TELEDYNE LECROY
Everywhereyoulook™
Protocol Solutions Group

Testing Zoned Namespace SSDs with SVF Pro/Enduro

WHITE PAPER

Rob Dobson, Director Product Marketing, OakGate Products

Date: Mar 2021

Summary

Several key SSD suppliers have Zoned Namespace solid-state drives under development and 2021 is primed to be the breakout year for ZNS SSDs.

In fact, customers are already using Teledyne LeCroy's industry-leading validation platform SVF Pro/Enduro to test their ZNS-ready drives.

SVF Pro/Enduro has three options for testing ZNS – Exerciser, Directed Tests, Automation. Our ZNS feature requirements were developed with the inputs from a number of our current customers.

Our full-featured validation software tests the majority of the NVMe ZNS Command Set 1.0 specification. Our ZNS functionality supports all the Zoned Admin Command Sets and Zoned I/O Command Sets defined in the current specification.

Introduction

To improve the performance and data placement capabilities of SSDs, the NVMe Technical Work Group developed Zoned Namespace (ZNS), a 2020 industry standardization that aligns the internals of SSDs with the Host. Zoned Namespaces was the result of a multi-year effort by the storage industry to exploit the latest understanding and developments of Open-Channel SSD architectures which hyper-scalers, all-flash array vendors, and large storage system vendors have been considering or currently use.

Zoned Namespace divides the logical address space of a namespace into zones. Each zone provides an LBA range that must be written sequentially and if written again must be explicitly reset. This allows the creation of namespaces that expose the boundaries of the device and offload management of mapping tables to the host. And it enables host software to direct I/O traffic to specific zones.

Challenges with Conventional Flash Storage

The challenge of using flash memory for a Solid-State drive is that all computers are built around the concept of how hard drives work**.** Flash memory doesn't behave like a hard drive**.** Solid-state Storage Devices (SSDs) consist of memory technology without moving parts, so they do not have the same behavior characteristics as magnetic and mechanically-based Hard Disk Drives (HDDs).

There are a **number of challenges** with using Flash Storage including:

- Data is written sequentially and must be erased before being re-written
- Data is written at the page level, but erased at the block level
- An SSD supports only a limited number of program/erase (P/E) cycles before memory cells start to fail
- If not managed, these issues can negatively impact SSD endurance and shorten its lifespan
- To offset these challenges, most SSDs include a flash translation layer (FTL) that handles life-extending tasks such as garbage collection and wear leveling

Zoned Namespace 101

As stated previously, Zoned Namespace divides the logical address space of a namespace into zones. Each zone provides an LBA range that must be written sequentially and if written again must be explicitly reset. This allows the creation of namespaces that expose the boundaries of the device and offload management of mapping tables to the host, which empowers host software to direct I/O traffic to specific zones. Figure 1 highlighted below, illustrates how a current conventional SSD controller and a ZNS SSD controller manage the same LBA space.

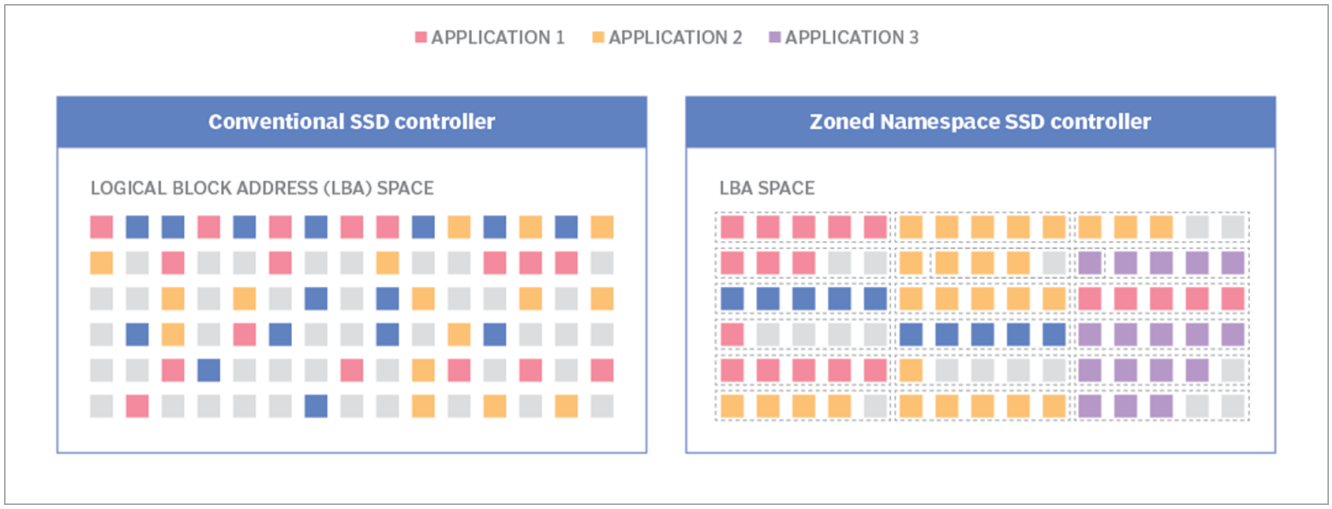

Figure 1 – Conventional vs ZNS SSD Controllers (image provided by TechTarget)

Zoned Namespace SSDs achieve improvements by exposing a namespace logical address space using zones which must be written sequentially and explicitly reset before rewriting. Each zone has a set of attributes reported in the Zone Descriptor data structure. Zone type is an attribute that defines the rules for reading and writing to a zone. Each zone also has an associated state machine – see Figure 2 highlighted below. And each state together with the zone type define the operational characteristics of each zone.

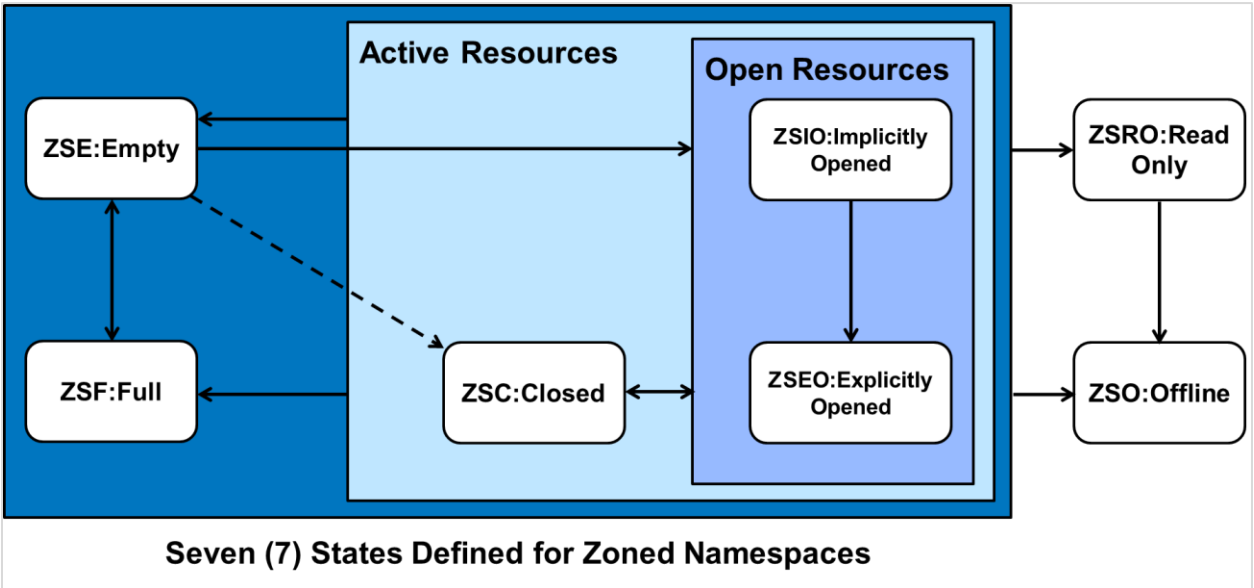

Figure 2 – Zoned Namespace State Machine

In the zoned storage where the host software is more responsible for data placement, the host needs to keep track of each zone's state. Zones in empty and full states are self-explanatory – either no data is written, or zone is at its full capacity. The read-only and offline states come to play only when a drive's flash is failing, and the drive is at the end of its life. Active zones are zones in implicitly open, explicitly open or closed states. The drive usually puts limitations on the number of active zones due to required extra tracking information.

Valid transitions between each state can be managed by host through direct commands or triggered by device with notification to host through an exception. The transition from empty to closed state, as indicated by the dotted line, is only valid when the zoned namespace is formatted with zone descriptor extension support.

Key Benefits of Zoned Namespace

ZNS NVMe solid-state drives provide **three key benefits** over traditional NVMe SSDs.

1. **Reduced write amplification**
   Write amplification is an undesirable effect associated with solid-state drives, where the actual amount of data physically written is a multiple of the logical amount intended to be written. The interface does not allow the host to understand the physical characteristics of the media within the SSD, and often writes inefficiently to the device.

   By using zoned namespaces, the device can expose a set of zones with the requirement that each zone must be written sequentially, matching the physical SSD media requirements. This allows the host and device to collaborate and place data efficiently on the SSD

2. **Reduced media over-provisioning**
   Conventional SSDs overcome the shortcomings of the narrow read/write interface using physical media over-provisioning. With over-provisioning, extra space has to be reserved for moving data around for garbage collection and to improve its efficiency.

   By using zoned namespaces, the host places data efficiently, and therefore, over-provisioning is no longer necessary. The lower over-provisioning improves the cost of the device as more storage space is exposed to the host.

3. **Reduced internal controller DRAM usage**
   A conventional SSD controller manages a fully-associative mapping table used to define the mapping of the host visible logical blocks to a physical address on the media. For performance reasons, the table generally is about 1GB in size per 1TB of media and is stored in the controller DRAM of the SSD.

   By using zoned namespaces, ZNS SSDs expose logical blocks through zones that can only be written sequentially. Consequently, the device controller can now reduce the size of the mapping table while still providing the same benefits of being able to move data internally, without notifying the host.

These reductions **should** lead to:

- Improved throughput and latency
- Potential savings due to decreased need for over-provisioning of NAND media

Testing Zoned Namespace Solid-state Drives with SVF Pro/Enduro

SVF Pro/Enduro has three options for testing ZNS – Exerciser, Directed Tests, Automation. The software requirements were developed with the inputs from a number of our current customers. The full-featured software tests the majority of the ZNS Command Set 1.0 specification.

The ZNS capabilities supports all the Zoned Admin Command Sets and Zoned I/O Command Sets defined in the current specification. Customers can utilize our advanced Enduro GUI and Exerciser module to completely test ZNS solid-state drives under a full I/O load. ZNS commands are also available through our command line interface, Directed Test API interface, and Python API interface for test automation.

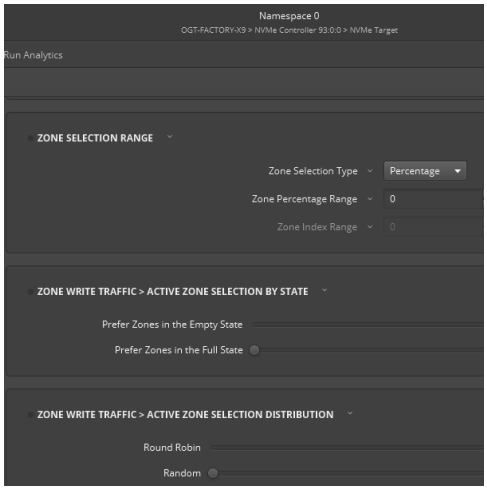

Figure 3 – Exerciser ZNS Settings

The SVF Pro/Enduro **Exerciser Module** is designed to stress the ZNS DUT by running concurrent I/O traffic. The user can easily generate high-performance, randomized traffic profiles with I/O profiles of ZNS test scenarios that would be extremely difficult to create manually. There are numerous new ZNS specific settings that allow users to choose range of zones, number of active zones, read and write distribution, and zone fill percentage for writes. See Figure 3 for an example of some of the ZNS parameters in the Exerciser Module.

The **Directed Test Module** provides a comprehensive testing for Zoned I/O commands and Zoned Admin Command Set. A Directed Test (DT) is a sequence of commands sent to the target with the expectation that the target returns a specific sequence of responses.

Our ZNS Directed Test Scripts test the correct positive and negative responses expected from the DUT. Additionally, the Directed Test module tests for exceeding active and open resources, zone state machine and specific to ZNS errors. Widgets can be defined to create a customized test and detailed logs list all test results. See Figure 4 for an example of the ZNS tests defined in the Directed Test module.

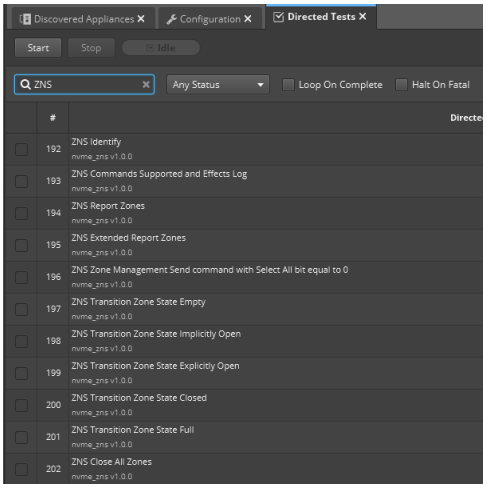

Figure 4 – ZNS Directed Tests

Users can also utilize the **Automation Module** to create specific or unique ZNS test cases. The test set can include multiple tests, actions, and loops. Each of the individual ZNS tests can run with configured ZNS Exerciser settings and/or run ZNS Directed Tests. The automation suite runs in a predetermined order without user intervention. See Figure 5 for an example of our Automation Module that allows customers to fully automate their ZNS test plan.

Key Features

- Customized Linux driver to handle ZNS state machine transition and sequential write requirement in each zone
- Multiple threaded I/O running in the zones of ZNS in parallel with high throughputs. Each zone can be tested using write, read, compare, and append as needed. Each zone will be reset at the start, and then later when finished at the end
- Configurable parameters and widgets to change settings to accomplish different results
- Namespace management for ZNS
- Negative testing through scripts to test all ZNS features

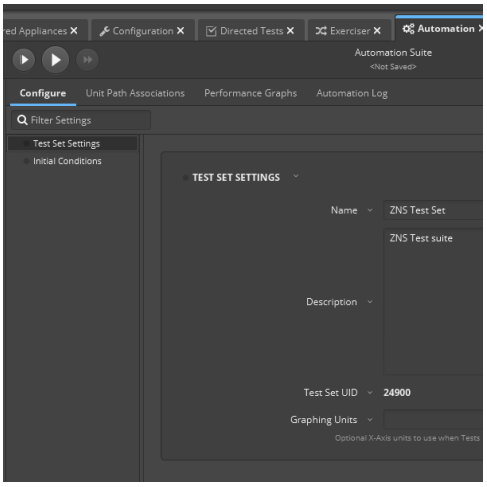

Figure 5 – Automation Module

Real-world ZNS Test Results in SVF Pro/Enduro

In the section below, we have provided some real-world test results highlighting a number of Zoned Namespace commands. Figure 6 below is an example of the SVF Pro Analyzer highlighting the correct drive response – NVME_ERR_CMD_SPEC_ZONE_BOUNDARY_ERROR – to issuing Write to Zone command with number of blocks going over zone capacity.

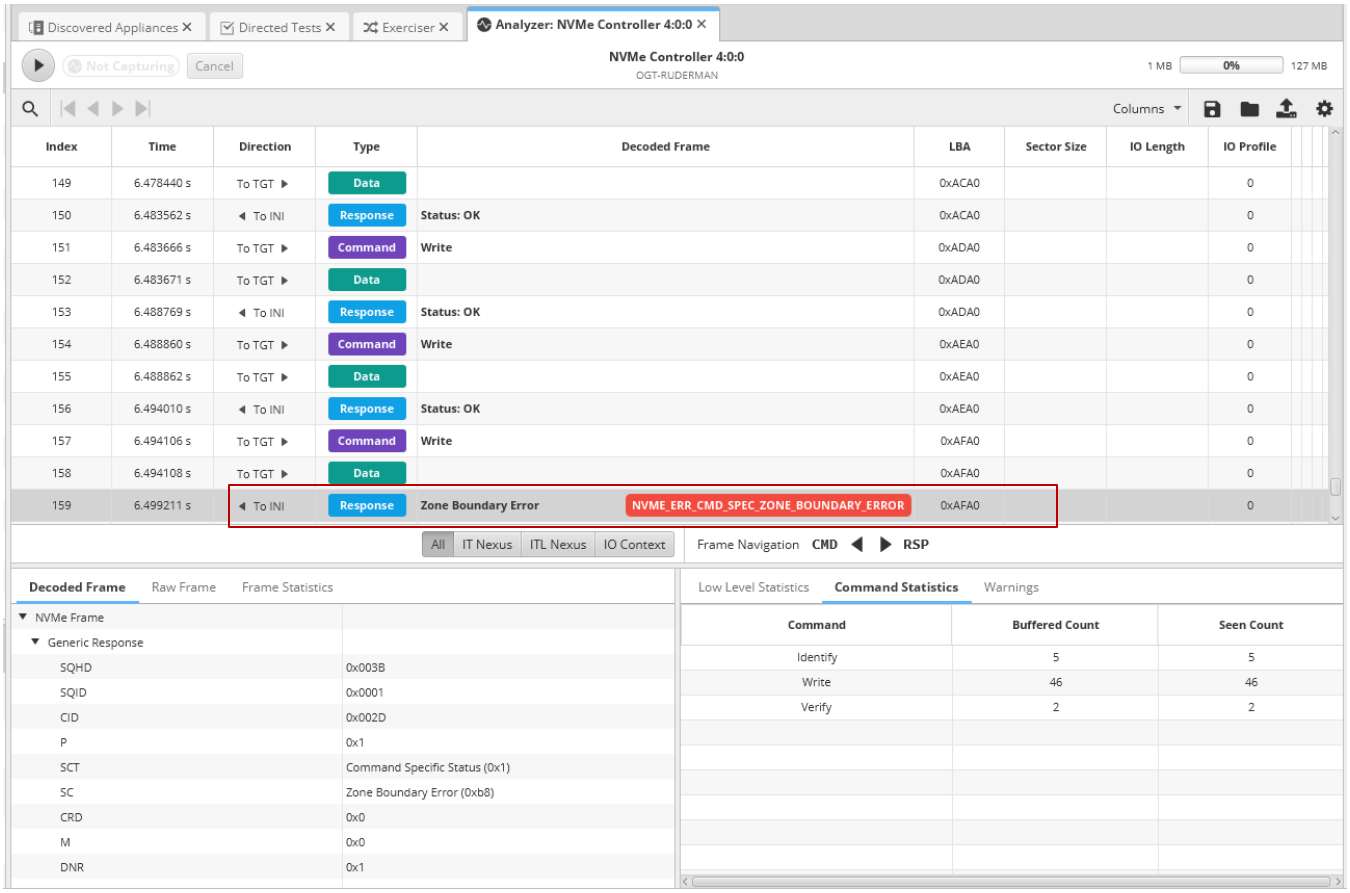

Figure 6 – Zone Boundary Error in Analyzer Module

Figure 7 is an example of the SVF Pro Analyzer with the correct drive response to issuing Zone Management Send commands.

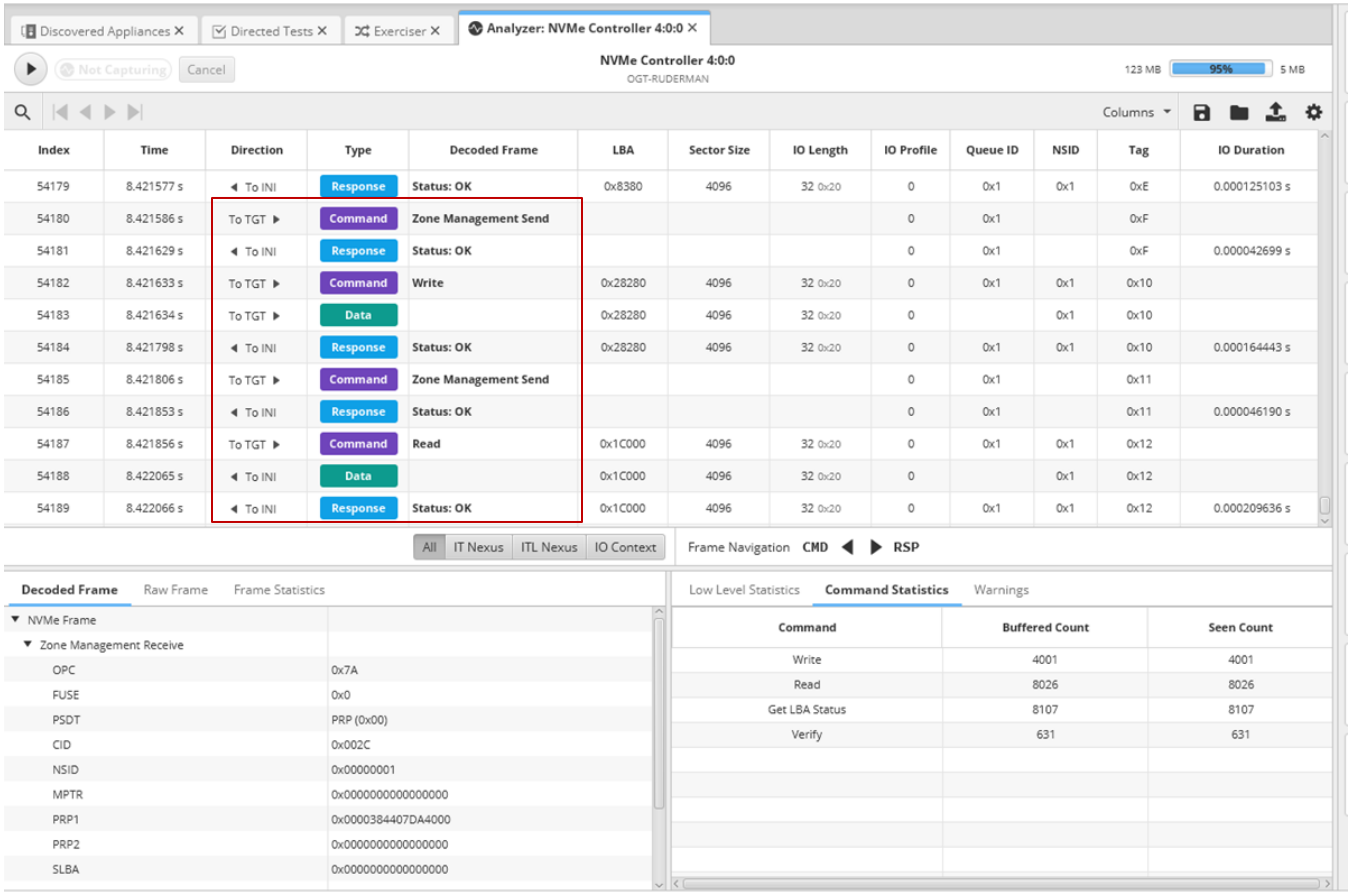

Figure 7 – Zone Management Send in Analyzer Module

Figure 8 is an example of ZNS Report Zones Directed Test output with Zone Receive Action Specific Field set to All Zones. As you can see zones are in different states, with the same zone type and zone capacity and write pointer pointing to different LBAs.

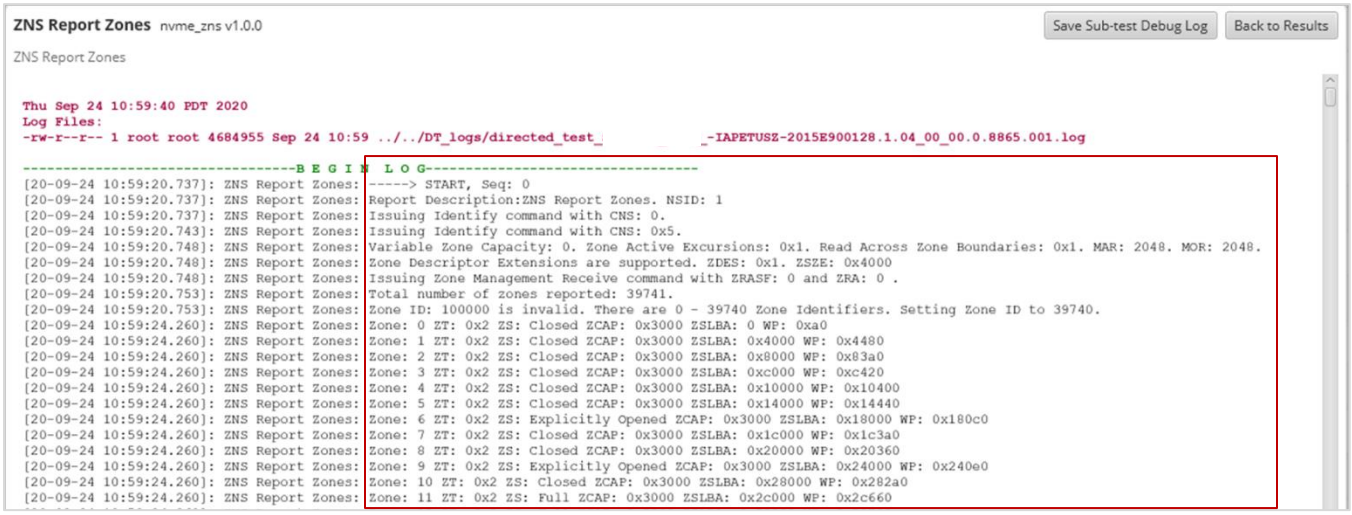

Figure 8 – ZNS Report Zones Directed Test Results

Figure 9 is an example of a ZNS Post-run Analytics chart highlighting a ZNS-enabled "prototype" SSD with high average write and read latencies. Our customer used the learnings to fix this in subsequent iterations of firmware.

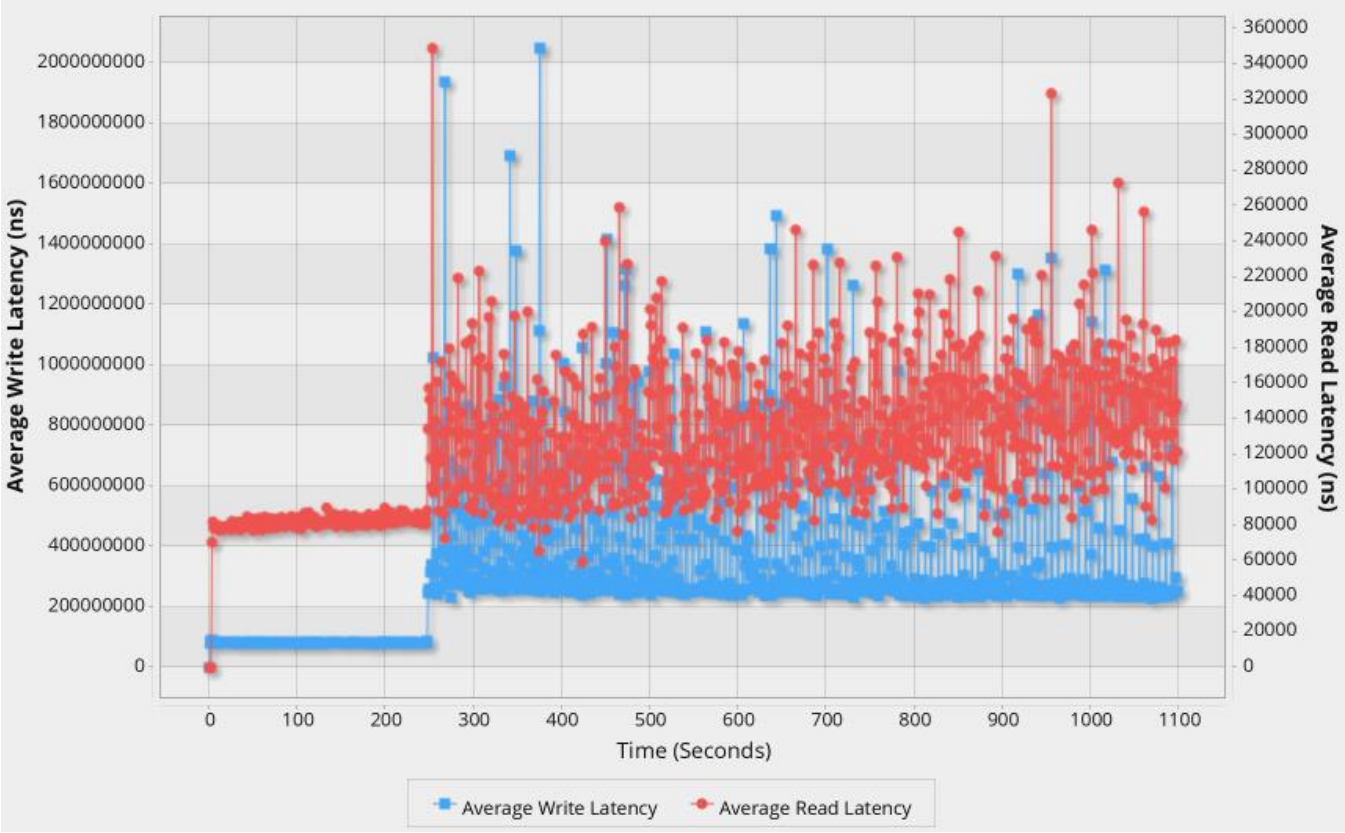

Figure 9 – Post Analytics Chart – Average Write and Read Latency

Summary

Teledyne LeCroy provides the Industry's most comprehensive Zoned Namespace support. The ZNS feature requirements were developed with input from a number of our current customers and are compliant with the NVMe ZNS Command Set 1.0 specification.

SVF Pro/Enduro offers **three options for testing ZNS** – Exerciser, Directed Tests and Automation.

1. With the Exerciser, easily generate high-performance, randomized traffic and I/O profiles for limitless ZNS test scenarios
2. With Directed Tests, run a complete set of "push-button" ZNS NVMe commands; write your own as well
3. With Automation, create customized, automated scripts that can run with configured ZNS Exerciser settings and/or run ZNS Directed Tests

Additionally, our **customized, proprietary driver talks directly to the ZNS SSD hardware** providing highest level of performance and low latency which provides a consistent and repeatable test performance

For more information and/or demonstration of SVF Pro/Enduro and other Teledyne LeCroy products, visit https://www.teledynelecroy.com/oakgate or contact us by email at oakgate_marketing@teledyne.com