

Western Digital  
and Percona deliver  
Ultrastar<sup>®</sup> DC ZN540  
Zoned Namespace SSD  
support for Percona<sup>®</sup>  
Server for MySQL<sup>®</sup>

## Introduction

Many cloud-scale businesses rely on processing huge volumes of on-line transactions. Many of those workflows are built on MySQL® server, the most popular open source transactional database.<sup>1</sup>

Modern SQL database deployments leverage SSD storage to achieve high transactional throughput and low latency responses. The introduction of the NVMe™ interface for SSDs unlocked the full potential of raw Flash throughput and latency for transactional SQL workloads.



Western Digital introduces the Ultrastar DC ZN540, its first generation of Zoned Namespace SSDs. Zoned Namespaces is an extension to the NVMe protocol that exposes how the SSD stores and organizes the data internally, bypassing the Flash Translation Layer in conventional SSDs. Zoned Namespace (ZNS) SSDs eliminate the internal re-organizing of the data, giving data management applications full control of which bits are stored and updated together.

For use cases that involve concurrent read and write operations this can lead to significantly higher storage efficiency, higher throughput, and lower latency data access.

To tap into the ZNS benefits, data management applications such as MySQL need to be enhanced to leverage the new NVMe ZNS commands. Western Digital has taken the lead in developing and contributing the necessary software components to make the Linux® kernel fully ZNS capable.<sup>2</sup> Moving up the stack, Western Digital has enabled RocksDB and MyRocks to run on a ZNS SSD natively and contributed the software sources to the respective open source communities. MySQL customers running the MyRocks storage engine can now benefit from ZNS capacity, throughput, and QoS advantages.

Western Digital partners with Percona to make those open source components available in the free to download and easy to install Percona Server for MySQL distribution.<sup>3</sup> Western Digital Ultrastar DC ZN540 customers can run Percona Server for MySQL natively, and get Percona support and consulting services for installing, configuring and performance tuning their data management environment.

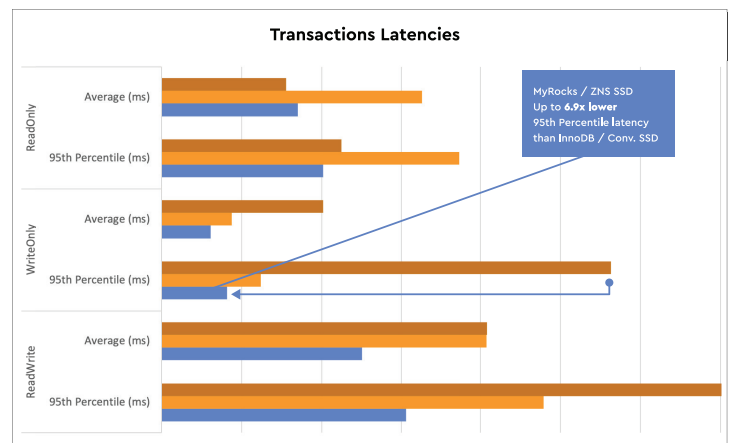
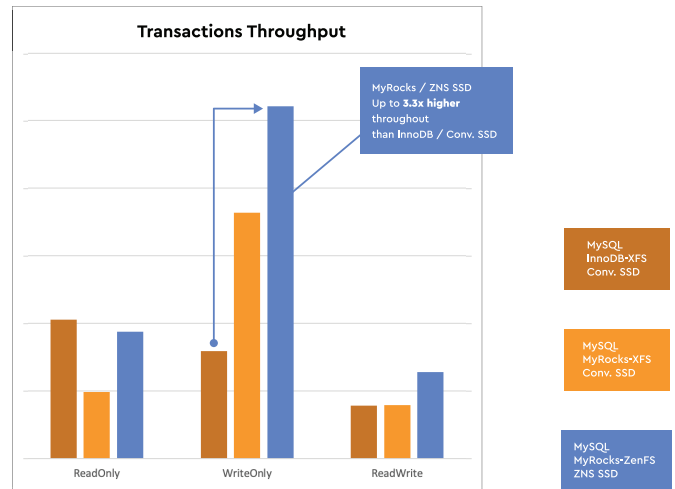
## ZNS performance benefits for MySQL

ZNS SSDs offer performance, latency, endurance and capacity benefits, but also move some of the data management tasks into the application layer. We measured the performance of MySQL server on a conventional and a ZNS SSD to validate that the ZNS benefits would still apply in the full MySQL application stack. The MySQL sever was exercised with a Read-heavy, Write-heavy and a mixed Read-Write sysbench benchmark.

Those benchmarks have been performed with:

- MySQL with InnoDB storage engine running on an Ultrastar DC SN540 conventional SSD
- MySQL with MyRocks storage engine running on an Ultrastar DC SN540 conventional SSD
- MySQL with MyRocks/ZenFS storage engine running on the ZN540 Zoned Namespace SSD

The measured benchmark results show that in the write heavy scenario, the ZNS equipped sever performed up to 3.3 times more transactions per second (TPS) compared to the default InnoDB based server, and 1.4 times higher TPS compared to RocksDB on a conventional SSD. With a mixed read/write load, the server performed 1.6 times higher TPS compared to both server configurations on conventional SSDs.



When considering latency in the write heavy scenario, the ZNS SSD equipped server completed the 95% fastest transactions with 6.9x lower latency than InnoDB and with up to 1.5x lower latency than RocksDB in a sever on conventional SSDs.

<sup>1</sup> <https://db-engines.com/en/ranking>

<sup>2</sup> <https://zonedstorage.io>

<sup>3</sup> <https://www.percona.com/downloads/>

## Enabling the MySQL software stack

The Zoned Namespace Command Set has been adopted in the 1.4 NVMe standard. This command set delivers a standards-based interface to read and write data to SSDs in a way that is very close to how the SSD physically organizes the data in the NAND flash cells. Hence, this is a major enhancement in the efficiency of the data path to the SSD.

For a detailed description of Zoned Namespace SSDs and their advantages, please refer to this earlier blog post: <https://blog.westerndigital.com/zns-ssd-ultrastar-dc-zn540-sampling/>

To enable the software stack, Western Digital contributed the NVMe ZNS kernel block support and tools, which got accepted and integrated in the Linux 5.9 kernel.<sup>4</sup>

For MySQL users to benefit from ZNS SSDs, the MySQL storage engine needs to implement the ZNS NVMe interface. The MyRocks storage engine, based on RocksDB, has log data structures that are well aligned with the NVMe ZNS sequential write requirement. Western Digital created ZenFS, a pluggable RockDB filesystem module optimized for writing to ZNS SSDs. This ZenFS filesystem has been contributed to the open source community.

Through this multi-layer software integration, Western Digital has created a MySQL software stack that natively benefits from the ZNS SSD performance and latency advantages.

## Full-stack software support from Percona

While Western Digital contributed all necessary software components required to make MySQL run on MyRocks with ZNS SSDs to the open source communities, we understand that enterprise customers might require full stack support for operating their mission critical MySQL environments.

Western Digital joins forces with Percona to offer customers full stack support for running Percona Server for MySQL on Western Digital ZN540 devices. Percona integrates MyRocks and ZenFS in their Percona Server for MySQL, and makes it freely available as open source download.

Western Digital customers can buy support from Percona to get support for installing, configuring, operating, and tuning the performance of their MySQL environments. In addition, Percona offers tools for backup, replication and creating clusters of MySQL servers.

---

<sup>4</sup> See Western Digital community site [zonedstorage.io](https://zonedstorage.io)

## Conclusion

We introduced the Western Digital Ultrastar DC ZN540 ZNS NVMe SSD and its use as a storage device for MySQL databases.

As demonstrated in a series of performance benchmarks, the use of the ZNS SSD can result in up to three times higher transaction throughput and six times lower latency response compared to using conventional SSDs in mixed or write heavy MySQL use cases. ZNS SSDs increase throughput and responsiveness for data intensive MySQL database applications and improve multi-tenancy isolation.

ZNS's SSDs give access to the full NAND capacity in the device, which results in up to 25% higher usable capacity compared to conventional SSDs. This efficiency, in addition to the sequential write requirement allow ZNS SSD's to endure 3.5 drive writes per day, compared to 2 and 0.8 DWPD durability for typical conventional SSDs.

The required software and tooling to use Western Digital ZNS SSDs have been contributed to the open source communities. Western Digital joins forces with Percona to make this full stack available as an open source Percona Server for MySQL distribution with extra enterprise tools. For customers that value full-stack enterprise support, Percona offers support packages that cover the use of Percona Server for MySQL with Western Digital ZNS SSDs.

For further details, please contact us at [zonedstorage@wdc.com](mailto:zonedstorage@wdc.com)