

---

**RESEARCH ARTICLE**

## MongoDB and Data Consistency: Bridging the Gap between Performance and Reliability

**Mukesh Reddy Dhanagari**

*Manager, Software Development & Engineering, Charles Schwab*

**Corresponding Author:** Mukesh Reddy Dhanagari, **E-mail:** [mukesh.dhanagari@schwab.com](mailto:mukesh.dhanagari@schwab.com)

---

**ABSTRACT**

MongoDB is a popular NoSQL database with high scalability, flexible schema management, and fast data performance. While this is similar to relational databases that require compliance with ACID principles, MongoDB takes an eventual consistency model instead, wherein even the partition tolerance is preferred over the consistency. This paper discusses MongoDB's placement concerning the CAP theorem; that is, it is a CP (Consistency and Stability) database, and it guarantees data reliability while at the same time, performance bottlenecks could be an issue for MongoDB because it happens on a single node by default when performing reads and writes. Tuning MongoDB for better performance allows one to distribute read operations over secondary nodes and, in turn, reduce the workload on the primary node. This, however, brings eventual consistency as depending on where you request the data from, it might not be completely up to date. The paper presents MongoDB's replication methods, read/write concerns, sharding strategies, indexing, caching, and concurrency control techniques. MongoDB is a genius in large-scale apps but lacks strict consistency in supporting financial transactions and managing healthcare data. The paper addresses that distributed database environments require adaptive consistency models and AI-driven optimization to bridge the performance and reliability gap.

**KEYWORDS**

MongoDB, NoSQL Databases, Data Consistency, Eventual Consistency, Performance Optimization

**ARTICLE INFORMATION**

**ACCEPTED:** 02 June 2024

**PUBLISHED:** 25 June 2024

**DOI:** 10.32996/jcsts.2024.6.2.21

---

### 1. Introduction

#### Introduction

The NoSQL database MongoDB stands among the most popular solutions because it delivers excellent scalability, flexible features, and fast response times. MongoDB implements a document-oriented methodology that enables developers to store data using BSON format, which resembles JSON but extends its capabilities through extra features. Through its schema-free data architecture, MongoDB delivers rapid development possibilities, making it the top selection for diverse systems that must adapt quickly, like real-time analytics projects, content systems, and IoT initiatives. Modern applications need faster and more scalable databases, so MongoDB established itself as an effective solution through its features for horizontal scalability, built-in replication, and distributed architecture capabilities. MongoDB is an efficient data processing system that is highly available thanks to these built-in features. However, these performance benefits from MongoDB systems create a major drawback because they compromise system speed and data reliability. MongoDB utilizes an eventual consistency approach instead of relational databases, which use strong consistency, so users achieve better performance while accepting minimal data reliability risks.




Figure 1: MongoDB Architecture

### Importance of Data Consistency and Performance in Databases

Database management relies heavily on data consistency because this process maintains correct and updated information synchronizing across various nodes. When using strongly consistent systems, data written into the system becomes immediately visible through all read operations without delay. This approach, common in relational databases, guarantees data accuracy but often at the cost of latency and performance bottlenecks. MongoDB achieves better performance through eventual consistency because it lets temporary inconsistencies exist during operation. This approach makes data consistent for all nodes, although instant synchronization between all nodes is not mandatory. The trade-off between consistency and performance allows distributed systems to become more available and tolerant of faults while supporting large-scale operations. Database systems demand high performance, especially for applications with fast responses and immediate processing requirements. Efficient data storage systems enhance data retrieval speed and write operations, resulting in faster responses for end users. MongoDB uses indexes, sharding, and replica sets to efficiently distribute data to multiple servers. The ongoing challenge is striking the right balance between consistency and speed because it represents a critical issue in current database architecture paradigms.

### Where Does MongoDB Fall in the CAP Theorem?

The Mongo DB illustrates how should trade off in Distributed DB like Cap Theorem. For a database, with network partitions, they can choose between consistency and availability. MongoDB is in the CP (Consistency and Partition Tolerance) group, so it provides consistency while handling network partitions usually with a lack of availability.

This is because MongoDB defaults to CP mode and writes and reads occur on the same node. For the strong consistency, it provides certain burden, and at the same time it puts a limitation on reaching high performance. Under high loads all operations take place on the primary node thus bottlenecks are possible and response times as well as overall throughput suffer.

### The Ongoing Debate: Relational (SQL) vs. NoSQL (MongoDB) Databases

The database sector exists between relational (SQL) databases, including MySQL and PostgreSQL, and NoSQL databases, including MongoDB and Cassandra. The ACID model (Atomicity, Consistency, Isolation, Durability) regulates transactions in relational databases, which maintain their structural format. Distributed scalability remains a challenge for those who use these database systems. The MongoDB database system and other NoSQL solutions focus on ensuring high performance and flexibility through limited consistency controls to achieve maximum resilience and system stability. MongoDB uses a BASE (Basically Available Soft-state Eventually Consistent) model to maintain system reliability and operational speed by compromising data consistency. The selection of SQL or NoSQL depends solely on the required usage type. Financial transactions and healthcare systems need strong consistency to operate effectively with relational databases. Applications requiring quick data inputs and outputs and large-scale expansion capabilities should choose NoSQL systems like MongoDB.

This article investigates how MongoDB manages its eventual consistency features while maintaining reliable performance. The evaluation examines how MongoDB maintains efficient data management alongside its consistency solutions, replication models, and performance improvement methods.

### Understanding Data Consistency in Databases

#### Definition and Importance of Data Consistency

Database management requires data consistency as an essential principle, which requires all database system elements to show precise, equal, and current data values (Silberschatz et al. 2011). Implementing consistency allows databases to show the latest valid information to all read operations after any data modification or update occurs. Data consistency remains crucial during distributed system operations with several simultaneous node or server transactions. Enterprise applications demand strict data consistency maintenance because it ensures appropriate decisions and financial operations and protects data security. Fleet management effectiveness requires analytical data accuracy to monitor assets, maximize operational efficiency, and improve communication networks within supply chains. All operational sectors, such as healthcare, e-commerce, and banking, operate with

consistent data to prevent operational issues and compliance breaks and maintain customer loyalty. NoSQL databases such as MongoDB tend to optimize performance and scalability more than ensure immediate consistency, unlike traditional relational databases, which maintain strict structured consistency protocols. Different consistency models emerged after organizations began deciding between achieving maximum performance or maintaining data consistency in systems. Strong consistency and eventual consistency emerged as the main categories.




Figure 2: What Is NoSQL? NoSQL Databases Explained

**Strong Consistency vs. Eventual Consistency**

Data consistency models establish procedures for data update distribution during system operations and decide when users can view changes. The primary consistency models are strong and eventually consistent, featuring their distinctive strengths and specific performance weaknesses.

- **Strong Consistency**

A transaction becomes permanent after submission, resulting in all database requests for that record displayed with the most current update, no matter which nodes process the request. Data integrity and reliability are ensured by this method, although it typically leads to reduced performance and diminished availability (Bowman, 2013). The RDBMS database solutions, including MySQL and PostgreSQL, implement strong consistency through their requirement for the ACID properties of Atomicity Consistency, Isolation, and Durability. The system employs two-phase commit procedures and locking controls to achieve strict synchronized states.

- **Eventual Consistency**

Systems implementing eventual consistency maintain temporary inconsistency until replicas automatically reach a consistent state during a specific period. Multiple distributed NoSQL systems like MongoDB, Apache Cassandra, and Amazon DynamoDB utilize this model because they emphasize high availability and horizontal scalability above immediate consistency. Large-scale applications, especially social media platforms, e-commerce websites, and logistics systems, benefit from eventual consistency since they need high throughput.

**Examples of Consistency Models in Relational vs. NoSQL Databases**

Visual design and target applications shape how both relational and NoSQL implementations execute their consistency rules.

**Relational Databases and Strong Consistency**

- Relational databases such as Oracle, SQL Server, and PostgreSQL adopt strong consistency through their implementation of transactions under the ACID model.
- The databases achieve data consistency through locking mechanisms, write-ahead logging (WAL) systems, and synchronous replication techniques.
- The relational database management system in banking facilitates account money transfers by executing debit and credit operations simultaneously or leaving no transactional changes if any operation fails.

**NoSQL Databases and Eventual Consistency**

- NoSQL databases, including MongoDB, Apache Cassandra, and Amazon DynamoDB, implement eventual consistency as a method to enhance speed and agility.
- The system provides a brief window where different nodes store different data versions before each node reaches consistency with the others.
- When users modify their social media profile picture on the platform, other users temporarily view the un-updated image until all servers reflect the new change. All users eventually see the change after its initiation, yet people might view outdated information during this period.

Organizations must evaluate their applications to determine whether strong or eventual consistency provides better value (Stufflebeam & Coryn, 2014). Applications dealing with mission-critical functions need strong consistency, yet eventual consistency delivers better performance when handling high user traffic, such as content delivery networks (CDNs) and recommendation engines.

**ACID vs BASE**

Databases uphold data consistency using two maintenance models named after their initials: ACID and BASE.

**ACID (Atomicity, Consistency, Isolation, Durability)**

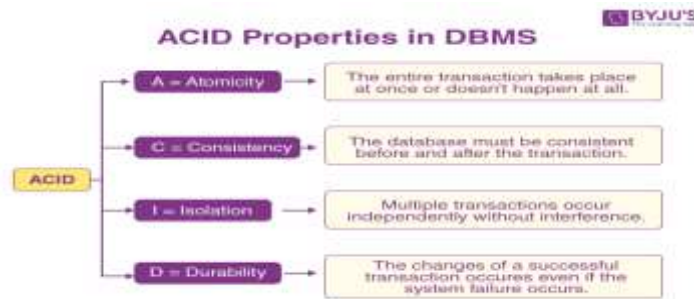


Figure 3: ACID Database Properties

- Transactional processes in relational databases depend on the ACID model, which guarantees reliable transaction execution.
- A transaction under Atomicity either achieves success or executes a complete rollback process.
- A transaction under this model ensures that valid states of the database transform automatically to new valid states.
- A transaction cannot disrupt other transactions when this feature is implemented.
- The durability element of the ACID model protects permanent transaction data storage against system failures.

The ACID model represents a vital requirement for maintaining data integrity, especially within banking and healthcare systems and legal record management operations. ACID makes booking flight tickets possible, properly maintaining seat availability to avoid double bookings.

**BASE (Basically Available, Soft-state, Eventually consistent)**

The BASE serves as a standard model in NoSQL databases, although it prioritizes performance and availability instead of demanding full consistency.

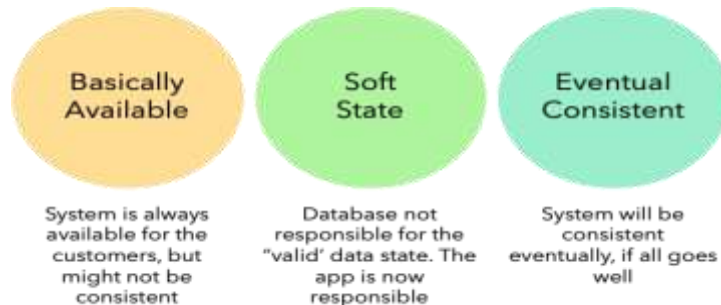


Figure 4: Database Selection & Design

- A system under BASIC operation stays active as long as parts of its infrastructure are operating intermittently.
- The system can accept unpredictable quantities of inconsistent data during the temporary period.
- The system eventually uses a consistent approach, making all nodes achieve identical statuses while operating in the background.
- BASE proves ideal for systems that need to scale rapidly because it supports the handling of little delayed data in cases involving content delivery networks, online retail, and IoT applications (Swamy & Kota, 2020).

**Tuning MongoDB to Use Secondary Nodes for Performance**

The default case is that MongoDB reads and writes on the same primary node. However, MongoDB can also be configured to use replica nodes (replica nodes) as read nodes to further improve performance. This is where the read requests are split onto all database cluster nodes, which can improve hardware usage and enhance performance.

**Benefits of Using Secondary Nodes for Reads**

- **Load Balancing:** Spreads the read operations across multiple nodes, reducing the workload on the primary node.
- **Better Resource Utilization:** Takes advantage of all cluster nodes instead of relying solely on the primary.
- **Faster Read Performance:** As multiple nodes handle queries, overall query execution time decreases.

**Limitation of Using Secondary Nodes for Reads**

Secondary nodes improve performance but use eventual consistency. While write operations continue to occur only on the primary node, replication to secondary nodes occurs only after a short period. This implies that users can read stale data by querying secondary nodes.

**How Eventual Consistency Negatively Affects Reliability**

When using eventual consistency, MongoDB can briefly see data that doesn't make it across all primary and secondary nodes. The problem stems from write operations, which are evaluated on the primary node and asynchronously propagated to the secondary nodes, which means that, at times, different nodes hold different versions of the data in this window of time.

**Problems Caused by Eventual Consistency**

1. **Data Inconsistency:** Different nodes may return outdated information, leading to confusion in real-time applications.
2. **Inaccurate Real-time Analytics:** Analytical applications requiring the latest data may suffer from outdated results due to replication delays.
3. **Race Conditions:** When applications read from a secondary node immediately after a write operation, they may receive outdated data, causing inconsistent application behavior.

Despite these challenges, businesses that prioritize performance over strong consistency can adopt **MongoDB's read preferences**, configuring applications to choose between consistency and performance based on their needs.

**MongoDB's Approach to Consistency**

The eventual consistency model that MongoDB supports as a NoSQL database keeps scalability and performance ahead of complete data compliance. MongoDB delivers data reliability through its replication mechanism, read/write concerns, and journaling features, whereas traditional relational databases (RDBMS) require strict ACID transactions to enforce data consistency. This part discusses how MongoDB attains eventual consistency while performing read/write tasks and maintaining durability attributes across distributed setups.

**How MongoDB Achieves Eventual Consistency**

The consistency model of MongoDB operates on an eventually consistent basis because all replicated datasets move toward equivalent states during the period. Several major aspects enable this outcome, including:

- MongoDB uses asynchronous replication for replica sets, which ultimately enables secondary nodes to synchronize with the primary node to achieve high availability.
- Through Write Concerns, developers can establish database interaction requirements for written data to advance, which enables them to strike the right balance between system speed and data preservation.
- Clients can control the freshness of the read operation by choosing between reading from primary nodes and, eventually, consistent secondary nodes.

MongoDB's eventual consistency model enables high throughput by allowing applications to function without transactional strictures; hence, it proves useful for real-time systems that include e-commerce platforms, social media feeds, and IoT systems (Kumar, 2019).

**Read and write concerns:**

The Read and Write Concerns in MongoDB system enable users to set precise rules about where each operation originates and where it is stored. The read-and-write concerns in MongoDB enable users to control the relationship between data accuracy and operational speed.

**1. Read Concern Levels**

The read concern specifies the current state duration that read operations must fulfill. The available levels include:



Figure 5: The "Majority" WriteConcern of MongoDB Replica Sets

- i. Local concerns allow reading the most current uncommitted data on the node serving the demand (Cao et al. 2021).
- ii. The available read option provides access to data on any node, although it might show outdated or inconsistent entries. The read performance reaches its peak while returning results that might already be outdated.
- iii. The read operation needs acknowledgment from most replica nodes to ensure fresh and efficient data retrieval.
- iv. Linearizable behavior ensures readers access the most recent committed version of data throughout the whole replica set. Critical financial transactions require this option because it represents the most resource-intensive solution, (Gill, 2018).

## 2. Write concern Levels

Different levels of write concern measurement combine with acknowledgment models as vital components for MongoDB. The method of confirming write operations in MongoDB is determined by write concerns. The different levels include:

- Through this approach ( $w=0$ ), the client writes data to MongoDB without requiring any confirmation. The write speed reaches maximum limits without guaranteeing data persistence.
- The write operation arrival confirmation occurs at the primary host per the acknowledged level ( $w=1$ ) before disk storage but without a storage guarantee.
- Operations with the  $w$ =majority parameter get recorded on multiple replica nodes until a majority approves them for better data uptime assurance.
- With  $j=true$  logging Journalled ( $j=true$ ), MongoDB implements a data-writing journal process before providing acknowledgment that enhances reliability when system crashes occur.
- The linearizable ( $w=all$ ) condition makes all nodes commit writes before returning confirmation. Strong consistency comes with a detrimental impact on system performance when using these settings.

The selection of read-and-write concerns follows the application requirements. Applications in the financial sector need the majority of readers along with journaled writing operations (Cortada, 2005), while real-time analytics systems rely on available reading and unacknowledged writing approaches.

### **Replica Sets and Their Impact on Consistency**

Replica Sets in MongoDB work as node clusters to provide high availability for database operations and backup functions. A standard replica set contains at least three parts:



Figure 6: How does replication work in MongoDB?

- **The primary node** maintains full write responsibilities while offering the latest system data to clients.
- **The secondary nodes** receive data asynchronously from the primary side and possess read capabilities if the system allows it.
- **The optional Arbiter Node** serves to perform primary node elections, yet it does not store any data.

Replica Sets maintain eventual consistency through their operation. The process of writing data begins exclusively on the primary node. Secondaries replicate the data asynchronously. Clients can read from secondaries for eventual consistency or obtain strong consistency from the primary. A new primary selection process occurs after failure to ensure continuous operational availability. Applications remain operational under all circumstances because replica sets provide fault tolerance (Guerraoui & Schiper, 1996).

### **Journaling and Durability in MongoDB**

The core durability feature of MongoDB relies on journaling. Journaling ensures permanent write recordation, which protects data from loss when a system crash occurs.



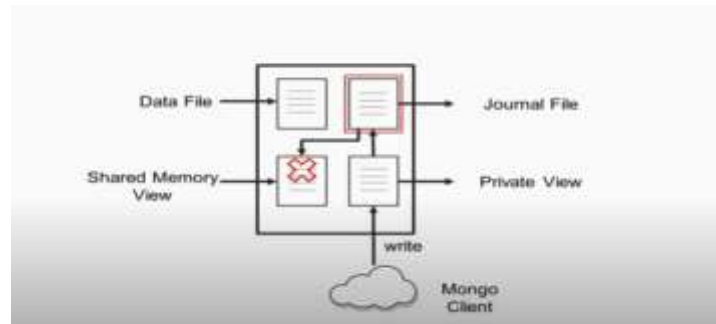


Figure 7: Journaling in MongoDB

- **How Journaling Works**

Each database operation starts by writing data to the journal file, which is then committed to the main database. MongoDB uses the journal entries to restore database consistency when a failure occurs. The journal system runs automatic checks at measured intervals (by default every 100 milliseconds) to strike a proper equilibrium between execution speed and system dependability.

- **Benefits of Journaling**

The crash recovery mechanism prevents data loss when the MongoDB server encounters unexpected crashes (Hafiz et al. 2023). Write performance boosts significantly since MongoDB stores data first in memory logs before automatic disk-writing operations increased reliability. A consistent write-ahead log is achievable through this feature, which maintains stability between distributed network nodes.

- **Journaling vs. Write Concerns**

The journaling system protects database integrity but depends on writing concerns to establish the safety features that determine data storage methods. For example: Strong durability results when enabling journaling ( $j=true$ ) and writing concerns set to ( $w=majority$ ). Applications without journaling ( $j=false$ ) choose faster data writes than reliable storage through write concern set to  $w=1$ . The use of journaling and majority writing has become essential for financial transactions to maintain data stability (Gomber et al. 2018).

### Performance Optimization in MongoDB

The NoSQL database MongoDB is widely popular due to its flexible design and strong ability to scale and handle high-performance operations. The database's performance reaches its best capacity through the utilization of sharding and horizontal scaling and combined strategies encompassing indexing, query optimization, caching, and concurrency control mechanisms. Multiple performance improvement methods enable MongoDB to maintain its speed, efficiency, and reliability to process large information volumes and rapid transaction sequences successfully.

#### Sharding & Horizontal Scaling

MongoDB implements sharding as its essential method for splitting data horizontally between multiple servers to improve performance. The system improves speed and reliability and expands capacities by splitting data into easier-to-manage sections.

- **How Sharding Works in MongoDB**

Each shard in the MongoDB implementation maintains an independent dataset range so that no one database instance handles an excessive workload. A client can retrieve data efficiently because the Mongos query router uses proper shard assignment.

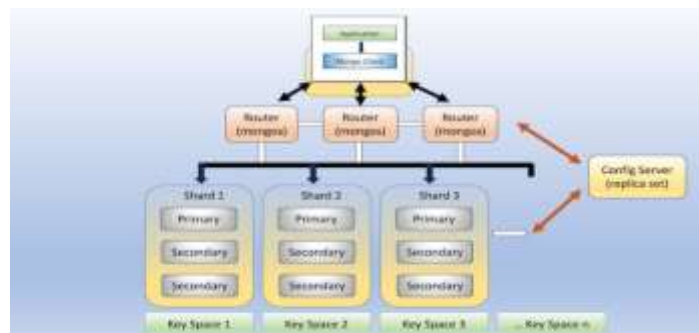


Figure 8: Understanding MongoDB's Sharding Architecture

- **Choosing the Right Shard Key**

Selecting an adequate shard key directly determines system performance levels (Benini & De Micheli, 2002). A carefully selected shard key can distribute data evenly across the clusters, reducing performance obstacles. Improper key selection creates hotspot conditions that cause certain shards to receive dramatically more queries than others.

- **Benefits of Sharding**

1. A distributed query system increases read-write performance since it avoids concentrating query processing on any node.
2. The distribution of duplicate data across shards results in better system reliability.
3. The system acquires elastic scalability through new node deployments, which boost performance measures but do not cause disruptions.

- **Challenges of Sharding**

Sharding technology increases system complexity when handling database data operations. Before rebalancing shards devices, users need to implement automated approaches and active monitoring since the process leads to a performance drop.

**Indexes and Query Optimization**

Index technology is the key to speeding up MongoDB performance by minimizing the amount of scanned data during search operations.

- **Types of Indexes in MongoDB**

1. Single-field indexes: Improve query performance for a single attribute.
2. Implementing compound indexes enables increased performance when executing search operations across multiple fields.
3. Multikey indexes optimize search operations that include arrays as part of the query terms.
4. Text indexes enable complete text search capabilities through their operations.
5. Hashed indexes function to maximize performance when dealing with hashed shard keys during querying.

- **Query Optimization Techniques**

When running Covered Queries, every field needed must be found inside a predefined index structure to block document-crawling operations. The usage of indexes must be optimized to prevent full collection scans in the query process. Relational Sorting benefits tremendously when indexes support the specified sorting structure.

- **Analyzing Query Performance**

Through its explain() feature, MongoDB enables developers to examine execution plans, which assists in finding bottlenecks.

- **Best Practices for Indexing**

1. The system requires indexing of active fields to achieve lower response times.
2. Deploy partial indexes to store selective documents within the system for better efficiency.
3. The application must implement TTL (Time-to-Live) indexes to manage data records with an expiration threshold.

**Caching Strategies in MongoDB**

The fundamental role of caching in MongoDB is to increase performance and efficiency by storing popular data requests in memory.




Figure 9: Caching a MongoDB Database with Redis — SitePoint

- **Built-in Caching with WiredTiger**

When running with MongoDB, Firestore provides an internal cache through its WiredTiger storage engine. The system automatically adjusts memory allocation during operation to avoid using the disk inappropriately.

- **External Caching Mechanisms**

The combination of Redis with MongoDB serves as a cache layer by storing data that is frequently accessed in memory. This approach improves the database's performance and retrieves data more quickly. Memcached serves as a solution for fast caching needs that support applications running through millions of read operations per second.

- **Query Result Caching**

The performance of MongoDB can be improved by creating a query result cache system that avoids repeat calculations (Bradshaw et al. 2019). Techniques include:

1. Implement aggregation framework caching as a solution for queries needing performance improvements.



2. Gas and N2O injections across the entire system for server-side operation.
3. Top pre-initialized data acts as instant serving resources for commonly requested information.
4. Balancing Cache Size and Consistency

Fast application performance results from caching, yet data staleness becomes a problem unless administrators take proper steps to manage it. Cache invalidation strategies guarantee every client user access to current information by minimizing database call frequency.

**Concurrency Control Mechanisms in MongoDB**

Multiple operations could run simultaneously through concurrent access control, which avoids data consistency and performance problems.

- **Locking Mechanisms in MongoDB**

The Optimistic Concurrency Control method functions well when conflicts appear infrequently. Clients perform updates by checking for conflicts before their changes become permanent. The Pessimistic Concurrency Control (PCC) serves data systems prone to frequent conflicts. The system locks records before updates so other processes cannot modify identical stored data.

- **Document-Level Locking**

Concurrent execution of read/write operations becomes possible through document-level locking features, which MongoDB employs instead of traditional table-level locking methods. The system encounters fewer conflicts because documents acquire locks prior to updates, which produces higher system performance.

- **Replica Set Synchronization**

Primary replication sets from MongoDB enable eventual consistency, allowing secondary nodes to display minor delays in their data update times. Developers can decide between two read preference configurations. To achieve higher performance, readers should access the secondary database nodes. The primary node is where administrators must read for maximum consistency. The Write Concern and Read Concern parameters determine how MongoDB acknowledges data operations (Mehmood et al. 2017). Write concerns parameters define the required level of acknowledgment that MongoDB needs before confirming write operations. Raising the write concern level makes systems more reliable, but their performance slows down. The read Concerns setting enables programs to adjust data freshness according to performance requirements.

- **Concurrency Bottlenecks and Solutions**

The system must prevent operations that consume more time than write operations. Establishing an asynchronous processing system for proper workload management. The system should use batch writes to enhance efficiency compared to individual inserts.

**Challenges of Balancing Performance and Reliability**

As a NoSQL database, MongoDB offers a scalable and high-performance alternative to traditional relational databases. Understandably, distributed systems present difficulties while maintaining both performance and reliability. The main obstacles in MongoDB usage consist of eventual consistency management risks, scaling obstacles, implementation difficulties in production environments, and practical MongoDB implementation examples. The optimization of MongoDB data consistency requires knowledge of these system issues when providing fast performance.

**Trade-offs in Eventual Consistency**

MongoDB uses the eventual consistency model because it prioritizes quick operations and available system access above instance-by-instance data conformity. However, the enhanced scalability comes at a price: It produces stale reads alongside delayed synchronization activities between nodes serving as replicas.

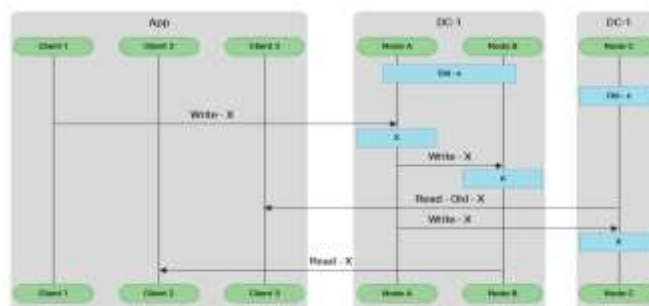


Figure 10: What is Eventual Consistency?

- **Latency in Data Synchronization:** The asynchronous replication mode in MongoDB enables secondary nodes to require time for change propagation between them (Zhang et al. 2015). All applications suffer from replication lag because it causes them to read outdated data, which becomes crucial in real-time applications that process financial transactions.

- **Inconsistency During Failovers:** The system selects one of its available secondary nodes to replace the failed primary nodes and function as the new primary node. When data completion fails to meet replication requirements before failure occurs, certain write operations become lost and produce divergent node data.
- **Conflict Resolution Complexity:** Application-level conflict resolution mechanisms and last-write-wins (LWW) require developers to maintain additional complexity for eventual consistency implementations.

Strong transactional guarantees in banking or healthcare environments might not find eventual consistency design acceptable because it does not provide the required durability standards.

**Scalability Concerns vs. Strict Consistency**

Because of its horizontal scalability features, businesses can use MongoDB to efficiently spread their data across multiple servers. However, enforcing ACID transactions commonly used in relational databases proves more difficult within distributed NoSQL systems (GC, 2016).

- **Write Performance vs. Durability:** The write concerns in MongoDB enable users to choose between performance-enhancing acknowledged writes or durable writes that demand majority acknowledgment confirmation. The application's performance decreases due to heightened latency because write concern levels aim for data reliability.
- **Sharding Overhead:** MongoDB uses sharding as the scaling method to distribute its data across multiple nodes. Maintaining strict consistency becomes challenging because multi-document transactions need to coordinate across various partitions.
- **Distributed Locking Challenges:** Because MongoDB implements distributed architecture instead of relational database locking systems, maintaining row-level consistency while preserving performance levels becomes complicated.

MongoDB, along with other NoSQL systems, trades off strict consistency guarantees for better fault performance and scalable infrastructure. Mission-critical programs need to establish proper ratios of write concerns to read consistency and transaction isolation measures to prevent data inconsistencies.

**Real-World Challenges in Using MongoDB in Production**

Different web applications currently use MongoDB as a database solution, but several production environments encounter major stability issues when achieving peak performance while maintaining system dependability.

- **Financial Transactions (Banking & Fintech):** Transactions need instant consistency because both double-spend and account balance errors must be avoided immediately. The delayed capacity for eventual consistency in updating balance information results in problems for fraud detection systems and reconciliation processes. The banking sector extends MongoDB by adding relational technology to maintain databases that require strict consistency (Giamas, 2022).
- **Healthcare Records & Compliance:** The strict requirements for EHR systems include data integrity and full audit capabilities. Eventual consistency in MongoDB creates risks that may affect medical decisions in healthcare due to noncompliance with HIPAA regulations. The implementation of multi-document transactions by Healthcare applications within MongoDB offers consistency but results in reduced performance speed.
- **E-Commerce and Real-Time Applications:** High-traffic systems need to use MongoDB clusters configured for sharding to support their expanding requirements. The system needs to maintain precise inventory management to prevent stocking errors that result in merchandise being sold out of stock. Extreme user response times result from the combination of read concerns set to majority level alongside write concerns set to acknowledged level in business applications. Due to its JSON document flexibility, businesses can utilize MongoDB despite PII data compatibility, yet they need to implement intensive access controls and encryption systems for security purposes.

**Case Studies of Businesses Using MongoDB for High Performance**

MongoDB is adopted by leading businesses because of its speed and scaling abilities, although these organizations integrate additional strategies to eliminate possible consistency-related problems.

- **eBay (E-commerce & Marketplace Analytics):** The system takes advantage of MongoDB to perform real-time analysis of customer activities. The system operates under eventual consistency for search indexing and recommendations but maintains order processing on its SQL-based systems.
- **Forbes (Media & Content Management):** The application uses MongoDB to process large quantities of articles with images and user-submitted content. The system emphasizes scalability more than consistency to provide content availability throughout multiple regions.
- **Uber (Geo-Tracking & Ride Matching):** MongoDB serves as the system for processing real-time location-based data. The system will automatically sacrifice instant data synchronization to achieve rapid match-ride performance and quick user responses. MongoDB's performance benefits remain successful because of the implementation of hybrid storage architectures combined with caching layers and customized consistency models by all these companies (Mehmood et al. 2017).



Figure 11: Real-Time Payments

### Implementation Methodology

The performance and consistency traits of MongoDB heavily depend on how users configure and deploy it. The following part offers a detailed guide to executing MongoDB effectively to achieve reliability, scalability, and performance efficiency. The following paragraphs present the required method for MongoDB deployment and detailed information about data consistency and system performance.

#### **Step 1: Setting up a MongoDB instance (local or cloud-based)**

The procedure for MongoDB implementation begins with a selection between running MongoDB on-site or using cloud resources.

**Local Installation:** The official MongoDB website provides a platform to download the MongoDB database. Use the correct installation method from the operating system to install MongoDB through its package manager. Implement configuration settings for the MongoDB server tool (mongod) before connecting it to the MongoDB shell application (mongo).

**Cloud-Based Deployment:** Service providers should adopt MongoDB Atlas because it operates as a fully managed cloud database solution. Develop a MongoDB cluster by selecting a relevant cloud provider between AWS, Azure, and Google Cloud Platform. Security configuration must include authentication settings, IP address list protection, and network access policies. Security measures for MongoDB data protection must be implemented after selecting an appropriate authentication and access control structure (Hoberman, 2014).

#### **Step 2: Configuring replica sets for redundancy and reliability.**

The MongoDB feature known as replica sets defines a critical aspect that enables high data availability and redundant storage.

**Setting Up a Replica Set:** Implement a MongoDB cluster with three nodes. The primary node takes the leading role and is backed by two secondary nodes. Apply `rs.initiate()` at the beginning to create the replica set. Add secondary nodes using `rs.add("secondary_node_address")`. A properly configured MySQL cluster with replica sets enhances consistency by enabling users to adjust different application demands for read and write operations.

#### **Step 3: Implementing read and write concerns for different use cases.**

The read-and-write concerns in MongoDB allow users to achieve performance goals together with data consistency.

**Read Concerns:** The read operation returns node data but does not ensure consistency levels.

- Available: Returns data from any available replica.
- A majority read operation checks whether data received a commit from more than half of the nodes.
- Linearizable: Guarantees the most recent version of the document.

#### **Write Concerns:**

This execution concern requires acknowledgment from the primary server following the processing of the written statement. The creation of written data is validated when written to most replica nodes using the "majority" clause. Write persistence is guaranteed through journaling when setting `{ j: true }`. The selection of read and write concerns is based on the specific application needs. Business transactions need a majority of written concerns to ensure strong consistency.

#### **Step 4: Setting up sharding for large-scale applications.**

Sharding is a MongoDB feature that spreads database data across multiple servers to deliver improved performance for extensive applications.

**Steps to Configure Sharding:**

1. Enable sharding on the database:

*Table 1: Enable sharding on the database*

```
bash
sh.enableSharding("database")
```

2. A shard key should be established to control how database information gets distributed among servers.

*Table 2: Define a shard key on how data is distributed among different shards*

```
db.myCollection.createIndex({ shardKey: 1 })
sh.shard collection("database.myCollection", { shardKey: 1 })
```

3. Form multiple shard clusters for the system and introduce them into this configuration.

*Table 3: Multiple shards created and added to the MongoDB cluster.*

```
sh.addShard("shard1.example.com:27017")
sh.addShard("shard2.example.com:27017")
sh.addShard("shard3.example.com:27017")
```

- **Benefits of Sharding:**

1. Horizontal Scaling: Distributes data across multiple machines.
2. The system executes queries on shards that correspond to their execution requirements.
3. Workload performance issues disappear because multiple servers handle an equal distribution of duties.
4. Proper shard key selection is the key requirement because range-based and hashed sharding methods are commonly used for this step.

**Step 5: Optimizing indexes and queries to boost performance.**

Enhancing indexes and queries is the fifth and final optimization step to increase operational speed. The efficiency of MongoDB query performance greatly depends on proper index optimization.

**Types of Indexes:**

1. Single Field Index: Improves lookup performance.
2. The Compound Index enables querying several fields within the database.
3. The TTL (Time-To-Live) Index has a built-in function to delete documents based on a predetermined time threshold.
4. Text Index provides users with a full-text search functionality.

**Optimizing Queries:**

Analyze query performance through the executionStats query parameter using the .explain() method. The use of correct index structures prevents complete collection scans from occurring. Apply projection with find({}, {field1: 1, field2: 1}) to retrieve specific fields alone. Indexes that have been optimized properly reduce both the execution time of database queries and system resource usage.

**Step 6: Use MongoDB Monitoring Tools for Performance Analysis**

Tracking MongoDB performance is an essential requirement for efficiency maintenance and problem-solving (Zhou et al. 2016).

**Built-in MongoDB Monitoring Tools:**

- Mongostat: Displays real-time performance metrics.
- The Mongotop tool delivers read-and-write operation information.
- MongoDB Atlas Monitoring: Cloud-based performance analytics.

**Best Practices for Performance Tuning:**

Use the Profiler to track slow database queries at consistent intervals. Performance enhancement requires optimized CPU usage, RAM, and disk I/O resources. A connection pooling system should be used to obtain efficient database connections. Monitoring allows organizations to detect operational slowdowns that lead to ongoing improvements of their databases.

**Best Practices for Using MongoDB**

The NoSQL database MongoDB delivers flexibility, scalability, and high-performance characteristics, which position it as an ideal solution for contemporary applications. The deployment of MongoDB depends on following best practices that protect data consistency, reliability, security, and performance and optimize performance. Whichever consistency model should be selected

depends on this section, along with practical guidelines for writing concern settings, tuning approaches, and security measures that protect data integrity.

**When to Use Strong Consistency vs. Eventual Consistency**

MongoDB implements eventual consistency, improving its database's performance and availability. Data accuracy needs sometimes demand the utilization of strong consistency, but eventually, consistency serves applications better (DeCandia et al. 2007).

**Strong Consistency Use Cases:** Financial applications that process money need immediate consistency to avoid paying one item twice or accidentally making wrong computations. In online commerce, real-time inventory tracking systems prevent business owners from selling products past their current stock levels. Medical records within the healthcare system must maintain consistent data to guarantee proper patient diagnosis and treatment results.

**Eventual Consistency Use Cases:** Social Media Feeds distribute updates across the system asynchronously, but users do not encounter performance issues. Systems that perform Big Data analytics benefit from delayed data processing to gain improved system performance. Content Delivery Networks (CDNs) use cached data for quick retrieval operations before consistency considerations occur.

Table 4: ways to achieve strong and eventual consistency in MongoDB applications.

Consistency Type	Use Cases	Priority
Strong Consistency	Financial Transactions, Inventory, Healthcare	Data Accuracy
Eventual Consistency	Social Media, Big Data Analytics, CDNs	System Performance

**How to Configure Write Concerns for Data Reliability**

Through MongoDB write concerns, users can determine to write acknowledgments while assuring data reliability. The implementation of write concern contains three main priority levels:

- No write acknowledgment is used in w: 0, which allows optimal performance but establishes unsafe conditions for data loss vulnerability.
- Primary nodes under MongoDB receive data acknowledgment through the w: 1 (Acknowledged) confirmation.
- W: The majority provides reliable performance by guaranteeing data writes to multiple nodes until they outnumber the total members of the replica set.
- A write concern at level w: Every node must confirm data transmission, achieving maximum data durability. The write concern should be set to the majority where journaling remains enabled.

Table 5: The Trade-Offs between various write concerns

Write Concern	Data Safety	Performance	Best For
w: 0 (Unacknowledged)	Low	High	Caching, Logging
w: 1 (Acknowledged)	Medium	High	General Applications
w: majority	High	Medium	Business-Critical Data
w: all	Very High	Low	Financial & Healthcare

**Performance Tuning Techniques for Large-Scale Applications**

The following optimization techniques need to be used when aiming to enhance MongoDB performance in substantial applications:

1. **Indexing:** The system will benefit from indexes on the frequently accessed fields because they boost read operation speed. The optimization of multiple fields can be achieved through compound index creation.
2. **Sharding:** Distributes large datasets across multiple servers for better performance. Apply hashed sharding methods when wanting an even distribution, and range sharding provides better performance for particular workloads.
3. **Connection Pooling:** Implementing connection pooling minimizes connection fees through database connection Reutilization instead of implementing brand-new instances.
4. **Query Optimization:** Avoid scanning the entire database by adopting indexed search queries. Projection techniques should enable data retrieval of only required fields while discarding complete document data.
5. **Write Optimization:** Enhancing write throughput occurs when data batches cut down the number of distinct operations used during write processes. Journaling is a durability mechanism because it prevents data loss through write persistence before acknowledgment.

**Security Measures to Protect Data Integrity in MongoDB**

MongoDB requires protection against unauthorized access and unauthorized data breaches, as well as protection against data inconsistencies (Goel & Ter Hofstede, 2021). Several best practices enable the preservation of data integrity as follows:

1. Authentication and Authorization: RBAC should be implemented to provide restricted data access permissions. Users need SCRAM (Salted Challenge Response Authentication Mechanism) as a secure authentication solution.
2. Encryption: Database servers must be protected through Transport Layer Security (TLS/SSL) to secure data transfers during communication. At-rest encryption protects all the stored data on company computer systems.
3. Backup and Recovery: The backup process requires scheduled execution through MongoDB's built-in MongoDB tool or cloud-based backups. Point-in-time recovery through PITR enables users to retrieve database states from different times.
4. Network Security: Limit MongoDB's network connectivity through the use of firewall parameters. Users should restrict remote access except for occasions when it is acutely necessary.
5. Auditing and Monitoring: MongoDB's audit logs function should be enabled to monitor unauthorized system access. Performance monitoring should be done through MongoDB Atlas and Prometheus platforms.

Implementing MongoDB's best practices maintains an equilibrium between system speed, dependability, and defense mechanisms. Organizations should use MongoDB to effectively combine its knowledge of strong or eventual consistency with optimized write concerns and performance optimization while implementing security measures to achieve maximum MongoDB benefits. Future MongoDB innovations will enhance its data consistency and integrity functionality to create a more powerful NoSQL solution.

**Future of MongoDB and Data Consistency**

**Trends in NoSQL and Distributed Database Systems**

NoSQL databases, together with distributed systems, develop modern data storage infrastructure that defines the current evolution of information management platforms. This data transformation demands a leading NoSQL solution, MongoDB, because the database leverages its capability to scale while ensuring availability and performance enhancement. The modern computational era requires databases to execute distributed transactions with efficient processing power and optimal performance results (Kallman et al. 2008). The market continues to welcome multi-model databases that unite key-value, document, graph, and columnar database management engine features. Combining database models in one system provides organizations with increased flexibility and enables them to merge their data needs into unified operational systems (Zicari, 1991). Serverless databases have become popular since they reduce database management challenges and provide automatic scalability features. Organizations can use MongoDB Atlas serverless instances to achieve automatic scaling benefits through their infrastructure provisioning systems. AI and machine learning represent important industry developments that influence database management. Artificial intelligence tools enhance index strategies together with query execution and data consistency by applying workload pattern predictions to modify automatically the database configuration parameters. Research into retrieval-augmented generation (RAG) methods enhances the accuracy of data retrieval within distributed computer systems.




Figure 12: Best Practices in Retrieval Augmented Generation

**Potential Improvements in MongoDB's Consistency Model**

The eventual consistency model has limitations for MongoDB, although it demonstrates strong scalability and performance capabilities. Implementing multi-document ACID transactions through MongoDB 4.0 enhanced database capabilities, but future performance optimization will resolve consistency-performance clashes. The main enhancement focuses on improving causal consistency features. Operations that use causal consistency models maintain order between distributed nodes, thus minimizing unexpected issues that eventual consistency models can produce. Future versions of MongoDB are expected to offer vector clocks, and conflict-free replicated data types to enhance consistency in worldwide replication. The upcoming development in MongoDB includes automated tools for consistency tuning (Schultz et al. 2019). AI-driven models running in the background would examine usage patterns to automatically modify read and write concerns, optimizing consistency levels without performance decay. Using this approach to maintain efficient resource allocation, developers would achieve higher consistency requirements in



critical areas. MongoDB will continue advancing its geo-partitioning technology to apply data consistency policies on smaller and more specific levels. Businesses with operations spanning different regions will achieve better consistency modeling through specific requirement-oriented optimizations.

**Predictions for the Future of Eventual Consistency vs. Strong Consistency**

Remote and distributed database technology will persist in influencing the ongoing debate between eventual consistency and strong consistency. Real-time systems like social media and IoT, together with content delivery networks, work better with eventual consistency, but financial transactions and critical data integrity require strong consistency. Upcoming database technologies will deploy adaptive consistency models because they let applications decide between strong consistency and eventual consistency through workload adjustment. The normative approach in distributed ledger technologies involves combining consensus-based protocols along with eventual consistency mechanisms. Blockchain-inspired consensus approaches will enable the creation of a connection between relational databases and NoSQL systems (Mijoska & Ristevski, 2020). The advancements in MongoDB will likely incorporate Byzantine Fault Tolerance (BFT) and other consensus-driven consistency protocols to achieve better reliability with similar performance levels. The leaders within NoSQL databases will face future directions from the increasing demand to achieve the perfect equilibrium between high performance, scalability, and dependable data storage models. Eventually, consistent systems will continue as the standard paradigm, but advancements in distributed computer technologies will develop novel smart consistency solutions.

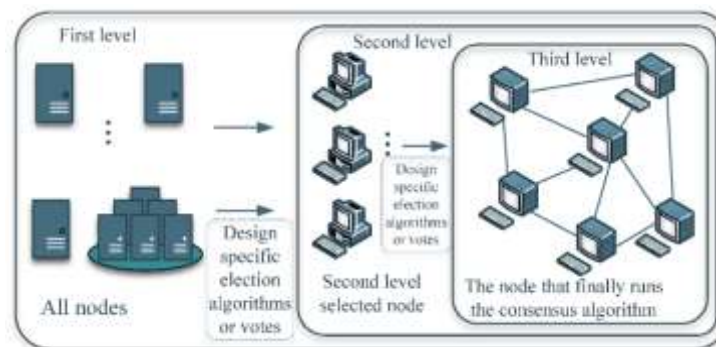


Figure 13: Byzantine Fault-Tolerant Consensus Algorithms: A Survey

**Conclusion**

The NoSQL database MongoDB establishes itself as a top solution that provides scalable, high-performance solutions for modern applications. The analysis of this piece has examined how MongoDB implements data consistency through eventual consistency methods and its performance versus reliability balance considerations. The BASE architecture enables MongoDB to provide high speed and availability because it functions without traditional database mechanisms that follow the strict ACID model for maintaining data integrity. The system can distribute data across multiple servers using sharding and replication to maintain high data availability. The speed increase through eventual consistency violates data synchronization between all nodes. MongoDB allows developers to select between performance and stronger consistency by configuring read and write concerns that match their application requirements. The research analyzed MongoDB's performance optimization methods, which include indexing, caching, and concurrency control, to enable efficient data retrieval and transaction processing.

How effectively MongoDB achieves performance and reliability depends on the system's needs according to the scenario and framework. Network systems requiring exact real-time data precision, like financial operations or medical applications, protect data accuracy through precise consistency models regardless of performance effect. Applications that prioritize system scalability and high availability find MongoDB's eventual consistency model to be quite successful. Organizations must evaluate their needs precisely before setting MongoDB configuration, which determines how they will manage speed versus data protection. Specifications for local, majority, and linearizable read concerns and acknowledged writes and majority writes concerns enable developers to set exact parameters for distributed data behaviors across systems. Implementing planned indexing methods combined with optimized queries and replication procedures produces performance enhancements to MongoDB that maintain its reliability levels.

Organizations should study their application needs to determine appropriate MongoDB configuration options when evaluating this database system. Developers must work towards improving data access speed by implementing indexing and using the built-in MongoDB profiler to monitor system performance and optimize access patterns. The infrastructure of businesses needs to enable MongoDB replication and sharding functions to sustain availability throughout periods of high demand. Authentication, access control, and encryption take center stage when it comes to data protection because security demands proper implementation. The implementations of MongoDB succeed despite its powerful flexibility and scalability based on appropriate

configuration strategies. Organizations can develop dependable and high-speed data applications by properly utilizing MongoDB consistency models and performance enhancement techniques.

**Funding:** This research received no external funding.

**Conflicts of Interest:** The authors declare no conflict of interest.

**Publisher's Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

## References

- [1] Benini, L., & De Micheli, G. (2002). Networks on chips: A new SoC paradigm. *computer*, 35(1), 70-78.
- [2] Bowman, S. (2013). Impact of electronic health record systems on information integrity: quality and safety implications. *Perspectives in health information management*, 10(Fall).
- [3] Bradshaw, S., Brazil, E., & Chodorow, K. (2019). *MongoDB: the definitive guide: powerful and scalable data storage*. O'Reilly Media.
- [4] Cao, W., Zhang, Y., Yang, X., Li, F., Wang, S., Hu, Q., ... & Tong, J. (2021, June). Polardb serverless: A cloud native database for disaggregated data centers. In *Proceedings of the 2021 International Conference on Management of Data* (pp. 2477-2489).
- [5] Cortada, J. W. (2005). *The digital hand: Volume II: How computers changed the work of American financial, telecommunications, media, and entertainment industries*. Oxford University Press.
- [6] DeCandia, G., Hastorun, D., Jampani, M., Kakulapati, G., Lakshman, A., Pilchin, A., ... & Vogels, W. (2007). Dynamo: Amazon's highly available key-value store. *ACM SIGOPS operating systems review*, 41(6), 205-220.
- [7] GC, D. (2016). A critical comparison of NOSQL databases in the context of ACID and BASE.
- [8] Giamas, A. (2022). *Mastering MongoDB 6. x: Expert techniques to run high-volume and fault-tolerant database solutions using MongoDB 6. x*. Packt Publishing Ltd.
- [9] Gill, A. (2018). Developing a real-time electronic funds transfer system for credit unions. *International Journal of Advanced Research in Engineering and Technology (IJARET)*, 9(1), 162–184. <https://iaeme.com/Home/issue/IJARET?Volume=9&Issue=1>
- [10] Goel, K., & Ter Hofstede, A. H. (2021). Privacy-breaching patterns in NoSQL databases. *IEEE Access*, 9, 35229-35239.
- [11] Gomber, P., Kauffman, R. J., Parker, C., & Weber, B. W. (2018). On the fintech revolution: Interpreting the forces of innovation, disruption, and transformation in financial services. *Journal of management information systems*, 35(1), 220-265.
- [12] Guerraoui, R., & Schiper, A. (1996, June). Fault-tolerance by replication in distributed systems. In *International conference on reliable software technologies* (pp. 38-57). Berlin, Heidelberg: Springer Berlin Heidelberg.
- [13] Hafiz, M. M., Zaman, A. K. M., & Shaf, M. S. (2023). *Preventing Data Loss using Raft Consensus Algorithm in a Decentralized Database System* (Doctoral dissertation, Department of Computer Science and Engineering (CSE), Islamic University of Technology (IUT), Board Bazar, Gazipur-1704, Bangladesh).
- [14] Hoberman, S. (2014). *Data Modeling for MongoDB: Building Well-Designed and Supportable MongoDB Databases*. Technics Publications.
- [15] Kallman, R., Kimura, H., Natkins, J., Pavlo, A., Rasin, A., Zdonik, S., ... & Abadi, D. J. (2008). H-store: a high-performance, distributed main memory transaction processing system. *Proceedings of the VLDB Endowment*, 1(2), 1496-1499.
- [16] Kumar, A. (2019). The convergence of predictive analytics in driving business intelligence and enhancing DevOps efficiency. *International Journal of Computational Engineering and Management*, 6(6), 118–142. <https://ijcem.in/wp-content/uploads/THE-CONVERGENCE-OF-PREDICTIVE-ANALYTICS-IN-DRIVING-BUSINESS-INTELLIGENCE-AND-ENHANCING-DEVOPS-EFFICIENCY.pdf>
- [17] Mehmood, N. Q., Culmone, R., & Mostarda, L. (2017). Modeling temporal aspects of sensor data for MongoDB NoSQL database. *Journal of Big Data*, 4(1), 8.
- [18] Mehmood, N. Q., Culmone, R., & Mostarda, L. (2017). Modeling temporal aspects of sensor data for MongoDB NoSQL database. *Journal of Big Data*, 4(1), 8.
- [19] Mijoska, M., & Risteviski, B. (2020). Blockchain Technology and its Application in the Finance and Economics.
- [20] Schultz, W., Avitabile, T., & Cabral, A. (2019). Tunable consistency in mongodb. *Proceedings of the VLDB Endowment*, 12(12), 2071-2081.
- [21] Silberschatz, A., Korth, H. F., & Sudarshan, S. (2011). Database system concepts.
- [22] Stufflebeam, D. L., & Coryn, C. L. (2014). *Evaluation theory, models, and applications* (Vol. 50). John Wiley & Sons.
- [23] Swamy, S. N., & Kota, S. R. (2020). An empirical study on system level aspects of Internet of Things (IoT). *IEEE Access*, 8, 188082-188134.
- [24] Zhang, Y., Yang, J., Memaripour, A., & Swanson, S. (2015, March). Mojim: A reliable and highly-available non-volatile memory system. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems* (pp. 3-18).
- [25] Zhou, L., Chen, N., & Chen, Z. (2016). A cloud computing-enabled spatio-temporal cyber-physical information infrastructure for efficient soil moisture monitoring. *ISPRS International Journal of Geo-Information*, 5(6), 81.
- [26] Zicari, R. (1991, January). A framework for schema updates in an object-oriented database system. In *Proceedings. Seventh International Conference on Data Engineering* (pp. 2-3). IEEE Computer Society.