

HPC Workflow on Shaheen

(VASP as an Example)

Zhiyong Zhu
KAUST Supercomputing Core Lab
2022/09/13

Outline

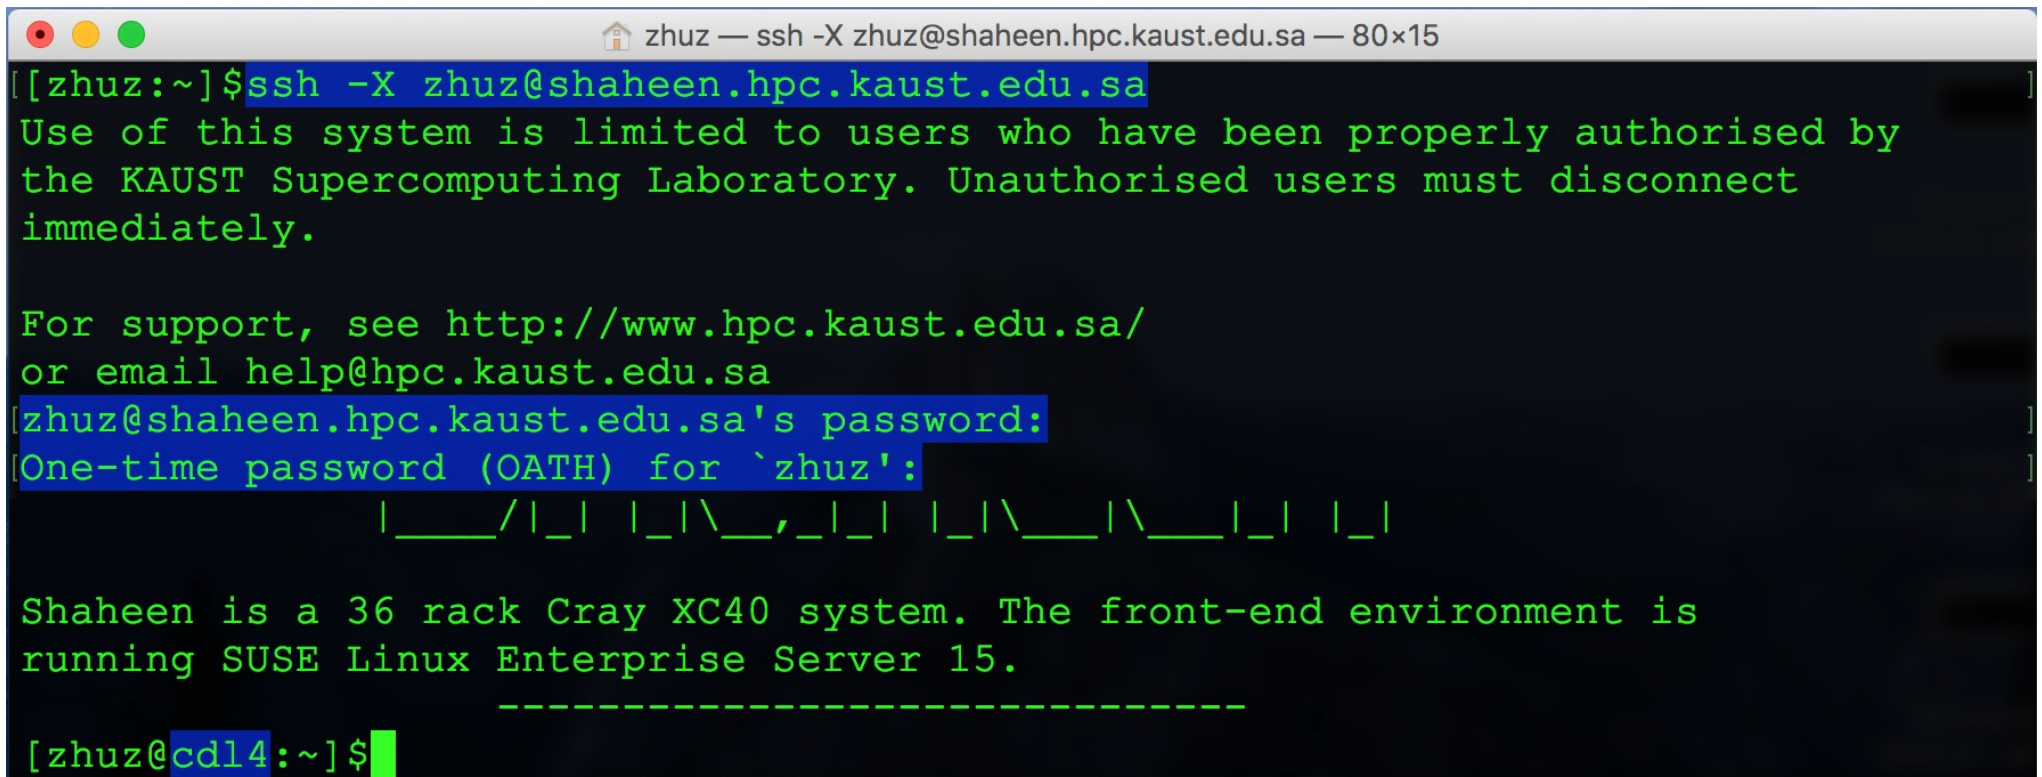
- Steps to Run
 - Use VASP as an example
 - The Vienna Ab-initio Simulation Package (VASP) is a computer program for atomic scale materials modeling, e.g. electronic structure calculations and quantum-mechanical molecular dynamics from first-principles (<https://www.vasp.at>).
- Tips

Steps to Run

- Login Shaheen
- Check Code Availability
- Working Directory
- Prepare Input Files for VASP
- Prepare Jobscript for Slurm Job Scheduler
- Job Submission using Slurm Commands
- Check Output Files

Login Shaheen

- Login
 - ssh -X [zhuz@shaheen.hpc.kaust.edu.sa](ssh://zhuz@shaheen.hpc.kaust.edu.sa)



```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz:~]$ ssh -X zhuz@shaheen.hpc.kaust.edu.sa
Use of this system is limited to users who have been properly authorised by
the KAUST Supercomputing Laboratory. Unauthorised users must disconnect
immediately.

For support, see http://www.hpc.kaust.edu.sa/
or email help@hpc.kaust.edu.sa
zhuz@shaheen.hpc.kaust.edu.sa's password:
One-time password (OATH) for `zhuz':
    |___/|_| |_\_,_| | |_\__|\__||_| | |
Shaheen is a 36 rack Cray XC40 system. The front-end environment is
running SUSE Linux Enterprise Server 15.
-----
[zhuz@cd14:~]$
```

Code Availability

- On Shaheen login node:
 - module avail
- In /sw/xc40cle7up03 software stack:
 - ls -l /sw/xc40cle7up03
- From our website:
 - www.hpc.kaust.edu.sa/app7up03

Code Availability

- module avail
 - module avail
 - module avail <code>/<version>

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:~]$ module avail
----- /sw/xc40cle7up03/modulefiles -----
abinit/8.10.3      git/2.35.0        p4vasp/0.3.30
abinit/9.6.2      gmp/6.1.2         pacchem/20200322
adf/2019.301      gnuplot/5.2.8    packmol/18.169
alamode/1.3.0     go/1.17.6        paraview/5.10.0
almabte/1.3.2    gollum2/2.0      periodic_nbo/20191008
amber/14          gpaw/19.8.1      perl/5.34.0
alps/6.6.6/-/1.0.3.1 2.25  gb91cd181.ar1 (default)
[[zhuz@cd14:~]$ module avail vasp/
----- /sw/xc40cle7up03/modulefiles -----
vasp/5.4.4 vasp/6.1.0
[[zhuz@cd14:~]$
```

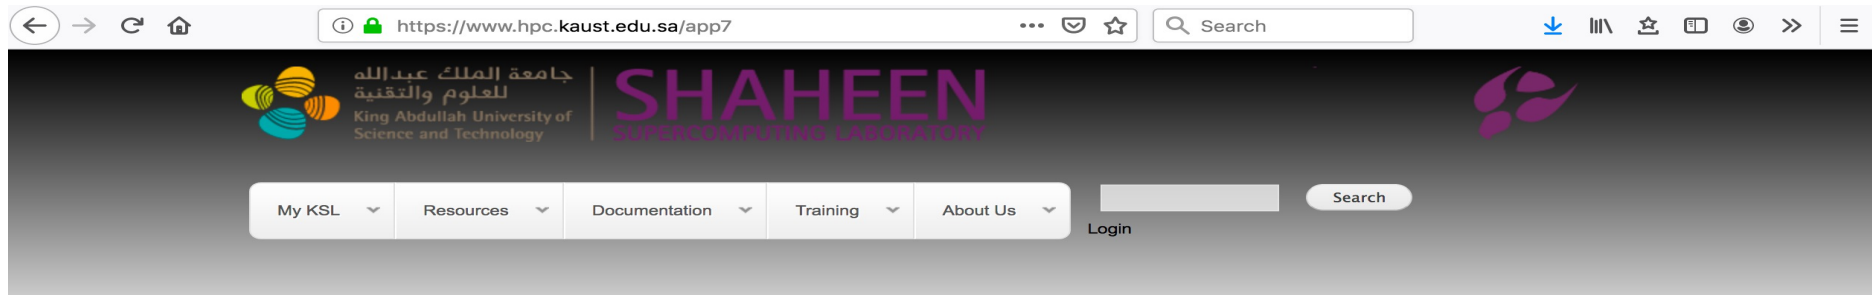
Code Availability

- `ls -l /sw/xc40cle7up03`

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cdl4:~]$ls /sw/xc40cle7up03/
abinit      critic2     jasper     octave     sod
adf         crystal14  java       octopus    sparskit2
adios       darshan    jdftx      oommf      spglib
admin-pe    dask       jmol       openbabel  swig
alamode     dftbplus   ktf        openfoam   zip
almabte     dftd4      lammps     openmpi    tbmodels
amber       dlpoly     lev00      openmx     tcl
ams         edmftf     libfabric  orca       tecplot
ansys       egsnrc     libint     ovito      tensorflow
antlr       eigen      libxc      p4vasp     thirdorder
arm-forge   elk        libxml2    pacchem    totalview
arm-reports elpa       lmod       package    towhee
ase         espresso   madagascar packmol     tramonto
atk         exciting   maestro    paraview   turbomole
```

Code Availability

- www.hpc.kaust.edu.sa
 - www.hpc.kaust.edu.sa/app7up03



[Home](#) » [Resources](#)

show all
Analysis
Computational Fluid Dynamics
Computational Materials Science
Debugging tools
Libraries
Machine learning
Mathematics
Molecular dynamics

Applications installed on Shaheen II / CLE 7

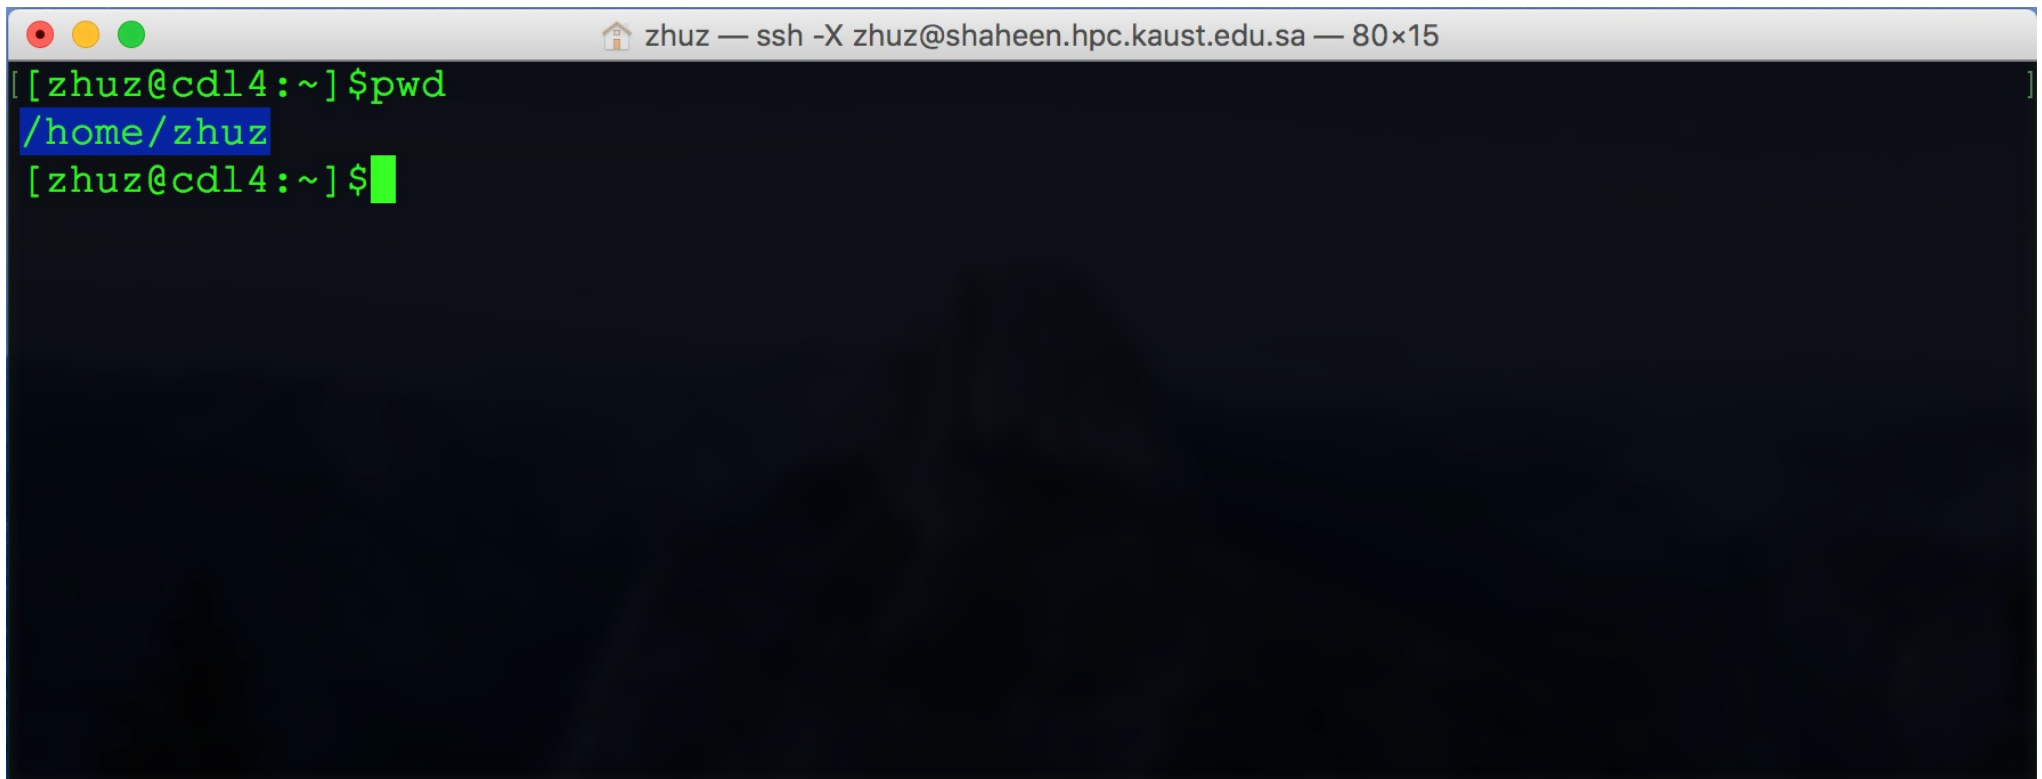
sort by name		sort by category	
f Analysis	ncl Analysis	ansys Computational Fluid Dynamics	converge Computational Fluid Dynamics
openfoam Computational Fluid Dynamics	abinit Computational Materials Science	alamode Computational Materials Science	almabte Computational Materials Science
berkeleygw Computational Materials Science	boltztrap Computational Materials Science	boltztrap2 Computational Materials Science	cif2cell Computational Materials Science
crystal14 Computational Materials Science	dmftw2k Computational Materials Science	edmftf Computational Materials Science	elk Computational Materials Science

Code Availability

- From our website:
 - www.hpc.kaust.edu.sa/app7
- On Shaheen login node:
 - module avail
- In /sw/xc40cle7 software stack:
 - ls -l /sw/xc40cle7
- Not found?
 - help@hpc.kaust.edu.sa

3 Different Working Directories

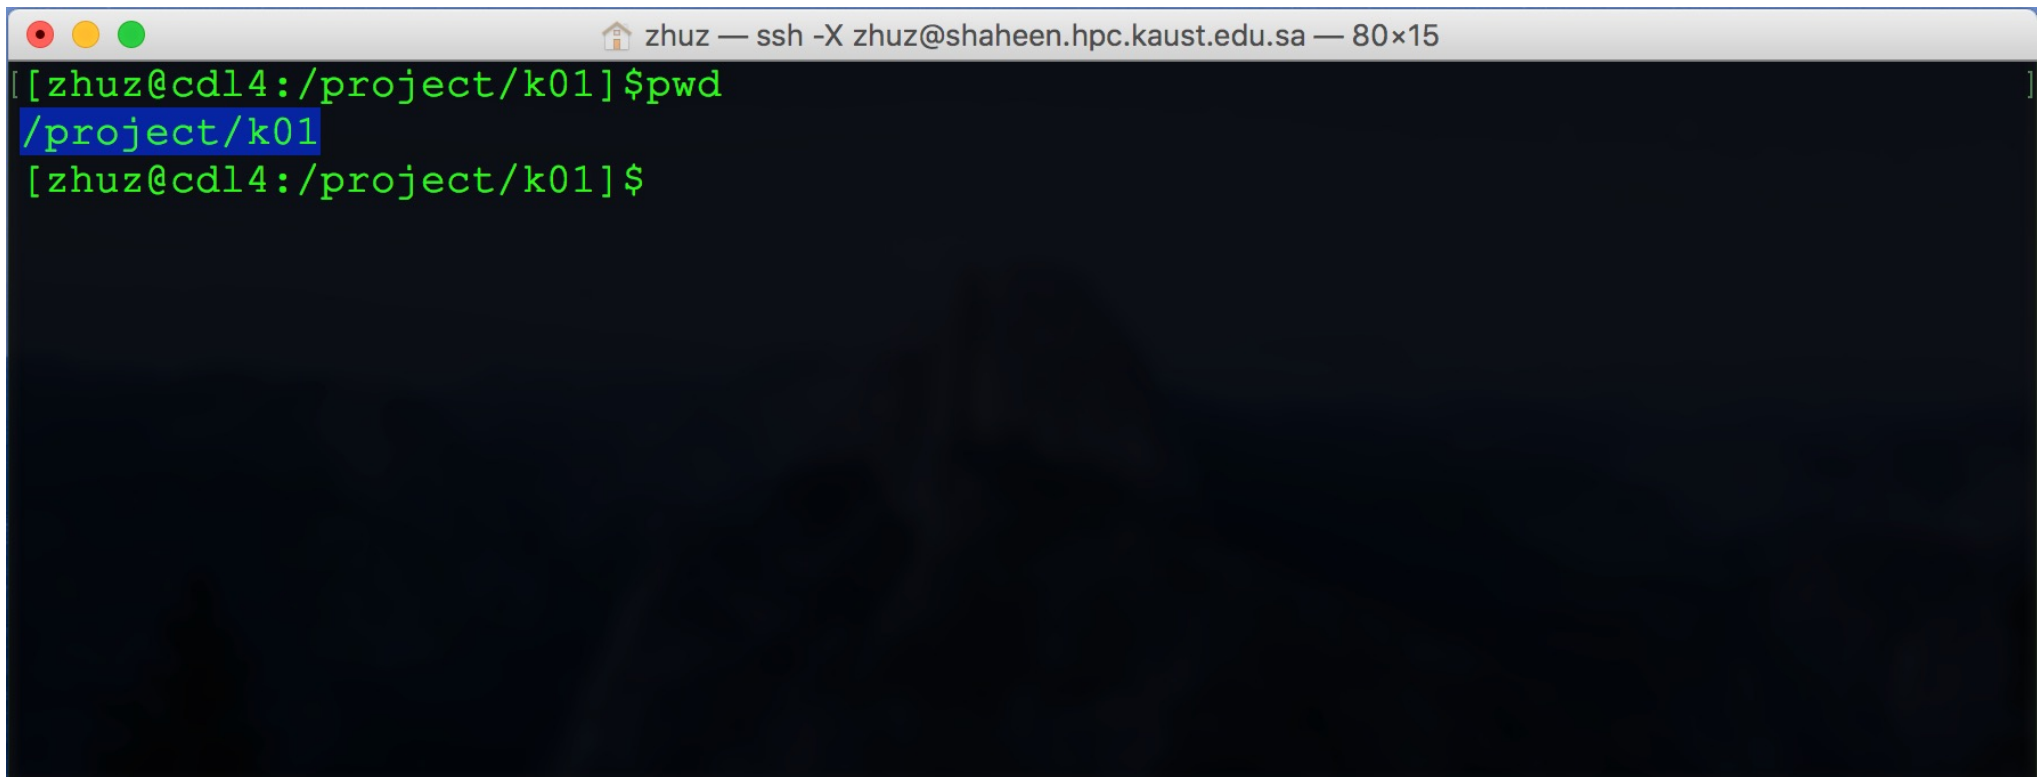
- /home
 - Very limited space; Not mounted on compute nodes (job submission will fail)



```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[zhuz@cd14:~]$ pwd
/home/zhuz
[zhuz@cd14:~]$
```

3 Different Working Directories

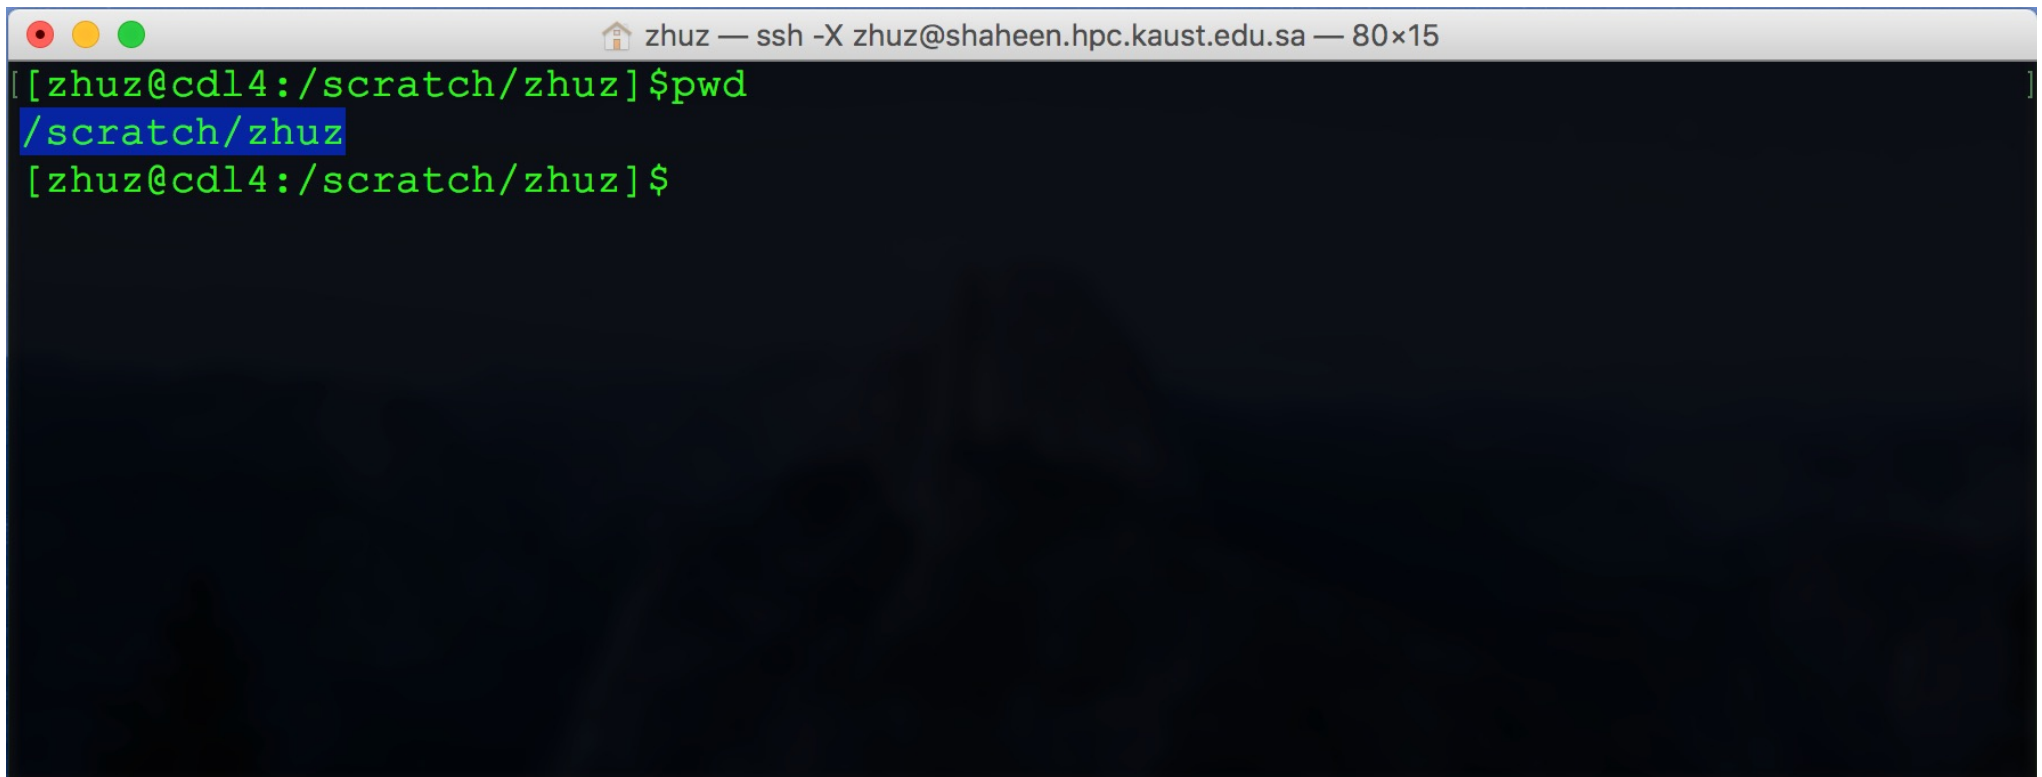
- /project/<projectname>
 - Read-only for compute nodes (job submission will fail); Used for data backup and data sharing



```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:/project/k01]$pwd
/project/k01
[zhuz@cd14:/project/k01]$
```

3 Different Working Directories

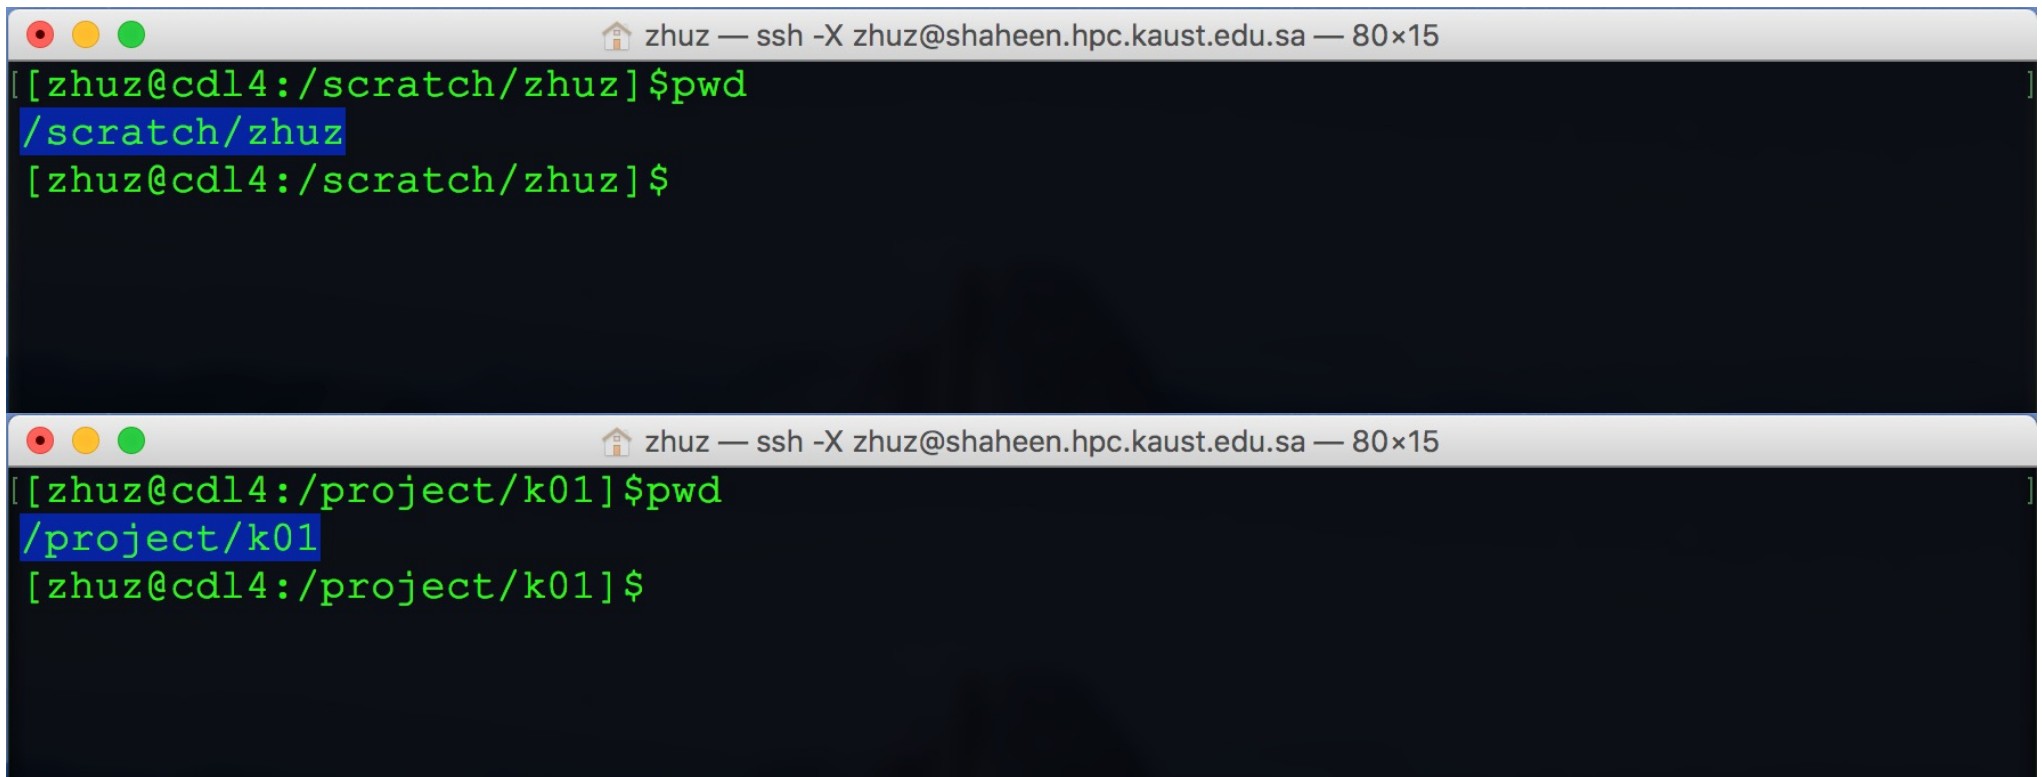
- /scratch/<username>
 - Almost unlimited space; Files older than 60 days are deleted automatically without warning



```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:/scratch/zhuz]$pwd
/scratch/zhuz
[zhuz@cd14:/scratch/zhuz]$
```

3 Different Working Directories

- Where to run? /scratch!
 - The only place to run
 - Remember to backup important data to /project



The image shows two terminal windows. The top window shows a user named 'zhuz' connected via SSH to 'shaheen.hpc.kaust.edu.sa'. The user is in the directory '/scratch/zhuz' and runs the 'pwd' command, which outputs '/scratch/zhuz'. The bottom window shows the same user in the directory '/project/k01' and runs the 'pwd' command, which outputs '/project/k01'. In both screenshots, the output of the 'pwd' command is highlighted in blue.

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:/scratch/zhuz]$pwd
/scratch/zhuz
[zhuz@cd14:/scratch/zhuz]$

zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:/project/k01]$pwd
/project/k01
[zhuz@cd14:/project/k01]$
```

Prepare Input Files

- Examples under Installation Folder
 - /sw/xc40cle7up03/code/ver/compilation/example
 - Inputs for application; Jobscript for Slurm

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:~]$cd /scratch/zhuz
[[zhuz@cd14:/scratch/zhuz]$mkdir vasp
[[zhuz@cd14:/scratch/zhuz]$cd vasp
[[zhuz@cd14:/scratch/zhuz/vasp]$cp /sw/xc40cle7up03/vasp/5.4.4/cle7_intel2021.4.0
/example/01/* ./
[[zhuz@cd14:/scratch/zhuz/vasp]$ls -l
total 612
-rw-r--r-- 1 zhuz g-zhuz      351 Feb 14 09:17 INCAR
-rw-r--r-- 1 zhuz g-zhuz       50 Feb 14 09:17 KPOINTS
-rw-r--r-- 1 zhuz g-zhuz   24378 Feb 14 09:17 POSCAR
-rw-r--r-- 1 zhuz g-zhuz  587541 Feb 14 09:17 POTCAR
-rw-r--r-- 1 zhuz g-zhuz    1088 Feb 14 09:17 z_jobs_shaheen
[[zhuz@cd14:/scratch/zhuz/vasp]$
```

Prepare Input Files

- VASP Input Files
 - Upload from local workstations (scp)
 - Modifying existing input files

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:/scratch/zhuz/vasp]$ls -l
total 612
-rw-r--r-- 1 zhuz g-zhuz    351 Feb 14 09:17 INCAR
-rw-r--r-- 1 zhuz g-zhuz     50 Feb 14 09:17 KPOINTS
-rw-r--r-- 1 zhuz g-zhuz  24378 Feb 14 09:17 POSCAR
-rw-r--r-- 1 zhuz g-zhuz 587541 Feb 14 09:17 POTCAR
-rw-r--r-- 1 zhuz g-zhuz   1088 Feb 14 09:17 z_jobs_shaheen
[[zhuz@cd14:/scratch/zhuz/vasp]$
```

Prepare Input Files

- Slurm Jobscrip
– SLURM directives

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:/scratch/zhuz/vasp]$more z_jobs_shaheen
#!/bin/bash
#SBATCH --partition=workq
#SBATCH --job-name="vasp"
#SBATCH --nodes=16
#SBATCH --time=4:00:00
#SBATCH --exclusive
#SBATCH --err=std.err
#SBATCH --output=std.out
#-----#
#export VASP_USE="" # This is the default, compiled with ScaLAPACK
#export VASP_USE=_omp # The version compiled with ScaLAPACK and threaded MKL
#export VASP_USE=_nbo # http://schmidt.chem.wisc.edu/nbosoftware
#export VASP_USE=_vaspsol # http://vaspsol.mse.ufl.edu/
#export VASP_USE=_vtst178 # http://theory.cm.utexas.edu/vtsttools/
```


Prepare Input Files

- Slurm Jobscrip
– Environments settings

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
#SBATCH --output=std.out
#-----#
#export VASP_USE="" # This is the default, compiled with ScaLAPACK
#export VASP_USE=_omp # The version compiled with ScaLAPACK and threaded MKL
#export VASP_USE=_nbo # http://schmidt.chem.wisc.edu/nbosoftware
#export VASP_USE=_vaspsol # http://vaspsol.mse.ufl.edu/
#export VASP_USE=_vtst178 # http://theory.cm.utexas.edu/vtsttools/
#export VASP_USE=_dftd4 # https://github.com/dftd4/dftd4_vasp
#export VASP_USE=_transopt # https://github.com/yangjio4849/TransOpt
#export VASP_USE=_z2pack # https://z2pack.greschd.ch/en/latest/tutorial/installa
tion.html
module load vasp/5.4.4
export OMP_NUM_THREADS=1
#-----#
echo "The job "${SLURM_JOB_ID}" is running on "${SLURM_JOB_NODELIST}
```

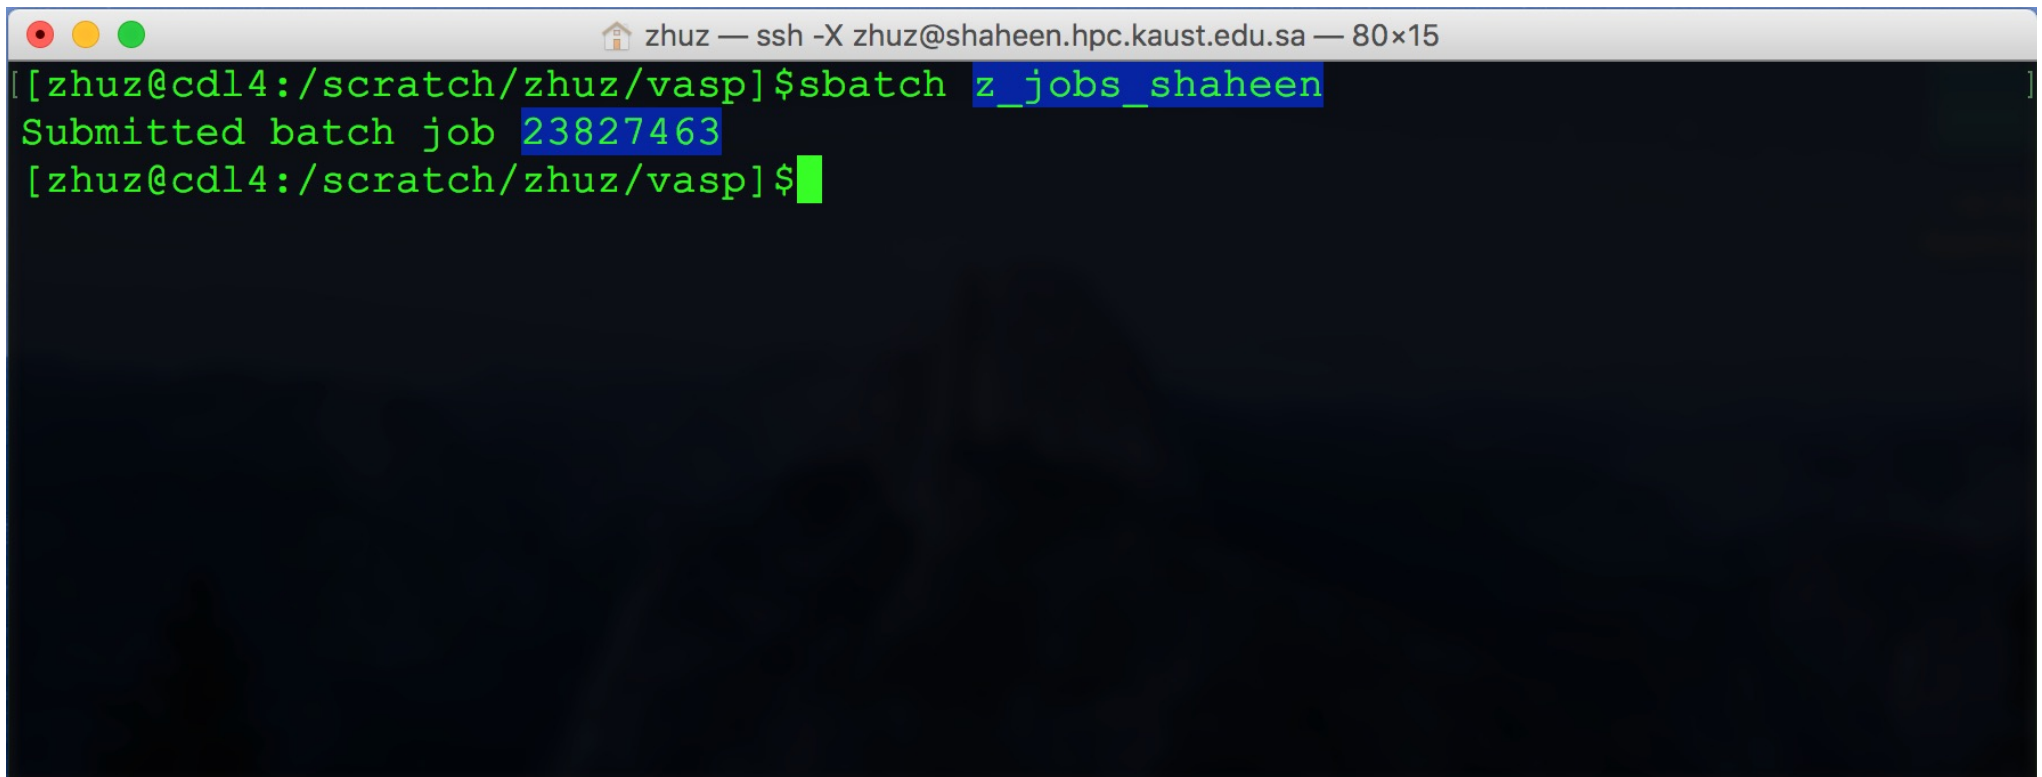
Prepare Input Files

- Slurm Jobscrip
– Commands to run

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
#export VASP_USE=_omp # The version compiled with ScaLAPACK and threaded MKL
#export VASP_USE=_nbo # http://schmidt.chem.wisc.edu/nbosoftware
#export VASP_USE=_vaspsol # http://vaspsol.mse.ufl.edu/
#export VASP_USE=_vtst178 # http://theory.cm.utexas.edu/vtsttools/
#export VASP_USE=_dftd4 # https://github.com/dftd4/dftd4_vasp
#export VASP_USE=_transopt # https://github.com/yangjio4849/TransOpt
#export VASP_USE=_z2pack # https://z2pack.greschd.ch/en/latest/tutorial/installa
tion.html
module load vasp/5.4.4
export OMP_NUM_THREADS=1
#-----#
echo "The job "${SLURM_JOB_ID}" is running on "${SLURM_JOB_NODELIST}"
#-----#
srun --ntasks=512 --hint=nomultithread ${VASP_HOME}/vasp_std
[zhuz@cd14:/scratch/zhuz/vasp]$
```

Job Submission

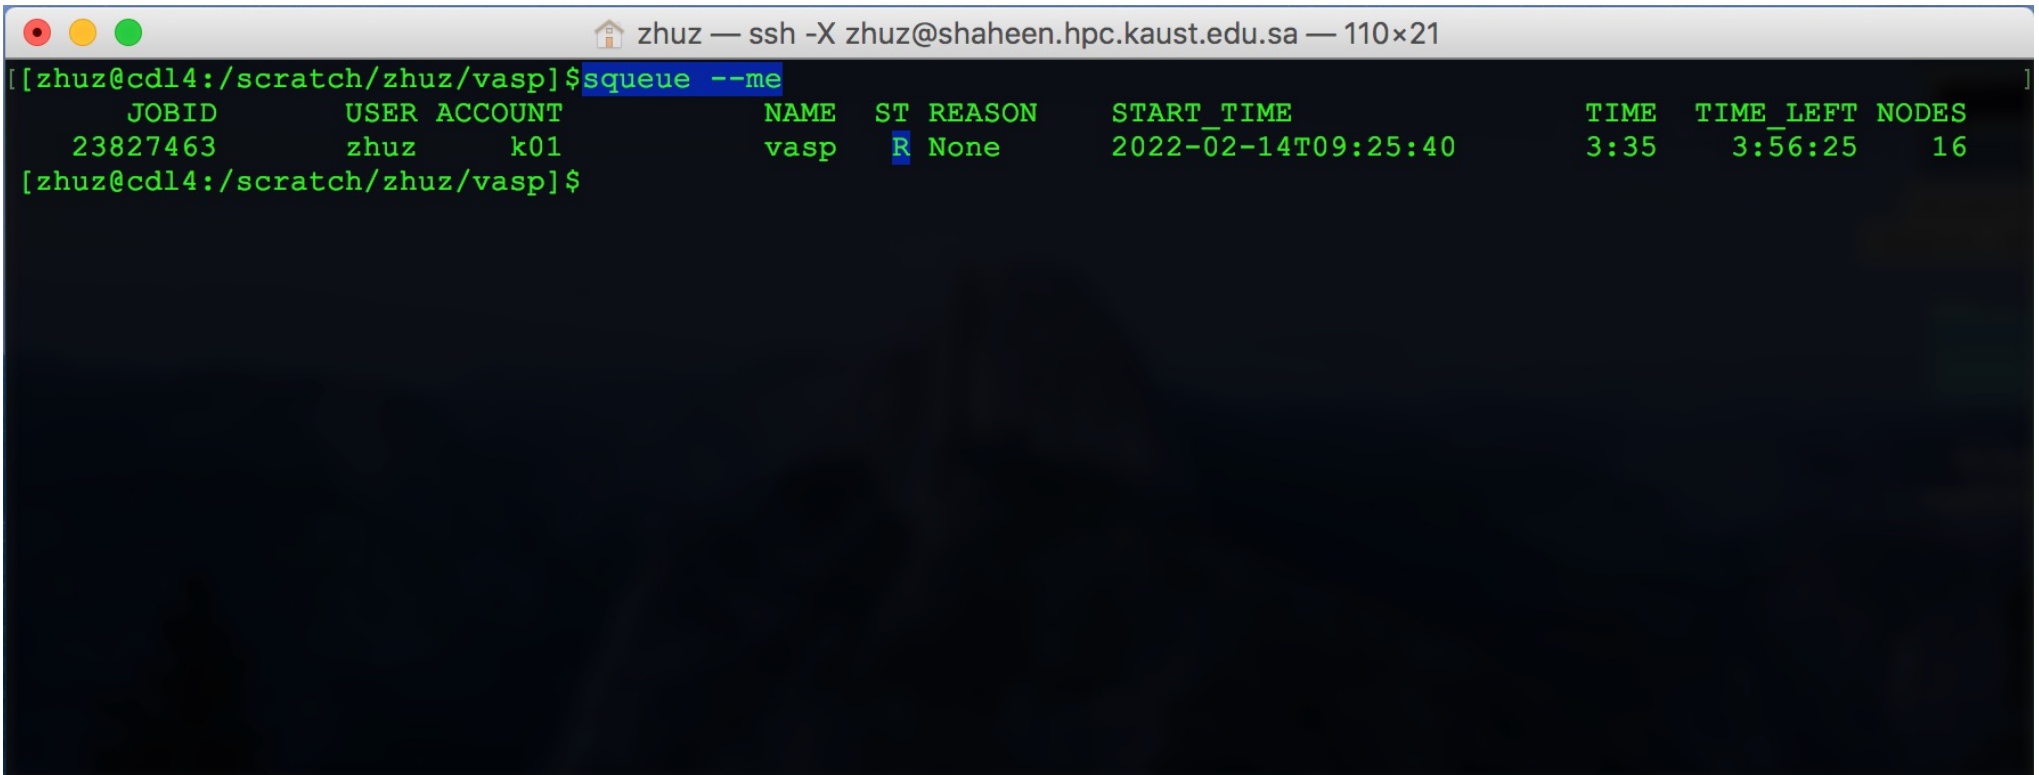
- sbatch
 - Submit jobs

A terminal window with a dark background and light green text. The window title bar shows 'zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15'. The terminal content shows the command 'sbatch z_jobs_shaheen' being executed, followed by the output 'Submitted batch job 23827463'. The prompt returns to '[zhuz@cd14:/scratch/zhuz/vasp]\$'.

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[zhuz@cd14:/scratch/zhuz/vasp]$sbatch z_jobs_shaheen
Submitted batch job 23827463
[zhuz@cd14:/scratch/zhuz/vasp]$
```

Job Submission

- `squeue`
 - Check job status

A terminal window with a dark background and light green text. The window title is "zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 110x21". The prompt is "[zhuz@cd14:/scratch/zhuz/vasp]". The command "squeue --me" has been entered and its output is displayed as a table with columns: JOBID, USER, ACCOUNT, NAME, ST, REASON, START_TIME, TIME, TIME_LEFT, and NODES. The output shows a single job with ID 23827463, user zhuz, account k01, name vasp, status R, reason None, start time 2022-02-14T09:25:40, time 3:35, time left 3:56:25, and 16 nodes.

```
[[zhuz@cd14:/scratch/zhuz/vasp]$ squeue --me
  JOBID      USER ACCOUNT      NAME  ST  REASON      START_TIME          TIME  TIME_LEFT  NODES
  23827463   zhuz   k01          vasp   R   None   2022-02-14T09:25:40  3:35   3:56:25   16
[zhuz@cd14:/scratch/zhuz/vasp]$
```

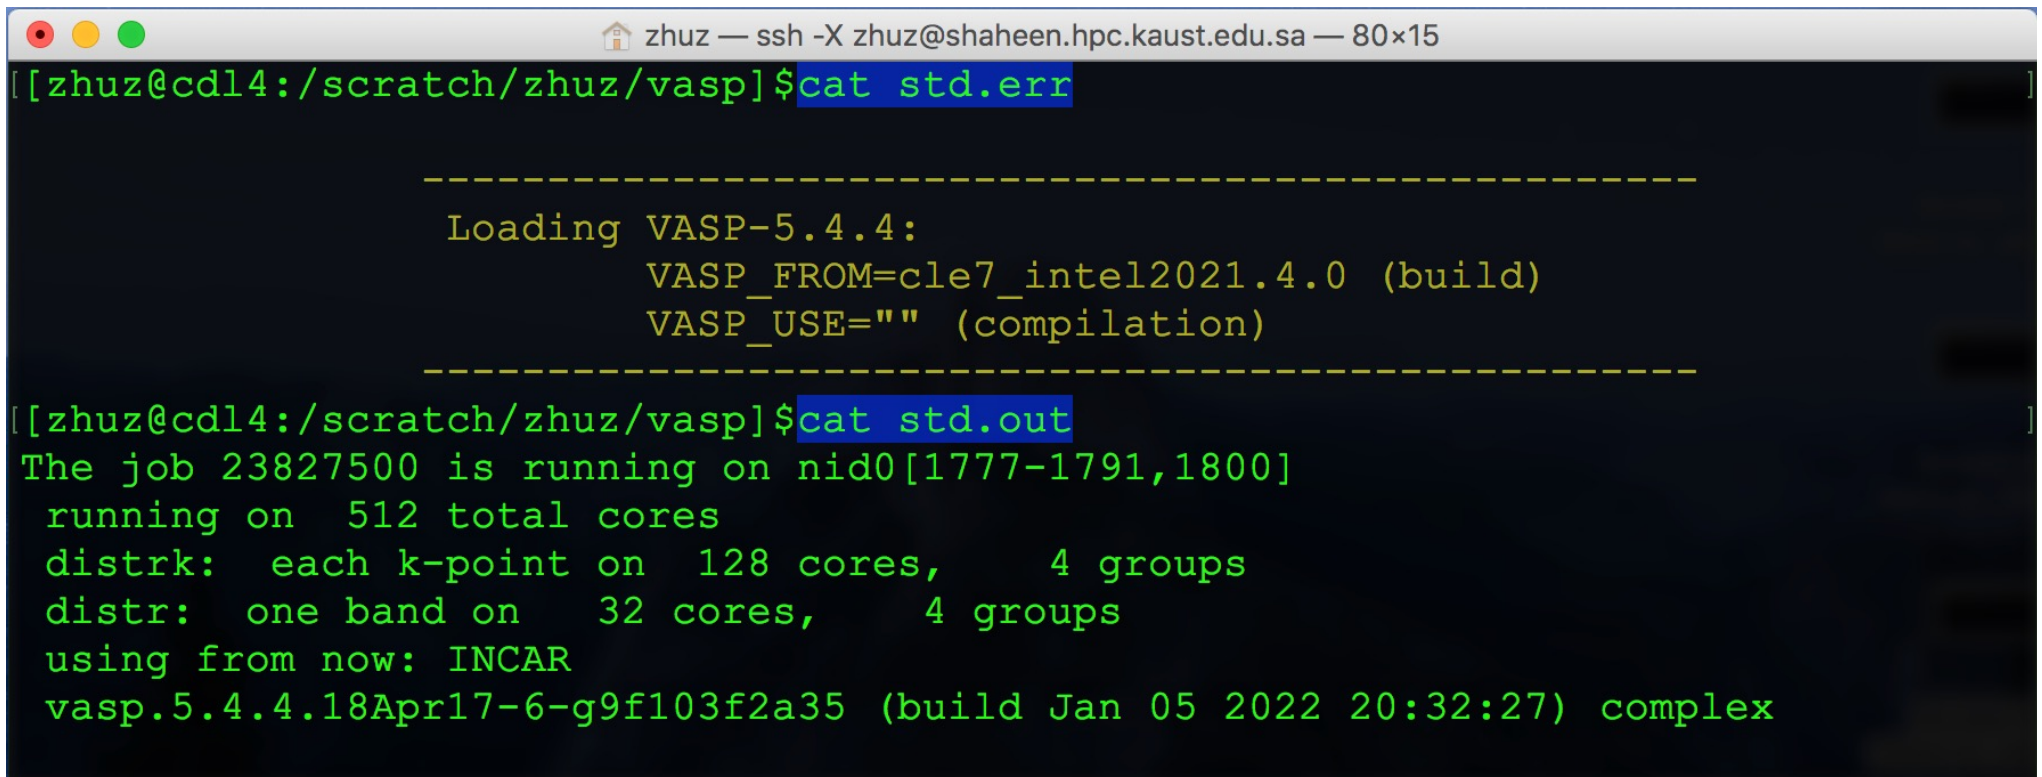
Job Submission

- scancel
 - Cancel jobs

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 110x21
[[zhuz@cd14:/scratch/zhuz/vasp]$squeue --me
  JOBID      USER ACCOUNT      NAME  ST REASON      START_TIME          TIME  TIME_LEFT  NODES
  23827463   zhuz    k01      vasp  R  None      2022-02-14T09:25:40  5:12   3:54:48   16
[[zhuz@cd14:/scratch/zhuz/vasp]$scancel 23827463
[[zhuz@cd14:/scratch/zhuz/vasp]$squeue --me
  JOBID      USER ACCOUNT      NAME  ST REASON      START_TIME          TIME  TIME_LEFT  NODES
[[zhuz@cd14:/scratch/zhuz/vasp]$
```

Check Output Files

- Errors or Not?
 - Standard outputs (std.out)
 - Standard errors (std.err)

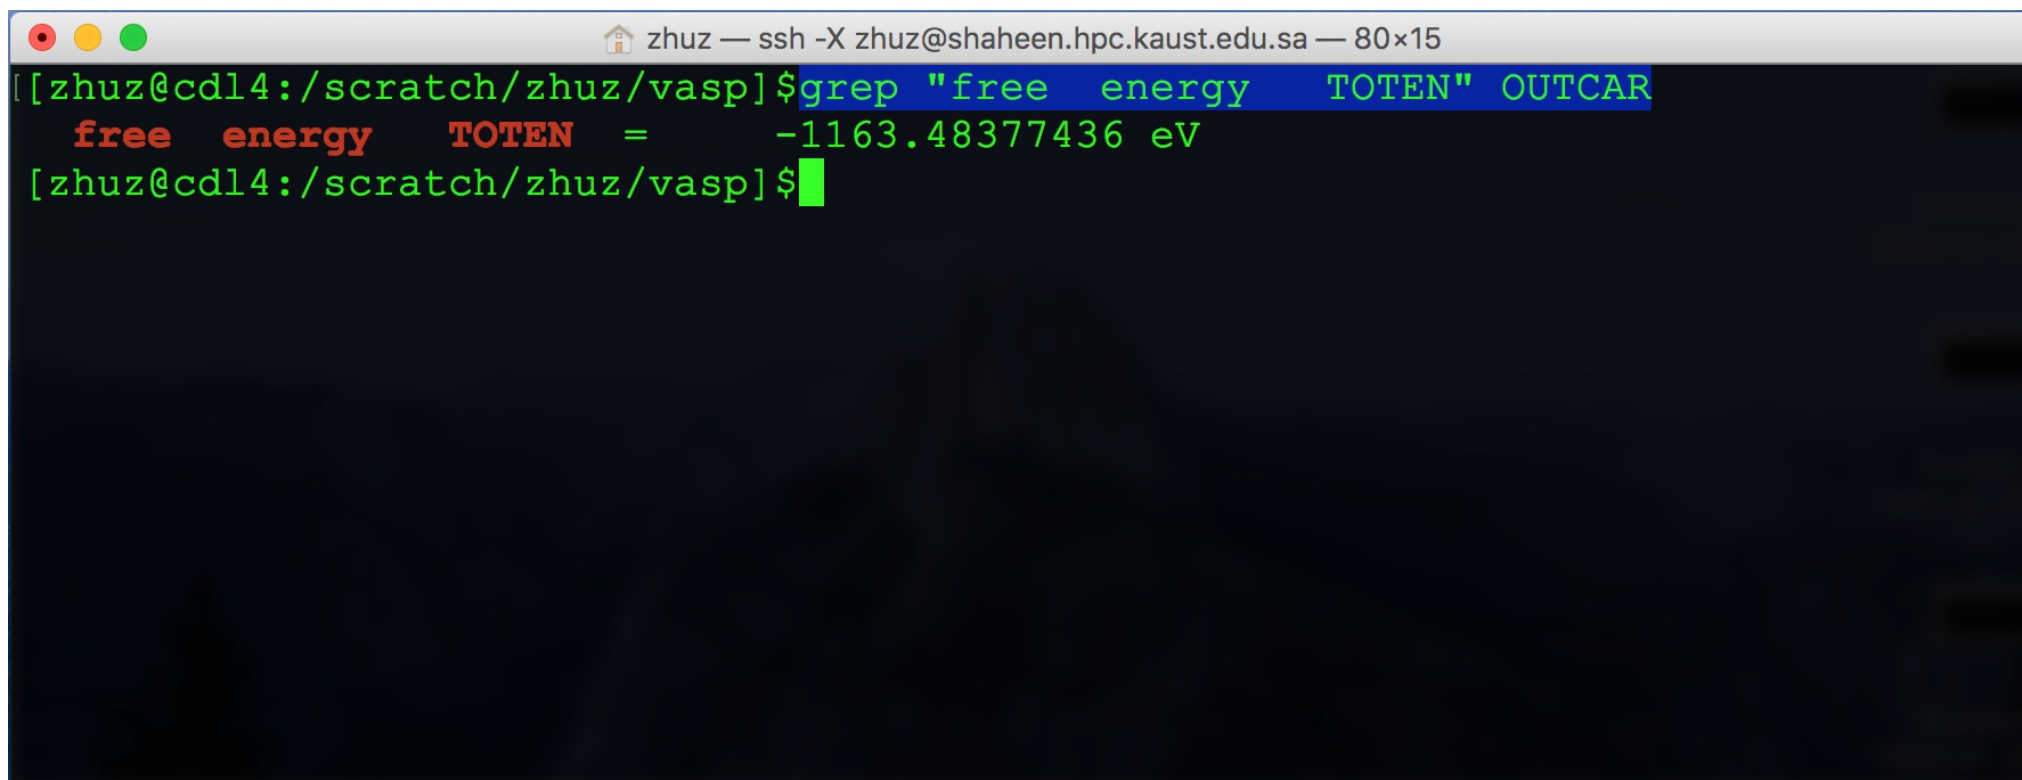


A terminal window screenshot showing the execution of VASP. The window title is "zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15". The prompt is "[zhuz@cd14:/scratch/zhuz/vasp]". The user enters "cat std.err", and the output shows VASP loading information: "Loading VASP-5.4.4: VASP_FROM=cle7_intel2021.4.0 (build) VASP_USE="" (compilation)". The user then enters "cat std.out", and the output shows job details: "The job 23827500 is running on nid0[1777-1791,1800] running on 512 total cores distrk: each k-point on 128 cores, 4 groups distr: one band on 32 cores, 4 groups using from now: INCAR vasp.5.4.4.18Apr17-6-g9f103f2a35 (build Jan 05 2022 20:32:27) complex".

```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:/scratch/zhuz/vasp]$ cat std.err
-----
Loading VASP-5.4.4:
      VASP_FROM=cle7_intel2021.4.0 (build)
      VASP_USE="" (compilation)
-----
[[zhuz@cd14:/scratch/zhuz/vasp]$ cat std.out
The job 23827500 is running on nid0[1777-1791,1800]
running on 512 total cores
distrk: each k-point on 128 cores, 4 groups
distr: one band on 32 cores, 4 groups
using from now: INCAR
vasp.5.4.4.18Apr17-6-g9f103f2a35 (build Jan 05 2022 20:32:27) complex
```

Check Output Files

- Results Analysis
 - Download to local workstations (scp)
 - Check it on-site



```
zhuz — ssh -X zhuz@shaheen.hpc.kaust.edu.sa — 80x15
[[zhuz@cd14:/scratch/zhuz/vasp]$grep "free energy TOTEN" OUTCAR
  free energy TOTEN = -1163.48377436 eV
[[zhuz@cd14:/scratch/zhuz/vasp]$
```

Tips

- Do not run directly on CDL login nodes
 - CDL login nodes are shared
- It won't work to submit jobs from /home and /project
 - /home is not seen on the compute nodes
 - /project is read-only on the compute nodes
- Backup important data from /scratch to /project (or /home, or your local computers)
 - Files in /scratch are not backed up, and are deleted automatically after 60 days
 - Do not confuse /scratch/project and /project

Tips

- Download and Upload (Mac & Linux)
 - scp -r LocalFiles [zhuz@shaheen.hpc.kaust.edu.sa:RemoteFolder](mailto:zhuz@shaheen.hpc.kaust.edu.sa)
 - scp -r [zhuz@shaheen.hpc.kaust.edu.sa:RemoteFiles](mailto:zhuz@shaheen.hpc.kaust.edu.sa) LocalFolder
- Download and Upload (Windows)
 - MobaXterm
 - <https://mobaxterm.mobatek.net>
 - Bitvise SSH Client
 - <https://www.bitvise.com/ssh-client-download>

Thank You!

help@hpc.kaust.edu.sa