

ADMIRE: Adaptive multi-tier intelligent data manager for Exascale

Description

One of the main drivers for building exascale HPC systems today is the growing need to process extremely large data sets. However, the flat storage hierarchies found in classic HPC architectures no longer satisfy the performance requirements of data-processing applications. Uncoordinated file access in combination with limited bandwidth make the centralised backend parallel file system a serious bottleneck. At the same time, emerging multi-tier storage hierarchies come with the potential to remove this barrier. However, maximising performance still requires careful control to avoid congestion and to balance computational and storage performance. Unfortunately, appropriate interfaces and policies for managing such an enhanced I/O stack are still lacking.

The main objective of the ADMIRE project is to establish this control by creating an active I/O stack that dynamically adjusts computation and storage requirements through intelligent global coordination, malleability of computation and I/O, and the scheduling of storage resources along all levels of the storage hierarchy. To achieve this, we will develop a software-defined framework based on the principles of scalable monitoring and control, separated control and data paths, and the orchestration of key system components and applications through embedded control points. Our software-only solution will allow for the throughput of HPC systems and the substantial increase of the performance of individual applications and for the consequent decrease in energy consumption by taking advantage of fast and power efficient node-local storage tiers using novel, European ad-hoc storage systems and in-transit/in-situ processing facilities. Furthermore, our enhanced I/O stack will offer quality-of-service (QoS) and resilience. An integrated and operational prototype will be validated with several use cases from various domains, including climate/weather, life sciences, physics, remote sensing, and deep learning.

Proyecto PCI2021-121952 financiado por MICIU/AEI /10.13039/501100011033 y por la Unión Europea NextGenerationEU/PRTR

Proyecto PCI2021-121952 financiado por:

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on 22 des 2024 - 06:04): <https://www.bsc.es/ca/research-and-development/projects/admire-adaptive-multi-tier-intelligent-data-manager-exascale>