

Puddlestore is a simple distributed file system, built on top of Zookeeper, a coordination system for distributed systems, and tapestry, our own distributed object location and retrieval system. Our implementation allows users to access data through multiple hosts using any number of remote clients, and allows for the data that they store to be replicated throughout the system. However, with multiple clients trying to access and modify data, it becomes essential that a secure file locking protocol is implemented. We created Zookeeper nodes to represent locks for each client, locking upon the opening of any file, and unlocking whenever that file is closed. By making these nodes ephemeral, we also prevented recoverable errors in the event that the servers crash. Overall, these changes allow our Puddlestore project to become more scalable and more consistent as more clients start to use the database.