

# Review last meeting



- Imbalance Factor
  - Part1: Normalized CoV
  - Part2: Sigmoid(Urgency)
- Monitoring Ceph
  - Monitoring migration
  - Construct Namespace
- Mixed Workload
  - Clients number (75 is nice)
  - Rule out AI

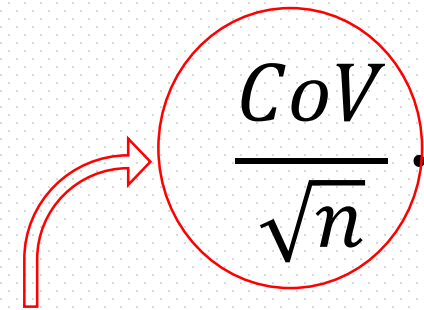
# Definition of IF

$$\frac{\sqrt{\sum_{i=1}^n (l_i - \bar{l})^2 / (n - 1)}}{\sqrt{n} \cdot \sum_{i=1}^n l_i / n} \cdot \frac{1}{1 + e^{5 - 10 \cdot \frac{\text{MAX}_l}{\bar{l} \cdot \text{Capacity}}} \cdot 100\%$$

# Definition of IF

$$\frac{\sqrt{\sum_{i=1}^n (l_i - \bar{l})^2 / (n - 1)}}{\sqrt{n} \cdot \sum_{i=1}^n l_i / n} \cdot \frac{1}{1 + e^{5 - 10 \cdot \frac{MAX_l}{\bar{l}} / Capacity}} \cdot 100\%$$




$$\frac{CoV}{\sqrt{n}} \cdot \frac{1}{1 + e^{5 - 10u}} \cdot 100\%$$

Is imbalance exist in the cluster?

# Definition of IF

$$\frac{\sqrt{\sum_{i=1}^n (l_i - \bar{l})^2 / (n - 1)}}{\sqrt{n} \cdot \sum_{i=1}^n l_i / n} \cdot \frac{1}{1 + e^{5 - 10 \cdot \frac{MAX_l}{\bar{l}} / Capacity}} \cdot 100\%$$

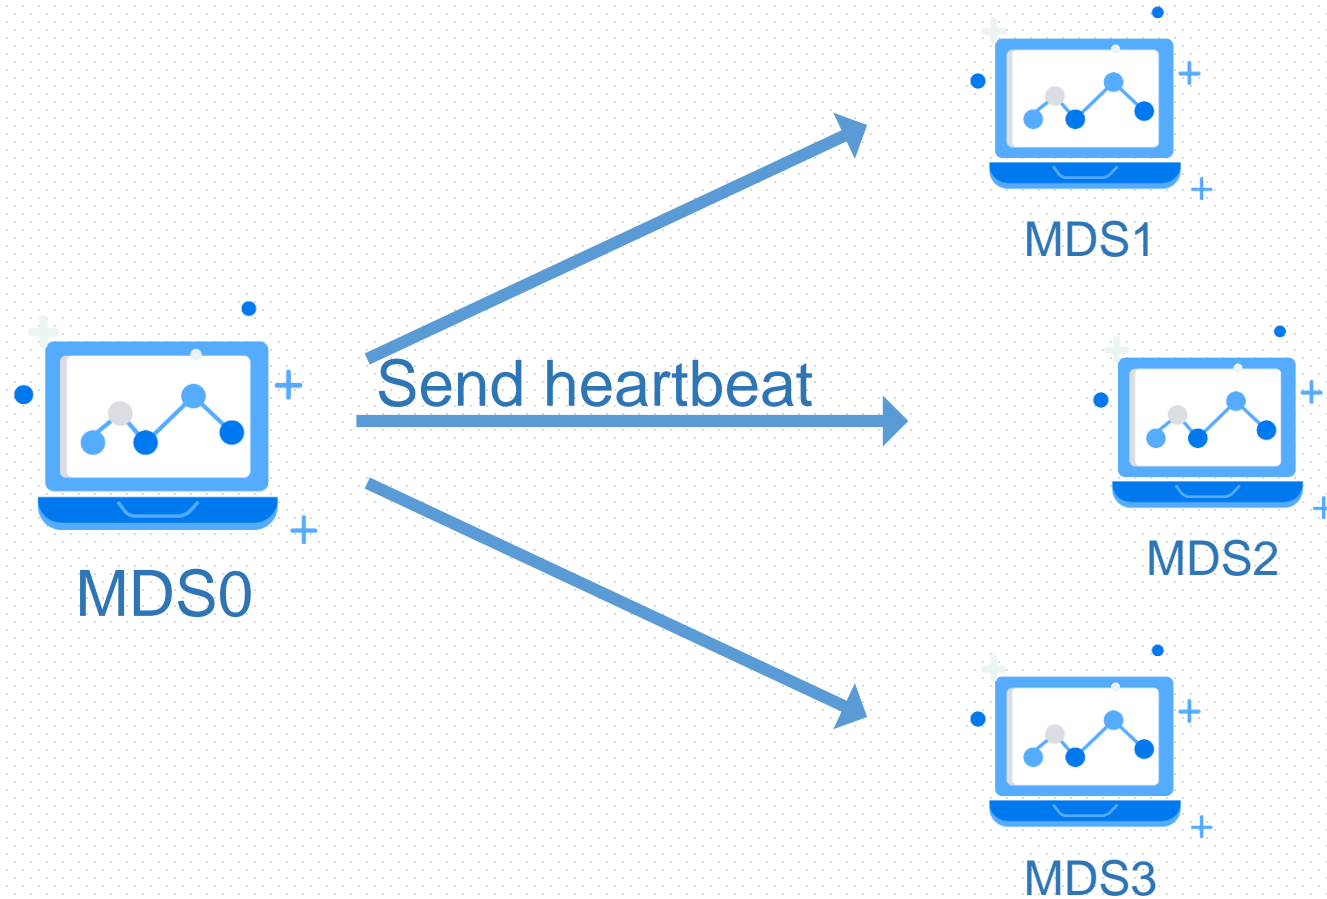


$$\frac{CoV}{\sqrt{n}} \cdot \frac{1}{1 + e^{5 - 10u}} \cdot 100\%$$

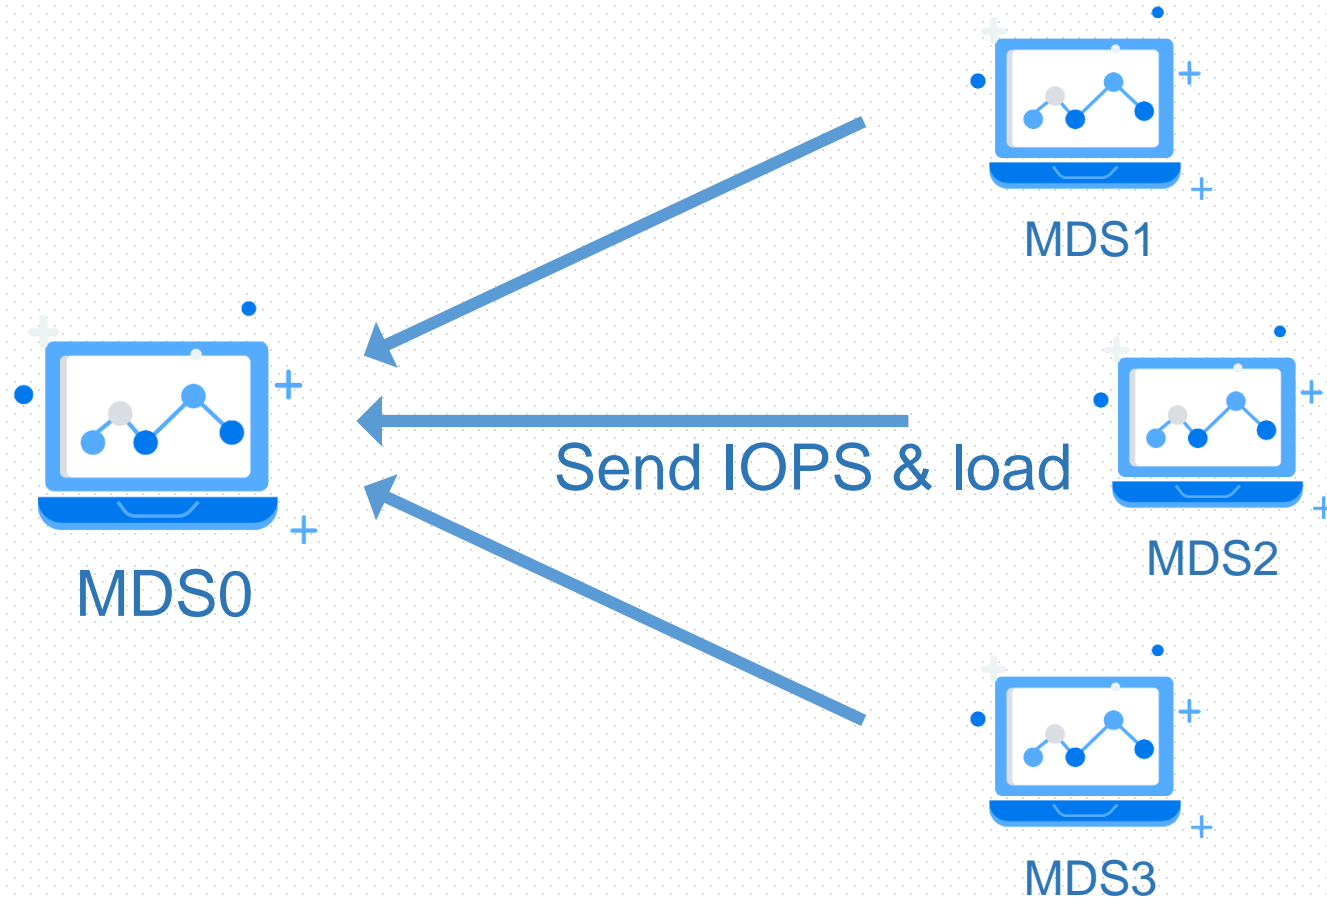


Is current imbalance harmful to the cluster?

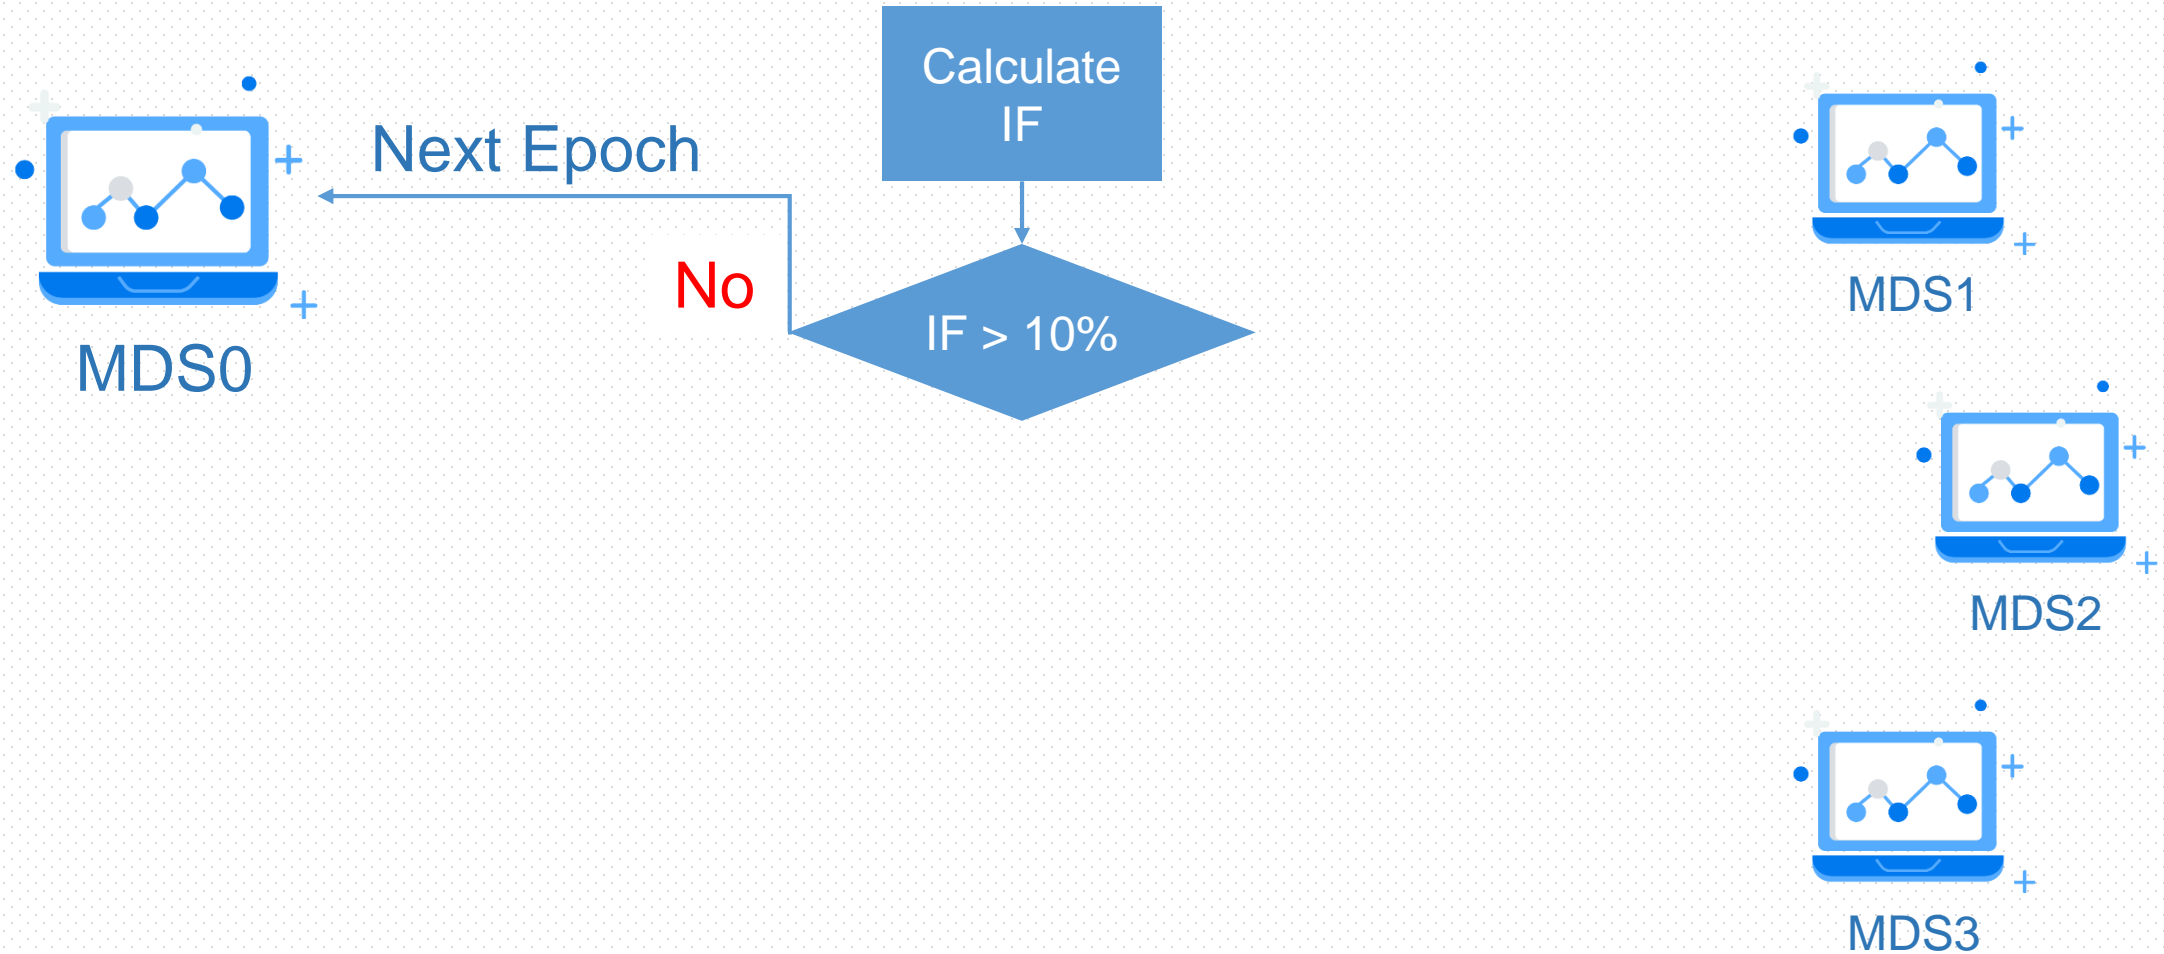
# Balancing process with IF enabled



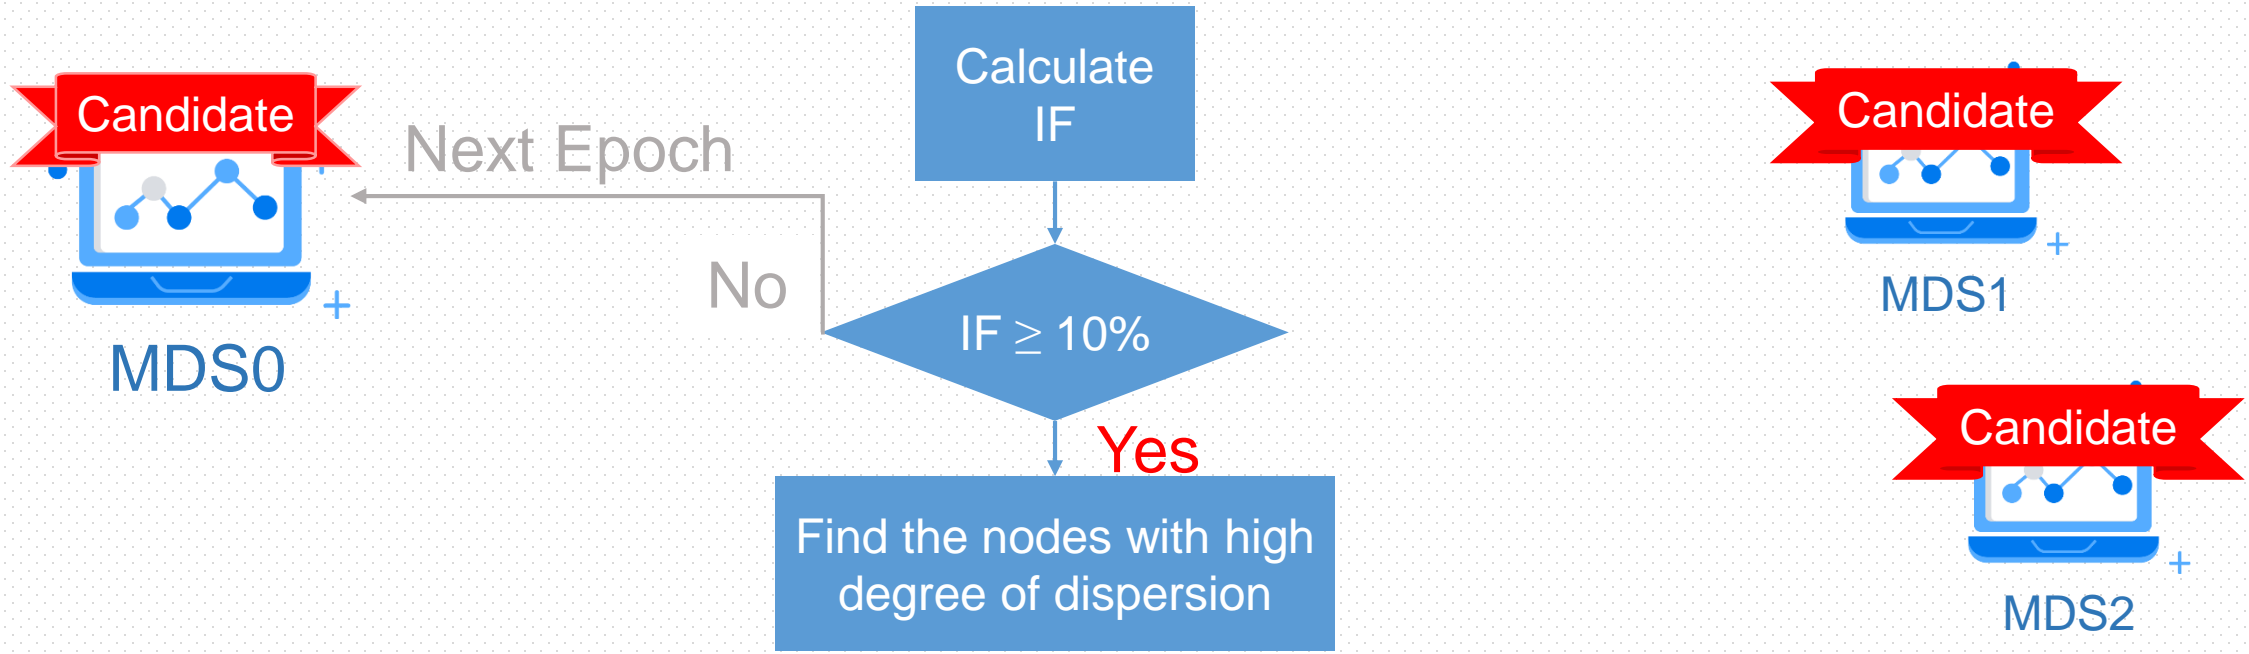
# Balancing process with IF enabled



# Balancing process with IF enabled



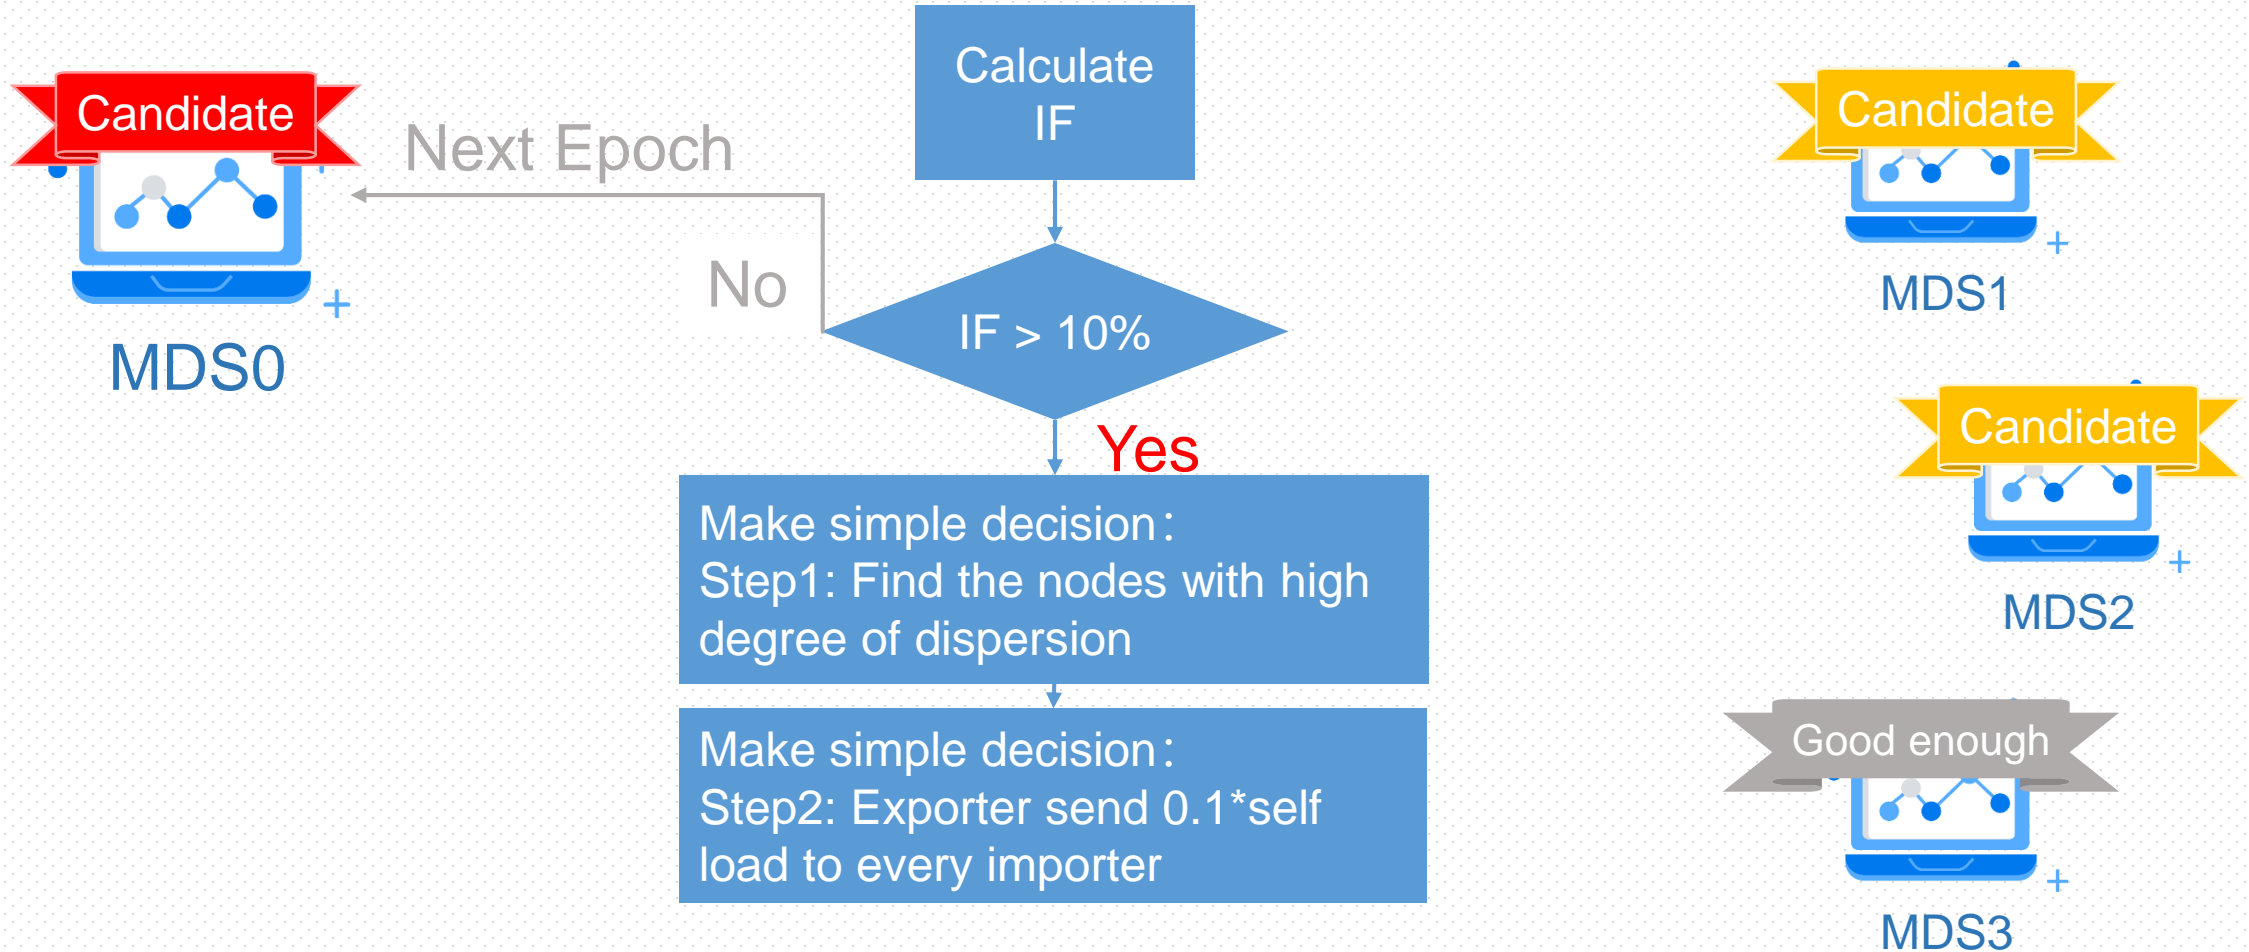
# Balancing process with IF enabled



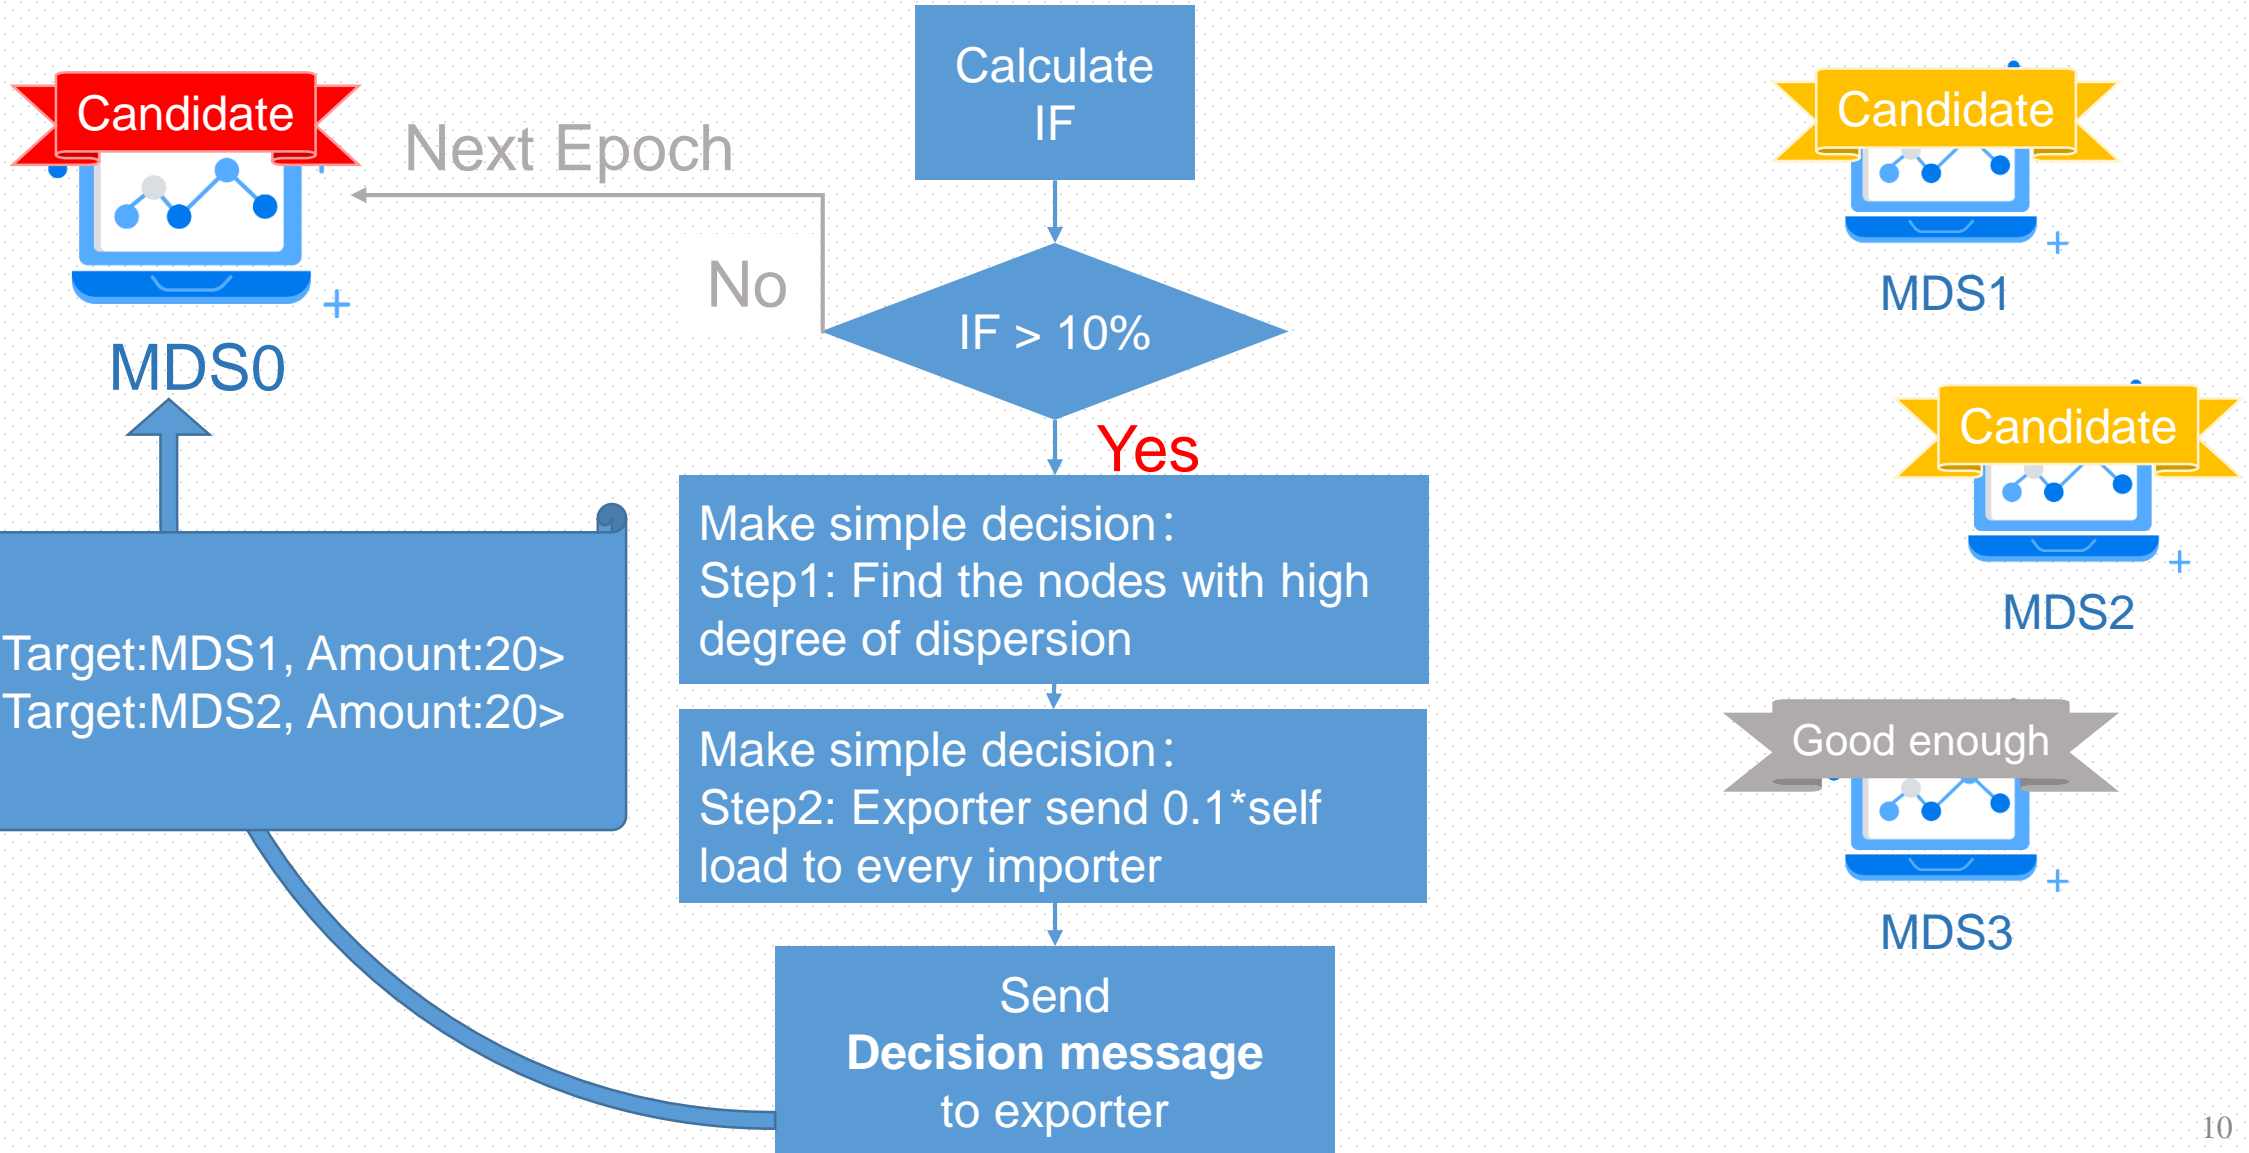
$$\text{Candidate: } \frac{\sqrt{(l-\bar{l})^2/(n-1)}}{\sqrt{n \cdot \bar{l}}} \cdot \frac{1}{1+e^{5-10 \cdot \frac{l}{\text{Capacity}}}} \cdot 100\% \geq \text{Threshold}(5\%)$$



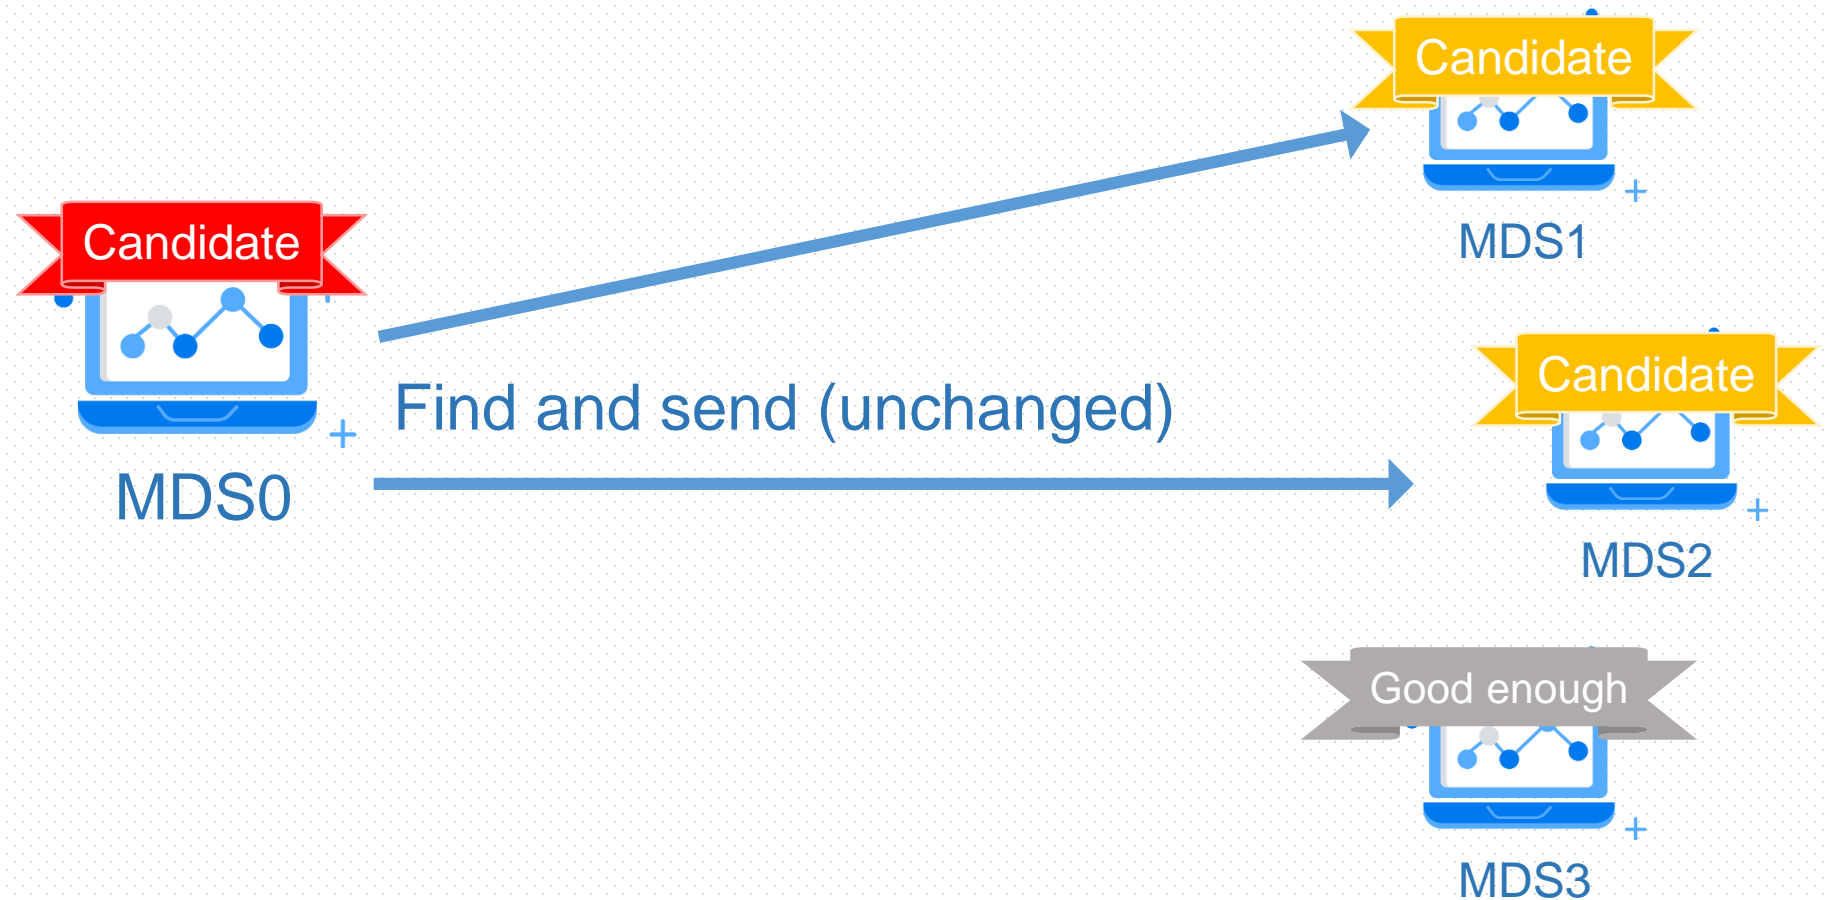
# Balancing process with IF enabled



# Balancing process with IF enabled





# Balancing process with IF enabled



<Target:MDS1,  
Amount:20>  
<Target:MDS2,  
Amount:20>


# Experiment Result: IOPS

- Mixed workload, 75 clients




# Experiment Result: Inodes

- Mixed workload, 75 clients



Original Ceph





IF enabled

# Revision of last meeting

Mixed workload without AI:


- No tar directories are migrated out from mds.0
- Shared Zipfian and web by other MDSs



# Revision of last meeting

Mixed workload without AI and Tar:

- mds.0 almost migrated everything
- mds.4 received most Zipfian directories
  - Setup:
    - 37 clients per workload (zipfian, web)
    - start at the same time



# Contents



---

- Re-experiment
- Implementation of Filebench-based Zipfian
  - Configuration scripts
  - Modification to executing scripts




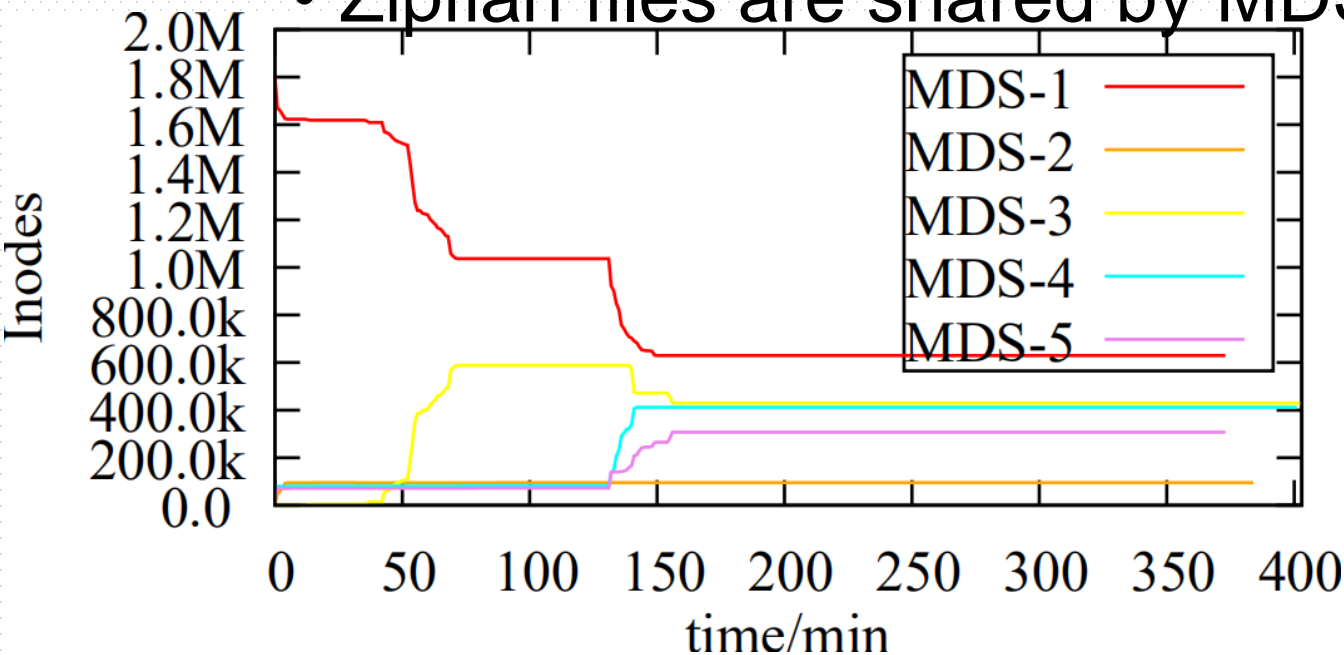
# Re-experiments

- Mixed 3 workloads (without AI):
  - 25 clients per workload (tar, zipfian, web)
  - start at the same time




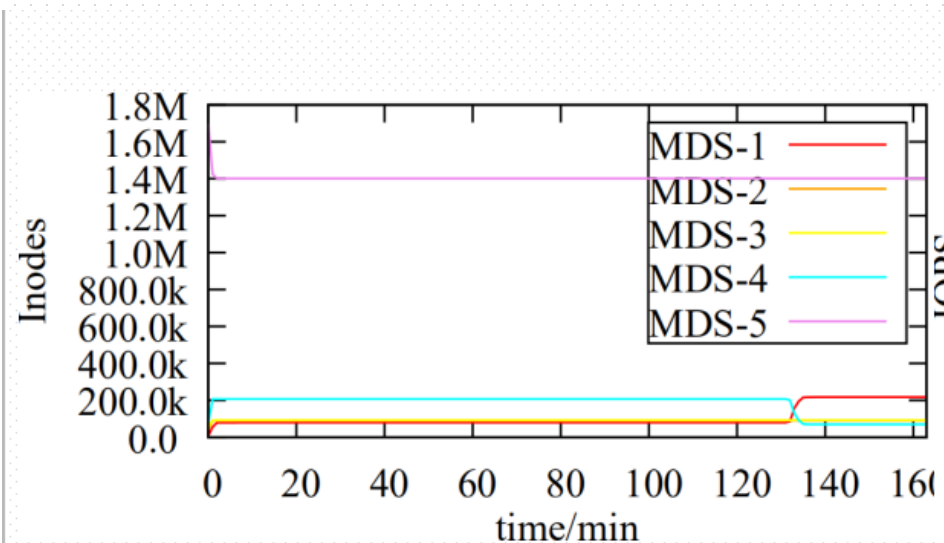


# Re-experiments

- Quite different after migration:
  - Tar directories *are* migrated (different from last time)
  - Nearly nothing on MDS-1 (MDS.0)
  - MDS-2 has only tar files, while MDS-3 has only web files. Other tar and web files are distributed in MDS-4 and MDS-5.
  - Zipfian files are shared by MDS-4 and MDS-5


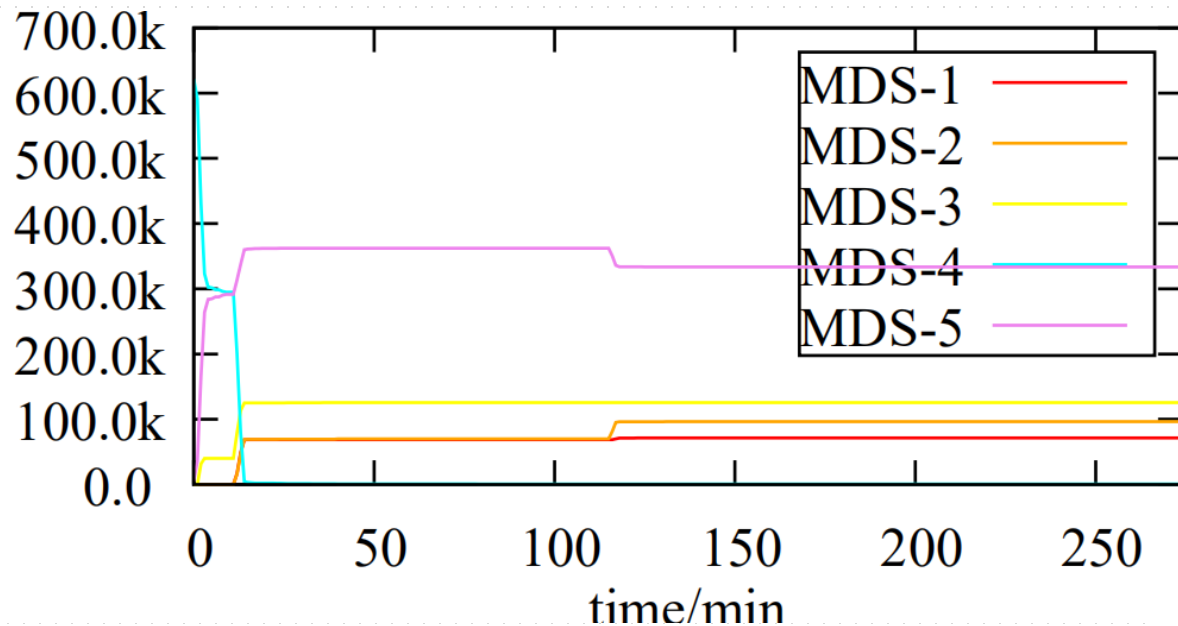


# Re-experiments



# Re-experiments

- Mixed 2 workloads (without AI, tar):
  - 37 clients per workload (zipfian, web)
  - start at the same time






# Re-experiments

- Mixed 2 workloads (without AI)
  - 37 clients per workload (zipfian, web)
  - start at the same time


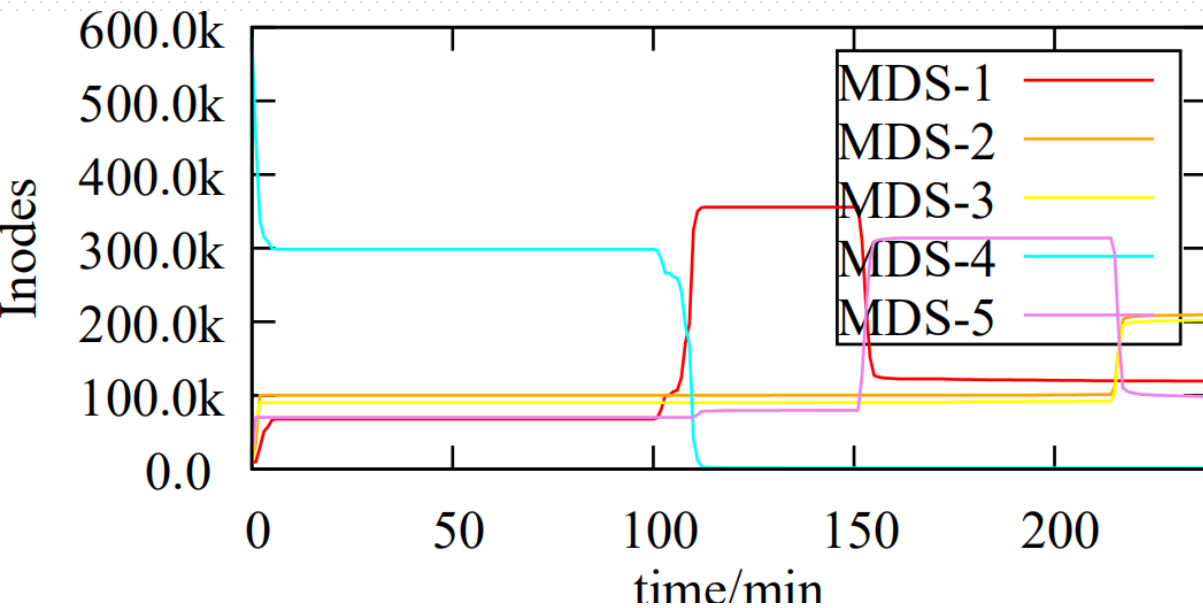


- Setup:
  - 37 clients per workload (zipfian, web)
  - start at the same time



# Re-experiments

- Mixed 2 workloads (without AI, tar) (**limited memory**):
  - 37 clients per workload (zipfian, web)
  - start at the same time



# Implementation of Filebench-based Zipfian

- Configuration:
  - Single round: For every 100 files (No. 1, 2, ..., 100), visit No.1 file 9801 times, with other files once
  - $9801 / (1 * 99) = 99:1$
- Starting scripts:
  - Filebench creates filesets at runtime. We need sync before running all workloads
  - Still unexpected bugs:
    - Interfere with other workloads?