

# Challenges and Strategies for Consolidating Data across Microservices

**Anju Bhole**

Independent Researcher, California, USA  
anjusbhole@gmail.com

## **Abstract**

The emergence of the microservices architecture has, by decoupling big systems into small manageable services, fundamentally changed the way we write and run large-scale applications. However, although this architecture provides great freedom as well as scale on demand, it also brings with it, severe problems in the area of maintaining coherence or aggregating data when working with services. In a microservices ecosystem, the spreading out of data often generates data consistency problems, sync difficulties and a lack of transactional integrity. This paper addresses these challenges and discusses some strategies to ameliorate them, focusing specifically on ensuring the consistency of data, increasing throughput and attaining scalable solutions. Drawing upon current research, practical experience in the industry, and new technologies that have become available such as event-driven architecture, service mesh and data synchronization mechanisms, this paper provides valuable insights for those who are moving to microservice-based systems. In addition, the paper introduces best practices, including use of such technologies to engage in event-driven communication as Apache Kafka (if you want data consistency) and adopting service mesh solutions as Istio (in order to better understand how to flow data and control the various service interactions). This paper provides an all-around strategy for solving data consolidation challenges in microservices environments.

**Keywords:** Microservices, Data Consolidation, Data Consistency, Scalability, Performance, Service Mesh, Event-Driven Architecture, Data Integration

## **Introduction:**

With organizations today increasingly demanding need to deploy large-scale applications that are agile, scalable and maintainable, the microservice architecture has rapidly won favor in the past few years. Microservices break monolithic systems down into small, independently deployable services delivering increased flexibility, faster development cycles, and improved fault isolation. Each service is responsible for a specific business capability and can be developed, deployed, or scaled independently of the others. This allows enterprises to respond quickly to market demands or changes in demand. The move to a distributed system of microservices, however, is not without its challenges particularly, in terms of managing and consolidating data across multiple services.

In a microservices-based architecture, each microservice owns its data. This data often lives in separate databases or storage systems associated with the services. While this approach supports the independence of services, it complicates the integration of data across services and creates issues

related to consistency, synchronization, and transactional integrity. In this decentralized environment, traditional methods such as monolithic databases and centralized data management systems no longer work. As a result, data consistency across services becomes increasingly complex, and the lack of effective strategies for data consolidation can lead to problems like data duplication, performance bottlenecks, and system failures.

This paper seeks to examine the challenges and impacts of these issues, as well as proposing strategies and tools which could systematically address or avoid them. It analyzes current academic results, industrial practice and cutting-edge technology in order to provide valid recommendations applicable to organizations faced with the issue of managing data within a microservice world. In addition, we discuss evolving technologies such as event-driven architecture, data synchronization mechanisms and the importance of service mesh solutions in ensuring seamless data integration and effective management across distributed services.

### **Research Aim:**

Throughout the course of this research, we strive to fully uncover the various challenges associated with consolidating and managing of data across microservices focusing on issues generally centered around data consistency, synchronization and scalability. Furthermore, we will assess the benefits of various strategies adopted to overcome these challenges and enhance data integration.

### **Research Objectives:**

1. To identify the key challenges in consolidating data across microservices.
2. To evaluate the effectiveness of current strategies for data synchronization and consolidation.
3. To examine emerging technologies and techniques that can assist in data management for microservices.
4. To propose a framework for consolidating data across microservices, considering scalability, performance, and consistency.

### **Research Questions:**

1. What are the main challenges organizations face when consolidating data across microservices?
2. How effective are current strategies for addressing these challenges?
3. What emerging technologies and methodologies can enhance data consolidation across microservices?
4. How can a framework for data consolidation across microservices be designed to ensure high performance and scalability?

### **Problem Statement:**

Microservices' increasingly rapid adoption introduced significant challenges in managing data consistency, integrity, and integration across disparate services. With each microservice owning its own data and relying on APIs for communication with other services, managing all-round data becomes more and more complex. Existing methods have trouble handling things like data redundancy, inconsistency and bottlenecks in performance as the number of services grows. In a

word, there is a need for effective techniques to consolidate data across microservices, ensuring consistency without compromising performance or scalability.

### **Literature Review:**

The acceleration in recent years of microservices has led to the present-day situation in which data being shared between multiple systems is one of main problems for modern software. Because these microservice systems, increasingly popular for their flexibility and scalability, have their distributed data growing more and complex to manage. This section provides a survey of different methods, methodologies and tools that have been proposed for dealing with various issues of data consistency, synchronization, integration in a microservices environment on the one hand; it also emphasizes the importance behind transaction management; and introduces fresh technologies such as service meshes and event-driven architectures.

### **Challenges in Data Consolidation across Microservices**

An inherently decentralized microservices architecture has each database managed by its related microservice. This approach enhances the independence and scalability of every service but introduces data management challenges. A big challenge is to ensure data consistency across various services that may all be working on different databases. In a monolithic system, data consistency is relatively easy to achieve because the database structure is centrally arranged. But in a microservices setup it becomes much more difficult as the data is distributed.

The core issues in data consolidation across microservices revolve around the following challenges:

1. **Data Consistency:** Data consistency is difficult to maintain across multiple services. Since each microservice is responsible for its own data, this becomes a problem of even greater magnitude when data has to be shared between them. In this case, traditional ACID (Atomicity, Consistency, Isolation, Durability) principles are often not directly applicable and way to achieve consistency in a distributed context without sacrificing performance is a major problem.
2. **Synchronization:** Maintaining data synchronized across microservices can be complex, especially when services are frequently updated or experience network issues. Inconsistent or delayed data synchronization can result in stale or incorrect data being presented to users.
3. **Transactional Integrity:** Handling transactions across distributed services is one of the most difficult aspects of data consolidation. In a monolithic application, a single database transaction can guarantee consistency, but in a microservices-based system, each service typically manages its own database and transaction boundaries, making it challenging to ensure atomicity and isolation.

### **Approaches to Address Data Consolidation Challenges**

To tackle these issues, researchers and practitioners have proposed several strategies and architectural patterns that enable better data management across microservices.

#### ***The Saga Pattern***

In the world of microservices, one popular way to deal with distributed transactions is the Saga pattern. A saga, a sequence of local transactions that are coordinated by means of events or commands. Each service involved in the saga performs its transaction and publishes an event to notify other services, which then respond accordingly. If any part of the saga malfunctions, then compensation actions are triggered to roll back the transaction and keep your data consistent.

Saga patterns can decouple services and also help attain data consistency through a variety of smaller, more manageable transactions. It supports eventual consistency, something that is good for many business processes where immediate consistency has no real importance. According to Sharma et al. (2020), the Saga pattern can greatly reduce the complexity of distributed transactions because it eliminates centralized coordination.

### ***Event-Driven Architecture (EDA)***

Event-driven architecture (EDA) is another important solution to solve the data synchronization problems in microservices. In EDA, services talk to each other asynchronously by publishing and subscribing to events. This approach allows microservices to respond to changes in state or data without being directly tied together. It enables services to fail more gracefully (i.e., without taking down other services) and supports decoupling.

Event-driven architecture improves scalability, flexibility and fault tolerance. Because services react only to events every now and then, they do not have to keep checking for updates, thus saving unnecessary load and improving performance. Event-driven systems also make it easier for us to manage state changes and to propagate these changes to other services. According to Pimentel et al. (2021), EDA can greatly simplify data synchronization by ensuring that microservices continue to be in sync regardless of whether they are actively communicating or not.

### ***Event Sourcing and CQRS***

The architectural pattern of event sourcing means storing the state of a system as a series of immutable events. Instead of simply reflecting the current state, event sourcing saves a record of every change that brought about the current state. This technology is good for rebuilding the state of microservices through the historical events and ensure consistency as well.

Event sourcing goes very well with Command Query Responsibility Segregation (CQRS), a design pattern that separates handling commands (which modify state) from queries, those that simply read state. Here, by applying CQRS with event sourcing, microservices can optimize the reading and writing of data; this results in more efficient data management. It reduces the complexity of managing distributed state changes and can improve performance by unloading data retrieval operations.

Both event sourcing and CQRS are popular in microservices environment because they separate the constituent parts of a system, which not only helps increase scalability and maintainability but makes them more suitable where the system state is complicated.

### **The Role of Service Meshes in Data Management**

Recently, Service mesh technologies have emerged as a critical tool for managing the inefficiencies that can arise from communication and data movement in microservices. Service meshes like Istio

and Linkerd add a dedicated infrastructure layer that manages service-to-service communication, allow fine control over features like traffic management, monitoring, service discovery, and security, which are necessary for the smooth performance of communication activities in microservice eco systems.

One of the key benefits of implementing a service mesh is its ability to govern and monitor the data flows between different services. That is crucial when managing data in a microservices system. By using a service mesh, organizations can enforce policies and set up monitoring mechanisms to ensure that data is transmitted in a reliable and secure manner, thereby improving the synchronization and aggregation of data.

### **Insights from Recent Studies**

We have recently started focusing our attention on better transaction management, ensuring fewer data errors between microservices. Adding to this, Sharma et al. (2020) point out that in managing distributed transactions among microservices we face very significant challenges. Its emphasis lies in consistency "mechanisms" like the saga pattern, and eventual consistency of data. In their study, Sharma et al. explore how microservices can manage long-running transactions while maintaining data integrity, without any negative effect on performance.

Pimentel et al. (2021) similarly discuss event-driven architectures for its potential benefits to data consolidation. They point out that use events as the major mode of communication between services can ease synchronization of data and lower dependency between services. This in turn reduces complexity in handling distributed data. They also consider the contribution event-driven systems make to reliable data exchange between services even when network failures or late communications are happening.

### **Emerging Tools and Technologies**

Moreover, in addition to the architectural patterns a number of new tools and technologies are making it possible to manage data consolidation effectively on microservices architectures. Technologies like Kafka, RabbitMQ and Amazon SNS/SQS can be used for messaging and event streaming so that services can respond as soon as data changes occur. These tools support efficient event propagation among services and help to maintain consistency without the need for tight coupling. Also, cloud-native solutions, such as Kubernetes and container orchestration tools play an important role in managing the deployment of microservices and ensuring that data synchronization mechanisms can scale. They provide the infrastructure required to deal with complex interactions between services and massive scale distributed data.

The consolidation of data across microservices is a challenging task that requires careful consideration of consistency models, synchronization mechanisms, and transaction management. Approaches such as the Saga pattern, event-driven architecture, and event sourcing, combined together with the application of service mesh technologies offer effective solutions to these challenges. New research underlines how important these strategies have been in maintaining integrity of data, scalability, and fault tolerance in systems based on microservices. As organizations continue to adopt microservices, further research and advancements in tools and

methodologies will be critical in refining data management strategies and enabling the successful deployment of large-scale microservice systems.

### **Research Methodology:**

The present research takes a qualitative approach, investigating the difficulties and methods of consolidating data from a microservice-oriented perspective. To this end, its main emphasis is to analyze existing literature, evaluate industry practices, and investigate new technology within the context of dealing with data synchronization and management in microservices oriented environments. By reviewing academic papers, industrial reports and case studies in depth, this study hopes to provide a comprehensive understanding of where we stand today on the issue of data consolidation in microservices.

A key component is a comparison, where different ways and policies adopted by organizations in addressing data integration and consistency issues of microservices architecture have been analyzed. This includes various strategies like event-driven architecture, the Saga pattern, and service mesh solutions, etc. The research then assesses these methods' effectiveness through real-world applications revealing the benefits, constraints and practical problems inherent. The research also includes the exploration of case studies. These show how organizations have deployed microservice architectures, what problems they encountered when dealing with data, and how they solved those problems. In addition to its in-depth analysis of these case studies, the work is supplemented with interviews from industry insiders and surveys carried out among businesses currently using microservices. These primary sources provide first-hand observations by professionals who are involved with services on a daily basis, revealing their approaches for integrating data across microservices and their perception of the best methods.

Secondary data is also a crucial part of this research. Academic papers, industry reports and best practices from large technology vendors, offer a strong basis for understanding the theoretical frameworks behind microservices and emerging trends in data management. By synthesizing both primary and secondary data, this research aims to provide a comprehensive view of the challenges and solutions in the area of data consolidation across services, making this much-needed bridge between academic research literature and that of quick moving industry developments.

### **Results and Discussion:**

The research identified several critical challenges in consolidating data across microservices, as well as strategies that have been found effective in overcoming these challenges. The findings emphasize the complexity of managing data consistency, synchronization, and transactional integrity across distributed services. This section discusses the key challenges identified, along with the strategies and emerging tools that address these issues.

#### **Key Challenges in Data Consolidation**

##### **1. Data Duplication:**

Data duplication is a common problem in microservices architectures. Since each microservice typically owns its own data and operates with its own database, there is often redundancy when the same data needs to be shared across services. This duplication can lead to inconsistencies and data management overhead, as changes made in one service need to be

propagated to others to maintain consistency. As a result, ensuring that each service works with the most current and accurate data becomes a significant challenge.

**2. Inconsistent State Across Services:**

In a microservices-based environment, services operate independently and maintain their own state. However, maintaining a consistent state across all services is difficult, particularly when data is being updated concurrently. For example, a change in one service may not be immediately reflected in other services, leading to inconsistencies that can cause errors in business logic, as well as in data retrieval and reporting.

**3. Transactional Guarantees:**

One of the most significant challenges in microservices architectures is implementing robust transactional guarantees. Traditional ACID (Atomicity, Consistency, Isolation, Durability) transactions, which work well in monolithic systems, do not scale effectively across microservices. This is because each service manages its own database and maintaining consistency across distributed transactions requires careful coordination. Ensuring that data changes are atomic and consistent across services while maintaining the performance and scalability of the system is a key issue.

**Effective Strategies for Data Consolidation**


Several strategies have been identified to address these challenges and enable effective data consolidation across microservices.

***Event-Driven Architectures (EDA)***

Event-driven architectures (EDA) have emerged as an effective solution to manage data synchronization across microservices. In an EDA, services communicate asynchronously by publishing and subscribing to events, which helps to decouple the services and improve the flexibility and scalability of the system. By using event-driven messaging systems like Kafka and RabbitMQ, microservices can react to changes in data without being tightly coupled.

Figure 1 below illustrates the flow of events in an event-driven architecture:

**Figure 1: Event-Driven Architecture Flow**



In this example, when Service A produces an event, Services B and C react by processing the event and updating their states accordingly. This decoupling of services ensures that microservices can stay in sync even without direct interactions, which in turn minimizes the risks of inconsistent data states.

***Eventual Consistency and the Saga Pattern***


In order to manage distributed transactions and maintain consistency of data across microservices, the Saga pattern is adopted in general. One primary reason is that unlike traditional distributed transactions, which require a single coordinator, the Saga pattern divides the transaction handling into a series of small, local transactions. If a transaction fails, compensating actions are triggered to undo any changes made during the transaction. This pattern supports eventual consistency models, that allow more scalable and a more robust system.

The Saga pattern gives long-running transactions a way to persist and remain recoverable over time without locking down resources too long. In contrast to immediate consistency across services, the eventual model ensures that all services will eventually reach a stable state and hence is suited well to many business processes which don't have strict consistency requirements.

**Service Mesh Technologies**

Emerging technologies, such as service meshes, increasingly play a key role in ensuring efficient communication between microservices but still addressing data management problem. Service meshes such as Istio provide a dedicated infrastructure layer to handle communication between services in a microservices architecture. Offering features such as traffic management, service discovery, load balancing, security and observability.

**Figure 2: Service Mesh Architecture**



As such, by using a service mesh like Istio, an enterprise can ensure data exchange between services is both safe and efficient. Istio enables fine-grained control over inter-service communication, including features like rate limiting and retries, which help in managing the data flow and minimizing latency. Moreover, service meshes provide full observability, enabling organizations to track and analyze the interactions between microservices, which is critical when consolidating data across distributed services.

**Cloud-Native Solutions: AWS and Kubernetes**

Cloud-native platforms have also increased the scalability of microservice architectures, making it easier to co-consolidate data between services. As a container orchestration tool, Kubernetes makes it possible for organizations to manage and scale microservices effectively.

Kubernetes enables organizations to use automated deployment and scaling to manage multiple containers across thousands of hosts, facilitating the deployment of microservices at scale without the overhead of having to manage infrastructure. A range of services from AWS e.g., Amazon RDS for database management, and Amazon SQS for messaging help unify data across microservices. AWS can further offer services like Serverless computing (AWS Lambda), which can trigger events based on changes in data and seamlessly interface between services.



The scalability provided by Kubernetes and AWS allows organizations to manage large-scale microservice architectures, ensuring that data consolidation strategies, such as event-driven architectures and service meshes, can be implemented effectively across a global infrastructure.

### **Discussion:**

The research highlights that consolidating data across microservices is a complex challenge due to data duplication, inconsistent states, and the difficulty of implementing distributed transactional guarantees. However, several strategies, such as event-driven architectures, the Saga pattern, and service meshes, have proven effective in addressing these challenges. Furthermore, cloud-native solutions like AWS and Kubernetes provide the scalability required to support large-scale microservice ecosystems. As organizations continue to adopt microservices, these strategies and tools will play a crucial role in ensuring efficient data management, helping businesses achieve improved performance, scalability, and resilience in their microservices architectures.

### **Conclusion:**

This paper provides a thorough examination of the challenges involved in consolidating data across microservices and presents several strategies that can effectively address these issues. As organizations increasingly adopt microservices architectures to enhance flexibility, scalability, and efficiency, managing data across distributed services has become a significant challenge. The decentralization of data, combined with the need for consistency and synchronization, makes it difficult to ensure that services operate on accurate and up-to-date information. This paper has explored several approaches, including event-driven architectures, the Saga pattern, and eventual consistency models, which provide viable solutions to these challenges.

Event-driven architectures, supported by tools such as Kafka and RabbitMQ, offer a way to decouple services while ensuring that they stay synchronized by reacting to changes in data. Additionally, the Saga pattern helps manage distributed transactions, providing a way to ensure eventual consistency without the need for centralized coordination. Service meshes like Istio further enhance the management of inter-service communication, improving data flow and reducing latency in large-scale environments. Together, these strategies form a robust framework for consolidating data across microservices.

As the adoption of microservices continues to grow, future research should focus on further refining these strategies and exploring new methodologies to improve data management. One promising direction for future research is the integration of artificial intelligence (AI) and machine learning (ML) to automate and optimize data synchronization processes. AI and ML could be employed to predict data patterns, detect inconsistencies, and autonomously handle data consolidation tasks, which would significantly improve the efficiency and scalability of microservices ecosystems. Exploring these advanced techniques will likely pave the way for even more sophisticated solutions in data management for microservices.

### **Future Scope of Research:**

Future research can explore the integration of advanced AI techniques for automating data consolidation tasks. Machine learning algorithms could be employed to predict data patterns and

automate synchronization between services. Additionally, further exploration into hybrid data management models, combining strong consistency with eventual consistency, could provide insights into more flexible and scalable solutions for large-scale applications.

### **References:**

- [1] Sharma, A., et al. (2020). "Managing Distributed Transactions in Microservices," *IEEE Transactions on Cloud Computing*, vol. 8, no. 3, pp. 1435-1447.
- [2] Pimentel, J., et al. (2021). "Event-Driven Architectures for Microservices: A Survey," *Journal of Software Engineering*, vol. 33, no. 4, pp. 350-367.
- [3] Fowler, M. (2019). "Patterns of Enterprise Application Architecture," Addison-Wesley.
- [4] Mernik, M., et al. (2020). "Challenges in Microservices Data Integration: A Literature Review," *Journal of Cloud Computing*, vol. 9, pp. 50-62.
- [5] Tzoumas, J., et al. (2018). "A Survey on the Use of Event-Driven Architectures in Microservices," *Proceedings of the 11th ACM International Conference on Distributed and Event-Based Systems*, pp. 125-133.
- [6] Behrens, A., et al. (2020). "Microservices and Data Consistency: Challenges and Approaches," *Journal of Cloud Computing*, vol. 8, no. 2, pp. 75-84.
- [7] Hassan, A., et al. (2020). "A Comparative Study of Service Mesh Technologies in Microservices Architectures," *IEEE Transactions on Services Computing*, vol. 13, no. 4, pp. 1023-1035.
- [8] Stepanov, S., et al. (2021). "Event Sourcing in Microservices: Benefits and Challenges," *Proceedings of the 2021 IEEE International Conference on Cloud Computing and Services Science*, pp. 206-213.
- [9] Vasiliadis, S., et al. (2021). "The Saga Pattern: Distributed Transactions for Microservices," *Springer International Publishing*, pp. 185-210.
- [10] Krishna, A., et al. (2020). "Analyzing the Impact of Service Mesh on Performance in Microservices Architecture," *Proceedings of the 2020 IEEE International Conference on Cloud Computing*, pp. 450-456.
- [11] Nunes, G., et al. (2022). "Using Kafka for Event-Driven Microservices," *IEEE Software*, vol. 39, no. 5, pp. 43-51.
- [12] Zhang, Y., et al. (2021). "Scalable Event-Driven Architectures in Microservices," *IEEE Access*, vol. 9, pp. 12580-12590.
- [13] He, L., & Zhang, P. (2021). "Distributed Data Management Techniques in Microservices," *Future Generation Computer Systems*, vol. 113, pp. 122-132.
- [14] Jiang, X., et al. (2022). "Challenges in Handling Data Consistency in Microservices-Based Systems," *International Journal of Cloud Computing and Services Science*, vol. 11, no. 2, pp. 91-104.
- [15] Nguyen, H., et al. (2020). "Integration of Microservices and Event-Driven Systems: A Systematic Review," *IEEE Transactions on Software Engineering*, vol. 46, no. 4, pp. 487-501.