This is a pre-print copy of our accepted paper in ISPASS'26

Characterization and Performance Analysis of Aging Effect on Enterprise HDDs

Ava Dezhban
Sharif University of Technology
Tehran, Iran
ava.dezhban@sharif.edu

Mohammadamin Ajdari*
HPDS Research
Tehran, Iran
m.ajdari@sharif.edu

Hossein Asadi**
Sharif University of Technology
Tehran, Iran
asadi@sharif.edu

***Abstract*—*Hard Disk Drives* (HDDs) are still a popular choice in data centers for cloud storage, due to lowest cost per terabyte and decent performance. Previous works have studied parameters related to disk failure and also proposed prediction models to predict the time that a disk fails. However, *the impact of disk aging on performance* has remained a big mystery. In this paper, we empirically study over 30 enterprise HDDs and analyze how different parameters related to aging (e.g., the amount of previous disk writes and power-on hours) affect the performance. We specifically provide six findings, summarized as four major insights. For example, we reveal that a new disk, after disk being used for over eight moths, shows up to 3.9× higher write latency. At the same time, no latency difference exists between disks older than two-years of age. Overall, a single used disk (with two-years of age) alongside three new disks in a RAID array worsen the I/O latency by 17%. Our findings can help storage architects and admins to (1) ensure high service-level-agreements for HDD-based storage systems, (2) detect used enterprise disks compared to fairly new disks even when disk S.M.A.R.T information are not available or have been altered deliberately to cause misleading information.**

***Index Terms*—Enterprise Disks, Enterprise Systems, Data Storage, Hard Disk Drive, Aging, Performance, SMART**

I. INTRODUCTION

In the modern era of big data, *Hard Disk Drives* (HDDs) remain vital in data centers due to their almost lowest cost per terabyte, high storage capacity, and decent performance. HDDs are widely used for large-scale data storage, with their usage over 90% of exabyte storage in cloud data centers, as reported by Seagate [3]. As another report by IDC published in 2024, HDDs continue to provide 80% of storage capacity in hyper-scale, and cloud data centers by 2028 [57]. Example large-scale HDD users are cloud providers such as Backblaze [1] and supercomputers such as Blue Brain 5 [2].

With such large amount of HDD usage, *in-depth understanding of how aging affects the performance of enterprise HDDs is still a mystery*. Such knowledge has two major benefits. First, it allows to tune the cloud architecture to ensure the desired performance of each HDD and provide better *service level agreements* (SLA). Second, it helps detecting used or refurbished disks or even fake disks compared to new original enterprise disks, which has become a challenge in recent years especially in some geographic regions. By using special tools, some sellers may reset the internal SMART[1] parameters of the disk to appear as brand new, while it has been used for a couple of years, and is thus undetectable for most end-users. If aging effects on HDD performance is demystified, it would provide hints for detecting fake and used disks very well.

Unfortunately, existing studies (a) aim for non-HDDs such as aging effects on *Solid-State Drives* (SSDs) [36], [37], (b) analyze HDDs but focus on failure rates (not aging effects on performance), (c) are neither HDDs nor analyze aging effects (e.g., performance behavior analysis under different workloads running persistent memory DIMMs [26], [55], [58] and emerging CXL memory prototypes [60]). Existing studies on HDDs have mainly focused on correlating disk failures with SMART parameters [21], [22], [35], [53], or predicting HDD failures using machine learning algorithms [6], [12], [30], [31], [33], [54]. Some studies analyze the parameters affecting HDD performance, but such studies *only* consider environmental parameters such as temperature [40] and noise [20]. Thus, (a) *none* of these works analyze how HDD performance changes due to its *aging* in few years of usage in data centers, and (b) these studies *only* target SATA HDDs while SAS HDDs are aimed more for enterprise environments.

In this paper, we empirically study over 30 enterprise SAS HDDs of the same HDD model ranging from new to six-years of age. We analyze how the following parameters affect the disk performance; (a) normal aging (i.e., the amount of hours the disk is ON), (b) amount of previous disk writes, (c) number of spin-up/spin-downs, and (d) the rate of activation for error correction logic. Our analysis leads to *six findings*, which can be summarized into *four major insights*:

First Insight. *The average random write latency of an enterprise HDD directly depends on the disk age (i.e., disk Power-on-Hours (PoH)): disk write latency significantly increases (on average by 1.8×) in the first eight months of operation (compared to brand new disk), another 9% increase in latency when disk PoH value increases from 8-month to 2 years, and after 2-years of age, the disk has almost constant write latency*

---

* Co-first author
** Corresponding author

[1] Self-Monitoring, Analysis, and Reporting Technology (SMART) is a monitoring system used in disk drives.

till the disk becomes faulty.

Second Insight*. The spin-up/spin-down operations (common for power-saving and other benefits in enterprise systems) and the amount of previous writes on enterprise HDDs have practically* ***no*** *effect on disk latency.*

Third Insight*. Disk aging and other types of disk usage (e.g., amount of previous writes) have almost no effect on disk random read latency and sequential read/write bandwidth (as long as disk is not faulty).*

Fourth Insight*. A single HDD with 2-years of age may worsen the write latency of a whole 4-disk RAID array (with other disks being new) by 17.1%. This signifies the importance of careful selection of disks (based on age) in RAID array creation in enterprise storage systems.*

In short, we offer the following **major contributions**:

- **The first characterization study of aging effects on performance of enterprise HDDs**. We analyze different parameters related to disk usage for over 30 SAS HDDs and show how they affect the disk performance.
- **Six findings on the behavior of enterprise SAS HDDs.** We reveal six novel findings on how aging affects the enterprise HDD performance.
- **Monitoring the performance behavior of an enterprise SAS HDD during almost a year period**. We analyze a single disk from day-one till 8 months of its usage and show how performance correlates with the disk age.

II. BACKGROUND

A. HDD Storage Why and Where

Trends. Recent reports by IDC [57] and Seagate [3] in 2024 show that HDDs continue to constitute over 80% of storage devices used in Cloud, and large scale deployments due to significantly lower cost than flash-based storage (i.e., SSDs), and providing decent performance for many applications. A report by Western Digital also indicates that HDDs stay a popular choice alongside other storage devices (e.g., SSDs) due to (a) exponential global data growth and the need to meet capacity demands, and (b) significantly lower cost of HDDs compared to SSDs while providing decent performance [13]. This reports also emphasizes that even though fast storage devices (e.g., SSDs) are required for keeping hot data in the cloud, HDDs are usually suitable for *warm* data (in addition to cold data).

Capacity-Oriented Enterprise HDDs. Enterprise HDDs used in large-scale systems are usually capacity-oriented. For example, almost all 317,000 HDDs of Backblaze data centers are capacity-oriented using 7200 *Rotations per Minute* (RPM) disks with 4TB-22TB capacity [8], [51]. We have listed the characteristics of enterprise HDDs of three major manufacturers in Table I. **First**, such disks have capacity starting from 1TB up to 22TB (in Toshiba MG Series), up to 24TB (in Seagate Exos series), and up to 32TB (in Western Digital Ultrastar HC series). **Second**, the RPM of the capacity-oriented enterprise disks, i.e., the rotational speed of platters in the disk per minute, are 7200 RPM. This is different compared to consumer-grade disks (commonly with 5400 RPM), and low-capacity server disks (with 10,000 RPM or 15000 RPM). **Third**, they are designed to be used for at least five years under excessive I/O loads. **Fourth**, enterprise disks have various data protection mechanisms (e.g., T10-DIF in SAS drives). T10-DIF is a feature that provides *Data Integrity Field* (DIF) defined by T10 technical committee, which pads 8-byte *Cyclic Redundancy Check* (CRC) to each 512-byte data sector. **Fifth**, these disks usually have small NAND or NOR flash cache to speed up random writes, and also small DRAM cache to accelerate both reads and writes. In this direction, the sequential bandwidth is typically around 250 MB/s, the random 4KB read IOPS is 100-250 IOPS, and random 4KB writes is 500-1000 IOPS. **Sixth**, the cost of enterprise HDDs (e.g., Seagate EXOS 10TB SAS HDD [38]) is on average below 30 dollars-per-TB as of today while that of enterprise SSDs (e.g., Samsung PM1653) is over 210 dollars-per-TB [49] (meaning enterprise HDDs are 7× cheaper).

Enterprise HDD model used in this paper. For our analysis, we select *an 8TB enterprise HDD model* from one of the three major vendors and use over 30 disks of the same model but with different amount read/writes, and aged (i.e., used) from *one day up to six years*. We specifically choose 8TB disks because such models are typically manufactured between 2017-2020 while there are enough number of 8TB disk samples in the field with different amount of usage. Disks with lower capacity (e.g., 2TB or 4TB HDDs) are less common in modern data centers. Disks with much larger capacity (e.g., with over 20TB) have been introduced recently, thus studying their aging (or usage) beyond one or two years is impossible in this study.

B. SMART Standard

Considering the high value of data stored on disks, mechanisms to monitor disk internal operations, tracking its read/write accesses and possible encountered errors is crucial to system admins. Almost all HDDs support the standard of *Self-Monitoring, Analysis, and Reporting Technology* (SMART), which implements internal counters to track various events about disk usage and provide warnings such as predictive disk failure states to prevent data loss in the near future.

Self-test. The SMART standard provides built-in self-test capabilities. These tests can be initiated by the system or manually triggered by administrators to assess the drive operational integrity. Common self-tests are *short self-test*, and *extended self-test*. The short self-test usually takes a couple of minutes and examines electro-mechanical components of the disk to detect possible issues. The extended self-test takes several hours and applies a thorough examination of the entire platter surface for possible detection of sector errors or additional mechanical issues.

Parameters. Standard parameters of SMART are different for SATA and SAS disks. Table II shows the SMART parameters for SAS disks. Some of such parameters may indicate faulty disks (e.g., high number of total uncorrected errors) or

Table I: Specification of Sample enterprise SAS HDDs from three major vendors. Specifications are mainly from public datasheets [42]–[47], [14]–[19], [52]. The performance numbers were extracted from model datasheets (e.g., for Seagate drives), or disk review websites (e.g., for random access performance in Western Digital drives [50], and Toshiba drives [11]). Some models do not provide reports on performance numbers, and even for available numbers, test conditions may be different. Thus, slight differences in performance numbers of different brands should not considered as better or worse performance. Instead, we observe that enterprise disks have similar order of performance: almost 250 MB/s sequential bandwidth, 100-250 IOPS for random reads, 500-1000 IOPS for random writes. *: *We intentionally hide the model and the HDD vendor name. These enterprise HDDs are from one of the three major vendors listed in the table.* RPM: Rotations Per Minute.

Brand	**Model Series**	**Capacity**	**RPM**	**Expected Operational Years (=Warranty)**	**Data Protection (on SAS Disks)**	**Seq. Performance (Sample model $\geq$ 4TB)**	**Random 4KB Reads/Writes (for a sample model)**
Seagate	EXOS	2TB–24TB	7200	5 Years	T10-DIF	270 MB/s	170 / 550 IOPS
Western Digital	Ultrastar HC	2TB–32TB	7200	5 Years	T10-DIF	274 MB/s	198 / 663 IOPS
Toshiba	MG	1TB–22TB	7200	5 Years	T10-DIF	285 MB/s	230 / 1000 IOPS
Our Choice (in this paper)	*	**8TB**	**7200**	**5 Years**	**T10-DIF**	**250 MB/s**	**160 / 500 IOPS**

Table II: SMART attributes of SAS hard disk drives. *:Parameters of interest in this paper.

SAS SMART's Parameters	**Definition**
ECC Fast/Delay*	# ECC invocations per logical block
Reread/Rewrite	# Block access retries
Total Errors Corrected	ECC + rereads/rewrites
Correction Algorithm Invocations	Total retries + failed recovery attempts
Total Uncorrected Errors	# uncorrected data error
Power-on-Hours*	# Hours the disk has been ON
Start-Stop Cycles	A tally of spindle start/stop cycles
Load-Unload Cycles*	# of spin-ups and spin-downs
Gigabytes Processed*	Total Read/Written Bytes
Non-Medium Error Count	# errors due to cable/controller/firmware issues

loose cable connectivity (e.g., non-medium errors). Four of such parameters frequently change in healthy disks (i.e., ECC, PoH, load-unload cycles, and gigabytes processed), which are more of our interest in this paper.

III. MOTIVATION

A. *Existing Studies on HDD Behavior Analysis*

Existing studies have mainly focused on how SATA HDDs fail and also proposing mechanisms to predict their failures (Table III). For example, some studies analyze how to correlate disk failure to SMART parameters [21], [22], [35], [53] and some others apply machine learning algorithms to predict the remaining time till disk failure [6], [12], [30], [31], [33], [54]. Some researchers have taken a different path and analyzed the effect of environmental factors such as temperature and noise on the HDD performance [20], [32], [34], [40]. Overall, *no studies have been published on how disk aging affects the performance of enterprise HDDs.* Such knowledge is not only important for SLA guarantee in data centers, but also helping detect possibly fake and refurbished disks used for a couple of years, but (sold in some regions by some providers) as brand new disks.

B. *Aging Effect on HDD Performance*

We observe that a brand new enterprise disk and an enterprise disk used for a couple years may show significantly different behavior and performance (Fig. 1). We choose two HDDs from the same model, one disk just recently unboxed, and the other one used for 2.05 years. We run a workload with uniform random 4KB writes on both HDDs separately. *We observe that the brand new disk shows an average write latency of 4.55ms, while the used disk shows 41.37ms latency, exhibiting over 9× difference in latency.* Such huge performance difference may negatively affect the cloud SLA guarantees. *Considering no previous studies on the effect of aging on HDD performance, especially no studies on SAS enterprise HDDs (even for failure analysis), we motivate the need for an in-depth characterization of enterprise SAS HDDs to reveal how aging correlates with performance.*

IV. MORE ON RELATED WORK

Many studies have pursued characterization of disk drives and memory systems from either performance or reliability perspective. Some have analyzed emerging devices at the time, and some have analyzed well-established devices with focus on how they are affecting the enterprise environments in a couple of years of their deployment. In the following, we go over major existing studies in these domains.

Table III: Summary of existing studies on characterization of memory and storage devices compared to our work. Meaning of symbols: ✓ : True, O : Partially true, × : False

Categorization	**Reference**	**Disk type**	**Aging Effect on Performance**	**SMART**	**Workload Effect on Performance**	**Failure or Error-Rate Analysis**
CXL Memory Analysis	[60]	CXL Prototype	×	N/A	✓	×
Optane Persistent Memory Analysis	[26] [58] [55]	Optane DIMM	×	N/A	✓	×
Deployed SSDs Analysis	[36] [37]	SSD	×	✓	×	✓
Temperature/Noise Effect on HDDs	[40] [20] [34] [32]	SATA HDD	×	×	O	✓
HDD Failure Analysis	[21] [35] [53] [22]	SATA HDD	×	✓	×	✓
HDD Failure Prediction	[31] [12] [54] [30] [6] [33]	SATA HDD	×	✓	×	✓
HDD Aging Analysis	**Our work**	**SAS HDD**	✓	✓	O	O

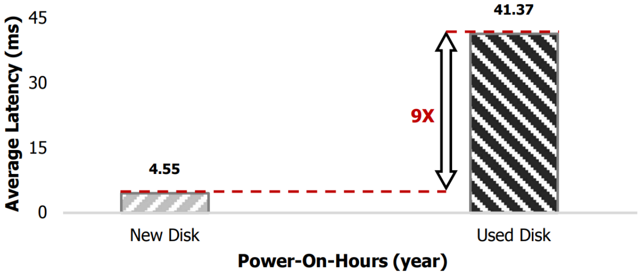

Figure 1: Average latency between new and used disks in the worst-case. New disk is just unboxed, and used disk is over 2-years old.

A. Analyzing HDDs

HDD Failure Analysis. Many previous studies have focused on analyzing disk failures, with most of them utilizing SMART reports from SATA disks [21], [25], [35], [41], [53]. Ref. [21] conducts a study on the SMART parameters of disks and demonstrates that some SMART data directly correlate with disk failures but are not sufficient alone to construct a predictive model for individual drive failures. Ref. [53] examines hardware errors over a four-year period. The research reveals that 80% of errors are related to HDD failures, with 29% of disk failures attributed to SMART failures. In RAIDShield [35], researchers examine six HDD models for up to 60-month period to assess the failure rates of each model. The study shows that disks with similar lifespans and higher reallocated sector count are more likely to fail.

HDD Failure Prediction. Predicting the failure of a disk or estimating the remaining life of a disk is another direction for existing work. Majority of work have focused on using machine learning algorithms for this purpose [6], [12], [30], [31], [33], [54]. Ref. [12] utilizes LSTM (Long Short-Term Memory) networks to indicate the remaining time to failure based on SMART data. Ref. [54] uses a method to analyze failure modes and proposes a novel anomaly prediction strategy based on Mahalanobis distance and SMART attributes, with a case study confirming its effectiveness.

Temperature and Noise Effect on HDDs. A couple of past research studies have examined the effect of temperature and noise on HDD performance [20], [32], [34], [40]. For example, Ref. [40] concentrates on the relationship between temperature, workload, and hard disk drive failures in large-scale data centers. It emphasizes the benefits of temperature optimization for improving hard disk drive reliability, modeling real-world data on temperature and failures, and underscores the importance of temperature-aware server design to enhance data center efficiency and reduce costs. Ref. [20] analyzes the frequency ranges and noise levels that negatively affect HDD performance. They conduct various tests across different noise ranges and use read/write speeds to measure HDD performance. They show that the most sensitive frequency range for HDDs is between 4 kHz and 10 kHz.

B. Characterizing SSDs and Memory Systems

SSD Analysis. A number of studies have characterized the performance and reliability of SSDs. For example, Grupp et. al [23] characterize a number of SSDs with different NAND chip technologies and show the SSD performance under different

I/O patterns, and how the SSD performance, along with its reliability (e.g., bit error rates) changes after applying many I/O operations. Some researchers have studied deployed SSDs in datacenters over a few year period. For example, Ref. [36] analyze SSDs deployed in enterprise storage systems over a two-year period with special focus on finding root causes that enterprise SSDs fail in data centers. Another study on deployed systems over a four-year period on NetApp data centers [37] focus on the amount of I/O operations (especially writes) on SSDs per year, how they affect SSD endurance, and revealing that using denser cells for SSDs is possible in production environments.Note that the findings in these studies are either in a completely different domain from our interest (e.g., about application I/O patterns) or only valid for SSDs (not HDDs), because HDDs' electro-mechanical operation is completely different from purely electronic operation of SSDs.

Analysis of Non-volatile Memories. A number of existing work have characterized recent non-volatile memory types (e.g., Optane DIMMs). Izraelevitz et. al [26] characterize the performance of Optane persistent memory under different operational modes for different workloads. Xiang et. al [58] also characterize Optane persistent memory but focus on performance behaviors of on-DIMM read/write buffering. Wang et. al [55] analyze Optane memory performance, and propose a simulation framework modeling Optane microarchitecture. Wei et. al [56] analyze the performance effect of combining *Remote Direct Memory Access* (RDMA) with optane memory compared to other memory types. Zeng el al [60] characterize a SAMSUNG *Compute Express Link* (CXL) memory prototype and show new insights on the latency and bandwidth of such devices under different access patterns.

C. Quality of Service (QoS) Guarantee

QoS guarantee is an important goal in storage systems, and many research papers have focused on this topic, yet *none* of them have analyzed HDD aging effects on performance. Most recent studies have analyzed the impact of internal SSD operations (e.g., garbage collection), and how they cause variations in performance and how to mitigate such problems [9], [10], [24], [27]–[29], [39], [48], [59], [61]. They propose internal SSD-aware I/O schedulers (e.g., in [27], [28]), new types of RAID architectures (e.g., in [48], [61]), or offloading complex SSD operations to operating system or application layer (as in open-channel SSDs or Zoned Namespaces SSDs [9], [10], [29], [59]). Some studies have shown how topology of connecting SSD enclosures may affect QoS (e.g., in [4]).

V. OUR METHODOLOGY FOR HDD CHARACTERIZATION

In this section, we first describe the hardware setup used for characterizations of enterprise HDDs. Next, we elaborate on six major categories of experiments on HDDs and the parameters that we analyze in this empirical study.

A. Hardware Setup for Experiments

We use a storage server with dual-socket Intel Xeon E5-2620 v4 CPUs, 32GB DRAM, and One 256GB Samsung 850 Pro SSD for the operating system (CentOS 7 with kernel updated to 5.4.52). We use over 50 enterprise SAS HDDs with 8TB capacity, all of which selected from the same major HDD vendor and same model (but with different amounts of disk usage). The HDDs are connected through a Broadcom MegaRAID 9361-8i card to motherboard (with RAID feature disabled for all individual HDD tests).

B. Overall Testing flow

For almost all of our experiments, we first target a specific parameter of the disk (e.g., the amount of terabytes of data previously written to disk). We run specific performance tests (e.g., random 4KB writes) on the each disk using FIO tool [7], and measure the latency in the current state. We then modify the target parameter and run the performance tests again to measure the effect. We repeat this procedure many times to ensure that we change a specific parameter enough to observe any possible performance changes. For example, we write over 500TB of data on a disk to compare its I/O latency to the case it had only 10TB written.

Isolation from External Factors. We isolate the external factors on our study to the extent possible, and carefully selected the sample of disks for our experiments. For example, *(a)* all the disks in our study are from the same data center with fairly similar temperature and cooling conditions. *(b)* We did not use the disks with peculiar behaviors (i.e., outliers) considering that they might have been physically damaged. *(c)* We ran each experiment 10 times on the same server and calculated the average performance to minimize the effect of any possible glitch or a sudden variation. *(d)* Multi-month experiment of one disk was on the same server, same disk slot, same system software and no major changes in the environmental factors (e.g., temperature or noise).

Performance Test Configuration.. For most performance tests, we run uniform random 4KB reads and uniform random 4KB writes on the disk, with number of I/O jobs set to 1, and I/O queue depth set to 16. We set the I/O rate to multiple fixed values from the minimum value to maximum value within the disk performance capability, and measure the I/O latency. Each performance test is run for 10 minutes, but For the sake of accuracy and minimizing anomalies, we repeat it 10 times. For selected number of tests, we also run sequential large 4MB reads/writes, but such tests do not show any new insight, thus we exclude their discussion in this paper.

C. Test Scenarios

Test Group #1: Effect of Power-On-Hours: The number of hours that the disk is ON (i.e., Power-On-Hours) is the first parameter we analyze. We evaluate this parameter in two ways.

a. Analysis of random write performance on over 30 HDDs: First, we use 30 HDDs with different PoH values (from PoH=one month to PoH=five years), and run performance tests at the maximum IOPS value. Each performance test lasts 10 minutes. We repeat it over 10 times and calculate the average latency. For trustable findings, we exclude disks with unusual behaviors and use 27 trustable disks. The unusual disk behavior may be due to partial physical damage during transportation from seller or movements between servers in data center during its usage.

b. Evaluating a disk from PoH=0 to PoH=8 months: We use a new disk (i.e., with PoH equals to zero) and run performance tests at its maximum IOPS. After every few weeks, repeat the performance test, and continue such periodic testing over the coarse of eight month. We record and analyze the measured latency during this eight-month period.

Test Group #2: Effect of Load-Unload Cycles: One of the seemingly important HDD parameters is load-unload cycles. Reducing the rotation rate of the disk and placing the disk head in a specific position (i.e., unload or spin-down) is effectively making the disk stand-by for power saving. To bring the disk in fully active mode, to handle I/O requests, or even due to some SMART reading requests, the load cycle (i.e., spin-up) should happen. In a real enterprise storage system, for disks not being assigned to an array, such state changing may frequently happen for disk SMART monitoring. If the storage system monitoring reads internal parameters of the disk every hour, in the five year recommended usage of the disk, around 44,000 load-unload cycles happen.

We select one of the SAS HDDs which has minimal load-unload cycles (near 1000). We then perform the stated reasonable number of spin-up and spin-downs, which may happen in the worst case scenario in a real storage system. For this reason, we run around 46,000 load-unloads, and after every 2,000 cycles, we run performance tests on the disk to measure the latency for specific IOPS values.

Test Group #3: Effect of Write Amount: We select *four SAS disks* with PoH ranging from almost zero to three-years old. We sequentially write on each disk up to 500 terabytes, and after every 50 terabytes, we measure the random write latency at different IOPS values. After executing such tests, and observing some changes in latency value at a specific IOPS, we select *14 disks with different PoH values* to run more fine-grained tests. In these tests, we write 50 terabytes on each of the 14 disks, and after every 10 terabyte, we measure the latency to detect the smallest variations with high accuracy.

Test Group #4: Effect of PoH and Write amount on Read Performance: We select four disks ranging from new to three years old (i.e., PoH=0, 0.5 year , 2 Years, 3 Years). We sequentially write on each disk up to 90 terabytes, and after every 10 terabytes, we measure the random read latency at different IOPS values.

Test Group #5: Effect of ECC: Every few writes on a disk may invoke error correction logic (ECC) and *one may suspect that when the disk has more error corrections, it has been used heavier and possibly suffers from lower performance*.

To evaluate such hypothesis, we run random write performance tests on 34 disks which have different PoH values (i.e., from almost zero to almost five-year old) and different amount of previously written data (from almost zero to few hundred terabytes). We measure the number of ECC invocations due to the random writes (i.e., the delta of change of ECC before and after the test). Next, we analyze whether any correlation exists between the disk aging and the number of ECC invocations.

Test Group #6: Overall Effect on Disk Array Performance: Disks are usually used in an array (known as RAID) for enterprise storage systems. Thus it is valuable to observe how the performance change of the disk affects the overall array performance. *Specifically, how does using the disks with different amount of usage (e.g., different PoH values) in the same array affect the performance?*

To answer the above question, we setup a RAID-10 array with three different configurations: (a) an array consisting of *4 new disks* (PoH is almost zero), (b) an array with *3 new disks and one used disk*, (c) an array with *4 used disks*. Next, we run random write tests on each array, and repeat the test 20 times for accuracy. Note that it is well-known that mirroring-based RAIDs (i.e., RAID-10) provide higher random I/O access performance than parity-based RAIDs (e.g., RAID-5/6) [5]. Thus, to evaluate the performance effect of disk array in cloud environments with mixture of applications (i.e., random I/O accesses), we select RAID-10 as our configuration.

VI. EXPERIMENTAL RESULTS

A. Effect of Disk Power-On-Hours

Finding 1*. The average random write latency of an enterprise HDD directly depends on the disk age (i.e., PoH): disk write latency significantly increases (an average of 1.8×) in the first eight months of operation, another 9% increase in latency when disk PoH value increases from 8-month to 2 years, and after 2-years, the disk has almost constant write latency (till becoming faulty).*

In an experiment over an eight-month period, we use a single enterprise HDD and monitor how its average random write latency increases as the PoH increases (Fig. 2). This experiment shows the disk has very low write latency of 10.08ms for a few weeks after it is unboxed (i.e., up to one month after insertion into the server for the first time). After three months, the disk write latency increases to 25.55ms (*2.5× worse*). When PoH reaches eight months, the disk write latency reaches the latency of 39.77ms (*3.9× worse*). *This experiment reveals the potential significant slowdown of a cloud/enterprise storage system even after eight months of normal usage* (or even being idle because PoH always increases).

In Fig. 3, we show more general view of how aging affects write latency by analyzing 27 enterprise HDDs with different ages (i.e., PoH values) ranging from one month to six years. We observe that *disks with less than two-month PoH have lowest write latency* of 12ms-24ms (average 21.25ms). Disks with over eight-months and less than two-years PoH experience much higher write latency (average 37.87ms). *Disks with over two years of age regardless of the exact age have similar average write latency* of 41.19ms. Note that we used the disks with at least one month of age in Fig 3 to reduce the significant variations in the write latency across disks when they are just unboxed. Thus latency values are slightly higher compared to what was seen in Fig 1 (i.e., over 10ms vs. 4.55ms) as brand new disk.

B. Effect of Load-Unload Cycles

Finding 2*. Conducting many spin-ups and spin-downs (i.e., load-unload) on enterprise HDDs, common in enterprise systems to minimize disk power consumption and possible lifetime increase, has almost no effect on disk performance.*

We choose an enterprise HDD with minimal previous usage amount and apply spin-ups and spin-downs heavily over 44,000 times, i.e., worst-case scenario of changing disk state every hour by the system monitoring in an enterprise system during the five-year expected lifetime of disk (Fig. 4). We observe that such load-unload operations which mainly involve some electro-mechanical operations have practically no effect on disk latency. Note that the maximum tolerable load-unload cycles specified by the enterprise HDD vendors is typically over 300,000. Thus even though we have heavily applied such operations, we have used less than 20% of disk expected capability, exhibiting a *definitive no effect on disk latency* for such operation during the disk lifetime.

C. Effect of Amount of Previous Disk Writes

Finding 3*. The amount of previous writes on disks older than six months has practically no effect on random write latency. For disks with less than one-month usage, the amount of writes cause some variations in latency at the maximum disk IOPS; however, such variations are NOT proportional to the amount of writes.*

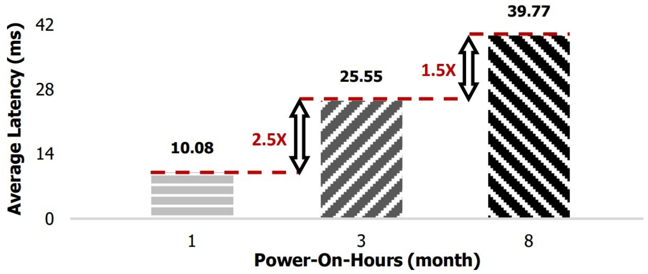

Figure 2: Monitoring average write latency of an enterprise HDD over eight months period [PoH=0 to eight months]

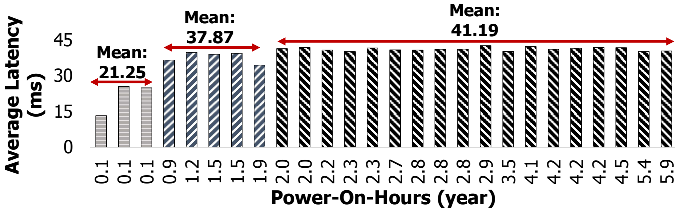

Figure 3: Observed trend of average write latency increase due to age difference in 27 enterprise HDDs

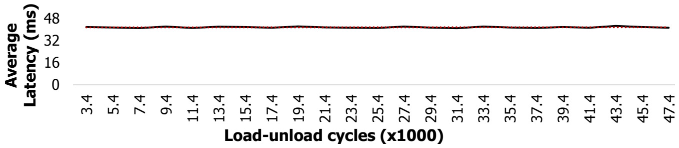

Figure 4: Showing no effect of Load-unload cycles on average latency of an enterprise HDD.

Fig. 5 shows the effect of sequential writes up to 500 terabytes on four disks with different PoH values. We observe that the latency for all IOPS values (i.e., 50 to 380 IOPS) practically does not change at all for disks older than six months. For a disk with very little usage (i.e., PoH value smaller than two months), the write latency is exactly same at most IOPS values. Only at highest disk IOPS (i.e., 380 IOPS) the write latency for this disk shows some variations.

To further analyze the variations, we select a couple of more disks and write 50TB on each, measuring the latency after each 10TB. For these tests, the focus is only on the maximum disk IOPS. We observe that similar to Fig. 5, the latency variations exist only for disks with little usage (Fig. 6). However, such latency variations are not monotonic. For example, after 10TB writes, the disk has write latency of 10.88ms (at maximum IOPS). Such latency gradually increases to 22.18ms after 30TB writes, then reduces to 10.81ms after 40TB, and slightly increase to 11.79ms. Such fluctuations in a disk with little usage (i.e., almost new disk) may be explained through bathtub curve. Bathtub curve is about high possibility of component (e.g., disk) failure or errors after a component is manufactured, and such instability exponentially decreases with usage or time till a stable state is reached. As Fig. 6 shows, disks with 1-year usage have much smaller variations, and disks with 2-years or more usage have no variations in latency. Overall, one might expect disk latency to increase if the amount of writes increases but our observations show such expectation is not valid (at least with decent write amounts). Possibly if the writes are significantly higher, and the disk becomes close to faulty state, latency increase would be obvious.

D. Effect of Various Parameters on Read Requests

Finding 4*. Typical aging (i.e., increase PoH) and the amount of writes to an enterprise HDD has no effect on read request performance.*

We have experimented with around 30 enterprise HDDs at different IOPS values within disk capability, however, for brevity, we have recorded and shown a sample of four disks with PoH value ranging from zero to 3 years old (Fig. 7). Our experiment reveals that the age of the disk, the amount of writes (shown up to 90TB) has absolutely no effect on random read latency. Although by increasing the I/O load (i.e., IOPS), the read latency naturally increases; such latency stays the same after writing around 100TB of data, and is also similar across disks with different ages (or PoH values).

E. Effect of Activating ECC Logic

Finding 5*. The number of ECC activations is fairly similar (with minor changes) across different enterprise HDDs with different ages; thus, an older disk with higher (worse) latency DOES NOT necessarily have higher rate of errors (or error corrections).*

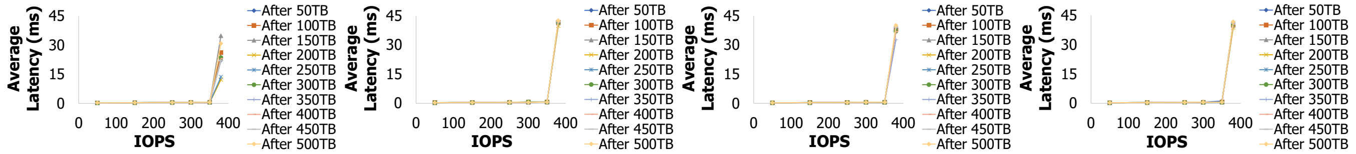

(a) New Disk (b) 6-Months Disk (c) 2-years Disk (d) 3-years Disk

Figure 5: Impact of sequential write on changes of average latency for four diks in different IOPS

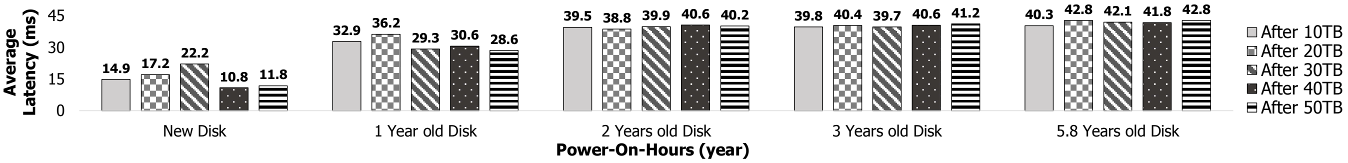

Figure 6: Correlation between sequential write and changes of average latency

(a) New Disk (b) 6-Months Disk (c) 2-years Disk (d) 3-years Disk

Figure 7: No visible effect of previous disk writes on random read access latency. Test runs on four disks with different ages and at various IOPS values.

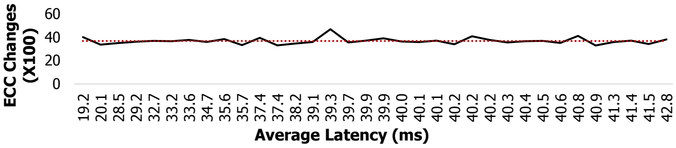

Figure 8: Correlation between ECC parameter and average latency

During each 10-minute performance test, we measure the ECC value change for over 30 enterprise HDDs with different ages, and observe no direct correlation between ECC activations and latency (Fig. 8). Regardless of disk age, and latency, we observe that the number of ECC activations for a 10-minute random write test on our enterprise HDDs is fairly constant (on average 4,000 times). Minor differences in ECC value exist in different HDDs. For example, a disk with minimal average latency of less than 20ms (which is typical of PoH less than 8 months old) show slightly even higher ECC activations than some disks with latency over 40ms, and vice versa. However, such minor ECC differences cannot be correlated to performance directly.

F. Overall Effect of Aging on Disk Array Performance

Finding 6*. A single HDD with 2-years of age may worsen the write latency of a whole array (with other disks being new) by 17.1%. This signifies the importance of careful selection of disks (based on age) in array creation in enterprise storage systems.*

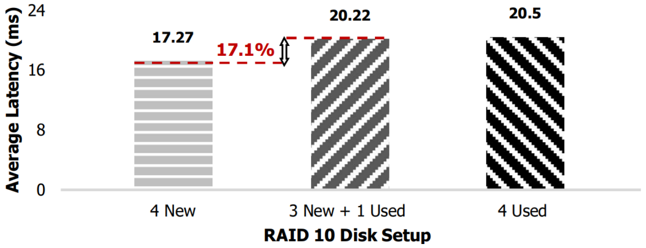

Figure 9: The change of average write latency due to disk-age in a RAID-10 array configuration

Our experiment on RAID-10 array with four disks reveals that with four new disks (i.e., PoH=0), the average write latency is 17.27ms. However, when a single 2-years old disk exists in the array, the latency increases by 17.1% to 20.22ms (Fig. 9). If all four disks are 2-years old, the latency is only 1.3% higher than the former case. This result emphasizes the need for storage system software (or even administrators) to carefully inspect the disk PoH when creating a RAID array to ensure maximum performance. Otherwise, by inserting a single old disk in an array, the latency for write-intensive workloads easily gets degraded by 17.1%.

VII. DISCUSSION

Scope of our findings. For fair comparison and accurate findings, we chose all HDDs to be from the exact same enterprise model and same vendor but with different amount of previous disk usage. Considering that enterprise SAS disks of the same class have exactly same technology and same features (with different capacity), we expect our findings to be applicable to the *same enterprise class from the same vendor with any capacity* (not just 8TB disks). Using other enterprise HDD models with different technology should still cause aging effects, but *how much* and *when* the increase on write latency happens could be different from our experimental results.

Implication of our findings on next-generation system designs. Our findings reveal that disk write latency is affected by aging. Therefore, the first and most important implication is the need to design *disk-age-aware write allocation policies*. File systems, RAID controllers, and any system software or hardware that handles data allocation and write scheduling may take into account our revealed aging profile of the disk, and based on the application QoS requirements, allocate data writes on proper disks. For example, data blocks of applications with strict latency requirements are allocated on younger disks.

VIII. CONCLUSION

We empirically studied over 30 enterprise HDDs and analyzed how different parameters related to aging affects disk performance. We revealed that only random write latency is affected by aging, and the highly effective parameter is disk being ON and working, *NOT* the amount of previous read/writes or spinup-spin-downs. We provided six findings following our experiments that guide system architects and system admins in data centers to better utilize disks to improve SLA guarantees.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their positive feedback. This research was generously funded by *High Performance Data Storage and Processing* (HPDS) Corp in Tehran, Iran.

REFERENCES

[1] Backblaze hard drive stat. https://www.backblaze.com/cloud-storage/resources/hard-drive-test-data, visited on July 21 2024.

[2] Blue brain 5 as an advanced supercomputer. https://www.cscs.ch/computers/blue-brain-5, visited on May 12 2024.

[3] The wide uses of hdds. https://www.seagate.com/blog/why-hdds-dominate-hyperscale-cloud-architecture/, July 6 2024.

[4] Mohammadamin Ajdari, Behrang Montazerzohour, Kimia Abdi, and Hossein Asadi. Empirical architectural analysis on performance scalability of petascale all-flash storage systems. *IEEE Computer Architecture Letters*, 23(2):158–161, 2024.

[5] Mohammadamin Ajdari, Patrick Raaf, Mostafa Kishani, Reza Salkhordeh, Hossein Asadi, and André Brinkmann. An enterprise-grade open-source data reduction architecture for all-flash storage systems. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 6(2):1–27, 2022.

[6] Preethi Anantharaman, Mu Qiao, and Divyesh Jadav. Large scale predictive analytics for hard disk remaining useful life estimation. In *IEEE Intl. Congress on Big Data (BigData Congress)*, pages 251–254, 2018.

[7] Jens Axboe. Flexible i/o tester (fio) github. https://github.com/axboe/fio, visited on September 18 2024.

[8] Backblaze. Backblaze drive stats: Hard drive reliability test data. https://www.backblaze.com/cloud-storage/resources/hard-drive-test-data, visited on October 20 2025.

[9] Matias Bjørling, Abutalib Aghayev, Hans Holmberg, Aravind Ramesh, Damien Le Moal, Gregory R Ganger, and George Amvrosiadis. ZNS: Avoiding the block interface tax for flash-based SSDs. In *USENIX Annual Technical Conference (ATC)*, pages 689–703. USENIX, 2021.

[10] Matias Bjørling, Javier Gonzalez, and Philippe Bonnet. LightNVM: The linux Open-Channel SSD subsystem. In *15th USENIX Conference on File and Storage Technologies (FAST)*, pages 359–374. USENIX, 2017.

[11] Simon Crisp. Toshiba mg10afa22te 22tb hdd review. https://www.kitguru.net/components/hard-drives/simon-crisp/toshiba-mg10afa22te-22tb-hdd-review/4/, visited on October 20 2025.

[12] Aniello De Santo, Antonio Galli, Michela Gravina, Vincenzo Moscato, and Giancarlo Sperlì. Deep learning for hdd health assessment: An application based on lstm. *IEEE Transactions on Computers*, 71(1):69–80, 2020.

[13] Western Digital. The longterm case for hdd storage: An analysis of the economic forces driving the future use of ssds and hdds in data centers. https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/collateral/white-paper/white-paper-the-long-term-case-for-hdd-storage.pdf, visited on October 20 2025.

[14] Western Digital. Ultrastar dc hc310 - 4tb. https://www.westerndigital.com/products/internal-drives/data-center-drives/ultrastar-dc-hc310-hdd?sku=0B36048, visited on October 20 2025.

[15] Western Digital. Ultrastar dc hc320 - 8tb. https://www.westerndigital.com/products/internal-drives/data-center-drives/ultrastar-dc-hc320-hdd?sku=0B36400, visited on October 20 2025.

[16] Western Digital. Ultrastar dc hc330 - 10tb. https://www.westerndigital.com/products/internal-drives/data-center-drives/ultrastar-dc-hc330-hdd?sku=0B42258, visited on October 20 2025.

[17] Western Digital. Ultrastar dc hc550 - 14tb. https://www.westerndigital.com/products/internal-drives/data-center-drives/ultrastar-dc-hc550-hdd?sku=0F38528, visited on October 20 2025.

[18] Western Digital. Ultrastar dc hc590 data center hard drive - 24tb. https://www.westerndigital.com/products/internal-drives/data-center-drives/ultrastar-dc-hc590-hdd?sku=0F59373, visited on October 20 2025.

[19] Western Digital. Ultrastar dc hc690 data center hard drive - 30tb. https://www.westerndigital.com/products/internal-drives/data-center-drives/ultrastar-dc-hc690-hdd?sku=WSH723200AL4201, visited on October 20 2025.

[20] Trinoy Dutta and Andrew R Barnard. Performance of hard disk drives in high noise environments. *Noise Control Engineering Journal*, 65(5):386–395, 2017.

[21] Pinheiro Eduardo. Failure trends in a large disk drive population. In *USENIX Conference on File and Storage Technologies. San Jose, USA*, 2007.

[22] Jon G. Elerath and Michael G. Pecht. Enhanced reliability modeling of RAID storage systems. In *The 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks DSN, 25-28 June, Edinburgh, UK*, pages 175–184, 2007.

[23] Laura M Grupp, Adrian M Caulfield, Joel Coburn, Steven Swanson, Eitan Yaakobi, Paul H Siegel, and Jack K Wolf. Characterizing flash memory: Anomalies, observations, and applications. In *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 24–33, 2009.

[24] Jian Huang, Anirudh Badam, Laura Caulfield, Suman Nath, Sudipta Sengupta, Bikash Sharma, and Moinuddin K Qureshi. {FlashBlox}: Achieving both performance isolation and uniform lifetime for virtualized {SSDs}. In *15th USENIX Conference on File and Storage Technologies (FAST 17)*, pages 375–390, 2017.

[25] Song Huang, Shuwen Liang, Song Fu, Weisong Shi, Devesh Tiwari, and Hsing-bung Chen. Characterizing disk health degradation and proactively protecting against disk failures for reliable storage systems. In *IEEE Intl. Conference on Autonomic Computing (ICAC)*, pages 157–166, 2019.

[26] Joseph Izraelevitz, Jian Yang, Lu Zhang, Juno Kim, Xiao Liu, Amirsaman Memaripour, Yun Joon Soh, Zixuan Wang, Yi Xu, Subramanya R Dulloor, Jishen Zhao, and Steven Swanson. Basic performance measurements of the intel optane dc persistent memory module. *arXiv preprint arXiv:1903.05714*, 2019.

[27] Jihun Kim, Joonsung Kim, Pyeongsu Park, Jong Kim, and Jangwoo Kim. Ssd performance modeling using bottleneck analysis. *IEEE Computer Architecture Letters*, 17(1):80–83, 2017.

[28] Joonsung Kim, Pyeongsu Park, Jaehyung Ahn, Jihun Kim, Jong Kim, and Jangwoo Kim. Ssdcheck: Timely and accurate prediction of irregular behaviors in black-box ssds. In *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 455–468. IEEE, 2018.

[29] Huaicheng Li, Martin L Putra, Ronald Shi, Xing Lin, Gregory R Ganger, and Haryadi S Gunawi. Ioda: A host/device co-design for strong predictability contract on modern flash storage. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles (SOSP)*, pages 263–279. ACM, 2021.

[30] Jing Li, Xinpu Ji, Yuhan Jia, Bingpeng Zhu, Gang Wang, Zhongwei Li, and Xiaoguang Liu. Hard drive failure prediction using classification and regression trees. In *44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN, Atlanta, GA, USA, June 23-26*, pages 383–394. IEEE Computer Society, 2014.

[31] Jing Li, Rebecca J Stones, Gang Wang, Xiaoguang Liu, Zhongwei Li, and Ming Xu. Hard drive failure prediction using decision trees. *Reliability Engineering & System Safety*, 164:55–65, 2017.

[32] Yi-Ju Liao and Jen-Yuan Chang. Identification of factors affecting disk drive's performance in data server by use of decision tree learning method. 06 2021.

[33] Fernando Dione S Lima, Francisco Lucas F Pereira, Lucas GM Leite, Joao Paulo P Gomes, and Javam C Machado. Remaining useful life estimation of hard disk drives based on deep neural networks. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2018.

[34] Nan Liu and David B Bogy. Temperature effect on a hdd slider's flying performance at steady state. *Tribology letters*, 35:105–112, 2009.

[35] Ao Ma, Rachel Traylor, Fred Douglis, Mark Chamness, Guanlin Lu, Darren Sawyer, Surendar Chandra, and Windsor Hsu. Raidshield: characterizing, monitoring, and proactively protecting against disk failures. *ACM Transactions on Storage (TOS)*, 11(4):1–28, 2015.

[36] Stathis Maneas, Kaveh Mahdaviani, Tim Emami, and Bianca Schroeder. A study of {SSD} reliability in large scale enterprise storage deployments. In *18th USENIX Conference on File and Storage Technologies (FAST 20)*, pages 137–149, 2020.

[37] Stathis Maneas, Kaveh Mahdaviani, Tim Emami, and Bianca Schroeder. Operational characteristics of ssds in enterprise storage systems: A large-scale field study. In *20th USENIX Conference on File and Storage Technologies (FAST 22)*, pages 165–180, 2022.

[38] NewEgg Online Shop. Seagate exos x18 st10000nm013g 10 tb hard drive - 3.5 inch internal - sas (12gb/s sas) - conventional magnetic recording (cmr) method - video surveillance system, storage system device supported - 72. https://www.newegg.com/p/1Z4-002P-024F1?srsltid=AfmBOoqNH3KhBUHYZ054su8qBi6jEHUtqts76xWDAhGGnVnkEMjnOJnD, visited on October 20 2025.

[39] Reza Salkhordeh, Kevin Kremer, Lars Nagel, Dennis Maisenbacher, Hans Holmberg, Matias Bjørling, and André Brinkmann. Constant time garbage collection in ssds. In *2021 IEEE International Conference on Networking, Architecture and Storage (NAS)*, pages 1–9. IEEE, 2021.

[40] Sriram Sankar, Mark Shaw, Kushagra Vaid, and Sudhanva Gurumurthi. Datacenter scale evaluation of the impact of temperature on hard disk drive failures. *ACM Transactions on Storage (TOS)*, 9(2):1–24, 2013.

[41] Bianca Schroeder and Garth A. Gibson. Understanding disk failure rates: What does an MTTF of 1, 000, 000 hours mean to you? *ACM Trans. Storage*, 3(3):8:1–8:31, 2007.

[42] Seagate. Exos 7e10 datasheet. https://www.seagate.com/files/www-content/datasheets/pdfs/exos-7e10-DS1957-6M-2104GB-en_IN.pdf, visited on October 20 2025.

[43] Seagate. Exos 7e8 datasheet. https://www.seagate.com/www-content/datasheets/pdfs/exos-7-e8-data-sheet-DS1957-1-1709US-en_US.pdf, visited on October 20 2025.

[44] Seagate. Exos x18 datasheet. https://www.seagate.com/files/www-content/datasheets/pdfs/exos-x18-channel-DS2045-1-2007GB-en_SG.pdf, visited on October 20 2025.

[45] Seagate. Exos x20 datasheet. https://www.seagate.com/files/www-content/datasheets/pdfs/exos-x20-channel-DS2080-2111GB-en_EM.pdf, visited on October 20 2025.

[46] Seagate. Exos x22 datasheet. https://www.seagate.com/content/dam/seagate/en/content-fragments/products/datasheets/exos-x22/exos-x22-DS2080-2209US-en_US.pdf, visited on October 20 2025.

[47] Seagate. Exos x24 datasheet. https://www.seagate.com/content/dam/seagate/en/content-fragments/products/datasheets/exos-x24/exos-x24-DS2080-2307US-en_US.pdf, visited on October 20 2025.

[48] Mingyao Shen, Suyash Mahar, Heewoo Kim, Joseph Izraelevitz, and Steven Swanson. Autossd: Cxl-enhanced autonomous ssds for low tail latency. In *Proceedings of the 34th International Symposium on High-Performance Parallel and Distributed Computing*, pages 1–13, 2025.

[49] NewEgg Online Shop. Samsung pm1653 7.68tb sas 24gbps 2.5in enterprise internal ssd mz-ilg7t60 new. https://www.newegg.com/p/1WK-002G-003J4?srsltid=AfmBOoq_c_odsVGoPcXOyKPh1BwI31LPjWHRl2SS6NofcQ7jZpvwdJ92, visited on October 20 2025.

[50] Lyle Smith. Western digital ultrastar dc hc590 review. https://www.storagereview.com/review/western-digital-ultrastar-dc-hc590-review, visited on October 20 2025.

[51] Backblaze Drive Stats Team. Backblaze drive stats for q2 2025 | hard drive failure rates. https://www.backblaze.com/blog/backblaze-drive-stats-for-q2-2025/, visited on October 20 2025.

[52] Toshiba. Mg series enterprise capacity hard drives. https://storage.toshiba.com/docs/enterprise-hdd-documents/HDD-MGEND.pdf, visited on October 20 2025.

[53] Guosai Wang, Lifei Zhang, and Wei Xu. What can we learn from four years of data center hardware failures? In *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 25–36. IEEE, 2017.

[54] Yu Wang, Qiang Miao, and Michael Pecht. Health monitoring of hard disk drive based on mahalanobis distance. In *2011 Prognostics and System Health Managment Confernece*, pages 1–8. IEEE, 2011.

[55] Zixuan Wang, Xiao Liu, Jian Yang, Theodore Michailidis, Steven Swanson, and Jishen Zhao. Characterizing and modeling non-volatile memory systems. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 496–508. IEEE, 2020.

[56] Xingda Wei, Xiating Xie, Rong Chen, Haibo Chen, and Binyu Zang. Characterizing and optimizing remote persistent memory with {RDMA} and {NVM}. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, pages 523–536, 2021.

[57] Adam Wright. Worldwide global storagesphere forecast, 2024-2028: Ai everywhere, but storage capacity remains a balancing act. https://my.idc.com/getdoc.jsp?containerId=US52312824&pageType=PRINTFRIENDLY, visited on October 20 2025.

[58] Lingfeng Xiang, Xingsheng Zhao, Jia Rao, Song Jiang, and Hong Jiang. Characterizing the performance of intel optane persistent memory: A close look at its on-dimm buffering. In *Proceedings of the Seventeenth European Conference on Computer Systems*, pages 488–505, 2022.

[59] Shushu Yi, Shaocong Sun, Li Peng, Yingbo Sun, Ming-Chang Yang, Zhichao Cao, Qiao Li, Myoungsoo Jung, Ke Zhou, and Jie Zhang. Biza: Design of self-governing block-interface zns afa for endurance and performance. In *Proceedings of the ACM SIGOPS 30th Symposium on Operating Systems Principles*, pages 313–329, 2024.

[60] Jianping Zeng, Shuyi Pei, Da Zhang, Yuchen Zhou, Amir Beygi, Xuebin Yao, Ramdas Kachare, Tong Zhang, Zongwang Li, Marie Nguyen, et al. Performance characterizations and usage guidelines of samsung cxl memory module hybrid prototype. *arXiv preprint arXiv:2503.22017*, 2025.

[61] Guangyan Zhang, Zican Huang, Xiaosong Ma, Songlin Yang, Zhufan Wang, and Weimin Zheng. {RAID+}: Deterministic and balanced data distribution for large disk enclosures. In *16th USENIX Conference on File and Storage Technologies (FAST 18)*, pages 279–294, 2018.