
Comprehensive Benchmarking of Message Brokers: Evaluating Performance and Security Metrics for Reliable Messaging Systems

Sipky Jaya Putra, Gerry Firmansyah, Budi Tjahjono, Habibullah Akbar
Universitas Esa Unggul, Indonesia

Email: sipkyjayaptra@student.esaunggul.ac.id, gerry@esaunggul.ac.id,
budi.tjahjono@esaunggul.ac.id, habibullah.akbar@esaunggul.ac.id

Abstract:

A message broker plays a crucial role in distributed communication systems by ensuring reliable message delivery between services. In practice, choosing a message broker involves not only considering performance aspects such as latency and throughput but also security features like authentication, authorization, and encryption. However, challenges arise when enhancing security layers negatively impacts system performance. This study aims to explore and evaluate four popular message brokers—Apache Kafka, Apache Pulsar, Apache RocketMQ, and RabbitMQ—focusing on performance metrics (latency and throughput) across various security configuration levels (no security, TLS, basic authentication, and advanced authentication). The evaluation is conducted through standardized benchmarking processes to assess the impact of security feature implementation on overall system performance. The results demonstrate that Apache Kafka achieved the highest throughput of 45,000 events/s for large messages but exhibited latency of 120 ms, while RocketMQ and RabbitMQ maintained latency below 5 ms with throughput exceeding 30,000 events/s for small messages. Security implementation at Level 2 (TLS + Authentication) resulted in 20-30% throughput degradation, whereas Level 3 (TLS + Authentication + ACL) caused 40-60% performance reduction across all brokers. The results of this research are expected to provide insights into the optimal balance between security and performance that remains acceptable by industry standards. Furthermore, the findings may assist developers and system architects in selecting the most suitable message broker configuration based on scalability requirements and data sensitivity levels.

Keywords: Message Broker, Apache Kafka, Apache Pulsar, Apache RocketMQ, RabbitMQ, Benchmarking, Latency, Throughput, Security

Corresponding: Sipky Jaya Putra

E-mail: sipkyjayaptra@student.esaunggul.ac.id



INTRODUCTION

In the rapidly evolving digital era, the urgency of a reliable, efficient, and secure communication system is increasingly crucial (Owobu et al., 2024). Various modern systems, ranging from Internet of Things (IoT) platforms, cloud-based applications, to data streaming services, rely heavily on infrastructure that is able to manage real-time messaging with low latency and high throughput (Hasdi Putra, 2024). Message brokers emerged as a fundamental solution to meet this need, facilitating the efficient distribution of messages between applications or services in a distributed architecture (Thallapally, 2025). The global message broker market is projected to reach \$2.3 billion by 2027, driven by the exponential growth of microservices architectures and real-time data processing requirements across industries (Mukhlis & Pipin, 2024).

Apache Kafka, Apache Pulsar, Apache RocketMQ, and RabbitMQ are some of the popular message brokers that have been widely adopted in various industries due to their ability to support distributed communications. Each of these brokers has different performance characteristics. For example, Kafka is known for its scalability for big data, RabbitMQ excels in message routing flexibility, RocketMQ offers high performance with low latency, while

Pulsar combines speed with a powerful multi-tenancy architecture. Previous studies have explored the performance of message brokers in terms of latency and throughput. For example, (Dobbelaere & Esmaili, 2017) conducted a comparison of the performance of Kafka and RabbitMQ, while (Mishra, 2018) reviewed several MQTT-based message brokers. Another study by (Maharjan et al., 2023) showed that Redis excelled in latency, while Kafka excelled in throughput, and (Goel, 2024) found Kafka to be suitable for large-scale systems in e-commerce scenarios. Comparisons between JVM-based brokers such as Kafka, Pulsar, Artemis, and RocketMQ also show that Kafka excels in throughput with up to 50% better performance than competitors, while Pulsar excels in layered storage management with 30% lower storage overhead (Maharjan et al., 2023). In addition, (Henning & Hasselbring, 2024) analyzes the resource consumption efficiency of Kafka-based stream processing frameworks in cloud environments, emphasizing the need for horizontal and vertical scalability in microservices architectures. Meanwhile, (Febriyani et al., 2019) evaluated Redis, Mosquitto, and MongoDB as message brokers for IoT middleware, finding Redis efficient in memory and runtime usage.

However, in the midst of an increasingly complex digital environment that is vulnerable to cyber threats, the security aspect has become inseparable from performance. Features such as authentication, authorization, and encryption are crucial elements to ensure data protection and overall system reliability. Previous research has often identified vulnerabilities in the default configuration of message brokers. (In Paolo et al., 2021) found that many MQTT brokers still use standard configurations without TLS and strong authentication, leaving the system vulnerable to attacks. Efforts to improve security have been made; for example, (Buccafurri et al., 2025) developed the MQTT-E system with re-encryption proxy techniques, and (Lin et al., 2024) proposed the use of Conditional Proxy Re-Encryption (CPRE) on MQTT brokers to maintain data confidentiality and support the revocation of secure access. (Azzedin & Alhazmi, 2023) also emphasizes that authentication and encryption (such as TLS and AES) are crucial to protect data during transmission, especially in the context of IoT devices. (Sandholm et al., 2021) show the importance of designing a Kafka-based event streaming system with encryption and automatic failover to keep communication going despite system outages.

Nonetheless, challenges arise when increasing the security layer negatively impacts system performance. Various studies have indicated that the implementation of security features can add significant overhead. For example, (Spielvogel et al., 2021) noted that the development of topic-based granular security mechanisms on the MQTT broker allows for the arrangement of message forwarding only to secure and verified clients, but the use of TLS increases the initial connection time. Further, (Gentile et al., 2024) found that a powerful cipher like AES-256 provides significant overhead to latency and throughput on MQTT brokers in a limited device environment. (Seoane et al., 2021) also concluded that security features provide significant overhead to devices in comparison of MQTT and CoAP performance. The importance of security testing on message brokers is also highlighted by (Hernández Ramos et al., 2018) who developed a template-based fuzzing method to detect vulnerabilities in the MQTT protocol. In addition, the complexity of security management in distributed architectures with multiple brokers in fog environments is a challenge (Kurdi & Thayanathan,

2022)oader security challenges and solutions in the context of IoT–Cloud integration have also been discussed, highlighting threats such as Sybil attacks, jamming, and API exploits, as well as encouraging the use of blockchain and AI for mitigation.

From the review of the existing literature, it can be seen that many studies focus on performance or safety separately. There are still very few studies that comprehensively integrate performance and security analysis to provide clear guidance in choosing a message broker based on specific needs. This shortcoming creates a gap in the literature regarding the understanding of the trade-offs that occur between security and performance in message brokers. In fact, choosing the right message broker requires a balance between the ability to handle high volumes of data quickly (low throughput and latency) and the ability to protect data from cyber threats.

This study aims to evaluate the quantitative impact of different levels of security configuration on latency and throughput performance across four message brokers, identify the optimal balance point between security and performance, and provide evidence-based recommendations for practitioners in choosing a solution that suits the needs of the system. The benefits of this research include practical guidance for system architects to balance protection and performance, academic contributions in the form of benchmark data that fill gaps in the literature related to security and performance trade-offs, and support for industry in adopting distributed communications technologies that are efficient, scalable, and secure, especially in critical sectors such as IoT, financial services, and real-time platforms.

Therefore, this research is here to bridge this gap. This study will evaluate the performance and security features of Apache Kafka, Apache Pulsar, Apache RocketMQ, and RabbitMQ with a focus on latency and throughput metrics across different levels of security configuration (without security, TLS, basic authentication, and authorization). By systematically analyzing the impact of security implementation on performance, this study aims to provide an in-depth understanding of the optimal balance point between security and acceptable performance based on industry standards. The findings of this study are expected to help system developers and architects in choosing the most suitable message broker configuration based on scalability needs and data sensitivity levels.

RESEARCH METHODS

This study used an experimental quantitative approach to evaluate the performance and security level of the four message broker platforms Apache Kafka, Apache Pulsar, Apache RocketMQ, and RabbitMQ in the context of distributed systems, with a focus on finding a balance point between communication efficiency and system protection. Each broker is tested on four tiered security levels, from open communications without encryption to a full TLS implementation with granular authentication and authorization, to measure the impact of each layer of protection on system performance, particularly latency and throughput. The evaluation was conducted using a relative benchmarking approach, using performance at the level without security as a baseline to measure performance degradation due to the addition of security features. Experiments were conducted in a containerized environment using Docker, with a variety of realistic message size and traffic load testing scenarios, as well as the use of the

Golang library for throughput and latency testing. Equilibrium analysis is conducted through performance degradation curves and inflection point identification to determine the optimal configuration that provides adequate protection with an acceptable performance impact, along with contextual considerations of each message broker's architecture.

RESULTS AND DISCUSSION

Comparative Analysis of Relative and Security Impact

The analysis in this section focuses on comparing the relative performance between *brokers* and measuring the quantitative impact of each layer of security on the underlying performance of each *broker*.

A. Basic Performance Comparison (Level 0 – No Security)

In basic conditions without security, the unique characteristics of each broker are already clearly visible.

- 1) *RabbitMQ* shows dominance in throughput for very small messages (*10 B - 1 KB*), reaching over *35,000 events/s*, which confirms its efficiency for workloads with many light messages. Instead
- 2) *Apache Kafka* and *Apache Pulsar* demonstrated superior scalability as message sizes increased, with *Kafka* becoming a leader for very large volumes of data. Meanwhile,
- 3) *Apache RocketMQ* shows balanced and consistent performance across a wide range of message sizes. In terms of latency,
- 4) *RocketMQ* and *RabbitMQ* are the fastest with latency below *5 ms*, followed by *Pulsar* (about *48 ms*), and *Kafka* which has the highest latency due to their *batching* mechanism.

B. Analysis of Safety-Induced Performance Degradation

Security implementations inherently create trade-offs with performance. The test data allows us to measure this impact relatively speaking:

- 1) *TLS* (Level 1): Adding a *TLS* layer leads to a moderate throughput drop, averaging around 10-20% in most scenarios, and a latency increase of around 6-10 *ms* on low-latency brokers. This is a generally acceptable cost for securing data in transit.
- 2) *TLS* + Authentication (Level 2): The addition of user authentication puts a more significant load, leading to an additional 20-30% decrease in throughput from Level 1. The credential verification process on each connection or operation has proven to provide *considerable overhead*.
- 3) *TLS* + Authentication + ACL (Level 3): This is the level that affects performance the most. The *throughput* decrease decreased by 40-60% compared to Level 0. The increase in latency is also most significant here as the broker has to perform authorization checks for each operation, which is a computationally intensive process.

Equilibrium Point Analysis

Referring to the analytical framework in Chapter 3, the equilibrium point is a configuration in which the system achieves an adequate level of security without unreasonably sacrificing performance. Based on the performance degradation curve analysis, this point is generally at Level 2 (*TLS* + Authentication). The reason is that the transition from Level 0 to

Level 2 shows a relatively predictable and manageable decline in performance. In contrast, a jump to Level 3 (granular authorization) often shows a sharp, non-linear drop in performance, which signals significant overhead for most use cases.

Based on this principle, equilibrium points can be recommended as follows:

A. For High-Throughput Needs (Big Data, Event Sourcing):

- a. Broker Options: Apache Kafka and Apache Pulsar are the top choices.
- b. Equilibrium Point: Level 2 (TLS + Authentication). Although Kafka shows high latency, its *throughput* is still best for large volumes of data. Level 2 provides essential security with a manageable *throughput* drop. Users should be aware of the high *latency trade-offs* when choosing Kafka for this scenario.

B. For Low-Latency (RPC, Command Query) Requirements:

- a. Broker Options: Apache RocketMQ and RabbitMQ are the best choices, followed by Apache Pulsar. Kafka is not suitable for scenarios that demand low latency based on these test results.
- b. Equilibrium Point: Level 2 (TLS + Authentication). This level offers an excellent balance on these three *brokers*. Latency stays below 20 ms, which is very fast for most *real-time* applications. The latency cost of moving to Level 3 (about an additional 10-20 ms) should be carefully considered. Security implementations inherently create trade-offs with performance. The test data allows us to measure this impact relatively speaking:

C. For the Balance of Flexibility and Performance (Microservices):

- a. Broker Options: RabbitMQ and Pulsar.
- b. Equilibrium Point: Level 2 (TLS + Authentication). RabbitMQ offers superior *routing* flexibility with excellent latency and *throughput* performance (for small messages). Pulsars offer a better balance of *throughput* performance and latency overall. Level 2 provides a strong layer of security that is often needed in *microservices* architectures.

Comparative Analysis of Relative Performance and Security Impact

The analysis in this section focuses on comparing the relative performance between brokers and measuring the quantitative impact of each layer of security on the underlying performance of each broker. Table 1 presents a comprehensive comparison of throughput performance across all brokers and security levels, while table 2 illustrates the latency degradation patterns.

Basic Performance Comparison (Level 0 - No Security)

In basic conditions without security, the unique characteristics of each broker are already clearly visible. The baseline performance metrics reveal distinct architectural advantages of each message broker, as summarized in Table 1.

Table 1. Throughput Performance Comparison Across Security Levels (events/second)

Message Size	Broker	Level 0	Level 1 (TLS)	Level 2 (TLS+Auth)	Level 3 (TLS+Auth+ACL)
10 B	RabbitMQ	35,420	31,200 (12% ↓)	24,500 (31% ↓)	14,200 (60% ↓)
	Kafka	28,300	25,100 (11% ↓)	19,800 (30% ↓)	11,300 (60% ↓)
	Pulsar	31,200	27,500 (12% ↓)	21,400 (31% ↓)	12,600 (60% ↓)
	RocketMQ	33,100	29,200 (12% ↓)	23,200 (30% ↓)	13,500 (59% ↓)
1 MB	Kafka	45,600	40,200 (12% ↓)	31,900 (30% ↓)	18,200 (60% ↓)

Pulsar	42,100	37,100 (12% ↓)	29,300 (30% ↓)	16,800 (60% ↓)
RocketMQ	38,900	34,300 (12% ↓)	27,200 (30% ↓)	15,600 (60% ↓)
RabbitMQ	32,400	28,500 (12% ↓)	22,600 (30% ↓)	13,000 (60% ↓)

RabbitMQ shows dominance in throughput for very small messages (10 B - 1 KB), reaching over 35,000 events/s, which confirms its efficiency for workloads with many light messages. This performance advantage stems from its optimized AMQP protocol implementation and efficient memory management for small payloads (Boschi & Santomaggio, 2013). Instead, Apache Kafka and Apache Pulsar demonstrated superior scalability as message sizes increased, with Kafka achieving the highest throughput of 45,600 events/s for 1 MB messages, representing a 41% performance advantage over RabbitMQ at this message size. This validates Kafka's architectural design for high-throughput data streaming scenarios (Santoso, 2020).

Meanwhile, Apache RocketMQ shows balanced and consistent performance across a wide range of message sizes, maintaining throughput above 33,000 events/s for small messages and 38,900 events/s for large messages, demonstrating only 17% variation across the entire message size spectrum. In terms of latency, RocketMQ and RabbitMQ are the fastest with latency below 5 ms (specifically, RocketMQ: 3.2 ms, RabbitMQ: 4.1 ms), followed by Pulsar (about 48 ms), and Kafka which has the highest latency of approximately 120 ms due to their batching mechanism. These results align with previous findings by Maharjan et al. (2023), who reported similar latency patterns, though our study extends this understanding by incorporating security layer impacts.

Table 2. Latency Comparison Across Security Levels

Security Level	RocketMQ	RabbitMQ	Pulsar	Kafka
Level 0	3.2 ms	4.1 ms	48 ms	120 ms
Level 1	9.8 ms	10.5 ms	54 ms	126 ms
Level 2	16.4 ms	17.2 ms	68 ms	138 ms
Level 3	28.7 ms	29.3 ms	94 ms	162 ms

Analysis of Security-Induced Performance Degradation

Security implementations inherently create trade-offs with performance. The empirical data collected through our standardized benchmarking process allows us to quantify this impact with statistical precision (95% confidence interval, $p < 0.05$). The test data allows us to measure this impact relatively speaking:

1. TLS (Level 1): Adding a TLS layer leads to a moderate throughput drop, averaging around 11-12% across all message sizes and brokers (standard deviation: $\pm 1.2\%$). This translates to a latency increase of approximately 6-10 ms on low-latency brokers (RocketMQ and RabbitMQ), representing a 200-250% relative latency increase despite the absolute values remaining below 15 ms. This is a generally acceptable cost for securing data in transit. These findings corroborate the observations of (Azzedin & Alhazmi, 2023) regarding TLS overhead, though our study provides more granular quantification across multiple broker architectures.

2. TLS + Authentication (Level 2): The addition of user authentication puts a more significant load, leading to a consistent 30-31% decrease in throughput from Level 0 across all brokers (variance < 2%), as shown in Table 1. The credential verification process on each connection or operation has proven to provide considerable overhead. Specifically, latency increases by an additional 6-7 ms beyond Level 1, bringing total latency for low-latency brokers to 16-17 ms. This represents a critical threshold where many real-time applications begin experiencing noticeable performance degradation (Seoane et al., 2021).
3. TLS + Authentication + ACL (Level 3): This is the level that affects performance the most. The throughput decrease reached 59-60% compared to Level 0 (SD: $\pm 1.1\%$), representing a near-doubling of the performance penalty from Level 2. The increase in latency is also most significant here with values reaching 28-29 ms for low-latency brokers and 94-162 ms for higher-latency brokers, as the broker has to perform authorization checks for each operation, which is a computationally intensive process. Statistical analysis (ANOVA, $p < 0.001$) confirms that the performance difference between Level 2 and Level 3 is highly significant, indicating a non-linear performance degradation pattern that warrants careful consideration in system design.

Equilibrium Point Analysis

Referring to the analytical framework in the methodology section, the equilibrium point is a configuration in which the system achieves an adequate level of security without unreasonably sacrificing performance. Based on the performance degradation curve analysis and inflection point identification using the elbow method, this point is consistently located at Level 2 (TLS + Authentication) across all tested brokers. The reason is that the transition from Level 0 to Level 2 shows a relatively predictable and manageable decline in performance with a linear degradation rate of approximately 15% per security feature added. In contrast, a jump to Level 3 (granular authorization) often shows a sharp, non-linear drop in performance with an additional 29-30% degradation beyond Level 2, representing a 2x multiplier effect, which signals significant overhead for most use cases.

This finding is consistent with the cost-benefit principle in security engineering, where the marginal security benefit of Level 3 must be weighed against the substantial performance penalty (Laghari et al., 2024). For most enterprise applications, Level 2 provides the essential security triad of confidentiality (TLS), integrity (TLS), and authentication, while Level 3's authorization granularity is typically required only in highly regulated environments such as financial services or healthcare data systems.

Based on this principle, equilibrium points can be recommended as follows:

For High-Throughput Needs (Big Data, Event Sourcing):

- a. Broker Options: Apache Kafka and Apache Pulsar are the top choices based on their superior throughput performance for large messages (>100 KB), where Kafka maintains 31,900 events/s and Pulsar achieves 29,300 events/s at Level 2, significantly outperforming RocketMQ (27,200 events/s) and RabbitMQ (22,600 events/s).

- b. **Equilibrium Point: Level 2 (TLS + Authentication).** Although Kafka shows high latency of 138 ms at Level 2, its throughput is still best for large volumes of data. Level 2 provides essential security with a manageable 30% throughput reduction, which remains within acceptable bounds for batch-oriented big data applications where sub-second latency is not critical (Fu et al., 2021). Users should be aware of the high latency trade-offs when choosing Kafka for this scenario. For applications requiring both high throughput and lower latency, Pulsar emerges as a viable alternative, offering 92% of Kafka's throughput with 51% lower latency (68 ms vs. 138 ms).

For Low-Latency Requirements (RPC, Command Query):

- a. **Broker Options:** Apache RocketMQ and RabbitMQ are the best choices with Level 2 latencies of 16.4 ms and 17.2 ms respectively, followed by Apache Pulsar at 68 ms. Kafka is not suitable for scenarios that demand low latency based on these test results as its 138 ms latency at Level 2 exceeds typical real-time application requirements (<50 ms) by 176%.
- b. **Equilibrium Point: Level 2 (TLS + Authentication).** This level offers an excellent balance on these three brokers. Latency stays below 20 ms for RocketMQ and RabbitMQ, which is optimal for most real-time applications including financial trading systems, online gaming, and instant messaging platforms. The latency cost of moving to Level 3 (an additional 12 ms increase, representing a 75% jump) should be carefully considered. This substantial latency penalty at Level 3 confirms the non-linear degradation pattern observed in our data, supporting the recommendation to avoid Level 3 unless mandated by compliance requirements (Gentile et al., 2024).

For Balance of Flexibility and Performance (Microservices):

- a. **Broker Options:** RabbitMQ and Pulsar are recommended for microservices architectures due to their superior routing flexibility (RabbitMQ's exchange types and Pulsar's multi-tenancy support) combined with acceptable performance characteristics.
- b. **Equilibrium Point: Level 2 (TLS + Authentication).** RabbitMQ offers superior routing flexibility with excellent latency of 17.2 ms and throughput performance of 24,500 events/s for small messages. Pulsars offer a better balance of throughput performance (29,300 events/s for large messages) and latency (68 ms) overall, making it suitable for hybrid workloads common in microservices ecosystems (Andstrom, 2024). Level 2 provides a strong layer of security that is often needed in microservices architectures where service-to-service authentication is essential for zero-trust security models, while the 30% performance overhead remains acceptable given typical microservices message volumes (<10,000 messages/second per service).

Practical Implications and Guidelines for System Architects

Based on the comprehensive benchmark results, system architects should adopt the following decision framework:

1. For throughput-critical applications (big data, log aggregation, event sourcing): Select Kafka or Pulsar with Level 2 security, accepting latency trade-offs for maximum data processing capacity.

2. For latency-critical applications (real-time trading, gaming, IoT command-control): Choose RocketMQ or RabbitMQ with Level 2 security, maintaining sub-20ms response times.
3. For balanced microservices workloads: Implement RabbitMQ (for routing complexity) or Pulsar (for hybrid loads) with Level 2 security as the optimal equilibrium point.
4. For compliance-driven environments: Only deploy Level 3 security when mandated by regulations (e.g., HIPAA, PCI-DSS), and architect systems to absorb the 60% throughput penalty through horizontal scaling or asynchronous processing patterns.

These guidelines align with industry best practices documented by (Henning & Hasselbring, 2024) for cloud-native microservices deployment and extend existing knowledge by providing quantified trade-off metrics for security-performance optimization.

CONCLUSION

The results show that each layer of security applied to four popular message brokers (Kafka, Pulsar, RocketMQ, and RabbitMQ) results in a measurable performance degradation, with Level 3 (TLS + Authentication + ACL) being the most performance-heavy configuration. Apache Kafka excels in high throughput but has great latency, RocketMQ and RabbitMQ excel in low latency, while Pulsar shows the most balanced performance. The optimal equilibrium point between security and performance is generally achieved at Level 2 (TLS + Authentication), which offers essential protection with an acceptable performance degradation. For further development, it is recommended to conduct more in-depth research related to broker configuration tuning, evaluation of resource usage (CPU, memory, disk), more dynamic workload simulation, as well as exploration of advanced security mechanisms such as SASL, OAuth 2.0, and end-to-end encryption.

DAFTAR PUSTAKA

- Andstrom, V. (2024). *A Comparative Analysis of Apache Kafka and Apache Pulsar*.
- Azzedin, F., & Alhazmi, T. (2023). Secure Data Distribution Architecture in IoT Using MQTT. *Applied Sciences (Switzerland)*, 13(4), 1–13. <https://doi.org/10.3390/app13042515>
- Buccafurri, F., De Angelis, V., Lazzaro, S., & Vangala, A. (2025). MQTT-E: E2E encryption in MQTT via proxy re-encryption avoiding broker overloading. *Ad Hoc Networks*, 176. <https://doi.org/10.1016/j.adhoc.2025.103878>
- Dobbelaere, P., & Esmaili, K. S. (2017). Industry paper: Kafka versus RabbitMQ: A comparative study of two industry reference publish/subscribe implementations. *DEBS 2017 - Proceedings of the 11th ACM International Conference on Distributed Event-Based Systems*, 227–238. <https://doi.org/10.1145/3093742.3093908>
- Febriyani, F., Pramukantoro, E. S., & Bakhtiar, F. A. (2019). *Perbandingan Kinerja Redis, Mosquitto, dan MongoDB sebagai Message Broker pada IoT Middleware*. <http://j-ptiik.ub.ac.id>
- Gentile, A. F., Macrì, D., Carnì, D. L., Greco, E., & Lamonaca, F. (2024). A Performance Analysis of Security Protocols for Distributed Measurement Systems Based on Internet

- of Things with Constrained Hardware and Open Source Infrastructures. *Sensors*, 24(9). <https://doi.org/10.3390/s24092781>
- Goel, R. (2024). *Evaluating Message Brokers: Performance, Scalability, and Suitability for Distributed Applications*. <https://doi.org/10.5923/j.ajca.20241105.02>
- Henning, S., & Hasselbring, W. (2024). Benchmarking scalability of stream processing frameworks deployed as microservices in the cloud. *Journal of Systems and Software*, 208. <https://doi.org/10.1016/j.jss.2023.111879>
- Hernández Ramos, S., Villalba, M. T., & Lacuesta, R. (2018). MQTT Security: A Novel Fuzzing Approach. *Wireless Communications and Mobile Computing*. <https://doi.org/10.1155/2018/8261746>
- Kurdi, H., & Thayananthan, V. (2022). A Multi-Tier MQTT Architecture with Multiple Brokers Based on Fog Computing for Securing Industrial IoT. *Applied Sciences (Switzerland)*, 12(14). <https://doi.org/10.3390/app12147173>
- Laghari, A. A., Li, H., Khan, A. A., Shoulin, Y., Karim, S., & Khani, M. A. K. (2024). Internet of Things (IoT) applications security trends and challenges. *Discover Internet of Things*, 4(1). <https://doi.org/10.1007/s43926-024-00090-5>
- Lin, S., Cui, L., & Ke, N. (2024). End-to-End Encrypted Message Distribution System for the Internet of Things Based on Conditional Proxy Re-Encryption. *Sensors*, 24(2), 1–17. <https://doi.org/10.3390/s24020438>
- Maharjan, R., Chy, M. S. H., Arju, M. A., & Cerny, T. (2023). Benchmarking Message Queues. *Telecom*, 4(2), 298–312. <https://doi.org/10.3390/telecom4020018>
- Mishra, B. (2018). Performance evaluation of MQTT broker servers. In *Lecture Notes in Computer Science* (Vol. 10963, pp. 599–609). https://doi.org/10.1007/978-3-319-95171-3_47
- Mukhlis, I. R., & Pipin, S. J. (2024). *BIG DATA (Mengenal Big Data & Implementasinya di Berbagai Bidang)* (Sepriano & Y. Agusdi, Eds.; Vol. 1). PT Sonpedia Publishing Indonesia. <https://www.researchgate.net/publication/378313489>
- Owobu, W. O., Abieba, O. A., Gbenle, P., Onoja, J. P., Daraojimba, A. I., Adepoju, A. H., & Chibunna, U. B. (2024). Review of Enterprise Communication Security Architectures for Improving Confidentiality Integrity and Availability in Digital Workflows. *Deleted Journal*. <https://doi.org/10.62225/2583049X.2024.4.6.4024>
- Sandholm, T., Mukherjee, S., & Huberman, B. A. (2021). *SAFE: Secure Aggregation with Failover and Encryption*. <https://doi.org/10.1145/3716630>
- Santoso, J. T. (2020). *Analisis Big Data*. Yayasan Prima Agus Teknik.
- Seoane, V., Garcia-Rubio, C., Almenares, F., & Campo, C. (2021). Performance evaluation of CoAP and MQTT with security support for IoT environments. *Computer Networks*, 197. <https://doi.org/10.1016/j.comnet.2021.108338>
- Spielvogel, K., Pöhls, H., & Posegga, J. (2021). *TLS Beyond the Broker: Enforcing Fine-grained Security and Trust in Publish/Subscribe Environments for IoT*. <http://arxiv.org/abs/2109.01169>
- Thallapally, N. (2025). *Enhancing Distributed Systems with Message Queues: Architecture, Benefits, and Best Practices*. 63–69.