metal dichalcogenide bilayers	Bekaert, Jonas	UAntwerp	Condensed Matter Theory (CMT), Department of Physics

Tier-1 Compute

Title	Applicant	Institution	Department
Probing the stability and flexibility of Interleukin-12 (IL-12) and extracellular IL-12 ligand-receptor complexes	Mehdipour, Ahmad Reza	UGent	Center for Molecular Modeling
Tuning properties of magnetic monolayers under strain and strain gradient	Pandey, Tribhuwan	UAntwerp	Department of Physics/ CMT group
Unsteady aerodynamic simulations of an airborne wind energy reference system in realistic flight and atmospheric conditions using computational fluid dynamics.	Pynaert, Niels	UGent	Department of Electromechanical, Systems and Metal Engineering
Iterative scheme for Collective variable discovery for Cesium Lead Iodide	Devoogdt, David	UGent	Center for Molecular Modeling / EA17
Superconductivity in intercalated bilayer kagome borophene	Soskic, Bozidar	UAntwerp	Condensed Matter Theory (CMT) - Department of Physics
Computational and Experimental Study of Grain Growth in Stainless Steels Containing Inclusion Particles	Gupta, Aman	KU Leuven	Nano- and Microstructure Design of Materials / Department of Materials Engineering (MTM)
Grain boundaries in UO2	Mayda Bacaksiz, Selma and Arts, Ine	UAntwerp	EMAT / Department of Physics
Investigating the effect of mutations on ABL-imatinib binding kinetics using a novel path-sampling methodology.	Vervust, Wouter	UGent	IBiTech – BioMMeda, department EA06
First principles analysis of stacking schemes and functional group mobility through molecular dynamics simulations of covalent organic frameworks	Vanlommel, Siebe	UGent	Center for Molecular Modeling / Department of applied physics
Coronal mass ejections and solar energetic particles: modelling their effects in the inner heliosphere.	Wijsen, Nicolas	KU Leuven	Centre for mathematical Plasma Astrophysics (CmPA), Dept. of Mathematics
Molecular Dynamics Simulations on UAMC-0001305 Warhead Derivatives to Theragnostically Target Fibroblast Activation Protein	Joep Wals	UAntwerp	Medicinal Chemistry Group
Computational Fluid Dynamics (CFD) study on the oxidation of Volatile Organic Compounds (VOCs) in a photocatalytic reactor.	Mohammad Rusydi Fatahillah, Yi Ouyang and Geraldine Heynderickx	UGent	Laboratory for Chemical Technology (LCT); Department of Materials, Textiles, and Chemical Engineering (EA11)
Membrane permeation study of a series of radical trapping agents using umbrella sampling simulations	De Winter, Hans	UAntwerp	Laboratory of Medicinal Chemistry / Department of Pharmaceutical Sciences
Advanced Document Image Processing using Large Scale Language Models	Tan LU	VUB	Digital Mathematics (DIMA), Department of Mathematics and Data Science
Direct Numerical Simulations of stably stratified Ekman layers over heterogeneous surfaces	Thijs Bon and Johan Meyers	KU Leuven	Turbulent Flow Simulation and Optimization, Department of Mechanical Engineering
Kinetic Modelling of Novel Solar Wind Observations	Luca Pezzini and Fabio Bacchini	KU Leuven	Centre for mathematical Plasma-Astrophysics (CmPA), Department of Mathematics
Large-eddy simulations of wind farms in stable boundary layers	Luca Lanzilao, Steven Vandenbrande and Johan Meyers	KU Leuven	Turbulent Flow Simulation and Optimization (TFSO) Research Group, Mechanical Engineering Departmen
Characterization of the flow field around a staple fiber yarn on microscale level using computational fluid dynamics	Axel Bral and Joris Degroote	UGent	Department of Electromechanical, Systems and Metal Engineering

Tier-1 Compute

Tier-1 Compute

Title	Applicant	Institution	Department
A new reaction route of initial stage methanol to olefins conversion via frustrated Lewis pair involving trimethyloxonium ion	Wei Chen and Veronique Van Speybroeck	UGent	Center for Molecular Modeling
Variational transition state theory investigation of the barrierless reactions involved in atomic-scale diamond growth mechanisms	Emerick Guillaume, Danny Vanpoucke, Luc Henrard and Ken Haenen	UHasselt	Material Physics (IMOMAF)
Next Level Flemish Speech Recognition - Flemish Speech Recognition At Scale	Jakob Poncelet	KU Leuven	Electrical Engineering ESAT – PSI, Speech Processing group
Unsteady computational fluid dynamics reactor design to accelerate the electrification of steam cracking.	Mike Bonheure, Yi Ouyang and Kevin Van Geem	UGent	Laboratory for Chemical Technology (LCT)-Department of Materials, Textiles and Chemical Engineering (EA11)
Theoretical DFT study of silicates' vibrational spectra in solution	Fileto Rodríguez and Frederik Tielens	VUB	Materials modelling group, ALGC, Chemistry department
Anatomy-aware representation learning of large medical images	Joris Wuts	VUB	ETRO (Department of Electronics and Informatics)
Multimodal Egocentric Vision	Gorjan Radevski and Dusan Grujicic	KU Leuven	PSI-VISICS / ESAT
Conformational dynamics of Covid Spike protein: Effects of glycosylation and nanobody interactions	Alen T. Mathew and Ahmadreza Mehdipour	UGent	Centre for Molecular Modelling
Rationalizing Phase Selection in Inorganic Zeolite Synthesis Through Extensive DFT Calculations	Jelle Vekeman and Toon Verstraelen	UGent	Centre for Molecular Modelling
Investigating Water Splitting Reaction on Covalent Organic Framework through Molecular Dynamics Simulations	Kuber Singh Rawat and Veronique Van Speybroeck	UGent	Centre for Molecular Modelling
Evaluating the adsorption properties of metal-organic frameworks for applications in olefin/paraffin separation and detection of volatile organic compounds in air.	Siddharth Ravichandran	UGent	Centre for Molecular Modelling
Modelling particle acceleration and transport at a 3-D coronal mass ejection-driven shock	Zheyi Ding, Zheyi Ding and Zheyi Ding	KU Leuven	Centre for mathematical plasma-astrophysics
Model Compression of Large Image Generative Models	Junyi Zhu, Han Zhou and Matthew Blaschko	KU Leuven	PSI / ESAT
Quantifying adsorption-diffusion barriers for unique adsorption sites on external zeolite surfaces	Jenna Mancuso, Massimo Bocus and Pieter Cnudde	UGent	Centre for Molecular Modelling
Investigating Metal-Phosphonate Properties in NU-1000 Structure through DFT Studies and Spectroscopy Simulations (Part of PHOSPORE Project)	Roberth Mateo Narvaez Adams, Frederik Tielens and Tom Hauffman	VUB	Materials Modelling Group, General Chemistry Research Group (ALGC),
Numerical simulation of re-entry vehicle aerothermodynamics using in-house CFD solver.	Vatsalya Sharma	KU Leuven	Centre for Mathematical Plasma Astrophysics (CmPA), Department of Mathematics
Developing SO-CASSI Energies for Te and Po molecules using Machine Learning	Okan Koeksal and Stefaan Cottenier	UGent	Department of Electromechanical, Systems and Metal Engineering and Center for Molecular Modeling
Calculating Diffusion Barriers of Hydrogen through Alloyed Iron Surfaces for the Exploration of the Hydrogen Embrittlement Effect in Pipeline Steel	Lukas Meier and Stefaan Cottenier	UGent	Center for Molecular Modelling
Characterization of wind farm wakes and weather conditions for offshore wind farms in the North Sea	Wim Munters, Alexandros Palatos-Plexides and Simone Gremmo	von Karman Institute for Fluid Dynamics	Department of Environmental and Applied Fluid Dynamics
Tier1 proposal - Power-2-Olefins: Electrified rotor-stator reactor for steam cracking – Geometry optimization using meta-models	Rejish Lal Johnson Samee Lal, Yi Ouyang and Kevin Van Geem	UGent	Department of Materials, Textiles and Chemical Engineering
Accurate free energy profile of trimethyloxoium ion evolution in initial stage methanol to olefins conversion by umbrella sampling	Wei Chen and Veronique Van Speybroeck	UGent	Center for Molecular Modelling
Validation of CFD liquid metal pool-type nuclear reactor thermal-hydraulics simulations by reduced scaled experiments	Silvania Lopes and Lilla Koloszar	von Karman Institute for Fluid Dynamics	Environmental and Applied Fluid Dynamics Department
Uncertainty quantification of the forces acting on a tube bundle in axial two-phase flow	Henri Dolfen and Joris Degroote	UGent	Department of Electromechanical, Systems and Meal Engineering
Machine learning potentials for an accurate modelling of alkene chemisorption in zeolites	Massimo Bocus, Pieter Cnudde and Sander Vandenhaute	UGent	Center for Molecular Modelling
Taming Dutch Llamas: Finetuning open-source LLMs on Dutch data	Bram Vanroy	KU Leuven	Computational and Formal Linguistics (ComForT)
Dynamo's in accretion disks: how did the magnetic field get there?	Bart Ripperda, Thomas Hertog, Lorenzo Küchler, Daniel Mayerson and Bert Vercnocke	KU Leuven	theoretical physics group, physics department
Cluster-based active learning to model spatially disordered metal-organic frameworks	Pieter Dobbelaere	UGent	Center for Molecular Modelling
Tier-1 Compute application: Hyperbolic contrastive learning for remote sensing images	Liang Zeng	KU Leuven	Geomatics, Faculty of Engineering Technology
Driven decay-less oscillations of loops in a stratified solar atmosphere.	Konstantinos Karampelas and Tom Van Doorsselaere	KU Leuven	CmPA / Department of Mathematics
Calculation of the Infrared spectra for the adsorption of citrate ions on hydroxyapatite surfaces and the carbonate-substituted hydroxyapatite surfaces by ab initio molecular dynamics simulations	Yuheng Zhao, Frederik Tielens and Ionut Tranca	VUB	ALGC material modelling group
Evaluating the diffusion properties of metal-organic frameworks for applications in olefin/paraffin separations.	Bernd Schmidt, Ruben Goeminne and Siddharth Ravichandran	UGent	Center for Molecular Modeling/ EA17
Llama-NL: a Dutch Llama-2-based, open-source Large Language Model	Veronique Hoste, Els Lefever, Thomas Demeester, Anthony Rathé and Matthieu Meeus	UGent	Language and Translation Technology Team (LT3)/ Internet and Data Science Lab (IDLab)
Control of magnetic states in a Ni-halide monolayer with in-plane strains: an ab initio study	Ali Ghojavand and Milorad Milosevic	UAntwerp	Condensed Matter Theory (CMT)/ Department of Physics
The effect of Mg on a Ni during CO2 methanation	Servaas Lips	UGent	Department of Materials, Textiles and Chemical Engineering
Discrete element method (DEM)-based mechanistic modelling of powder flow in a gravimetric feeder for continuous pharmaceutical production)	Luz Naranjo and Ashish Kumar	UGent	Pharmaceutical Engineering Research Group (PharmaEng), Department of Pharmaceutical Analysis
Radiation General Relativistic magnetohydrodynamics simulations of formation of black hole accretion disk coronae in active galactic nuclei	Bart Ripperda, Fabio Bacchini and Matthew Liska	KU Leuven	Centre for mathematical Plasma Astrophysics (CmPA)
Automated Food Intake Monitoring with Multi-Modality Sensors and Deep Learning	Chunzhuo Wang, Hans Hallez and Bart Vanrumste	KU Leuven	ESAT-STADIUS and the e-Media Research Lab
Conformational dynamics of membrane proteins	Alen T. Mathew and Ahmadreza Mehdipour	UGent	Center for Molecular Modelling (Department of Applied Physics)
Computational characterization of interlayer exciton states in multilayer transition metal dichalcogenides	Cem Sevik and Milorad Milosevic	UAntwerp	Condensed Matter Theory Group / Physics
Long Read Sequencing for the detection of methylome changes in patients with developmental disorders	Erika Souche, Benjamin Huremagic and Joris Vermeesch	KU Leuven	Laboratory of Cytogenetics and Genome Research

Tier-1 Compute

Title	Applicant	Institution	Department
Understanding Solar Energetic Particle Discrepancies and Space Weather Forecast Uncertainties	Nicolas Wijsen and Christine Verbeke	KU Leuven	Centre for mathematical Plasma Astrophysics (CmPA)
Deep neural network study on satellite imagery and Flemish biological valuation map	Mingshi Li	KU Leuven	PSI-VISICS / ESAT
Application Tier-1 Hortense, López Figueiras, Emilio vsc46149	Emilio López Figueiras and Jure Oder	von Karman Institute for Fluid Dynamics	Aeronautics and Aerospace Department
Magnons in moiré heterobilayers	Maarten Soenen	UAntwerp	Condensed Matter Theory (CMT)
Application form: Compute component of the Flemish Tier-1 supercomputing platform	Tingyu Qu and Wei Sun	KU Leuven	Language Intelligence and Information Retrieval Lab, Declarative Languages and Artificial Intelligence section / Department of Computer Science
Exploring the design space of Solid Micellar Catalysts	Sara Santos, Konstantijn Rommens and Dengxin Yan	UGent	Department of Materials, Textiles and Chemical Engineering
Simulating Ruddlesden-Popper defects in CsPbI3 due to Zn and Cd doping	Tom Braeckevelt, David Devoogdt and Mieke De Schepper	UGent	Center for Molecular Modeling / EA17
Theoretical characterization of black TiO2	Selma Mayda Bacaksiz	UAntwerp	EMAT / Department of Physics
Design of a microkinetic model for Fischer-Tropsch synthesis on Ni at CO saturation coverages	Konstantijn Rommens, Sara Santos and Dengxin Yan	Universiteit Gent	Department of Materials, Textiles and Chemical Engineering
Global 3D GR-PIC simulations of non-thermal flares in the magnetosphere of Kerr black holes	Ileyk El Mellah and Fabio Bacchini	KU Leuven	CmPA
Development and applications of efficient methods to study conformational transitions in biomedically relevant proteins	Kenno Vanommeslaeghe	VUB	Analytical Chemistry, Applied Chemometrics and Molecular Modelling (FABI)

Tier-1 Data

Title	Applicant	Institution	Department
PiDGiN: Future proof pathology for predictive medicine and disease prognosis based on tumor heterogeneity	Melvin Geubbelmans, Christel Faes, Dirk Valkenborg, Jari Claes, Esther Wolfs, Kim Nijsten, Michiel Thomeer and Sandrina Martens	UHasselt	Data Science Institute (DSI) – BIOMED – LCRC-ZOL
AIDefSpace, Using Artificial Intelligence to defend telecommunications and satellite positioning systems from the interference of space weather events	Giovanni Lapenta and Ekaterina Dineva	KU Leuven	"Center for Mathematical Plasma Astrophysics / Departement Wiskunde"
Tier-1 Data Application form: Adaptation Genomics	Steven Van Belleghem	KU Leuven	Eco-Evolutionary Genomics / Biology
Understanding genetic variation in Parkinson's disease through long-read and single-cell multi-omics analyses	Stein Aerts	VIB–KU Leuven	"Laboratory of Computational Biology, VIB Center for AI & Computational Biology, VIB-KU Leuven Center for Brain & Disease Research and Department of Human Genetics KU Leuven"

Tier-1 Cloud

Title	Applicant	Institution	Department
"Hosting clinical data in a harmonized data model to participate in the European Health Data and Evidence Network (EHDEN)"	Marcel Parciak, Lotte Geys, Liesbet Peeters, Ashkan Pirmani and Hamza Khan	UHasselt	Research group in Biomedical Data Sciences / Biomedical research institute (BIOMED) and Data Science Institute (DSI)
Annotation database for oligogenic variant analysis.	Tom Lenaerts and Emma Verkinderen	VUB	Artificial Intelligence lab / Interuniversity Institute of Bioinformatics Brussels (IB)2 (Vakgroep Computerwetenschappen)
Tier-1 Data Application form: Adaptation Genomics	Steven Van Belleghem	KU Leuven	Eco-Evolutionary Genomics / Biology
Understanding genetic variation in Parkinson's disease through long-read and single-cell multi-omics analyses	Stein Aerts	VIB–KU Leuven	Laboratory of Computational Biology, VIB Center for AI & Computational Biology, VIB-KU Leuven Center for Brain & Disease Research and Department of Human Genetics KU Leuven

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 Brussels
T 02 512 91 10

Published by
Dr. Hans Willems
Secretary General FWO

Text
Leen Van Rentergem, Jan Ooghe, Ingrid Barcena
Annie Cuyt, Stefan Becuwe
Johan Van Camp, Ewald Pauwels
Geert Jan Bex
Stefan Weckx, Ward Poelmans
Caroline Volckaert

Creation and Design
Rafal Al-Tekreeti

Published in **July 2024**

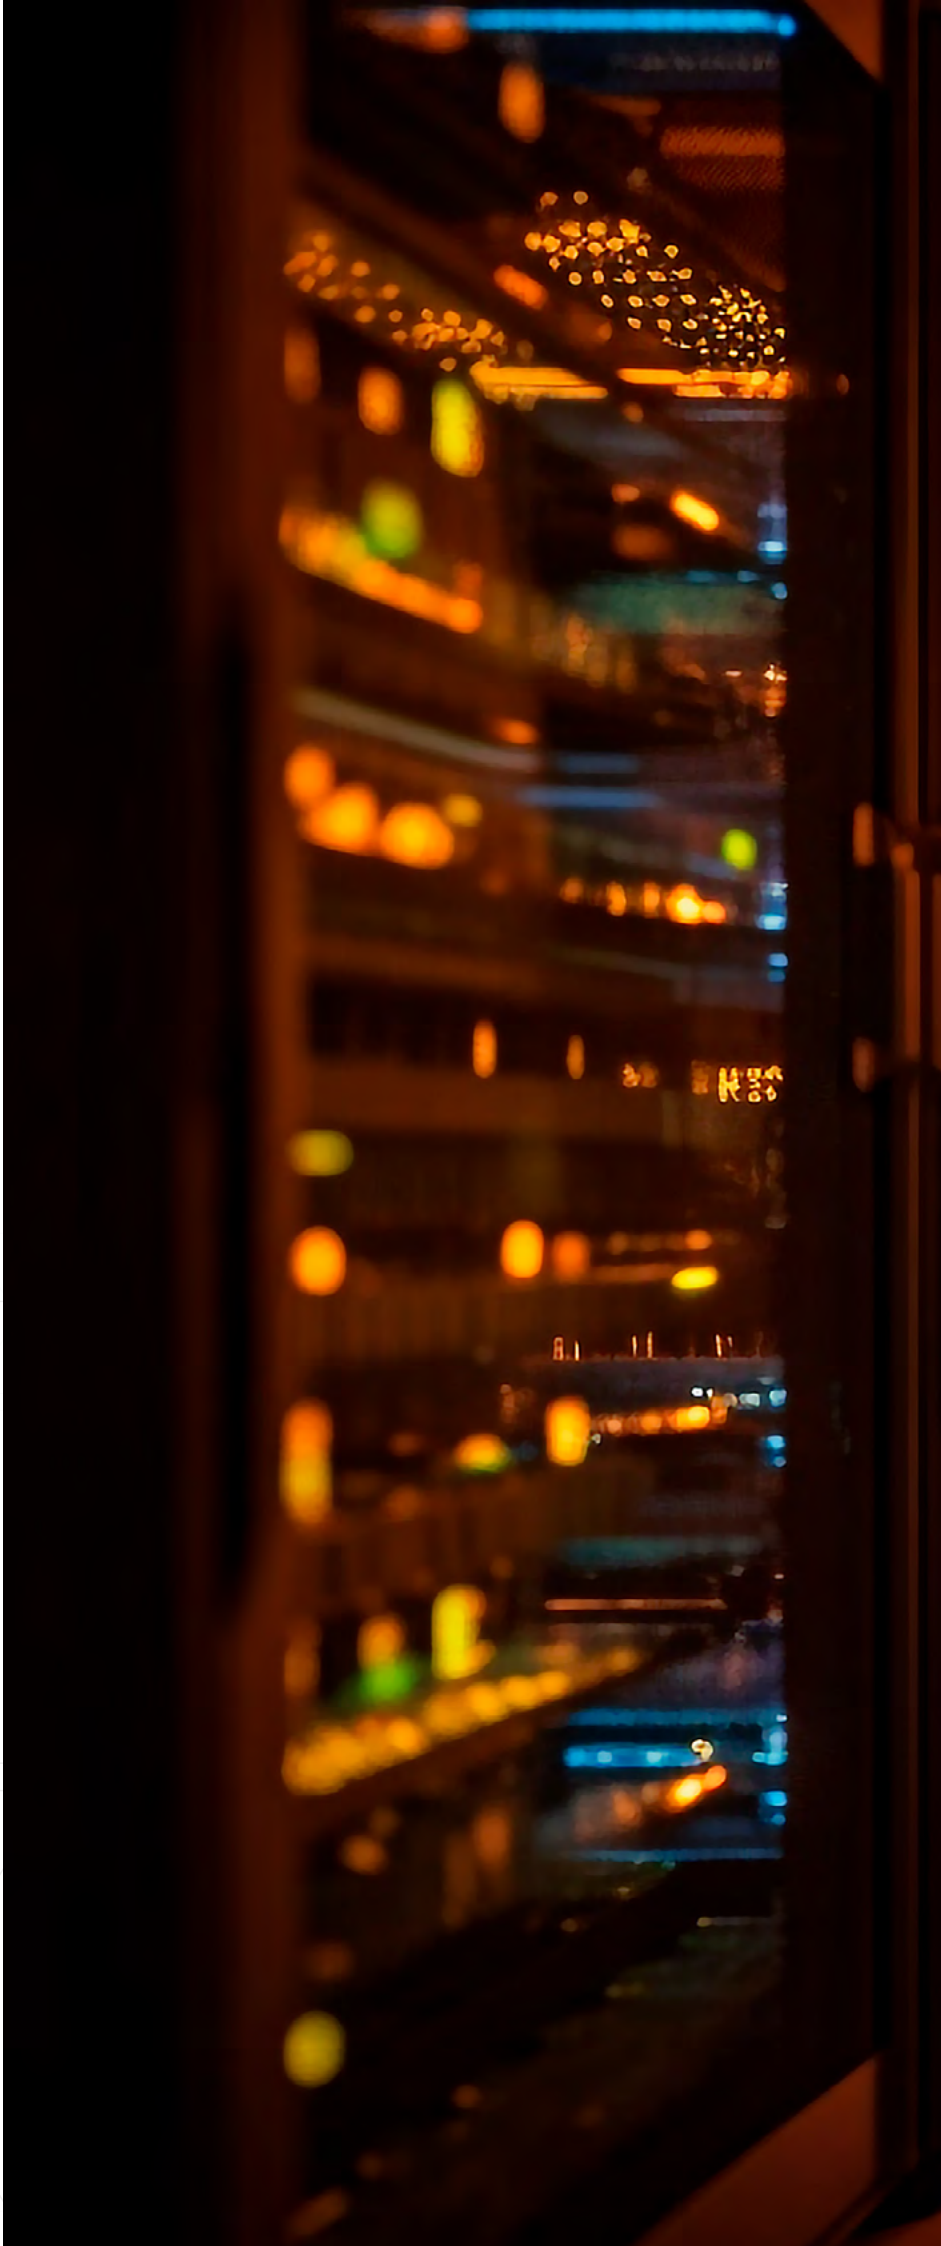

vscentrum.be