



HPC-UGent pilot kickoff meeting



March 14th 2018

<http://users.ugent.be/~kehoste/hpcugent-pilot-kickoff-20180314.pdf>



hpc@ugent.be

<http://ugent.be/hpc>



Pilot users for skitty & victini



- members of `gpilot` user group (invitation only)
- experienced users of existing HPC-UGent Tier-2 infrastructure
- different research domains & applications
- mailing list: hpc-pilot-users@lists.ugent.be
 - used by HPC-UGent team to contact pilot users (status updates, etc.)
 - can be used by pilot users to get in touch with each other

Technical details: skitty



- **72** workernodes, each with:
 - 2x18 Intel Skylake cores
 - **192GB RAM memory (~185GB usable) => ~5GB/core**
 - 1TB local disk + 240GB SSD (caching)

- total: **2,592 cores**
- **EDR Infiniband** interconnect
- **fast access to shared filesystems (GPFS)**
- OS: CentOS 7.4

intended usage:
- multi-node jobs

will replace *delcatty*
(~ summer 2018)



Technical details: victini



- **96** workernodes, each with:
 - 36 cores (2x18 Intel Skylake)
 - **96GB RAM memory (~92GB usable) => ~2.5GB/core**
 - 1TB local disk + 240GB SSD (caching)
- total: **3,456 cores**
- **10 Gb Ethernet** interconnect
- **slow(er) access to shared filesystems (NFS)**
- OS: CentOS 7.4




intended usage:

- **single-node/core jobs**

replacement for *raichu*
(already retired)

future *default* cluster

Technical details: Intel Skylake



- latest server processor generation by Intel
- 2x 18-core Intel Xeon Gold 6140 (36 cores per node)
- 2.3GHz base clock, 3.7GHz max clock (turbo)
18MB L2 cache + 24.75MB L3 cache
- *in theory* 2x higher floating-point performance thanks to AVX-512 (vs AVX2 in Intel Haswell)
 - ~10% speedup for integer applications



Differences with existing Tier-2



- **SLURM** as resource manager (instead of Torque/PBS)
 - Torque/PBS is not sustainable due to support issues
 - wrappers are in place to make this switch **transparent** to users
 - `qsub`, `qstat`, `qalter`, `qdel` commands should still work
 - `#PBS` header lines in job scripts should still work
 - `$PBS_*` environment variables should still be defined
- some (technical) differences w.r.t. how home directories are set up (but you shouldn't notice anything different)



Getting access



- submit jobs from HPC-UGent Tier-2 login nodes
 - only for members of `gpilot` user group
- `module swap cluster/.skitty`
or
`module swap cluster/.victini`
→ hidden cluster modules!
- **if you compile any software yourself,
make sure you do that *on the workernodes!***
(login nodes: Intel Sandy Bridge, `skitty/victini`: Intel Skylake)



Scientific software



- request software installations via new form:
<https://www.ugent.be/hpc/en/support/software-installation-request>
- **only with `foss/2018a` and/or `intel/2018a` toolchains** (or compatible)
 - (or more recent toolchains in the future, of course)
 - very recent compilers are required because of Intel Skylake architecture
 - strong preference for *latest* software versions



Scientific software



- currently available (see `module avail`):
 - ABAQUS, CP2K, FLUENT, MATLAB, molmod, OpenFOAM
- work in progress:
 - AmberTools, Gaussian, gnuplot, ISCE, packmol, phonopy, StaMPS, TAMkin, VASP, yaff



mympirun



- was ported to SLURM (or will be very soon)
- should just work as before

```
module load software/1.2.3
```

```
module load vsc-mympirun
```

```
mympirun ...
```

don't specify a version,
always use latest!

- new: support for `mympirun --dry-run`



Expectations from pilot users



- **testing** usage of new clusters & provided software
- **comparing** with existing Tier-2 (golett, swalot) & Tier-1 clusters
 - re-run jobs you have run before & compare performance
 - try to make comparisons 'honest' (same or similar node/core count)
- **scaling tests** for parallel software
 - don't hold back, try large runs!
- **report back** findings (& problems) to hpc@ugent.be



Attention points (1/2)




- **be very critical** w.r.t. (scientific) results obtained during pilot phase
 - Intel Skylake & AVX-512 are known to cause inaccuracies
 - ~~double~~ triple-check your results!
 - report problems to `hpc@ugent.be` so we can mitigate if needed
- start with small experiments & re-running stuff you've run before
- gradually scale up, run new things when you're more confident



Attention points (2/2)



- **pilot clusters should work like existing Tier-2 clusters**
 - there should be no need to change workflow/job scripts (other than `module load` statements)
 - *if you need to change something to get it to work (well), let us know!*
- **pilot clusters may become unavailable on (very) short notice**
 - down for maintenance to resolve problems or install updates
 - unexpected downtime due to software/hardware/DC problems



Known issues



- only `skitty` is available for now
- no mails from jobs (yet)
- `#!/bin/bash -l` needed (?) in jobs to get `module` command
- `mympirun` in SLURM jobs not tested extensively (4.1.0-beta)


*These problems will hopefully be resolved soon,
updates will be sent to `hpc-pilot-users` mailing list.*



Timeline (preliminary)



- *March 14th 2018 (today):* access for pilot users
- *ASAP:* all requested software available
- *May 2018:* follow-up meeting for skitty/victini pilot phase
- *summer 2018:* skitty & victini go into production
- *Fall 2018:* switch default cluster to victini + retire delcatty cluster
- *mid 2019:* switch to SLURM on all Tier-2 clusters



Problems or questions?



- contact hpc@ugent.be
- make it clear in e-mail subject that it's related to pilot clusters
- provide clear problem description
 - what did you expect to work, what went wrong
 - mention relevant error messages, job IDs, etc.
 - mention location of output files in your account (please don't send them in attachment)
 - exact steps to reproduce the problem