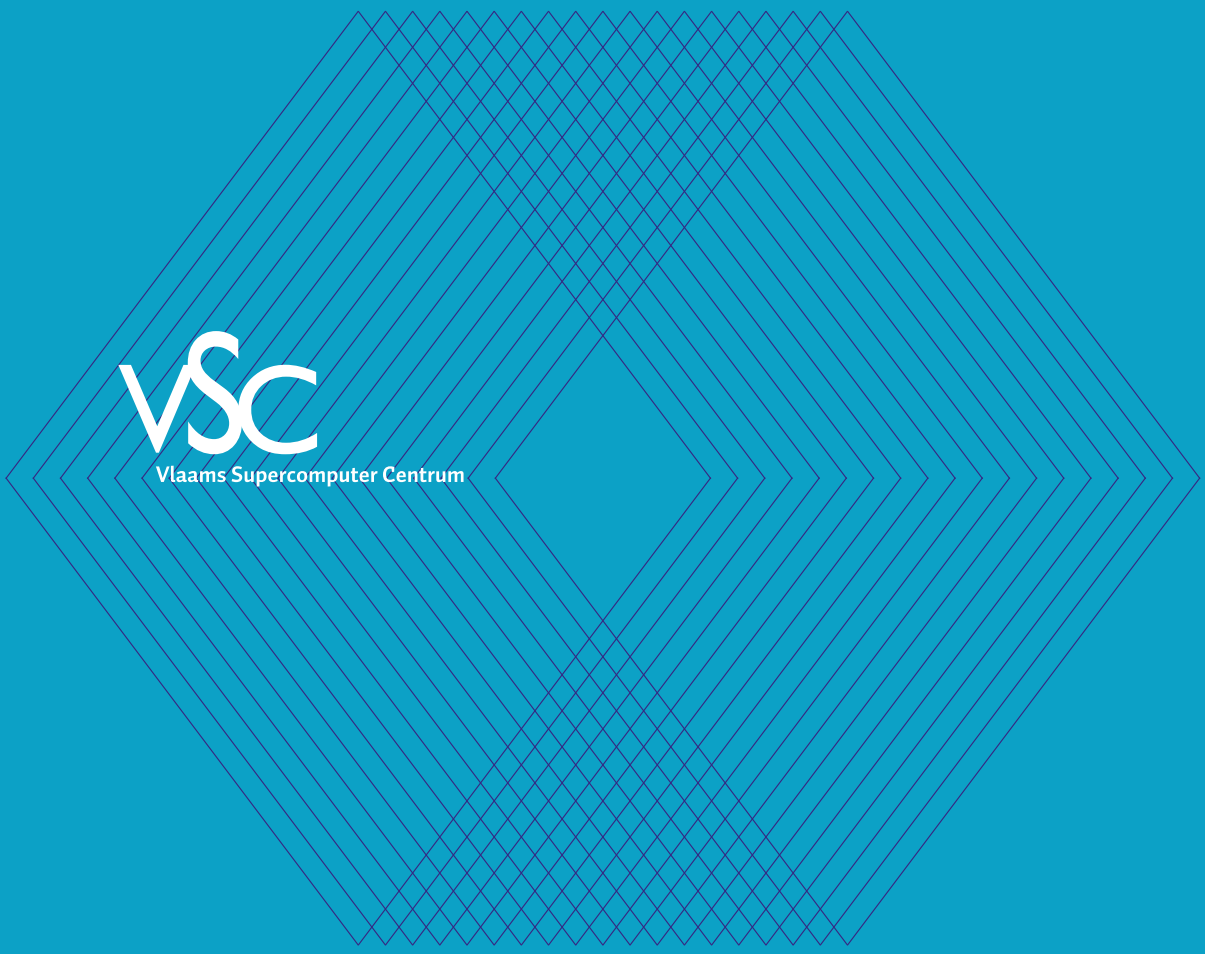

Annual report 2016

www.vscentrum.be

Table of Contents

Preface

For the VSC, the partnership between the Flemish universities and the Research Foundation Flanders (FWO), 2016 arrived as a year full of challenges in a new environment. Having worked for years under the inspiring auspices of the Hercules Foundation, the VSC is now managed by the FWO.

The inauguration of a new supercomputer at KU Leuven in October 2016 was yet another milestone for the VSC. It illustrates the commitment of the Flemish government to continuing investing in state-of-the-art infrastructure for research and innovation. A Tier-1 supercomputer is actually an incentive for academic and industrial researchers to initiate challenging research in R&D.

In 2016, particular attention was also paid to the expansion of the Tier-2 infrastructure at the Flemish universities and additional staff resources were provided for training and user support. The VSC ensures that researchers can easily migrate between the university clusters and the Tier-1 supercomputer so that the most suitable computer is used at all times. In addition, the number of training courses on the use of the Tier-1 and the Tier-2 was expanded for researchers from Flemish universities, strategic research centres and other public institutions of knowledge, but also for companies.

For the allocation of computing time on the Tier-1, the Board of Directors approved a slightly modified regulation, whereby a panel of international experts was again appointed for the evaluation of the applications. In 2016, applications were assessed on three occasions. A total of 198,490 node days were allocated to 74 projects. From 2016, academic researchers are no longer charged part of the costs, which had had a positive impact on the number of submitted applications.
The VSC Users Committee identified also in 2016 users' needs and formulated proposals to improve user service.

Special attention was paid to collaboration with companies. The Industrial Board was re-appointed in 2016 and developed initiatives to highlight the advantages of a collaboration with the VSC: professional support, customised training as required, but above all embedding within an academic environment.
In brief, 2016 was a busy year, but also one that allows both the VSC and the FWO to look forward to a bright future!

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

About VSC

The Flemish model in a nutshell

In the European model for High Performance Computing (HPC), a distinction is made between three levels: the processing capacity which research institutes have at their disposal (Tier-2), the processing capacity which goes beyond the needs and costs of an individual institution and is provided at regional or country level (Tier-1), and the super powerful processing infrastructure (Tier-0).

At the start of this century, Flanders lagged behind in HPC. The universities had invested in local processing capacity and the government had allocated punctual funding, but an overall vision was lacking. Acting on the 'Advice of the Royal Flemish Academy of Belgium for Science and the Arts (KVAB) on High Performance Computing in Flanders', the Flemish Supercomputer Centre was founded.

The Flemish Supercomputer Centre (VSC) is a partnership between the five Flemish universities and their associations: Antwerp University Association, University Association Brussels, Ghent University Association, KU Leuven Association, and the Limburg Association of Higher Education. This consortium brings together know-how in scientific and technical computing (including high performance computing, high throughput computing, cloud computing and data processing) in Flanders. It houses infrastructure in four hubs: UAntwerp, Vrije Universiteit Brussel, UGhent and KU Leuven. The VSC is managed by the FWO, Research Foundation Flanders.

At the end of 2012, the first Tier-1 supercomputer, housed at UGhent, was taken into service. It is accessible to public research institutions such as universities, strategic research centres and institutions for post-initial education, and companies in Flanders. In October 2016, the second Tier-1 supercomputer, housed at KU Leuven, became operational. However, investing in a Tier-1 is only sensible if, at the same time, each university can be assured of its own Tier-2 capacity.

For the expansion of HPC in Flanders, the Flemish consortium model offers important benefits and furthermore reduces the costs. The Tier-1 and Tier-2 computers are connected together via BELNET, the federal research network, and applications can migrate to the machine that is best suited to them. The employees are appointed at the five Flemish universities but form an integrated team that is responsible for the training and the support of users. For specialised support, each institution can make use of the specialist, irrespective of where he or she is employed. Hardware and software alone are not enough. It is at least equally important to ensure there are sufficient competent employees who can train and support the researchers. High processing capacity is after all used by researchers from the most diverse disciplines. These are people who are familiar with their professional discipline but not necessarily with the advanced techniques that are necessary to make optimal use of software on large computers.

The universities also invest in HPC infrastructure and the VSC can appeal to the central services of these institutions. In addition, embedment within an academic environment creates opportunities for cooperation with industrial partners.

Finally, the VSC monitors the international developments and the possibilities that applications on Tier-0 machines offer to Flemish researchers. For example, Belgium is a member of PRACE (Partnership for Advanced Computing in Europe), a project of European strategic importance that is included in the roadmap of the European Strategy Forum for Research Infrastructure (ESFRI).

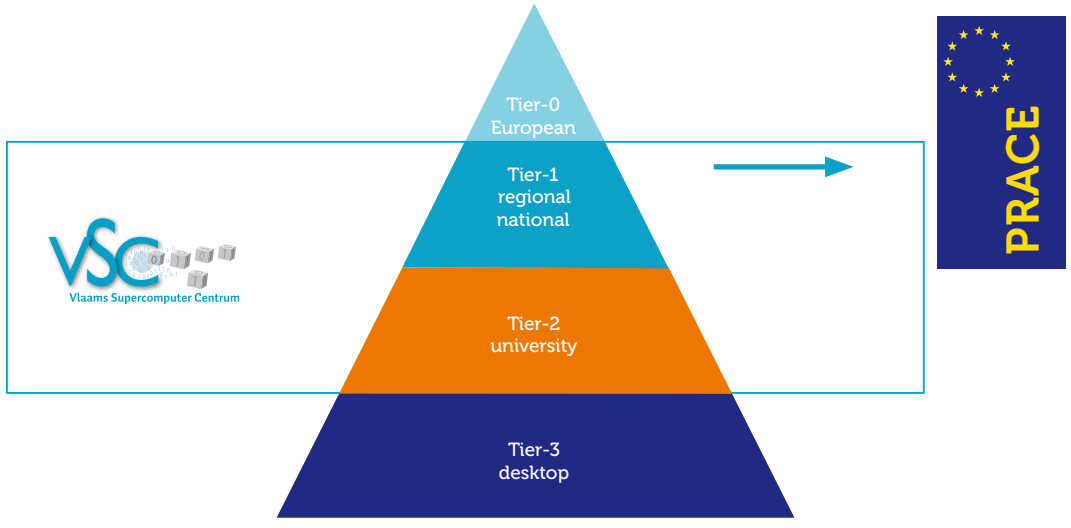

Mission and Objectives of VSC

Mission of the VSC

The VSC encourages the use of scientific and technical computing in the Flemish academic and industrial landscape. To this end it provides infrastructure, training and services. In addition, VSC acts as a lever to promote the importance of scientific and technical computing and its added value to society.

Objectives of the VSC

The strategic objectives can be summarised as follows:

- Provide the academic and industrial landscape with access to diversified ICT infrastructures that is tailored to the needs of scientific/technical computing.
- Provide a common user environment on the computing infrastructure which is embedded in the operation of the four hubs.
- Provide support to its users allowing them to take their research or development in the field of scientific/technical computing to a higher level.
- Inform about the latest developments in scientific/technical computing and its potential added value.
- Promote the exchange of ideas and expertise among public research institutes and industry.
- Offer a diversified and coordinated training programme throughout the VSC consortium to encourage the use of scientific/technical computing among new and existing users.
- Engage and actively participate in international initiatives such as PRACE and Horizon2020 and collaborate with other centres specialising in scientific/technical computing.

VSC Advisory Bodies

To ensure an effective operation of the VSC, a number of advisory bodies were created.

VSC Steering Group

The VSC Steering Group, which consists of FWO staff and persons in charge of the coordination of the HPC activities at the respective universities, is responsible for the substantive operational management of the VSC. The Group meets every month. During these meetings, operational agreements are made and strategic advice is given to the FWO Board of Trustees.

VSC Steering Group 2016

institution/organisation	representative(s)	
KU Leuven	Leen Van Rentergem	Jan Ooghe
University of Ghent	Ewald Pauwels	
University of Antwerp	Annie Cuyt	Stefan Becuwe
VUB	Rosette Vandenbroucke	Stefan Weckx (from 01/12)
University of Hasselt	Geert Jan Bex	
FWO	Caroline Volckaert	Bart Van Beek

VSC Industrial Board

To develop the cooperation with companies but also with the non-profit sector, an Industrial Board was installed.

- The Industrial Board serves as a communication channel between the VSC and the Flemish industry and aims to reinforce the exchange of ideas and expertise among public research institutes and industry.
- The Industrial Board 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.
- This board thus assists in devising a policy to increase the involvement of Flemish industry in HPC and to stimulate collaboration with the VSC, as well as to effectively promote the use of supercomputing in the development of innovative products and services.

On 19 October 2016, the FWO Board of Trustees appointed the VSC Industrial Board for a term of two years, with the following composition:

Members of VSC Industrial Board 2016

Mia Vanstraelen	IBM
Charles Hirsch	Numeca
Herman Vander Auweraer	LMS Siemens
Benny Westaedt	Van Havermaet
Marc Engels	Flanders Make
Marcus Drosson	Umicore
Sabien Vulsteke	Bayer Crop Science
Birgitta Brys	Wordline

Mrs Mia Vanstraelen was appointed as chairperson for the first operating year. For the subsequent years, the chairperson will be appointed from among the members.

VSC Users' Committee

For advising on user needs and formulating proposals to improve the service, including user training, the FWO Board of Trustees installed a Users' Committee of at least seventeen members, on the recommendation of the university-university college associations, the strategic research centres and the Flemish Government. In addition, each university has its own users' committee.
The members were appointed for a term of three calendar years starting on 26 June 2014. To keep in touch with the needs of new users, it is desirable that a number of members can be replaced at periodic intervals. When the activities were transferred to the FWO in early 2016, all members were asked whether they were prepared to continue their terms under the new structure. In response to this, the representation from iMinds was changed and Jan Fostier was designated to become a member of the Users' Committee from 2016. For UGhent, Marie-Françoise Reyniers takes the place of Veronique Hoste.

Members VSC Users' Committee 2016

nominated by	Name	Institution
Flemish government	Francisco Hernandez	VLIZ
Association Leuven	Dirk Roose	KU Leuven
	Nicole van Lipzig	KU Leuven
	Nele Moelans	KU Leuven
Association Ghent	Veronique Van Speybroeck	Ghent University
	Dirk Van den Poel	Ghent University
	Veronique Hoste / Marie-Françoise Reyniers	Ghent University
Association Antwerp	Michele Giugliano	University of Antwerp
	Wouter Herrebout	University of Antwerp
Association Brussels	Frank De Proft	VUB
Association Limburg	Niel Hens	University of Hasselt
Strategic Research Center	Clemens Mensink	VITO
	Steven Maere	VIB
	Piet Demeester / Jan Fostier	iMinds
	Wilfried Verachtert	IMEC

Funding of Tier-1 and Tier-2

In 2016, the Tier-2 infrastructure was co-financed by the FWO for an amount of €4,270,800. FWO also invested €776,000 in staffing and energy costs for the first Tier-1 and in the start-up of the second Tier-1 and made a one-time investment of €5,5000,000 in the hardware of the second Tier-1.

These resources were used for funding of:

- staff costs for operation of the first Tier-1 (€190,000) and start-up of the second Tier-1 (€190,000);
- staff costs for training and support of users of both Tier-1 and Tier-2 (€1,425,000); (For this amount, the five Flemish universities can subsidise in total the equivalent of 15 FTEs.)
- energy costs of the Tier-1 (€216,000) and part of the energy costs of the second Tier-1 (€180,000);
- investment and operating costs for the Tier-2 infrastructure (€2,845,800);

(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.

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.

Summary Tier-1 and Tier-2 funding in 2016

Heading	KU Leuven	UHasselt	UGhent	VUB	UAntwerp	Total	Tier-1a	Tier-1b
Staff in FTE	5	1	4	2	3	15	2	2
Staff in Euro	€ 475,000	€ 95,000	€ 380,000	€ 190,000	€ 285,000	€ 1,425,000	€ 190,000	€ 190,000
Hercules key 2014	0.4316	0.0324	0.3187	0.1008	0.1165	1.000		
Operation Tier-2	€ 1,228,247	€ 92,204	€ 906,956	€ 286,857	€ 331,536	€ 2,845,800		
Operation Tier-1 (energy)							€ 216,000	€ 180,000
Belnet-connection	€ 35,000							
total 2016	**€ 1,738,247**	**€ 187,204**	**€ 1,286,956**	**€ 476,857**	**€ 616,536**	**€ 4,270,800**	**€ 406,000**	**€ 370,000**

The second Flemish Tier-1

The first Flemish Tier-1 supercomputer was taken into service in the first half of 2013 and was in urgent need of replacement.

The second Flemish Tier-1 supercomputer, representing an investment of €5,500,000 was inaugurated at the data center of KU Leuven at the end of October 2016. The company NEC was selected to build the machine through a public tender procedure.
The second supercomputer has a processing power of more than 600 TFlop/s, the equivalent of 2000 fast PCs. This machine thus has a capacity that is three times that of the first Flemish supercomputer. The second Tier-1 supercomputer was equipped with the latest Intel processors. Also the memory, the internal network and the storage capacity were adapted, allowing the supercomputer to efficiently perform complex computations. At time of its inauguration in October 2016, the new supercomputer was one of the 200 fastest computers in the world.

The supercomputer in Leuven will support research into, for instance, renewable sources of energy or the development of new materials and medication. The supercomputer will also enable researchers to create even more detailed climate models or to map the climate on other planets. The supercomputer also opens up new horizons for fundamental research, which is of critical importance in the medium term.

Flemish minister Philippe Muyters: *"The Flemish Government aims to focus consistently on new investments in research and innovation. Our research institutions are among the best in the world. We can only maintain and strengthen that top position if we keep investing. The supercomputer can play a crucial role in a wide range of fields. It fills me with pride to see that the Flemish Government has fully played its role in bringing this about."*

From Tier-1a to Tier-1b

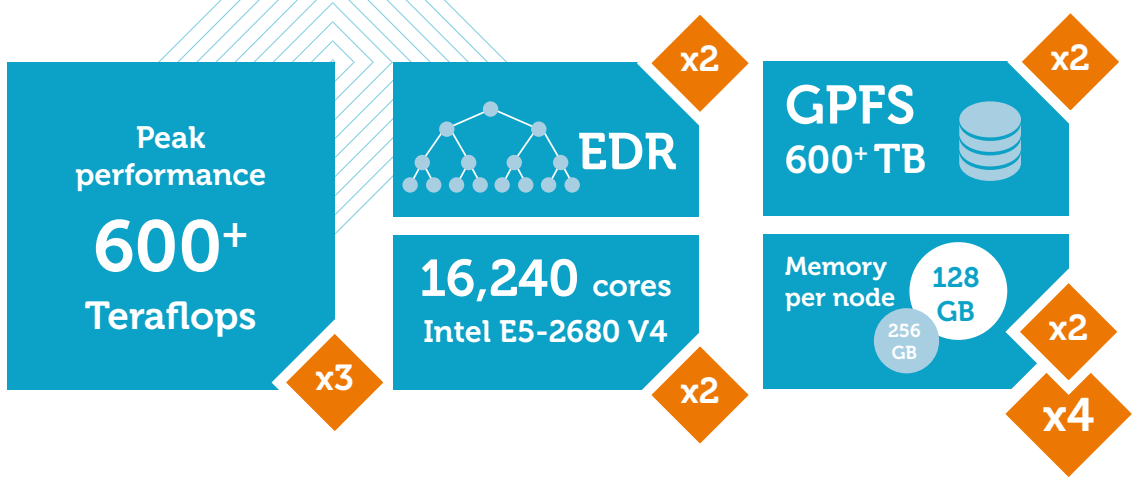

Tier-1 Infrastructure

Tier-1 at University of Ghent (UGhent)

The first Flemish Tier-1 went into production in 2013. The purchase of this supercomputer was financed with FFEU funding. UGhent 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 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 this too, became inactive.

After a cost analysis, it was decided to reconfigure the still usable hardware of this machine as an on-demand cluster, yet without the same service level as that of a fully-fledged Tier-1 or Tier-2 cluster. Parts of the hardware can, per chassis, be rented for a fee by researchers of Flemish knowledge institutions or industrial users. This requires a substantial reconversion effort, which will be launched as soon as all researchers have retrieved their data from the storage. Further procedures for use and rental will be determined upon completion of the reconversion.

The VSC has gone through a whole learning process in the operation of this machine. Although UGhent took care of all the technical maintenance and modification aspects, all VSC sites collaborated intensely in appointing an allocation board, drawing up, and repeatedly adjusting, a project application procedure and regulations, training and stimulating users to initiate Tier-1 projects, etc.. This learning experience also allowed the transition from the 'old' to the 'new' Tier-1, at the end of 2016, to go very smoothly.

Flanders' first Tier-1 machine was also used very well. The following graph provides an overview of the use of machine since it went into production. A total of 144 projects, corresponding to 7,163,200 core days of computing time, were run on this Tier-1. This would correspond to over 9,800 years of computing on a conventional laptop (with 2 cores). On this Tier-1, also 435,000 core days of computing time were reserved for companies in the sectors of renewable energy, pharmaceuticals, medical technology, electrical engineering and material technology.

The presence of a Tier-1 system has thus resulted in a mind shift among Flemish researchers in academics and industry. A Tier-1 induces them to think more out of the box and to take on greater challenges in scientific research or within R&D.

As shown by the graph below, molecular modelling was by far the most popular scientific field in the research projects run on Tier-1 muk, followed by climate modelling.

CPU core hours used on Tier-1 @ UGhent

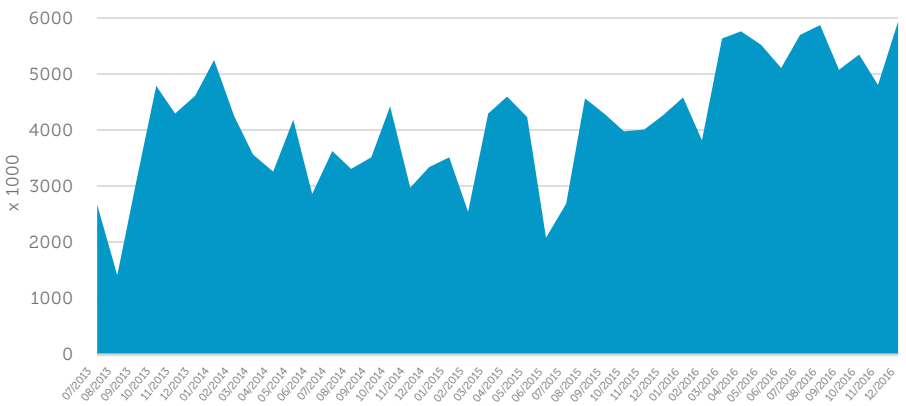

Figure 1
CPU core hours used on Tier-1 @ UGhent

Computing time used on Tier-1 @ UGhent (muk) by scientific field

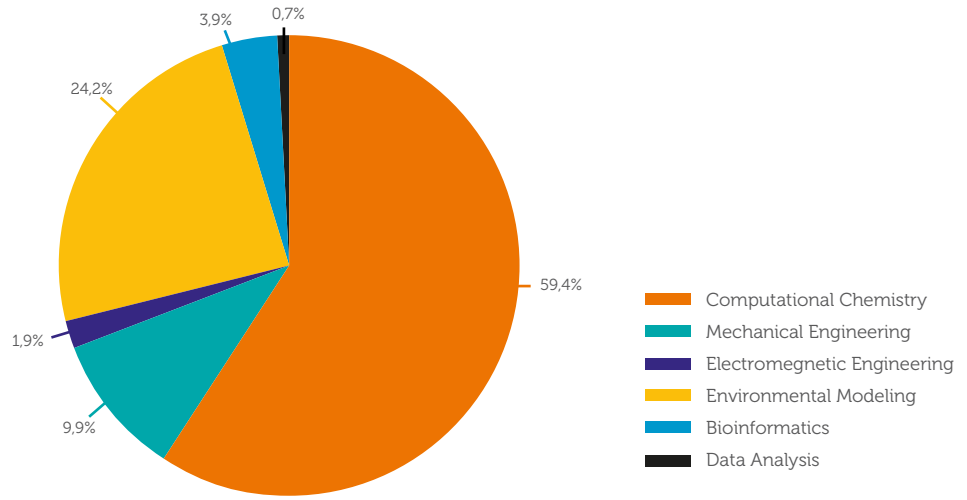

Figure 2
Computing time used on Tier-1 @UGhent (muk) by scientific field

Tier-1 at KU Leuven

The implementation of the second Tier-1 was a major project in 2016. It was divided up into a number of phases:

- installation and acceptance
- test period
- production phase

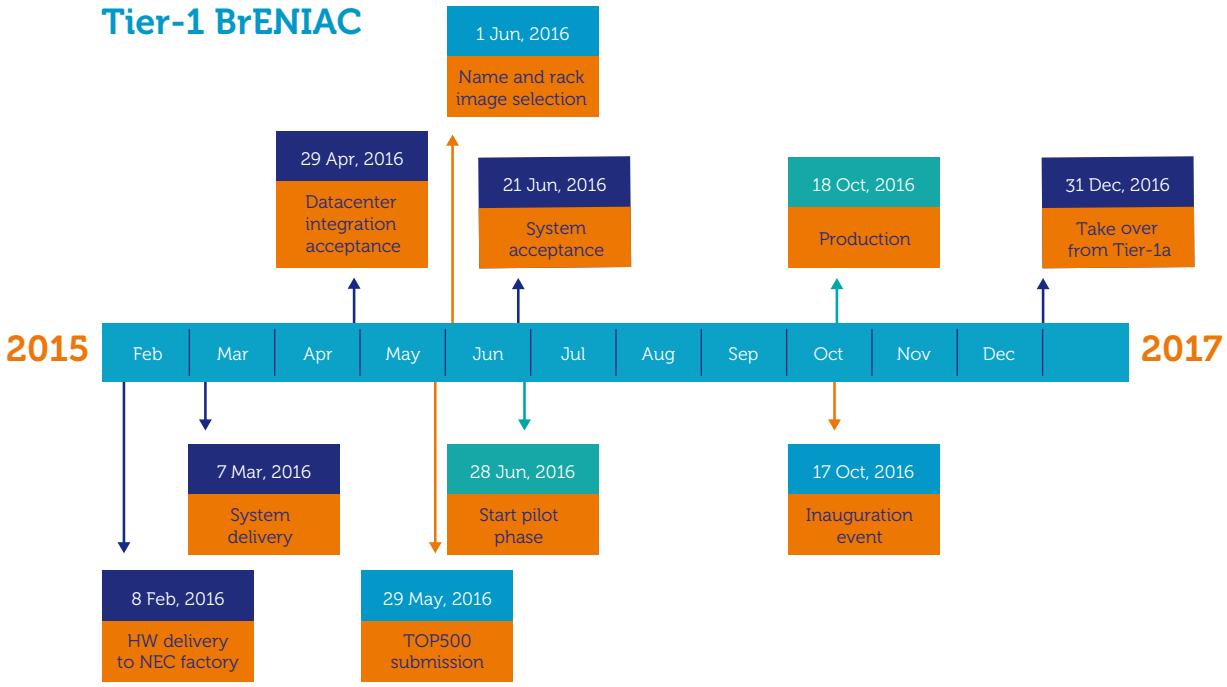

Installation and implementation

The machine components were delivered at the data center on 7 March 2016. In an initial phase, modifications were made to the data center to adapt the cooling and power supplies to the new Tier-1. On 29 April 2016 all work had been completed and accepted. In the meantime, work on the assembly of the machine had been started. After the installation of the basic system software, the machine was fully integrated into the VSC HPC environment. Next, tests were carried out in accordance with the specifications. A Linpack benchmark could be completed in time to officially register the machine in the Top 500. The machine was placed 196th with a score of 548 TFlop/s. The system was accepted on 21 June 2016. This new Tier-1 has twice the number of computational units as its predecessor, allowing three times as many calculations to be carried out. The faster network (by factor of 2) also allows more computing work to be done in parallel.

Test period

To introduce the users to the new supercomputer and also to test it in a near-production environment, a pilot phase was launched. All Tier-1 projects that were active on the first Tier-1 supercomputer in June 2016, were allocated 1000 node days on the new Tier-1. In addition, all researchers were able, via an open call, to submit a project for the allocation of computing time on the new Tier-1 for the purpose of checking software installations and performing benchmarks in preparation for new Tier-1 project applications. A total of 23 projects were allocated computing time. The participants in the pilot phase came from different research fields, adequately reflecting the actual use.

In preparation for the pilot phase, info sessions were organised at all the universities. At these info sessions all VSC users were introduced to BrENIAC. They were explained the architecture and configuration of the machine, which was important to successfully migrate the work from the first Tier-1 to the second Tier-1. A common user environment is used within VSC, so that the switchover went smoothly. During the pilot phase, the most commonly used software was installed and tested, and minor configuration adjustments were made. On the whole, the system was stable and ready for production. Some pilot users reported performance gains of up to 30% or codes that still scaled efficiently on the faster infrastructure for a larger number of cores (up to double the amount).

Production period

On 17 October 2016, the machine was officially inaugurated in the presence of Minister Muyters. The inauguration was given the necessary press coverage, thereby bringing HPC and VSC to the attention of the general public. The inauguration marked the end of the pilot phase and the start of the production phase of BrENIAC. The 3rd call for projects for computing time in 2016 was the first to be fully allocated on BrENIAC. A total of 27 projects for a total of 90,000 node days were submitted. For the first application round on the new supercomputer - the third call in 2016 - 64,000 node days were available. This meant that projects had to be rejected already from the first call.

For the next application rounds, less computing time will be available, as projects of 2 active application rounds will be active each time. Consequently, in future an even higher percentage of Tier-1 project applications will have to be rejected.

The utilisation of the machine shows that the full capacity is already being used (cf. graph n° 4). The new investment is fully utilised. The rise in the number of project applications and requested computing time already suggests that the needs of researchers are not yet saturated.

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, a strategic research centre, 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. Four foreign experts reside on this board:

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

Mrs Caroline Volckaert of the FWO takes care of the secretariat. The HPC coordinators of the Flemish universities are 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 users 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 2016. 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 2016, 10 Starting Grants were awarded and 4 Exploratory Access projects were ongoing.

Starting Grants 2016

Jos Teunissen	VUB	Computational Chemistry
Wilfried De Corte	University of Ghent	Data Analysis
Tom Ruttink	ILVO	Bioinformatics
Wim Delva	University of Ghent	Bioinformatics
Charlotte Vets	University of Antwerp	Computational Chemistry
Stefan Knippenberg	University of Hasselt	Computational Chemistry
Tatiana Woller	VUB	Computational Chemistry
Laura van Bergen	VUB	Computational Chemistry
Ehsan Moravveji	KU Leuven	Environmental Modeling
Jose Guillermo Rivera de la Cruz	University of Ghent	Computational Chemistry

Exploratory Access 2016

FEOPS	Medical engineering
Bekaert nv	Mechanical engineering
Siemens	Mechanical engineering
Umicore	Materials Design

Approved Tier-1 applications

All applications approved in 2016, grouped per evaluation moment, are listed below.

01 February 2016

Applicant	Host institution	Title	Science field	Allocated computing time (node days)	Allocated SCRATCH storage (TB)
Reyniers Pieter	University of Ghent	Computational Fluid Dynamics based run length simulations of enhanced 3D reactors (I)	Mechanical Engineering	4,000	2
Michael Sluydts	University of Ghent	Shape tuning of CdSe nanostructures by ab initio determination of the anisotropic growth mechanism	Computational Chemistry	4,752	0.602
Simon Bailleul	University of Ghent	Ab initio molecular dynamics study on the role of water in the reaction mechanism during methanol conversion in H-SAPO-34	Computational Chemistry	4,880	0.656
Thierry De Meyer	University of Ghent	pKa calculation of Bromothymol Blue and modified derivatives	Computational Chemistry	3,180	0.047
Nasrin Sarmadian	University of Antwerp	The role of impurities in the gas sensing mechanism of the ZnO (10-10) surface	Computational Chemistry	4,173	0.264
Chiara Caratelli	University of Ghent	Investigating active sites in hydroxylated and dehydroxylated UiO-66 for catalysis of Oppenauer-type oxidation	Computational Chemistry	2,110	0.334
Kristof De Wispelaere	University of Ghent	Dynamical first principle modelling of ethene oligomerization in Ni-SSZ-24	Computational Chemistry	3,232	0.021
Wilfried De Corte	University of Ghent	Robustness and Sensitivity of Pareto-optimal Selection Designs	Data Analysis	3,000	0.5
Giovanni Lapenta	KU Leuven	Kinetic SimulationS of the Magnetospheres of Mercury and the Earth (KISSMME)	Environmental Modeling	5,000	1
Piet Termonia	University of Ghent	UGhent EURO-CORDEX climate runs	Environmental Modeling	4,500	32
Pieter Verhees	University of Ghent	Development of a one dimensional model for the evaporation of hydrocarbons based on 3D CFD simulations	Mechanical Engineering	1,500	0.9
Sam De Waele	University of Ghent	Error estimates for ab initio predictions of surface energy and work function	Computational Chemistry	1,165	2.5
Geoffrey Pourtois	IMEC	Modeling of quantum transport of two-dimensional nanoelectronic devices	Computational Chemistry	1,000	0.281
Johan Meyers, Wim Munters	KU Leuven	Optimal control of wind farm boundary layers	Environmental Modeling	4,320	9.21
Johan Meyers, Dries Allaerts	KU Leuven	Simulation of shallow and strongly stratified wind-farm boundary layers	Environmental Modeling	4,960	1.3
Oriana De Vos	University of Ghent	Study of oxygen diffusion through cell membranes	Computational Chemistry	1,125	0.54
Matthias Vandichel	University of Ghent	Unraveling promoting elements in selective CO-oxidation	Computational Chemistry	3,744	1.488

06 June 2016

Applicant	Host institution	Title	Science field	Allocated computing time (node days)	Allocated SCRATCH storage (TB)
Knippenberg Stefan	University of Hasselt	Atomistic picture of fluorescent probes into lipid bilayer membranes: probing lipid rafts through (non-) linear absorption spectroscopy	Computational Chemistry	1,040	0.52
Vanpoucke Danny	University of Hasselt	Breathing behavior of flexible Metal-Organic Frameworks with MIL-47/53 topology	Computational Chemistry	4,970	0.48
Pourtois Geoffrey	IMEC	Modeling of quantum transport of two-dimensional low power nano-electronic devices	Computational Chemistry	4,800	0.281
Amini Mozhgan	University of Antwerp	Interaction of halides with Au-surfaces	Computational Chemistry	4,105	0.2
Sarmadian Nasrin	University of Antwerp	Project extension: The role of impurities in the gas sensing mechanism of the ZnO (10-10) surface	Computational Chemistry	4,528	0.306
Dabaghmanesh Samira	University of Antwerp	Cr2O3 as a p-type transparent conducting oxide: alloying with sulfur and selenium	Computational Chemistry	3,999	0.2
Jaeken Jan	University of Ghent	Ab initio prediction of temperature-dependent acoustic anisotropy of iron in the Earth's inner core	Computational Chemistry	3,871	1
Saniz Rolando	University of Antwerp	Electronic properties of defects at the surface of CuInSe2	Computational Chemistry	3,024	0.018
Piet Termonia	University of Ghent	UGhent EURO-CORDEX climate runs	Environmental Modeling	3,600	34
De Wispelaere Kristof	University of Ghent	Dynamical first principle modelling of zeolite dealumination in H-SSZ-13	Computational Chemistry	3,624	0.156
Cnudde Pieter	University of Ghent	DFT study of reaction paths in zeolite-catalyzed 2-hexene cracking	Computational Chemistry	4,536	1
Xia Chun	KU Leuven	Solar prominence eruption and the blending of solar wind and magnetosphere	Environmental Modeling	5,000	4
Vets Charlotte	University of Antwerp	Chirality controlled growth of carbon nanotubes on bi-metallic catalysts: the role of thermodynamics	Computational Chemistry	4,176	1
De Vos Arthur	University of Ghent	Defect engineering in UiO-66: How linker defects affect the electronic structure.	Computational Chemistry	4,116	6.468
Reyniers Pieter	University of Ghent	Computational Fluid Dynamics based run length simulations of enhanced 3D reactors (II)	Mechanical Engineering	4,000	2
Meyers Johan	KU Leuven	Controller design for wind-farm boundary layers	Mechanical Engineering	4,960	1.12
Meyers, Johan	KU Leuven	Simulation of fully finite wind-farm boundary layers	Mechanical Engineering	5,000	1.66
Lambrechts Diether	KU Leuven	Optimization of long-read sequencing mapping to discover biomarkers for cancer immunotherapy.	Bioinformatics	4,500	10
Vandenbrande Steven	University of Ghent	Insight into the stacking of 2D COFs from ab initio and force-field calculations	Computational Chemistry	2,200	0.065
Moors Samuel	VUB	Model-guided design of novel cyclic peptidomimetics as binder of the humanbβ2 adrenergic rexeptor	Computational Chemistry	1,892	0.115

3 October 2016

Applicant	Institution	Title	Department	Nodedays awarded	Diskspace awarded (TB)
Wilfried de Corte	University of Ghent	Robustness and Sensitivity of Pareto-optimal Selection Designs under Various Types of Candidate Withdrawal	Department of Data Analysis, Faculty of Psychology and Educational Sciences	3255	0.025
Giovanni Lapenta and Jorge Amaya	KU Leuven	xPic3D : Next generation Particle-in-Cell algorithms for exascale HPC architectures	Centrum voor Plasma-Atrofysica, Department Wiskunde	1900	1
Kurt Lejaeghere, Stefaan Cottenier and Veronique Van Speybroeck	University of Ghent	Assessing the accuracy of an efficient meta-GGA functional for property predictions of elemental solids	EA17 / Center for Molecular Modeling	640	0.2
Michael Sluydts, Veronique Van Speybroeck and Stefaan Cottenier	University of Ghent	High-throughput screening of an unknown quaternary crystal space	Center for Molecular Modeling	4056	0.86
Nasrin Sarmadian and Dirk Lamoen	University of Antwerp	The role of humidity in the CO gas sensing mechanism of the clean and impurity modified ZnO surfaces	EMAT / Departement Fysica	2374	0.204
Ehsan Moravveji and Conny Aerts	KU Leuven	Asteroseismic Grid for Massive Pulsating and Rotating Stars	Institute of Astronomy	3719	11
Krisztina Feher	University of Ghent	Free energy of binding calculations on the complexes formed by peptidoglycan based immunestimulators and the Macrophage Mannose Receptor	NMRstr group, Department of Organic and Macromolecular Chemistry	100	0.18
Dimitrios Millas, Bart Ripperda and Rony Keppens	KU Leuven	Modeling outflows and particle evolution in relativistic astrophysics, solar eruptive flares, and solar wind-magnetosphere interaction	CmPA, Department of Mathematics	5000	6
Geoffrey Pourtois	IMEC	Modeling of the impact of non-idealities on the performances of two-dimensional low power nanoelectronic devices	Modeling Simulation & Physics (MSP)/ NCAIS/UPM/PT	3554	0.16
Ruben Demuynck, Veronique Van Speybroeck and Jelle Wieme	University of Ghent	Construction of ab initio free energy profile for metal-organic frameworks	Center for Molecular Modeling	4200	0.14
Pieter Reyniers, David Van Cauwenberge, Yu Zhang, Laurien Vandewalle and Jens Dedeyne	University of Ghent	Computational Fluid Dynamics based run length simulations of enhanced 3D reactors (III)	Department of Chemical Engineering and Technical Chemistry (EA12) Laboratory for Chemical Technology	4000	1
Julianna Hajek, Kristof De Wispelaere, Chiara Caratelli and Veronique Van Speybroeck	University of Ghent	Ab initio umbrella sampling simulations of dehydration of UiO-66	EA17, Center for Molecular Modeling	3972	0.05
Jelle Wieme, Steven Vandenbrande and Veronique Van Speybroeck	University of Ghent	Computational exploration of the free energy profile of guest-free M(bdp) (M=Co,Fe) (bdp2-=1,4-benzenedipyrazolate)	EA17, Center for Molecular Modeling	2940	0.52

Simon Bailleul, Pieter Cnudde and Veronique Van Speybroeck	University of Ghent	Benchmark study of ab initio molecular dynamics simulations for the methylation of HMB	EA17, Center for Molecular Modeling	4812	0.05
Wim Munters and Johan Meyers	KU Leuven	Optimal control of Lillgrund wind farm	Turbulent Flow Simulation & Optimization (TFSO) Research Group Mechanical Engineering Department	4300	2.31
Dries Allaerts and Johan Meyers	KU Leuven	Simulation of the GABLS3 diurnal cycle	Turbulent Flow Simulation & Optimization (TFSO) Research Group Mechanical Engineering Department	4320	1.05
Piet Termonia, Steven Caluwaerts, Pieter De Meutter, Michiel Vanginderachter, Rozemien De Troch, Rafiq Hamdi and Lesley De Cruz	University of Ghent	CORDEX.be high-resolution climate runs with ALARO-0	Department Physics and Astronomy	1400	16
Sam De Waele and Stefaan Cottenier	University of Ghent	Error estimates for ab initio predictions for point defects in the Fe-N system	Center for Molecular Modeling (CMM)	3022	2.902
Jayson Gutiérrez Betancur and Steven Maere	University of Ghent VIB	Modelling gene regulatory network evolution through gene duplication from mechanistic perspective	Evolutionary Systems Biology lab, Dept. of Plant Biotechnology and Bioinformatics (Ghent University) / Plant Systems Biology Dept. (VIB)	3240	0.18
An Ghysels and Oriana De Vos	University of Ghent	Molecular dynamics simulations of oxygen transport through membranes	Center for Molecular Modeling	700	0.54
Rolando Saniz and Dirk Lamoen	Universiteit Antwerpen	Finite size corrections to charged defect formation energies	Condensed Matter Theory / Departement Fysica	2449	0.018

Graphs on the use of Tier-1

CPU core hours used on Tier-1 muk (UGhent)

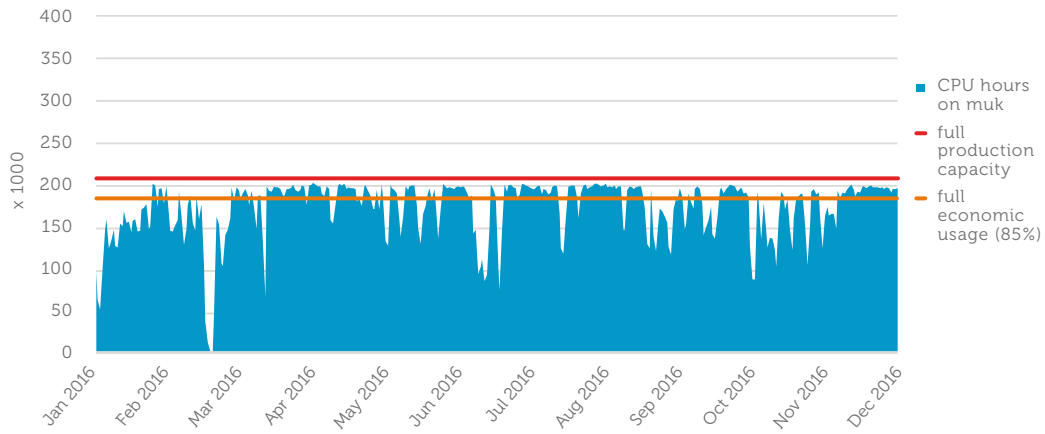

Figure 3
CPU core hours used on Tier-1 muk (UGhent)

CPU core hours used on Tier-1 BrENIAC (KU Leuven)

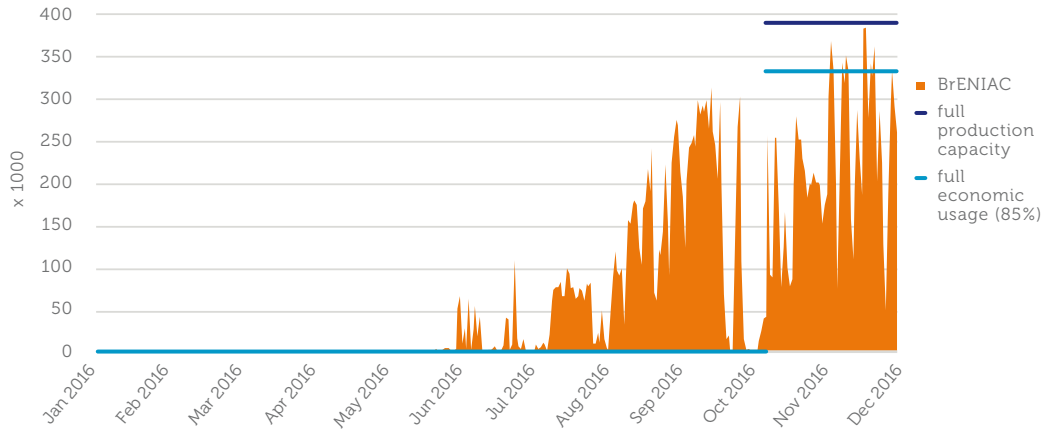

Figure 4
CPU core hours used on Tier-1 BrENIAC (KU Leuven)

The above graphs provide a picture of the usage history of both Tier-1, muk and BrENIAC in 2016. In both graphs, the theoretical maximum capacity (full production capacity) is indicated, but 85% of this figure is considered as 'fully economic' use of the cluster.

Tier-1 muk was very well used in 2016, with an average utilisation rate of 85% and peaks up to 99%. Variations in use are still visible, depending on the computing efforts of the users, but there were no low temporary utilisation rates. The dip in the period 22 to 26 February 2016 reflects a final series of necessary maintenance operations which required a short downtime.

The pilot phase of Tier-1b ran from July to the end of October. Upon completion, the system was emptied. The production period started with the projects that were allocated in the 3rd allocation round of 2016. At that time both Tier-1 clusters were active. After two weeks the system already reached a utilisation rate of over 80%. It demonstrates that all researchers got off to a good start on the system, the software stacks were available following the tests in the pilot phase and no technical problems occurred.

Tier-2 infrastructure

Available Tier-2 infrastructure

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

In March 2016, the *ThinKing* thin node cluster was expanded with 96 nodes in the Hasswell section. This increased the capacity by 90 TFlops.
In 2016, work was started on the renovation of the shared GPFS scratch infrastructure. This took place in a number of phases. The purpose was to:

- Replace the DDN 10K system, which had reached the end of its life cycle
- Increase the capacity to support also the operation of a future cluster
- Decouple the archive system from the GPFS cluster file system to achieve greater flexibility.

This should, in future, allow the archive system to remain available when working on the GPFS scratch file system and also provides more options for making the archive system available on other computers.

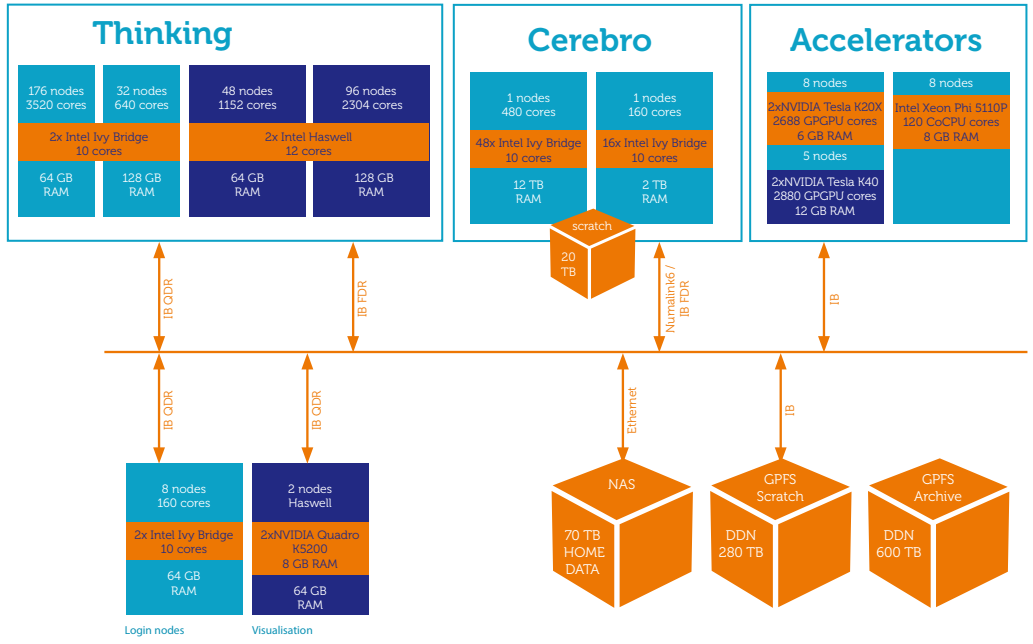

Figure 5 - Tier-2 infrastructure KU Leuven - UHasselt

University of Ghent

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

- 5 clusters
- 226 TF
- 11,328 CPU cores
- 50 TB memory

The Tier-2 infrastructure is built up of various clusters, serving specific characteristics. In the course of 2016, cluster swalot (purchased from Dell) - a new MPI cluster to facilitate multi-node jobs - went into production. A new shared storage (for 3 PB) will be taken into production in the course of 2017. In addition, the Tier-2 infrastructure will be expanded with a new MPI cluster for the purpose of replacing the standard *delcatty* cluster and *raichu* cluster.

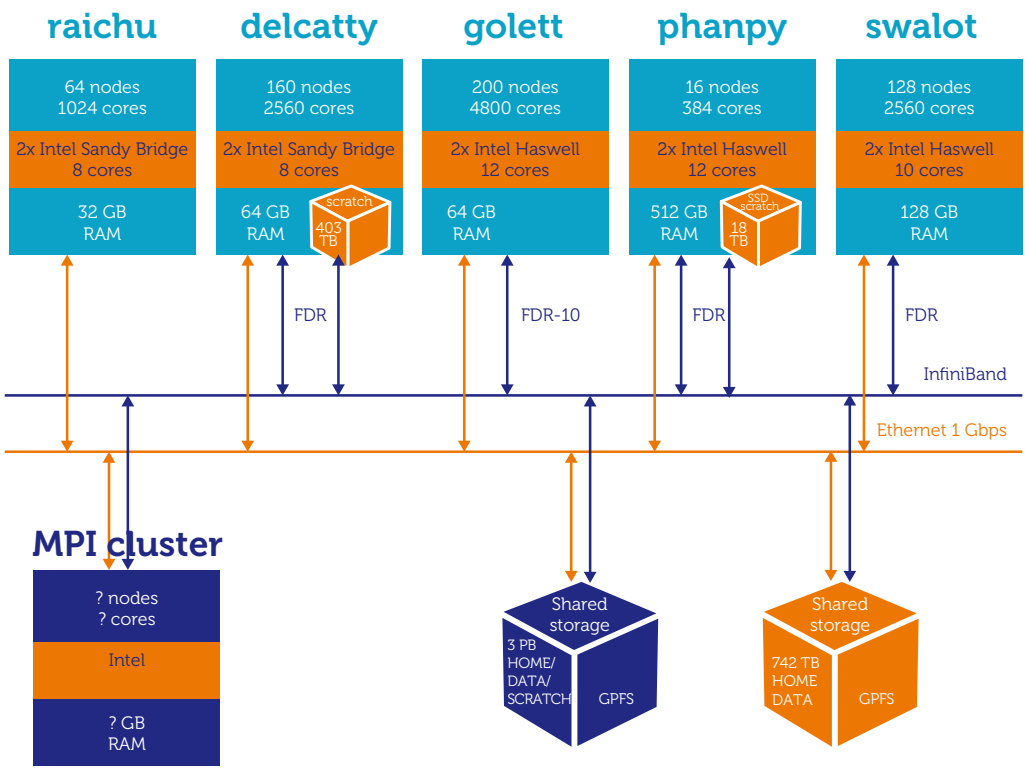

Figure 6 - Tier-2 infrastructure UGhent

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 (Turing and Hopper), divided into 4 partitions
- 90 TF
- 4992 CPU cores
- 12 TB memory

In 2016 a tender was issued to replace the oldest of the two clusters, Turing. Leibniz, the new cluster, will be put into use in the course of 2017.

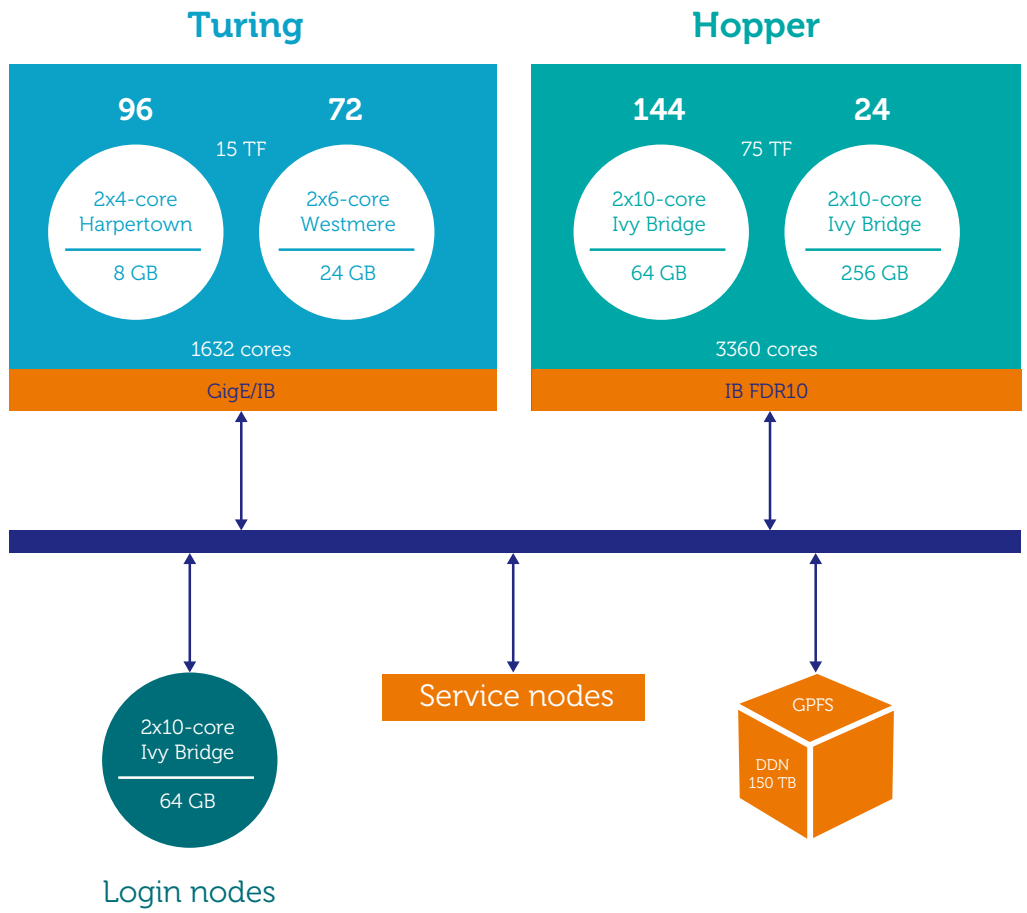

Figure 7 - Tier-2 infrastructure University of Antwerp

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. In the course of 2016, the old AMD Shanghai nodes were taken out of production, and 28 new Intel Broadwall nodes were added, including 1 high memory node with 1.5 Tb memory and 4 Tb storage space, in order to meet user demand.

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.

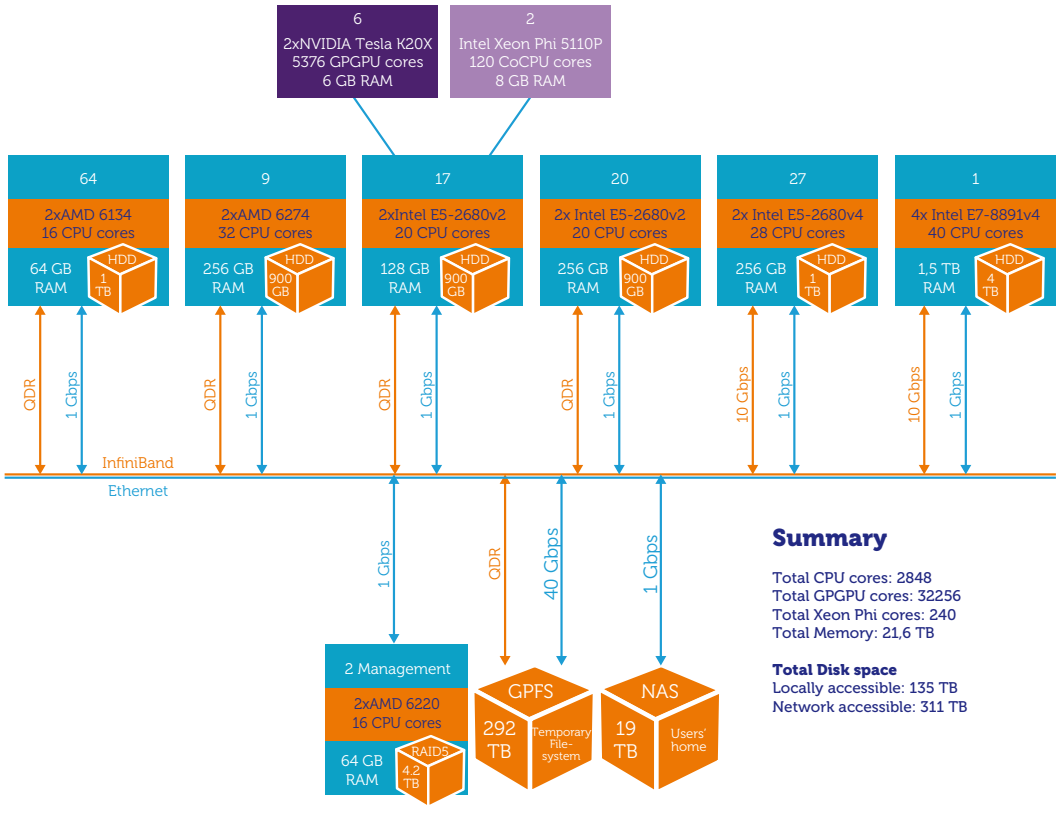

Figure 8 - Tier-2 infrastructure VUB

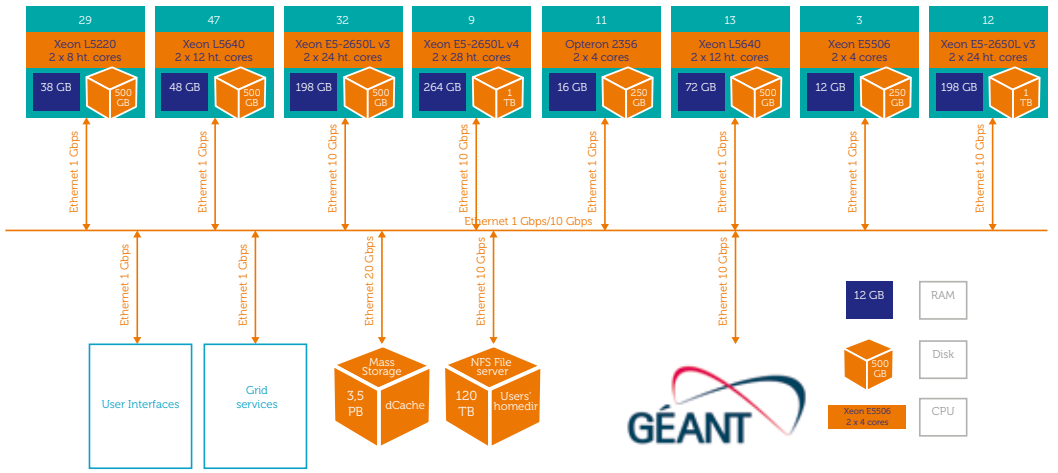

Figure 9 - Grid infrastructure VUB

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

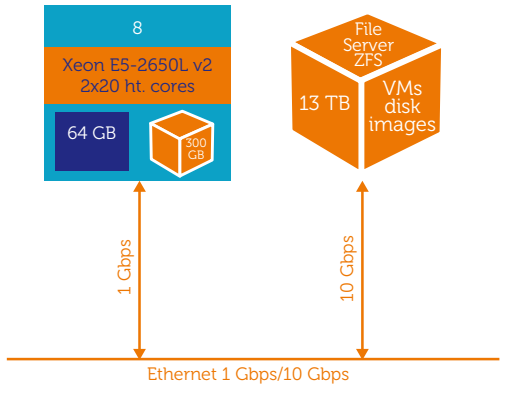

Figure 10 - VUB Cloud infrastructure

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.

CPU core hours: All VSC clusters

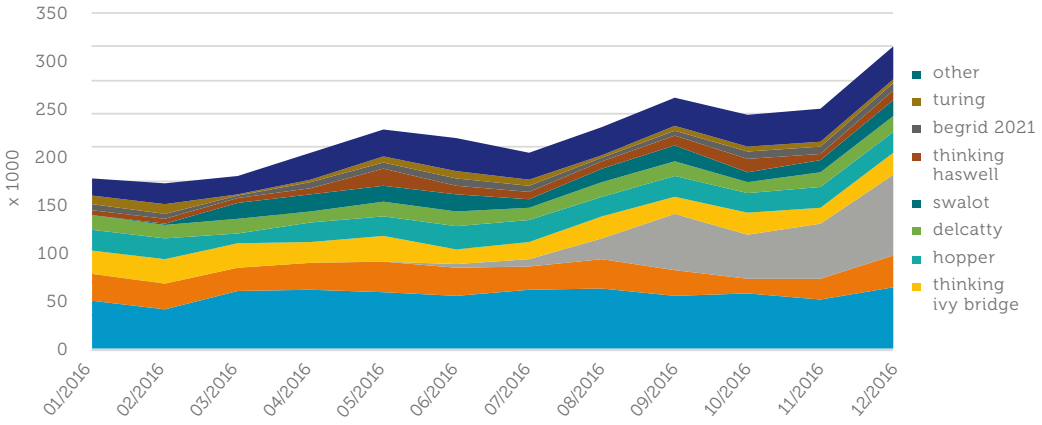

Figure 11
Use expressed in core hours of all VSC clusters

Figure 11 shows the cumulative use of the Tier-1 (muk and BrENIAC) and Tier-2 infrastructure. There is a noticeable increase between the start and the end of the year. At the end of 2016, 78% more computing time was used. This was due mainly to the use of the new Tier-1 BrENIAC. Next year the old Tier-1 muk will no longer be included in the reports, but the investments at Tier-2 level will ensure that the total number of core hours will continue to rise.

CPU hours used on thin node clusters KU Leuven - UHasselt

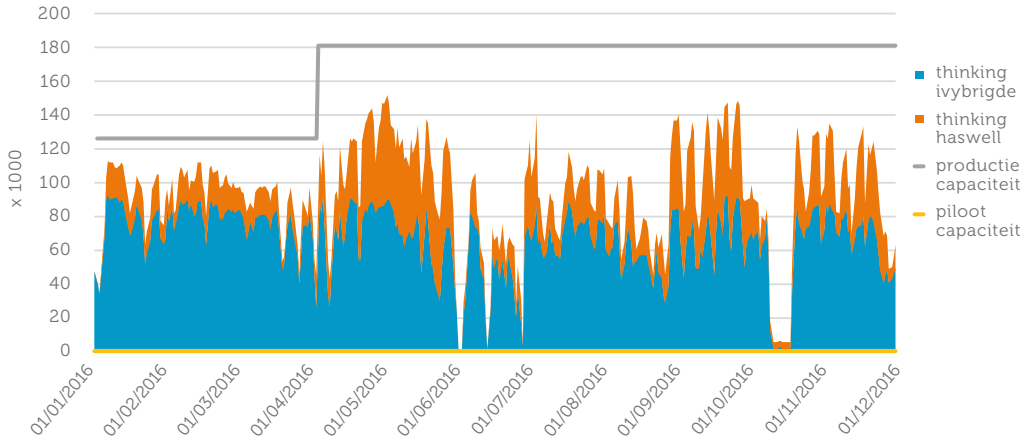

Figure 12
Core hours used on the KU Leuven/U Hasselt thin node cluster

Figure 12 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 cluster consists of a partition with Ivy Bridge processors and a partition with Haswell processors which were extended in 2016. This is clearly reflected in the increase in computing time used in May. From end June, operations were carried out to replace the shared parallel file system. A first planned downtime took place in June. However, due to unforeseen compatibility problems, the downtime took longer and the modifications could not be completed. A solution was found together with the supplier. The final commissioning was prepared between July and October. To this end, a rack in the cluster was reserved for testing and data migrations. In mid-November, the file system was completely renewed during a longer downtime. In this operation, data were retrieved from the existing file system, which required a longer period of time.

Core hours used on shared memory cluster KU Leuven - UHasselt

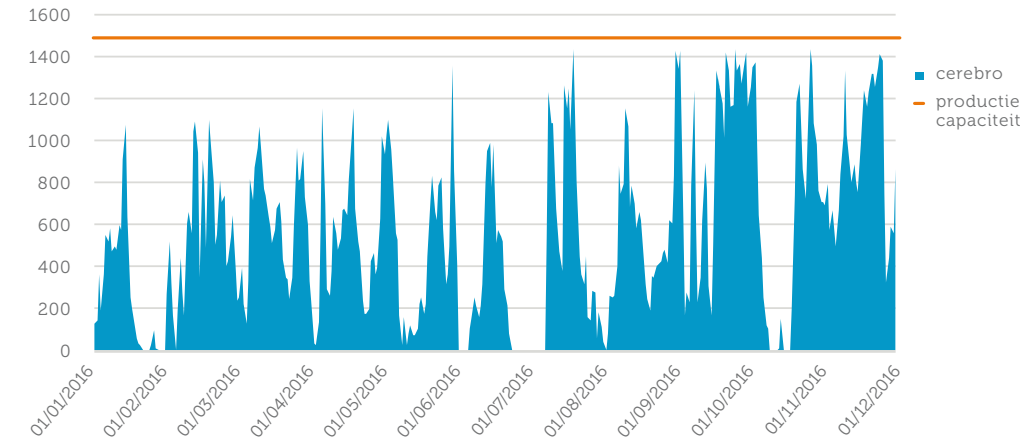

Figure 13
Use of shared memory

Core hours used on GPU cluster KU Leuven - UHasselt

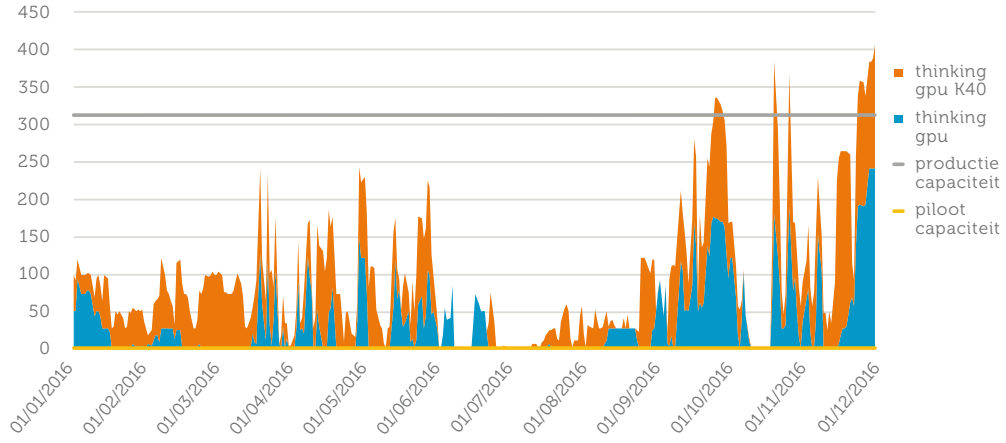

Figure 14
Use of accelerators

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. This is because frequently the complete node is used because of the memory needed or because of the accelerators; and this without all CPU cores in the machine necessarily being used.

The shared memory machine is used less intensively than the thin node cluster, but also regularly shows good utilisation. It is used mainly in the fields of chemistry and bioinformatics. The latter field in particular made intensive use of the shared memory machine in 2016.

CPU hours used on computing infrastructure UGhent

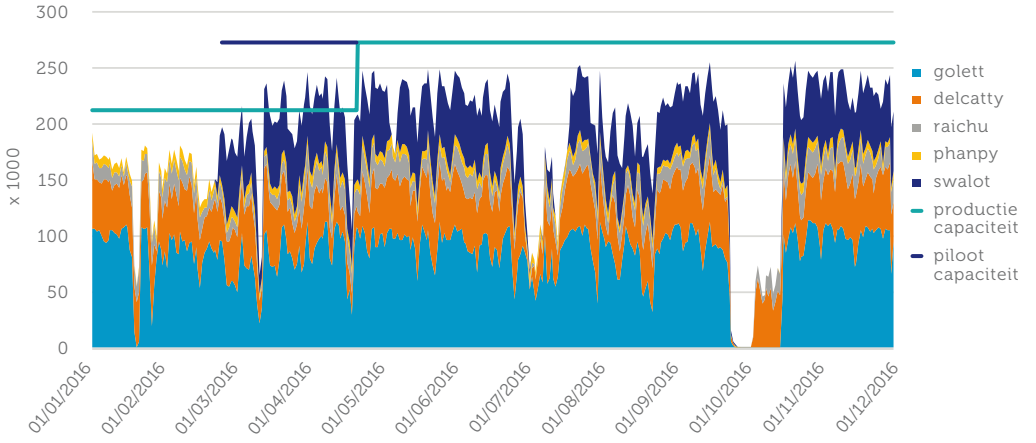

Figure 15
Use of the University Ghent infrastructure

Figure 15 shows the cumulative use of all computing clusters the University of Ghent in 2016.

The commissioning of the new swalot cluster led to an increase in production capacity of 29%. Following a two-month pilot phase, in which the machine was extensively tested under high load from a number of power-users, the new machine was made available to all users. The users adapted almost immediately to the new situation and initiated several challenging computational projects: the average use of the clusters remained virtually unchanged around 77% in 2016.

Unfortunately, however, the central SCRATCH storage, which is shared by all Tier-2 compute nodes and is crucial for e.g. MPI computing tasks, was unavailable for two weeks in the period October-November. As a result, the computing clusters could not be used either (especially for MPI tasks), although the delcatty and raichu clusters had been released exclusively for single-node tasks already after one week. Several factors contributed to this 'perfect storm' unavailability:

- almost simultaneous failure of 12 discs in the SCRATCH storage system causing unavailability of the file system, with risk of data corruption;
- errors in the management tools of the parallel GPFS file system causing serious delays, as these errors first had to be patched by the software developer;
- on 27 October, the 'Dirty Cow' security flaw in Linux required the HPC-UGhent team to give first priority to updating and patching the login and worker nodes;
- a brief power outage on the same day shuts down all worker nodes, and the team again loses valuable time to restart and repair all these nodes;
- hardware failure of a storage controller in early November again delaying the final release of the SCRATCH partition, yet without comprising the integrity of the file system.

In the end, data loss was avoided due to the efforts of the HPC-UGhent team.

Some minor dips in the above graph indicate downtimes resulting from power failures or scheduled maintenance work. Earlier this year we introduced the regular use of rolling updates whereby necessary software updates are gradually applied to the nodes of the clusters, as soon as they become available. This reduces unavailability to the user.

CPU hours used on thin node clusters UAntwerp

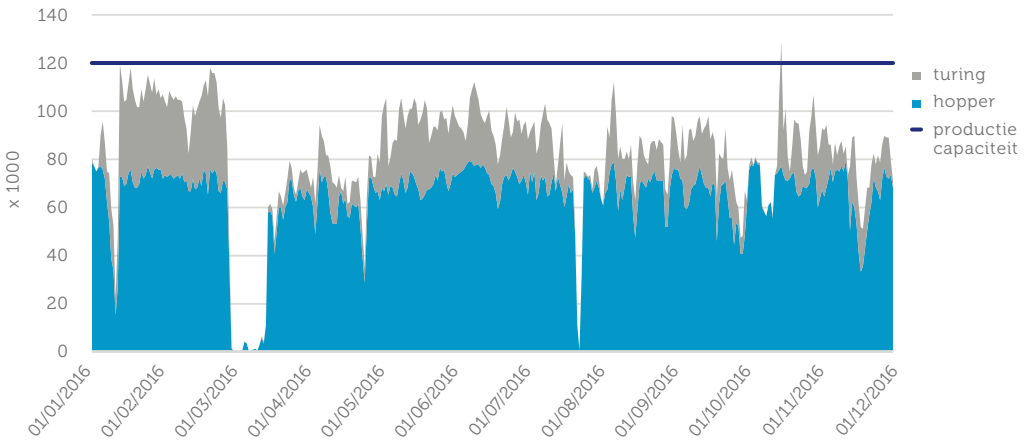

Figure 16
Use of infrastructure UAntwerp

The graph of figure 16 shows the cumulative use of the Turing and Hopper clusters at the University of Antwerp. A number of discontinuities can be seen on the graph.
These include works that were carried out on the data center infrastructure in March and on the on-campus electrical infrastructure in August. There were also a few short interruptions due to power outages in January and May. At the end of October there was an interruption because of the "Dirty Cow" security patch. For Turing this patch was available only later, so this cluster was also slightly longer out of service.
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.

CPU hours used on thin node clusters VUB

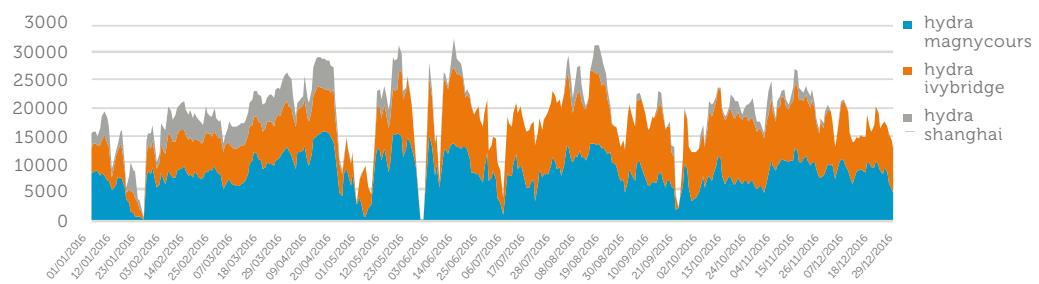

Figure 17
Use of the VUB infrastructure

In 2016, only the Magnycours and Ivy Bridge nodes were online, the new Intel Broadwell nodes were in a pilot phase because the network connectivity had to be optimised. There were two planned, short downtimes, for the migration to the new storage (May 2016) and for a power supply switchover (from 5.5 kV to 11 kV, in September 2016). There were also two short unplanned downtimes, due to a technical failure of the cooling system in the data center (January 2016) and a problem with the storage (June 2016).

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. 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.

CPU hours used on the BEGrid clusters VUB

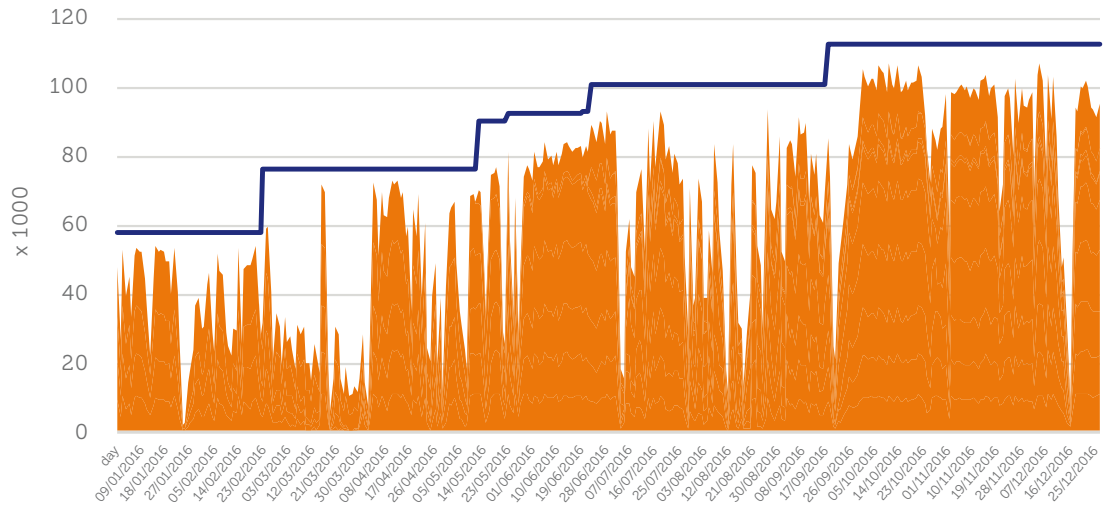

Figure 18
Use of BEGrid cluster at VUB (blue line: production capacity)

BEgrid 19, BEgrid 20-21 and BEgrid 22 are the three parts of the BEgrid cluster at the VUB. In the course of 2016, additional worker modes were added in several phases and the standalone cluster for the IceCube project was integrated into the BEgrid cluster. There were a few planned downtimes, including for planned electricity works (January and July 2016), as well as a few unplanned downtimes, also a result of problems with NFS (December 2016).

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 University of Ghent 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 concise summary is given below.

At UAntwerp, computation was performed by two UGhent researchers: almost 3000 jobs in 73 node days.
At VUB, computation was performed by one UAntwerp researcher: 11 jobs in 8 node days.
At UGhent, computation was performed by various researchers from other institutions: one UAntwerp user (55,928 core hours), 4 VUB users (8,672 core hours) and 14 KULeuven users (180,165 core hours). This computing time (a total of 244,765 core hours) was provided free of charge.
At KU Leuven, computation was performed by 1 VUB researcher on the GPU nodes for 4 node hours, 3 UAntwerp researchers on Thinking and Cerebro for 3.2 node days and by UGhent researchers spread over the systems as follows: 1 researcher on Cerebro for 411 node days, 8 researchers on ThinKing for 33.8 node days, and 4 researchers on the GPU systems for 58.7 node days.
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:

UGhent	23.19 %
UAntwerp	13.04 %
VUB	11.63 %

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

The majority of researchers use the so-called "glide in" mechanism. This involves the submission of pilot jobs, which, once they are active on a worker node, fetch the payload somewhere else. For the calculation of the above percentages, only the computing time of the pilot jobs, and not the individual payloads, has been taken into account. Moreover, in the event of grid, a user workflow may be spread over different sites in different countries, so that the above percentages reflect only a portion of the actually used computing time.

Projects

In this section, we present an overview of the projects worked on in 2016.

Website and marketing

On the website, a large part of the user portal section was reviewed and improved where appropriate. In addition, suggestions from an external review were further implemented. An analysis of the website's statistics shows that the most popular page in the Dutch section is "What is a supercomputer", apart from the contact details. In the English section, the list is topped by the user portal, the system status page, and the training courses.

In the marketing area, particular attention was paid to the launch of BrENIAC, the second supercomputer hosted at the KU Leuven. All relevant information about BrENIAC was also put on the website.

Code of conduct

Since all VSC sites are interconnected and cross-site computation is available to the users, a line of conduct has to be developed for these users. This code of conduct is supplementary to the rules that already apply at the respective institutions. The document sets out the purpose and scope, unauthorised use, responsibilities of VSC users, supervision and monitoring, and sanctions in case of non-compliance. This code of conduct will be finalised in the course of 2017.

Access for non-academic users

Non-academic users (e.g. in industry, public authorities, etc.) increasingly use the VSC computing infrastructure. In this project, the access procedure for such users was optimised. A clear and simple procedure now allows users to request a VSC user-id and gain access to the VSC infrastructure in just a few clicks. Following a few modifications to the VSC account page, the results of this project went into production in September 2016.

User support

User support comprises several components:

- answering questions from users (helpdesk);
- meetings with users / specific support;
- training and outreach.

The 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 University of Ghent (for the first Tier-1) or 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.

	KU Leuven / University of Hasselt	UGhent	UAntwerp	VUB
Tier-2 + grid				
Accounts	722	141	46	121
Software	596	42	49	159
Other	741	634	300	517
Tier-1				
Accounts	36	3		
Software	37	1		
Other	59	124		
Total	**2191**	**945**	**395**	**797**

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 to 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 2016, for example, there were workshops for a Biomechanics group and a group within chemical engineering techniques.

In 2016 specific actions were taken for various groups in the broad research area of bioinformatics. A new staff member was hired to undertake these actions. One of these groups is the Genomics Core of UZ Leuven. It provides technological and intellectual support of new genomic technologies in support of several research groups. In 2016 the group organised 2 workshops to better introduce HPC to a large group of users. This workshop was supported by the VSC staff. For the genome analyses we assisted in setting up the software pipelines on the cluster. The sequencing equipment is advancing rapidly, so that more computing time is required for analyses. The use of the VSC HPC computing power should make it possible to sustain this growth. The support continues in 2017.

Furthermore, specific software environments, GenePattern and Jupytherhub for use by different bioinformaticians, were set up. This allows bigger computational tasks to be performed on the cluster via a web interface used within the particular field. This contributes to making large computer power better accessible to these user groups.

The Plasma Astrophysics Working Group is involved in the Virtual Space Weather Modelling Centre (VSWMC). As part of a European project, work is in progress on an integrated system to combine results of different models in order to obtain better predictions. This implied, on the one hand, assistance with the installation and configuration of the software allowing for this integration and, on the other hand, also with installation of new prediction models (Euphoria) used within this framework.

Actions were also undertaken within human sciences. For the ECOOM Research Centre, specific support was provided for the performance of Hadoop and Spark analyses on the cluster. The Quantitative Lexicology and Variation Linguistics (QLVL) working group has a software pipeline for performing analyses on large databases with words. We assisted in transferring this pipeline to the cluster.

For the University of Hasselt, the emphasis was on the attraction of new user groups. Efforts on behalf of the Faculty of Business Economics were continued and paid off. The Logistics research group was the second largest user group in terms of computing time in 2016. The Faculty of Engineering Sciences also received the necessary attention, and now has two research groups using the VSC infrastructure.

For the CenStat group a new software tool was developed that allows commonly used workflows to run efficiently on the own infrastructure, but also on Amazon EC2.

University of Ghent

Several specific user meetings were organised with the participation of researchers from different research groups of the University of Ghent:

- Marinazzo Lab, Department of Data Analysis, Faculty of Psychology and Educational Sciences
- Cancer Research Institute Ghent
- Department of Reproduction, Obstetrics and Herd Health, Faculty of Veterinary Medicine
- Center for Molecular Modeling, Faculty of Sciences & Engineering and Architecture
- Laboratory of Functional Plant Biology, Faculty of Sciences
- Data Science Lab, Faculty of Engineering and Architecture
- Research Group Evolutionary Morphology of Vertebrates, Faculty of Sciences
- Data Mining and Modelling for Biomedicine (Dambi) research group, Faculty of Medicine and Health Sciences & VIB
- Theoretical Nuclear and Statistical Physics Group, Faculty of Sciences
- Evolutionary Systems Biology lab, department of Plant Systems Biology, Faculty of Sciences & VIB
- Department of Flow, Heat and Combustion Mechanics, Faculty of Engineering and Architecture
- Department of Applied Bioscience, Faculty of Bioscience Engineering
- Department of Animal Production, Faculty of Bioscience Engineering

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

- 10 May 2016, 30 third-year students industrial engineering electronics-ITC, as part of the advanced computer systems course.
- 10 October 2016, guided tour for Zeus working group UGhent

In addition, user meetings were organised for (potential) users from outside the University of Ghent:

- Centre of Medical Genetics Ghent, UZ-Gent
- Gynaecological Clinic, UZGent
- Multimedia Technologies department, iMinds
- KMI
- ILVO

On 12 & 13 September 2016, the HPC-UGhent team co-organised the BeneLearn 2016 conference, the annual machine learning conference of Belgium and the Netherlands, for an interuniversity and international audience. This event took place within the framework of the Big Data specialisation of UGhent within VSC, as the focus at this conference was very much on 'large-scale data analysis'. There were over 100 participants, including a significant number from about 15 companies.

University of Antwerp

On the one hand, we deal with questions from existing users and try to organise the computational work as optimally as possible, and on the other hand, we try to make the VSC known to other researchers to convince them of 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 2016 we put researchers from Applied Engineering Sciences on their way towards HPC. They are now frequent users of the infrastructure. In addition, other research groups from departments that are already active on the infrastructure, such as Biology, Physics, Medicine and Health Sciences, and Mathematics, are now regular users. The number of users thus continues to grow in both width and depth. Some examples of specific support:

- Following on the Tips & Tricks sessions on "code modernisation", we gave advice on the further development of software. In a second case, we made a complete analysis of the code with suggestions for improvements and the expected gains. With another researcher we discussed how to find optimal solutions to specific linear problems.
- Development of a job monitor to track the performance of jobs. This allows us to provide users with target tips for more efficient use of the machine and to examine how we can solve frequently occurring problems or to identify programs that need extra attention.
- Preliminary work to verify the feasibility of a joint project with a company.
- Cooperation in the context of a master's thesis: visualization of results generated with FINE-Marine by means of TurboVNC
- Cooperation in the context of a master's thesis: shared memory parallelisation with OpenMP and parallel calling of an external program.
- 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".

Staff

Since the infrastructure of the VSC (Tier-2 and Tier-1 infrastructure) is installed in the various university data centres, 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, 15 FTE are subsidised. On the other hand, 2 additional FTE are subsidised at the institution that accommodates the Tier-1 supercomputer. To this end, 2 FTE were allocated for the operation of the first Tier-1 and 2 FTE for the second Tier-1, in 2016.

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
University of Ghent	4
University of Antwerp	3
VUB	2
University of Hasselt	1
KU Leuven	5
Total	**15**

Institution	Number of subsidised FTEs for Tier-1 operation
University of Ghent	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 30.35 FTE for HPC operation and support. One person assigned to the SESAME Net project works under direct supervision of the FWO. 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 VTE deployed for HPC	Number of employees involved with HPC operation and support
University of Ghent	10.5	16
University of Antwerp	4.35	8
VUB	3.75	8
University of Hasselt	1	1
KU Leuven	9.75	16
Total	**29.35**	**49**

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 centre. They install the infrastructure in the data centre. 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 centre.

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 part of their time promoting scientific computing and HPC. They do this by giving info sessions within the academic context, but also by visiting 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
University of Ghent	5	7	7	6	3	3
University of Antwerp	1	3	5	3	2	1
VUB	2	4	4	2	0	1
University of Hasselt			1	1	1	1
KU Leuven	2	7	5	5	7	3
FWO			1	1		1
Total	**10**	**21**	**23**	**18**	**14**	**10**

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
4 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
7 persons with academic experience in this field

Bioinformatics
1 person with academic experience in this field

Staff List

Name	Institution	% employed for HPC operations and support
Stefan Becuwe	University of Antwerp	100%
Franky Backeljauw	University of Antwerp	100%
Bert Tijskens	University of Antwerp	100%
Kurt Lust	University of Antwerp	100%
Koen Decauwsemaecker	University of Antwerp	10%
Muriel Dejonghe	University of Antwerp	10%
Herwig Kersschot	University of Antwerp	10%
Annie Cuyt	University of Antwerp	5%
Leen Van Rentergem	KU Leuven	20%
Herman Moons	KU Leuven	10%
Ingrid Barcena	KU Leuven	100%
Jan Ooghe	KU Leuven	100%
Martijn Oldenhof	KU Leuven	100%
Mag Selwa	KU Leuven	100%
Alexander Vapirev	KU Leuven	70%
Jo Vanvoorden	KU Leuven	100%
Jo Vandeginste	KU Leuven	100%
Tom Leuse	KU Leuven	100%
Yorick Poels	KU Leuven	15%
Rudy Rys	KU Leuven	60%
Tom Vanmierlo	KU Leuven	60%
Sofie Pieraerd	KU Leuven	10%
Tom Vanhout	KU Leuven	20%
Philip Brusten	KU Leuven	10%
Wouter Depypere	University of Ghent	100%
Stijn De Weirdt	University of Ghent	100%
Alvaro Simon Garcia	University of Ghent	100%
Andy Georges	University of Ghent	100%
Ewan Higgs	University of Ghent	100%
Kenneth Hoste	University of Ghent	100%
Ewald Pauwels	University of Ghent	100%
Jens Timmerman	University of Ghent	100%
Kenneth Waegeman	University of Ghent	100%
Danny Schellemans	University of Ghent	10%
Johan Van Camp	University of Ghent	30%
Luk Claes	University of Ghent	20%
Wim Waeyaert	University of Ghent	20%
Werend Brantegem	University of Ghent	10%
Bruno Cardon	University of Ghent	10%
Dieter Roefs	University of Ghent	50%
Ward Poelmans	FWO	100%
Geert Jan Bex	University of Hasselt	100%
Stéphane Gérard	VUB	100%
Balázs Hagató	VUB	100%
Rosette Vandenbroucke	VUB	50%
Olivier Devroede	VUB	25%
Johan D'Hondt	VUB	20%
Peter Van Rossem	VUB	20%
Dirk Heyvaert	VUB	30%
Philippe Leemans	VUB	30%

Outreach to Flemish companies

One of the missions entrusted to the VSC is informing companies and institutions from the non-profit sector about the added value of HPC in the development and optimisation of services and products. Another mission is to support them with the use of this new technology.

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 2016, two Flanders-based companies made use of the Tier-1 in Ghent. The agreements for use were contractually concluded between FWO, the University of Ghent and the company involved. Four other companies used the Tier-1 in an explorative context. These companies are active in the sectors of renewable energy, pharmaceuticals, medical technology, electrical engineering and material technology. The figure below provides an (anomised) view of the use of computing time on Tier-1 muk by the industry in 2016. From mid-2016, industrial activity virtually came to a halt, as some companies were already switching to Tier-1 BrENIAC, which had in the meantime become available in the pilot phase, whereas others moved their computing work to Tier-2 clusters within the VSC network which had become more powerful than Tier-1 muk.

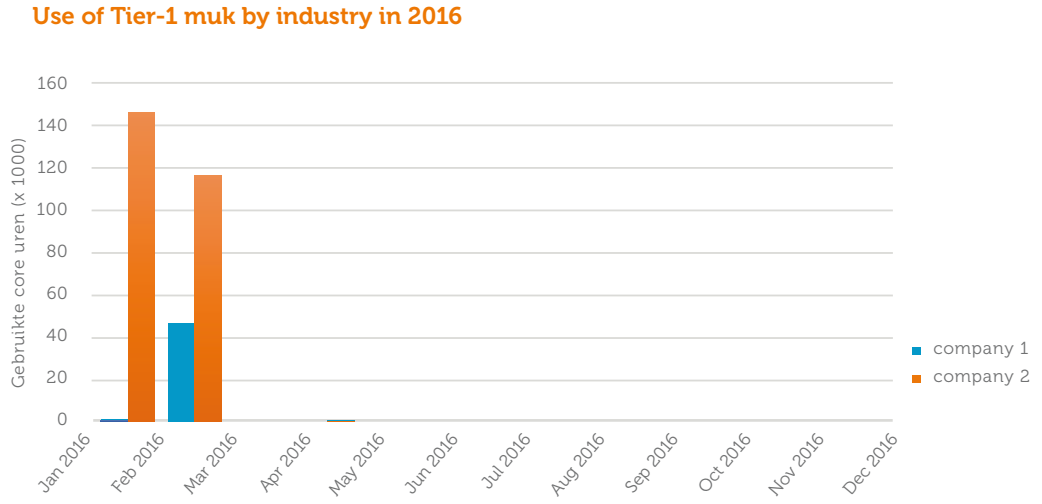

Figure 19
Use of Tier-1 muk by industry in 2016

Outreach to companies and other knowledge institutions

A new edition of the VSC Industry Day was prepared and scheduled for early April 2016. However, in the aftermath of the bomb attacks in Brussels on 23rd March, several keynote speakers had to cancel their travel plans. It was therefore decided to postpone the event and to consider the inauguration of the new Tier-1 BrENIAC as the optimal moment to bring supercomputing and the VSC's activities under the attention of Flemish industry on a larger scale.
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:

- 3E, 8 January and 20 April 2016 (UGhent)
- FEOPS, 16 March 2016 (UGhent)
- Feasibility study for spinoff, 25 March and 27 May 2016 (UAntwerp)
- IMDC, 13 May 2016 (UAntwerp)
- UMICORE, 16 September 2016 (UGhent)
- Outreach to IOF business developers @ UGhent, 26 January 2016 (UGhent)
- Diabatix, 13 April 2016 (KU Leuven)

VSC was also actively present at various networking events for the promotion of supercomputing (see p.46 for an overview of all participations). Furthermore, several industry use cases were prepared and posted on the VSC website, and disseminated through the SESAME Net channel, amongst others.

Training

Organisation of training programmes

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 when one does not 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 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 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.

A summary of the planned training courses is also included in each issue of VSC Echo, and a thematic issue on training is published once a year.

Training programme

An overview of the available training, listed per month, is given below:

February

- Python introduction, part I, 8 & 9 February, 8 hours, KU Leuven/University of Hasselt
  level: intermediate
- Linux introduction, 9 February, 7 hours, KU Leuven/University of Hasselt
  level: introductory
- Matlab, part 1: intermediate, 11, 12, 18 & 19 February, 16 hours, KU Leuven
  level: introductory
- HPC@KU Leuven introduction, 12 February, 6 hours, KU Leuven/University of Hasselt
  level: introductory
- C introduction, 16, 17, 23 & 24 February, 12 hours, KU Leuven/University of Hasselt
- HPC@KU Leuven introduction, 19 February, 6 hours, KU Leuven/University of Hasselt
  level: introductory
- Python introduction, part II, 22 & 23 February, 8 hours, KU Leuven/University of Hasselt
  level: intermediate
- HPC@UGhent introduction, 23 February, 6 hours, University of Ghent
  level: introductory
- supercomputers for starters, 25 February, 3.5 hours, University of Antwerp
  level: introductory
- Matlab part 2: intermediate, 25 & 26 February, 8 hours, KU Leuven
  level: introductory

March

- HPC@UAntwerp introduction, 3 March, 4 hours, University of Antwerp
  level: introductory

April

- Introduction to Linux, 13 April, 6 hours, University of Ghent
  level: introductory
- Message Passing Interface (MPI), 20 April, 7 hours, University of Ghent
  trainer: prof. Dr. Jan Fostier (UGhent)
  organisation: VSC, doctoral schools
  level: intermediate
- HPC@UGhent introduction, 26 April, 6 hours, University of Ghent
  level: introductory
- Introduction to Scientific Visualization, 28 April, KU Leuven/University of Hasselt
  organisation: VSC, doctoral schools
  trainer: Hamish Carr (Leeds University, UK) level: introductory

May

- High performance Python, 2 May, 4 hours, KU Leuven/University of Hasselt
  level: intermediate
- OpenFOAM, 3 & 4 May, 14 hours, University of Ghent
  trainer: prof. dr. Hrvoje Jasak (Wikki Ltd., Zagred University, Croatia)
  organisation: VSC, doctoral schools
  link to videos: https://www.youtube.com/playlist?list=PLqxhJj6bcnY9RoIgzeF6xDh5L9bbeK3BL
  level: specialist
- Python course, 9, 10, 18 & 19 October, 16 hours, KU Leuven/University of Hasselt
  level: introductory
- Lunchbox session Tier-2@KU Leuven, 12 May, 2 hours, KU Leuven/University of Hasselt
  level: introductory
- High performance Python, 25 May, 4 hours, KU Leuven/University of Hasselt
  level: intermediate
- Introduction to multithreading and OpenMP, 26 & 27 May, 14 hours,
  trainer: dr. Reinhold Bader (LRZ, Garching, Germany)
  organisation: VSC, doctoral schools
  level: intermediate
- Getting started with HPC, 25-27 May and 1-3 June, 36 hours, University of Ghent
  organisation: Prof. Peter Dawyndt, doctoral schools
  level: introductory

June

- Tier-1 Info session, 13 June, 2 hours, University of Hasselt
  level: introductory
- Visualization with VTK and Visit, 16 & 17 June, 14 hours, KU Leuven/University of Hasselt
  trainer: prof. dr. Hank Childs (University of Oregon, US)
  organisation: VSC, doctoral schools
  level: intermediate
- Tier-1 info session, 23 June, 2 hours, University of Ghent
  level: introductory
- Tier-1 Info session, 24 June, 2 hours, KU Leuven
  level: introductory
- Tier-1 Info session, 24 June, 2 hours, University of Antwerp
  level: introductory
- Tier-1 Info session, 27 June, 2 hours, Vrije Universiteit Brussel
  level: introductory
- HPC@UGhent introduction, 3 June, 6 hours, University of Ghent
  level: introductory
  on demand for the Data Analysis department (Faculty of Psychological and Educational Sciences)

September

- HPC@UGhent introduction, 16 September, 6 hours, University of Ghent

level: introductory

October

- Introduction to Fortran, 3, 4 & 5 October, 24 hours, University of Ghent
  trainer: dr. Reinhold Bader (LRZ, Garching, Germany)
  organisation: VSC, doctoral schools
  level: introductory
- Linux introduction, 6 October, 4 hours, KU Leuven/University of Hasselt
  level: introductory
- Linux introduction, 10 & 11 October, 8 hours, University of Antwerp
  level: introductory
- Matlab, 13, 14, 20, 21, 27 & 28 October, 24 hours, KU Leuven
  level: introductory
- HPC@KU Leuven introduction, 11 October, 6 hours, KU Leuven/University of Hasselt
  level: introductory
- supercomputers for starters, 17 October, 3.5 hours, University of Antwerp
  level: introductory
- HPC@UAntwerp introduction, 18 October, 4 hours, University of Antwerp
  level: introductory
- Linux for HPC, 18 October, 4 hours, KU Leuven/University of Hasselt
  level: intermediate
- Leveraging Big Data Tools on HPC with HOD, 6 hours, University of Ghent
  level: intermediate

November

- HPC Tips & Tricks 4: Code modernisation - Two real world case studies and Intel tools for finding performance bottlenecks, 7 November, 2 hours, University of Antwerp
  level: intermediate
- Linux introduction, 9 November, 6 hours, KU Leuven
  level: introductory
- Python as a second language, 7 & 8 November, 14 & 15 November, 16 hours, KU Leuven/University of Hasselt
  level: introductory
- Introduction to Linux, 18 November, 6 hours, University of Ghent
  level: introductory
- Linux for beginners, 18 November, 6 hours, Vrije Universiteit Brussel
  level: introductory
- Scaling your data analysis in Python with Pandas and Dask, 6 hours, University of Ghent
  trainer: dr. Joris Van den Bossche (Continuum Analytics, VUB)
  organisation: VSC, doctoral schools
  level: intermediate
- Introduction to High Performance Computing, 28 November, 6 hours, Vrije Universiteit Brussel (VUB)
  level: introductory
- MPI course, 28 & 29 November, 15 hours, VSC, location: KU Leuven
  based on PATC course by dr. Rolf Rabenseifner (HLRS, Stuttgart, Germany)
  level: intermediate

December

- Debugging techniques, 1 December, 4 hours, KU Leuven
  level: introductory
- Lunchbox session storage overview HPC@KU Leuven, 3 December, 2 hours, KU Leuven/University of Hasselt
  level: introductory
- OpenMP, 5 December, 7 hours, VSC location: KU Leuven
  based on PATC course by dr. Rolf Rabenseifner (HLRS, Stuttgart, Germany)
  level: intermediate
- Linux@UHasselt introduction, 9 December, 4 hours, University of Hasselt
  level: introductory
- Version control with git, 13 December, 3 hours, University of Hasselt
  level: introductory
- HPC@UHasselt, 16 December, 4 hours, University of Hasselt
  level: introductory
- Fortran for programmers, 19 December, 4 hours, KU Leuven/University of Hasselt
  level: intermediate

It is interesting to note that alongside the annually varying topics with training courses by external people, visualizations and OpenFOAM in 2016, there are also new, in-house training courses being developed each year. The aim is to have a training programme that is as complete and balanced as possible.

Communication and events

Inauguration BrENIAC

In 2016, in Leuven, BrENIAC, the latest HPC Tier-1 infrastructure in Flanders, was installed with the support of the Flemish government. This supercomputer is the successor of Muk, the Tier-1 which was hosted at UGhent for the past four years. This time, too, the VSC will manage the entire infrastructure. KU Leuven will act as host.
The name 'BrENIAC' assigned to the Leuven-based cluster, has a special meaning. It is a contraction of the words 'Brain' and 'ENIAC', one of the first electronic computers for general use in the world, released in 1946, and programmed by 6 women.
A solemn inauguration was organised on Monday 17 October 2016 at the KU Leuven in the presence of Flemish Minister Philippe Muyters, Rector of KU Leuven Rik Torfs, General Manager of KU Leuven Koenraad Debackere, Director ICTS of KU Leuven Annemie Depuydt and around 140 guests from the various Flemish university associations involved, industrial partners, and the FWO.

Photo 1
From left to right: Christel Maeyaert (Director Research Coordination), Annemie Depuydt (Director ICTS), Hans Willems (FWO), Liliane Schoofs (Vice Rector Research Policy) and Leen Van Rentergem (Head of Department Facilities for Education, Research, Communication and Cooperation)

The event itself consisted of two parts: an academic session and the official inauguration. During the academic session, the guests were offered a varied programme with the following speakers and topics:

- Introduction: Professor Dirk Roose (chairman of the VSC Users Committee)
- Session 1: The virtual spectrometer: theory supports experiment by professor Wouter Herrebout (University of Antwerp)
- Session 2: Why climate science needs high-performance computing: Demonstrations, prospects and challenges by professor Nicole van Lipzig (KU Leuven)
- Session 3: The good, the bad and the ugly: Chemical engineering meets high-performance computing by professor Kevin Van Geem (University of Ghent)
- Session 4: High-performance computing in health care by dr. ir. Wilfried Verachtert (Imec)

This was followed by the actual inauguration with short speeches by:

- Professor Koenraad Debackere, General Manager KU Leuven
- Mr. Kimihiko Fukuda, Executive Vice President, NEC Corporation
- Minister Philippe Muyters, Flemish Minister for Employment, the Economy, Innovation and Sport
- Professor Rik Torfs, rector KU Leuven

The event was closed by Minister Muyters who had the honour of giving the official starting orders by starting up a time lapse video that showed the installation of the new infrastructure. Subsequently, everyone was invited to raise a glass to the new Tier-1.

The entire event was largely covered by the press which had also been invited for a guided tour in the data centre in Heverlee, and was given the opportunity to conduct interviews and take the first shots of the latest supercomputer. As a result, the supercomputer could be seen on VRT, VTM, ROB and Kanaal Z, and was written about in De Morgen, De Tijd, Datanews, Zdnet and later also TechPulse.

Since the event was also the starting signal of a more large-scale campaign around HPC, we had also designed a campaign image and campaign video for BrENIAC, which were disclosed during the inauguration. The campaign image is a creation by Joris Snaet, in-house cartoonist of KU Leuven and also cartoonist at De Standaard, and adorns the complete front of the supercomputer. The campaign video is a production of Limel, the video producer of KU Leuven. Let us hope that, for the next four years, we can give BrENIAC and supercomputing in general the attention it deserves!

Participation in national, international events, conferences and workshops

Organised by VSC

- 21th Quattor Workshop – University of Ghent
  22-24 March 2016 – Brussels and after 23/3 Ghent
  http://www.quattor.org/meeting/2016/03/22/bruxelles-workshop.html
- BeneLearn 2016 – University of Ghent
  12-13 September 2016 – KULAK, Kortrijk
  https://www.kuleuven-kulak.be/benelearn/
- Big Data & Data Science devroom – University of Ghent
  Devroom (30/01) within FOSDEM 2016, co-organised by HPC-UGhent
  30-31/01/2016 – Brussels
  https://archive.fosdem.org/2016/schedule/track/hpc,_big_data_and_data_science

Organised by other organisations

- BeGeo 2016
  3 March 2016 – Brussels
  Contribution: talk 'Hanythingondemand - easily creating on-the-fly Hadoop clusters (and more) on HPC systems' (University of Ghent)
- HPC Advisory Council meeting
  21-23 March 2016 – Lugano, Switzerland
  Contribution: EasyBuild tutorial (University of Ghent)
  http://insidehpc.com/2016/03/easybuild/
- iMinds The Conference
  28 April 2016 – Brussels
- Spectrum Scale user group
  17-18 May 2016 – London, UK
- Big N2N Annual Symposium 2016
  19 May 2016 – Ghent
- Contribution: invited talk 'HPC-UGhent: empowering researchers with supercomputing' (UGhent)
  http://www.bign2n.UGhent.be/node/644
- Elixir Belgium training meeting
  28 September 2016 – Brussels
- BRUCON 2016 security conference
  27-28 October 2016 – Brussels
- 22nd Quattor workshop
  4-6 October 2016 – London, UK
- OpenNebulaConf 2016
  24-26 October 2016 – Barcelona, Spain
- Contribution: talk 'OpenNebula Sunstone integration with FreeIPA using Single Sign On' (UGhent)
  http://2016.opennebulaconf.com
- Meeting with the Taipei Representative Office in the EU and Belgium (University of Ghent)
  2 March 2016 – Ghent
- Day on cooperation with SCK.CEN organised by University of Antwerp
  11/02/2016 – Antwerp
- HTCondor - ARC Workshop at ALBA Synchrotron
  29 February – 4 March 2016 – Barcelona, Spain
- Data Innovation Summit (DIS 2016)
  23/03/2016 – Brussels
- EGI Conference 2016
  6-8 April 2016 – Amsterdam, Netherlands
- IOFday organised by University of Hasselt
  17/05/2016 – Hasselt
- Vortech mini symposium, VUB
  14 June 2016 – Brussels
- ISC 2016
  20/06-22/06 2016 – Frankfurt, Germany
- Digital Infrastructures for Research 2016
  27-30 September 2016 – Krakow, Poland
- Belnet Networking Conference 2016
  25 October 2016 – Brussels
- HP Cast
  11-12 November – Utah USA
- Supercomputing 2016
  13-17 November – Utah USA

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.

VSC ECHO

In 2016 the VSC Echo, the newsletter of the VSC, was published in mid-January and in late October. The January 2016 issue, VSC Echo 9, was focused on the amended regulations for Tier-1 applications, the planned training programmes, and announced the new Tier-1 supercomputer. The autumn issue, VSC Echo 10, was traditionally devoted to the upcoming training programmes and events, with attention being also paid to the training provided by CÉCI, the Walloon counterpart of the VSC.

Each VSC Echo also contains news about the use of Tier-1. Thus, information was provided about the submission of project proposals for the use of computing time on Tier-1, about the closing dates for submitting proposals etc.. The VSC Echo is sent electronically to more than 2,000 addresses and a paper version is available at events, meetings etc. as well as on the VSC website.

International collaboration

PRACE

The VSC is, via the Belgian membership, part of the European PRACE project. 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 and ranked. The highest ranked projetcs receive the requested computing time. Flemish researchers are made aware of the existence of these calls through notifications on the VSC website and announcements in the VSC Echo. PRACE also offers interesting training courses that are promoted by the VSC.

The cost model for PRACE is currently being revised at European level. A new financing model, PRACE 2.0, was elaborated 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.

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 enables researchers from Flanders to use this European computing infrastructure. It is used primarily by high-energy physicians from the universities of Antwerp, Ghent and Brussels. The VUB research group "Icecube", for example, decided to use this grid infrastructure to replace their own infrastructure. Solid, a new international research group with researchers from the University of Ghent, Antwerp 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 easy data sharing with the participating universities. 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.

SESAME Net

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.
One of the particular characteristics of SESAME Net is the consortium: it consists of a mix of national and regional HPC groups supplemented with private partners. There are 14 partners from 12 EU countries: HPC-Wales (UK), ICHEC (Ireland), Fraunhofer institute (Germany), PSNC (Poland), GRNET (Greece), CESGA (Spain), IT4Innovations (Czech Republic), Yotta (Croatia), UVT (Romania), IICT (Bulgaria), RBI (Croatia), FWO/VSC (Belgium) and Vilnius University (Lithuania). The project was officially launched on 1 June 2015 and runs for a period of 2 years.
The project reached cruising speed in 2016. The activity of VSC is concentrated in Work Package 3 (Competency Centre Best Practice) and Work Package 4 (Awareness Raising). VSC has, for example, shared a best practice on EasyBuild and HanythingOnDemand with all SESAME Net partners. Conversely, HPC Wales and ICHEC, for example, have shown how they deal with commercial users and given insight into their SLAs. For Work Package 4 it proved particularly useful to share experiences with the partners. It appears, for example, that virtually all HPC Centres have difficulty in finding SMEs that are interested in supercomputing and ready to adopt this methodology.
In June 2016, the interim review of the SESAME Net project was held in Ghent. The review panel expressed some concerns about the direct impact of the project, but after some discussion the project was given the green light to continue. In September 2016, the project coordinator, HPC Wales, announced its decision to withdraw from the project. At that time, HPC Wales had already ceased after the budget was discontinued by the government. Although Bangor University had taken over the SESAME Net tasks from HPC Wales, it finally decided that the project was no longer a priority for them. Croatian partner Yotta/Arctur took over the coordinator role with unanimous approval. Under their impulse a new objective was added to the project: a self-assessment tool for SMEs (HPC4SME). The purpose of this tool is to assess the HPC potential of an SME prior to any further talks. The tool will be operational in early 2017.
Ward Poelmans (FWO) is temporarily in charge of these project resources (up to and including December 2016) and responsible for all SESAME Net-related tasks. In this connection, he proactively contacts potentially interested companies, and monitors them when they effectively make the switch and use the VSC services.
Website: http://sesamenet.eu

Contacts with other HPC centres

The VSC regularly maintains contact with other HPC centres not only to be informed of developments in HPC, but also to exchange best practices. In fact, many HPC centres worldwide use various hard- and software solutions for resolving recurrent problems. Exchanging such information is of paramount importance for the operation of each HPC centre.

- Rutherford Appleton Laboratory (UK)
  6-8 December 2016 – Ghent
  Cooperation and consultation on Quattor
- Texas Advanced Computing Center (TACC, USA)
  25-26 January 2016 – Ghent
  1-2 February 2016 – Ghent
  23-28 June 2016 – Ghent
  Cooperation on Lmod and XALT
  Presentation by Robert McClay from XALT to all VSC admins (28 June)
- Uppsala Multidisciplinary Center for Advanced Computational Science (UPPMAX, Sweden)
  9 May 2016 – Ghent
  Consultation with Prof. Anders Hast
- SurfSARA (Netherlands)
  20 October 2016 – Ghent
  Consultation on visualization
- OKAN University (Turkey)
  3 May 2016 – Ghent
  Networking

The Grid team (VUB) also has contacts with colleagues from Louvain-La-Neuve and with people from many other HPC/HTC/Cloud sites in Europe through its participation in EGI and projects like Long-tail of Science, FedCloud, etc.

Several international contacts took place in the context of the cooperation on EasyBuild (http://hpcUGhent.github.io/easybuild/). This is a 'killer tech' software solution for installing scientific software on HPC systems in an easy, efficient and reproducible manner. Initially, this software was developed within the VSC. Today, however, many HPC centres throughout the world have come to appreciate the importance of this open source tool and are increasingly

participating in the further development and maintenance of the code. Within the EasyBuild network there are regular contacts with JSC (Germany), CSCS (Switzerland), TACC (USA), Ottawa Hospital Research Institute (Canada), The Francis Crick Institute London (UK).
The following specific events were organised around EasyBuild:

- 1st EasyBuild User Meeting – UGhent
  27-29 January 2016 – Ghent
  https://github.com/hpcUGhent/easybuild/wiki/1st-EasyBuild-User-Meeting
- 11th EasyBuild hackathon – UGhent
  24-25 March 2016 – CSCS, Lugano, Switzerlandhttps://github.com/hpcUGhent/easybuild/wiki/11th-EasyBuild-hackathon
- 'Introduction to EasyBuild' for delegation of Cuban universities (University of Informatic Sciences UCI - Havana, Central University of Las Villas - Santa Clara, University of Santiago de Cuba)
  9, 14 November 2016 – Ghent

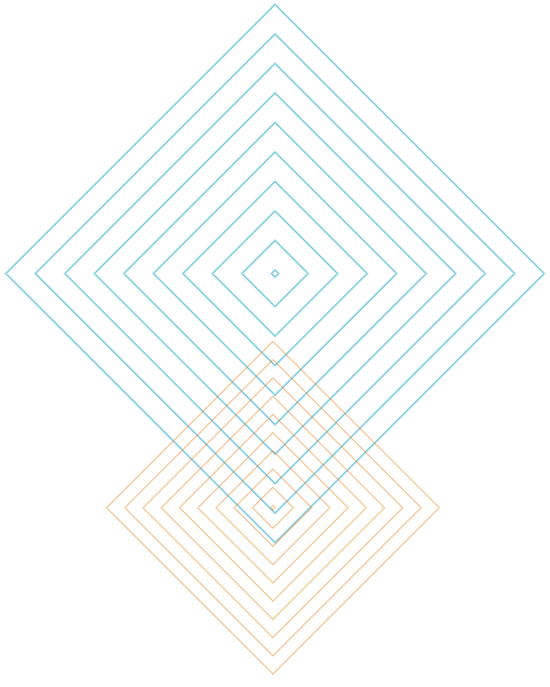

VSC Success Stories

Dr. Wim Thiery – Prof. Nicole Van Lipzig (Department of Earth and Environmental Sciences, KULeuven): Climate change intensifies night-time thunderstorms over Lake Victoria

Lake Victoria in East Africa will become a hotspot for hazardous thunderstorms due to climate change. This is shown by an international study led by Wim Thiery and Nicole van Lipzig from the Department of Earth and Environmental Sciences, KU Leuven, published in Nature Communications.
Lake Victoria is divided among Uganda, Kenya, and Tanzania. With a surface close to 70,000 km2, it is the biggest lake in Africa. The lake is also a notoriously dangerous place for the 200,000 people who go fishing there at night. The International Red Cross estimates that between 3,000 and 5,000 fishermen per year lose their lives in violent storms on the lake. That Lake Victoria can be so stormy at night is related to the circulation in the atmosphere above the enormous water surface. Wim Thiery: "During the day, a breeze develops flows from the cool water towards the warm land. At night, we see the opposite: the land breeze flows away from the cooling land towards the warmer lake. As the lake is shaped like a circle, these land breezes from all directions converge above the lake. Add evaporation to this cocktail and you get a lot of storms, rain, wind, and waves." This pattern was confirmed by analyses of NASA satellite observations in the period from 2005 to 2013. To examine how climate change will affect this process, climate studies were carried out using an advanced computer model, implemented with COSMO-CLM². For this, the non-hydrostatic regional climate model COSMO-CLM version 4.8, the "Community Land Model version 3.5 (CLM3.5)" and the "Freshwater Lake model (FLake)" were linked. Assuming greenhouse gases continue to increase, the extreme precipitation will rise twice as fast as that over the surrounding land, resulting in even more night-time thunderstorms. Superstorms that now occur once every 15 years, will occur almost yearly by the end of the century. Further examination should result in an improvement of the warning system to reduce the vulnerability of local fishermen.

Prof. Daniele Marinazzo (UGhent): modeling brain dynamics in health and disease using supercomputers

Marinazzo research lab at Department of Data Analysis, University of Ghent
Prof. Daniele Marinazzo leads a research lab within the Department of Data Analysis, based in the Faculty of Psychological and Pedagogical Sciences at Ghent University.

His group focuses on methodological and computational aspects of neuroscience research. Specifically, they examine how the brain works and how the collective dynamics and interplay of brain regions give rise to function.

Marinazzo's lab is part of the multi-disciplinary research platform 'Neuroscience' at Ghent University.

Computational research and applications
The human brain can be considered as a complex network, a so-called 'connectome', from which all function and behaviour can in principle be inferred. A computational model of such a connectome is therefore a very powerful tool, with predictive power and many medical and clinical applications.

Prof. Marinazzo's group develops new techniques to determine to what extent brain areas are structurally connected or functionally related, based on the data provided by neuro-

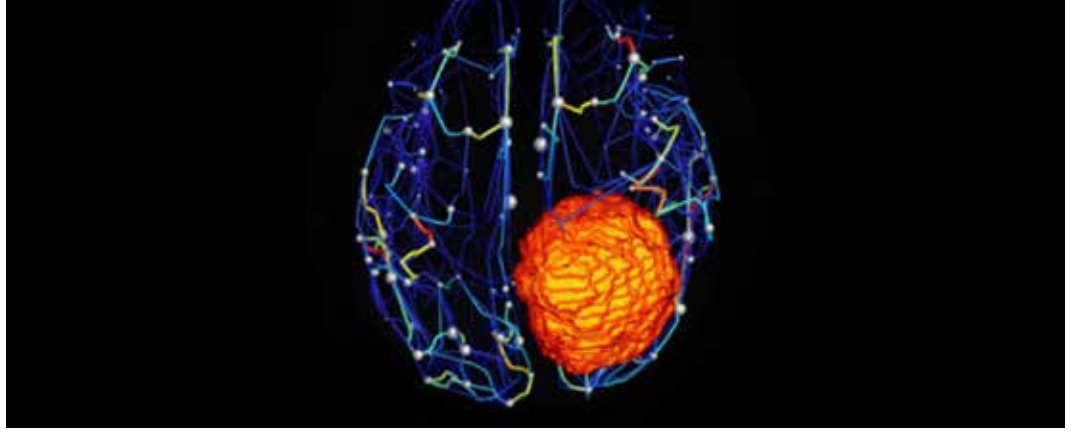

imaging methods such as magnetic resonance imaging (MRI). They employ data-driven and biological dynamical causal modeling, in addition to model-free theoretical frameworks, such as transfer entropy.

Thanks to these developments, and with the aid of supercomputers, Marinazzo's research group can construct a comprehensive whole-brain map of the structural and functional connections within the brain of an individual patient.

Dynamical network analysis of these 'brain maps' in patients with epilepsy has led to an unprecedented characterisation of the epileptic condition, enabling a more objective diagnosis and monitoring of recovery or progression processes over time. Furthermore, it opens up the possibility to predict the onset of epileptic seizures in a patient, and teach patients how to identify the onset and take remedial treatment.

Virtual brain models can also be used to study the mechanisms underlying brain diseases, and carry the potential to directly support treatments. Hannelore Aerts, one of Prof. Marinazzo's PhD students, recently applied these novel techniques to examine global brain dynamics in brain tumour patients. The final aim of the project is to assist neurosurgeons in presurgical planning, by simulating the effects of tumour resection on brain dynamics and functioning using a virtual brain model.

VSC facilitates computational research

"To derive the connectivity architecture of brain regions from recorded neuroimaging data, considerable computing power is required" says Prof. Marinazzo.

"But also our analyses need to run on supercomputers. The availability of several compute clusters at Ghent University and within the VSC network enables us to swiftly test several models at the same time. It really speeds up our research."

"Finally, the user-specific approach and fast-response helpdesk of the HPC-UGhent team is a tremendous help."

Links

- Department of Data Analysis (http://www.da.UGhent.be/index.html)
- Prof. Daniele Marinazzo's lab (http://users.UGhent.be/~dmarinaz/index.html

Reference

- 'Brain networks under attack: robustness properties and the impact of lesions', Hannelore Aerts, Wim Fias, Karen Caeyenberghs and Daniele Marinazzo, Brain, 2016, 139, 3063-3083. (https://www.ncbi.nlm.nih.gov/pubmed/27497487)
- 'Modeling global brain dynamics in brain tumor patients using The Virtual Brain', Hannelore Aerts and Daniele Marinazzo, Front. Neurosci. Conference Abstract: 12th National Congress of the Belgian Society for Neuroscience. doi: 10.3389/conf.fnins.2017.94.00101(http://www.frontiersin.org/10.3389/conf.fnins.2017.94.00101/event_abstract)

Flanders Hydraulics Research

Flanders Hydraulics Research is the centre of expertise of the Department of Mobility and Public Works of the Flemish Government, focusing on waterways, coastal and water-related structures built by humans. It advises the Flemish Government on the navigability and safety of waterways at the Belgian Coast. Flanders Hydraulics Research examines the impact of human activity on water systems and the consequences for sludge movement, morphological developments, nature, and accessibility for shipping. To enable it to carry out its tasks and develop innovative solutions, it has access to a wide range of physical model installations, in situ measuring instruments and numerical models.
Numerical models describe in 1D, 2D or 3D the tidal and wave-driven currents, sediment transports and soil changes based on the shallow water equations (simplified Navier-Stokes equations for free-surface flows). At the same time, 3D CFD models are used for complex flows in and around structures, but also for ships' hulls and rudders. In addition to the in-house computing capacity, the VSC infrastructure is used to compute larger models. One of these models is the Scaldis model. This is a 3D finite element model built using the TELEMAC-MASCARET suite of solvers. The model comprises the entire Belgian coast and estuary of the river Scheldt (roughly 900,000 elements horizontal and 5 β-layers, totalling about 4,500,000 elements). The model is currently applied to study the effects on the tide and sludge movement of potential future human interventions in the Upper Scheldt. To this end, different combinations of possible interventions are simulated and analysed. On the VSC cluster, 9 nodes are typically used per model run (180 cores). Each scenario is simulated for a 3-month period. The computing time on the VSC cluster is on average 11 days per scenario.

Quote: "Access to the VSC infrastructure enables us to efficiently calculate a large number of potential scenarios and analyse them for effects on tide, flow, nature and shipping, shaping our advice to the Flemish Government."

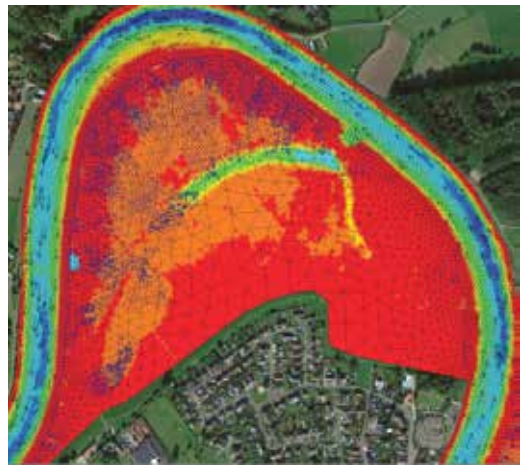

Detail of the Scaldis model: Ebb tide in Scheldt and Reduced Tidal Area (GGG) Bergenmeersen. The vectors indicate the flow, the colour, the depth of the river bed. The complete model covers the entire coast and Oosterschelde and the complete estuary from the mouth to Melle. Right: an aerial photograph of the area at storm surge (6/6/2017 - Saint Nicholas storm)

Diabatix – Making efficient thermal design readily accessible

Diabatix is a Flemish startup that focuses on the design, simulation and optimization of refrigeration components. This encompasses cooling components for lasers, (power) electronics, electric car batteries, internal combustion engines, etc.. Manufacturers work together with Diabatix to make their machines, appliances and installations more efficient, compact and environmentally friendly. Diabatix offers an innovative, easily-accessible platform solution, which assists every engineer to design critical thermal applications in a faster and better way. The generative design method used removes the typical hurdles of trial-and-error-engineering, leading to disruptive cooling performance enhancements up to even 30%.
The platform that Diabatix offers runs on proprietary software. It gives the customer easy access to an in-depth analysis of the complex physics of heat transfer in liquids, solids and gas flows during design. In this way, customers need not be experts in thermal design, but still get access to optimized scientifically-sound design.
The calculations behind the design, simulations and optimizations require a lot of computing power, for which the VSC compute clusters are very useful. With the aid of HPC, Diabatix is able to master more complex thermal applications in a faster and more profound manner. Diabatix founders Lieven Vervecken and Joris Coddé: "During our doctoral research, we have gained experience with the use of HPC, also in the VSC. This allows us to process thermal design studies on hundreds of processors, which typically makes the computation 50 times shorter."

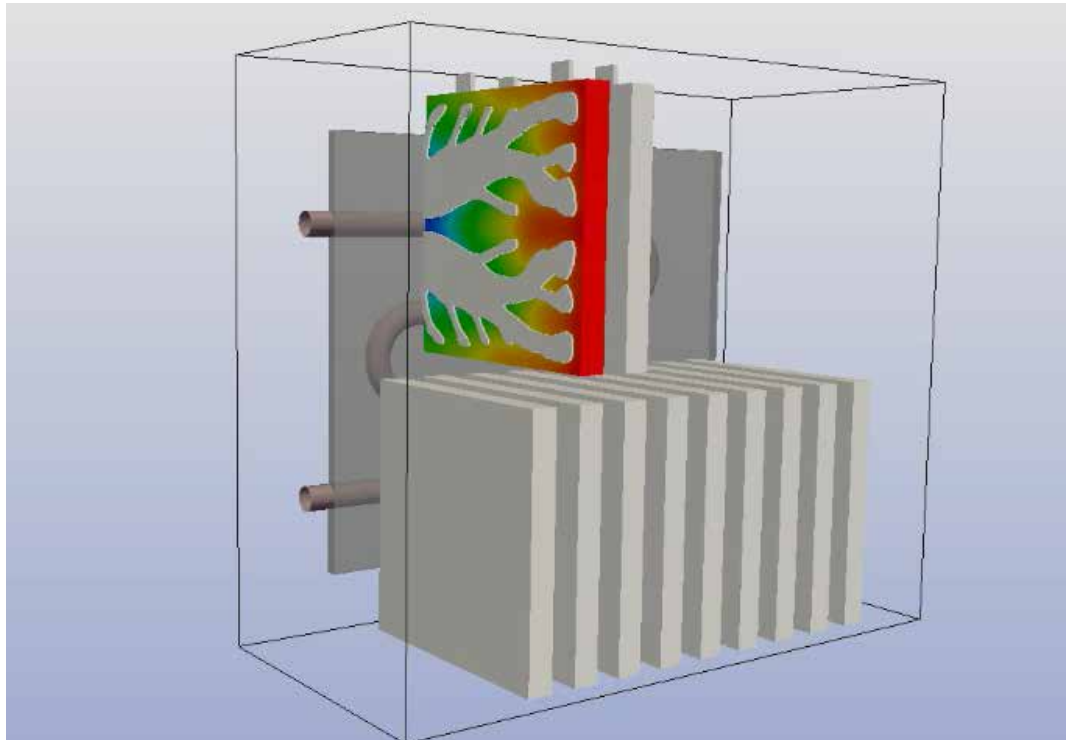

Heat emitter for indirect battery cooling

Colophon

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

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

Editor-in-chief
dr Hans Willems
Secretary-General FWO

Text
Leen Van Rentergem, Jan Ooghe
Annie Cuyt, Stefan Becuwe
Ewald Pauwels, Ward Poelmans
Geert Jan Bex
Rosette Vandenbroucke
Caroline Volckaert
Stefan Weckx

Creation and design
Billie Bonkers

Published in june 2017

fwo

Opening
new
horizons