

# Architectures of Distributed Storage Systems in the Context of High-Performance Digital Infrastructure in the United States

Terletska Khrystyna<sup>1</sup>

<sup>1</sup>Bachelor's degree, Department of Social Communication and Information Activities, Lviv Polytechnic National University, Lviv, Ukraine

**Abstract**— The article focuses on the analysis of storage architectures in distributed systems used in high-load digital infrastructures in the USA. It examines key technological solutions such as LSM and B-trees, column-oriented storage formats, indexing strategies, and read/write balancing techniques. Special attention is given to hybrid architectures combining local and cloud components, as well as multi-regional replication and erasure coding mechanisms. The study systematizes engineering approaches to designing scalable and fault-tolerant storage systems that meet modern requirements for data availability and performance in a globally interconnected digital environment.

**Keywords**—Distributed storage, cloud architecture, LSM trees, indexing, replication, data coding.

## I. INTRODUCTION

Modern digital infrastructures that operate in the context of rapid data growth and the need for high availability of services place special requirements on storage systems. Distributed data storage systems have become an important component of the architecture of high-performance computing platforms. The rapid evolution of cloud and in-memory solutions encourages a rethink of traditional storage models in favor of hybrid and adaptive architectures.

The aim of this study is to systematically analyze the architectural solutions, indexing strategies, and trade-offs that underlie modern distributed storage systems used in high-load digital systems. Special attention is paid to considering ways to balance the performance of read and write operations, storage efficiency, and scalability in a multicenter infrastructure. The research focuses on identifying theoretical patterns and engineering approaches that contribute to optimizing the operation of distributed storage in a cloud-based environment.

This paper discusses key architectural storage models, including LSM and B-trees, SSTables and column formats, as well as high-throughput indexing strategies. Methods of data compression and filtering, their impact on system performance, as well as the principles of multi-regional replication and remote coding are studied. This study allows us to generalize existing approaches and formulate recommendations for building efficient storage architectures in distributed computing.

## II. THE MAIN PART. STORAGE MODELS AND ARCHITECTURAL-LEVEL TRADEOFFS

Distributed storage systems are complex multi-layered architectures, each with its own unique characteristics and applicability depending on the type of workloads and performance requirements. One of the most common models is the use of LSM-trees (Log-Structured Merge Trees), which allows you to efficiently process records in conditions of intensive write operations. This model is used in many modern

distributed storage systems, such as RocksDB and Apache Cassandra, which allows you to achieve significant improvements in throughput for workloads with high write intensity, while minimizing read delays. However, its weak side is the need for frequent data merging, which can cause additional costs for disk I/O operations [1].

The opposite of this architecture is B-trees, which are often used in traditional relational databases and provide high efficiency when processing queries with a high read frequency. In systems such as MySQL or PostgreSQL, B-trees allow you to maintain a balance between read and write speeds, minimizing the need for data merging. However, under high loads associated with large amounts of data and frequent updates, B-trees can face problems associated with the high cost of write operations, which makes them less suitable for some types of high-performance systems [2].

Care needs to be exercised with those architectures that rely on columnar storage schemes, such as in data storage systems like Apache HBase or Google Bigtable. These storage schemes allow for working with massive amounts of data, supporting high read rates and aggregation efficiently across columns. For example, in US cloud services such as Google Cloud Bigtable, columnar data storage model enables the efficient processing of queries with large record counts and aggregation by key fields, which is significant for applications that require fast analytics and real-time data processing.

Each of these architectures has its advantages and disadvantages, and the choice between them is a matter of the specifics of use. The trade-off is that have to sacrifice write and read performance, system scalability, and the possible costs of operational processes such as data merging and storage. The final choice of architecture is a matter of the type of application and anticipated loads, which in turn determine the design of the storage system.

## III. INDEXING STRATEGIES AND IMPACT ON OPERATION DELAYS

The efficiency of a distributed data storage system largely depends on the chosen indexing strategy, as it has a direct

impact on the latency of read and write operations. One of the most common indexing methods is the use of B+ trees, which provide a balance between write performance and search performance (fig. 1).

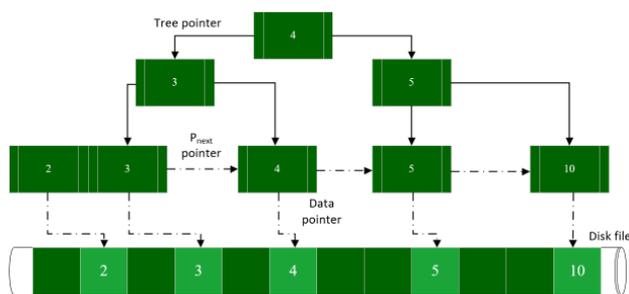


Fig. 1. B+ tree example [2]

Like a B-tree, a B+ tree has a fixed number of keys in each node and is balanced so that all leaf nodes are at the same level. But in B+ trees, all keys are stored in leaf nodes, and internal nodes only hold pointers to other nodes. They are sorted-key data structures that allow you to perform search, insert, and delete operations in time  $O(\log N)$ , where  $N$  is the number of elements in the tree. This strategy is applied actively in various databases, for example, Apache HBase, where B+ trees provide the capability for efficient querying under high load conditions with fast access to the data stored in a distributed system.

To minimize latency and optimize throughput when handling enormous amounts of data, the architecture of high-speed interfaces is an important consideration. The use of such architectures accelerates data transfer among nodes in distributed systems, thus significantly reducing the response time when executing indexing and search queries [3]. It is worth noting that high-speed interfaces can become crucial in ensuring effective interaction of system components, which is necessary for real-time data processing and ensuring high performance at high loads.

Another valuable indexing method is Skip Lists, which are a less complex and more general alternative to B+ trees. They provide similar search performance, but with lower overhead in maintaining the data structure. This makes them an appealing choice for high-throughput, low-latency systems like ScyllaDB, which uses skip lists to perform well under large datasets. Unlike B+ trees, skip lists allow you to adapt faster to data structure changes, which could be crucial in dynamic systems, for example, when working with data on US cloud platforms such as Amazon DynamoDB [4].

Adaptive Radical Trees (ART) can be suitable for systems with dense queries that require fast insertion and search. These trees are more helpful than standard search trees because they change their shape dynamically according to the load. ART trees can significantly reduce search time, which is especially important for massive distributed databases such as Google Bigtable, where the read and write speed is an important consideration in managing millions of requests per second.

Each indexing strategy has its own trade-offs and attributes that must be examined while designing the system. Directly dependent upon the data structure, operation frequencies, and data amounts is the impact on the operation delays, hence

requiring careful choice of the appropriate indexing architecture in relation to the special load and storage requirements.

#### IV. BALANCING READ AND WRITE OPERATIONS

One of the major tasks while creating distributed storage systems is balancing read and write operations, which is essential for ensuring high performance under different workloads. To balance them, a variety of I / O optimization techniques are widely used, including the use of caching. It can significantly reduce the number of read operations by providing quick access to recently used data. In databases such as Apache Cassandra and ScyllaDB, data is often cached in RAM, which avoids normal disk access and reduces read latency [5].

In addition, caching-aware indexing plays an important role in write-intensive operations. In systems focused on high write speeds, such as RocksDB, indexing is often performed in a multi-layered structure, where data is first written to memory and then asynchronously written to disk. This allows you to optimize writes and minimize their impact on read performance. However, the choice between incremental and full data merges also has a significant impact on the latency of operations. Using incremental merges reduces the response time, but requires additional load on the system, while full merges provide a cleaner data structure, but can cause load peaks during their execution [6].

Special attention should be paid the compression mechanism to data compression mechanisms, which can also significantly affect the balancing of operations. For example, cloud storage systems such as Amazon S3 use compression based on algorithms such as Zlib, which reduces the amount of data stored and speeds up the transfer of information between nodes. However, using compression can slow down write operations, as each new block of data requires additional calculations to compress before saving. Therefore, the correct compression setting, depending on the type of data and the specific load, plays a key role in optimizing the entire system.

#### V. HYBRID AND CLOUD SOLUTIONS: NEW ARCHITECTURAL PARADIGMS

In recent years, there has been a significant increase in interest in hybrid storage architectures that combine the advantages of on-premises systems with cloud solutions. This approach allows organizations to optimize their work with large amounts of data, providing both high performance and flexibility in resource management (table 1).

One of the most prominent examples of hybrid architectures is the integration of in-memory solutions such as RocksDB with cloud object storage systems such as Amazon S3 or Google Cloud Storage (GCS). This allows you to achieve high-speed data access while simultaneously scaling storage in the cloud, depending on business needs [7].

Hybrid architectures also include the use of local data warehouses for rapid access and real-time processing, with cloud storage being utilized for long-term storage and worldwide availability of data. This hybrid is widely used in systems such as ScyllaDB, where memory usage is optimized for rapid writes and reads, and cloud storage scalability enables

it to easily adapt to changing data volumes and fault tolerance requirements.

TABLE I. Risks and benefits of hybrid architecture

Category	Benefits	Potential Risks
Security	Ability to store and process sensitive data within the local perimeter	Security system complexity increases when integrating with the public cloud
Performance	Use of local resources for critical operations reduces latency	Uneven load distribution may require advanced balancing mechanisms
Cost	Flexible resource allocation reduces infrastructure and licensing expenses	Increased costs related to support and synchronization between components
Flexibility	Quick scalability and adaptability to changing business requirements	Higher architectural complexity requires advanced IT skills
Management	Centralized control over environments and policies	Need for additional orchestration and monitoring tools

Including cloud solutions within on-premises storage technology decreases infrastructure cost and complexity management. Clouds like Amazon DynamoDB provide high availability of data and automatic replication along with auto-scaling, which is vitally essential for enabling smooth run during high load scenarios like in case of transactional processing systems or real-time analysis. With such situations, cloud-based solutions can be flexed to meet needs by either expanding or contracting resources according to immediate needs, while on-premises systems take care of low latency and high performance upon processing mission-critical data.

#### VI. DATA DISTRIBUTION, REPLICATION, AND ENCODING

Distributed storage systems must possess good replication and coding mechanisms in order to attain fault tolerance, availability, and consistency of data for multi-regional architectures. Replication of data allows you to duplicate data over different nodes, hence reducing the likelihood of data loss due to failure of individual parts of the system. One of the key aspects of replication is consistency vs. availability tradeoff: using the concept of quorum, as in the Amazon DynamoDB system, you can provide high data availability, at the cost of strong consistency, resulting in a fault-tolerant and scalable system.

Also, replication not only increases availability, but it also turns out to be essential for speeding up read operations since requests can be satisfied from multiple replicas, thereby spreading the load across nodes. Large cloud providers like Google Cloud Storage use multi-region replication for global data availability, lower latency, and better fault tolerance. Moreover, special concern is given to technologies such as blockchain that are utilized in managing IT infrastructure so as to augment the level of security. While typical replication technology does not deploy the decentralized approach to storing information, wherein copies of data (or «blocks») are not changeable and are cryptography-encoded after writing, blockchain utilizes just that. This can actually improve the level of data security and guarantee their immutability in distributed computing [8].

But duplicating data requires more resources, and therefore storage costs more. As a result, storage is reduced by employing deletion encoding methods like Reed-Solomon coding. It enables you to make efficient use of network bandwidth and reduce the storage of copies for data without affecting their integrity and recoverability upon failure. This process is utilized widely within cloud systems for long-term storage of data with minimal replication cost.

The application of replication and drop-coding to distributed data warehouses is able to effectively improve their reliability and responsiveness using very low latency and high availability, which becomes especially important to high-traffic applications such as real-time transaction processing systems and analytics.

#### VII. CONCLUSION

Distributed storage systems keep evolving as they adapt to increasing pressure on fault tolerance, scalability, and performance. The availability of various architectural solutions – such as LSM trees, B+ trees, columnar storage formats, and hybrid approaches combining in-memory technologies with cloud storage – enables the customization of read and write operations as well as the optimization of resource utilization. Indexing, compression, and balancing of operations techniques impacts have become necessary in delivering high performance to various workloads.

Replication and coding techniques, for instance, multi-region replication and deletion coding, provide effective means of ensuring data availability and integrity. Such techniques and new architectural paradigms like cloud solutions create new opportunities for the design of high-performance, scalable and cost-effective storage systems. In the future, research must persist in the exploration of the trade-offs between the different models and resource optimization for efficient data upkeep in the paradigm of global distributed computing environments.

#### REFERENCES

- [1] J. Liu, F. Wang, D. Mo, S. Luo. “Structural Designs Meet Optimality: Exploring Optimized LSM-tree Structures in A Colossal Configuration Space,” *Proceedings of the ACM on Management of Data*, vol. 2, issue 3, pp. 1–26, 2024.
- [2] What is B+ Tree meaning / Geeks for Geeks // URL: <https://www.geeksforgeeks.org/what-is-b-plus-tree-b-plus-tree-meaning/> (date of access: 27.03.2025)
- [3] R. Garifullin. “Development and implementation of high-speed frontend architectures for complex enterprise systems,” *Cold Science*, issue 12, pp. 56–63, 2024.
- [4] R. Sajid, M. Gohar, H.Z. Zaidi, Z. Mubeen. “Securing DynamoDB: In-Depth Exploration of Approaches, Overcoming Challenges, and Implementing Best Practices for Robust Data Protection,” *Journal of Computing & Biomedical Informatics*, vol. 7, issue 2, 2024.
- [5] R. Aoyagi, K. Takahashi, Y. Shimomura, H. Takizawa. “Combining Lossy Compression with Multi-Level Caching for Data Staging over Network,” *In 2024 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pp. 212–221, 2024.
- [6] W. Xu, Y. Qiu, C. Jiang, J. Liu, X. Wang, S. Chen. “Merging and cooperative caching based HDFS small files access performance optimization,” *In 2024 4th International Conference on Communication Technology and Information Technology (ICCTIT)*, pp. 211–219, 2024.
- [7] S. Pandey, A. Basu. “H-Rocks: CPU-GPU accelerated Heterogeneous RocksDB on Persistent Memory,” *Proceedings of the ACM on Management of Data*, vol. 3, issue 1, pp. 1–28, 2025.

- [8] A. Dudak, A. Israfilov. “Application of blockchain in IT infrastructure management: new opportunities for security assurance,” *German International Journal of Modern Science*, issue 92, pp. 103–107, 2024.