REGIONALES RECHENZENTRUM ERLANGEN [RRZE]

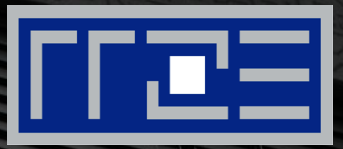

HPC from the home office

Tips & tricks on working efficiently from remote

HPC Services, RRZE

General Notes

- Universities are connected via DFN: Read DFN newsticker! https://www.dfn.de/newsticker-covid19/
- Connection to FAU may be very slow (esp. Deutsche Telekom, see above)
- Switching between IPv4 and IPv6 can help in some rare cases (e.g. **`ssh -4 /ssh -6, ping -4 / -6`**)

SSH remote connections

- Linux & Mac: **`ssh`**, **`scp`**, **`sftp`**, and **`rsync`** are standard commands
- Windows
  - **`putty`**, **`winscp`**, https://mobaxterm.mobatek.net/
  - What about the Linux subsystem for Windows10?

- Only cshpc can be reached by IPv4 and IPv6 from the Internet
- Meggie frontends can be reached by IPv6
- All other HPC hosts only from within the University (i.e. VPN or some gateway host)

Ways to avoid VPN:
SSH connection as SOCKS v5 proxy

- Use SSH connection SOCKS5 proxy on your localhost:1080 :
  - **`ssh hpcUser@cshpc.rrze.fau.de -D1080`**
  - Firefox/Thunderbird: go to edit/preferences (about:preferences) scroll down to: Network Settings: manual proxy configuration: SOCKS host: localhost or 127.0.0.1 Port: 1080

Ways to avoid VPN:
SSH connection to tunnel to your workstation

- `ssh hpcUser@cshpc.rrze.fau.de -L5555:your.workstation.uni-erlangen.de:22`
- `ssh workstationUser@127.0.0.1:5555`
- rsync via ssh tunnel:

  `rsync -e "ssh -p 5555" wsUser@127.0.0.1: ....`

(re)connecting to the command line

- `screen` – installed on cshpc and all fronend nodes
- `tmux` – installed on cshpc and all fronend nodes

- `mosh` – similar to SSH, with additional features meant to improve usability for mobile users (roaming and intelligent local echo); installed on cshpc

REGIONALES RECHENZENTRUM ERLANGEN [RRZE]

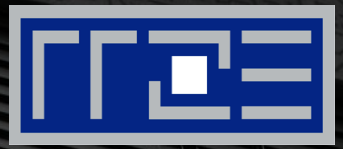

Managing remote sessions with tmux

HPC Services, RRZE

Overview

- What is “tmux”?
  - “tmux” is a text-based window manager that can handle multiple shell sessions from a single session
  - Sessions can be given names for easy access
  - User can detach from the screen session anytime and re-attach later
  - Typical usage model:

$ ssh screenhost

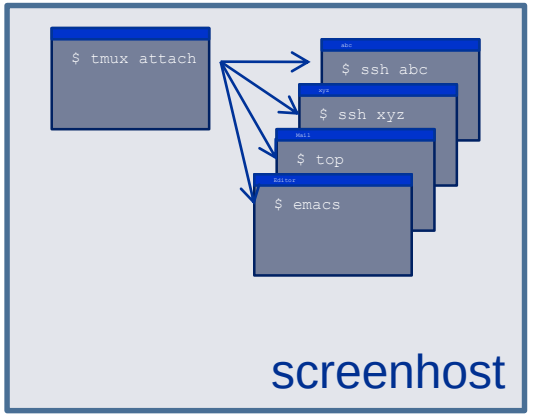

Starting, attaching, and detaching a session

- Starting a new tmux session from a normal shell
  `$ tmux [new [-s <session_name>]]`
- Attaching to a running but unattached screen session
  `$ tmux attach-session [-t <ID>]`
- Listing my existing sessions on this host
  `$ tmux ls`
- Detach from session and go back to shell
  `<cmd> d` (default **`<cmd>: ctrl-b`**)

Guess what...:
`$ tmux kill-session [-t <ID>]`

Dealing with windows

A “window” is a “screen” with a shell running (think “tab”)

create new	**<cmd> c**
select from list	**<cmd> w**
go to next/previous	**<cmd> n/p**
find by name	**<cmd> f**
rename	**<cmd> ,**

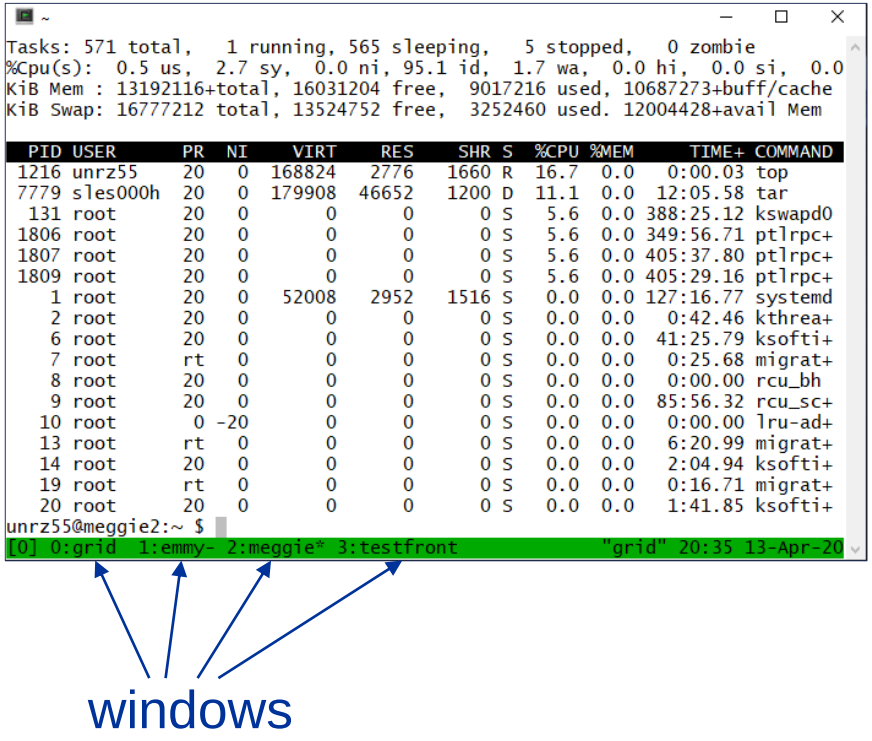

Dealing with panes

A “pane” is a sub-window (think “split screen”)

vertical split	`<cmd> %`
horizontal split	`<cmd> "`
move to next	`<cmd> o`
tog. pane zoom	`<cmd> z`
display clock	`<cmd> t`

resize pane `<cmd> <ctrl>-{←,→,↓,↑}`

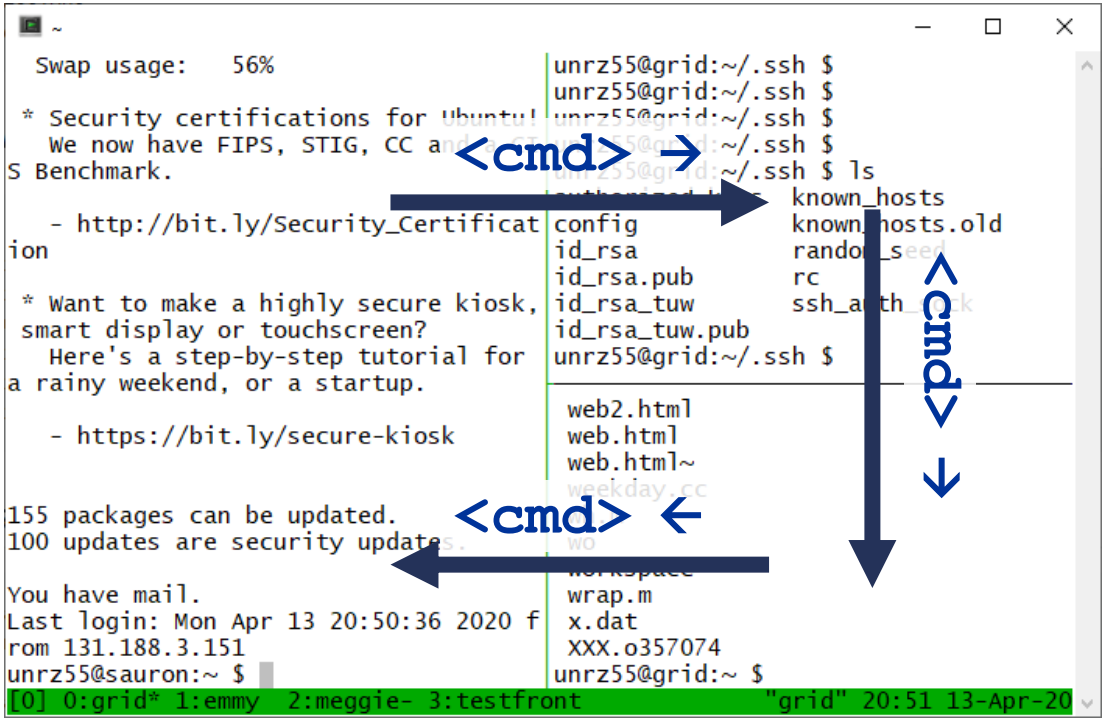

Copy mode

Move around in a pane, view scrolled-out data, copy stuff, and paste

- First, **`set-option -g mode-keys emacs`** in **`~/.tmux.conf`**

- Some important keys

enter copy mode	**`<cmd> ]`**
start selection	**`C-space`**
copy to tmux buffer	**`M-w`**
leave copy mode	**`ESC`**
paste tmux buffer	**`<cmd> [`**

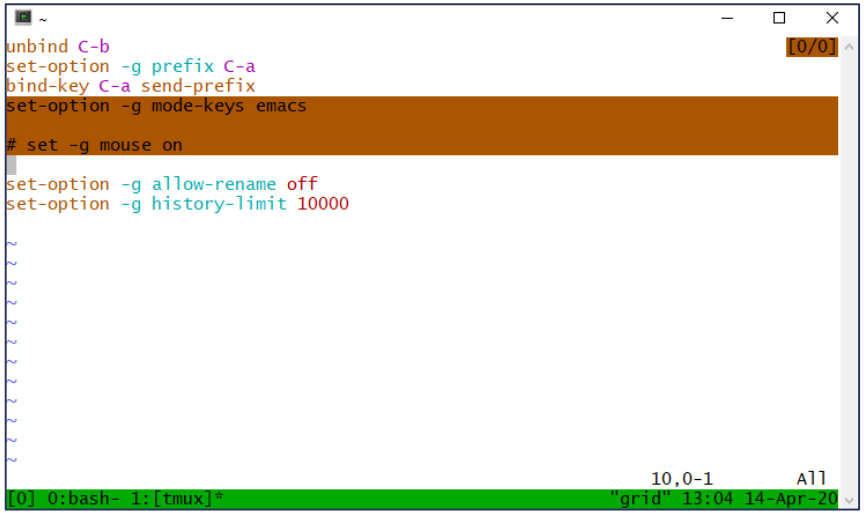

The config file

- `~/.tmux.conf` is read at server restart or via an explicit command:
  `<cmd> : source-file .tmux.conf`

- Example `~/.tmux.conf`:

```
# Make C-a the cmd key
unbind C-b
set-option -g prefix C-a
# Send C-a by hitting C-a twice
bind-key C-a send-prefix
# use emacs controls for copy mode
set-option -g mode-keys emacs
# Mouse control (I don't like that)
# set -g mouse on
# Disallow window renaming
set-option -g allow-rename off
# 10k lines in pane buffer
set-option -g history-limit 10000
```

REGIONALES RECHENZENTRUM ERLANGEN [RRZE]

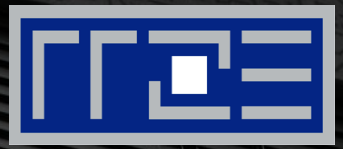

Working with remote X11

HPC Services, RRZE

Working with remote X11 (I)

- ssh -X usually is a bad idea on high-latency links
- MobaXterm can provide an X server for Windows user

Working with remote X11 (II)

- Preferred solution: use terminal server connection with Nomachine NX to cshpc
  https://www.anleitungen.rrze.fau.de/hpc/dialogserver/
- You'll get a standard Linux desktop on cshpc.
- However, keep in mind that cshpc is shared by many users and can quickly become a bottle neck.
- Clients are available for Windows, Mac, Linux (x86), and Raspbery Pi.

Working with remote X11 (III)

- x2go.org can be an alternative to Nomachine on cshpc if you have a Linux PC at the University (installation: root access mandatory)
- x2go has Windows clients, but does **not** support OpenGL. Thus, even some Linux desktops may not work smoothly.
- inbuilt proxy mechanism (use proxy server for SSH connection)
- session type "single application" (Terminal=Xterm) can be used to open just a single terminal w/o any desktop environment. (~ screen with X)

Working with remote X11 (IV)

- x2goclient can mount local directory on remote host
  (Session preferences: Shared folders)
- try different connection speed and compression settings
  (Session preferences: Connection)

Working with remote X11 (V)

- VNC is an other alternative due to its specialized transmission protocol which only sends “framebuffer updates” using a Client/Server architecture
- Optimized implementation for 3D workloads: https://www.turbovnc.org/
  (but any other VNC implementation should also do; clients are available for diverse operating system)
- Can also be used for simple screen sharing (including limited guest access)
- Traffic can be tunneled through SSH

```
remote% vncserver :1 -localhost
 local% ssh -N -T -L 5901:<server's address>:5901 &
 local% vncviewer localhost:5901
```

(TightVNC and TurboVNC also have a built-in -via option)

Working with remote OpenGL

- VirtualGL (native or in combination with VNC) might be an option.
- Currently **not** implemented on the HPC systems for general use, but a dedicated node can be provisioned on short notice if needed.

- See https://www.hpc.rrze.fau.de/kundenbereich/campustreffen/HPC-Campustreffen-2014-02-11.pdf (slide 19-22) for more info. However, keep in mind that the property `:virtualgl` is currently not active by default as described above

Last resort: Xpra with HTML5 client

1. Login to Woody and run
   `qsub -q xpra /apps/xpra/job.pbs`
2. Query your node using **`qstat -rn1`**
3. On your local system without VPN
   `ssh -J <HPC_ACCOUNT>@cshpc.rrze.uni-erlangen.de`
   `-L 10101:localhost:10101 <HPC_ACCOUNT>@<PBS_NODE>`
   or with VPN
   `ssh -L 10101:localhost:10101 <HPC_ACCOUNT>@<PBS_NODE>`

- Open **`http://localhost:10101`** with a web browser; use your HPC account/password to login.
- Max 2 queued jobs per user; max 4h per job

REGIONALES RECHENZENTRUM ERLANGEN [RRZE]

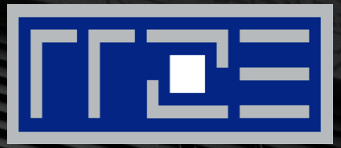

Transferring data

HPC Services, RRZE

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG

Accessing / transferring data on $FASTTMP

- $FASTTMP is not mounted on cshpc
  - Copy data to any /home/* directory ☹
  - Tunnel scp/rsync to cluster frontend through cshpc

    ```
    % ssh -L 1234:emmy:22 hpc123@cshpc
    % scp -P 1234 hpc123h@localhost:/elxfs/…  .
    ```

    (with **ssh** it's **-p**, but **-P** for **scp**)
  - Or better (simpler to use) with **~/.ssh/config**

    ```
    Host emmy-via-cshpc
      User hpc123   HostName emmy
      ProxyCommand ssh hpc123@cshpc.rrze.fau.de -W emmy:22
    ```
  - **scp emmy-via-cshpc:/elxfs/… .**

REGIONALES RECHENZENTRUM ERLANGEN [RRZE]

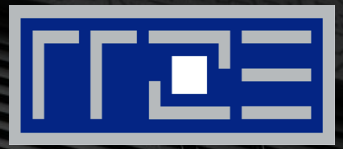

Mounting remote file systems with SSHFS

HPC Services, RRZE

Remote file system access

Problem: How to access remote files via an SSH connection?

- Windows: WinSCP
- Linux: Midnight Commander
- macOS:
  - CyberDuck
  - ForkLift3

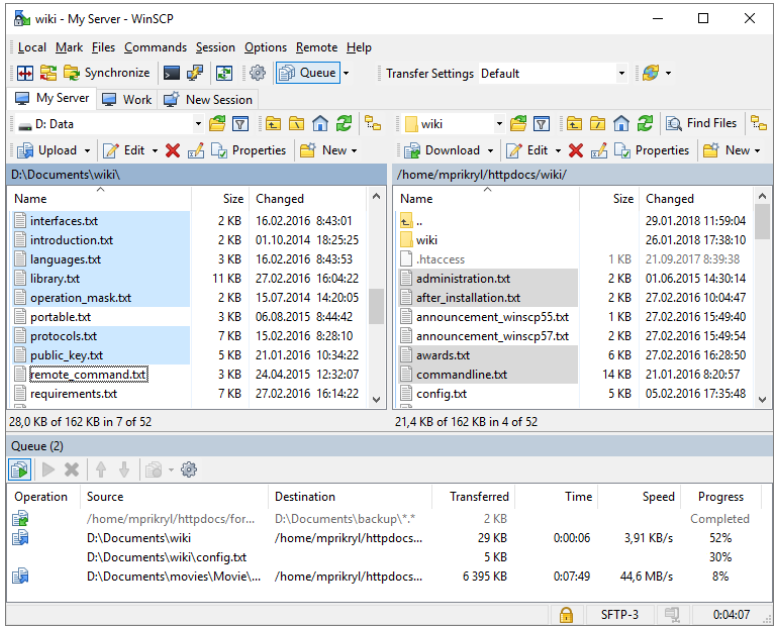

SSHFS

For macOS:
Fuse for macOS and plugin SSHFS

- Usage

```
sshfs [user@]host:[dir] mountpoint [options]
```

- Example:

```
$ mkdir ~/mnt
$ sshfs unrz55@cshpc.rrze.fau.de:papers mnt \
        -o reconnect,idmap=user
```

try to reconnect

Map remote UID to local

SSHFS

… and instant magic:

```
$ df -k
Filesystem                          1K-blocks      Used Available Use% Mounted on
udev                                  4061680         0   4061680   0% /dev
[...]
unrz55@cshpc.rrze.fau.de:papers     102687672  55132232  42296176  57% /home/gh/mnt
```

Unmounting:

`$ fusermount -u ~/mnt` (Linux)

`$ diskutil umount ~/mnt` (macOS)

REGIONALES RECHENZENTRUM ERLANGEN [RRZE]

Final words

HPC Services, RRZE

General notes on HPC utilization (I)

- All HPC clusters currently have high load.
- Therefore, expect some waiting time in the queue and don’t ask us “why is my job not yet running despite being submitted some hours ago”.
- If jobs stick in the route queue you either submitted too many jobs for the execution queue (which is o.k.) or you made an error specifying node properties, walltime (too large or not separated by comma), etc.

General notes on HPC utilization (II)

- Always check the performance for your jobs, e.g. using https://www.hpc.rrze.uni-erlangen.de/HPC-Status/job-info.php
- Try to reduce the number of nodes requested per job, as large jobs usually have less efficiency and cause overhead due to emptying nodes

Questions during session

- Convert text files between Windows and Linux:
  - **`dos2unix <file>`** (in-place)
  - **`dos2unix <inputfile> <outputfile>`**
- Use putty for xpra? Yes, that's possible
- https://www.supercomputing-akademie.de/
  - Organized by HLRS Stuttgart
  - High quality content