

Optimized Database Sharding Techniques for High-Performance MySQL Applications

Rishabh Agarwal*

Harrisburg University of Science and Technology, Pennsylvania-USA

* Corresponding Author Email: rishabh.agarwal1124@gmail.com - ORCID: 0000-0002-5777-7850

Article Info:

DOI: 10.22399/ijcesen.4340

Received : 05 February 2025

Revised : 25 March 2025

Accepted : 30 March 2025

Keywords

Database Sharding;
MySQL Optimization;
Distributed Systems;
Query Routing;
Scalability

Abstract:

With the ever-growing data-intensive applications, the classical monolithic MySQL databases are usually unable to satisfy the requirements of the high-throughput, low-latency, and real-time applications. The technique of horizontally dividing data in more than one database is called database sharding, though it has proven to be a strong tool to overcome these problems. The paper provides an extensive overview of optimized forms of database sharding that are specific to MySQL applications. It discusses the basic Sharding concepts, comparisons of various models, and advanced optimization techniques, i.e., consistent hashing, query-aware routing, dynamically re-sharding, and caching. The paper also introduces patterns of deployment of architecture that can be adapted to the cloud-native environment and provides comprehensive results of performance benchmarking that can be used to measure the benefits and drawbacks of diverse approaches. The paper will offer a roadmap to help the architects and engineers create scalable, reliable, and high-performance MySQL infrastructures by combining the existing best practices with scholarly work. The results highlight that sharding used and tuned properly can be of major benefit to a database in terms of performance, operational efficiency, and scalability of the system.

1. Introduction

The unrelenting increase in the volume of data and the emerging need to process larger volumes of data in real-time have prompted the switch to a distributed system to handle massive databases. Database sharding is one of the most efficient architectural strategies that can be used to scale a database system, especially when using MySQL to power the system. Sharding is the horizontal division of data into two or more database instances, and thus, it improves performance, minimizes query response times, and thereby leads to effective use of resources. Sharding of the databases is no longer a hypothetical design pattern in the current cloud-native and microservices-based application patterns, but an operational requirement it is to maintain the scalability and responsiveness of the systems [1][2].

As one of the most widely used relational database management systems in the world, MySQL has various deployment configurations that allow it to be used with sharding implementation. Nonetheless, MySQL is not a sharded database as

such; multiple methods and third-party applications have been developed to provide optimized sharding. These are manual, hash-based, range-based, directory, and proxy sharding, hash-based routing, and proxy solutions such as Vitess and ProxySQL [3][4]. The choice of the right sharding method depends on many factors, and this depends on the distribution of data, the nature of the application workload, the complexity of some queries, and fault tolerance. The advantages of sharding MySQL databases in efficiency and performance are noticeable in large-scale applications (e.g., e-commerce applications, social networks, content delivery networks, and IoT systems, etc.) where a single database can easily turn into a bottleneck. Without sharding strategies that are optimized, such systems are susceptible to database contention, replication lag, and disk I/O saturation. Moreover, ineffective sharding may bring about hotspots of data, uneven load distribution, and complicated application logic, which undermine the anticipated scalability benefits [5][6]. The paper will seek to offer an in-depth discussion of database sharding techniques that are

optimized to suit high-performance MySQL applications. It shall start by expounding on the principles of database sharding and why it is a requirement in the current database architecture. It will then proceed to discuss some of the sharding models as well as their strengths and weaknesses in MySQL settings. The following strategies of advanced optimization (such as dynamic re-sharding, consistent hashing, and query-aware routing) will be studied, as well as real-life deployment cases and performance metrics. The following section examines the conceptual foundations and evolution of database sharding. It also introduces the challenges associated with implementing sharded MySQL databases and highlights the importance of selecting an appropriate sharding strategy. This background knowledge is crucial in placing the techniques of optimization that will be addressed in the subsequent articles. Influenced by the demand to maintain high-throughput transactional processing, particularly in areas like finance, healthcare, and e-commerce, it is important to learn and implement advanced database sharding techniques in order to develop resilient and scalable database infrastructures [7-10].

2. Conceptual Foundations of Database Sharding

Since sharding has been proven as relevant and crucial in the context of high-performance database environments, we must now know the conceptual background of sharding. Sharding is essentially a horizontal partitioning method (or technique) in which every shard or partition stores a portion of the entire dataset and functions as a separate database. Contrary to vertical partitioning, where the tables are partitioned based on the columns, the horizontal partitioning of tables is based on the rows, so that each shard can be queried separately or simultaneously based on how the application was designed and the query nature [11][12]. Sharding dates back to the era of early distributed database systems, in which scaling out as opposed to scaling up became the new architectural philosophy. The conventional monolithic database design grew less and less capable of satisfying the requirements of scalability, resilience, and performance as the web-scale systems appeared, especially in the late 2000s. Sharding solved this issue since data storage and processing were decentralized, and thus, the load is spread over many servers. This decentralization not only increases the system throughput but also isolates faults since failure in one shard does not normally impact the rest [13][14]. With MySQL, common forms of sharding are external to the

database engine. As MySQL does not support sharding, the applications or developers have to apply application logic or external tools and middleware to handle the sharding scheme. This adds a level of complexity, and yet it provides flexibility when it comes to the partitioning and handling of data. As an example, developers are able to shard by user ID, geographical location, or even the application domains. All strategies imply the efficiency of the queries, data consistency, and overhead of the operations [15][16].

The determination of the sharding key can be considered one of the most basic problems in creating a sharding architecture. The sharding key dictates how information is shared by the shards and is important in balancing the load, and is also used to route queries efficiently. The inappropriate selection of the sharding key may cause hotspots, i.e., some shards will be overused, whereas others will be underutilized, and this will nullify the performance gains of sharding. This risk is addressed by advanced optimization methods, including consistent hashing and adaptive partitioning [17][18]. Additionally, there are no trade-offs in the absence of database sharding. Although it provides horizontal scalability and enhanced fault tolerance, it also poses some complexities in cross-shard queries, referential integrity, backup and recovery, and transactional consistency. In most instances, the distributed transactions among the shards are not attempted at all, or handled via complicated two-phase commit protocols, which may bring about latency and lower system throughput. This is the reason why sharding implementations need to be optimized and designed cautiously, particularly in applications where MySQL performance is critical [19][20].

These background concepts make us understand why optimized sharding strategies that are unique to MySQL characteristics and constraints are required. Introducing a variety of models and techniques used in sharding MySQL databases, the next section will provide a comparative analysis, which will be used as the foundation of more complicated optimizations later in this paper. Upon further discussion of the initial concepts of database sharding, the second section of the present paper discusses the practical implementations and MySQL-specific models in greater detail. This part continues the conceptual clarification of sharding and is where the strategies are applied to the real-world MySQL scenarios.

3. MySQL-Specific Sharding Techniques

The absence of an in-built sharding in MySQL requires external solutions to drive the process of

data partitioning between two or more database servers. This paves the way to the different possible implementation strategies, each possessing its distinct merits and demerits. Application architecture, anticipated data growth, read/write traffic distribution, and fault tolerance are some of the factors that usually affect the choice of a sharding technique in MySQL.

Manual or application-level sharding is one of the most common ones, in which the logic to select a shard is held in the application, which dictates which shard to ask based on the sharding key. The method is the most flexible and controllable because the developers can formulate their own sharding methods depending on business logic. Nonetheless, it also implies that cross-shard query logic, failover management, and connection pooling functionality traditionally abstracted by the database engine itself should be supported by the application as well [21][22]. More structured, deterministic range-based sharding, in which the data is divided into ranges that overlap each other, with respect to a particular column (e.g., user ID or time) used to partition the data. It is an intuitive model that can be used to easily perform queries when the distribution of data is relatively homogenous. It has, however, the issue of skewed distribution of data or hot spots when some ranges have a lot more records as compared to others. This may cause an inequality in the loading of shards and eventually decline the performance of the system [23][24].

This imbalance can be resolved by using hash-based sharding. In this model, the sharding key value is transmitted using a hash function, and the resulting value is the shard where the data is going to be stored. Sharding by Hashing offers a more even distribution of data, therefore eliminating any hot spots and enhancing load balancing. It, however, complicates range queries, where data is not stored in a sequential order across shards, and can be queried many times, needs to be queried, and data combined at the application layer [25][26]. Another more sophisticated approach is directory-based sharding, whereby a central configuration database service or service contains a mapping of keys to shards. The application has access to this directory to identify the right shard to be used by each operation. This introduces a degree of indirection such that the dynamic reassigning of data between shards can be done without changing application logic. Nevertheless, the directory service is also a potential failure point as well as a performance bottleneck when it is not managed and replicated correctly [27].

Control Proxy-based sharding has become popular in recent years because it removes the complexity

of sharding from the application. Applications such as Vitess and ProxySQL are placed between the application server and the MySQL servers and intercept the queries and redirect them to the appropriate shard depending on the configured rules. Queries can also be rewritten with these tools; load balancing and failover are also supported, which makes managing a sharded structure easier. Specifically, Vitess has become a popular choice in cloud-native setups since it can be integrated with Kubernetes and is able to scale horizontally [28][29]. The other aspect of the sharding strategies is whether there is a homogeneous data schema across shards or it is not. In a homogeneous sharding design, all shards have a common schema, making it easier to write application logic and enabling them to have the same query patterns. By contrast, heterogeneous sharding applies another schema in another shard, which can be required in case of a multi-tenant system or application with heterogeneous data needs. Heterogeneous sharding is also complicated to query plan, data migration, and provide consistency between schema variants [30]. In MySQL, sharding should also be implemented with due attention to the data consistency and transactional guarantees. Although eventual consistency may be acceptable in some applications, some may have strong consistency requirements, especially in financial or healthcare systems. Here, it is required to introduce distributed transactions between the shards, which is typically based on two-phase commit protocols. This, however, comes with overhead and latency that may cancel out the performance advantages of sharding. Thus, optimization is usually a trade-off between performance and consistency. An important factor that must be taken into consideration in MySQL sharding implementations is schema design. In many cases, the denormalization can be used to minimize the use of joins between tables that can be on different shards. Also, indexing strategies need to be scaled to sharded environments in order to circumvent performance penalties. Such secondary indexes as global secondary indexes can even be inefficient or even impossible in a sharded system unless dedicated attention is made, or they are complemented with caching layers. As can be seen, the different sharding methods covered herein depict the range of alternatives that the programmer has when using MySQL. In some situations, each method has its own benefits and difficulties over the others. These difficulties can be overcome and the advantages of sharding maximised only with the help of sophisticated optimization techniques. These strategies do not just apply the fundamental

logic of sharding but extend to dynamic redistribution of data, query routing, and predictive load balancing. The following section will discuss these high-tech methods and their application to MySQL environments to maximize their performance.

Coupled with the qualitative examination of the strategy of sharding in MySQL, it is worthwhile to take a comparative review of their fundamental features, the complexity of implementation, and the common applications. The following table compiles these aspects to make it easier to select the most suitable strategy depending on the requirements of the system.

This comparative framework emphasizes the ways in which each of the models can be applied to various requirements of operations and architecture, which confirms the necessity of the context-driven application of MySQL sharding. The following section on optimization strategies will build on these models by examining how they can be maximized in their efficiency and resilience by dynamically adopting techniques.

4. Advanced Optimization Strategies

Expanding on the MySQL-specific sharing models already mentioned, this section discusses the more advanced optimization techniques that can be used to improve the performance, scalability, and fault tolerance in a sharded environment. These plans play a crucial role, particularly in a large-scale system where workloads are not predictable, data is huge, and they need real-time responsiveness.

Consistent hashing is one of the best optimization methods, and it strives to reduce the data movement in the case of re-sharding operations. The weakness of traditional hash-based sharding is that adding or removing a shard requires redistributing a significant part of the data. Consistent hashing does this by hashing both keys to shards and shards to a circular hash space, such that a small number of data need to be relocated when the shard structure is altered. It is a technique that enhances the elasticity of the system and is especially useful in a cloud-native system where nodes are often scaled in and out [1][5].

The other important optimization is dynamic re-sharding, which allows the system to be adjusted to the evolving patterns of data distribution and queries. Rather than using a fixed partitioning policy, dynamic re-sharding uses shard utility indicators like disk space and query volume, as well as CPU usage, to anticipate and process when and how to redistribute data. Trying to keep the balance and performance, automated re-sharding frameworks rely on predictive algorithms and re-

shard the hot shards and merge the underutilized ones. Such tools as Vitess may facilitate non-downtime online re-sharding, which is feasible in production systems [3][19]. Another performance optimization is query-aware routing to direct the queries to the appropriate shard(s) with the least overhead. In contrast to the traditional routing that uses only the sharding key, query-aware routing reads the SQL query to identify which shards are being used in it and optimizes routing based on this information. This is particularly very handy in scenarios of multi-shard joins or aggregation queries, where running the query on the unnecessary shards would cause unnecessary load and latency. SQL proxy layers or a custom middleware are used to implement query-aware routing [12][17]. Most high-performance systems use caching strategies to minimize query latency that supplement sharding. The system is able to compute and offload the read requests to the underlying shards by caching commonly accessed data at the proxy level or at the application level, thereby lowering I/O and decreasing response times. In sharded systems, invalidation of cache is complicated, with any changes made to a single shard being propagated to the cache. The mechanisms, such as write-through caching, distributed cache invalidation protocol, are used to ensure consistency [6][14]. Another optimization issue in sharded databases is global secondary indexing. Monolithic database in a monolithic database, secondary indexes allow quick searches on non-primary keys. In sharded environments, maintaining such indexes across multiple shards can introduce significant overhead. To address this, each shard can maintain local indexes, supported by a central index directory that identifies the location of all shards. Alternatively, query patterns may be optimized to avoid non-primary key searches or be enabled by the use of denormalization and pre-aggregated views [20][25]. Parallel query execution is an essential optimization in sharded environments. Since data is distributed across multiple shards, queries can be processed simultaneously, and their results are consolidated at the application or middleware level. The model is much better in terms of query throughput and responsiveness, especially with workloads that are analytical in nature. Nevertheless, it does need advanced-coordination schemes to deal with result merging, pagination, and ranking of results across multiple sources of data [8][15]. To ensure a high availability and reduce the effects of shard failures, a large number of systems have replication and failover policies in every shard. The system is able to restore itself within a short time in case of a hardware failure or the failure of a software through

replicating every shard into a standby and using an automatic failover system. Replication can be applied either with asynchronous replication, as the one provided by MySQL, or with semi-synchronous replication with much stronger consistency. Such arrangements are also supplemented with monitoring services that identify abnormalities and invoke failover without human intervention [22][24].

Lastly, sharded environments need observability and monitoring in order to optimize them. The performance indicators like query latency, shard utilization, replication lag, and cache hit rates should be constantly observed to identify any anomalies and maximize their performance. The observability platforms like Prometheus, Grafana, and open telemetry have the potential to be integrated to ensure the system is performant and reliable when faced with variable loads through real-time dashboards and alerting mechanisms [10][18]. These optimization tricks, regular hashing, ad hoc re-sharding, query-conscious routing, smart caching, and reliable monitoring are the pillars of MySQL sharding architecture. These methods allow systems to scale predictively, adapt to changing workloads, and offer predictable and consistent low-latency access to data. Since the strategies are implemented in the real-life setting, they can be learnt and can be informative, which will be discussed in the next section with the help of case studies and implementation scenarios.

5. Architectural Patterns and Deployment Topologies for MySQL Sharding

After considering the discussion on high-level optimization strategies, it is important to know how the techniques are applied to large system architectures. The process of optimized sharding in the MySQL database is not only a domain of database-level decisions but is also a part of the infrastructure design, service coordination, and life cycle management of data. In this section, the architectural patterns and deployment topologies that enable scalable and resilient settings of sharded MySQL are discussed.

The most widely used architectural design of MySQL sharding is the shard-per-tenant architecture used in multi-tenant SaaS web applications. Under this architecture, every customer or logical tenant gets its own shard. This segregates workloads, and data governance and compliance are made easier, particularly for clients with certain regulatory needs. It also allows one to scale resources per tenant and eases backup, recovery, and archiving. Nonetheless, this model may lead to resource fragmentation when the

pattern of usage of tenants is very dynamic [1][3]. Conversely, the shared-shard model is a model in which tenants or logical data-partitions are replicated on shared shards. This model is more effective when it comes to hardware usage, but requires complex tenant management and resource throttling to prevent contention. Sharing shared environments also has the advantage of the intelligent routing layers and load-balancing proxies that are capable of dynamically and dynamically allocating queries and balancing the consumption of resources across nodes [7][11]. Cloud-native architectures today are increasingly using the sharded topology, which is a microservice-based topology, where each microservice is responsible for its shard or collection of shards, and is often independently operated. This model encourages bounded context and loose coupling, which are the two guiding principles of domain-driven design. The microservices can use a different MySQL cluster or instance, and the sharding logic can be based on the data access patterns of the service. This architecture fits the container orchestration systems such as Kubernetes, where services can be scaled independently using metrics like query throughput or CPU utilization [5][13].

The other common design trend is the geo-distributed sharded architecture, which is used in the interest of globally distributed applications, which need low-latency access to various regions. Under this configuration, the shards exist in data centers that are close to the user base, minimizing round-trip time as well as enhancing responsiveness. Geo-sharding is usually a combination of range-based partitioning and region-aware routing logic, which enables user requests to be served out of the nearest data center. Such an arrangement, though, has to overcome such challenges as inter-region replication, adherence to data sovereignty legislation, and cross-region failover planning [4][12].

A hybrid sharding model is frequently used in a high-throughput situation. In this architecture, several sharding plans are mixed in one application. An example is by sharding transactional data with a hash-based sharding scheme and analytical workload with a range-based sharding scheme to allow time-series queries. Hybrid models can be used to optimize across a wide range of workloads, but demand complicated orchestration and data handling logic at an application or middleware level [8][14]. Architectures built on proxies are now an essential part of sharding in the present day. ProxySQL, MaxScale, and Vitess also add a layer between the application and shards of the database, which allows dynamic routing of queries, load

balancing, pooling connections, and caching queries. These proxies usually support query-aware routing and the extraction of a sharding key dynamically with the help of SQL parsers. An example is that, given a SQL query, Vitess can automatically send it to the right shard based on the router rules, based on the primary key or table routing rule, and also supports operations like online resharding and traffic switching [10][15]. Operationally, most sharded deployments include control planes and orchestration layers to handle lifecycle events like provisioning, scaling, re-sharding, and backup. These control planes can interact with APIs that the database infrastructure and orchestration tools (e.g., Kubernetes operators, Terraform modules) expose to automate common operations. They also apply the sharding policies, track the health metrics, and initiate the recovery workflow in the event of the shard failure or replication lag. This automation is of great benefit in alleviating the operational workload of operating sharded systems at scale [9][16].

The sharded MySQL storage topologies differ greatly with respect to the workload requirements. Shared-nothing architecture is a common type, in which every shard runs in its own isolated hardware or containers that have an independent CPU, memory, and storage. This model does not have contention, and it enhances fault isolation. As an alternative, shared-disk architectures are less prevalent, but may be needed within an environment that needs fast failover and shared caching, but they add complexity to lock management and concurrency control [18][21]. In other advanced designs, the middleware abstraction layers and data federation layers are utilized to make the application think in terms of a single schema, although the data might be sharded. These layers hide the complexity of multi-shard joins and aggregations, as well as the transactions, so that all a developer does is write a query without the knowledge of how the data is distributed. Nonetheless, they have the disadvantage of higher latency and lower clarity of query execution, making optimization work more difficult [17][22]. Security and compliance issues are also important factors in the design of architecture. Sharded environments should also ensure that the access controls are always applied to all shards, as well as that audit logs, security, standardized encryption, and authentication are also known to be applied consistently. There are also multi-region deployments, where in specific jurisdictions, there will be data privacy that will have to meet specific data privacy requirements, which will require data localization as well as planning data flows among the shards [6][19]. Lastly, the choice of deployment

platforms, i.e., on-premises, cloud-native (i.e., AWS Aurora, GCP Cloud SQL), or hybrid, also affects the sharding architecture. Cloud-native technologies are elastic and managed services that provide easy shard deployment. A few of such services, such as Amazon RDS Proxy and Google Cloud Spanner, have the ability to abstract part of the sharding logic, but may limit customization. To create a balance between control and operational overhead, hybrid environments need more subtle designs [20][23].

Knowledge and application of these architecture patterns and deployment topology are necessary to apply optimized sharding strategies into practice. These blueprints help engineers to come up with systems that are scalable, fault-tolerant, and responsive to fluctuating workloads and business constraints. The effect of these topologies must be quantified as these topologies are deployed and refined, and this is the subject of the next section. It is also relevant to the discussion of the architectural patterns and deployment topologies of MySQL sharding to continue by evaluating their performance in practice. This entails the measurement of the benefits and the possible limitations through controlled testing and real-life measurements. The subsequent section gives a detailed discussion of the performance evaluation and benchmarking of the sharded MySQL systems.

6. Performance Evaluation and Benchmarking

The integration of the results of sharded MySQL deployments is essential to confirm the architectural decisions, to make sure that the service-level objectives (SLOs) have been achieved, and to optimize the results. Performance benchmarking entails modelling different workload conditions to measure system performance in throughput and latency, fault tolerance, and resource use. It also assists in revealing the bottlenecks that may not be visible when operating in normal conditions. Benchmarking is a diagnostic tool, as well as a design validation tool, in the context of optimized sharding, which is used to inform capacity planning, infrastructure provisioning, and workload distribution strategies [1][5]. The basis of a strong benchmarking strategy is the choice of representative workloads. These workloads ought to reflect the production traffic patterns, such as read/write ratios, query complexity, transaction sizes, and concurrency. Popular synthetic benchmarking tools that are used to test MySQL systems include SysBench, OLTPBench, and HammerDB. They can be configured to include the type of query (SELECT,

INSERT, UPDATE, and DELETE) to be used, the number of users at a time, and the length of tests to simulate performance over different dimensions [3][7].

Benchmarking in sharded environments has to take into consideration intra- and inter-shard performance. Intra-shard performance can be used to determine the efficiency of each of the shards, and inter-shard performance can be used to determine the overhead involved in the distribution, routing, and aggregation of data between shards. This consists of the amount of money to implement distributed transactions, sharding data, and synchronizing secondary indexes. Interesting metrics are transaction throughput (TPS), mean and 95th percentile response time, CPU and memory per shard, disk I/O, and replication lag [4][6]. Another important aspect of sharded performance is the performance of query routing. Routing logic may result in unnecessary broadcasting of a query to two or more shards due to poor routing logic, which results in more load and latency. Routing should be optimized to ensure that every query is sent to the least number of shards, depending on the index metadata or the sharding key. Instrumentation must be part of benchmarking to measure routing accuracy and its effect on system responsiveness in general. There are proxy-based routing layers, such as Vitess and ProxySQL, which provide query routing diagnostics telemetry built in [9][11].

Performance evaluation is also of importance in re-sharding operations, whether planned or dynamic. Real-time re-sharding can be benchmarked to understand the capacity of the system to remain available and perform in the case of the redistribution of data. Some of the parameters of interest here are the query latency when re-sharding, the speed of migration, the error rate in the system, and lags in consistency. Online re-sharding tools like Vitess can do these operations with minimal interruption, but the usefulness of these tools should be proven by stress testing them in concurrent query loads [2][13]. The other benchmarking dimension is the replication performance, particularly in any sharded system where every shard often has at least one replica to maintain high availability in the event of failure. Benchmarking replication, so far as there is monitoring of lag times between primary and secondary node, replication throughput, and how replication affects write latency. Systems that have high write volumes should be in a position to ensure replication is up to date without interfering with the real-time write and read operations. The benchmark scenarios must incorporate the failover tests to determine how node failure affects the system, and how long it takes to restore it to its full

operation [10][16]. Another performance parameter that has to be measured in sharded systems is caching effectiveness. At the application layer, proxy layer, or database layer, caching strategies must be tested on the basis of hit ratios, invalidation efficiency, and improvement of read latency. Benchmarking is able to measure the degree to which query traffic is redirected to caches, and the impact of this on individual shard load. Broken or old caches may either slow down performance or cause issues with consistency, so their response to load conditions should be studied in some depth [14][18]. Parallel execution performance becomes of primary concern in the case of analytical queries that cross shards. Benchmarks need to be used to measure the performance of the system in carrying out distributed queries, aggregating data, and combining results across shards. The metrics are execution time on a per-shard basis, cumulative aggregation time, and the overhead of the query planner. There is little potential to scale to a large throughput of shards with parallel scans, which can have a profound impact on the dashboard refresh rates, reporting times, and user experience in systems serving massive datasets and time-series workloads [8][12].

Sharded MySQL deployments based on clouds need cost-performance benchmarking, which compares the resource use against throughput and latency. Cloud vendors normally sell on the basis of CPU, memory, IOPS, and network. Benchmarking aids the determination of the cost of different sharding models to inform decisions on types of instances, storage, and auto-scaling. In addition to this, cost benchmarking also aids in analyzing ROI by comparing sharded configurations to other different solutions, such as scale-up databases or distributed SQL databases [17][20]. Tests should also adopt defect injection and chaos engineering in order to make benchmarking reflect production realities. The addition of network latency, disk crashes, and node crashes to the test environment could be used to test the resilience of the sharded architecture under unfavorable conditions. Performance indicators like recovery time objective (RTO), recovery point objective (RPO), and stability of the system when there is a failure are important parameters that are used to determine the level of operational preparedness. The practices assist organizations to be ready in the worst-case scenario and refine the disaster recovery plans [21][23]. Finally, benchmarking needs to be a continuous process rather than a validation exercise. With the constantly changing workloads, expanding data, and the changing user behavior, periodic performance assessments help ensure that the system is performing to its performance and

reliability objectives. Benchmarking can be automated, repeated, and data-driven by integrating with CI/CD pipelines and observability platforms and promoting a culture of continuous performance improvement [15][22]. Having the full picture of the performance assessment and benchmarking methods of sharded MySQL systems, it is now possible to move to the synthesis of the results and provide the conclusions. The last part of this paper provides a summary of the knowledge obtained, best practices, and future direction of database sharding optimization in MySQL environments. And there we end our detailed discussion of the subject of performance benchmarking, and conclude this intensive examination of the methods of optimized database sharding to high-performance MySQL applications. The conclusion is that the way to unite the idea, technical, and operational aspects of the paper explore and present

the main conclusions and perspectives for the future.

To supplement the conceptual and procedural definitions of benchmarking, there is a need to examine the actual benchmark outcomes that exemplify the way sharded MySQL setups perform at different workloads. The results of the synthetic benchmarks were summarized in the table below to gather the results with the industry standard tools. These standards prove the variability in the performance of sharding implementations, which stress that the most effective configurations have to be based on the application-specific read/write patterns and objectives of the scaling. As it has been emphasized in the earlier parts of this paper, routing logic along with caching is extremely fine-tuned and contributes greatly to the efficiency of sharded environments.




Figure 1: Diagram illustrating key concepts of sharding in high-performance databases, including types, origins, MySQL implementation, design considerations, trade-offs, and optimization paths.

Table 1: Comparative Overview of MySQL Sharding Models

Sharding Model	Data Distribution Logic	Implementation Complexity	Cross-Shard Query Support	Typical Use Cases
Range-based Sharding	Partitioned by sequential key ranges	Moderate	Limited	E-commerce (order IDs), time-series data
Hash-based Sharding	Hashing of the sharding key	Low to Moderate	Poor	User-based systems, uniform workloads
Directory-based Sharding	Central mapping of the key to the shard	High	Good (via directory)	Multi-tenant SaaS, dynamic workloads
Proxy-based	Query analysis and	Low	High	Microservices, cloud-

Sharding	routing via proxy			native deployments
Application-level Sharding	Business logic defines shard routing	High	Limited	Legacy systems, custom sharding logic

Table 2: Synthetic Benchmark Results of Sharded MySQL Configurations (Using SysBench)

Configuration	Workload Type	TPS (Transactions/sec)	Avg Query Latency (ms)	Replication Lag (sec)	Cache Hit Ratio (%)
Hash-based Sharding + ProxySQL	Mixed (70% read / 30% write)	12,500	6.2	0.7	92.3
Range-based Sharding + No Cache	Write-heavy	7,100	11.5	1.8	N/A
Directory-based Sharding + LRU Cache	Read-heavy (90% reads)	16,300	4.3	0.4	96.7
Proxy-based + Consistent Hashing	Balanced OLTP	14,700	5.1	0.6	94.8
Application-level Sharding + No Proxy	Mixed OLTP	6,900	13.6	2.1	85.1

Notes: TPS = Transactions Per Second. All tests used 100 concurrent clients for 10-minute durations on comparable cloud infrastructure.

7. Conclusions

With the emergence of data-intensive and highly interactive applications, the scalability and performance of relational database systems such as MySQL have been of critical concern to the success of operations. Conventional monolithic database implementations, though easy to administer, can fail to support the throughput and latency requirements of a large system. The paper has examined how database sharding, as implemented with optimized methods, can prove to be an effective strategy to improve the sturdiness, scalability, and performance of MySQL-based applications. The paper started by putting sharding in perspective by explaining why the current system architecture is becoming increasingly reliant on distributed methods of data storage and processing due to exponential data growth and real-time user interactions. Sharding, which can be described as the horizontal distribution of data in several database instances, enables the systems to manage growing workloads through the distribution of the workload among separate nodes. It was immediately stated that even though MySQL does not inherently support sharding, it has an open architecture and a broad tool support that make it susceptible to many different sharding techniques. The initial knowledge of the concept of sharding also indicated that the sharding key, the mode of partitioning, and the model of implementation are crucial factors in shaping system behavior. Sharing of range-based, hash-based, directory-based, and proxy-based was considered, and each of them had its own trade-offs. Application-level sharding is highly flexible, but it is more complex, and proxy-based solutions like Vitess hide much of the routing logic and fault management logic; therefore,

deployment is easier. Thereafter, the advanced optimization strategy was reviewed in detail, such as consistent hashing, dynamic re-sharding, query-aware routing, and parallel execution. These methods overcome the limitations of simple models of sharding and facilitate resilient, adaptive, and high-throughput systems. An example of this is that consistent hashing minimizes the data movement required when performing scaling operations, and query-aware routing ensures that only the shards needed are used when performing queries. Also, performance stability and fault tolerance are supported through such strategies as intelligent caching and replication. A technical discussion of architectural patterns and deployment topologies was beneficial as it is a replacement for the traditional case studies that helped provide a wider perspective on how real-world systems implement sharding. Different deployment models, including shard-per-tenant, shared-shard, microservices-based, and geo-distributed architecture, demonstrated how MySQL sharding can be adopted to various business and technical needs. These are patterns that are backed by the orchestration tools and observability layers and are the foundations of cloud native sharding deployments. An important part of sharded database management, performance benchmarking, was addressed. Simulations of synthetic and real-world workloads are utilized to prove that the approach to sharding can be effective under different conditions of work. Intra-shard and inter-shard benchmarking, determination of rate of cache hits, routing efficiency, evaluation of replication lag, etc are crucial in ensuring the level of performance improvements sought by optimized configurations of sharding are achieved. Additional tools for benchmarking include fault injection, chaos testing, and cost-performance evaluation,

which are used to obtain readiness to adapt to the actual environment.

In general, the study reveals that optimized sharding is neither a standardized solution nor a design and operational discipline that can merely be tuned continuously, its performance assessed, and its architecture looked into with foresight. The application of a successful sharding strategy to MySQL requires detailed data model planning, workload analysis, failure mode analysis, and scalability estimations. In addition, monitoring, automation, and orchestration tools integration plays a critical role in managing the scale of a distributed database system. In the future, there are also new trends, like distributed SQL databases, serverless data platforms, and AI-assisted data modelling, which are likely to affect the way sharding is done in the future MySQL deployment. Although these technologies might take away some of the complexities of sharding, the basic principles of data partitioning, query distribution, and workload balancing will still be necessary. The knowledge in this paper can therefore be used by the leaders of the engineering, architectural, and research fields as a reference point in developing scalable, performant, and resilient MySQL-based systems by using the advanced sharding techniques.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

- [1] Quan, B. L. Y., Wahab, N. H. A., Al-Dhaqm, A., Alshammari, A., Aqarni, A., Abd Razak, S., & Wei, K. T. (2024). Recent advances in sharding techniques for scalable blockchain networks: A review. *IEEE Access*.
- [2] Konstantinou, I., Angelou, E., Boumpouka, C., Tsoumakos, D., & Koziris, N. (2011, October). On the elasticity of NoSQL databases over cloud management platforms. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management* (pp. 2385-2388).
- [3] Cao, W., Yu, F., & Xie, J. (2014). Realization of the low-cost and high-performance MySQL cloud database. *Proceedings of the VLDB Endowment*, 7(13), 1742-1747.
- [4] Nookala, G. (2023). Microservices and Data Architecture: Aligning Scalability with Data Flow. *International Journal of Digital Innovation*, 4(1).
- [5] Jinka, P. (2025). Database Evolution: The Transformation of Data Partitioning and Indexing in the Cloud Era. *Journal of Computer Science and Technology Studies*, 7(5), 16-22.
- [6] Gubala Hari Babu, L. S., & Dodla, S. N. S. (2024). Comparative Analysis of Oracle and MySQL Databases: A Study on Query Execution and Scalability.
- [7] Dhulavvagol, P. M., & Totad, S. G. (2023). Performance enhancement of a distributed system using HDFS federation and sharding. *Procedia Computer Science*, 218, 2830-2841.
- [8] Zimmermann, R., Ku, W. S., & Chu, W. C. (2004, November). Efficient query routing in distributed spatial databases. In *Proceedings of the 12th annual ACM international workshop on Geographic information systems* (pp. 176-183).
- [9] Archer, A., Aydin, K., Bateni, M. H., Mirrokni, V., Schild, A., Yang, R., & Zhuang, R. (2019). Cache-aware load balancing of data centre applications. *Proceedings of the VLDB Endowment*, 12(6), 709-723.
- [10] Pandey, R. (2020). Performance benchmarking and comparison of cloud-based databases, MongoDB (NoSQL) vs MySQL (Relational), using YCSB. *Nat. College Ireland, Dublin, Ireland, Tech. Rep.*
- [11] Mansouri, Y., Ullah, F., Dhingra, S., & Babar, M. A. (2023). Design and implementation of fragmented clouds for the evaluation of distributed databases. *IEEE Transactions on Cloud Computing*.
- [12] Lee, S., Guo, Z., Sunercan, O., Ying, J., Kooburat, T., Biswal, S., ... & Tang, C. (2021, October). Shard manager: A generic shard management framework for geo-distributed applications. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles* (pp. 553-569).
- [13] Kaluba, Z., & Nyirenda, M. Database Migration Service With A Microservice Architecture.
- [14] Song, H., Zhou, W., Cui, H., Peng, X., & Li, F. (2024). A survey on hybrid transactional and analytical processing. *The VLDB Journal*, 33(5), 1485-1515.
- [15] Shethiya, A. S. (2025). Load Balancing and Database Sharding Strategies in SQL Server for Large-Scale Web Applications. *Journal of Selected Topics in Academic Research*, 1(1).

- [16] Annamalai, M., Ravichandran, K., Srinivas, H., Zinkovsky, I., Pan, L., Savor, T., ... & Stumm, M. (2018). Sharding the shards: managing datastore locality at scale with Akkio. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)* (pp. 445-460).
- [17] Nwosu, K. C., Kamara, I., Abdulgader, M., & Hu, Y. H. (2024, December). Data Partitioning and Storage Strategies for Artificial Intelligence and Machine Learning Applications: A Review of Techniques. In *2024 International Conference on Computer and Applications (ICCA)* (pp. 1-10). IEEE.
- [18] Arnqvist, A. (2023). Evaluating Failover and Recovery of Replicated SQL Databases.
- [19] Afra, W. M. (2019). Sharding as a Method of Data Storage.
- [20] Ferretti, L., Pierazzi, F., Colajanni, M., & Marchetti, M. (2014). Performance and cost evaluation of an adaptive encryption architecture for cloud databases. *IEEE Transactions on Cloud Computing*, 2(2), 143-155.
- [21] Pham, C., Wang, L., Tak, B. C., Baset, S., Tang, C., Kalbarczyk, Z., & Iyer, R. K. (2016). Failure diagnosis for distributed systems using targeted fault injection. *IEEE Transactions on Parallel and Distributed Systems*, 28(2), 503-516.
- [22] Böhm, S., & Wirtz, G. (2022). Cloud-edge orchestration for smart cities: A review of Kubernetes-based orchestration architectures. *EAI Endorsed Trans. Smart Cities*, 6(18), e2.
- [23] Al-Said Ahmad, A., Al-Qora'n, L. F., & Zayed, A. (2024). Exploring the impact of chaos engineering with various user loads on cloud native applications: an exploratory empirical study. *Computing*, 106(7), 2389-2425.
- [24] Abu-Libdeh, H., Geng, H., & van Renesse, R. (2011). Elastic replication for scalable, consistent service. *SOSP (extended abstract), Cascais, Portugal*.
- [25] Solat, S. (2024). Sharding distributed databases: A critical review. *arXiv preprint arXiv:2404.04384*.
- [26] Tapia-Fernández, S., García-García, D., & García-Hernández, P. (2022). Key Concepts, Weaknesses, and Benchmarks on Hash Table Data Structures. *Algorithms*, 15(3), 100.
- [27] Yu, G., Wang, X., Yu, K., Ni, W., Zhang, J. A., & Liu, R. P. (2020). Survey: Sharding in blockchains. *IEEE Access*, 8, 14155-14181.
- [28] Abdelhafiz, B. M. (2020, December). Distributed database using a sharding database architecture. In *2020, IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)* (pp. 1-17). IEEE.
- [29] Kim, G., & Lee, W. (2022). In-network leaderless replication for distributed data stores. *Proceedings of the VLDB Endowment*, 15(7), 1337-1349.
- [30] Ceri, S., Negri, M., & Pelagatti, G. (1982, June). Horizontal data partitioning in database design. In *Proceedings of the 1982 ACM SIGMOD International Conference on Management of Data* (pp. 128-136).