High Performance Research Computing
DIVISION OF RESEARCH

TEXAS A&M UNIVERSITY
Oceanography

Porting CESM to Grace

May 25, 2021
4th Annual Texas A&M Research Computing Symposium

Abishek Gopal
Assistant Research Scientist
iHESP, Texas A&M Oceanography
Texas A&M High Performance Research Computing

Acknowledgements

- **Francis Dang (HPRC)** helped diagnose the numerous issues running HR-CESM on > 400 nodes.
- **Jim Edwards (NCAR & iHESP)** helped optimize HR-CESM performance for Grace
- **Lisa Perez & Michael Dickens (HPRC)** helped figure out the Hierarchical NMS in EasyBuild.
- **Qiuying Zhang (iHESP)** helped with porting CESM v2 and testing model runs on Grace.

Porting the Community Earth System Model (CESM)

~2M lines of Fortran code developed over 20+ years*

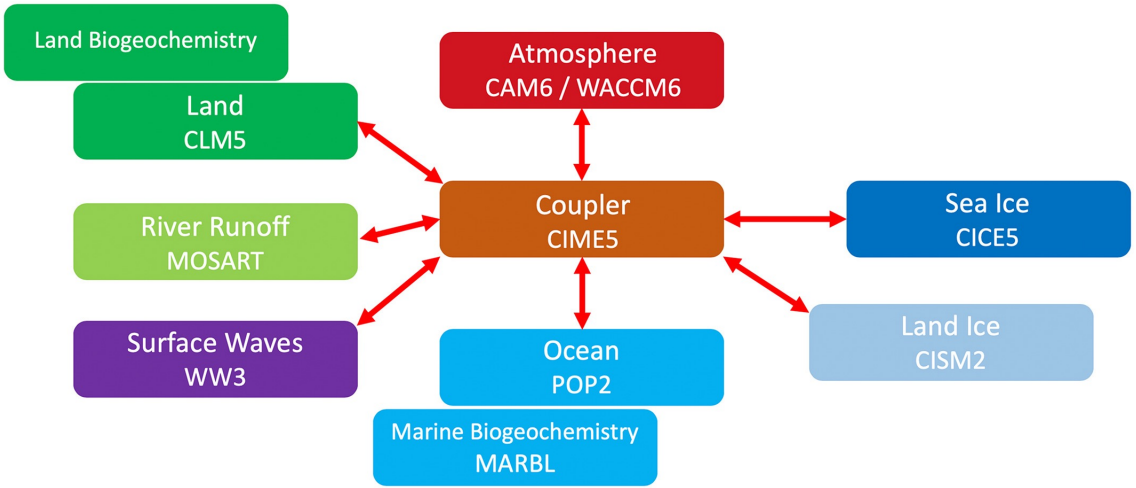

https://www.cesm.ucar.edu/

Relevant CIME configuration

1. config_machines.xml
2. config_compilers.xml
3. config_batch.xml

*Allison Baker, SC 2017 workshop

Building the CESM software dependencies - I

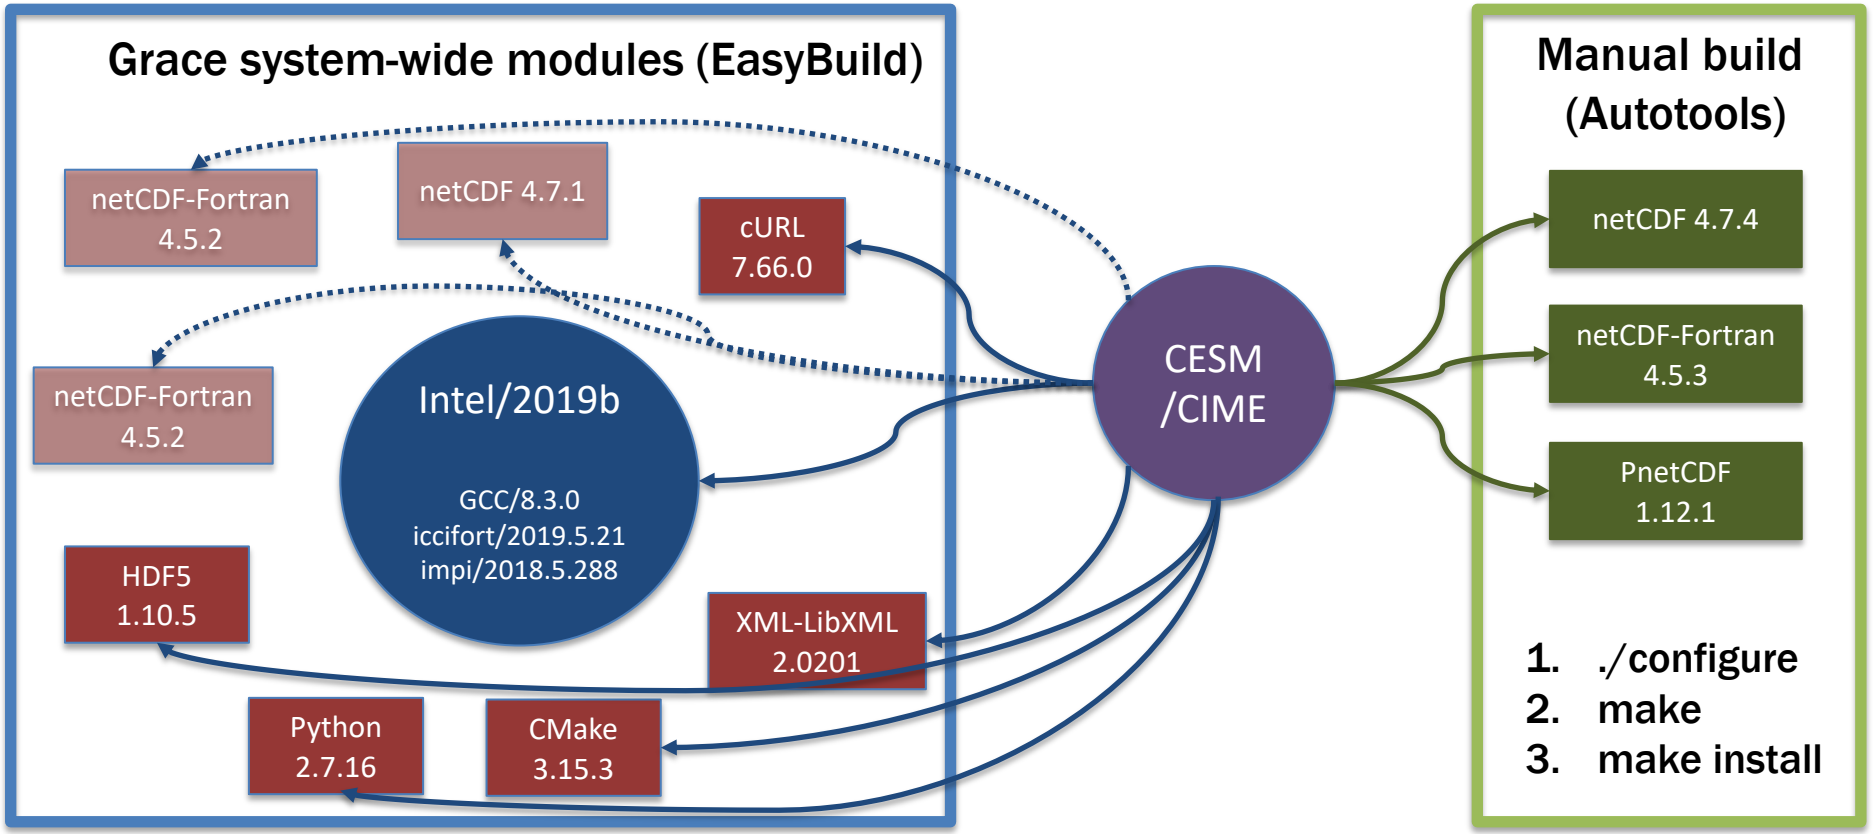

Building the CESM software dependencies - II

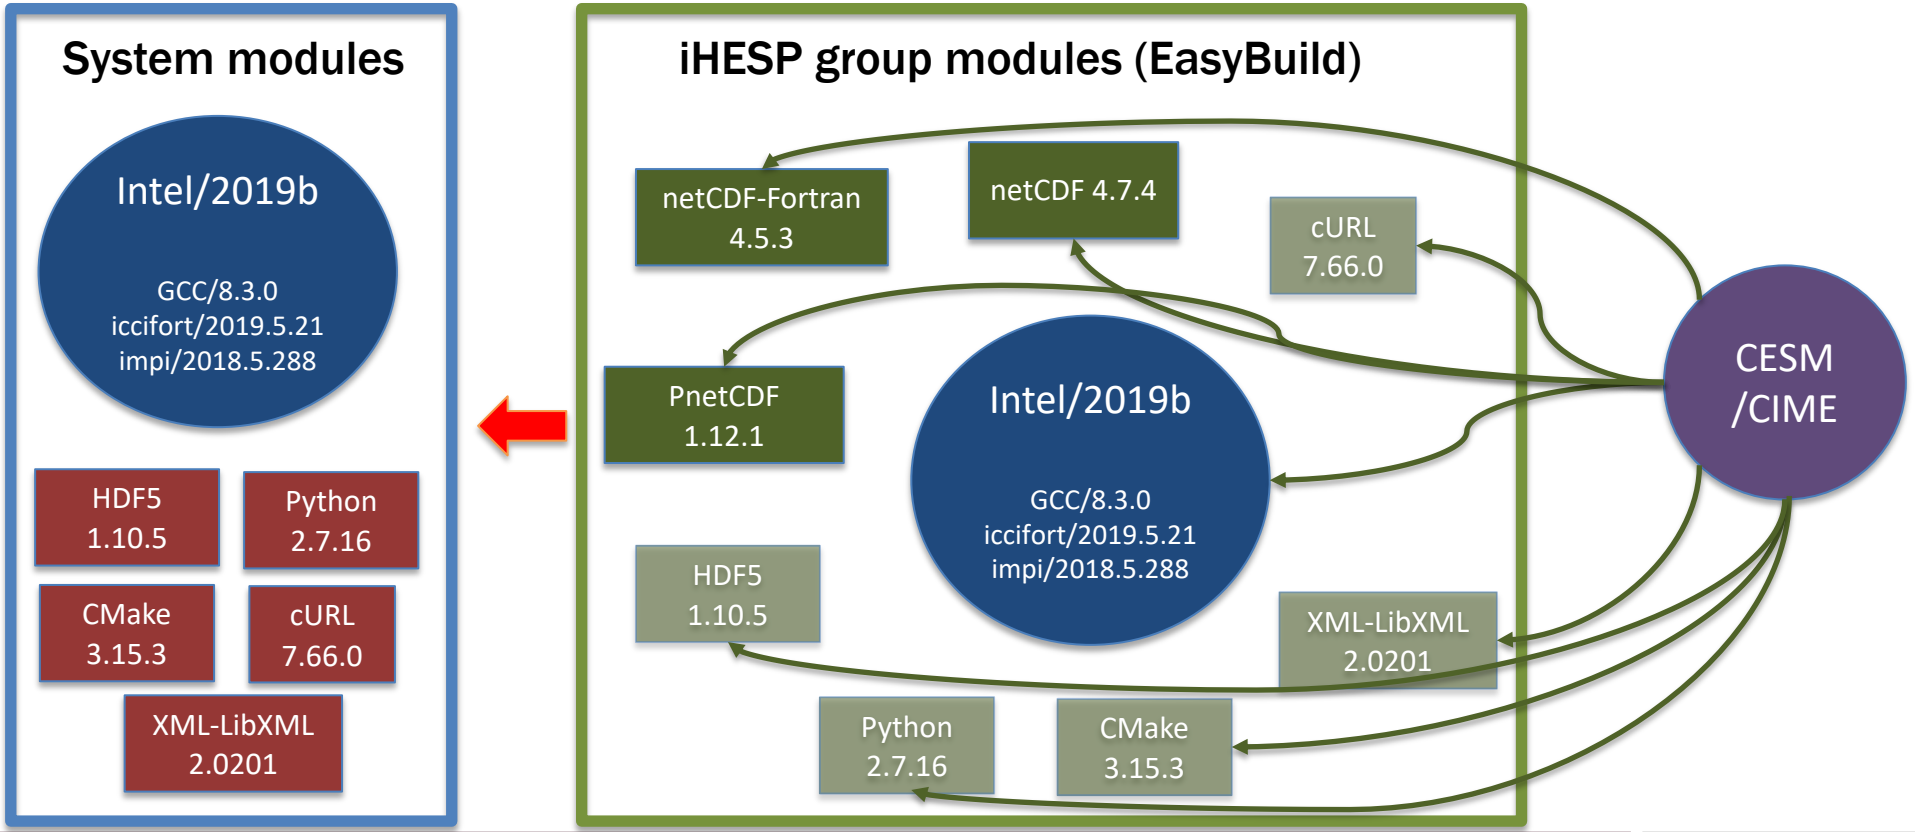

Summary of steps involved

1. Test low-res, fully-coupled CESM configuration on 128 nodes

CPU -> Task pinning issues.
Runs very slow

2. Upgrade to SLURM v20.11.3 and retry (thanks to Francis Dang, HPRC)

3. Test high-res, fully-coupled CESM configuration on 128 nodes

4. Scale up case to 400 & 800 nodes

UCM connect: REQ
RETRIES EXHAUSTED

Summary of steps involved

5. Rebuild netCDF libraries to ensure parallel IO & retry ✗

6. Test on GPFS-mounted drive (instead of Lustre) ✗

7. Upgrade toolchain + dependencies to intel/2020b and retry ✓

8. Optimize HR-CESM on Grace for 400 & 800 node runs (Thanks to Jim Edwards, NCAR)

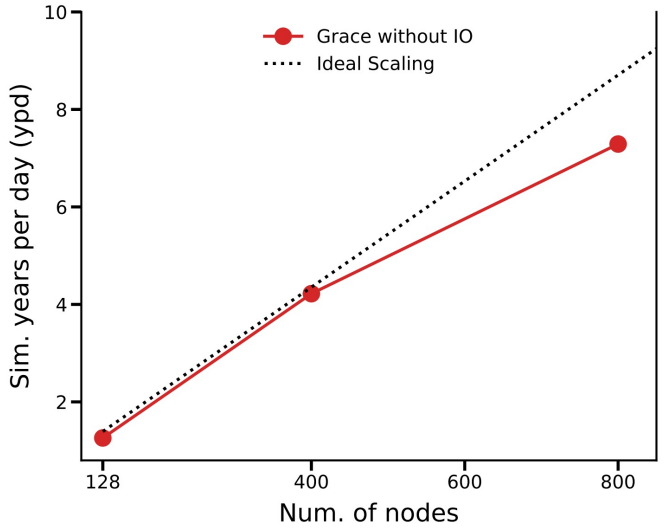

Next steps

8. Compare IO speeds on iHESP (GPFS) drive vs Scratch (Lustre) drive

9. Validate the correctness of the porting using the CESM Ensemble Consistency Test tools https://github.com/NCAR/PyCECT

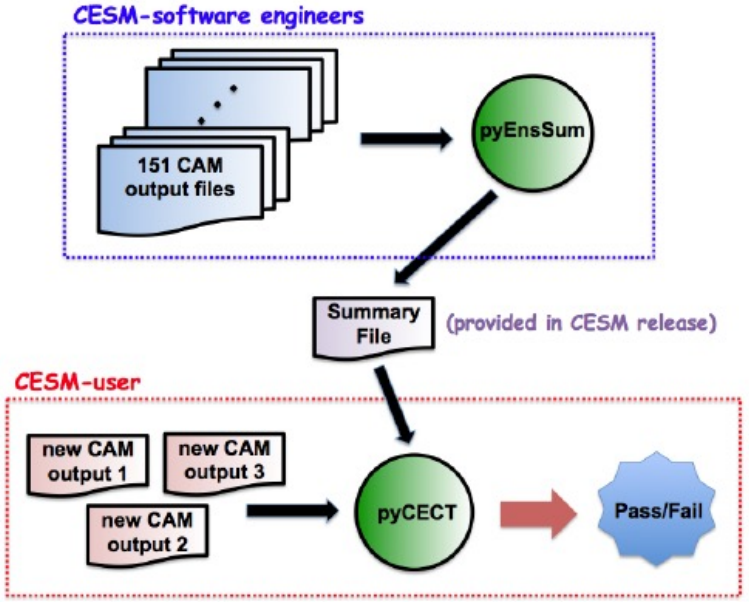

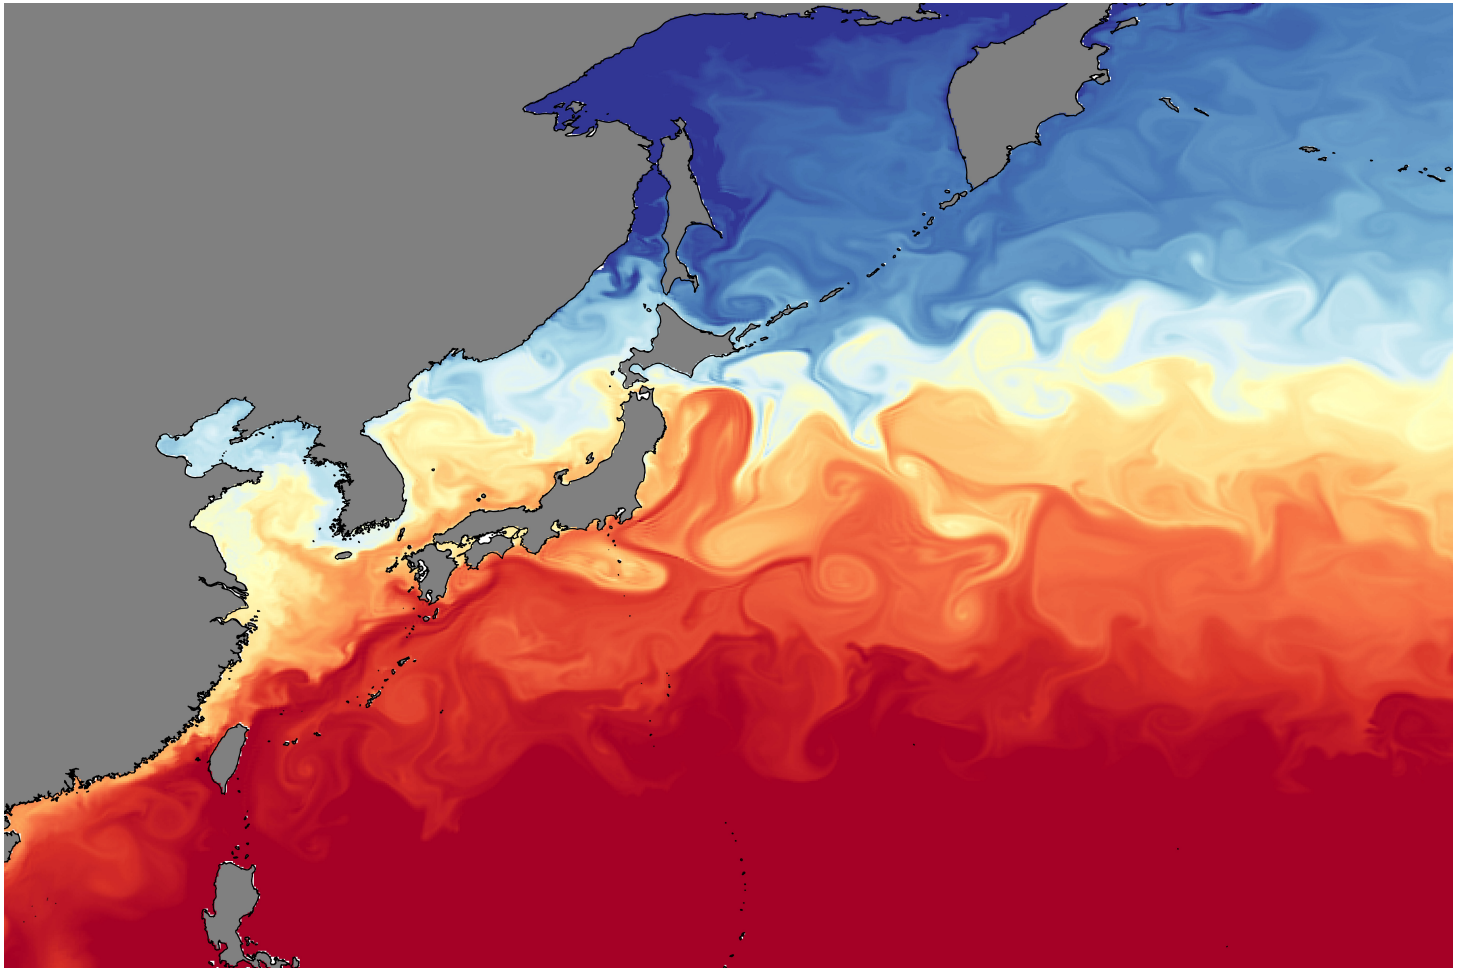

Credit: Dapeng Li, iHESP

Intel/2020b toolchain

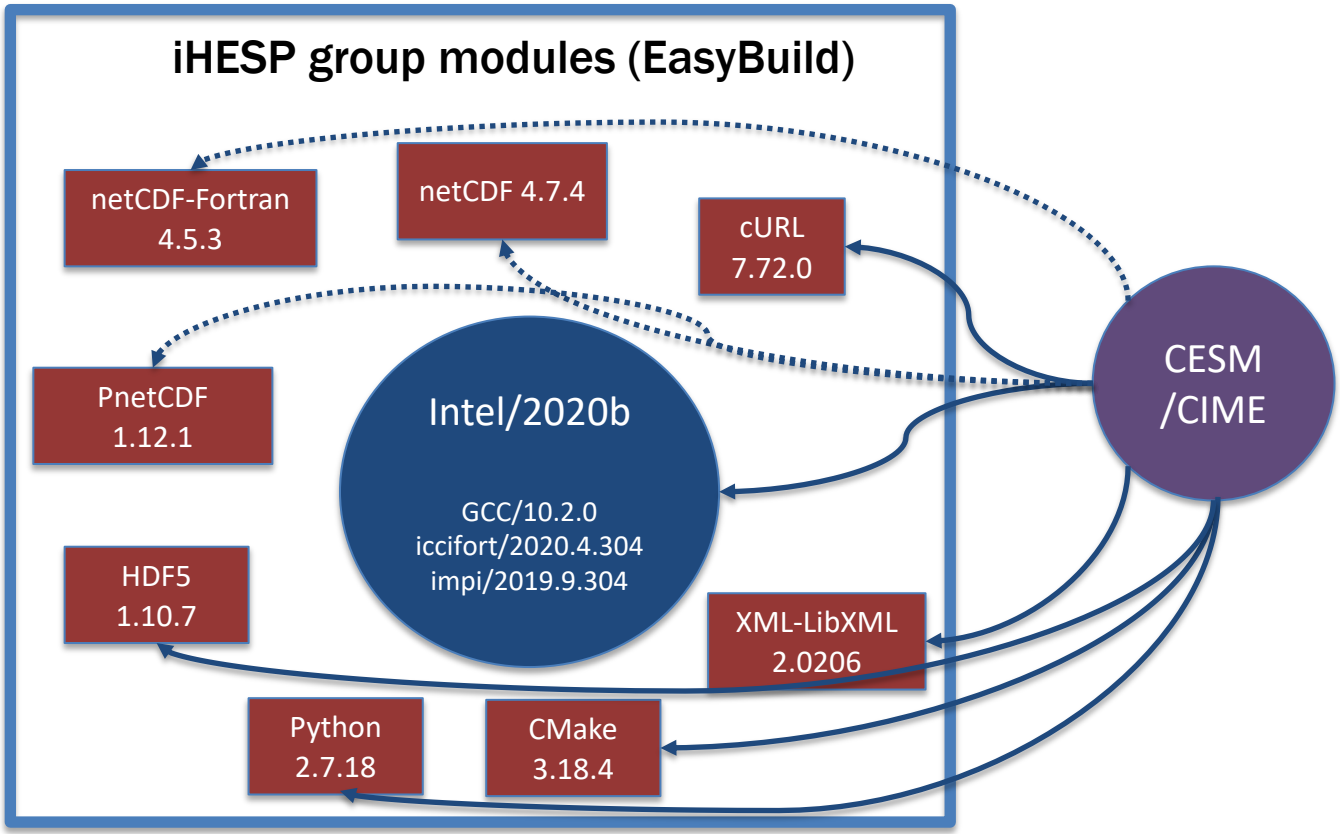

CESM Software Dependencies

- UNIX style operating system such as CNL, AIX or Linux
- python >= 2.7
- perl 5
- subversion client (version 1.8 or greater but less than v1.11) for downloading CAM, POP, and WW3
- git client (1.8 or greater)
- Fortran compiler with support for Fortran 2003
- C compiler
- MPI (although CESM does not absolutely require it for running on one processor)
- NetCDF 4.3 or newer.
- ESMF 5.2.0 or newer (optional).
- pnetcdf 1.7.0
- CMake 2.8.6 or newer