White Paper
Data Center
December 2023

Intel® Optane™ SSDs Deliver 42% Higher IOPS for Mixed VMware vSAN 8 Workloads Compared to NVMe NAND SSDs

The SSDs' excellent value for mixed vSAN 8 workloads combines higher IOPS performance with 29% lower latency and 35% more performance per watt

Authors

Caiti Dunn
Cloud Solutions Engineer
Intel Optane Group

Ken LeTourneau
Cloud Solutions Architect
Intel Optane Group

Table of Contents

Executive Summary 1
Introduction 2
A Closer Look at Today's vSAN 2
Intel Optane SSDs Provide Value for Both Existing and New vSAN Deployments 2
Testing Methodology 2
High, Predictable Performance for vSAN OSA 3
Accelerate Everyday Workloads 3
Continue to Meet SLAs Even When Write Pressure Escalates 5
Endurance Considerations 7
Conclusion 7
Learn More 7

Executive Summary

Quickly moving data between the CPU and storage is a key metric for data center efficiency and also underpins business outcomes such as fast time to insight, customer experience and quality of service. Intel tested the effectiveness of vSAN 8 cache performance, monitoring IOPS, performance/watt and latency by comparing an Intel® Optane™ SSD to a 3 drive writes per day (DWPD) NVMe NAND SSD.

The Intel Optane SSD excels at the most common enterprise workloads (42% more IOPS than a NVMe NAND SSD with 29% lower latency and 35% more performance/watt) as well as in the write pressure tests (83% more IOPS with 45% lower latency and 72% more performance/watt).[1] What's more, the Intel Optane SSD's 100 DWPD endurance rating makes it a better choice for write-heavy scenarios like the vSAN cache.

With the volume of data across the world doubling every two years (currently sitting at about 64 zettabytes[2]), storage systems like vSAN need to be able to handle more and more storage operations. Intel Optane SSDs represent a smart data center investment:

- High and predictable performance with low latency.
- Excellent read performance in the face of growing write pressure.
- Bidirectional capability that enables simultaneous reads and writes.
- High endurance.

Intel® Optane™ SSDs are a great choice for the vSAN 8 cache tier.

Introduction

In today's hyper-competitive business environment, enterprises need their data center workloads to run as efficiently as possible. Investing in the right infrastructure components, such as solid-state drives (SSDs) and CPUs, can contribute to the two primary categories of data center efficiency:

- **High performance.** Compute-hungry AI is being woven into nearly every application in almost every industry. A deluge of data threatens to drown AI and analytics applications; however, the faster that workloads are completed, the more insights an enterprise can uncover to drive business decisions and value. Modern enterprises need high I/O operations per second (IOPS) and low latency for a variety of other workloads as well. Examples include digital twins in manufacturing, automated network security, genomics research, as well as periodic high-storage-stress situations like substantial workload migrations and the deployment of a new workload with supporting TBs of data. All these have a common denominator: the need for speed.
- **Cost effectiveness.** IT budgets remain flat, but expenses continue to rise. Getting the most value out of every infrastructure investment means considering SSD endurance as well as an infrastructure's performance/watt. Increased performance/watt enables enterprises to achieve the same performance for less power, thereby saving on energy bills. Better endurance also translates to fewer drive replacements. In a nutshell, choosing the right SSD can substantially lower total cost of ownership (TCO).

A Closer Look at Today's vSAN

Many enterprises—more than 30,000 according to VMware—have turned to hyperconverged infrastructure (HCI) powered by VMware vSAN to help lower TCO, accelerate operations and pave a path to hybrid cloud.[3] Starting with vSAN 8, enterprises can choose between two storage system configurations:

- **vSAN Original Storage Architecture (OSA)** is the traditional two-tier storage with caching/buffering and capacity storage.
- **vSAN Express Storage Architecture (ESA)** was released with vSAN 8 and is a flat tier with no caching.

OSA is a proven and robust solution for tiered storage and remains the predominant deployment model for vSAN.[4] Workloads running on vSAN OSA can significantly benefit from Intel® Optane™ SSDs in the cache tier, due to these drives' unique characteristics:

- The world's fastest data center SSD[5], with high and predictable performance[6] and low latency.[7]
- Excellent read performance even as write pressure increases.[8]
- Impressive bidirectional capability that enables simultaneous processing of both reads and writes while maintaining low latency.[9]
- Amazing endurance as measured by drive writes per day (DWPD)—up to 16x higher than a NAND SSD.[10]

Intel Optane SSDs Provide Value for Both Existing and New vSAN Deployments

- **Upgrading existing vSAN deployments.** For organizations with an older version of vSAN, upgrading to vSAN 8 with OSA using existing Intel Optane SSDs can help them scale their workloads while continuing to meet performance service-level agreements (SLAs).
- **New vSAN 8 deployments.** Organizations that are new to vSAN can optimize TCO by purchasing Intel Optane SSDs for the OSA cache tier. This approach results in the best of both worlds—great TCO and great performance.

For more information about Intel Optane SSDs, read the white paper, "Why Intel Optane SSDs Are a Better Option than NAND Flash SSDs in the Storage Cache Tier." The following sections demonstrate the performance, latency and performance/watt advantages of using Intel Optane SSDs with vSAN OSA, compared to a 3 DWPD TLC NAND SSD, for both a mixed workload and a sequential write workload. Note that both tested SSDs are NVMe-based.

Testing Methodology

The testing workload was HCIBench. This is an automation developed by VMware that manages multiple FIO instances across VM guests and produces cluster-wide storage metrics. HCIBench is an industry-recognized I/O benchmark for HCI.

The tests were run for two types of workloads, at various outstanding I/O operations (OIOs) and block sizes:

- 70/30 random read/write at 32 KB block size. VMware and Dell have identified this workload as an approximation of the most common I/O pattern observed in vSAN deployments.
- 100% sequential write. This workload provides a write pressure test.

Tests ran on a four-node vSAN OSA cluster; the only delta between clusters was the use of the Intel Optane SSD P5800X or a 3 DWPD TLC NAND SSD as cache drives.[11] Clusters featured eight VMs per node, and each VM was configured with four virtual disks. OIOs were distributed evenly among all disks in all VMs. For example, a 4K 70/30 random read/write workload with 128 OIOs ran one OIO per virtual disk across all VMs in the cluster (32 VMs × 8 nodes × 4 virtual disks per node = 128 OIOs).

For each FIO profile/workload, the test consisted of three passes of 80 minutes. The results of each pass were then averaged.

High, Predictable Performance for vSAN OSA

An easy way to understand storage cache performance is to use the analogy of cars crossing a bridge between the fictional locations of CPU Town and Dataville. Consider NAND NVMe SSDs as a two-lane bridge; Intel Optane SSDs are a six-lane bridge with a higher speed limit. More lanes enable more cars to cross the bridge in the same amount of time. A higher speed limit means the cars can go faster. For bridge commuters, the benefits of such a scenario can include getting to the office more quickly (enhancing productivity) or getting home faster (spending more time with family). Of course, storage operations don't have jobs or a family, but more efficient I/O—more IOPS and lower latency—prevents CPUs from sitting idle, burning cycles waiting for data.

The analogy can extend to fuel efficiency as well. For example, cars with great gas mileage can travel the same distance as a car with lower gas mileage, but burn less fuel to get there. Or, conversely, they can go farther using the same amount of fuel. In the data center, "high fuel efficiency" translates to enhanced performance/watt, which helps organizations meet their sustainability and cost-efficiency goals.

Overall, testing proved that Intel Optane SSDs can keep up with steadily growing IOPS demand and can sustain full read and write bandwidth simultaneously, while NAND SSDs cannot. More efficient I/O frees up CPU cycles for other work, while also enabling enterprises to achieve better quality of service during periods of sustained write pressure. Refer to subsequent sections for detailed results and discussion. See the endnotes for system configuration details.

Accelerate Everyday Workloads

After reviewing a sample of more than 200 vSAN clusters with more than 200 VMs running a variety of enterprise applications, VMware and Dell concluded that the average vSAN I/O workload is approximately 70/30 random read/write with a block size of 32 KB. Because these workloads are mixed read/write, they require an elevated degree of responsiveness from the storage system. Figures 1 through 3 illustrate how the Intel Optane SSD P5800X performs compared to an NVMe NAND drive across all metrics.

Throughput

For mixed workloads, the commuter analogy can be extended to a double-decker bridge versus a bridge with dedicated in-bound lanes in the morning and dedicated out-bound lanes in the evening. The double-decker bridge can carry traffic in both directions at once—a far more efficient design for commuting.

Figure 1 shows IOPS throughput for both drives for a mixed workload. Across a variety of OIOs, the Intel Optane SSD P5800X consistently outperforms the NAND drive by up to 42%, because Intel Optane SSDs can handle simultaneous read and write operations with fewer internal drive collisions bogging down performance.

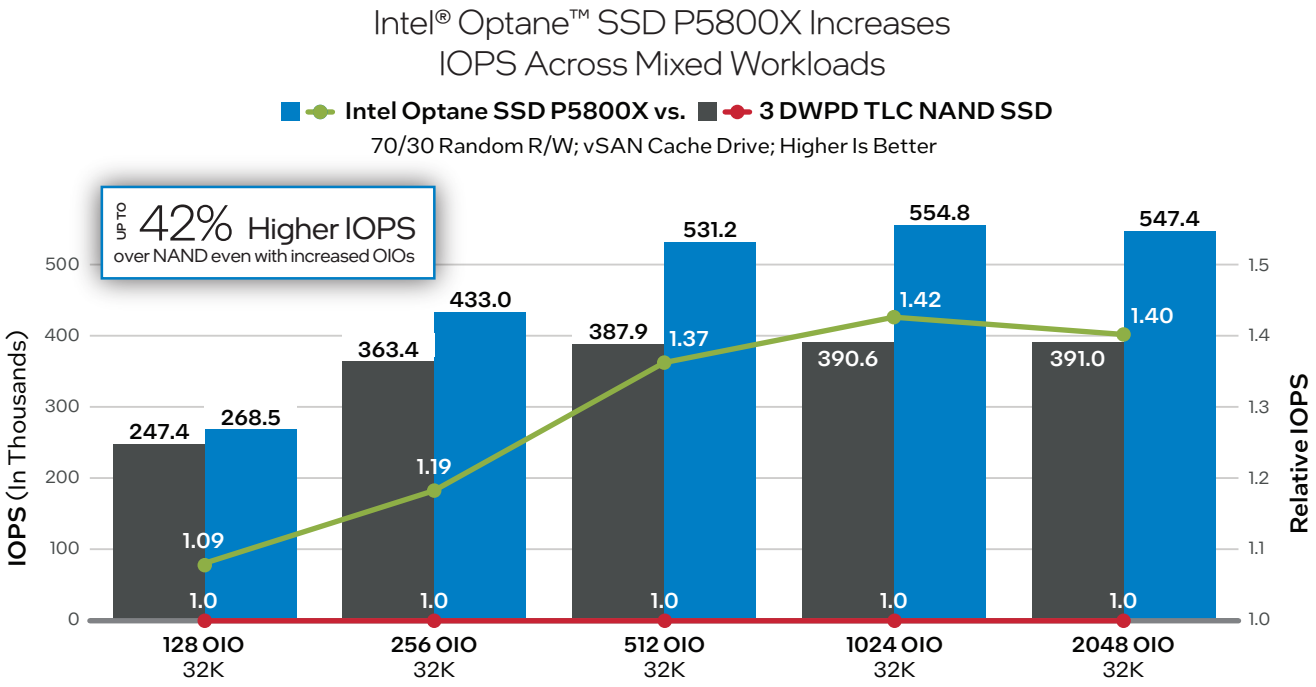

Figure 1. Intel® Optane™ SSD P5800X can achieve up to a 42% increase in IOPS for mixed workloads.

Latency

Figure 2 shows average latency for both drives for the same mixed workload as shown in Figure 1 (on previous page). Across a variety of OIOs, the Intel Optane SSD P5800X consistently outperforms the NAND drive by up to 29%.

IOPS per Watt

Figure 3 shows performance/watt for both drives for the same mixed workloads as shown in Figures 1 and 2. Across a variety of OIOs, the Intel Optane SSD P5800X consistently outperforms the NAND drive by up to 35%.

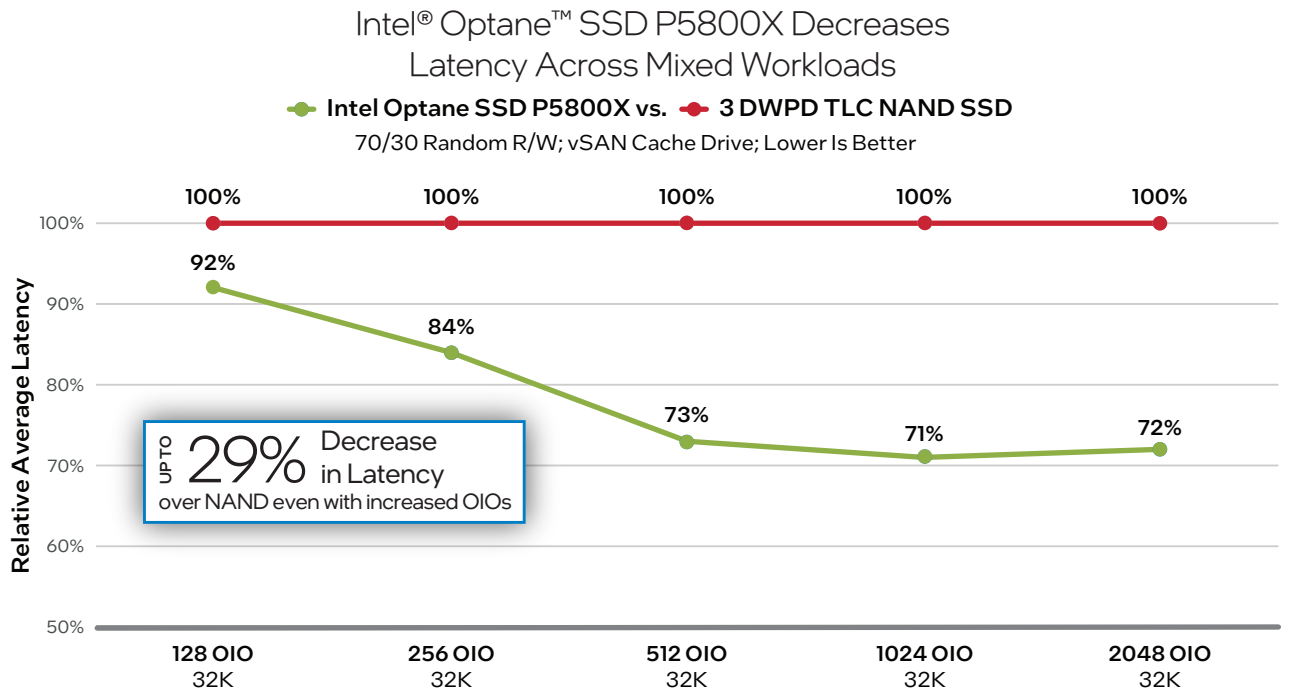

Figure 2. Intel® Optane™ SSD P5800X can reduce mixed-load latency by up to 29%.

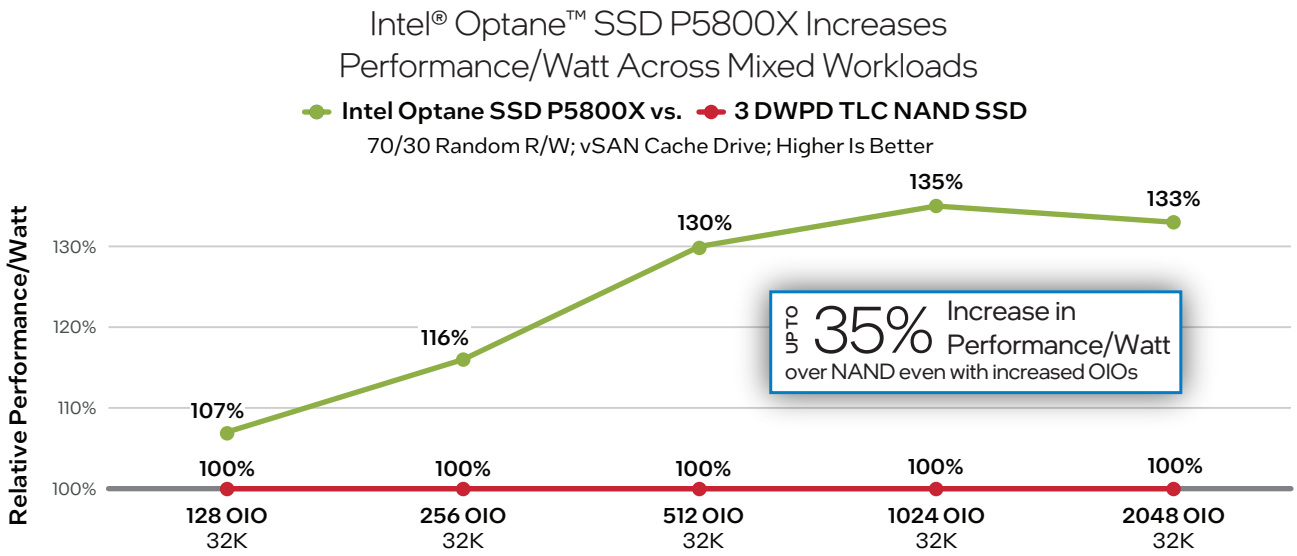

Figure 3. Intel® Optane™ SSD P5800X delivers up to a 35% increase in performance/watt for mixed workloads.

Continue to Meet SLAs Even When Write Pressure Escalates

Illustrated in the previous section, Intel Optane SSDs deliver value for workloads typically encountered in an enterprise environment. Intel Optane SSDs' value shines through even more in periods of sustained write activity, such as data backups, media transcoding and VM migrations. Sustained write pressure causes the cache drive to be unable to absorb additional writes. When this happens, vSAN begins destaging the write cache—that is, it moves the data from the cache layer to the capacity layer. This destaging results in simultaneous writes and reads to the cache drive, which conventionally results in lower performance and higher latencies due to garbage collection on NAND SSDs.

An SSD with higher write throughput adds meaningful value to the vSAN two-tier architecture because all writes must go through the cache device. Unlike NAND SSDs, Intel Optane SSDs are a write-in-place media; therefore, they provide a significantly higher write throughput (measured in MB per second) and lower latency than NAND SSDs. Higher throughput enables higher scalability and the ability to deliver predictable quality of service even when faced with prolonged periods of write-heavy workloads.

Throughput

Figure 4 shows IOPS throughput for both drives for a sequential write workload. Across a variety of OIOs and block sizes, the Intel Optane SSD P5800X consistently outperforms the NAND drive by up to 83%.

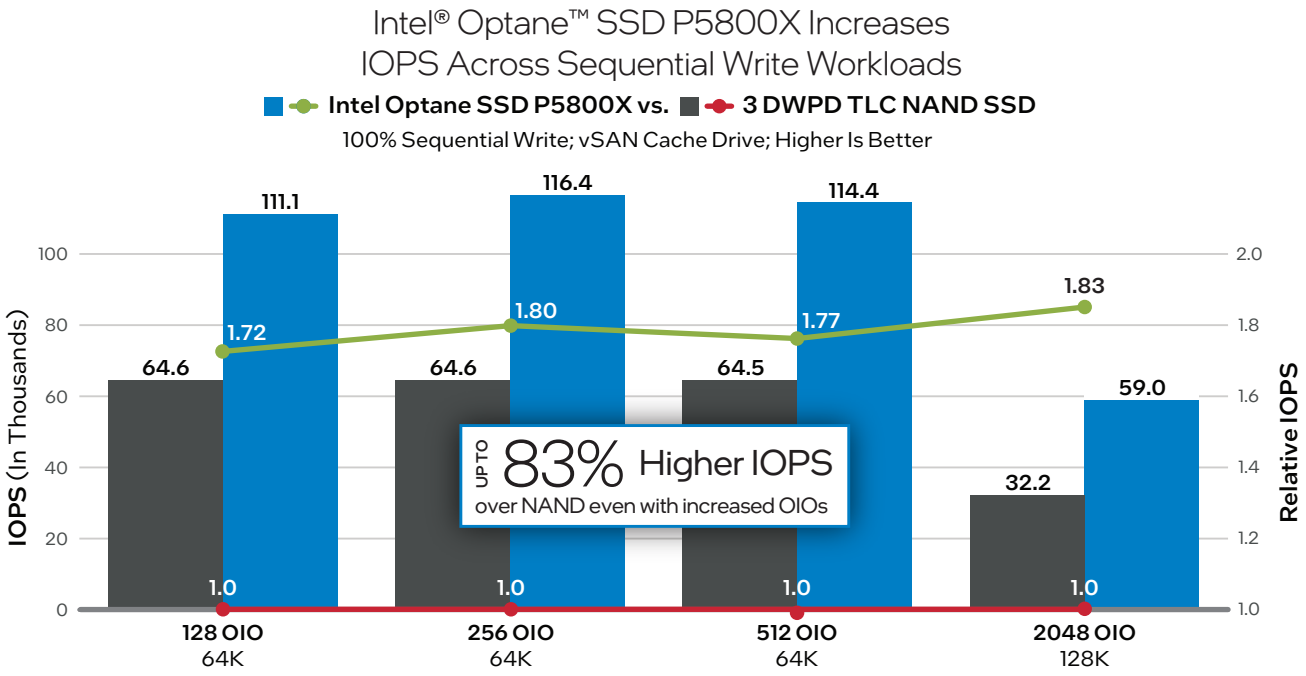

Figure 4. Intel® Optane™ SSD P5800X can achieve up to an 83% increase in IOPS for sequential write workloads.

Latency

Figure 5 shows latency for both drives for a sequential write workload. Across a variety of OIOs, the Intel Optane SSD P5800X consistently outperforms the NAND drive by up to 45%.

IOPS per Watt

Figure 6 shows performance/watt for both drives for a sequential write workload. Across a variety of OIOs, the Intel Optane SSD P5800X consistently outperforms the NAND drive by up to 72%.

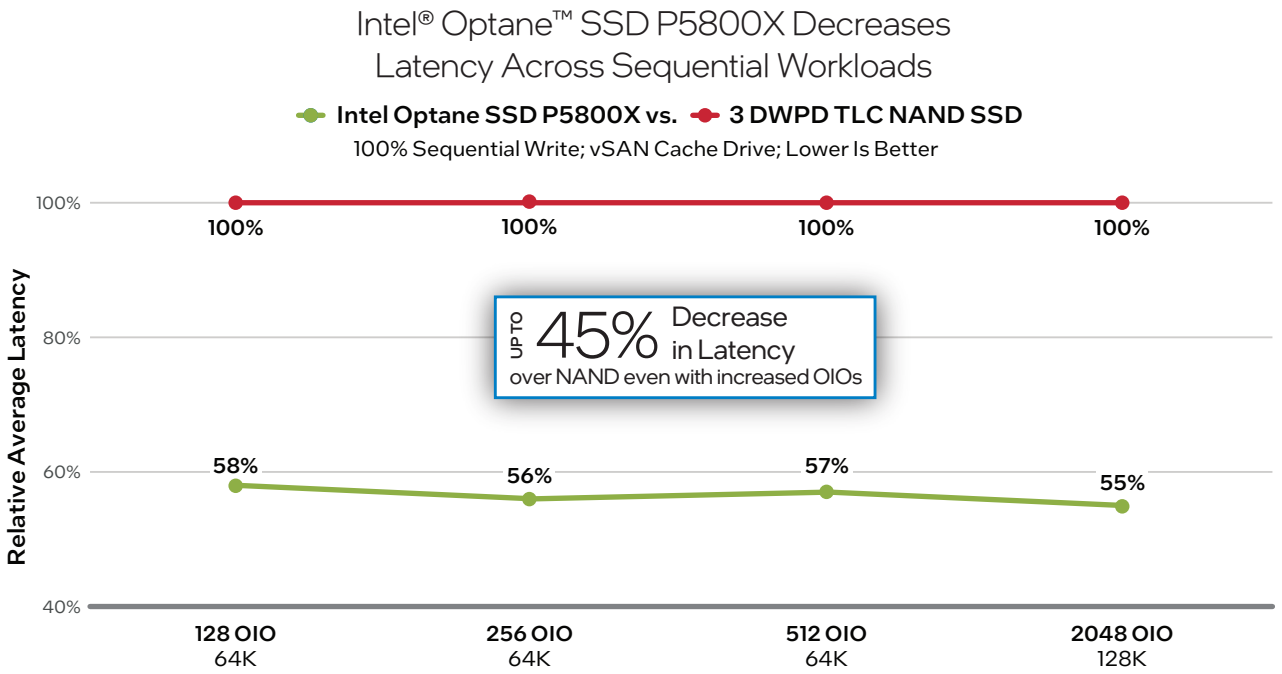

Figure 5. Intel® Optane™ SSD P5800X can reduce sequential write latency by up to 45%.

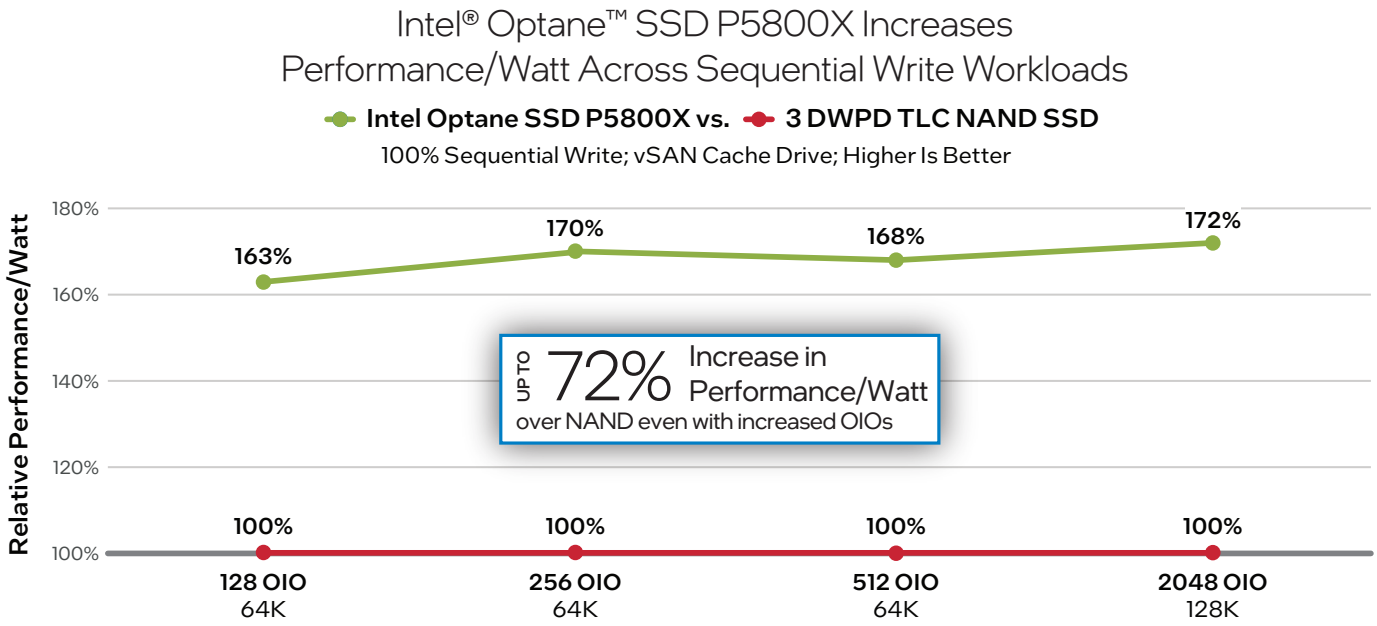

Figure 6. Intel® Optane™ SSD P5800X delivers up to a 72% increase in performance/watt for sequential write workloads.

Endurance Considerations

The test results presented in this white paper clearly show that Intel Optane SSDs are an ideal choice for the vSAN cache tier due to predictable high throughput, low latency and improved performance/watt compared to a NVMe NAND SSD. Another characteristic of Intel Optane SSDs that enterprises should consider is their high endurance.

To continue our commuter analogy, think of SSD endurance in terms of how many cars can travel the bridge before it needs to be shut down for maintenance. A bridge with higher endurance will need to be maintained less frequently.

The Intel Optane SSD P5800X is rated at 100 DWPD.[12] This type of endurance is advantageous for today's higher VM and I/O density due to the introduction of new CPU generations and intense hybrid cloud I/O workloads. Cache drive failure can cause the entire vSAN disk group to become unavailable, which increases exposure to data loss. It also negatively affects operations, because all drives in the disk group must be removed and subsequently recreated—which can be time-consuming for large-scale implementations.

When evaluating drive endurance, it is important to consider the total write load an application places on the cache drive. Table 1 compares the endurance of a 1.6 TB NAND SSD and the 800 GB Intel Optane SSD P5800X, showing the total I/O activity that can be supported for five years.

Another aspect to think about is trapped capacity. Typically, higher-endurance NAND SSDs, except for SLC drives, have a larger capacity than vSAN will use for caching. Older versions of vSAN could use only 600 GB of a cache device for write buffering; vSAN 8 can use up to 1.2 TB. However, many higher-endurance NAND SSDs have capacities of several TBs, meaning the excess capacity is wasted when the drive is used as a cache drive. Intel Optane SSDs are just the right size, with no trapped capacity going to waste.

Table 1. Endurance As It Relates to I/O Activity

Drive	Endurance DWPD/TBW	Maximum Continuous Writes Supported (MB/s)	Maximum Continuous 70/30 32 KB IOPS Supported
1.6 TB NAND SSD	3/8,600	56	3,000
800 GB Intel® Optane™ SSD P5800X	100/146,000	948	50,785

TBW = terabytes written

Conclusion

Whether it's a new or existing vSAN OSA deployment, enterprises that value high performance and performance/watt, low latency, predictable throughput and high endurance should choose Intel Optane SSDs for the vSAN cache tier. Our test results are as follows:

- **Mixed workloads.** 42% more IOPS than a NVMe NAND SSD with 29% lower latency and 35% more performance/watt).
- **Sequential write workloads.** 83% more IOPS with 45% lower latency and 72% more performance/watt).[13]

The advantages of Intel Optane SSDs are clear, whether the workload is mixed or 100% sequential write. And with a rating of 100 DWPD, Intel Optane SSDs are far less likely to fail—even with aggressive write loads—than a NAND SSD rated at 3 DWPD. As the tsunami of data continues to surge across every industry, Intel Optane SSDs are the go-to solution to keep up with storage demands.

Learn More

- Intel® Optane™ SSDs
- Intel® Optane™ SSD P5800X product brief
- What's New with vSAN 8
- VMware and Intel® Partner Alliance
- Intel and VMware Partnership

For more information, visit intel.com/optane.

[1] Testing by Intel as of October 18, 2023.

Common configuration: 2x Intel® Xeon® Platinum 8480+ processor (56 cores, 2.0 GHz), 512 GB memory (16x 32 GB DDR5 4800 MT/s), Intel® Hyperthreading Technology = ON, Intel® Turbo Boost = ON, capacity tier: 6x Solidigm D7-P5510 SSD (formerly Intel SSD D7-P5510) 3.5 TB, boot drive: 1x 119.2G CL1-3D128-Q11 NVMe SSSTC 128 GB, network = 2x Intel® Ethernet Controller X710 for 10GBASE-T, 2x Intel® Ethernet Controller E810-C for QSFP, BIOS = 3A11.uh, microcode = 0x2b000111, VMware ESXi 8.0U2, HCIBench 2.8.2

NAND cache tier: 2x Solidigm D7-P5600 SSD (formerly Intel® SSD D7-P5600) U.2 1.6 TB

Intel® Optane™ cache tier: 2x Intel Optane SSD P5800X 800 GB

[2] Rivery, August 2023, "Big Data Statistics: How Much Data Is There in the World?"

[3] VMware, July 2022, "The Big Payoff: Customers' Perspectives on the Benefits of VMware vSAN."

[4] BrightTalk, October 2023, "Modernizing Your vSAN Infrastructure for Next-Gen Workloads."

[5] Claim [14] at https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-ssd-p5800x-series/

[6] Claim [7] at https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-ssd-p5800x-series/

[7] Claim [5] at https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-ssd-p5800x-series/

[8] Claim [18] at https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-ssd-p5800x-series/

[9] Claim [4] at https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-ssd-p5800x-series/

[10] A NAND-based 800 GB drive with an endurance rating of 3 DWPD will last five years only if the sustained write workload is lower than 2,400 GB per day or 28.4 write MB/s. Put another way, even a light sustained I/O load such as 450 32 KB write IOPS activity would fill the drive three times per day, assuming the use of RAID 1 as data protection. In contrast, an Intel® Optane™ SSD P5800X 400 GB drive has an endurance rating of 100 DWPD. It can sustain 40,000 GB written per day or 470 MB/s written per day for five years. That equates to a 16x increase in endurance.

[11] See endnote 1.

[12] https://www.intel.com/content/www/us/en/products/sku/201859/intel-optane-ssd-dc-p5800x-series-1-6tb-2-5in-pcie-x4-3d-xpoint/specifications.html

[13] See endnote 1.

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

 1223/KLET/KC/PDF