

Data Storage Concepts for pre-exascale Earth System Model Output handling

Karsten Peters-von Gehlen, et al

Abstract for NHR Data Storage Workshop, 22.9.23

Output data volumes stemming from simulations performed with state-of-the art numerical Earth System Models are currently approaching scales on the order of several PetaBytes per multi-year simulation periods. With the immanent need for longer, climate time scale simulations, the climate model output volume needing to be handled by scientists is most likely to continue increasing beyond what can be efficiently handled by current data storage concepts employed at HPC centers. With the DKRZ being a topical IT service provider for the simulation-based climate science community, current developments are targeted to continue providing custom-tailored data handling services for climate scientists.

From a hardware perspective, DKRZ's operational storage infrastructure consists of ~130 PB of disks (lustre file system) and a 300+ PB tape archive (StrongLink HSM). Further, in-house cloud storage resources (~1 PB) are available for data sharing with external partners, very much similar to anonymous FTP. Data stored on the lustre file system can be globally disseminated via federated data infrastructures, e.g. ESGF, and the tape archive is used in three different ways (operational project storage, long-term project storage, CoreTrustSeal certified long-term WDCC archival).

To meet the upcoming data challenges in simulation-based climate science, DKRZ is specifically working on methods to incorporate of the in-house tape archive (300+ PB) more into the day-to-day workflow of climate scientists in order to provide a seamless interplay of multiple storage tiers.

In this presentation, we will present the current status of high-volume data handling at DKRZ, illustrate the ongoing process of designing and implementing a domain-specific data handling solution at DKRZ and give an overview of the currently evolving data handling concepts geared towards exascale applications, e.g. digital twins.