

ANNUAL REVIEW HPC-UGENT

2017



Department ICT, Infrastructure office

E hpc@ugent.be
T +32 9 264 4716

Campus Sterre S9
Krijgslaan 281, 9000 Gent

<https://www.ugent.be/hpc>

CONTENTS

CONTENTS	2
1 About HPC-UGent	3
1.1 Our mission	3
1.2 Our vision	3
1.3 Personnel	3
2 Infrastructure	4
2.1 Overview	4
2.2 Usage	5
3 Development and maintenance	11
3.1 Number of JIRA issues resolved in 2017, grouped per main component	11
3.2 Projects finalized in 2017	12
3.3 Github commits in 2017 by HPC-UGent staff, per repository	13
4 Training and support	14
4.1 Training overview and evaluations	14
4.2 Lectures and community meetings	17
4.3 Helpdesk	18
4.4 User evaluation	19
5 Outreach	23
5.1 Within Ghent University	23
5.2 To policy makers, industry and general public	23
5.3 Within international HPC community	25
6 Budget	26
7 Tier-1 usage	27
8 User in the spotlight	28

1 ABOUT HPC-UGENT

**In scientific computing*, computers are used to solve complex problems.
(*aka: supercomputing or high-performance computing - HPC)**

1.1 Our mission

HPC-UGent provides centralised scientific computing services, training, and support for researchers from Ghent University, industry, and other knowledge institutes.

HPC-UGent is part of the central ICT department of Ghent University, and is a strategic partner of the [Flemish Supercomputer Center \(VSC\)](#).

1.2 Our vision

HPC-UGent offers a professional scientific computing environment that is stable, user-friendly, and serves the diverse purposes of researchers from Ghent University, industry and other research institutions.

We provide a structural training curriculum for new and advanced users, and provide supporting course material.

We present a supercomputing portfolio that is well known within Ghent University and beyond, and we establish ourselves in the international community via contributions to centralised solutions, such as [EasyBuild](#).

1.3 Personnel

The HPC-UGent team in 2017 consists of 8 people:
Alvaro Simon Garcia, Andy Georges, Ewald Pauwels, Jens Timmerman, Kenneth Hoste, Kenneth Waegeman, Stijn De Weirdt, Wouter Depypere.

Tasks include:

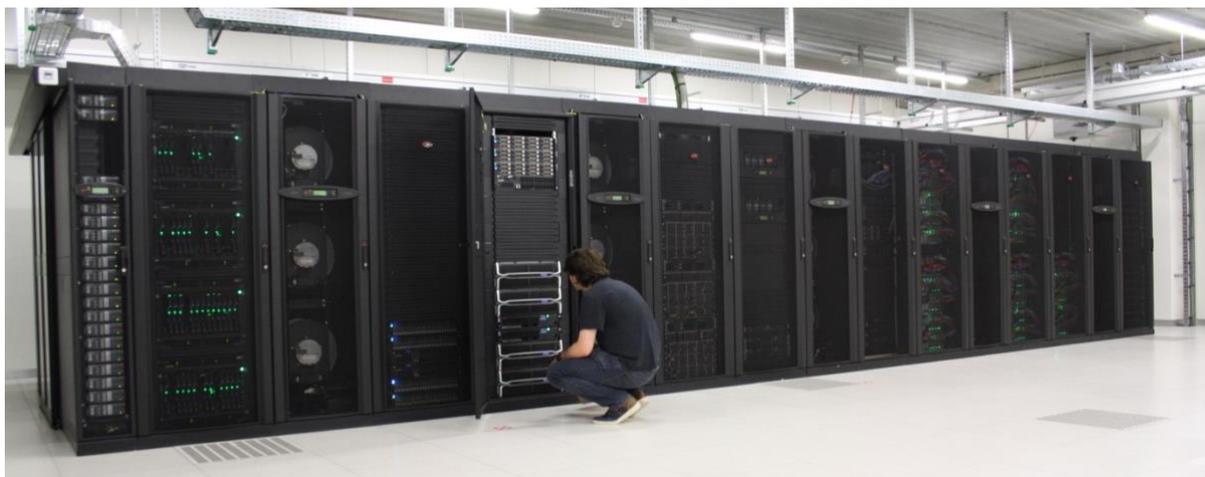
- User support
- Training
- Infrastructure installation and upkeep (both hardware and software)
- Outreach and marketing
- Collaboration with other supercomputing centers

2 INFRASTRUCTURE

2.1 Overview

The Ghent University compute infrastructure (Tier-2) consists of several specialised clusters, jointly called Stevin, hosted in the S10 datacenter.

In 2017, a new DATA and SCRATCH storage platform was taken in production, totalling 2 PB in size. Orders were also placed for two new compute clusters, which will be taken in production in 2018.



2.1.1 Compute clusters

Cluster name	#nodes	CPU per node	Memory per node	Local disk per node	Network interconnect
Raichu	56	2 x 8-core Intel E5-2670 (Sandy Bridge @ 2.6 GHz)	32 GB	400 GB	Gb Ethernet
Delcatty	151	2 x 8-core Intel E5-2670 (Sandy Bridge @ 2.6 GHz)	64 GB	400 GB	FDR InfiniBand
Phanpy	16	2 x 12-core Intel E5-2680v3 (Haswell-EP @ 2.5 GHz)	512 GB	3 x 400 GB (SSD, striped)	FDR InfiniBand
Golett	196	2 x 12-core Intel E5-2680v3 (Haswell-EP @ 2.5 GHz)	64 GB	500 GB	FDR-10 InfiniBand
Swalot	128	2 x 10-core Intel E5-2660v3 (Haswell-EP @ 2.6 GHz)	128 GB	1 TB	FDR InfiniBand

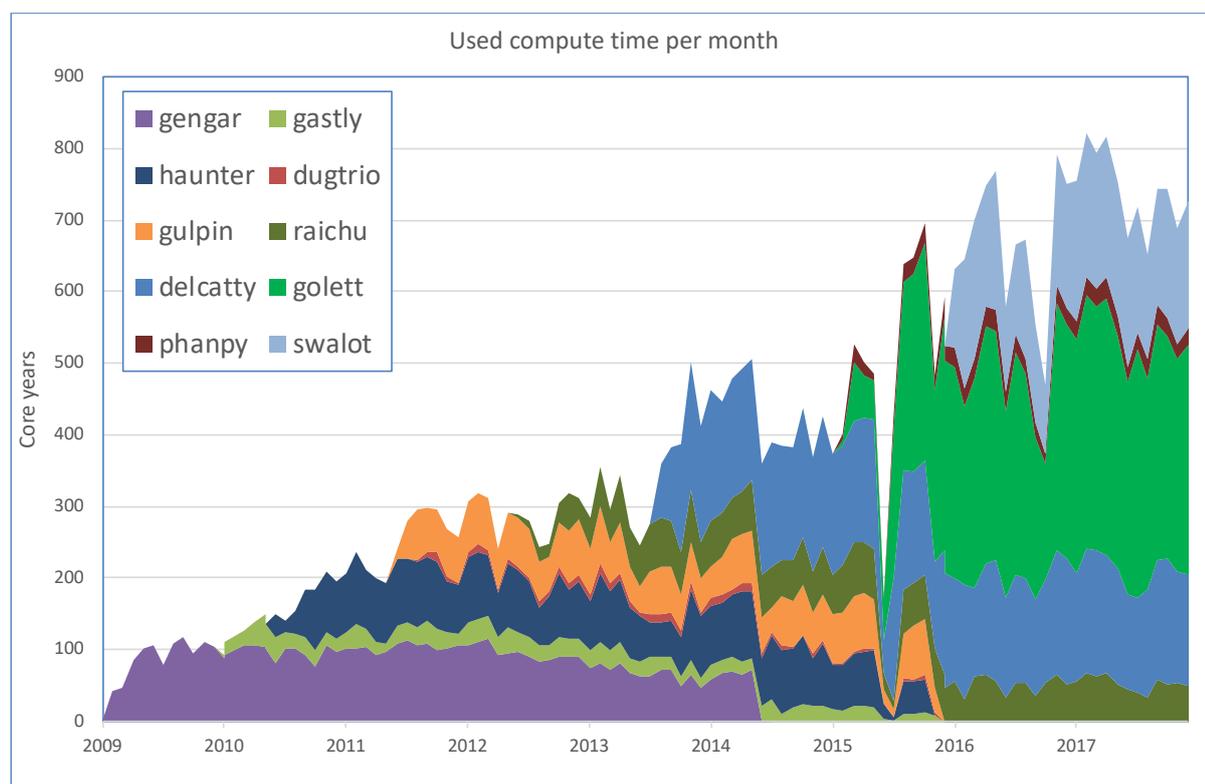
2.1.2 Shared storage

Partition	Size
<code>\$VSC_HOME</code>	35 TB
<code>\$VSC_DATA</code>	702 TB (can grow to 1 PB)
<code>\$VSC_SCRATCH</code>	1 PB
<code>\$VSC_SCRATCH_KYUKON</code>	
<code>\$VSC_SCRATCH_PHANPY</code>	35 TB SSD

2.2 Usage

Modern CPUs contain many cores that are capable of running a computational task. In the graphs below, the used or available compute time on the compute clusters is typically expressed in *core years*, where 1 core year would correspond to the work done by one core in one year of time. (One core hour is the work done by one core in one hour of time.)

2.2.1 Historical perspective on used compute time

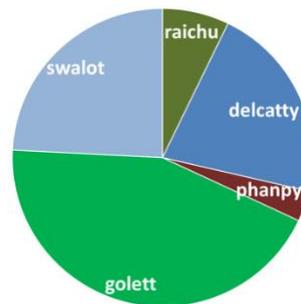


This plot gives a historical perspective on the amount of compute time that is used every month, coloured according to the cluster that does the work. On average, every month 7.000.000 core hours (corresponding to 800 core years) of compute time are produced. All clusters together, since the start in 2009, have produced more than 36 billion core hours. This is equivalent to 42.000 years of compute work on one core.

2.2.2 Consumed compute time in 2017

2.2.2.1 Consumed compute time per compute cluster

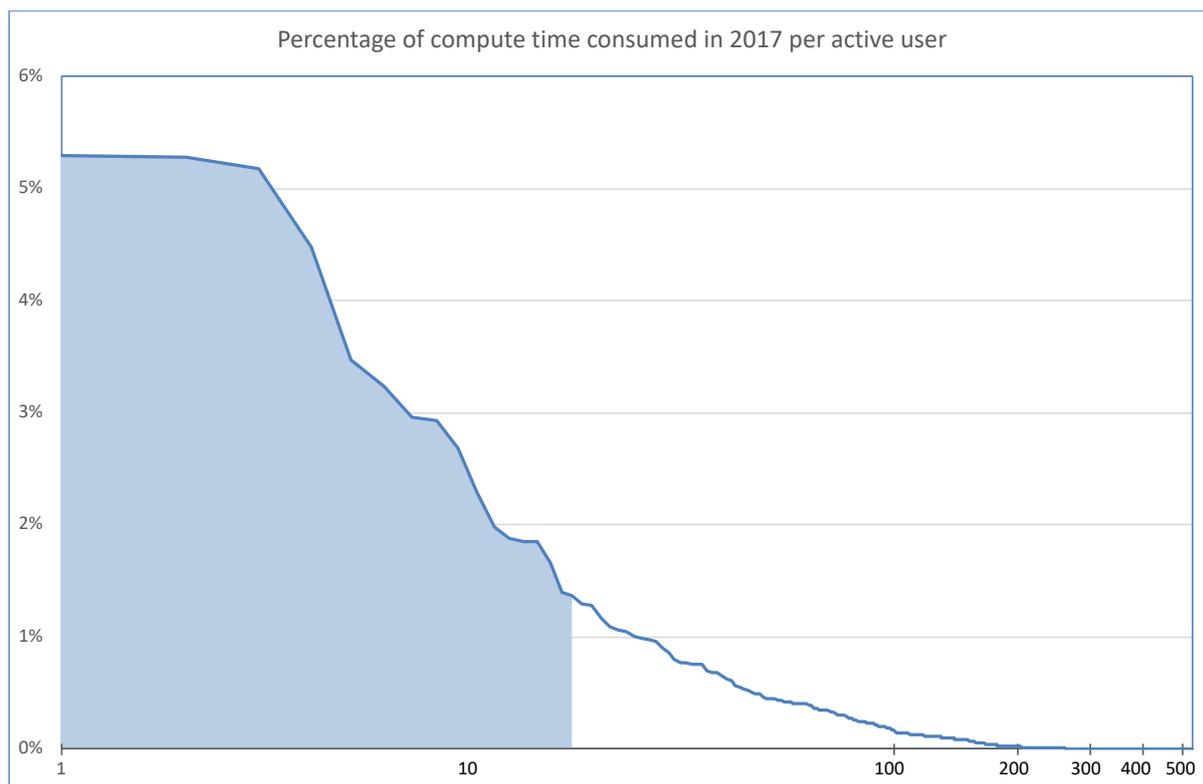
Cluster name	Compute time consumed (in core years)	Effective use percentage
Raichu	652	73%
Delcatty	1942	80%
Phanpy	301	78%
Golett	3974	84%
Swalot	2191	86%
Total	8408	77%



In 2017, a total of 73.656.948 core hours has been consumed on the Tier-2 compute clusters of Ghent University. This corresponds to 8.408 core years.

The effective use percentage expresses how much of the theoretically available compute power in one year ($\#nodes \times \#cores/node$) was used. Down-times were not taken into account, so the percentages represent a lower bound. Effective use percentages are quite high for typical HPC systems and indicate good economic usage of the infrastructure.

2.2.2.2 Consumed compute time versus active users



In 2017, 531 persons actively used the Tier-2 compute clusters of Ghent University. The user base typically contains a limited number of power users in addition to regular users with a lower usage profile. The plot above graphs per user (x axis) what percentage of all compute time produced in 2017 this user consumed. The blue area indicates that 50% of all Tier-2 compute time was consumed by 17 power users.

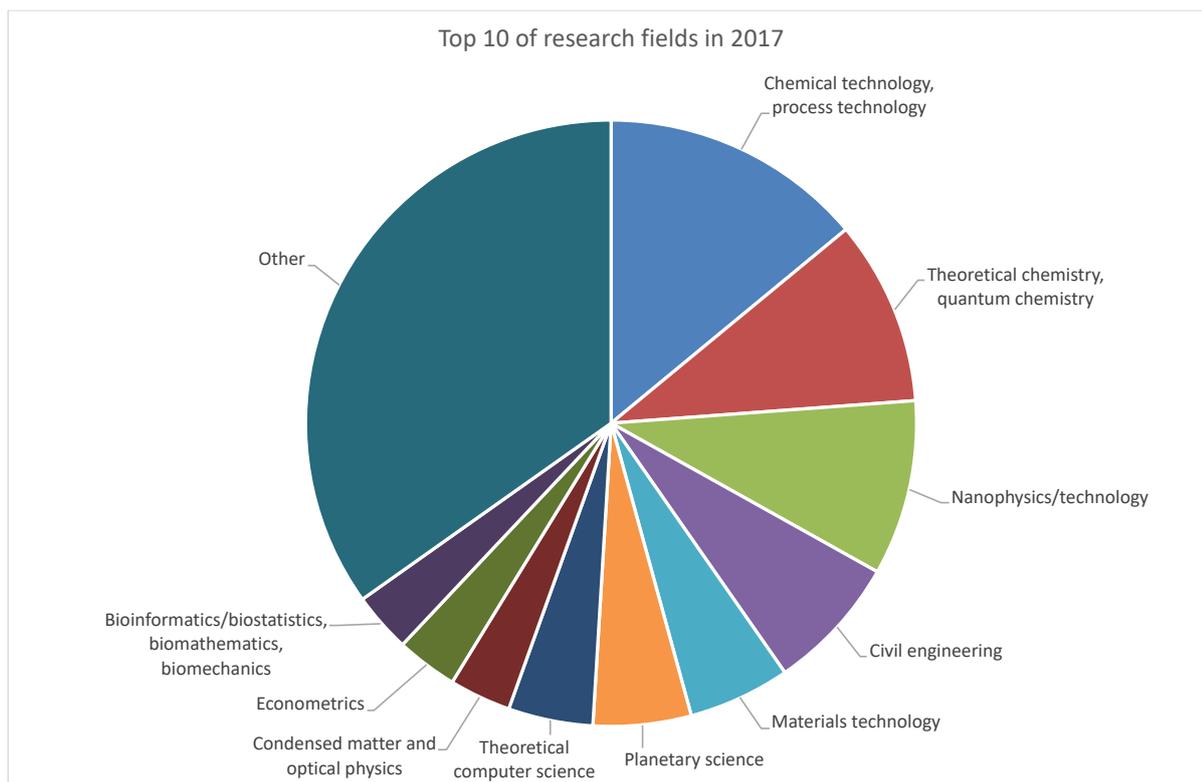
Everyone can get access to the Tier-2 compute clusters of Ghent University. The conditions that apply depend on the affiliation of the researcher. Researchers affiliated with Flemish university associations (constituting the Flemish Supercomputer Center – VSC) get free access. Other Flemish or federal research institutes can get access for their researchers on a contract basis, with 1500 euro worth of compute time given out free of charge per year. Industry can buy compute time in a pay-what-you-use model.

Breakdown of consumed compute time by affiliation	
UAntwerpen	0.240%
VUB	0.013%
UGent	99.480%
KULeuven / UHasselt	0.258%
Other research institutes	0.008%
Industry	0.001%
Total	100.000%

Master or Bachelor students enrolled in a Flemish university association can also get access to the Tier-2 compute clusters of Ghent University. Several teachers effectively rely on the infrastructure for training purposes, stimulating several students to become a user. Master students often rely on the infrastructure to perform research included in their Master thesis.

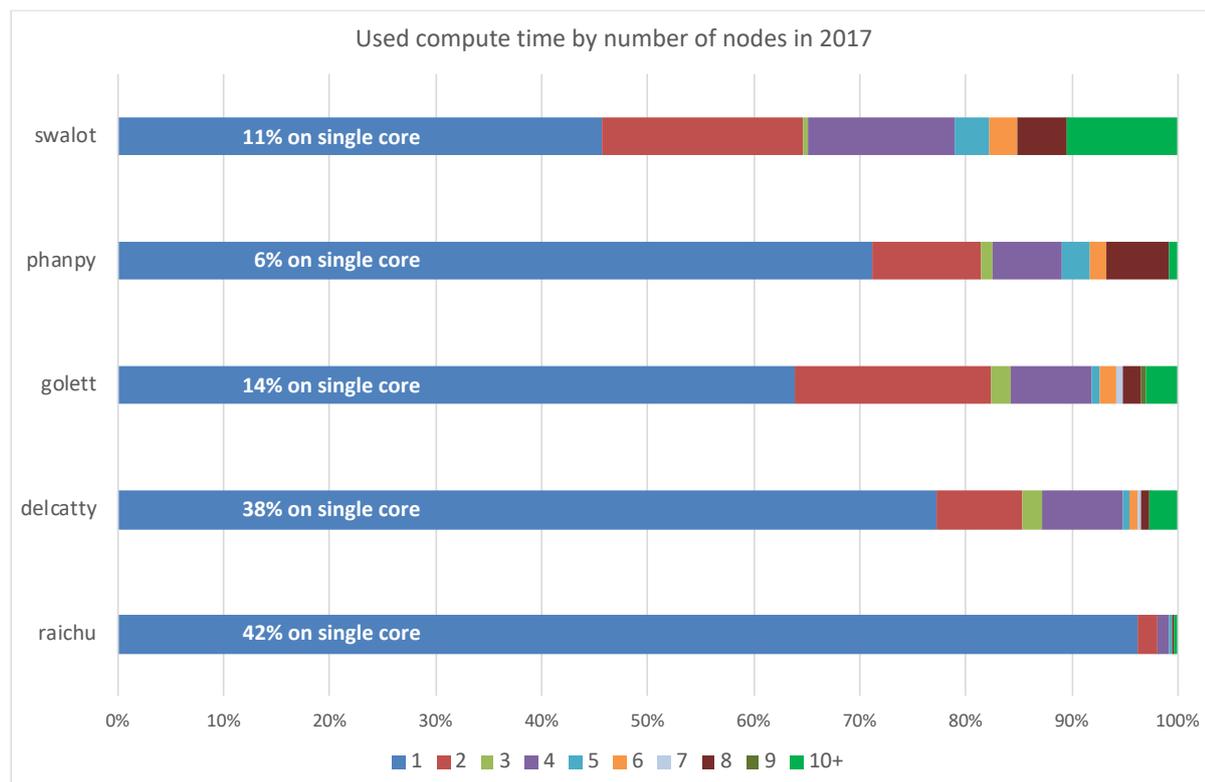
Breakdown of consumed compute time by category	
Ma/Ba students	13%
Researchers	87%
Total	100%

2.2.2.3 Consumed compute time versus research field



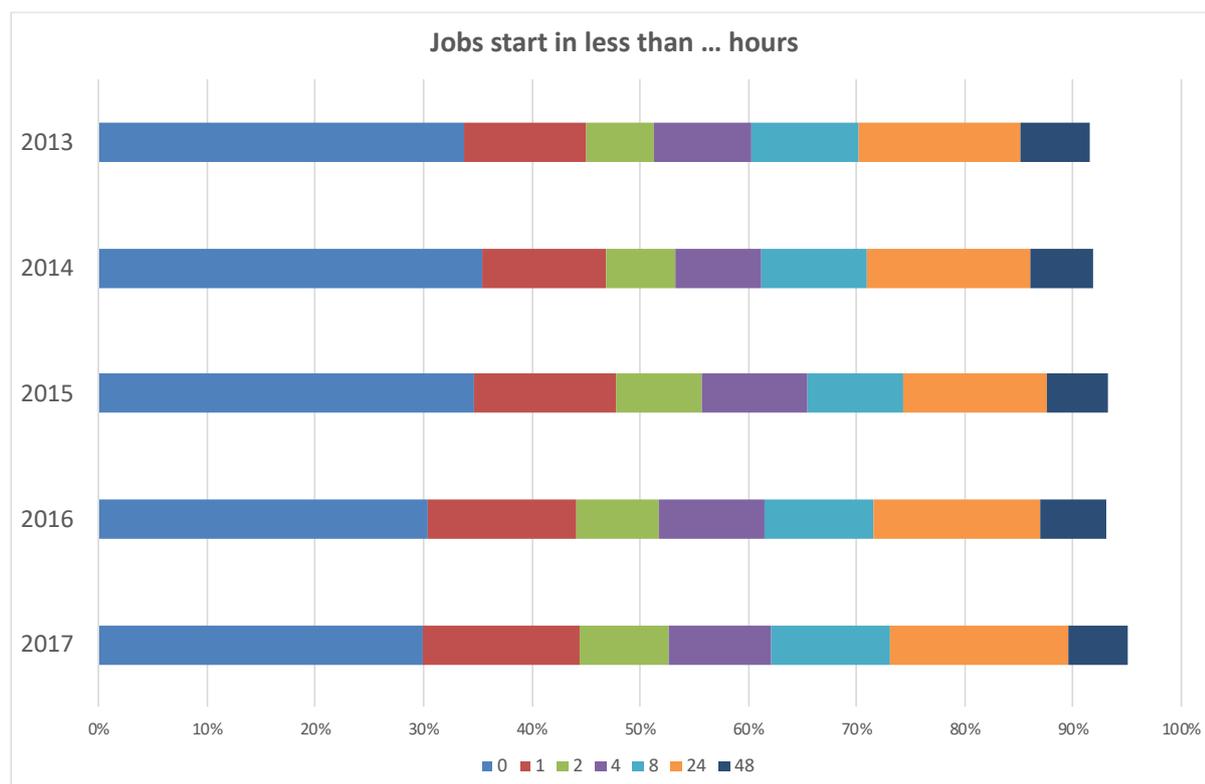
These are the research fields that users enter when requesting their account to access HPC-UGent resources. No doubt, there is overlap between certain fields and some users have not listed a research field at all. As such, the above distribution merely gives an indication of the top research fields that actively use scientific computing.

2.2.2.4 Versus number of nodes requested



Computational tasks on the Tier-2 compute clusters vary in the amount of resources that they require. This plot correlates the number of nodes that is requested in jobs with the amount of compute time that they use. Single-node use of the clusters is popular, but only a minor proportion of those jobs only use a single core; within-node parallelization is largely applied.

2.2.3 Queue time



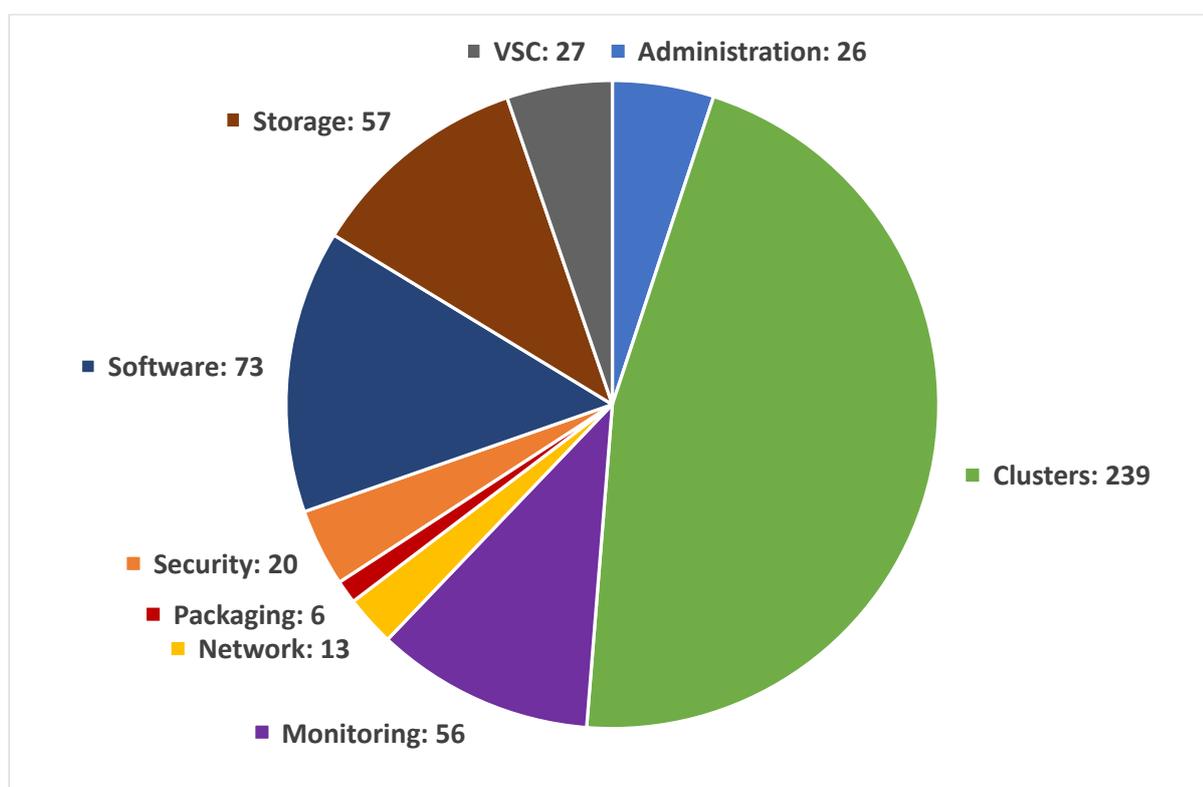
Since the Tier-2 compute clusters are shared resources, there is unavoidably a queue to use these resources. On the clusters, this is managed by a scheduler (i.e. software) that determines which task goes first depending on requested resources (number of cores/nodes, memory, compute time, ...), historical usage of the user (e.g. very active in last few weeks versus not used Tier-2 in a while) and cluster usage. The scheduler aims for fair share of allocated compute time between all users.

The graph above gives an indication on how long jobs typically spend waiting in the queue before they effectively get the resources they requested. In 2017, 30% of jobs started immediately (i.e. no waiting time at all), over 50% of jobs started less than 2 hours after submission. 90% of jobs started in less than 24 hours. This pattern has roughly remained the same since 2013.

3 DEVELOPMENT AND MAINTENANCE

To maintain compute infrastructure capabilities, HPC-UGent continuously performs maintenance works. On hardware, such as installing new components or replacing broken ones, and on software, ranging from the operating system to higher-level software that enables services directly available to the end user. To report, plan and follow-up hardware and software bugs, issues or developments, these are managed with the aid of a tracker, JIRA. Larger developments, grouping many individual or interrelated issues are grouped as 'Epics'. These effectively track the main development and maintenance projects that HPC-UGent executes. Github services are used to facilitate software development and Jenkins services for continuous integration. The number of 'commits' by HPC-UGent staff to the various repositories is indicative of the effort spent over the last year.

3.1 Number of JIRA issues resolved in 2017, grouped per main component



3.2 Projects finalized in 2017

Issue #	Title	Summary
HPC-4351	Testing and benchmarking of clusters	Set of scripts to automatically submit benchmarking and performance test jobs to clusters. Typically used after maintenance or unexpected downtimes.
HPC-4848	OpenNebula cluster phase II	Further roll out OpenNebula on testbed cluster, used to test new implementations, updates, etc. before applying to production clusters.
HPC-5169	Quattor documentation	Clean up and update of the documentation. Quattor is used for the automated installation, configuration, and management of clusters.
HPC-5821	Maintenance window Tier1 2016Q3	The end of the final maintenance work to the former Tier1 platform muk included measures preparing its decommissioning: alerting users, offloading storage, wiping storage, hardware clean-up.
HPC-5950	Tier2 site update 2016Q4	Rolling updates in the period Jan-Feb 2017, replacing full downtime that was originally planned in the Nov-Dec 2016 timeframe. Rolling updates, minimal downtime for users.
HPC-5959	Tier2 network revamp	Restructuring of internal network that connects all Tier2 clusters.
HPC-6199	Tier2 site update 2017Q2	General maintenance of all Tier2 clusters and servers. Rolling updates, minimal downtime for users.
HPC-6226	Replacement of HOME, DATA and SCRATCH	Main preparatory work related to new storage platform for HOME/DATA/SCRATCH.
HPC-6390	Rebuild muk into a cloud pilot	Decommissioned hardware from the former Tier1 platform muk was reconverted into a cloud pilot platform, preparing ground for the VSC 'Tier1 supercomputing as a service' project of 2018.
HPC-6578	Tier2 site update 2017Q3	General maintenance of Tier2 storage platforms - full downtime. Production stage of new HOME/DATA/SCRATCH storage volumes.
HPC-6580	An HPC-UGent booth highlighting added value of supercomputing to broad audience	Preparatory work for a booth on the 'Iedereen UGent!' event on Sunday 8 October 2017. Booth where everyone can experience hands-on supercomputing.

3.3 Github commits in 2017 by HPC-UGent staff, per repository

github.com/hpcugent	
Repository	#commits
vsc-accountpage	311
vsc-mympirun	94
vsc-jobs	52
vsc_user_docs	28
vsc-manage	27
vsc-accountpage-clients	17
hanythingondemand	14
vsc-install	14
csb	12
vsc-utils	11
vsc-base	8
vsc-filestystems	6
mympingpong	5
vsc-zk	1

github.ugent.be/hpcugent	
Repository	#commits
quattor	1685
icinga-checks	53
vsc-quattor	30
oncall-service	28
pbsmon2php-client	15
vsc-cluster-modules	13
sync-ldap-collector	11
documents	10
generatemotd	7
pbsmon2php-server	7
pbsacc2db	5
vsc-postgres	5
pbsmon2php-users	1
reposnap	1

github.com/easybuilders	
Repository	#commits
easybuild-easyconfigs	1090
easybuild-framework	379
easybuild-easyblocks	184
easybuild	129
eb-singularity	1

github.com/quattor	
Repository	#commits
configuration-modules-core	224
CAF	46
maven-tools	39
aii	33
configuration-modules-grid	28
pan	20
release	20
ncm-ncd	13
ncm-lib-blockdevices	11
CCM	11
template-library-core	8
quattor.github.com	7
ncm-cdispd	5
cdp-listend	2
aquilon	1

4 TRAINING AND SUPPORT

4.1 Training overview and evaluations



4.1.1 Introduction to HPC @ Ghent University

Trainer: Kenneth Hoste (HPC-UGent)

Date	#participants
23/01/2017	27
03/04/2017	12
17/05/2017	22
06/09/2017	21
10/11/2017	24

Satisfaction scores ¹	
Training content	98%
Lecturer	94%
Hands-on	61%

Selected suggestions for improvement and remarks:

- Improve referral to documentation
- Work on a common problem with all students together
- Focus on a specific application in particular
- The support we received afterwards was excellent. Give the guy and his team a pay raise.
- More trainers during the hands-on session

¹ These satisfaction scores indicate positive response of evaluation respondents on questions "The presentation gave me all the information I wanted" (Yes or Mostly), "The lecturer presented well" (Strongly Agree or Agree) and "The hands-on session was in accordance with my expectations" (Strongly agree or Agree).

4.1.2 Introduction to Linux

Trainer: Kenneth Hoste (HPC-UGent)

Date	#participants
15/03/2017	25
13/10/2017	28

Satisfaction scores ¹	
Training content	100%
Lecturer	90%
Hands-on	81%

Selected suggestions for improvement and remarks:

- Reduce the speed a bit
- Reduce length of lecture and increase the tempo
- Provide to participants excercises to complete beforehand
- Split up the course in two days
- Very good session, I liked it! Thank you!

4.1.3 Specialist Workshops in Scientific Computing – Introduction to multithreading and OpenMP

Trainer: Reinhold Bader (Leibniz Supercomputing Centre LRZ, Germany)

Date	#participants
18-19/04/2017	13

Satisfaction scores ²	
Training content	100%
Lecturer	90%

Selected suggestions for improvement and remarks:

- Nice alteration of theory and excercises
- Important tutorial for many problems that we solve during our PhD work
- Well-organized, good location, high-level professional teaching the course
- Some exercises I found too complicated

4.1.4 Specialist Workshops in Scientific Computing – Introduction to MPI

Trainer: Jan Fostier (INTEC department of Information Technology, UGent)

Date	#participants
05/05/2017	11

Satisfaction scores ²	
Training content	100%
Lecturer	100%

Selected suggestions for improvement and remarks:

- More real scenario applications
- Additional hands-on exercise session
- Everything was great

² These satisfaction scores indicate positive response of evaluation respondents on questions “In general I judge this course as” (Very good or Good) and “The lecturer has the specialized expertise required to teach this course” (Fully agree or Agree).

4.1.5 Specialist Workshops in Scientific Computing – Introduction to Machine Learning Algorithms

Trainer: Morris Riedel (Juelich Supercomputing Centre JSC, Germany & School of Natural Sciences and Engineering of the University of Iceland)

Date	#participants
23-24/11/2017	38

Satisfaction scores ²	
Training content	100%
Lecturer	100%

Selected suggestions for improvement and remarks:

- Very good facilities prepared for the course (HPC, room, screens)
- Good combination of theory and hands-on
- Good infrastructure (HPC)
- No coding required, only applications
- Knowledge of Linux and HPC commands required beforehand
- Additional hands-on session with exercises

4.1.6 Specialist Workshops in Scientific Computing – Deep Learning using a Convolutional Neural Network

Trainer: Morris Riedel (Juelich Supercomputing Centre JSC, Germany & School of Natural Sciences and Engineering of the University of Iceland)

Date	#participants
30/11-01/12/2017	38

Satisfaction scores ²	
Training content	100%
Lecturer	100%

Selected suggestions for improvement and remarks:

- The course was very inspiring
- Good speaker, good coverage of topic
- I would have preferred more days of seminars

4.1.7 Suggestions from participants for new training sessions

- Parallelization on GPU
- OpenCL
- Parallel Programming and Big Data for HPC
- Fortran
- Neural networks
- Machine learning
- Object based image analysis
- Databases like mySql or similar
- Text mining
- Statistical data analysis
- More applications for Earth Sciences

4.2 Lectures and community meetings

4.2.1 Singularity: Containers in HPC

Lecturer: Michael Bauer (University of Michigan & GSI national laboratory Darmstadt, Germany)

Date	#participants
7/02/2017	21

4.2.2 Tier1 information session

Learn more about the VSC Tier-1 infrastructure and how to write a good proposal.

Date	#participants
4/09/2017	14

4.2.3 OpenFOAM user meeting

This user meeting was organized to bring together various research groups using the popular open source CFD software OpenFOAM on the VSC infrastructure. They showcased their work, and helped propagate collaboration with best practices for using this software on HPC clusters.

Lecturers:

- Peter De Jaeger (Bekaert)
- Brecht Devolder (Department of Civil Engineering, UGent and KU Leuven)
- Pieter Reyniers & Laurien Vandewalle (Laboratory for Chemical Technology, UGent)
- Joris Degroote (Department of Flow, Heat and Combustion Mechanics, UGent)

Date	#participants
13/09/2017	20

Satisfaction scores ³	
Meeting content	100%
Hands-on	93%

Selected feedback:

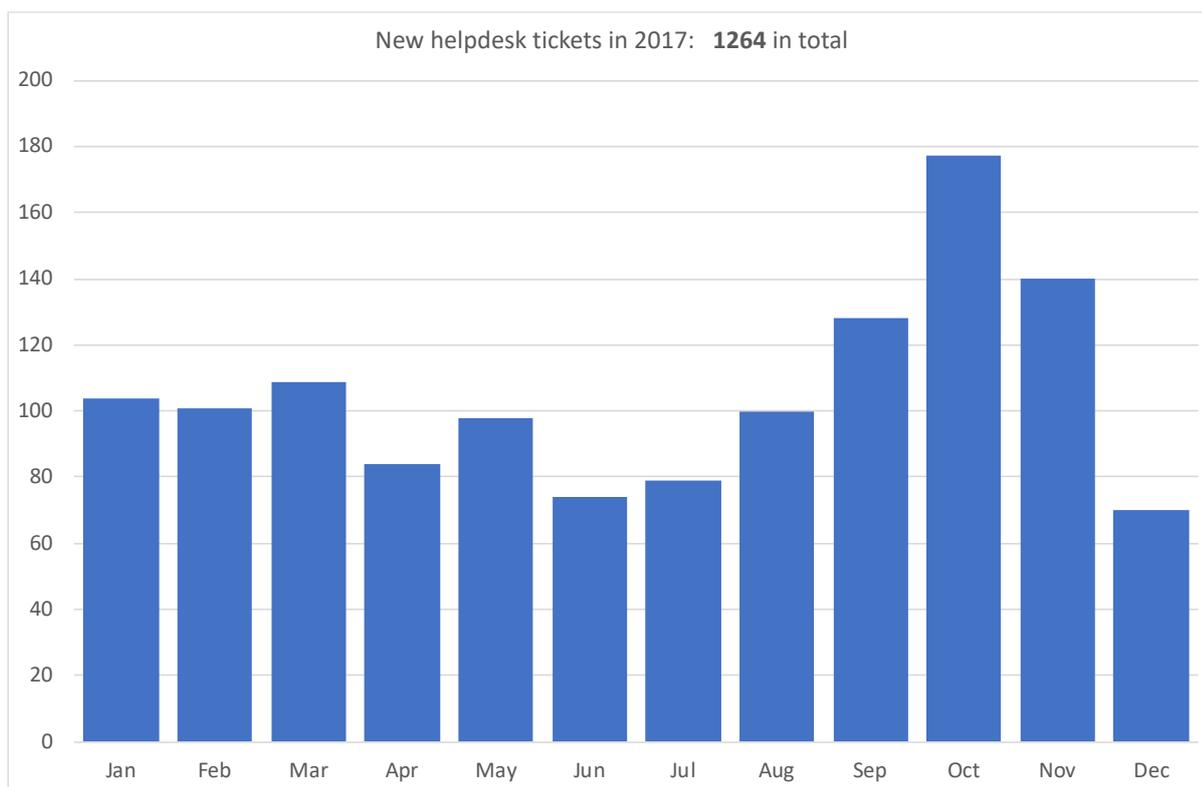
- The diversity of the application domains was an advantage, in particular the example from industry.
- Hands-on was super good, for beginning developers and experienced developers to have another perspective on the development structure.
- A nice initiative and hopefully it can become an annual event. Thx!!
- The hands-on sessions should be more OpenFOAM oriented, instead of how to use it at UGent (that can be learned during the use of OpenFOAM).

³ These satisfaction scores indicate positive response of evaluation respondents on questions "Overall, the OpenFOAM user meeting was useful/interesting to me" (Yes or Mostly) and "The hands-on sessions were useful to me" (Yes or Mostly).

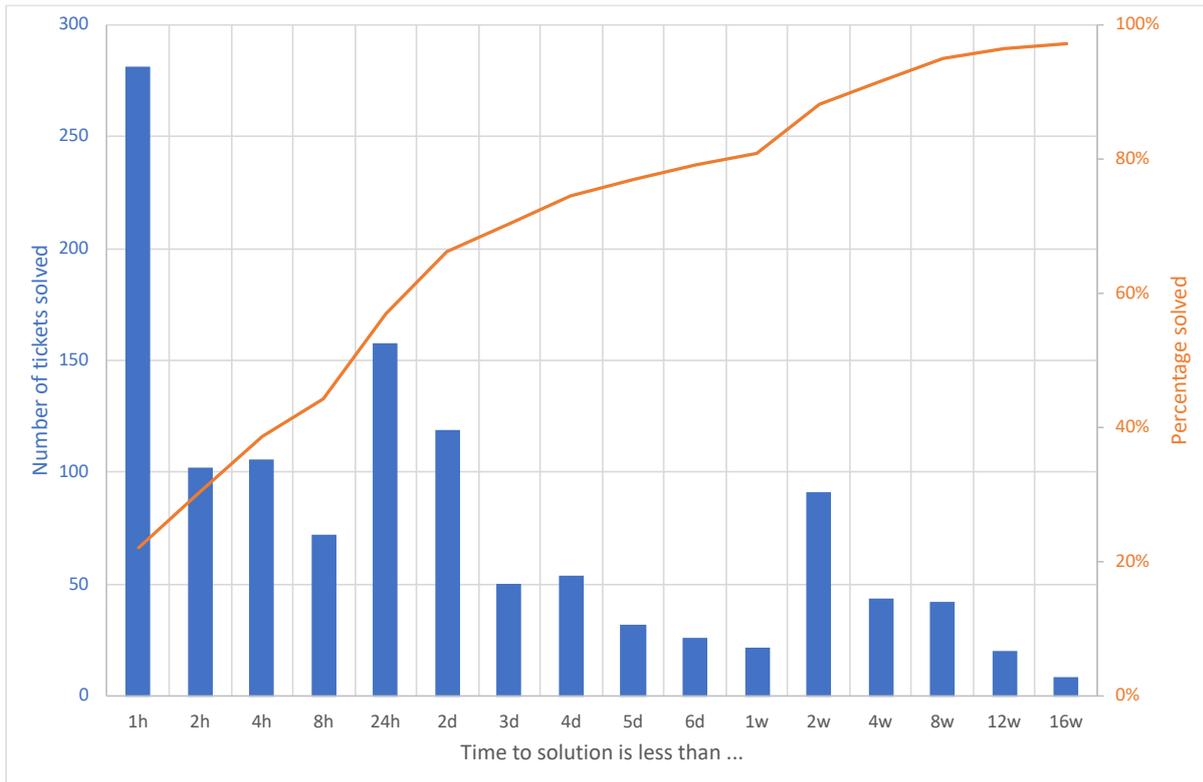


4.3 Helpdesk

4.3.1 Number of new tickets per month



4.3.2 Average time to resolution



4.4 User evaluation

In order to improve HPC-UGent services, all current users were asked to complete a user survey. This survey was anonymous, short (average completion time 5 minutes) and in all 104 users responded. A selection of questions and responses is shown.

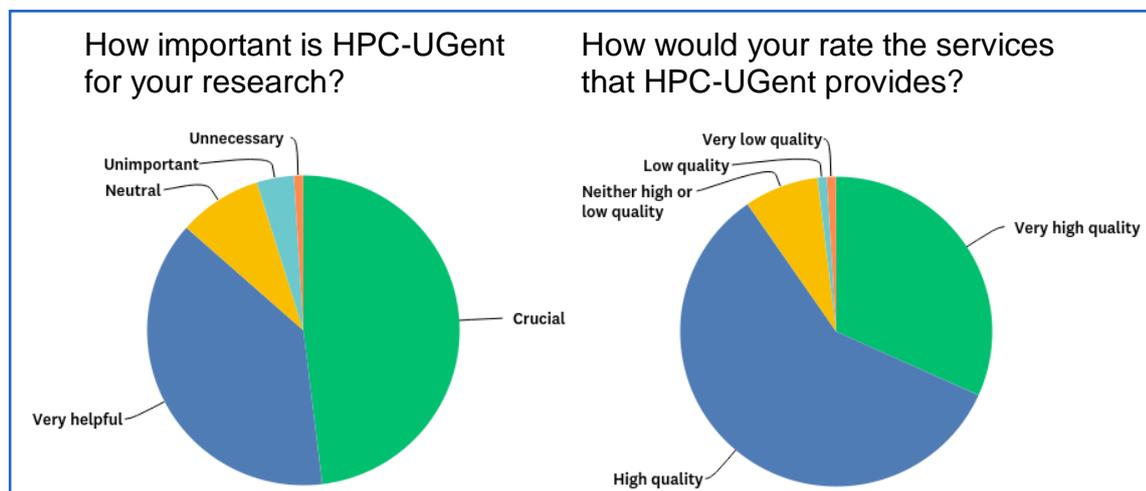
How often do you use the HPC-UGent scientific computing infrastructure?		
	#responses	%
On a daily basis	41	39%
Usually once per week	22	21%
Once per month	22	21%
1-2 times per year or less	19	18%

How important is HPC-UGent for your research?		
	#responses	%
Crucial - I can't do my research without	50	48%
Very helpful - It allows me to do my research at a faster pace and at a higher level	40	38%
Neutral	9	9%
Unimportant - I can just as well do my research in another way	4	4%
Unnecessary - I don't need it at all to do my research	1	1%

How would you rate the services that HPC-UGent provides?

(compute clusters, login nodes, training, user support, website)

	#responses	%
Very high quality	33	32%
High quality	61	59%
Neither high or low quality	8	8%
Low quality	1	1%
Very low quality	1	1%



Is there a particular service that stands out or you care to comment about?

Service	#positive comments	#negative comments
User support	16	
Training	8	
Infrastructure	6	5
Software	5	1
Documentation	4	1

"The clusters in Ghent are still the best managed systems I've been working with. Keep up the nice work guys!"

"The Helpdesk is extremely supportive and helpful! Thank you!"

"The consistent use of the system status page in the User Portal of the VSC website is much appreciated. It is good to have this page as the single point of information for cluster related updates and during upsets."

How could we further improve HPC-UGent services?	
Topic	#suggestions
Documentation	21
<ul style="list-style-type: none"> • More structured wiki/website that is SPOI • Documentation better geared towards new users <ul style="list-style-type: none"> ○ Checklist for starters ○ Different clusters + storage locations ○ Frequently used commands • More examples: scripts, software • Updated documentation • Newsletter 	
User experience	21
<ul style="list-style-type: none"> • Shorter queue times • Way to estimate queue time • Longer wallclock time • Feedback on efficiency of specific job 	
Infrastructure	17
<ul style="list-style-type: none"> • Larger compute power • GPU • Debug infrastructure, remote visualization 	
Data	10
<ul style="list-style-type: none"> • More storage • Easier sharing of data <ul style="list-style-type: none"> ○ Between VOs ○ Public • Staging in/out data • Archive data of inactive users 	
Training	7
<ul style="list-style-type: none"> • More 	
Security	6
<ul style="list-style-type: none"> • Ftp connectivity • Node sharing, data visible to other users 	
Policy	4
<ul style="list-style-type: none"> • Collaboration within VSC • Keep services free of charge 	
Software	3
<ul style="list-style-type: none"> • Form to request software install/update 	
User support	3
<ul style="list-style-type: none"> • More HPC-UGent staff 	

Would your research benefit from specific IT hardware or services that HPC-UGent currently does not provide?		
	#responses	%
No – The current compute platform suffices	52	51%
I don't know	29	28%
Yes – I need specific hardware/services	22	21%
<ul style="list-style-type: none"> • Bigdata cluster with fast I/O • Large, shared storage (many TB) • Very high memory machine (1 TB) • GPU • More cores / node • Other compilers than Intel • Longer wallclock time • Personnel to support coding design • Additional personnel for software installs 		

5 OUTREACH

5.1 Within Ghent University

Date	Event
6/02/2017	Information session FWO call 2017 for Research Grants, Research Projects and SBO
9/05/2017	HPC introduction and UGent datacenter tour for students of course 'Computer Systems' (UGent)
8/11/2017	UGent datacenter tour for students of faculty of economics and business administration (UGent)
5/12/2017	HPC introduction and UGent datacenter tour for students Bachelor of Science in Engineering Technology (UGent)

5.2 To policy makers, industry and general public

Date	Event
24/01/2017	Visit to Ghent University datacenter of MEPs from European Parliamentary Commission ITRE (Industry, Technology, Research and Energy), European Commission and representatives of European network and lobby organisation in the field of R&D&I. Gent, Belgium
8/03/2017	Presentation on board of directors of BiR&D (Belgian Industrial Research and Development association) <i>Supercomputing for academics & industry - Flemish Supercomputer Center (VSC)</i> Vilvoorde, Belgium http://birdbelgium.com/
27/04/2017	TechBoost! 2017 <i>Innovations in Big Data Analytics</i> Gent, Belgium https://aiq.ugent.be/event-2017/techboost-voor-ingenieurs
8/10/2017	Iedereen UGent <i>Try out the supercomputer! Experiment market</i> Gent, Belgium https://iedereenugent.be/en/activities/try-out-the-supercomputer
10/10/2017	HPC introduction and UGent datacenter tour for CESAER Task Force Research Infrastructures (TFRI) Gent, Belgium http://www.cesaer.org/en/projects/task-force-research-infrastructures-tfri

Highlighted event:

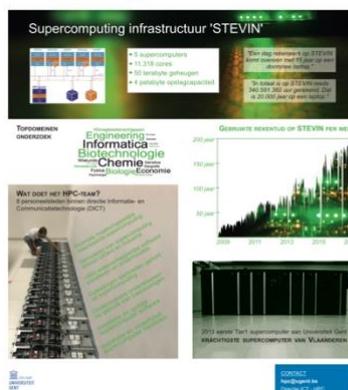
Try out the supercomputer! - Experiment market

(at the occasion of the 'Iedereen UGent' event, 8 October 2017)

"Ghent University has been hosting the Tier-1 Supercomputer since 2012. Have you always wondered what a supercomputer is, how it works and what you can do with it? This is your chance! See a supercomputer live at work, enjoy an interactive demo yourself to calculate currents or predict the weather, and discover what you can do with 200 years of calculating time. A unique chance to have a taste of high performance computing!"

HPC-UGent specifically constructed a portable (only 300 kg) mini-cluster 'ditto', on which 3 interactive cases from UGent research groups were presented:

- Atmospheric Physics - Predict the climate for your region for the next 80 years
<https://www.ugent.be/we/physics-astronomy/en/research/atmophys>
- Fluid Mechanics - Tweak your drone so that it flies further
<http://www.floheacom.ugent.be>
- Cancer Research Institute Ghent - Determine the cancer type for an unknown patient
<https://www.crig.ugent.be>



5.3 Within international HPC community

Date	Event
4-5/02/2017	FOSDEM'17 Co-organisation of <i>HPC, Big Data and Data Science</i> devroom Brussels, Belgium https://archive.fosdem.org/2017/schedule/track/hpc,_big_data_and_data_science
8-10/02/2017	2nd EasyBuild User Meeting Lecture <i>EasyBuild: past, present & future</i> Tutorial <i>Implementing easyblocks</i> Jülich Supercomputing Centre, Germany https://indico-jsc.fz-juelich.de/event/30
21-22/03/2017	23rd Quattor workshop Annecy, France https://www.quattor.org/meeting/2017/03/21/annecy-workshop.html
25-27/04/2017	PRACE 2017 Spring School - System Administration and Data/Computational Services for Scientific Communities <i>Modern Scientific Software Management using EasyBuild& co</i> The Cyprus Institute, Nicosia, Cyprus https://events.prace-ri.eu/event/601
17-18/05/2017	SESAMENET partner meeting Barcelona, Spain https://sesamenet.eu
13-15/06/2017	iRODS User Group Meeting 2017 Utrecht, the Netherlands https://irods.org/ugm2017/
15-16/06/2017	HPC Knowledge Meeting '17 <i>5 years of getting scientific software installed using EasyBuild</i> San Sebastián, Spain http://www.hpckp.org/index.php/annual-meeting/hpckp-17
23-24/10/2017	OpenNebula Conf 2017 <i>Transforming an Old Supercomputer into a Cloud Platform</i> Madrid, Spain http://2017eu.opennebulaconf.com
22-27/10/2017	SPLASH 2017 <i>SIGPLAN Most Influential OOPSLA Paper Award for 'Statistically rigorous Java performance evaluation' by Andy Georges, Dries Buytaert, Lieven Eeckhout, OOPSLA 2007.</i> Vancouver, Canada https://2017.splashcon.org/attending/splash-awards
13-14/11/2017	EasyBuild hands-on meeting SurfSARA, Amsterdam, the Netherlands
30/11-1/12/2017	Digital infrastructures for research 2017 Brussels, Belgium https://www.digitalinfrastructures.eu/di4r2017/about

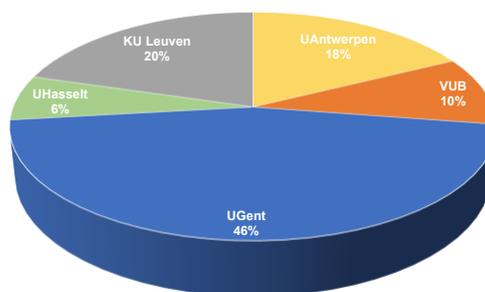
6 BUDGET

Budget line	UGent budget statement	Income statement
Central compute infrastructure	1.200.000	
Central compute infrastructure (wages)	675.000	
Central compute infrastructure training and events	50.000	
FWO subsidy for personnel (Tier1+2)		665.000
FWO subsidy for investment and operations		1.198.150

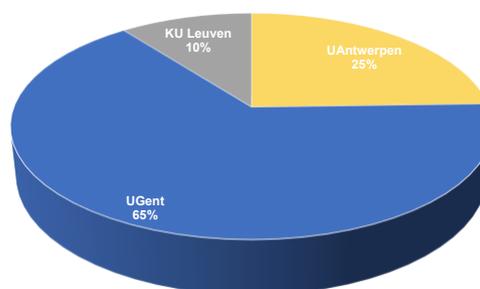
7 TIER-1 USAGE

Every year, three calls are open to apply for access to the Tier-1 system of the Flemish Supercomputer Center (VSC). HPC-UGent stimulates and assists researchers as much as possible to apply for access. The tables and graphs below give insight in the number of proposals (#proj) and the awarded compute time (#nd, expressed in Tier1 nodedays), listed per institute.

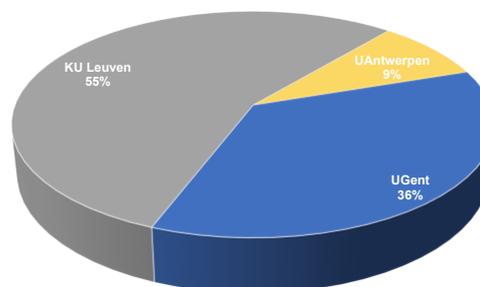
Call 1		Requested		Granted	
6/02/2017					
Institute	#proj	#nd	#proj	#nd	
UAntwerpen	3	9216	2	8544	
VUB	3	7233	2	4733	
UGent	14	45933	8	22153	
UHasselt	2	7780	1	3100	
KU Leuven	8	23889	2	9884	



Call 2		Requested		Granted	
5/06/2017					
Institute	#proj	#nd	#proj	#nd	
UAntwerpen	3	14709	3	14709	
VUB	0	0	0	0	
UGent	13	40801	12	39001	
UHasselt	0	0	0	0	
KU Leuven	7	14269	6	6269	



Call 3		Requested		Granted	
2/10/2017					
Institute	#proj	#nd	#proj	#nd	
UAntwerpen	2	8177	1	4928	
VUB	0	0	0	0	
UGent	7	20331	7	20331	
UHasselt	1	2100	0	0	
KU Leuven	13	37206	9	31234	



Percentage of awarded compute time, listed per institute per call

8 USER IN THE SPOTLIGHT

Hannes Mareen (IDLab, UGent-imec) uses supercomputing to help fight video piracy.



IDLab – Internet Technology and Data Science Lab, Ghent University - imec
Hannes Mareen is a PhD researcher at IDLab, a core research group of imec whose research activities are embedded in Ghent University and University of Antwerp. IDLab is involved in various research areas, including multimedia processing, machine learning and data mining, wireless networking, ... and actively targets the most daunting challenges industry is facing to deliver the digital transformation, connecting everything and extracting high value from data.

Computational research and applications

“Tackling digital piracy is a major concern to the film industry. Not only is it estimated to have a multi-billion-euro economic impact, piracy also deters content creators to invest in new content, resulting in less choice for consumers.” says Hannes Mareen.

“In my research, I have devised a method to identify digital pirates who illegally share videos on the internet.

When a video is legally sent to a customer or reviewer who is allowed a preview, a unique fingerprint or watermark is hidden inside the video. This way, when the video is illegally distributed online, it can be traced back to the malicious customer/reviewer.”

“I have developed a novel method which ensures that the watermark is invisible and impossible to remove, by altering the small errors that are made by the encoder that compresses the video to a manageable storage size. By deliberately introducing one small change during the encoding process, a unique combination of tiny errors is spread over the entire video. Although the collection of introduced errors is huge, they are so subtle that viewers cannot detect or recognize this watermark, let alone remove it.

But when such a video would be illegally shared, the unique combination of encoding errors in it can be easily read out. In this way, the watermark can be used to identify the source of the illegal activity.

Furthermore, a clever encoder ensures that adding a watermark in this fashion is fast, and so is useable on large scale. I am currently refining this technique as part of my PhD grant strategic basic research with FWO.”

Awards

Hannes performed this research as part of his Master’s dissertation ‘A novel video watermarking approach based on implicit distortions’, supervised by prof. Peter Lambert, Glenn Van Wallendael and Johan De Praeter.

He received multiple awards for this work:

- Best Poster award at the UGent Faculty of Engineering and Architecture Research Symposium (FEARS) 2018.
- Winner of Agoria award 2017 (as part of ‘de Vlaamse Scriptieprijs’). This award is given to a graduate student with the best Master's thesis concerning technology.

HPC-UGent facilitates computational research

Hannes relied on the vast computational resources of HPC-UGent to test new compression algorithms, study efficient detection methods and develop his novel watermark technique. Overall, he used close to 422.000 core hours of compute time on the Tier-2 clusters of Ghent University to complete his Master’s thesis. This would correspond to using a current-day laptop for more than 12 years, non-stop.

Links

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

Publications

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