

servers are like cattle



ZooKeeper

...

Because coordinating distributed systems is a zoo

ZooKeeper




A distributed coordination service for distributed applications

ZooKeeper is...

A Distributed Filesystem

- Distributed processes coordinate via a shared hierarchical namespace
- Information is stored in *znodes*
- Data is kept in-memory for high throughput and low latency




Basic API:

- Create
- Delete
- Exists
- Get Data
- Set Data
- Get Children
- Sync

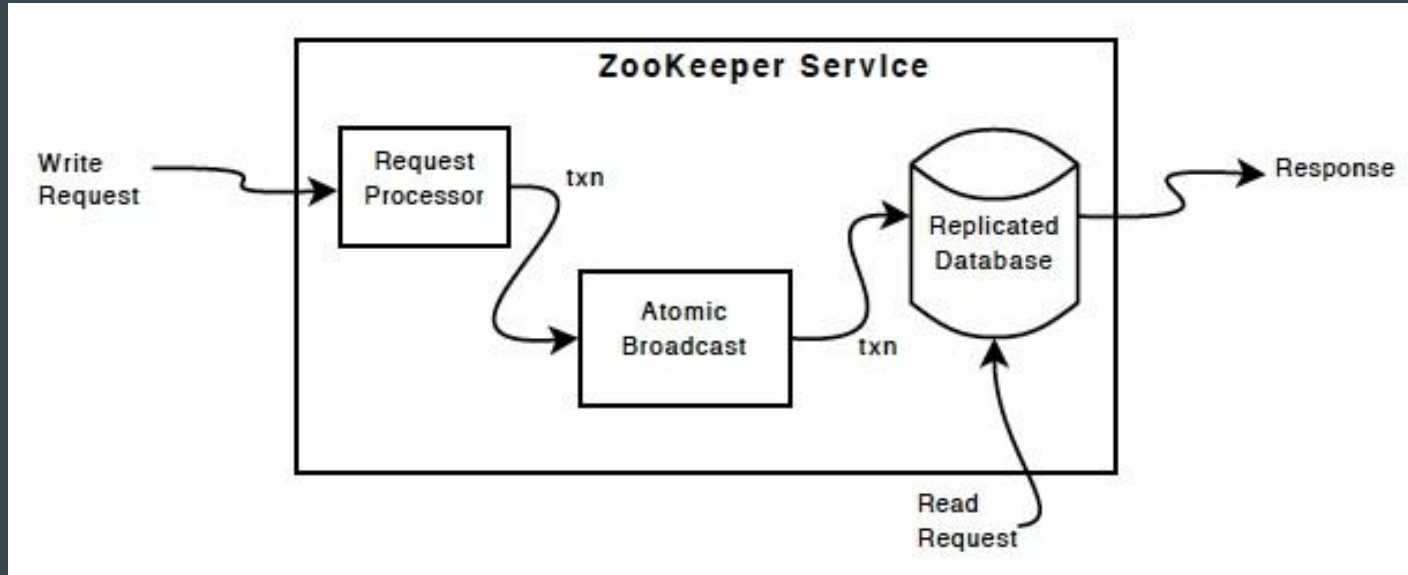
ZooKeeper is...

Replicated

- The ZooKeeper service is intended to be replicated across a set of hosts
- Each server caches system state and transaction logs for durability



ZooKeeper Components



```
ls /  
time = 3 msec  
/: rc = 0  
      zookeeper  
time = 3 msec  
_
```

```
create /foo
Creating [/foo] node
Watcher CHILD_EVENT state = CONNECTED_STATE for path /
[/foo]: rc = 0
        name = /foo

ls /
time = 10 msec
/: rc = 0
    foo
    zookeeper
time = 10 msec
```

```
get /foo
```

```
time = 47988 msec
```

```
/foo: rc = 0
```

```
value_len = 3
```

```
new
```

```
Stat:
```

```
ctime = Wed Sep 28 13:56:10 2016
```

```
czxid=5
```

```
mtime=Wed Sep 28 13:56:10 2016
```

```
mzxid=5
```

```
version=0          aversion=0
```

```
ephemeralOwner = 0
```

```
set /foo bar
Watcher CHANGED_EVENT state = CONNECTED_STATE for path /foo

get /foo
time = 703084 msec
/foo: rc = 0
  value_len = 3
bar
Stat:
  ctime = Wed Sep 28 13:56:10 2016
  czxid=5
  mtime=Wed Sep 28 14:08:09 2016
  mzxid=7
  version=1      aversion=0
  ephemeralOwner = 0
```

Use in Distributed OpenNetVM

- Coordinate instances of the manager
- Contain a mapping of Services to machines running that service
- Keep a mapping of MAC addresses for manager instances
- Allow individual instances to make decisions affecting global system state

< I'm a server >

