SOLID STATE STORAGE TECHNOLOGY CORPORATION

LITE-ON Storage is now Solid State Storage Technology Corporation

Zoned Namespace

A Layout for Efficiency, Endurance, and Performance

Introduction

Zoned Namespace (ZNS) solid state drives (SSDs) are those that utilize the Zoned Namespace Command Set as per the Non-Volatile Media Express (NVMe™) specification. The first (1.0) revision of the ZNS Command Set was ratified within the NVMe™ 1.4a specification, by the NVM Express Organization, in 2020. The ZNS standard is the direct result of industry experience gained over many years working with flash, including previous implementations using Open-Channel SSDs. The primary goal is to wed both the software and hardware ends of storage together in a common interface. This allows the host and SSD to communicate and work together with data placement, improving the efficiency, endurance, and performance of flash media.

The Structure of Zoned Namespaces

Conceptually ZNS shares many characteristics with shingled magnetic recording (SMR) hard disk drive (HDD) technology. SMR works by overlapping magnetic tracks, called shingles, to improve capacity, with the downside that writes are generally done sequentially to avoid a performance hit, making rewriting or overwriting more difficult. This sort of “zoned” configuration fits much better with the way flash works, with block-sized granularity and the benefits of sequential versus random write operations. Specifically, the SSD is broken up into pre-defined zones, each of which is written sequentially with a write pointer for application-specific data placement. The ZNS specification further defines zone size and capacity, using the logical block and flash erase block sizes to align the zones for efficient use.

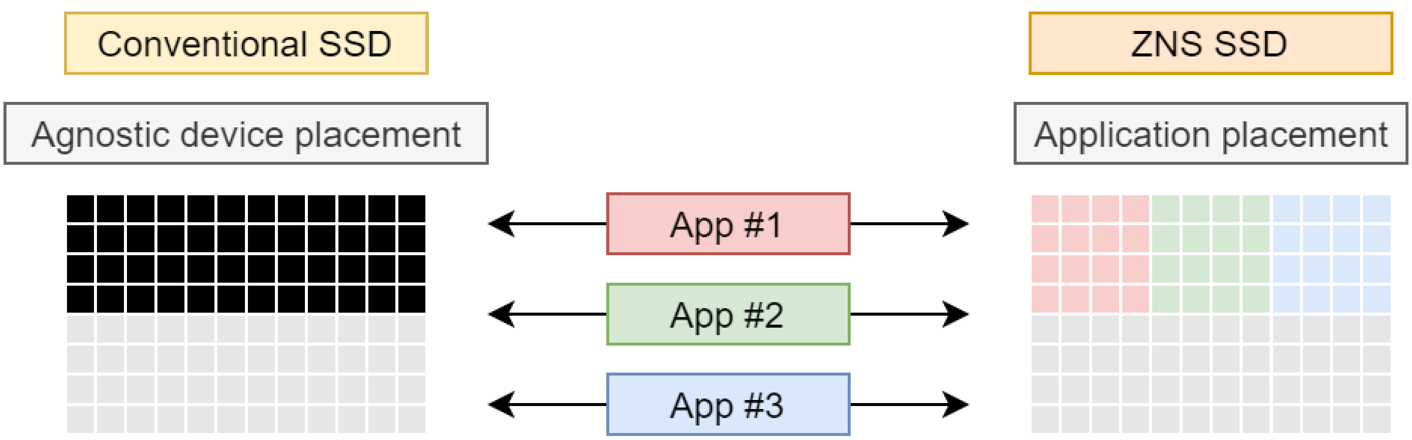

Conventional vs. ZNS SSD Data Placement

The ZNS specification also contains implementations for active zones and Zone Append write operations. Active zones are an additional limit imposed on the SSD controller to make sure applications properly reset or finish active zones before selecting new zones for writing data. This is a form of resource management to give more control to the host. Within the NVMe™ specification, the SSD controller has the ability to reorder host commands but this can lead to poor performance within zones if the zone is limited to one outstanding operation or a queue depth of one (QD1). ZNS introduces the Zone Append command which sets a write position within a zone using logical block granularity with a change to the pointer position as data is written; through this, it is possible to remove the queue depth restriction while still benefiting from the reordering of write requests.

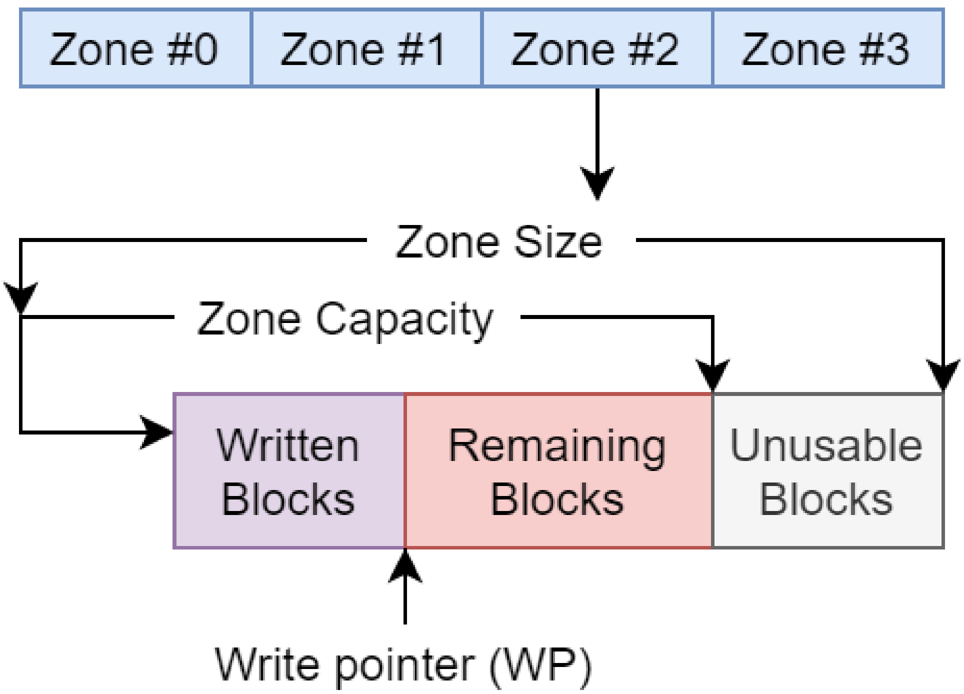

Structure of a Zone

Benefits for Solid State Storage

ZNS offers manifold benefits for flash arrays that touch primarily on the areas of efficiency, endurance, and performance. When developing the specification, the peculiarities of SSDs were considered with the goals of reducing write amplification, reducing excessive over-provisioning (OP), reducing the DRAM load of the flash translation layer (FTL), and reducing quality of service (QoS) variability. All of these relate, to some extent, to maintenance or garbage collection (GC) which when done inefficiently will increase flash wear, reduce performance consistency, and generally make for less cost-effective use of the storage.

Write amplification, or the amount of writes actually imposed on the flash for a given piece of data, varies based on a ratio known as the write amplification factor (WAF). Certain operations, such as random writes, have higher write amplification than others, such as sequential writes. By using zones that configure all writes to be sequential, ZNS greatly reduces the WAF. This improves performance by requiring fewer writes while also being able to use the Zone Append to benefit from re-ordering of write requests. Endurance and drive lifespan is therefore improved, not least because unnecessary rewrites through GC mechanisms can be avoided by proper initial data placement.

This in turn also improves performance and the consistency of response time or latency. By reducing the amount of random I/O, the DRAM load is also reduced. DRAM on SSDs is most often used to store metadata, or data about data, such as mapping or addressing information for data. With simple zone pointers and sequential writes, addressing is greatly simplified through contiguous placement. This is especially true as ZNS allows for block-based granularity which reduces maintenance overhead as flash erases at the block level. Further, with predictable writes the necessity for over-provisioning – that is, the reservation of some space for incoming writes – is reduced or eliminated which increases effective capacity without a corresponding loss in endurance.

Integration and Adoption

As ZNS is still pretty new, organizations are deciding how to adopt and integrate the technology into their storage arrays. This often requires new devices or firmware that support NVMe™ 1.4a with a corresponding update to the software stack – this means an updated operating system (OS). Linux has strong kernel support for zoned block devices including for co-existence with legacy devices. Beyond the hardware and OS, application support is also required to make the best use of ZNS – this means developers have to work towards support for the specification so that the host-controller relationship is leveraged effectively. Luckily, there are plenty of resources available from multiple organizations under the Zoned Storage Initiative.

Summary

The Zoned Namespace Command Set offers flexibility for solid state storage by more intelligently “zoning” space around flash-friendly logical block sizes. Workloads are written in an organized, streamlined manner to reduce wear on the flash without a loss in performance. Traditional flash roadblocks, such as high DRAM use, massive over-provisioning, and inconsistent latency from garbage collection, are greatly alleviated through the use of zones. ZNS offers both active zones for host control and the Zone Append command to effectively reorder write requests. This offers improved efficiency, endurance, and performance for supported devices.

All-Flash Arrays (AFA) moving forward will adopt the NVMe™ 1.4(a) specification including ZNS with OS and application support for a unified, standardized architecture. Hardware costs will be reduced through higher effective capacities and less demand for DRAM, with greater endurance through the reduction of write amplification. Performance in terms of latency is improved with streamlined garbage collection and more consistent response times. Through this, technology developed for HDDs can be more efficiently applied to SSDs with incredible improvements for every aspect of the storage.

ABOUT US

A subsidiary of KIOXIA Corporation, **Solid State Storage Technology Corporation** is a global leader in the design, development, and manufacturing of digital storage solutions. We offer a comprehensive lineup of high-performance customizable SSDs for the Enterprise, Industrial, and Business Client markets. With various form factors and interfaces, our SSD solutions help businesses simplify their storage infrastructures accelerating variable workloads, improving efficiency, and reducing total cost of ownership.

Learn more at www.ssstc.com