

Article

Toward Ultra-Low Latency SSDs: Analyzing the Impact on Data-Intensive Workloads

Insoon Jo

Center for Creative Convergence Education, Hanyang University, Seoul 04763, Republic of Korea; insoonjo@hanyang.ac.kr

Abstract: The recent trend in hyper-scale computing has shown that applications are getting hungrier for low, consistent latency for predictably fast service. In accordance with this change, storage vendors have innovated their SSDs to keep ultra-low latencies while delivering high IOPS. Even though these SSDs can provide significant value to a wide variety of workloads, their practical impact on performance has not been discussed much. In this context, our study aims to empirically examine the impact of emerging ultra-low latency SSDs on workload performance, particularly in comparison to the latest conventional SSDs. We conduct benchmarking using a diverse set of data-intensive workloads to comprehensively assess the performance of both ultra-low latency SSDs and conventional SSDs. Our research provides an in-depth performance analysis of the examined SSDs, shedding light on their comparative strengths and weaknesses. The results of our study contribute valuable insights into the practical implications of adopting ultra-low latency SSDs across various workloads, guiding future decisions in storage technology.

Keywords: low-latency SSDs; data-intensive workloads; hyper-scale computing



Citation: Jo, I. Toward Ultra-Low Latency SSDs: Analyzing the Impact on Data-Intensive Workloads. *Electronics* **2024**, *13*, 174. <https://doi.org/10.3390/electronics13010174>

Academic Editors: Pradeep Kumar Singh, Zoltán Illés and Fausto Pedro García Márquez

Received: 15 November 2023
Revised: 19 December 2023
Accepted: 28 December 2023
Published: 30 December 2023



Copyright: © 2023 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Achieving lower and more consistent latency has always been the goal of storage solutions. In the pursuit of minimizing the latency between processors and storage devices, new technologies based on connectivity options such as remote direct memory access (RDMA) [1] and non-volatile memory express (NVMe) [2] have entered the market. RDMA has facilitated direct memory access between computers, improving data transfer efficiency with minimal CPU involvement. NVMe, a protocol optimized for swift communication between a computer's storage devices and the CPU over PCIe, also has significantly reduced data transfer latency. These technological advancements were particularly noteworthy in addressing the historical lack of striking progress in storage device development and media over recent decades. Since 2015, device vendors have been announcing innovations for next-generation storage solutions [3–7]. The key features of these products are ultra-low and consistent latencies that have never been seen before. Thus, the industry has fueled the development of non-volatile memories (NVMs)—Intel and Micron's 3D XPoint [5,8] and Samsung's Z-NAND [9]—that are faster than any conventional NAND flash.

Storage, software, and service vendors have shown great expectations for new NVM technology, which they would benefit from, in the exploration of a tool to meet market expectations. One of the recent trends in hyper-scale computing is that data center workloads demand predictable and low latency for their even service quality [10,11]. Another trend is the evolution of networking; 5G networks target their minimum latency as less than 1 ms in the air [12]. The latency of a device would be more transparent without concealment in this environment and the stable ultra-low latency characteristic of SSDs fits perfectly into these trends.

We sense that there is a strong desire in the technical community for an empirical analysis of their performance because they have the unexamined potential to deliver much

better performance than conventional SSDs. In order to accurately judge such potential, one must answer the following questions: (1) How do ultra-low latency (ULL) SSDs distinguish themselves from conventional SSDs? What are interesting use cases for them? (2) What kind of workloads can benefit most from using them? (3) Are they over-specified? What will be their positions in the market? What may be challenges that prevent them from reaching their full potential?

This paper presents an empirical comparative study of ULL and traditional storage devices to address the questions above. To date, it is the first comprehensive performance analysis reported. In the context of both storage and application performance, we examine ULL SSDs using representative data-intensive workloads and compare their performance to those of state-of-the-art NVMe SSDs. We intend to cover evolving solutions for hyper-scale service as well as traditional mission-critical applications, therefore finding that the performance gain varies considerably with each workload.

2. Ultra-Low Latency SSDs

We define ULL SSDs as a group of SSD technologies that feature lower latencies than conventional NAND solutions (hereinafter referred to as ‘CONV’). The key defining requirement for them is very low latencies at both chip and device levels. Table 1 lists the details of ULL SSDs that have been announced so far. The latest SSDs of datacenter-class (DC) and enterprise-class (EP) are also included for comparison with them.

Table 1. Specifications of conventional and ultra-low latency SSDs.

	DC SSD	EP SSD	Optane SSD	Z-SSD
Features				
Model No.	PM963 [13]	PM1725a [14]	DC P4800X [5]	SZ985 [15]
Capacity (TB)	0.96/1.92/3.84 [#] /7.68	0.8/1.6/3.2/6.4 [#]	0.375 [#]	0.8 [#]
Memory	TLC V-NAND	TLC V-NAND	3D XPoint	Z-NAND
Active/idle power (W)	7.5/2.5	23/7.7	18/5	9/2.5
Endurance (DWPD)	1.3 for 3 years	5 for 5 years	30 for 3 years	30 for 5 years
Latency—Typical				
Typical (R/W)	79/22 [*]	90/20	10/10	12–20/16
Latency—QoS: 99.999%				
QD1 (R/W)	185/69 [*]	179/490 [*]	60/100	50/50
QD16 (R/W)	294/940 [*]	258/1736 [*]	150/200	100/350
QD1 (70/30 mix)	3088/141 [*]	4384/217 [*]	74/74 [*]	86/47 [*]
QD16 (70/30 mix)	3568/124 [*]	7520/860 [*]	115/115 [*]	153/588 [*]
IOPS—4KB Random				
Max. read	430 K	800 K(×4)/1.8 M(×8)	550 K	750 K
Max. write	40 K	160 K(×4)/170 K(×8)	500 K	170 K
Max. 70/30 mix	118 K [*]	295 K [*]	500 K [*]	310 K [*]
Bandwidth (GB/s)				
Sequential read	2	3.3/6.4	2.4	3.2
Sequential write	1.2	2.95/3	2	3.0

[#] SSD capacity tested, ^{*} Measured value.

2.1. Characteristics

ULL SSDs exhibit distinct characteristics that set them apart from conventional SSDs, showcasing impressive random I/O performance, high endurance, and an exceptional quality of service (QoS), even in mixed pattern workloads. The key disparities are highlighted below:

- **Ultra-low read latencies.** ULL SSDs specify random read and write latencies of less than 20 μ s for 4 KB. Notably, their read latency significantly outperforms CONV, surpassing DC/EP SSDs by up to nine times, as indicated in Table 1. This substantial advantage is particularly evident in read cache performance, a primary usage scenario

for SSDs. While ULL SSDs do not exhibit a drastic improvement in write latency compared to CONV, this is less critical for most SSD applications. Storage system vendors commonly leverage battery-backed non-volatile RAM (NVRAM) to enhance write performance while ensuring permanence [16].

- **Outstanding QoS.** ULL SSDs excel in quality of service, with a remarkable latency profile, especially under mixed pattern workloads. The 99.999% read latency of ULL SSDs, at QD1, is 50–60 μ s, representing only 1/3 of CONV's latency. This superiority becomes even more pronounced under mixed pattern workloads (70/30 mix), where ULL SSDs demonstrate 115–153 μ s latency, contrasting with DC/EP SSDs that exhibit millisecond-level latency. Importantly, ULL SSDs officially declare their 99.999% latency as part of device specifications, underscoring the vital role of storage device QoS in ensuring reliable and predictable performance for mission-critical systems.
- **Strong endurance.** ULL SSDs boast a substantial endurance advantage over CONV, especially evident in the case of 3D XPoint, with endurance surpassing NAND flash by 1000 \times . Additionally, Z-NAND's endurance is significantly improved by employing fewer bits per cell [9]. As indicated in Table 1, ULL SSDs offer up to 30 drive writes per day (DWPD), surpassing CONV by 10–40 times. This endurance advantage is particularly advantageous for high-demand storage applications such as cache, tiered layers, and OS paging areas, mitigating concerns about limited write endurance, especially in read/write mixed pattern workloads with frequent data swap-outs.

2.2. Use Cases

ULL SSD's key features are (1) consistent low latency and (2) strong endurance. The use cases associated with such characteristics can be categorized into two groups: fast storage and memory extender. As illustrated in Figure 1, fast storage use includes replacing conventional data/log devices with ULL SSDs.

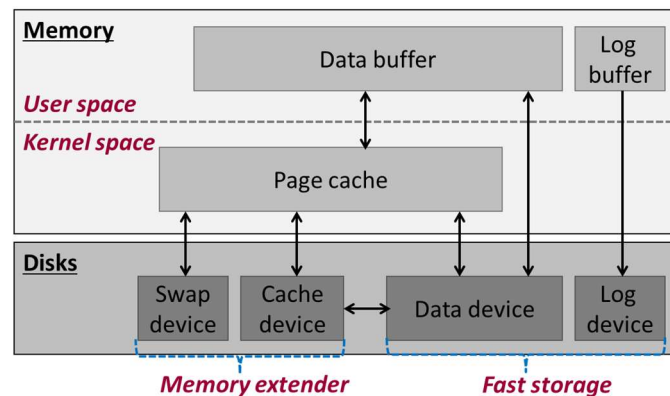


Figure 1. Use cases of ULL SSDs.

2.2.1. Predictably Fast Storage

In data centers, latency QoS matters in that extremely high levels of QoS such as five nines are demanded and are actually being provided for mission-critical workloads. However, keeping low latencies while delivering high IOPS has been the most challenging task for CONV. High IOPS requires high QD for massive internal parallelism [17], while low QD should be maintained for low latencies. On the other hand, the consistent low latency feature enables ULL SSDs to keep a low queue depth relatively easily and not suffer the long tail problem [18]. Data center storage systems will benefit most from this unique feature to enable predictably fast services.

2.2.2. Memory Extender

Storage tiering for price-to-performance has several implementations which can be classified into two groups: swap and cache. Swap [19] uses SSDs as part of virtual memory.

Also, cache usage (e.g., disk caching [20–22]) caches disk data to SSDs mainly to make random access faster. The problem is that high endurance use cases (e.g., excessive swap, paging, or key-value operations) cause the rapid wear of such SSD memory extenders. In contrast, the up to $40\times$ boosted endurance as well as the consistently low latency of ULL SSDs makes it practically attractive to adopt them for a cache tier.

3. Performance Characterization

This section presents the performance characteristics of ULL SSDs that are obtained from running a variety of workloads, most of which are typically found in data centers. Our goal is to reveal what aspects of them affect application performance and QoS and by how much. The observed characteristics are presented in two parts, storage performance and application-level performance.

3.1. Experimental Setup

We conducted our experiments on a Dell PowerEdge R730 server, featuring a dual Intel Xeon(R) CPU E5-2699@2.30GHz (18 threads per socket) and 256 GB of DRAM. The server ran on a 64-bit Ubuntu 18.04 with kernel version 4.6.3. The specifications of the examined SSDs are provided in detail in Table 1.

To facilitate a comparative analysis, we selected ultra-low latency SSDs from two prominent manufacturers, Intel and Samsung. Specifically, our focus was on Intel's Optane SSD, showcasing cutting-edge technology, and Samsung's Z-SSD, highlighting the advancements made by Samsung in the same domain. Our primary goal was to assess and compare the performance of these SSDs against the high-end legacy counterparts chosen from the data center and enterprise categories, namely, Samsung's PM963 and PM1725a.

The rationale behind selecting PM963 and PM1725a lies in their status as high-end legacy SSDs, offering robust solutions for data center and enterprise environments. PM963 is anticipated to showcase the prowess of data center-class SSDs, emphasizing elements like the sustained performance and reliability crucial for data center operations. Conversely, PM1725a is expected to exemplify the advanced features and performance characteristics that define enterprise-class SSDs, demonstrating innovation and capabilities tailored for enterprise-level storage requirements. The intention is to highlight and leverage the distinctive strengths inherent in each product.

Through this meticulous selection process, our research aims to provide a comprehensive understanding of the comparative performance, strengths, and nuances of ultra-low latency SSDs from Intel and Samsung, alongside other high-end legacy SSDs.

3.2. Storage Level Performance

3.2.1. Mixed Use

In investigating ULL SSD performance, we measured both IOPS and QoS using a synthetic 70/30 read/write mix, generated with the FIO tool. This mix reflects a common real-world workload, providing a practical evaluation of the SSD's capabilities under typical usage patterns and contributing to a more comprehensive understanding of its efficiency in handling the diverse data access patterns encountered in operational environments. A high IOPS at low QD and outstanding QoS under mixed patterns are known as their most distinctive features, and we wanted to verify the two using real systems. First, we observed that a high IOPS at low QD does not clearly distinguish ULL SSDs from CONV. As shown in Figure 2a, the former achieve their maximum IOPS at a low QD of 16. However, EP SSD also hits a peak at QD32, while DC SSD is still not near its maximum with a QD of less than 64. In contrast, the QoS under a mixed workload is a unique characteristic of ULL SSDs. Figure 2b shows a 99.999% QoS distribution that clearly categorizes ULL SSD and CONV into two sections. Interestingly, EP SSD shows a QoS distribution even worse than DC SSD. A high IOPS at high latencies is typical in CONV in that a high throughput is achieved by massive parallelism that sacrifices latencies [17].




Figure 2. Mixed-use (4 KB random 70/30 read/write) (a) IOPS and (b) 99.999% latency.

3.2.2. With User-Level Libraries

User-level device drivers address the issue of using costly kernel interrupt and context switches by instead polling within the user space. As shown in Figure 3, their effect becomes more noticeable at a lower QD than the ULL SSDs target. By employing user-level drivers, the IOPS performance of Optane SSD and Z-SSD were boosted up to 210% and 16%, respectively. This shows that the increased potential of ULL SSDs will be realized by developing a user space code or lightening kernel-level libraries, whose good examples can be the user-level drivers above.

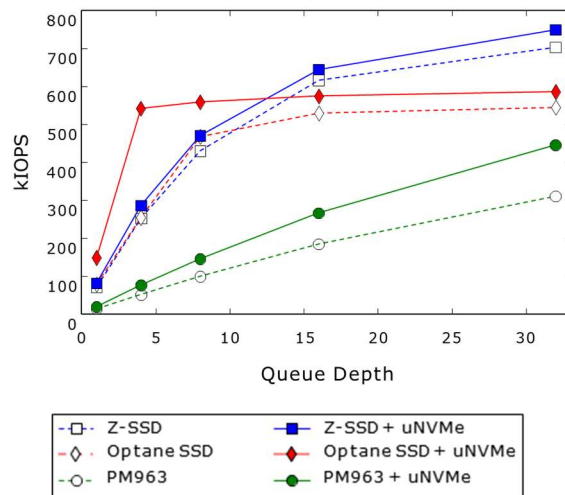


Figure 3. IOPS with the user-level device driver.

3.3. Application Level Performance

In order to characterize ULL SSDs in the context of application performance, we ran benchmark workloads that would exhibit a wide range of behaviors. Applying the use case categories in Section 2.2, we divided them as shown in Table 2. We aimed to use widely recognized tools that are publicly available and commonly adopted in the industry as well as academia. Each serves a specific purpose in testing our system across diverse scenarios. For SQL-related assessments, we relied on PostgreSQL (PgSQL), an open-source relational database system esteemed for its robustness and compliance with standards.

Table 2. Target workloads.

Use Case	Workloads
Fast storage–Data device	SQL: PostgreSQL NoSQL: Aerospike(KV), RocksDB(KV), MongoDB(Docu.), and FlashX(Graph) Distributed storage: Ceph Write stress test: WPT
Fast storage–Log device	SQL: PostgreSQL NoSQL: RocksDB
Memory extender–Cache device	KV cache: Memcached and Fatcache Disk cache: Bcache Read-only cache: RocksDB
Memory extender–Swap device	PMbench

In the dynamic realm of NoSQL databases, our choices include Aerospike, RocksDB, MongoDB, and FlashX, encompassing a range of data models commonly adopted in various applications. In scenarios demanding distributed storage, we turned to Ceph—an established open-source distributed storage system known for its scalability and reliability. To stress-test our system under high-stress write scenarios, we employed the Write Stress Test (WPT) tool which aligns with the industry practices for evaluating system robustness. For log devices, we incorporated RocksDB, an open-source embedded key-value store, tailored to assess the specific performance requirements of this context. In the realm of caching, our toolkit included Memcached and Fatcache for key-value caching, Bcache for disk caching, and RocksDB for read-only caching scenarios. To evaluate the potential of persistent memory as a swap device, we enlisted PMbench—an open-source benchmark designed for this specific purpose. Each of these selections was driven by a commitment to using reputable, open-source solutions that are readily accessible within the broader community. Evaluation results corresponding to each workload are illustrated in Figure 4.

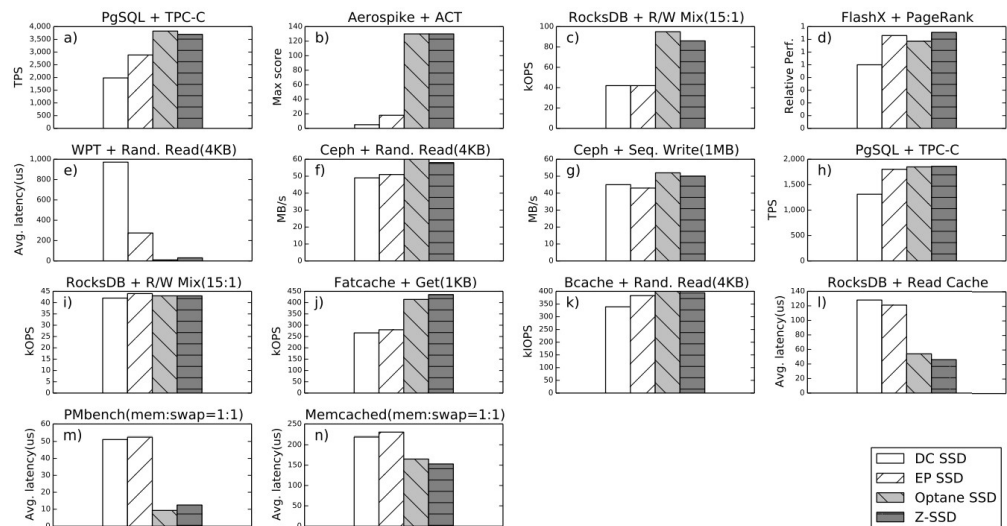


Figure 4. Application-level results with SSDs across different use cases: (a–g) Data Device, (h,i) Log Device, (j–l) Cache Device, (m,n) Swap Device; (a) TPC-C Throughput on PostgreSQL, (b) Aerospike ACT Score, (c) R/W Mix Throughput on RocksDB, (d) PageRank Performance on FlashX, (e) Random Read Latency on WPT, (f) Random Read Throughput from Rados on Ceph, (g) Sequential Write Throughput from Rados on Ceph, (h) TPC-C Throughput on PostgreSQL, (i) R/W Mix Throughput on RocksDB, (j) Get Throughput with Fatcache, (k) Random Read Throughput with Bcache, (l) Average Latency with SSD as Read Cache, (m) Average Latency with PMBench Swap, (n) Average Latency with Memcached Swap.

3.3.1. Performance as Data Device

The results obtained from using ULL SSDs and CONV as data storage are illustrated in Figure 4a–g.

The performance gain obtained by ULL SSDs. How much ULL SSDs affect the overall performance of each workload varies, but they boosted performance up to $26\times$. Their effect as data storage is determined by (1) how perfectly the workload characteristics fit into them and (2) how complex the legacy software is. Among all the workloads we ran, the most suitable ones for ULL SSDs are the Aerospike Certification Tool (ACT) test (Figure 4b) and Dell-EMC WPT (Figure 4e) in that they require a high throughput and QoS simultaneously. Indeed, they passed $7.2\text{--}26\times$ higher (heavier) ACT load levels than CONV. Furthermore, during the WPT run, Optane SSD and Z-SSD took $95\times$ and $31\times$ lower random read latencies than the DC SSD, ACT, and WPT model I/O patterns in many real-time database servers and virtual machine environments, respectively. Overwhelming ACT records and the latency gain from ULL SSDs prove that they will perfectly fit into real-time mixed pattern workloads. Though they did not obtain as much as ACT, the maximum increases in the throughputs seen with PostgreSQL and RocksDB reached up to $1.9\times$ and $2.3\times$, respectively. Even though omitted, MongoDB and Cassandra showed performance improvements similar to RocksDB. This is an interesting point. Prior work forecasted that the drop-in replacement for remarkably fast NVMe SSDs would yield a $1.4\times$ speedup at most [23]. Different from this prediction, a drop-in replacement of CONV with ULL SSDs brought an approximately $2\times$ performance gain.

Remaining issues and challenges. Despite such impressive performance gains, ULL SSDs still have price-to-performance and capacity issues as data storage; however, they are relatively small compared to DC/EP SSDs. Moreover, 3D XPoint-based SSDs are roughly $4\text{--}5\times$ more expensive than CONV and half cheaper than DRAM [24]. The real challenge is that their performance advantages at the storage level can be more or less diminished by legacy software. Regarding the read-while-writing test we ran on RocksDB (Figure 4c), we think it was bounded by typical read latencies. Not only is the 65 MB/s write load negligible for the target SSDs, but RocksDB also includes in-memory bloom filters and indices that lower the need for a high IOPS. Indeed, DC/EP SSDs have similar typical latencies and achieve similar throughputs. In this case, even though an over $2\times$ increase in throughput is considerable, the up to $9\times$ latency advantage of ULL SSDs prevented them from reaching their potential due to the complexity of legacy database software. The worst case occurred when running the Rados benchmark [25] on Ceph (Figure 4f,g); this complex object storage system made the performance of all SSDs uniformly low. A promising way to mitigate the high price-to-performance and severe performance drop explained above is to develop a kernel/user-level software environment that can extract more performance benefits from ULL SSDs.

IOPS asymmetry. In terms of IOPS, EP SSDs may be comparable or rather superior to ULL SSDs. As shown in Figure 4d, when running PageRank on the FlashX graph database, EP SSD and Z-SSD delivered a $1.06\text{--}1.1\times$ higher performance than Optane SSD. This is because the benchmark is affected mainly by read IOPS. The read/write IOPS of Optane SSD is almost symmetric, which is a unique feature. On the other hand, the ratio of read/write IOPS of EP SSD and Z-SSD is in the range of 4 or $4\text{--}5$ to 1 (Table 1). Each has its merits and faults, i.e., SSDs with read/write IOPS symmetry can deliver almost constant throughput for any ratio of read/write mix; in contrast, workloads driven by read IOPS (e.g., PageRank on FlashX in Figure 4d) benefit more from SSDs with read/write IOPS asymmetry.

3.3.2. Performance as Data Device

Unremarkable impact on logging. Figure 4h,i shows how ULL SSD-based logging affects the overall throughput. PostgreSQL obtained a $1.4\times$ higher throughput, while RocksDB had no noticeable gain. In the case of PostgreSQL logging, EP SSDs and ULL SSDs delivered almost the same throughput. For this workload, we can consider that the throughput was

a key to the overall performance rather than latencies. Also, the logging throughput was enough to affect performance, but not that large enough to saturate SSDs, except DC SSD. Therefore, ULL SSD-based log storage would not be attractive for this throughput-oriented use. In the case of RocksDB, it is very interesting that a far faster log storage did not contribute to throughput at all. A single write thread on RocksDB seemed to generate a small amount of log streams that even DC SSD could process quickly. Different from fast storage use, using ULL SSDs as log storage may require more careful consideration.

3.3.3. Performance as Cache Device

Anticipated advancement as a system accelerator. DRAM is still cost-bearing and volatile as well as has challenges like power consumption and operational complexity in scaling architecture. Using SSD as a cache device has a higher capacity per dollar and lower power consumption per byte, without degrading random read latency beyond network latency. The results show that ULL SSDs achieved a 1.5–1.6 \times throughput improvement compared to using DC/EP SSDs (Figure 4j–l). The workload in Figure 4j was bounded by typical latencies taken to read 1 KB. Similarly, using ULL SSDs as persistent read cache (Figure 4l) decreased the latencies by a factor of 2.2–2.8 \times . Compared to the prior results, random reads when using ULL SSDs as a Bcache drive showed a somewhat lowered IOPS achievement (1.1–1.2 \times) due to the low cache hit ratio (Figure 4k). Compared to data storage use cases, cache usage is bounded by latencies, and, thus, ULL SSDs beat CONV by a factor of up to 2.8 \times .

3.3.4. Performance as Swap Device

Strongly recommended for swap partition. The swap space on a disk has become vital in virtualized OS environments because hypervisors adopt the ‘Dynamic Memory’ feature for their virtual machine [26,27]. Dynamic Memory offers the flexible management of limited resources but can cause serious competition at the same time. ULL SSDs can be perfect candidates for swap partitions. In order to evaluate ULL SSDs as swap devices, we ran PMbench [19] and memcached [28] with a swap on them. The swap workload from PMbench was affected by typical latencies, and, thus, ULL SSDs required 4.1–5.6 \times lower latencies for file access (e.g., Figure 4m). Also, when running a more practical memcached workload, the average time taken to reach values was up to 1.5 \times lowered, as shown in Figure 4n. Therefore, the latency-bounded swap workload turned out to benefit considerably from ULL SSDs.

4. Discussion

Based on our results and experiences, we make several observations and remarks.

Utility of ULL SSDs. ULL SSDs are similar in low latency profiles for read throughputs but have different specifications at the same time—write latency and sequential throughput—as shown in Table 1. Z-SSD has a high throughput and lower latency, mainly in read, while Optane SSD shows relatively stable performance in both read and write throughput. As shown in most of our experiments, there is hardly a difference in the application-level performance due to the relatively weak write throughput: only WPT shows a little effect from it. As Table 1 shows, Z-SSD has both read/write throughputs close to the interface limit of PCIe Gen 3 \times 4 while demonstrating low-latency characteristics similar to ULL SSDs. Therefore, Z-SSD effectively operates applications that require a low latency and high throughput simultaneously. On the other hand, Optane SSD seems to have plenty of room to optimize in terms of throughput. As we found in Section 3.2, the throughput of the device is still the most important factor for application performance, even though the device is a ULL SSD.

Opening of SCM era. The market readily classifies ultra-low latency (ULL) SSDs as storage class memory (SCM), positioning them as the long-awaited ‘gap-filler’ between DRAM and conventional storage (CONV) [29]. SCM represents a category of non-volatile memory that combines the speed characteristics of dynamic random-access

memory (DRAM) with the persistence of traditional storage. ULL SSDs have emerged as pioneering devices [30,31], fulfilling the expectations for latency, capacity, and permanency inherent to SCM [32,33]. In a role similar to NAND flash in the past, ULL SSDs function as system accelerators through caching or in higher storage layers above NAND, thereby contributing to an overall improvement in system speed. Despite a price that is 3–10× that of CONV, this underscores their significance [34]. However, there is another view that the opening of a new era would be slackened by the enterprise storage server market, which is known to be very conservative and to never pick up unconfirmed media [35].

Readiness of applications and systems. As we observed in Figure 4b, ULL SSDs can handle tens of times of workloads more than CONV in applications that require the extremely short read latencies in the read/write mixed conditions represented by ACT. However, most of the application performances of ULL SSDs are much better than those of CONV. This implies two important present conditions of the market: first, traditional applications and systems have been ultimately optimized for traditional storage and devices. Second, the emerging application for hyper-scale services requires the innovation of storage and devices. These are getting more definite with the gap between performance with a user-level library and performance with a kernel-level library, as shown in Section 3.1. Applications and systems need to evolve into another optimized form to maximize the benefit of ULL SSDs, powered by a lightened OS and driver for low latency.

5. Related Work

ULL SSDs have opened the storage-class memory (SCM) era, which is about a decade-old concept [33,36]. Vendors have also developed other SCM solutions that feature ultra-low latency as well as DRAM-like accessibility, e.g., 3D XPoint products with DIMM form factors [37]. Although there is already NVDIMM-N in the market, it is not enough to be called SCM due to its limited capacity, which is less than 32 GB [38]. NVDIMM-P [6] belongs to this category and can compete against 3D XPoint-based products. Specifically, NVDIMM-P unites DRAM and NVM on DIMMs and offers more density and lower prices than DDR4 solutions along with 3D XPoint-based DIMMs. The software can benefit from this byte-addressable SCM (i.e., an in-memory database like SAP HANA) which gives them their enthusiastic welcome [39]. OS and software-defined storage also participate in this campaign [40,41], so we need to stay tuned for upcoming events around SCM.

6. Conclusions

This paper is the first empirical comparative study to characterize ULL SSDs in the context of both storage and application performance. They have been shown to have distinctive storage features, such as consistent ultra-low latencies and high endurance, which make them the best fit for storage for real-time mixed pattern workloads and consistently fast memory extenders. However, we believe ULL SSDs still have more performance potential that should be realized, and, thus, demand kernel and user-level libraries and code to be carefully designed and optimized for them.

Funding: This work was supported by the National Research Foundation of Korea grant funded by the Korean government (MSIT) (RS-2023-00250918).

Data Availability Statement: Data are contained within the article.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Recio, R.; Metzler, B.; Culley, P.R.; Hilland, J.E.; Garcia, D. A Remote Direct Memory Access Specification. *RFC* **2007**, *5040*, 1–66.
2. NVM Express. Available online: <http://www.nvmexpress.org> (accessed on 30 December 2023).
3. Cutress, I. Samsung Shows Off A Z-SSD: With New Z-NAND. 2017. Available online: <https://www.anandtech.com/show/11206/samsung-shows-off-a-z-ssd> (accessed on 30 December 2023).
4. Mellor, C. Say Hello to Samsung and Netlist's Flash-DRAM Grenade: HybriDIMM. 2016. Available online: https://www.theregister.co.uk/2016/08/08/samsung_and_netlist_hybridimm/ (accessed on 30 December 2023).

5. Pirzada, U. Intel Launches Optane SSD DC P4800X, Enters Revenue Quarter For 3D XPoint, Storage Roadmap Detailed. 2017. Available online: <http://wccfttech.com/intel-optane-ssd-dc-p4800x-revenue-quarter-3d-xpoint-roadmap/> (accessed on 30 December 2023).
6. Shilov, A. JEDEC: DDR5 to Double Bandwidth over DDR4, NVDIMM-P Specification Due Next Year. Available online: <http://www.anandtech.com/show/11238/ddr5-to-double-bandwidth-over-ddr4-specification-due-next-year> (accessed on 30 December 2023).
7. Tallis, B. Micron Announces QuantX Branding For 3D XPoint Memory (UPDATED). 2016. Available online: <https://www.anandtech.com/show/10556/micron-announces-quantx-branding-for-3d-xpoint-memory> (accessed on 30 December 2023).
8. Intel and Micron Produce Breakthrough Memory Technology. 2015. Available online: <https://www.intc.com/news-events/press-releases/detail/324/intel-and-micron-produce-breakthrough-memory-technology> (accessed on 30 December 2023).
9. Tallis, B. Samsung at Flash Memory Summit: 96-Layer V-NAND, MLC Z-NAND, New Interfaces. 2017. Available online: <https://www.anandtech.com/show/11703/samsung-at-flash-memory-summit-96layer-vnand-mlc-znand-new-interfaces> (accessed on 30 December 2023).
10. Dean, J.; Barroso, L.A. The Tail at Scale. *Commun. ACM* **2013**, *56*, 74–80. [CrossRef]
11. Khan, F. The Cost of Latency. 2015. Available online: <https://www.digitalrealty.com/resources/articles/the-cost-of-latency> (accessed on 30 December 2023).
12. NGMN 5G White Paper. Available online: https://www.ngmn.org/wp-content/uploads/NGMN_5G_White_Paper_V1_0.pdf (accessed on 30 December 2023).
13. Samsung SSD PM963. 2016. Available online: https://www.compuram.de/documents/datasheet/Samsung_PM963-1.pdf (accessed on 30 December 2023).
14. Samsung PM1725a NVMe SSD. 2018. Available online: https://download.semiconductor.samsung.com/resources/brochure/Brochure_Samsung_PM1725a_NVMe_SSD_1805.pdf (accessed on 30 December 2023).
15. Ultra-Low Latency with Samsung Z-NAND SSD. 2017. Available online: <https://semiconductor.samsung.com/resources/brochure/Ultra-Low%20Latency%20with%20Samsung%20Z-NAND%20SSD.pdf> (accessed on 30 December 2023).
16. Woods, M. Optimizing Storage Performance and Cost with Intelligent Caching. 2010. Available online: <https://api.semanticscholar.org/CorpusID:16054646> (accessed on 30 December 2023).
17. Watanabe, K. UFS 2.0 NAND Device Controller with SSD-Like Higher Read Performance. In Proceedings of the Flash Memory Summit, Santa Clara, CA, USA, 5–7 August 2014.
18. Xu, Y.; Musgrave, Z.; Noble, B.; Bailey, M. Bobtail: Avoiding Long Tails in the Cloud. In Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation, Lombard, IL, USA, 2–5 April 2013; pp. 329–341.
19. Yang, J.; Seymour, J. Pmbench: A Micro-Benchmark for Profiling Paging Performance on a System with Low-Latency SSDs. In Proceedings of the Information Technology-New Generations: 14th International Conference on Information Technology, Las Vegas, NV, USA, 10–12 April 2017; pp. 627–633.
20. Flashcache. Available online: <https://en.wikipedia.org/wiki/Flashcache> (accessed on 30 December 2023).
21. Stec-Inc/Enhanceio. Available online: <https://github.com/stec-inc/EnhanceIO> (accessed on 30 December 2023).
22. Ali, A.; Rose, C. Bcache and Dm-Cache: Linux Block Caching Choices in Stable Upstream Kernel. 2014. Available online: <https://www.scribd.com/document/484146870/LinuxBlockCaching> (accessed on 30 December 2023).
23. Xu, Q.; Siyamwala, H.; Ghosh, M.; Suri, T.; Awasthi, M.; Guz, Z.; Shayesteh, A.; Balakrishnan, V. Performance analysis of nvme ssds and their implication on real world databases. In Proceedings of the SYSTOR '15: Proceedings of the 8th ACM International Systems and Storage Conference, Haifa, Israel, 26–28 May 2015.
24. Mearian, L. Micron Reveals Marketing Details about 3D XPoint Memory QuantX. 2016. Available online: <https://www.computerworld.com/article/3104675/data-storage/micron-reveals-3d-xpoint-memory-quantx.html> (accessed on 30 December 2023).
25. Pustina, L. Ceph Object Storage as Fast as It Gets or Benchmarking Ceph. 2014. Available online: <https://www.codecentric.de/wissens-hub/blog/ceph-object-storage-fast-gets-benchmarking-ceph> (accessed on 30 December 2023).
26. Guo, F. Understanding Memory Management in VMware vSphere 5. 2011. Available online: https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/mem_mgmt_perf_vsphere5.pdf (accessed on 30 December 2023).
27. Hyper-V Dynamic Memory Overview. 2014. Available online: [https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-r2-and-2012/hh831766\(v=ws.11\)](https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-r2-and-2012/hh831766(v=ws.11)) (accessed on 30 December 2023).
28. Fitzpatrick, B.; Vorobey, A. Memcached—A Distributed Memory Object Caching System. 2011. Available online: <https://memcached.org/> (accessed on 30 December 2023).
29. Evans, C. HPE Demos 3PAR with Intel Optane (3D-XPoint). 2016. Available online: <https://community.hpe.com/t5/Around-the-Storage-Block/3PAR-and-Storage-Class-Memory-From-Ludicrous-Speed-to-Plaid/ba-p/6900913> (accessed on 30 December 2023).
30. Arnold, R. 3PAR 3D Cache. 2016. Available online: <https://d8stadude.com/2016/12/16/3par-3d-cache-3dxpoint-2/> (accessed on 30 December 2023).
31. Optane Now Available on IBM Cloud, but with Limited Uses. 2017. Available online: <https://www.techerati.com/the-stack-archive/data-centre/2017/08/02/optane-now-available-on-ibm-cloud-but-with-limited-uses/> (accessed on 30 December 2023).
32. Burr, G.W. Towards Storage Class Memory. 2013. Available online: https://s3.us.cloud-object-storage.appdomain.cloud/res-files/922-SCMandMIEC_overview_12Feb2013.pdf (accessed on 30 December 2023).

33. Freitas, R.F.; Wilcke, W.W. Storage-class memory: The next storage system technology. *IBM J. Res. Dev.* **2008**, *52*, 439–447. [CrossRef]
34. Sliwa, C. Intel Optane SSD Starts Shipping to Select Customers. 2017. Available online: <https://www.techtarget.com/searchstorage/news/450415272/Intel-Optane-SSD-starts-shipping-to-select-customers> (accessed on 30 December 2023).
35. Shah, A. HPE Is Bringing Optane Storage to Unix Servers. 2017. Available online: <https://www.networkworld.com/article/963442/hpe-is-bringing-optane-storage-to-unix-servers.html> (accessed on 30 December 2023).
36. Burr, G.W.; Kurdi, B.N.; Scott, J.C.; Lam, C.H.; Gopalakrishnan, K.; Shenoy, R.S. Overview of candidate device technologies for storage-class memory. *IBM J. Res. Dev.* **2008**, *2*, 449–464. [CrossRef]
37. Operating Conditions for Intel Optane Persistent Memory. 2023. Available online: <https://www.intel.com/content/www/us/en/support/articles/000032853/memory-and-storage/intel-optane-persistent-memory.html> (accessed on 30 December 2023).
38. Sainio, A.; Martinez, M. NVDIMM-N Cookbook: A Soup-to-Nuts Primer on Using NVDIMM. 2016. Available online: https://www.snia.org/sites/default/files/tutorials/FMS2016/Sainio-Martinez_NVDIMM_Cookbook_Tutorial_080816.pdf (accessed on 30 December 2023).
39. Ferron-Jones, M. Intel: Persistent Memory, Based on 3D XPoint, Gets First Public Demo. 2017. Available online: <https://www.storagenewsletter.com/2017/05/22/intel-persistent-memory-based-on-3d-xpoint-gets-first-public-demo/> (accessed on 30 December 2023).
40. Wheeler, R. Persistent Memory and Linux: New Storage Technologies and Interfaces. 2013. Available online: https://events.static.linuxfound.org/sites/events/files/eeus13_wheeler.pdf (accessed on 30 December 2023).
41. Joergensen, C. Storage Spaces Direct with Persistent Memory. 2019. Available online: <https://techcommunity.microsoft.com/t5/storage-at-microsoft/storage-spaces-direct-with-persistent-memory/ba-p/425881> (accessed on 30 December 2023).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.