# Welcome to the CFD-Workshop Day

15.6.2021

Lewin Stein

# Welcome

Please ask questions anytime → unmute + chat

Materials/all slides → www.hlrn.de/doc/display/PUB/CFD+Day+2023

# Agenda

**10am morning session**

CFD project planning on a supercomputer, parallelization concepts, scaling tests
[Lewin Stein - NHR@ZIB]

Running OpenFOAM on the GPU
[Ogaja - NHR@GWDG]

**12.30pm lunch break (1h)**

**13.30pm afternoon session**

Flow solver CODA –
How to simulate aircraft on GPUs
[Immo Huismann – DLR]

Ansys Fluent - nozzle flow (CPU vs GPU)
[Lewin Stein - NHR@ZIB]

# CFD project planning on a supercomputer

15.6.2021

Lewin Stein

ZIB
ZUSE INSTITUTE BERLIN

# Start small if you plan big

Preparation on a test-account minimizes delay-risk of big projects

- ✓compile beforehand (check system environment: flags, modules, library paths)
- ✓scalability test of your software → recommended for project proposals, “sweet spot” detection
- ✓plan data workflow (size/file-number/temporal limits of home-, local-, (ssd)work-, perm-storage)

Minimal examples are speed boosters of your big project

- ✓easy debugging
- ✓provide “unit tests”
- ✓full cycle check (compile, run, evaluate results, optimize code)
- ✓enable others (us) to help you
- ✓fit in test-queue (fast feedback)
- ✓interactive execution (even faster feedback)

```
salloc -p standard96:test
ssh $SLURM_NODELIST
```

# Your time required for data mangement

**Standard96 CLX node**

RAM = 362GB

local SDD = 2TB

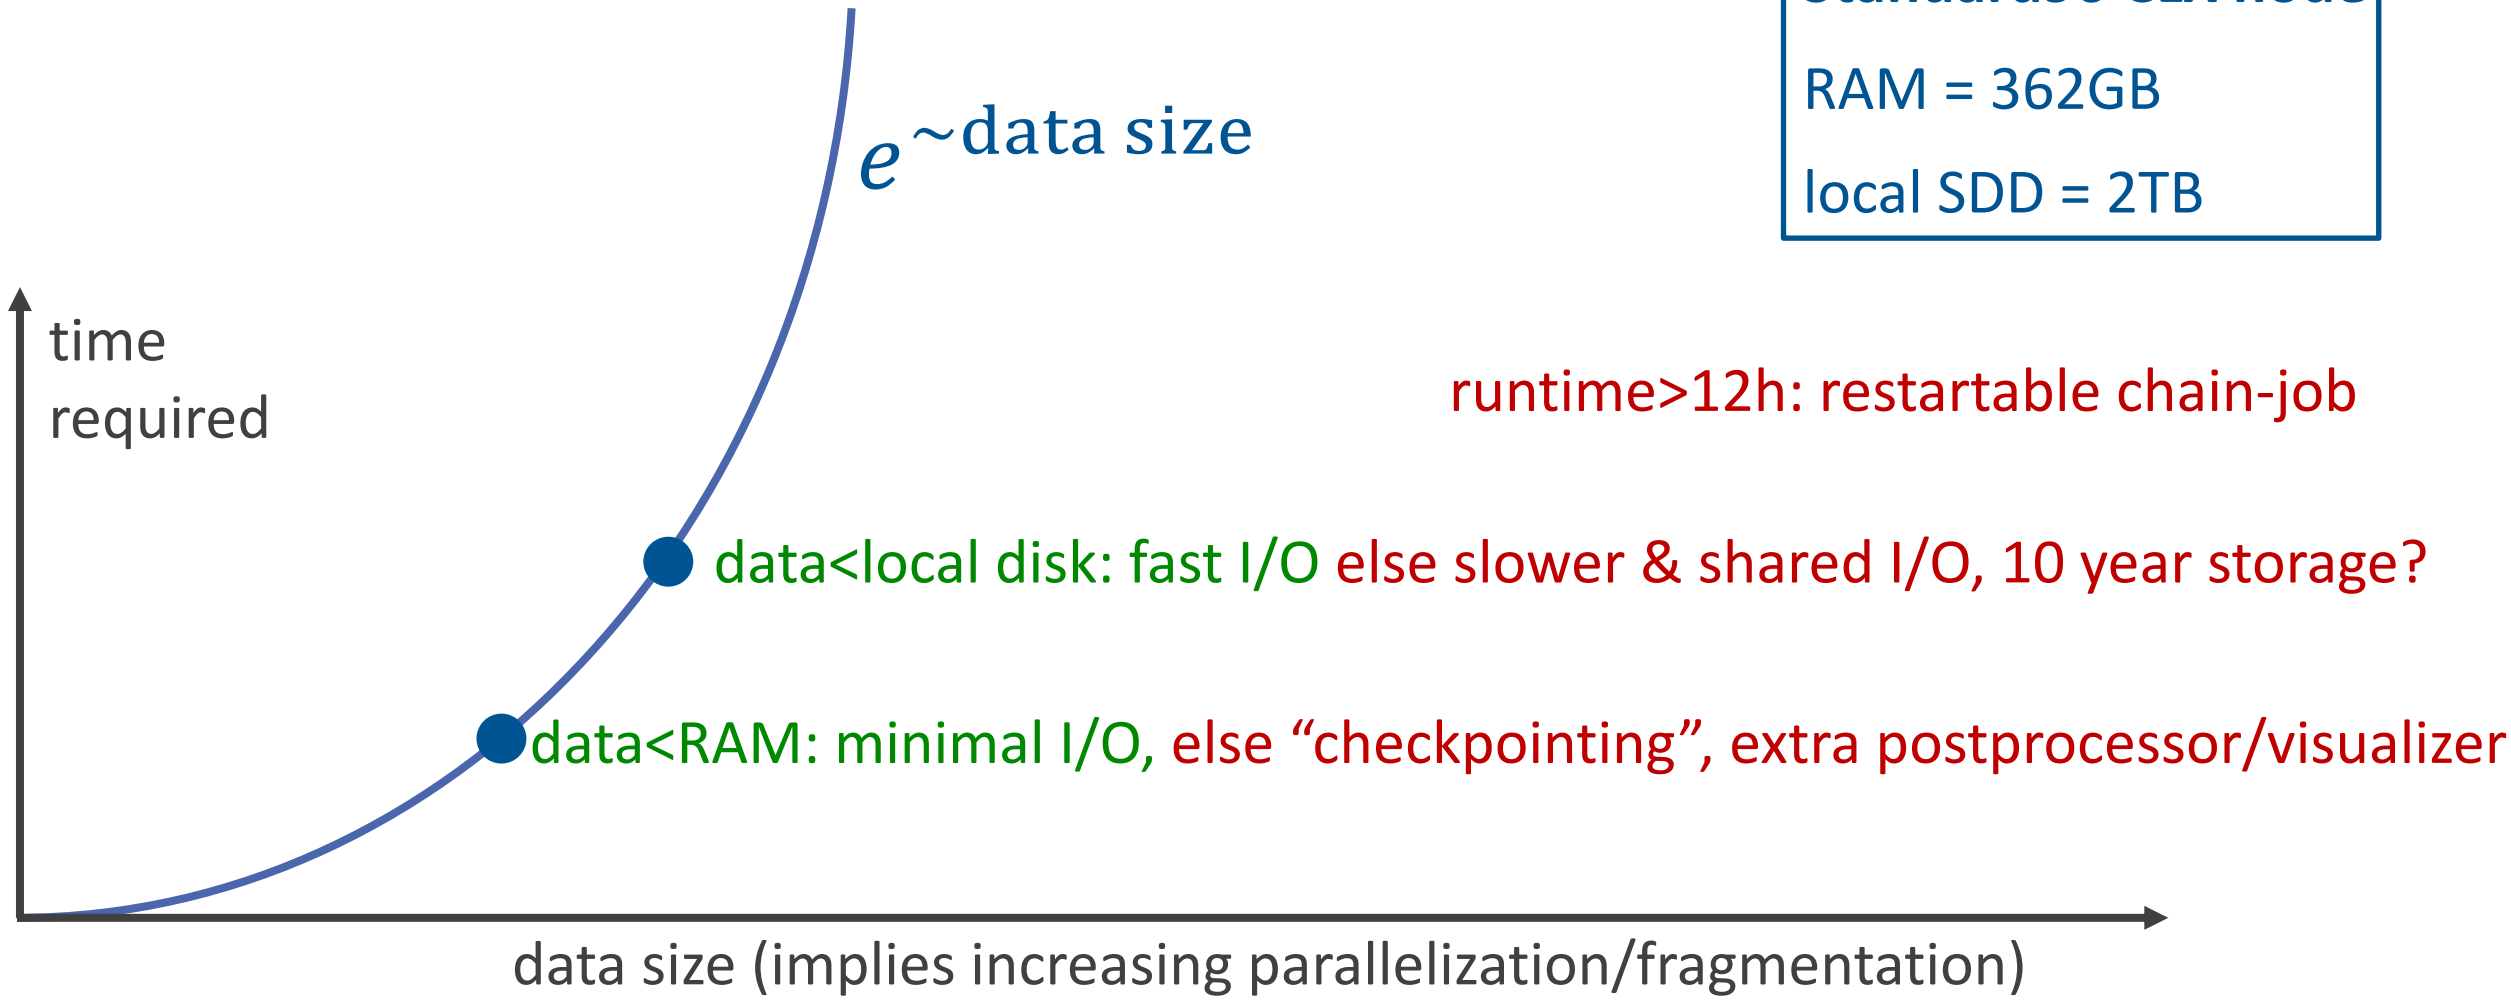

# Critical supercomputer resources

"A supercomputer is a device for turning compute-bound problems into I/O-bound problems."
Ken Batcher

uncritical / independent: local operations on compute-nodes

moderate / shared: cpu load of login-nodes

critical / scarce + shared: many file-I/O on work (Lustre)

→ Stay below 10 file access per cpu-core per minute.

(as temporal average of your job)

# Numerical Precision

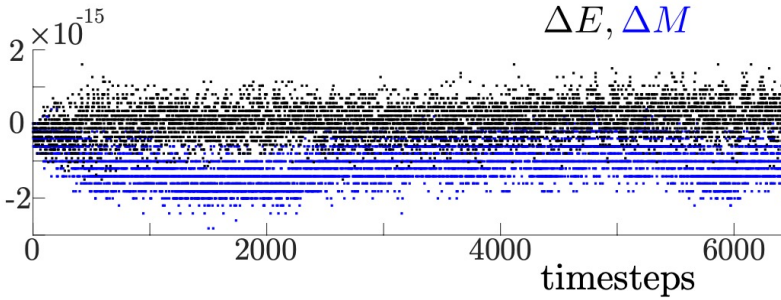


```
single precision:       S EEEEEEEE FFFFFFFFFFFFFFFFFFFFFFF
(32 bit)                0 1      8 9                       31
```

→ 7-8 decimal digits

```
double precision:       S EEEEEEEEEEE   FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
(64 bit)                0 1         11 12                                                      63
```

→ 15-16 decimal digits

```
rounding error:     ambient pressure 10000.000 Pa
(32 bit)          + human voice          + 0.025 Pa
                                   = 10000.02 or 10000.03
```

Only use double if needed by physical problem (range of system scales $>10^6$).

# Optimization hints

- optimize most expansive lines first “Pareto principle”
- serial performance (underestimated) is as important as scalability
- high impact, low effort: tic&toc, `htop` to check memory/affinity (before VTune, TotalView)
- use libraries whenever suitable “standing on the shoulders of giants”
  `module avail` → MPI, HDF5/NetCDF, FFTW, BLAS, …

# Access to & Application for Lise & Emmy

Has your principal investigator (PI) an address at a public German university?

Open your test account any time at
https://zulassung.hlrn.de
(75k up to 300k cpu-hours per quarter)

Yearly application deadlines of full projects are: 28.1, 28.4, 28.7, 28.10.

You will get feedback within 8 weeks and cpu-h credited at: 1.4, 1.7, 1.10, 1.1.

More details? Look here:
www.hlrn.de/doc/display/PUB/Application+Process

How to access our GPU? Look here:
www.hlrn.de/doc/display/PUB/GPU+A100+Partition

# Access to & Application for Lise & Emmy

https://jards.nhr-verein.de

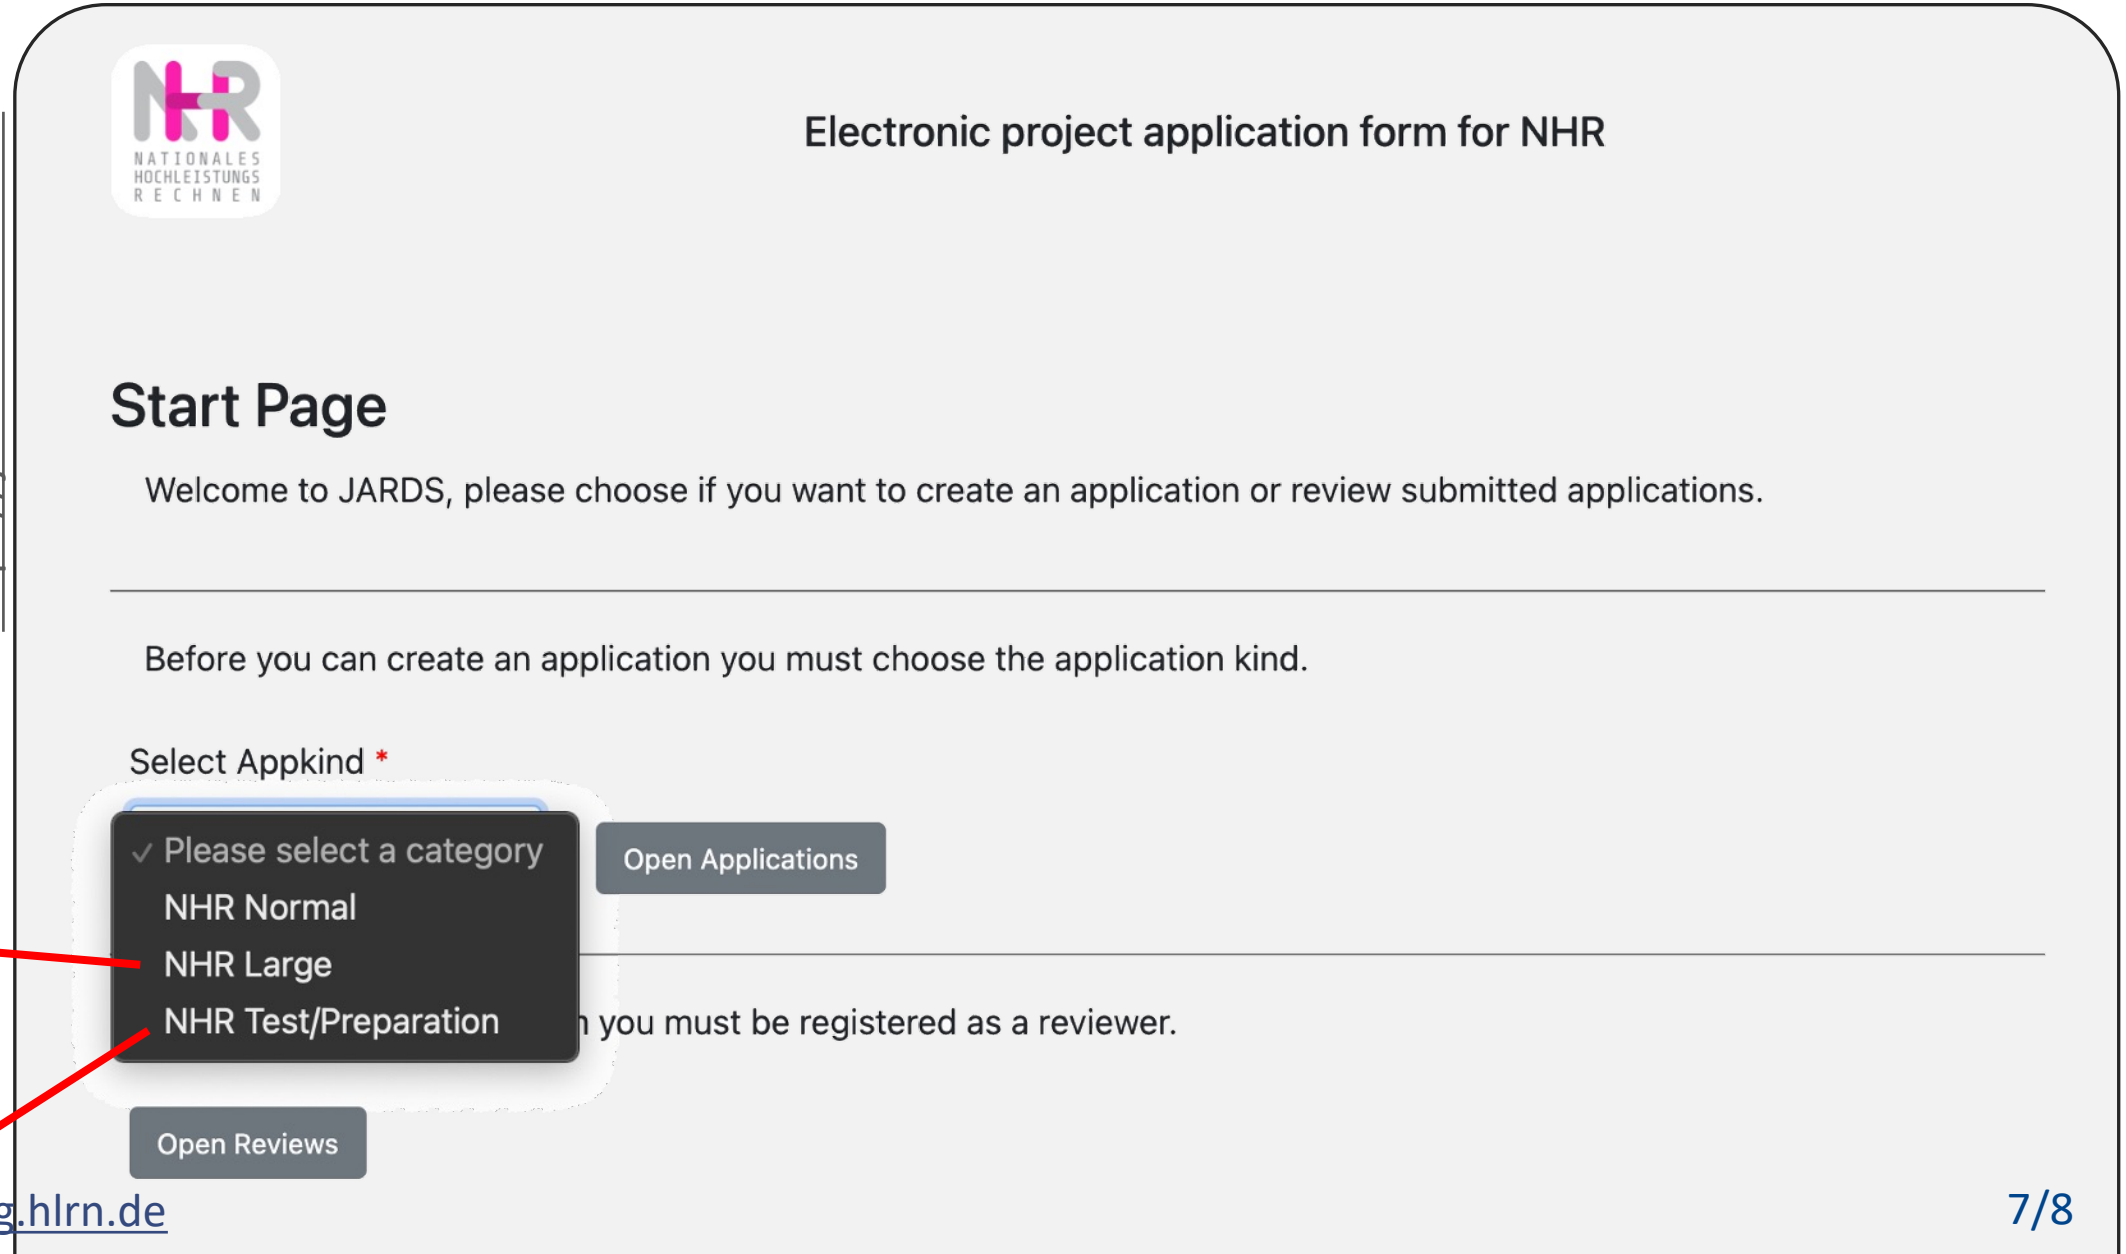


Large project?
20 M. core-h/year
or more…

Not functional yet
use https://zulassung.hlrn.de

# Hints for a successful application

Provide arguments why exactly N runs are necessary.

Proof that your specific case scales well on our specific machine.

Demonstrate your HPC experience.

Read: www.hlrn.de/doc/display/PUB/Project+proposal