# The Case for Benchmarking Control Operations in Cloud Native Storage

## 12th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage '20)

Alex Merenstein[1], Vasily Tarasov[2], Ali Anwar[2], Deepavali Bhagwat[2], Lukas Rupprecht[2], Dimitris Skourtis[2], and Erez Zadok[1]

[1]Stony Brook University; [2]IBM Research - Almaden




IBM

Stony Brook University

# Outline

- **Introduction**
- Storage Control Operations
- Impact of Storage Control Operations
- Benchmark Design
- Conclusion

# New Trends in Clouds

- Cloud native software
  - Often container based
  - Microservice architectures
  - Frequent scaling and updates
- Cloud native storage
  - Used by applications, not systems
  - Automated management
  - Container Storage Interface (CSI) provides standard interface


https://landscape.cncf.io

# Benchmarking's Blind Spot

- Storage challenges
  - Choosing a storage provider
  - Evaluating different storage configurations
- Current benchmarks (e.g., fio[1], pgbench[2], NoSQLBench[3])
  - I/O operations
  - Metadata operations
  - **Storage control operations**

1. https://fio.readthedocs.io/en/latest/index.html
2. https://www.postgresql.org/docs/current/pgbench.html
3. https://www.datastax.com/blog/2020/03/nosqlbench



IBM

Stony Brook University

# Outline

- Introduction
- **Storage Control Operations**
- Impact of Storage Control Operations
- Benchmark Design
- Conclusion

# Storage Control Operations

- Storage control operations
  - Creating volumes, attaching volumes, snapshotting, resizing, etc.
  - Volumes: single unit of storage provisioned by a storage provider
- More frequent in cloud native environments
- Existing benchmarks do not generate storage control operations

# Increasing Number of Storage Control Operations

- Some companies have increased deployments from 2–3 × week to 150 × day[1]
- On one platform, 54% of containers ran for ≤5 minutes and hosts ran a median of 30 containers[2]
  - On a 20 nodes cluster, that results in a rate of one container creation per second

1. https://www.weave.works/technologies/going-cloud-native-6-essential-things-you-need-to-know
2. https://sysdig.com/blog/sysdig-2019-container-usage-report/



IBM

Stony Brook University

# User Creates Container Requiring Storage

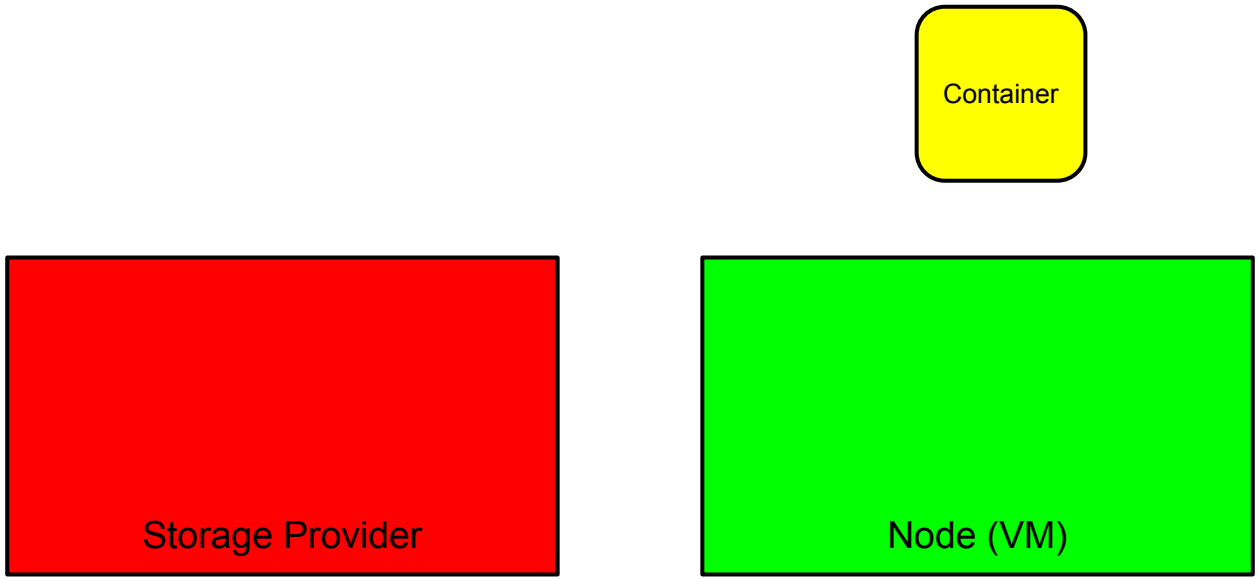

# Create Volume (Storage Control Operation #1)

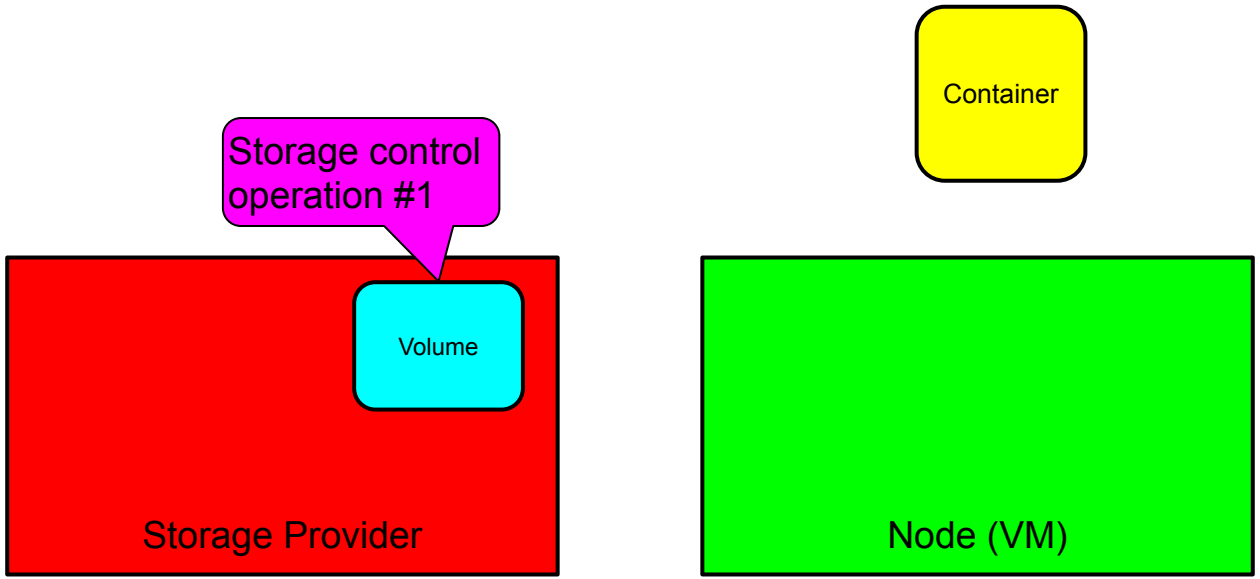

# Container Scheduled on Node

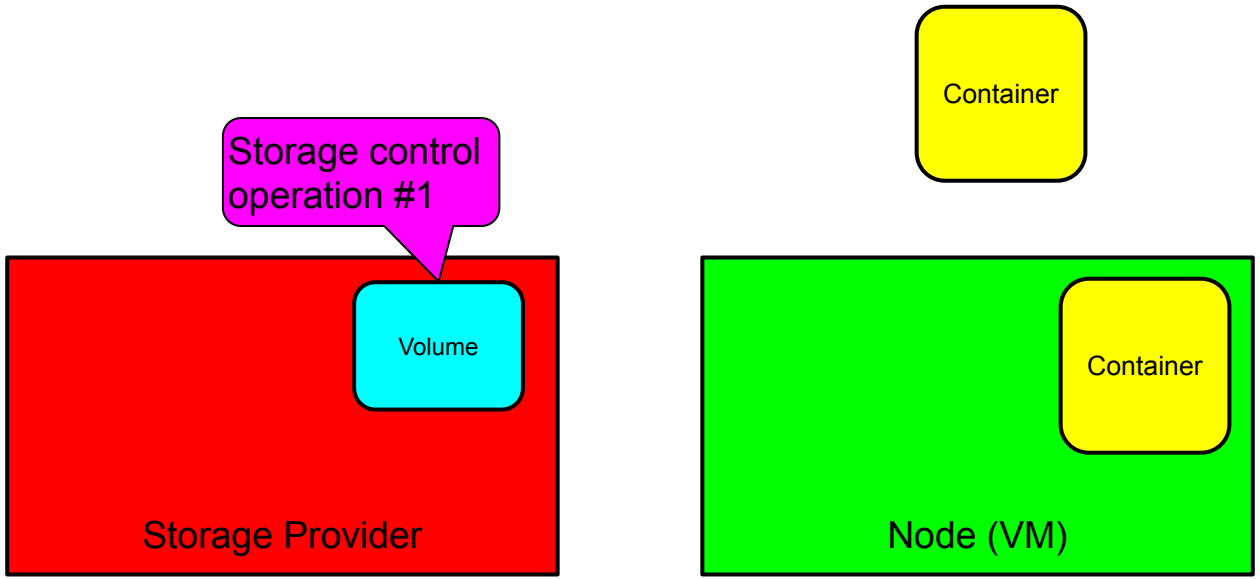

# Volume Mounted on Node (Storage Control Operation #2)

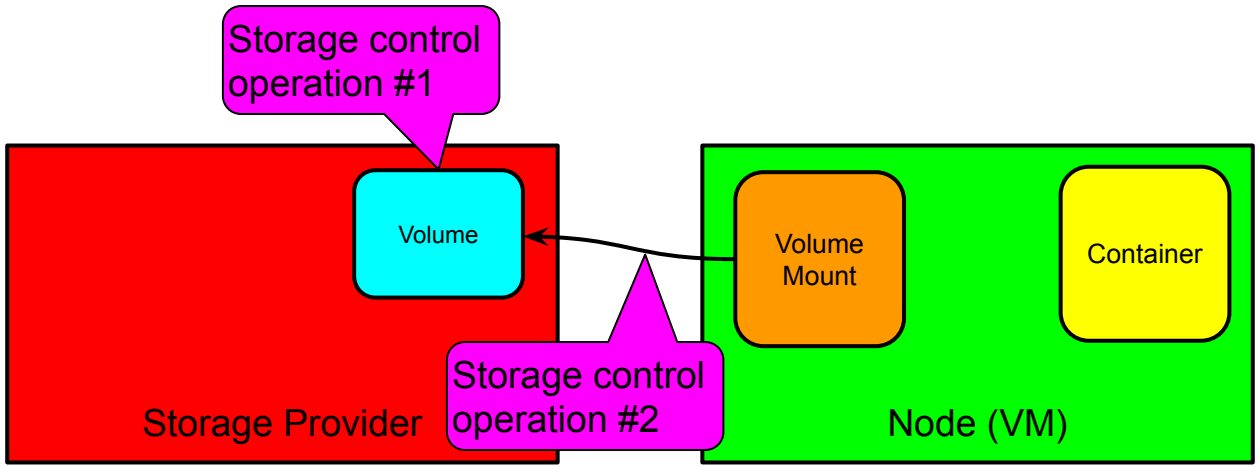

# Volume Attached to Container (Storage Control Operation #3)

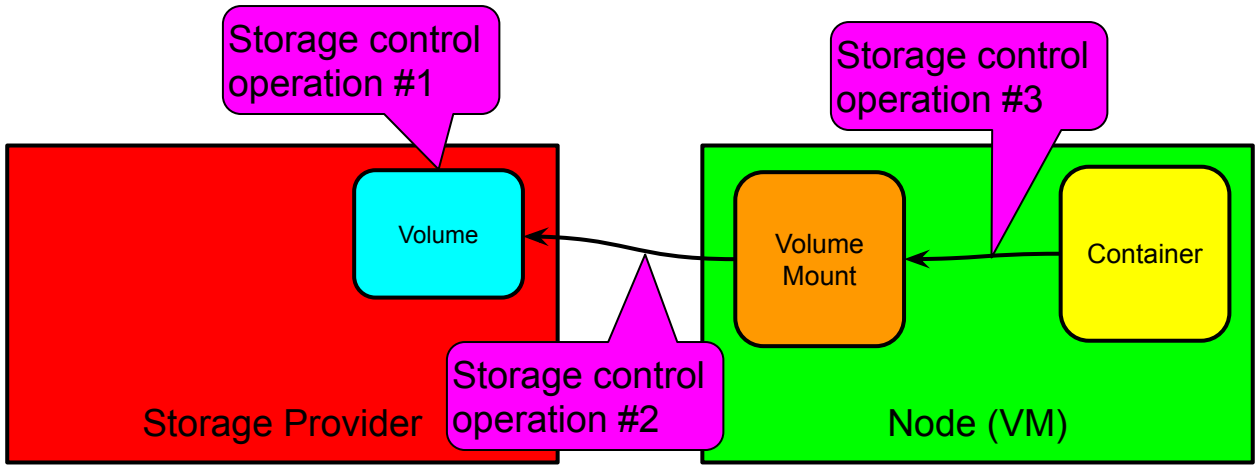

# Outline

- Introduction
- Storage Control Operations
- **Impact of Storage Control Operations**
- Benchmark Design
- Conclusion

# Impact of Storage Control Operations

- Experiment 1: creating and attaching volumes
  - Do storage providers have different performance characteristics when executing these operations?
- Experiment 2: snapshots with concurrent workload
  - Can storage control operations impact other workloads?
  - Is the level of impact different across different storage providers?

# Experimental Setup

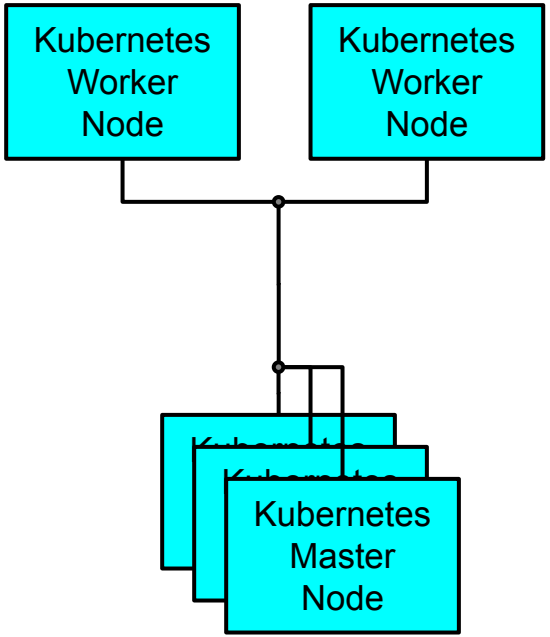


Kubernetes with three masters in high availability configuration and two workers nodes


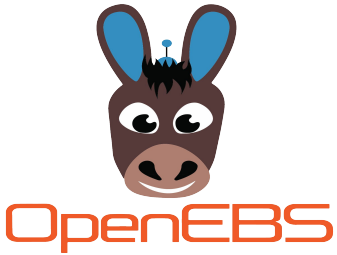



Three different-by-design storage providers

# Experiment 1: Volume Creation and Attachment

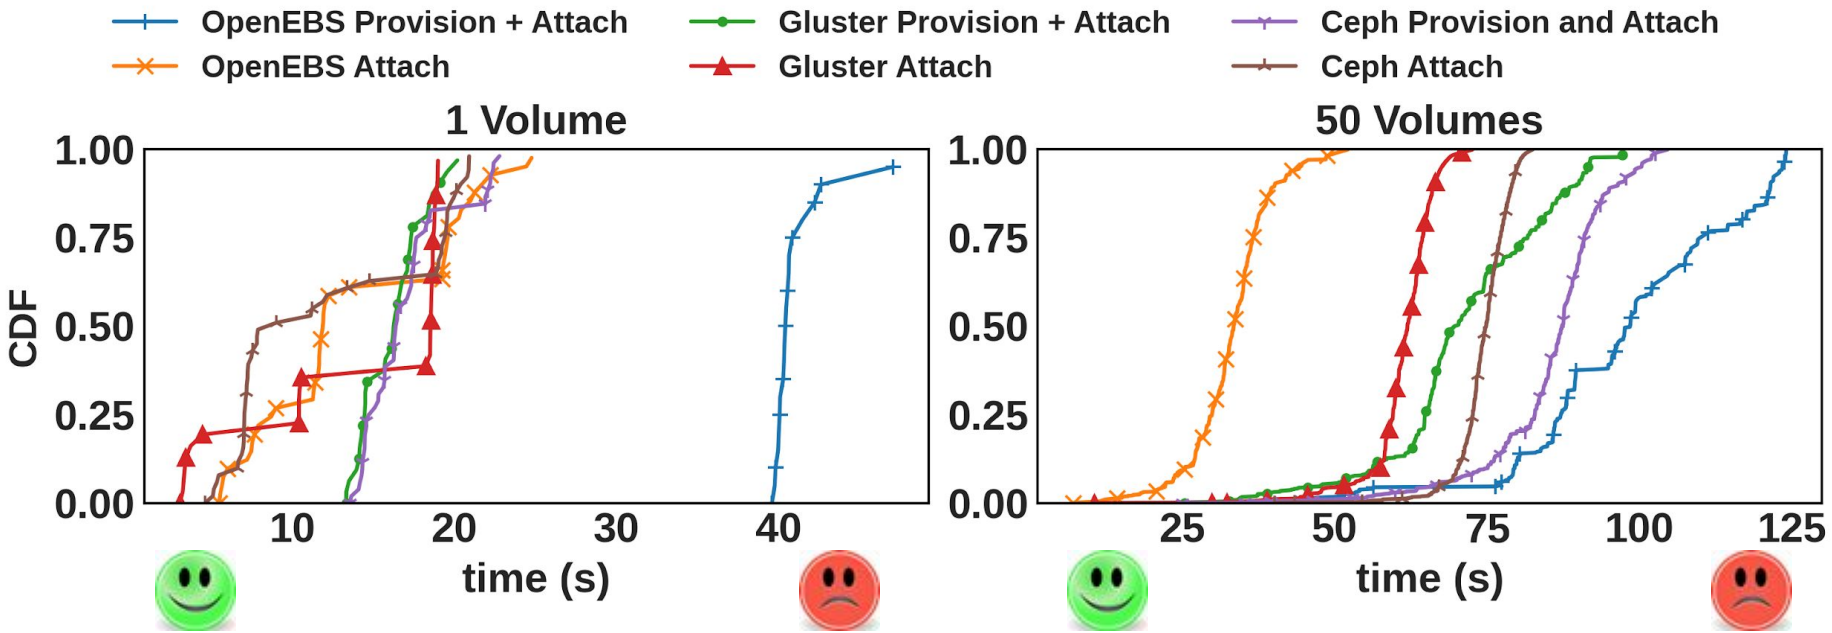

# Experiment 1: Volume Creation and Attachment

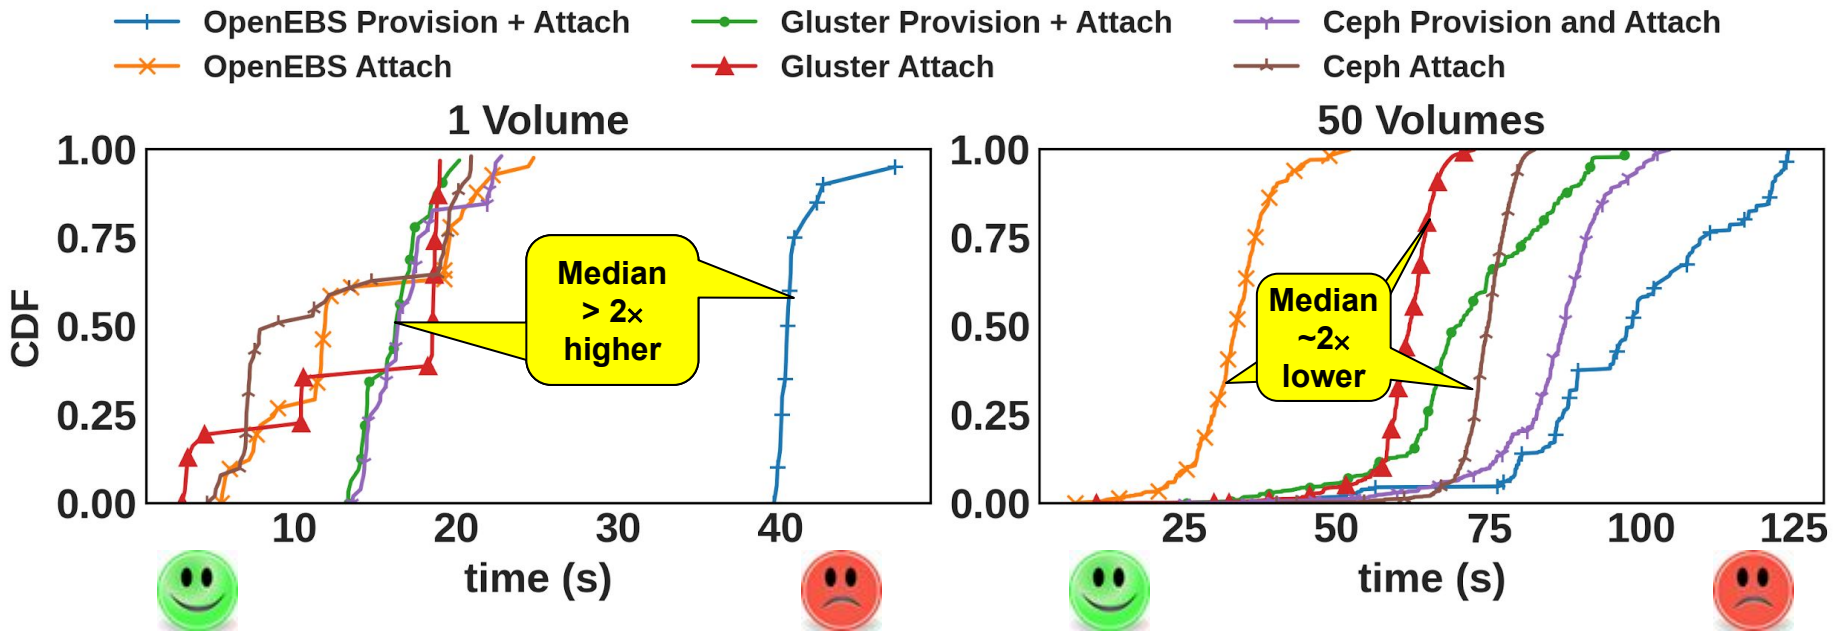

# Experiment 1: Volume Creation and Attachment

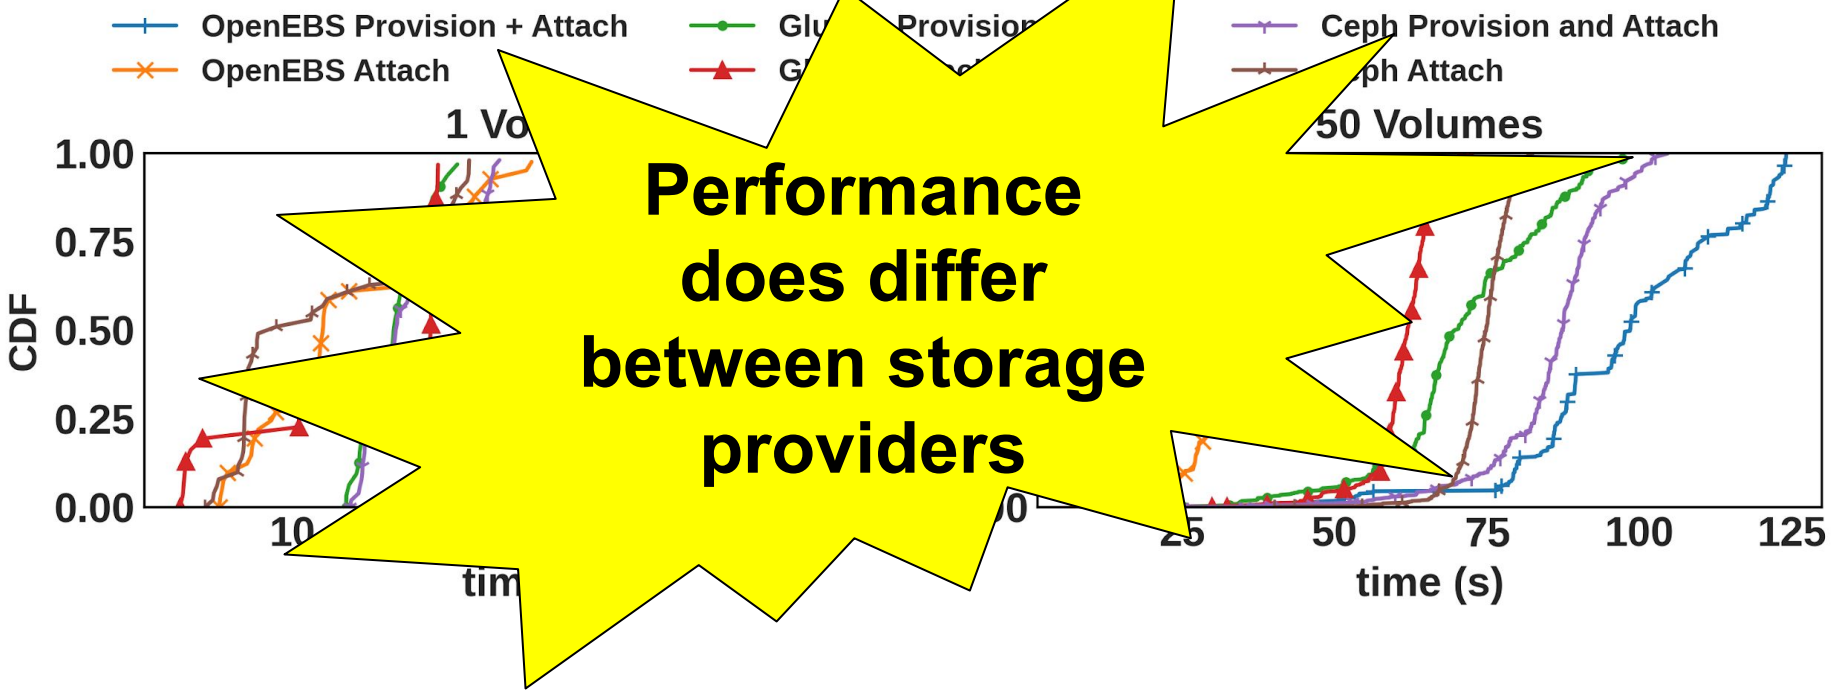

# Experiment 2: Snapshotting

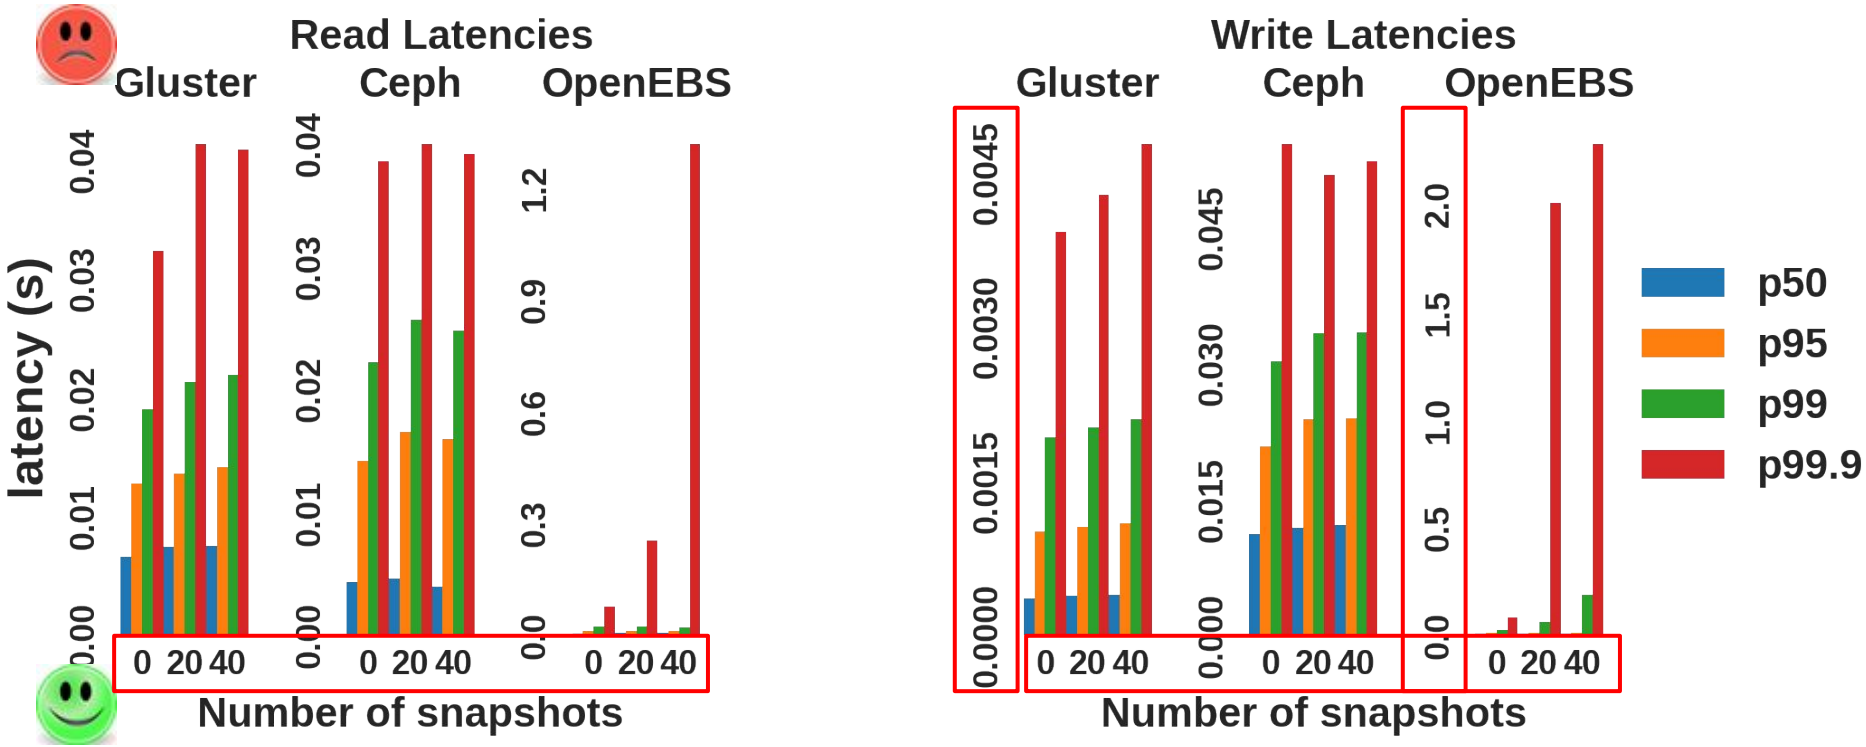

# Experiment 2: Snapshotting

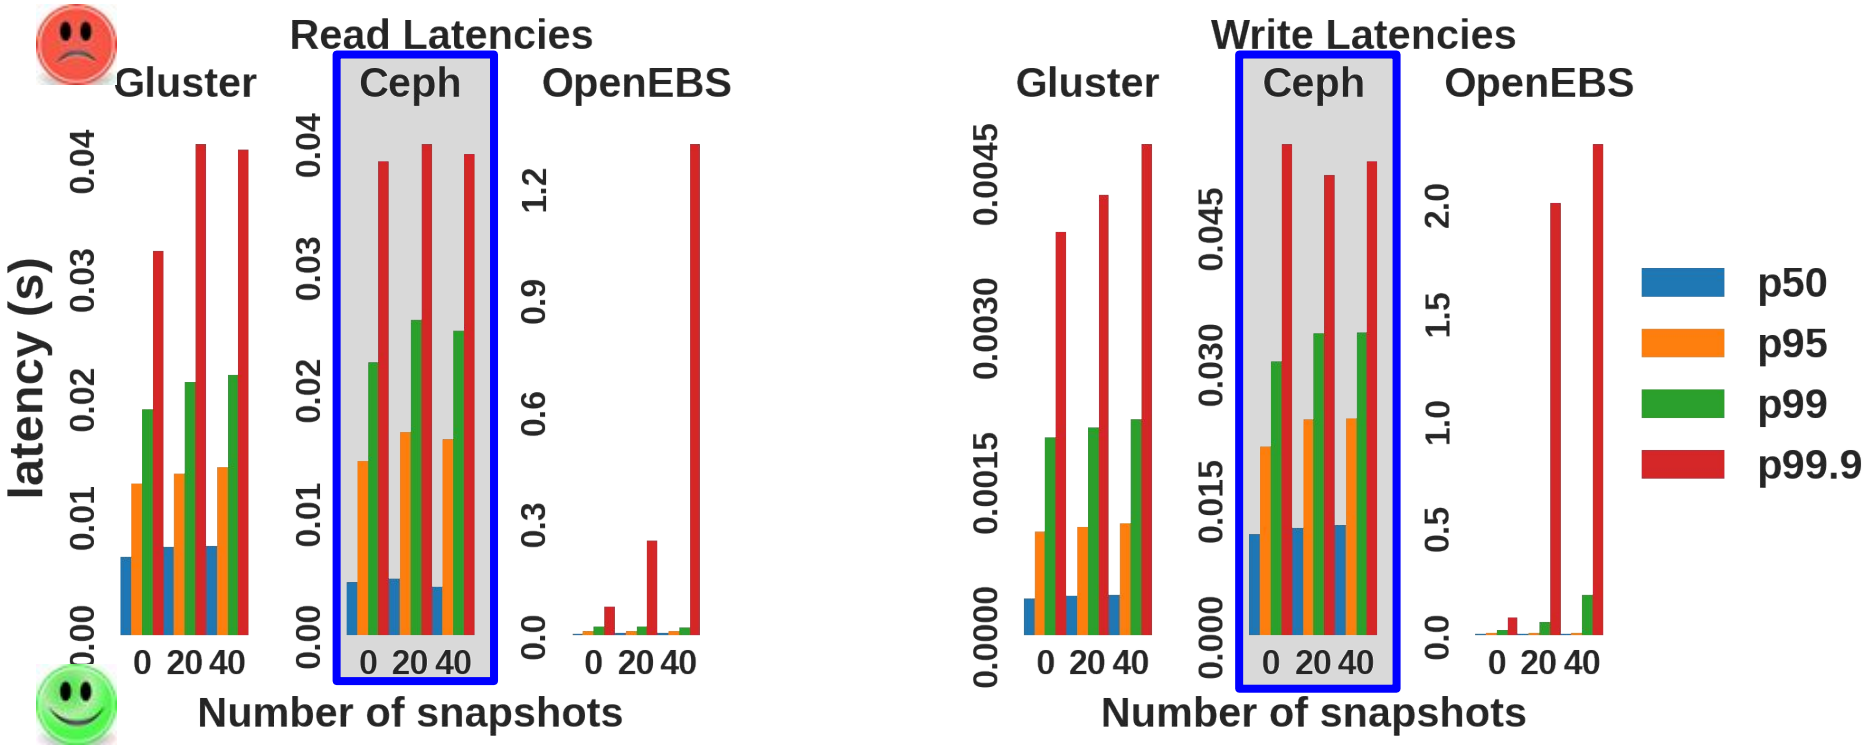

# Experiment 2: Snapshotting

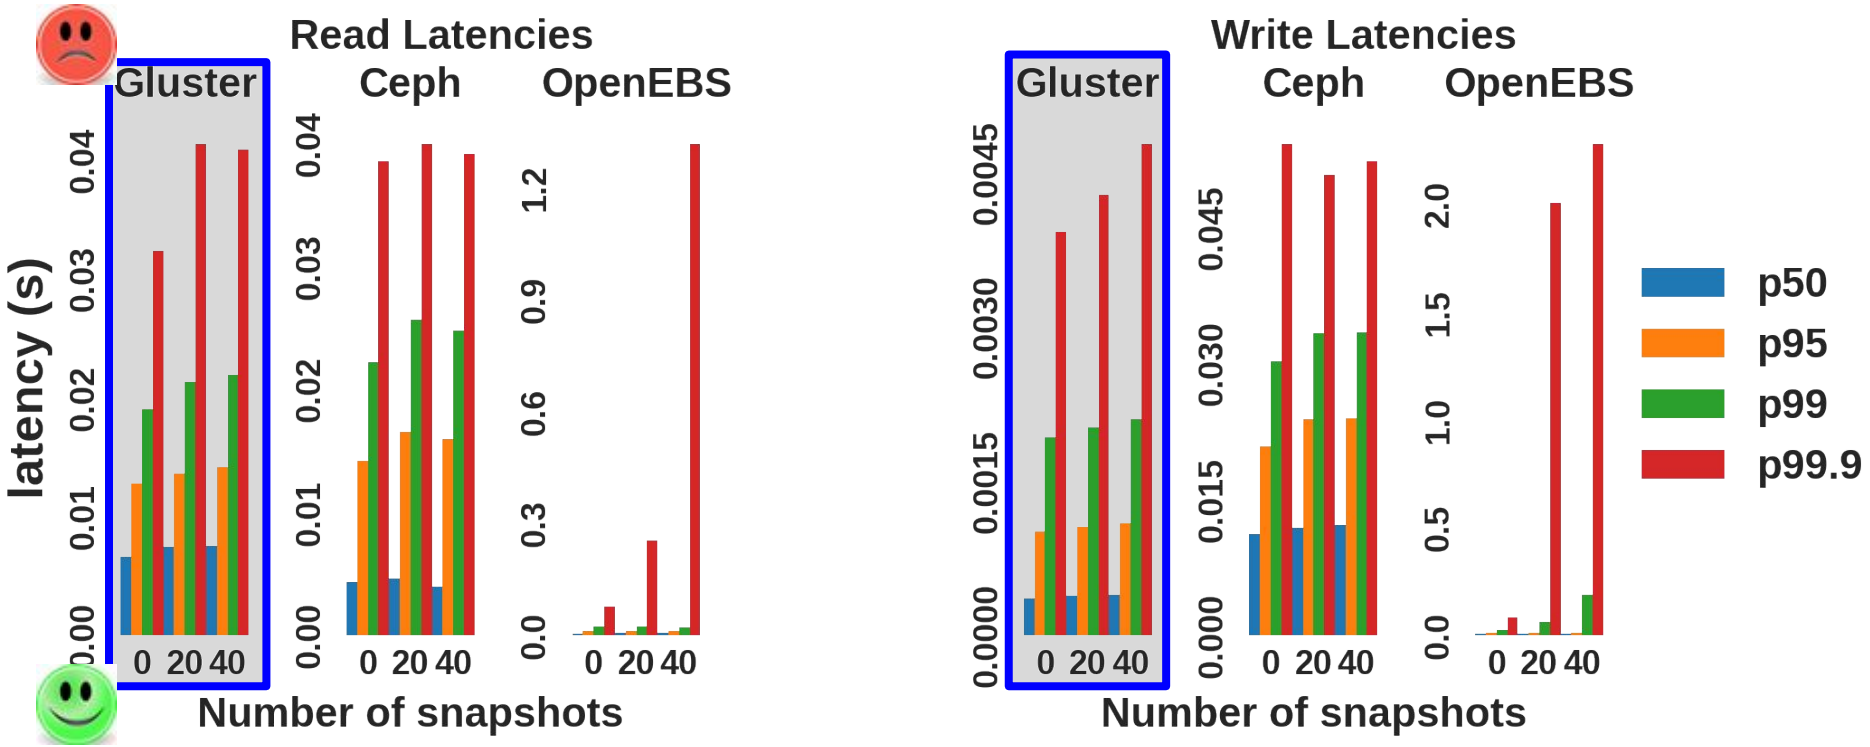

# Experiment 2: Snapshotting

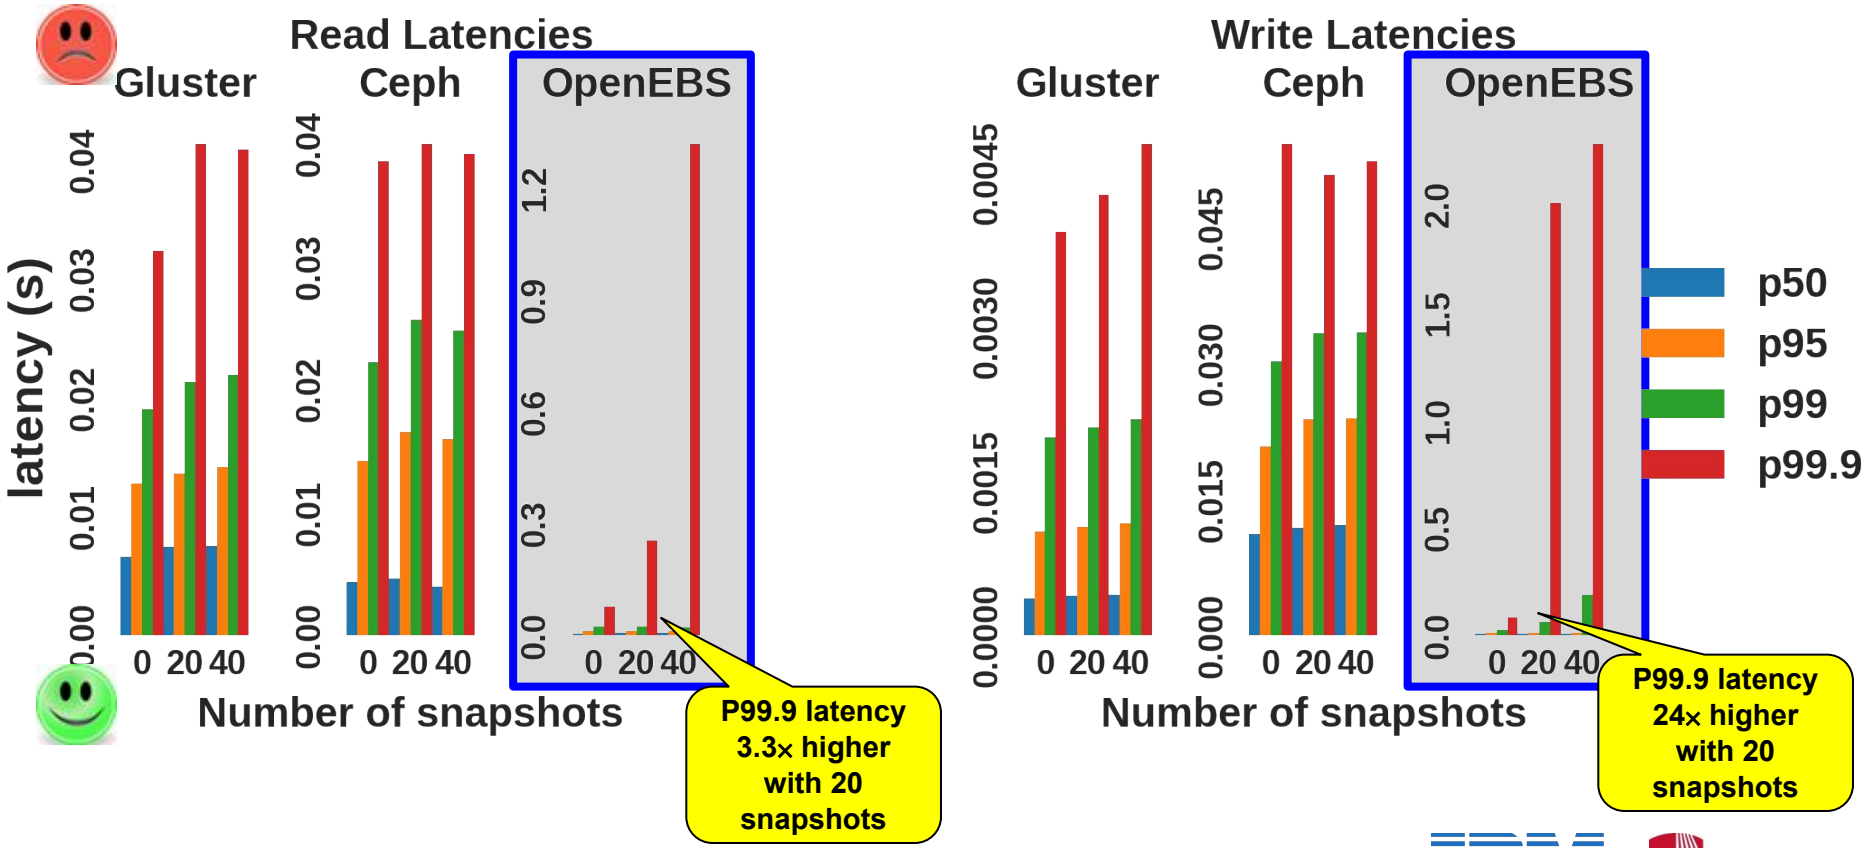

# Experiment 2: Snapshotting

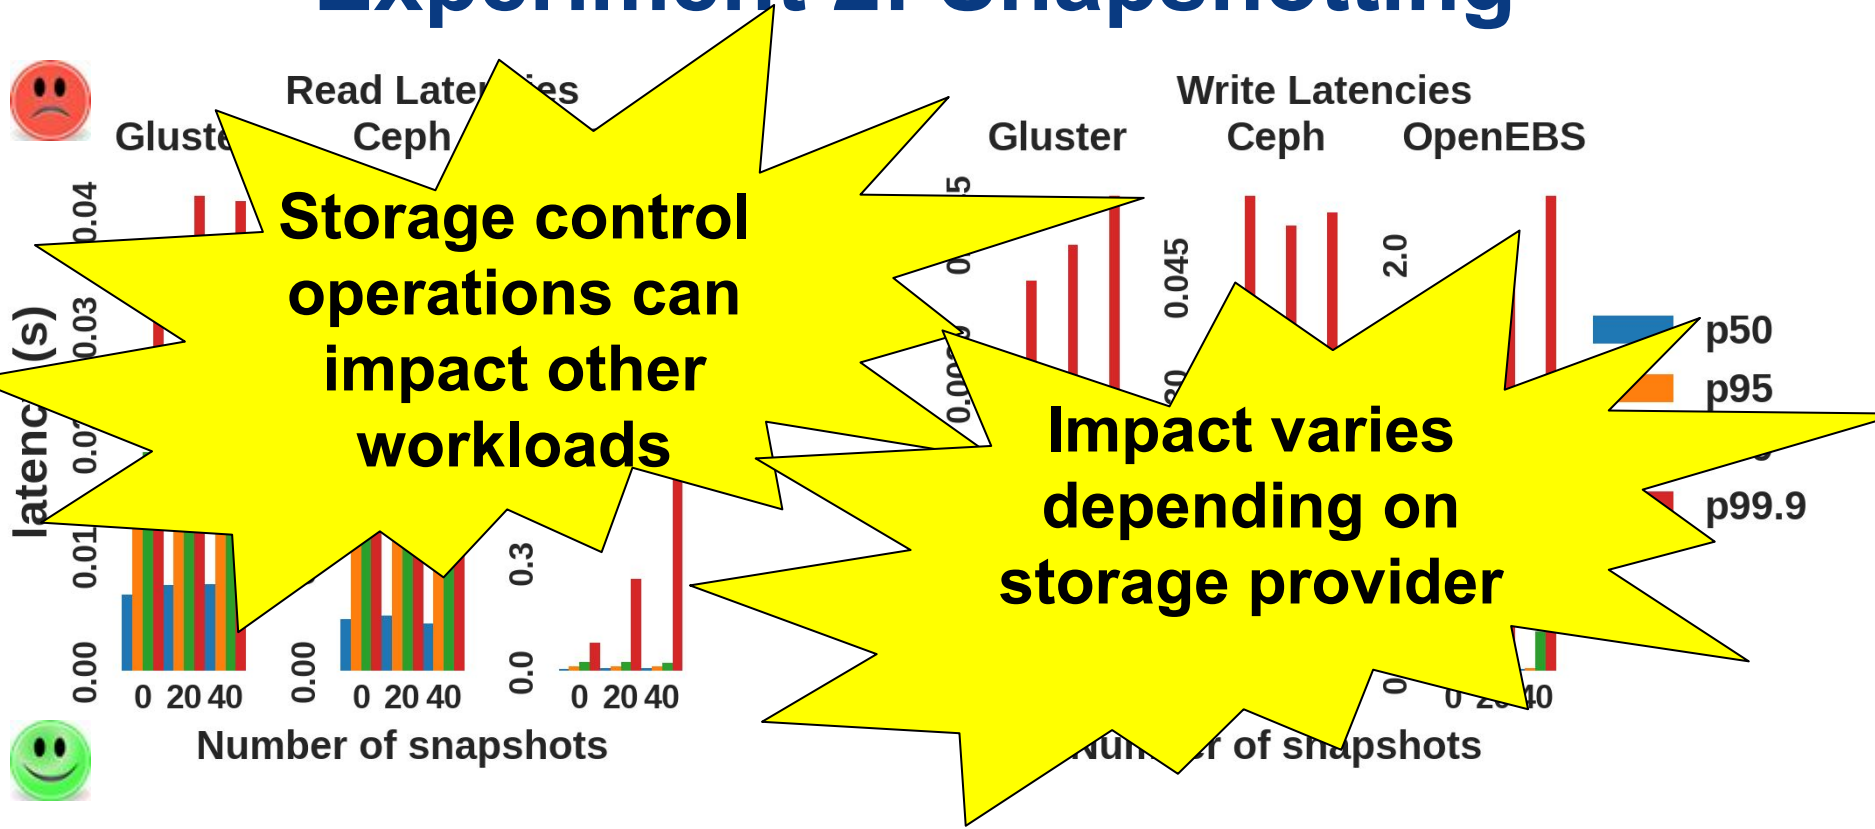

# Outline

- Introduction
- Storage Control Operations
- Impact of Storage Control Operations
- **Benchmark Design**
- Conclusion

# Benchmark Design Requirements

## Workload

1. Create I/O and storage control workloads
2. Specify complex & realistic storage control workloads
3. Use existing tools for I/O workloads
4. Include QoS targets

## Useability

1. Enable reproducibility
2. Be easy to use

## Result Measurement & Visualization

1. Measurement should be decoupled from I/O generation
2. Results should be aggregated in clear, actionable manner
3. Metrics collection should have low overhead

# Proposed Design

1. Benchmark Controller: creates I/O workload containers and executes control operations

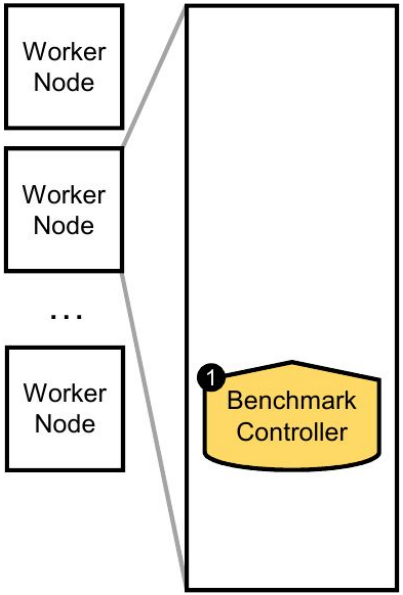

# Proposed Design

2. User creates Benchmark object
3. Benchmark objects: custom object type, created by users to define a benchmark

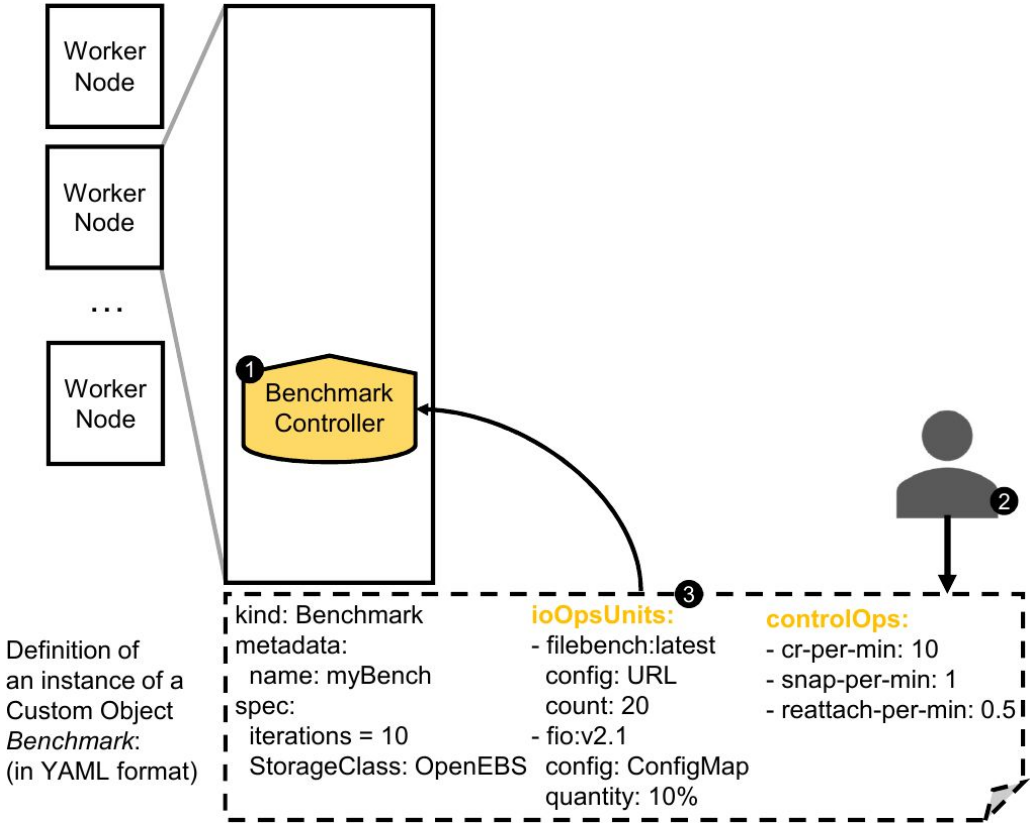

# Proposed Design

4. I/O containers: created by Benchmark Controller to run I/O workload
5. Container image repository: I/O workloads can be created using existing I/O benchmarking tools such as fio or filebench

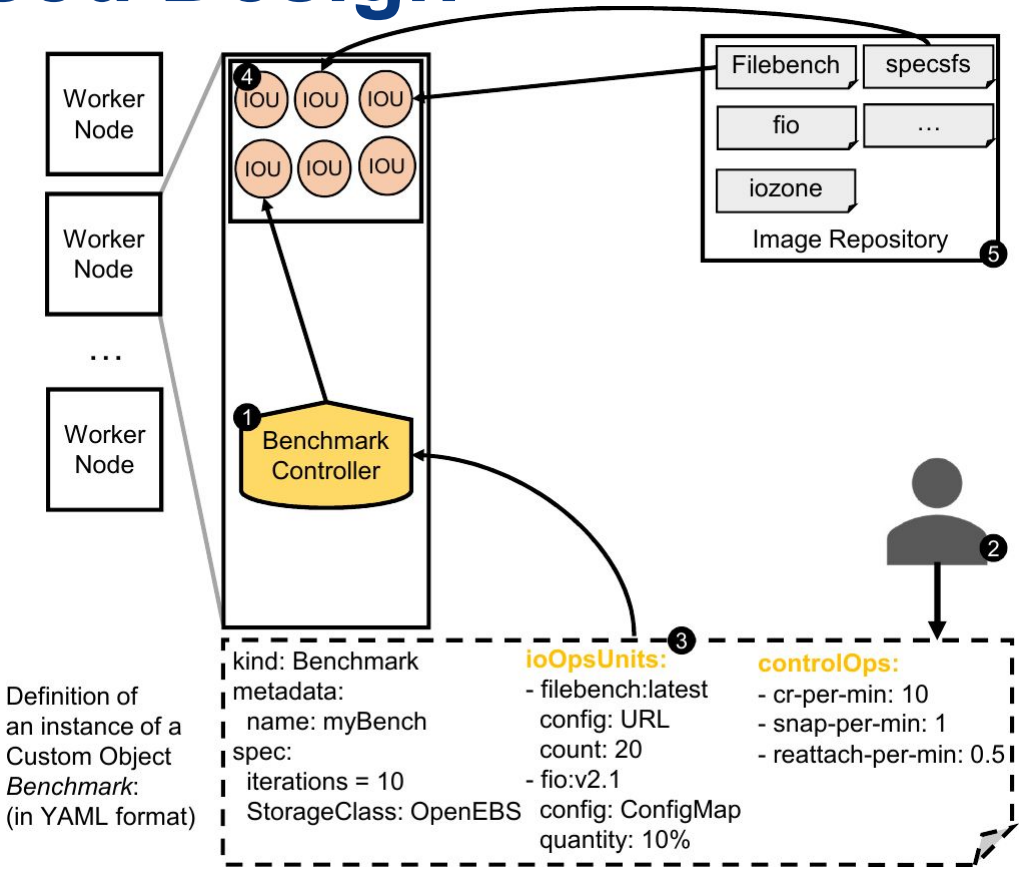

# Proposed Design

6. The Benchmark Controller executes Control Operation workloads by acting directly on PVs and PVCs

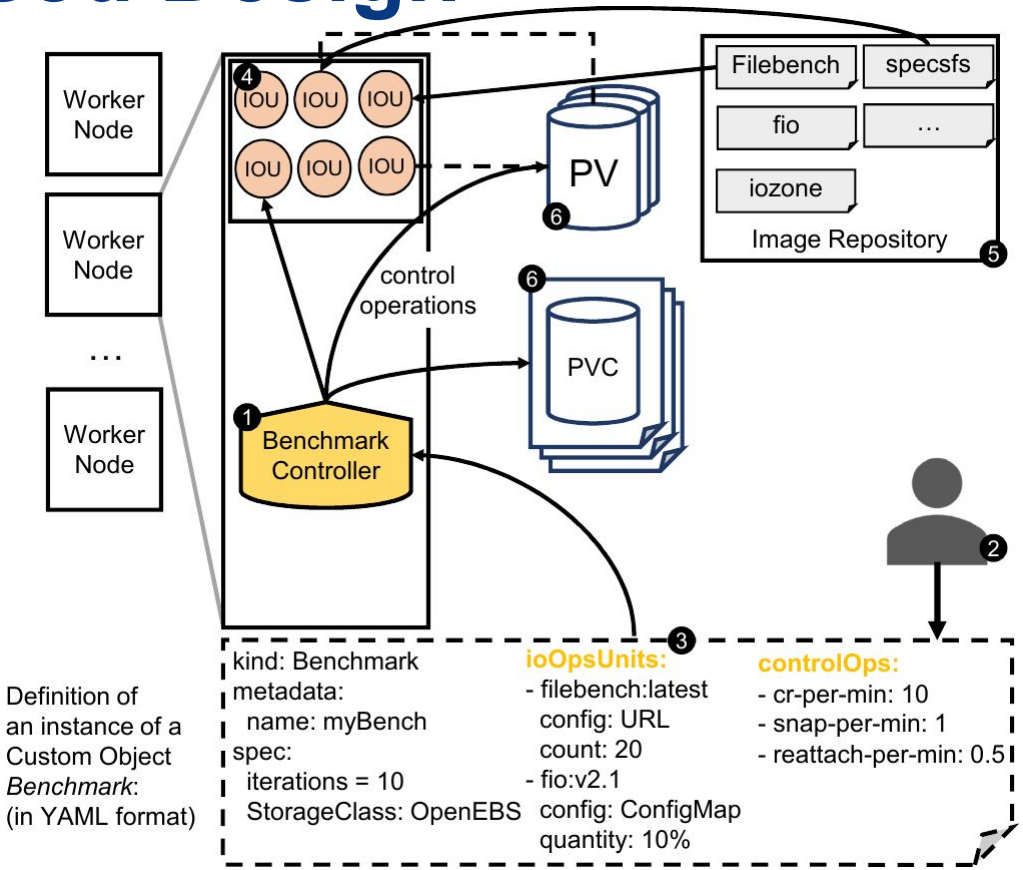

# Proposed Design

7. The volumes used by the benchmark are provisioned by the storage provider specified in the Benchmark object

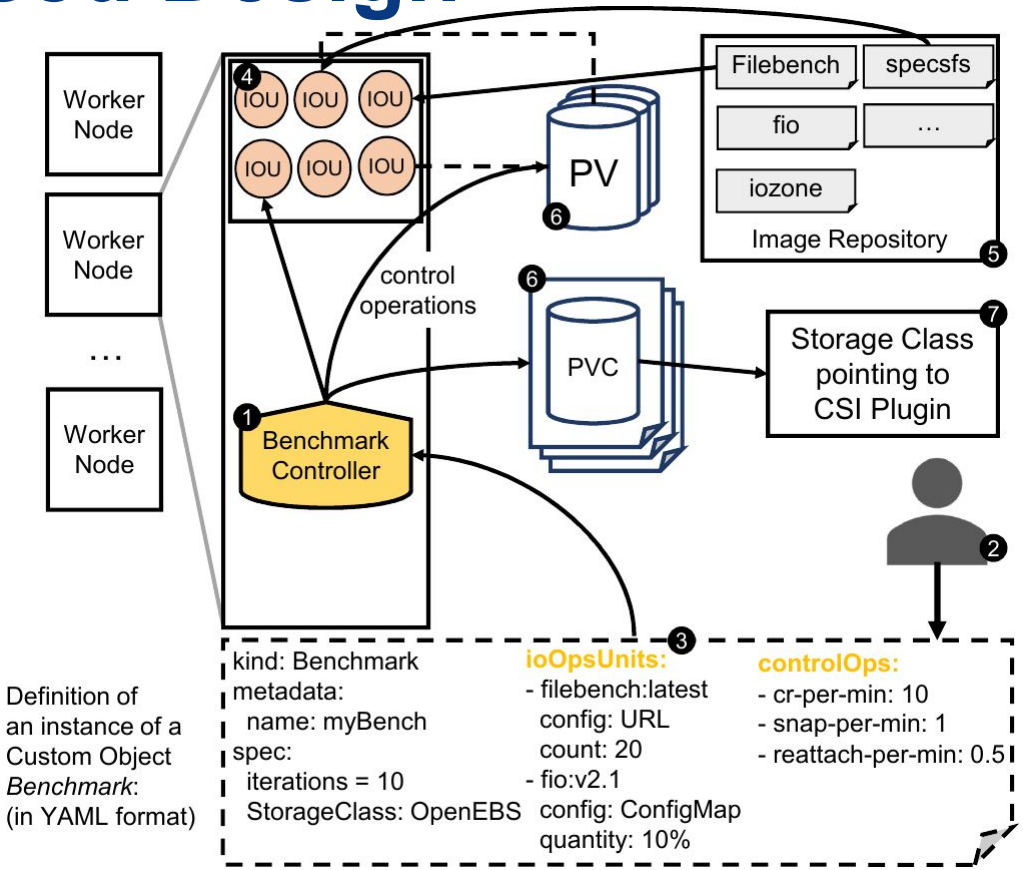

# Proposed Design

8. Results and metrics are collected and can be analyzed and visualized using tools such as ELK or Grafana

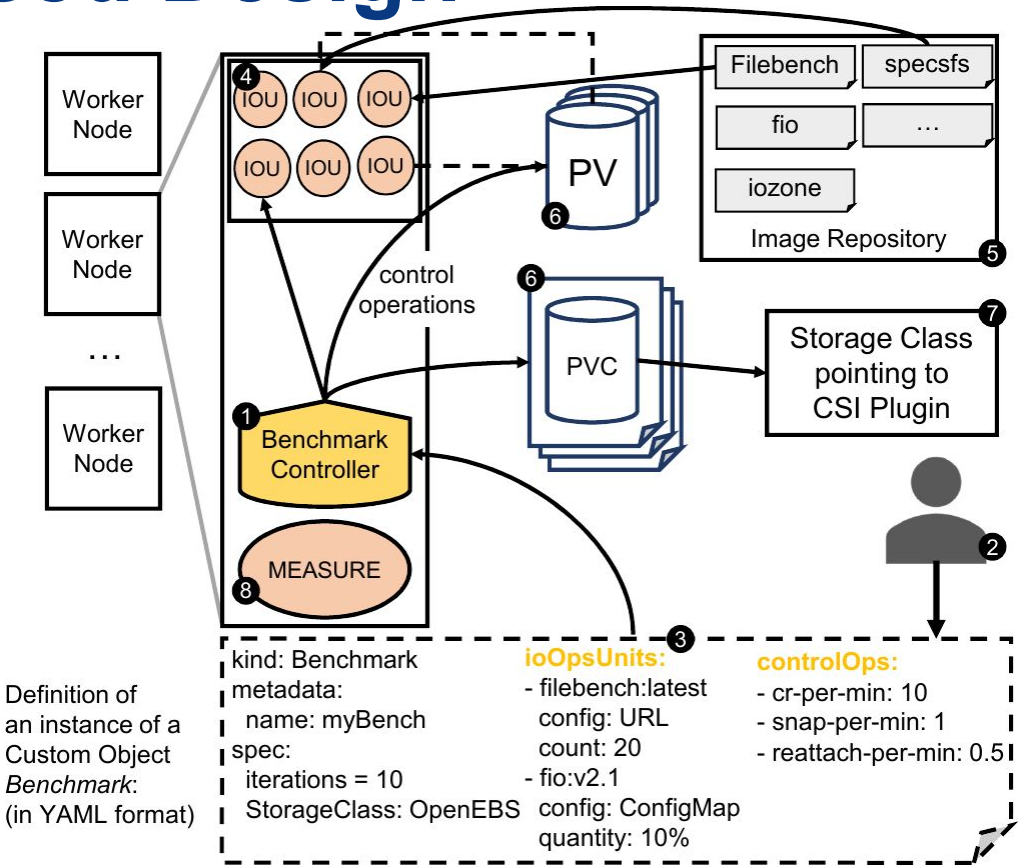

# Outline

- Introduction
- Storage Control Operations
- Impact of Storage Control Operations
- Benchmark Design
- **Conclusion**

# Conclusion

- New benchmark is needed to support cloud native workflows
- Proposed nine requirements and an initial design for such a benchmark
- Looking for community input, especially for storage control operation rates

# The Case for Benchmarking Control Operations in Cloud Native Storage

Alex Merenstein[1], Vasily Tarasov[2], Ali Anwar[2], Deepavali Bhagwat[2], Lukas Rupprecht[2], Dimitris Skourtis[2], and Erez Zadok[1]

[1]Stony Brook University; [2]IBM Research - Almaden

Contact: mmerenstein@cs.stonybrook.edu

# Thank You
# Q&A