

# Apache ZooKeeper

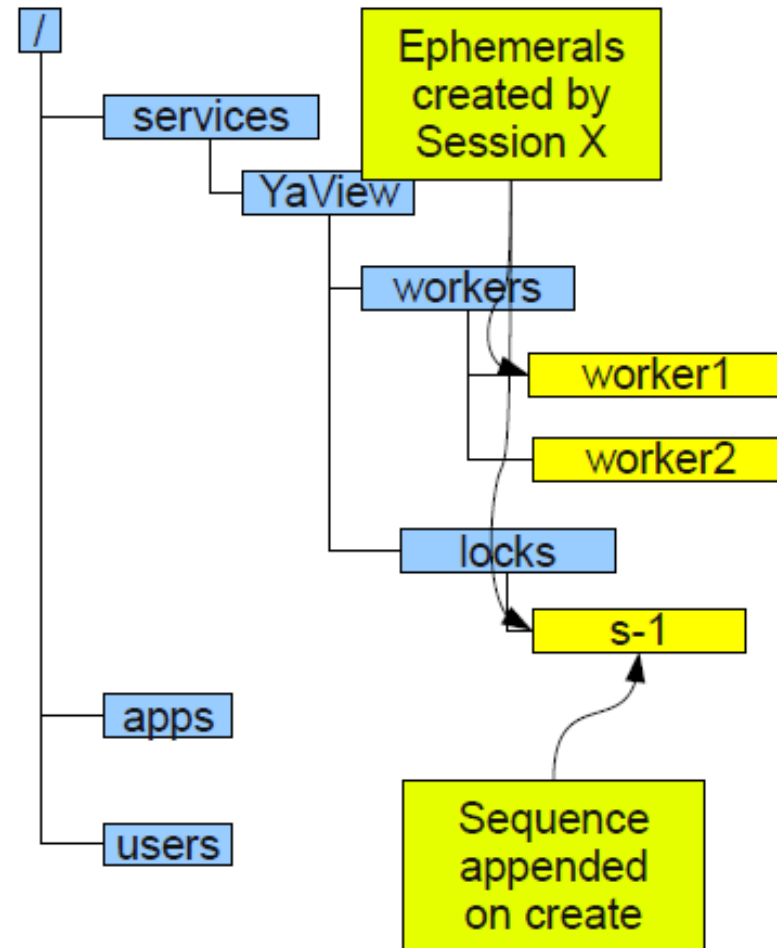
Presented by Rui Zhang

# What is ZooKeeper

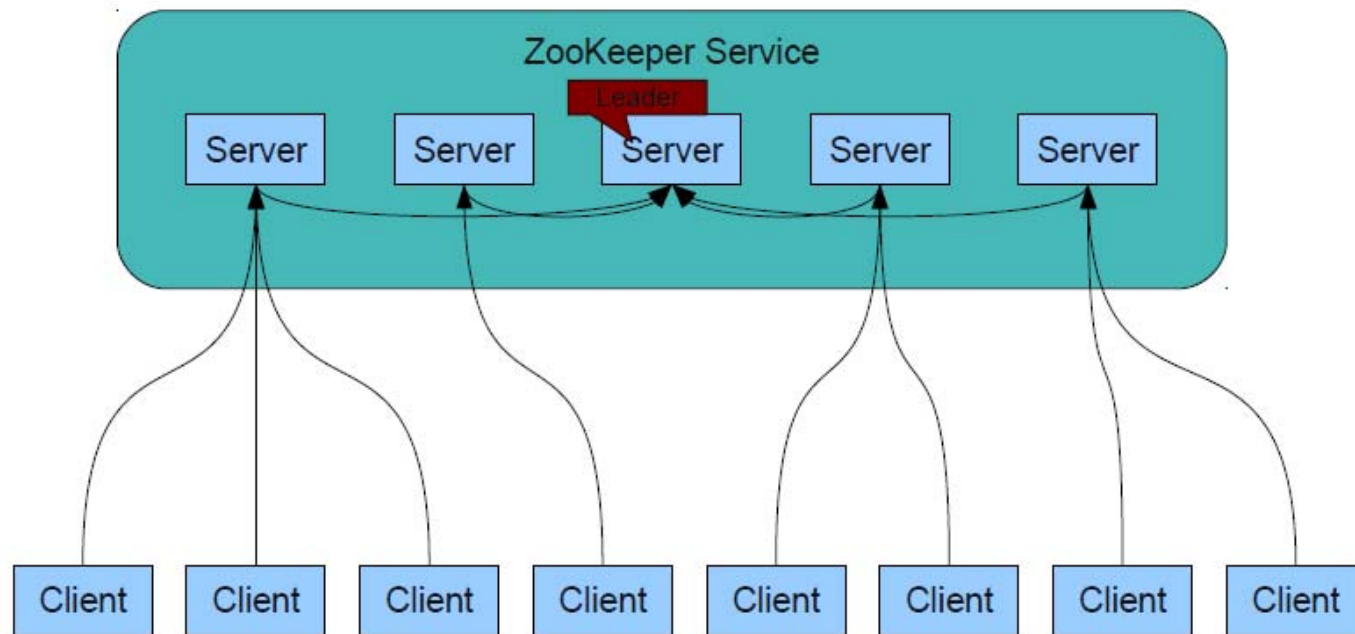
- A general wait-free coordination service
  - membership, leader election, lock...
- Provide order guarantees for operations
  - Linearizable writes
  - FIFO client order
- Change notifications

# Data Model

- Z-Node
- Hierarchical name spaces
- Flags
  - Ephemeral
  - Sequential

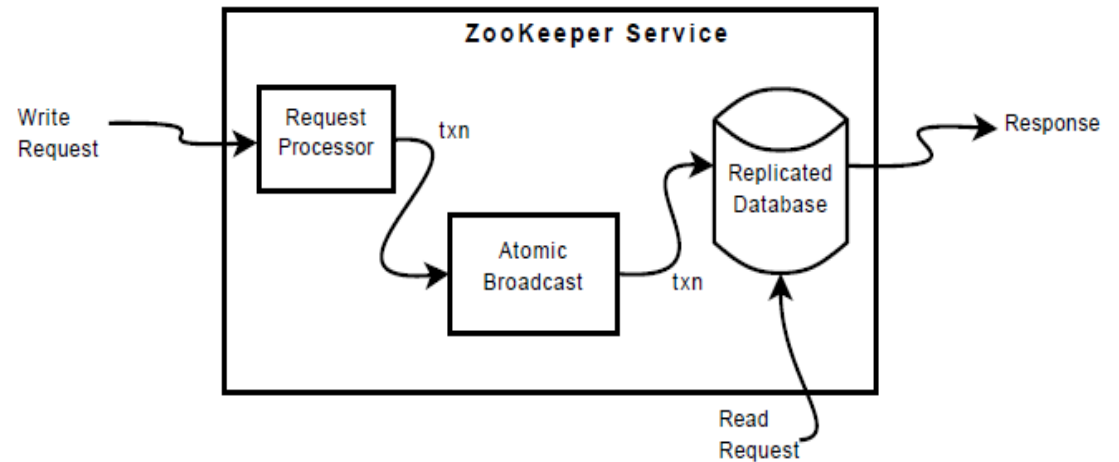


# Implementation



# Implementation

- Request Processor
- Atomic Broadcast
  - Zab Protocol
- Replicated Database

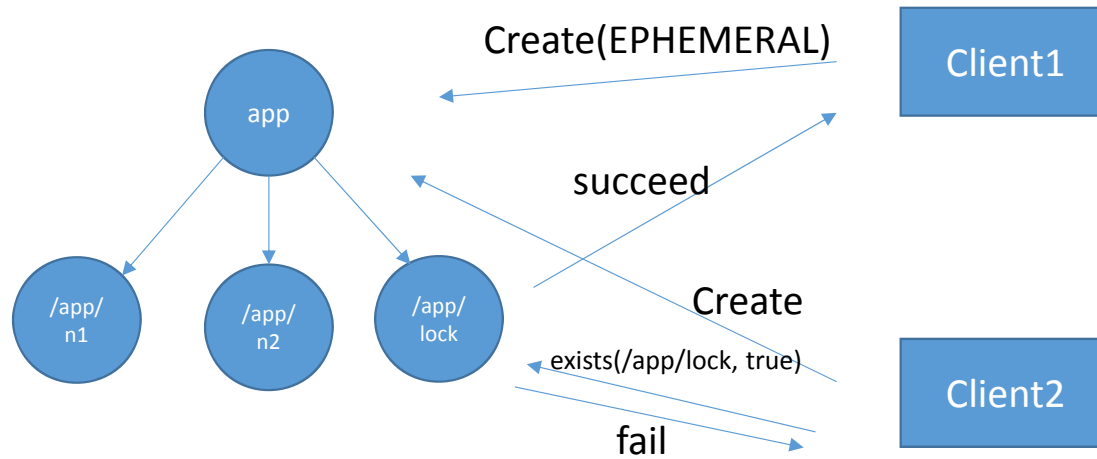


# APIs

- Create(path, data, flags)
- Delete(path, version)
- Exists(path, watch)
- getData(path, watch)
- setData(path, data, version)
- getChildren(path, watch)
- Sync(path)

# Simple Example

- Lock



# Reference

- Hunt, Patrick, et al. "ZooKeeper: wait-free coordination for internet-scale systems." *Proceedings of the 2010 USENIX conference on USENIX annual technical conference*. Vol. 8. 2010.



Thanks!