

Introduction to Levante

Please note all questions here. Put your name in front of the question such that we might contact you later, if we didn't answer directly at the meeting.

section 1: General System Overview

section 2: First Steps

- (Joan Saladich) Why a user's public key for remote login to Levante must be updated every 6 weeks? Could it be every +6 months instead?
 - (Hendryk Bockelmann, DKRZ) This is for security reasons. After a major incident at many data centres last year, in which, among other things, unprotected ssh keys were stolen, some sites introduced 2FA. DKRZ has decided against this approach (as many users find it too inconvenient) and is initially trying to minimise the risk by asking the users to frequently change their ssh keys.
- (Hernan Campos) How do I setup an SSH key on levante. Is the system the same as on mistral (web interface)?
 - (Daniel Heydebreck) It is the same procedure as on mistral.
 - (Alexander Voigt) see also: <https://gitlab.dkrz.de/help/ssh/index.md#generate-an-ssh-key-pair> (<https://gitlab.dkrz.de/help/ssh/index.md#generate-an-ssh-key-pair>)

Editors

- (Katrin Nissen) Is there a editor which can be used on Levante (apart from vi)?
 - (Carsten Lemmen) I am using VSCode with Remote SSH extension
 - (Rene Redler) module load emacs; emacs &
 - (Flo) Spyder (see <https://docs.dkrz.de/blog/2022/spyder.html> (<https://docs.dkrz.de/blog/2022/spyder.html>))
 - (Alexander Voigt) You could mount Levantes file system with sshfs and use your favorite (locally installed) editor
- (E.S.) Is GVIM (a very light weight X11 version of vim) available? If not, can it be made available? It would make cp/paste/clipboard easier than in pure terminal-based vim.
 - (Alexander Voigt) yes
 - (Flo) see <https://docs.dkrz.de/blog/2022/libGL.html> (<https://docs.dkrz.de/blog/2022/libGL.html>) if you get an libGL.so.1 error.
- (C.S.) Would it be possible to get joe installed?
 - (Alexander Voigt) joe is on our list of requested software
- (FatemeH Chegini) Would it be possible to also have meld?
 - (Alexander Voigt) meld is on our list of requested software. Meld has so far proven difficult because of multiple GUI toolkit dependencies.
- (Katrin Nissen) Will R be available on Levante (Jupyter Notebook on Levante)?
 - (Patrick Jöckel) "module load r/4.1.2-gcc-11.2.0" (is that what you need?)

- (Quan Liu) How to use R in jupyterhub with levante?
 - this is good to be installed later and documented here: <https://docs.dkrz.de/doc/software%26services/jupyterhub/kernels.html> (<https://docs.dkrz.de/doc/software%26services/jupyterhub/kernels.html>)

section 3: Usage Model

- (Thomas Kleinen): Currently, quite a few libraries need to be installed by users of MPIESM – will those be installed at system level by some point?
 - (Thomas Kleinen): Check Enrico's notes on how to get MPIESM running (<https://palmod.gitlab-pages.dkrz.de/modelperformance/recommendedsetups/levante>) (<https://palmod.gitlab-pages.dkrz.de/modelperformance/recommendedsetups/levante>) – mainly different versions of netcdf / hdf and similar. It DOES make support much more difficult, if all users compile their own libraries. (hdf5/1.8.22-openmpi-4.1.2-intel-2021.5.0, netcdf-c/4.4.0-openmpi-4.1.2-intel-2021.5.0, netcdf-fortran/4.4.3-openmpi-4.1.2-intel-2021.5.0, libaec/1.0.0-intel-2021.5.0)
 - (Thomas Jahns) I don't see any reason to not use the pre-installed netcdf etc. For python2 I must say that's not planned to be supported because there is no longer any upstream maintenance.
 - (Kalle Wieners) When using an up-to-date mkexp Version, python2 should no longer be needed
 - (Thomas Jahns) All in all the write-up appears to be from a time when almost no software was installed yet. I suggest to re-visit this in light of currently available software tree.
 - (Thomas Kleinen) Thanks very much for that assessment – we'll give it another try, then.
- (Thomas Kleinen): Will ksh appear at some point, or should we invest the time to convert all scripting to bash?
 - ksh is already there and you should invest in porting to bash perhaps, ksh is no longer under development and it is unclear if it will eventually be dropped by Linux distributors (Thomas Jahns)
- (Lukas Pilz) What is the difference between the two similar packages for each library (e.g. for netcdf-c%intel@2021.5.0: 2k3cmu3 and 7dq6g2h)?
 - (Dominik Zobel): In this case the first library was built with OpenMPI support, the second with IntelMPI.
 - (Irina Fast): try: `spack find -dp /2k` and `spack find -dp /7dq`
- (C. Danek) So when my project has zero node hours left, I cannot run a single cdo cmd on a large file (i.e. as it is always possible on mistralpp)?
 - Traditionally that just means you spend more time in the queue waiting for your allocation.
 - If you only have projects that have no compute resources allocated (i.e. no resources to begin with), please contact us, and we can probably find a solution.
- (C. Danek) Can you please explain the main difference between spack manager vs module? E.g. when I search some sw, what is the first place to look for it?
 - (Lukas Pilz) I would look with `spack find -pv1c` because you get much more info on the software and how it was built

- (S.Hagemann) In the talk it was mentioned that one should use the option `-march=core-avx2` when the intel compiler is used. Does this also applies for all model compilations where the intel openmpi compiler is used, e.g. `mpifort -march=core-avx2 -further_flags ...?`
 - yes, `mpifort` is just a wrapper to have the proper MPI environment already in the compiler/linker command; you still have to specify for which architecture you want to produce code

section 4: AOB

- (Joan Saladich) Will Levante has the ".json" sub-catalogs for CMIPs, Cordex and ERA5 with the same JSON' structure & pathname as in Mistral (e.g. from Mistral `"/pool/data/Catalogs/mistral-cmip6-cmip.json"` to `"/pool/data/Catalogs/levante-cmip6-cmip.json"`)?
 - (Stephan Kindermann, DKRZ) All intake catalogs will also be available for Levante (with adapted unified catalog names) from the `/pool/data/Catalogs` entry point, for further info see the tutorials site at <https://data-infrastructure-services.gitlab-pages.dkrz.de/tutorials-and-use-cases/tutorials.html> (<https://data-infrastructure-services.gitlab-pages.dkrz.de/tutorials-and-use-cases/tutorials.html>)
- (Joan Saladich) Do you plan to add new variables for ERA5 dataset in Levante's catalog (e.g. wind components at 100m.)?
 - (Angelika Heil, DKRZ) We add new ERA5 variables if there is a general need across the DKRZ-user community. Under certain conditions, we can also fulfil individual requests. This depends on the additional resource requirements in terms of additional memory space and manpower.

Random questions

- (Anon) I wonder how you prevent users from accidentally blocking the main node with expensive jobs (Just asking for a friend... ;)
 - (DKRZ) we will be using `cgroups` to limit the resources

Python

- (Johannes Bieser) Does loading python modules affect the intel compiler in any way?
 - (Thomas Jahns) Some parts of the Intel compiler tools use python, but the pure compiler tool chain should be unaffected. Otherwise since conda has its own development environment, it's recommended to not load the `python3` module when building software vs. the DKRZ software.
 - (Daniel Heydebreck; copied from other comments) The order in which modules are loaded might be important (e.g. python bringt own openmpi and pfortran compiler). Therefore, the `python3` module might overload the paths of content of the intel compiler module.
- Anything on python2 python3 incompatibilities? (I figure there will by no py2 on levante)
 - yes, no python2
- (Carsten Lemmen) We had a problem with CMake FindNetCDF not recognizing the compiler-versioned `netcdf` but the one that came with Mamba/python (workaround with loading python first does not work)
- (Lars Buntmeyer): Do you recommend to use `mamba` ? i get some permission warnings, but it seems to work...
 - (Friedemann Reum): Works flawlessly here so far.

- (Lars Buntmeyer): Do you consider `nb_conda_kernels` in the base environment? It's tedious to install conda kernels manually for the Jupyterhub...
 - (Thomas Jahns) Is there already an ongoing inquiry to support@dkrz.de (<mailto:support@dkrz.de>) or beratung@dkrz.de (<mailto:beratung@dkrz.de>)?
 - No, not from my side!
- Using JupyterHub with your own kernels:
 - <https://docs.dkrz.de/doc/software%26services/jupyterhub/kernels.html#use-your-own-kernel> (<https://docs.dkrz.de/doc/software%26services/jupyterhub/kernels.html#use-your-own-kernel>)

Git

- (Thomas Kleinen): What's the recommended way to run git? (private ssh keys often required, security concern) My understanding had been that we should NOT, under any circumstances, store private ssh keys in mistral home directories due to the security issues mentioned above. Was I wrong in this understanding? Can we keep them there, as long as they are password-protected?
 - (Thomas Jahns) Exactly: only unencrypted is problematic. The encrypted ones do not pose similar risks. We force key rotation, because it's impossible to tell from our end if the key was ever stored in unencrypted form.
 - It is usually sensible to use a different ssh-key for different application/machines. E.g. one for git on levante, one for accessing levante from your laptop...

WRF

- (Sergey Osipov): I know that WRF was tested on Levante and you did a good progress tuning it. Could you please publish that on the docs for the future reference.
 - (Thomas Jahns) We don't actually have an expert on WRF. The porting was done by one colleague, who does not have experience running the model.
 - (Sergey Osipov) Still, Irina shared some great compilations scripts. Many users will find it very useful. You will get less questions later :)
 - (Thomas Jahns) If there is sufficient interest we will assist users in organizing a break-out group for WRF at the user workshop. Consider attending the 'Schneesturm' meeting on 12th of May 2022.
 - (Sergey Osipov) That is very helpful, thank you! How do I join?
 - (Thomas Jahns) Details can be found at <https://luv.dkrz.de/news/1722/> (<https://luv.dkrz.de/news/1722/>) (German only, unfortunately)
 - (Friedemann Reum) Some guidance on optimally compiling WRF would be great! I got it running, but don't know whether it could be sped up...

Interactive python

- (Sergey Osipov): Not everyone runs Jupyterhub. Pycharm, for example, offers great handling of the remote python interpreters via ssh. But I have to manually allocate the interactive node and to setup ssh tunnel to connect to the remote python interpreter. Would be great to streamline the setup procedure. I have some scripts as a show case, which you could probably improve a lot. @email: Sergey.Osipov@mpic.de (<mailto:Sergey.Osipov@mpic.de>)
 - I guess this is not a widely used application
 - Definitely, I agree. Just a suggestion.
 - VS Code offers free remote development over ssh...
 - Free is not an issue. The issue is the manual setup required every time. The compute nodes are behind the firewall. And connecting via ssh twice (login, then interactive) is usually not supported, hence the tunneling.

- (Thomas Jahns) I can provide some information on ssh jump hosts that should make this simpler.

SSH

- (E.Sch., erik DOT research attt gmx DOT us) Can you provide a complete list of server-side supported SSH options? Meant for configuring my local .ssh/config file when ssh to levante. For example: supported ssh versions (I used to have an issue with a DWD HPC login to certain other HPC due to DWD only offering an too old ssh version...), X11 timeouts, connect timeout, controlmaster, serverAlive..., ssh -X vs ssh -Y issues, e.g. at hpc Jülich there are issues like ncview and gvim not working at all or crashing with ssh -X. Yet, I don't like doing ssh -Y for security reasons...
 - (Thomas Jahns) ssh version requirements are subject to change: we remove compromised ciphers from time to time
 - (Thomas Jahns) The ssh X11 forwarding does not usually see problems at our end but rather on client system or network transit. We ask to submit problems to support@dkrz.de (<mailto:support@dkrz.de>) or beratung@dkrz.de (<mailto:beratung@dkrz.de>) with suitable amount of details.