

packems

command line tool for packing and
archiving of data



Max-Planck-Institut
für Meteorologie

Idea

- support arbitrary structure of data
- packing in .tar files of given target size
- ensure file integrity - no split files
- optionally store .tar files on HSM system
- add information to easily restore/retrieve files
- optionally removal of processed files

Design

- computation of file to tar mappings by file sizes
- use *make* to actually process the files
 - parallelization, re-run
- allow immediate use of queuing system

Packing

- file groups get their own set of .tar files
- grouping by name (time series) or time (restart files)
- distribution of groups over .tar files by size (largest first), by order (acc. to grouping), round-robin

Archiving

- expects Kerberos (or netrc) authentication
- stops if files already exist or local and remote sizes differ

Testing

- MPI-ESM 1.2 LR historical runs
 - 9407 files, total 4.5 TiB → avg. 496 MiB, max. 10 GiB, min. 7 kiB
- Full prepost/compute/compute2 node on mistral, 18 parallel *make* jobs
- files and packing on the same file system
- group by name, distribute by size; with removal, archiving
- 100 GiB target size

Results

- packing created 54 .tar files, 108 - 61 GiB, 1 x 9 GiB
- elapsed time 01:25 h
- total transfer rate ~925 MiB/s
- avg. transfer rate pftp ~87 MiB/s x 18 jobs → 1.6 GiB/s

Open

- compression of data is handled separately
- seamless operation difficult to continue after errors