

# Journal of Artificial Intelligence, Machine Learning and Data Science

<https://urfpublishers.com/journal/artificial-intelligence>

Vol: 1 & Iss: 1

Research Article

## SAGA and CQRS Implementation Techniques for Distributed Transaction Management

Sriram Ghanta\*

**Citation:** Ghanta S. SAGA and CQRS Implementation Techniques for Distributed Transaction Management. *J Artif Intell Mach Learn & Data Sci* 2018 1(1), 3203-3208. DOI: doi.org/10.51219/JAIMLD/sriram-ghanta/650

Received: 02 June, 2018; Accepted: 18 June, 2018; Published: 20 June, 2018

\*Corresponding author: Sriram Ghanta, MTS III Consultant, India

**Copyright:** © 2018 Ghanta S., This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

### ABSTRACT

Modern distributed systems particularly those built using microservices must coordinate business operations that span multiple independently deployed services, each maintaining its own database and runtime environment. Achieving reliable transactional behavior across these boundaries, without introducing tight coupling or centralized control, presents a fundamental architectural challenge. Traditional ACID-based distributed transaction mechanisms like two-phase commit (2PC) depend on global locks and synchronous communication, which degrade performance, limit elasticity and reduce system availability as services scale horizontally. To address these limitations, system architects increasingly rely on the Saga pattern, which breaks a long-running business process into a series of autonomous local transactions coordinated either through a central orchestrator or through decentralized event choreography, with compensating actions to handle failures. In parallel, Command Query Responsibility Segregation (CQRS), combined with Event Sourcing, provides a scalable way to handle data consistency by separating write operations from read models and persisting state changes as immutable events. Together, Saga and CQRS/Event Sourcing offer complementary strategies that enable distributed applications to maintain integrity, support eventual consistency and achieve high levels of resilience and scalability in complex transactional workflows.

**Keywords:** Distributed transactions, Microservices, Saga pattern, CQRS, Event sourcing, Asynchronous messaging, Compensating transactions, Eventual consistency

### 1. Introduction

The rapid adoption of microservices architecture has fundamentally reshaped the design of modern enterprise systems. Unlike traditional monolithic applications that rely on a centralized relational database, microservices-based systems are decomposed into small, independently deployable services, each responsible for its own data and business logic. This decomposition enhances modularity, scalability and maintainability, enabling organizations to evolve their systems more flexibly and respond to changing business needs. However, distributing functionality and data ownership across multiple services introduces significant challenges in maintaining

transactional integrity. Business operations frequently span multiple microservices, but traditional distributed transaction protocols most notably two-phase commit (2PC) are ill-suited to this environment. These protocols require strict coordination, synchronous communication and long-held locks, which degrade system responsiveness, limit scalability and weaken the inherent fault-tolerant characteristics expected from microservices.

The Saga pattern originally proposed for managing long-running transactions in distributed database systems, offers a more suitable alternative for microservices. Instead of relying on global locks, Saga decomposes a large business transaction into a series of smaller, autonomous local transactions coordinated through

messages or events. If any part of the workflow fails, previously completed steps are reversed using compensating transactions, achieving eventual consistency without requiring a centralized transaction manager. Complementing this, the Command Query Responsibility Segregation (CQRS) pattern frequently paired with Event Sourcing provides a robust mechanism for handling state changes in distributed environments. CQRS separates write operations (commands) from read operations (queries), allowing each side to scale independently and enabling denormalized models optimized for performance.

Event Sourcing further strengthens this model by persisting all state changes as immutable events, improving auditability, reproducibility and system transparency. Together, CQRS and Event Sourcing address the challenges of maintaining consistent and reliable state across microservices while preserving scalability and high availability. This article examines how Saga and CQRS/Event Sourcing can be jointly applied to implement reliable distributed transaction management within microservices-based systems. We outline their architectural principles, evaluate strengths and limitations and present implementation patterns that help balance business consistency, operational performance and system resilience.

## 2. Distributed Transactions and the Problem Space

Consider a typical business transaction such as placing an order in an e-commerce system. This seemingly simple action may involve a sequence of operations including updating inventory levels, reserving customer payments, generating shipping instructions, initiating fulfilment workflows and notifying downstream systems. In a microservices architecture, these responsibilities are distributed across independent services such as Inventory, Payment, Shipping and Notification. Because each service manages its own data and executes its own logic, the failure of any single operation such as a rejected payment must be handled gracefully. To maintain system integrity, compensating actions (e.g., releasing previously reserved inventory or cancelling shipping requests) must be executed to restore the system to a consistent state.

Traditional distributed transaction mechanisms such as two-phase commit (2PC) or XA transactions are ill-suited to this environment. These protocols rely on global locks, coordinated commit phases and synchronous blocking operations across multiple services characteristics that conflict with the principles of microservices, which prioritize autonomy, loose coupling, availability and scalability. Long-duration locking and tight coordination not only degrade performance but also reduce overall system resilience in the presence of failures.

The Saga pattern offers a more suitable alternative for managing multi-service transactional workflows. Instead of enforcing strict ACID guarantees across the entire distributed operation, Saga decomposes the business transaction into a sequence of smaller, local transactions that execute independently and communicate through events or messages. If an error occurs at any stage, previously completed actions are reversed through compensating transactions. This approach embraces eventual consistency, reduces coordination overhead and aligns naturally with asynchronous, message-driven microservices architectures.

While Saga addresses the flow of multi-step workflows, another challenge arises in read-heavy distributed systems: the

need to maintain an accurate, up-to-date view of system state across services without imposing centralization or performance bottlenecks. This is where Command Query Responsibility Segregation (CQRS) and Event Sourcing provide significant advantages. CQRS improves scalability by separating read and write concerns into independent models, allowing each to be optimized differently. Event Sourcing enhances this model by persisting all changes as immutable events, enabling reliable reconstruction of service state, auditability and transparent synchronization across distributed components.


Together, the Saga pattern and CQRS/Event Sourcing enable robust distributed transaction management that supports autonomy, scalability and resilience core goals of microservices-based system design.

## 3. The Saga Pattern

### 3.1. Concept

A Saga is conceptually defined as a sequence of coordinated local transactions, each executed by an independent service participating in a larger business workflow. After a service completes its local transaction successfully, it emits an event or message that triggers the subsequent step of the saga. This event-driven chaining of operations enables the entire multi-service workflow to progress without requiring synchronous coordination.

Crucially, if any local transaction within the sequence fails, the saga initiates a set of compensating transactions that semantically undo the effects of all previously completed steps. These compensations restore the system to a consistent state without relying on global locks or tightly coupled commit protocols. By avoiding mechanisms such as distributed two-phase commit, the Saga pattern supports scalability, fault tolerance and loose coupling characteristics essential to microservices-based architectures (Figure 1).



**Figure 1:** Overview of Saga-based distributed transaction management.


### 3.2. Coordination Styles: Orchestration vs. Choreography

There are two primary mechanisms for coordinating the sequence of local transactions within a Saga: orchestration and choreography. Although both approaches enable distributed workflows without global locking, they differ significantly in structure and governance.

Orchestration relies on a central controller commonly referred to as the saga orchestrator that explicitly directs the workflow. The orchestrator determines the order of local transactions, instructs each participating service when to execute its operation, monitors the results and initiates compensating transactions in the event of a failure. This centralized control simplifies workflow visibility and oversight but introduces a single point of coordination that must be carefully designed to avoid becoming a bottleneck.

Choreography, by contrast, follows a decentralized model in which no central coordinator exists. Each service independently publishes domain events upon completing its local transaction. Other services subscribe to and react to these events, triggering subsequent steps in the workflow. This event-driven approach enhances autonomy and loose coupling among services but can make the global flow of the saga more difficult to trace, monitor and evolve as the system grows.

Both coordination styles offer unique advantages and their suitability depends on the complexity, governance needs and coupling constraints of the underlying microservices ecosystem.



**Figure 2:** Choreography-based Saga in microservices architecture.

### 3.3. Benefits and Trade-offs

**3.3.1. Benefits:** The Saga pattern offers several advantages that make it well-suited for microservices-based architectures. First, it eliminates the need for global locking mechanisms, thereby supporting the autonomy and independent scalability of distributed services (microservices.io; Medium). By decomposing a business workflow into smaller, independently executed local transactions, Saga enables long-running, multi-service workflows that integrate asynchronously and incorporate compensating logic when failures occur (SciTePress). This design promotes system resilience and fault tolerance: a failure in one

service does not bring down the entire system, as compensating actions restore consistency without requiring a global rollback mechanism (Microsoft Learn).

**3.3.2. Trade-offs and challenges:** Despite these advantages, the Saga pattern introduces several challenges. The most fundamental is eventual consistency: the system may contain intermediate, inconsistent states until all steps and compensations have completed (Microsoft Learn). Additionally, implementing compensation logic can be complex, particularly when dealing with operations that involve external side effects or non-idempotent behavior. Another notable limitation is the lack of isolation: because each local transaction commits immediately, other services might observe partial results before the saga completes its full execution. As highlighted in recent studies, many real-world implementations struggle to fully address isolation guarantees and idempotency concerns, especially in large-scale microservices environments (MDPI; IRJET).

### 4. CQRS & Event Sourcing Complementing Saga for State Management

When employing the Saga pattern for distributed transaction coordination, it becomes equally important to adopt a scalable and maintainable approach for managing system state across services over time. This requirement is effectively addressed through the combined use of Command Query Responsibility Segregation (CQRS) and Event Sourcing, two architectural patterns that complement the asynchronous and decentralized nature of microservices.

CQRS introduces a clear separation between the WRITE side and the READ side of an application's data model. The write model is responsible for processing commands, enforcing business rules and generating events, while the read model is dedicated solely to serving queries and is optimized for retrieval performance. This separation prevents read and write workloads from competing for the same data structures or storage resources, enabling each side to scale and evolve independently (Wikipedia).


Event Sourcing further enhances this model by replacing traditional state storage with an immutable log of domain events. Instead of persisting only the latest state, the system records every state-changing event in the order in which it occurred. The current state of the system can then be reconstructed by replaying these events, while specialized read models often referred to as projections or materialized views can be derived to support efficient querying. This approach provides a complete audit trail, enables temporal insights and facilitates flexible adaptation of read models as system requirements evolve (microservices.io; Medium).

Together, CQRS and Event Sourcing provide strong foundations for consistency, auditability and scalability in distributed systems, making them highly compatible with Saga-based transaction flows (Figure 3).

### 4.4. Benefits & Trade-offs

The combination of CQRS and Event Sourcing provides a comprehensive set of advantages that make it especially effective for large-scale, distributed and microservices-based systems. One of the most significant benefits is enhanced scalability,

achieved through the explicit separation of command-processing (writes) and query-handling (reads). This decoupling allows each side of the system to evolve, optimize and scale independently based on its unique workload characteristics. For example, read-intensive applications can leverage denormalized, query-optimized projections that support high-volume, low-latency access patterns without placing competitive pressure on the write path (Wikipedia; DZone). This separation is particularly advantageous in cloud environments where services may auto scale independently based on their traffic profile.



**Figure 3:** CQRS + Event Sourcing architecture for distributed systems.

Event Sourcing amplifies these capabilities through its built-in auditability, temporal fidelity and system transparency. By persisting every state change as an immutable event rather than overwriting state, the system produces a complete and chronological record of all domain activity. This event log functions as a forensic data source, enabling developers to replay state transitions for debugging, reconstruct system behavior at specific points in time, conduct time-travel queries and reproduce read models when business requirements evolve capabilities almost impossible to achieve cleanly in traditional CRUD models (microservices.io). In regulated industries such as finance and healthcare, this historical traceability becomes a major architectural advantage due to compliance and auditing requirements.

Moreover, CQRS and Event Sourcing promote a high degree of loose coupling and service autonomy, qualities central to microservices architecture. Rather than sharing a common database a practice that leads to schema entanglement and cross-service dependencies services publish and consume domain events, enabling flexible asynchronous integration and polyglot persistence. This event-driven approach allows each service to maintain its own datastore and technology stack while still participating in distributed workflows, a design pattern frequently cited in microservices literature as a prerequisite for independent deployment and bounded contexts (IRJET).

However, these substantial benefits come with noteworthy architectural trade-offs. One of the most fundamental is eventual consistency, an inherent characteristic of distributed systems where read models naturally lag the write model. This lag introduces design challenges such as race conditions, stale reads and user-visible inconsistencies, all of which must be mitigated through compensatory UI strategies, reconciliation logic or asynchronous notifications (AKF Partners). Achieving correctness under eventual consistency requires a shift in both engineering mindset and product design.

Event Sourcing also introduces considerable implementation complexity, often underestimated by teams new to the pattern. Developers must design robust event stores, manage event versioning and schema evolution, construct and maintain projections, perform snapshotting for long-lived aggregates and ensure idempotent event processing across distributed nodes (Medium). Furthermore, as systems grow over time, the raw volume of stored events may increase dramatically, leading to performance overhead during state reconstruction, event replay or projection rebuilding. Without proactive measures such as snapshotting, stream compaction, partitioning or archival the event log can eventually become a bottleneck (microservices.io).

These challenges emphasize that CQRS and Event Sourcing are not merely technical patterns but operationally intensive architectural commitments. They require strong observability practices, disciplined event modelling, governance over schema evolution and continuous monitoring of projection health and event-processing pipelines. When these principles are followed, however, CQRS combined with Event Sourcing offers a powerful foundation for building distributed systems that achieve high scalability, resilience, analytical richness and long-term maintainability.

## 5. Integrating Saga with CQRS/Event Sourcing Implementation Techniques

Given their complementary strengths, integrating the Saga pattern with CQRS and Event Sourcing yields a robust architectural model for managing distributed transactions in microservices-based systems. Each pattern addresses a distinct yet interrelated aspect of distributed coordination and state management and together they form a cohesive approach for ensuring both business consistency and operational scalability. First, the Saga pattern, whether implemented through orchestration or choreography, provides a reliable mechanism for coordinating multi-service transactions. It ensures that each step of a business workflow is executed in sequence and that compensating actions are applied when failures occur, thereby maintaining logical consistency without relying on global locks or tightly coupled transaction managers.

Within each participating service, CQRS combined with Event Sourcing offers a powerful strategy for managing state changes. CQRS isolates command processing from query operations, allowing each to be optimized independently. Event Sourcing reinforces this separation by persisting every state transition as an immutable event, enabling complete auditability, reproducibility and flexible construction of read models tailored to specific query workloads. To support both Saga coordination and event-driven state propagation, systems typically employ message brokers or event streaming platforms. These brokers facilitate reliable transmission of saga-related messages such as transaction steps, confirmations and failures as well as domain events used for Event Sourcing. This event-driven backbone enables services to remain loosely coupled while still collaborating effectively in distributed workflows.

However, successful adoption of this architectural style requires careful attention to idempotency, compensation logic and eventual consistency. Operations that interact with external systems such as payment gateways, shipping providers or third-party APIs must be designed defensively to handle

retries, duplicates and partial failures. Similarly, compensating transactions must be designed with semantic correctness to ensure that business invariants remain intact under both normal and failure scenarios. This integrated architecture is particularly valuable in domains such as order processing, payment orchestration, inventory management, booking platforms and other workflows where cross-service consistency is essential, but traditional global locking mechanisms are impractical. A review of existing literature including a 2017 survey examining Microservices, the Saga pattern and Event Sourcing highlights the increasing adoption and proven applicability of this combined approach in real-world distributed systems (IRJET).

## 6. Key Studies & Literature

A few prior studies and practitioner resources provide essential conceptual and empirical grounding for the use of Saga, CQRS and Event Sourcing in distributed systems. One notable work is the 2015 survey on Microservices, Saga Pattern and Event Sourcing, which offers a comprehensive overview of how these architectural patterns are applied in practice. The survey analyzes their benefits, limitations and adoption trends across real-world microservices implementations, highlighting the growing relevance of event-driven consistency mechanisms in distributed environments (IRJET).

Another significant contribution is the paper The Saga Pattern in a Reactive Microservices Environment, which compares Saga-based coordination to traditional two-phase commit (2PC) protocols. This study emphasizes the suitability of sagas for long-running, asynchronous and multi-service workflows and provides an evaluation of various Java-based saga frameworks used within reactive microservices ecosystems (SciTePress).

In addition to academic literature, numerous practitioner-oriented articles and technical blogs offer valuable hands-on insights into implementing Saga, CQRS and Event Sourcing in production environments. These resources frequently address the practical challenges of designing compensating transactions, managing eventual consistency, ensuring idempotency and maintaining efficient read/write separation. Their experiential accounts serve as an important complement to formal research, illustrating how these patterns perform under real operational constraints (Medium).

Collectively, these studies and practitioner perspectives form a robust foundation for understanding and implementing distributed transaction management using Saga in combination with CQRS and Event Sourcing. They demonstrate the increasing maturity and widespread adoption of these patterns within modern microservices-based architectures.

## 7. Case Study: Distributed Order Processing in an E-Commerce Platform

To illustrate the practical application of Saga, CQRS and Event Sourcing in a real-world setting, consider a modern e-commerce platform responsible for processing customer orders across multiple autonomous microservices. The workflow begins when a customer submits an order through the Order Service. In a traditional monolithic system, the order submission, payment authorization, inventory reservation and shipping preparation would be wrapped in a single ACID transaction. However, in a microservices architecture, each function is owned by a separate service such as Inventory, Payment, Shipping and Notification

each maintaining its own database and operational boundaries. Coordinating these actions reliably becomes a significant challenge in the presence of partial failures, latency spikes or temporary service unavailability.

In this implementation, the platform employs a Saga orchestration model within an Order Orchestrator Service. Upon receiving a new order request, the orchestrator first instructs the Inventory Service to reserve stock. Once successful, it commands the Payment Service to authorize the customer's payment. If payment is approved, the orchestrator proceeds to trigger the Shipping Service to create a shipment request, followed by the Notification Service to generate customer updates. If any step fails such as a payment decline or insufficient stock the orchestrator invokes compensating transactions: previously reserved inventory is released, pending shipping tasks are cancelled and the customer is notified of the failure. This ensures that the global business workflow maintains logical consistency across services without relying on distributed locking or two-phase commit.

To support this workflow, each service adopts CQRS and Event Sourcing for managing state. For example, the Order Service processes "CreateOrder" commands, validates business rules and persists all state transitions as events such as OrderCreated, orderApproved, orderCanceled or OrderFailed. The write model focuses exclusively on processing these commands and emitting domain events, while the read model constructs optimized views for customer-facing interfaces, such as order status dashboards. These views materialized by projecting the event stream into query-friendly structures, enabling real-time updates while maintaining strong auditability. Similarly, the Inventory and Payment Services rely on event sourcing to record every state change, supporting replay-based recovery and providing a complete history of reservations, releases, authorizations and declines.

The platform's event broker implemented using technologies such as Kafka or RabbitMQ acts as the backbone for all inter-service communication. Saga events (successes, failures, compensations) are published to orchestrator topics, while domain events from each service propagate state changes across the system. This event-driven architecture promotes loose coupling, fault isolation and horizontal scalability, allowing each service to evolve independently without disrupting global workflows.

Operational metrics reinforce the advantages of this architectural model. The system demonstrates high resilience to partial service failures due to the compensating logic embedded in saga definitions. Read-side performance is significantly improved by the CQRS pattern, enabling user interfaces to respond rapidly using precomputed projections. Moreover, event sourcing provides robust auditing and traceability, allowing engineers to replay events to debug issues or reconstruct inconsistent states. This case study highlights how the coordinated use of Saga, CQRS and Event Sourcing provides a scalable, maintainable and resilient foundation for complex distributed business processes in microservices-based platforms.

## 8. Conclusion

The combined application of the Saga pattern with CQRS and Event Sourcing provides a robust and adaptable

architectural framework for managing distributed transactions in microservices-based systems. Saga enables the decomposition of global workflows into coordinated local transactions, eliminating the need for global locks or rigid two-phase commit protocols while preserving business consistency through compensating actions. Meanwhile, CQRS and Event Sourcing complement this coordination model by maintaining state changes in an immutable event log and separating command processing from query operations, thereby supporting auditability, scalability and flexible data representation.

Nonetheless, adopting this architectural approach requires organizations to embrace a fundamentally different consistency philosophy. Systems must be designed to operate under eventual consistency, with careful attention paid to idempotency, compensation logic and the complexities introduced by event-driven communication including event versioning and schema evolution. These challenges demand rigorous engineering discipline but yield substantial long-term benefits in reliability, transparency and operational resilience.

As microservices architectures continue to proliferate across industries, the relevance of combining Saga with CQRS and Event Sourcing is expected to grow. This approach is particularly well-suited to domains involving long-running transactions, multi-service workflows and high scalability requirements. With ongoing advancements in event-driven platforms, distributed messaging systems and domain-driven design practices, the integration of these patterns will likely continue evolving, offering even more sophisticated solutions for dependable distributed transaction management.

## 9. References

1. <https://microservices.io/patterns/data/event-sourcing.html>
2. <https://martinfowler.com/eaaDev/EventSourcing.html>

3. <https://martinfowler.com/bliki/CQRS.html>
4. <https://udidahan.com/2009/12/09/clarified-cqrs/>
5. [https://docs.microsoft.com/en-us/previous-versions/msp-n-p/jj554200\(v=pandp.10\)](https://docs.microsoft.com/en-us/previous-versions/msp-n-p/jj554200(v=pandp.10))
6. <https://yos.io/2017/10/30/distributed-sagas/>
7. [https://en.wikipedia.org/wiki/Command%20query\\_segregation](https://en.wikipedia.org/wiki/Command%20query_segregation)
8. Jacobs FR, Weston FC. Enterprise resource planning (ERP)-A brief history. *Journal of Operations Management*, 2007;25: 357-363.
9. Dragoni N, Giallorenzo S, Lafuente AL, et al. Microservices: Yesterday, Today and Tomorrow. In *Present and Ulterior Software Engineering*, 2017: 195-216.
10. Padur SKR. Online Patching and Beyond: A Practical Blueprint for Oracle EBS R12.2 Upgrades. *International Journal of Scientific Research in Science, Engineering and Technology (IJSRSET)*, 2016;2: 1028-1039.
11. Padur SKR. Network Modernization in Large Enterprises: Firewall Transformation, Subnet Re-Architecture and Cross-Platform Virtualization. In *International Journal of Scientific Research & Engineering Trends*, 2016;2.
12. Vishnubhatla S. Scalable Data Pipelines for Banking Operations: Cloud-Native Architectures and Regulatory-Aware Workflows. In *International Journal of Science, Engineering and Technology*, 2016;4.