

# ANNUAL REVIEW HPC-UGENT

2019



Department ICT, Infrastructure office

E [hpc@ugent.be](mailto:hpc@ugent.be)  
T +32 9 264 4716

Campus Sterre S9  
Krijgslaan 281, 9000 Gent

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

# CONTENTS

<b>CONTENTS</b>	<b>2</b>
<b>1 About HPC-UGent</b>	<b>3</b>
1.1 Our mission	3
1.2 Our vision	3
1.3 Personnel	3
<b>2 Infrastructure</b>	<b>4</b>
2.1 Overview	4
2.2 Usage	6
<b>3 Development and maintenance</b>	<b>10</b>
3.1 Number of JIRA issues resolved in 2019, grouped per main component	10
3.2 Projects finalized in 2019	10
3.3 Github PRs in 2019 by HPC-UGent staff, per repository	11
<b>4 Training and support</b>	<b>12</b>
4.1 Training overview and evaluations	12
4.2 Lectures and community meetings	15
4.3 Helpdesk	17
4.4 User evaluation	19
<b>5 Outreach</b>	<b>25</b>
5.1 Within Ghent University	25
5.2 Within international HPC community	25
<b>6 Budget</b>	<b>27</b>
<b>7 Tier-1 usage</b>	<b>28</b>
<b>8 User in the spotlight</b>	<b>29</b>

# 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 at the start of 2019 consists of 7 people:

Álvaro Simón García, Andy Georges, Ewald Pauwels, Kenneth Hoste, Kenneth Waegeman, Stijn De Weirdt, Wouter Depypere.

In the course of 2019, two people joined the team: Balázs Hajgató (February) and Bart Verheyde (August).

Tasks include:

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

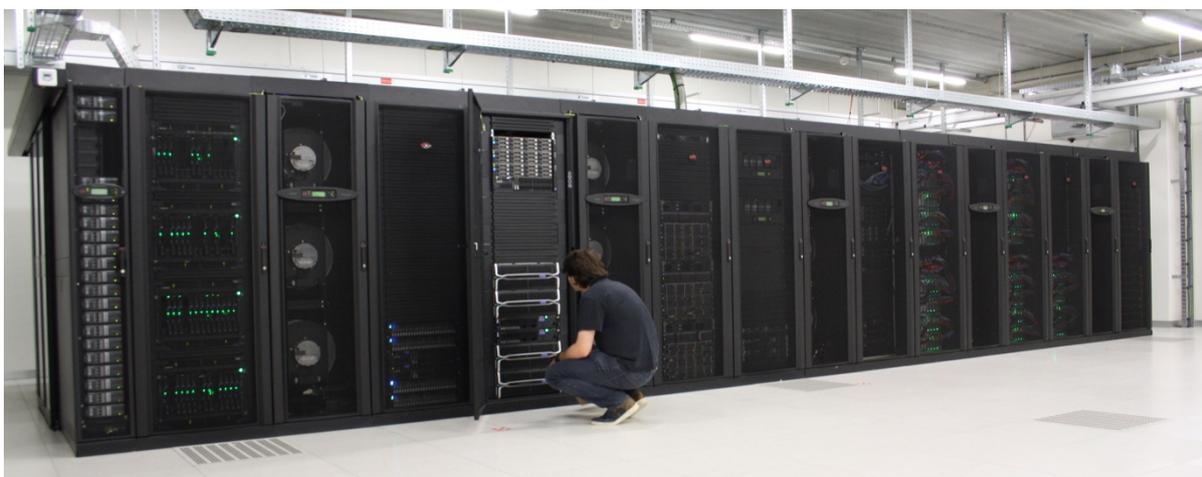
## 2 INFRASTRUCTURE

### 2.1 Overview

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

In 2019, several changes to the infrastructure were applied:

- the SLURM job scheduler was introduced on all clusters
- cluster delcatty was decommissioned
- a new GPU cluster, named joltik, was installed and put into pilot phase.



#### 2.1.1 Compute clusters

Cluster name	#nodes	CPU / GPU per node	Usable memory per node	Local disk per node	Network interconnect
<b>Delcatty</b>	124	2 x 8-core Intel E5-2670 (Sandy Bridge @ 2.6 GHz)	64 GB	400 GB	FDR InfiniBand
			Decommissioned in two stages: first part 26 Aug, second part 17 Oct		
<b>Phanpy</b>	16	2 x 12-core Intel E5-2680v3 (Haswell-EP @ 2.5 GHz)	480 GiB	3 x 400 GB (SSD, striped)	FDR InfiniBand
<b>Golett</b>	200	2 x 12-core Intel E5-2680v3 (Haswell-EP @ 2.5 GHz)	53 GiB	500 GB	FDR-10 InfiniBand
<b>Swalot</b>	128	2 x 10-core Intel E5-2660v3	116 GiB	1 TB	FDR InfiniBand

		(Haswell-EP @ 2.6 GHz)			
<b>Skitty</b>	72	2 x 18-core Intel Xeon Gold 6140 (Skylake @ 2.3 GHz)	177 GiB	1 TB	EDR InfiniBand
<b>Victini</b>	96	2 x 18-core Intel Xeon Gold 6140 (Skylake @ 2.3 GHz)	88 GiB	1 TB	10 Gb ethernet
<b>Joltik</b>	10	2 x 16-core Intel Xeon Gold 6242 (Cascade Lake @ 2.8 GHz)  4x NVIDIA Volta V100 GPUs (32GB GPU memory)	256 GiB	800 GB SSD	Double EDR InfiniBand

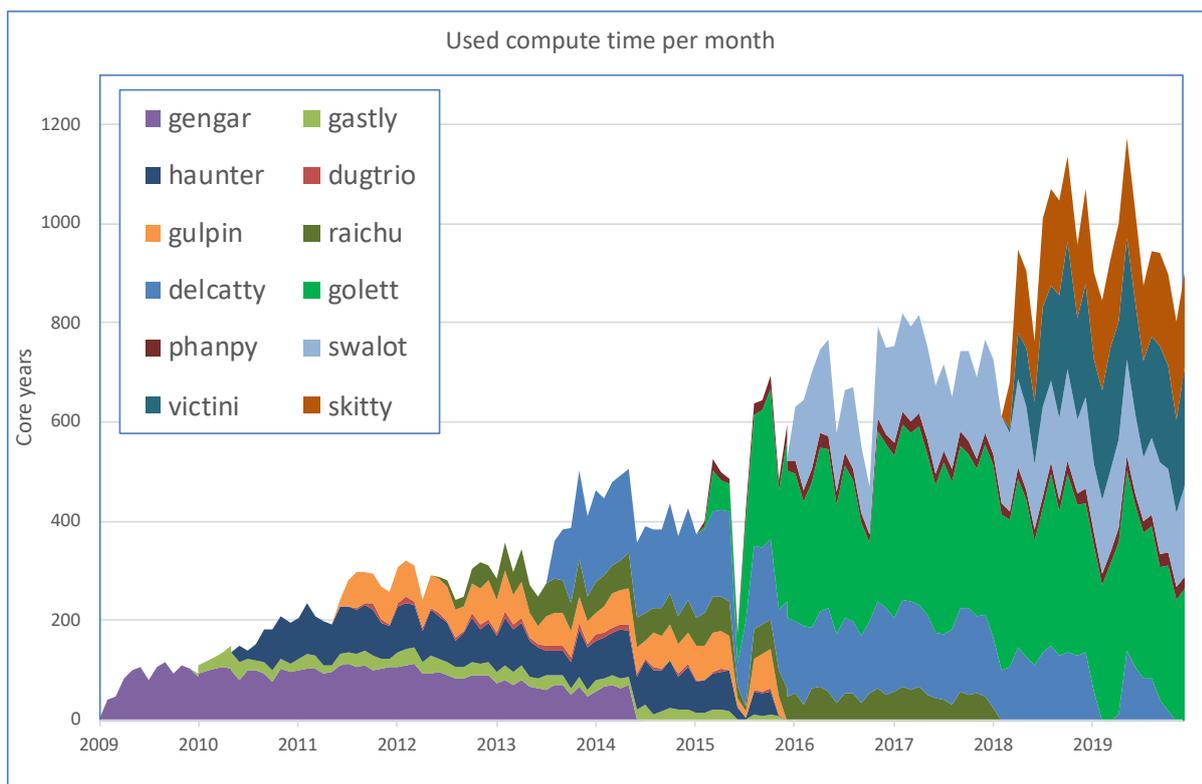
### 2.1.2 Shared storage

Partition	Size
<b>\$VSC_HOME</b>	51 TB
<b>\$VSC_DATA</b>	1.8 PB
<b>\$VSC_SCRATCH</b>	1.9 PB
<b>\$VSC_SCRATCH_KYUKON</b>	
<b>\$VSC_SCRATCH_PHANPY</b>	39 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 (365 days). One core hour is the work done by one core in one hour of time.

### 2.2.1 Historical perspective on used compute time

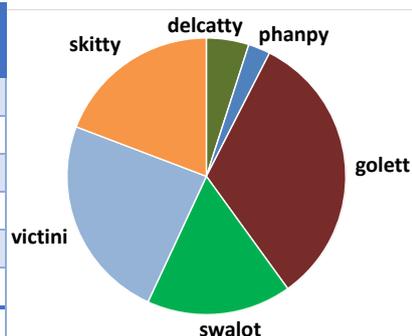


This plot gives a historical perspective on the amount of compute time that is used every month, coloured according to the cluster that does the work. All clusters together, since the start in 2009, have now produced more than 560 million core hours. This is equivalent to 64.000 years of compute work on one core.

## 2.2.2 Consumed compute time in 2019

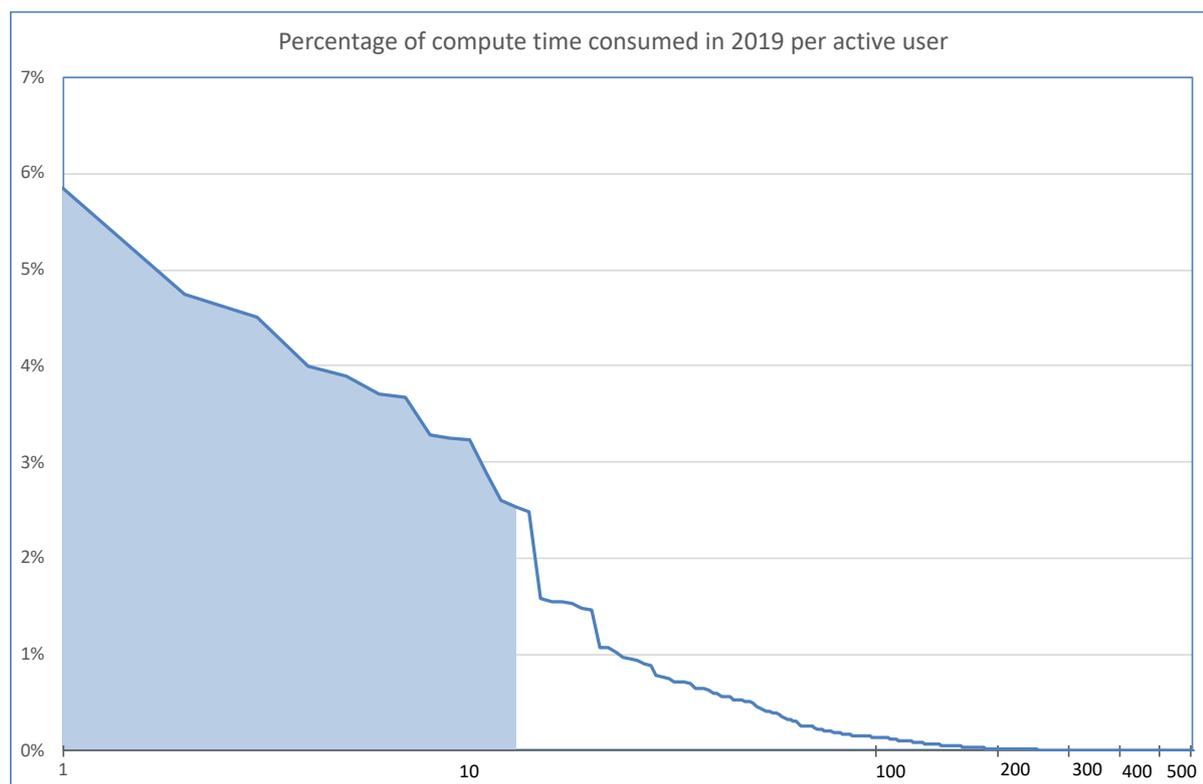
### 2.2.2.1 Consumed compute time per compute cluster

Cluster name	Compute time consumed (in core years)	Effective use percentage
<b>Delcatty</b>	<i>(until October 2019)</i> 548	33%
<b>Phanpy</b>	297	77%
<b>Golett</b>	3590	75%
<b>Swalot</b>	1957	76%
<b>Victini</b>	2650	77%
<b>Skitty</b>	2199	85%
<b>Total</b>	<b>11242</b>	<b>73%</b>



In 2019, a total of 98.478.810 core hours has been consumed on the Tier-2 compute clusters of Ghent University. This corresponds to 11.242 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.

### 2.2.2.2 Consumed compute time versus active users



In 2019, 608 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 2019 this user consumed. The blue area indicates that 51% of all Tier-2 compute time was consumed by 14 power users.

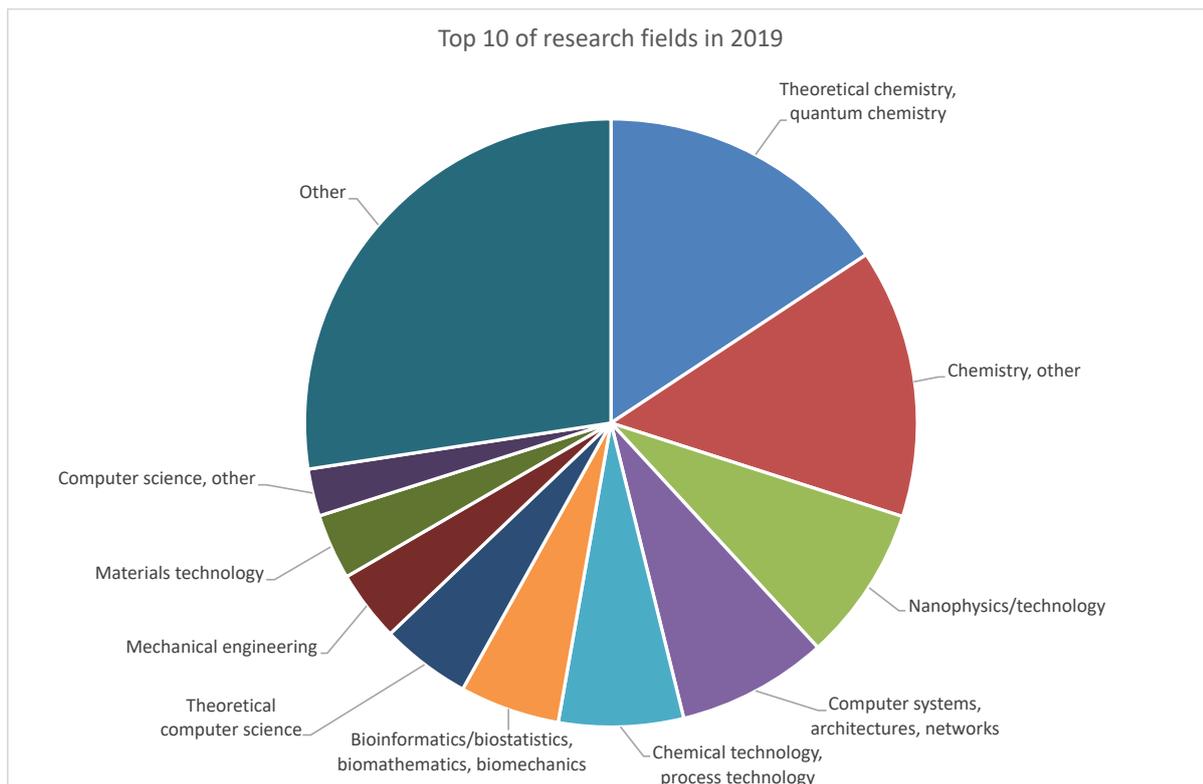
Everyone can get access to the Tier-2 compute clusters of Ghent University. The conditions that apply depend on the affiliation of the researcher. Researchers affiliated with Flemish university associations (constituting the Flemish Supercomputer Center – VSC) get free access. Other Flemish or federal 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.

<b>Breakdown of consumed compute time by affiliation</b>	
<b>UAntwerpen</b>	1.730%
<b>VUB</b>	0.045%
<b>UGent</b>	97.881%
<b>KULeuven / UHasselt</b>	0.181%
<b>Other research institutes</b>	0.162%
<b>Industry</b>	0.000%
<b>Total</b>	<b>100.000%</b>

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.

<b>Number of students/researchers versus breakdown of consumed compute time by category</b>		
<b>Ma/Ba students</b>	77	0.001%
<b>Researchers</b>	531	99.999%
<b>Total</b>	<b>608</b>	<b>100.000%</b>

### 2.2.2.3 Consumed compute time versus research field

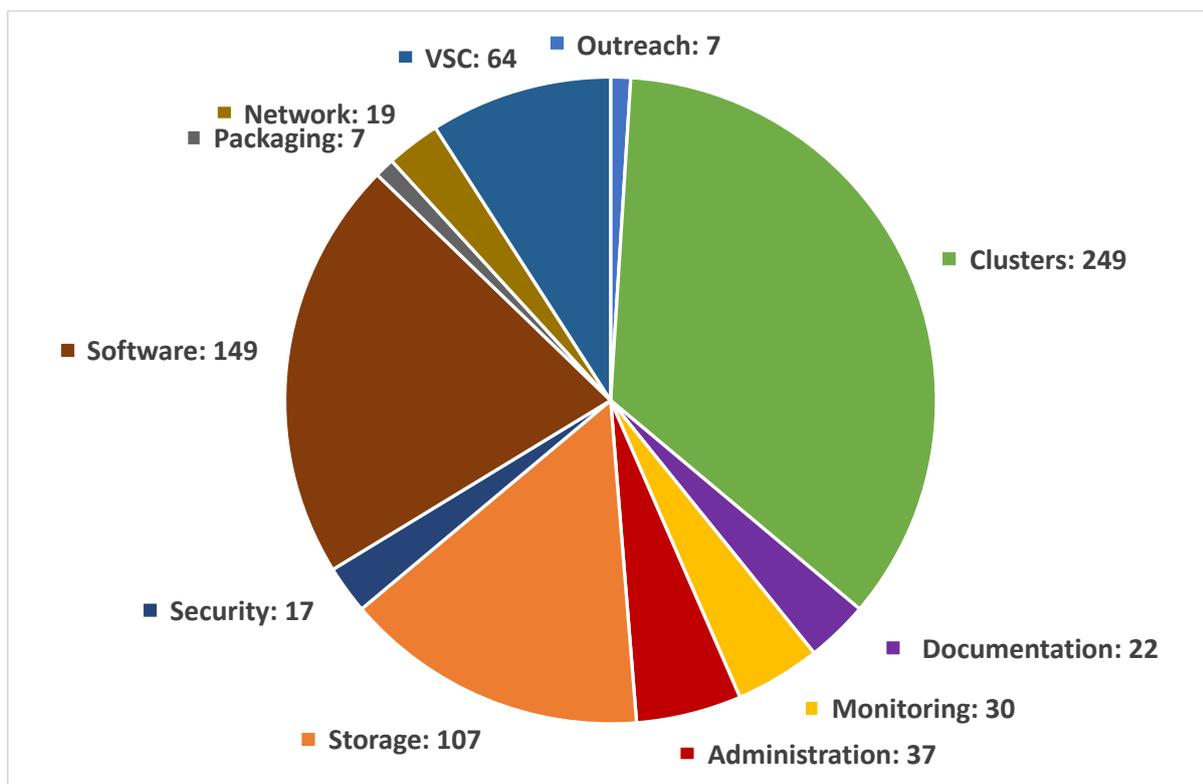


These are the research fields that users enter when requesting their account to access HPC-UGent resources. No doubt, there is overlap between certain fields and some users have not listed a research field at all. As such, the above distribution merely gives an indication of the top research fields that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

### 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 'pull requests' (PRs) 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 2019, grouped per main component



#### 3.2 Projects finalized in 2019

Issue #	Title	Summary
HPC-7372	Tier2 site update 2019Q2	General maintenance of all Tier2 clusters and servers.
HPC-7470	All things slurm	Slurm job scheduler was implemented on all clusters.

### 3.3 Github PRs in 2019 by HPC-UGent staff, per repository

github.ugent.be/hpcugent	
Repository	#PRs
android-systems-locator	3
clusterbuildrpm-server	12
discovery	1
documents	10
flexlm	1
generatemotd	4
icinga-checks	23
legacy-rpms	48
oncall-service	16
pbsacc2db	1
pbsmon2php-server	4
quattor-host-mngmt-tools	9
reposnap	4
rundeck	8
service_sanity	2
sync-ldap-collector	1
vsc-backup	2
vsc-burstbuffer	12
vsc-cloud	13
vsc-cluster-modules	11
vsc-data	1
vsc-freeipa	4
vsc-host-tools	8
vsc-ood	29
vsc-profiles	2
vsc-project	3
vsc-quattor	15
vsc-rpms	2
vsc-testing	2

quattor	690
slurm-prepilogue	4
slurm-spank-talamini	3
slurm-torque-wrappers	24
slurm	11
spank-private-tmp	1
vsc-accountpage-clients	5
vsc-accountpage	46
vsc-administration	11
vsc-base	10
vsc-config	22
vsc-fileystems-quota	3
vsc-fileystems	11
vsc-install	29
vsc-jobs	14
vsc-ldap-extension	1
vsc-manage	7
vsc-modules	6
vsc-mympirun	6
vsc_user_docs	27
vsc-utils	6
vsc-zk	1
yoda-packages	2

github.com/easybuilders	
Repository	#PRs
easybuild-framework	218
easybuild-easyblocks	130
easybuild-easyconfigs	728

github.com/hpcugent	
Repository	#PRs
csub	1
django-wayf	2
eb_inuits	26
gpfsbeat	1
hanythingondemand	3
hnormalise	4
jobcli	34
Lmod-UGent	5
nhc	2
openstack-templates	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
9/01/2019	20
20/03/2019	8
27/03/2019	23

Satisfaction scores <sup>1</sup>	
Training content	100%
Lecturer	96%
Hands-on	83%

Trainer: Ewald Pauwels (HPC-UGent)

Date	#participants
11/10/2019	30

Satisfaction scores <sup>1</sup>	
Training content	100%
Lecturer	100%
Hands-on	76%

Selected suggestions for improvement and remarks:

- Extra focus on non-data outputs (e.g. Python pickle for a model)
- More (applied) exercises
- The visit to the datacenter was a cool addition!
- I enjoyed the hands-on session where we could discuss problems that would have taken me so much longer to solve on my own.
- A good preparation in advance for participants would be basic Linux youtube videos to watch.
- Overall great and engaging session! Maybe it'd be nice to have a longer course with more hands-on examples.

<sup>1</sup> 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: Christopher Peeters (ITGilde)

Date	#participants
25/01/2019	20
04/11/2019	15

Satisfaction scores <sup>1</sup>	
Training content	94%
Lecturer	81%
Hands-on	47%

Selected suggestions for improvement and remarks:

- The pdf on Linux is very good.
- Difficult to keep following the lecture. It felt for me like the lecturer was shooting all Linux commands to me, just one after the other.
- Sometimes there was the lack of some basic explanation to understand/follow the lecture properly.
- The last part of the course with additional things you can do with Linux was brushed over. Maybe there should be an intermediate level Linux course?
- Introducing how to apply for an HPC account and how to use HPC should be covered in the 'Introduction to HPC' training, not in this Linux course.

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

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

Date	#participants
2-3/04/2019	19

Satisfaction scores <sup>2</sup>	
Training content	90%
Lecturer	100%

Selected suggestions for improvement and remarks:

- The professor is an expert and attended all students that asked questions including during exercises, helping to understand the different approaches.
- Very good pace of the lecture.
- It needs strong background in computer science.

### 4.1.4 Specialist Workshops in Scientific Computing – Introduction to MPI

Trainer: Jan Fostier (INTEC, IDLab, UGent & imec)

Date	#participants
24/04/2019	10

Satisfaction scores <sup>2</sup>	
Training content	100%
Lecturer	100%

Selected suggestions for improvement and remarks:

- The lecturer is not only an expert in this field, he also has a clear way of explaining the matter in a way that is particularly satisfying to follow: the

<sup>2</sup> 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).

structure was well-thought of and the pace was perfect for an introductory workshop.

- Very good lecturer, very clear presentation, exercises had right difficulty curve.
- Since this doctoral school was so interesting, I wouldn't have minded another day with more exercises.
- I would have liked to have had the opportunity to use MPI through a Python library, since that is what I will use in practice.

#### 4.1.5 Specialist Workshops in Scientific Computing – Introduction to CP2K

Trainers:

- Juerg Hutter (University of Zurich, Switzerland)
- Marcella Iannuzzi (University of Zurich, Switzerland)
- Matt Watkins (University of Lincoln, UK)
- Tiziano Mueller (University of Zurich, Switzerland)

Date	#participants
11-13/03/2019	24

Satisfaction scores <sup>2</sup>	
Training content	100%
Lecturer	88%

Selected suggestions for improvement and remarks:

- As the developers of CP2K were present, there was ample of opportunity to ask specialized and detailed questions about the code and workflow relevant to your own research.
- The course required full knowledge of physics and chemistry.
- The room often got quite hot, maybe a larger room would help.
- The exercise sessions could be organised in a more efficient way, maybe with some additional theory.

#### 4.1.6 Scientific Python

Trainer: Geert Jan Bex (UHasselt, VSC)

Date	#participants
16/01/2019	33
26/06/2019	27

Satisfaction scores <sup>2</sup>	
Training content	89%
Lecturer	87%

Selected suggestions for improvement and remarks:

- Very enthusiastic and highly skilled teacher.
- Very nice and extensive training material.
- Excellent preparation of the teacher.
- Course level was at times too basic and at times too specific.
- There weren't many exercises opportunities.

#### 4.1.7 Suggestions from participants for new training sessions

- Upcoming high performance (low-level) programming languages.
- Object-oriented programming
- High level automation
- Hands-on for data storage approaches
- Performance and optimization issues and techniques
- Data processing
- Statistical analysis
- Python for data science
- Python for natural language processing
- Cython
- Numba
- Interfacing\_C\_C++\_Fortran with python
- Parallel computing with python
- Working with spark on the HPC environment
- A C++ course
- Advanced Pandas
- Jupyter notebook courses for teaching purposes
- Visualizing and post-processing of results
- Geometry optimizations in computational chemistry
- More specialized python courses for physical scientists, behavioral scientists, etc.
- Machine Learning (e.g. scikit-learn, hyper parameter/grid search on the HPC environment)
- Advanced scikit-learn
- Machine learning with Python
- Deep learning (e.g. keras, pytorch, tensorflow,...) on GPUs

## 4.2 Lectures and community meetings

### 4.2.1 Tech talks on scientific computing

Speakers and topics:

- “Towards self-managed, re-configurable streaming dataflow systems”  
Vasia Kalavri (ETH Zürich, Switzerland)
- “The convergence of HPC and BigData - What does it mean for HPC sysadmins?”  
Damien François (UC Louvain, Belgium)
- “HPC on OpenStack”  
Petar Forai (Vienna Biocenter, Austria)
- “System Administration at Fred Hutchinson Cancer Research Center”  
John F. Dey (Fred Hutchinson Cancer Research Center, US)
- “ReFrame: A Regression Testing and Continuous Integration Framework for HPC systems”  
Vasileios Karakasis (CSCS, Switzerland)

Date	#participants
4/02/2019	20

#### 4.2.2 Tier1 information session

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

Date	#participants
10/05/2019	2
27/09/2019	6

#### 4.2.3 PRETREF project workshop

Workshop on the use of open source software (OpenFOAM) for running CFD simulations on HPC systems, organized by the PRETREF project community (<http://www.pretref.ugent.be/>). HPC-UGent helped in the organization of this workshop.

Featured lecture “Introduction to HPC-UGent” by Stijn De Weirdt.

Date	#participants
16/10/2019	24

On average, 91% of participants rated the lectures of this workshop as very or mostly useful/interesting.

#### 4.2.4 BUCCS summer school

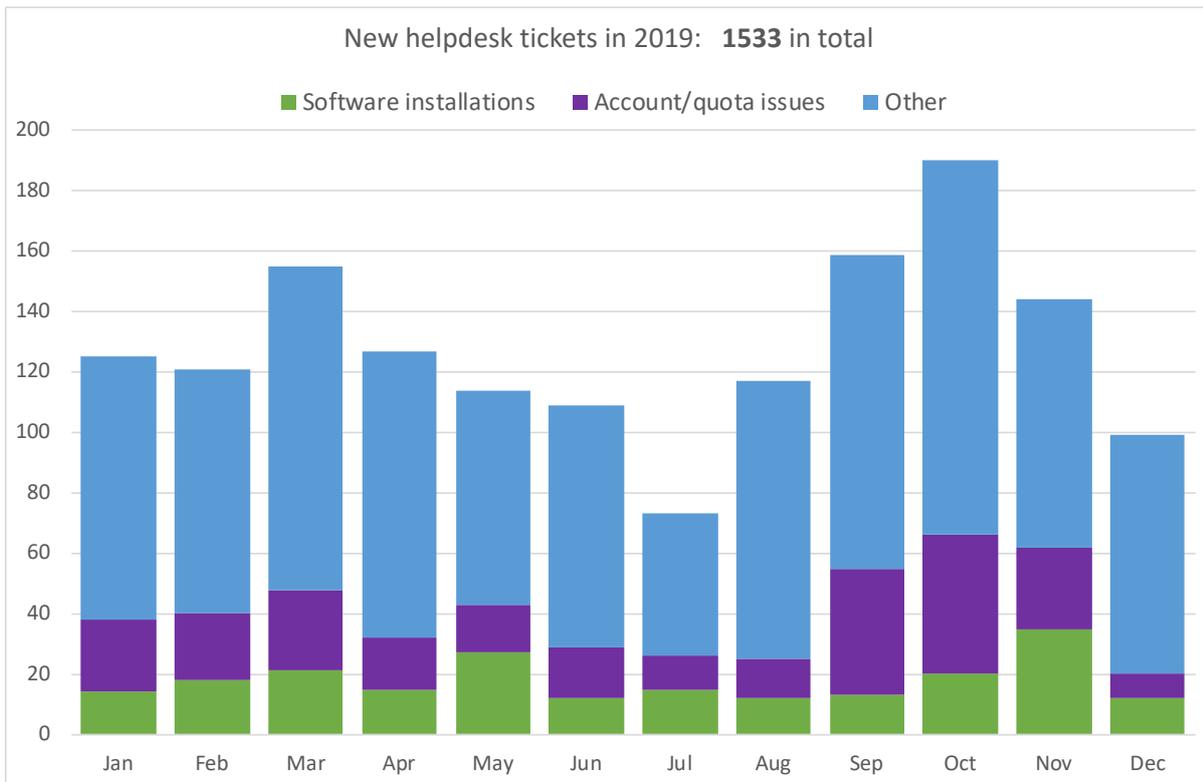
The Bucharest Urban Climate Summer School (2019 - third edition) focused on urban climate monitoring and modelling, urban ecology and biometeorology. It was co-organized by the Hydro-Climate Extremes Lab (H-CEL) of Ghent University (<https://www.ugent.be/bw/environment/en/research/h-cel>). Computational resources for the hands-on sessions throughout the workshop were provided by HPC-UGent.

Date	#participants
2-6/09/2019	18

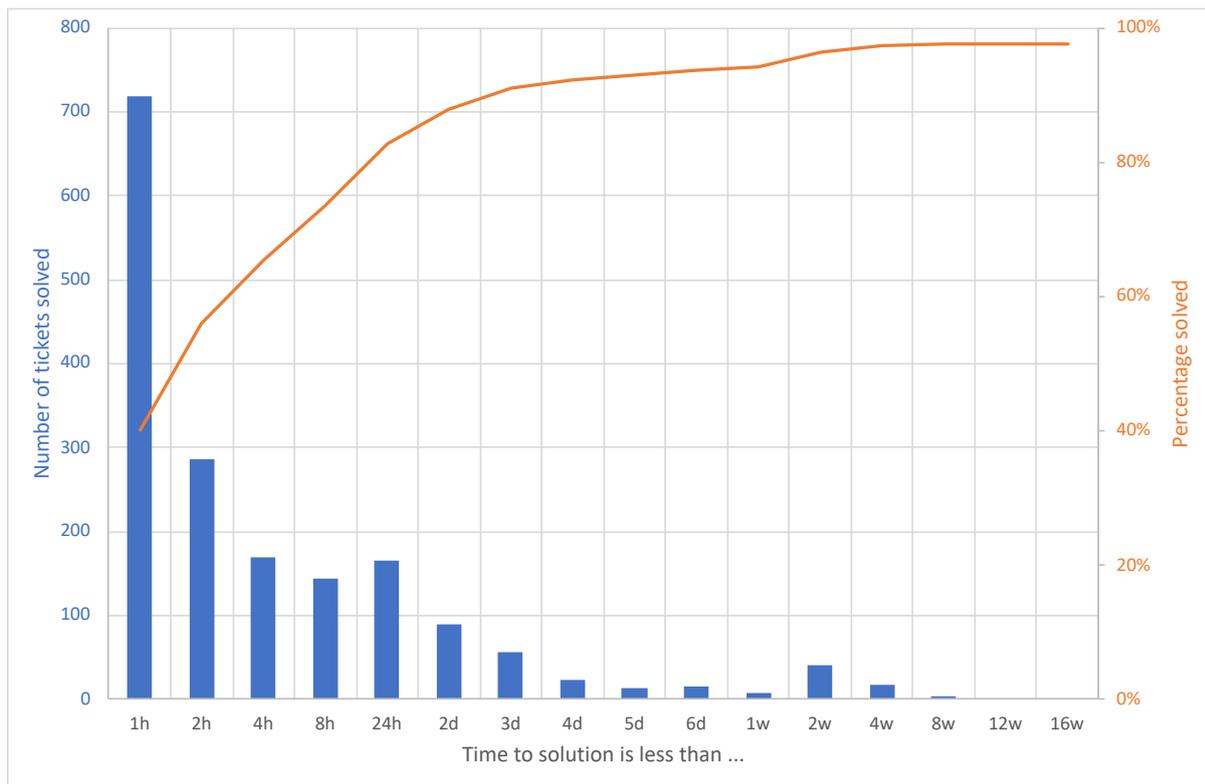


## 4.3 Helpdesk

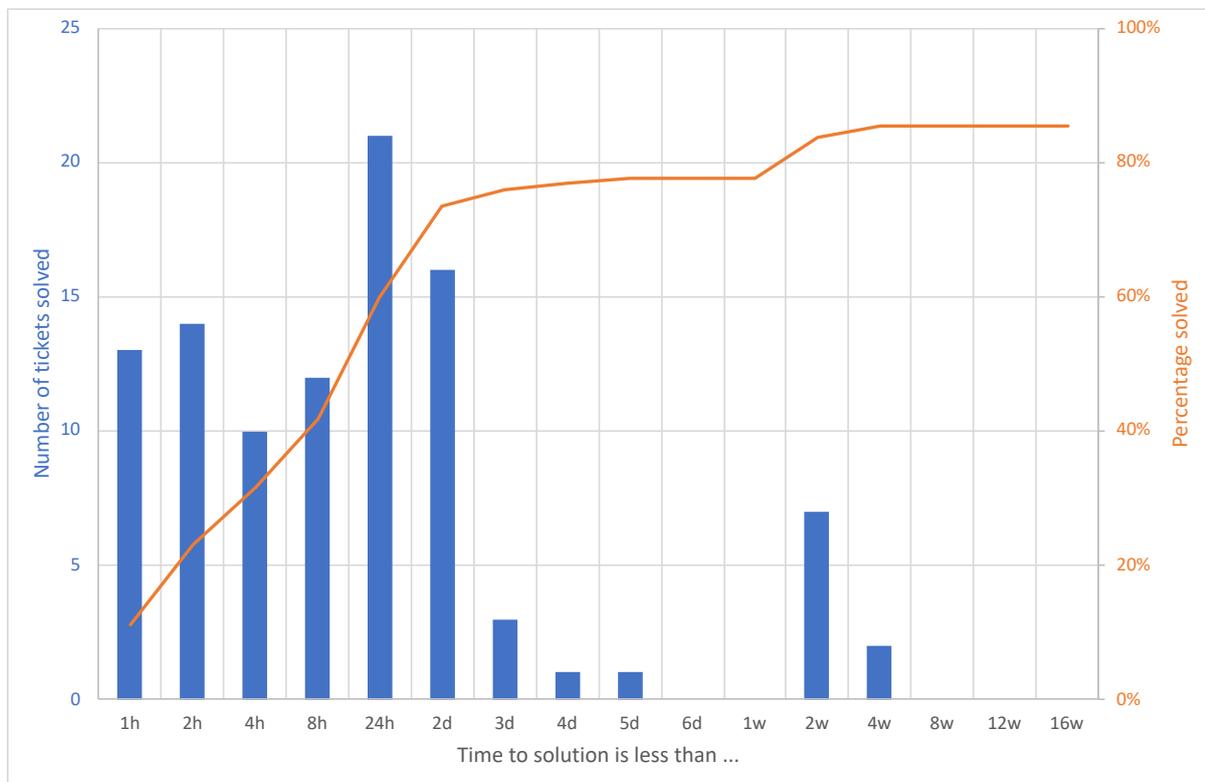
### 4.3.1 Number of new tickets per month



### 4.3.2 Average time to resolution – overall



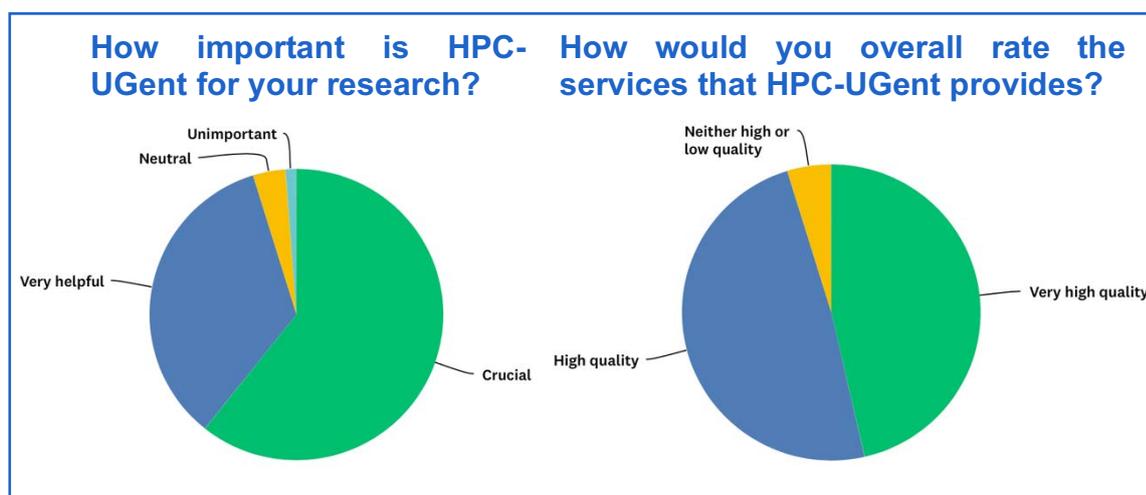
### 4.3.3 Average time to resolution – software installation requests



## 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 3 minutes) and in all 84 users responded. A selection of questions and responses is shown.

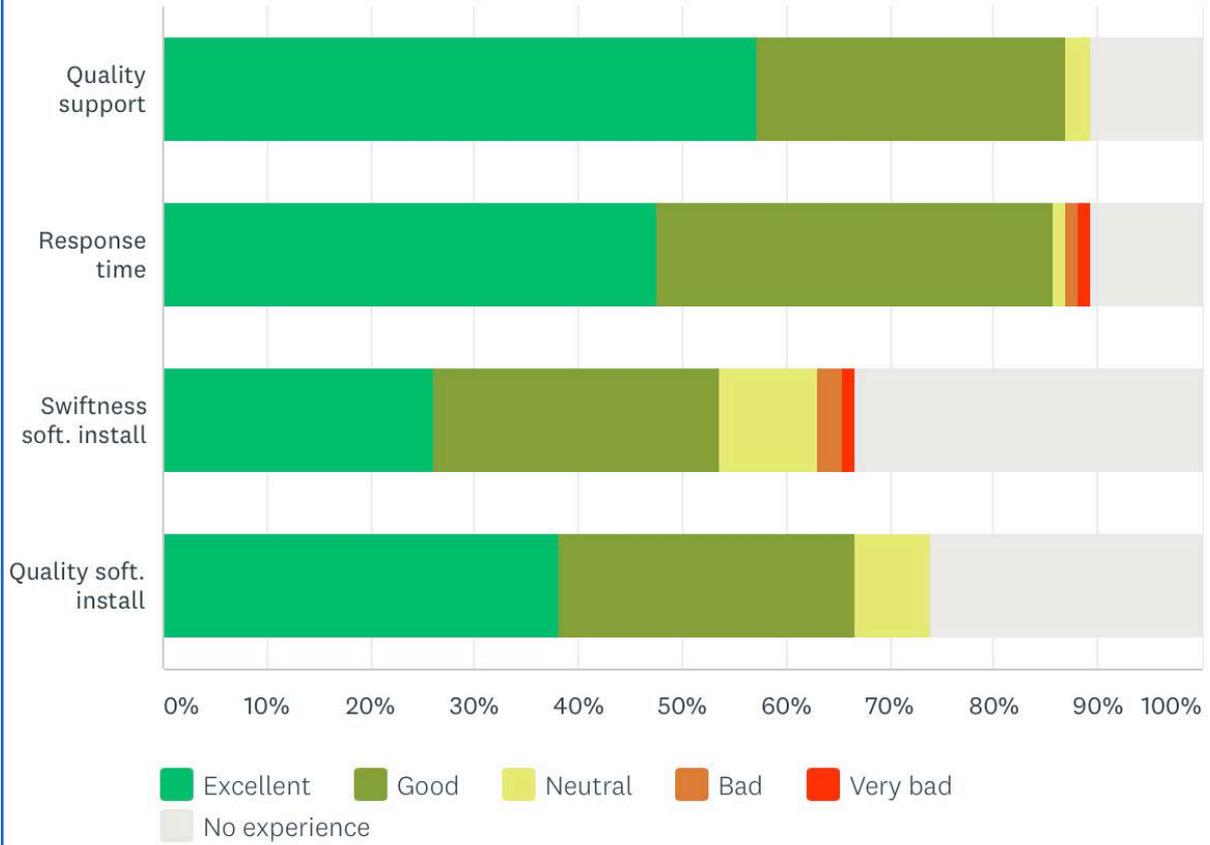
How often do you use the HPC-UGent scientific computing infrastructure?		
	#responses	%
On a daily basis	30	36%
(More than) once per week	30	36%
(More than) once per month	21	25%
1-2 times per year or less	3	4%



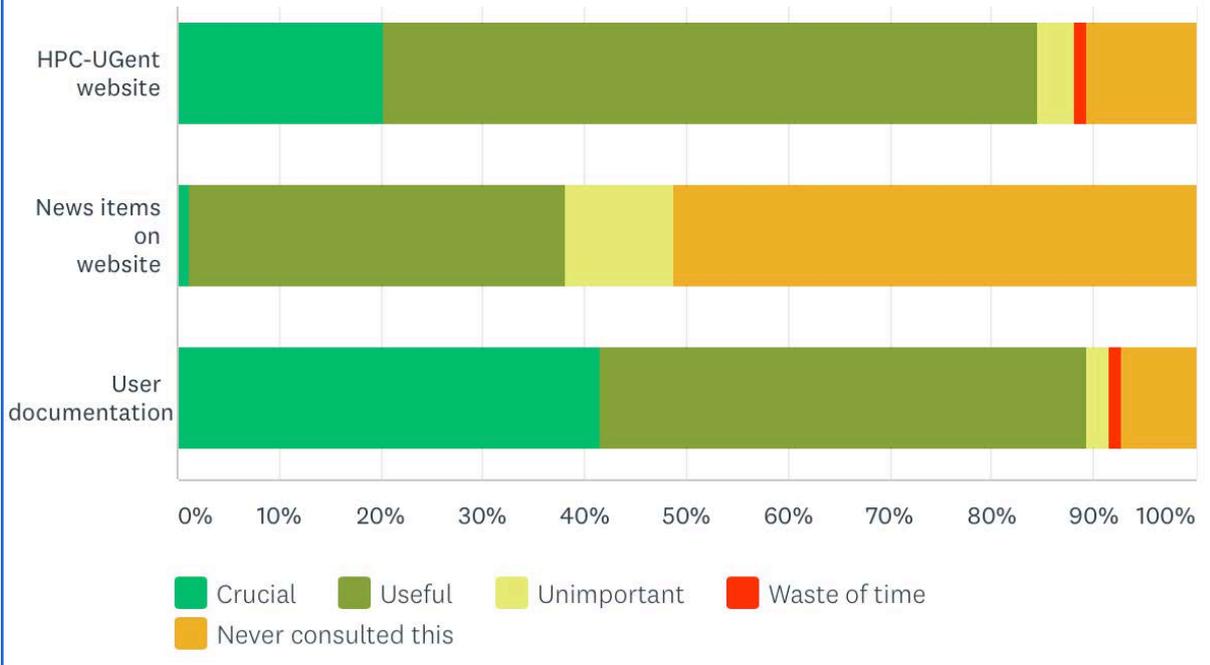
How important is HPC-UGent for your research?		
	#responses	%
Crucial - I can't do my research without	51	61%
Very helpful - It allows me to do my research at a faster pace and at a higher level	29	35%
Neutral	3	4%
Unimportant - I can just as well do my research in another way	1	1%
Unnecessary - I don't need it at all to do my research	0	0%

How would you overall rate the services that HPC-UGent provides?		
	#responses	%
Very high quality	39	46%
High quality	41	49%
Neither high or low quality	4	5%
Low quality	0	0%
Very low quality	0	0%

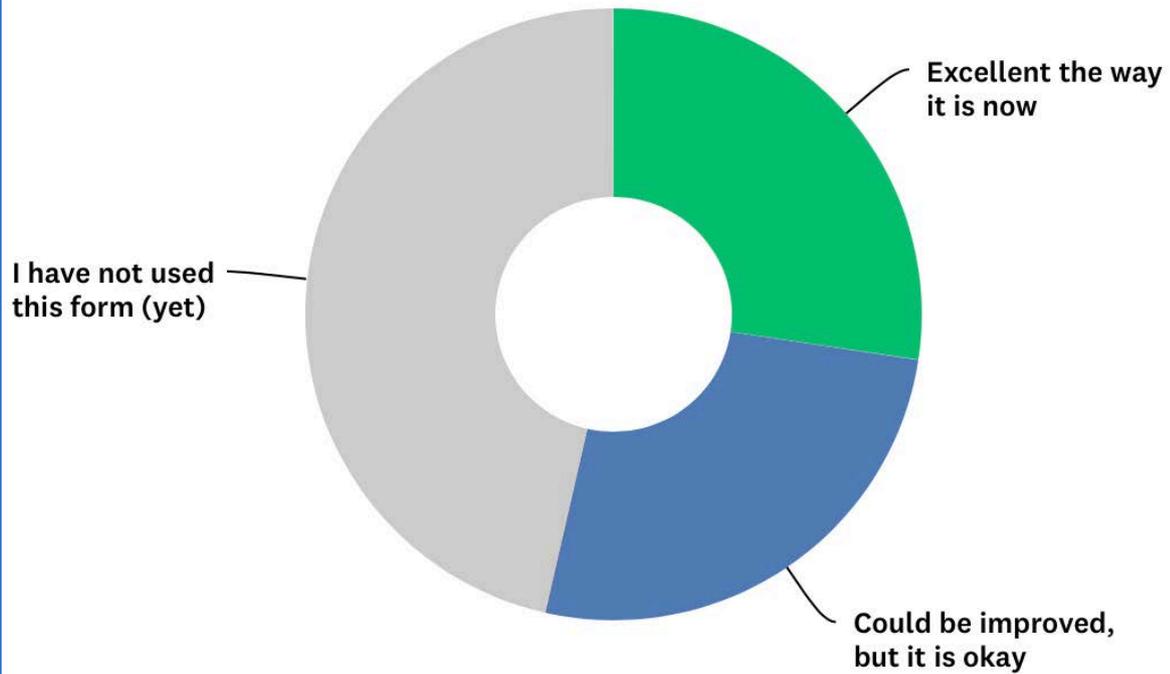
**Please rate the following aspects related to HPC-UGent user support**



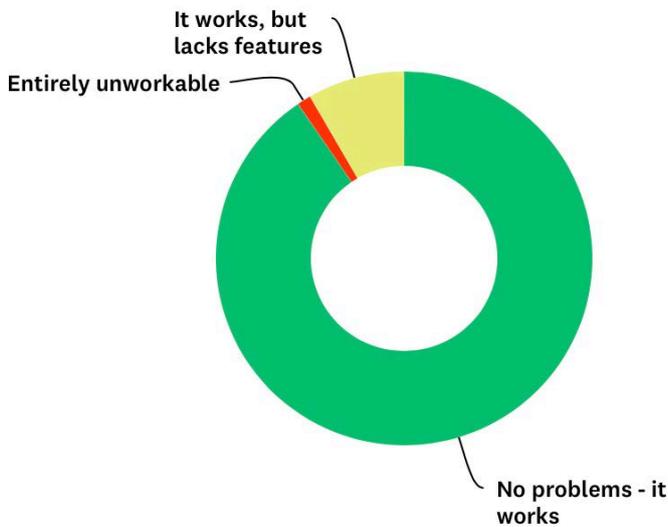
**Please rate the following HPC-UGent documentation sources**



### How do you like the form for software installation requests?



### Please rate your experience with the new SLURM job scheduler.



#### Requested features

Info on (number of) submitted jobs by other users

Info on own relative priority

Job arrays

Emails mentioning directory job was started in

Some previous commands, like 'qstat -q'

No need to select queue, 72h default

## Would your research benefit from specific IT hardware or services that HPC-UGent currently does not provide?

	#responses	%
No – The current compute platform suffices	47	56%
I don't know	27	32%
Yes – I need specific hardware/services	10	12%

- (More) GPUs
- (More) storage space
- More memory per cpu
- Support for compiling software
- CPLEX IBM library already installed
- Optimization of Abaqus solver settings for HPC
- Smoother usage of software such as QGIS

## How could we further improve HPC-UGent services?

### Documentation

- Library of scripts
- Expand FAQ and guidelines for starters
- How to set file permissions in a VO group
- More elaborate OpenFOAM tips in documentation (e.g. \$PBS\_O\_WORKDIR vs. \$WORKDIR, example of compiling new library/solver)

### User experience

- Shorter queue times
- Longer wall-clock time
- Better GUI, or more visual approach/interface for (new) users
- Way to estimate queue time
  - See number of jobs/corehours requested per cluster by all users
  - See how many cores are still available per node (in pbsmon)
  - 'queue -start' proved to be 80% reliable for jobs requiring 72h & +8 nodes
- Less downtime

### Infrastructure

- More data storage on SCRATCH/DATA
- More nodes
- More GPUs

### Training

- Regular based training courses
- Working with industry for certain domains
- Training about GPUs

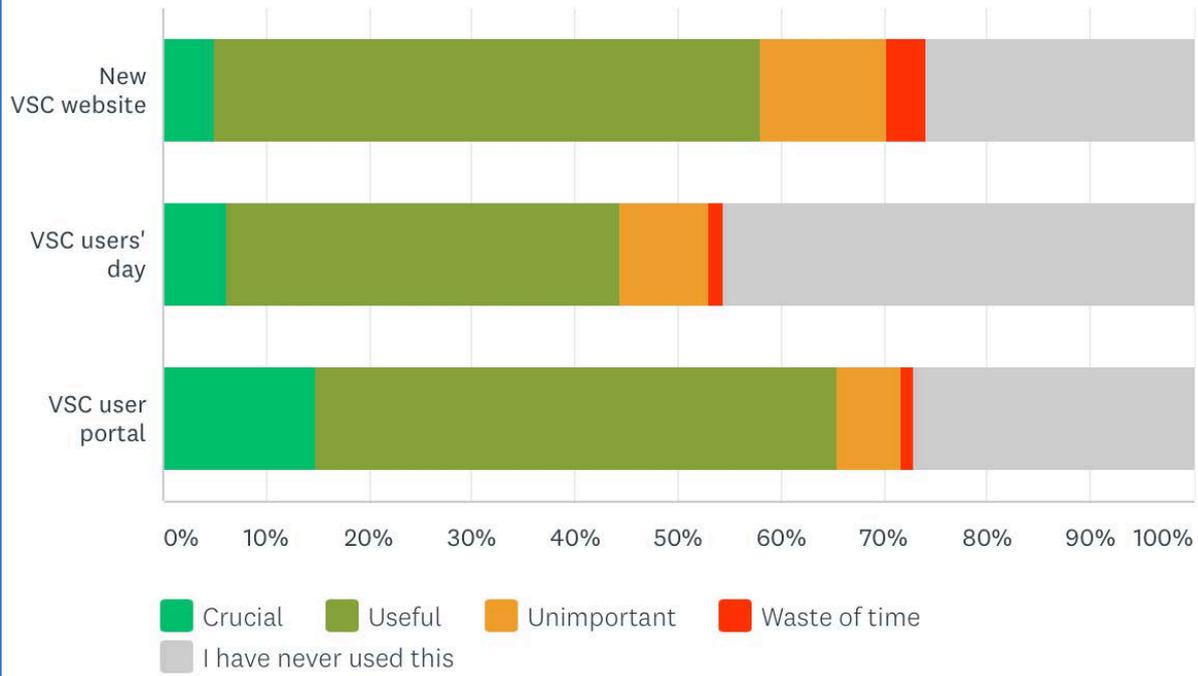
### Software

- Docker

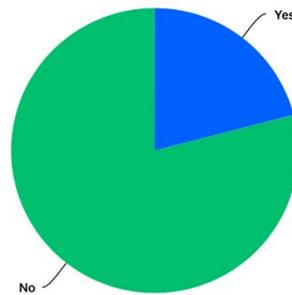
### User support

- Speed up software installation
- Rate employees or issue resolution

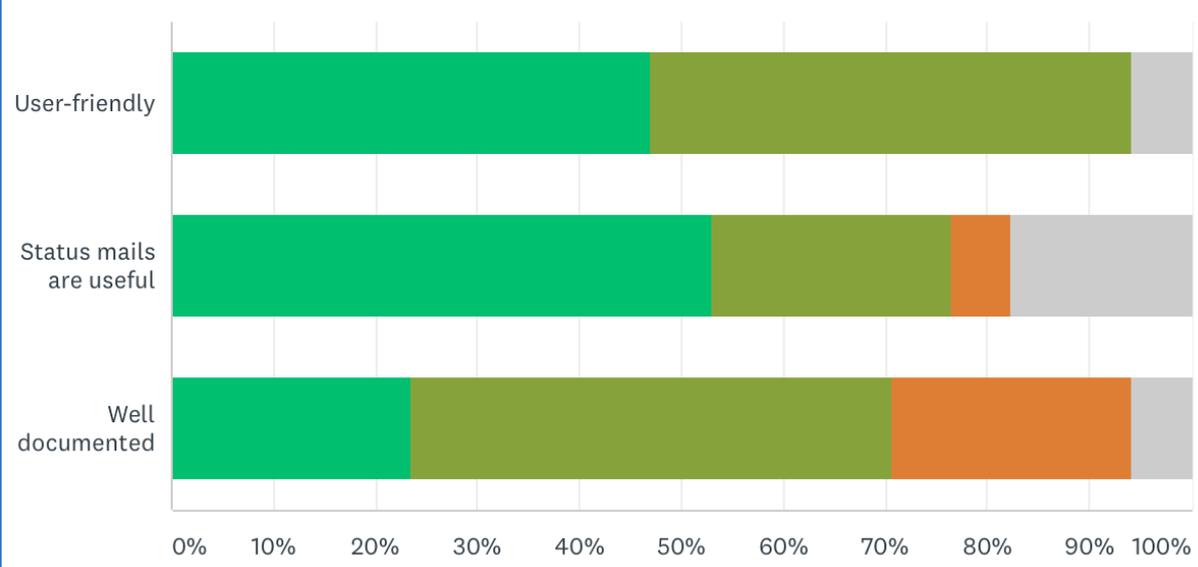
**Please rate the following VSC initiatives**



**Have you ever used the Tier-1 compute infrastructure?**



**Rate these aspects of the Tier-1 creditsystem (used for project allocations)**



### **User encouragements**

*"We are super super happy - thank you all @HPC!"*

*"You are doing an excellent job, please keep up the good work!"*

*"Thanks for the great job you are doing."*

*"Generally excellent service!"*

*"You're doing a great job as it is! Thanks for all the good work of the past year!"*

## 5 OUTREACH

### 5.1 Within Ghent University

Date	Event
16/01/2019	HPC introduction and UGent datacenter tour for students in field of accounting/informatics on Big Data (Bernardustechnicum, Oudenaarde)
5/02/2019	Information session FWO call 2019 for Research Grants, Research Projects and SBO
30/04/2019	HPC introduction and UGent datacenter tour for students Bachelor of Science in Engineering Technology (UGent)
14/05/2019	HPC introduction and UGent datacenter tour for students in field of 'Technology and networks' (Scholengroep Leiepoort, Campus Sint-Vincentius, Deinze)
9/09/2019	Research day Digital Humanities <a href="https://www.lwresearchday.ugent.be">https://www.lwresearchday.ugent.be</a>

### 5.2 Within international HPC community

Date	Event
30/01-1/02/2019	4 <sup>th</sup> EasyBuild User Meeting UCL, Louvain-la-Neuve, Belgium <a href="https://github.com/easybuilders/easybuild/wiki/4th-EasyBuild-User-Meeting">https://github.com/easybuilders/easybuild/wiki/4th-EasyBuild-User-Meeting</a>
2-3/02/2019	FOSDEM'19 Co-organisation of <i>HPC, Big Data and Data Science</i> devroom Brussels, Belgium <a href="https://fosdem.org/2019/schedule/track/hpc_big_data_and_data_science/">https://fosdem.org/2019/schedule/track/hpc_big_data_and_data_science/</a>
25/04/2019	CÉCI Scientific Meeting Université Libre de Bruxelles, Belgium Keynote presentation " <i>Getting scientific software installed using EasyBuild</i> " <a href="http://www.ceci-hpc.be/scientificmeeting.html">http://www.ceci-hpc.be/scientificmeeting.html</a>
12-14/06/2019	HPC Knowledge Meeting 2019 & Singularity training 2019 Barcelona, Spain <ul style="list-style-type: none"><li>• EasyBuild tutorial at Singularity containers workshop</li><li>• Keynote presentation "10 Years of EasyBuild and The Road Ahead"</li></ul> <a href="https://2019.hpckp.org">https://2019.hpckp.org</a>

Highlighted event:

### Research day Digital Humanities

(9 September 2019, De Krook, Ghent)

The UGent faculty of Arts and Philosophy organized a research day focused on 'digital humanities'. With this event, the faculty clearly indicated its desire to focus on the digital literacy of its staff in the coming years.

HPC-UGent presented a poster and collected feedback from participants.

DEPARTMENT ICT  
HPC-UGent

## SCALE UP YOUR RESEARCH – SUPERCOMPUTING @ UGENT

**SUPERCOMPUTERS**  
Help solve complex problems in everyday life, research and industry

Does this sound familiar?  
 "I regularly leave my laptop running to do a computational analysis"  
 "I analyze only 10 out of 10.000 samples because it takes too long on my computer"  
 "I don't have enough storage space to do my data analysis"

CLIPPING: 2019

**Supercomputing infrastructure 'STEVIN'**

- 6 supercomputers
- 15.760 cores
- 68 TeraBytes of memory
- 2 PetaBytes of shared storage space

24 hours of work on entire STEVIN corresponds to 22 years on an average laptop

**HPC-UGent team**

- Infrastructure installation and upkeep (hardware & software)
- User support
- Training
- Outreach + marketing
- Collaboration with other supercomputing centers
- Help with Tier1 project proposal

Work in terminal: steep learning curve, but very powerful

We are currently working on a cloud interface to suit specific user demands

Contact:  
 hpc@ugent.be  
 www.ugent.be/hpc

@HPCUGent  
 High Performance Computing UGent

GHENT UNIVERSITY

Link: <https://www.lwresearchday.ugent.be>

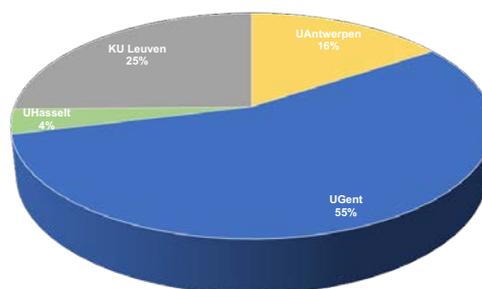
## 6 BUDGET

Budget line	UGent budget statement	Income statement
Central compute infrastructure	1.200.000	
Central compute infrastructure (wages)	880.000	
Impulse funding Tier1 staff	63.000	
FWO recurrent subsidy for personnel (Tier1+2)		665.000
FWO recurrent subsidy for investment and operations		1.247.557
FWO project subsidy Tier1 operational management team 2019 – staff		63.333,33
FWO project subsidy Tier1 cloud 2019 – staff		95.000
FWO project subsidy Tier1 cloud 2019 – investment and operations		435.000
	euro	euro

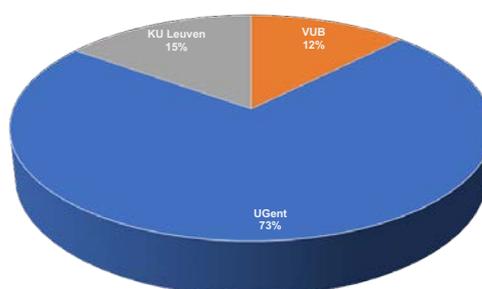
## 7 TIER-1 USAGE

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.

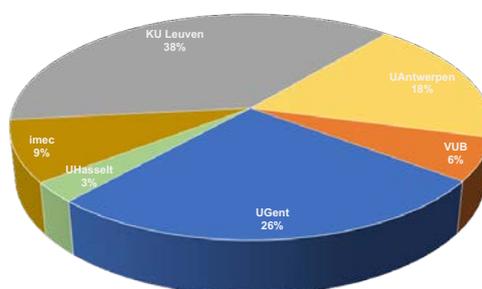
Call 1		Requested		Granted	
4/02/2019					
Institute	#proj	#nd	#proj	#nd	
UAntwerpen	3	13441	3	13441	
VUB	1	3200	0	0	
<b>UGent</b>	<b>14</b>	<b>47592</b>	<b>14</b>	<b>47592</b>	
UHasselt	1	3050	1	3050	
KU Leuven	4	21719	4	21719	



Call 2		Requested		Granted	
3/06/2019					
Institute	#proj	#nd	#proj	#nd	
UAntwerpen	0	0	0	0	
VUB	3	5906	3	5906	
<b>UGent</b>	<b>10</b>	<b>34469</b>	<b>10</b>	<b>34469</b>	
UHasselt	0	0	0	0	
KU Leuven	3	7162	3	7162	



Call 3		Requested		Granted	
7/10/2019					
Institute	#proj	#nd	#proj	#nd	
UAntwerpen	6	32120	4	19494	
VUB	2	6572	2	6572	
<b>UGent</b>	<b>8</b>	<b>28577</b>	<b>8</b>	<b>28577</b>	
UHasselt	1	3150	1	3150	
KU Leuven	10	42782	9	40807	
imec	2	9500	2	9500	



*Percentage of awarded compute time, listed per institute per call*

## 8 USER IN THE SPOTLIGHT

Dirk Seidensticker uses supercomputing to study the ceramic vessels and their use in the past based on archaeological contexts in Central Africa. He is an archaeologist and junior postdoctoral researcher within the BantuFirst research project (ERC) at the Languages and Cultures department of Ghent University (faculty of Arts and Philosophy).



*Dirk Seidensticker during excavations in the Equator province of the Democratic Republic of the Congo (Photo: K. Jungnickel 2015)*

The HPC-UGent systems are a crucial part of this research as it involves the conversion of a vast collection of scanned images, which were extracted from the literature.

“A key part of archaeological research revolves around establishing chronological sequences of material remains found during surveys and excavations. In my study area, the equatorial rainforests of Central Africa, the most abundant category of finds are ceramic vessels or sherds of those. Besides their decoration, shape, and proportions, these vessels are seen as central criteria for groupings and so-called style-traditions. While my research shifts focus on the technology by which potters produced these vessels in the past, another aspect that was equally ignored thus far is the intended use and function of different types of vessels based on their shapes and shifts of these types throughout the past millennia. This involves a standardization of all individual classification systems, called typologies, that have been used in order to describe the shape of ceramic vessels in different regions.”

“Within the past two decades, the use and function of ceramic vessels have been discussed for various regions and time-frames throughout the world. Most of these studies used measurements like the width of the mouth or the height of a vessel and followed traditional morphometric approaches in order to detect groups within their datasets. While these approaches have the advantage that some of these measurements can also be taken on heavily fragmented inventories, the general shapes are mostly reduced to only about half a dozen data points per vessel. In my research, I rely on the more robust approach of outline analysis in order to quantify the shape of the vessels.”

“While a geometric morphometric analysis on outlines can be calculated easily using various R-packages on any modern office computer or laptop, getting usable input data mostly involves long and repetitive preparations of images in a raster graphics editor. In contrast, I use the R-package shapAAR that was developed by Martin Hinz (Bern) as part of the Initiative for Statistical Analysis in Archaeology Kiel (ISAACK). This package includes functions to process scanned drawings of objects, like ceramic vessels, and extract the outline using an active contour algorithm. Especially the

detection of the outline takes quite substantial CPU time and needs to be performed for each individual drawing. This kind of task, I figured, was perfect for the HPC infrastructure of UGent. As of now, the shapAAR package is still in development, and in collaboration with the main developer, I try to even more simplify the workflow. Running the script on the HPC servers further enables me to run multiple permutations per image in order to find optimal parameters for the active contour algorithm. These parameters are linked to the individual style of drawing and the quality of the reproduction.”

“This research is extremely novel to my study area, as it makes use of existing documentation rather than regarding additional fieldwork as the only means of generating new knowledge. While I am currently building a database of all the drawings I scanned and collected during the past month, as well as the necessary metadata, a reanalysis of the ceramics found at the necropolis of Campo, in southern Cameroon, showed substantial improvements compared to the analogue and thus subjective classification system. In the next step, we will analyse the entire dataset, containing more than 500 drawings from about 75 archaeological sites within the Congo Basin. At a later stage, all available drawings of vessels from Central Africa shall be included. Based on such a vast dataset, which spans multiple regions and time-slices, patterns of preferences in the usage of particular shapes in certain contexts can be deduced.”

“As part of a proposed FWO junior postdoctoral fellowship, I intend to study the transfer of knowledge of past potters’ communities and how differences in intended use affected technological approaches towards making pottery. The unique possibilities of relying on the HPC infrastructure of UGent will be most valuable in that effort.”

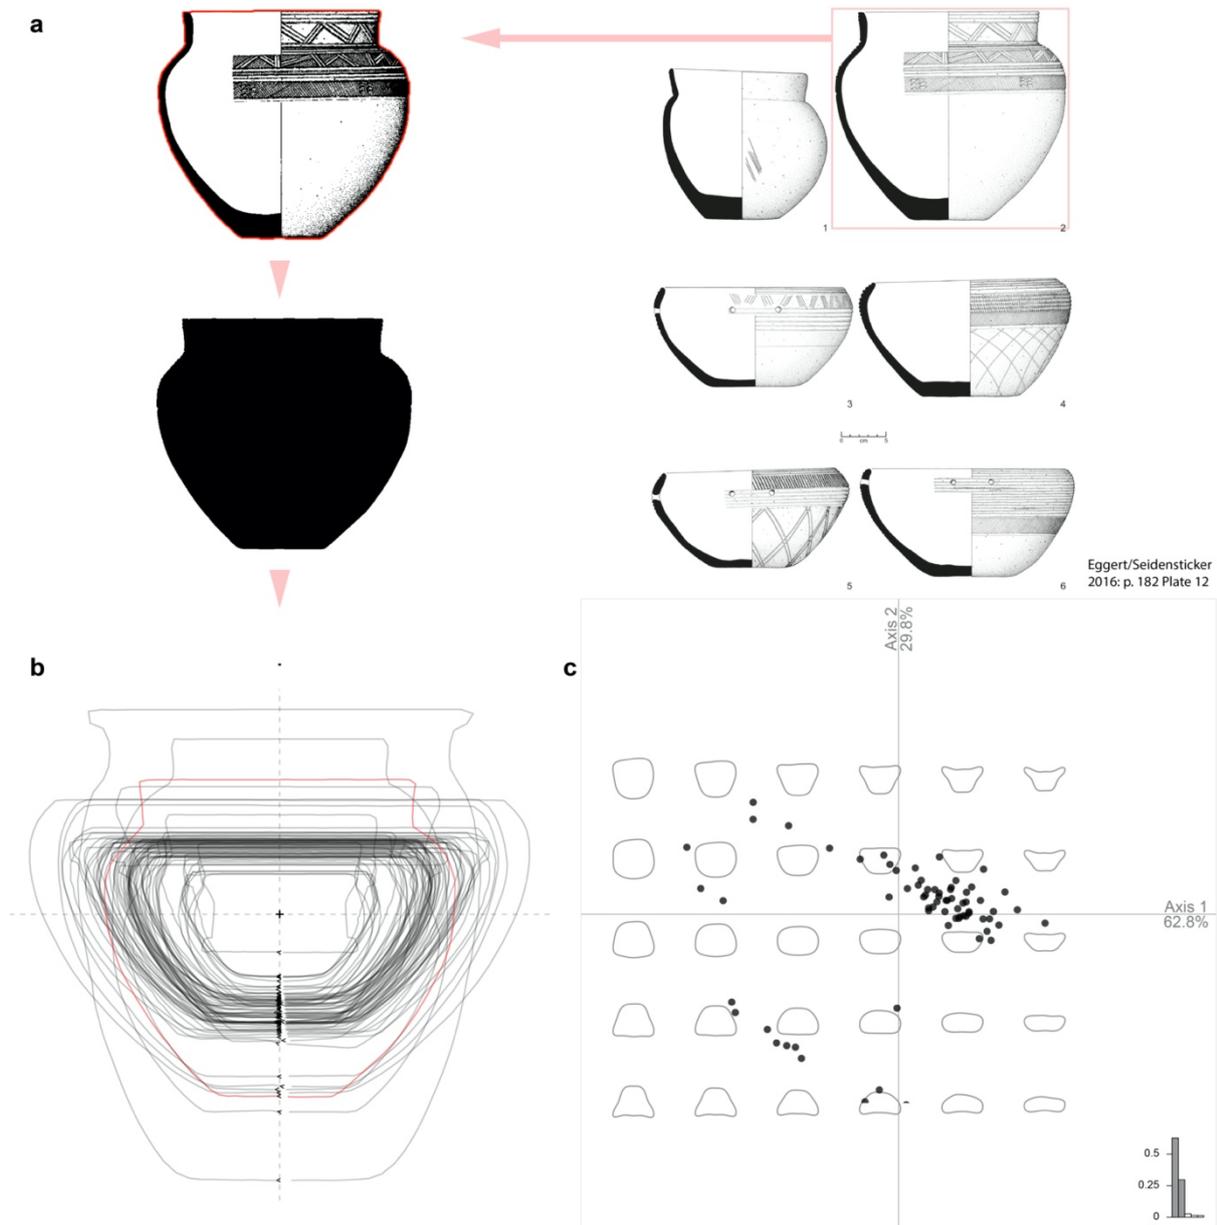
## Links

- BantUGent – UGent Centre for Bantu Studies (<https://research.flw.ugent.be/en/bantugent>)
- BantuFirst Project (<https://www.bantufirst.ugent.be>)
- The shapAAR R-package was used to detect the outlines of individual vessels from scanned drawings (<https://github.com/ISAAKiel/shapAAR>)
- The Momocs R-package was used for the geometric morphometric analysis (<http://momx.github.io/Momocs/>)

## Publications

The initial analysis of the features and finds excavated at the necropolis at Campo have been published in a monograph Dirk co-authored:

*Manfred K. H. Eggert and Dirk Seidensticker*  
*Campo: Archaeological Research at the Mouth of the Ntem River (South Cameroon)*  
*Africa Praehistorica, Vol. 31, 2016*  
*Cologne: Heinrich-Barth-Institut*  
<http://www.uni-koeln.de/hbi/bookshop/buecher/campo2016.htm>



Eggert/Seidensticker  
2016: p. 182 Plate 12

The outline of 75 ceramic vessels from the early iron age Necropolis is detected automatically on the HPC servers through the *shapAAR* R-package using an active contour algorithm (a). A morphometric analysis enabled a new classification system based on elliptic Fourier transformations and Principal Component Analyses using the *Momocs* R-package (b-c).