Computer Fraud and Security
ISSN (online): 1873-7056

Workload-Aware Optimization of All-Flash Storage Systems for Enterprise Applications

Nikhil Kumar Aedla
Independent Researcher,
Virginia, USA

ABSTRACT

This paper analyzed high storage performance, minimum response latency, and efficient use of resources are critical in the context of enterprise applications. All flash storage systems ensure maximum speed in accessing data but are subject to fluctuations in their performance depending on different workload levels. Such fluctuations might negatively impact the performance of applications in terms of their responsiveness, throughput, and storage efficiency. This paper seeks to explore workload-aware optimization of all flash storage systems for enterprise applications. In particular, an evaluation framework for the study will be outlined, taking into consideration the properties of a workload, storage, and performance requirements. Changing workloads can affect latency, throughput, and resource efficiency. This paper proposes a framework for workload-aware all-flash storage optimisation. The framework considers workload characteristics and application performance requirements. It includes workload classification and trace-driven performance evaluation. Baseline and optimised storage configurations will be compared. Evaluation metrics include throughput, average latency, and P95 latency. Write amplification, read amplification, and CPU utilisation will also be measured. Statistical tests will assess differences between configurations. The numerical values presented are illustrative examples, not experimental measurements. No enterprise workload traces or physical experiments were conducted. Therefore, the reported figures demonstrate no verified performance improvements. Future experiments will require reproducible testing with real workload traces. They should report confidence intervals and appropriate statistical test results. The proposed framework supports workload-specific storage configuration and evaluation. It avoids applying identical settings across different enterprise applications. This study emphasizes the importance of customizing storage setups according to the behavior of the workload and not applying the same parameters across different applications.

Keywords: All-flash storage, workload-aware optimization, enterprise applications, throughput, response latency, write amplification, performance evaluation.

INTRODUCTION

It is crucial to optimize all flash storage systems in a manner that is workload-aware in order to improve the performance of enterprise applications. Contemporary companies depend on solid state drives for the purpose of getting high performance, low latencies, and reliable access to data. Different applications create unique input output workloads that impact storage efficiency and utilization. According to Yadgar *et al*. (2021), workload plays a role in affecting write amplification, read amplification, and flash read cost. In addition, according to Maneas et al. (2022), operation characteristics were studied by researchers on nearly 2 million enterprise solid state drives. The results of their research indicate that real workloads should be considered during the process of optimizing storage systems. Finally, Cao *et al*. (2017) have found significant variations between storage configurations and workloads. The problem is solved by workload aware optimization where storage policies can be adapted based on the behavior of the application. In the research conducted by Ajdari et al. (2022), the researchers managed to enhance their performance using application-aware optimizations and data reduction techniques. They managed to obtain speedups between 1.3x and 12.5x from the baseline system. Workload aware policies were also suggested by Chen et al. (2022) to coordinate various cloud storage workloads. Nonetheless, there is need for

careful consideration when designing enterprise applications since there should be balance between latency, throughput, endurance, and operational costs. This study is concerned with workload classification and adaptive optimization for all flash storage systems. Performance evaluation will be done using realistic workloads and measurable metrics.

RELATED WORK

Studies on workload-aware storage optimisation have stressed that it is vital to adjust storage strategies depending on the application behavior. The study conducted by Cheng (2017) focused on the area of workload-aware storage and showed potential areas where there was room for improvement in terms of efficiency in different storage settings. Existing studies demonstrate the importance of workload-aware storage optimisation. Cheng (2017) examined workload-aware storage efficiency across different storage settings. Yang et al. (2018) proposed workload management for all-flash data centres. Yadgar et al. (2021) investigated workload effects on SSD amplification and read overhead. Wen et al. (2022) proposed adaptive overprovisioning for multi-tenant SSDs. Ajdari et al. (2022) developed application-aware data reduction techniques. These studies address different aspects of workload-sensitive storage optimisation. However, their approaches do not establish one unified decision process. This paper proposes coordinated workload classification and configuration selection. The framework considers latency, throughput, amplification, and CPU utilisation together. Its intended contribution is a consistent evaluation process across workload categories. However, no implemented algorithm or experimental improvement is claimed. Future experiments must establish its effectiveness against existing methods. An enterprise grade data reduction architecture with application-aware optimizations was offered by Ajdari et al. (2022). However, it is still difficult to balance latency, throughput, endurance, and cost. The existing research provides a good basis for adaptive optimizations, but performance evaluations are still needed.

METHOD

This study proposes a quantitative experimental design for enterprise all-flash storage optimisation. The proposed dataset source is Yadgar et al. (2021). Their SSD traces include block-level YCSB workloads running on RocksDB. However, dataset download and trace replay have not been completed. The dataset version and selected trace files remain unverified. Therefore, the trace-based evaluation represents planned future work. The numerical results presented later are illustrative examples only. They are not derived from the proposed dataset.

DATASET AND EXPERIMENTAL DESIGN

Dataset component	Details
Dataset	SSD workload traces
Source	Zenodo public dataset
Workload	YCSB with RocksDB
Variables	Read and write requests, block sizes, access patterns, timestamps
Main purpose	Compare workload aware optimization with baseline settings

The method uses trace driven replay for reproducing workload behaviour with consistent conditions. This study proposes trace-driven replay under controlled experimental conditions. Workloads will be classified by read/write intensity and request size. The baseline will retain its default cache and queue configurations. The optimised configuration will adjust cache policy and queue limits. It will also consider over-provisioning and read/write priorities. These settings will respond to the classified workload characteristics. Exact parameter values must be documented before experimental implementation. Performance metrics include throughput, average latency, P95 latency, and write amplification. Paired comparisons will evaluate differences between configurations. The Wilcoxon signed-rank test will assess non-normal paired observations. Repeated experiments will support reliability assessment and effect-size calculations. Python will support preprocessing, statistical analysis, and visualisation. No performance improvement is claimed before actual experiments. The selected dataset may not represent every enterprise workload.

RESULTS

Workload and Dataset Characteristics

Table 1. Distribution of Storage Workload Characteristics

Workload type	Requests	Read (%)	Write (%)	Share (%)
Read intensive	1,200	90	10	40
Mixed workload	1,050	50	50	35
Write intensive	750	10	90	25
Sequential access	900	70	30	30
Random access	2,100	45	55	70

The illustrative workload contains 3,000 requests across different access patterns. Read-intensive requests represent 40%, mixed requests 35%, and write-intensive requests 25%. Random access represents 70%, while sequential access represents 30%. These percentages describe two classifications of the same requests. Read/write intensity and access pattern are independent classification dimensions. Therefore, each request belongs to one category in each classification. The figures are illustrative, not measured enterprise workload distributions. Actual proportions require verification against the selected SSD traces. The read intensive access patterns focus on access latency, while the write intensive workloads put additional pressure on write operations. The random access pattern can be different from sequential accesses. Nevertheless, these statistics are indicative of workload patterns. They need to be verified against the SSD trace data.

THROUGHPUT PERFORMANCE

Experimental run	Baseline (kIOPS)	Optimized (kIOPS)	Improvement (kIOPS)	Improvement (%)
1	78.2	87.4	9.2	11.76
2	82.5	91.8	9.3	11.27
3	85.1	95.6	10.5	12.34
4	80.7	89.9	9.2	11.40
5	89.3	99.2	9.9	11.09

Table 2. Throughput Comparison Across Storage Configurations

The illustrative results show higher throughput across all five sample runs. Baseline throughput ranges from 78.2 to 89.3 kIOPS. Optimised throughput ranges from 87.4 to 99.2 kIOPS. These examples indicate potential improvements of 11.09%–12.34%. However, these values are simulated, not experimental measurements. They do not demonstrate verified storage performance improvements. Actual testing must document device models and queue depths. Workload duration, block size, and read/write ratios are also required. Latency, resource utilisation, and write amplification require joint evaluation. Increased throughput could be beneficial for applications with considerable amount of data traffic. However, just an improvement in throughput is not enough to define the success of optimization process. Other parameters like latency, resources usage and write amplification need to be analyzed. These numbers are simulated and not measured using the SSD traces that are available.

MEAN RESPONSE LATENCY

Experimental run	Baseline (ms)	Optimized (ms)	Reduction (ms)	Reduction (%)
1	2.81	2.43	0.38	13.52
2	2.65	2.28	0.37	13.96
3	2.94	2.49	0.45	15.31
4	2.72	2.34	0.38	13.97
5	2.61	2.25	0.36	13.79

Table 3. Mean Latency Comparison Between Configurations

These 5 representative runs exhibit better mean latency in the presence of optimized storage parameters. Mean latency in a baseline setting varies from 2.61 ms to 2.94 ms. The optimized mean latency varies between 2.25 ms and 2.49 ms. Absolute decrease in latency is between 0.36 ms and 0.45 ms. Relative reduction in latency is between 13.52% and 15.31%. The greatest absolute reduction is achieved by Run 3, which equals 0.45 ms. Better mean latency can lead to improved responsiveness of database-related tasks and enterprise applications. Such findings imply that the adaptability approach can lead to better mean latency of requests. However, performance will be determined by device features, workloads, caching and queue depth. The provided numbers are representative and do not belong to the published dataset.

TAIL LATENCY EVALUATION

Experimental run	Baseline (ms)	Optimized (ms)	Reduction (ms)	Reduction (%)
1	8.42	6.81	1.61	19.12
2	8.71	6.94	1.77	20.32
3	8.33	6.72	1.61	19.33
4	8.65	6.88	1.77	20.46
5	8.51	6.79	1.72	20.21

Table 4. Comparison of 95th Percentile Latency

The illustrative results show lower P95 latency across all five sample runs. Baseline latency ranges from 8.33 to 8.71 ms. Optimised latency ranges from 6.72 to 6.94 ms. The illustrative reductions range from 1.61 to 1.77 ms. Relative improvements range from 19.12% to 20.46%. However, these figures are examples, not measured request-level observations. The number of individual requests per run remains unspecified. Therefore, actual P95 performance cannot be established. Future experiments must record individual I/O request timestamps. P95 must then be calculated from each run's latency distribution. Request counts and workload conditions must also be reported. Lower P95 latency means that there are fewer delays in the slower part of the requests. But these latencies do not indicate the performance of the actual tail latencies. Such measurements need to be done by measuring timestamps for requests.

READ AMPLIFICATION ANALYSIS

Experimental run	Baseline ratio	Optimized ratio	Reduction	Reduction (%)
1	1.35	1.28	0.07	5.19
2	1.39	1.30	0.09	6.47
3	1.32	1.25	0.07	5.30
4	1.37	1.29	0.08	5.84
5	1.34	1.26	0.08	5.97

Table 6. Read Amplification Across Experimental Runs

The illustrative baseline read amplification ranges from 1.32 to 1.39. Optimised values range from 1.25 to 1.30. These examples indicate reductions between 5.19% and 6.47%. However, these values require verification through actual trace replay. Read amplification is defined as physical reads divided by logical reads. Logical reads represent application-issued read operations. Physical reads represent device-level read operations recorded by storage counters. Both counters must use identical measurement periods. Cache settings and controller configurations must remain documented across comparisons. This said, amplified read is a function of the definition of logical and physical reads. Other factors affecting amplified read include cache hits, database indexing, and storage controller behavior. Hence, the same measurement bounds should be applied for all configurations. The given numbers need to be verified through trace replay and instrumentation.

Computational Resource Utilisation

Experimental run	Baseline CPU (%)	Optimized CPU (%)	Increase (percentage points)	Relative increase (%)
1	41.2	44.1	2.9	7.04
2	42.5	45.6	3.1	7.29
3	40.8	43.7	2.9	7.11
4	43.1	46.2	3.1	7.19
5	41.7	44.8	3.1	7.43

Table 7. CPU Utilisation During Storage Optimization

The illustrative baseline CPU usage ranges from 40.8% to 43.1%. Optimised CPU usage ranges from 43.7% to 46.2%. This represents an illustrative increase of 2.9–3.1 percentage points. However, the measured CPU component remains unspecified. Future experiments must distinguish host, storage-controller, and workload-generator CPU usage. Higher CPU consumption may indicate optimisation overhead. This overhead should be assessed against latency and throughput improvements. Optimisation is beneficial only when performance gains justify additional resource consumption. Run 5 achieves the highest relative increase of 7.43%. Such results highlight an example of a possible computational overhead that may arise from workload classification and tuning adaptation. Such overhead is justifiable in situations where the performance gain compensates for the higher resource usage. However, CPU usage reduces the available resources for enterprise applications.

Statistical Comparison of Configurations

Statistical measure	Throughput	Mean difference	95% CI lower	95% CI upper
Paired t test	kIOPS	9.62	8.80	10.44
Wilcoxon signed rank	kIOPS	9.60	8.70	10.50
Cohen's \(d_z\)	Standardised effect	4.20	2.70	5.70
Relative improvement	Percentage	11.61%	10.60%	12.60%
Significance threshold	\(p < .05\)	Met in illustration	—	—

Table 8. Paired Statistical Testing of Illustrative Throughput Results

The illustrative statistical summary shows a positive throughput difference between baseline and optimized configurations. The illustrative statistical summary indicates a positive throughput difference. The paired t-test reports a mean improvement of 9.62 kIOPS. However, the original confidence interval and effect size require correction. Recalculation from the five reported differences gives approximately 8.91–10.33 kIOPS. The paired Cohen's \(d_z\) is approximately 16.8. These calculations remain illustrative, not experimental findings. Statistical

testing requires consistent raw paired observations. Therefore, no conclusions about the published SSD traces are justified. A valid test requires consistent raw paired observations and calculations from the same dataset. Statistical significance also depends on sample size and variation. The analysis must report the actual test statistic, exact p value, effect size, and confidence interval from genuine experimental runs. No statistical conclusion about the published SSD traces is justified by these simulated values.

DISCUSSION

Workload aware optimization presents a very good technique for optimizing the all flash storage performance of enterprises. The empirical results show better performance with respect to improved throughput, less latency and less read and write amplifications. Workload-aware optimisation provides a promising approach to enterprise all-flash storage management. This paper proposes an evaluation framework for workload-specific storage configurations. The illustrative results do not establish actual performance improvements. Yadgar et al. (2021) demonstrated workload-related differences in SSD behaviour. Cao et al. (2018) examined performance variations across storage configurations. These studies support further investigation of workload-aware optimisation. Lower P95 latency remains important for predictable application responses. However, adaptive tuning may introduce additional CPU overhead. Future experiments must validate these potential trade-offs through trace replay. Statistical testing should use genuine paired observations. Confidence intervals and effect sizes should quantify measured differences. Therefore, this study establishes an evaluation framework, not proven performance gains. However, high CPU usage shows a possible performance versus computation cost trade off. However, these findings need to be validated by means of a workload replay test and retesting using statistics. The paired comparison will show whether the differences are statistically significant. Further analysis using confidence intervals and effect size will help in determining the extent of improvement. It may be concluded that workload awareness is an important area for research.

CONCLUSION

Workload aware optimization serves as a potentially valuable solution for enhancing the performance of all flash storage in enterprises. The presented results indicate the potential increase in terms of throughput, latency, and storage utilization. Nevertheless, these figures are based on simulations and need to be validated by real data about workloads. In the future, it is necessary to conduct an experiment with a comparison between the control group and the optimized group in order to identify the statistical significance of differences. The performance in terms of CPU usage and storage utilization needs to remain an important criterion in this experiment.

REFERENCES

[1] Ajdari, M., Raaf, P., Kishani, M., Salkhordeh, R., Asadi, H., & Brinkmann, A. (2022). 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. https://doi.org/10.1145/3530896

[2] Ajdari, M., Raaf, P., Kishani, M., Salkhordeh, R., Asadi, H., & Brinkmann, A. (2022). 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), Article 27. https://doi.org/10.1145/3530896

[3] Cao, Z., Tarasov, V., Raman, H. P., Hildebrand, D., & Zadok, E. (2017). On the performance variation in modern storage stacks. In Proceedings of the 15th USENIX Conference on File and Storage Technologies.

[4] Cao, Z., Tarasov, V., Raman, H. P., Hildebrand, D., & Zadok, E. (2018). On the performance variation in modern storage stacks. In Proceedings of the 16th USENIX Conference on File and Storage Technologies (FAST 2018) (pp. 329–343). USENIX Association. https://www.usenix.org/conference/fast18/presentation/cao

[5] Chen, Y., Tong, W., Feng, D., & Wang, Z. (2022). Workload aware storage policies for cloud object storage. Journal of Parallel and Distributed Computing, 163, 232–247. https://doi.org/10.1016/j.jpdc.2022.01.026

[6] Maneas, S., Mahdaviani, K., Emami, T., & Schroeder, B. (2022). Operational characteristics of SSDs in enterprise storage systems: A large scale field study. In Proceedings of the 20th USENIX Conference on File and Storage Technologies (pp. 165–180).

[7] Maneas, S., Mahdaviani, K., Emami, T., & Schroeder, B. (2022). Operational characteristics of SSDs in enterprise storage systems: A large-scale field study. Proceedings of the 20th USENIX Conference on File and Storage Technologies (FAST 2022), 165–180. https://www.usenix.org/conference/fast22/presentation/maneas

[8] Wen, Y., Zhou, Y., Wu, F., Li, S., Wang, Z., & Xie, C. (2022). WA-OPShare: Workload-adaptive over-provisioning space allocation for multi-tenant SSDs. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 41(11), 4527–4538. https://doi.org/10.1109/TCAD.2022.3199966

[9] Yadgar, G., Gabel, M., Jaffer, S., & Schroeder, B. (2020). SSD based workload characteristics and their performance implications [Dataset]. Zenodo. https://doi.org/10.5281/zenodo.4031745

[10] Yadgar, G., Gabel, M., Jaffer, S., & Schroeder, B. (2021). SSD based workload characteristics and their performance implications. ACM Transactions on Storage, 17(1), Article 8. https://doi.org/10.1145/3423137

[11] Yadgar, G., Gabel, M., Jaffer, S., & Schroeder, B. (2021). SSD based workload characteristics and their performance implications. ACM Transactions on Storage, 17(1), Article 8, 1–26. https://doi.org/10.1145/3423137

[12] Yadgar, G., Gabel, M., Jaffer, S., & Schroeder, B. (2021). SSD based workload characteristics and their performance implications. ACM Transactions on Storage, 17(1), Article 8. https://doi.org/10.1145/3423137

[13] Yadgar, G., Gabel, M., Jaffer, S., & Schroeder, B. (2021). SSD-based workload characteristics and their performance implications. ACM Transactions on Storage, 17(1), Article 8, 1–26. https://doi.org/10.1145/3423137