

PuddleStore: A Distributed File System with Synchronous Locking

Ryan Hecht, with Peter Hahn

Capstone under CS1380 Distributed Systems with Theo Benson

Our Capstone Project revolved around creating PuddleStore, a distributed file system based on the Berkeley Paper, OceanStore in 2011. Having a file system that allows for distant clients to access, read and write to a single system is something that has many applications in the sharing of information in our modern age. That being said, it comes with several challenges, making sure that our algorithms are fault tolerant, resistant to failing servers and data replicas. We address this by using a Distributed Object Location and Retrieval system (Tapestry) to deal with the file blocks specifically, as well as Zookeeper to maintain consistent information about the metadata associated with a given file. In addition, in order to make sure that multiple accesses by different clients would not interfere with each other, we implemented a distributed file locking system to regulate the writing of all files.