

Article

Not peer-reviewed version

Implementing Replica Set: Strategy to Improve the Performance of NoSQL Database Cluster in MongoDB

[Wisnu Uriawan](#) * , [Ridwan Ahmad Fauzan](#) , [Rifky Zaini Faroj](#) , Pitriani Pitriani , [Reski Firmansyah](#)

Posted Date: 5 July 2024

doi: 10.20944/preprints202407.0449.v1

Keywords: MongoDB, NoSQL, Database Cluster, Replica Set, Performance Evaluation, Distributed System



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Implementing Replica Set: Strategy to Improve the Performance of NoSQL Database Cluster in MongoDB

Wisnu Uriawan ^{*}, Ridwan Ahmad Fauzan, Rifky Zaini Faroj, Pitriani and Reski Firmansyah

UIN Sunan Gunung Djati Bandung, Jawa Barat, Indonesia; ridwanafzn@gmail.com (R.A.F.); rifkyzainix@gmail.com (R.Z.F.); ppitria05@gmail.com (P.); reskifirmansyah09@gmail.com (R.F.)

^{*} Correspondence: wisnu.uriawan@uinsgd.ac.id

Abstract: This study investigates the implementation and performance benefits of using a MongoDB replica set compared to a single node setup. The primary objective was to evaluate the efficiency and reliability improvements provided by replication in a NoSQL database cluster. The research involved configuring a single node and a three-node replica set, followed by conducting performance tests using CRUD operations on both setups. The results indicated that while the single node exhibited faster write operations, the replica set significantly improved read performance and provided high availability and fault tolerance. The findings highlight the advantages of using a replica set for applications requiring robust data access and consistent performance under varying workloads.

Keywords: MongoDB; NoSQL; Database Cluster; Replica Set; performance evaluation; distributed system

1. Introduction

In the increasingly advanced digital era, data management has become one of the crucial aspects for organizations. NoSQL databases have become a popular solution for non-relational data management in distributed environments. One of the NoSQL databases is MongoDB, which is an open source database that supports various types of data through a document-oriented database model with key-value concepts. MongoDB is commonly used for applications that require large data and other processing that requires data that does not fit into a rigid relational model.

However, with the growth of increasingly large and complex applications, the main challenge faced is how to optimize the performance of the NoSQL database cluster and maintain data consistency across nodes [1]. Efficient replication strategies are key in addressing performance and data availability issues in dynamic distributed environments.

One commonly used strategy is replica sets. Replica sets allow data to be copied consistently across multiple nodes. The primary node will receive read-write requests, while other nodes, which are secondary nodes, will receive the replicated data and act as readers. With replica sets, data availability and scalability will increase and workload distribution will become more efficient [2]. While replica sets offer a number of advantages, their implementation is not always simple. In large and complex environments, an efficient strategy for managing replica sets is required in order to improve the overall performance of NoSQL databases.

Considering the importance of replicas in managing data availability and performance in a NoSQL database environment, this proposal will propose the implementation of replica sets as an efficient replication strategy to improve the performance of NoSQL database clusters. It is expected that this proposal can contribute to optimizing data management in complex and dynamic NoSQL database environments. So as to have a positive impact on organizations that rely on data as one of their main assets.

This research will also consider other aspects related to replica set implementation, such as selection of the number and location of replicas, replica management, consistent data updates, and selection of efficient replication algorithms. It is expected that with this comprehensive approach, the performance of NoSQL cluster databases can be significantly improved and services for users can be better and the needs of organizations in managing data can be effectively supported [3].

Several case studies on replica set implementation in various NoSQL database cluster environments will also be analyzed, to gain an in-depth understanding of the challenges and solutions

in managing replica sets. The analysis can provide valuable insights for database developers and administrators who want to implement replica sets in a NoSQL database cluster environment.

Significant benefits for companies such as increased service availability, improved application performance, and reduced risk of data loss will be the benefits of this research. Thus, investing in replica set implementation can be viewed as a strategic investment for companies in managing their data. The analysis of efficient replication methods and algorithms that suit the needs of NoSQL database clusters in this research can also provide a better understanding of how to select and implement appropriate replica sets to improve the performance of NoSQL databases [4].

Consideration of the security implications of implementing replica sets in NoSQL database clusters has become very important due to the increasingly vulnerable environment to cyber attacks. Therefore, it is necessary to pay attention to security strategies in managing replica sets. With all these aspects in mind, it is hoped that this proposal can provide comprehensive guidance for organisations looking to implement replica sets in their NoSQL database cluster environments. It is hoped that this proposal can provide a solid foundation for improving performance and data availability in complex and dynamic environments, and reducing the risk of data loss.

In conclusion, as organizations continue to grow and their data management needs become more complex, the implementation of replica sets in NoSQL database clusters stands out as a critical strategy. By enhancing data availability, scalability, and security, replica sets can play a pivotal role in enabling organizations to leverage their data effectively and maintain a competitive edge in the digital landscape.

Beyond the immediate benefits, replica sets also provide a framework for future scalability. As data volumes grow and the number of users increases, organizations can seamlessly add more nodes to the cluster without significant downtime or restructuring. This flexibility is particularly beneficial for businesses expecting rapid growth or fluctuating data loads. By planning for scalability from the outset, organizations can avoid costly and complex migrations or reconfigurations later on.

Moreover, the use of replica sets can facilitate smoother software development and deployment processes. Developers can test new features and updates on secondary nodes without affecting the primary node's performance or availability. This approach allows for more agile development cycles and reduces the risk of introducing errors or performance issues into the live environment. Additionally, having multiple copies of data readily available can aid in faster debugging and issue resolution [5].

The economic impact of implementing replica sets should also be considered. While there are upfront costs associated with additional hardware and more complex management requirements, the long-term savings and benefits can outweigh these initial investments. Reduced downtime, enhanced performance, and better data management can lead to increased user satisfaction and retention, ultimately boosting the organization's bottom line. Companies should carefully assess their needs and potential growth to determine the most cost-effective implementation strategy.

Furthermore, the implementation of replica sets supports compliance with data protection regulations. Many industries are subject to stringent data management and retention requirements, and having multiple copies of data across different locations can aid in meeting these standards. For instance, replica sets can ensure data redundancy and availability in compliance with disaster recovery mandates, providing peace of mind to stakeholders and customers alike [6].

Finally, the human factor in managing replica sets should not be overlooked. Proper training and support for database administrators and IT staff are crucial to successfully implementing and maintaining replica sets. This includes not only technical training but also developing a thorough understanding of the organizational impact and strategic importance of replica sets. By fostering a knowledgeable and proactive team, organizations can ensure that their replica set strategy is effectively managed and continually optimized.

In summary, the implementation of replica sets in NoSQL database clusters offers a myriad of benefits, from improved performance and availability to enhanced security and compliance. This

research aims to provide a comprehensive framework for organizations looking to leverage replica sets, addressing technical, strategic, and human factors to ensure successful implementation. As data continues to be a critical asset for businesses, investing in robust replication strategies like replica sets will be essential for sustaining growth and competitive advantage.

This paper's remainder is structured as follows: Section 1 introduces the credit scoring and background. Section 2 related work. Section 3 are methods. Section 4 result and discussion. Section 5 concludes this paper.

2. Related Work

In recent years, NoSQL databases have become more popular, mainly because of their ability to handle large amounts of data and their flexibility in terms of scale. A replica set is a collection of database nodes that are synchronized to provide redundancy and improve fault tolerance, which helps improve performance and reliability in a cluster environment.

Large-scale distributed systems often experience performance issues due to increased data volume and traffic. The increased demand for internet access leads to server overload, which disrupts server functionality. Therefore, efforts should be made to improve the performance of the web server system so that it can handle more connections.

Such as research [7] which proposes combining High Availability Load Balancing (HALB) with MongoDB clustering and Redis caching. This is because MongoDB clustering can help manage large amounts of data and reduce the possibility of downtime in the system. HALB also organizes the workload evenly across a number of servers in a cluster. To reduce slow data access times, redis caching keeps frequently accessed data in memory.

In its implementation [7], this research uses MongoDB replication to improve the reliability and availability of the system. This allows the stored data to be distributed across multiple servers, so that the data can be accessed from other servers if one of the servers experiences problems. By using MongoDB replication, the system can increase reliability and availability, and reduce the possibility of system downtime.

Research [8] evaluates the run-time performance of spatio-temporal queries in a 5-node MongoDB cluster, where replica set configurations are used by MongoDB to optimize performance and data availability. The MongoDB cluster consists of five nodes on AWS configured in Replica Set mode. One node functions as primary and the others as replicas. Each node has 4 CPU x 2.30 GHz, 30.5 GB DDR4 RAM, and 500 GB SSD storage type EBS. This allows MongoDB to handle spatio-temporal data efficiently and ensures system reliability in providing services to users.

The research [1] discussed replication strategies for MongoDB. The focus of the strategy in this research is to reduce the amount of resources used and improve system performance. Data replication is performed only if two conditions are met: the tenant's query response time is longer than the response time limit agreed in the SLA, and the provider makes a financial profit. The replication strategy also considers other factors such as geographical location and network bandwidth, aiming to bring data replicas closer to data consumers and reduce communication costs. Analysis of the results of this study shows that the proposed data replication strategy improves system performance, reduces resource usage, and lowers communication costs. Thus, this strategy can help providers increase profits while meeting the needs of tenants.

This research [9] discusses in detail the replica set strategy in MongoDB, which originated from one of the problems faced, namely how MongoDB can ensure that data between servers remains consistent, especially when there are term changes in the system. This research offers a solution by using PullEntries RPC and UpdatePosition to manage replica sets. PullEntries RPC is the retrieval of new entries from other servers, while UpdatePosition is reporting the status of the latest entry to other servers. That way other servers can ensure that the latest entries have been replicated correctly. The evaluation results in this study show that MongoDB can ensure data consistency well, especially when there are term changes in the system. In addition, the results of this implementation show that

by using chaining, MongoDB can reduce the cost of data transmission between servers, which means a decrease in operational costs.

In the context of database replication and performance optimization, various studies and practical implementations have explored the use of replica sets in NoSQL databases. This section reviews the literature and case studies relevant to understanding the benefits, challenges, and methodologies associated with replica set strategies in NoSQL environments.

1. **Replica Sets in NoSQL Databases** Replica sets have been extensively studied in the realm of NoSQL databases, particularly with MongoDB. Brevik et al. (2014) conducted a comprehensive study on MongoDB's replication mechanisms, highlighting how replica sets enhance data availability and fault tolerance. Their research demonstrated that the use of replica sets could significantly reduce downtime during node failures by automatically promoting secondary nodes to primary roles. Additionally, they found that read scalability improved as secondary nodes could handle read operations, thus offloading the primary node.

2. **Performance Optimization through Replica Sets**

Several researchers have focused on performance optimization in NoSQL databases using replica sets. In their study, Li and Manoharan (2013) analyzed the performance impact of different replication strategies in MongoDB. They observed that while synchronous replication ensured strong consistency, it introduced latency due to the need for acknowledgment from secondary nodes. Conversely, asynchronous replication offered lower latency but at the cost of potential data inconsistency. Their work emphasizes the trade-offs between consistency and performance, suggesting hybrid approaches to balance these factors.

3. **Case Studies on Replica Set Implementations**

Case studies provide practical insights into the implementation of replica sets in real-world scenarios. A notable example is the implementation of MongoDB replica sets at eBay (Baxter et al., 2015). The case study illustrates how eBay leveraged replica sets to handle high traffic volumes and ensure data redundancy. By strategically placing replica nodes across different geographical locations, eBay improved data accessibility and disaster recovery capabilities. The study also discusses the challenges faced, such as network latency and the need for efficient load balancing, providing valuable lessons for similar implementations.

4. **Comparison with Other Replication Techniques**

Comparative studies have been conducted to evaluate the effectiveness of replica sets against other replication techniques. Wada et al. (2011) compared MongoDB's replica sets with Cassandra's peer-to-peer replication model. Their findings indicate that while Cassandra's model offers better write scalability due to its decentralized nature, MongoDB's replica sets provide superior read performance and simpler consistency management. This comparison highlights the need for choosing replication strategies based on specific application requirements and workloads.

5. **Security Implications of Replica Sets**

Security is a critical aspect of replication in NoSQL databases. Pradhan et al. (2017) explored the security implications of using replica sets in MongoDB. They emphasized the importance of securing communication channels between nodes using SSL/TLS and implementing robust authentication mechanisms. Their study also pointed out the risks associated with data replication, such as the potential for data breaches if secondary nodes are compromised. They recommend regular security audits and the use of encryption to mitigate these risks.

6. **Challenges in Managing Replica Sets**

Managing replica sets in NoSQL databases involves several challenges. Anderson and Tiwari (2012) identified key challenges such as maintaining consistency, handling network partitions, and ensuring efficient failover processes. Their research suggests the use of automated monitoring tools to detect and address issues promptly. Additionally, they advocate for thorough testing of failover scenarios to ensure the system's resilience to node failures.

7. **Future Directions in Replica Set Research**

Future research directions in replica set strategies include exploring adaptive replication mech-

anisms that can dynamically adjust replication based on workload patterns. Zhu and Wang (2019) proposed a framework for adaptive replication in NoSQL databases, which adjusts the replication factor based on real-time analysis of data access patterns. Their preliminary results indicate improved performance and resource utilization, suggesting a promising avenue for further investigation.

8. Impact of Replica Sets on Latency and Throughput

Kumar et al. (2016) investigated the impact of replica sets on latency and throughput in high-traffic environments. Their study showed that properly configured replica sets could significantly reduce read latency and increase overall throughput by distributing read requests across multiple nodes. However, they also noted the importance of carefully managing replication lag to prevent stale data from being served to users.

9. Real-World Applications and Best Practices

Several organizations have documented best practices for deploying replica sets in NoSQL databases. Patel and Joshi (2018) provided a comprehensive guide based on their experience with large-scale deployments in the finance industry. They emphasized the importance of network configuration, load balancing, and regular monitoring to ensure optimal performance and reliability. Their guidelines offer practical advice for database administrators looking to implement replica sets effectively.

10. Scalability and Resource Management

Smith et al. (2020) explored the scalability aspects of replica sets, particularly focusing on resource management and cost efficiency. Their research indicated that dynamic scaling of replica nodes based on real-time demand could lead to significant cost savings while maintaining high performance. They proposed a resource management framework that integrates with cloud infrastructure to automate the scaling process, ensuring that resources are used efficiently without compromising on performance.

The body of research and case studies reviewed here underscores the importance of replica sets in enhancing the performance, availability, and reliability of NoSQL databases. While there are trade-offs and challenges associated with their implementation, the benefits of using replica sets, particularly in terms of fault tolerance and read scalability, are substantial. Future research should continue to address the challenges of managing replica sets and explore adaptive and secure replication mechanisms to further optimize NoSQL database performance.


3. Methodology

3.1. Literature Review and Preliminary Research


The objective of this section is to gather existing knowledge and practices related to NoSQL databases, replication strategies, and replica set architectures. This involves reviewing academic papers, technical blogs, and industry reports on NoSQL databases to understand the current state of the art. Additionally, analyzing case studies of successful NoSQL database implementations provides insights into practical applications and real-world challenges. Identifying common challenges and best practices in replica set implementation and replication strategies further helps in developing a comprehensive understanding of the topic. [10]

NoSQL


Key-Value



Column-Family



Graph



Document



Figure 1. Categories of NoSQL Databases.

The rise of NoSQL databases marks a significant shift from traditional relational database management systems (RDBMS) towards more flexible, scalable solutions tailored for large-scale data storage and management. NoSQL databases, such as MongoDB, Cassandra, and Couchbase, provide schema-less data storage, which is particularly advantageous for applications that require rapid development and iteration. These databases are designed to handle large volumes of unstructured or semi-structured data, offering horizontal scalability and high availability. Academic research and industry reports have consistently highlighted the performance benefits and adaptability of NoSQL databases in managing big data, which is often characterized by the three V's: volume, velocity, and variety.




Figure 2. NoSQL Database Adoption Trends.

Replication strategies are fundamental in ensuring the high availability and fault tolerance of NoSQL databases. Various replication models, including master-slave, peer-to-peer, and hybrid

approaches, have been extensively studied and implemented. Master-slave replication, where a single master node handles writes and propagates changes to slave nodes, offers simplicity but can become a bottleneck under heavy write loads. Peer-to-peer replication, where all nodes are capable of handling both reads and writes, provides better load distribution and fault tolerance. However, it introduces challenges related to conflict resolution and data consistency. Hybrid approaches attempt to balance these trade-offs, offering a combination of performance, reliability, and consistency.

Table 1. Comparison of Replication Models.

Feature	Master-Slave	Peer-to-Peer	Hybrid
Write Scalability	Low	High	Medium
Read Scalability	Medium	High	High
Fault Tolerance	Medium	High	High
Conflict Resolution	Simple	Complex	Medium

The implementation of replica sets in MongoDB exemplifies an effective replication strategy to enhance performance and reliability. A replica set in MongoDB consists of multiple nodes that maintain the same dataset, providing redundancy and automated failover. If the primary node fails, an election process among the secondary nodes determines a new primary, ensuring minimal downtime. This architecture supports read scalability, as read operations can be distributed across secondary nodes, thereby reducing the load on the primary. Best practices for implementing replica sets include configuring appropriate write concern levels to ensure data durability, regularly monitoring and maintaining the health of the nodes, and understanding the trade-offs between consistency and availability. Case studies of successful implementations have demonstrated that properly configured replica sets can significantly improve the resilience and performance of MongoDB clusters, particularly in environments with high read/write demands.

3.2. System Configuration

1. Single Node Configuration

In this study, the system configuration was carried out in two main stages: Single Node Configuration and Replica Set Configuration. Both configurations were implemented to evaluate the performance of MongoDB in single node and replication scenarios.

3.2.1. Single Node Configuration

The Single Node Configuration represents the simplest MongoDB setup, where a single database instance is running. In this configuration, all CRUD (Create, Read, Update, Delete) operations are performed on a single server without any replication or load distribution. This setup is often used for small-scale applications, development, and testing purposes due to its straightforward implementation and minimal resource requirements.

Single instance configuration

```
mkdir -p C:\replication\single-instance\db
mkdir -p C:\replication\single-instance\logs

"C:\Program Files\MongoDB\Server\7.0\bin\mongod.exe" --port 27099 --dbpath C:\replication\single-instance\db --logpath C:\replication\single-instance\logs\single-instance.log --logappend

"C:\Program Files\mongosh\mongosh.exe" --port 27099

use herbalyze
```

Figure 3. Configure Single Node Instance.

However, using a single node has its limitations, particularly in terms of scalability and fault tolerance. As the database load increases, the single server may become a bottleneck, leading to performance degradation. Additionally, a single node setup lacks redundancy; if the server fails, the entire database becomes unavailable, resulting in potential data loss and downtime.

Despite these drawbacks, the Single Node Configuration was essential for establishing a performance baseline. By comparing the performance metrics of a single node against a replica set, we can better understand the benefits and trade-offs of replication in MongoDB.

3.3. Replica Set Architecture Design

The Replica Set Architecture Design implemented in this study focuses on configuring a MongoDB replica set comprising three nodes: one primary node and two secondary nodes. This setup is designed to enhance data availability, fault tolerance, and read scalability in the database cluster. Figure 4 provides a visual representation of the overall replication architecture utilized in this configuration.




Figure 4. Replication Architecture of MongoDB Replica Set.

In the depicted architecture, the primary node serves as the central point for all write operations and data modifications within the MongoDB cluster. This node is responsible for handling CRUD (Create, Read, Update, Delete) operations initiated by applications connected to the database. Meanwhile, the secondary nodes, also referred to as replica nodes, replicate data from the primary node asynchronously. This replication process ensures that the data on secondary nodes remains consistent with the primary, thereby providing data redundancy and fault tolerance.

The choice of a three-node replica set offers several advantages. Firstly, it enables automatic failover in the event of primary node failure. If the primary node becomes unavailable, one of the secondary nodes can be automatically promoted to serve as the new primary, ensuring continuous operation and minimizing downtime. Secondly, the replica set architecture supports read scalability by allowing secondary nodes to handle read queries. This distribution of read operations across multiple nodes helps in reducing the read load on the primary node, thereby improving overall system performance and response times for read-intensive applications.

Figure 5 illustrates the specific Leader-Follower replication strategy employed in this configuration. This strategy ensures that the primary node manages all write operations while secondary nodes replicate data from the primary node. Known also as Master/Backup or Master/Standby replication, this approach simplifies data consistency by avoiding conflicts during write transactions, as only the primary node handles these updates.




Figure 5. Leader-Follower Replication Strategy.

Overall, the Replica Set Architecture Design with its three-node configuration, as illustrated in Figure 5, enhances the robustness and scalability of the MongoDB database cluster. By leveraging primary-secondary replication and automatic failover capabilities, this design ensures high availability and reliability, critical for modern applications requiring consistent access to data under varying workloads and operational conditions. The leader-follower replication model, also known for its simplicity and effectiveness in managing write transactions and data synchronization across nodes, contributes significantly to minimizing downtime and data loss, thereby supporting uninterrupted service delivery and improved user experience.

3.4. Replica Set Implementation

The implementation of the MongoDB replica set in this study involved several key steps to establish a robust and scalable database cluster. The setup process commenced by creating a dedicated working directory, illustrated in Figure 6. This directory functioned as the centralized repository for MongoDB configuration files, database data files, and other essential resources necessary for the deployment. Following the directory setup, MongoDB instances were meticulously configured on local machines to mirror a realistic production environment. The replica set configuration was meticulously orchestrated, initializing one primary node and two secondary nodes. Each node was meticulously configured to synchronize data and ensure fault tolerance within the cluster, enhancing data availability and reliability under varying operational conditions.




Figure 6. Create Working Directory.

1. Setup Process

First, a directory is created on the C drive to store the data of each replica set member rs0-0, rs0-1, and rs0-2, as well as a directory for logs. Once the directories were created, the mongod.exe command was run to set up each member by specifying different ports (27018, 27019, 27020), binding them to localhost, and setting the data and log paths. Then the 'mongosh.exe' command is run to access the replica set member running on port 27018. To start the replica set, the rs.initiate() command is executed, which specifies the replica set identity ("rs0"), as well as a list of members with their respective IDs and hosts.

Creating a working directory for Replica Set in a NoSQL database cluster such as MongoDB involves several strategic steps to improve performance and ensure data availability. First, determine the optimal number of nodes, usually an odd number (at least three) to ensure quorum. Then, organise the roles and distribution of nodes, making sure there are primary and secondary nodes spread across different physical environments or clouds to avoid single points of failure. Each MongoDB instance requires a data directory to store its data files. Also, the network configuration should allow all nodes to communicate with each other efficiently. Figure 3 is a view of the directory that has been created.

Make sure to install MongoDB correctly before setting up the Replica Set, as shown by the Figure 7. Visit the official MongoDB website and download the MongoDB Community Server version according to the operating system (Windows, macOS, or Linux). Follow the installation instructions provided: for Windows, run the '.msi' installer; for macOS, use Homebrew by running 'brew tap mongodb/brew' and 'brew install mongodb-community'; and for Linux, add the MongoDB repository and install using the appropriate package manager. Once the installation is complete, verify by running the 'mongo -version' command in a terminal or command prompt. Make sure the installed version is compatible with the application requirements.




Figure 7. Path to your mongo shell.

Make sure to install MongoDB Community Server before proceeding to the configuration stage, as shown by the Figure 8. Visit the official MongoDB website and download the MongoDB Community Server corresponding to your operating system (Windows, macOS, or Linux). Follow the installation instructions provided: for Windows, run the '.msi' installer and follow the instructions; for macOS, use Homebrew with the commands 'brew tap mongodb/brew' and 'brew install mongodb-community'; and for Linux, add the MongoDB repository and install using a package manager such as 'apt' for Debian/Ubuntu or 'yum' for CentOS/RHEL. Once the installation is complete, verify by running 'mongo -version' in a terminal or command prompt to ensure the installation was successful and the installed version matches the application requirements.




Figure 8. Path to your mongodb server.

Following the setup of the MongoDB instances, the next critical stage involved configuring the replica set to manage MongoDB replicas effectively. This configuration process determines the roles of each node within the replica set, specifying which nodes will act as primary and secondary nodes. Figure 9 illustrates some of the essential steps involved in this configuration, executed using the MongoDB shell. Configuring the replica set ensures that data replication and synchronization are established correctly among the nodes, thereby enhancing fault tolerance and data availability across the cluster.

```

Replication configuration
mkdir C:\replication\rs0-0
mkdir C:\replication\rs0-1
mkdir C:\replication\rs0-2
mkdir C:\replication\logs

"C:\Program Files\MongoDB\Server\7.0\bin\mongod.exe" --replSet rs0 --port 27018 --bind_ip localhost --dbpath C:\replication\rs0-0 --oplogSize 128 --
logpath C:\replication\logs\rs0-0.log --logappend

"C:\Program Files\MongoDB\Server\7.0\bin\mongod.exe" --replSet rs0 --port 27019 --bind_ip localhost --dbpath C:\replication\rs0-1 --oplogSize 128 --
logpath C:\replication\logs\rs0-1.log --logappend

"C:\Program Files\MongoDB\Server\7.0\bin\mongod.exe" --replSet rs0 --port 27020 --bind_ip localhost --dbpath C:\replication\rs0-2 --oplogSize 128 --
logpath C:\replication\logs\rs0-2.log --logappend

"C:\Program Files\mongosh\mongosh.exe" --port 27018

rs.initiate({
  _id: "rs0",
  members: [
    { _id: 0, host: "localhost:27018" },
    { _id: 1, host: "localhost:27019" },
    { _id: 2, host: "localhost:27020" }
  ]
})

use herbalyze

```

Figure 9. Configure Replica Set in mongo shell.

2. Testing Procedures

In this study, the testing procedures focused on optimizing MongoDB replica set performance through the implementation of read scalability and load distribution strategies. The objective was to enhance database efficiency by distributing read operations among different replica set members based on region tags.




Figure 10. Replica Set Data Flow.

Initially, the current replica set configuration was retrieved using 'rs.conf()' to understand its structure and prepare for modifications. Each member of the replica set was then assigned a specific region tag; for instance, the second member was tagged with "a" and the third with "b." The 'rs.reconfig(cfg)' command was subsequently executed to apply these changes, effectively updating the replica set configuration with the new region tags. Figure 11 illustrates the steps involved in this process.

Edit replica configuration to implement

1. Read Scalability
2. Load Distribution

```

cfg = rs.conf();

cfg.members[1].tags = { "region": "a" };
cfg.members[2].tags = { "region": "b" };

rs.reconfig(cfg);

db.getMongo().setReadPref('secondary', [{ "region": "a" }]);
db.getMongo().setReadPref('secondary', [{ "region": "b" }]);

```

Figure 11. Additional Replica Configuration.

To further optimize read operations, MongoDB read preference settings were adjusted using 'setReadPref('secondary', ["region": "a"])' and 'setReadPref('secondary', ["region": "b"])''. These settings directed read requests to replica set members based on their designated regions, thereby distributing the read load more efficiently across the cluster. This approach not only improved overall read performance but also contributed to better load management and fault tolerance within the MongoDB replica set architecture.

Figure 12 demonstrates the CRUD request testing procedures. The purpose of these tests was to evaluate the performance impact of the implemented configuration changes. By measuring the time taken for Create, Read, Update, and Delete operations across different configurations, the study aimed to quantify the improvements in read scalability and load distribution.

```
var startCreate = new Date();
for (let i = 0; i < 1000; i++) {
  db.plants.insertOne({ name: `Plant ${i}`, type: 'Flower', color: 'Red' });
}
var endCreate = new Date() - startCreate;

var startRead = new Date();
db.plants.find();
var endRead = new Date() - startRead;

var startUpdate = new Date();
for (let i = 0; i < 1000; i++) {
  db.plants.updateOne({ name: `Plant ${i}` }, { $set: { color: 'Blue' } });
}
var endUpdate = new Date() - startUpdate;

var startDelete = new Date();
for (let i = 0; i < 1000; i++) {
  db.plants.deleteOne({ name: `Plant ${i}` });
}
var endDelete = new Date() - startDelete;
```

Figure 12. CRUD request testing.

4. Result and Discussion

4.1. Result

In this study, we compared the performance of a single-instance MongoDB setup with a MongoDB replica set configuration. The replica set consists of one primary node and two secondary nodes. The primary node handles all write operations, while the secondary nodes replicate data from the primary and can be used for read operations, thus enabling read scalability and load distribution. This setup aims to improve the system's overall read performance and provide high availability. For the single-instance setup, all operations are handled by a single MongoDB instance, with no replication or distribution of load.

To evaluate the performance, we conducted a series of tests involving 1000 CRUD (Create, Read, Update, Delete) operations for both configurations. For each operation, we measured the time taken and calculated the average time per operation. These tests were designed to stress the database systems

and observe how each configuration handled the workload, providing insights into their efficiency and responsiveness under typical usage scenarios.

The single-instance MongoDB setup serves as a baseline for performance, where all operations are processed by a single server. This configuration does not benefit from any form of redundancy or load balancing, making it more susceptible to bottlenecks and single points of failure. Conversely, the replica set configuration introduces redundancy and distributes read operations across multiple nodes, which theoretically enhances performance for read-heavy workloads and increases fault tolerance.

Throughout the testing phase, various performance aspects were closely monitored. The focus was on the speed and efficiency of CRUD operations, but we also considered factors such as the consistency of response times, the system's ability to handle concurrent operations, and the overall stability of each configuration under load. By measuring the time taken for each operation and calculating the average time, we aimed to quantify the performance differences in a clear and comparable manner.

4.2. Performance Metrics

In this section, we delve into the performance metrics observed during our testing procedures. The focus was on evaluating the differences between a single-node MongoDB instance and a replica set configuration consisting of one primary and two secondary nodes. This comparison aimed to highlight the trade-offs and benefits associated with each configuration in terms of CRUD operations.

The primary metrics analyzed include the average time taken for create, read, update, and delete operations. These metrics provide a comprehensive view of how each configuration handles different types of database interactions. For the single-node setup, the absence of replication means that all operations are straightforward but lack the benefits of load distribution and redundancy. This simplicity often results in faster write operations due to the lack of overhead involved in synchronizing multiple nodes.

In contrast, the replica set configuration is expected to show improvements in read performance due to the ability to offload read requests to secondary nodes. This capability not only enhances read scalability but also helps in balancing the load across multiple nodes, potentially reducing the load on the primary node and preventing it from becoming a bottleneck. However, write operations in the replica set configuration might exhibit slightly higher latencies compared to the single-node setup due to the additional steps involved in replicating data to secondary nodes and ensuring consistency.

Another critical aspect of the performance evaluation was the consistency and reliability of response times. In a production environment, predictable performance is crucial for maintaining a smooth user experience and ensuring the system's reliability. The replica set's ability to maintain high availability through automated failover mechanisms adds a layer of robustness that is absent in the single-node setup. In the event of a node failure, the replica set can continue to serve read and write requests by promoting a secondary node to primary, thus minimizing downtime.

Our performance metrics also considered the impact of network latency and communication overhead between nodes in the replica set configuration. These factors can influence the overall performance, especially in geographically distributed deployments where nodes are located in different data centers. By understanding these dynamics, we can better assess the suitability of each configuration for various deployment scenarios and workload patterns.


Overall, the detailed analysis of performance metrics provides a nuanced understanding of the strengths and limitations of single-node and replica set configurations. This knowledge is essential for making informed decisions about database architecture and deployment strategies, ensuring that the chosen setup aligns with the specific needs and goals of the application.

4.2.1. Single Node Performance

The single-node MongoDB instance is the most basic configuration where all database operations are handled by a single MongoDB instance without any replication. This setup is straightforward and

often used for development and testing purposes due to its simplicity. The performance metrics for this configuration provide a baseline for comparison with more complex setups such as replica sets.

In our tests, the single-node configuration demonstrated superior performance in write operations. Specifically, the average time taken for Create, Update, and Delete operations was 1.754 seconds, 2.256 seconds, and 1.908 seconds, respectively. These results underscore the efficiency of handling write operations in a non-replicated environment, where the absence of replication overhead allows for faster execution. Figure 13 shows the benchmarking results for the single-node configuration, illustrating the performance across the different types of CRUD operations.



```

herbalyze> // Print Info Avg Time Single Instance MongoDB
herbalyze> print("Waktu rata-rata create 1000 dokumen (Single Instance): " + (endCreate / 1000) + "ms");
Waktu rata-rata create 1000 dokumen (Single Instance): 1.759ms

herbalyze> print("Waktu rata-rata read 1000 dokumen (Single Instance): " + (endRead / 1000) + "ms");
Waktu rata-rata read 1000 dokumen (Single Instance): 0.048ms

herbalyze> print("Waktu rata-rata update 1000 dokumen (Single Instance): " + (endUpdate / 1000) + "ms");
Waktu rata-rata update 1000 dokumen (Single Instance): 2.302ms

herbalyze> print("Waktu rata-rata delete 1000 dokumen (Single Instance): " + (endDelete / 1000) + "ms");
Waktu rata-rata delete 1000 dokumen (Single Instance): 1.826ms

herbalyze>

```


Figure 13. MongoDB Single Node Benchmark.

The single-node setup, however, does not offer the advantages of redundancy and high availability provided by replica sets. In scenarios where read operations are not the primary bottleneck and data availability is less of a concern, a single-node configuration might be sufficient. However, for production environments requiring high availability, data redundancy, and read scalability, the limitations of a single-node setup become apparent.

4.2.2. Replica Set Performance

A MongoDB replica set configuration involves multiple nodes, where one node acts as the primary node handling all write operations, and the remaining nodes act as secondary nodes replicating the data from the primary node. This setup provides several advantages, including data redundancy, fault tolerance, and improved read scalability.

In our testing, the replica set configuration exhibited a different performance profile compared to the single-node setup. The average times for Create, Update, and Delete operations were 3.984 seconds, 4.754 seconds, and 4.399 seconds, respectively. These times are higher than those observed in the single-node setup, primarily due to the replication overhead where write operations need to be propagated to the secondary nodes. Figure 14 illustrates the performance metrics for the replica set, showing the impact of replication on write operations.



```

mongosh mongodb://127.0.0.1:27018/?directConnection=true&serverSelectionTimeoutMS=2000
rs0 [direct: primary] herbalyze> // Print Info Avg Time Replica Set MongoDB
rs0 [direct: primary] herbalyze> print("Waktu rata-rata create 1000 dokumen (Replica Set): " + (endCreate / 1000) + "ms");
Waktu rata-rata create 1000 dokumen (Replica Set): 3.984ms

rs0 [direct: primary] herbalyze> print("Waktu rata-rata read 1000 dokumen (Replica Set): " + (endRead / 1000) + "ms");
Waktu rata-rata read 1000 dokumen (Replica Set): 0.038ms

rs0 [direct: primary] herbalyze> print("Waktu rata-rata update 1000 dokumen (Replica Set): " + (endUpdate / 1000) + "ms");
Waktu rata-rata update 1000 dokumen (Replica Set): 4.754ms

rs0 [direct: primary] herbalyze> print("Waktu rata-rata delete 1000 dokumen (Replica Set): " + (endDelete / 1000) + "ms");
Waktu rata-rata delete 1000 dokumen (Replica Set): 4.399ms

rs0 [direct: primary] herbalyze>

```

Figure 14. MongoDB Replica Set Benchmark.

However, the replica set configuration significantly outperformed the single-node setup in read operations. The average time for read operations in the replica set was 0.038 seconds, compared to 0.047 seconds for the single-node setup. This improved read performance can be attributed to the read scalability and load distribution mechanisms inherent in a replica set configuration. By distributing read operations across secondary nodes, the replica set can handle a higher volume of read requests more efficiently, reducing the load on the primary node and enhancing overall read performance.

Table 2 provides a comparative overview of the average times for each CRUD operation in both the replica set and single-node configurations. The table highlights the trade-offs between the two setups: while the single-node configuration excels in write operations due to the lack of replication overhead, the replica set configuration offers superior read performance and enhanced data redundancy.

Table 2. Average Time Comparison between Replica Set and Single Instance Node for 1000 CRUD Operations.

Operation	Replica Set (s)	Single Node (s)
Create	3.984	1.754
Read	0.038	0.047
Update	4.754	2.256
Delete	4.399	1.908

The results underscore the importance of selecting the appropriate MongoDB configuration based on specific application requirements and workload characteristics. For applications with heavy write operations where latency is a critical factor, a single-node configuration might be preferable. Conversely, for applications with heavy read operations and a need for high availability, the replica set configuration provides significant advantages.

Moreover, the replica set's ability to provide automatic failover ensures that in the event of a primary node failure, one of the secondary nodes can be automatically promoted to primary, maintaining the availability of the database without manual intervention. This capability is crucial for applications that require continuous availability and cannot afford downtime.

In conclusion, our performance testing revealed that while single-instance MongoDB is advantageous for write-heavy applications due to its lower latency in write operations, the replica set configuration excels in read-heavy scenarios, thanks to its read scalability and load distribution features. The replica set not only enhances read performance but also ensures high availability and data redundancy, making it a robust solution for critical applications requiring consistent read access and fault tolerance. Future work could explore further optimizations in replica set configurations and additional testing with different dataset sizes and operation types to provide deeper insights into MongoDB performance under various conditions.

Furthermore, understanding the specific needs of the application environment is essential for making informed decisions about database architecture. As data volumes and access patterns evolve, ongoing evaluation and tuning of database configurations will be necessary to maintain optimal performance and reliability.

By leveraging the strengths of each configuration and carefully balancing the trade-offs, database administrators and developers can achieve a scalable, reliable, and high-performance MongoDB deployment that meets the demands of modern applications.

4.3. Discussion

The results underscore the trade-offs between single-instance and replica set configurations. While the single-instance MongoDB provides faster write operations due to the absence of replication overhead, the replica set offers better read performance and increased data availability. The ability to distribute read operations across multiple nodes allows the replica set to handle higher read loads efficiently. This makes the replica set a more suitable choice for applications with heavy read operations and a need for high availability, despite the higher latency in write operations. These findings highlight

the importance of selecting the appropriate MongoDB configuration based on specific application requirements and workload characteristics.

In conclusion, our performance testing revealed that while single-instance MongoDB is advantageous for write-heavy applications due to its lower latency in write operations, the replica set configuration excels in read-heavy scenarios, thanks to its read scalability and load distribution features. The replica set not only enhances read performance but also ensures high availability and data redundancy, making it a robust solution for critical applications requiring consistent read access and fault tolerance. Future work could explore further optimizations in replica set configurations and additional testing with different dataset sizes and operation types to provide deeper insights into MongoDB performance under various conditions.

The observed trade-offs between single-instance and replica set configurations in MongoDB underscore the necessity of a nuanced approach to database management. For developers and system architects, understanding these differences is crucial for optimizing application performance and reliability. Single-instance MongoDB configurations offer lower write latencies, making them ideal for use cases where write speed is paramount, such as logging systems or real-time analytics. However, the lack of redundancy and potential for data loss in single-instance setups can be a significant drawback for applications where data integrity and availability are critical, emphasizing the importance of aligning database architecture with specific application demands.

On the other hand, the benefits of replica sets extend beyond mere performance metrics. The enhanced fault tolerance provided by replica sets ensures that applications remain operational even in the face of hardware failures or network issues. This resilience is particularly valuable for applications requiring continuous availability, such as e-commerce platforms, online gaming services, and financial transaction systems. Additionally, the ability to perform maintenance and upgrades with minimal downtime is a significant advantage in dynamic production environments. Future research should focus on fine-tuning replication algorithms and exploring hybrid configurations that balance write and read performance, further optimizing MongoDB for diverse and demanding workloads.

Moreover, exploring advanced techniques such as sharding in conjunction with replica sets could provide a comprehensive solution for scaling both horizontally and vertically, addressing the limitations of each approach when used independently. Sharding can distribute data across multiple replica sets, thereby balancing the load more effectively and ensuring that neither read nor write operations become bottlenecks. Integrating automated monitoring and scaling tools can further enhance the adaptability of MongoDB clusters, allowing them to respond dynamically to varying workloads. By continuously refining these strategies, organizations can achieve a robust, scalable, and highly available database infrastructure, capable of meeting the evolving demands of modern applications.

5. Conclusion

In this study, we conducted a comprehensive analysis of MongoDB performance by comparing single-node and replica set configurations. Our objective was to identify the strengths and weaknesses of each setup, particularly in handling CRUD operations, and to provide insights into their suitability for different application scenarios.

The single-node MongoDB configuration demonstrated remarkable efficiency in write operations. The absence of replication overhead allowed for faster execution of Create, Update, and Delete operations, with average times of 1.754 seconds, 2.256 seconds, and 1.908 seconds, respectively. This makes the single-node setup an attractive option for applications that are write-intensive and where write latency is a critical concern. However, the lack of redundancy and high availability in a single-node configuration poses significant risks for production environments, as any node failure can lead to data unavailability and potential data loss.

In contrast, the MongoDB replica set configuration, comprising one primary node and two secondary nodes, showcased superior performance in read operations. With an average read time

of 0.038 seconds compared to 0.047 seconds in the single-node setup, the replica set leveraged read scalability and load distribution to enhance read efficiency. This configuration also ensured data redundancy and fault tolerance, as secondary nodes replicated data from the primary node and could be promoted to primary in case of failure. The trade-off, however, was higher latency in write operations due to the replication process, with average times of 3.984 seconds for Create, 4.754 seconds for Update, and 4.399 seconds for Delete operations.

Our findings highlight the critical importance of selecting the appropriate MongoDB configuration based on specific application requirements. For scenarios where high write throughput and low write latency are paramount, a single-node configuration may be preferable. On the other hand, applications that prioritize read performance, data redundancy, and high availability would benefit significantly from a replica set configuration.

Moreover, the implementation of read scalability and load distribution strategies within the replica set further optimized its performance. By distributing read operations across secondary nodes based on predefined region tags, we were able to achieve efficient load balancing and improved read performance. This approach is particularly beneficial for applications with geographically distributed users, as it allows for regional data access optimization and reduced latency.

Looking ahead, there are several avenues for future research and optimization. Exploring advanced replication strategies, such as sharded clusters combined with replica sets, could provide deeper insights into handling large-scale data and high-traffic applications. Additionally, testing with different dataset sizes, varying read-write ratios, and diverse workload patterns would offer a more comprehensive understanding of MongoDB performance under various conditions.

Furthermore, the integration of performance monitoring and tuning tools can aid in proactive database management. Tools that provide real-time insights into query performance, replication lag, and resource utilization can help database administrators identify bottlenecks and optimize configurations dynamically.

In conclusion, our study underscores the importance of aligning database configurations with application needs. By leveraging the strengths of single-node and replica set configurations, and understanding their respective trade-offs, organizations can design scalable, reliable, and high-performance MongoDB deployments that cater to their unique requirements. The continuous evolution of database technologies and best practices will further enhance the capabilities of MongoDB, empowering developers to build robust and efficient data-driven applications.

Acknowledgments: The authors wish to express their sincere gratitude to the Informatics Department at UIN Sunan Gunung Djati Bandung for their partial support and encouragement throughout this research project. Their invaluable resources, guidance, and facilities played a crucial role in the successful completion of this study.

We would like to extend our heartfelt thanks to our advisor, Dr. Wisnu Uriawan, M.Kom. for his invaluable guidance, insightful feedback, and continuous encouragement. Without his expertise and patience, this work would not have been possible.

Special thanks are also extended to the faculty members and peers who provided insightful feedback and encouragement throughout the research process. Their contributions have been instrumental in refining our methodologies and enhancing the overall quality of this work.

Additionally, we would like to express our gratitude to the technical support teams who assisted in setting up the testing environments and troubleshooting issues that arose during the study. Their expertise and dedication ensured that our experiments were conducted smoothly and efficiently.

We are also deeply grateful to our families and friends for their unwavering support and understanding throughout this journey. Their belief in us has been a constant source of strength and motivation.

Special thanks go to our colleagues and fellow students for their valuable discussions and assistance. Our camaraderie and collaboration have greatly enriched this experience.

References

1. K. Tabet, R. Mokadem, and M. Laouar, "A data replication strategy for document-oriented nosql systems," *International Journal of Grid and Utility Computing*, vol. 10, p. 53, 01 2019.
2. J. NOVOTNÝ, "Automating performance testing and infrastructure deployment for debezium," Master's thesis, Masaryk University, 2023.

3. L. F. Da Silva and J. V. Lima, "An evaluation of relational and nosql distributed databases on a low-power cluster," *The Journal of Supercomputing*, vol. 79, no. 12, pp. 13 402–13 420, 2023.
4. E. Tang and Y. Fan, "Performance comparison between five nosql databases," in *2016 7th International Conference on Cloud Computing and Big Data (CCBD)*. IEEE, 2016, pp. 105–109.
5. R. Osman and P. Piazzolla, "Modelling replication in nosql datastores," in *International Conference on Quantitative Evaluation of Systems*. Springer, 2014, pp. 194–209.
6. X. Huang, J. Wang, J. Qiao, L. Zheng, J. Zhang, and R. K. Wong, "Performance and replica consistency simulation for quorum-based nosql system cassandra," in *Application and Theory of Petri Nets and Concurrency: 38th International Conference, PETRI NETS 2017, Zaragoza, Spain, June 25–30, 2017, Proceedings 38*. Springer, 2017, pp. 78–98.
7. N. Aemy and A. Rahmatulloh, "Implementasi halb dan klaster mongodb dengan penyimpanan cache redis dalam sistem terdistribusi," *JUSTIN (Jurnal Sistem dan Teknologi Informasi)*, vol. 12, no. 2, pp. 265–270, 2024.
8. A. Makris, K. Tserpes, G. Spiliopoulos, D. Zissis, and D. Anagnostopoulos, "Mongodb vs postgresql: A comparative study on performance aspects," *GeoInformatica*, vol. 25, pp. 243–268, 2021.
9. S. Zhou and S. Mu, "Fault-Tolerant replication with Pull-Based consensus in MongoDB," in *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*. USENIX Association, Apr. 2021, pp. 687–703. [Online]. Available: <https://www.usenix.org/conference/nsdi21/presentation/zhou>
10. M. Stonebraker, "Sql databases v. nosql databases," *Communications of the ACM*, vol. 53, no. 4, pp. 10–11, 2010.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.