

# VASP Licensing and Compilation Guide for HPC Clusters


Last updated 1 May, 2025 • 2 min read

- [VASP licensing](#)
- [Compiling VASP on the cluster](#)
- [Creating a loadable VASP module on the cluster](#)
- [Using VASP on the cluster](#)
  - [Multithread job](#)
  - [GPU job](#)

## VASP licensing

Teams that would like to run VASP would need to purchase and manage their own VASP license.

VASP licensing is tied to each team that purchases a license and the license is not allowed to be shared among different teams, unless a team authorizes other users within their designated license.

 HPC does not manage the VASP licensing.

After the VASP license is purchased, a team will be granted/assigned the installation files for the version of VASP the license supports.

Teams can place the VASP installation files within their designated team shared directories on HPC.

More information for the licensing is available here: [FAQs](#) .

VASP license registration form is available here:

[Get a license](#)

## Compiling VASP on the cluster

Reading the comments in the VASP makefiles, we can compile VASP with GCC, the Intel compiler, the AMD AOCC compiler, or NVIDIA's NVHPC compiler. We will use the Intel compiler and VASP 6.3.2 in this example:

```
> # Create a directory for our VASP project.
1 mkdir /path/to/shared/directory/location/vasp
2 cd /path/to/shared/directory/location/vasp
3 mkdir vasp-6.3.2
4 cd /path/to/shared/directory/location/vasp/vasp-6.3.2
5 mkdir buildfiles
6 cd buildfiles
7
8
9 # Transfer the associated Linux VASP tarball for 6.3.2 to the
10 # Use any file transfer program or means to transfer the tarball
11 #FileZilla, WinSCP, a direct copy on HPC using cp or rsync, etc.
12
13 rsync -a --progress vasp.6.3.2.tgz /path/to/shared/directory/location/vasp/
14
15 # Unpack the source.
16 tar -xvf vasp.6.3.2.tgz
17
18
19 # Load the Intel compiler and hdf5/1.13.2-ics version dependencies
20 # The libraries can change depending on the compiler being used
21 # and if dependencies were installed locally under a team's shared
22 #
23 # List available Intel compiler versions:
24 module avail intelc
25 # Get rid of any other modules that might interfere with our
26
27 module purge
28 module load intel/oneapi/2022.3 hdf5/1.13.2-ics
29
30 # Compile the VASP program.
31 cd vasp.6.3.2
32
33 # Compile by copying over a pre-set makefile.include file from
34 cp arch/makefile.include.intel_omp .
35
36 # Edit the makefile.include file that got copied with needed
37 nano or vi makefile.include.intel_omp
38
39
```

Here is an example of the contents that can be used with an Intel compiler build of VASP, settings may change.

```
> # Default precompiler options
1 CPP_OPTIONS = -DHOST="LinuxIFC" \
2               -DMPI -DMPI_BLOCK=8000 -Duse_collective \
3               -DscalAPACK \
4               -DCACHE_SIZE=4000 \
5               -Davoidalloc \
6               -Dvasp6 \
7
```

```

8         -Duse_bse_te \
9         -Dtbdyn \
10        -Dflock_dblbuf
11
12 CPP      = fpp -f_com=no -free -w0 -march=core-avx2 $*$(FL
13
14 FC        = mpiifort
15 FCL       = mpiifort -qmkl=sequential
16
17 FREE      = -free -names lowercase
18
19 FFLAGS    = -assume byterecl -w
20
21 OFLAG     = -O2
22 OFLAG_IN  = $(OFLAG)
23 DEBUG     = -O0
24
25 OBJECTS   = fftmpi.o fftmpi_map.o fftw3d.o fft3dlib.o /gpf
26 OBJECTS_01 += fftw3d.o fftmpi.o fftmpi.o
27 OBJECTS_02 += fft3dlib.o
28
29 # For what used to be vasp.5.lib
30 CPP_LIB    = $(CPP)
31 FC_LIB     = $(FC)
32 CC_LIB     = icc
33 CFLAGS_LIB = -O
34 FFLAGS_LIB = -O1
35 FREE_LIB   = $(FREE)
36
37 OBJECTS_LIB = linpack_double.o
38
39 # For the parser library
40 CXX_PARS   = icpc
41 LLIBS      = -lstdc++
42
43 ##
44 ## Customize as of this point! Of course you may change the p
45 ## part of this file as well if you like, but it should rarel
46 ## necessary ...
47 ##
48
49 # When compiling on the target machine itself, change this to
50 # relevant target when cross-compiling for another architectu
51 VASP_TARGET_CPU ?= -march=core-avx2
52 FFLAGS      += $(VASP_TARGET_CPU)
53
54 # Intel MKL (FFTW, BLAS, LAPACK, and scaLAPACK)
55 # (Note: for Intel Parallel Studio's MKL use -mkl instead of
56 FCL        += -qmkl=sequential
57 MKLR00T    ?= /gpfs/sharedfs1/admin/hpc2.0/apps/intel/oneapi/
58 LLIBS      += -L$(MKLR00T)/lib/intel64 -lmkl_scalapack_lp64 -
59 INCS       =-I$(MKLR00T)/include/fftw
60
61 # HDF5-support (optional but strongly recommended)
62 CPP_OPTIONS+= -DVASP_HDF5
63 HDF5_ROOT  ?= /gpfs/sharedfs1/admin/hpc2.0/apps/hdf5/1.13.2-i
64 LLIBS      += -L$(HDF5_ROOT)/lib -lhdf5_fortran
65 INCS       += -I$(HDF5_ROOT)/include

```

```
66
67 # For the VASP-2-Wannier90 interface (optional)
68 #CPP_OPTIONS += -DASP2WANNIER90
69 #WANNIER90_ROOT ?= /path/to/your/wannier90/installation
70 #LLIBS += -L$(WANNIER90_ROOT)/lib -lwannier
```

After the `makefile.include` file has the build configuration set, save the file.

Submit an interactive SLURM job to a compute node to perform the build:

```
> srun -N 1 -n 126 --partition=general --pty bash
1
```

Wait for a node to assign to the job, once assigned, perform the build by running the **make** command.

Example make command to build with 12 cores and specify to build all of the VASP executables.

```
> make -j12 all
1
```

More information on the make command for the VASP build available on VASP's Wiki page here:

[Installing VASP.6.X.X - VASP Wiki](#)

If there are no errors, VASP should build successfully.

After building, if the `vasp_std`, `vasp_gam`, and `vasp_ncl` executables are not available, run the following make command:

```
> make install
1
```

## Creating a loadable VASP module on the cluster

Create a module file for VASP that so that we can conveniently load VASP and its dependencies. The name that you choose for your module file is important as that is what module uses to reference it. We will make our name different by adding the "-mine" suffix to help separate it from the system installed vasp.

```

>
1  mkdir -p /path/to/shared/directory/location/mod/vasp
2  cd /path/to/shared/directory/location/mod/vasp
3  vi 6.3.2
4
5
6
7
8

```

```

>  #%Module1.0
1  ## vasp modulefile
2  ##
3  proc ModulesHelp { } {
4
5      puts stderr "\tAdds vasp/6.3.2 to your environment"
6  }
7
8  module-whatis    "Adds vasp/6.3.2 to your environment"
9
10
11 # Throw an error if any of these modules are loaded.
12 module load pre-module
13
14 module load intel/oneapi/2022.3
15 module load zlib/1.2.12-ics
16 module load hdf5/1.13.2-ics
17
18 conflict vasp
19
20 setenv          MOD_APP          vasp
21 setenv          MOD_VER          6.3.2
22 set            prefix            /gpfs/sharedfs1/path/to/VASP/
23 prepend-path    PATH             $prefix/locationWhereVaspBinariesAre
24
25 module          load post-module
26
27
28
29

```

If you are interested, in learning about module files you can read `man modulefile`

Finally, make sure that `module` knows to look in your `~/mod` directory for your module files by setting the `MODULEPATH` environmental variable:

```

>
1  nano ~/.bashrc # Add the lines below.
2
3
4
5
6

```

```

>
1 1 # My modules
2 2 source /etc/profile.d/modules.sh
3 3 MODULEPATH=/path/to/shared/directory/location/mod:${MODULEF
4
5
6
7
8

```

Reload your ~/.bashrc file in your current shell:

```

> source ~/.bashrc
1 # Finally Now we can load and run our VASP module
2 module load vasp/6.3.2
3 which vasp
4 vasp -h
5

```

## Using VASP on the cluster

All jobs on the cluster must be submitted through the SLURM scheduler using `sbatch`. Please read the [SLURM Guide](#) for more details. The preferred way to run `VASP` jobs is by specifying your associated VASP INCAR file to the vasp executable that a team installs within their team's shared directory. Please note that if your job uses many cores or a lot of memory it is better to submit to reserve a whole compute node (126 cores)

### Multithread job

To submit a job that uses 126 computational threads on one node, create a submission script vaspMP.sh:

```

> #!/bin/bash
1 #SBATCH -N 1
2 #SBATCH -n 126
3 #SBATCH -p general
4
5 module load vaspmodule
6 export UCX_TLS=tcp,self,sysv,posix
7 source /gpfs/sharedfs1/admin/hpc2.0/apps/intel/oneapi/2022.3/
8
9 mpirun vasp_std <restofcommandhere>
10
11

```

Then submit the script by:

```

> sbatch vaspMP.sh
1

```

## GPU job

If you are running on a GPU node, feel free to reduce the cores and allocate resources as needed.



Do not forget to allocate a GPU card to the job submission by using the `#SBATCH --gres=gpu:X` SLURM submission header

You can replace the X above for a GPU job submission with the number of GPU cards.

Here is an example to request 1 GPU card and 62 cores to a GPU node.

```
> #!/bin/bash
1 #SBATCH -N 1
2 #SBATCH -n 62
3 #SBATCH -p general-gpu
4 #SBATCH --gres=gpu:1
5
```