Flanders
State of the Art

ANNUAL REPORT 2017

VLAAMS SUPERCOMPUTER CENTRUM

Innovative Computing for a Smarter Flanders

vscentrum.be

COLOPHON

The Flemish Supercomputer Centre (VSC) is a virtual centre making supercomputer infrastructure available for both the academic and industrial world. This centre is managed by the FWO in partnership with the five Flemish university associations.

Research Foundation - Flanders
Egmontstraat 5
1000 Brussels
T 02 512 91 10
info@vscentrum.be
www.vscentrum.be

Published by
dr. Hans Willems
Secretary General FWO

Texts
Leen Van Rentergem
Jan Ooghe
Annie Cuyt
Stefan Becuwe
Ewald Pauwels
Geert Jan Bex
Stefan Weckx
Ward Poelmans
Caroline Volckaert

Creation and design
Pantarein Publishing
Rafal Al-Tekreeti

Published in August 2018

PREFACE

In 2017, the Flemish Supercomputer Centre (VSC), the partnership between the Flemish universities and the FWO, was running at full throttle, providing HPC infrastructure and support to the broad research community.

BrENIAC, the Flemish Tier-1 machine, was available virtually without any downtime in 2017, and the utilisation rates show that its use is well established among researchers. For the evaluation of the applications for Tier-1 computing time, the international panel was expanded with an extra member. Applications were evaluated on three occasions. A total of 54 applications were approved and 164,886 node days allocated.

The user committee organised a successful Users Day with national and international speakers, workshops and a well-received poster session. The event attracted more than 120 participants.

The Flemish universities continued to expand their Tier-2 capacity and additional staff resources were provided for training and user support. The VSC in fact ensures that researchers can easily migrate their applications between the university clusters and the Tier-1 supercomputer so that the most suitable infrastructure is used at all times.

Furthermore, training courses on the use of Tier-1 and Tier-2 are open not only to researchers from Flemish universities, strategic research centres and other public institutions of knowledge, but also to companies.

Attention was also paid to informing companies and other institutions of knowledge about the possibilities provided by the VSC. The Industrial Board developed initiatives to highlight the advantages of a collaboration with the VSC: professional support, customised training where appropriate, but above all the embedding within an academic environment.

However, the developments go very fast. Challenges within the research landscape require adequate responses. The VSC remains alert to the needs of its users. That is why, starting from 2018, a new Tier-1 supercomputing platform "Supercomputing as a service" will be developed to provide an integrated Tier-1 compute, data, and cloud infrastructure service. At the end of 2017 the VSC received 30 million euros for this purpose. It illustrates the commitment of the Flemish government to continuing investing in state-of-the-art infrastructure for research and innovation.

The VSC is thus assured of a bright yet challenging future!

KU Leuven: Leen Van Rentergem, Jan Ooghe
UAntwerp: Annie Cuyt, Stefan Becuwe
UGent: Ewald Pauwels
UHasselt: Geert Jan Bex
VUB: Stefan Weckx, Ward Poelmans
FWO: Caroline Volckaert, Bart Van Beek

ABOUT VSC

Introduction

The support and operation of High Performance Computing (HPC) in Flanders is organised within the Flemish Supercomputing Centre, VSC. The VSC is a consortium of the 5 Flemish universities. Its objective is to offer HPC infrastructure and support to the broad research community in Flanders. The VSC also aims, through targeted initiatives, to inspire, inform and support researchers from the universities, strategic research centres and the industry to incorporate computational models into their business. In addition, as a service provider, the VSC provides a range of training programmes aimed at promoting the use of the infrastructure. The VSC is managed by the FWO (Research Foundation - Flanders).

Funding of Tier-1 and Tier-2

In 2017, the Tier-2 infrastructure was co-funded by the FWO for a total sum of € 5,606,000, with € 3,706,000 investments in hardware and € 1,900,000 in staff. This means that the FWO supported 5 additional FTEs as compared to previous years. In addition, the FWO invested € 380,000 in staff for the two Tier-1s and the energy cost (€ 360,000) of the second Tier-1.

These resources were used for funding of:

- staff costs for the operation of the first and second Tier-1 (€ 389,000);
- staff costs for training and support of users of both Tier-1 and Tier-2 (€ 1,900,000);
  (For this amount, the five Flemish universities can subsidise in total the equivalent of 20 FTEs.)
- energy costs of the second Tier-1 (€ 360,000);
- investment and operating costs for the Tier-2 infrastructure (€ 3,706,000).

  (The universities have mainly used this last amount for additional investments in Tier-2. These institutions finance the energy and operating costs of the Tier-2 from their own resources.)

The table below presents a summary of the distribution of these amounts over the five Flemish universities.

Summary Tier-1 and Tier-2 funding in 2017

Heading	KU Leuven	UHasselt	UGent	VUB	UAntwerp	Subtotal Tier-2	Tier-1a	Tier-1b	Total
staff (FTE)	5	1	4	2	3	15	2	2	19
Staff (Euro)	€ 475.000	€ 95.000	€ 380.000	€ 190.000	€ 285.000	€ 1.425.000	€ 190.000	€ 190.000	€ 1.805.000
Extra staff	€ 95.000	€ 95.000	95.000	€ 95.000	€ 95.000	€ 475.000			€ 475.000
Total staff	**€ 570.000**	**€ 190.000**	**€ 475.000**	**€ 285.000**	**€ 380.000**	**€ 1.900.000**			**€ 2.280.000**
Hercules key 2015	0,4209	0,0341	0,3233	0,1011	0,1206	1,000			
Operation Tier-2	**€ 1.559.855**	**€ 126.375**	**€ 1.198.150**	**€ 374.677**	**€ 446.944**	**€ 3.706.000**			**€ 3.706.000**
Energy Tier-1								€ 360.000	€ 360.000
Belnetcon-nection	€ 35.000					€ 35.000			€ 35.000

For the allocation of the subsidies, the FWO concluded an agreement with each university in which it was agreed that the spending of the allocated resources can be spread over two budgetary years. The allocated resources must be justified with receipts and a spending report must be submitted about the use to which the funding has been put, which should include information about the use of the Tier-1 and Tier-2 infrastructure.

CASE

WILL COMPUTERS MAKE ANIMAL TESTS OBSOLETE?

Prof. dr. Hans De Winter
Department of Pharmaceutical Sciences,
University of Antwerp

Supercomputers are playing an ever-increasing role in the search for new and improved medicines. Partly thanks to the increasing use of computer simulations, the number of animal tests in pharmaceutical research has fallen significantly over the past ten years," says Prof. Dr. Hans De Winter from University of Antwerp.

Across the pharmaceutical industry and academia, researchers are working day in day out on improved and new drugs. To simulate the interactions between a potential drug and the protein linked to a certain syndrome, the "computational chemist" uses algorithms that are mainly based on Newton's second law. These simulations - known as "molecular dynamics calculations" in technical jargon - provide researchers with a detailed picture of the kinetics and thermodynamics of the interactions within the model system.

The importance of computing power

Until ten years ago, the use of molecular dynamics simulations in the research for new drugs was virtually non-existent because the computing power was not available. In those days, researchers had to rely mainly on test animals and costly experiments. "Today we run simulations capable of simulating biological kinetics within 1 second," says Prof. Dr. Hans De Winter from University of Antwerp. "This is made possible by fast computer clusters such as the Flemish Tier-1 system and the CalcUA cluster at the University of Antwerp, coupled with the growing use of graphics processors and more efficient parallelization algorithms of molecular dynamics software. Processes that are sequential by nature are split up into a large number of parallel computer simulations. Using Markov state models and their statistical analysis, the results of these parallel simulations enable us to calculate kinetic and thermodynamic parameters describing the process." Some results of this research were recently published in *Nature Communications and Scientific Reports*.

Towards a world without animal tests?

Could this research herald the end of animal testing? Hans De Winter: "This seems utopian, at least in the short to medium term. The gap between the test tube and the complexity of an animal is still too wide. However, compared with ten years ago, the number of animal tests in pharmaceutical research has already been reduced significantly, also due to the growing use of computer simulations. We will never be able to avoid animal tests altogether, just like clinical tests on humans, but the aim is to reduce them as far as possible in the long term."

Prospects

The use of computer simulations within pharmaceutical research is constantly evolving. "In the coming years, new algorithms, in combination with ever-faster computer systems, will continue to raise the bar for drug research," Hans De Winter concludes.

TIER-1 INFRASTRUCTURE

Tier-1 at UGent

The first Flemish Tier-1 went into production in 2013. The purchase of this supercomputer was financed with FFEU funding. UGent was responsible for the technical operation of this machine, and also financed the accommodation. The maintenance contract for this machine expired in November 2016, and the supercomputer was officially decommissioned on 31 December 2016, a few months after the new Tier-1 supercomputer at KU Leuven had gone into production. All users of the 'old' Tier-1 were timely and repeatedly informed about the end of term arrangements and given the opportunity to retrieve their data from the storage during the first two months of 2017 before they, too, became inactive.

Following a cost-benefit analysis, it was decided to reconfigure the still usable hardware of this machine to an on-demand cloud cluster, as part of a pilot setup in preparation for a 'new Tier-1 supercomputing as a service' concept that will be presented by VSC in 2018. After all researchers had retrieved their data from the storage, comprehensive reconversion work was initiated, including:

- rework of Tier-1a network connectivity within data center
- clearing of defective hardware
- consolidation of still operational hardware
- reconfiguration of all switches, management nodes and servers
- installation on-demand rollout
- configuration maintenance
- pilot user support

The goal of this pilot set-up is to gain experience within the VSC and to experiment with this technology for the purpose of better identifying user needs and to allow potential challenges to be identified as soon as possible. The graph below shows the consumption by pilot users in 2017.

Also after 2017, the decommissioned Tier-1a continues to be useful. In 2018, upon completion of the pilot phase, this machine will be made available to end users as cloud infrastructure in production.

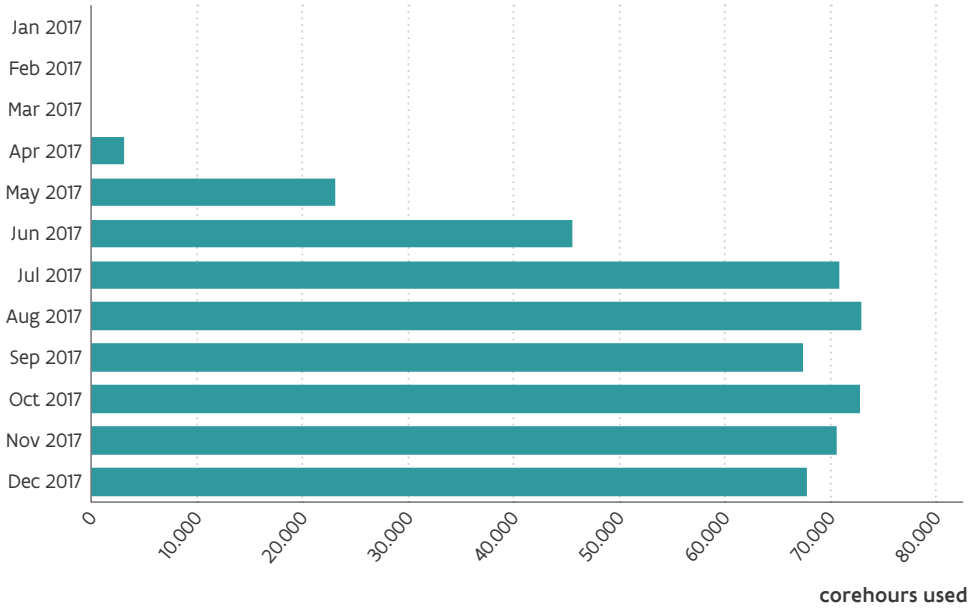

Figure 1 - Tier-1 usage at UGent in 2017 - compute

Tier-1 at KU Leuven

BrENIAC, the second Flemish Tier-1 machine went into production in October 2016.

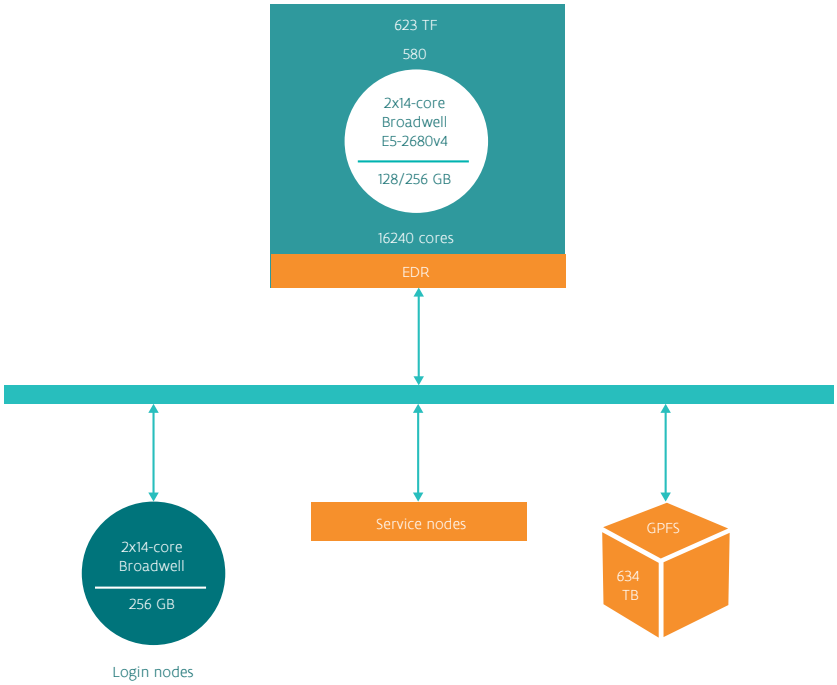

Figure 2 - Tier-1 BrENIAC

The machine ran stable in 2017. The time graph shows that the machine was available virtually without any downtime in 2017. There was one short, unplanned downtime due to an overall power outage.

There were 3 project rounds in which a total of 77 applications for Tier-1 computing time were submitted. Of these, 54 were approved. Projects of 4 different allocation rounds were active in 2017 (2016-3, 2017-1, 2017-2 and 2017-3). A lower occupancy rate was recorded in May and to a lesser extent in September. In both cases, this coincided with the end of an allocation period during which only 1 allocation round was active for 2 months. It was therefore decided to extend the allocation periods to 8 months, so that projects of two allocation rounds are continuously active. This results in a better spreading of the computing time. Also in 2017, the allocated computing time was used up almost completely (2016-3 93%, 2017-1 96%, 2017-2 91%).

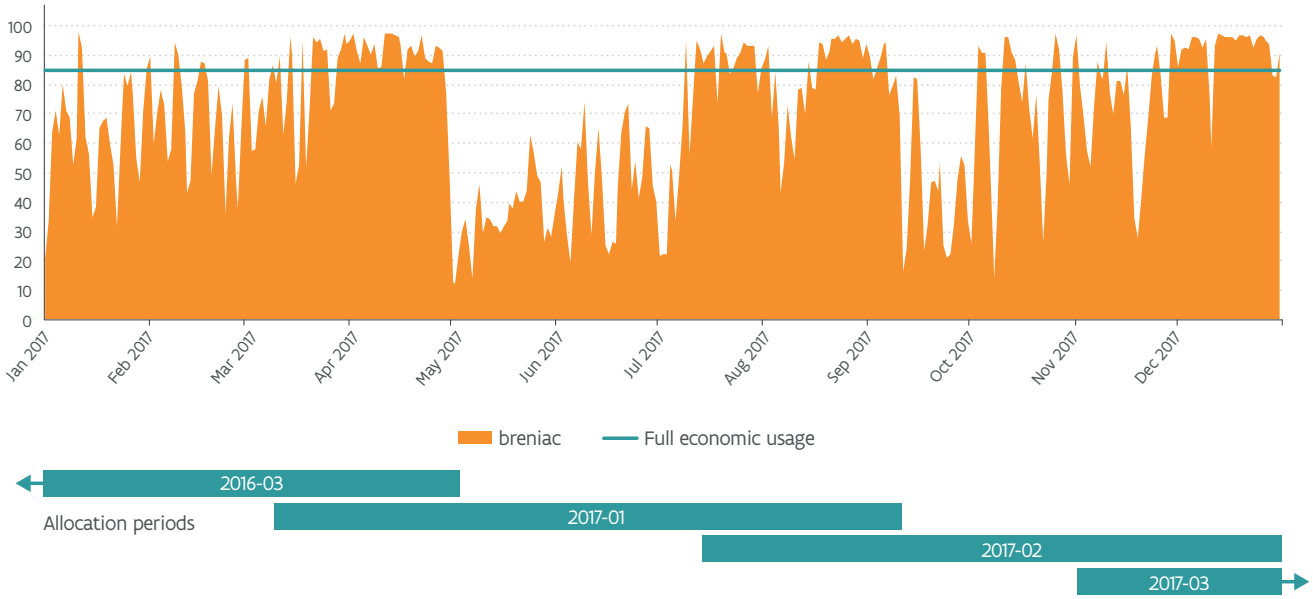

Figure 3 - BrENIAC % CPU Used 2017

The Tier-1 users were surveyed after 2017. 30% of the users responded to the survey. These 30% participated in 40 of the 73 projects that were active in 2017.

The survey showed that overall satisfaction was high:

How satisfied are you overall with the use of Tier-1?	
Very Satisfied	**45 %**
Satisfied	**50 %**
Neutral	**5%**

Many researchers need the Tier-1 infrastructure to analyse normal, day-to-day research issues.

How important are the Tier-1 resources for your research?	
Access to the VSC Tier-1 is essential for our research, our normal workload exceeds the capacity of the current Tier-2 environment.	**77 %**
Access to the VSC Tier-1 Access infrastructure is necessary to resolve specific research issues, the normal workload can be handled on Tier-2.	**23 %**

95% of the respondents indicate they will continue submitting Tier-1 projects in the future. The use of Tier-1 shows that it is well established within the existing research community. The following distribution shows that the different fields of research are well represented. Just as with HPC clusters abroad, chemistry, physics and the engineering fields are well represented. It will, of course, remain important in the future to cover as many research fields as possible.

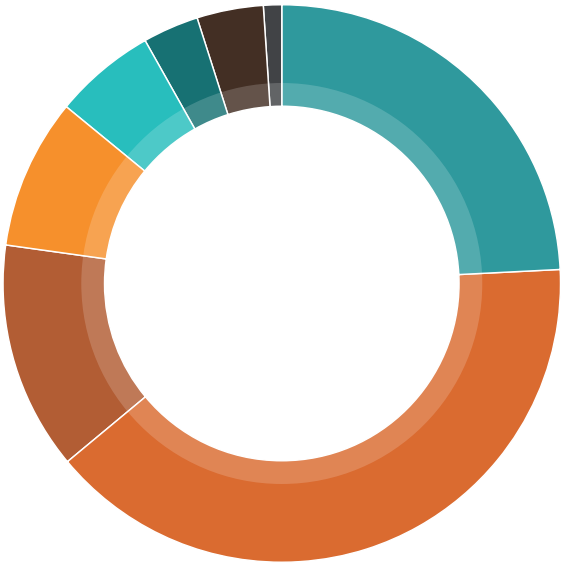

24% > Technology
40% > Molecular Modelling
13% > Physics
9% > Astronomy and astrophysics
6% > Life sciences
3% > Psychology
4% > Chemistry
1% > Earth sciences

Figure 4 - Used time by science field

Computing on the Tier-1

ALLOCATION OF TIER-1 COMPUTING TIME

There are a number of ways in which researchers can apply for computing time on Tier-1, as stipulated in the regulations.

For researchers associated with a university, an SOC, or equivalent research institution, the following access channels are available:

- Starting Grant
  - Maximum 100 node days of computing time
  - Can be applied for at any time, with quick turnaround time
  - To try out Tier-1 and perform benchmarks or software tests, as preparation for a full-scale project application
  - Free of charge
- Project Access
  - For allocations of 500-5000 node days of computing time
  - Project applications describe
    the overall scientific project;
    the consortium of users that will carry out the computations;
    the funding channel;
    the computational tasks to be carried out (technical);
    the software to be used;
    where appropriate, the scientific results obtained with previous Tier-1 project allocations.
  - Project applications can be submitted at all times, but are evaluated on 3 occasions during the year by the Tier-1 Allocation Board.
  - Free of charge

For the evaluation of Tier-1 project applications a 'Tier-1 Allocation Board' was installed. From 2017, the Evaluation Commission was extended with one member and now includes five foreign experts:

- Walter Lioen, chairman (SURFsara, the Netherlands);
- Derek Groen (Computer Science, Brunel University London, UK);
- Sadaf Alam (CSCS, Switzerland);
- Nicole Audiffren (Cines, France);
- Gavin Pringle (EPCC, Edinburgh University, Scotland).

Mrs Caroline Volckaert of the FWO provides the secretariat. The HPC coordinators of the Flemish universities may be invited to participate as observers in the meetings of the Tier-1 Allocation Board.

The board evaluates the applications and decides whether the requested computing time can be allocated in whole, in part, or not at all.

From 2016, no more costs will be charged for Tier-1 projects allocated to academic researchers. Irrespective of the (limited) amount of the cost charged, this nonetheless constituted a clear obstacle for researchers and prevented less experienced research groups from gaining access to new, challenging but high-risk research on the Tier-1 supercomputer.

For industrial researchers, there are also two access channels for Tier-1 computing time:

- Exploratory Access
  - Maximum 100 node days of computing time
  - To try out the Tier-1 user environment or to perform benchmarks or software tests.
  - Free of charge
- Full Access
  - Companies can conclude an agreement with the Tier-1 housing institution and FWO to purchase computing time
  - Full cost charging of consumed computing time and used storage

The rates at which industrial users can buy Tier-1 computing time were laid down in the Access Regulations 2017. In addition, industrial researchers can also gain access to Tier-1 as part of a research project that is carried out in collaboration with a public research institution, e.g. a Flemish university.

TIER-1 STARTING GRANTS / EXPLORATORY ACCESS

In 2017, 40 Starting Grants were awarded and 2 Exploratory Access projects were ongoing.

Starting Grants 2017

Applicant	Host institution	Science field
Tatiana Woller	VUB	Theoretical chemistry
Stefan Knippenberg	UHasselt	Biophysics
Bercx Marnik	UAntwerp	Nanophysics/technology
Camila Scolini	KU Leuven	Astronomy and astrophysics
Christine Verbeke	KU Leuven	Astronomy and astrophysics
Stefan Weckx	VUB	Bioinformatics/biostatistics
Bonan Bertrand	KU Leuven	Earth sciences
Darko Stosic	UAntwerp	Physics
Charlotte Vets	UAntwerp	Theoretical chemistry
Emmanuel Chané	KU Leuven	Astronomy and astrophysics
Chrils Ulens	KU Leuven	Neurobiology
Lucas Delcou	UGent	Mechanical engineering
Laurent DeMoerloose	UGent	Civil engineering
Matthieu LeRoy	KU Leuven	Astronomy and astrophysics
Niels Souverijns	KU Leuven	Environmental science
Ileyk Elmellah	KU Leuven	Astronomy and astrophysics
Siegfried Cools	UAntwerp	Unknown
Kirit Makwana	KU Leuven	Astronomy and astrophysics
dimitrios millas	KU Leuven	Astronomy and astrophysics
Juan Jose Gutierrez Sevillano	UGent	Physical chemistry
Krisztina Fehér	UGent	Chemistry
Hui Zhao	KU Leuven	Bioinformatics/biostatistics
Toon Verstraelen	UGent	Theoretical chemistry

Diego Gonzalez	KU Leuven	Astronomy and astrophysics
Norbert Magyar	KU Leuven	Astronomy and astrophysics
David Gobrecht	KU Leuven	Astronomy and astrophysics
Cole Johnston	KU leuven	Astronomy and astrophysics
Francesco Contino	VUB	Mechanical engineering
Elisabetta Boella	KU Leuven	Astronomy and astrophysics
Pierre De Buyl	KU Leuven	Theoretical physics
Dries Allaerts	KU Leuven	Mechanical engineering
Athanasios Vitsas	KU Leuven	Mechanical engineering
Panos Tsirikoglou	VUB	Mechanical engineering
Wim Munters	KU Leuven	Mechanical engineering
Maarten Vanloo	KU Leuven	Geography
Wenzhi Ruan	KU Leuven	Astronomy and astrophysics
Hans De Winter	UAntwerp	Bioinformatics/biostatistics
Abhishek Dutta	KU Leuven	Materials technology
Jan Turek	VUB	Chemistry
Frederik Tielens	VUB	Chemistry

Exploratory Access

Applicant	Company
Marcus Drosson	Umicore
Gareth Linsmith	ApheaBio

APPROVED TIER-1 APPLICATIONS

All applications approved in 2017, grouped per evaluation point, are listed below.

06 February 2017

Applicant	Host institution	Department	Title	Allocated computing time (node days)	Allocated SCRATCH storage (TB)
Danny Vanpoucke	UHasselt	Materials physics (IMOMAF)	Breathing behavior of flexible mixed metal Metal-Organic Frameworks	3100	0,85
Jolan Wauters	UGent	Department of Flow, Heat and Combustion Mechanics	Optimization of a dual feather wing tip geometry using CFD	2460	2,5
Pieter Reyniers, David Van Cauwenberge, Laurien Vandewalle	UGent	Laboratory for Chemical Technology	Computational Fluid Dynamics based design of a novel reactor technology for the Oxidative Coupling of Methane	2000	1
Gilberto Santo	UGent	Department of Flow, Heat and Combustion Mechanics	Computational Fluid Dynamics simulation of wind turbines	2379	0,175

Wilfried de Corte	UGent	Department of Data Analysis, Faculty of Psychology and Educational Sciences	Targeting Key Features that Determine the Robustness and Sensitivity of Pareto-optimal (PO) Selection Designs	4044	0,025
Oriana De Vos, An Ghysels	UGent	Center for Molecular Modeling	Effect of periodic boundary conditions on simulations of oxygen transport through membranes	1260	0,02
Arthur De Vos, Veronique Van Speybroeck, Kurt Lejaeghere	UGent	Center for Molecular Modeling	Electronic properties of 3D nitrogen-containing Covalent organic Frameworks from First-Principles Simulations	4140	2,71
Sven Rogge, Ruben Demuynck, Steven Vandenbrande, Veronique Van Speybroeck	UGent	Center for Molecular Modeling	Full quantum mechanical study of the influence of functionalization and temperature on phase transformations in metal-organic frameworks	2450	0,245
Jelle Wieme, Kurt Lejaeghere, Steven Vandenbrande, Veronique Van Speybroeck	UGent	Center for Molecular Modeling	Assessing the accuracy of hybrid functionals for the relative stability of a flexible MOF	3420	0,21
Jonas Bekaert, Milorad Milosevic, Bart Partoens	UAntwerp	Condensed Matter Theory / Physics Department	First-principles study of the superconducting properties of ultrathin transition metal dichalcogenides	4802	0,25
Marnik Bercx	UAntwerp	EMAT / Department of Physics	Study of redox reactions in Li-Rich layered oxides	3742	0,06
Johan Meyers, Ali Emre Yilmaz	KU Leuven	Turbulent Flow Simulation & Optimization (TFSO) Research Group Mechanical Engineering Department	LES-Based Optimal Control Studies of Wind Farms with Advanced Turbine Models	4964	90
Johan Meyers, Dries Allaerts, Athanasios Vitsas	KU Leuven	Turbulent Flow Simulation & Optimization (TFSO) Research Group Mechanical Engineering Department	Simulation of operational offshore wind farms for comparison with experimental data	4920	7,5
Samuel Moors, Tatiana Woller	VUB	Department of Chemistry, group ALGC	Modeling aromatic chlorination reactions with ab initio molecular dynamics	3233	0,021
Stefan Weckx	VUB	Research Group of Industrial Microbiology and Food Biotechnology, Faculty of Sciences and Bioengineering Sciences	Metagenomic data analysis to unravel food fermentation processes	1500	1

05 June 2017

Applicant	Host institution	Department	Title	Allocated computing time (node days)	Allocated SCRATCH storage (TB)
Charlotte Vets, Erik Neyts	UAntwerp	Research group Plasmant, Department of Chemistry	Adsorption energies of carbon nanotubes on bimetallic catalysts: impact for chirality-selective growth	4980	1
Jelle Wieme, Aran Lamaire, Veronique Van Speybroeck	UGent	Center for Molecular Modeling	Understanding the high-pressure behavior of a flexible nanoporous material	2430	0,81
Arthur De Vos, Veronique Van Speybroeck, Kurt Lejaeghere	UGent	Center for Molecular Modeling	Electronic properties of 2D nitrogen-containing Covalent organic Frameworks from First-Principles Simulations	2128	1,5
Oriana De Vos, An Ghysels	UGent	Center for Molecular Modeling	Simulating oxygen transport through membranes with various lipid compositions	985	0,005
Sam De Waele, Stefaan Cottenier, Kurt Lejaeghere	UGent	Center for Molecular Modeling	Analysis of the off-stoichiometry of Fe3N1+y with Density-Functional Theory	3456	4,378
Darko Stosic, Milorad Milosevic	UAntwerp	Condensed Matter Theory / Physics Department	Large-scale Ginzburg-Landau simulations of Single-Photon Detectors	5000	1
Ruben Demuynck, Sven Rogge, Veronique Van Speybroeck	UGent	Center for Molecular Modeling	Studying the mechanical and thermal stability of UiO-66 with a full quantum mechanical description	4320	0,27
Michael Sluydts, Stefaan Cottenier, Veronique Van Speybroeck	UGent	Center for Molecular Modeling	Determining temperature-dependent formation energies for point defects in Ge using the HSE06 functional [resubmit]	4895	4,096
Pieter Cnudde, Simon Bailleul, Kristof De Wispelaere, Veronique Van Speybroeck	UGent	Center for Molecular Modeling	Ab initio study on the stability of cracking intermediates	4920	0,081
Steven Vandenbrande, Toon Verstraelen, Veronique Van Speybroeck	UGent	Center for Molecular Modeling	Ab initio calculation of the Henry constant of methane in (functionalized) UiO-67	2400	1,008
Ileyk El Mellah, Jannis Teunissen	KU Leuven	Centre for mathematical Plasma Astrophysics / Department of Mathematics	Wind accretion on to compact objects in Supergiant X-ray binaries	1468	0,8

Kristof De Wispelaere, Veronique Van Speybroeck	UGent	Astrophysics / Department of Mathematics	Mechanistic investigation of the early stages of the methanol-to-hydrocarbons conversion	3640	0,9
Chiara Caratelli, Julianna Hajek, Veronique Van Speybroeck	UGent	Center for Molecular Modeling	Probing the strength of basic sites on UiO-66 using pKa calculations [Resubmit]	4450	0,9
Dimitrios Millas, Bart Ripperda, Rony Keppens	KU Leuven	CmPA, Department of Mathematics	Outflows and particle evolution in relativistic astrophysics (OPERA) IV	1250	2,5
Frank Peelman, Ewald Pauwels	UGent/VIB	VIB-UGent Center for Medical Biotechnology	In silico study of lymphoma-related MyD88 mutations and their effect on protein actions and mechanism.	1677	0,015
Christine Verbeke, Camilla Scolini, Stefaan Poedts	KU Leuven	CmPA, Department of Mathematics	Modeling the evolution of interplanetary Coronal Mass Ejections: inclusion of a magnetic flux rope and coupling to a magnetospheric model	1144	2,6
Pieter Reyniers, Laurien Vandewalle	UGent	Department of Materials, Textiles and Chemical Engineering (EA11); Laboratory for Chemical Technology	Computational Fluid Dynamics based design of a novel reactor technology for the Oxidative Coupling of Methane (II)	3700	2
Jonas Bekaert, Milorad Milosevic, Bart Partoens	UAntwerp	Condensed Matter Theory / Physics Department	Ab initio study of the influence of atomic defects and strain on superconductivity in ultrathin transition metal dichalcogenides	4729	0,25
Eliot Boulanger	KU Leuven	Theoretical and Computational Chemistry Group / Chemistry Department	Linking Protein Motion and Function: A Combined Molecular Dynamics and QM/MM Study	1000	1
Alexandra Gossart, Niels Souverijns, Nicole van Lipzig, Jan Ooghe	KU Leuven	Regional Climate Studies / Department of Earth and Environmental Sciences	Long-term hindcast climate simulation over Antarctica, using the coupled COSMO-CLM model	572	0,1
Matthieu Leroy	KU Leuven	CmPA, Department of Mathematics	The Kelvin-Helmholtz instability at the magnetosphere : impact of the environmental parameters on the penetration of high energy plasma	835	1

02 October 2017

Applicant	Host institution	Department	Title	Allocated computing time (node days)	Allocated SCRATCH storage (TB)
Norbert Magyar	KU Leuven	CmPA (Centre for mathematical Plasma Astrophysics)	Wave heating and turbulence in the solar atmosphere	2500	0,5
Diego Gonzalez Herrero	KU Leuven	CmPA (Centre for mathematical Plasma Astrophysics)	Study of magnetic cusp properties with Particle In Cell simulations.	2200	1
Chun Xia	KU Leuven	CmPA (Centre for mathematical Plasma Astrophysics)	Formation and eruption of solar prominences	1700	1,2
Wilfried De Corte	UGent	Department of Data Analysis, Faculty of Psychology and Educational Sciences	A Sample to Population Cross-Validation Approach to Assess the Robustness and Sensitivity of Pareto-optimal (PO) Selection Designs	3888	0,025
Michael Sluydts	UGent	Center for Molecular Modeling	Discovering the temperature dependence of charged defect concentrations in Ge using HSE06.	3887	3,072
Pieter Reyniers	UGent	Department of Materials, Textiles and Chemical Engineering (EA11), Laboratory for Chemical Technology	Computational fluid dynamics based design of a novel reactor technology for the oxidative coupling of methane (III)	3700	1
Jolan Wauters	UGent	Department of Flow, Heat and Combustion Mechanics	Robust Game Theory based optimization of a wing fence geometry using CFD	3600	1
Jonas Bekaert	UAntwerp	Condensed Matter Theory / Physics Department	The quest for superconductivity in atomically thin noble metals	4928	0,218
Jelle Wieme	UGent	Center for Molecular Modeling	Benchmarking elastic properties of a metal-organic framework	1296	0,158
Simon Bailleul	UGent	Center for Molecular Modeling	Enhanced sampling study of the methylation of ethene, propene and trans-2-butene.	3300	0,0795

Johan Meyers	KU Leuven	Turbulent Flow Simulation & Optimization (TFSO) Research Group Mechanical Engineering Department	Annual energy production (AEP) estimates of Cabauw and Fino1 based on large eddy simulations	4992	3,4
Johan Meyers	KU Leuven	Turbulent Flow Simulation & Optimization (TFSO) Research Group Mechanical Engineering Department	Optimal dynamic induction and yaw control of wind farms in the atmospheric boundary layer	4912	6,35
Johan Meyers	KU Leuven	Turbulent Flow Simulation & Optimization (TFSO) Research Group Mechanical Engineering Department	Simulation of the Lillgrund offshore wind farm for loads analysis	4560	10,545
Eliot Boulanger	KU Leuven	Theoretical and Computational Chemistry Group / Chemistry Department	Linking Protein Motion and Function: A Combined Molecular Dynamics and QM/MM Study	4000	1
Diether Lambrechts	KU Leuven/VIB	Laboratory of Translational Genetics / Faculty of Medicine	Optimization of long-read sequencing mapping to discover biomarkers for cancer immunotherapy	5255	15
Oriana De Vos	UGent	Center for Molecular Modeling	Simulating oxygen transport through membranes at various temperatures	660	0,003
Elisabetta Boella	KU Leuven	CmPA / Department of Mathematics	First principles simulations of magnetic reconnection in the solar corona	0	1
Kirit Makwana, Bart Ripperda, Dimitros Millas, Ronny Keppens	KU Leuven	CmPA, Department of Mathematics	Fluid to kinetic modeling of the magnetic island coalescence problem	1115	0,86

The Tier-1 Supercomputing platform from 2018

Apart from the structural financing for Tier-2, the Flemish government decided at the end of 2017 to earmark an amount of € 30 million for the realisation of a Flemish Tier-1 supercomputing platform from 2018 onwards. With the implementation of a new platform for HPC "Supercomputing as a service", the VSC also seeks to attract new users and to provide a more integrated service to existing users.

The new platform will comprise three linked environments, notably a Tier-1 compute infrastructure, a Tier-1 data infrastructure and a Tier-1 cloud infrastructure. This should make the HPC infrastructure as a whole more deployable and accessible to all users. This impulse financing will further broaden the use and increase confidence in the presence of such a HPC infrastructure in Flanders. If also the support can be further developed in the coming years, this financial injection will certainly promote the embedding of computational research.

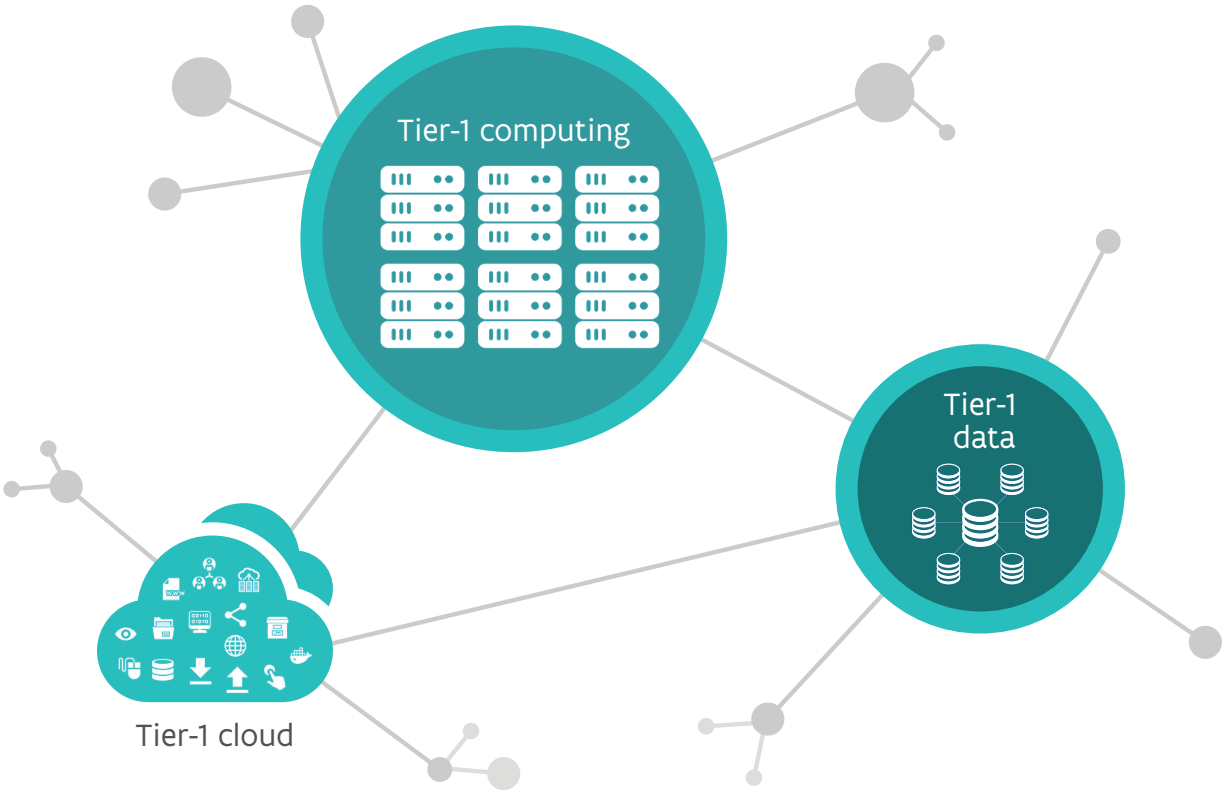

Figure 5 - Tier-1-supercomputing-platform

CASE

INNOVATIVE WATERMARK TRACKS DOWN DIGITAL PIRATES

PhD researcher Hannes Mareen
IDLab, imec-Ghent University

The economic impact of digital piracy is estimated at several billion euros. Researcher Hannes Mareen from IDLab (imec – Ghent University) has conceived a method for identifying digital pirates who illegally share videos.

Digital piracy is a major concern for the film industry. In addition to its economic impact, it also deters filmmakers from investing in new content. "This eventually results in less choice for consumers," says Hannes Mareen from IDLab

Mareen has devised a method to better track down digital pirates. If the makers provide a customer or reviewer (who is authorised, for example, to review a preview) with a legal video, a unique fingerprint or watermark is added to the video. This makes it possible to subsequently identify the individual who illegally distributed the video online.

Unique watermark

Hannes Mareen: "I have developed a new method that makes the watermark invisible and prevents it from being removed. When the encoder compresses the video to a manageable storage size, it deliberately introduces changes during the encoding process. The result is a unique combination of small errors distributed across the entire video. The number of errors is enormous, but they are so subtle that viewers cannot recognise the watermark, let alone remove it. If the video is shared illegally, a computer can easily read out the unique combination of encoding errors and thereby identify the infringer."

Furthermore, a smart encoder ensures that the watermark can be added rapidly, making the method suitable for large-scale use. Hannes Mareen: "I am currently further refining this technique as part of my PhD grant strategic basic research of the FWO."

422,000 computer hours

For the development of his new watermarking technique, Hannes used the VSC's extensive computational capacity. He tested new compression algorithms and studied efficient detection methods. He spent a total of almost 422,000 core hours of computer time on the Tier-2 clusters of the VSC UGent hub. This corresponds to continuous use of a modern laptop for more than 12 years.

Awards

Hannes carried out the research as part of his Master's thesis "A novel video watermarking approach based on implicit distortions", under the supervision of Prof. Peter Lambert, Glenn Van Wallendael and Johan De Praeter (UGent).

He received several awards for his research:

- Best Poster Award at the research symposium of the UGent Faculty of Engineering and Architecture (FEARS) 2018.
- Winner of the Agoria Award 2017 (as part of the Flemish Dissertation Award). This award is given to the best Master's thesis in engineering.

LINKS

IDLab (http://idlab.ugent.be/)
Hannes Mareen wins 10th Agoria prize (https://www.scriptieprijs.be/nieuws/hannes-mareen-ugent-winnaar-tiende-editie-agoriaprijs)

PUBLICATIONS

"Piraten gekielhaald", EOS Magazine 2018, p. 58-59 (http://hdl.handle.net/1854/LU-8557646)
"A novel video watermarking approach based on implicit distortions", ICCE2018, p. 543-544 (http://hdl.handle.net/1854/LU-8546255)

TIER-2 INFRASTRUCTURE

Operation and use

This section will provide an overview of the Tier-2 infrastructure available within the various Flemish universities. Its use will also be illustrated.

KU LEUVEN AND UNIVERSITY OF HASSELT

For the Tier-2 infrastructure, KU Leuven and the University of Hasselt work together.

The infrastructure consists of:

- 2 clusters / 7 partitions
- 244 TF
- 8256 CPU / 71808 accelerator cores
- 44 TB memory

The new cluster on the scheme was purchased in 2017, but its implementation is scheduled for 2018.

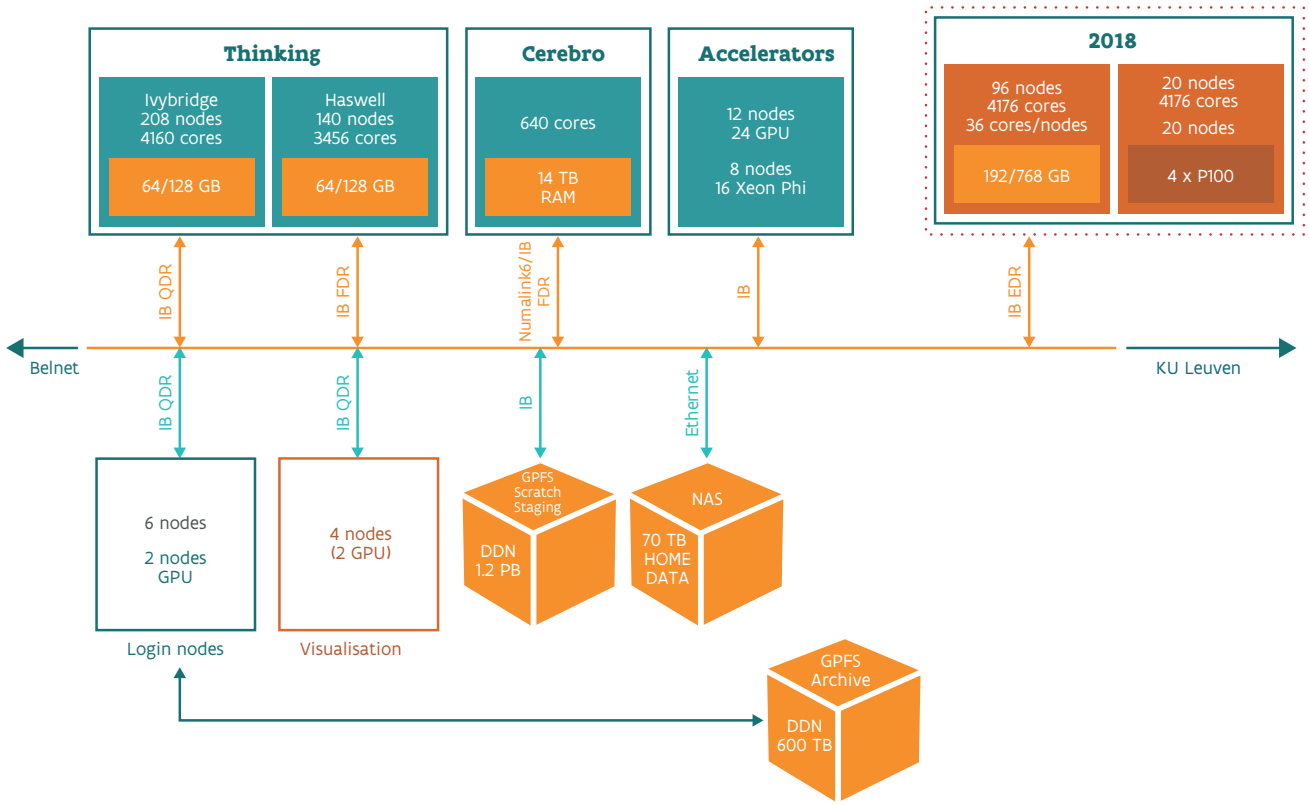

Figure 6 - Tier-2 infrastructure KU Leuven - UHasselt

In 2017, no material adjustments were made to the Tier-2 infrastructure. However, a purchase procedure was conducted for a completely new Tier-2 cluster. HPE submitted the most advantageous tender. The most important features of the new cluster are:

- 307 TFlops/s peak performance
- 10 nodes with more memory (768 GB)
- 20 nodes with 4 P100 nVivia GPUs each
- EDR interconnect
- Integration with the existing parallel file system

The new machine was ordered in October 2017 and will be commissioned in 2018.

GHENT UNIVERSITY

UGent has been investing for several years in the development of a powerful infrastructure. Today it consists of:

- 5 clusters
- 218 TF
- 10,960 CPU cores
- 49 TB memory

The Tier-2 infrastructure is built up of various clusters, serving specific characteristics. A new shared DATA and SCRATCH storage (for 2 PB) was taken into production in the course of 2017. In addition, an order was placed to expand the Tier-2 infrastructure with two new clusters, which will come into production in 2018.

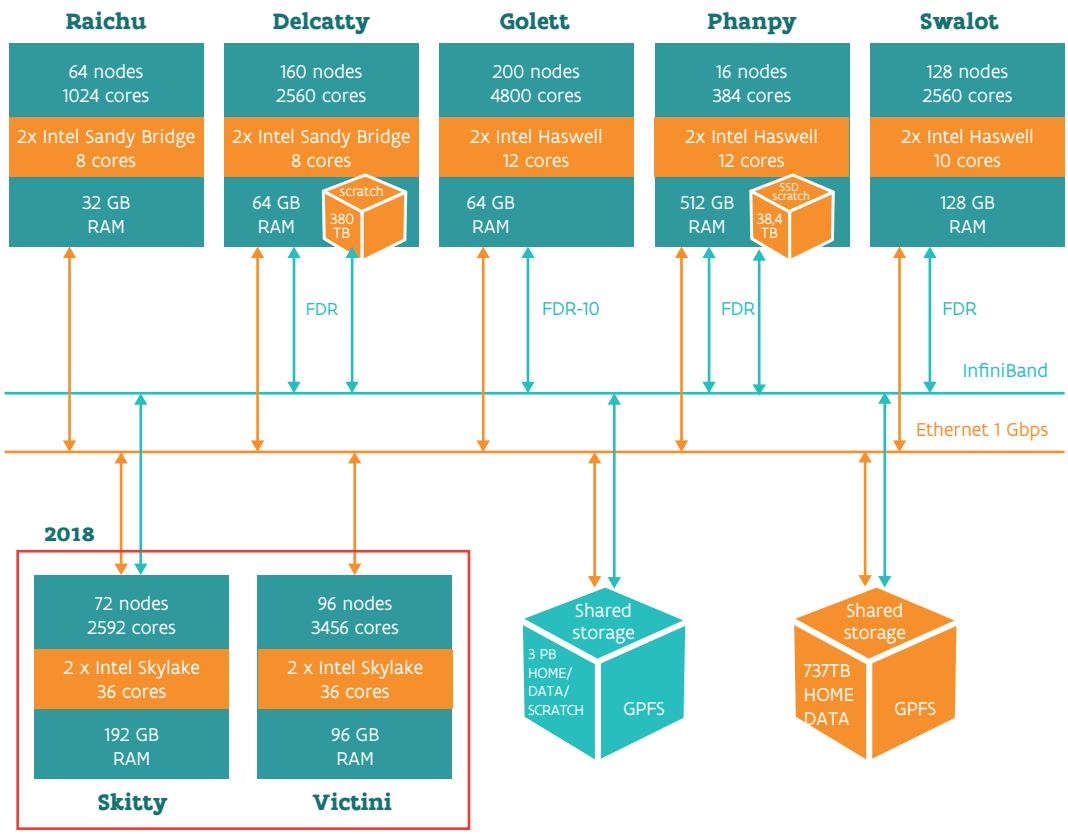

Figure 7 - Infrastructure UGent

UNIVERSITY OF ANTWERP

For the University of Antwerp, large processing capacity for research is a strategic priority. The Tier-2 infrastructure consists of:

- 2 clusters (Hopper and Leibniz), divided into 4 partitions
- 239 TF
- 7616 CPU cores
- 36 TB memory.

Leibniz, the latest cluster, was commissioned in the course of 2017.

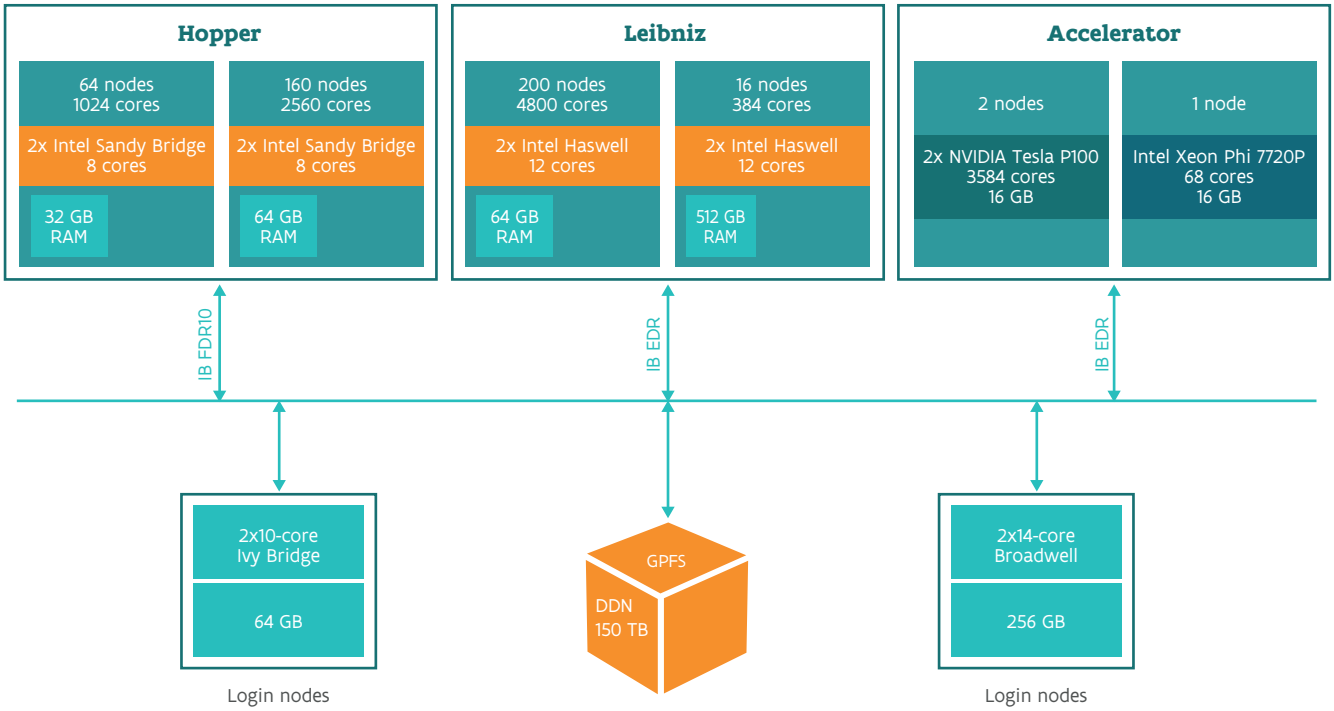

Figure 8 - Tier-2 infrastructure UAntwerp

VRIJE UNIVERSITEIT BRUSSEL (VUB)

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

- 1 cluster / 6 partitions
- 16 TF
- 2848 CPU cores / 32256 GPGPU cores
- 21.6 TB memory

The VUB decided to implement all expansions within the same Hydra environment, which is more efficient for both the users and the management team. This results in a more heterogeneous cluster that meets the specific needs of different research groups.

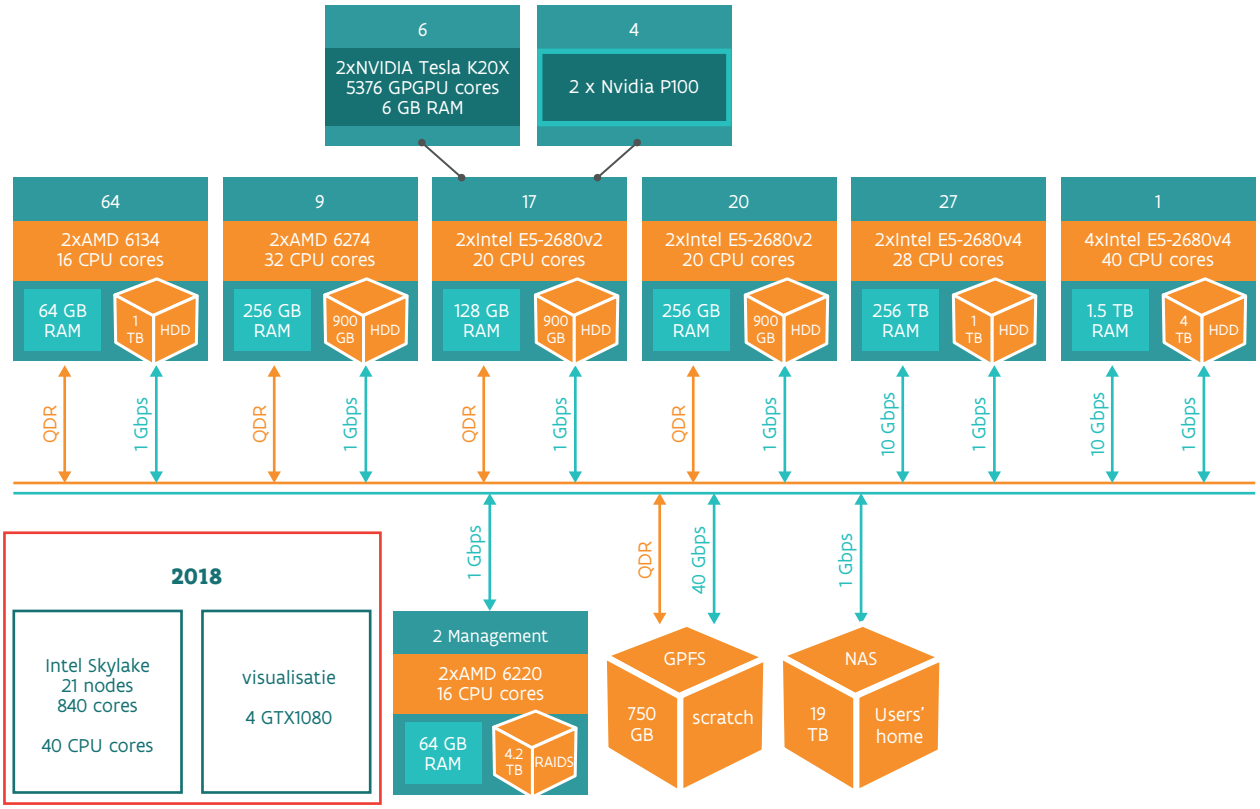

Figure 9 - TIER-2 infrastructure VUB

In addition to its own Tier-2 infrastructure, the VUB - together with the ULB - manages the grid infrastructure, which is used, among other things, for processing data that are collected during experiments with the Large Hadron Collider (HPC) at the CERN, but also within the Flemish research community.

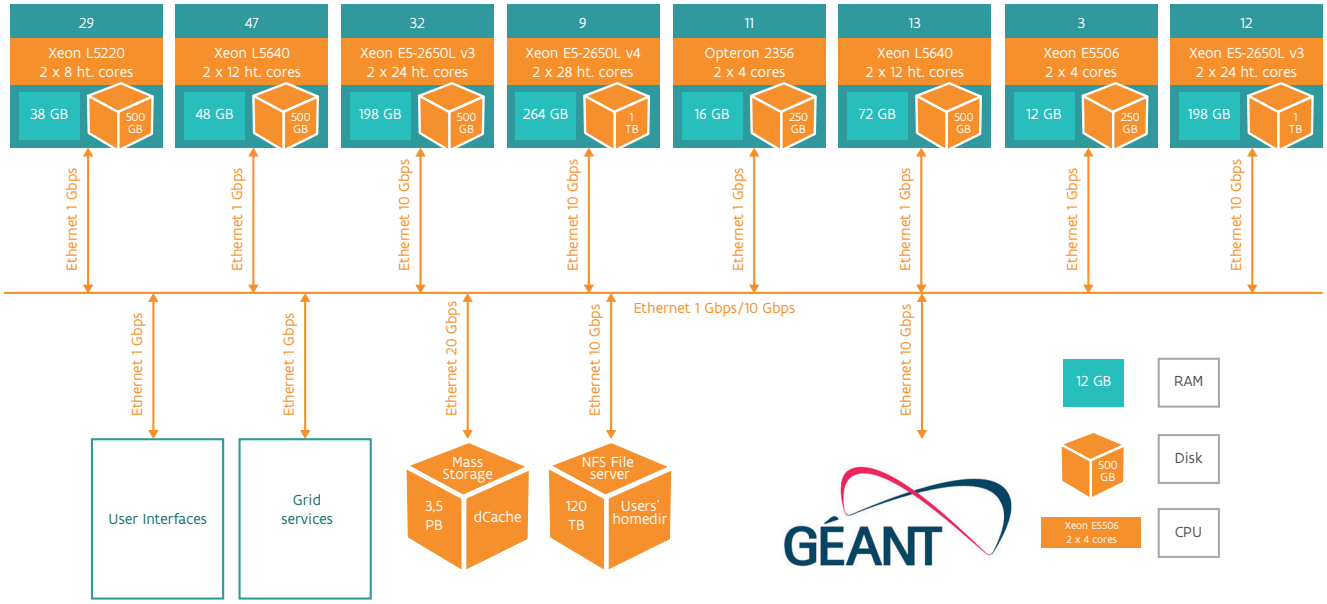

Figure 10 - Grid infrastructure VUB

Finally, the VUB has its own test setup for cloud infrastructure.

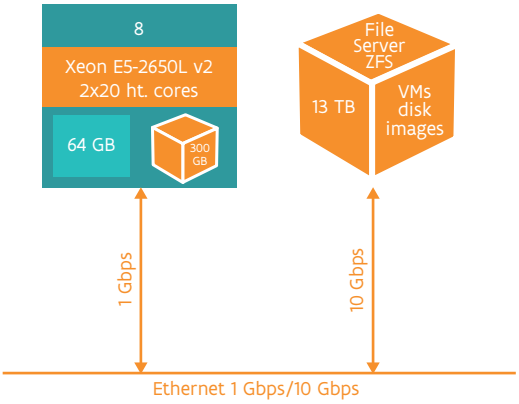

Figure 11 - VUB cloud infrastructure

Operation and use

Having described the infrastructure, this section presents an overview of the use of the Tier-1 and the Tier-2.

For monitoring the use, the VSC has a central XDMoD infrastructure that collects all data from the various clusters and generates the necessary overviews.

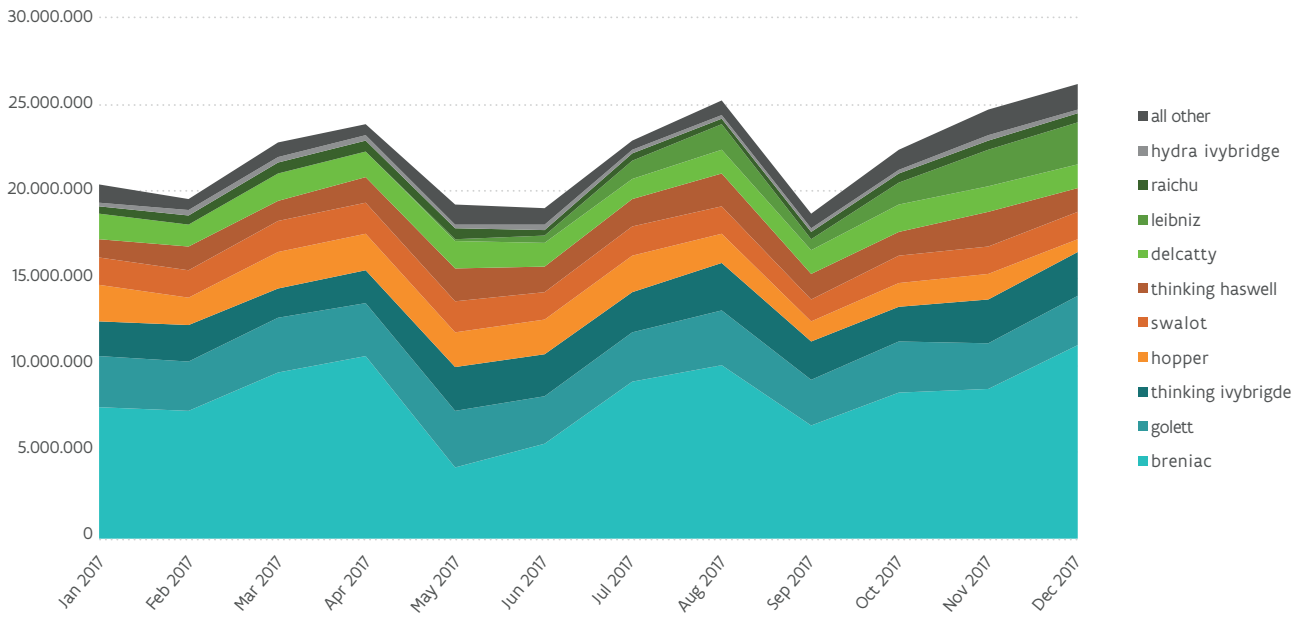

Figure 12 - Use expressed in core hours of all VSC clusters

Figure 12 shows the cumulative use of the Tier-1 and Tier-2 infrastructure. A significant increase is visible. The total number of core hours continued to rise in 2017 following the deployment of the new Tier-1 BrENIAC and the renewal of the Tier-2 infrastructure at the various institutions.

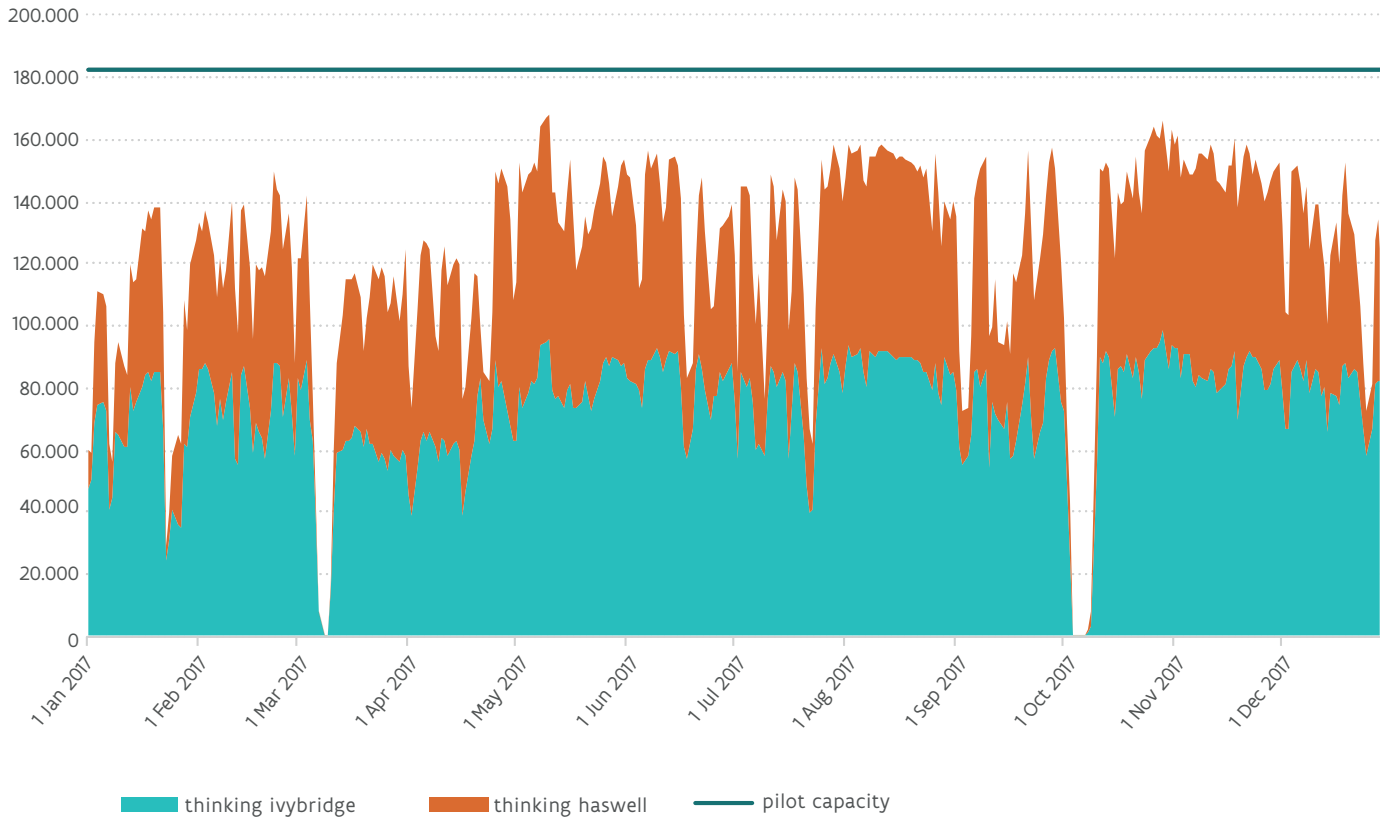

Figure 13 - Core hours used on the KU Leuven/U Hasselt thin node cluster

Figure 13 gives an overview of the use of the KU Leuven infrastructure. The graph shows the cumulative use of the various thin node clusters at KU Leuven. The graph shows 2 breaks. Both are related to works on the parallel file system. This system is intensively used in most compute jobs, as a result of which the cluster cannot be kept operational during such maintenance. Additional storage had been purchased at the end of 2016, which was installed in March 2017. The aggregated bandwidth was increased and the capacity expanded to 1.1PB. This makes the storage system ready to support a second cluster in the future. In October, the storage system's operating parameters had to be updated.

For the rest of the year, the cluster had a high occupancy rate.

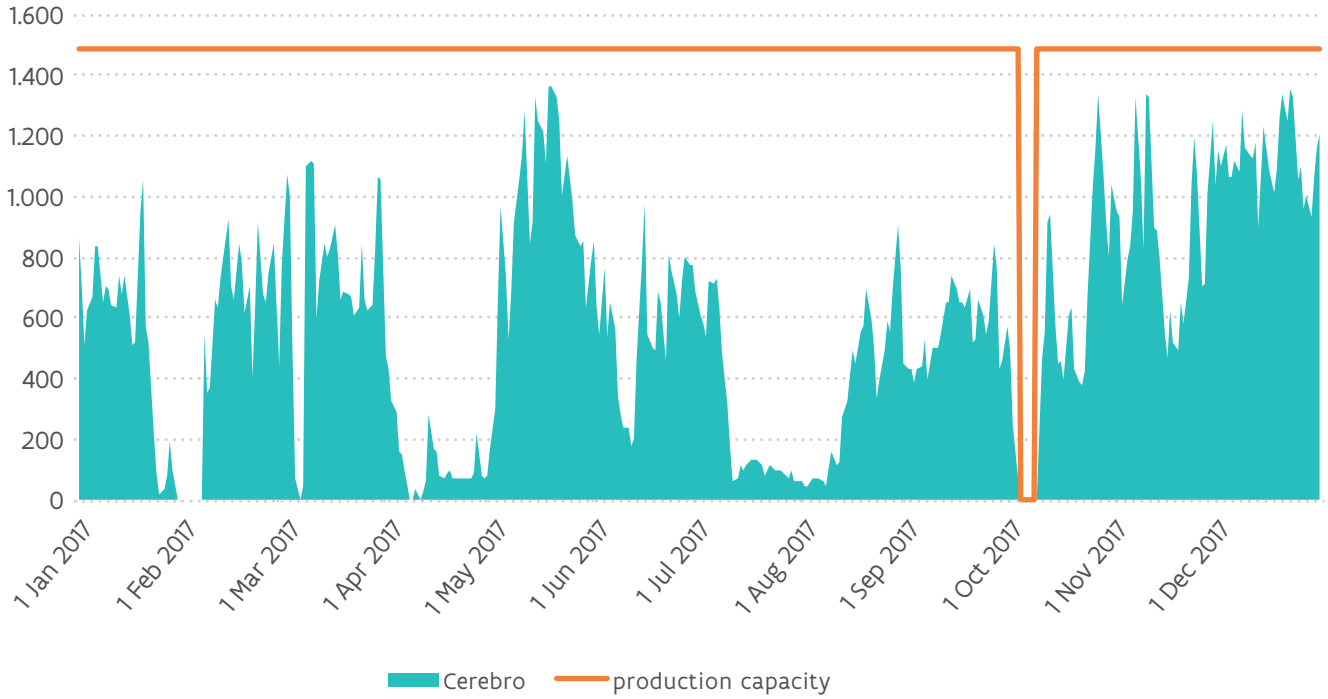

Figure 14 - Use of shared memory

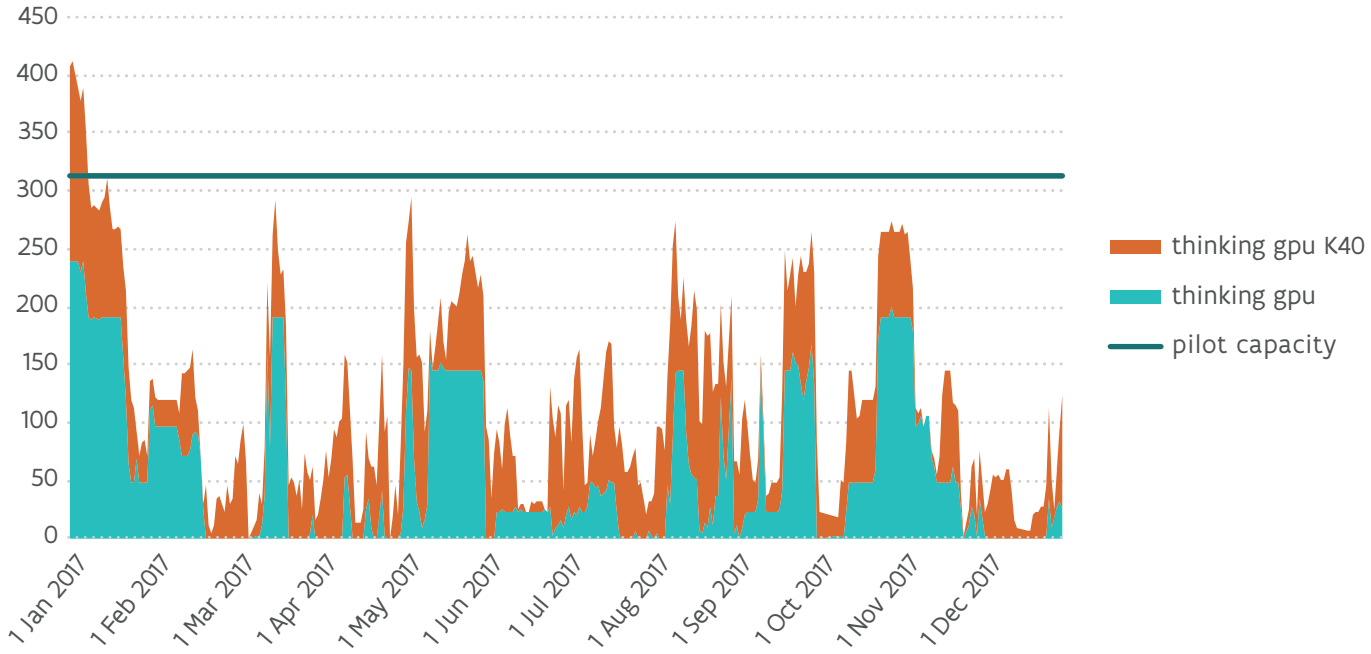

Figure 15 - Use of accelerators

In addition to the thin node processing cluster, KU Leuven also has a shared memory machine and a configuration with accelerators. The use of these machines is expressed in node hours and not in core hours. The reason for this is that the complete node is often used because of the required memory or the accelerators. Cerebro recorded typical usage peaks. The major users of this machine are chemists and bioinformatics. Both have a number of softwares that require more memory than is available on the standard compute nodes.

The GPU nodes were used in development for data analyse codes, on the occasion of a hackaton organised by astronomers, and several training programmes. In 2017, users who might benefit from GPUs continued to compute CPUs because the total capacity of the available CPUs is much higher than that of the GPUs. As demand grows for more machine learning applications, GPU usage may also be expected to increase.

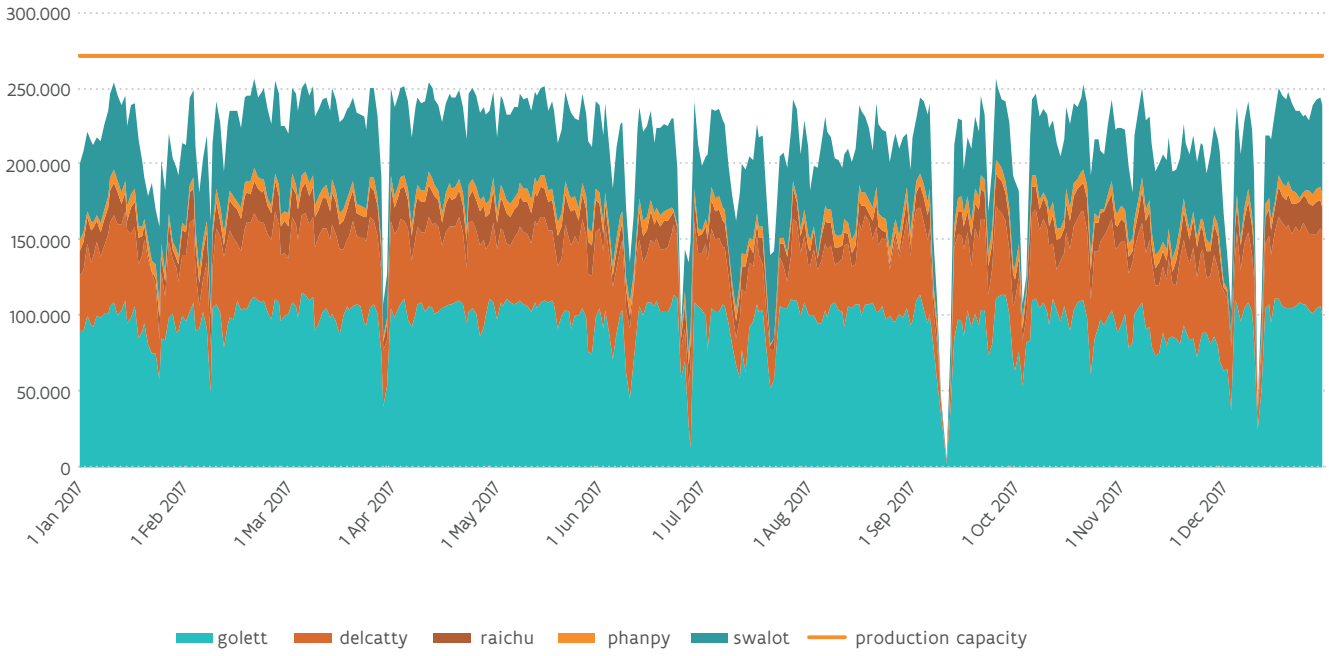

Figure 16 - Use of the infrastructure

Figure 16 shows the cumulative use of all computing clusters the Ghent University in 2017.

In 2017, a total of 73,656,948 core hours were consumed on the Tier-2 computing clusters of the Ghent University. This would correspond to 8,408 years of computing on a single core.

The average effective usage percentage for all clusters amounted to 77% in 2017. This percentage indicates how much of the theoretically available computing power was used in one year (excluding downtimes, meaning that the actual value is higher). The effective usage percentage is rather high as compared to typical HPC systems and indicative of good economical use of the infrastructure.

Dips in the above graph indicate downtimes resulting from power failures or scheduled maintenance work.

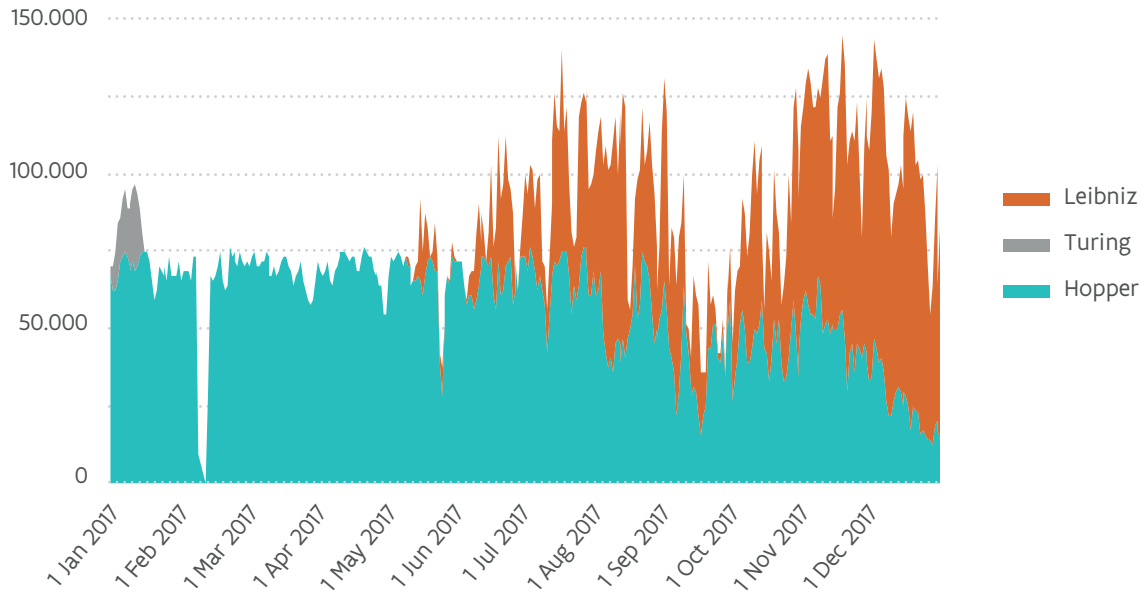

Figure 17 - Use of infrastructure UAntwerp

The graph of Figure 17 shows the cumulative use of the Turing cluster (decommissioned at the end of January 2017) and Hopper and Leibniz clusters at the University of Antwerp. The graph shows a number of breaks.

In February, the storage system was upgraded to prepare the installation of Leibniz. As a result, Hopper was temporarily unavailable. In the second quarter, Leibniz became available for pilot use. In the last quarter, the Hopper nodes were gradually switched to the same operating system as used for Leibniz in order to simplify maintenance. Although the "renewed" nodes were immediately available for the users, this is not visible in the graph.

Aside from these interruptions, the clusters are very well used. As a result of the reporting method (related to so-called hyperthreading), a utilisation of more than 100% is sometimes shown.

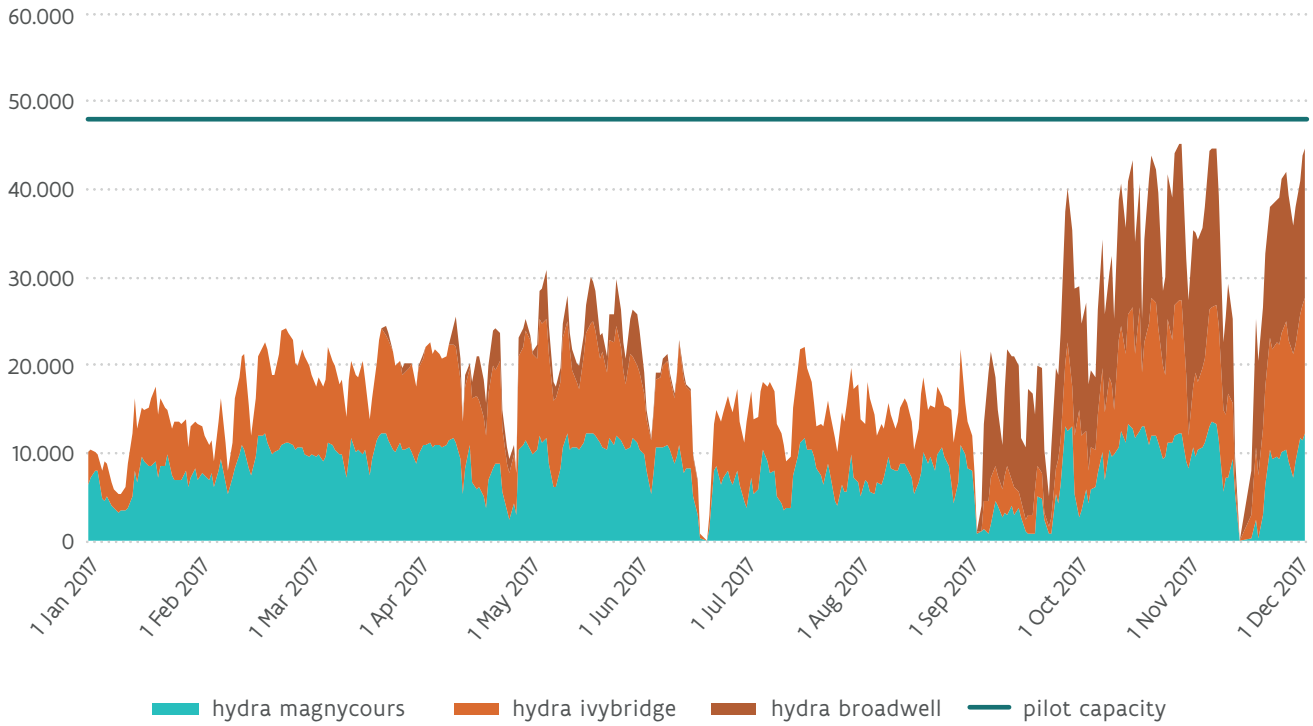

Figure 18 - Use of the VUB infrastructure

In 2017, the new Intel Broadwell nodes were officially launched following the resolution of the network problems. This involved one hour of downtime. In addition, there were 2 scheduled maintenance windows in July and December, one including the complete replacement of the network within Hydra by a more powerful version. There was also a short unscheduled downtime due to a cooling failure in the month of May.

The Hydra cluster runs both jobs that use all processing cores in a node and jobs that do not use all cores but require the complete node because of memory requirements. Furthermore, a number of nodes were reserved for interactive work by a number of research groups. Such a mix produces a distorted picture of the utilisation level by core hours or node days, which is the standard reporting method for the other universities, which is why the total number of available core hours is omitted in the graph.

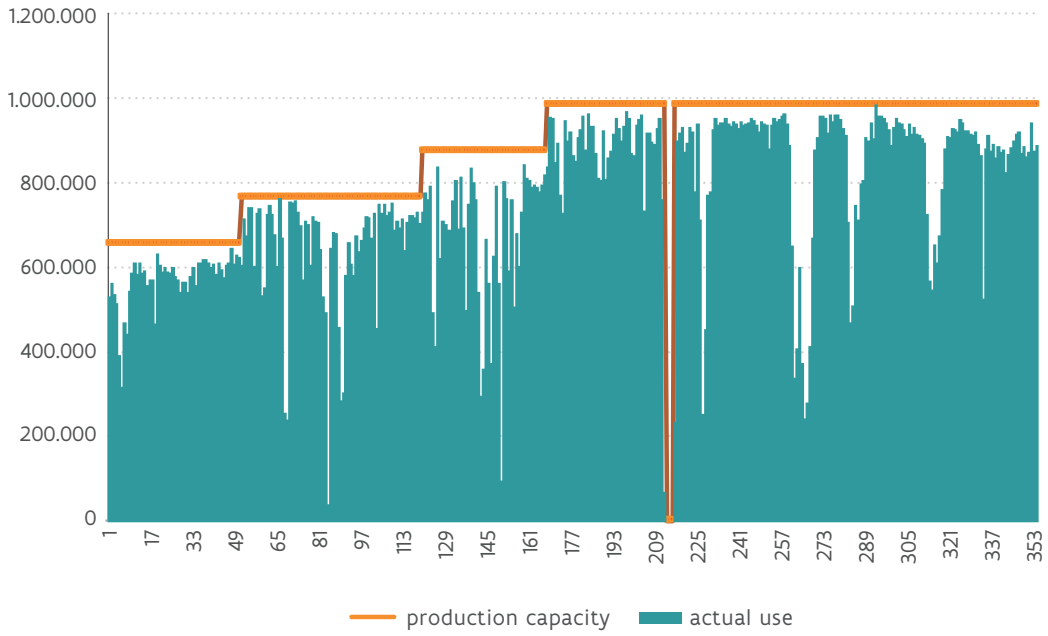

Figure 19 - Use of BEGrid cluster at VUB (blue line: production capacity)

BEgrid 20-27 and BEgrid 47-50 are the constituent parts of the BEgrid cluster at the VUB. . Additional worker nodes were added in several phases in the course of 2017. In August, there was a planned downtime for routine maintenance work. There were also a number of unplanned downtimes, notably due to problems with the file server (in July) and cooling issues (in May).

Allocation of Tier-2 computing time

Each university has its own procedure for the allocation of computing time on the Tier-2 infrastructure and may or may not charge a small portion of the costs to the academic researcher. For industrial/external users, all consumed computing time is always charged in full. To gain access to one of the Tier-2 clusters in the four VSC hubs (Antwerp, Brussels, Ghent, Leuven), the user must have a VSC userid, which can be requested at https://account.vscentrum.be. This website and the database also centralises all user information across the institutions, such as storages quota, membership of user groups, virtual organisations, etc.

Researchers of the University of Antwerp and its association have full free access to the Tier-2 infrastructure. Research groups can, however, make a financial contribution on a voluntary basis.

Researchers at the VUB can work on the HYDRA cluster after they have been granted access by the computer centre. The grid cluster is available on demand from the infrastructure administrator. Use of the Tier-2 infrastructure is free of charge.

Researchers of the Ghent University and its association have full free access to the local Tier-2 infrastructure. Research groups can, however, make a financial contribution on a voluntary basis, with a (slightly) higher fair share as direct return-on-invest.

The clusters of KU Leuven/University of Hasselt use a credit accounting system that is incorporated in the scheduling software. New users receive computing time in order to become accustomed to the system and to carry out first tests. In this way, the barrier for researchers to switch to the Tier-2 infrastructure is kept as low as possible. Subsequently, credits can be applied for via a simple procedure and at minimum cost. The credits distribute the available computing time over various projects and have an empowering effect. When a computing job is carried out, the project to which the credits are to be charged, is specified. The principal researcher is the project manager. He/she can grant researchers access to the computing time and also monitor the used computing time.

The use of central accounts allows users to compute also at other sites, taking into account the conditions that apply at those sites. A brief summary of cross-site computing is given below.

Cross-site computing

Site	Overview	Compute time in core hours					
		UAntwerp	VUB	UGent	KU Leuven/ UHasselt	Other research institutes	Industry
@KU Leuven/ UHasselt	52 users/665 node days	88.346	0	165.835	47.678.343	1.149.130	431.070
@UGent	21 users/47 core years	190.879	10.651	7.900.2512	204.544	6.147	0
@UAntwerp	0 users/0 node days						
@VUB	6 users/886 jobs/2826 days	140.787		0	0	0	0

This so-called cross-site use is closely monitored.

In addition to the cross-site use of the Tier-2 infrastructure, the grid infrastructure, managed by VUB/ULB, is also intensively used by researchers from different institutions: VUB, UAntwerp and UGent.

The remaining computing time on the grid infrastructure is used by ULB and UCL researchers.

User support

User support comprises several components:
- answering questions from users (helpdesk);
- meetings with users / specific support;
- training and outreach. This last item is discussed in the sections "Training" and "Outreach to Flemish companies".

ANSWERING QUESTIONS FROM USERS

An overview of the tickets that reach the helpdesk is given below. There is no central VSC helpdesk. Each institution answers all questions and queries from its own users (i.e. users that have not requested an account at the relevant institution), both regarding the own Tier-2 infrastructure and the central Tier-1, but also from external users using the VSC infrastructure. Where necessary, for questions relating to the Tier-1, the helpdesk of the Ghent University (for the first Tier-1) or that of KU Leuven (for the second Tier-1) is contacted. With respect to questions, a distinction is made between:

- questions about accounts;
- questions about software;
- other questions.

The table below provides an overview of the number of tickets handled, by category and by institution.

Tier-2 + grid	KU Leuven UHasselt	UGent	UAntwerp	VUB
Accounts	1040	191	35	180
Software	168	180	96	157
Other	604	535	257	610

Tier-1	KU Leuven UHasselt	UGent	UAntwerp	VUB
Accounts	12	0		
Software	21	0		
Other	183	24		

MEETINGS WITH USERS / SPECIFIC SUPPORT

On the one hand, we try to involve as many researchers as possible in the HPC story by examining whether and how they can make the switch from their desktop to the HPC infrastructure or how they can use their own desktop more efficiently. For existing users this can also involve helping them make the switch from Tier-2 to Tier-1 and possibly Tier-0. On the other hand, we try to provide specific support to researchers.

Some examples:
- optimising existing work flows;
- analysing/optimising code;
- providing input to the writing of research projects.

In addition, user meetings take place at each institution in which a delegation of the users is represented. Here we make - per institution - a selection of the above-mentioned support.

KU Leuven and University of Hasselt
Daily support deals with questions regarding accounts, basic use of the cluster, and software installations. These questions come from users from groups who have been using the cluster for quite some time. Making the VCS infrastructure known to new research groups is an ongoing effort. However, also within groups that already use the cluster, specific actions may help to promote cluster use. In this way, more computational work can be performed more efficiently and in a shorter period as compared to a local infrastructure such as workstations and desktops. New researchers are familiarised with the use of the cluster in the regular introduction sessions or through a one-on-one consultation. The latter are very important in helping new researchers get off to a successful start. After a two-hour consultation, the researcher will in most cases have gathered enough information to be productive on the cluster, provided he/she has some prior Linux and HPC knowledge. If a new group wants to start on the cluster, specific workshops are organised.

In 2017, a number of specific support projects were continued.

For the Genomics Core, the whole genome pipelines were implemented, including new versions of GATK and Halvade. The Genomics Core researchers, together with HPC support, organised additional info sessions for their researchers.

For the Nucleomics Core, gene pattern was set up. This is a web application that allows specific analyses to be carried out for the analysis of gene expressions, including sequence variations, copy number, and proteomics analysis. Analyses that do not require much computing power can be performed on the server, whereas more intensive analyses are computed by the HPC cluster.

For the Computational Biology laboratory, support was provided to allow the power of the HPC to be used and to enable more interactive work on the same data. Specific workflows consisting of interactive work with short CPU loads involving only a limited capacity of a node, can be run on shared nodes. It is meaningful to perform this work also within the HPC environment since both the software and the data are available within this environment.

The Plasma Astrophysics research group is involved in the Virtual Space Weather Modelling Centre (VSWMC). Assistance was provided with the installation and configuration of the software and with the installation of new prediction models (Euphoria) used within this framework.

A new group was set up within the Department of Soil and Water Management. This group uses the software framework supported by NASA: Land Information System (LIS) and Land Data Assimilation Systems (LDAS) and also additional GIS software for visualizations. This was set up on Tier-2 and the new researchers were familiarized with the cluster.

For the Clinical and Epidemiological Virology laboratory, an assembly pipeline was developed, which combines different modules for research on viruses. The use of HPC provides researchers with access to many more computational cores with sufficient RAM, which may significantly shorten the lead time of their research. Additionally, for some modules they can simultaneously use different computational nodes via Hadoop.

The Department of Astronomy and the Department of Geography received specific support for the preparation of Tier-1 project applications.

Furthermore, specific sessions were organised in support of Astrohack, a hackaton organised by Data Science Ghent at the Ghent University with the GPUs at KU Leuven. Throughout the hackaton support was provided for participants who wanted to use the HPC.

For the Von Karmann Institute (VKI), a training programme on introduction to the use of the HPC cluster was developed. The cluster can be used to assess performance on new HPC architectures.

For the Department of Electrical Engineering, preparatory sessions were organised. This group uses and manages a heterogeneous set of different GPU machines. With the new Tier-2 which was purchased in 2017, considerable GPU capacity is now available within the VSC.

In 2017, KU Leuven organised the NEC User Group meeting; the KU Leuven Tier-1 is a NEC machine. The NUG brings together other NEC users to share experiences and to provide feedback to the supplier, also for new developments. The meeting was attended by 40 or so participants from various European countries and Japan.

For the University of Hasselt, the emphasis was on the attraction of new user groups. The Faculty of Engineering Science received the necessary attention, resulting in an increase in the number of researchers using the VSC infrastructure, both for research and for education. The Centre for Environmental Sciences (CMK) also received extra support.

Ghent University
Throughout the year various user meetings were organised in response to specific questions from researchers.

To convince as many (potential) users as possible of the importance and added value of supercomputing, the VSC was presented internally and several guided tours were organised in the UGent data center:

- 6 February 2017, Information session FWO call 2017 for Research Grants, Research Projects and SBO
- 9 May 2017, general presentation and tour for students of the Computer Systems study programme (UGent)
- 8 November 2017, tour for students of the Faculty of Economics and Business Administration (UGent)
- 5 December 2017, general presentation and tour for students Bachelor of Science in Engineering Technology (UGent)

Also in 2017, the first OpenFoam users conference was held (13/09/2017). The purpose of this user meeting was to bring together different research groups using OpenFoam, the popular Open Source CFD software, on the VSC infrastructure. Participants presented their work and shared best practices for the use of this software on HPC clusters. Participating research groups:

- Bekaert
- Department of Civil Engineering, UGent and KU Leuven
- Laboratory for Chemical Technology, UGent
- Department of Flow, Heat and Combustion Mechanics, UGent

University of Antwerp
On the one hand, we deal with questions from existing users and try to organise their computing work as optimally and efficiently as possible by working proactively. A properly maintained software stack and supporting tools are vital in this respect. On the other hand, we inform other researchers of the existence of the VSC to highlight the potential benefits of using the central infrastructure. To this end, we target specific researchers/research groups.

Every year we also organise two intro sessions, which, since 2016, consist of 3 parts: "Linux introduction", "Supercomputers for starters" and "HPC introduction". It is necessary not only to be able to work with the environment, but also to have access to the necessary background knowledge. In addition, due attention is paid to keeping the documentation on the VSC website up-to-date.

In 2017, we put researchers from the Institute of Tropical Medicine and VITO on their way towards HPC. In addition, other research groups from departments already active on the infrastructure, such as Biology, Bioengineering Science, Physics, Medicine and Health Sciences, Mathematics and Applied Engineering Science, are now regular users.

Some examples of specific support:

- Lunch seminar on supercomputing for the Biomina group.
- Supervision of a Master thesis: "From sequential MATLAB to parallel Julia: a toolbox study".
- Support with applications for Tier-1 computing time.
- Support with (inter-university) project applications.

In addition to training within the VSC, courses are also organised under the regular programme: "Scientific computing environments" and "(Parallel) programming".

The University of Antwerp has an active user group which was formed in 2006 and meets twice a year. The user group consists of representatives from 12 groups and disciplines.

Vrije Universiteit Brussel (VUB)
In addition to the monitoring of existing users at Tier-2 and Tier-1 level, we focused on the active identification of new potential Tier-1 users, first of all encouraging them to apply for starting grant for computing time. For the Tier-2 level, the focus remained on researchers mainly from human sciences for whom the use of Tier-2 infrastructure would constitute a great step. Because of the specific, Window-based software, these researchers were offered a solution within the cloud environment. The VUB has a HPC users committee that meets every two months and provides feedback from the VSC, while monitoring the needs for HPC. The users committee includes members from all faculties, from the computing centre and from the research policy department. The following courses are organised twice a year: "Introduction to Linux" and "Introduction to the use of HPC at the VUB". In addition, user meetings were organised with a number of groups to identify their specific needs and/or issues with HPC.

- Department of General Chemistry
- Department of Mechanical Engineering, Combustion and Robust optimization
- VIB-Center for Structural Biology
- Department of Electronics and Informatics
- BRIGHTcore
- Interuniversity Institute of Bioinformatics

Staff

Since the infrastructure of the VSC (Tier-2 and Tier-1 infrastructure) is installed in the various university data centers, the staff are also employed at the various universities.

FUNDING

On the one hand, each university needs staff for the operation of the Tier-2 infrastructure and the support of the end-users. For this, 20 FTE are subsidised. On the other hand, funding goes to the institution that houses Tier-1 supercomputer infrastructure. To this end, 2 FTE were allocated for the operation of Tier-1a and 2 FTE for Tier-1b, in 2017.

Since each university has its own employee policy and applies different remuneration principles, each FTE is paid a fixed amount of € 95,000.

Institution	Number of subsidised FTE for Tier-2 operation and support
UGent	5
UAntwerp	4
VUB	3
UHasselt	2
KU Leuven	6

Institution	Number of subsidised FTEs for Tier-1 operation
UGent	2
KU Leuven	2

EFFECTIVE STAFF DEPLOYMENT

Operating, maintaining and supporting the users of the various Tier-2 and Tier-1 configurations requires on the one hand more manpower than is provided for within the funding. On the other hand, a broad range of expertise is required that cannot be accumulated within one limited team. To address this, HPC technicians and support staff can make use of other experts who work in the ICT departments of the various universities. The universities deploy together 31.9 FTE for HPC operation and support. We should point out here that, because of an institution-wide regulation, the FTE who are deployed at the University of Antwerp and are not financed by the FWO, can only be included for a prescribed and maximum fraction of the deployment in this table.

Institution	Number of FTEs deployed for HPC	Number of heads involved with HPC operation and support
UGent	9,5	15
UAntwerp	4,35	8
VUB	4,31	11
UHasselt	1,16	3
KU Leuven	12,6	19
Totaal	31,91	56

PROFILES

In order to operate and support a HPC infrastructure effectively, various ICT profiles are required. In broad lines, these profiles are as follows:

Infrastructure administrators
These persons are responsible for integrating the HPC infrastructure in the data center. They install the infrastructure in the data center. They also install and manage the specific storage that is coupled to the HPC. In addition, they install and manage the internal network of the HPC and they connect the HPC into the university network and the Intranet. They are responsible for the security of the infrastructure and the daily monitoring of it. They participate in purchasing decisions on the integration of the infrastructure in the data center.

System administrators
These persons are responsible for installing and administering the basic software on the HPC machines. This extends from the operating system to the scheduling software. They are responsible for the daily monitoring of HPC systems. They develop software for the efficient management of the HPC systems. They participate in purchasing decisions on the architecture of the HPC machine and the management software.

User support staff
These persons are responsible for basic user support. They man the first-line helpdesk and help users on their first acquaintance with the machine. They install the user software and help users efficiently use the machine. They are responsible for the documentation and provide basic training.

Academic or advanced support staff
These persons are responsible for optimisation of the user software on the HPC machines. For this they frequently work for a longer period with one user and give advanced training. These persons are responsible for adjusting the scheduler software so that it answers the needs of the users. They are responsible for setting up actions aimed at raising awareness of the added value of HPC and attracting new users. They participate in the purchasing decision on the user requirements and the benchmarks.

Project managers
These persons have the immediate management over the HPC teams or manage larger HPC projects. They are responsible for embedding the local HPC in the environment of the VSC. They are responsible for the coordination of HPC initiatives and user groups within their own institution. They are responsible for reporting to the funding providers. They have the final responsibility for purchasing files.

Outreach
Many persons involved in the VSC spend a portion of their time on promoting scientific computing and HPC. They do this by giving info sessions within the academic context but also by visiting the companies. They arrange for press coverage or promote scientific computing and VSC at conferences and through other activities. In addition, they maintain the website and provide material that can be used at the various events.

The number of persons involved in these various activities is given below.

	Infrastructure administration	System administration	Basic user support	Academic or advanced support	Management	Outreach
UGent	5	7	6	5	3	2
UAntwerp	1	3	5	3	2	1
VUB	3	4	5	4	3	3
UHasselt	0	0	2	2	1	2
KU Leuven	5	6	6	6	5	7
Total	14	20	24	20	14	15

ADVANCED SUPPORT

In order to provide advanced and academic support to the end-user, expertise in the domain is generally a plus. It is, however, impossible to accumulate domain expertise within each institution for a broad range of specific areas. That is why it is good to retain an overview of the various expertises that are present within the VSC so that users can call on advanced expertise across the institutions should they need it.

Computational Chemistry
4 persons with academic experience in this field

Physics
6 persons with academic experience in this field

Engineering
2 persons with academic experience in this field

Mathematics (Numerical Methods)
4 persons with academic experience in this field

Computer science
9 persons with academic experience in this field

Bioinformatics
4 persons with academic experience in this field

STAFF LIST

Name		Percentage active within the framework of HPC
Wouter Depypere	UGent	100%
Stijn De Weirdt	UGent	100%
Alvaro Simon Garcia	UGent	100%
Andy Georges	UGent	100%
Kenneth Hoste	UGent	100%
Ewald Pauwels	UGent	100%
Jens Timmerman	UGent	100%
Kenneth Waegeman	UGent	100%
Danny Schellemans	UGent	10%
Johan Van Camp	UGent	30%
Frédéric De Leersnijder	UGent	20%
Wim Waeyaert	UGent	20%
Werend Brantegem	UGent	10%
Bruno Cardon	UGent	10%
Dieter Roefs	UGent	50%
Balázs Hagató	VUB	100%

Samuel Moors	VUB	50%
Stéphane Gérard	VUB	100%
Ward Poelmans	VUB	80%
Dirk Heyvaert	VUB	15%
Philippe Leemans	VUB	20%
Eddy Haulet	VUB	10%
Olivier Devroede	VUB	25%
Johan D'Hondt	VUB	15%
Steven Opstaele	VUB	10%
Stefan Weckx	VUB	5%
Stefan Becuwe	UAntwerp	100%
Franky Backeljauw	UAntwerp	100%
Bert Tijskens	UAntwerp	100%
Kurt Lust	UAntwerp	100%
Koen Decauwsemaecker	UAntwerp	10%
Muriel Dejonghe	UAntwerp	10%
Herwig Kersschot	UAntwerp	10%
Annie Cuyt	UAntwerp	5%
Herman Moons	KU Leuven	20%
Leen Van Rentergem	KU Leuven	30%
Jan Ooghe	KU Leuven	100%
Martijn Oldenhof	KU Leuven	100%
Mag Selwa	KU Leuven	100%
Alexander Vapirev	KU Leuven	100%
Ingrid Barcena	KU Leuven	100%
Ehsan Moravveji	KU Leuven	60%
Jo Vanvoorden	KU Leuven	100%
Jo Vandeginste	KU Leuven	100%
Tom Leuse	KU Leuven	100%
Yorick Poels	KU Leuven	100%
Peter Veraedt	KU Leuven	50%
Rudy Rys	KU Leuven	40%
Tom van Mierlo	KU Leuven	100%
Tom Vanhout	KU Leuven	20%
Sofie Pieraerd	KU Leuven	10%
Els Veraverbeke	KU Leuven	20%
Philip Brusten	KU Leuven	10%
Rafal Al-Takreeti	UHasselt	8%
Geert Jan Bex	UHasselt	100%
Alexander Jaust	UHasselt	8%

OUTREACH TO FLEMISH COMPANIES

From the outset, the VSC's services and infrastructure have been available for use by researchers not directly involved in the VSC consortium. However, until this day, technical access modalities for users external to the university associations have not always been straightforward. The successful conclusion of the VSC project "Access for non-academic users" (2017) will put an end to that.

At the same time, an increased demand from outside the VSC consortium is noticeable, which is without any doubt linked to the increased visibility of the VSC in recent years. In the context of the implementation of the Tier-1 supercomputing platform, this demand may be expected to grow even further.

VSC services and infrastructure are currently being used from various institutions. This is not always limited to the consumption of computing time and associated standard user support (including training), but occasionally also involves a more comprehensive service, such as specific software optimisation.

Services to companies

The range of services that VSC supply to companies are built on four pillars: consultancy, research collaboration, training and computing power.

Consultancy
VSC experts analyse the specific needs of the company and examine how supercomputing can offer added value for this company. The VSC offers a free intake interview to ascertain which benefits HPC can offer the company and which services in the VSC network best satisfy these needs.

Research collaboration
The VSC acts as HPC contact and can channel the question from a company for research collaboration (e.g. in the context of an R&D project).

Within the network, the VSC can bring the company into contact with a suitable top-level research partner in the Flemish academic landscape and act as intermediary.

Training
The VSC regularly organises basic training events that are also open to industrial users. Topics include Linux, (parallel) programming languages and paradigms, code optimisation, but also application-driven training such as materials science, computational fluid dynamics etc. Customised end-user training can also be provided.

Computing power
Companies can, for internationally competitive prices, purchase computing time on the state-of-the-art supercomputing infrastructure within the VSC network. Users receive step-by-step support for starting computational tasks and a special helpdesk is ready to solve user-related questions and to assist the users in the installation of specific software.

Use of Tier-1 by companies

In 2017, three Flanders-based companies made use of the Tier-1 in Leuven. The user agreements were contractually concluded between FWO, KU Leuven and the company involved. Two other companies used the Tier-1 in an explorative context. For one of them, this was a test to evaluate the migration from the old to the new Tier-1. The company subsequently decided to sign a user agreement. The companies are active in the sectors of renewable energy, agricultural technology and material technology.

The figure below provides an (anomised) view of the use of computing time on Tier-1 BrENIAC by the industry in 2017.

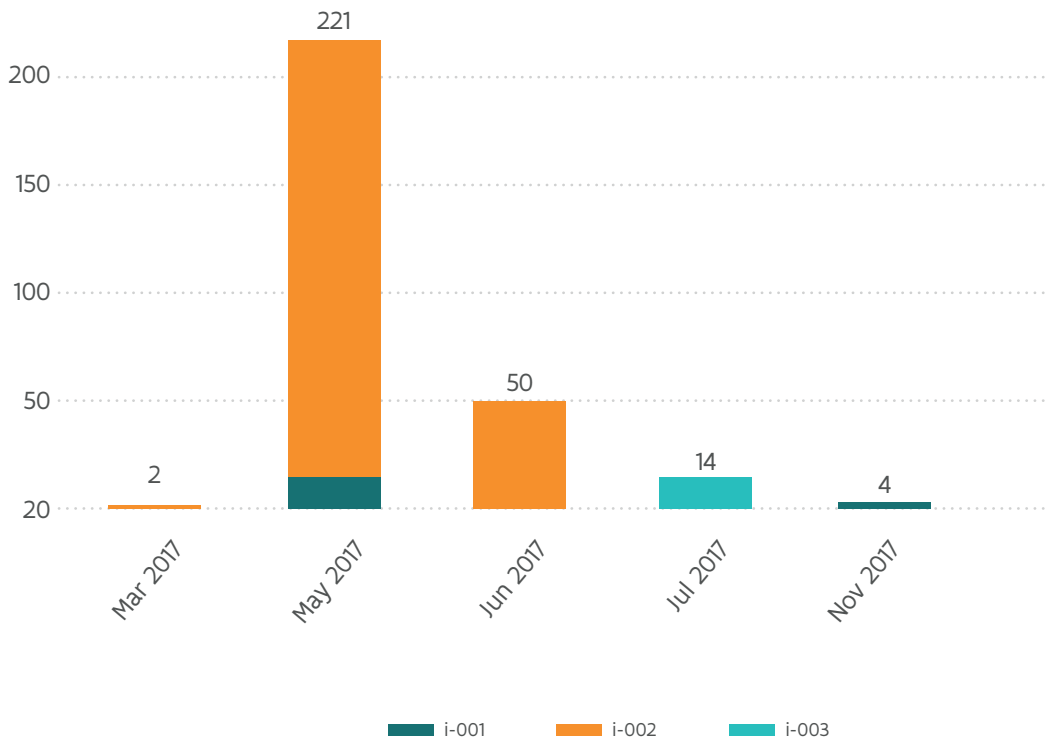

Figure 20 - Tier-1 BrENIAC usage by industry in 2017

Use of Tier-2 by external users

Not only researchers of institutions of the VSC consortium use the Tier-2 infrastructure. The different categories of VSC users are listed below, each time indicating which institutions are currently already active (or were active in 2017), via which VSC partner, and the kind of infrastructure or services used.

STRATEGIC RESEARCH CENTRES

Institution involved	VSC partner	Services provided by VSC
VITO	UHasselt KU Leuven UAntwerp	Tier-2 computing time Software optimisation and parallelisation Tier-2 computing time
VIB	UGent KU Leuven UAntwerp	Tier-2 computing time Tier-2 computing time Tier-2 computing time
imec	UGent KU Leuven UAntwerp	Tier-2 computing time Tier-2 computing time Tier-2 computing time
Flanders Make	KU Leuven	Tier-2 computing time

OTHER RECOGNISED FLEMISH SCIENTIFIC INSTITUTIONS

Institution involved	VSC partner	Services provided
ITG (Institute of Tropical Medicine)	UAntwerp	Tier-2 computing time
ILVO (Research Institute for Agriculture, Fisheries and Food)	UGent	Tier-2 computing time
INBO (Research Institute Nature and Forest	UGent	Tier-2 computing time

FEDERAL SCIENTIFIC INSTITUTIONS

Institution involved	VSC partner	Services provided
IRCEL (Belgian Interregional Environment Agency)	UHasselt & KU Leuven	Tier-2 computing time Software optimisation and parallelisation
RMI (Royal Meteorological Institute of Belgium)	VUB KU Leuven	Tier-2 computing time Tier-1 computing time (in partnership with UGent)
RBINS (Royal Belgian Institute of Natural Sciences)	UGent	Tier-2 computing time
RMCA (Royal Museum for Central Africa)	UGent	Tier-2 computing time
VKI (Von Karman Institute)	KU Leuven	Tier-2 computing time (explorative)

FLEMISH GOVERNMENTAL BODIES

Institution involved	VSC partner	Services provided
WL (Hydraulics Research)	KU Leuven	Tier-2 computing time
VMM (Flanders Environment Agency)	UGent	Tier-2 computing time (explorative)

FEDERAL GOVERNMENT BODIES

Institution involved	VSC partner	Services provided
Federal Police	KU Leuven	Tier-2 computing time

COMPANIES

Institution involved	VSC partner	Services provided
International Marine and Dredging Consultants (IMDC) n.v.	UAntwerp	Tier-2 computing time
Diabatix	KU Leuven	Tier-2 computing time
3E	UGent	Tier-2 computing time
Bekaert	UGent	Tier-2 computing time
Inbiose	UGent	Tier-2 computing time

VSC Industrial Board

The Industrial Board acts as communication channel between the VSC and Flemish industry. The VSC offers large-scale scientific/technical computing (HPC) to the Flemish research community and industry. The Industrial Board can strengthen the exchange of ideas and expertise between institutions of knowledge and industry. It develops initiatives to inform companies and non-profit organisations about the added value of HPC in the development and optimisation of services and products, and promotes the services provided by the VSC to companies, such as consultancy, research collaboration, training and computational power.

In 2017, the Industrial Board launched the Start to VSC project. A step-by-step plan to be implemented with new industrial users, was drawn up. It consists of an initial contact between the industrial partner and the VSC, backed up by an FAQ document for the users and a checklist for the IT staff. The first is used to develop a business use case. The second already provides an overview of the technical specifications, allowing the IT department to assess the security of the connection, data transfer and data storage. When the assessment of the business case is positive, a contract to be based on a standard contract, can be drafted. The use case is initially fully implemented, and the experience gained is used to organise training for a larger group of users within the company. This should perpetuate the use of HPC within the company. In support of the step-by-step plan a number of tools need to be developed: a standard contract, FAQ for users, a checklist for IT departments, a document including the major safety specifications, a list of software used, an overview of use cases. This Start to VSC project was already implemented partly in 2017 and will be further developed in 2018.

Publication to companies and other knowledge institutes

At the same time, the focus on the direct contacting of companies to boost their interest in supercomputing and the VSC services, continued unabated. Exchanges of best practices between the HPC centres of the SESAME Net consortium in fact showed that these smaller types of events, where personal contacts with industry are central, are better suited to convincing companies of the added value of supercomputing. Several user meetings and kick-start events were organised, mostly on-site.

VSC was also actively present at several networking events to promote supercomputing. Furthermore, several industry use cases were prepared and posted on the VSC website, and disseminated through the SESAME Net channel, amongst others.

Specific outreach events to industry and policy-makers:

- 24/01/2017, general HPC presentation and visit to UGent data center by MEPs of European Parliamentary Commission ITRE (Industry, Technology, Research and Energy), European Commission and representatives from European network and lobby associations in the field of R&D&I.
- 08/03/2017, presentation of VSC at board of directors of BiR&D (Belgian Industrial Research and Development association) – "Supercomputing for academics & industry - Flemish Supercomputer Centre (VSC)" – Vilvoorde (http://birdbelgium.com/)
- 27/04/2017, booth of VSC at TechBoost! 2017 "Innovations in Big Data Analytics" – Ghent (https://aig.UGent.be/event-2017/techboost-voor-ingenieurs)

CASE

COMPUTER SIMULATIONS HELP UNDERSTAND THE SPREAD OF INFECTIOUS DISEASES

Prof. dr. Niel Hens
Centre for Statistics, I-BioStat, University of Hasselt
Centre for Health Economics and Modelling Infectious Diseases, University of Antwerp

How do infectious diseases spread and how can they be effectively halted? This is being studied by Prof. Dr. Niel Hens from Hasselt University. With his research, he informs decision-makers on how the spread of infectious diseases can be reduced or prevented.

For his research Niel Hens developed statistical and mathematical models to predict and better understand the transmission of diseases. In addition, he studies the effectiveness of interventions such as vaccination and quarantine. For this study, the Tier-1 and Tier-2 infrastructure was used. "My colleagues and I needed this computing power to simulate and compare multiple vaccination scenarios. For example, should we vaccinate children against flu as a way to protect elderly people? As many as 5600 vaccination scenarios were run on the VSC infrastructure for this study. The conclusion: the vaccination of children to protect elderly people is meaningful but prohibitively expensive at this point in time."

Chickenpox

Niel Hens also used the supercomputer to develop new models: "To map the spread of the varicella-zoster virus – also known as chickenpox – we have developed a model that maps immunological mechanisms within the host, as well as the infection between hosts. Such modelling is necessary to better understand the spread of the varicella-zoster virus."

Finally, Niel Hens used new computational techniques for solving partial differential equations (PDE). Example: solving PDEs on graphics processing units (GPUs). "These calculations allow the integration of PDEs into a statistical model to estimate significant transmission parameters based on data. In brief, without the VSC supercomputer infrastructure, this research would not have been feasible," says Niel Hens.

RSV and whooping cough

Studies in this context include the development of real-time disease spread models, study of the impact of vaccination on flu, study of respiratory syncytial virus (RSV) and whooping cough, development of methods to better estimate the risk of measles outbreaks... "Moreover, we seek to improve the set-up of studies aimed at collecting data on the spread of infectious diseases."

TRAINING

The VSC spends the necessary time on supporting and training researchers who make use of the infrastructure. It is important that calculations can be executed efficiently because this increases the scientific competitive position of the universities in the international research landscape.

Training organised by the VSC is intended not only for researchers attached to Flemish universities and the respective associates, but also for the researchers who work in the Strategic Research Centres, the Flemish scientific research institutes and the industry.

The training can be placed into four categories that indicate either the required background knowledge or the domain-specific subject involved:

1 Introductory
2 Intermediate
3 Advanced
4 Specialist courses & workshops

Introductory courses are intended for all users of the infrastructure and are highly recommended for those who do not yet possess the required skills. These sessions are conducted by the local VSC staff. This offers researchers the opportunity to become acquainted with the people who answer the questions submitted to the helpdesk. This helps remove the impersonal and anonymous character that is part and parcel of email traffic.

To follow the sessions at the intermediate level, one must have followed the introductory courses so that the required background knowledge has been obtained. These sessions are also more specific in the subjects they deal with. The majority of these courses are intended for users who develop software themselves, either for computing-intensive applications, or for pre and post processing of data. This training is more specialised and intensive than the introductory training, and is therefore not provided at every VSC site. Users are therefore encouraged to attend the training sessions at a different site.

Training at the advanced level requires even more experience and is more domain-specific than the intermediate training. For these courses, VSC hires external instructors. Often they are connected to a PRACE Advanced Training Centre (PATC) or are from the industry. Only two or three such courses are organised each year. Some courses, however, do not fit in with any of the three above-mentioned levels. This may be because they are specific to a certain field or because they encompass everything from introductory to advanced level.

The VSC also offers customised training programmes, both for research groups and for institutions of knowledge, and industry. These usually involve standard training, supplemented with discipline-specific modules specially developed for this purpose.

Existing training modules are naturally brought up-to-date in line with recent developments, but each year there are also new training subjects being selected and developed.

The training offer is published via the VSC website so that the information is available to all interested parties. The announcements are further distributed among the users of the infrastructure via internal mailing lists. Targeted mailings highlight specific training courses that could prove useful to a limited target group or to potential users.

This infographic gives a brief overview of the training offer for 2017.

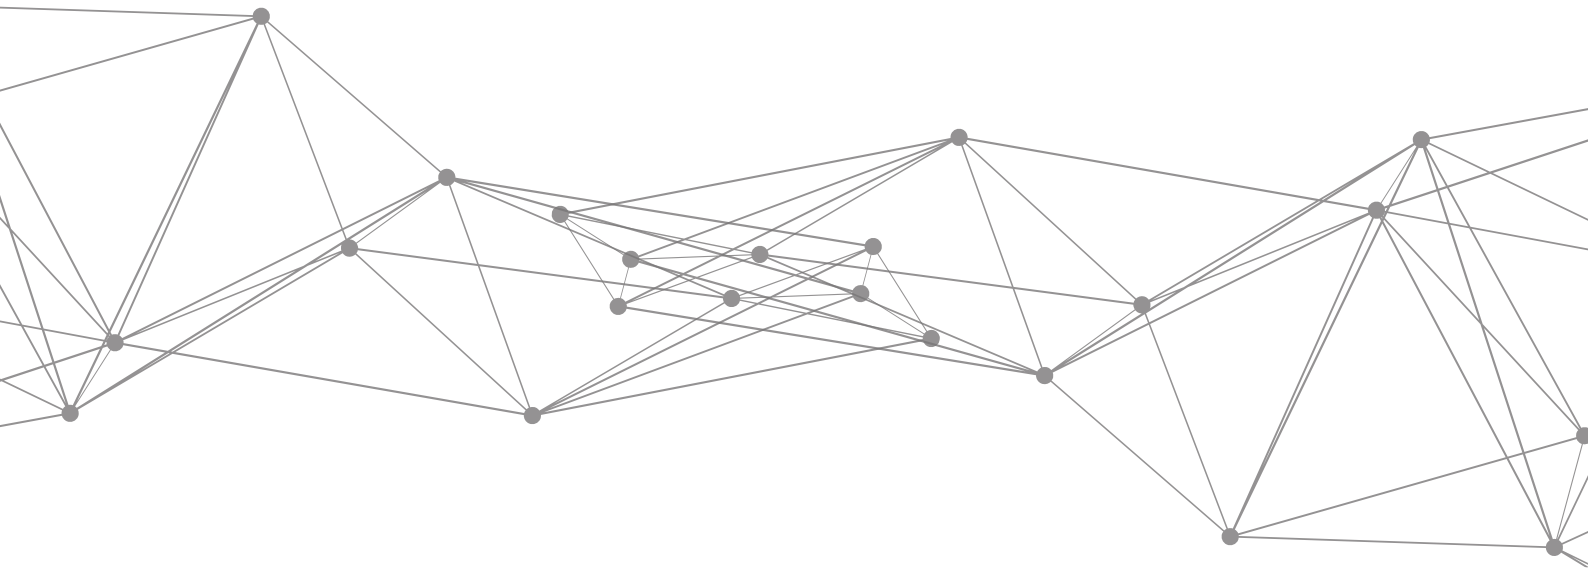

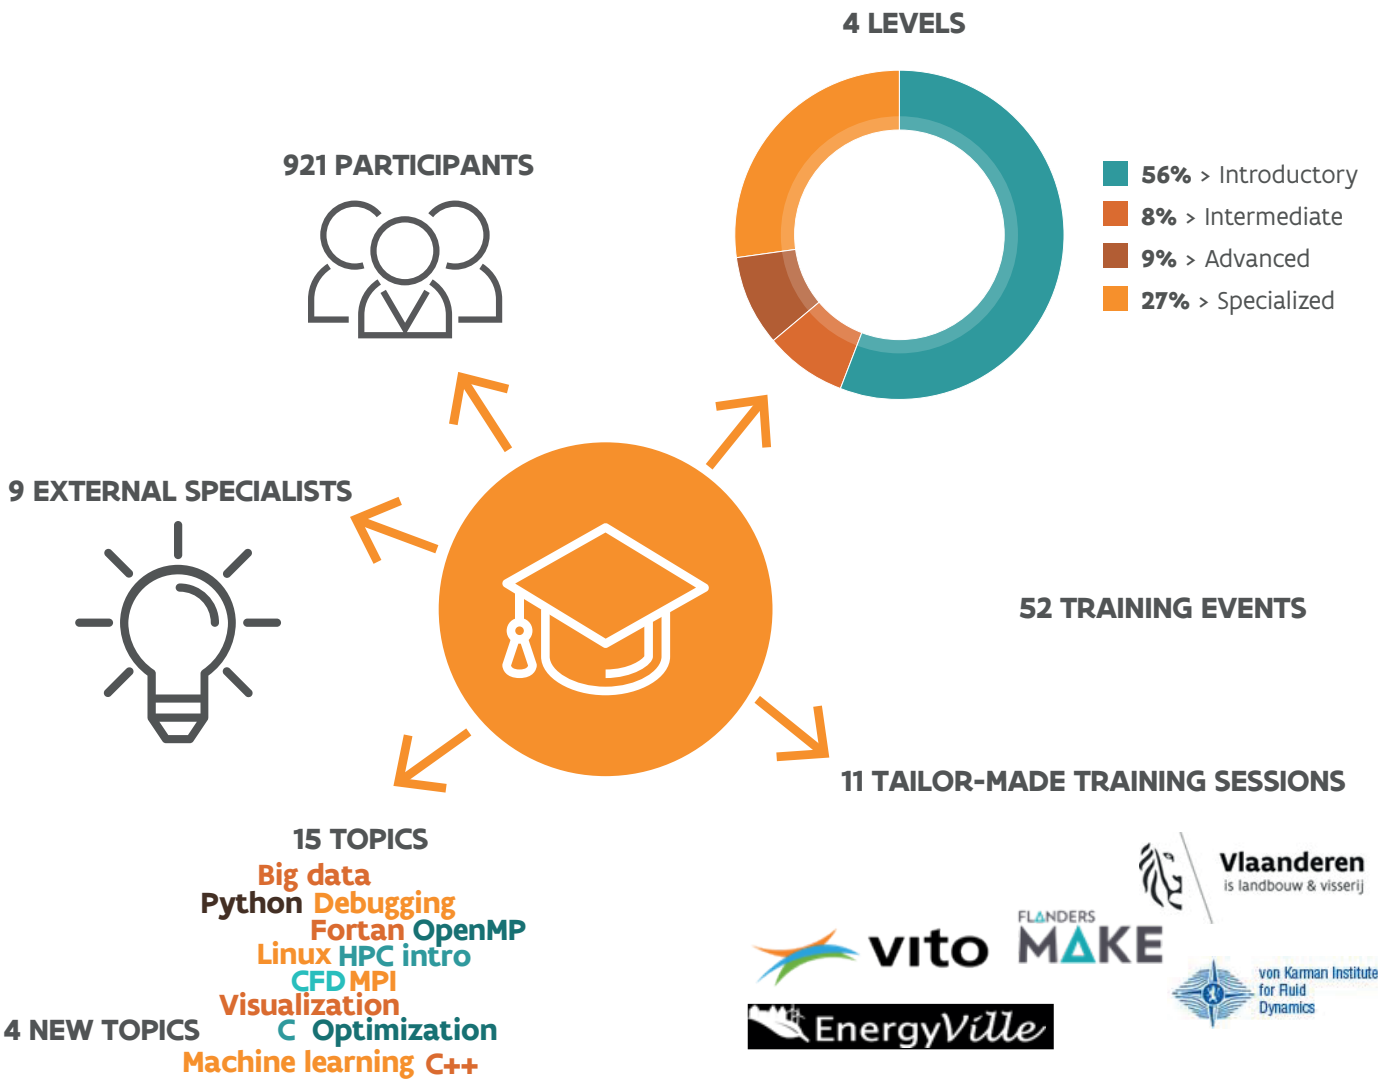

Figure 21 - Training offer 2017

CASE

LARGE-SCALE MEDICAL IMAGE ANALYSIS FOR DIAGNOSIS AND PATIENT FOLLOW-UP

Researcher Mathias Polfliet
Department of Electronics and Informatics (ETRO), VUB

Monitoring tumour growth, comparing joint movements before and after surgery: group-based processing of radiological images provides a wealth of information. Researcher Mathias Polfliet from VUB uses the Tier-1 and Tier-2 infrastructure for his research into the scalability of medical image processing.

Terabytes of data

Until recently, the size of a radiological image was a few megabytes. Today, sequences may consist of several gigabytes for a single acquisition of a single patient's data. Analyses of groups of patients even require terabytes of medical images. To process all these images into useful medical information, the scalability of algorithms and machines capable of handling such large amounts of data are critical.

Within the Department for Electronics and Informatics (ETRO) of the Vrije Universiteit Brussel (VUB) and under the supervision of Prof. Jef Vandemeulebroucke, Mathias Polfliet conducts research into scalable medical image processing. For this, he uses the Tier-1 and Tier-2 infrastructure at the VUB and works together with the Erasmus Medical Center in Rotterdam, the largest hospital in the Benelux.

Image capture

The research is focused mainly on the simultaneous capture of multiple images. Mathias Polfliet: “Image capture allows the detection of minimal changes between different images. In a medical context, this can be used for many applications, for example, to monitor tumour growth, identify anatomical variations in the population, or detect suspect lesions compared to previous scans of the same patient.”

What does this mean for the future?

Group-based image capture provides a wealth of diagnostic information. Mathias Polfliet: “Thus, for example, the age of brains can be estimated based on an MRI scan, or movements of joints can be compared before and after surgery. This could provide quantitative evidence as to whether the surgery is a success or not. However, the high computational cost is still an obstacle for many research groups. By making our scalable software open-source as part of a larger package (https://github.com/SuperElastix/elastix), we are hoping to convince more people of our way of working.”

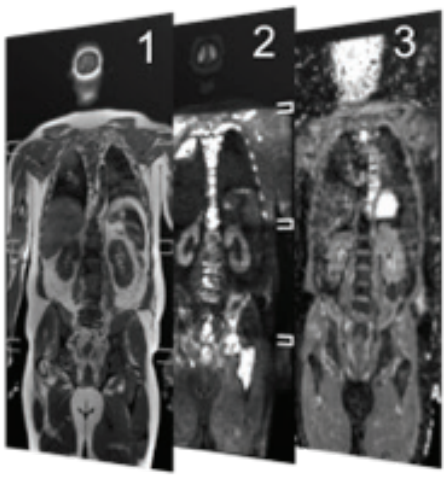

EVENTS

VSC Users Day 2017

On 2 June 2017, the VSC Users Committee organised the third VSC Users Day at the Paleis der Academiën in Brussels. More than 120 participants had registered for the varied programme. Both the morning and afternoon sessions were built around the same pattern: a plenary session followed by a choice from 3 workshops. The following topics were discussed in the workshops: "VSC for starters", "Profiler and debugging", and "Programming GPU's" or "Feedback from Tier-1 Allocation Board".

The international speaker was dr. Achim Basermann of the German Aerospace Center with a lecture on "High Performance Computing with Aeronautics and Space Applications".

The day ended with a poster session during which researchers presented their HPC related research, preceded by a spitfire session of one-minute presentations. This approach also explained why so many participants were clearly interested in the numerous posters. Discussions were continued while enjoying a refreshing drink.
The completed evaluation forms were unanimously positive.

Participation in national, international events, conferences and workshops

ORGANISED BY VSC

- Singularity: Containers in HPC
  7 February 2017, UGent

ORGANISED BY OTHER ORGANISATIONS

- FOSDEM'17
  Co-organisation of "HPC, Big Data and Data Science" devroom
  4-5 February 2017, Brussels
  https://archive.fosdem.org/2017/schedule/track/hpc,_big_data_and_data_science
- ELIXIR Belgium Launch
  9 February 2017, VIB, Gent
- Presentation "Supercomputing for academics & industry - Flemish Supercomputer Center (VSC)" on board of directors of BiR&D (Belgian Industrial Research and Development association)
  8 March 2017, Vilvoorde
  http://birdbelgium.com

- 23nd Quattor workshop
  21-22 March 2017, Annecy, France
  https://www.quattor.org/meeting/2017/03/21/annecy-workshop.html
- CECI users day
  21 April 2017, Louvain-la-Neuve
- PRACE 2017 Spring School - System Administration and Data/Computational Services for Scientific Communities
  Lecture "Modern Scientific Software Management using EasyBuild & co"
  25-27 April 2017, The Cyprus Institute, Nicosia, Cyprus
  https://events.prace-ri.eu/event/601
- TechBoost! 2017 - Innovations in Big Data Analytics
  27 April 2017, Gent
  https://aig.UGent.be/event-2017/techboost-voor-ingenieurs
- EGI Conference
  9-12 May, Catania, Italy
- SESAME Net partner meeting
  17-18 May 2017, Barcelona, Spain
  https://sesamenet.eu
- NEC users group
  22-24 May, Leuven
- iRODS User Group Meeting 2017
  13-15 June 2017, Utrecht, the Netherlands
  https://irods.org/ugm2017/
- ISC 2017
  18-22 June, Frankfurt, Germany

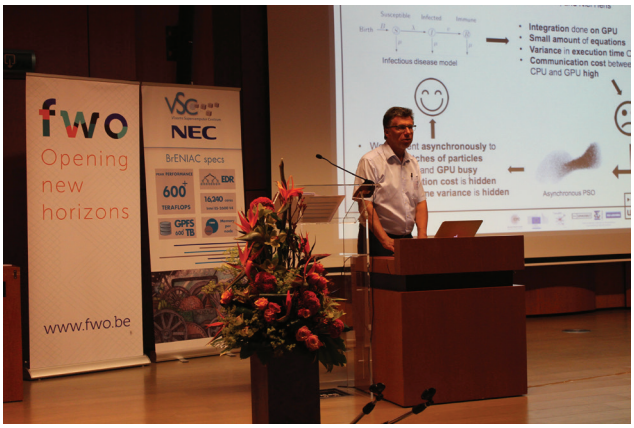

- HPC Knowledge Meeting '17
  Lecture "5 Years of Getting Scientific Software Installed Using EasyBuild"
  15-16 June 2017, San Sebastián, Spain
  http://www.hpckp.org/index.php/annual-meeting/hpckp-17
- Iedereen UGent – "Try out the supercomputer!" experiment market
  8 October 2017, Gent
  https://iedereenUGent.be/en/activities/try-out-the-supercomputer
- 2nd EasyBuild User Meeting
  Lecture "EasyBuild: past, present & future"
  Tutorial "Implementing easyblocks"
  8-10 October 2017, Jülich Supercomputing Centre, Germany
  https://indico-jsc.fz-juelich.de/event/30
- OpenNebula Conf 2017
  Lecture "Transforming an Old Supercomputer into a Cloud Platform"
  23-24 October 2017, Madrid, Spain
  http://2017eu.opennebulaconf.com
- SPLASH 2017
  22-27 October 2017, Vancouver, Canada
  https://2017.splashcon.org/attending/splash-awards
- Supercomputing 2017
  12-17 November 2017, Colorado, USA
- EasyBuild hands-on meeting
  13-14 November 2017, SurfSARA, Amsterdam, the Netherlands
- Discover 2017 (HP/HPE)
  27-29 November, Madrid, Spain
- Digital Infrastructures for Research
  30 November-1 December 2017, Brussels

By annual tradition, the VSC sponsored the Flemish Programming Competition. This competition aims to increase name recognition among students and staff members, especially of association partners. This year the VSC had its own info stand at the event.

INTERNATIONAL COLLABORATION

PRACE and EuroHPC

The VSC is, via the Belgian membership, part of the PRACE happening. PRACE offers the possibility of using Tier-0 computing time. Access to Tier-0 computing time is organised through calls for submissions of project proposals. These proposals are assessed according to 'excellent science' norms. The projects that are ranked highest receive the requested computing time. The 15th call included a researcher from KU Leuven, member of a consortium (led by CNRS) which was allocated 15 million core hours on the French Curie supercomputer. In addition, under the Preparatory Access programme, a researcher from UGent was allocated a total of 200,000 core hours for "Code scalability testing" on the Italian Marconi and French Curie supercomputers. Flemish researchers are made aware of the existence of these calls through notices posted on the VSC website. PRACE also offers interesting training courses that are promoted by the VSC.

A new financing model, PRACE 2.0, has been elaborated since 2016, which is based on a cost of € 3.3m, to be distributed among the members. For Belgium, the standard membership was chosen.

These contributions should cover the costs for the so-called high-level support teams that will provide support and are located on the Tier-0 sites. In this way, Flemish researchers will continue to have access to large-scale supercomputer infrastructure that transcends the capabilities of a country/region.

To put Europe on the exascale map by 2022/2023, the EU created EuroHPC. This programme covers not only the installation of exascale supercomputers, but also the development of the necessary technology (processors, architecture, tools and programs). On 23 March 2017, ministers of 7 European countries signed the EuroHPC declaration. Belgium followed suit on 9 June 2017. In 2018 and 2019 it will become clear whether and how the different countries will participate in the EuroHPC programme.

http://www.prace-ri.eu
http://www.eurohpc.eu

EGI

The VSC is actively present at the European Grid Infrastructure event (EGI). The VUB grid-cluster is part of EGI as a resource centre. Through the support of international virtual organisations, it thus enables researchers from Flanders to use this European computing infrastructure. It is used primarily by high-energy physicians from UAntwerp, UGent and VUB. The grid-cluster supports the international collaboration "IceCube", with participant researchers from VUB and UGent. "SoLid", a new international research group with researchers from UGent, UAntwerp and VUB, amongst others, studies nutrino-oscillations at very short distance from the core of a reactor that was started up in the SCK-CEN in Mol. With support from the VUB-grid team, this research group uses the EGI grid middleware which facilitates data sharing with the participating universities. The WeNMR project is also an active user of the grid-cluster. VUB is active in the EGI core activity to support the "long tail of science" (LToS), an initiative to support individual researchers and smaller research groups, which, unlike large laboratories and partnerships, have no access to specific computing and storage infrastructure. From VUB efforts were made to join the EGI FedCloud project as a certified site for the purpose of making the federated cloud infrastructure available to all VSC users.

https://www.egi.eu

SESAME Net

The VSC is partner in the SESAME Net project, that was awarded to the consortium within the Horizon 2020 programme of the EU. SESAME Net stands for 'Supercomputing Expertise for SmAll and Medium Enterprise Network' and its main objectives are: supporting, expanding and promoting a network of HPC knowledge and HPC experience in Europe. It is also responsible for the dispersal of best practices around HPC usage by the industry. The primary target group are SMEs.

Partners of the SESAME-network

The project was officially launched on 1 June 2015 and completed on 31 May 2017. During the final project review by the EC, the reviewers were very enthusiastic about the advancement and achievements of this project, including:

- development and launch of the HPC4SME self-assessment tool
- publication and dissemination of "EU Roadmap for SME to uptake the HPC"
- participation in various major events (Digital Day Rome, HPC Forum Stuttgart, HPC Summit Barcelona, Hannover Messe, and numerous other local events)

Although officially ended, all current SESAME Net partners have agreed to continue with their network and community activities.

http://sesamenet.eu

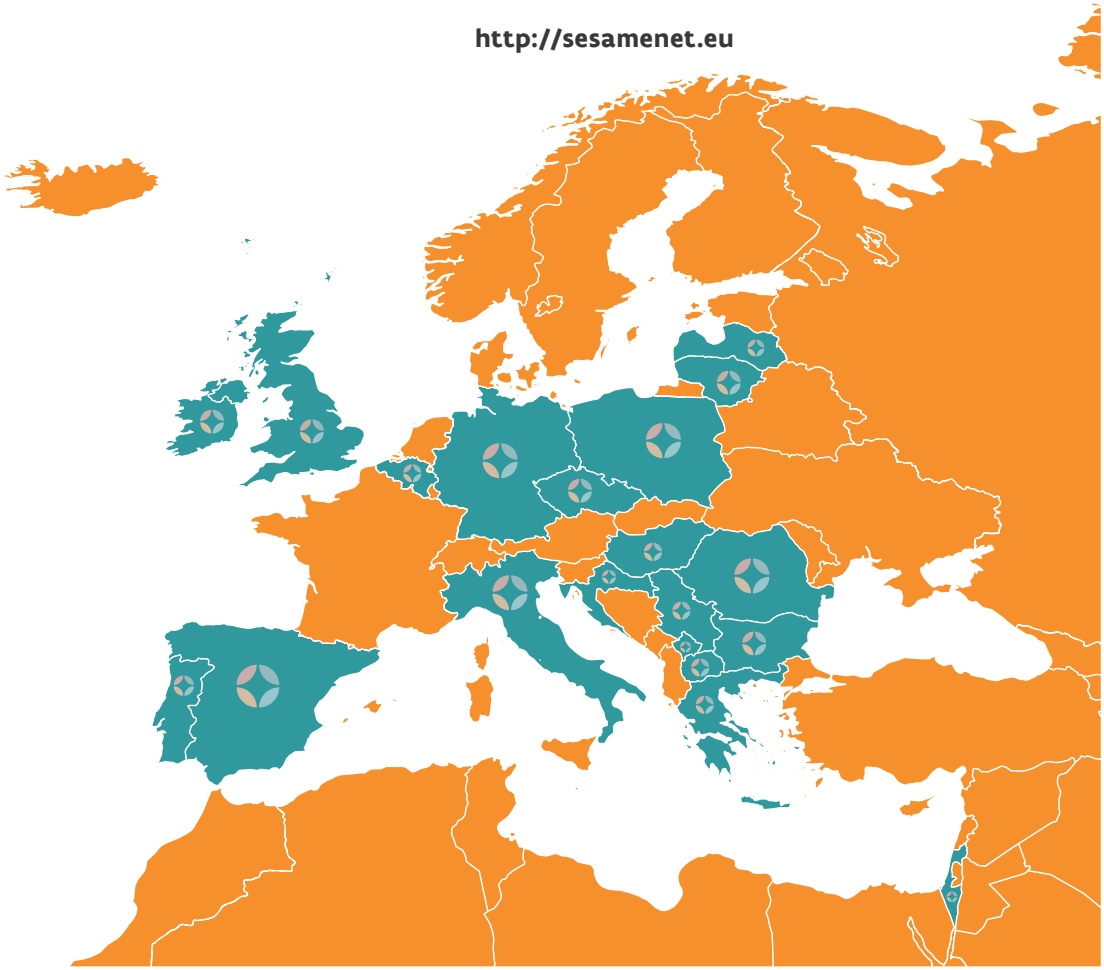

Figure 22 - Partners van SESAME Net

CASE

BETTER PREDICTING THE EFFECTIVENESS OF IMMUNOTHERAPY

Dr. Hui Zhao
Laboratory for Translational Genetics, KU Leuven

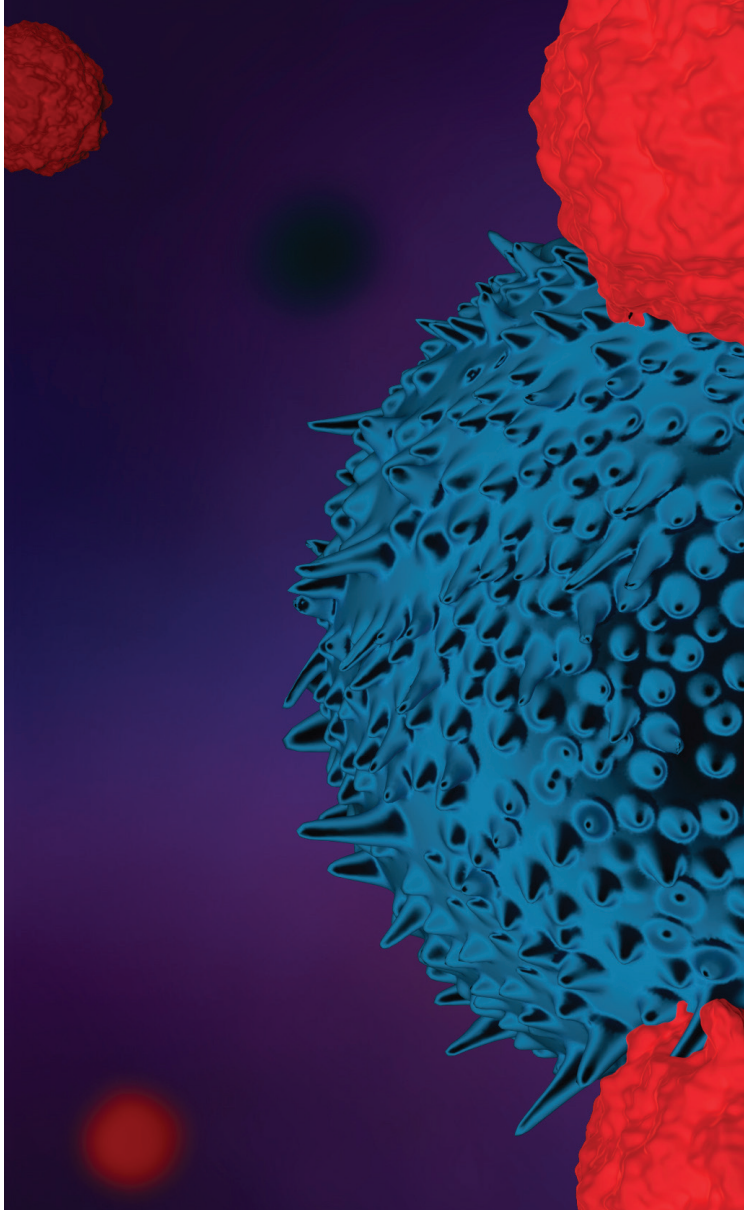

Immunotherapy is increasingly used in the treatment of cancer. Dr. Hui Zhao from KU Leuven is exploring ways of better predicting the success of this therapy.

Immunotherapy stimulates the body's natural defence system, enabling our immune system to attack and destroy cancer cells. Various kinds of cancers such as lung and skin cancers respond very well to this therapy. However, immunotherapy is not equally successful in all patients and is, moreover, very expensive. Dr. Hui Zhao from the Laboratory for Translational Genetics of professor Diether Lambrechts (KU Leuven): "In order to avoid unnecessary costs and side effects, it's essential to be able to predict the effect of the therapy. Using appropriately selected biomarkers and better technologies to detect them, we're looking for new methods of better predicting the efficiency of the therapy."

Terabytes

Hui Zhao' s research more specifically aims to identify genes that are most suitable as biomarkers for immunotherapy. "We focus on detecting differences in gene expression of immune-related genes in tumour tissue and healthy tissue," Hui Zhao explains.

In addition, Hui Zhao used long-read sequencing to detect retroviral sequences in the tumour. Hui Zhao: "Data of different cancer types are taken from TCGA (The Cancer Genome Atlas). The analysis of one sample takes about 24 hours on a Tier-1 node. Many samples are needed to obtain a statistically relevant assessment. We used around 7000 samples in our study. One sample has a size of 5 to 10 gigabytes and requires about twice as much space during processing, so that a large amount of computing time and many terabytes of storage are needed. These data are eventually compared with data from patients who respond positively to immunotherapy."

New techniques, more refined analyses

In the future, researchers will increasingly use new sequencing technologies to study the expression profiles of individual cells. Hui Zhao: "This means that on a piece of tumour tissue with 5000 cells, we will have to perform 5000 analyses on the individual cells. The analysis of an individual cell requires slightly less computing time and data than one tumour, but since the analysis will be reiterated 5000 times, we expect we will need hundred times as much computing time and storage space. But most important of all: these analyses will yield much more accurate predictions and lead to better treatments."

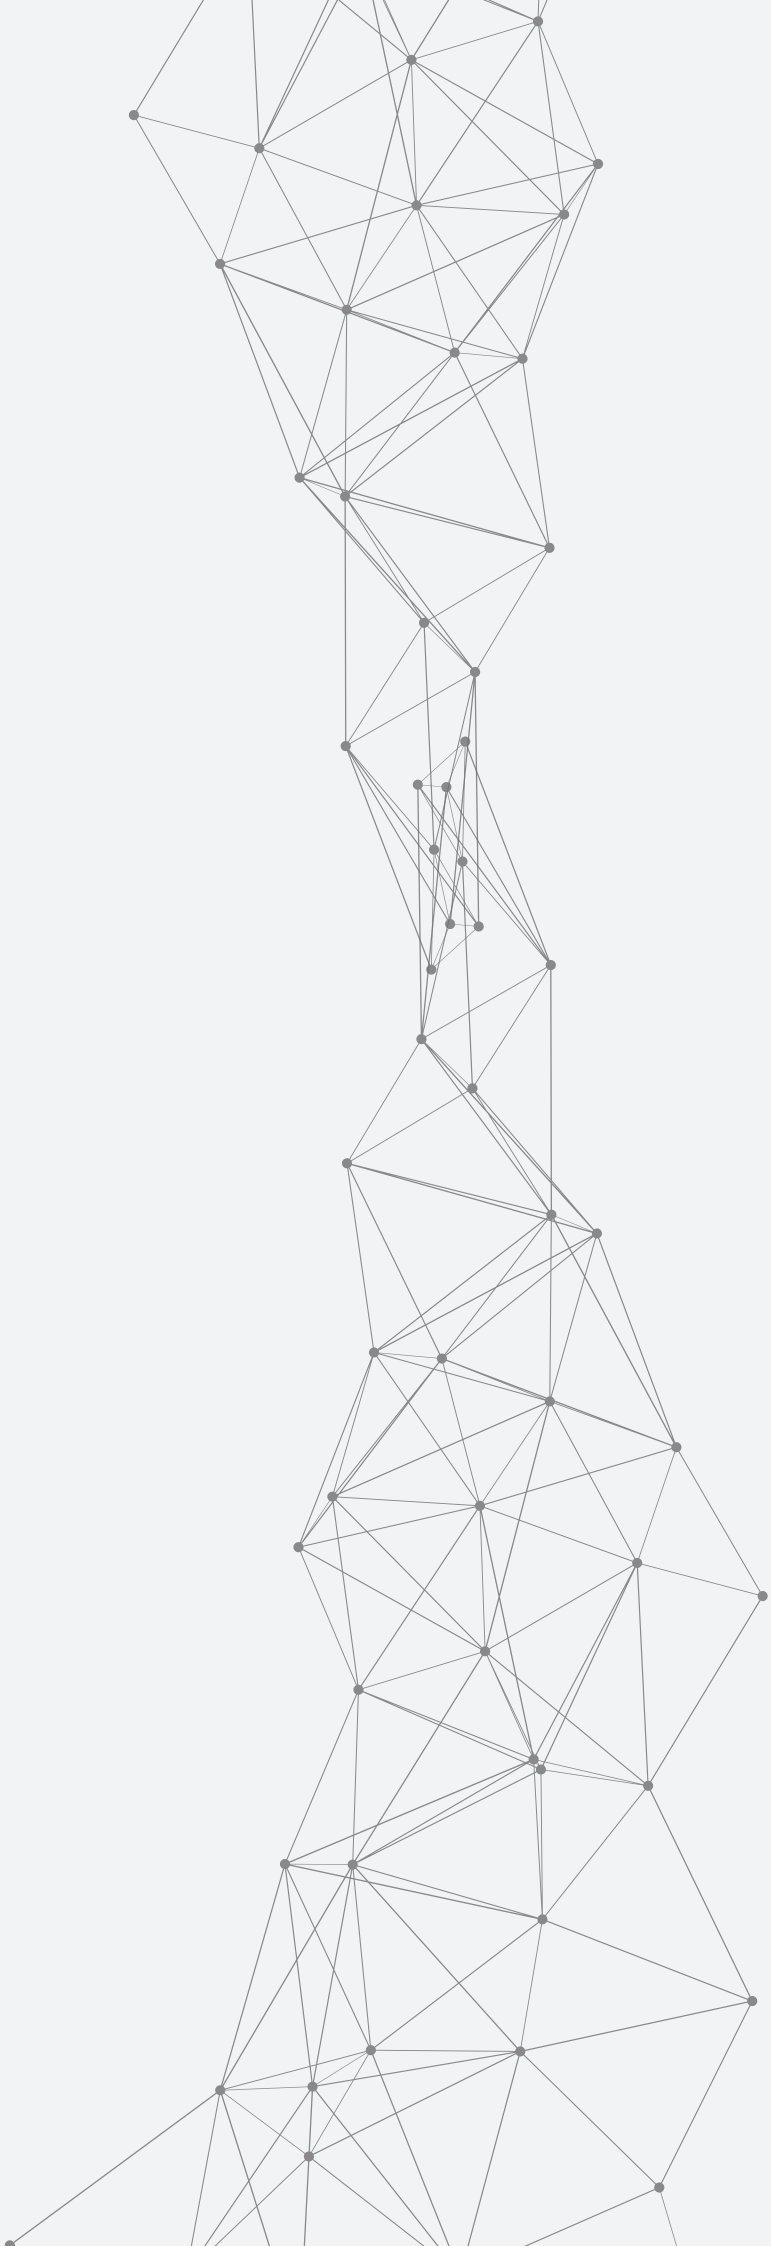

FWO
Research Foundation - Flanders
Egmontstraat 5
1000 BRUSSEL
vscentrum.be