

Among the broad variety of challenges that arise from HPC and HPDA workloads, data movement is of paramount importance, especially on coming exascale systems featuring multiple tiers of memory and storage. While the focus has, for years, been primarily on optimizing computations, the importance of improving data handling on such architectures is now well understood. As optimization techniques can be applied at different stages (operating system, run-time system, programming environment, and so on), a middleware providing a uniform and consistent data awareness becomes necessary. In this presentation, we introduce a novel memory- and data-aware middleware called Maestro, designed for data orchestration, by a Horizon 2020 project consortium of major European actors.

Maestro attempts to fully bridge the missing gaps in data-awareness at several levels of the software stack and across the entire memory storage hierarchy. By being accessible to both applications and systems software, an ecosystem of memory- and data-awareness could emerge, where data products can be passed amongst software components, with a degree of confidence in knowing that the receiving party can interpret the appropriate metadata for orchestration purposes. Similarly, the memory awareness aspects of Maestro are fully portable, so an application exploiting one memory system could do the same when run on a new hardware platform. Maestro will enable key technologies in systems software that we believe are implausible in HPC without memory and data-awareness, such as data-centric execution and intelligent workflow optimisation of HPC and HPDA. By advancing the ease-of-use of complex memory hierarchy, Maestro will ease the transition to memory/storage convergence and make the lives of programmers easier. This will help attract a wider community of users to HPC systems. Maestro will seek to establish the middleware as a programming framework used to increase user productivity.

By solving workflow-level technical challenges and by enabling systems software features, Maestro will move European exascale-focused application areas such as numerical weather prediction, climate modelling and materials science closer to making exascale a reality. In particular, Maestro has been successfully tested on an I/O-intensive benchmark of weather-data production and post processing, where performance improvements of up to 2.5 times were obtained. This test case was provided by the European Centre for Medium-Range Weather Forecasts (ECMWF). The dramatic projected growth in the data-production and data-movement rate required by scientific progress is a significant limiting factor on ECMWF's ambition to provide ever more accurate weather forecasts. Maestro offers a credible potential solution to this data congestion. It facilitates restructuring the forecast-data-generation and post-processing parts of ECMWF's operational workflow by minimising unnecessary transfers of data between different levels of the memory and storage hierarchy. It is more careful when managing and controlling the use of the available and upcoming memory and storage technologies – and thus may support and underpin scientific developments in the coming years and decades.