

The background of the slide is a photograph of a large, modern building with a complex facade of glass and metal panels, likely the LRZ building. The image is overlaid with a semi-transparent blue filter. A dark blue horizontal bar is positioned across the middle of the image, containing the title and date text.

Introduction to Multiuser Cluster Systems at LRZ

April, 13th 2023

- **Aim:** provide an introduction to multiuser cluster systems in general and to those operated at the Leibniz Supercomputing Centre (LRZ), specifically
- You will probably benefit the most if you're not yet familiar with the LRZ HPC/HPDA/HPAI infrastructure, but plan to work with these systems in the future
- A majority of systems will be covered in more detail in dedicated sessions later this week

By the end of today's workshop, you should have a general understanding of multiuser HPC/HPDA/HPAI cluster systems and the basic skills to successfully interact remotely with such systems at LRZ

User perspective

User Perspective: System Access



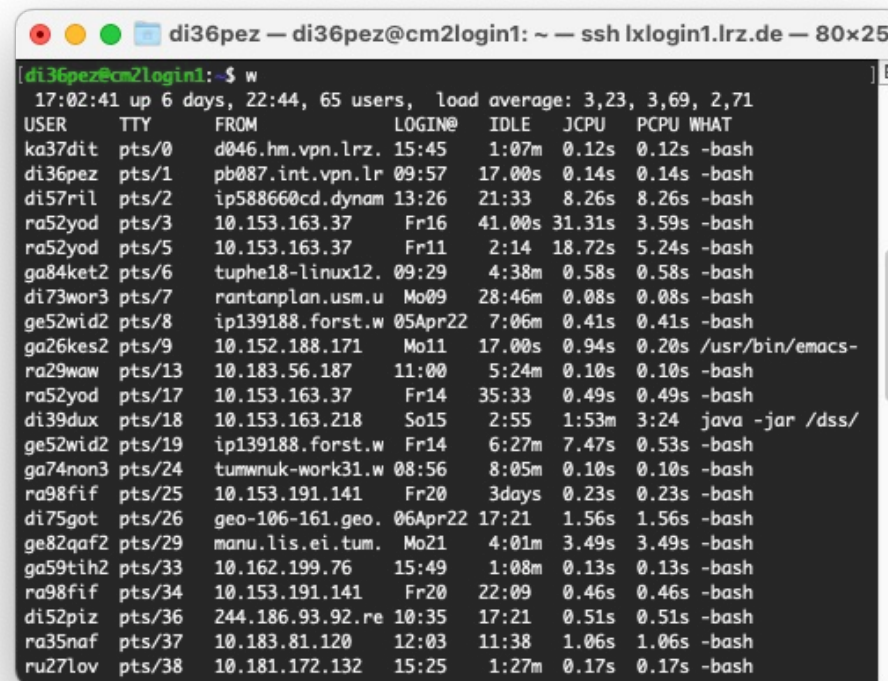
```
di36pez — di36pez@cm2login1: ~ — ssh lxlogin1.lrz.de — 80x25
[BADWLRZ-CMS9145:~ di36pez$ ssh lxlogin1.lrz.de
Last login: Tue Apr 12 09:57:48 2022 from pb087.int.vpn.lrz.de
Welcome to the CoolMUC2 Infiniband cluster, one of the Linux cluster systems
operated by Leibniz Supercomputing Centre (LRZ).

Please do not run any extensive computational programs on login nodes.
Instead, please submit SLURM batch scripts for production jobs, and SLURM
interactive shells for testing and short-running programs.
Misuse of the interactive resources will lead to violating accounts being
blocked from access to the cluster.
!!! Please note in particular this pertains to specifying invalid !!!
!!! mail addresses in SLURM scripts, please read          !!!
https://doku.lrz.de/display/PUBLIC/Available+SLURM+clusters+and+features
-----
Documentation: https://doku.lrz.de/display/PUBLIC/Linux+Cluster
Messages/System Status:
                    https://doku.lrz.de/display/PUBLIC/High+Performance+Computing
-----
Mar 18, 2022: Please note the announcement for a scheduled maintenance
of all cluster systems at https://www.lrz.de/aktuell/ali00936.html
-----
spack/release/22.2.1 22.2.1 release linux-sles15
intel-mpi: using intel wrappers for mpicc, mpif77, etc

di36pez@cm2login1:~$
```


User Perspective: Environment & Workspace


- These are systems shared by many users, i.e. other people will be working on the same (login) node at the same time.
- Be aware of your surroundings and considerate of your fellow colleagues!



```
di36pez@cm2login1: ~ — ssh lxlogin1.lrz.de — 80x25
di36pez@cm2login1:~$ w
17:02:41 up 6 days, 22:44, 65 users,  load average: 3,23, 3,69, 2,71
USER      TTY      FROM             LOGIN@   IDLE   JCPU   PCPU WHAT
ka37dit   pts/0    d046.hm.vpn.lrz. 15:45    1:07m  0.12s  0.12s -bash
di36pez   pts/1    pb087.int.vpn.lr 09:57    17.00s 0.14s  0.14s -bash
di57ril   pts/2    ip588660cd.dynam 13:26    21:33  8.26s  8.26s -bash
ra52yod   pts/3    10.153.163.37    Fr16    41.00s 31.31s 3.59s -bash
ra52yod   pts/5    10.153.163.37    Fr11     2:14  18.72s 5.24s -bash
ga84ket2  pts/6    tuphe18-linux12. 09:29     4:38m 0.58s  0.58s -bash
di73wor3  pts/7    rantanplan.usm.u Mo09    28:46m 0.08s  0.08s -bash
ge52wid2  pts/8    ip139188.forst.w 05Apr22 7:06m  0.41s  0.41s -bash
ga26kes2  pts/9    10.152.188.171   Mo11    17.00s 0.94s  0.20s /usr/bin/emacs-
ra29waw   pts/13   10.183.56.187    11:00     5:24m 0.10s  0.10s -bash
ra52yod   pts/17   10.153.163.37    Fr14    35:33  0.49s  0.49s -bash
di39dux   pts/18   10.153.163.218   So15     2:55  1:53m 3:24  java -jar /dss/
ge52wid2  pts/19   ip139188.forst.w Fr14     6:27m 7.47s  0.53s -bash
ga74non3  pts/24   tumwnuk-work31.w 08:56     8:05m 0.10s  0.10s -bash
ra98fif   pts/25   10.153.191.141   Fr20     3days 0.23s  0.23s -bash
di75got   pts/26   geo-106-161.geo. 06Apr22 17:21   1.56s  1.56s -bash
ge82qaf2  pts/29   manu.lis.ei.tum. Mo21     4:01m 3.49s  3.49s -bash
ga59tih2  pts/33   10.162.199.76    15:49     1:08m 0.13s  0.13s -bash
ra98fif   pts/34   10.153.191.141   Fr20    22:09   0.46s  0.46s -bash
di52piz   pts/36   244.186.93.92.re 10:35    17:21   0.51s  0.51s -bash
ra35naf   pts/37   10.183.81.120    12:03    11:38  1.06s  1.06s -bash
ru27lov   pts/38   10.181.172.132   15:25     1:27m 0.17s  0.17s -bash
```

User Perspective: Environment & Workspace

- You don't have administrative rights on these systems, i.e. no root access.
- This may be in contrast to your local machine and certain usage patterns and/or expectations may therefore not apply, e.g.,
 - you will not be able to use the `sudo` command
 - you're prohibited from making system-wide modifications
 - disk access is restricted to your home directory (and possibly other storage areas accessible to your account, e.g., your DSS containers)
- That said, your home (directory) is your castle – there, anything goes!



- If available on the system, modules allow for the dynamic modification of environment variables, e.g., they provide a flexible way to access various applications and libraries available on the system
- List the currently active modules (loaded by default):
`$ module list`
- Search for available modules:
`$ module available <module>` or
`$ module av <module>`
- Get more information about a specific module:
`$ module show <module>`
- Use `$ module load <module>` to apply the changes of a module to the environment


- Short of manually compiling all the binaries you need (of course, this is an option, too), you can resort to using package managers to set up the software/applications you need
- This will have to take place in your home directory, obviously
- **Spack** (<https://spack.io/>) is “a package manager for supercomputers, Linux, and macOS. It makes installing scientific software easy”. It is “supporting multiple versions, configurations, platforms, and compilers”.
It is used to set up (most of) the software available on the Linux Cluster and SuperMUC-NG. This setup can be extended by users.



- **Conda** (<https://conda.io>) is “a package, dependency and environment management for any language – Python, R, Ruby, Lua, Scala, Java, JavaScript, C/ C++, FORTRAN, and more”.
- **pip** (<https://pip.pypa.io>) is “the package installer for Python. You can use it to install packages from the Python Package Index and other indexes”.

Make sure to install packages to the home directory instead of the system-wide default location:

```
$ pip install --user <package>
```



User Perspective: OS-level Virtualization, Containers

- Isolated **user space** instances, called containers, allow programs running inside to only see the container's contents and devices assigned to the container.
- Thus, the environment inside a container can essentially be modified freely, typically **providing (encapsulated) root privileges**
- The most prominent container runtime, Docker, is typically not available on multiuser systems, but you will encounter alternatives
 - Charliecloud (<https://hpc.github.io/charliecloud/>)
 - Enroot (<https://github.com/NVIDIA/enroot>)
- Containers imposes no noticeable overhead, i.e. there should be no performance impact and parallelization, GPU access, etc. should – if set up correctly – work as expected
- Containers are UDSS: User Defined Software Stacks: you're basically independent from the environment created by system administrators, but you will only receive limited support for the environment created instead (inside the container).

