Hard Disk Drive Failure Analysis and Prediction: An Industry View

Zach Miller, Olusiji Medaiyese, Madhavan Ravi, Alex Beatty, Fred Lin
Meta Platform Inc.
zachmiller, siji, madravi, abeatty, fanlin {@meta.com}

***Abstract*—Storage media devices are fundamental to Meta's hardware infrastructure, which supports a diverse family of applications such as Facebook, Instagram, and WhatsApp. Understanding the factors that impact the reliability of storage devices is important for setting application expectations on specifications such as throughput, latency, and read/write success rate. Improving hardware reliability helps us meet those expectations.**

In this paper, we examine the impact that age and workload have on the annualized failure rate (AFR) of Hard Disk Drives (HDDs), one of the most used types of storage devices for Meta's applications. We analyze the correlation based on data collected from our production hardware fleet. In our datacenter environment, we observe that HDD AFR increases as either age or lifetime cumulative workload increases. We discuss the difference between the AFR curves and the projections that manufacturers make using statistical modeling. Additionally, we use a decision tree-based predictive machine learning (ML) model, XGBoost, for analyzing the correlation between the SMART (Self-Monitoring, Analysis, and Reporting Technology) metrics and the health of HDDs. Through this study, we observe that age and workload-related SMART parameters are most correlated to the health of a drive based on the trained ML model. More so, we identify that the difference of SMART metrics over a 30-day time window could improve the prediction performance for HDD failures.

***Index Terms*—hard disk drive, reliability, machine learning, annualized failure rate, workload, SMART metric, storage**

I. Introduction

Meta deploys large-scale distributed storage services across datacenters. Storage applications are often categorized based on the type and temperature of the data stored: hot, warm, and cold data. At Meta, we have an exabyte-scale distributed file system, known as Tectonic [1]. Tectonic has tenants that include a warm Binary Large Object (BLOB) storage tier and a data warehouse tier. The warm BLOB tier is used for external media storage (photos, videos, documents), and internal application data (traces, heap dumps, logs) [2]. The data warehouse tier is designed to store data analytics for business intelligence, and objects such as massive map-reduce tables, snapshots of the social graph, and AI training data and models. Both tiers run on specialized storage servers containing Hard Disk Drives (HDDs), also known as Just-a-Bunch-of-Drives (JBODs) [3]. In industry, HDDs are widely used as either a boot device or a data device. The HDDs discussed in this paper are used as data devices. In our infrastructure, one compute module facilitates concurrent I/O across all HDDs in a JBOD.

If an HDD fails and our tooling determines that the correct remediation is to replace that HDD, an online repair process is followed. For online repairs, a hot-swap is issued to replace the failed HDD, all without stopping I/O to the other healthy drives in the server. Once the failure is validated, a datacenter operator hot-swaps the drive with a new HDD.

To account for drive failures, encoding schema and data replication are used in Meta's infrastructure. Tectonic relies on a combination of data replication and erasure encoding, such as Reed-Solomon (RS), to reconstruct data [4]. Servers are physically distributed across datacenters to minimize the localization of data in the same failure domain, e.g. HDDs in the same RS encoding group that may be in the same server, rack, or datacenter. This ensures continuous data availability to applications when HDD failures occur. In the case of a perfect storm of events that result in HDD capacity loss beyond the reconstruction thresholds of the RS encoding, data would become inaccessible or ultimately lost.

HDDs are a unique type of component in datacenters, as they are composed of both mechanical and electrical elements [5]. Drive heads, storage media, electrical components, and mechanical components can all lead to HDD failures [6]. Datacenter environments can influence HDD failures through agitators such as thermal interactions, rotational/acoustic vibrations (RV/AV), and device handling [6], [7].

Previous studies investigated factors that influence HDD failures [8]–[10]. In one study, HDD data sets were used to analyze the impact of workload, age, and temperature on HDD failures [8]. The result was inconclusive regarding the correlation between the factors and failures. The authors also analyzed self-monitoring, analysis, and reporting technology (SMART) metrics to determine the probability of HDD failures; they established that using only SMART metrics is limited in prediction capabilities.

The scale we operate at makes repairs challenging to manage, and the variety of failure types triggered by different services makes failure predictions challenging. Our storage architecture is complex, and reliability requirements can vary across applications. Each application has its own tolerance for latency, internal constraints for failing HDDs, complex workloads, and varying environmental conditions.

Understanding the conditions that HDDs are exposed to in our fleet helps us optimize our HDD deployments and remediate failures. Hence, the motivation for this paper is to provide insights from Meta's industry perspective to supplement prior research explored in this space. The contributions of this work are summarized as follows:

- We present two correlation studies to evaluate the effects of HDD age on AFR and HDD cumulative workload on AFR. The studies leverage a large production data set sampled from Tectonic's storage fleet.
- We use a machine learning model (XGBoost algorithm) to carry out a feature analysis. Through the modeling, we identify the most impactful SMART metrics that are indicative of a drive's health. Our ML model setup helps us establish that the variation in SMART metric readings over a 30 day window may indicate potential failure in HDDs.

The remainder of this paper is organized as follows: in Section II we present the storage hardware and telemetry setup within our datacenters. In Section III, we present our case studies on the impact of age and workload on AFR. In Section IV, we present a study where an ML model is used for both SMART metric based feature analysis and failure analysis. We conclude the paper and scope future work in Section V.

II. Background

A. HDDs in a Large-Scale Production Environment

The HDDs sampled in this analysis are classified as nearline HDDs: high-capacity, high reliability drives designed for cloud and datacenter applications. The studies sample HDDs used as data devices across Tectonic and its tenants. We deploy HDDs from different vendors, models, and vintages in our production fleet. In this paper we provide insight based on a sampled subset of the HDD population.

B. Hardware Health Telemetry within Data Centers

SMART is a set of industry standard health metrics for HDDs that can be used to monitor and indicate drive health [7], [11]. In this study, we use the daily collected SMART data to determine the power-on-hours and cumulative workload for each HDD. We improve the accuracy of power-on-hours by removing the time the drives were powered on during the manufacturing process, but were not actually running in Meta's environment. When we perform the SMART metric-based feature analysis, SMART attributes are captured for studying their potential influence and weight on failures.

In our storage servers, we use micro-services to collect and log health data from HDDs in our fleet, including HDD SMART data. These services pull the respective health data and upload it into the data warehouse tier as unstructured data. The logs are parsed through a data pipeline into a table and presented as structured data for analysis.

C. HDD Failure Mechanisms

HDDs have different failure modes resulting from complex design mechanisms. The failures can be categorized into one of the following groups: HDD connection loss, operational timeout, or I/O operational failure. The failures can be attributed to read/write head degradation, contamination, media defects, electrical/mechanical elemental failures, and interface/connectivity loss. These types of failures can all lead to hard failures within drives [6], [12], [13].

In our fleet, we need to detect and remediate HDD failures. We run a health monitoring micro-service, Machinechecker, on a regular cadence to detect hardware failures [14]. Machinechecker invokes our repair process to initiate remediations when it detects unhealthy HDDs.

D. Statistical Modeling of Hard Disk Drive Failures

In order to maintain proper balance between application performance and stability, we rely on the HDD vendors to create robust, reliable storage devices. The most important part of the process is progressively scaling up their manufacturing as they pass specific reliability thresholds. To achieve this goal, vendors rely on a theoretical model that assesses how AFR changes over time, better known as the bathtub curve [15]. HDD manufacturers have different types of accelerated tests to determine potential failure modes at each stage of field usage. The beginning of the bathtub curve balances the effectiveness of manufacturing screens compared to their cost impact. The middle of the bathtub curve expects consistent, stable AFR. The end of the bathtub curve predicts a sharply increasing failure rate as the device wears out to the end of its deployed service life. Not all HDDs are manufactured equivalently, AFR during deployment will vary as vendors implement continuous improvements. Additionally, each HDD model has a different margin available in its subsystems.

Once the vendor and Meta successfully qualify an HDD product for Meta's storage applications, there is another important theoretical model that must be considered: the derating curve [16]. The purpose of a derating curve is to assess how given stressors to a component affect the failure rate of a product over time. For HDDs, the derating curves factor in temperature and workload as stressors. The HDD vendor creates boundary conditions for safe usage of the product and publishes these specifications in their product manual.

E. Dimensions that affect Hard Disk Drive Failures

There are many factors that can influence the reliability of HDDs. We identify and discuss four contributing factors that affect HDD reliability: age, workload, temperature, and interference from vibrations.

Age of HDDs and their associated failures can be statistically modeled with bathtub curves and derating curves [17], [16]. These projections are designed to estimate failure rates over time, however, the projections may not match the actual trends observed in our large-scale production environment. For instance, in a previous industry study, failure projections from statistical models were shown to be inaccurate when compared against field data [17]. Additional studies also suggest that age is a factor in HDD failure rates [18], [9].

Workload is another factor that can stress HDDs. Workload recommendations are defined in each HDD vendor specification. These are representative of the maximum workload that the vendor tests HDDs to in their validation processes before mass production [19]–[21]. How workload is defined varies from vendor to vendor. Additionally, vendor models assume a consistent mix of workload over the life of a drive.

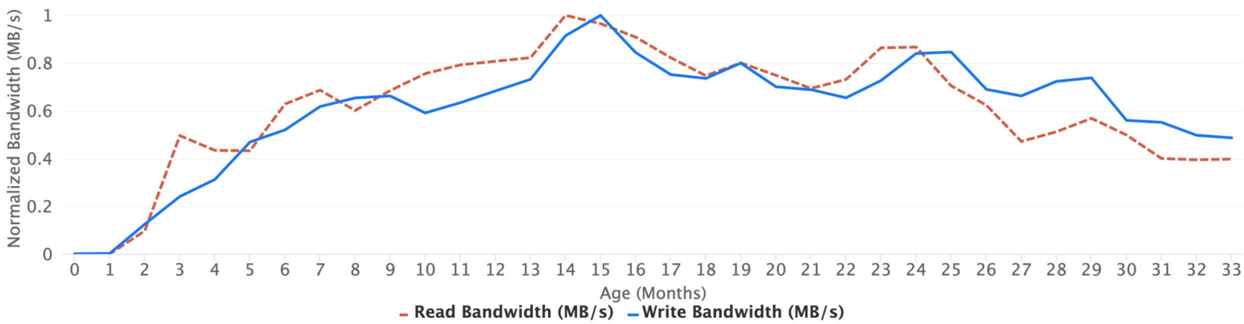

Fig. 1. Read and write bandwidth for a single HDD over the course of its production usage.

In our applications, drives experience different workloads over their lifetime. For example, when we first provision and induct racks into a service, drives must be filled with data. This results in a heavier write to read workload ratio. Applications may also change encoding schema or BLOB sizes, which can impact the stress on the HDD. Fig. 1 shows how workload could vary across the lifetime of an HDD in production. The data is normalized for both read and write workloads to remove manufacturer specific information.

Temperature is an agitator to HDDs, and a previous study showed HDD AFR can double when a drive operates at a temperature of 55C rather than 40C [22]. Operational temperatures are taken into account when designing for power and cooling in the server hardware to optimize HDD reliability. Our servers use a Baseboard Management Controller (BMC) to monitor the temperature across all components. The BMC controls fan speeds, adjusting the airflow within a server's chassis to provide optimal operational temperatures.

Rotational and Acoustic Vibrations can disrupt HDD operations, leading to HDD failures. HDD servo mechanisms provide a feedback loop to correctly position their read/write heads [23]. The servo design compensates for mechanical interference that can arise from the vibrations of fans and other HDDs in the system [18]. Our server hardware design must take this into account to dampen interference.

Temperature and RV/AV interference are studied when we design server hardware. HDD age and workload are opportunistic candidates to examine in this paper, as they are both factors that cannot easily be accounted for when designing servers. In Section III. B, we present how HDD AFR trends could be different based on the data observed in our production environment and the statistical models provided by HDD manufacturers. In Section III. C we present the workload distributions between healthy and unhealthy HDDs.

F. HDD Annualized Failure Rates

Previous studies show that AFR trends can vary across the industry. In one external study, AFR ranged from 1.7% for HDDs in year one of operation, 8.6% AFR in year three, and approximately 7% AFR in year five [8]. Another study, which modeled failures by hazard rates, observed that HDDs do not follow the traditional bathtub curves, rather an increasing-decreasing-stabilized pattern that occurs in early deployments [9]. These studies show variation across industry observed HDD failure rate modeling and failure rate field behaviors.

III. Impact of Age and Workload on HDD AFR

A. Data Setup

To determine the impact of age on AFR, we sample data from three HDD models in our fleet. Each HDD model has been in our production fleet for an extended period of time. The sampled data is a subset of the entire install base for that HDD model. To represent the impact of AFR on age of the HDD, we present our data as a normalized AFR.

To determine the impact of workload on AFR, we sample a subset of all HDD models in the fleet from the last one year. We characterize the HDDs as either healthy or unhealthy. Healthy HDDs are those in production at the time of sampling this data. Unhealthy HDDs are those that have been hot-swapped. The cumulative workload collected from unhealthy drives is collected at the time of the HDD failure.

B. Impact of Age on Annualized Failure Rate

In the first study, we aim to understand the trends that the impact of age has on HDD AFR in our production fleet. We examine the relationship between age and AFR for three drive models, referred to as M_1, M_2, and M_3. It should be noted that we sample data from subsets of the total install base for each of the three HDDs across our storage infrastructure. The exact time in production is removed to prevent any vendor identifiable information. Normalized AFR is presented to remove manufacturer specific information. The data is normalized by scaling AFR to a range, as defined in the following equation:

$$\mathrm{A}FR_{\text{normalized}} = \frac{AFR_{\text{current}} - AFR_{\text{min}}}{AFR_{\text{max}} - AFR_{\text{min}}}, \tag{1}$$

AFR_{current}, AFR_{min} and AFF_{max} represent the true AFR values, the minimum AFR value in the data set, and the maximum AFR value in the data set.

Fig. 2 shows the AFR trend for M_1. From the data, we see an increase in normalized AFR from its initial deployment at an age of 0 months to an age of 10 months. After 10 months, we observe that the normalized AFR stabilizes and remains constant. This continues until we hit an inflection point at an age of 24 months. After this age, we observe continued AFR growth. Fig. 3 presents the AFR trend for M_2, which observes a similar non-linear AFR trend to that in M_1.

Fig. 4 shows the AFR trend for M_3. We observe that different HDD models may or may not exhibit different trends across the fleet. The trend could be similar or could vary depending on the HDD model. In M_3 the major difference is that the AFR continues to trend upwards, and does not reach a steady-state AFR by the age of 24 months.

From the analysis across the three HDD models, we observe that there is a non-linear regression as each HDD ages. We observer unique AFR trends within Meta's environments and applications, and each HDD model may or may not follow similar AFR behavior. These trends are not reflective of the traditional statistical models across industry. Statistical models project that HDD failure rates increase in early-life, stabilize to constant failure rates, and increase again at the end of life. The data presented in our study does not suggest this behavior happens across all models or vintages. The following differences are observed when our data is compared against statistical models in industry:

- Based on the typical bathtub curve modeling [15], we would expect to see higher early life failure rates at the beginning of the hardware lifetime, followed by a steady-state where the failure rate remains relatively constant, followed by an increase in failure rate towards the end of the hardware lifetime.
- In our environment, however, we expect that most of the early life failures should have been screened and removed by the manufacturers. So we do not see the high failure rate at the beginning of the hardware lifetime. In addition to that, our failure curves start with an increase then reach the steady-state, which is also different from the typical bathtub curve for which we would expect to see the steady-state failure rate right after the initial high failure.
- After the steady-state failure period, we observe additional AFR growth. This AFR growth is before the portion of the HDD life cycle where we expect to hit the end of life, and at these ages, we would expect the AFR to remain constant rather than increase.

C. Impact of Workload on Annualized Failure Rate

In this study, we compare the cumulative workload distributions from healthy and unhealthy HDDs. The data is collected by randomly sampling a subset of healthy and unhealthy HDDs in the production fleet from the last one year. To remove manufacturer specific information, we normalize the workloads. This preserves the workload distributions between healthy and unhealthy HDDs.

Fig. 5 presents the data for our workload distribution across healthy and unhealthy drives in the fleet. The boxes show the P75 and P25 workloads, while the middle lines show the median workloads. The lines above and below the boxes show the minimum workloads (workloads below P25) and maximum

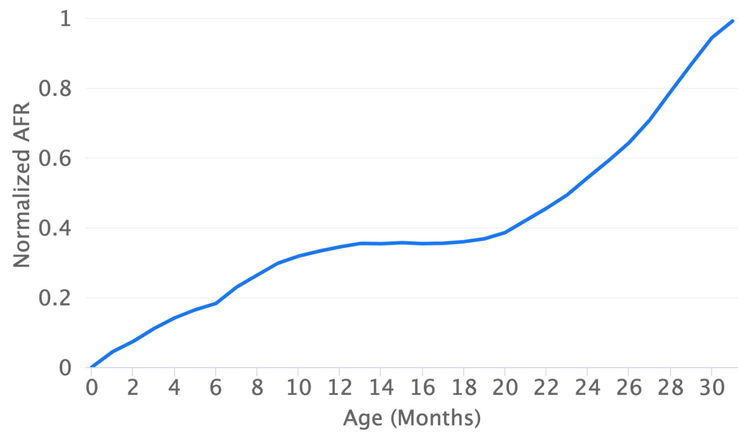

Fig. 2. M_1's annualized failure rate (AFR) over age.

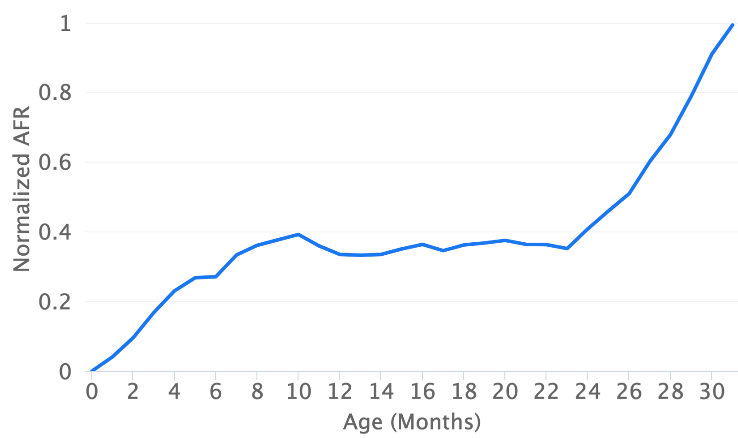

Fig. 3. M_2's annualized failure rate (AFR) over age.

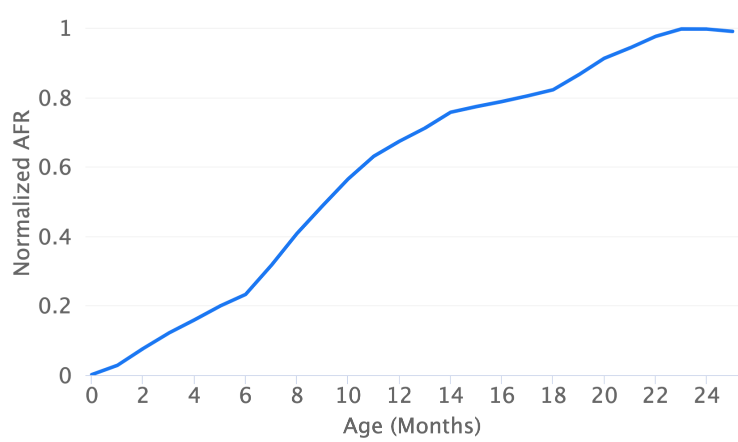

Fig. 4. M_3's annualized failure rate (AFR) over age.

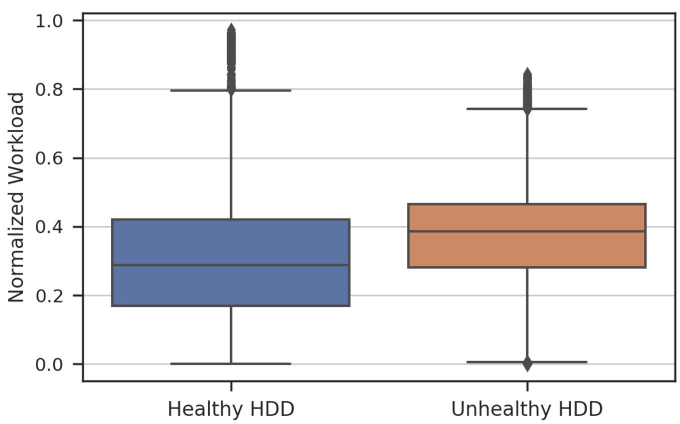

Fig. 5. Distribution of workloads between healthy and unhealthy HDDs.

workloads (workloads above P75). The markers above and below those lines represent the workload outliers.

The first observation is that the median workload for unhealthy HDDs is 1.5x greater than the median workload for healthy HDDs. The second observation is that amongst healthy and unhealthy drives, there is an overlap across the interquartile ranges of workload. For our data, a statistical test (z-test) confirmed a p-value less than 0.05. This provides evidence to reject the null hypothesis. Hence, the mean workloads between healthy and unhealthy drives are statistically significant.

IV. Feature study using a machine learning model

Our next step is to examine how SMART metrics can assist in understanding the underlying trends or contributing factors to HDD failures within our fleet. To achieve this, we use the XGBoost algorithm to predict these failures, and we formulate the problem as a binary classification task. In other words, we train a classifier to identify which drives are healthy or unhealthy based on their SMART metrics. After the classifier is trained, we analyze how each feature affects the performance of the classifier.

Our data for health metrics, repairs, and failures of the HDD-based systems are retained for multiple years. For this analysis, we use the daily collected data from over 53,000 failed HDDs from our Tectonic storage fleet. We use the SMART readings from one day and 30 days before the drives failed and categorize them as unhealthy drives. Additionally, we selected 54,000 healthy drives randomly from our hardware fleet as healthy drives. By combining the failed drives and the 54,000 healthy drives that were randomly sampled, our sample size for this study amounts to over 107,000 drives.

A. Training Setup

An XGBoost classifier is used as the modeling algorithm to classify unhealthy and healthy drives. We consider three different sets of features and each set of features is referred to as a configuration. Each configuration is used to train a separate model. The three configurations of features used are:

- Model A: SMART readings of drives taken one day before a failure event.
- Model B: SMART readings of drives taken 30 days before a failure event.
- Model C: The difference/delta between the SMART readings of a drive taken one day and 30 days before a failure.

B. Prediction Results

The evaluation of each model is done using two categories of test sets that include:

- Test data I: Test data derived from a 20% random selection of all the drives (i.e., over 107,000 drives) in the data set. We split the data into a training set and a test set. 80% of the drives were used for training the model, while the 20% subset was held out for evaluating the model's performance on unseen data. This approach helps to prevent overfitting of the model to the training data and provides a more accurate estimate of its generalization performance on new data. It is worth noting that this test data is taken from the same time window as the training set. This means that this test data is taken from the same period of time as the data used to train the model. This approach ensures that the test data is similar in nature to the training data, allowing for a fair evaluation of the model's ability to generalize to new data.
- Test data II: Test data based on all the unhealthy and healthy drives in production, but from a different time window than the training set. The purpose of this test data is to assess the model's ability to generalize to a live production environment. By testing the model on a different time window and in a real-world setting, we can verify that the model's performance is consistent and reliable across different contexts. This is important because the model should be able to accurately predict the health of drives in production, regardless of the time period or other factors that may impact the data. Therefore, the use of this test data is critical in assessing the practical application of the model in a real-world scenario.

Table I shows the models' performance metric using precision and recall across two test datasets from the different time windows. Precision is the ratio of the number of true positive predictions to the total number of positive predictions, where a true positive prediction is when the model correctly predicts the positive class. On the other hand, recall is the ratio of the number of true positive predictions divided by the total number of actual positive instances. We define the positive as the unhealthy drives and the negative as the healthy drives.

From the result shown in table I, the performance of our ML models is limited across the three setups, especially when using test data that is sourced from a time window different from the one used in training (i.e., test data II). The precision of the models significantly dropped to 0.28%, 0.26% and 1.11% for Model A, B and C respectively. This shows that the trained models did not generalize well to new datasets from future time windows. The precision is particularly low because of the imbalanced nature of the problem, *i.e.* there are much more negative samples than positive samples so the false

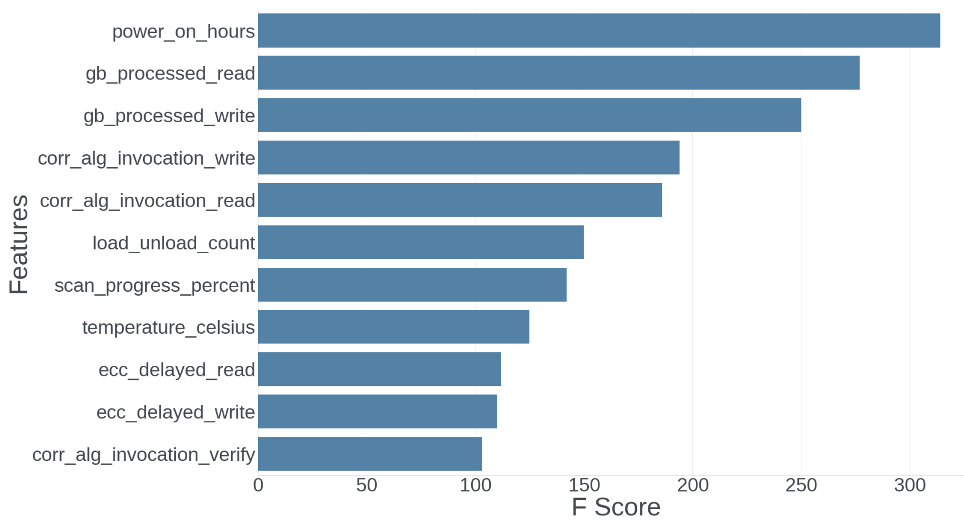

Fig. 6. The feature importance of Model C based on the number of times each feature (i.e., F-score) is used to split the data across all trees in the model. It depicts the features that have greater than 100 F-score.

positive rate tends to be high. This is because it is expected that the ratio of unhealthy to healthy drives will be minimal. However, Model C outperforms Model A and B in terms of precision and recall when the test and training sets are from the same or different time window. Model C has a precision of 1.11% compared to Model A and B when the test set is from a different time window. This suggests that utilizing the change in SMART metrics over a period of 30 days may enhance the capability to predict HDD failures.

With the low precision, the models are not ready for deployment in production considering the engineering overhead of the false positives. For a model to work in production, we need to evaluate the saving from the correctly predicted failures (i.e., true positives), and the cost of incorrectly predicted failures (i.e., false positives). This evaluation is important to ensure that the overall benefits of the predictions outweigh the engineering overhead of the false positives.

TABLE I
PERFORMANCE MEASURE OF THE XGBOOST CLASSIFIERS USING DIFFERENT LEVELS OF TEST DATA.

Model	Precision (%)		Recall (%)	
	Test data I	Test data II	Test data I	Test data II
A	88.16	0.28	86.70	94.39
B	83.84	0.26	85.47	94.16
C	**97.59**	**1.11**	**97.52**	**97.76**

C. Feature Analysis

The feature importance for Model C is demonstrated in Fig. 6, where features are ranked based on their significance in classifying (i.e., F-score) the drives. F-score measures the significance of a feature in a model, and is determined based on the frequency each feature is used for splitting the data in all the trees of the XGBoost model. Only the top ranked features are shown in Fig. 6.

It can be observed that the most important feature is the power-on-hours, suggesting that the age of the HDD in operation is indicative of identifying a faulty drive in our fleet. Furthermore, the workload is another important feature, with read and write workload serving as pointers in determining the health of a drive in our fleet.

V. CONCLUSION

The studies presented in this paper show the influence that workload and age have on HDD AFRs in Meta's production hardware fleet. The first study finds that AFR increases as HDD age increases in the fleet. Through the study, we find that the AFR across HDD models may or may not be similar to one another. Additionally, HDD AFR trends across Meta's storage infrastructure can be different from the statistical models presented within industry. The second study finds that there is a difference in the workload distributions between healthy and unhealthy HDDs. The median cumulative workload is 1.5x higher for unhealthy HDDs than healthy HDDs. However, there is an overlap in the interquartile ranges of workload between healthy and unhealthy HDDs.

In addition to our fleet analytic studies, we present a feature study using an ML algorithm to find the correlation between HDD failures and SMART metrics. We observe that age and workload related SMART metrics are most correlated to HDD failures. We could not obtain a production worthy ML model by using the SMART metrics. However, in the future, we plan to extend this effort to include additional HDD metrics, such as extended SMART logs, server level metrics, and application performance metrics.

ACKNOWLEDGMENT

We would like to acknowledge and thank Chintu Abraham, Yashar Bayani, Kevin Drucker, Sarah Hanna, Xun Jiao, Raman Kaur, Jim Mcclatchey, Manish Modi, Daniel Moore, James Munninghoff, Sriram Sankar, Neha Tibrewal, Matt Xiao, Wilson Yee, and Gloria Zhu for their support and contributions to this effort.

References

[1] S. Pan, T. Stavrinos, Y. Zhang, A. Sikaria, P. Zakharov, A. Sharma, S. Shankar P, M. Shuey, R. Wareing, M. Gangapuram, G. Cao, C. Preseau, P. Singh, K. Patiejunas, JR Tipton, E. Katz-Bassett, W. Lloyd, "Facebook's Tectonic Filesystem: Efficiency from Exascale", 19th USENIX Conference on File and Storage Technologies, 2021.

[2] S. Muralidhar, W. Lloyd, S. Roy, C. Hill, E. Lin, W. Liu, S. Pan, S. Shankar, V. Sivakumar, L. Tang, S.v Kumar,"Facebook's Warm BLOB Storage System", 11th USENIX Symposium on Operating Systems Design and Implementation, Broomfield, Colorado, 2014.

[3] JBOD *https://www.techtarget.com/searchstorage/definition/JBOD*

[4] K.V. Rashmi, N. Shah, D. Gu, H. Kuang, D. Borthakur, K. Ramchandran, "A Solution to the Network Challenges of Data Recovery in Erasure-coded Distributed Storage Systems: A Study on the Facebook Warehouse Cluster", USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage), 2013

[5] Chen PM, Lee EK, Gibson GA, Katz RH, Patterson DA. RAID: High-performance, reliable secondary storage. ACM Computing Surveys (CSUR). 1994 Jun 1;26(2):145-85.

[6] L. Meng, C. Zhang, M. Dong and B. Hou, "Analysis of Degradation Characteristics of Hard Disk Drives in Accelerated Aging Tests and Exploration of PHM Methods," 2022 IEEE International Conference on Sensing, Diagnostics, Prognostics, and Control (SDPC), Chongqing, China, 2022, pp. 162-165, doi: 10.1109/SDPC55702.2022.9915903.

[7] J. Zhao et al., "Disk Failure Early Warning Based on the Characteristics of Customized SMART," 2020 19th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Orlando, FL, USA, 2020, pp. 1282-1288, doi: 10.1109/ITherm45881.2020.9190324.

[8] Pinheiro, Eduardo, Wolf-Dietrich Weber, and Luiz André Barroso. "Failure trends in a large disk drive population." (2007).

[9] Jon G. Elerath and Sandeep Shah. Server class disk drives: How reliable are they? In Proceedings of the Annual Symposium on Reliability and Maintainability, pages 151 – 156, January 2004.

[10] I. Manousakis, S. Sankar, G. McKnight. Environmental conditions and disk reliability in free-cooled datacenters. 14th USENIX conference on file and storage technologies (FAST 16), Santa Clara, CA, 22–25 February 2016, pp.53–65.

[11] Lu, Sidi, Bing Luo, Tirthak Patel, Yongtao Yao, Devesh Tiwari, and Weisong Shi. "Making disk failure predictions smarter!." In FAST, pp. 151-167. 2020

[12] S. Huang, S. Fu, Q. Zhang and W. Shi, "Characterizing Disk Failures with Quantified Disk Degradation Signatures: An Early Experience," 2015 IEEE International Symposium on Workload Characterization, Atlanta, GA, USA, 2015, pp. 150-159, doi: 10.1109/IISWC.2015.26.

[13] J. Ieamsaard, F. E. Sandnes and P. Muneesawang, "An Area-Based Prior Value Method for Detection of Micro Contamination in Hard Disk Drives," in IEEE Access, vol. 5, pp. 10017-10023, 2017, doi: 10.1109/ACCESS.2017.2706498.

[14] F. Lin, M. Beadon, H. D. Dixit, G. Vunnam, A. Desai and S. Sankar, "Hardware Remediation at Scale," 2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W), Luxembourg, Luxembourg, 2018, pp. 14-17, doi: 10.1109/DSN-W.2018.00015.

[15] A. Chattopadhyay, "Hard Disk Drive - Reliability Overview", HEPIX Fall 2013 Workshop, 2013

[16] K.Bindu Madhavi, BH.Sowmya, B. Sandeep, M.Lavanya, "Derating Analysis for Reliability of Components", International Journal of Engineering Research and General Science Volume 2, Issue 5, 2014

[17] Z. He, H. Yang and M. Xie, "Statistical modeling and analysis of hard disk drives (HDDs) failure," 2012 Digest APMRC, Singapore, 2012, pp. 1-2.

[18] S. Kamarthi, A. Zeid and Y. Bagul, "Assessement of current health of hard disk drives," 2009 IEEE International Conference on Automation Science and Engineering, Bangalore, India, 2009, pp. 246-249, doi: 10.1109/COASE.2009.5234105.

[19] *Exos X18 Product Manual*, Seagate, 2021, 100865853, Rev. C, https://www.seagate.com/files/www-content/product-content/enterprise-hdd-fam/exos-x18-channel/_shared/en-us/docs/100865853c.pdf

[20] *Enterprise Capacity HDD MG09 Series Product Manual*, Toshiba, 2021, Rev 00, https://storage.toshiba.com/docs/enterprise-hdd-documents/p_l_mg09aca16tay_210202072836_1_0_0e.pdf

[21] *Hard Disk Drive Specifications Ultrastar DC HC550*, Western Digital, 2021, Rev 1.3, https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/product/data-center-drives/ultrastar-dc-hc500-series/product-manual-ultrastar-dc-hc550-sata-oem-spec.pdf

[22] S. Sankar, M. Shaw and K. Vaid, "Impact of temperature on hard disk drive reliability in large datacenters," 2011 IEEE/IFIP 41st International Conference on Dependable Systems Networks (DSN), Hong Kong, China, 2011, pp. 530-537, doi: 10.1109/DSN.2011.5958265.

[23] X. Chen, *Hard Disk Drive Control*, Mechanical Systems Control Lab, Berkeley, 2016, *https://msc.berkeley.edu/research/disk.html*