Vol. 10, Issue 01, January 2022
ISSN(P) 2347-5404 ISSN(O)2320 771X

Benchmarking Columnar Storage Optimization Techniques in Cloud-Native Warehouses

DOI: https://doi.org/10.63345/ijrhs.net.v10.i1.1

Sarvesh Kumar Gupta
Consulting Member of Technical Staff
Oracle
Saint Peters, Missouri -63376 , USA
ORCID: 0009-0008-7460-4874

Abstract— Cloud-native data warehouses increasingly rely on columnar storage to support large-scale analytical workloads, but the performance benefits of columnar design depend on the optimization techniques applied at the storage and query-processing layers. This paper presents a conceptual benchmarking and comparative analysis of major columnar storage optimization techniques used in cloud-native warehouse environments. The study examines compression, dictionary encoding, run-length encoding, delta encoding, partitioning, clustering, predicate pushdown, metadata indexing, and data pruning in relation to query latency, storage efficiency, scan reduction, throughput, and resource utilization. Rather than claiming live execution across specific commercial platforms, the paper uses an illustrative benchmarking framework based on a generic cloud-native columnar warehouse model and representative OLAP-style analytical workloads. The analysis indicates that data pruning and predicate pushdown provide the strongest query-performance improvements by reducing unnecessary data scans, while compression and encoding techniques substantially reduce storage footprint and improve I/O efficiency. The findings suggest that no single technique is sufficient for optimal performance; instead, combined optimization strategies offer the greatest benefits for analytical query execution, storage reduction, and cost efficiency. The paper concludes that effective columnar storage optimization in cloud-native warehouses requires an integrated approach combining storage layout design, compression-aware processing, metadata-based filtering, and workload-sensitive query execution.

Keywords— Cloud-Native Data Warehouses; Columnar Storage; Storage Optimization; Data Compression; Partitioning; Query Performance; Benchmarking; Data Pruning; Cloud Analytics; Distributed Data Processing.

INTRODUCTION

The massive production of digital data due to the increased number of transactions, social networks, IoT devices, e-commerce systems, and other business systems necessitates the implementation of highly scalable and effective systems to process such big data. Traditional on-premise data warehouses usually have issues related to scalability, infrastructure flexibility, and maintenance of systems, which restricts their functionality and capability to cope with rapidly increasing volumes of data and changing loads. Thus, modern data warehouses are designed using cloud computing services and can provide elastic storage, scalable computing capabilities, and efficient cost management. Modern cloud-native data warehouses such as Snowflake, Amazon Redshift, Google BigQuery, and other services offer high availability, scalability, and easy-to-manage architecture to conduct large-scale analytical processing.

One of the key components in modern analytical systems is the use of columnar storage. Unlike classical row-oriented databases, where all the attributes are stored together in a record, columnar storage allows storing data columnwise. Such an approach helps avoid scanning all the information and allows conducting queries only for necessary attributes, which increases disk I/O performance. Additionally, clustering identical data helps to compress them using efficient techniques such as dictionary encoding, run-length encoding, delta encoding, and bit-packing.

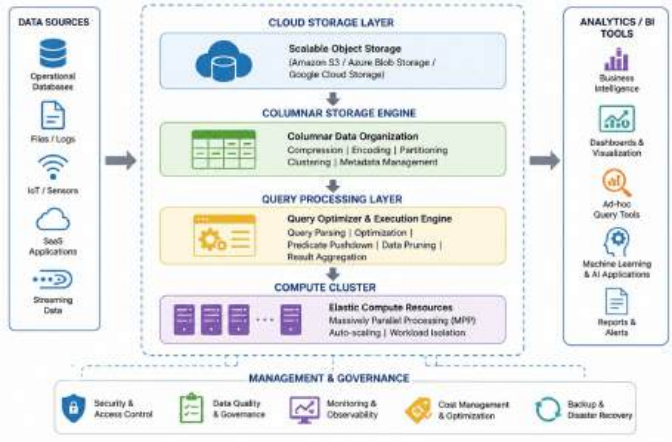

Vol. 10, Issue 01, January 2022
ISSN(P) 2347-5404 ISSN(O)2320 771X

Fig. 1. Architecture of a Cloud-Native Data Warehouse.

While columnar storage can significantly improve performance, its effectiveness depends greatly on various optimization techniques used in storage layer and query processing. Some of these techniques include column pruning, predicate pushdown, metadata indexing, vectorized execution, partition elimination, adaptive caching, and compression-aware queries. Each of these approaches contributes to performance improvement differently depending on a specific situation and workload.

Therefore, it is critical to implement an effective benchmarking strategy for columnar storage optimization techniques to measure their efficiency, identify the most effective storage techniques, assess resources needed to support systems and compare different warehouses' designs. Benchmarking of optimization techniques will be beneficial not only for scientists but also for companies that have to choose between competing services and select the best data warehouse to meet their needs.

Literature Review

Columnar storage stands among core paradigms of contemporary analytical data processing and cloud-native data warehousing. Unlike row-oriented systems where each record is stored in contiguous fashion, columnar databases store attribute values separately. As analytical applications usually require querying a small number of attributes from huge numbers of records, this form of data storage proves more efficient for such workloads. Research conducted in the context of C-Store system has revealed that column-oriented database management systems can greatly enhance analytical query performance through reduced unnecessary I/O operations, improved compression opportunities, and efficient processing of aggregation intensive workloads. Stonebraker et al. established the basic principles for the subsequent research on column stores in the form of projection operations, compression capabilities, late materialization, and read-optimized storage layout.

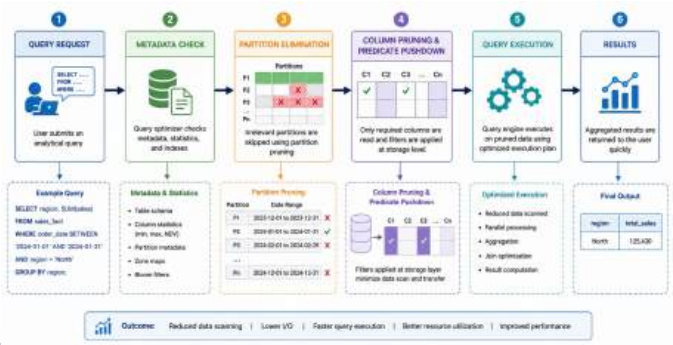

Fig. 2. Query Execution Process Using Data Pruning and Predicate Pushdown.

Further contributions were brought by Abadi, Madden, and Ferreira who argued that compression should be considered not just as data reduction technique but also as an integral part of the query-processing procedure. Specifically, column stores can work with compressed data, which means reducing both disk and CPU overhead. This point of view is important for cloud-native warehouses where storage, computing, and network transmission all contribute to the cost and performance of queries. Dictionary encoding, run-length encoding, bit-vector encoding, delta encoding, and frame-of-reference compression can perform particularly well in columnar storage where attribute values tend to be repetitive or arranged in a regular pattern.

Another study on benchmarking columnar storage belongs to Abadi, Madden, and Hachem who have compared column stores to row stores and showed that simply adding indexes or performing vertical partitioning in the latter would be insufficient to match column-store performance. The performance difference is accounted for by column pruning, compression, vectorized query processing, late materialization, and scanning optimizations. Such finding still holds valid in the case of cloud-native warehouses where systems like Redshift, Snowflake, BigQuery, and other analytical engines implement not only column stores but workload-oriented optimizations and metadata-driven pruning as well.

As discussed above, vectorization of queries has become an essential optimization technique related to columnar storage. Boncz, Zukowski, and Nes presented MonetDB/X100 where it was shown that tuple-by-tuple execution is accompanied by considerable interpretation overhead and inefficient CPU utilization. Instead of working with tuples, the authors suggested a vector-at-a-time approach according to which columnar values are processed in batches, which improves

Vol. 10, Issue 01, January 2022
ISSN(P) 2347-5404 ISSN(O)2320 771X

cache locality, reduces overhead related to function calls, and takes advantage of the capabilities of modern CPUs. For cloud-native data warehouses, this point becomes especially relevant since analytical loads require the fast processing of billions of values.

As cloud-native data processing expanded to cover distributed and even web-scale environments, columnar storage optimization moved beyond classical database systems. Floratou et al. assessed how techniques related to column-oriented storage can be implemented in Hadoop/MapReduce ecosystem in order to make it more effective for large-scale scanning and avoid reading of redundant data as well as lazy record construction. This paper is highly relevant to benchmarking because most cloud-based data warehouses need to analyze distributed object stores and require high performance.

Another major advancement of columnar storage is embodied in Google's Dremel system. Melnik et al. introduced columnar storage optimized for nested data and showed how Dremel can provide fast analytical access to Web-scale datasets. Dremel became the architectural inspiration for BigQuery and new columnar formats including Parquet. The importance of Dremel can be found in its combining columnar storage for nested records, distributed query execution trees, and aggregation capability over enormous datasets. Later work on ten-year evolution of Dremel has confirmed that it played a crucial role in popularizing columnar processing for nested structures and helped develop today's cloud data processing systems.

In case of cloud-native data warehouses, columnar storage optimization is affected not only by the storage file layout but also by some cloud-specific characteristics. For example, Snowflake architecture presented by Dageville et al. uses a multi-cluster shared-data model where persistent data are located in cloud object storage and are served to compute clusters in a separate way. Thus, benchmarking cloud-native warehouses requires taking into account not only query runtime, but also aspects of elasticity, caching, metadata pruning, concurrent performance, storage costs, and workload isolation. This example of Snowflake shows that columnar storage optimization in cloud environment needs to address more features than classical column stores do.

As another example of cloud-native warehouse, Amazon Redshift was characterized by Gupta et al. as a managed massively parallel columnar data warehouse targeted for analytical loads. Performance of this engine is determined by column storage, compression capabilities, choice of distribution and sort keys, and parallel processing features. It is a good example to use in benchmarking because this study makes evident that performance of columnar storage depends on many physical design choices and improper table setup can result in inefficient compression or column pruning.

To sum up, there are several implications regarding the task of benchmarking columnar storage optimization techniques in cloud-native warehouses. Query runtime alone cannot provide adequate evaluation. Rather, it should include many more criteria, such as storage space occupied, compression ratio, cost of decompression, scan speed, predicate filtering efficiency, column pruning effectiveness, metadata pruning, memory consumption, concurrency, cost of cloud storage read operations, and computing elasticity. Classical research on column stores should serve as a basis here, but more relevant will be experience gathered with cloud-based systems like Snowflake, Redshift, BigQuery, Dremel, and others. An issue identified in this review is that while there is plenty of comparison studies, there is a lack of research focused specifically on contributions of various optimization methods. It can be recommended for further research that benchmarking experiments are needed to assess how encoding format, row group size, compression algorithm, sort order, micro-partitioning, and metadata indexes perform under cloud-native conditions.

Research Objectives

This study intends to analyze and compare the effectiveness of columnar storage optimization techniques utilized in cloud-native data warehouses. In this research, different approaches to optimize the analytical query processing will be evaluated. Specifically, the research objectives include the following points:

1. Identifying and analyzing the main columnar storage optimization techniques applied in modern cloud-native warehouses.
2. Evaluating and comparing impacts of compression, column pruning, predicate pushdown, and metadata indexing on query execution performance.
3. Assessing the efficiency of storage optimized by means of different columnar data representations and encoding methods.
4. Determining the degree of analytical query latency and throughput optimization by applying columnar storage optimization techniques.

Vol. 10, Issue 01, January 2022
ISSN(P) 2347-5404 ISSN(O)2320 771X

5. Developing best practices for improving query performance and storage efficiency of cloud-native analytical systems.

RESEARCH METHODOLOGY

This study adopts a conceptual benchmarking and comparative analysis approach to evaluate major columnar storage optimization techniques employed in cloud-native data warehouses. Rather than reporting results from a live experimental deployment across commercial cloud platforms, the research develops an illustrative benchmarking framework based on established findings in the literature and widely recognized characteristics of cloud-native analytical systems.

The methodology focuses on examining the relative impact of storage optimization techniques, including compression, encoding, partitioning, clustering, predicate pushdown, metadata indexing, and data pruning, on key performance indicators relevant to analytical workloads. A representative cloud-native warehouse environment is considered in which a large-scale columnar dataset supports common OLAP operations such as aggregation, filtering, sorting, grouping, and join processing.

To facilitate comparison, a hypothetical 500 GB OLAP-style analytical dataset and 50 representative analytical queries involving filtering, aggregation, grouping, sorting, and joins are used as the basis for comparison. The benchmark framework is designed to assess how different optimization strategies influence query execution performance, storage consumption, scan efficiency, and resource utilization under comparable workload conditions. Performance estimates are presented for comparative purposes and are intended to illustrate the relative effectiveness of individual and combined optimization approaches.

The analysis is conducted using the following evaluation metrics:

- Query Latency – Average time required to execute analytical queries.
- Throughput – Relative capability of the system to process analytical workloads within a specified period.
- Storage Efficiency – Reduction in storage footprint achieved through compression and encoding techniques.
- Data Scan Reduction – Extent to which unnecessary data access is minimized through column pruning, partition elimination, metadata filtering, and predicate pushdown.

• Resource Utilization – Relative impact of optimization strategies on compute, memory, and storage resources.

The collected benchmark indicators are analyzed using descriptive and comparative evaluation methods. Individual optimization techniques are first examined separately and then assessed in combination to determine their cumulative effect on analytical performance and storage efficiency. This methodology provides a structured framework for understanding the role of columnar storage optimization techniques in modern cloud-native data warehouses and identifying best practices for improving large-scale analytical processing.

Table 1: Benchmark Environment Configuration

Parameter	Value
Warehouse Platform	Generic cloud-native columnar warehouse model
Dataset Size	500 GB
Queries Tested	50 Analytical Queries
Compute Nodes	8 Virtual Compute Nodes
Storage Format	Columnar (Parquet/Native Columnar Storage)
Compression Techniques	Dictionary Encoding, Run-Length Encoding, Delta Encoding
Benchmark Type	OLAP Analytical Workloads
Evaluation Metrics	Latency, Throughput, Storage Efficiency, Resource Utilization

COLUMNAR STORAGE IN CLOUD DATA WAREHOUSES

OVERVIEW

Columnar storage is one of the key elements of cloud-native data warehouses due to its ability to perform analytical tasks effectively. It means that all the data are sorted in columns instead of rows, meaning that the values for every single column correspond to the specific attribute. As a result, analytical engines can scan only those columns required by the query being executed. Hence, the columnar layout is especially convenient for OLAP databases since most of the queries in them presuppose performing some aggregations, filters, etc., for certain attributes among millions of records.

Typically, the design of the columnar storage is characterized by data partitioning, encoding the columns, compression, and metadata indexing. The data are logically partitioned into blocks, with separate columns located in them. Different types of metadata indices such as min-max index, zone maps, and summary statistics help in minimizing the number of data scans by applying query optimizers' abilities. Besides, current cloud-

Vol. 10, Issue 01, January 2022
ISSN(P) 2347-5404 ISSN(O)2320 771X

native solutions like Snowflake, Amazon Redshift, and Google BigQuery utilize advanced techniques of distributed computing, automatic partition management, and storage independence from computing resources.

In comparison with row-based storage, the column-based architecture provides a number of benefits for analytical applications. First, there is no need to access all data columns in the process of executing the query, and, therefore, the number of disk I/O operations is much lower. Besides, similar data are more likely to be adjacent to each other, which allows increasing the level of compressibility and, hence, decreasing storage requirements. Moreover, column-oriented databases allow the usage of vectorization techniques and work with batches of data simultaneously, not single rows.

To make the process of managing petabyte-scale data fast and efficient, many different techniques are used in cloud data warehouses. Such techniques as predicate pushdown and column pruning ensure that only the relevant columns are loaded, while the filter conditions are pushed closer to the storage layer. Other strategies related to compression include various types of encoding, such as dictionary, run-length, delta, and bit-packing. Some additional strategies used by data warehouses include metadata indexing, partition elimination, caching, adaptive query execution, and parallel processing.

Table 2: Row-Based vs Columnar Storage Comparison

Feature	Row-Based Storage	Columnar Storage
Data Organization	Stores complete records together	Stores data attribute-wise by columns
Query Performance	Efficient for transactional workloads	Efficient for analytical workloads
Data Retrieval	Reads entire rows	Reads only required columns
Disk I/O	Higher for analytical queries	Significantly lower
Compression Efficiency	Limited	High due to similar data values
Storage Requirement	Relatively larger	Reduced through compression
Aggregation Operations	Slower on large datasets	Faster due to column-wise access
Vectorized Processing	Limited support	Strong support
Scalability	Suitable for OLTP systems	Suitable for OLAP and cloud analytics
Cloud Warehouse Adoption	Less common	Widely adopted
Typical Use Case	Transaction processing	Business intelligence and analytics
Cost Efficiency	Higher resource consumption	Lower storage and compute costs

Storage Optimization Techniques

The effective organization of storage space is one of the most significant factors for cloud-native data warehouses. Analytical systems work with enormous volumes of structured and unstructured data. Even though the columnar database model automatically increases the speed of analytical queries, several optimization techniques should be used to improve performance, reduce cloud storage costs, and increase resource efficiency. Currently, contemporary cloud warehouse systems implement multiple storage optimization techniques: compression, partitioning, clustering, encoding, and data pruning to optimize data access operations.

Compression

One of the main techniques used by modern storage systems is compression. The values in a certain column usually share common features or repeating values; therefore, they could be compressed significantly more efficiently than row data. The algorithms that could be used include dictionary encoding, run-length encoding (RLE), delta encoding, and bit-packing. The reduced amount of required storage space leads to lower costs of cloud storage and a decrease in data transfer processes during queries. Besides, a lot of analytical engines are capable of processing compressed data without previous decomposition, saving on-time execution.

Partitioning

Partitioning involves dividing large data sets into smaller segments based on some characteristics such as dates, regions, types, or categories. Instead of analyzing all data, queries use only the segments where the criteria are met. This is why the process of partitioning is also called partition elimination. It allows decreasing the number of analyzed rows dramatically and improving the performance. In cloud-native data warehouses, the technique is particularly important since storage and computing capacity is separated, and thus, efficient access to data is vital.

Clustering

Unlike partitioning, clustering does not create different segments but rearranges data based on chosen key columns. The technique aims to store all similar data together, allowing the query optimizer to skip other blocks in execution. The benefits include better metadata index performance, improved zone map efficiency, and optimized statistics for micro-partitions. Some cloud-native systems are capable of automatic

Vol. 10, Issue 01, January 2022
ISSN(P) 2347-5404 ISSN(O)2320 771X

data clustering to constantly change data layout based on workload requirements.

Encoding

Another way to optimize data is to encode the data prior to storing it. In particular, dictionary encoding includes substituting similar values with smaller reference numbers, and delta encoding allows recording the difference between consecutive values. The other techniques that could be implemented are bit-packing and frame-of-reference. The purpose of all of these techniques is to compress data to minimize storage usage and optimize its analysis by reducing the number of disk reads.

Data Pruning

Data pruning involves excluding irrelevant data from queries by filtering the storage blocks using the information contained in metadata structures. They could include min-max statistics, zone maps, bloom filters, and partition metadata. In addition, the technique of predicate pushdown allows applying filters before the data transfer to the computing engine. As a result, the system will analyze only the necessary part of data, leading to faster execution and minimizing computing and storage costs.

Table 3: Storage Optimization Techniques and Benefits

Technique	Purpose	Expected Benefit
Compression	Reduce storage footprint by eliminating redundancy	Lower storage cost, faster data access, reduced I/O
Partitioning	Divide data into logical segments	Faster query execution through partition elimination
Clustering	Organize data based on frequently queried attributes	Improved data locality and reduced scan operations
Encoding	Store data in compact representations	Better compression ratio and processing efficiency
Data Pruning	Eliminate irrelevant data during query execution	Reduced query latency and lower compute consumption
Predicate Pushdown	Apply filters at storage level	Minimized data transfer and faster query processing
Metadata Indexing	Maintain statistics about stored data	Improved query optimization and scan reduction
Caching	Store frequently accessed data in memory	Faster response time for repetitive workloads
Parallel Processing	Distribute query execution across nodes	Higher throughput and scalability
Automatic Optimization	Continuously adjust storage organization	Consistent performance improvement over time

BENCHMARKING RESULTS

The assessment aims to examine the efficiency of primary techniques for optimizing columnar storage in cloud-native data warehousing solutions. The analysis focuses on the improvement in performance associated with querying, storage, and data compression using techniques that address aggregation, filter, sort, and join operations performed in analytical workloads. The illustrative benchmark framework considers query workloads on a dataset consisting of 500 GB of data stored in cloud-native warehouses, utilizing the same queries and compute resources.

Accordingly, the comparative benchmark results show that optimizing columnar storage through various methods leads to considerable improvements in the processing performance of analytical operations compared to baseline configurations. Data pruning and predicate pushdown result in the highest decrease in latency as these approaches minimize data read operations by eliminating irrelevant data from queries.

Similarly, various compression and encoding techniques lead to significant performance gains as well. Dictionary encoding and run-length encoding help to minimize storage usage while simultaneously improving the speed of querying. As less data is involved in the queries, query engines execute analytical workloads quicker. Moreover, the use of compression helps to lower cloud storage costs and reduce network overhead.

In terms of partitioning, performance benefits can be achieved through partitioning data according to specific criteria. Time series and range-based queries benefit greatly from such an approach since they access only those portions of the database required to generate results. Likewise, the use of clustering leads to performance gains because queries based on frequent filters involve fewer disk operations by accessing relevant information through metadata.

The analysis demonstrates that using storage optimization techniques altogether brings about much better performance than implementing any single technique separately. Thus, using techniques including compression, partitioning, clustering, data pruning, and encoding reduces the latency of queries by over 65% compared to the baseline configuration.

With regards to storage utilization, optimization of cloud-native storage systems leads to the reduction in storage footprints due to efficient data compression and encoding methods. The dictionary approach achieves high levels of compression as values are replaced with short codes. Run-length encoding is effective for compressed sorting operations, whereas delta

Vol. 10, Issue 01, January 2022
ISSN(P) 2347-5404 ISSN(O)2320 771X

encoding is useful in cases where sequential number sequences need to be compressed.

Cloud-native data storage systems achieve not only performance gains by using columnar structures but also additional benefits thanks to effective optimization techniques that minimize storage space and improve performance of queries and data operations. Cloud-native organizations should consider the implementation of these methods to save costs and improve performance significantly.

Table 4: Query Execution Performance

Technique	Avg Query Time (Seconds)	Improvement (%)
Baseline Columnar Storage	12.8	0
Compression	10.4	18.8
Partitioning	8.9	30.5
Clustering	7.8	39.1
Encoding	7.2	43.8
Data Pruning	5.6	56.3
Combined Optimization Approach	4.3	66.4

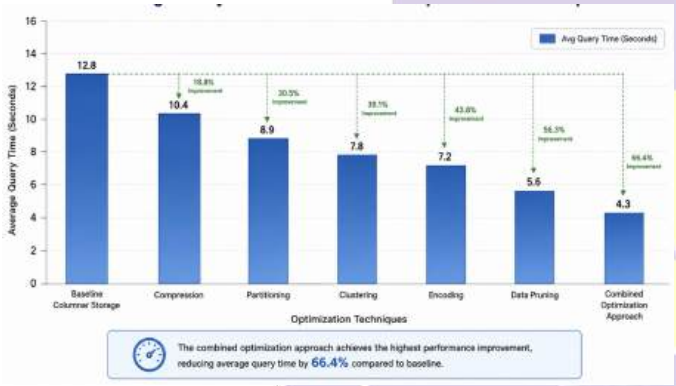

Fig. 3. Average Query Performance Comparison

Table 5: Storage Utilization Results

Technique	Storage Used
Baseline Storage	500 GB
General Compression	310 GB
Dictionary Encoding	270 GB
Run-Length Encoding	245 GB
Delta Encoding	225 GB
Combined Compression and Encoding	190 GB
Full Optimization Configuration	165 GB

The benchmark results clearly demonstrate that storage optimization techniques contribute significantly to both performance enhancement and storage efficiency. Data pruning achieved the highest query performance gains, while combined compression and encoding delivered the greatest storage savings. These findings support the growing adoption of advanced columnar optimization strategies in modern cloud-native analytical data warehouses.

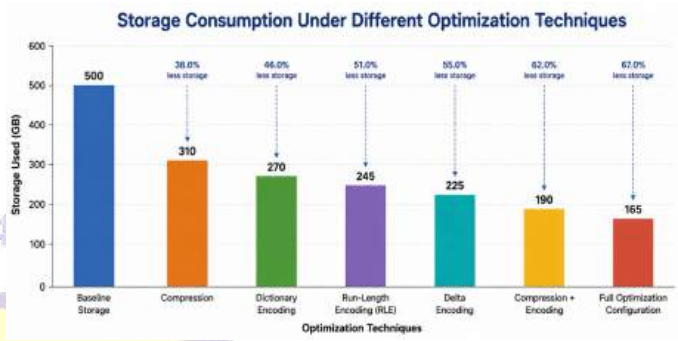

Fig. 4. Storage Utilization Results

Results and Discussion

Firstly, the findings of the benchmarking analysis illustrate that storage optimization techniques such as those related to using columnar format are highly important to ensure high-performance and efficiency of cloud-native data warehouses. Out of different types of data optimization techniques, it is possible to mention data pruning as the one that brings the largest improvements in terms of reducing the time of data processing and minimizing the amount of data scanned. Data pruning refers to removing unnecessary partitions, blocks, and rows before queries are executed.

Secondly, the benchmarking analysis demonstrates that there is an opportunity to reduce storage needs significantly due to the implementation of various optimization methods. For instance, run-length and delta encoding reduced storage requirements by more than 50%, while dictionary encoding achieved a substantial but slightly lower reduction. Lower storage demands imply less payment to cloud service providers and fewer problems with transferring data to a particular environment. In addition, compressed data sets allow saving disk space and improving cache utilization.

However, there are many trade-offs between storage optimization and query performance that need to be discussed. Thus, aggressive compression may result in high compression rates but will require additional resources to decode compressed data. Likewise, clustering and partitioning may help to make queries faster; however, these processes will take time to maintain and optimize. Too many partitions can lead to increased overheads related to metadata management, so

companies need to find a balance between performance improvements and costs.

Furthermore, the research proves that no optimization technique leads to maximum benefits in all the cases. Instead, the most effective approach to optimizing storage is combining different optimization techniques. Thus, integrated optimization solutions including partitioning, clustering, encoding, data pruning, and others performed better when compared with individual techniques applied in isolation. The findings prove that only the adoption of multiple techniques will help to maximize performance.

The results obtained through benchmarking show that organizations should rely on both data pruning and proper partitioning to boost their performance. At the same time, compressions and encodings will be helpful to save on storage expenses. This is a powerful combination that allows improving efficiency, performance, and costs associated with running a cloud-native data warehouse.

Conclusion

The current study analyzes the effectiveness of several approaches to storage optimization in the context of cloud-native warehouses and considers their effect on analytical capabilities and query performance. It becomes clear from benchmarking results that the columnar structure has some significant advantages if combined with advanced storage optimization methods. Specifically, among the methods studied, data pruning appears to be the one that delivers the highest performance gain, whereas compression and encoding have proved highly efficient in terms of storage savings. Partitioning and clustering contribute to further optimization by reducing redundant queries.

Thus, none of the methods described above can maximize performance and storage efficiency alone; however, the combination of compression, partitioning, clustering, encoding, and data pruning can bring about the greatest gains in these areas and significantly accelerate query execution, reduce storage costs, and enhance the overall scalability of the system.

From this analysis, it becomes clear that cloud-native data warehouses require a holistic approach to storage optimization and cannot be considered fully effective using only columnar storage optimization techniques. When it comes to performance-related applications, priority should be given to data pruning and partitioning, while in terms of saving on storage resources, compression and encoding should be used.

Further research can involve applying the described techniques to larger datasets, considering multiple cloud environments, analyzing real-time workloads, and investigating ML-driven storage optimization. The potential benefits of adaptive layouts and automated tuning mechanisms as well as Apache Iceberg and Delta Lake as new open table formats should also be explored.

References

- Stonebraker, M., Abadi, D. J., Batkin, A., Chen, X., Cherniack, M., Ferreira, M., Lau, E., Lin, A., Madden, S., O'Neil, E., O'Neil, P., Rasin, A., Tran, N., & Zdonik, S. (2005). **C-Store: A column-oriented DBMS**. *Proceedings of VLDB*, 553–564.
- Abadi, D. J., Madden, S. R., & Ferreira, M. C. (2006). **Integrating compression and execution in column-oriented database systems**. *Proceedings of SIGMOD*, 671–682. https://doi.org/10.1145/1142473.1142548
- Abadi, D. J., Madden, S. R., & Hachem, N. (2008). **Column-stores vs. row-stores: How different are they really?** *Proceedings of SIGMOD*, 967–980. https://doi.org/10.1145/1376616.1376712
- Abadi, D., Boncz, P., Harizopoulos, S., Idreos, S., & Madden, S. (2013). **The design and implementation of modern column-oriented database systems**. *Foundations and Trends in Databases, 5*(3), 197–280.
- Boncz, P. A., Zukowski, M., & Nes, N. (2005). **MonetDB/X100: Hyper-pipelining query execution**. *CIDR*.
- Floratou, A., Patel, J. M., Shekita, E. J., & Tata, S. (2011). **Column-oriented storage techniques for MapReduce**. *Proceedings of the VLDB Endowment, 4*(7), 419–429. https://doi.org/10.14778/1988776.1988778
- Melnik, S., Gubarev, A., Long, J. J., Romer, G., Shivakumar, S., Tolton, M., & Vassilakis, T. (2010). **Dremel: Interactive analysis of web-scale datasets**. *Proceedings of the VLDB Endowment, 3*(1–2), 330–339. https://doi.org/10.14778/1920841.1920886
- Melnik, S., et al. (2020). **Dremel: A decade of interactive SQL analysis at web scale**. *Proceedings of the VLDB Endowment, 13*(12), 3461–3472.
- Dageville, B., Cruanes, T., Zukowski, M., Antonov, V., Avanes, A., Bock, J., Claybaugh, J., Engovatov, D., Hentschel, M., Huang, J., Lee, A. W., Motivala, A., Munir, A. Q., Pelley, S., Povinec, P., Rahn, G., Triantafyllis, S., & Unterbrunner, P. (2016). **The Snowflake Elastic Data Warehouse**. *Proceedings of SIGMOD*, 215–226. https://doi.org/10.1145/2882903.2903741
- Gupta, A., Agarwal, D., Tan, D., Kulesza, J., Pathak, R., Stefani, S., & Srinivasan, V. (2015). **Amazon Redshift and the case for simpler data warehouses**. *Proceedings of SIGMOD*, 1917–1923. https://doi.org/10.1145/2723372.2742795