



Contribution ID: 22

Type: **Oral presentation**

## Utilizing Heterogeneous Storage Infrastructures via the Earth-System Data Middleware

*Thursday, 26 September 2019 15:00 (30 minutes)*

The Earth-System Data Middleware (ESDM) is a performance-aware middleware that builds upon a data model similar to NetCDF and utilises a self-describing on-disk data format for storing structured data. ESDM allows to employ multiple (shared and local) storage systems concurrently and explicitly supports heterogeneous storage infrastructures. From the user/application perspective, the complexity of multiple-storage tiers is hidden, and data fragmentation controlled by a configuration file.

Performance measurements running on Mistral (DKRZ) show that multiple Lustre file systems can be used efficiently with the current ESDM version. Several configurations were tested generating performance of 200 GB/s. The results were compared to a best-case benchmark using IOR that achieves 150 GB/s. Using tmpfs locally, ESDM was able to achieve 2 TB/s on 500 nodes showing that non-volatile memory could also be well utilized. ESDM also provides a backend to utilise the Kove XPD in-memory shared storage.

While the tooling is not yet completed, ESDM already integrates with NetCDF and can be used as a drop-in replacement for typical use-cases without changing anything from the application perspective. While our current version utilises the manual configuration by data-centre experts, the ultimate long-term goal is to employ machine learning to automatise the decision-making and reduce the burden for users and experts.

### Keywords

NetCDF, ESDM, ESIWACE

**Primary authors:** Dr KUNKEL, Julian (University of Reading); Dr PEDRO, Luciana (University of Reading); Prof. LAWRENCE, Bryan (NCAS)

**Presenter:** Dr KUNKEL, Julian (University of Reading)

**Track Classification:** NEXTGenIO Workshop on applications of NVRAM storage to exascale I/O