



**GHENT
UNIVERSITY**



Introduction to HPC-UGent

May 27th 2019

<https://www.ugent.be/hpc/en/training/materials/2019/introhpcugent>

hpc@ugent.be

<https://ugent.be/hpc>



Vlaanderen
is computing

What is High Performance Computing?

“High Performance Computing” (HPC) is computing on a “supercomputer”, a system at the frontline of contemporary processing capacity – particularly in terms of size, supported degree of parallelism, network interconnect and (total) available memory & disk space.

A computer cluster consists of a set of loosely or tightly connected computers that work together so that in many respects they can be viewed as a single system.



(a.k.a. “supercomputing”)

What is High Performance Computing?


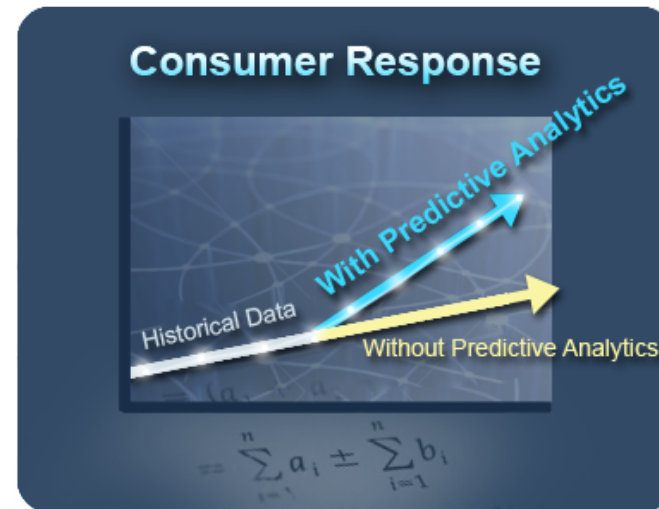




harness power of multiple interconnected cores/nodes/processing units



Everyday applications of supercomputing



Scientific applications of supercomputing

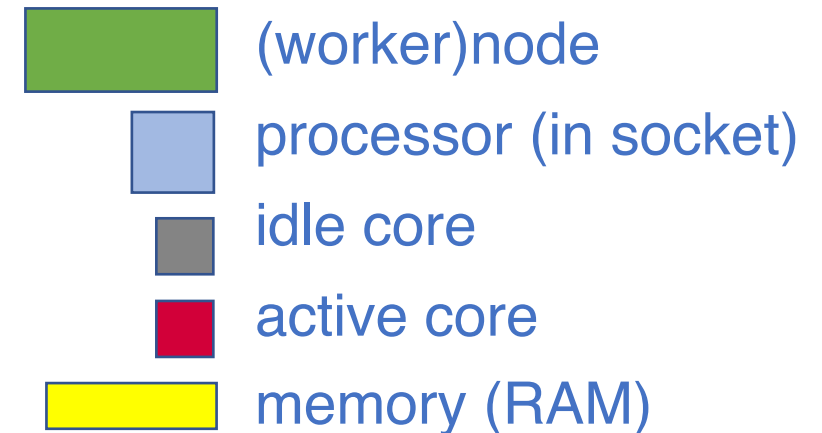



Cores, CPUs, processors, sockets, (worker)nodes

Modern servers, also referred to as (worker)nodes in the context of HPC, include one or more sockets, each housing a multi-core processor (next to memory, disk(s), network cards, ...).

A modern (micro)processor consists of multiple CPUs or cores that are used to execute computations.

example: workernode
with two 16-core
processors running
a single core job



(not included in picture:
local disk, network cards, ...)

Parallel vs sequential software

In parallel software, many calculations are carried out simultaneously.

This is based on the principle that large problems can often be divided into smaller tasks, which are then solved concurrently (“in parallel”).


Example: OpenFOAM can easily use 160 cores at the same time to solve a CFD problem

Parallel programming paradigms:


OpenMP for shared memory systems (multithreading) -> on cores of a single node

MPI for distributed memory systems (multiprocessing) -> on cores of multiple nodes

OpenMP software can
use multiple or all cores
in a single node




MPI software
can use (all) cores
in multiple nodes



Parallel vs sequential programs


Sequential (a.k.a. serial) software does not do calculations in parallel, i.e. it only uses one single core of a single workernode.



This type of software does not run faster by just throwing cores at it...

But, you can run multiple instances at the same time!

e.g., you can run a Python script 100 times on 100 cores to quickly analyse 100 datasets



HPC-UGent

hpc@ugent.be

Part of ICT Department of Ghent University

Our mission

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

Our core values

Empowerment - Centralisation - Automation - Collaboration

HPC-UGent: staff



Stijn De Weirdt
technical lead



Ewald Pauwels
team lead



Kenneth Hoste
user support & training



Wouter Depypere
sysadmin, hardware



Andy Georges
sysadmin, tools



Kenneth Waegeman
sysadmin, storage



Balázs Hajgató
sysadmin, tools

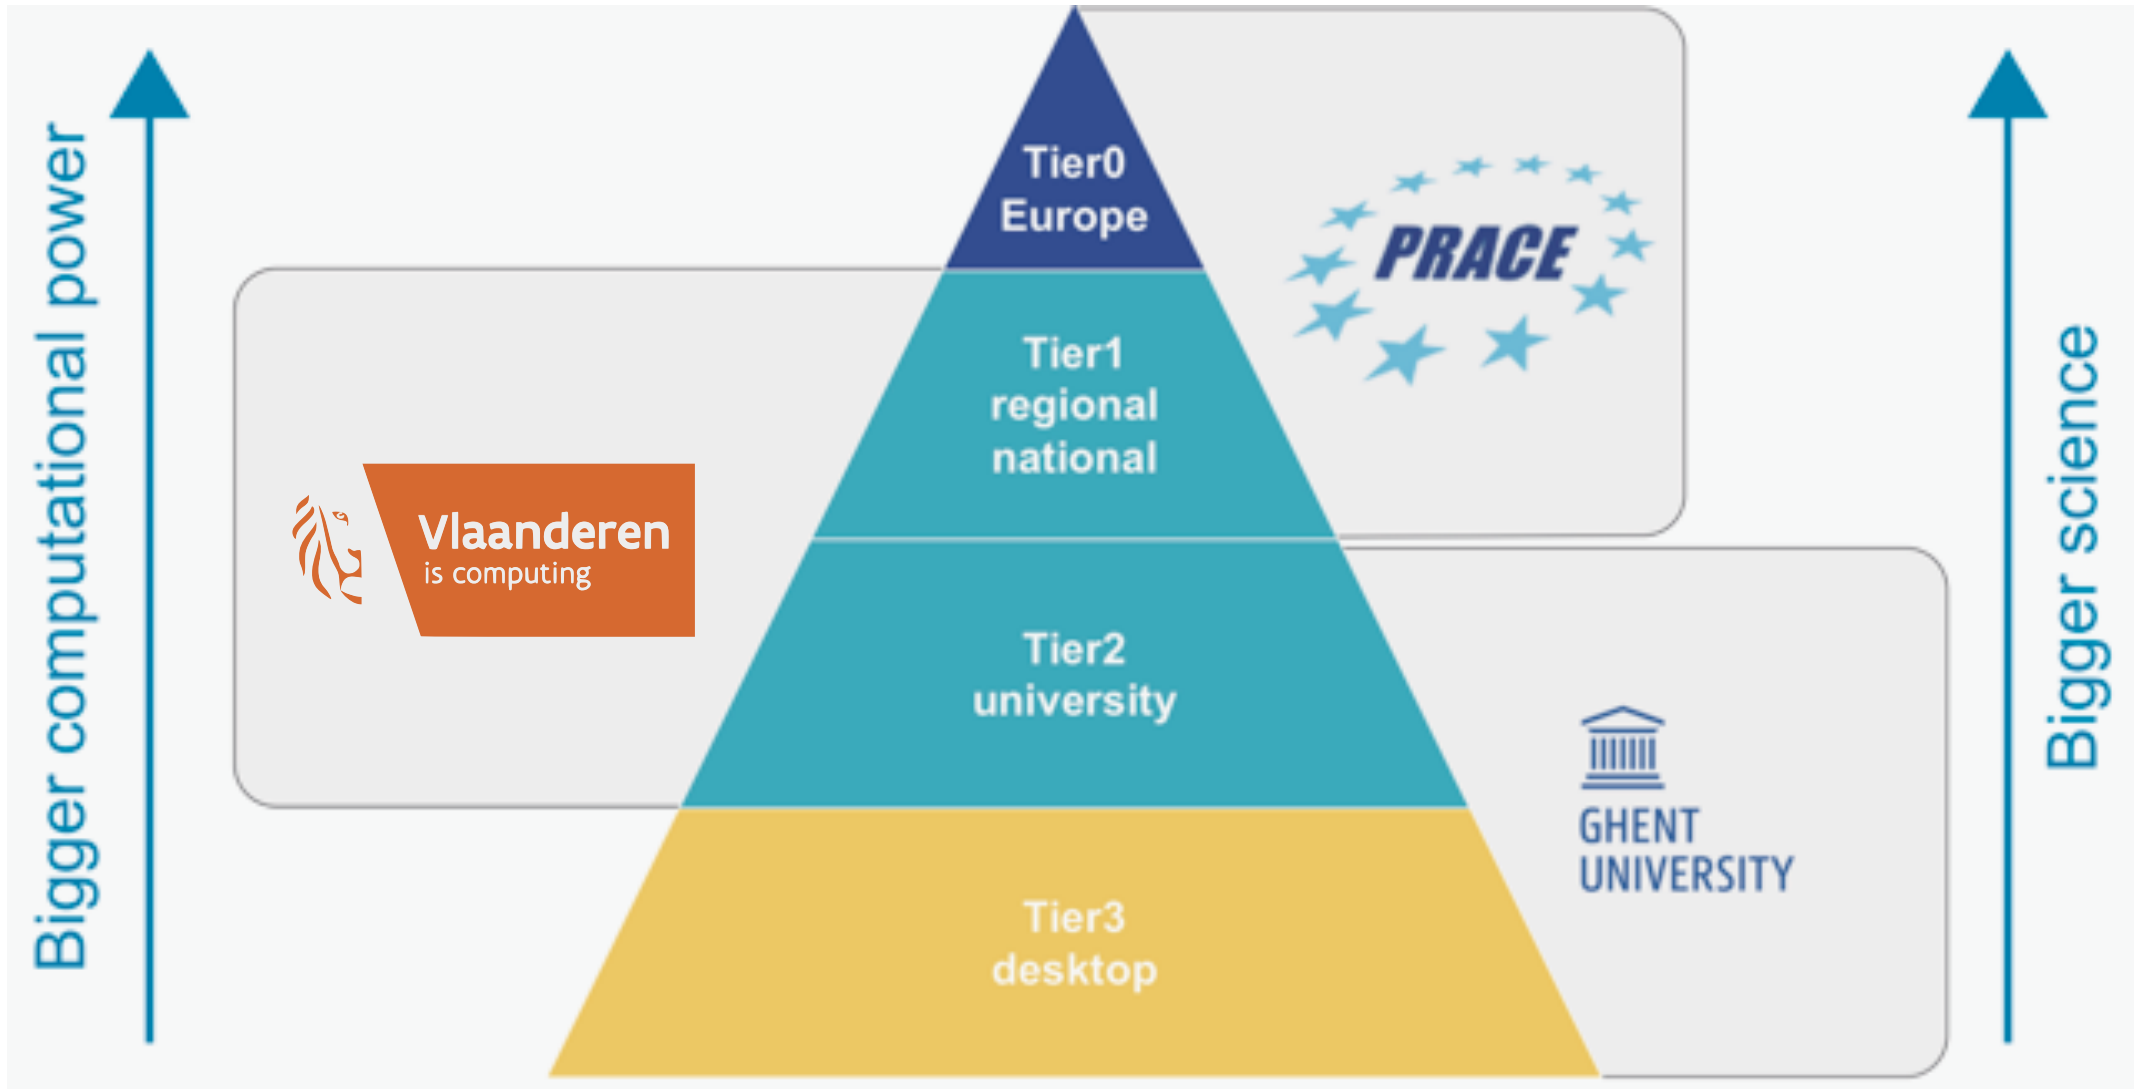


Álvaro Simón García
cloud, user support

Centralised hardware in the UGent datacenter at campus Sterre (building S10)



Centralised hardware




HPC-UGent Tier-2 (STEVIN): central investments



1548 - 1620
°Bruges

STEVIN
HPC
infrastructure

Total investment in HPC-UGent compute infrastructure



Financing by:



Vlaanderen
is computing



HPC-UGent
users

HPC-UGent Tier-2 (STEVIN)

<https://www.vscentrum.be/infrastructure/hardware/hardware-ugent>



Compute clusters

6 Tier-2 clusters

> 600 workernodes, > 15,000 cores

	#nodes	CPU	Mem/node	Diskspace/node	Network	
	delcatty	123	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	3x 400 GB (SSD, striped)	FDR InfiniBand
	golett	200	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
	skitty	72	2 x 18-core Intel Xeon Gold 6140 (Skylake @ 2.3 GHz)	192 GB	1 TB 240 GB SSD	EDR InfiniBand
	victini*	96	2 x 18-core Intel Xeon Gold 6140 (Skylake @ 2.3 GHz)	96 GB	1 TB 240 GB SSD	10 GbE

VSC Tier-2 infrastructure

Vlaams Supercomputer Centrum
(Flemish Supercomputer Center)

<https://www.vscentrum.be/en/access-and-infrastructure/tier-2>

Antwerp University association

Brussels University association
+ Grid specialization

Ghent University association
+ Big Data specialization

KU Leuven association
Limburg association University-Colleges
+ Shared memory, accelerator specialization (GPU)



Vlaanderen
is computing



VSC Tier-1 – BrENIAC (@ KUL)

For up to date information, see:

<https://www.vscentrum.be/en/access-and-infrastructure/tier-1>

Hardware

- 580 computing nodes (16,240 cores in total)
 - Two 14-core Intel Xeon processors (Broadwell, E5-2680v4)
 - 128 GiB RAM (435 nodes) or 256 GiB (145 nodes)
- EDR InfiniBand interconnect
 - High bandwidth (11.75 GB/s per direction, per link)
 - Slightly improved latency over FDR
- Storage system
 - Capacity of 634 TB
 - Peak bandwidth of 20 GB/s

Extension is being installed currently, which should bring total compute power to ~1.5 petaflop.

- 408 additional workernodes,
each with 2x Intel Skylake 14-core processors
- double the scratch storage volume



VSC Tier-1 – BrENIAC (@ KUL)

For academics (all Flemish research centers):

- Free of charge
- Starting Grant (100 node days)
 - <https://www.vscentrum.be/en/access-and-infrastructure/tier1-starting-grant>
 - Fill in application form, send it to hpc@ugent.be
- Project access (500-5000 nodedays)
 - 3 evaluation moments per year
 - Application form and more info
<https://www.vscentrum.be/en/access-and-infrastructure/project-access-tier1>
- Don't hesitate to contact hpc@ugent.be for help!



VSC Tier-1 – BrENIAC (@ KUL)

For industry:

- Exploratory access (100 node days)
 - Free of charge
 - Contact hpc@ugent.be
- Contract access
 - FWO/UGent/company contract
 - Payed usage (~13 euro / node / day)
 - Contact hpc@ugent.be



Getting a VSC account




- See Chapter 2 in HPC-UGent tutorial
- <https://www.vscentrum.be/en/access-and-infrastructure/requesting-access>
- All users of AUGent can request a VSC account
 - Researchers & staff
 - Master/Bachelor students (after motivation of ZAP)
- VSC account be used to use HPC infrastructure on all VSC sites
- Subscribed to hpc-announce and hpc-users mailing lists
- Beware of using HPC for teaching/exam purposes!
 - No guarantee on HPC availability (power outage/maintenance)
 - Have a backup plan at hand
 - Advisable teaching/exam formula: project work

Managing your VSC account



- You can manage your VSC account via the VSC account page:

<https://account.vscentrum.be>



View Account	Edit Account	View Groups	New/Join Group	Edit Group	New/Join VO	View VO	Edit VO	Reservations	Log Out
---------------------	--------------	-------------	----------------	------------	-------------	---------	---------	--------------	---------

View account

General information

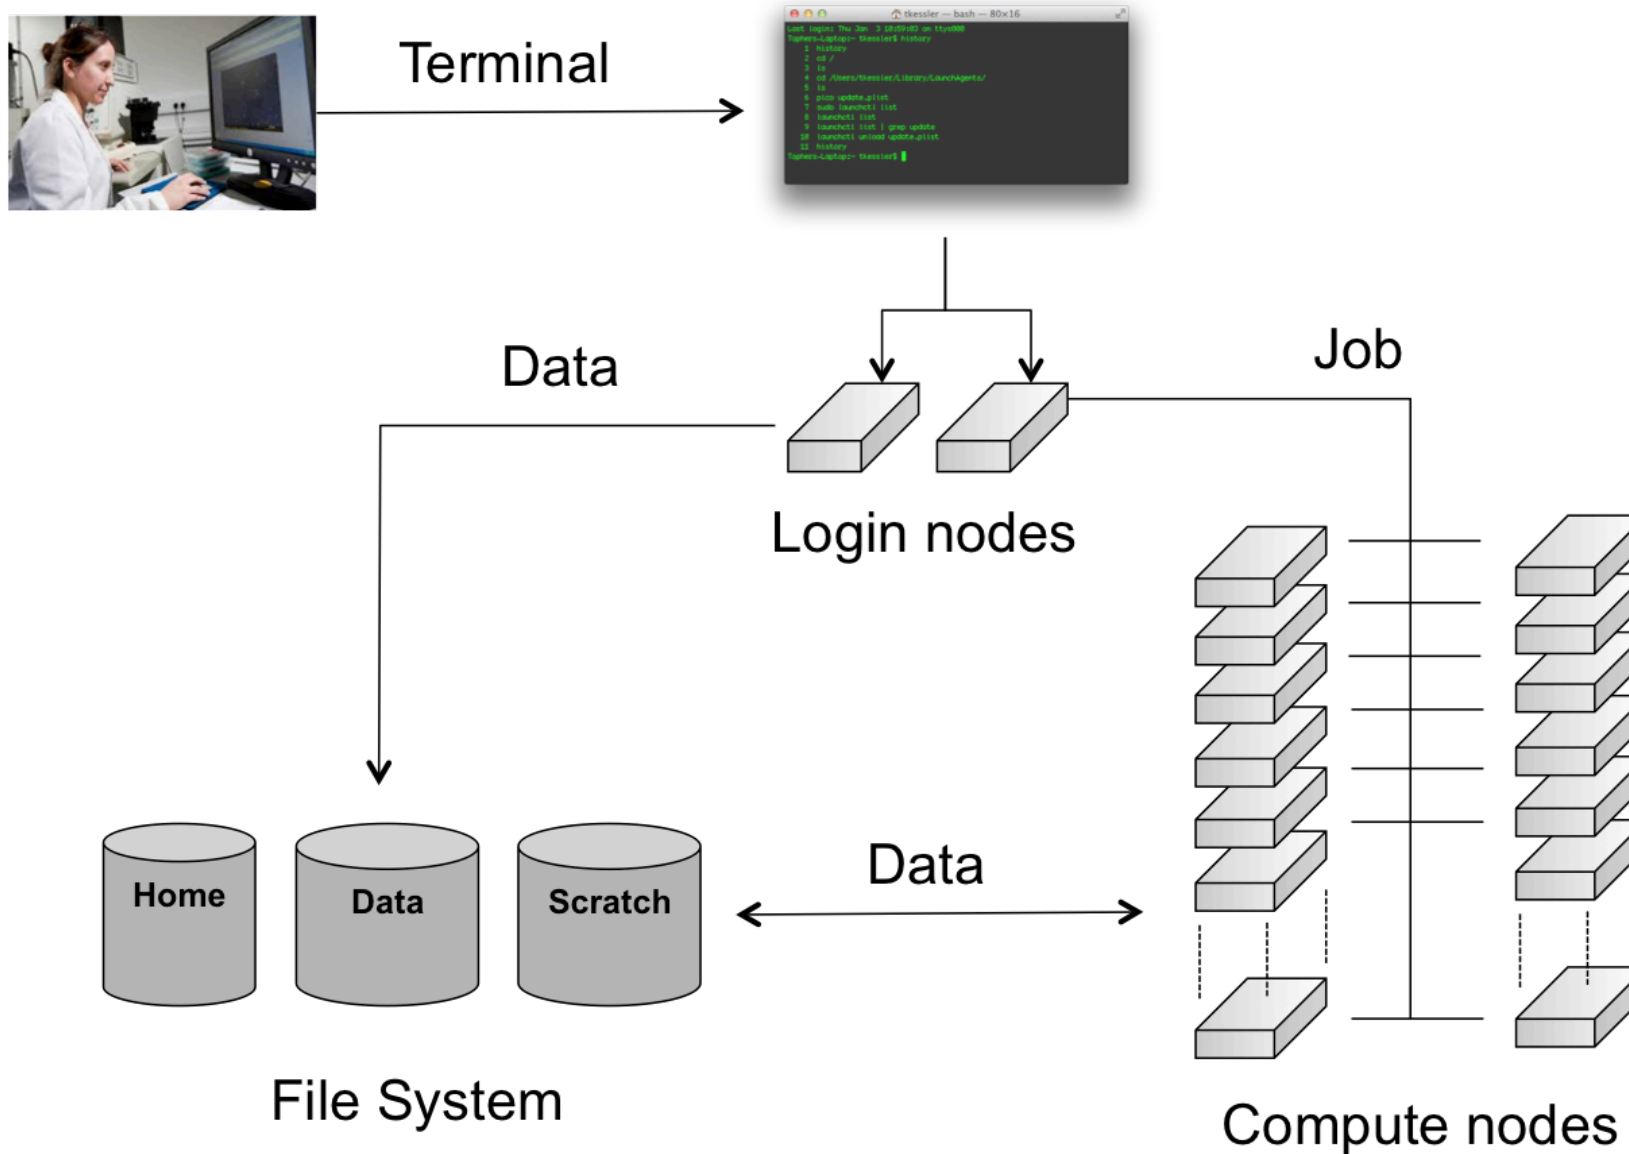
Uid: vsc40023

Institute: Gent

Workflow on HPC infrastructure

1. Connect to login nodes
2. Transfer your files
3. (Compile your code and test it)
4. Create a job script
5. Submit your job
6. Be patient
 - Your job gets into the queue
 - Your job gets executed
 - Your job finishes
7. Move your results

High-level overview of HPC-UGent infrastructure



Connected to an HPC-UGent login node

```
► ssh vsc40023@login.hpc.ugent.be
Last login: Tue Jan  8 19:29:07 2019 from gligarha01.gastly.os

STEVIN HPC-UGent infrastructure status on Tue, 08 Jan 2019 19:20:01

  cluster - full - free - part - total - running - queued
            nodes nodes free  nodes jobs      jobs
-----
delcatty   2    0    0   125   N/A    N/A
golett    71    0  128   200   N/A    N/A
phanpy    15    1    0    16   N/A    N/A
swalot    46    0   42   128   N/A    N/A
skitty    63    0    1    72   N/A    N/A
victini   57    0   32    96   N/A    N/A

For a full view of the current loads and queues see:
http://hpc.ugent.be/clusterstate/
Updates on maintenance and unscheduled downtime can be found on
https://www.vscentrum.be/en/user-portal/system-status

-bash-4.2$ hostname
gligar05.gastly.os
-bash-4.2$ █
```

What is a job script?

```
#!/bin/bash  
echo "hello world"
```

A job (shell) script is a text file that specifies:

- the resources that are required by the calculation
(number of nodes/cores, amount of memory, how much time, ...)
- the software that is used for the calculation
(via `module load` commands)
- the steps that should be done to execute the calculation
(starting from `$HOME`), specified as shell commands, typically:
 - 1) staging in of input files
 - 2) running the calculation
 - 3) staging out of results

Job scripts: full example (single-core job)

```
#!/bin/bash
#PBS -N count_example          ## job name
#PBS -l nodes=1:ppn=1         ## single-node job, single core
#PBS -l walltime=2:00:00      ## max. 2h of wall time

module load Python/3.6.6-intel-2018b
# copy input data from location where job was submitted from
cp $PBS_O_WORKDIR/input.txt $TMPDIR
# go to temporary working directory (on local disk) & run
cd $TMPDIR
python -c "print(len(open('input.txt').read()))" > output.txt
# copy back output data, ensure unique filename using $PBS_JOBID
cp output.txt $VSC_DATA/output_${PBS_JOBID}.txt
```

Job scripts: full example (multi-node job)

```
#!/bin/bash
#PBS -N mpi_hello          ## job name
#PBS -l nodes=2:ppn=all    ## 2 nodes, all cores per node
#PBS -l walltime=2:00:00   ## max. 2h of wall time

module load intel/2018b
module load vsc-mypirun

# go to working directory, compile and run MPI hello world
cd $PBS_O_WORKDIR
mpicc mpi_hello.c -o mpi_hello
mypirun ./mpi_hello
```

Demo: qsub, qstat, qdel

- Submit job scripts from a login node to a cluster for execution using qsub:

```
$ module swap cluster/golett
$ qsub example.sh
12345.master19.golett.gent.vsc
```

- An overview of the active jobs is available via qstat:

```
$ qstat
```

Job id	Name	User	Time Use	S	Queue
-----	-----	-----	-----	-	-----
12345.master19	example	vsc40000	07:39:30	R	long

- To remove a job that is no longer necessary, use qdel:

```
$ qdel 12345
```

Job scheduling

- All our clusters use a fair-share scheduling policy.
- No guarantees on when job will start, so plan ahead!
- Job priority is determined by:
 - historical usage
 - aim is to balance usage over users
 - infrequent/frequent users => higher/lower priority
 - requested resources (# nodes/cores, walltime, memory, ...)
 - large resource request => lower priority
 - time waiting in queue
 - queued jobs get higher priority over time
 - user limits
 - avoid that a single user fills up an entire cluster

Software installations

To submit a request for software installation:

<https://www.ugent.be/hpc/en/support/software-installation-request>

Always include:

- software name and website
- location to download source files
 - or make install files available in your account
- build instructions (if you have them)
- a simple test case with expected output
 - including instructions on how to run it

Requests may take a while to process; make the request sooner rather than later!

Documentation & training

- Documentation is available at:
 - <https://www.vscentrum.be/en/user-portal>
 - <https://www.ugent.be/hpc/en/support/documentation.htm>
 - HPC tutorial, basic Linux tutorial
- Training sessions - <https://www.vscentrum.be/en/education-and-trainings>
 - upcoming sessions in Ghent (see also <https://www.ugent.be/hpc/en/training/training>)
 - **Scientific Python - Wed June 26th 2019**
see also <https://www.ugent.be/hpc/en/training/materials/2019/scipy>

Questions, problems, getting help

Don't hesitate to contact HPC-UGent support: hpc@ugent.be

Always include:

- VSC login id
- clear description of problem (or question)
- location of job script and output/error files in your account
 - don't send them in attachment, we prefer to look at it 'in context'
- job IDs, which cluster

Preferably use your UGent email address.

Alternatives:

- short meeting (for complex problems, big projects)
- [hpc-users mailing list](#)



**GHENT
UNIVERSITY**



Introduction to HPC-UGent

May 27th 2019

<https://www.ugent.be/hpc/en/training/materials/2019/introhpcugent>

hpc@ugent.be

<https://ugent.be/hpc>



Vlaanderen
is computing