Performance Characterization of NVMe Devices with Zoned Namespaces (ZNS)

Krijn Doekemeijer (k.doekemeijer@vu.nl), Nick Tehrany, Balakrishnan Chandrasekaran, Animesh Trivedi
VU Amsterdam

1 Data center storage

Data center storage:
- Digitally stored data will reach **> 1 yottabytes** in 2030
- High performance requirements
- Data centers use **NVMe flash** storage

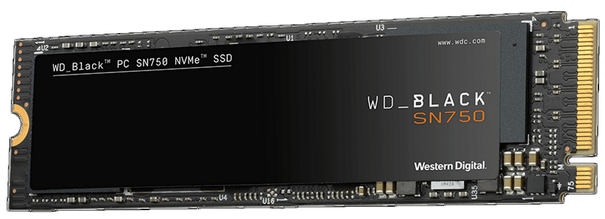

Problem:
- NVMe flash has **unstable** write performance

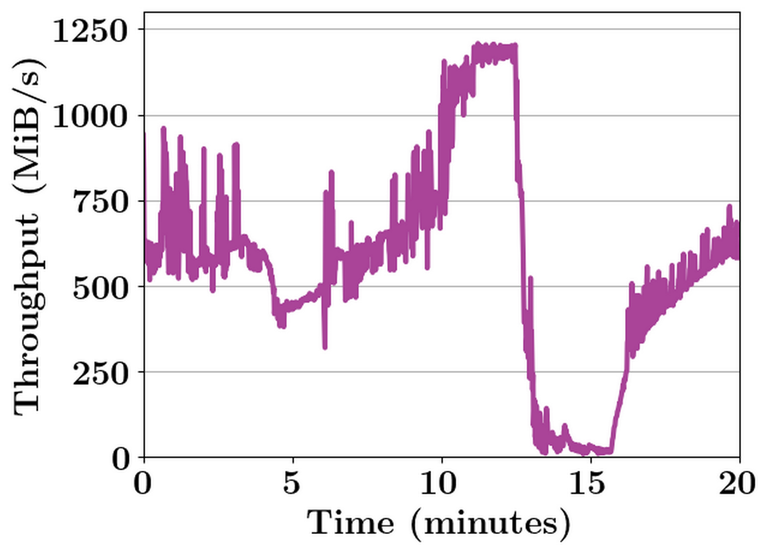

2 ZNS SSDs

- A new industry-backed interface
- **stable write performance**

- I/O issued to **zones** instead of **blocks**

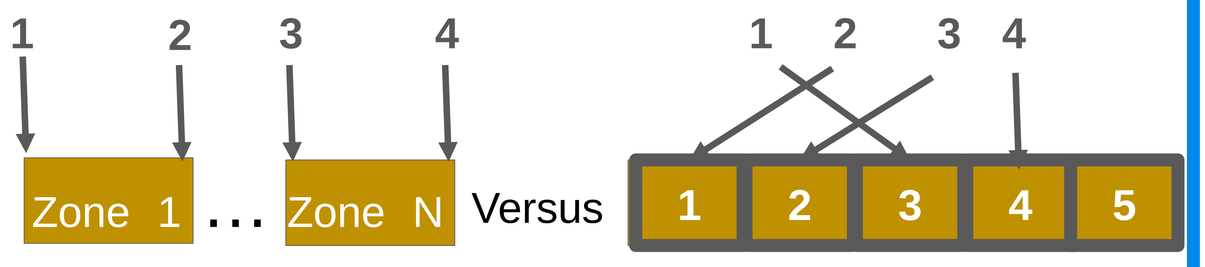

Problem
- **Operational performance properties not known!**
- 4 **new** operations (**Reset, finish, open, close)**
- I/O is issued to large zones (**GiBs**!)
- Does not allow random writes, requiring rethinking scaling

We need a performance characterization!

3 ZNS characterization

We demonstrate 2 of our results, read our paper for the rest!
- Both **append** and **writes** saturate device bandwidth

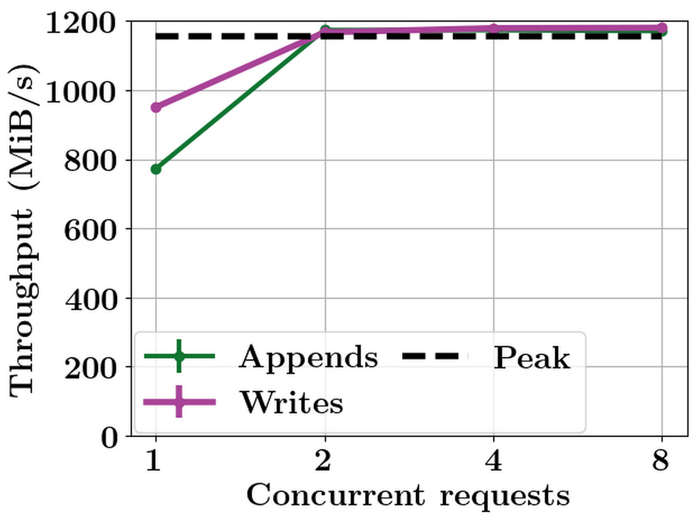

- ZNS new **reset** and **finish** operations are expensive!

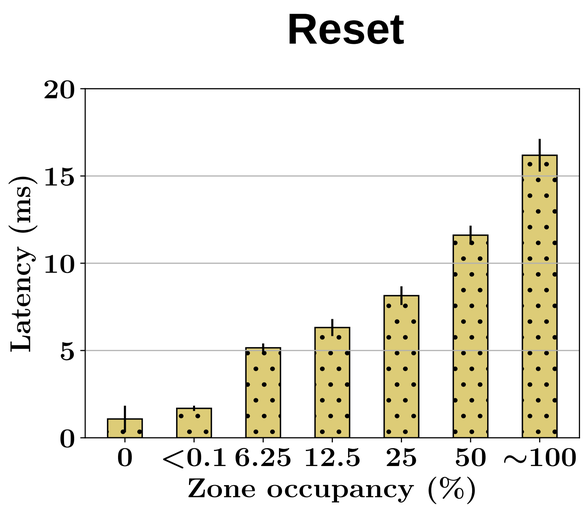

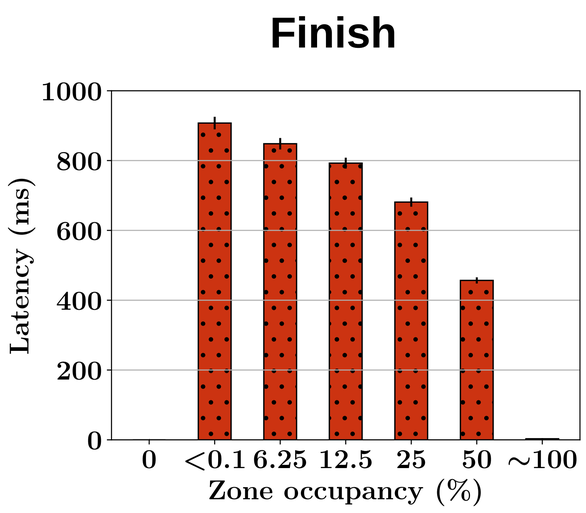

4 Our recommendations

1. **Use ZNS for high-performance stable storage**
2. Use ZNS **writes** for low-latency writes
3. Use ZNS **append** operation to scale writes
4. Avoid ZNS **reset** and **finish** operations

5 What next?

We have shown ZNS on the micro-level:
- Extend to applications
- Evaluate ZNS in a networked setup
- Showcase our recommendations in application design
  - File systems
  - KV-stores

RocksDB