



Good morning slides - Thursday




- Everybody able now to access all required resources?
- Lmod will be default on Hawk → everybody should use it
- qstat shows your own jobs only (due to DSGVO ...)
→ use `batchstat (-t rome)` instead
- Path to MPT F08 modules should be in env vars now
- `rome1` will have NPS=1 in the next hours
- Slides available at `login1:/tmp/slides_porting_workshop`
and workshop homepage

If your code is ready to use ...

... it's time for optimization by means of low-hanging fruits:

- Start with measuring baseline
- Try different numbers of active cores (cf. scripts in the Wiki)
- Try different
 - Compilers
 - MPIs
 - BLAS/LAPACK (MKL, AOCL, PGI)
- Try compiler flags suggested in <http://www.prace-ri.eu/IMG/pdf/Best-Practice-Guide-AMD.pdf> (replace znver1 by znver2!)
- Investigate NUMA behavior
- Adapt MPI parameters
- Use 32bit versions of libraries if sufficient
- Try link time and profile-guided optimization
- Try different
 - NUMA nodes per socket
 - Transparent huge pages settings
 - NUMA balancing settings



- There might be runtime variations
→ if required use median of multiple runs in order to establish baseline and assess benefit of optimisations
- Please focus on time stepping loop, neglect setup and clean-up phases (by means of Tool APIs)
- Please state improvements in terms of „serial speedup“, i.e.
$$S = \frac{t_{old}}{t_{new}}, \quad \text{e.g. } 1.2x$$
- Also look on *absolute* runtimes consumed by a specific routine / group / etc., not only *fraction* of walltime!
- From time to time, you should validate your optimizations w.r.t. your production job (size)