



Contribution ID: 4

Type: **Presentation**

RobinHood: storing and querying a filesystem's metadata for fast access

Friday, 13 March 2026 09:00 (25 minutes)

As supercomputers are becoming faster and faster, so does their data output. Since the regularly accessed data must be stored and available quickly to users, it is important to put it on fast storage systems. However, these tend to have a low capacity, meaning we must be able to choose the data which should remain on those types of storage systems, and which can be placed on slower but more capacitive systems. As such, it is important to be able to accurately know the state of a filesystem at any point, but using the conventional means provided by the operating system for this, for instance to do filesystem traversals, can be time consuming if done regularly. Moreover, these operations impose a heavy load on the filesystem, making it slower. To counter these problems, we created a suite of tools called RobinHood that aims to mirror a filesystem in a database, and use the latter to define policies that will manage data placement according to their usage.

Author: VALERI, Yoann (CEA)

Presenter: VALERI, Yoann (CEA)

Session Classification: Session C

Track Classification: Session II: Status and Development Trends