




ZooKeeper Tutorial

Part 3

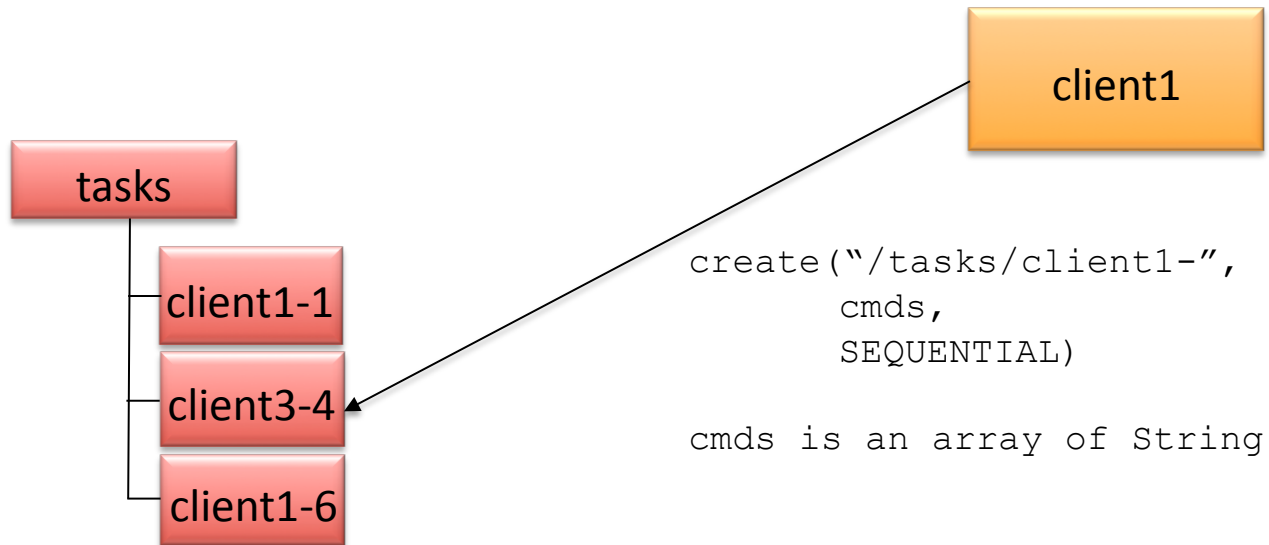
How it really works

Master/Worker System

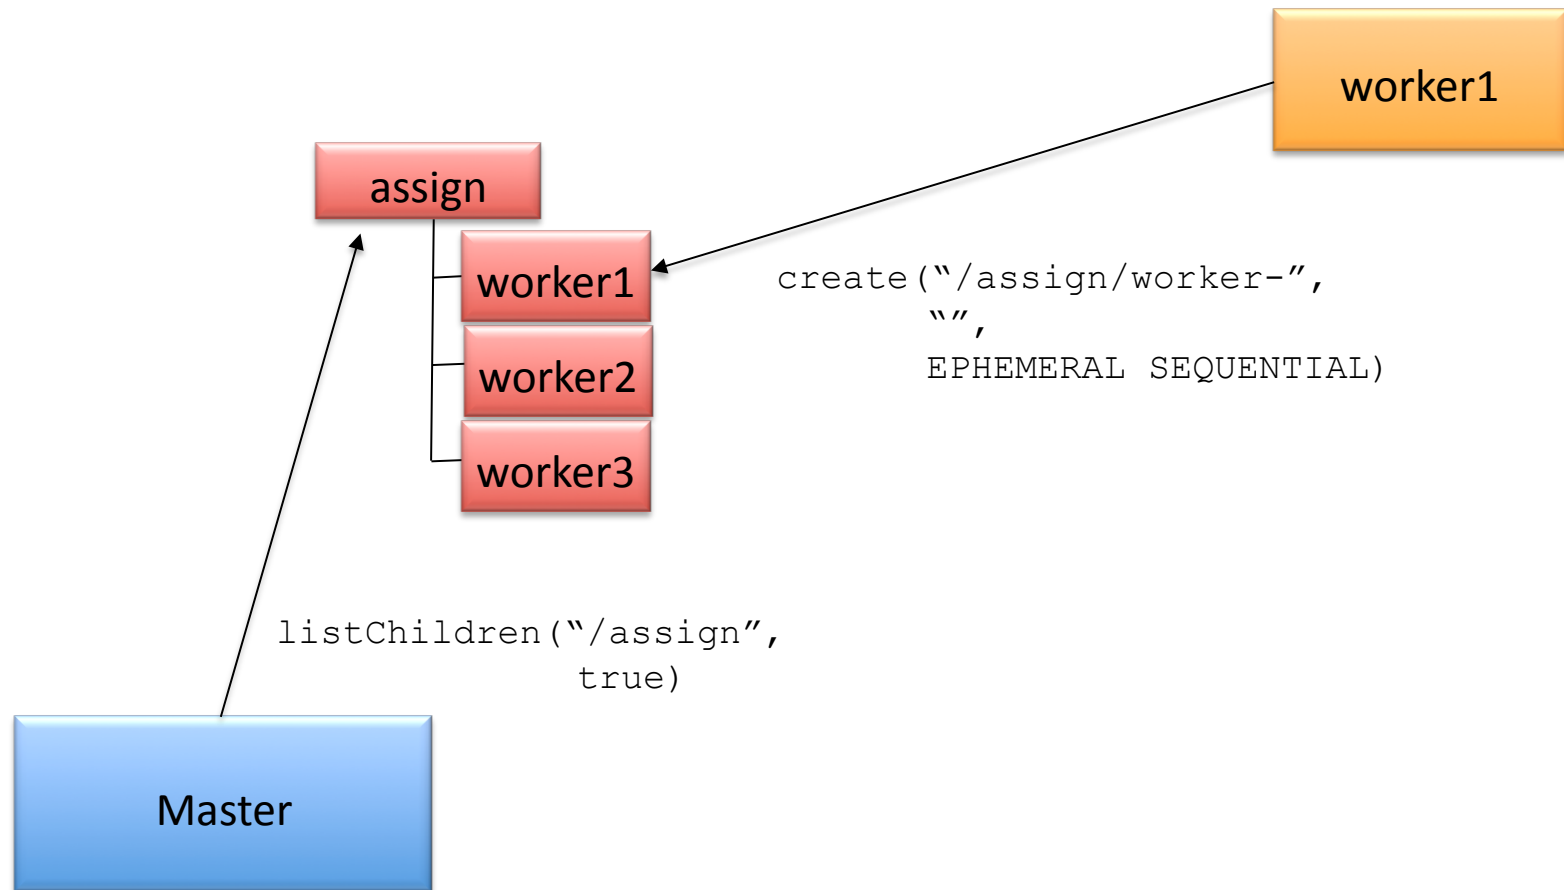
- Clients
 - Queue tasks to be executed
 - Monitor the tasks
- Masters
 - Assign tasks to workers
- Workers
 - Get tasks from the master
 - Execute tasks




Task Queue



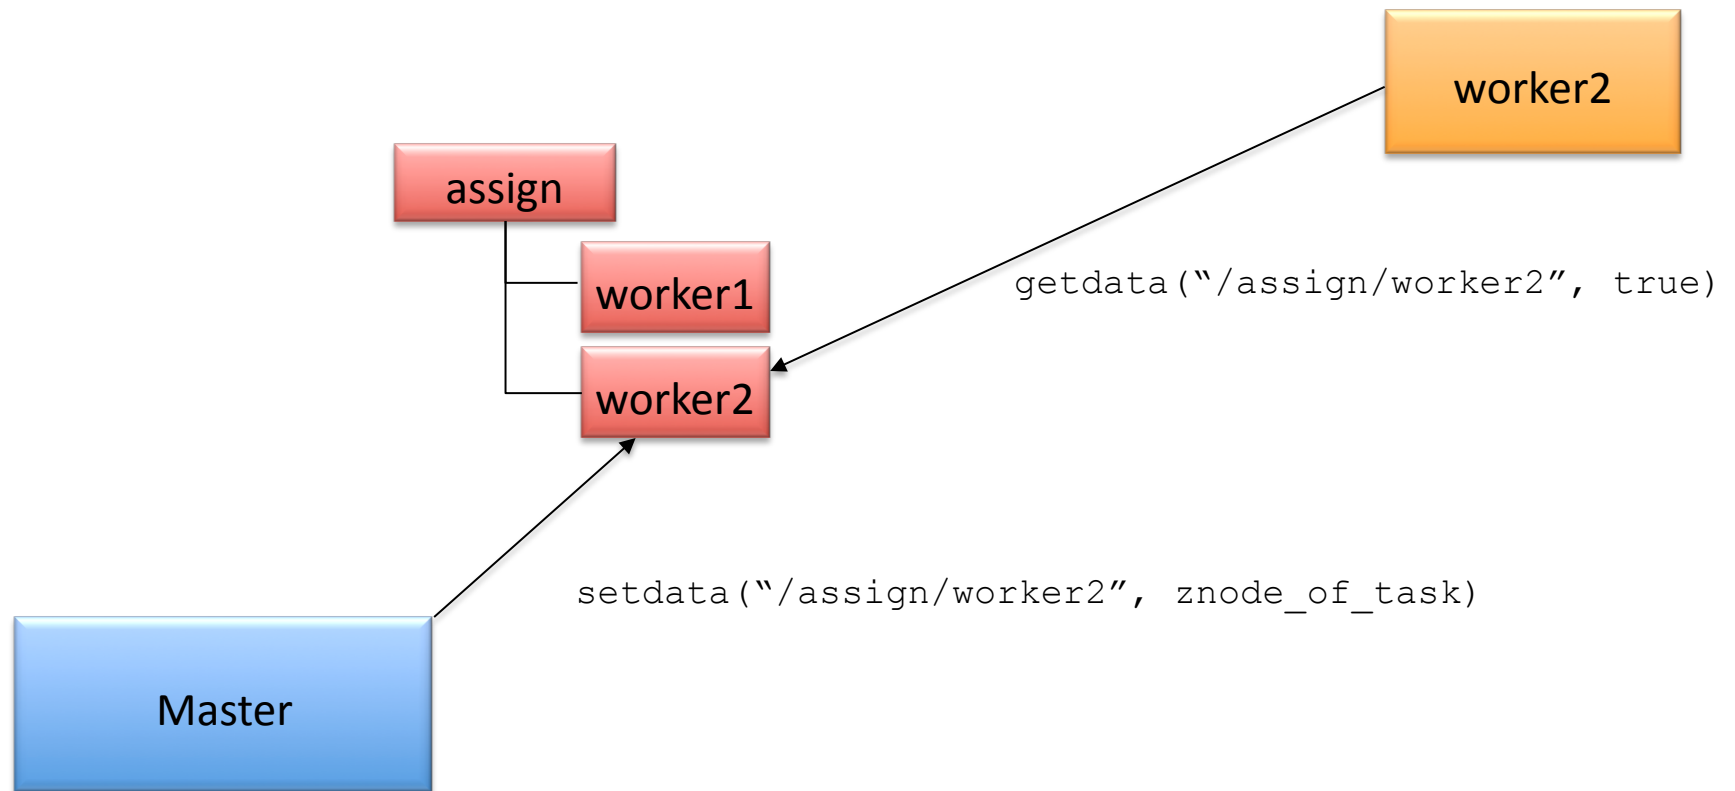
Group Membership



Leader Election



Configuration



Connecting to ZooKeeper

- Everyone has their own ZooKeeper address and auth info.

- Try connecting to ZooKeeper with the CLI.

```
java -jar zookeeper-3.3.2-fatjar.jar client zkaddr
```

- Use addAuth command to authenticate
- Try out some commands
 - Create znodes for `/servers`, `/tasks`, `/assign`



Worker Processing

- Create a session
- Create the “worker” ephemeral znode
- Watch for the assign znode
- Deal with the watches
 - Processing the assignment
 - Update status in the task
 - Delete assignment znode when finished
 - What do to with SessionExpired





Code on your own or follow
together

Client Processing

- Create a session
- Create a task as a child of the `/tasks` znode
- Watch the status child of the `/tasks` znode





Code on your own or follow
together

Master Processing

- Create a session
- Do leader election using master znode
- Watch the worker list
- Watch the task queue
- Watch the assignment queue
- Deal with the watches
 - Deal with workers coming and going
 - Assign new tasks
 - Watch for completions





Code on your own or follow
together

Give it a try...

- Start up the master
- Start up a worker
- Try submitting a command
- Queue up a bunch of sleep 100
- Add more workers
- Try killing a worker
- Try killing the master. Did take over work?

