



## CS561 Spring 2026 - Research Project

**Title:** *Analyzing zone size in ZNS SSDs*

**Background:** ZNS SSDs follow the zoned namespace standard where the storage is exposed to the host as zones. The number and size of the zones depend on the manufacturer. The most common commercially available ZNS SSDs are either large-zone or small-zone. The size of the zone presents several tradeoffs. Smaller zone can be reclaimed faster but the intra-zone parallelism may be low as the small zone covers a smaller number of parallel units in the SSD [2,3]. Although the host application can still achieve high bandwidth by striping the writes across the zones. A larger zone size may cause more space amplification, but intra-zone parallelism is higher as the zone spans all the parallel units making the most use of parallelism. Since ZNS SSDs allow a limited number of zones to be open at once, the host can issue the **zone FINISH** command. On the device-side this operation leads to *dummy* writes to the rest of the zone, causing device-level write amplification (DLWA) and interference with host I/O [1]. In small-zoned SSDs, the **zone FINISH** operation may have less impact since there is a smaller number of blocks to write as opposed to the large-zoned SSDs that may cover the full stripe of blocks across all parallel units.

**Objective:** The objective of the project is to benchmark ZNS SSDs with both device benchmarks (like FIO) and applications (like RocksDB) to understand the impact of the zone size on throughput, space-amplification and write amplification.

### Steps:

- Understand the architecture of the ZNS SSD and configuration options of ConfZNS++ emulator [1]
- As an initial experiment run intra-zone and inter-zone experiments where you can measure the write throughput in small-zone and large-zone SSD. Intra-zone experiment would evaluate the scalability of writes within a zone by increasing the request size. Inter-zone experiment would measure the scalability of writes across zones and would run multiple threads concurrently across several zones
- Design experiments that would evaluate space amplification and write amplification using RocksDB

**Responsible mentor:** *Teona Bagashvili*

### References

- [1] Doekemeijer, Krijn, et al. "Exploring I/O Management Performance in ZNS with ConfZNS++." *Proceedings of the 17th ACM International Systems and Storage Conference*. 2024.
- [2] Song, Inho, et al. "Confzns: A novel emulator for exploring design space of zns ssds." *Proceedings of the 16th ACM International Conference on Systems and Storage*. 2023.
- [3] Bjørling, Matias, et al. "{ZNS}: Avoiding the block interface tax for flash-based {SSDs}." *2021 USENIX annual technical conference (USENIX ATC 21)*. 2021.