Solution brief

Data Center
Intel® Optane™ SSDs

Proven performance, speed and endurance for key workloads

Open-E demonstrates the value of Intel® Optane™ SSDs for read and write cache.

75%

Intel® Optane™ SSD combined with RAM delivers up to **75%** the performance of all-RAM for less than half the price per GB.[2]

20x

Intel® Optane™ SSD DC P4800X delivers up to **20x** more writes per day than typical NAND SSD cache.[7]

Executive Summary

Data storage software provider, Open-E, evaluated the use of Intel® Optane™ SSDs for read and write cache, using its JovianDSS application. The team found that using these devices as a Zettabyte File System (ZFS) Intent Log (ZIL), the storage volume delivered high performance, higher, consistent IOPs and lower latency compared to traditional NAND SSDs.[1] As a read cache (Level 2 Adaptive Replacement Cache) when combined with RAM devices, Intel Optane SSDs delivered up to 75 percent of the performance for less than half the price per gigabyte of a RAM-only configuration.[2]

Open-E now recommends Intel Optane SSDs to clients for their cache devices thanks to their combination of performance, low latency and high endurance. They effectively support customers' common workloads, including databases, virtualization, online transaction processing (OLTP) and virtual desktop infrastructure (VDI). They also provide a cost-effective caching solution, with lower cost and greater performance per dollar for read cache than DRAM and greater performance per dollar for write cache than NAND.

Solution Benefits

- **High, consistent IOPs.** Intel® Optane™ SSDs deliver high, consistent IOPs for ZFS Intent Log (ZIL), without losing speed over time as traditional NAND does.
- **Low latency.** Intel Optane SSDs maintain low, consistent latency for ZIL while traditional NAND devices increase latency over time.
- **High quality of service (QoS).** Low latency means more write operations can be handled in a lower and narrower latency time period, helping improve overall QoS.
- **Cost-effective performance.** Used as L2ARC, Intel Optane SSDs combined with RAM devices deliver up to 75 percent of the performance of all-RAM configurations for less than half the price per gigabyte.[2]
- **High endurance.** Intel® Optane™ SSD DC P4800X delivers up to 60 drive writes per day (DWPD), approximately 20x more than typical NAND SSDs used for caching.[7]

Customer Need

Data is increasingly at the heart of business operations, processes and performance for every business, regardless of its size or the industry in which it operates. This means that an organization's approach to managing, optimizing and using data can have real impact on the bottom line.

Data storage software provider, Open-E, works with customers across Europe and the United States every day to help them address this challenge. "We see two main requirements – the need to handle large data volumes, and the need to do it cost effectively," explains Krzysztof Franek, CEO and president of Open-E. "Many of our clients are in the media and entertainment industry, so for them large data capacity is key. This is true also of healthcare organizations with lots of big medical imaging files. Meanwhile government and education organizations are focused on ensuring they have a really cost-effective way to store and protect their data."

Open-E must support both these customer requirements. This involves regularly evaluating available storage technologies against customers' storage, backup & recovery, and business continuity needs. It must then identify the most appropriate configurations to optimize software performance.

Therefore, when the opportunity arose to try a new caching approach for the most frequently or recently accessed (hot) data, Open-E was happy to take it. Caching is an important focus, as it can be a significant bottleneck, especially for open-source platforms like Zettabyte File System (ZFS). This underpins Open-E's JovianDSS application, which is designed to support businesses of any size, running enterprise software-defined storage (SDS) environments. In order to ensure optimal application performance and customer workload support, Open-E must carefully evaluate the caching technologies it recommends.

Solution Evaluation

Open-E wanted to determine whether using Intel® Optane™ SSDs in the cache layer would help improve performance for its customers' workloads. The NAND SSDs commonly used for write caching, although they perform well initially, tend to lose performance very quickly and require the support of an internal RAM buffer. Their lifetime can be limited and high-endurance NAND drives can be expensive. Meanwhile, DRAM devices typically used for read caching are difficult to scale up as data accesses increase due to their relatively high price.

Open-E ran a series of tests to evaluate the performance of the Intel Optane SSD as a cache device for both write and read, managed using its JovianDSS application. The first evaluation compared a common installed-base configuration such as the Intel® DC SSD P3700 with Intel Optane SSDs used for ZIL cache in Open-E JovianDSS. It demonstrated that the Intel Optane devices maintained a higher, consistent level of IOPs over time, while that of the NAND SSDs decreased (see Figure 1).[3]

There was a difference in latency too. The Intel Optane SSDs delivered low, consistent latency through the test, while that of the NAND SSDs increased over time (see Figure 2).[4] This is particularly important for applications that write large numbers of small files, especially in databases. Low latency is also critical for virtualization, online transaction processing (OLTP) and virtual desktop infrastructure (VDI) workloads.

The lower, more stable latency of the Intel Optane SSDs helps improve overall quality of service (QoS). It means that more write operations can be handled in a lower and narrower latency time period (see Figure 3).[5]

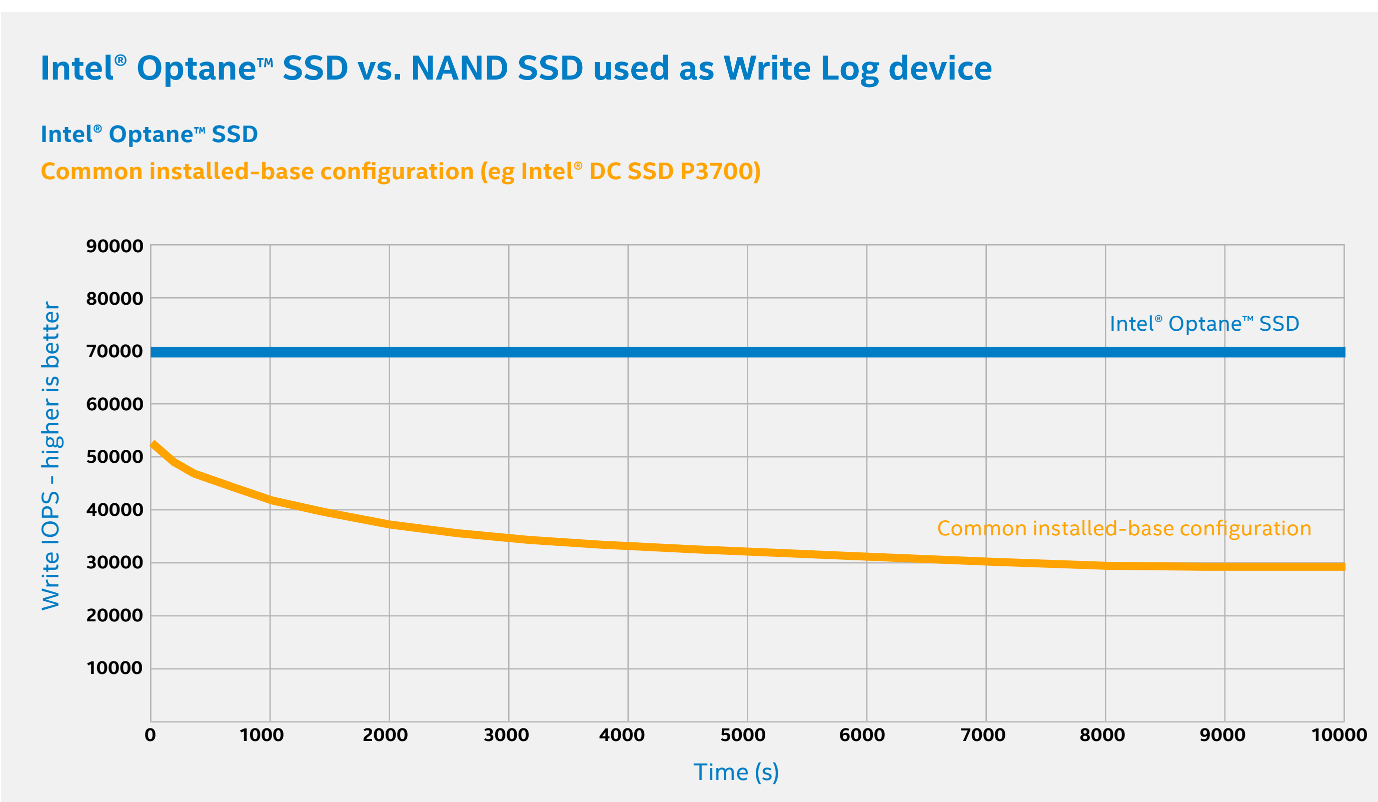

Figure 1. NAND SSD used as a write cache slows down over time.[3]

Latency for Intel® Optane™ SSD as a Write Log in Open-E JovianDSS

Intel® Optane™ SSD

Common installed-base configuration (eg Intel® DC SSD P3700)

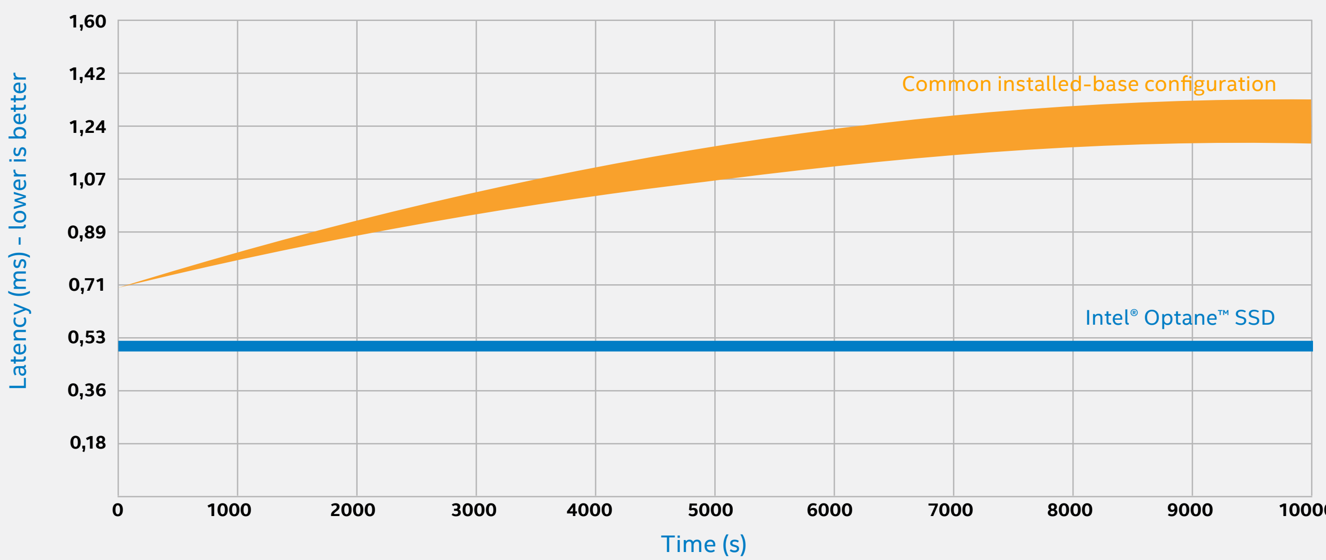

Figure 2. Intel® Optane™ SSDs deliver low, consistent latency while NAND SSD latency increases over time.[4]

QoS for random write with Intel® Optane™ SSD as a ZFS Write Log

Intel® Optane™ SSD

Common installed-base configuration (eg Intel® DC SSD P3700)

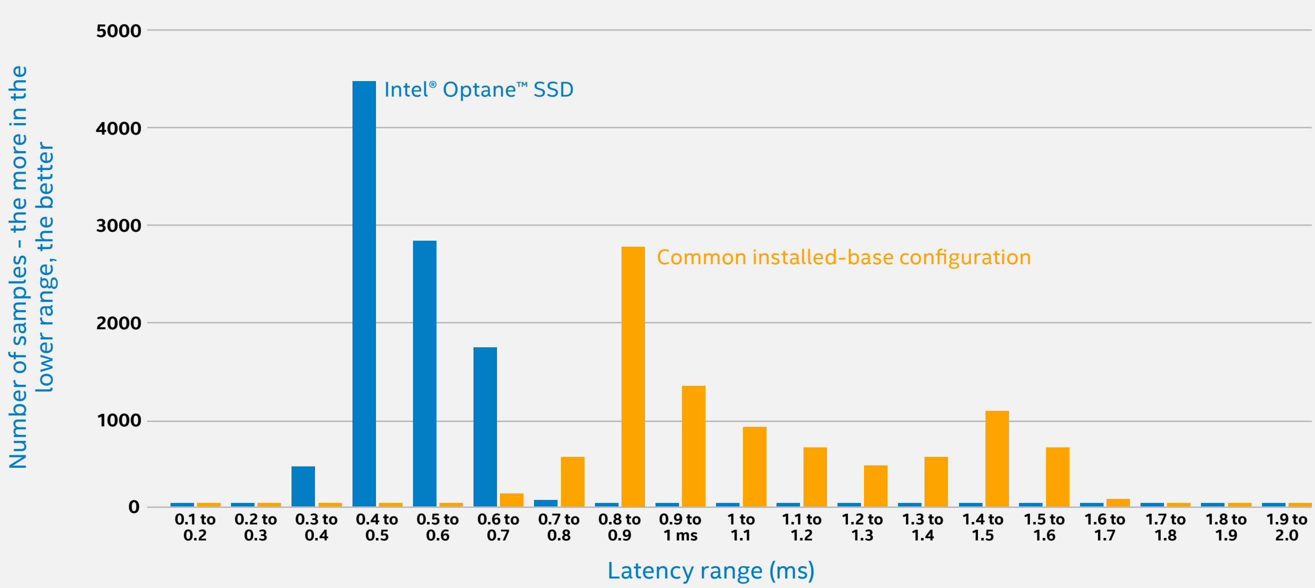

Figure 3. Quality of Service for random write improves with Intel Optane SSDs compared to NAND SSDs.[5]

The performance of Intel Optane SSDs for read caching was also evaluated. In this test, a RAM-only configuration was run against a second configuration that combined RAM (ARC) with Intel Optane SSDs (L2ARC). The combination including Intel Optane SSDs was procured for less than half the price per gigabyte than the RAM-only option, but still delivered 75 percent of the performance.[6] This makes using Intel Optane SSDs a compelling and cost-effective option for scaling read cache capabilities (see Figure 4).

Business Value

The performance, latency and endurance benefits Open-E has identified in using Intel Optane SSDs translate to solid business value for its customers. “Price and performance are critical for everyone,” says Franek. “We work hard to identify the technologies that will give our customers optimal price performance. For caching, there’s often a perception that you should use specific (NAND) devices, and that you need to overprovision to overcome their limitations. However, we have shown that using Intel Optane SSDs offers consistent, high performance and low latency. Furthermore, Intel Optane SSD devices come in smaller capacities, like 100GB, which is plenty for a write cache functionality to accompany NAND SSDs. This means our customers can get the performance they need, without having to overprovision. This isn’t possible with NAND devices, which start at 200GB today.”

By helping its clients avoid overprovisioning, Open-E is continuing its commitment to building cost-effective solutions. “Our software is inexpensive, but we also want to ensure our customers are getting value for money across their storage resources,” says Franek. While using Intel Optane SSDs for the write cache helps avoid overprovisioning, using it as a read cache also helps reduce server costs as it can be a replacement for more costly error-correcting code (ECC) RAM.

Another advantage of using Intel Optane SSDs is their high endurance. For example, the Intel® Optane™ SSD DC P4800X delivers up to 60 drive writes per day (DWPD), approximately 20x more than typical caching NAND SSDs[7]. For Open-E’s clients, this durability translates to reliability – to being able to ‘set and forget’ a device, confident that it will continue to perform as needed. From Open-E’s perspective, this is also beneficial as it reduces the burden on the team running support services.

Intel Optane SSDs used as a write log also help to lower the total cost of other storage tiers. Hard drives and NAND SSDs used for capacity storage benefit from less frequent writes, extending their lifetimes.

Open-E also recommends that its customers use Intel® technology in their storage pools – specifically, Intel® QLC 3D NAND SSDs. These large-capacity drives provide a cost-effective flash alternative to traditional hard disk drives (HDDs), offering performance and capacity at a reasonable price point. Their large capacity and relatively low power consumption also help Open-E’s customers reduce their overall power use and carbon footprint. Not only does this translate to total cost of ownership (TCO) savings, it also helps organizations align with industry and governmental regulation around IT power consumption.

Random read performance: ARC (RAM) vs. ARC + L2ARC based on Intel® Optane™ SSD

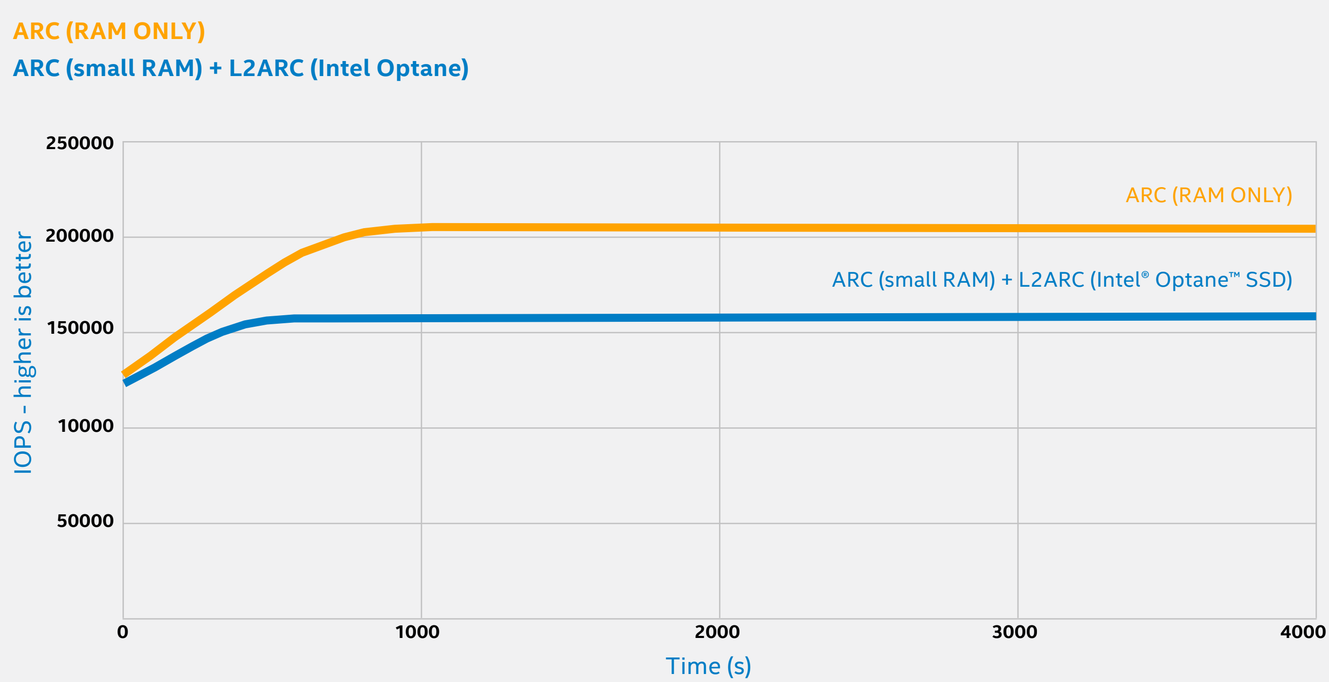

Figure 4. Intel® Optane™ SSD is slower compared to RAM but at time of testing was less than half the cost of ECC RAM.[6]

Conclusion

The team at Open-E has already implemented solutions for multiple customers using Intel Optane SSDs for caching. "We've also done a lot of work to tune up our software for this hardware so that we can help our customers achieve greater results than our competitors using the same platform," Franek concludes. "We always recommend our clients use this technology now, and when we show them the data from our evaluations, they see why."

Learn More

- **Webpage:** Open-E JovianDSS Advanced Caching with Intel® Optane™ SSDs
- **Webpage:** Intel® Optane™ SSDs for Data Center

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

About Open-E

Open-E, Inc. develops IP-based storage management software, with headquarters in the United States and Europe. Founded in 1998, the company sells its line of storage management software through a world-wide network of system integrators and resellers. The Open-E JovianDSS and Open-E DSS V7 line of products enjoy a reputation for best-in-class performance, flexibility, reliability, scalability and return-on-investment. Learn more at: www.open-e.com/

[1] Test was performed locally on the operating system. Random write workload was generated by fio tool and results - IOPS, Latency and QoS, were read from kernel diskstats (iostat). Additionaly, ZFS arcstats were used to monitor that 100% cache was used during the test. **Test data:** Test data size: 200GB; Workload pattern: 100% random write; Used IO engine: libaio; Outstanding IO: 256; Number of threads: 32; Test block size: 4KB; Test tools: fio-2.1.14+iostat+arcstat. **Software & storage configuration:** Open-E JovianDSS ver: 1.0 up28 (b37311); Pool architecture: 12x single disk (no redundancy); Write log (ZIL): single NVMe disk; Read cache (L2ARC): none; Read cache (ARC): 144GB (75% of 192GB RAM); zvol size: 200GB; zvol block size: 4KB; zvol provisioning: thick; **Hardware specification:** Server model: TAROX ParX R2242i G6 server; CPU: Intel® Xeon™ Silver 4208R CPU @ 2.20GHz; RAM: 192GB; Capacity storage: 12x Intel® SSDPE2NV153T8; SSD/NVMe (write log): Intel® Optane™ DC SSD P4800X / Intel(R) Optane(TM) DC SSD P3700; SSD/NVMe (read cache): None

[2] Test was performed locally on the operating system in May 2020. Random write workload was generated by fio tool and results - IOPS, Latency and QoS, were read from kernel diskstats (iostat). Additionaly, ZFS arcstats were used to monitor that 100% cache was used during the test. **Test data:** Test data size: 200GB; Workload pattern: 100% random write; Used IO engine: libaio; Outstanding IO: 256; Number of threads: 32; Test block size: 4KB; Test tools: fio-2.1.14+iostat+arcstat. **Software & storage configuration:** Open-E JovianDSS ver: 1.0 up28 (b37311); Pool architecture: 12x single disk (no redundancy); Write log (ZIL): single NVMe disk; Read cache (L2ARC): none; Read cache (ARC): 144GB (75% of 192GB RAM); zvol size: 200GB; zvol block size: 4KB; zvol provisioning: thick; **Hardware specification:** Server model: TAROX ParX R2242i G6 server; CPU: Intel® Xeon™ Silver 4208R CPU @ 2.20GHz; RAM: 192GB; Capacity storage: 12x Intel® SSDPE2NV153T8; SSD/NVMe (write log): Intel® P4800 / Intel® P3700; SSD/NVMe (read cache): None

[3] Test was performed locally on the operating system. Random write workload was generated by fio tool and results - IOPS, Latency and QoS, were read from kernel diskstats (iostat). Additionaly, ZFS arcstats were used to monitor that 100% cache was used during the test. **Test data:** Test data size: 200GB; Workload pattern: 100% random write; Used IO engine: libaio; Outstanding IO: 256; Number of threads: 32; Test block size: 4KB; Test tools: fio-2.1.14+iostat+arcstat. **Software & storage configuration:** Open-E JovianDSS ver: 1.0 up28 (b37311); Pool architecture: 12x single disk (no redundancy); Write log (ZIL): single NVMe disk; Read cache (L2ARC): none; Read cache (ARC): 144GB (75% of 192GB RAM); zvol size: 200GB; zvol block size: 4KB; zvol provisioning: thick; **Hardware specification:** Server model: TAROX ParX R2242i G6 server; CPU: Intel® Xeon™ Silver 4208R CPU @ 2.20GHz; RAM: 192GB; Capacity storage: 12x Intel® SSDPE2NV153T8; SSD/NVMe (write log): Intel® P4800 / Intel® P3700; SSD/NVMe (read cache): None

[4] Test was performed locally on the operating system. Random write workload was generated by fio tool and results - IOPS, Latency and QoS, were read from kernel diskstats (iostat). Additionaly, ZFS arcstats were used to monitor that 100% cache was used during the test. **Test data:** Test data size: 200GB; Workload pattern: 100% random write; Used IO engine: libaio; Outstanding IO: 256; Number of threads: 32; Test block size: 4KB; Test tools: fio-2.1.14+iostat+arcstat. **Software & storage configuration:** Open-E JovianDSS ver: 1.0 up28 (b37311); Pool architecture: 12x single disk (no redundancy); Write log (ZIL): single NVMe disk; Read cache (L2ARC): none; Read cache (ARC): 144GB (75% of 192GB RAM); zvol size: 200GB; zvol block size: 4KB; zvol provisioning: thick; **Hardware specification:** Server model: TAROX ParX R2242i G6 server; CPU: Intel® Xeon™ Silver 4208R CPU @ 2.20GHz; RAM: 192GB; Capacity storage: 12x Intel® SSDPE2NV153T8; SSD/NVMe (write log): Intel® P4800 / Intel® P3700; SSD/NVMe (read cache): None

[5] Test was performed locally on the operating system. Random write workload was generated by fio tool and results - IOPS, Latency and QoS, were read from kernel diskstats (iostat). Additionaly, ZFS arcstats were used to monitor that 100% cache was used during the test. **Test data:** Test data size: 200GB; Workload pattern: 100% random write; Used IO engine: libaio; Outstanding IO: 256; Number of threads: 32; Test block size: 4KB; Test tools: fio-2.1.14+iostat+arcstat. **Software & storage configuration:** Open-E JovianDSS ver: 1.0 up28 (b37311); Pool architecture: 12x single disk (no redundancy); Write log (ZIL): single NVMe disk; Read cache (L2ARC): none; Read cache (ARC): 144GB (75% of 192GB RAM); zvol size: 200GB; zvol block size: 4KB; zvol provisioning: thick; **Hardware specification:** Server model: TAROX ParX R2242i G6 server; CPU: Intel® Xeon™ Silver 4208R CPU @ 2.20GHz; RAM: 192GB; Capacity storage: 12x Intel® SSDPE2NV153T8; SSD/NVMe (write log): Intel® P4800 / Intel® P3700; SSD/NVMe (read cache): None

[6] Test was performed localy on the operating system. Random read workload was generated by fio tool and results (IOPS) were read from kernel diskstats (iostat). Additionaly, ZFS arcstats were used to monitor that 100% cache was used during the test. (1) ARC only test; (2) ARC+L2ARC test. **Test data:** Test data size: 100GB; Workload pattern: Random read; Used IO engine: libaio; Outstanding IO: 8; Number of threads: 8; Test block size: 4KB; Test tools: fio-2.1.14+iostat+arcstat. **Software & storage configuration:** Open-E JovianDSS ver: 1.0 up28 (b37311); Pool architecture: 12x single disk (no redundancy); Write log (ZIL): Single NVMe disk, 375GB; Read cache (L2ARC): Single NVMe disk, 375GB – (2); Read cache (ARC): 144GB (75% of 192GB RAM) – (1) 48GB (75% of 64GB RAM) – (2); Zvol size: 1TB; Avol block size: 4KB; Zvol provisioning: thick. **Hardware specification:** Server model: TAROX ParX R2242i G6 Server; CPU: Intel® Xeon™ Silver 4208R CPU @ 2.20GHz; RAM: 192GB – (1) 64GB – (2); Capacity storage: 12x Intel® SSDPE2NV153T8; SSD/NVMe (write log): Intel® P4800; SSD/NVMe (read cache): Intel® P4800

[7] 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-dc-p4800x-brief.html.