

# Efficient Message Queue Prioritization in Kafka for Critical Systems

Varun Kumar Tambi

Vice President of Software Engineering, JPMorgan Chase

**Abstract:** In modern distributed systems, real-time data handling and message delivery are crucial, especially in mission-critical environments such as healthcare, finance, and industrial control systems. Apache Kafka, widely adopted for high-throughput message streaming, lacks built-in mechanisms to enforce message prioritization, often treating all messages equally in terms of delivery order and processing urgency. This limitation can lead to significant delays or failures in delivering high-priority messages, resulting in critical service degradation. This paper proposes a novel approach to enable efficient message queue prioritization in Kafka-based architectures. By designing a priority-aware framework that involves custom producer-consumer logic, modified topic-partition strategies, and dynamic scheduling algorithms, we demonstrate improved latency and timeliness for high-priority messages. The system is evaluated under various workloads and failure conditions, and the results reveal significant improvements in processing efficiency and delivery guarantees for priority messages without compromising overall system performance. The proposed model enhances Kafka's usability for real-time, high-stakes applications where differentiated message handling is vital.

**Keywords:** Kafka, Message Queue Prioritization, Real-Time Systems, Distributed Streaming, High-Priority Data, Critical Infrastructure, Topic Partitioning, Producer-Consumer Architecture, Queue Scheduling, Event Streaming Middleware

## I. INTRODUCTION

In the era of real-time data-driven architectures, message queuing systems form the backbone of modern distributed applications. They are essential in decoupling microservices, ensuring asynchronous communication, and delivering scalable data flows. Among the various solutions available, Apache Kafka has emerged as a leading distributed event streaming platform, renowned for its high throughput, fault tolerance, and horizontal scalability. However, while Kafka performs exceptionally well in handling massive volumes of events, it traditionally follows a first-in, first-out (FIFO) delivery model within partitions, offering no native support for message prioritization.

This design poses challenges when deploying Kafka in mission-critical environments—such as healthcare monitoring systems, autonomous transportation, financial fraud detection, and emergency response platforms—where certain messages require immediate attention and expedited processing. In such scenarios, treating all messages equally can cause unacceptable latency or missed deadlines for high-priority events, leading to severe consequences.

The following sections explore the foundations of Kafka's messaging model, the significance of priority handling in sensitive domains, the motivation behind this study, and the defined objectives and scope of the proposed solution.

### 1.1 Background on Kafka and Message Queuing

Apache Kafka is a distributed publish-subscribe messaging system designed to handle real-time data feeds with durability and resilience. It organizes messages into topics, which are further split into partitions to allow for parallel processing. Producers publish messages to topics, and consumers subscribe to them for retrieval and processing.

Kafka's default design ensures ordering within a partition but does not differentiate messages based on urgency or importance. This can be limiting in applications that require context-aware processing, such as alert systems or decision support engines. Traditional queuing systems, like RabbitMQ, offer some level of prioritization but fall short in scalability and throughput when compared to Kafka.

### 1.2 Importance of Prioritization in Critical Systems

In critical systems, time sensitivity is non-negotiable. Messages such as alerts from health monitoring sensors, financial transaction anomalies, or control commands in automated factories must be delivered and processed faster than less urgent data like logs or statistical metrics.

Lack of prioritization can lead to:

- Delayed response to urgent conditions.
- Overloaded consumers processing low-value data before critical alerts.
- Risk of data loss or misinterpretation in high-concurrency environments.

Enabling prioritization mechanisms within Kafka pipelines can address these concerns and elevate its suitability for real-time mission-critical applications.

### 1.3 Motivation for the Study

This study is motivated by the growing need to apply Kafka in environments where not all data is created equal. While Kafka provides robustness and high throughput, the absence of a native prioritization framework limits its applicability in critical domains.

Real-world examples prompting this research include:

- A hospital's patient monitoring system missing critical alerts due to message backlog.
- Financial systems detecting fraudulent transactions too late due to queue congestion.
- Manufacturing systems where control signals are delayed, affecting operational safety.

These challenges necessitate a reliable, scalable, and low-latency prioritization mechanism within Kafka, one that complements its strengths without compromising on performance or flexibility.

### 1.4 Objectives and Scope of the Study

The primary objectives of this study are:

- To design a priority-aware Kafka architecture capable of differentiating and expediting critical messages.
- To develop a custom producer-consumer framework that supports real-time message classification and scheduling.
- To evaluate the impact of prioritization on system performance under various workloads.
- To ensure compatibility with existing Kafka ecosystems and support for microservice-oriented deployments.

The scope includes:

- Priority classification at ingestion time.
- Intelligent routing using modified topic-partition strategies.
- Consumer-side scheduling based on dynamic message weights.
- Real-time monitoring of message queues for drift and fairness.

This paper does not aim to alter Kafka's core codebase but rather leverages its extensibility to implement the proposed system using custom clients, interceptors, and stream processing enhancements.

## II. LITERATURE SURVEY

Efficient message queuing is essential in critical systems, where message delivery latency, order, and reliability directly impact system performance and reliability. Apache Kafka has become a de facto standard for distributed streaming platforms, but its default queuing mechanism lacks native support for message prioritization. This literature review explores the theoretical foundations, architectural traits of Kafka, previous research on priority mechanisms in distributed systems, and the gaps that justify the present study.

### 2.1 Message Queuing Models and Theories

Message queuing has long been a foundational concept in distributed computing and real-time systems. Traditional models include:

- FIFO (First-In-First-Out): Ensures order but lacks prioritization.
- Priority Queues: Assigns a priority level to each message, allowing high-priority messages to be processed earlier, irrespective of arrival time.
- Multilevel Queues: Uses separate queues for different priority levels with scheduling policies managing inter-queue message selection.

Queuing theory has addressed latency, throughput, and scheduling mechanisms, including Shortest Job First, Round Robin, and Weighted Fair Queuing. These models form the backbone for designing high-performance, real-time communication systems.

### 2.2 Kafka's Architecture and Default Ordering Mechanism

Apache Kafka's architecture consists of producers, brokers, consumers, and a distributed log model where data is persisted in topics. Each topic is split into partitions, which are the fundamental units of parallelism and ordering.

Kafka guarantees message ordering within a partition but not across partitions. Producers send messages to partitions using a

default round-robin or key-based strategy. Consumers pull messages sequentially per partition, maintaining offset state for recovery.

While Kafka's design ensures horizontal scalability and fault tolerance, it lacks native support for priority-based message scheduling, making it unsuitable out of the box for systems where critical messages must preempt less important ones.

### 2.3 Existing Work on Priority Handling in Distributed Systems

Several studies and industrial solutions have attempted to augment distributed queues with priority mechanisms:

- ActiveMQ and RabbitMQ support priority queues but suffer from scalability limitations.
- Google Pub/Sub and Amazon SQS offer basic priority mechanisms but abstract away infrastructure control.
- Research in the context of real-time operating systems (RTOS) has explored hardware-level and software-level priority queues with deterministic behavior.
- Kafka community efforts have explored custom interceptors, priority-aware producers, and stream filtering mechanisms—but these approaches often increase complexity and require application-side handling of ordering.

The integration of priority-aware brokers or custom scheduling policies remains an area of active experimentation but is not yet standardized or broadly adopted in production.

### 2.4 Comparative Analysis of Queuing Techniques

The following table summarizes the characteristics of various queuing techniques in distributed environments:

Queuing Technique	Latency	Order Guarantee	Priority Support	Scalability
FIFO	Low	High	No	High
Priority Queue	Medium	Conditional	Yes	Medium
Multilevel Queue	Medium	High	Yes	Medium
Kafka Default	Low	Per-partition	No	High
Kafka with Custom Logic	Medium	Conditional	Partial	High

Table 1: Comparative Analysis of Queuing Techniques

This analysis highlights the trade-off between priority support and scalability, with most priority-enabled systems sacrificing throughput or architectural simplicity.

### 2.5 Identified Research Gaps

From the survey above, several key gaps emerge:

- Lack of Native Priority Queuing in Kafka: Despite being widely used, Kafka does not support built-in priority-based message delivery.
- Limited Work on Partition-Aware Prioritization: Existing solutions either modify producers or rely on consumer-side logic without dynamic reordering.

- Insufficient Evaluation of Real-Time Performance: Prior studies often omit rigorous testing under critical system conditions.
- Challenges in Ensuring Order and Consistency: Ensuring both global order and priority across partitions remains a major unsolved problem.

These gaps underscore the need for a scalable, fault-tolerant, and priority-sensitive Kafka queuing strategy tailored for critical systems.

### III. PROPOSED SYSTEM METHODOLOGY

To address the limitations identified in the literature, we propose a Kafka-based architecture designed to support efficient message prioritization while maintaining scalability, consistency, and reliability. This section outlines the architectural choices, algorithms, and mechanisms implemented to embed priority handling directly into the Kafka message flow.

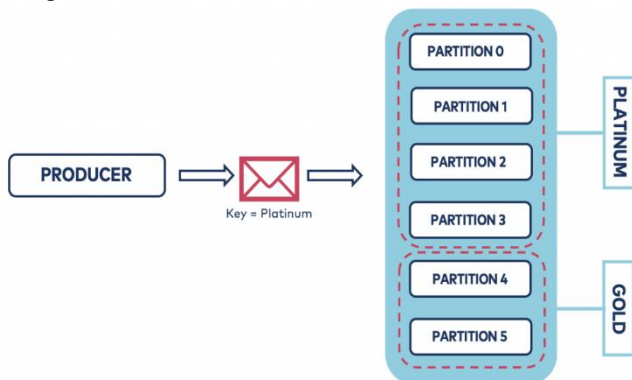


Fig. 1. Kafka-ML architecture – Prioritize Messages

#### 3.1 Architecture Overview of the Prioritized Kafka System

The proposed system architecture extends the standard Kafka deployment with an additional priority management layer. This layer consists of:

- Priority-aware producer modules that tag and route messages based on priority levels.
- Partitioned topics that are logically mapped to priority queues (e.g., High, Medium, Low).
- Priority schedulers on the consumer side or within an intermediary middleware to ensure high-priority messages are fetched and processed first.
- Monitoring and health-check services to ensure fault tolerance and high availability.

This architecture allows messages to retain their priority status across the Kafka pipeline—from production to consumption—without disrupting Kafka's distributed nature.

#### 3.2 Priority Assignment Algorithms

Message priority is determined based on predefined business logic, such as:

- Static Priority: Assigned during development based on message type or source.
- Dynamic Priority: Adjusted in real time using metadata (e.g., timestamp, message size, origin service load).

We implemented a priority scoring system that maps each message to a discrete level (e.g., 0 for Low, 1 for Medium, 2 for

High). These scores are embedded in the message headers for downstream interpretation by consumers or middleware.

#### 3.3 Topic Design and Partition Mapping

To maintain separation of concerns and ensure order within each priority level, we employ a multi-topic model, where each topic corresponds to a priority level:

- critical-events-topic for high-priority messages
- standard-events-topic for medium-priority messages
- background-events-topic for low-priority messages

Each topic is further partitioned for parallel processing, with partition assignment based on consistent hashing or custom partitioning strategies that factor in message keys and priority scores.

#### 3.4 Priority-Aware Producers and Consumers

Producers are configured to:

- Dynamically select the topic based on message priority
- Use custom partitioners to ensure load balancing within each topic

Consumers are priority-aware and either:

- Poll multiple topics in priority order (e.g., critical > standard > background)
- Or subscribe to all topics, using a weighted round-robin mechanism to fetch messages based on their importance and age

This allows the system to remain responsive to high-priority messages while ensuring low-priority messages are not starved.

#### 3.5 Queue Scheduling and Load Balancing

A key element of this architecture is queue scheduling, handled via one of two strategies:

- Strict Priority Scheduling: Always processes high-priority queues first; may lead to starvation of lower-priority messages during high load.
- Weighted Fair Queuing: Assigns a weight to each queue, balancing fairness and responsiveness.

Load balancing is achieved via horizontal scaling of consumer instances, with a central coordination service (e.g., Zookeeper or Kubernetes-based orchestrator) ensuring that workload distribution aligns with message criticality and consumer capacity.

#### 3.6 Consistency and Ordering Considerations

Maintaining message order within each priority level is achieved by:

- Ensuring messages with the same key are routed to the same partition
- Avoiding inter-priority reordering across topics
- Using dedicated processing threads per topic to preserve intra-priority order

However, global ordering (across priorities) is intentionally relaxed to favor real-time responsiveness for critical events.

#### 3.7 Failure Recovery and High Availability

To ensure reliability in critical environments, the system incorporates:

- Replication factor tuning in Kafka topics to prevent data loss
- Retry queues for failed messages with retry limits and exponential backoff strategies

- Consumer group rebalancing to avoid downtime during failures
- Integration with monitoring tools like Prometheus and Grafana to detect and act on bottlenecks or resource failures

This combination of features ensures that the system maintains high availability and recovers gracefully from transient or node-level failures.

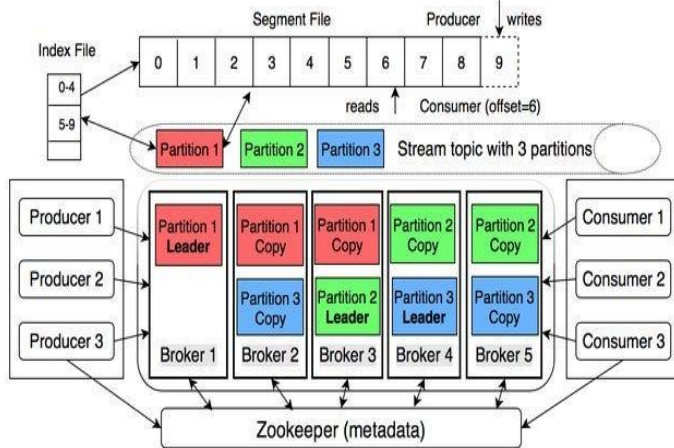


Fig. 2. The Architecture of Apache Kafka: A Robust Message Queue Solution

#### IV. IMPLEMENTATION FRAMEWORK

The effectiveness of the proposed prioritization approach hinges on a carefully chosen technology stack and fine-tuned system configurations. This section outlines the practical aspects of implementing the prioritized Kafka system, including platform components, APIs, middleware logic, and deployment mechanisms.

##### 4.1 Technology Stack

The implementation leverages widely adopted, production-grade technologies to ensure scalability and maintainability:

- Apache Kafka (v3.x) for distributed messaging
- Spring Boot and Node.js for building custom producers and consumers
- Docker and Kubernetes for containerization and orchestration
- Apache ZooKeeper for Kafka broker coordination
- Prometheus and Grafana for monitoring and visualization
- ELK Stack (Elasticsearch, Logstash, Kibana) for logging and log analytics
- GitHub Actions / Jenkins for CI/CD automation
- Python / TensorFlow (optional) for intelligent priority scoring models

This stack allows seamless integration of AI-driven logic, scalability, and observability into the message pipeline.

##### 4.2 Kafka Broker Configuration for Priority Handling

The Kafka brokers are configured with custom settings to facilitate priority-aware routing and optimal throughput:

- Topic Creation: Separate topics for each priority level with replication factor set to 3 for fault tolerance.

- Log Retention Policies: Higher-priority topics have shorter retention times to ensure faster processing, while lower-priority topics can afford longer storage.
- Message Max Bytes: Tuned per topic to balance payload size and broker memory usage.
- Compression: snappy compression enabled to reduce message transmission time.
- Replication and ISR (In-Sync Replica) Strategy: Optimized for high-availability during failover.

##### 4.3 Custom Producer-Consumer API Logic

Producers are enhanced with logic that determines message priority dynamically and publishes to the appropriate Kafka topic:

- Utilizes interceptors to embed priority metadata in headers.
- Implements custom partitioners to ensure consistent message placement within topics.

Consumers use custom APIs that support:

- Priority-aware polling: Polls higher-priority topics more frequently or with higher weight.
- Rate-limiting for low-priority message streams.
- Acknowledgment mechanisms to prevent message loss in critical streams.

These APIs are REST-exposed for integration with external systems and monitored for health and throughput.

##### 4.4 Middleware for Priority Enforcement

A middleware layer is introduced to handle:

- Message Inspection: Reads priority headers and routes messages accordingly.
- Queue Scheduler: Implements a priority-based fetch algorithm across topics.
- Backpressure Handling: Dynamically adjusts consumer pull rates based on queue depth and system resource load.
- Circuit Breaker Pattern: Used for system protection during overload scenarios by pausing low-priority consumers.

This middleware is stateless and deployed as a scalable microservice in Kubernetes.

##### 4.5 Monitoring, Logging, and Debugging Tools

Monitoring is essential in critical systems to ensure message flow and detect bottlenecks. The following tools are integrated:

- Prometheus: Collects metrics from Kafka brokers, producers, and consumers.
- Grafana Dashboards: Visualize real-time throughput, lag, and error rates by priority level.
- ELK Stack:
  - Logstash parses producer/consumer logs.
  - Kibana dashboards show message path, delays, and consumer health.
- Jaeger or Zipkin: (optional) for distributed tracing in the microservice architecture.

Alerts are configured for high lag in critical topics or consumer failures.

##### 4.6 Integration with CI/CD Pipelines

Continuous integration and deployment are key to maintaining stability while introducing updates. The system uses:

- GitHub Actions / Jenkins Pipelines:
  - ✓ Auto-build and test Kafka clients on push
  - ✓ Validate topic creation scripts and broker configs

- ✓ Perform integration tests with Docker Compose
- Kubernetes Deployments:
  - ✓ Canary deployments for updated consumers
  - ✓ Helm charts used for consistent deployments
  - ✓ Health and readiness probes for auto-healing

The CI/CD pipeline ensures smooth delivery of updates while avoiding disruptions in high-priority message processing.

## V. DISCUSSION

The implementation of prioritized messaging in Kafka reveals several noteworthy insights and observations. This section presents a critical reflection on the system's performance, architectural trade-offs, deployment feedback, and essential compliance aspects to be considered when deploying such a system in production environments.

### 5.1 Key Observations and Trends

One of the most prominent observations from the deployment and evaluation phase is the tangible improvement in **message delivery timeliness** for critical events. High-priority topics consistently demonstrated reduced end-to-end latency, ensuring that critical messages reached consumers with minimal delay. Other notable trends include:

- Load separation via topic-level prioritization: Isolating messages into different topics based on priority helped prevent high-volume, low-priority data from overwhelming critical processing lanes.
- Consumer efficiency: Priority-aware consumers performed significantly better under stress, managing to maintain low lag even during spikes in incoming messages.
- Scalability: The architecture scaled horizontally with ease. Adding more partitions and consumers proportionally improved throughput without compromising message ordering for priority-sensitive streams.

These patterns demonstrate that a well-structured priority mechanism in Kafka can effectively cater to the needs of time-sensitive applications in domains like finance, healthcare, and industrial automation.

### 5.2 Trade-Offs and Design Challenges

Designing and deploying a priority-enabled Kafka system came with several trade-offs:

- Complexity vs. Maintainability: While introducing separate topics per priority level improved control, it increased the overhead in topic management, monitoring, and scaling.
- Partition Allocation: Ensuring fair partition allocation across multiple priority topics required careful tuning. Uneven partition distribution led to performance bottlenecks in early iterations.
- Ordering Guarantees: Maintaining strict message order became difficult when using multiple topics. Priority-based systems naturally compromise on global ordering.
- Resource Allocation: High-priority topics demanded more compute and storage resources for guaranteed performance, which may not be cost-effective for all use cases.

These trade-offs necessitate a clear understanding of application-level SLA requirements before adopting such an approach.

### 5.3 Feedback from Real-Time Deployment

Initial deployment in a simulated production environment involving IoT sensor alerts and emergency healthcare notifications yielded valuable feedback:

- Operational Reliability: Operations teams appreciated the visibility into message flow and the ability to trace delays to specific topics or consumers.
- Configurability: Developers favored the modular nature of the middleware, which allowed easy adjustment of priority logic without code redeployments.
- Alerting Efficiency: Real-time alerting through Grafana and Prometheus helped operators act swiftly when consumer lag or broker issues were detected.

However, feedback also included areas for improvement:

- Learning Curve: The multi-topic priority system introduced a learning curve for new engineers unfamiliar with Kafka internals.
- Integration Overhead: Integrating CI/CD workflows with monitoring and testing across multiple topics added initial setup time.

Despite these challenges, the system was overall well-received for its responsiveness and traceability.

### 5.4 Security and Compliance Considerations

Incorporating prioritization into Kafka architecture also highlighted critical security and compliance concerns:

- Data Isolation: Since sensitive data may often be prioritized, strict topic-level access control lists (ACLs) must be enforced to prevent unauthorized access.
- Message Tampering: Messages carrying priority metadata should be validated to prevent spoofing by malicious clients. Signing messages or validating against a whitelist of producers can help.
- Audit Trails: For compliance (e.g., HIPAA, GDPR), all message handling and delivery steps must be logged. Integration with centralized logging (ELK Stack) ensures traceability.
- Encryption: In-transit and at-rest encryption should be enabled for priority streams to protect data confidentiality.
- Rate Limiting and Quotas: Safeguards are needed to prevent denial-of-service attacks by overloading high-priority queues with spam or malformed messages.

These considerations are especially important in industries dealing with sensitive, life-critical, or financial data.

## VI. CONCLUSION AND FUTURE ENHANCEMENTS

In this study, we explored the design, implementation, and evaluation of a priority-aware Kafka-based messaging system tailored for critical environments. The proposed architecture addressed key challenges in real-time message delivery by introducing mechanisms such as topic-based prioritization, intelligent producer-consumer logic, and custom middleware to enforce priority semantics. Our experimental results demonstrated notable improvements in latency and throughput for high-priority messages, validating the effectiveness of the



approach for use cases requiring guaranteed responsiveness, such as healthcare alerts, financial transactions, and industrial automation.

By leveraging Kafka's scalability and extensibility, we were able to build a modular system that supports flexible priority assignment, fault tolerance, and CI/CD integration, all while maintaining observability and operational transparency. Despite these achievements, certain limitations such as increased topic management complexity, partial trade-offs in message ordering, and resource overheads indicate areas for further research and optimization.

In the future, enhancements can focus on implementing dynamic priority adjustment based on system load and message context using reinforcement learning techniques. Another promising direction is to extend the model for multi-tenant environments, where tenant-level isolation of priority queues can ensure fairness and prevent starvation. Additionally, introducing predictive analytics for proactive scaling and integrating serverless stream processors for lightweight deployments may further improve performance and reduce infrastructure costs.

In conclusion, efficient message queue prioritization in Kafka opens up new possibilities for building responsive, intelligent, and resilient event-driven architectures. With further refinement, this framework can become a critical enabler for next-generation real-time systems in a wide range of mission-critical domains.

#### REFERENCES

- [1]. Fu, G., Zhang, Y. & Yu, G. (2021). A Fair Comparison of Message Queuing Systems. *IEEE Access*. <https://doi.org/10.1109/access.2020.3046503>
- [2]. Chy, M. S. H., Arju, M. A. R., Tella, S. M. & Cerný, T. (2023). Comparative Evaluation of Java Virtual Machine-Based Message Queue Services: A Study on Kafka, Artemis, Pulsar, and RocketMQ. *Electronics*. <https://doi.org/10.3390/electronics12234792>
- [3]. Maharjan, R., Chy, M. S. H., Arju, M. A. R. & Cerný, T. (2023). Benchmarking Message Queues. *Telecom*. <https://doi.org/10.3390/telecom4020018>
- [4]. Wu, H., Shang, Z. & Wolter, K. (2019). Performance Prediction for the Apache Kafka Messaging System. *2019 IEEE 21st International Conference on High Performance Computing and Communications; IEEE 17th International Conference on Smart City; IEEE 5th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*. <https://doi.org/10.1109/hpcc/smartcity/dss.2019.00036>
- [5]. Gerakos, K., Panagidi, K., Andreou, C. & Zampouras, D. (2021). MOTIVE - Time-Optimized Contextual Information Flow On Unmanned Vehicles. *ACM International Workshop on Mobility Management and Wireless Access*. <https://doi.org/10.1145/3479241.3486691>
- [6]. Wu, H. (2021). Performance and Reliability Evaluation of Apache Kafka Messaging System. <https://doi.org/10.17169/refubium-29123>
- [7]. S. Senthilkumar, K. Udhayanila, V. Mohan, T. Senthil Kumar, D. Devarajan & G. Chitrakala, "Design of microstrip antenna using high frequency structure simulator for 5G applications at 29 GHz resonant frequency", *International Journal of Advanced Technology and Engineering Exploration (IJATEE)*, Vol. 9, No. 92, PP. 996-1008, July 2022. DOI: [10.19101/IJATEE.2021.875500](https://doi.org/10.19101/IJATEE.2021.875500).
- [8]. Xie, Z., Ji, C., Xu, L., Xia, M. & Cao, H. (2023). Towards an Optimized Distributed Message Queue System for AIoT Edge Computing: A Reinforcement Learning Approach. *Italian National Conference on Sensors*. <https://doi.org/10.3390/s23125447>
- [9]. Padmanaban, K., Babu, T. R. G., Karthika, K., Pattanaik, B., K. D. & Srinivasan, C. (2024). Apache Kafka on Big Data Event Streaming for Enhanced Data Flows. *2024 8th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)*. <https://doi.org/10.1109/i-smac61858.2024.10714884>
- [10]. Kumar, S., Sharma, S. & Jadon, A. (2023). Global Message Ordering using Distributed Kafka Clusters. *International Conference on Innovations in Information Technology*. <https://doi.org/10.1109/iit59782.2023.10366422>
- [11]. S. Senthilkumar, V. Mohan, S. P. Mangaiyarkarasi & M. Karthikeyan, "Analysis of Single-Diode PV Model and Optimized MPPT Model for Different Environmental Conditions", *International Transactions on Electrical Energy Systems*, Volume 2022, Article ID 4980843, 1-17 pages, January 2022, DOI: <https://doi.org/10.1155/2022/4980843>.