



Contribution ID: 18

Type: **Short talk**

Bringing Elasticity to HPC Storage and Data Services

Thursday 23 March 2023 16:50 (10 minutes)

Elasticity, or the ability to adapt a system to a dynamically changing workload, has been a core feature of Cloud Computing storage since its inception more than two decades ago. In the meantime HPC applications have mostly continued to rely on static parallel file systems to store their data. This picture is now changing as more and more applications adopt custom data services tailored to their needs, including in transit analysis systems, staging areas for coupling, and transient file system aggregating on-compute-node storage capacity. As a result, it is increasingly important to incorporate the concept of elasticity within HPC so that these new data services can dynamically adapt in response to time-varying application requirements.

In this talk, we will present current efforts from the Mochi team to tackle the challenges of data service elasticity. Mochi is an R&d100-awarded collection of building blocks for developing composable HPC data services. It enjoys a growing community of users and contributors, with applications that increasingly need Mochi to support elasticity. We will highlight the different levels at which elasticity can be implemented, from low-level thread-scheduling decisions, to scaling across nodes with data migration. We will show the technical challenges, as well as opportunities from the AI domain to enable self-adapting data services.

JLESC topic

Primary author: DORIER, Matthieu (Argonne National Laboratory)

Presenter: DORIER, Matthieu (Argonne National Laboratory)

Session Classification: Short Talks on Workflows, I/O and Frameworks

Track Classification: I/O, storage and in-situ processing