

EGU2020-18121

<https://doi.org/10.5194/egusphere-egu2020-18121>

EGU General Assembly 2020

© Author(s) 2022. This work is distributed under the Creative Commons Attribution 4.0 License.



Performance gains in an ESM using parallel ad-hoc file systems

Stefan Versick, Ole Kirner, Jörg Meyer, Holger Obermaier, and Mehmet Soysal

Karlsruhe Institute of Technology, Steinbuch Centre for Computing, Eggenstein-Leopoldshafen, Germany

Earth System Models (ESM) got much more demanding over the last years. Modelled processes got more complex and more and more processes are considered in models. In addition resolutions of the models got higher to improve weather and climate forecasts. This requires faster high performance computers (HPC) and better I/O performance.

Within our Pilot Lab Exascale Earth System Modelling (PL-EESM) we do performance analysis of the ESM EMAC using a standard Lustre file system for output and compare it to the performance using a parallel ad-hoc overlay file system. We will show the impact for two scenarios: one for today's standard amount of output and one with artificial heavy output simulating future ESMs.

An ad-hoc file system is a private parallel file system which is created on-demand for an HPC job using the node-local storage devices, in our case solid-state-disks (SSD). It only exists during the runtime of the job. Therefore output data have to be moved to a permanent file system before the job has finished. Quasi in-situ data analysis and post-processing allows to gain performance as it might result in a decreased amount of data which you have to store - saving disk space and time during the transfer of data to permanent storage. We will show first tests for quasi in-situ post-processing.