



# Exploring I/O Management Performance in ZNS with ConfZNS++


Krijn Doekemeijer (k.doekemeijer@vu.nl)<sup>1</sup>, Dennis Maisenbacher<sup>2</sup>, Zebin Ren<sup>1</sup>, Nick Tehrany<sup>3</sup>, Matias Bjørling<sup>2</sup>, Animesh Trivedi<sup>4</sup>  
<sup>1</sup>VU Amsterdam, <sup>2</sup>Western Digital, <sup>3</sup>BlueOne Business Software LLC, <sup>4</sup>IBM Research Europe

## 1 Data center storage

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




- Problems with NVMe flash:**
- I/O-only, hides internal I/O management from host OS...
  - Unstable write performance
  - Scaling up by buying more SSDs is unsustainable...




## 2 Meet ZNS SSDs

- A solution to NVMe's instability?:**
- Expose internal I/O management to the host OS




- Achievable stable write performance!



- Remaining problems:**
- Does I/O management interfere with I/O performance?
  - How do we mitigate this interference?
  - I/O management is specific to SSD designs...
  - Building new SSDs to test designs is expensive...
- We need a tool to explore interference and its impact**

## 3 I/O management interference

- We observed I/O management to **always** interfere with I/O:
- For example, on reads



- Interference is due to SSD resource contention (see paper!)
- Host-issued, so now the OS is responsible...

- We observe interference reductions of up to 56.9% with our I/O management optimizations on the host:
- ZINC: ZNS Interference-aware NVMe Command Scheduler
  - Softfinish: a finegrained I/O management operation

## 4 ConfZNS++ emulator

- How to designs OS software with I/O management in mind?:
- Evaluate performance on various ZNS designs!

- Our solution is **ConfZNS++**:
- the first function-accurate emulator for I/O management
  - Supports 7+ I/O management designs

Function-accurate operation support in ZNS emulators:

Emulator	I/O	Reset	Finish	Mapping
FEMU	✗	✗	✗	✗
NVMeVirt	✓	—	✗	✗
ConfZNS	✓	—	✗	✗
<u>ConfZNS++</u>	✓	✓	✓	✓

## 5 Take-away messages

- SSDs have internal I/O management operations
- I/O management operations interfere with I/O
- ZNS SSDs expose this management to the OS
- ConfZNS++ allows exploring this management (interference) in an emulator for OS software

**We are hiring a scientific programmer to research VectorDBs for NVMe (12 months, MSc)! Please email us if you are interested in working with our team.**