

Title:

Puddlestore and Distributed Locking with Zookeeper

Abstract:

PuddleStore is a distributed file system designed to provide continuous access to persistent information. Users can create files and directories through the use of Zookeeper (a distributed coordination service hosted by Apache) and a Distributed Object Location and Retrieval (DOLR) service called Tapestry for storing key-value pairs. The data blocks of the file system are stored across the nodes of the Tapestry network, while the file metadata and directory structure are stored in Zookeeper. In order to handle conflicting updates, we implemented coarse grained file locking using Zookeeper to prevent file corruption and to help maintain strong consistency.