

Solution Brief

Data Center
Intel® Optane™ SSDs



Bringing Record-Breaking Performance to the Mainstream

StarWind uses Intel® Optane™ SSDs to enhance cache performance in hyperconverged infrastructure solution.



StarWind HCI solution running Intel® Optane™ SSDs as writeback cache achieved over **26.8** million IOPs.¹

Executive Summary

As part of its commitment to staying ahead of customer challenges in the virtualization and hyperconverged infrastructure (HCI) space, StarWind regularly experiments with new technologies and use cases. In one evaluation, it ran an extensive series of tests to assess the performance of Intel® Optane™ SSDs in its hyperconverged appliance (HCA) solution and identify practical use cases for the product. The results demonstrated record-breaking performance¹ of 26.8 million input/output operations per second (IOPs) when Intel® Optane™ SSDs were used as a write-back cache.²

StarWind is now incorporating Intel Optane SSDs into the next generation of its HCA solution to bring the performance benefit to its customers. The combination of high performance³, high endurance⁴, and cost effectiveness makes them a good fit for customers' remote and branch office environments.

Solution Benefits

- **Record-breaking performance:** StarWind HCI solution running Intel® Optane™ SSDs as writeback cache achieved over 26.8 million IOPs.¹
- **High endurance:** Intel® Optane™ SSD DC P4800X delivers up to 60 drive writes per day (DWPD), significantly reducing the risk of cache device wear out.⁴
- **Cost effective:** Using Intel Optane SSDs for cache can help reduce total cost of ownership (TCO) because of the endurance and lower caching capacities required.
- **Replicable configuration:** Similar performance could be obtained using most hypervisors with iSCSI initiators, StarWind Virtual SAN and Intel Optane SSDs as cache.

Customer Need

Storage virtualization specialist, StarWind, is committed to providing its customers large and small with enterprise-grade HCI solutions at a reasonable price. The majority of StarWind's customers are small to medium-sized businesses (SMBs), or remote or branch offices (ROBO) of larger enterprises. They often have limited IT expertise internally, so rely on StarWind to provide an HCI infrastructure—including hardware, software, network and storage—that will plug in and 'just work'. "Our solutions should be like the foundations of our customers' office buildings," says Andrii Tuzov, director of business development at StarWind. "They should do their job without anyone ever noticing they're there."

With technology constantly evolving, the company must be proactive in identifying new capabilities and use cases that could help optimize its offerings. “We’ve seen a lot of change in the hardware market,” says Tuzov. “For example, it took time for people to move from SATA to SAS storage solutions, but now companies are accelerating to non-volatile memory express (NVMe). This means we need to help our customers with large NVMe environments get the most out of that hardware. Many software stacks today struggle to do that, so we wanted to experiment with NVMe technologies to find the optimal hardware and software configurations for our HCI solutions.”

Solution Evaluation

StarWind is a long-term user of Intel® technology in its solutions and is proactive in evaluating new use cases on an ongoing basis. The team wanted to assess Intel Optane SSDs’ potential role in its HCA, an HCI platform of hardware and software, built using cost-effective, industry-standard components. “The low latency of Intel Optane SSDs, even when under heavy load, makes them a good candidate to support caching for data-intensive workloads,” says Tuzov. “We wanted to put the devices to the test to see how they would benefit our customers needing high performance in the data center.”

The StarWind team upgraded the CPU option of a typical StarWind HCA configuration, to attempt record-breaking HCI performance. The 12-node cluster used Supermicro Super Server platforms powered by Intel® Xeon® Platinum 8268 processors. Primary storage was provided by Intel® SSD D3-S4510 drives and write-back cache by Intel® Optane™ SSD DC P4800X drives.

A Microsoft Hyper-V hypervisor ran on the cluster, with StarWind VSAN software for software-defined storage and StarWind Command Center for management of the whole appliance. The testing was conducted using the Microsoft DISKSPD benchmark, and virtual machines deployed to stress the storage system, orchestrated using Microsoft VM Fleet.

The team evaluated the performance of Intel Optane SSDs in multiple scenarios, including raw storage, cache, and NVMe-over-Fabric (NVMe-oF) cache for a variety of read/write combinations. Configuring Intel Optane SSDs as a write-back cache delivered the strongest results (see Figure 1).¹ When running 100 percent reads and focusing on maximizing IOPs, the cluster delivered over 26.8 million IOPs.

“These results are really a breakthrough in performance for a production configuration,” says Tuzov. “We see this being highly beneficial for our clients using innovative software stacks in need of high-performance storage. These include those in fields like academic research, computer-aided design or oil and gas exploration.”

Business Value

Following this evaluation, StarWind sees potential to broaden use of Intel Optane SSDs as cache, beyond high-performance use cases. Similar performance could be obtained using any hypervisor with iSCSI initiators, StarWind VSAN and Intel Optane SSDs as cache. StarWind therefore plans to use the technology in the next generation of its HCA platform so customers of all sizes and industries can benefit from it.

“We’ve proved the value of this approach, and now we’re ready to take it to our customers,” says Tuzov. “Many of the organizations we work with, such as those in the public sector, typically have limited resources to invest in their IT so they tend to rely on out-of-the-box solutions. We pride ourselves on giving them customized, architect-fitted solutions and white-glove treatment without a high price tag. Being able to offer this type of high-performance cache technology is another example of that approach.”

Cost effectiveness is also important for StarWind’s customers, and StarWind feels Intel Optane SSDs offer an advantage here too. “The average recommended caching size is 10 percent of your total capacity,” explains Tuzov. “But with Intel Optane SSDs, we’ve found that just 5 percent can be enough, depending on the working set. This means price per gigabyte can be significantly reduced, regardless of what storage media is running underneath it.”

The recommended caching size of 10 percent cache to 100 percent capacity is generally applicable to slower media like HDDs, where both read and write performance are lower. NVMe NAND SSDs provide higher read performance, meaning that only write operations need to be cached, making smaller cache sizes adequate. In all-flash configurations though, additional factors need to be considered when sizing necessary write cache capacity. The required performance, workload profile, and endurance of the drives used as a write cache all need to be taken into account. For example, the lower the endurance characteristics of NAND SSD drives used for write caching, the larger capacity will be required for the given scenario. Due to the exceptional endurance of Intel Optane SSDs, much smaller write cache capacities

| RUN | PARAMETERS | RESULT ¹ |
|--|----------------------------------|---------------------|
| Maximize input/output operations per second (IOPs), all-read | 4 kB random, 100% read | 26,834,060 IOPS |
| Maximize IOPS, read/write | 4 kB random, 90% read, 10% write | 25,840,684 IOPS |
| Maximize IOPS, read/write | 4 kB random, 70% read, 30% write | 16,034,494 IOPS |

Figure 1. Summary of StarWind benchmark results for Intel® Optane™ SSDs as cache in 12-node hyperconverged infrastructure (HCI) cluster

are needed compared to NAND SSDs even in cases when write-intensive workloads require sustainable high levels of performance.

This high endurance also makes Intel Optane SSDs a good fit for organizations using StarWind's solutions for remote or branch office locations. For example, the Intel Optane SSD DC P4800X delivers up to 60 DWPD, about 20x more than many NAND SSDs.¹ This means that they deliver longevity as well as performance – an important advantage for companies operating multiple, geographically distant locations. “This all supports the goal of our solutions running unnoticed in the background,” says Tuzov. “We and our customers can be confident that the technology will perform well, last a long time, and be cost effective. Also, there's no need to compromise as you get all these advantages in the one device.”

Conclusion

StarWind has already begun exploring next steps with its customers. “We've had interest in the configuration we used in the evaluation for high-performance workloads,” Tuzov says. “We're in discussions about that while we also develop our next-generation appliance using Intel Optane SSDs for our mainstream customers.”

He concludes: “We've worked with Intel for a long time and look forward to continuing to use Intel's latest technology innovations in future releases of our HCI solutions.”

Find the solution that's right for your organization. Contact your Intel representative or visit www.intel.com/optane

Learn More

- Hyper-Converged Infrastructure Production Record: 26.8 million IOPS. 100% of the theoretical limit
- Intel® Optane™ SSDs for Data Center

About About StarWind

StarWind specializes in hyperconvergence and storage virtualization. It provides a unique blend of simplicity, performance, and affordability, yet with the ultimate in flexibility. Its focus is to help companies create sustainable and effective IT infrastructures with minimum effort.

Learn more at: www.starwindsoftware.com/



¹ The previous HCI record was 13.7 million IOPS: <https://techcommunity.microsoft.com/t5/storage-at-microsoft/the-new-hci-industry-record-13-7-million-iops-with-windows/ba-p/428314>.

² **Configuration details:** Platform: Supermicro SuperServer 2029UZ-TR4+; CPU: 2x Intel® Xeon® Platinum 8268 Processor 2.90 GHz. Intel® Turbo Boost ON, Intel® Hyper-Threading ON; RAM: 96GB; Boot and Storage Capacity: 2x Intel® SSD D3-S4510 Series (240GB, M.2 80mm SATA 6Gb/s, 3D2, TLC); Write-Back Cache Capacity: 4x Intel® Optane™ SSD DC P4800X Series (375GB, 1/2 Height PCIe x4, 3D XPoint™). The latest available firmware installed; RAW capacity: 5.7TB; Usable capacity: 5.4TB; Working set capacity: 1.32TB; Networking: 2x Mellanox ConnectX-5 MCX516A-CCAT 100GbE Dual-Port NIC; Switch: 2x Mellanox SN2700 Spectrum 32 ports 100 GbE Ethernet Switch; StarWind iSCSI Accelerator and Microsoft iSCSI Initiator, StarWind Loopback Accelerator. Operating System: Windows Server 2019 Datacenter Evaluation version 1809, build 17763.404. Hype-V Role installed, MPIO and Failover Cluster features enabled. StarWind VSAN version 8.0.0.12996.

Testing conducted May 2019 by StarWind Software Incorporated. For more details see here: www.starwindsoftware.com/iscsi-shared-storage-with-wb-cache-enabled

³ See endnote 1.

⁴ **Source – Intel.** Capacity per form factor: Half-height, half-length (HHHL) add-in card (AIC): 375GB, 750GB, 1.5TB, 2.5"x15mm; Small form factor U.2: 375GB, 750GB, 1.5TB; P4801X: U.2: 100GB, 200GB, 375GB. Form factor add-in card (AIC), half-height, half-length (HHHL), low-profile; U.2 2.5", 15mm; M.2 110mm. Interface: PCIe 3.0 x4, NVMe. Latency (typical) R/W: <10/12µs. Quality of Service (QoS): 99.999%: 4KB random, queue depth 1, read/write <60/100µs; 4KB random, queue depth 16, R.W <150/200µs. Throughput: 4KB random, queue depth 16, read/write up to 550/550k IOPS; 4KB random, queue depth 16, mixed 70/30 read/write up to 500k IOPS. Endurance (JESDS219 workload): 30 DWPD: 375GB – 20.5 PBW; 750GB – 41 PBW; 60 DWPD: 100GB – 10.9PBW; 200GB – 21.9 PBW; 375GB – 41.0 PBW; 750GB – 82.0 PBW; 1.5TB – 164 PBW. Power: Enhanced power-loss data protection. P4800X: Active/idle – up to 18W / 7W; P4801X: Active/idle – up to 11W / 3W. Endurance ratings available at <https://www.intel.com/content/www/us/en/solid-state-drives/optane-ssd-dcp4800xbrief.html>.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

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

Your costs and results may vary.

Results have been estimated or simulated. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

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

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others. 1020/AS/CAT/PDF Please Recycle 344026-002EN