

Adaptive Tier Selection for NetCDF/HDF5

Jakob Luettgau*, Eugen Betke*, Olga Perevalova‡, Julian Kunkel*, Michael Kuhn‡
 *German Climate Computing Center (DKRZ), ‡Universität Hamburg (UHH)

Abstract

Scientific applications on supercomputers tend to be I/O-intensive. To achieve portability and performance, data description libraries such as HDF5 and NetCDF are commonly used. Unfortunately, the libraries often default to suboptimal access patterns for reading/writing data to multi-tier distributed storage. This work explores the feasibility of adaptively selecting tiers depending on an application's I/O behavior.

Overview

The contributions presented in this work are:

- A proof of concept prototype implementation demonstrating the benefit of adaptive tier selection on a real system.
- An architecture for I/O middleware beyond adaptive tier selection for more intelligent data placement from user space.

Opportunities using HDF5 Virtual Object Layer

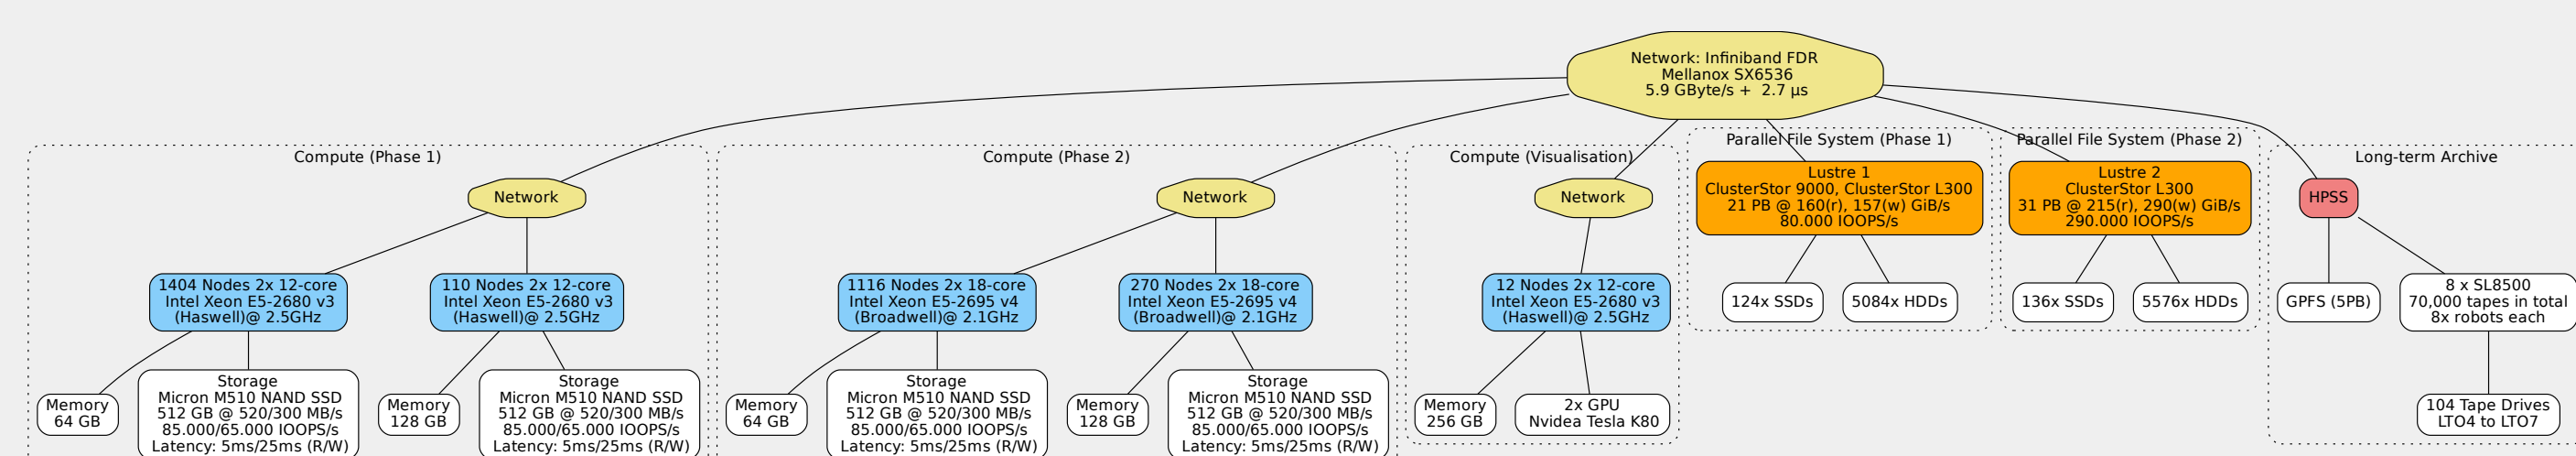
Hierarchical Data Format (HDF5): HDF5 is an open source, hierarchical, and self-describing format that combines data and metadata. Advantages of this format make it widely used by scientific applications.

Virtual Object Layer (VOL): The VOL is an abstraction layer in the HDF5 library with the purpose of exposing the HDF5 API to applications while allowing to use different storage mechanisms. The VOL intercepts all API calls and forwards those calls to plugin object drivers. Additionally, external VOL plugins are supported to allow third-party plugin development.

Plugin for Separate Metadata Handling: A VOL plugin was developed to handle data and metadata separately. For adaptive tier selection, this is necessary to keep track of alternating data sources but it also offers additional opportunities. Generated simulation data is routinely published, but at the moment not automatically catalogued. With the VOL plugin, it would be possible to extract a dataset description to make it available for search in a catalogue as the dataset is written.

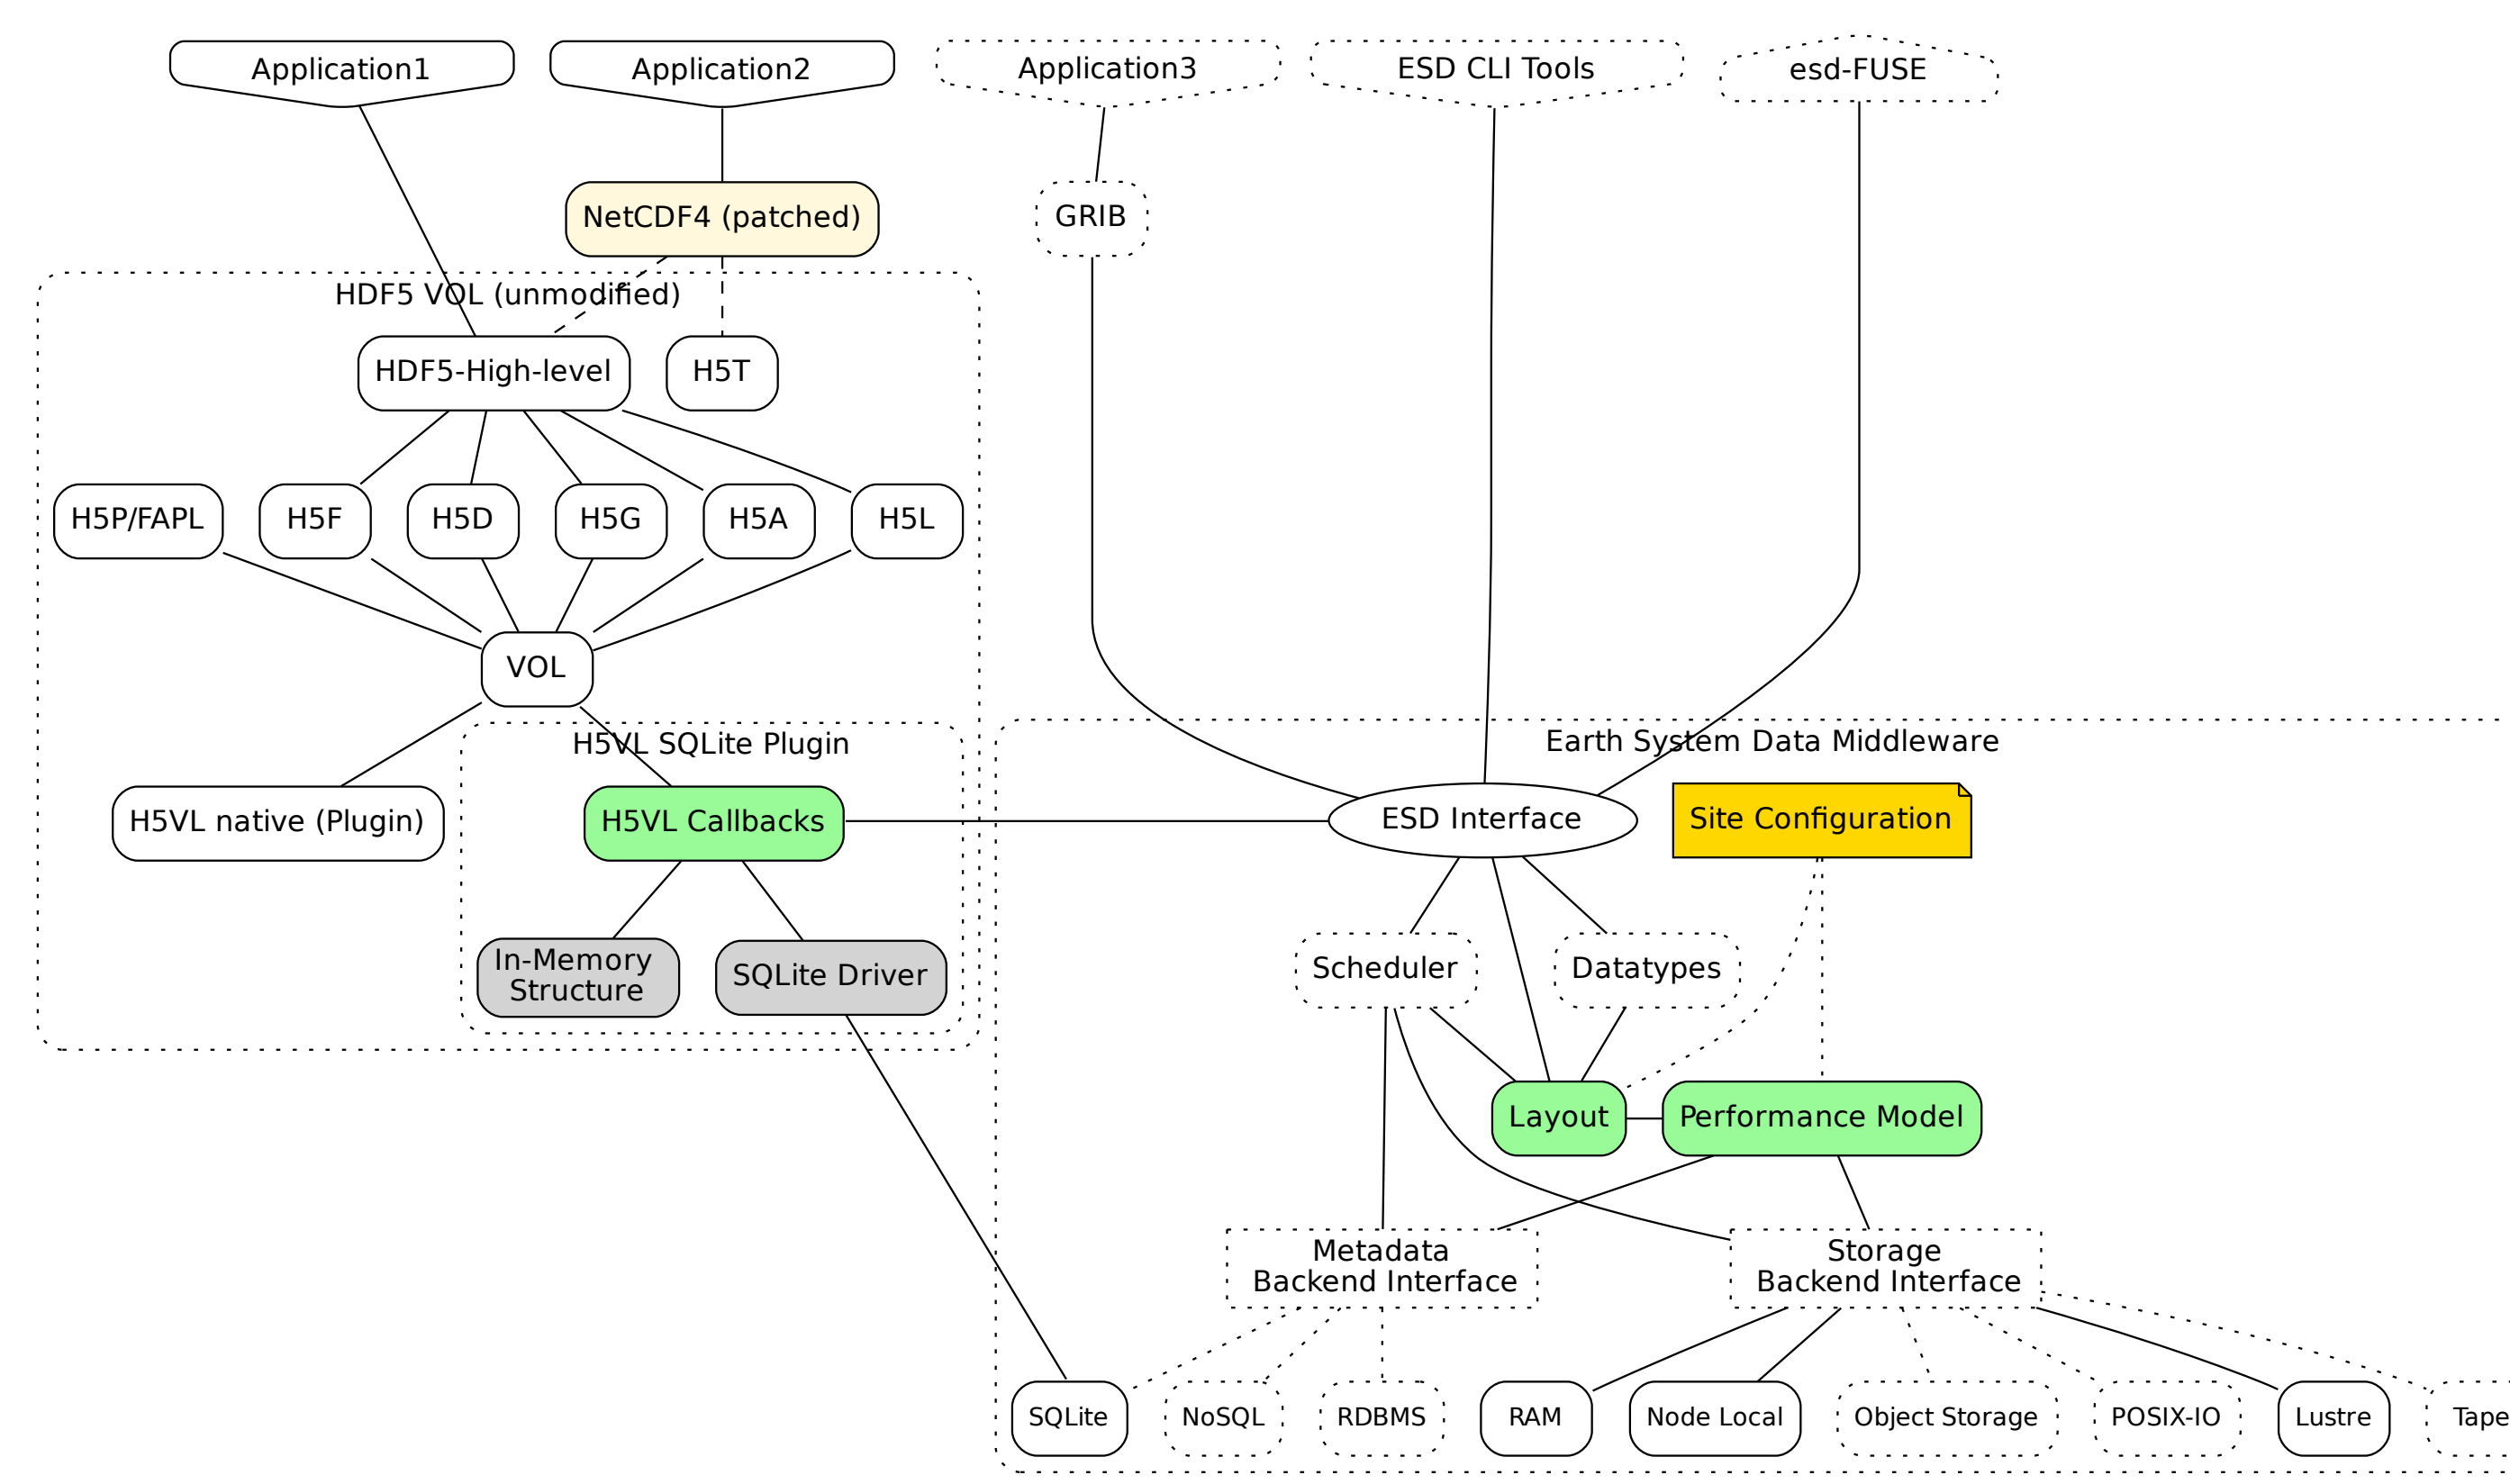
Mistral Supercomputer

Performance evaluation was carried out on the Mistral supercomputer. The computer, listed #38 in the Top500 (June 2017), is located at the German Climate Computing Center (DKRZ) in Hamburg and exclusively serves the climate research community. The current site configuration is as follows:



Architecture for Adaptive Tier Selection

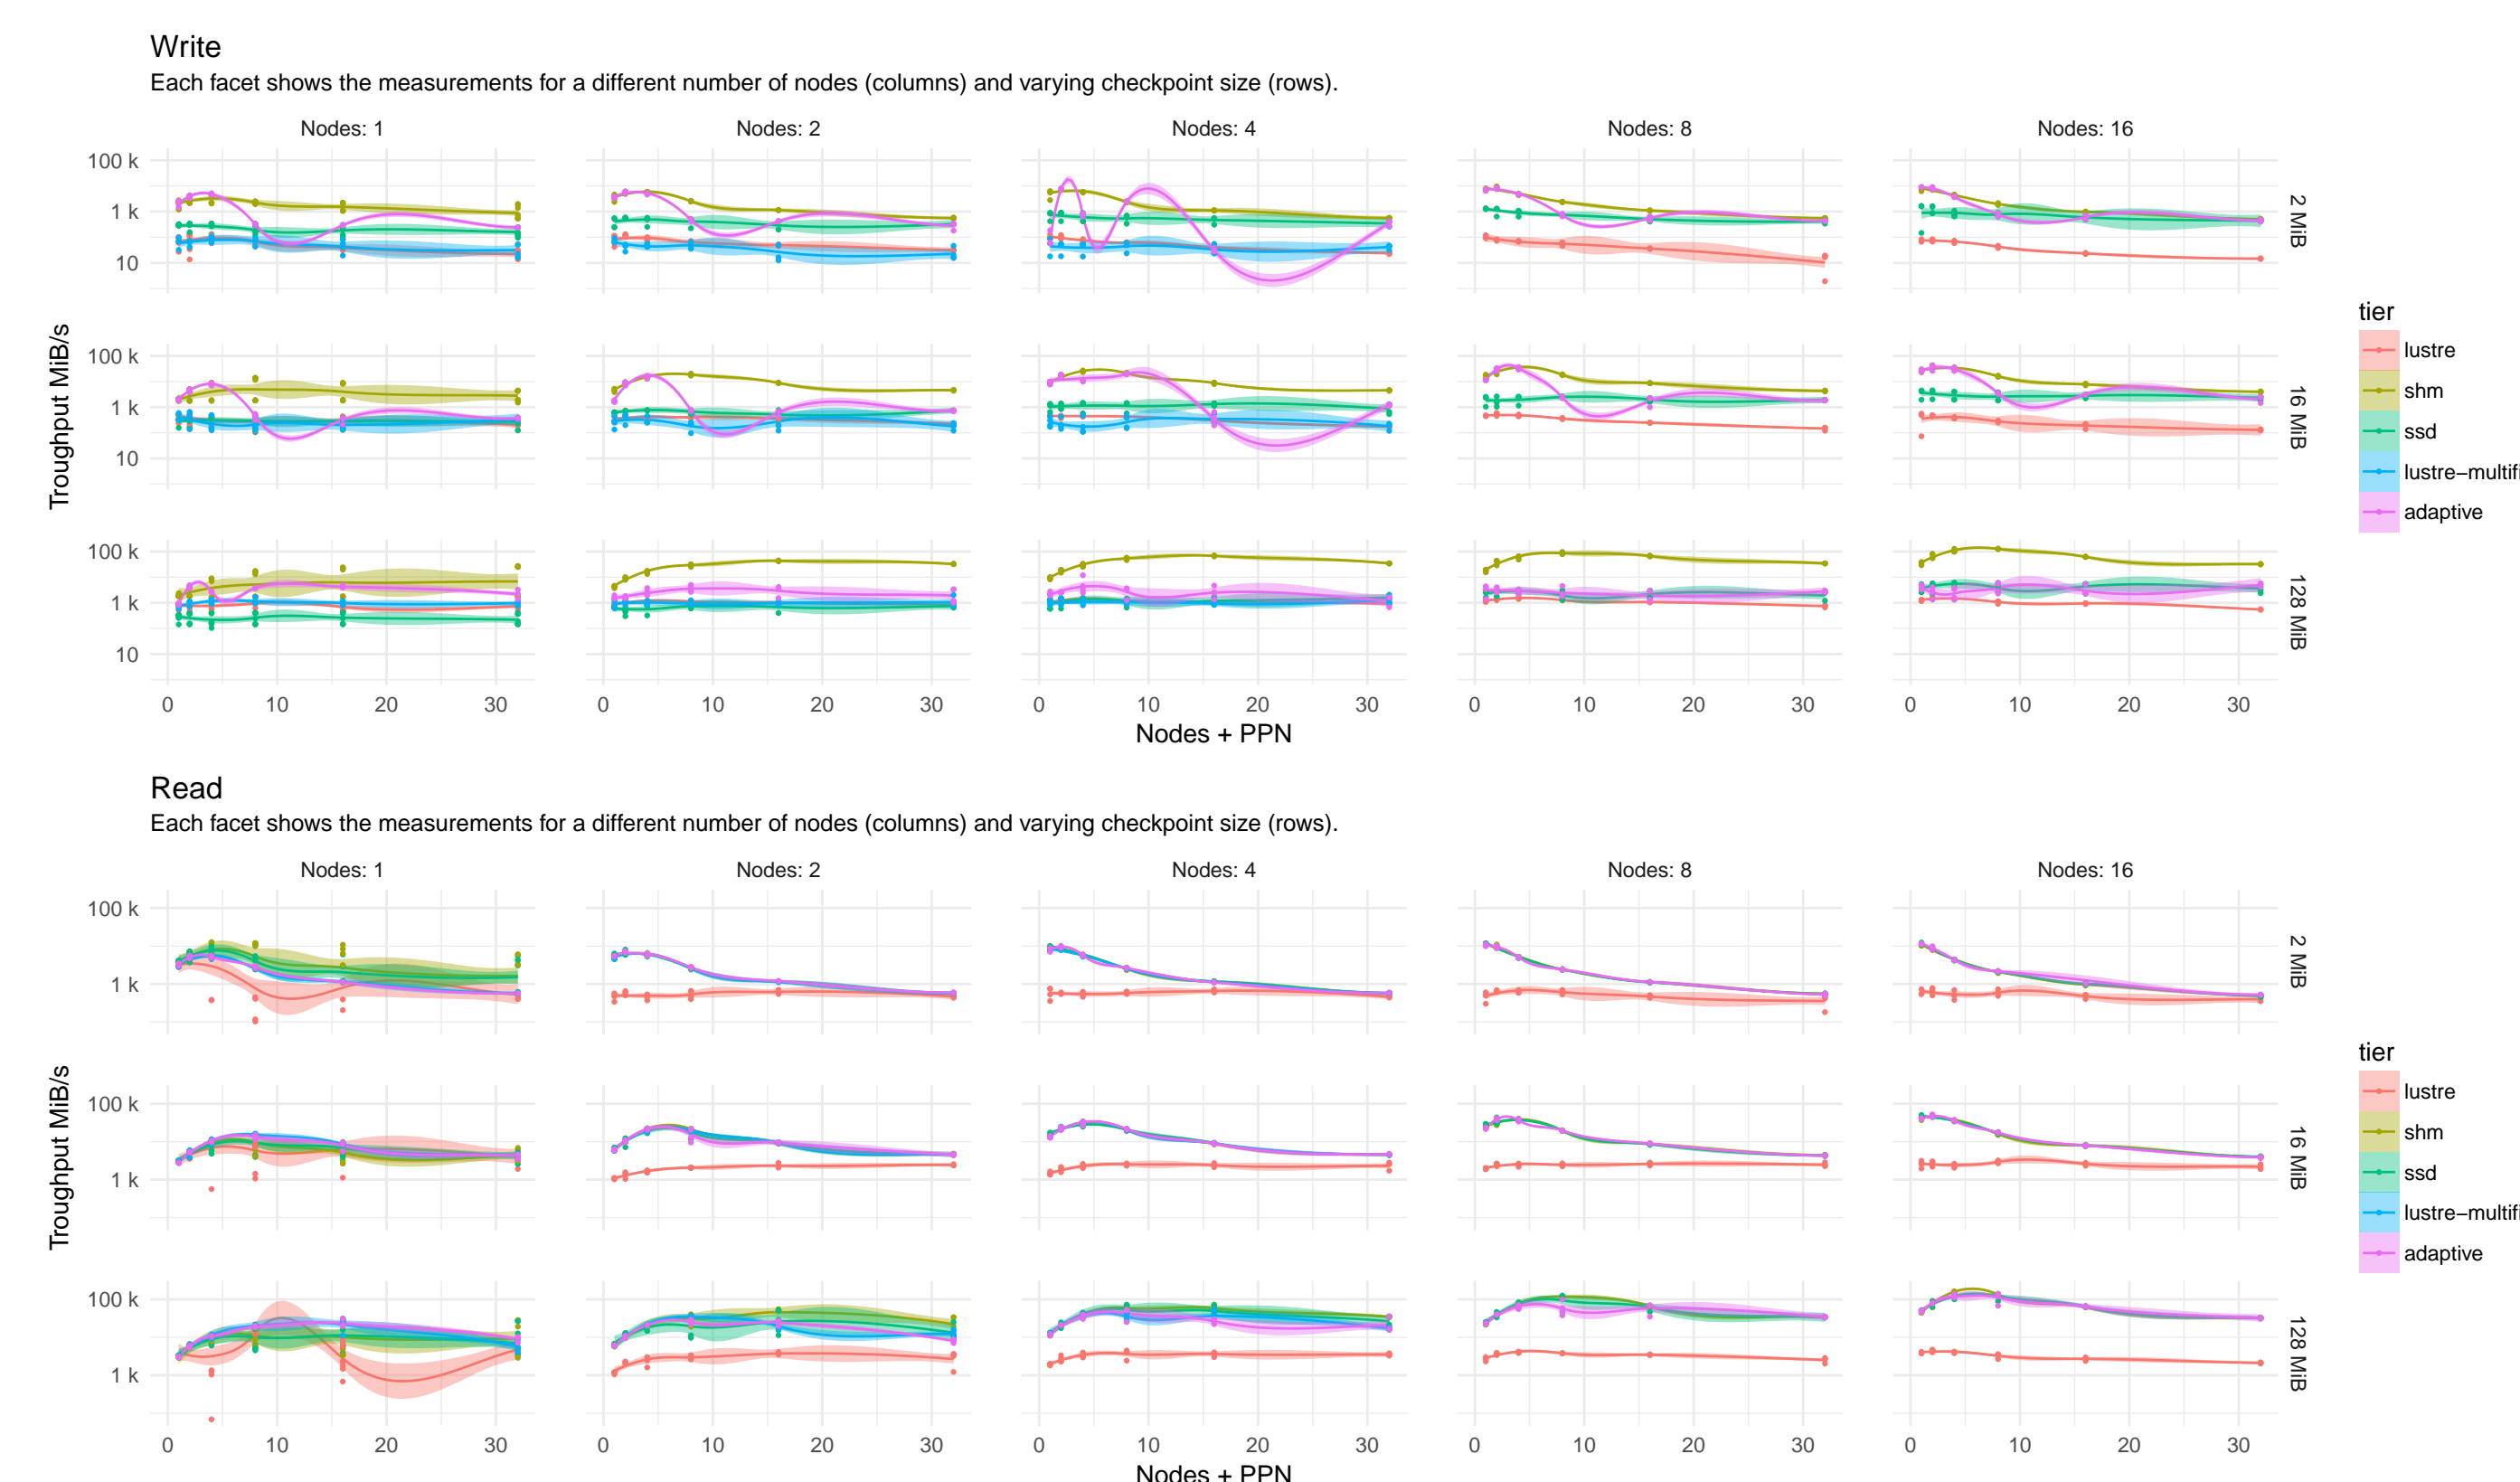
Adaptive tier selection is realized as depicted in the architecture illustration below. The proof of concept decision component accounts for runtime information made available by SLURM, MPI and HDF5: 1) domain decomposition and 2) the domain description of a dataset. The tier selection policies are based on benchmark measurements obtained at an earlier time.



Performance Evaluation

The following plots show throughput of each tier for READ and WRITE in comparison to the performance when a VOL plugin that adaptively selects the most appropriate tier which is not necessarily the fastest:

- Shared memory: Small random I/O and in expectation of burst buffers.
- Local SSDs: For medium random I/O not shared with other nodes.
- Parallel file system: When performing large sequential file I/O.



Notice the flexibility achieved with policies e.g., as measurable for {Nodes=4, Size=2MiB}. Read cases for local storage must adopt the same policy used for writing.



A main goal of the Centre of Excellence in Simulation of Weather and Climate in Europe (ESiWACE) is to improve efficiency and productivity of numerical weather and climate simulations on high-performance computing platforms by supporting the end-to-end workflow of global Earth system modeling in HPC environments. Part of the project is the development of a middleware for Earth system data featuring:

- Access to shared data with different APIs
 - NetCDF4, HDF5 or GRIB
- Data layouts optimized for data centers
 - Advanced data placement optimizing for cost and performance
 - Support for different backends: object storage, file systems


Summary and Future Work

Adaptive tier selection promises to be a viable approach for performance optimization of I/O performance. As storage systems become more heterogeneous in the wake of burst buffers and non-volatile memory, I/O middleware can help to avoid exposing unnecessary complexity to users. As a result, in many cases no changes to an application are necessary. In future work, the decision component should automatically extract tier selection rules from benchmark measurements and observed access patterns. The integration of various storage tiers is continued as part of the ESiWACE project. In particular, the following backends deserve further exploration:

- Object storage mappings for short term storage of working sets
- Tape and other nearline storage for affordable long-term archival

NetCDF Benchmark

NetCDF Performance Benchmark Tool (NetCDF-Bench) was used to recreate a typical checkpoint restart workload. NetCDF-Bench was developed to measure NetCDF performance on devices ranging from notebooks to large HPC systems. It mimics the typical I/O behavior of scientific climate applications and captures the performance on each node/process. Details: <https://github.com/joobog/netcdf-bench>



Acknowledgments

The ESiWACE project received funding from the EU Horizon 2020 research and innovation programme under grant agreement no. 675191. Disclaimer: This material reflects only the authors' view and the EU commission is not responsible for any use that may be made of the information it contains.