



Contribution ID: 18

Type: **not specified**

## **Data Storage Institute: Speeding-up large-scale storage with non-volatile memory**

*Wednesday, 10 June 2015 10:40 (15 minutes)*

CERN EOS distributed, scalable, fault-tolerant filesystem serves data to 100s-1000s of clients. To speed-up the service, it keeps and updates a large and fast data Catalog in RAM that can reach a footprint of 100+ GBs. When a EOS server node crashes the Master node must rebuild a new coherent Catalog in memory from disk logs, a long operation that disrupts the activity of the clients.

We propose to design a persistent, version of the EOS Catalog to be stored in the new Non-Volatile Memory that will remain consistent after faults. Upon restart, the EOS server process will find the Catalog in NVM and immediately resume serving all the clients, skipping the slow reconstructions from disk-based logs altogether.

**Presenters:** YONG, Khai Leong (DSI); RUOCCO, Sergio (DSI)