More IOPS for Less:
Exploiting Burstable Storage in Public Clouds

Hojin Park, Greg Ganger, George Amvrosiadis

Carnegie Mellon University

Burstable storage service

- Cloud storage type that has **“burst credits”**
  - Burst credit provides an increase in throughput, over baseline IOPS for a limited time

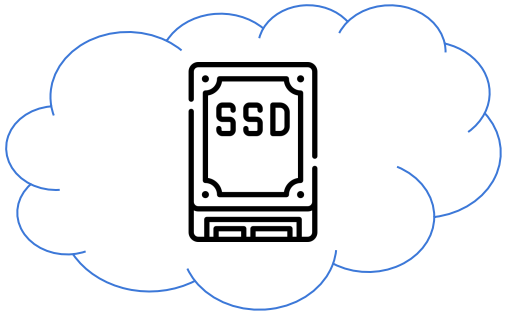

E.g., General Purpose SSD on AWS

Capacity : 100 GiB
Baseline perf. : 300 IOPS
Burst perf. : up to 3,000 IOPS

AWS Case Study: Burstable storage service

E.g., Using 100 GiB of General Purpose SSD

0 - 20 min: storage perf. $\leq$ baseline perf. (300 IOPS)

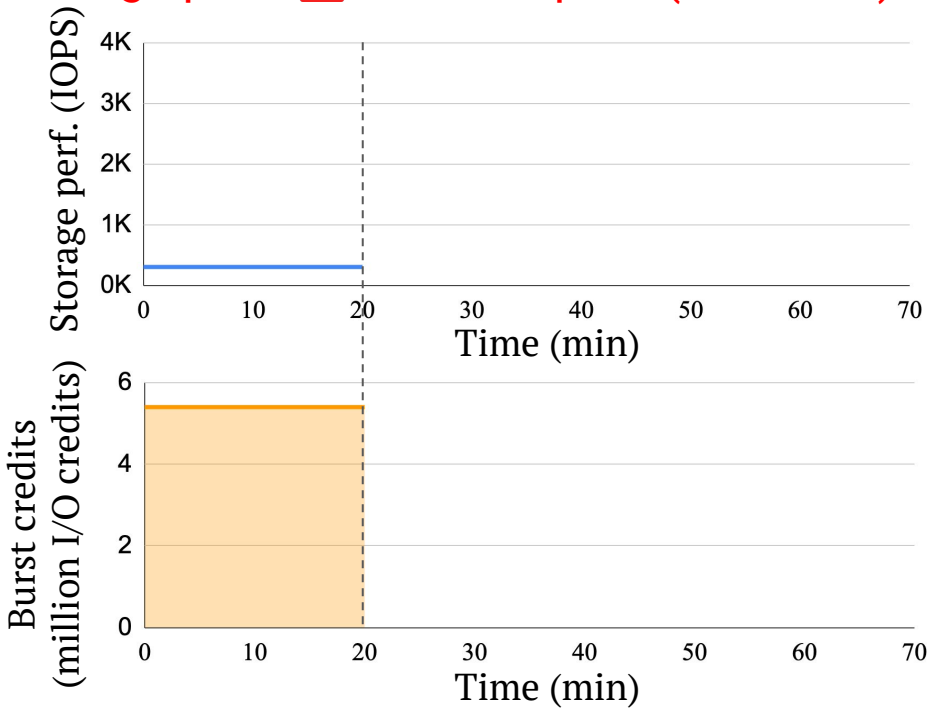

Burst credit bucket

5.4 million I/O credits

AWS Case Study: Burstable storage service

E.g., Using 100 GiB of General Purpose SSD

20 - 53 min: storage perf. ≥ baseline perf. (use burst credits)

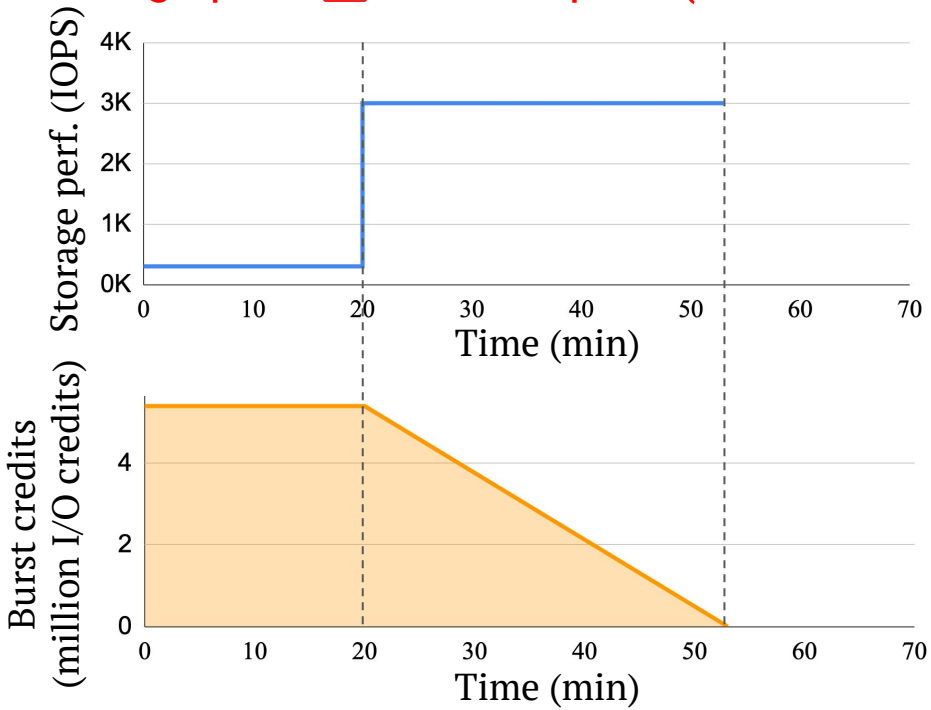

Burst credit bucket

-2700 I/O credits/sec

Lasts for 33 min

AWS Case Study: Burstable storage service

E.g., Using 100 GiB of General Purpose SSD

53 - 70 min: storage perf. is limited by baseline perf. (no burst credits left)

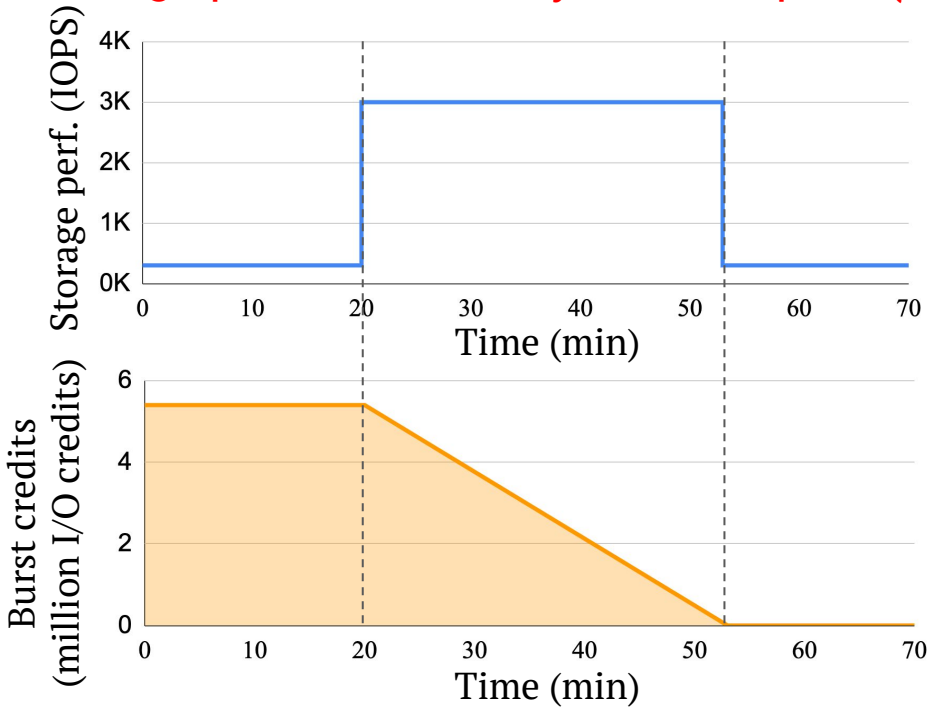

Burst credit bucket

No I/O credits left

Two properties of burstable storage service

New volume starts with **full** burst credits

Burst credits are assigned **per volume**

Two properties of burstable storage service

New volume starts with full burst credits

Burst credits are assigned per volume

Approach: if no credit left, replace with a new volume

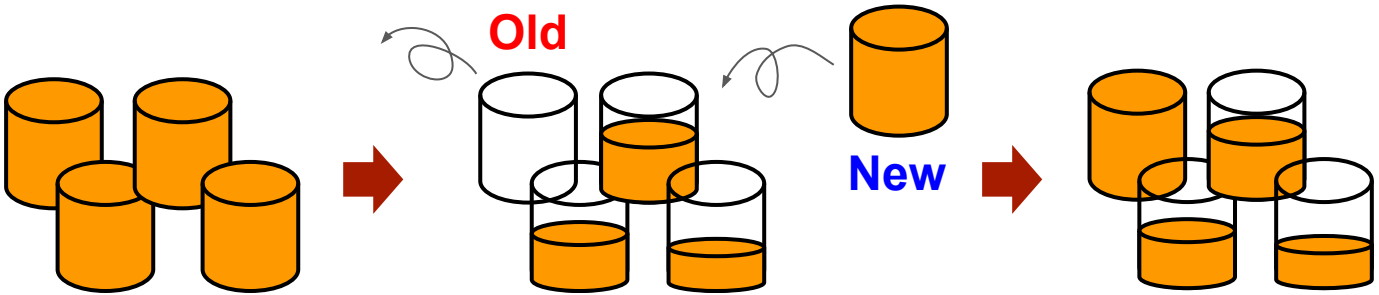

Two properties of burstable storage service

New volume starts with full burst credits

Burst credits are assigned per volume

Approach: split a volume into multiple smaller volumes

Two properties of burstable storage service

New volume starts with full burst credits

Burst credits are assigned per volume

Approach: split a volume into multiple smaller volumes

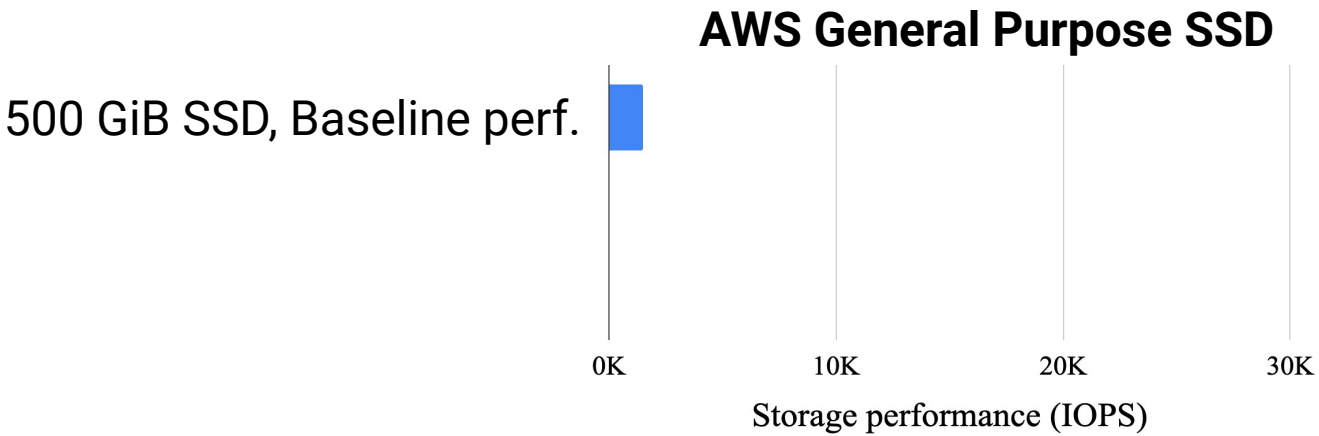

Two properties of burstable storage service

New volume starts with full burst credits

Burst credits are assigned per volume

Approach: split a volume into multiple smaller volumes

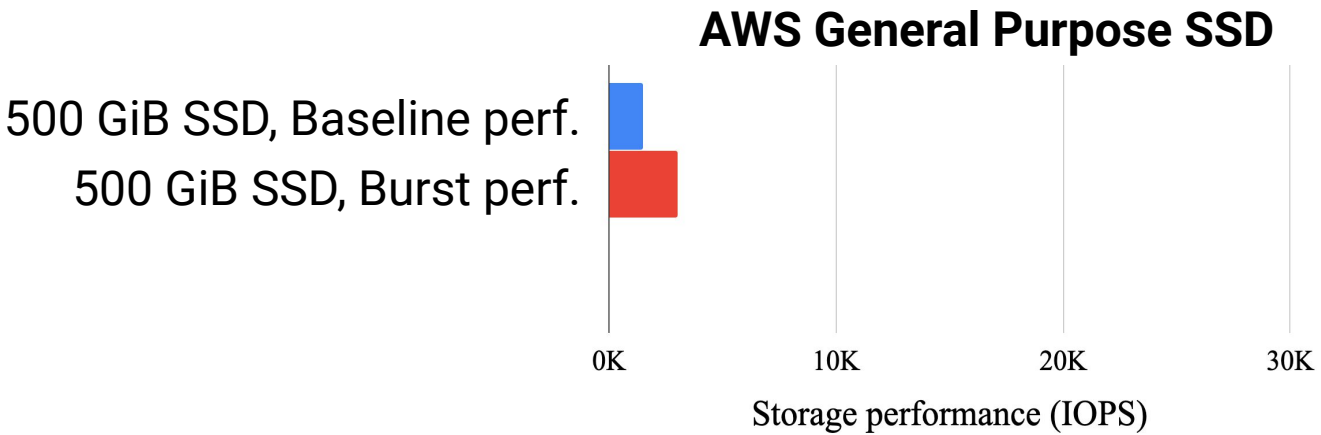

Two properties of burstable storage service

New volume starts with full burst credits

Burst credits are assigned per volume

Approach: split a volume into multiple smaller volumes

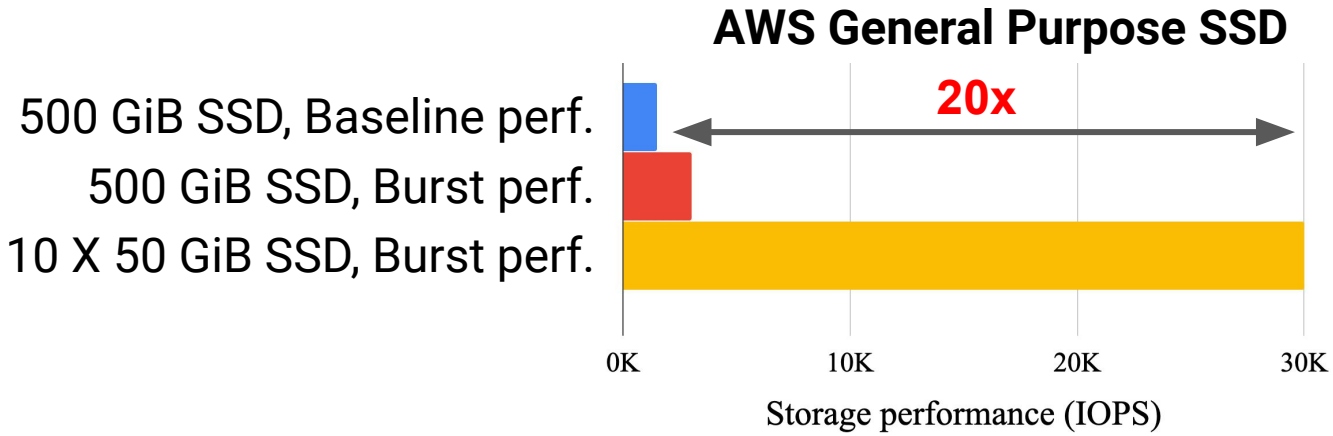

Two properties of burstable storage service

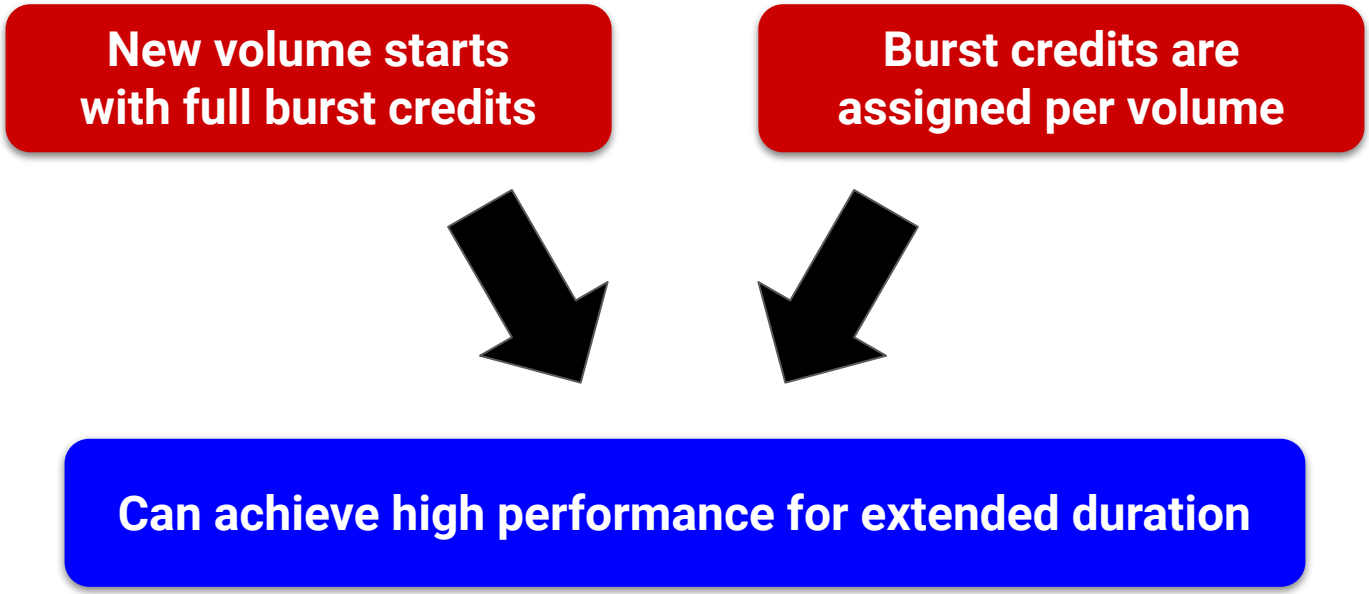

Use cases: exploiting burstable storage

1. Persistent data storage
2. SSD caching (write-through policy)
3. Ephemeral data storage

Use cases: exploiting burstable storage

1. Persistent data storage
2. SSD caching (write-through policy)
3. Ephemeral data storage

See our paper for more details

Use case. Persistent data storage - Design

Step 1. Exploit aggregated throughput of N burstable volumes

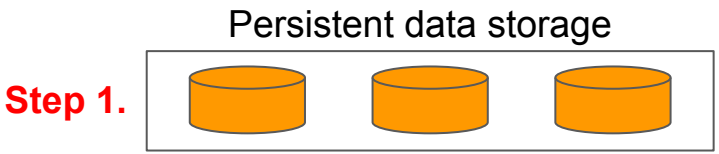

Use case. Persistent data storage - Design

Step 1. Exploit aggregated throughput of N burstable volumes

Use case. Persistent data storage - Design

Step 2. When burst credits run low, create new burstable volumes

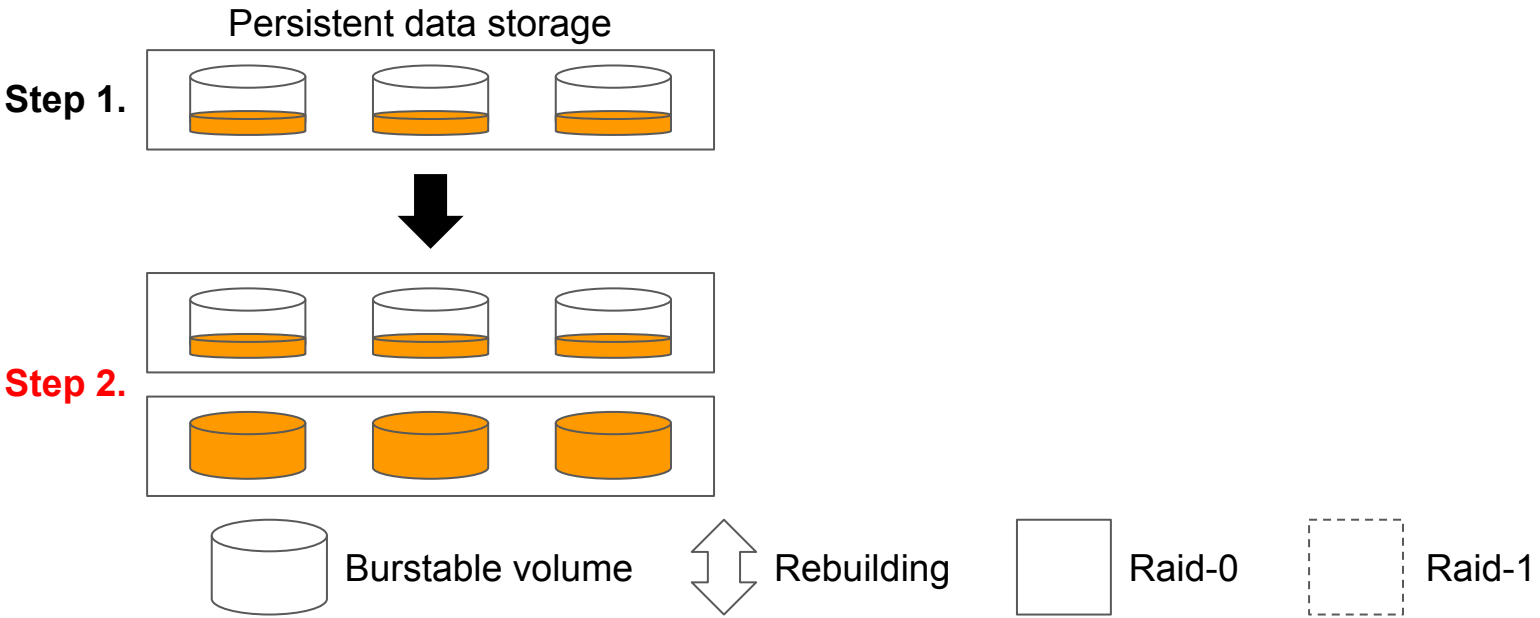

Use case. Persistent data storage - Design

Step 3. Rebuild between two volumes in each RAID-1 pair

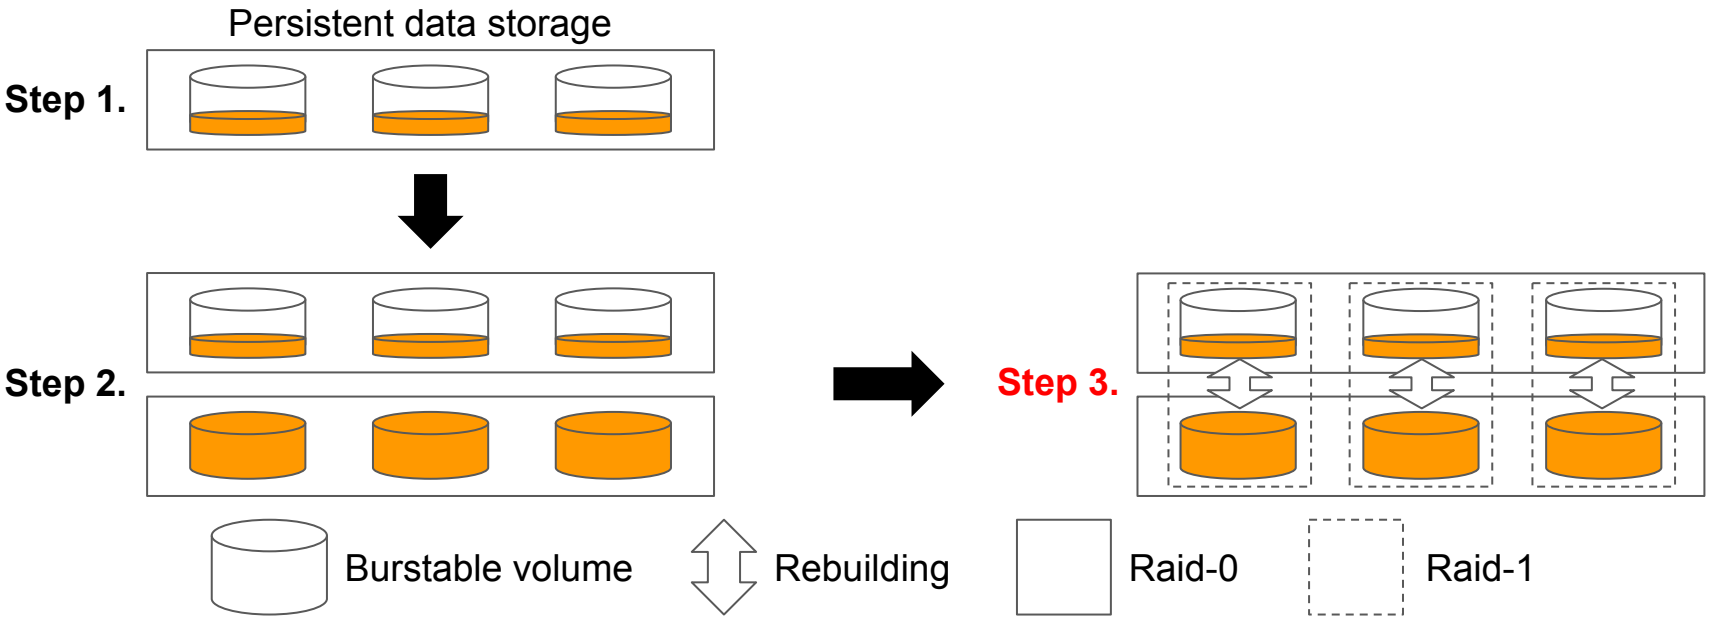

Use case. Persistent data storage - Design

Step 4. Remove the old volumes and use the new burstable volumes

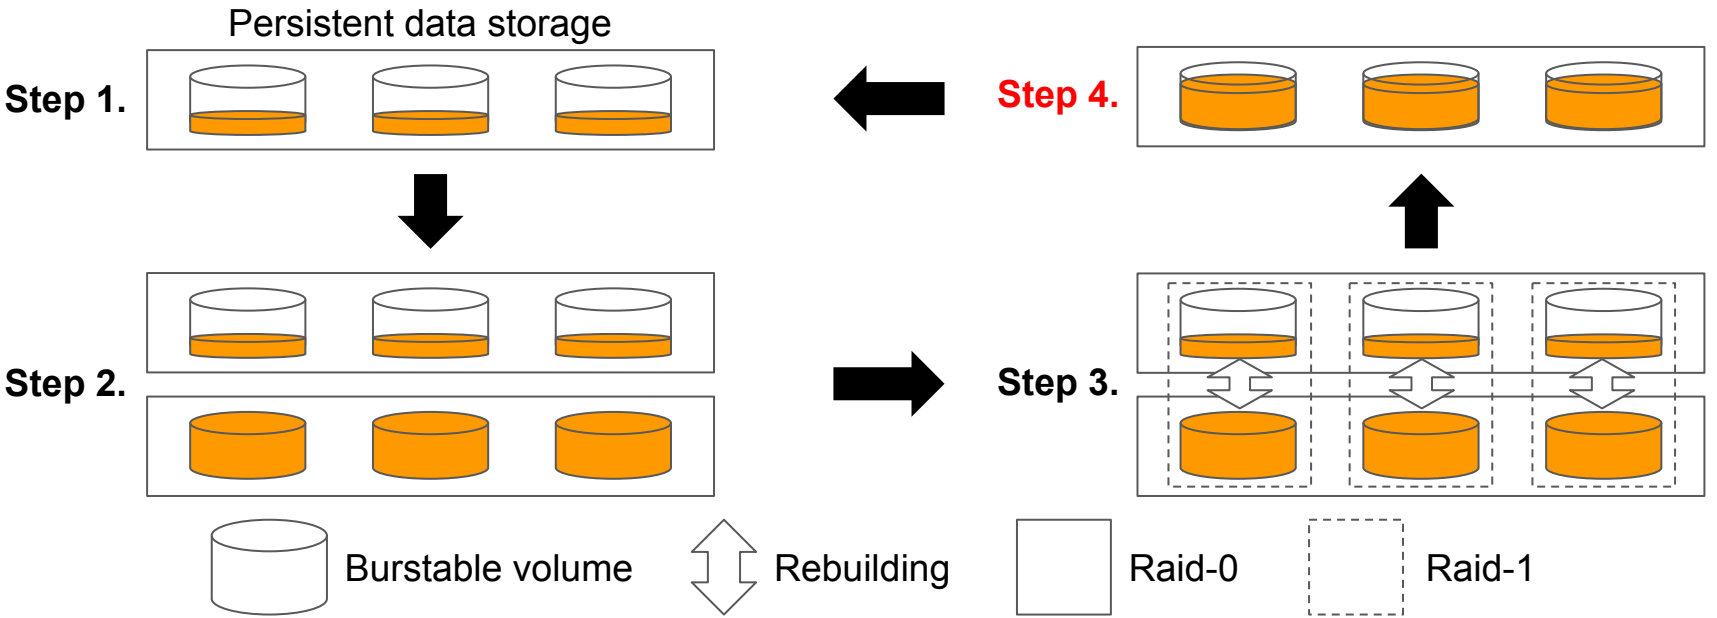

Evaluation: Persistent data storage

1. **How much performance gain we can obtain?**

2. **How much data migration overhead is incurred?**

Evaluation: Persistent data storage

Evaluation Setup

- Evaluated on AWS
  - Instance type : c5.9xlarge
  - Storage type : General Purpose SSD (gp2)

- Performance comparison (Flexible I/O Tester, FIO)

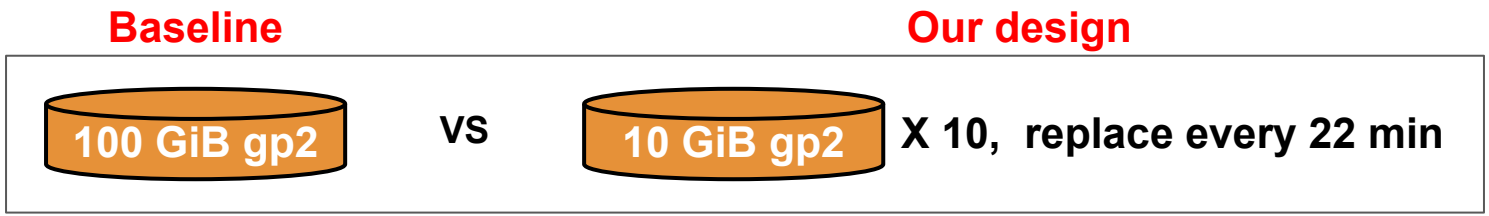

Evaluation: Persistent data storage

Performance comparison between Baseline vs Our design

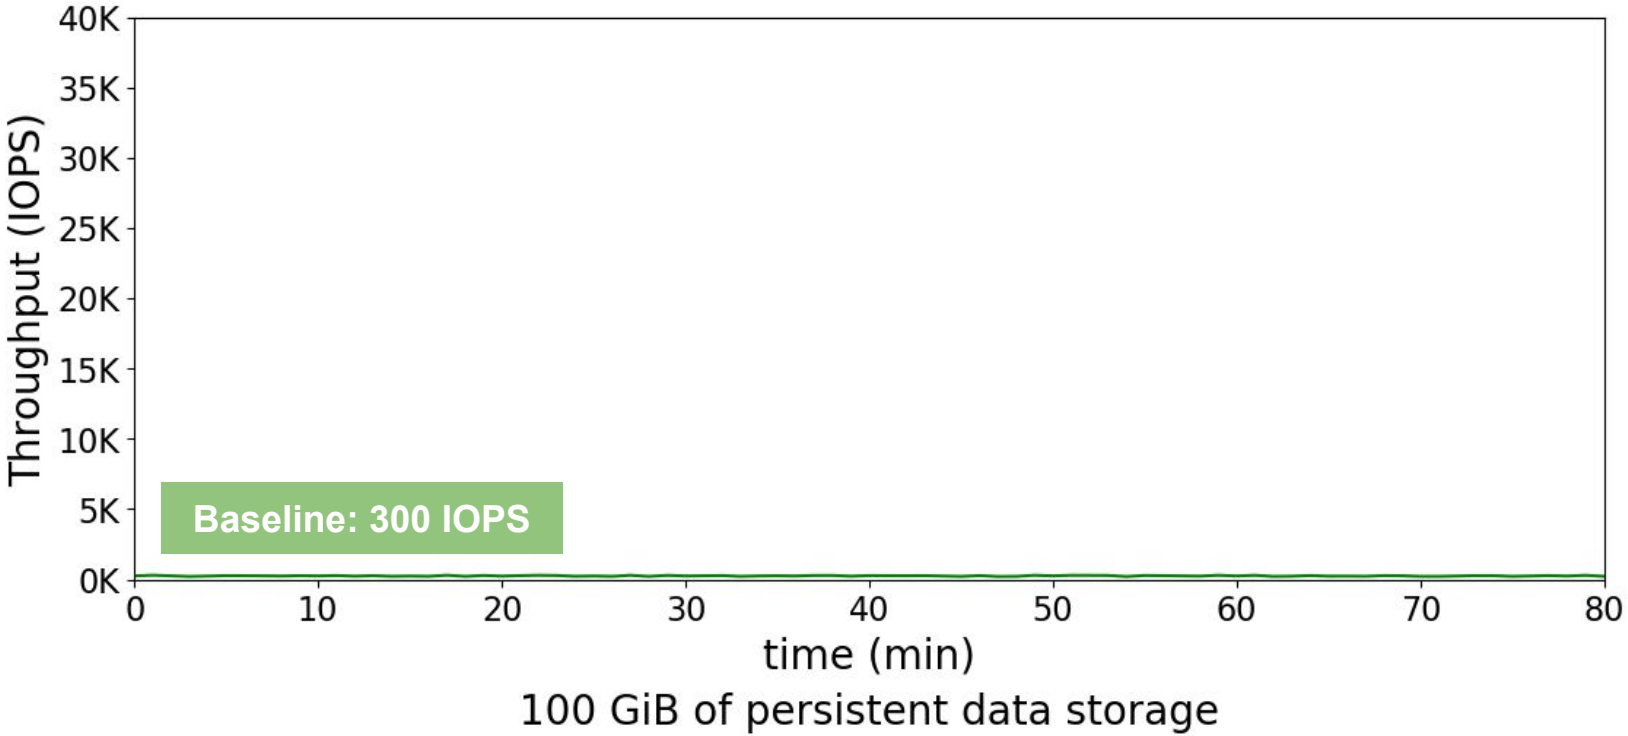

Evaluation: Persistent data storage

Performance comparison between Baseline vs Our design

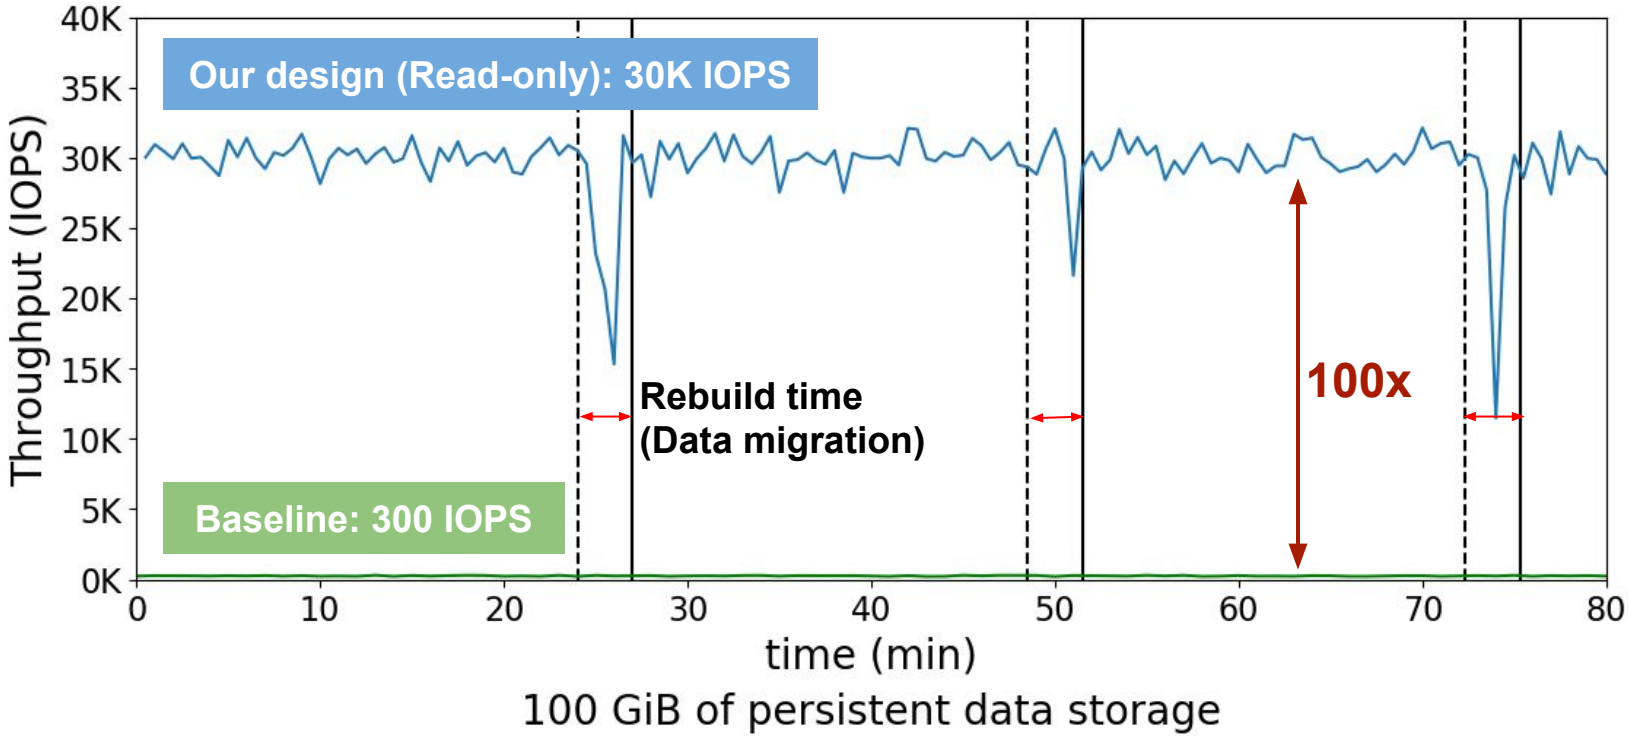

Evaluation: Persistent data storage

Performance comparison between Baseline vs Our design

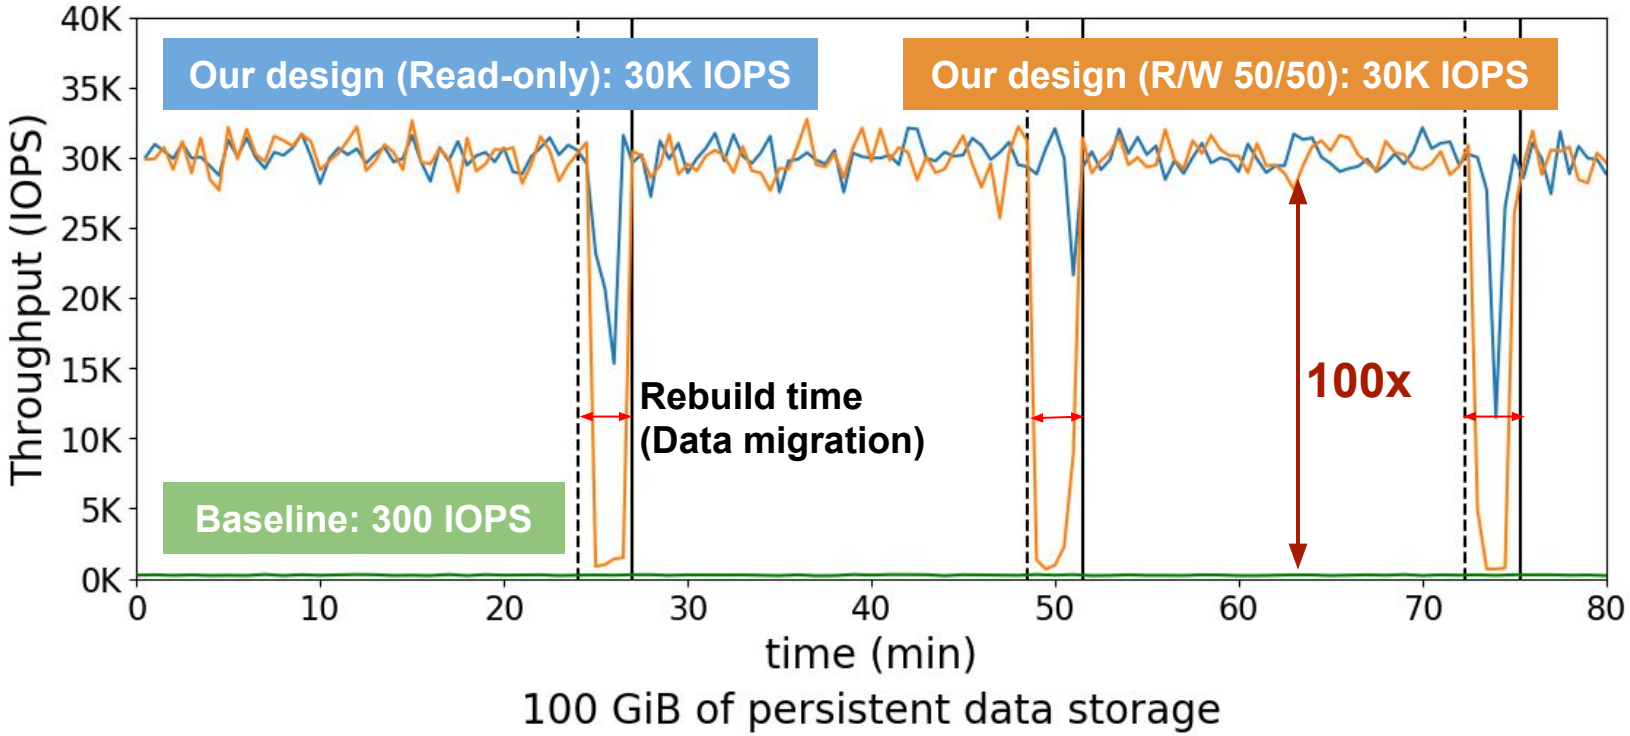

Implications: Persistent data storage

- For 10% additional cost, can get 100X higher throughput
  - 93~99X higher, when accounting for rebuild overhead

- Performance degradation during rebuild time can be mitigated
  - Using more number of smaller capacity volumes
  - Replacing a single volume at a time

Conclusion

- Properties of burstable volumes
  - Every new burstable volume starts with a full burst credits
  - Burst credits are assigned per volume.

- Interchanging small burstable volumes achieves up to 100x higher throughput, at 10% higher price

- Check out our paper for more details!
  - Building an SSD cache using burstable volumes
  - Using burstable volumes as ephemeral storage

Discussion topics

1) Are there any other suitable applications that can take advantage of burstable storage service?

2) How should cloud storage pricing and SLAs be designed to address these unintentional opportunities?

Conclusion

Hojin Park
hojinp@andrew.cmu.edu

- Properties of burstable volumes
  - Every new burstable volume starts with a full burst credits
  - Burst credits are assigned per volume.

- Interchanging small burstable volumes achieves up to 100x higher throughput, at 10% higher price

- Check out our paper for more details!
  - Building an SSD cache using burstable volumes
  - Using burstable volumes as ephemeral storage