




University  
of Manitoba



Digital Research  
Alliance of Canada

---

# Software Environments on High Performance Computing clusters: Software stacks and containers

*UofM-Spring-Workshop 2023*  
*May 17<sup>th</sup>-19<sup>th</sup>, 2023*

*Ali Kerrache*  
*HPC Analyst*



- ★ Software distribution on HPC clusters:
  - Local installation
  - modules
- ★ Software stacks on Grex
  - GrexEnv
  - CCEnv
- ★ Containers:
  - Singularity
  - **apptainer**

# Software layers

**User layer:** Python packages, Perl and R modules, home made codes, ...

User

**Software stacks:** modules for Intel, PGI, OpenMPI, CUDA, MKL, high-level applications. Multiple architectures (sse3, avx, avx2, avx512)

RSNT

**Nix or gentoo:** GNU libc, autotools, make, bash, cat, ls, awk, grep, etc.

**Gray area:** Slurm, Lustre client libraries, IB/OmniPath/InfiniPath client libraries (all dependencies of OpenMPI) in Nix {or gentoo} layer, but can be overridden using PATH & LD\_LIBRARY\_PATH.

Sys. Admin

**OS:** kernel, daemons, drivers, libcuda, anything privileged (e.g. the sudo command): always local. Some legally restricted software too (VASP).



# Why modules?

## ★ Why modules?

- Control different versions of the same program.
- Avoid conflicts between different versions and libraries.
- Set the right path to each program or library.



## ★ Useful commands for working with modules:

- module **list**; module **avail**
- module **spider** <soft>/<version>
- module **load** soft/version; module **unload {rm}** <soft>/<version>
- module **show** soft/version; module **help** <soft>/<version>
- module **purge**; module --force **purge**
- module **use** ~/modulefiles; module **unuse** ~/modulefiles

```
[someuser@bison]$ module list  
Currently Loaded Modules:  
1) GreXEnv (S)
```

Where:

S: Module is Sticky, requires --force to unload or purge



## ★ Grex environment [default]: GrexEnv

- no module loaded by default.
- use `module spider <name of the software>` to search for modules
- **Compilers:** {GCC, Intel}, MKL, PETSc, ... etc.
- Gaussian, ANSYS, MATLAB, ... etc.

## ★ The Alliance (Compute Canada) environment [optional]: CCEnv

- Switch to CCEnv; load a standard environment; choose the architecture[`sse3`, `avx2`, `avx512`], use `module spider <soft>`

`module load CCEnv`

`module load StdEnv/2020`

`module load arch/avx512`

`module load StdEnv/2020 gcc/9.3.0 geant4/10.7.3`

```
module load StdEnv/2016.4
```

```
module load arch/sse3
```

```
module load nixpkgs/16.09 gcc/5.4.0 geant4/10.05.p01
```



```
----- /opt/lmod/grex/7.6/modulefiles/Core -----
admixture/1.3.0          intel/14.0.2.144      ncl_ncarg/6.4.0      uofm/adf/2019.305-impi
admixture/1.23          (D) intel/15.0.5.223      ncl_ncarg/6.5.0      uofm/adf/2020-impi
ant/1.10.11             intel/2017.8          nodejs/4.4.7         uofm/adf/2020.103-impi
circos/0.69-6          intel/2019.5          (D) nodejs/8.12.0       uofm/adf/2021-impi
cmake/3.12.3           intel/2020.4          nodejs/13.2.0        (D) uofm/adf/2021.102-impi
cmake/3.14.0           j/j903                openbabel/2.3.2      uofm/adf/2021.106-impi (D)
cmake/3.16.9           jags/3.4.0            openbabel/2.4.1      uofm/cfx/15.0
cmake/3.23.2           (D) jags/4.0.0          openbabel/3.0.0      uofm/cfx/16.2
cns/1.3                 jags/4.3.0           (D) openbabel/3.1.1    (D) uofm/cfx/18.2
eigen/3.3.7            java/jdk7u25          ovito/2.9.0          uofm/cfx/19.2
fastqc/0.11.9          java/jdk7u45          ovito/3.0.0-dev502   (D) uofm/cfx/20.1
gaussian/g09.b01.unlim java/jdk8u5           pandoc/2.9.2.1       uofm/cfx/20.2
gaussian/g09.b01       java/jdk8u66          perl/5.14.4          (D) uofm/cfx/21.1
gaussian/g09.e01.unlim java/jdk8u92          (D) perl/5.22.1       (D) uofm/feko/2021.2
gaussian/g09.e01       java/jdk13.0.1        perl/5.28.1          uofm/gaussian/g03
gaussian/g16.b01       julia/1.3.0-bin       php/5.6.40           uofm/gaussian/g09.e01
gaussian/g16.c01.avx2.unlim julia/1.5.4-bin       php/7.3.12          (D) uofm/gaussian/g09 (D)
gaussian/g16.c01.avx2 julia/1.6.1-bin       python/2.7.12-miniconda uofm/mathematica/11.0
gaussian/g16.c01       (D) julia/1.7.0-bin    python/3.6-miniconda (D) uofm/matlab/R2014A
gcc/4.8                 libccerf/1.4          settarg              uofm/matlab/R2015B
gcc/5.2                 lmod                  singularity/3.5.2    uofm/matlab/R2017A
gcc/7.4                 (D) ls-prepost/4.7.13  smrtlink/6.0.0.47841 uofm/matlab/R2019B
gcc/9.2                 mcr/mcr               stata/14.2-fagfs     uofm/matlab/R2020B2 (D)
gcc/11.2                mkl/10.3.11          (D) stata/15.0-fagfs   (D) uofm/starccm/16.06.010
git-lfs/3.2.0           mkl/11.1.0           tbb/14               uofm/starccm/17.02.008-R8 (D)
git/2.21.0              mkl/2019.5           tbb/2019.5          (D) uofm/umcfid/2.4
gnuplot/5.2.7           molden/5.9            trimmomatic/0.39     vina/1.1.2
go/1.10.4               multiwfn/3.8-gui      uofm/adf/2016-impi-test vmd/1.9.3
go/1.11.5               multiwfn/3.8-nogui   (D) uofm/adf/2016-impi vncworkspace/1.1
go/1.12.12              nbo/6.0              uofm/adf/2017-impi  vtune/2019.4
go/1.13                 nbo/7.0              (D) uofm/adf/2017.114-impi vtune/2019.5 (D)
go/1.13.3               (D) ncl_ncarg/6.2.1   (D) uofm/adf/2018dev-impi wine/3.0
intel/12.1.5.339        ncl_ncarg/6.3.0      uofm/adf/2019-impi  xtb/6.5.0-bin
----- /opt/lmod/stacks -----
CCEnv ($) GrexEnv ($,L)
```

- module avail
- module spider python
- module spider java
- module load gcc ompi
- module avail
- module spider <soft>
- module spider <soft>/<ver>
- module show <soft>
- module purge

If not available:  
 → contact support  
[support@tech.alliancecan.ca](mailto:support@tech.alliancecan.ca)



# Modules on the Alliance clusters

- ★ StdEnv/2016.4
  - nixpkgs/16.09 intel/2016.4 gcc/5.4.0 openmpi/2.1.1
  - MKL, Boost, HDF5, NetCDF, PETSc, ...
- ★ StdEnv/2018.3
  - nixpkgs/16.09 intel/2018.3 gcc/7.3.0 openmpi/3.1.2
  - MKL, Boost, HDF5, NetCDF, PETSc, ...
- ★ StdEnv/2020
  - gentoo/2020 intel/2020.1.217 gcc/7.3.0 openmpi/4.0.3
  - FlexiBLAS, Boost, HDF5, ...
- ★ StdEnv/2023
  - In preparation by RSNT.



```
[~@yak ~]$ cvmfs_config probe  
Probing /cvmfs/cvmfs-config.computeCanada.ca... OK  
Probing /cvmfs/soft.computeCanada.ca... OK  
Probing /cvmfs/restricted.computeCanada.ca... OK
```

```
[~@yak ~]$ ls -l /cvmfs/  
cvmfs-config.computeCanada.ca  
restricted.computeCanada.ca  
soft.computeCanada.ca
```

```
[~@yak ~]$ module load CCEnv  
[~@yak ~]$ module load arch/avx512  
[~@yak ~]$ module load StdEnv/2020  
[~@yak ~]$ module spider geant4  
[~@yak ~]$ module spider geant4/11.1.0  
[~@yak ~]$ module load StdEnv/2020 gcc/9.3.0  
geant4/11.1.0
```



# cvmfs on cedar

```
[~@cedar1: ~]$ cvmfs_config probe
Probing /cvmfs/atlas-condb.cern.ch... OK
Probing /cvmfs/atlas-nightlies.cern.ch... OK
Probing /cvmfs/atlas.cern.ch... OK
Probing /cvmfs/grid.cern.ch... OK
Probing /cvmfs/icecube.opensciencegrid.org... OK
Probing /cvmfs/ref.galaxy... OK
Probing /cvmfs/ref.muggic... OK
Probing /cvmfs/restricted.computeCanada.ca... OK
Probing /cvmfs/sft.cern.ch... OK
Probing /cvmfs/snoplus.egi.eu... OK
Probing /cvmfs/soft.computeCanada.ca... OK
Probing /cvmfs/soft.galaxy... OK
Probing /cvmfs/soft.muggic... OK
Probing /cvmfs/sft-nightlies.cern.ch... OK
```

```
[~@cedar1: ~]$ ls /cvmfs/
atlas.cern.ch
cvmfs-config.computeCanada.ca
soft.computeCanada.ca
atlas-condb.cern.ch
soft-dev.computeCanada.ca
atlas-nightlies.cern.ch
data.rsnt.computeCanada.ca
oasis.opensciencegrid.org
ref.galaxy
soft.muggic
ref.muggic
sphenix.opensciencegrid.org
bio-test.data.computeCanada.ca
—
—
```



- ★ **Alternative for running software:** difficult to build from source
- ★ Possibility to **convert Docker** images to **singularity**.
- ★ **Singularity installed on all clusters** {no Docker for security reasons}
- ★ **Build the image:**
  - module load singularity
  - singularity build qiime2-2021.11.sif docker://quay.io/qiime2/core:2021.11
- ★ **Run the code via singularity:**

```
singularity exec -B $PWD:/home -B /global/scratch/someuser:/outputs \  
-B /global/scratch/someuser/path/to/inputs:/inputs <path to qiime2-2021.11.sif> \  
qiime feature-classifier fit-classifier-naive-bayes \  
--i-reference-reads /outputs/some_output_feature.qza \  
--i-reference-taxonomy /outputs/some_output_ref-taxonomy.qza \  
--o-classifier /outputs/some_output_classifier.qza
```

# Singularity/Apptainer

```
[~@yak ~]$ module load singularity
[~@yak ~]$ singularity build qiime2-2023.3.sif docker://quay.io/qiime2/core:2023.2
INFO: Starting build...
2023/05/15 14:46:02 info unpack layer:
sha256:3f4ca61aafcd4fc07267a105067db35c0f0ac630e1970f3cd0c7bf552780e985
....
INFO: Creating SIF file...
INFO: Build complete: qiime2-2023.3.sif

[~@yak ~]$ module load apptainer
[~@yak ~]$ apptainer build qiime2.sif docker://quay.io/qiime2/core:2023.2
```



**Resources:** [Github](#), [DockerHub](#), SingularityHub, Aptainer.

**Singularity examples:** <https://github.com/singularityware/singularity/tree/master/examples>

- ★ **Documentation:** <https://singularityware.github.io/user-guide.html>
- ★ **DockerHub:** <https://hub.docker.com/explore/>
- ★ **SingularityHub:** <https://www.singularity-hub.org/>
- ★ **Apptainer:** <https://apptainer.org/docs/>

<https://docs.alliancecan.ca/wiki/Singularity/en>  
<https://um-grex.github.io/grex-docs/>

**Access to Singularity:**

- ★ **Connect to cluster:** Grex, cedar, graham, beluga or narval:
- ★ **Load a module:** module load singularity
- ★ **Build the image:** convert the image from Docker to Singularity
- ★ **Note:** You may need to use your own Linux machine or VM to build the image

# Singularity: deprecated message

---

```
[~@beluga2 ~]$ module load singularity
```

-----  
There are messages associated with the following module(s):  
-----

**singularity/3.8:**

This module is deprecated and will be removed in the future due to security reasons.

**Please move to [apptainer/1.1](#)** or, if you use encrypted containers, [apptainer-suid/1.1](#).

Ce module est marqué pour être retiré et sera supprimé dans le futur pour des raisons de sécurité.

**Utilisez plutôt [apptainer/1.1](#)**, ou si vous utilisez des conteneurs cryptés, [apptainer-suid/1.1](#).

-----



- ★ **Tools and basic commands:** Operating system and nix or gentoo layer
- ★ Modules provided from local software stack or from cvmfs
- ★ Restricted software:
  - Restricted cvmfs: ORCA, CFOUR, DL\_POLY, Maker, ...
  - Installed under /opt/software {site license}: Gaussian, VASP, ...
- ★ Modules provided as a collection:
  - StdEnv/2016.4 {Available but not maintained}
  - StdEnv/2018.3 {Available but not maintained}
  - StdEnv/2020 {up to date and maintained}
  - **Preparing for the next environment:** in progress.
- ★ Local installation under user's account: R, Python, Perl, Julia, ...
- ★ Containers: singularity and apptainer

*Thank you for your attention*

*Any question?*