

cesnet
metacentrum
.....



Grid service MetaCentrum

Jiří Vorel

MetaCentrum User Support

vorel@cesnet.cz meta@cesnet.cz

April 13 2023

Prague, The National Library of Technology



■ MetaCentrum is

- ... The National Grid Infrastructure (NGI) <https://metacentrum.cz>
- ... the activity of the CESNET association <https://metavo.metacentrum.cz>
- ... a provider of **computational resources, application tools** (commercial and free/open source) and **data storage** <https://docs.metacentrum.cz/>
- ... **free of charge** <https://docs.metacentrum.cz/access/terms/>
 - Users "pay" by Acknowledgement in their research publications

■ MetaCentrum is available for

- ... employees and students from Czech universities, the Czech Academy of Science, non-commercial research facilities, etc. <https://metavo.metacentrum.cz/en/myaccount/pubs>
- ... industry users (only for non-profit and public research)

■ CPU

- ~45,000 CPU cores (x86_64) available
- **Intel**, AMD (minority); **Debian 11**, CentOS 7 (minority)
- Typically 32/64 CPU, up to 1 TB RAM (400-700 GB)
- Special machines (up to 504 CPU, 10 TB RAM, CentOS 7)
- Elixir-CZ nodes up to 3 TB RAM

<https://docs.metacentrum.cz/advanced/pbs-options/>

<https://www.elixir-czech.cz/>

■ GPU

- 16 clusters, more than 400 GPU cards
- NVIDIA A10, A40, A100, RTX A4000, Tesla*, GeForce*
- **Max. 45 GB memory**

<https://docs.metacentrum.cz/advanced/gpu-comput/>

■ Fill out and submit the registration form

<https://metavo.metacentrum.cz/en/application/index.html>

- Select your organisation (click on the eduID logo)
- Use your institutional username and password
- Fill out the form and create a strong password
- Users must extend MetaCentrum membership from the beginning of each calendar year (typically during January)
- MetaCentrum users obtain access to CERIT-SC resources automatically



■ Read our documentation, FAQ and tutorial for beginners

<https://docs.metacentrum.cz/>

<https://wiki.metacentrum.cz/wiki/> - will be deprecated

<https://docs.metacentrum.cz/access/account/>

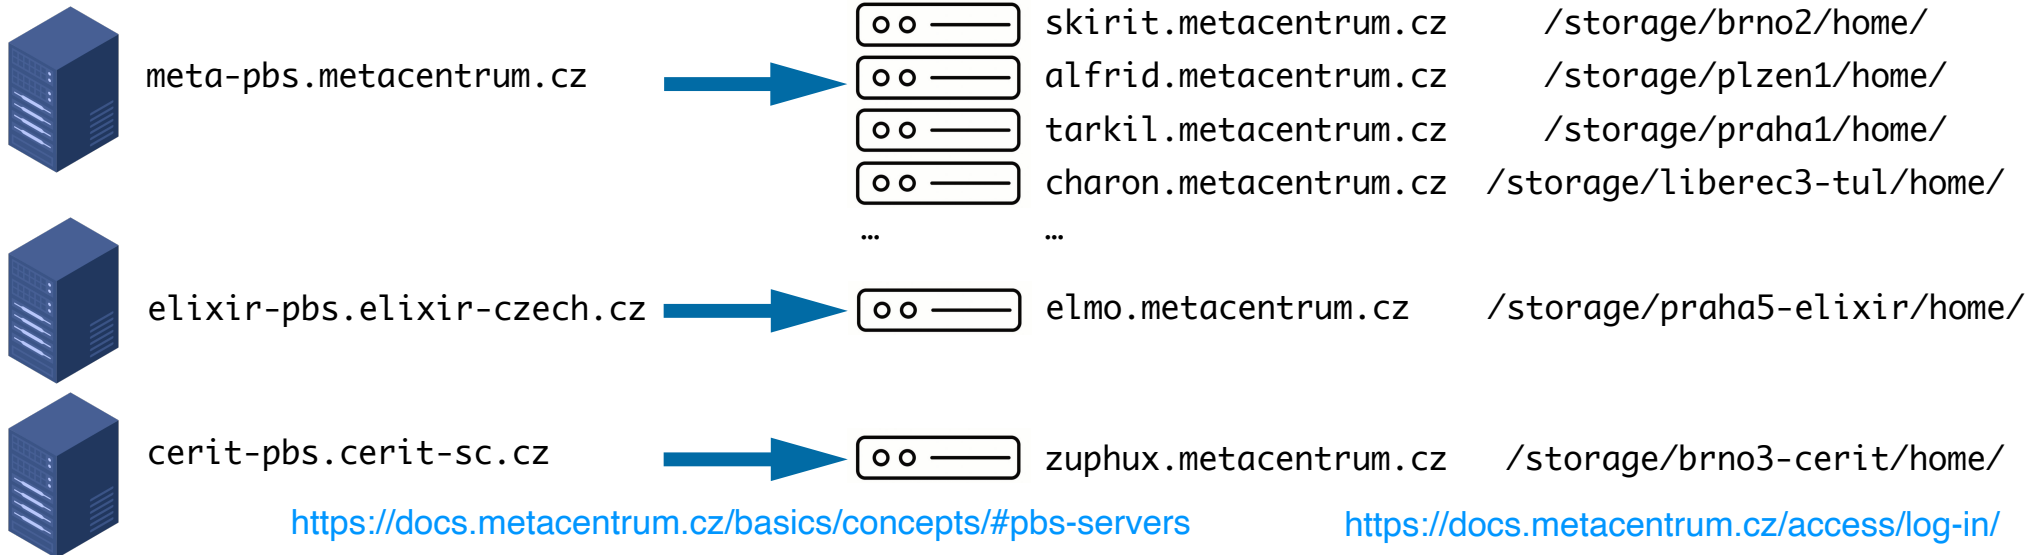
<https://docs.metacentrum.cz/troubleshooting/faqs/>

- Gateway to the entire grid infrastructure
- Accessible via **ssh with a password** (ssh tickets are not supported)
- Frontends submit jobs to PBS servers
- Frontends are relatively small virtual machines mainly for writing scripts for batch jobs, checking applications and user data etc.

- **Do not run long and/or demanding calculations directly on frontends!**
(rather start an interactive job)

- Frontend servers usually have different home directories
- Command line interface <https://docs.metacentrum.cz/basics/concepts/#frontends-storages-homes>
<https://docs.metacentrum.cz/advanced/grid-infrastruct/#frontends-and-storages>

- Ten frontends (+ one alias) submit jobs to three PBS servers
- PBS (Portable Batch System) is a software that performs job scheduling
- **Frontend servers can have different home directories**
- All user home directories are available from all frontends



- Data is stored on a few independent storages; the capacity is not infinite
- **All storages are accessible through all frontends and computational nodes**
- Storages have quotas for the total volume of data and the number of files
- MetaCentrum storage capacities are dedicated mainly to data in active usage

NFS4 server	adresář - directory	velikost - capacity	zálohovací třída - back-up policy
storage-brno1-cerit.metacentrum.cz	/storage/brno1-cerit/	1.8 PB	2
storage-brno2.metacentrum.cz	/storage/brno2/	306 TB	2
storage-brno11-elixir.metacentrum.cz	/storage/brno11-elixir/	313 TB	2
storage-brno12-cerit.metacentrum.cz	/storage/brno12-cerit/	3.4 PB	2
storage-budejovice1.metacentrum.cz	/storage/budejovice1/	44 TB	3

<https://docs.metacentrum.cz/advanced/mount-storages/>

```
name@my_pc:~$ ssh vorel@perian.metacentrum.cz
```

```
vorel@perian.metacentrum.cz's password: ← Type a password
```

```
vorel@perian:~$ pwd
```

```
/storage/brno2/home/vorel ← Where am I?
```

```
vorel@perian:~$ cd /storage/plzen1/home/vorel
```

```
vorel@perian:~$ pwd
```

```
/storage/plzen1/home/vorel ← Hmm, I forgot that my data is on different storage...
```

```
name@my_pc:~$ ssh vorel@minos.metacentrum.cz
```

```
vorel@minos.metacentrum.cz's password: ← Direct access to the same storage through different frontend
```

```
vorel@minos:~$ pwd
```

```
/storage/plzen1/home/vorel
```

<https://docs.metacentrum.cz/access/log-in/>

- SSH keys for logging into frontends **are not fully supported**. We want to "force" you to generate a Kerberos ticket (with a limited lifetime) by typing the password

```
jirivorel@MacBook ~$ ssh vorel@nympha.metacentrum.cz
```

```
vorel@nympha.metacentrum.cz's password:
```

```
(BULLSEYE)vorel@nympha:~$ klist
```

```
Credentials cache: FILE:/tmp/krb5cc_1597_LTYWlt
Principal: vorel@META
```

Type a password

klist command prints the status of issued tickets

Issued	Expires	Principal
May 6 11:22:55 2022	May 6 21:22:55 2022	krbtgt/META@META
May 6 11:22:55 2022	May 6 21:22:55 2022	afs/ics.muni.cz@META
May 6 11:22:55 2022	May 6 21:22:55 2022	krbtgt/ZCU.CZ@META
May 6 11:22:55 2022	May 6 21:22:55 2022	afs/zcu.cz@ZCU.CZ

```
(BULLSEYE)vorel@nympha:~$ ssh halmir1
Linux halmir1.metacentrum.cz 5.10.0-13-amd64 #1 SMP Debian 5.10.106-1+zs1 (2022-03-28) x86_64
Last login: Thu Apr 21 09:54:05 2022 from elmo2-4.hw.elixir-czech.cz
(BULLSEYE)vorel@halmir1:~$
```

```
(BULLSEYE)vorel@nympha:~$ klist
klist: No ticket file: /tmp/krb5cc_1597_rw50KaLk0H
(BULLSEYE)vorel@nympha:~$ qsub -I -l select=1:ncpus=1:mem=5gb:scratch_local=1gb -l walltime=1:00:00
No Kerberos credentials found.
(BULLSEYE)vorel@nympha:~$ ssh halmir1
vorel@halmir1's password:

(BULLSEYE)vorel@nympha:~$ kinit ← kinit command generates new tickets
vorel@META's Password:
```

- Kerberos is a network authentication protocol
- You can have the Kerberos ticket issued on your personal computer. During the validity of the ticket, you can log in to every frontend, compute node or storage without entering a password again

- **A new version of modules has been recently deployed**
- Each software (in a specific version) is prepared as an individual module
- The module file, after activation (`module add`), will load the main application, all dependencies and all needed libraries
- Modules are shown based on the compatibility with the OS and CPU architecture
- MetaCentrum contains >3,000 modules; **a license agreement may be required**
- Available modules can be listed on the frontends (`module ava`)

```
(BULLSEYE)vorel@nympha:~$ module -h
Modules Release 5.1.0 (2022-04-30)
Usage: module [options] [command] [args ...]
```

<https://docs.metacentrum.cz/software/modules/>

<https://docs.metacentrum.cz/basics/concepts/#modules>

```
(BULLSEYE)vorel@nympha:~$ module ava Mum
----- /packages/run/modules-5/debian11avx512 -----
mummer/  mumps/

Key:
modulepath  directory/
(BULLSEYE)vorel@nympha:~$ module ava *mme*
----- /packages/run/modules-5/debian11avx512 -----
glimmerhmm/  hmmer/  mummer/  py-snowballstemmer/

Key:
modulepath  directory/
(BULLSEYE)vorel@nympha:~$ module ava mummer/
----- /packages/run/modules-5/debian11avx512 -----
mummer/3.23  mummer/4.0.0beta2

Key:
modulepath
```

case insensitive

wildcard asterisk

- Pay attention to library prefixes (mainly Perl, Python, R)

```
(BULLSEYE)vorel@nympha:~$ module ava r
```

```
----- /packages/run/modules-5/debian11avx512 -----
r-abind/          r-coda/          r-fastmap/       r-gplots/        r-lpsolve/       r-prettyunits/
r-amap/           r-codetools/    r-fastmatch/     r-gridextra/     r-lubridate/     r-processx/
r-annotate/      r-colorspace/   r-filelock/      r-gtable/        r-magrittr/      r-prodlim/
```

```
(BULLSEYE)vorel@nympha:~$ module ava py
```

```
----- /packages/run/modules-5/debian11avx512 -----
py-absl-py/      py-datrie/       py-isodate/      py-opt-einsum/
py-alabaster/    py-debugpy/     py-jedi/         py-packaging/
py-anyio/        py-decorator/   py-jinja2/      py-pandas/
```

```
(BULLSEYE)vorel@nympha:~$ module ava perl
```

```
----- /packages/run/modules-5/debian11avx512 -----
perl-acme-damn/    perl-devel-symlump/  perl-gdgraph/    perl-libwww-perl/    per
perl-algorithm-diff/  perl-encode-locale/  perl-gdtextutil/  perl-libxml-perl/    per
perl-alien-build/    perl-error/          perl-graph/      perl-list-moreutils-xs/  per
```

- HW resources (CPUs, GPUs, RAM, scratch, walltime,...) are reserved by PBS
- Detailed documentation: https://wiki.metacentrum.cz/wiki/About_scheduling_system
- It requires some experience (keep some reserve)
- Helper tool for qsub command assembly

Go to metavo.metacentrum.cz -
Current state - Personal view - **Qsub assembler**

(Stav zdrojů - Osobní pohled
Sestavovač qsub)

<https://metavo.metacentrum.cz/pbsmon2/person>

Personal view

This page shows a personal view of the PBS system for the user **vorel**.

Jobs of user "vorel"

user	job count					CPU count				
	total	queued	running	completed	other	total	queued	running	completed	other
vorel	0	0	0	0	0	0	0	0	0	0

Links

- list of my jobs
- **personal view of storages**
- **qsub assembler**

Click on it...

qsub -l walltime=1 : 0 : 0 -q default@meta-pbs.metacentrum.cz \

-l select=1 :ncpus=1 :ngpus=0 :mem= 400 mb :scratch_ local = 400 ✓ mb
gb

cluster ...

city ...

SPECfp2017 per core ...

other resources ...

:arch=

:biocev=

:centos=

:cgroups=

:cluster=

:cpu_flag=

:cpu_vendor=

:cuda_version=

:gpu_cap=

:host=

:hyperthreading=

:infiniband=

:luna=

:os=

:osfamily=

:pbs_server=

:pruhonice=

:scratch_shm=

:uv=

:vestec=

:vnode=

CPU performance

:spec=7.4

✓ :spec=8.0

✓ :spec=9.1

✓ :spec=10.4

Click on it...

Find machines matning the resource specification

■ And you will see...

selection from command line

```
qsub -l walltime=24:0:0 -q default@meta-pbs.metacentrum.cz -l select=1:ncpus=8:mem=100gb:scratch_ssd=50gb
```

selection in shell script

```
#!/bin/bash
#PBS -q default@meta-pbs.metacentrum.cz
#PBS -l walltime=24:0:0
#PBS -l select=1:ncpus=8:mem=100gb:scratch_ssd=50gb
#PBS -N my_awesome_job
```

Result

OK

The requirement is 1 machine, and 93 such machines are free, out of 289 machines matching the requirements. The job for it.

Machines available right now

adan1 (32 CPU, 187.6 GIB RAM, 697.6 GIB HDD)	adan2 (32 CPU, 187.6 GIB RAM, 783.6 GIB HDD)	adan3 (16 CPU, 171.6 GIB RAM, 766.6 GIB HDD)	adan5 (32 CPU, 187.6 GIB RAM, 744.6 GIB HDD)	adan6 (32 CPU, 187.6 GIB RAM, 705.4 GIB HDD)
--	--	--	--	--

- For safety reasons, it is not possible to implement direct submission to PBS
- Be aware of inappropriate combinations of resources
- The PBS can also accept inappropriate jobs (but the execution will never start)

Personal view

This page shows a personal view of the PBS system for the user **vorel**.

Jobs of user "vorel"

user	job count					CPU count				
	total	queued	running	completed	other	total	queued	running	completed	other
vorel	0	0	0	0	0	0	0	0	0	0

Links

- list of my jobs
- personal view of storages
- qsub assembler

úloha	server	CPU	vyhrasz. paměť	použitá paměť	jméno	uživatel	CPU čas	čas běhu	stav	fronta
	meta	1	1gb	0b	STDIN				Q - ve frontě	q_1d@meta-pbs.metacentrum.cz
požadované prostředky	1:ncpus=1:mem=1gb:scratch_local=1gb:os=debian10:mpiprocs=1:ompthreads=1									
vytvořena	čtvrtek 1. prosince 2022 14:05:23									
způsobilá k běhu	čtvrtek 1. prosince 2022 14:05:23									
poslední změna stavu	čtvrtek 1. prosince 2022 14:06:24									
komentář	Not Running: Insufficient amount of resource: os									

- Not all visible queues are suitable for direct use
- Explore the `-q` option of the `qsub` assembler

The screenshot shows the `qsub` assembler interface with various options like `-l walltime=1:0:0` and `-q` circled in red. A list of queues is displayed, with colored arrows pointing to specific entries:

- Blue arrow:** Points to `default@meta-pbs.metacentrum.cz` and `default@cerit-pbs.cerit-sc.cz`. Description: Queues for jobs requesting up to 720 hours.
- Yellow-green arrow:** Points to `even@meta-pbs.metacentrum.cz`. Description: GPU jobs up to 24 hours on MetaCentrum nodes.
- Orange arrow:** Points to `gpu_long@meta-pbs.metacentrum.cz`. Description: GPU jobs up to 336 hours on MetaCentrum nodes.
- Green arrow:** Points to `global@meta-pbs.metacentrum.cz` and `global@cerit-pbs.cerit-sc.cz`. Description: Queues prioritising jobs requesting more than 500 GB RAM.
- Black arrow:** Points to `gpu@cerit-pbs.cerit-sc.cz`. Description: GPU jobs up to 24 hours on CERIT-SC nodes.
- Red arrow:** Points to `phi@cerit-pbs.cerit-sc.cz`. Description: Nodes with Intel Xeon Phi 7210.
- Purple arrow:** Points to `uv@cerit-pbs.cerit-sc.cz`. Description: SMP machines with OS CentOS 7.

Queue default@meta-pbs.metacentrum.cz

Default queue (routing)

<https://docs.metacentrum.cz/basics/concepts/#queues>

The queue is routing, it delivers jobs depending on their walltime to the following queues:

queue	Priority	time limits	jobs				max CPUs per user
			queued	running /max	completed	total	
q_2h@meta-pbs.metacentrum.cz	50	0 - 02:00:00	381	0 /	5676	6058	2000
q_4h@meta-pbs.metacentrum.cz	50	02:00:01 - 04:00:00	1001	1057 /	12072	18078	
q_1d@meta-pbs.metacentrum.cz	50	04:00:01 - 24:00:00	2270	100 /	4153	6536	4000
q_2d@meta-pbs.metacentrum.cz	50	24:00:01 - 48:00:00	126	11 /	150	287	1000

uv_bio@cerit-pbs.cerit-sc.cz	🔒	31	00:00:01 - 96:00:00	0	0 /
uv_small@cerit-pbs.cerit-sc.cz	🚫	30	00:00:01 - 96:00:00	20	12 /
fireprot_devel@cerit-pbs.cerit-sc.cz	🔒	30	0 - 336:00:00	0	0 /

global@elixir-pbs.elixir-czech.cz		50	0 - 48:00:00
elixircz@elixir-pbs.elixir-czech.cz	📄	0	0 - 720:00:00

do not submit to the queue directly, use a routing queue instead

reserved for: leontovyc_roman simekmilos vorel

```
#!/bin/bash
#PBS -q default@meta-pbs.metacentrum.cz
#PBS -l walltime=24:0:0
#PBS -l select=1:ncpus=8:mem=100gb:scratch_ssd=50gb
#PBS -N my_awesome_job
#PBS -m e

# test if a scratch directory exists
# variable SCRATCHDIR is set automatically
test -n "$SCRATCHDIR" || { echo >&2 "Variable SCRATCHDIR is not set!"; exit 1; }

# set a DATADIR variable
DATADIR=/storage/brno12-cerit/home/vorel/data/

# copy input file "data.fa" to the scratch directory
cp $DATADIR/data.fa $SCRATCHDIR

# move into the scratch directory
cd $SCRATCHDIR

# load a module for your application
module add blast-plus/blast-plus-2.12.0-gcc-8.3.0-ohlv7t4

# run the calculation
# do not forgeto to use reserved CPUs by '-num_threads' flag
# variable PBS_NCPUS is a number of CPUs requested for the entire job
blastp -query data.fa <other_parameters> -num_threads $PBS_NCPUS -out results.txt

#copy results
cp results.txt $DATADIR

# clean the scratch directory
clean_scratch
```

- Define HW resources (**-l**), queue (**-q**), and walltime (**-l**), set the job name (**-N**) and email alert (**-m**)
- You can define as many variables as you want
- Available modules can be listed by command **module avail *key_word*** on any frontend
- The scratch directory will be cleaned automatically

<https://docs.metacentrum.cz/basics/jobs/>

- The opposite of batch jobs (waiting for the user's input...)
- Best choice for test calculations (which should not be run directly on frontends)
- An interactive job is requested by the qsub command with the `-I` (uppercase "i") option

<https://docs.metacentrum.cz/basics/jobs/>

```
(BUSTER)vorel@skirit:~$ qsub -I -l select=1:ncpus=4:mem=50gb:scratch_local=30gb -l walltime=1:00:00
qsub: waiting for job 11405230.meta-pbs.metacentrum.cz to start
qsub: job 11405230.meta-pbs.metacentrum.cz ready
```

```
vorel@zenon31:~$ cd $SCRATCHDIR
vorel@zenon31:/scratch.ssd/vorel/job_11405230.meta-pbs.metacentrum.cz$ module add orca/orca-5.0.1-intel-19.0.4-bnofsgq
vorel@zenon31:/scratch.ssd/vorel/job_11405230.meta-pbs.metacentrum.cz$ module list
Currently Loaded Modulefiles:
  1) metabase                2) openmpi/openmpi-4.0.4-intel-19.0.4-gpu-xri6uan  3) orca/orca-5.0.1-intel-19.0.4-bnofsgq
vorel@zenon31:/scratch.ssd/vorel/job_11405230.meta-pbs.metacentrum.cz$
vorel@zenon31:/scratch.ssd/vorel/job_11405230.meta-pbs.metacentrum.cz$ ...time for coffee...
-bash: ...time: command not found
vorel@zenon31:/scratch.ssd/vorel/job_11405230.meta-pbs.metacentrum.cz$ orca < input > output
```

- Temporary storage on physical computing nodes
- **Very intensive operations can cause network overload and the slowdown of central storage** (/storage/city/...)
- Copy the input data into the scratch directory on a dedicated machine
- Variable SCRATCHDIR is set automatically
- Faster, more stable

```
qsub -l select=1:ncpus=1:mem=4gb:scratch_local=10gb -l walltime=1:00:00
```

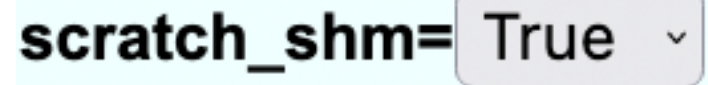
```
cp my_input_data.txt $SCRATCHDIR
```

```
...
```

```
cp $SCRATCHDIR/my_results.txt /storage/city/home/user_name/
```

<https://docs.metacentrum.cz/advanced/grid-infrastruct/#scratch-storages>

- MetaCentrum offers four types of scratch
 - `scratch_local` <https://docs.metacentrum.cz/basics/concepts/#scratch-directory>
 - on every node, HDD, default
 - `scratch_ssd`
 - fast SSD, typically smaller in volume, not everywhere
 - `scratch_shared`
 - network volume which is shared between all clusters in a given location, not everywhere
 - `scratch_shm`
 - scratch held in RAM, very fast, on every node
 - boolean type (True/False), limited by mem parameter (:mem=XYgb)

A screenshot of a configuration interface showing the parameter `scratch_shm` set to `True`. The text is displayed in a light blue box with a dropdown arrow on the right side.

`scratch_shm= True`

- Do not forget to clean the scratch directory when your calculation is done or has been killed by PBS
- You can do it **manually** after each finished job (but it won't be very pleasant) or **activate utility** `clean_scratch`

```
trap 'clean_scratch' TERM EXIT
cp my_input_data.txt $SCRATCHDIR
cd $SCRATCHDIR
...
...
cp my_results.txt /storage/city/home/... || export CLEAN_SCRATCH=false
```

<https://docs.metacentrum.cz/advanced/job-tracking/#trap-the-term>

■ GPU calculations

<https://docs.metacentrum.cz/snippets/GPU-clusters-table/>

- 16 clusters, more than 400 GPU cards, max. 45 GB memory
- Maximum si **eight** GPU cards on a single machine, typically two or four
- Three dedicated GPU queues
 - `gpu@meta-pbs.metacentrum.cz` (up to 24 hours)
 - `gpu_long@meta-pbs.metacentrum.cz` (up to 336 hours)
 - `gpu@cerit-pbs.cerit-sc.cz` (up to 24 hours)
- Jobs can migrate between PBS servers
- cuda versions 11.2 and 11.4

```
qsub -l walltime=4:0:0 \  
-q gpu@meta-pbs.metacentrum.cz \  
-l select=1:ncpus=1:ngpus=1:mem=10gb:scratch_local=20gb
```


- Each GPU calculation (`ngpus=1`) needs at least one CPU (`ncpus=1`)
- Remember that the newest GPU is NOT the best for all jobs
- One GPU card per job is enough for novices
- GPU card can not be shared and is entirely dedicated to one calculation
- GPU calculations can be monitored on the same computation nodes by `nvidia-smi` or `nvidia-smi` command
- In most cases is not wise to target one specific cluster (e.g. `:cl_adan=True`), select a smaller set of machines using the parameters:
 - `gpu_mem=20gb` (minimum amount of memory on card)
 - `gpu_cap=cuda80` (compute capability)
 - `cuda_version=11.4` (cuda version)

<https://docs.metacentrum.cz/advanced/gpu-comput/>

cesnet
metacentrum
.....




**Common issues and
how to deal with them**

Tips & tricks



- Do not use frontends, copy data directly on storage, use compressed files (.tar, .zip, .gz, etc.)
- SFTP client for Windows users (WinSCP, FileZilla, CyberDuck)



```
scp my_data.gz vorel@skirit.metacentrum.cz:\  
/storage/praha5-elixir/home/vorel
```



```
scp my_data.gz \  
vorel@storage-praha5-elixir.metacentrum.cz:~
```



<https://docs.metacentrum.cz/data/data-within/>

<https://docs.metacentrum.cz/advanced/mount-storages/>

<https://docs.metacentrum.cz/data/data-within/#moderate-data-handling>

- MetaCentrum storage capacities are dedicated mainly to data in active usage
- Unnecessary data should be removed or moved to CESNET Storage Department for long-term archiving



- MetaCentrum users can use the following archive

/storage/du-cesnet/home/user_name/V0_metacentrum-tape_tape-archive/

- And for backup (12 months cycle)

/storage/du-cesnet/home/user_name/V0_metacentrum-tape_tape/

<https://du.cesnet.cz/en/start>

<https://docs.metacentrum.cz/data/backup-archive/>

- Users can install the software on their own (in home directories)
- No restrictions; do not violate the license terms and conditions or/and our rules
- Users do not have sudo rights and can not right outside of the home directory

```
(BULLSEYE)vorel@skirit:~$ apt-get install package_name
E: Could not open lock file /var/lib/dpkg/lock-frontent - open (13: Permission denied)
E: Unable to acquire the dpkg frontend lock (/var/lib/dpkg/lock-frontent), are you root?
```

- Python (pip, venv)

```
(BULLSEYE)vorel@skirit:~$ module ava py-pip/
----- /packages/run/modules-5/debian11avx512 -----
py-pip/19.3-intel-19.0.4-hudzomi  py-pip/21.3.1-gcc-10.2.1-mjt74tn
```

- R packages (with lib="user/path")

```
(BULLSEYE)vorel@skirit:~$ module ava r/
----- /packages/run/modules-5/debian11avx512 -----
r/2.14.0      r/3.1.1      r/3.5.1-gcc      r/4.0.2-intel-19.0.4-5vzfhtq
r/3.0.1      r/3.2.3-intel  r/3.6.2-gcc      r/4.1.1-intel-19.0.4-ilb46fy
r/3.0.3      r/3.3.1-intel  r/4.0.0-gcc      r/4.1.1-intel-19.0.4-xrup2b3
r/3.1.0      r/3.4.0-gcc   r/4.0.2-aocc-2.2.0-q43q56w  r/4.1.3-gcc-10.2.1-6xt26d1
r/3.1.0shlib r/3.4.3-gcc   r/4.0.2-aocc-2.2.0-zrf6vyw  r/4.2.1-intel-19.0.4-d3gtjq7
```

<https://docs.metacentrum.cz/software/install-software/>

- **Pre-compiled binaries** can be directly downloaded/copied into \$SCRATCHDIR

```
(BULLSEYE)vorel@skirit:~$ qsub -I -l select=1:ncpus=1:mem=5gb:scratch_local=1gb -l walltime=1:00:00
qsub: waiting for job 14986173.meta-pbs.metacentrum.cz to start
qsub: job 14986173.meta-pbs.metacentrum.cz ready

(BULLSEYE)vorel@elmo3-1:~$ cd $SCRATCHDIR
(BULLSEYE)vorel@elmo3-1:/scratch/vorel/job_14986173.meta-pbs.metacentrum.cz$ wget -q https://www.drive5.com/downloads/usearch11.0.667_i86linux32.gz
(BULLSEYE)vorel@elmo3-1:/scratch/vorel/job_14986173.meta-pbs.metacentrum.cz$ gunzip usearch11.0.667_i86linux32.gz
(BULLSEYE)vorel@elmo3-1:/scratch/vorel/job_14986173.meta-pbs.metacentrum.cz$ chmod u+x usearch11.0.667_i86linux32
(BULLSEYE)vorel@elmo3-1:/scratch/vorel/job_14986173.meta-pbs.metacentrum.cz$ ./usearch11.0.667_i86linux32
usearch v11.0.667_i86linux32, 4.0Gb RAM (791Gb total), 112 cores
(C) Copyright 2013-18 Robert C. Edgar, all rights reserved.
https://drive5.com/usearch
```

- **Perl** (cpanm) libraries
- **Conda/Miniconda/Mamba/Micromamba** package managers
- Do your **own compilations** (GCC, Intel oneAPI, AOCC for AMD CPUs, Open MPI, CUDA for GPU support, CMake, ...)

- Computing nodes and frontends have limited quotas (~1 GB) for writing out of the scratch and home directory
- Exceeding this quota will cause the termination of the process
- The most common problems are caused by:
 - Write to /tmp
 - Very large stdout and stderr streams

```
export TMPDIR=$SCRATCHDIR
```

```
my_app < input ... 1>$SCRATCHDIR/stdout 2>$SCRATCHDIR/stderr
```

- Utility `check-local-quota` can be executed on each node

- `go_to_scratch` utility
 - Realtime monitoring of running jobs; redirection to computational node
`go_to_scratch job_ID@PBS_server_full_name`
- `qextend` utility <https://docs.metacentrum.cz/advanced/pbs-options/#qextend>
 - Walltime, which could be reserved by PBS, is limited to 720 hours
 - Users are allowed to prolong their jobs in a limited number of cases
`qextend job_ID@PBS_server_full_name additional_walltime_hh:mm:ss`
- `pbs-get-job-history` utility <https://docs.metacentrum.cz/advanced/job-tracking/#finished-jobs>
 - Complex information about current or historical jobs
`pbs-get-job-history job_ID@PBS_server_full_name`

- Text files created on Windows PS use more characters for the termination of a lines
- This format can not be read by Unix-like systems
- Individual lines are not recognised
- Utility dos2unix can fix the line terminators
- Typical error:
 - `'\r': command not found , EXIT^M: invalid signal specification`

https://owasp.org/www-community/vulnerabilities/CRLF_Injection

```
[vorel@zuphux ~]$ file example.txt
example.txt: ASCII text, with very long lines, with CRLF line terminators
[vorel@zuphux ~]$ dos2unix example.txt
dos2unix: converting file example.txt to Unix format ...
[vorel@zuphux ~]$ file example.txt
example.txt: ASCII text, with very long lines
[vorel@zuphux ~]$
```

Command file determines the type of a file

Problem detected

- We keep OS Debian up-to-date on our nodes
- We upgraded completely from Deb10 (BUSTER) to Deb11 (BULLSEYE)
- However, some libraries may be missing in the new system...

```
gmx_mpi: error while loading shared libraries: libevent_core-2.1.so.6:  
cannot open shared object file: No such file or directory
```

- Therefore we provide universal modules with these missing libraries

```
(BULLSEYE)vorel@skirit:~$ module ava debian*  
----- /packages/run/modules-5/debian11avx512  
debian7/  debian8/  debian9/  debian10/  
  
Key:  
modulepath  directory/  
(BULLSEYE)vorel@skirit:~$ module add debian10/compat  
(BULLSEYE)vorel@skirit:~$ ls /software/debian-compat/debian10/lib/
```

- Users can still use other (older) modules...

- From the point of view of performance (necessary PBS hardware requirements to run every single job), an ideal job is running for at least 30 minutes
- Startup overhead may be a significant part of the whole processing time
- Aggregate short jobs into bigger groups with longer walltime

`-l walltime=00:30:00` (and more)

- Optimise your calculations (hardware usage)
- Reservation of too many resources decreases your fairshare score and reduces the priority for your future jobs

<https://wiki.metacentrum.cz/wiki/Fairshare>

- You can increase your fairshare score by acknowledgement to MetaCentrum in your publications

<https://docs.metacentrum.cz/access/terms/#acknowledgements-and-publications>

- Effectivity can be checked on the computation node by standard Linux tools (top, htop) or on metavo.metacentrum.cz web portal

- There is no reason to be afraid to use MetaCentrum
- You can find plenty of information and instructions in our documentation

<https://docs.metacentrum.cz/>

<https://docs.metacentrum.cz/troubleshooting/faqs/>

- If you are lost - send an email <https://docs.metacentrum.cz/contact/write-request/>

meta@cesnet.cz

- If grid infrastructure does not fulfil your expectations, maybe the MetaCentrum Cloud service would be a better choice

<http://cloud.metacentrum.cz/>

cesnet
metacentrum
.....



THANK YOU FOR YOUR ATTENTION

meta@cesnet.cz vorel@cesnet.cz