

# APACHE KAFKA FOR EVENT-DRIVEN ARCHITECTURE

MOHAMMED ANAS

Department of Computer Applications  
RV College of Engineering  
Bengaluru, Karnataka, India

Dr. B.H. CHANDRASHEKAR

Department of Computer Applications  
RV College of Engineering  
Bengaluru, Karnataka, India

**Abstract**—In the rapidly evolving landscape of software development, the shift from monolithic architectures to microservices and Event-Driven Architecture (EDA) has introduced new complexities and challenges in system communication. This paper explores the strategic role of Apache Kafka in addressing these challenges, focusing on its applications in event tracking, retry mechanisms, asynchronous processing, and inter-service communication. By analyzing Kafka's advantages over traditional messaging systems, we highlight its scalability, durability, and low-latency characteristics that make it a preferred choice for modern distributed systems. Additionally, the paper delves into real-world industry practices, offering insights into how leading tech companies integrate Kafka into their infrastructures to enhance reliability and resilience. We also examine the common challenges encountered when deploying Kafka, such as rebalancing issues, consumer group errors, and message delays, providing practical solutions to these problems. The findings presented here underscore Kafka's significance as a key enabler of robust and scalable EDA in contemporary software systems.

**Keywords**—Apache Kafka, Event-Driven Architecture, Microservices, Asynchronous Processing, Retry Mechanism, Distributed Systems, Messaging Systems, Scalability, Fault Tolerance, System Communication.

## I. INTRODUCTION

The world of software technology is slowly moving towards replacing monolithic systems with microservices. This approach has proven to be more systematic and maintainable to build sophisticated systems. Even the tech giants are investing huge resources into transforming their decade old systems to usher in this new era. Similarly, we should consider how communications are set up between systems, be it monolithic(with external services) or microservices. More so if there are many upcoming projects on making the concept of communication between systems more resilient and robust. One such tool is Kafka, which can be implemented as an Event-Driven Architecture approach for communication.

### A. Different Communication Approaches

Following are some of the techniques and approaches to make sure the conversation happens between two services/systems.

First is the API communication, which is the most commonly used approach for any kind of communication. All communication performed through HTTP and REST, comes under this category. Basically, API communication is synchronous, but later by utilizing webhooks a callback message can be sent to the client asynchronously.

Second is the Event-Driven Architecture. As the name suggests this is the type of communication that takes place only when a desired state is achieved or an event is triggered. This decouples the components by making them react to events and reduce the dependency on the server. The usages of message queues, pub/sub model, event streams, etc fall under being Event-Driven Architecture.

Third is by using a Shared Database. This type of communication is used rarely where any information transferred to other systems is done through inserting and retrieving from a database that is common to both the systems. But this technique has quite a bit of drawbacks such as limiting to the communicated message size, compromising of the database since it is shared and others.

### B. Importance of Event-Driven Architecture(EDA)

Event-Driven Architecture is a valuable and efficient model in any software development, as it increases the scope processing. EDA will help in following areas:

- Decoupling of components
- Asynchronous processing
- Realtime processing
- Robust, Resilient and Fault Tolerance
- Scalability

### C. Role of Kafka

Kafka plays an important role in enabling EDA. Kafka was initially designed for distributed streaming, which has now become a core component in building real-time pipelines and distributed stream processing applications. Kafka facilitates EDA by providing a durable, scalable, and distributed log of events. This makes kafka more ideal(subjective) for EDA compared to others. The reason kafka is a better EDA tool is mainly due to the Architectural design of Kafka with its components.

### D. Objectives

The objective of the work is to provide the following:

- Discuss different implementations of EDA using Kafka, including asynchronous processing, handling callbacks, retry mechanisms, and other key features.
- Identify use cases in industry practices where Kafka is utilized effectively.
- Develop solutions to common production problems with Kafka, such as rebalancing, consumer group errors, and multiple message consumption issues.

## II. USE CASE

### A. Event Tracking

Event tracking is a crucial aspect of modern applications, where the need to monitor and analyze user interactions across various services is paramount. Kafka provides a robust solution for centralizing and aggregating events happening across multiple microservices. By utilizing Kafka, all relevant events, such as click events, page landings, and user interactions, can be collected into a central analytics topic. This topic then streams the data to a monitoring system, which processes and displays the information on a dashboard in an analyzed format. This centralized event tracking mechanism enables real-time monitoring and analytics, offering insights into user behavior, system performance, and potential bottlenecks.

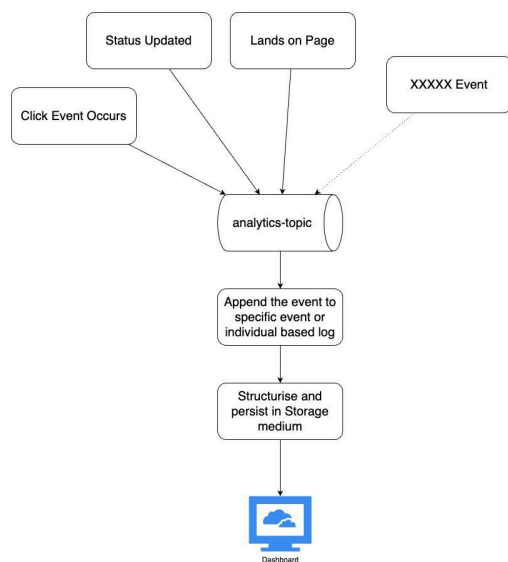


Fig: Event Analytics using Kafka

### B. Retry Mechanism

Handling failures in distributed systems is a challenging task, and manual intervention for retries can lead to inefficiencies and downtime. Kafka's retry mechanism is designed to overcome these limitations by automating the retry process for failed operations, such as API calls or other processing tasks. The proposed approach involves catching exceptions in a Kafka consumer that needs to retry processing after a specific time. The failed messages are then sent to a retry topic, where the system checks the retry count and the appropriate time for the retry.

If the retry count reaches zero, the message is sent to a dead-letter topic for further analysis. If the retry time is reached, the message is sent back to the original topic for reprocessing. If not, the message is placed in a delay-topic to wait until the retry time is met. This automated retry mechanism enhances system reliability and reduces the need for manual intervention, ensuring that operations are eventually completed or properly handled in case of repeated failures.

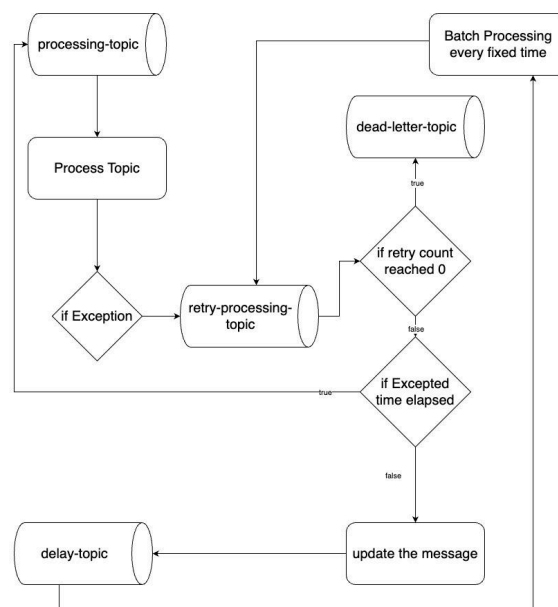


Fig: Retry Mechanism flow

C. Asynchronous Processing

Kafka excels in enabling asynchronous processing, which is essential for decoupling tasks and improving system responsiveness. In this proposed application, when an execution is triggered, Kafka is used to send a message to a process-topic, allowing the task to be executed asynchronously. Instead of waiting for the process to complete, the system can continue other operations. Typically, a database is used as a medium to track the completion status of the process. When the database is updated to reflect a pass or fail state, Kafka can trigger an appropriate response that updates the client. This method allows for efficient processing of tasks in parallel, reducing delays and improving the overall performance of the system.

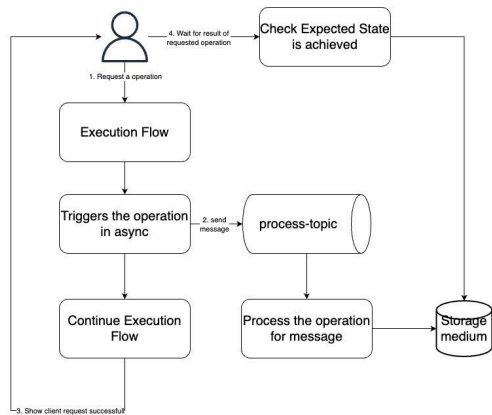


Fig: Asynchronous process implementation

D. Communication across Services

Kafka's ability to facilitate communication across services is one of its most significant advantages. This application extends the concept of asynchronous processing by enabling services to communicate with one another through Kafka topics.




Fig: Retry Mechanism flow

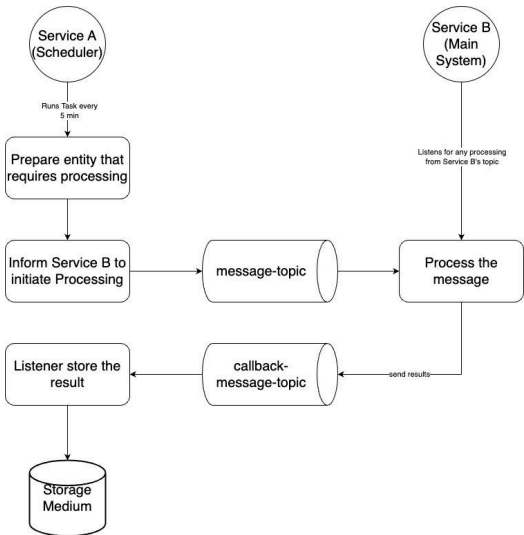


Fig: Two-way Communication using Kafka

In the proposed one-way communication scenario, a scheduler service sends a message to a Kafka topic that requires processing by another service. The scheduler then remains inactive while the processing service consumes the message and executes the required task.

In the two-way communication scenario, after processing the message, the processor sends a response back to the scheduler through a callback-topic. The scheduler, listening to the callback-topic, receives the response and can take further action. This setup allows for loose coupling between services, enabling them to interact without direct dependencies, thereby improving the modularity and scalability of the system.

III. COMPARISON

When comparing Kafka to other tools or methods for implementing Event-Driven Architecture (EDA) and microservices communication, several distinct advantages make Kafka a preferred choice. Here's a breakdown of why Kafka is often chosen over other alternatives:

A. Scalability

Kafka is designed to handle a massive amount of data. It's great for large-scale systems because you can keep adding more brokers to your Kafka cluster, which increases its capacity. This means Kafka can handle millions of messages per second without breaking a sweat. Traditional message brokers like RabbitMQ or ActiveMQ are solid, but they might struggle when dealing with the kind of high throughput that Kafka can manage easily.

### *B. Durability & Reliability*

Kafka uses something called a distributed log, which basically means it keeps all the messages that get published and makes sure they're delivered reliably to the subscribers. Kafka also has a replication feature that ensures your data is safe even if some of the nodes fail. Other systems like RabbitMQ also try to ensure durability, but they usually do it through message acknowledgment and persistence, which can slow things down. Kafka's approach is more straightforward and efficient, making sure no data is lost.

### *C. High Throughput with Low Latency*

Kafka is built for speed. Its design allows it to deliver a huge number of messages very quickly and with minimal delay. This is largely thanks to how Kafka partitions data so that consumers can read from these partitions independently, reducing bottlenecks. Traditional message queues might not be as fast, especially when they're trying to ensure that messages are durable and acknowledged. Kafka's setup just tends to be quicker and more efficient overall.

### *D. Decoupling of services*

One of the best things about Kafka is how it naturally decouples producers and consumers. This means that producers can send data to Kafka without needing to worry about who's going to consume it, and consumers can process that data whenever they're ready. This kind of decoupling is really important in microservices architectures because it helps keep services independent. While other message brokers like RabbitMQ or ActiveMQ can also do this, Kafka tends to make it easier and requires less complex configurations to achieve the same level of independence.

### *E. Replayability*

A unique feature of Kafka is its ability to retain messages for a set amount of time, which allows consumers to replay those messages if needed. This is super helpful for recovering from errors, auditing, or handling late-arriving consumers. Most traditional message brokers delete messages as soon as they're consumed, which means once a message is processed, it's gone for good. This limits the ability to recover from mistakes or reprocess data later. Kafka doesn't have that problem, which makes it much more flexible.

### *F. Community and Ecosystem support*

Kafka has a huge and active community, along with a strong ecosystem of tools and connectors. This makes it easier to integrate Kafka with other systems and get support when you need it. Other message brokers have their own communities too, but Kafka's ecosystem is often more extensive,

which can make a big difference if you're working in a complex, distributed environment.

## **IV. CHALLENGES & SOLUTIONS**

While Kafka is a powerful tool for implementing Event-Driven Architecture (EDA) and facilitating microservices communication, it's not without its challenges. However, understanding these challenges and how to address them can ensure smooth and effective Kafka deployments. Below are some of the common challenges faced when working with Kafka, along with practical solutions to overcome them:

### *A. Handling Consumer Groups*

Consumer groups are a core concept in Kafka, allowing multiple consumers to divide the work of processing messages. However, issues like consumer lag, rebalancing problems, and consumer group coordination failures can arise, leading to inconsistent message processing or delays.

Solution:

- **Monitoring and Alerts:** Implement robust monitoring for consumer lag and consumer group status using tools like Prometheus, Grafana, or Kafka's own JMX metrics. Setting up alerts can help you quickly identify and address issues before they impact the system.
- **Rebalancing Tuning:** Adjust the session timeout and heartbeat interval settings for consumers to ensure that rebalancing occurs smoothly. This involves finding the right balance between quick rebalancing and stability in consumer group operations.
- **Sticky Assignor:** Use the sticky partition assignor to minimize the number of reassignments during rebalancing, which can reduce the disruption caused by rebalancing events.

### *B. Topic Pollution*

In Kafka, messages might not always be processed immediately due to network latency, slow consumers, or backpressure in the system. Additionally, creating multiple topics to handle various scenarios, such as retries or dead-letter queues, can lead to topic pollution, making management more complex.

Solution:

- **Backpressure Handling:** Implement a retry mechanism that uses delay topics or exponential backoff strategies to manage retries without overwhelming the system. You can also use rate limiting to control the flow of messages into the system.
- **Topic Management:** Establish naming conventions and topic management policies



to prevent unnecessary topic creation. Use compacted topics or partitions efficiently to reduce the overhead of managing multiple topics.

- DLQ and Retry Topics: Set up Dead Letter Queues (DLQ) and retry topics effectively, ensuring that messages that fail processing can be retried or handled later without affecting the main data flow.

#### C. Data Consistency

Kafka's distributed nature can sometimes lead to data consistency issues, especially when dealing with duplicate messages or ensuring exactly-once delivery semantics. This can be critical in financial transactions or other sensitive data processing scenarios.

Solution:

- Idempotent Producers: Enable idempotence for Kafka producers to ensure that duplicate messages are not produced in the event of retries or network failures.
- Transactional Messaging: Use Kafka's transactional APIs to ensure that a series of writes are either committed together or not at all, which is crucial for maintaining consistency.
- Consumer Side Deduplication: Implement deduplication logic on the consumer side using unique message keys or hashes to ensure that duplicate messages are not processed multiple times.

#### D. Rebalancing Issues

Rebalancing is necessary for distributing partitions among consumers, but it can cause interruptions in processing and lead to inconsistent performance, especially in systems with frequent rebalancing.

Solutions:

- Minimize Rebalancing Frequency: Tune Kafka's `session.timeout.ms` and `max.poll.interval.ms` settings to reduce unnecessary rebalancing events. By extending these intervals, you can minimize the frequency of rebalances.
- Use Cooperative Rebalancing: Consider using Kafka's cooperative rebalancing protocol, which allows consumers to rebalance partitions incrementally rather than all at once, reducing the disruption caused during rebalancing.
- Monitor and Optimize Consumer Lag: Keep an eye on consumer lag and optimize consumer configurations to ensure that all consumers within a group are processing messages efficiently, thereby reducing the need for frequent rebalancing.

## V. CONCLUSION

In this paper, we explored the robust capabilities of Apache Kafka as a foundational tool for implementing Event-Driven Architectures (EDA) and facilitating seamless communication in microservices environments. Kafka's design as a distributed, scalable, and fault-tolerant messaging system makes it uniquely suited to handle the demands of modern, complex software systems.

We began by examining various applications of Kafka, such as event tracking, retry mechanisms, asynchronous processing, and service communication. These applications demonstrate how Kafka can centralize and streamline the flow of information across systems, enhance the reliability of processes, and ensure that services remain decoupled yet effectively coordinated.

Furthermore, we delved into the advantages of Kafka over other messaging systems, particularly its scalability, durability, low latency, and ability to maintain high throughput under heavy loads. Kafka's ability to handle large-scale data streams with minimal latency, coupled with its powerful features like replayability and idempotence, positions it as a leading choice for organizations looking to build resilient, real-time systems.

However, working with Kafka is not without its challenges. Issues such as consumer group errors, message delays, rebalancing, and ensuring data consistency require careful consideration and proactive management. We discussed practical solutions to these challenges, highlighting strategies like monitoring, backpressure handling, and leveraging Kafka's built-in features like idempotent producers and transactional messaging.

In conclusion, Kafka's strengths in handling event-driven communication and its versatility in addressing a wide range of use cases make it an indispensable tool for modern software architecture. While it comes with its own set of challenges, the benefits of implementing Kafka, especially in large-scale, distributed systems, far outweigh the difficulties. By understanding Kafka's architecture and applying best practices, organizations can build more robust, scalable, and efficient systems, ultimately driving greater innovation and resilience in their technological infrastructure.

## REFERENCES

- [1] Peddireddy, K. (2023). Streamlining Enterprise Data Processing, Reporting and Realtime Alerting using Apache Kafka. International Symposium on Digital Forensics and Security (ISDFS), 2023, 1-4.
- [2] Pelle, I., Szöke, B., Fayad, A., Cinkler, T., and Toka, L. (2023). A Comprehensive Performance Analysis of

- Stream Processing with Kafka in Cloud Native Deployments for IoT Use-cases. NOMS IEEE/IFIP Network Operations and Management Symposium, 2023, 1-6.
- [3] Adila, R., Nusantara, A. B., and Yuhana, U. L. (2023). Optimization Techniques for Data Consistency and Throughput Using Kafka Stateful Stream Processing. International Seminar on Research of Information Technology and Intelligent Systems (ISRITI), 2023, 480-485.
- [4] Sayar, A., Arslan, Ş., Çakar, T., Ertuğrul, S., and Akçay, A. (2023). High-Performance Real-Time Data Processing: Managing Data Using Debezium, Postgres, Kafka, and Redis. Innovations in Intelligent Systems and Applications Conference (ASYU), 2023, 1-4.
- [5] Vyas, S., Tyagi, R. K., Jain, C., and Sahu, S. (2022). Performance Evaluation of Apache Kafka – A Modern Platform for Real Time Data Streaming. International Conference on Innovative Practices in Technology and Management (ICIPTM), 2022, 465-470.
- [6] Srijiith, K. B. R., N, G., and M. R, A. (2022). Inter-Service Communication among Microservices using Kafka Connect. IEEE International Conference on Software Engineering and Service Science (ICSESS), 2022, 43-47.
- [7] Hugo, Å., Morin, B., and Svantorp, K. (2020). Bridging MQTT and Kafka to support C-ITS: a feasibility study. IEEE International Conference on Mobile Data Management (MDM), 2020, 371-376.
- [8] Wu, H., Shang, Z., and Wolter, K. (2019). TRAK: A Testing Tool for Studying the Reliability of Data Delivery in Apache Kafka. IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW), 2019, 394-397.
- [9] Alaasam, A. B. A., Radchenko, G., and Tchernykh, A. (2019). Stateful Stream Processing for Digital Twins: Microservice-Based Kafka Stream DSL. International Multi-Conference on Engineering, Computer and Information Sciences (SIBIRCON), 2019, 804-809.