

---

# Designing Resilient Systems: The Power of Event Sourcing in AWS

Vigneshwaran Manivelmurugan\*

---

## Abstract

This article provides an in-depth exploration of the design and implementation of distributed event-driven systems using the event sourcing pattern within modern cloud-native architectures. It examines the transformative benefits of this approach, including enhanced scalability, reliability, and data consistency. The article also addresses the challenges and practical considerations associated with transitioning from legacy systems to cloud-based solutions, with a particular emphasis on AWS. Through detailed architectural breakdowns, service integration methodologies, and real-world use cases, it demonstrates how event sourcing can optimize performance, ensure data integrity, and support diverse operational needs in an increasingly dynamic technological landscape.

---

### Keywords:

AWS;  
Event Driven Architecture.  
Event Sourcing.  
Serverless.  
Microservice.

Copyright © 2025 International Journals of Multidisciplinary Research Academy. All rights reserved.

---

### Author correspondence:

Vigneshwaran Manivelmurugan,  
Senior Lead Software Engineer, Capital One, Richmond, Virginia, USA  
Email: manivelvicky@gmail.com

---

## 1. Introduction

Event sourcing is a fundamental architectural pattern in distributed systems that emphasizes capturing all changes to an application's state as a sequence of immutable events. Unlike traditional approaches that overwrite state with updates, event sourcing retains every change in a durable event log, allowing the system to reconstruct its state at any point in time. This design offers inherent advantages such as traceability, flexibility, and the ability to handle complex workflows with asynchronous communication. Before diving into the technical details, it is important to understand how event sourcing fits into the broader context of cloud-native architectures and why it is particularly valuable in modern application development.

In the era of cloud-native development, organizations increasingly migrate their legacy systems to modern distributed architectures to meet demands for scalability, performance, and cost-effectiveness. One such architectural approach is event-driven systems, where components communicate asynchronously through events. Event sourcing, a subset of this paradigm, captures every change to an application's state as a sequence of immutable events. This pattern not only supports operational efficiency but also enables traceability, auditability, and flexibility for downstream systems.

Transitioning to this approach is not without challenges. Splitting monolithic systems into microservices connected by message queues may inadvertently introduce complexities. Without a well-thought-out design, this can lead to inefficiencies and data inconsistencies. This article focuses on how event sourcing can address these issues, providing practical insights into building robust, distributed systems on AWS.

## 2. High-Level Design and Architecture

At its core, the event sourcing pattern segregates event producers from event consumers, ensuring clean separation of concerns and maintaining data integrity. In a cloud-native setup, this involves leveraging services like DynamoDB or Aurora RDS for event storage, coupled with messaging systems like Kinesis or DynamoDB Streams for reliable event propagation.




Figure 1. *High level system design*

Let us examine each component in detail.

### Client

Clients initiate interactions, whether through web browsers, mobile applications, or IoT devices. These requests set the event pipeline in motion. The reliability of the system hinges on robust API design, ensuring clients can operate seamlessly under varying conditions.

### Event Source

The event source microservice serves as the gateway for client requests. It validates and processes input data before recording it in the event store. Deployed on AWS Lambda, it integrates with API Gateway for scalable, secure request handling. The event source focuses solely on its domain logic, delegating downstream responsibilities to other components.

### Event Store

The event store acts as the system's source of truth, persisting events in an immutable log. In AWS, services like DynamoDB or Aurora PostgreSQL provide high availability and fault tolerance. For instance, DynamoDB Streams offers change data capture (CDC) capabilities, enabling downstream consumers to react to state changes. Aurora PostgreSQL's triggers and logical replication allow similar functionality for relational use cases.

Efficient event storage design includes considerations like partition and sort key selection to optimize query performance and scalability. Enabling multi-region replication via DynamoDB Global Tables or Aurora Global Database further enhances system reliability.

### Messaging System

The messaging system is the bridge between the event store and event handlers. AWS provides multiple options, such as DynamoDB Streams, Kinesis Data Streams, or SQS, for reliable event delivery. These services ensure durability and scalability, handling large event volumes without compromising performance. For high-throughput scenarios, Kinesis offers fine-grained control over partitioning and data retention.

### Event Handlers

Event handlers consume events and perform domain-specific processing. AWS Lambda, triggered by DynamoDB Streams or Kinesis, is a popular choice for serverless event handling. Alternatively, ECS, Fargate, or on-premises systems can use the Kinesis Client Library (KCL) for custom processing logic.

Handlers often execute parallel workflows, enabling modular and independent processing. For instance, in an e-commerce system, one handler might validate payment details while another checks inventory. By decoupling these responsibilities, the architecture achieves greater flexibility and maintainability.

## 3. Advantages of Event Sourcing

- **Auditability:** The event log provides a complete history of state changes, enabling detailed traceability and compliance with regulatory requirements.
- **Scalability:** By decoupling event producers and consumers, event sourcing allows horizontal scaling of components to handle varying workloads.
- **Resilience:** Immutable events ensure data integrity, even in the face of failures or partial outages. Systems can replay events to recover or rebuild state.
- **Flexibility:** Downstream systems can subscribe to events for real-time or batch processing, supporting diverse use cases without impacting upstream operations.
- **Time Travel:** The event log enables reconstruction of system state at any point in time, facilitating debugging and ability to replay historical event if needed.

#### **4. Design Considerations and Challenges**

When adopting the event sourcing pattern in a cloud-native environment, several design considerations and potential challenges must be addressed to ensure a robust, scalable, and maintainable architecture. These aspects often overlap, requiring a holistic approach to implementation.

##### **Data Modeling**

The choice of partition and sort keys is critical for optimizing query performance. These keys should align with the application's access patterns. For example, in a customer-centric application, the partition key could be the customer ID, while the sort key might be a timestamp to organize events chronologically. Proper indexing strategies and secondary indexes should also be planned to accommodate diverse query requirements efficiently.

##### **Managing Data Consistency**

Event sourcing inherently embraces eventual consistency, which can be challenging for systems requiring strict guarantees. In asynchronous architectures, there is often a delay between state changes and their visibility to consumers, leading to stale data. To manage this, implement idempotent processing in event handlers to ensure duplicate events do not cause unintended side effects. Mechanisms like conflict resolution strategies and compensating transactions can reconcile discrepancies and ensure business processes remain reliable.

##### **Storage Optimization**

An ever-growing event log can lead to unmanageable storage costs and degraded performance. Strategies like archiving older events to cheaper storage solutions (e.g., S3) or implementing periodic snapshots of the system's state can mitigate this issue. Snapshots allow for efficient state reconstruction without replaying the entire event history.

##### **Debugging and Troubleshooting Distributed Systems**

Distributed systems introduce complexity in debugging and troubleshooting due to their asynchronous nature and multiple components. Identifying the root cause of issues can be challenging without comprehensive visibility into event flows. Implementing robust logging and tracing frameworks is essential. Tools like AWS X-Ray and CloudWatch Logs can provide detailed insights into event processing pipelines, enabling developers to visualize the flow of events, detect bottlenecks, and identify failed components efficiently.

##### **Fault Tolerance and Recovery**

Fault tolerance is vital in distributed systems. Use multi-region replication with DynamoDB Global Tables or Aurora Global Database to ensure data availability during regional outages. Regularly back up event stores and test recovery mechanisms to prepare for unexpected data loss scenarios. These practices enhance system reliability and ensure continuity during failures.

##### **Handling Event Duplication**

Event duplication is a common challenge in distributed systems, typically arising from retries or network issues. Duplicate events can cause unintended side effects, such as processing the same transaction multiple times. To mitigate this, event handlers should be designed to be idempotent, ensuring that repeated processing of the same event does not alter the system's state. Including unique identifiers in event metadata allows handlers to verify whether an event has already been processed, ensuring duplicates are safely ignored.

##### **Schema Evolution**

Over time, the structure of events often needs to change as system requirements evolve. This introduces the risk of incompatibilities between new and existing event formats. Adopting a robust versioning strategy for events is essential. Embedding version information in event metadata or using a schema registry can ensure backward compatibility. These practices allow systems to process both old and new event formats seamlessly, minimizing disruption during updates.

##### **Mitigating Cold Start Latency**

Serverless architectures, particularly those using AWS Lambda, can suffer from cold start latency, which affects system responsiveness. This is especially problematic in latency-sensitive applications. To address this, optimize Lambda configurations by enabling provisioned concurrency, which keeps instances warm and reduces cold start times. For tasks that require consistent low-latency performance, consider using containers, such as AWS Fargate, to process events.

By addressing these design considerations and challenges together, organizations can unlock the full potential of event sourcing while ensuring their systems remain efficient, scalable, and maintainable.

## 5. Implementation Example: Auto Insurance Quote System

To illustrate event sourcing, consider an auto insurance application that processes quote requests. When a client submits their details via a web interface, the system triggers an API call to the event source microservice, deployed using AWS Lambda and exposed through API Gateway. The microservice validates the request and enriches the data, then writes it to a DynamoDB table configured with a suitable partition and sort key (e.g., customer ID and timestamp).

The DynamoDB table acts as the event store. DynamoDB Streams, enabled on the table, captures change events such as inserts, updates, or deletions and ensures they are available for downstream processing. A Lambda function listens to these streams, extracting events as they occur. Each event contains critical details such as a unique identifier and event metadata to maintain traceability.

The Lambda function, acting as an event handler, performs domain-specific processing. For instance, one function may verify the vehicle's VIN to retrieve make and model details, while another gathers the driver's accident history from external services. These independent handlers operate concurrently, ensuring modularity and scalability. After processing, the results are consolidated by a fulfillment service that calculates the final insurance quote. This quote is then stored using projections to create a highly optimized read model. Projections aggregate events into a consolidated view, enabling efficient querying for specific use cases. For example, a projection service could generate a "customer quotes" table, which maps each customer ID to their respective quotes and statuses. This table is stored in DynamoDB or Amazon RDS to facilitate low-latency queries for customer-facing applications.

Projections not only enhance query performance but also allow real-time updates of quote information in customer dashboards or mobile applications. By leveraging this pattern, the architecture separates the read and write concerns, ensuring that operational systems and analytics workloads do not interfere with each other. Projections can also power notifications, enabling features like sending personalized quote updates via email or SMS when specific thresholds or conditions are met.

This design ensures scalability, modularity, and fault tolerance while maintaining data integrity. By leveraging AWS-native services like DynamoDB, Lambda, and API Gateway, and integrating projections effectively, the system achieves an efficient, cost-effective architecture that meets the demands of modern cloud-native applications.

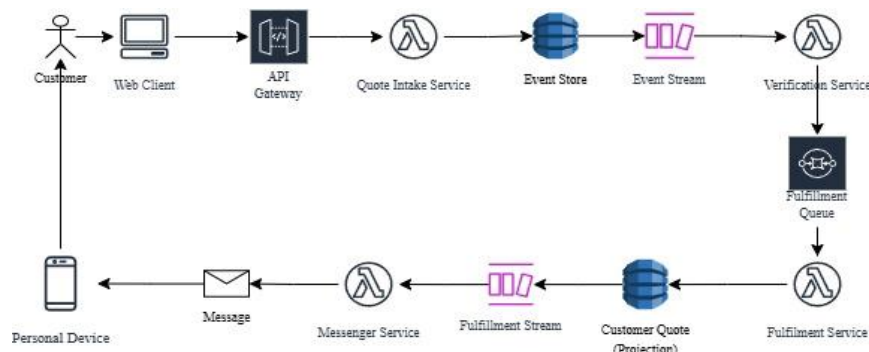


Figure 2. Low level system design

## 6. Conclusion

Event sourcing provides a powerful architectural pattern for building distributed event-driven systems in modern cloud-native environments. By capturing every state change as an immutable event, it ensures auditability, traceability, and flexibility while enabling robust fault tolerance and scalability. The design allows for asynchronous workflows, decoupling upstream and downstream systems, which can handle large-scale workloads with high reliability.

While the pattern introduces challenges such as managing eventual consistency, event duplication, and storage growth, these can be effectively mitigated with thoughtful implementation strategies. Leveraging AWS-native services like DynamoDB, Kinesis, and Lambda simplifies the implementation, offering built-in solutions for high availability, observability, and performance optimization.

The example of an auto insurance quote system demonstrates how event sourcing can address real-world problems, delivering a scalable, modular, and efficient solution. As organizations continue to migrate to cloud-native infrastructures, adopting event sourcing offers a pathway to resilient, maintainable systems that can meet the growing demands of modern applications.

## References

- [1] Evans, Eric. *Domain-Driven Design: Tackling Complexity in the Heart of Software*. Addison-Wesley, 2003.
- [2] Fowler, Martin. "Event Sourcing." Martin Fowler. <https://martinfowler.com/eaaDev/EventSourcing.html>
- [3] AWS Documentation. "Using DynamoDB Streams." Amazon Web Services. <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html>
- [4] AWS Documentation. "Using AWS Lambda with DynamoDB Streams." Amazon Web Services. <https://docs.aws.amazon.com/lambda/latest/dg/with-ddb.html>
- [5] Pivotal Software. "Event-Driven Microservices." Pivotal Blog. <https://tanzu.vmware.com/content/blog/event-driven-microservices-why-and-how>
- [6] Hohpe, Gregor, and Bobby Woolf. *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions*. Addison-Wesley, 2003.
- [7] AWS Documentation. "Modernize legacy databases using event sourcing and CQRS with AWS DMS" <https://aws.amazon.com/blogs/database/modernize-legacy-databases-using-event-sourcing-and-cQRS-with-aws-dms/>
- [8] Amazon Web Services. "Best Practices for Architecting Resilient Streaming Data Applications." <https://aws.amazon.com/architecture/Event>