

Pool Migration

Tom Sollers & Bill Scales, IBM



Good morning everyone my name is Tom Sollers, Im a graduate who just joined IBM last September and since then I have been working on pool migration alongside Bill Scales (gesture to bill) and others





The Problem

The Problem

Workloads change

Clusters grow

Requirements shift

Pools are created years ago with fixed properties

Data stays locked in outdated configurations

The current solution is:

Manual,

Slow,

Requires downtime,

And increases the chance to introduce errors


What is Pool Migration?




What is Pool Migration?




What is Pool Migration?




What is Pool Migration?



What is Pool Migration?




What is Pool Migration?



The background of the slide is a light gray silhouette of the London skyline, including the Tower Bridge, the Gherkin, Big Ben, and St. Paul's Cathedral. A large, semi-transparent version of the red circular icon from the logo is overlaid on the skyline. At the bottom, there is a decorative wavy line in shades of red and white.

Why do we want Pool Migration?

Why do we want Pool Migration?




Pool Migration allows you to non-disruptively migrate all the objects from one pool to a new pool
This enables you to change the fixed properties of a pool

How does Pool Migration work



How does Pool Migration work?



How does Pool Migration work?



```
ceph osd pool create blue_train --migrate-from-pool red_train
```


```
ceph status
```

```
cluster:  
  id:      326a7d6d-d044-4c0e-b553-7c19ea97ead8  
  health: HEALTH_WARN  
          Pool migration: 196/998 objects migrated (19.639%), 1 pool migrating, 32 pgs migrating
```

How does Pool Migration work?

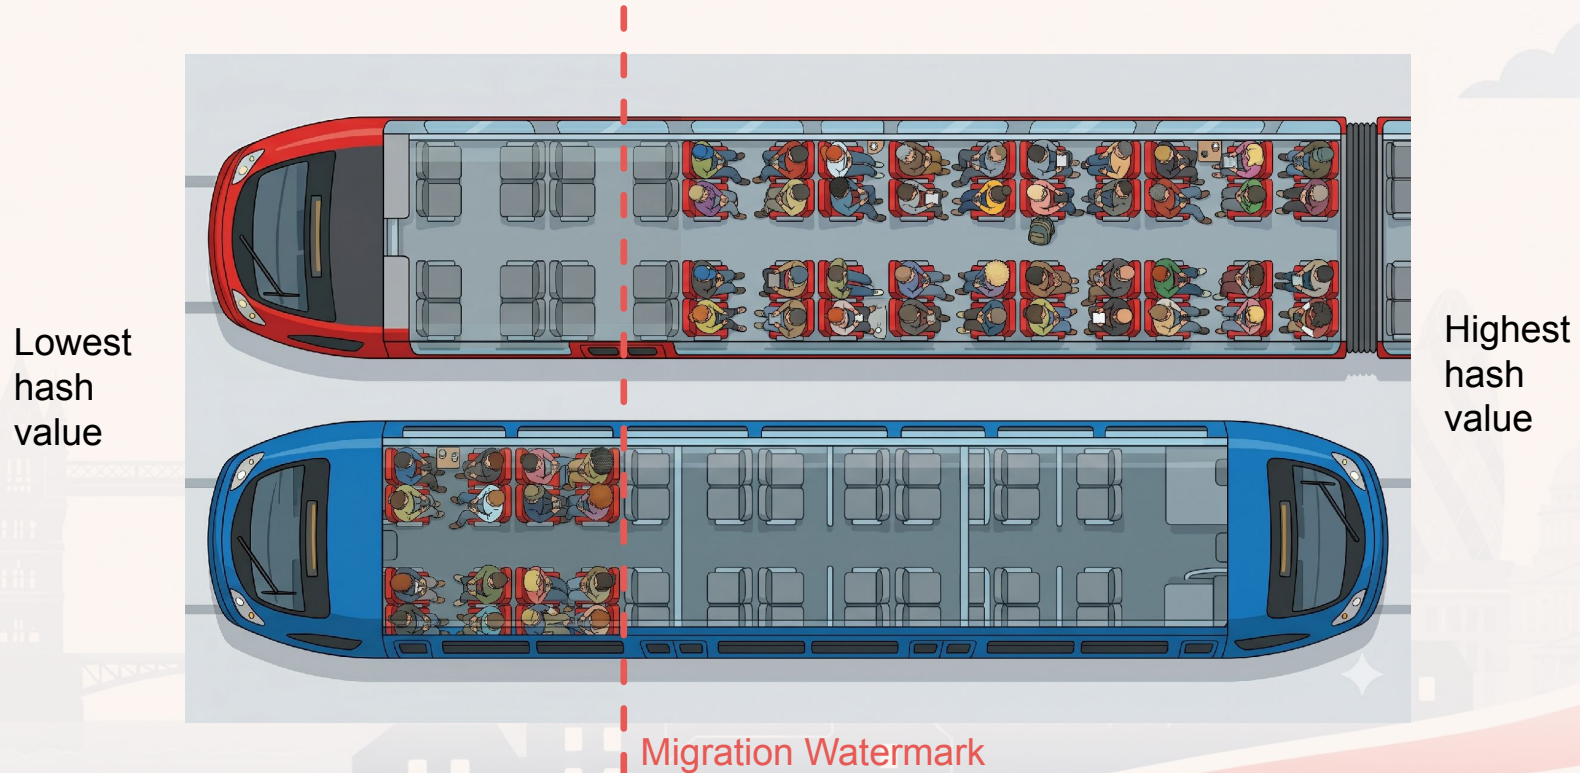
```
ceph osd pool create blue_train --migrate-from-pool red_train
```

Lowest
hash
value




Highest
hash
value


How does Pool Migration work?



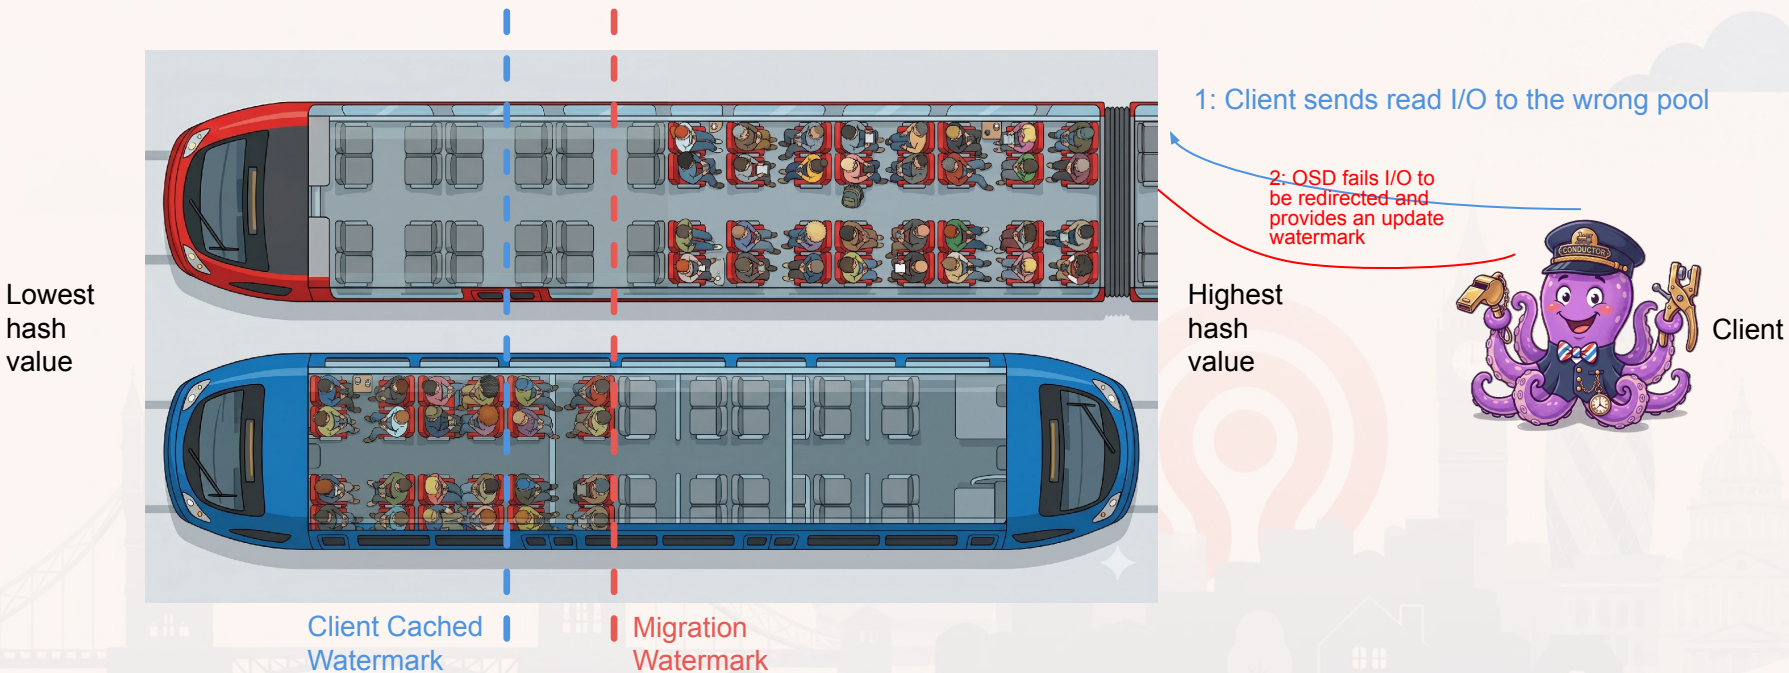
How does Pool Migration work?



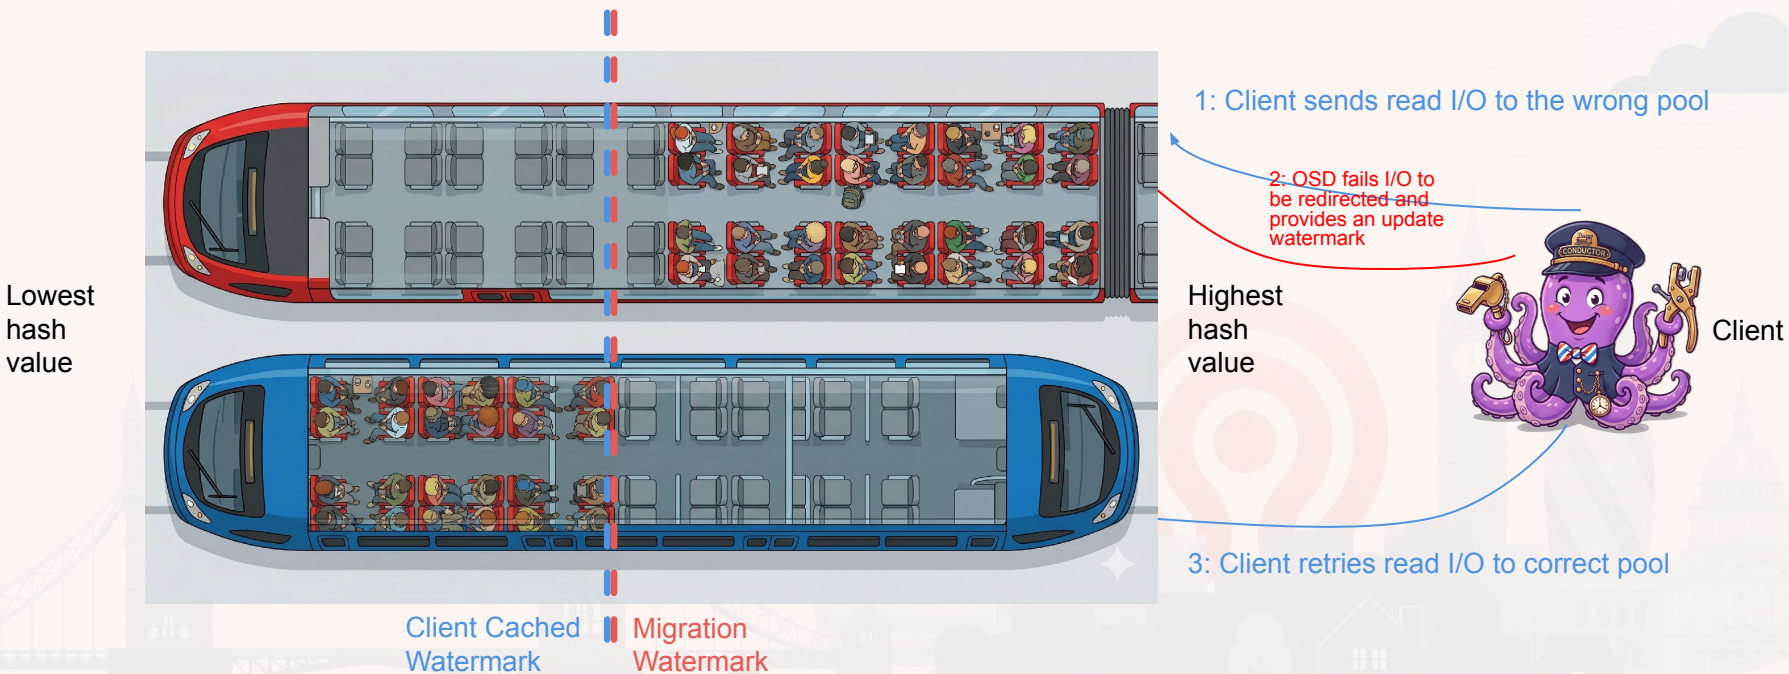
How does Pool Migration work?



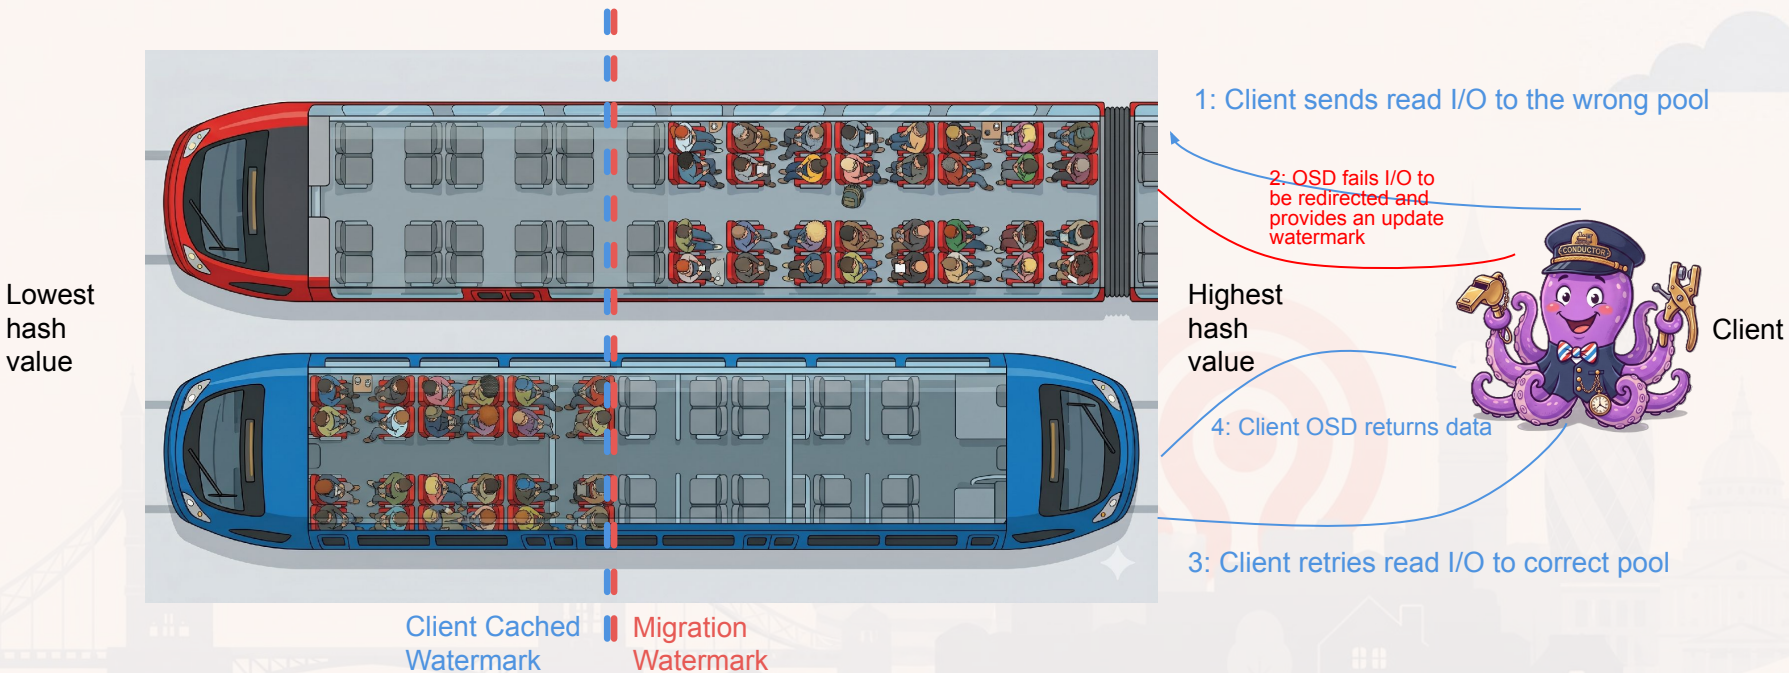
How does Pool Migration work?



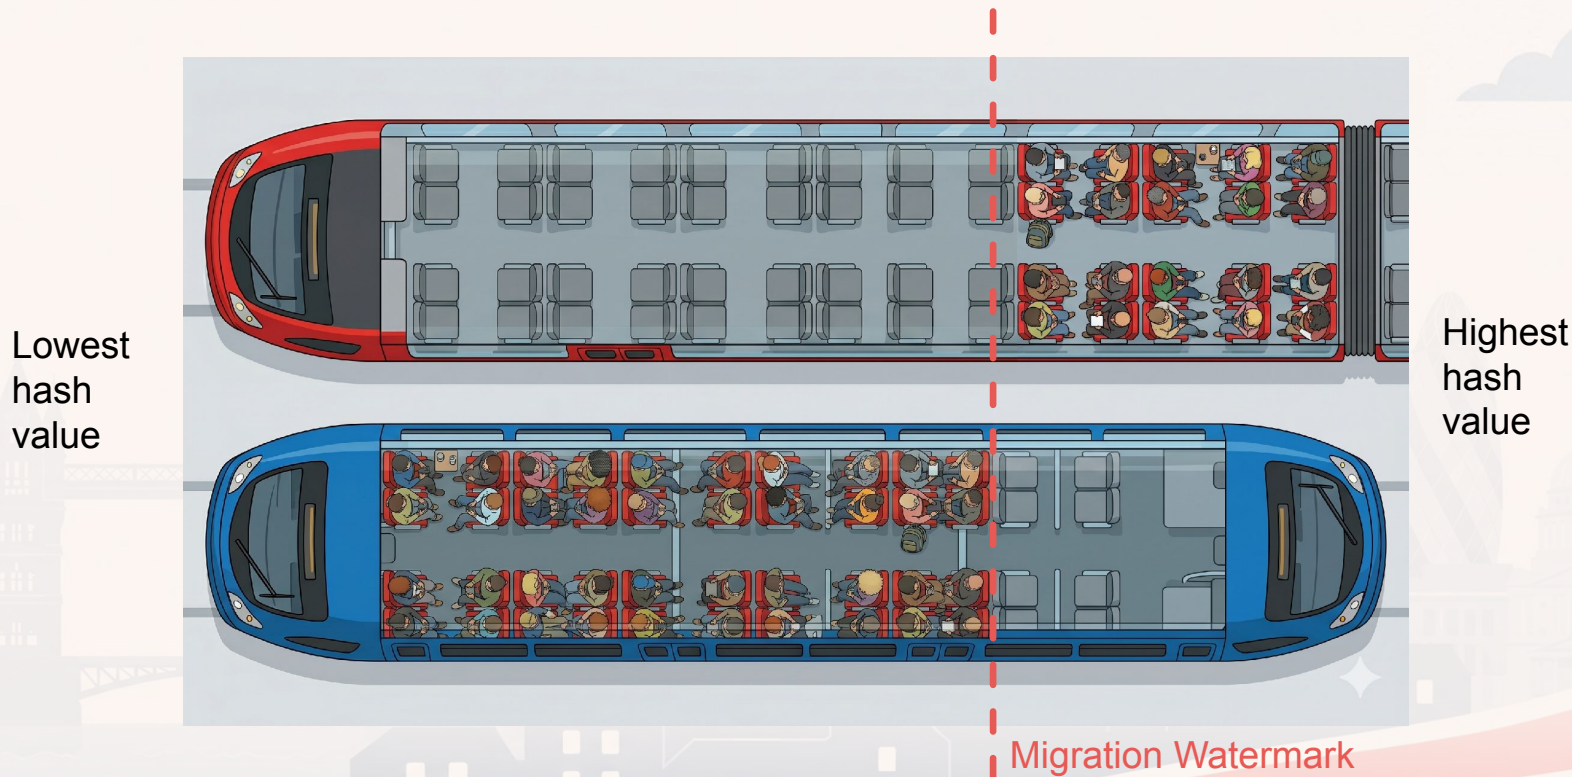
How does Pool Migration work?



How does Pool Migration work?




How does Pool Migration work?




How does Pool Migration work?

Lowest
hash
value




Highest
hash
value

How does Pool Migration work?




How does Pool Migration work?




How does Pool Migration work?

Multiple Placement Groups

Target pool has more placement groups



Target pool has less placement groups





Current Status

Current Status



Supported Actions consist of:

- Object, attribute and omap migration
- Snapshots – no restrictions during a migration
- Watch/Notify operations
- Object listing

Limitations:

- No cache tiering support
- No OSD deduplication
- Crimson support not yet implemented
- Some validity checks for migrations not yet implemented

