



## **Outbox Pattern for Reliable Distributed Systems**

**Akshay Pratinav**

Staff Software Engineer, Intuit, USA.

**Shibashis Mishra**

Senior Engineering Manager, Adobe, USA

### **Abstract**

Modern distributed systems increasingly rely on asynchronous communication through event streams and message brokers. Ensuring reliable, exactly-once event publication while maintaining transactional integrity across heterogeneous services remains a core challenge. The Outbox Pattern has emerged as a robust architectural mechanism to guarantee atomicity between database state changes and event emission. This paper presents an in-depth analysis of the Outbox Pattern, design considerations, implementation strategies, performance implications, and its applicability in cloud-native, microservice-based architectures. It also discusses advanced variants including idempotent consumers, deduplication techniques, and its role within event-driven frameworks.

### **Keywords:**

Outbox Pattern, Inbox-Outbox Pattern, Event Router Pattern, Microservices, Event-Driven Architecture, Eventual Consistency, Exactly-Once Processing, At-Least-Once Delivery, Distributed Systems, Transactional Integrity, Atomicity, Dual-Write Problem, Idempotency, Deduplication, Change Data Capture (CDC), Polling, Message Broker, Kafka, Debezium, Database Transaction Log, Write Amplification, Financial Systems, E-commerce, Serverless Architectures.

---

**How to cite this paper:** Akshay Pratinav, Shibashis Mishra. (2025). Outbox Pattern for Reliable Distributed Systems. *ISCSITR-International Journal of Cloud Computing (ISCSITR-IJCC)*, 6(6), 18–25.

**DOI:** [http://www.doi.org/10.63397/ISCSITR-IJCC\\_2025\\_06\\_06\\_002](http://www.doi.org/10.63397/ISCSITR-IJCC_2025_06_06_002)

**URL:** [https://iscsitr.com/index.php/ISCSITR-IJCC/article/view/ISCSITR-IJCC\\_2025\\_06\\_06\\_002/ISCSITR-IJCC\\_2025\\_06\\_06\\_002](https://iscsitr.com/index.php/ISCSITR-IJCC/article/view/ISCSITR-IJCC_2025_06_06_002/ISCSITR-IJCC_2025_06_06_002)

**Published:** 1<sup>st</sup> December 2025

**Copyright** © 2025 by author(s) and International Society for Computer Science and Information Technology Research (ISCSITR). This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>  **Open Access**

---

## 1. Introduction

Distributed applications often require the consistent propagation of state changes to downstream systems. Typical use cases include audit logs, cache invalidation, search indexing updates, and communication across microservices. Directly publishing messages to a broker within a database transaction can compromise reliability due to the absence of distributed ACID semantics. The Outbox Pattern offers a solution by decoupling data modification from event publication using a dedicated “outbox” table or log within the service’s database. This technique ensures that database updates and event generation occur atomically, thus preventing message loss or duplication.

## 2. Background and Motivation

### a. Event-Driven Microservices

Microservices commonly use event streams (e.g., Kafka, Kinesis, Pulsar) to ensure loose coupling and eventual consistency. A core requirement is that every state change generating an event must be delivered reliably.

### b. The Dual-Write Problem

Systems that write to a database and publish messages sequentially face the dual-write problem, where failures between operations lead to:

- Missing events
- Inconsistent state between services
- Phantom or duplicated actions

### **c. Outbox Design Principle**

The Outbox Pattern eliminates dual-writes by recording outbound events as part of the same database transaction that modifies the business data. A separate process reliably reads and publishes these events.

## **3. Architecture of the Outbox Pattern**

### **a. Components**

1. Primary Database: Stores both business data and the Outbox table.
2. Outbox Table: A durable, append-only log of events awaiting publication.
3. Publisher/Relay: A background worker or change data capture (CDC) pipeline that reads outbox records and publishes them to a message broker.
4. Message Broker: Ensures asynchronous delivery to downstream services.
5. Consumer Services: Apply updates idempotently.

### **b. Transactional Workflow**

1. Application performs local business operations.
2. As part of the same ACID transaction, an Outbox event is inserted.
3. The relay process reads pending events and publishes them.
4. Upon successful publishing, the event record is marked or deleted.

## **4. Implementation Strategies**

### **a. Polling-Based Outbox**

- A scheduled job polls the outbox table for unprocessed rows.
- Pros: Simple, database-native
- Cons: Higher latency, potential contention

### **b. Change Data Capture (CDC) Outbox**

- Tools like Debezium or native cloud CDC services stream outbox table changes.
- Pros: Low latency, scalable, minimal DB impact
- Cons: Operational complexity

### **c. Exactly-Once vs At-Least-Once Delivery**

The pattern typically ensures at-least-once publication, requiring idempotency on consumers using techniques like:

- Event UUIDs
- Deduplication tables
- Consumer-side idempotent logic

## 5. Performance and Scalability Analysis

### a. Database Load

Outbox events increase write amplification. Strategies to control load:

- Partitioned tables
- TTL/archival
- Batching publishes

### b. Latency Considerations

Polling intervals influence end-to-end propagation delays, whereas CDC offers sub-second latencies.

### c. High Availability

In multi-region designs, Outbox Pattern pairs well with:

- Region-local brokers
- Active-passive failover

## 6. Advanced Patterns

### a. Inbox–Outbox Pattern for End-to-End Exactly-Once Workflows

The Inbox–Outbox Pattern extends the classical outbox mechanism by incorporating an inbox table on the consumer side. It prevents duplicate processing of events even when brokers deliver messages more than once.

#### 1. Architecture

- Producer Outbox: Records events as part of the producer's transaction.
- Consumer Inbox: Records a message ID for every processed event before executing business logic.
- Idempotent Handler: Ensures repeated message deliveries do not re-execute side effects.

#### 2. Benefits

- Achieves application-level exactly-once processing, even when the messaging layer

only guarantees at-least-once delivery.

- Suitable for financial and inventory systems where reprocessing can cause monetary inconsistencies.

## **b. Outbox with Change Data Capture (CDC) for Low-Latency Streaming**

Traditional polling-based outbox processes may introduce latency or cause table contention. A CDC-based outbox uses database log tailing mechanisms such as Debezium, Oracle GoldenGate, or AWS DMS.

### **1. Architecture**

- Log-Based CDC: Reads from the DB transaction log; ideal for MySQL binlog, PostgreSQL WAL.
- Trigger-Based CDC: Uses DB triggers to push changes into an event stream.

### **2. Benefits**

- Sub-second latency from write to event publish
- Minimal impact on OLTP workloads
- Automatically handles batching and parallel event deliveries

## **c. Transactional Outbox + Event Router Pattern**

In complex microservice ecosystems, services may emit multiple types of events that must be routed to different downstream systems.

### **1. Architecture**

The outbox stores:

- Event payload
- Event type
- Target topic/stream
- Routing metadata (region affinity, priority flags)

A dedicated Event Router service reads the outbox/CDC log and dispatches events accordingly.

### **2. Benefits**

- Clean decoupling of event generation and distribution logic
- Simplifies multi-broker setups (e.g., Kafka in active region + Kinesis in passive region)

## **Use Cases**

The Outbox Pattern is broadly applicable across several classes of distributed systems

where transactional integrity and reliable event propagation are essential. This section explores key domains where the pattern provides significant architectural benefits.

#### **a. Financial systems and payment platforms**

It demands strict consistency, auditability, and fault-tolerance because mistakes directly translate into monetary loss and regulatory risk. The Outbox Pattern is especially valuable here: when a user initiates a payment or transfer, the core system updates balances, ledger entries, and transaction state in the database while atomically recording an “outbox” event in the same transaction. A separate, reliable dispatcher then publishes these events to downstream services such as fraud detection, accounting, settlement, and regulatory reporting pipelines. This prevents classic dual-write problems where the database updates but the event is never sent (or vice versa), enabling accurate reconciliation, robust audit trails, and dependable integration with external gateways and partners, even under partial failures, network partitions, or service restarts.

#### **b. E-commerce and retail platforms**

These are digital systems that enable businesses and individuals to buy and sell goods or services online, usually through websites or mobile apps that handle everything from product discovery to checkout. They typically provide tools for listing products, managing inventory, processing payments securely, calculating taxes and shipping, and handling orders and returns. On the customer side, these platforms offer features like search and filtering, personalized recommendations, reviews, and multiple payment and delivery options to make shopping easy and convenient. On the business side, they often integrate with marketing, analytics, and customer support tools so sellers can understand buyer behavior, optimize pricing and promotions, and scale operations. Well-known examples include marketplaces like Amazon, eBay, and Etsy, as well as branded online stores powered by platforms like Shopify, Magento, or WooCommerce.

#### **c. Serverless Architectures and Event-Oriented Integrations**

Serverless architectures and event-oriented integrations center on building applications as collections of small, stateless functions and services that are triggered by events rather than running continuously on dedicated servers. In a serverless model, the cloud provider manages provisioning, scaling, and infrastructure operations, so developers focus on writing function code that responds to events like HTTP requests, message queue

updates, database changes, file uploads, or scheduled timers, paying only for actual execution time. Event-oriented integrations connect these functions and services via event buses, message queues, and pub/sub systems, enabling loosely coupled components that can evolve independently and scale automatically under variable load. This style encourages fine-grained, reactive workflows and simplifies spiky or unpredictable traffic patterns, though it also introduces challenges around observability, cold starts, debugging distributed flows, and designing idempotent, reliable event processing.

## 7. Open Challenges and Future Work

Despite its maturity, challenges remain:

- Clean-up automation for high-throughput services
- Reducing relay operational complexity
- Combining Outbox with emerging serverless or agentic AI-driven architectures
- Versioning and schema evolution for event payloads

Future research may explore autonomous, self-healing outbox pipelines using agentic AI to detect and remediate event inconsistencies.

## 8. Conclusion

The Outbox Pattern remains a foundational element for reliable event-driven distributed systems. It provides strong transactional guarantees without requiring distributed transactions and significantly reduces the risk of state divergence. Its applicability spans microservices, serverless integrations, and multi-region Disaster Recovery. The pattern's integration with CDC systems and its growing synergy with AI-driven operations make it increasingly relevant in modern cloud architectures.

## Reference

- [1] Talaver, O.V. and Vakaliuk, T.A., 2023. Reliable distributed systems: review of modern approaches. *Journal of Edge Computing* [Online], 2(1), pp.84–101. Available from: <https://doi.org/10.55056/jec.586>
- [2] P. Mundhenk, A. Hamann, A. Heyl and D. Ziegenbein, "Reliable Distributed Systems," 2022 Design, Automation & Test in Europe Conference & Exhibition (DATE), Antwerp, Belgium, 2022, pp. 287-291, doi: 10.23919/DATE54114.2022.9774734.

- [3] Vedant Agarwal. (2025). Designing Resilient Distributed Systems: Fault Tolerance Strategies and Insights. *International Journal of Research in Computer Applications and Information Technology (IJRCAIT)*, 8(1), 1102-1113
- [4] Ageed, Z. S. ., & Zeebaree, S. R. M. . (2024). Distributed Systems Meet Cloud Computing: A Review of Convergence and Integration. *International Journal of Intelligent Systems and Applications in Engineering*, 12(11s), 469–490. Retrieved from <https://ijisae.org/index.php/IJISAE/article/view/4468>
- [5] Jeanneau, É., Rodrigues, L., Arantes, L. et al. An autonomic hierarchical reliable broadcast protocol for asynchronous distributed systems with failure detection. *J Braz Comput Soc* 23, 15 (2017). <https://doi.org/10.1186/s13173-017-0064-9>
- [6] Tushar Deepak Chandra and Sam Toueg. 1996. Unreliable failure detectors for reliable distributed systems. *J. ACM* 43, 2 (March 1996), 225–267. <https://doi.org/10.1145/226643.226647>
- [7] van Steen, M., Pierre, G. & Voulgaris, S. Challenges in very large distributed systems. *J Internet Serv Appl* 3, 59–66 (2012). <https://doi.org/10.1007/s13174-011-0043-x>
- [8] S. Srinivasan and N. K. Jha, "Safety and reliability driven task allocation in distributed systems," in *IEEE Transactions on Parallel and Distributed Systems*, vol. 10, no. 3, pp. 238-251, March 1999, doi: 10.1109/71.755824.
- [9] Ehsan Bazgir, Tasmita Tanjim Tanha, Anwarul Azim Bhuiyan, Ehteshamul Haque, Md Shihab Uddin . Analysis of Distributed Systems. *International Journal of Computer Applications*. 186, 48 ( Nov 2024), 16-21. DOI=10.5120/ijca2024924136
- [10] Yadav, P.K., Bhatia, K., and S. Gulati, "Reliability driven Soft Real-time Fuzzy Task Scheduling in Distributed Computing Environment", *Advances in Intelligent and Soft Computing*, Vol 130, pp. 207–214, (2012). [https://doi.org/10.1007/978-81-322-0487-9\\_21](https://doi.org/10.1007/978-81-322-0487-9_21)