

# Leveraging Event Streaming to Decouple Legacy ERP Systems from Modern Commerce Platforms

Viplove Goswami

goswamiviplove@gmail.com

## Abstract:

The ongoing demand for a single enterprise resource planning (ERP) system creates an obstacle to the flexibility needed for digital commerce. While reliable for internal admin functions, these legacy systems stable, but with rigid architectures can't satisfy the sub-second latencies and extreme degrees of scalability of today's e-commerce ecosystems. The paper discusses how event streaming can be strategically used to decouple these legacy cores from agile cloud-native commerce platforms. Organizations can achieve the benefits of temporal and spatial decoupling by switching from synchronous request-response patterns to asynchronous event-driven architectures (EDA) so services can grow at their own pace and scale. The work investigates patterns like log-based Change Data Capture (CDC) and the Transactional Outbox pattern to ensure data integrity and atomicity without requiring invasive changes to the legacy application code. The Strangler Fig pattern is additionally analyzed in this paper as a less operationally risky phased migration strategy when compared to a Big Bang replacement. Evidence from large-scale implementation showcases improved system resilience, improved developer productivity, and improved deployment frequency as result set. The study suggests that event streaming is a vital component of digital transformation allowing to retain preexisting value while unlocking new innovation potential.

**Keywords:** Event Streaming, Enterprise Resource Planning, Decoupling, Microservices, Change Data Capture, Strangler Fig Pattern, E-commerce Scalability.

## INTRODUCTION

The architecture of modern enterprise computing is marked by an underlying tension between the need for stability over time and the need for real-time response. ERP systems have been the single source of truth for organizational data, including finance, procurement, human resources, and supply chain for decades. This controversy is rooted in the ERP systems. In the 1960s and 1970s, as centralized mainframe computers computerized their core banking and administration functions, banks introduced their first internal systems. Although monolithic designs have inherent advantages (the integration of operations and common security controls), they are a massive impediment to the digitalisation of international trade because of their rigidity.

The latest trends in e-commerce involve much higher volumes of transactions across the globe, coupled with an accelerated pace of customer experiences with personalized expectations. The requirements can no longer be satisfied by traditional ERP systems, tightly coupled, synchronous architectures. When it comes to e-commerce sites, traffic peaks at certain seasons or during sale events. These monoliths are the first point of failure and application downtimes. Research into software architecture trends found that 87% of enterprises have started to move away from monolithic software architectures to distributed designs to ensure mission-critical availability.

Modernization complications arise because legacy systems hold decades of business data and regulatory history, making it impossible to replace them overnight. It has been seen that “Big Bang” migrations offer a direct replacement but can be risky, resulting in spiralling costs, corrupted data, and business interference. Consequently, architects are moving toward decoupling strategies that allow new-age commerce apps to interact with legacy systems via event streams.

In this paper we will explore event streaming as the modern enterprise’s “central nervous system” that moves things away from synchronous, brittle integrations and instead towards resilient event-driven ecosystems. We investigate how technologies such as Apache Kafka make it possible to extract data from legacy silos and propagate it to cloud-native services in near real time. By exploring techniques such as Change Data Capture (CDC) and the Strangler Fig pattern, the research outlines a framework that organizations can employ to modernize their technical infrastructure and retain operational continuity.

### **THE STRUCTURAL LEGACY: ERP SYSTEMS IN A DIGITAL WORLD**

Enterprise technology has evolved to a point where organizations have now been left with monolithic applications. And, they present considerable operational and financial expenses to firms. The design of these systems occurred at a time of low software complexity and little demand for distributed systems. In these situations, the “all-in-one” package that the monolith provided gave you a simple deployment model. It also provided transaction management in a simpler way. However, the downsides of these architectures have become sharp as business requirements changed.

Traditionally, ERP systems were created as independent applications for various departments which were later integrated via human or faulty interfaces. Increase, and quite often, ad-hoc, has led to a rigid IT landscape demanding high maintenance and being overly rigid. Updating these collection systems is not simply a technical exercise. It is fundamental to being competitive because failure to periodically update them to changing market conditions can mean lost competitiveness.

Monolithic ERP systems have highly coupled modules which is one of their most serious limitations. As business logic, data management, and user interfaces are often contained in one code, even a small change made in one of the components may cause ripple effects across the entire system. The redevelopment of the entire application, as a result of this interdependence, leads to downtime and a slowdown in the development cycle.

It costs a lot of money to maintain legacy ERP. Most organizations spend a lot of their IT budget on maintenance and operating costs than spending it on innovations like cloud computing. As original developers retire and technology frameworks age, the price of experts continues to rise, in turn increasing costs again. Additionally, legacy systems face performance bottlenecks that involve poor hardware and inefficient data processing that do not scale easily to changing demand.

When it comes to e-commerce, their architectural deficiencies translate into business risk. Legacy ERP systems lack the capabilities to talk to modern-day commerce systems, causing the wrong inventory availability messages, order processing delays, and incorrect financials. Real time data synchronization is mostly not achieved leading to a bad customer experience during peak shopping periods. According to research, monolithic systems experience an average downtime of 4.2 hours per month during peak loads of over 10,000 concurrent users and have an average cost per failure of \$5,600 per minute.

### **DECONSTRUCTING THE MONOLITH: TECHNICAL AND ECONOMIC DRIVERS**

The way physical exams are coded is adversely affected by several factors, including changes to CC and MCC coding. As transactional volumes increase and consumer expectations for speed and reliability grow,

traditional infrastructures are becoming inadequate to keep up. With the need for global scale, cost efficiency and regulatory compliance, enterprises are increasingly adopting hybrid cloud models that combine the security of private cloud services with the flexibility of public cloud services.

Technical debt is a main economic force for renewal. The examination of enterprise architectures has revealed that organizations with large monolithic applications spend as much as 75% of their development resources on technical debt and maintenance. This leaves them with just a quarter of their capacity for the innovation needed to drive new business value. Shifting to microservices and event-driven architectures can help reduce this debt because services can then be developed, updated, and scaled independently.

Decoupling is another essential condition for higher system resilience. In a monolithic architecture, if there is an error in any one module, the complete system fails. By breaking applications into smaller, independently deployable services introduced, organizations allow for fault isolation which means one area breaking will not impact the whole platform. It is a critical requirement for mission-critical applications in the banking and healthcare sector, where high availability is key.

Another important technical driver is the move to decentralized data management. Monolithic systems usually use a single shared database, which causes data contention and integrity problems as the application gets bigger over time. Decoupling allows each service to own its data domain, thereby minimizing interference and enabling the choice of information storage solutions that best suit the requirements of a given service. The shift presents challenges in ensuring data consistency in a distributed system, requiring eventual consistency models and complex integration patterns.

From the organization's point of view, moving away from monolithic architectures helps embrace DevOps practices and continuous integration/continuous delivery (CI/CD) pipelines. When teams own specific service boundaries, both the quality of the software and its delivery speed improves. When teams are aligned to work towards the same goals, empirical research shows it leads to more deployments as well as faster recovery from incidents than traditional teams.

### **EVENT-DRIVEN ARCHITECTURE: THE DECOUPLING PARADIGM**

Event-driven architecture refers to a software architecture pattern that centers around the production, detection, consumption of, and reaction to events. Events enable components of various systems to communicate with one another without waiting for server responses. This is contrary to the common synchronous models of communication like Remote Procedure Calls (RPC) or REST APIs. DevOps encourages a more seamless and responsive communication between microservices as the complexity and scalability of applications increase.

An event here refers to a significant occurrence or a change in state. It may be like Order Placed or Payment Processed and so on. In a system that uses events to trigger actions, the producers create the events and publish them onto an event stream while the consumers subscribe to the relevant streams and react to the events at their own pace. This mechanism breaks the link between a producer and a consumer to operate independently and asynchronously.

Since Apache Kafka was introduced seven years ago, it has become the de facto standard in the enterprise for implementing event streaming backbones. Through the publish-subscribe model of Kafka, it has become easier for event-driven microservices to communicate and scale horizontally to handle bursts in data flow. Differing from just a queue, Kafka is a distributed, fault-tolerant commit log, providing a high throughput, low-latency delivery, and durable storage of events. This design enables consumers to process and reprocess event streams as they require, helping to decouple legacy systems resiliently.

System bottlenecks are significantly reduced if EDA is implemented. Research shows replacing synchronous request-response patterns with asynchronous event streams can cut bottlenecks by up to 78%. The ability to separate complex workflows into simple events allows scaling them depending on the demand which improves efficiency. In an e-commerce platform, for example, events can be used to start the payment verification order processing in parallel with allocating inventory and coordinating shipping. However, the adoption of EDA also introduces new complexities, particularly regarding eventual consistency and distributed transaction management. In an asynchronous system, there is no immediate guarantee that all parts of the system have seen the latest state. While this is often acceptable for business logic, it requires careful planning to ensure long-term data integrity. Organizations must implement robust conflict resolution strategies and reconciliation mechanisms to handle discrepancies that may arise due to the asynchronous nature of event updates.

### **MECHANISM OF DATA LIBERATION: CHANGE DATA CAPTURE**

Most of the time, the first step in decoupling a legacy ERP system is extracting real-time data without disrupting the existing operations of the monolith. Change Data Capture (CDC) has become a crucial technology to identify and capture changes in a database so that they can be sent to other systems in a timely manner. CDC alters a static repository into streams that can react to events, allowing enterprises to enhance the value of their legacy data through new technologies.

There are different methods of implementing CDC. For example, a pull-based method utilizes row versioning, timestamps, and one or more database triggers. Pull-based approaches may add extra overhead to the source system, as the target system must periodically check for changes. Log-based CDC is often the preferred method due to its efficiency and non-intrusiveness. This approach extracts updates from the transaction logs an ordered and sequential record of operations rather than the tables themselves.

The log-based CDC process consists of various technical stages. A database transaction is created like a new customer being created or order status changing on his database and written into the transaction log. A CDC engine continuously monitors these log files and processes the raw log data into a structured format, sending it to target systems through messaging queues or streaming platforms like Apache Kafka. This method provides dependable low latency of roughly 10 ms while having a negligible effect on the source database's performance.

The CDC ensures consistency of events sent on the modern e-commerce platforms with changes made to the legacy system.

When organizations treat all changes to databases as events, they can hydrate modern databases, leverage operational use cases, and power real-time analytics without the risk and cost of ripping out mission-critical legacy databases. For example, in an e-commerce migration, CDC can be used to synchronize inventory data between a legacy warehouse system and a new cloud-native inventory service. This is so that applications facing customers have accurate information.

Implementing CDC also facilitates the use of advanced architectural patterns like the Transactional Outbox. In this scenario, the outbox pattern ensures that the capture of a database transaction and the publication of the corresponding event are treated as a single atomic operation. This prevents common integration problems where a database update succeeds but the notification to downstream systems fails, leading to data inconsistencies.

### **RELIABLE DISTRIBUTED STATE: THE TRANSACTIONAL OUTBOX PATTERN**

Ensuring reliability and consistency is one of the core challenges in event-driven systems. In distributed environments, it is difficult to guarantee that an update to a database and the publication of an event happen

atomically—a problem known as the "dual-write" issue. If a service updates its database but crashes before it can publish the event, other services will not be aware of the change, leading to a loss of data integrity. The Transactional Outbox pattern provides an elegant solution to this problem by recording outbound events as part of the same database transaction that modifies the business data.

In this pattern, the application is designed to write events into a dedicated "outbox" table within the same relational database used for its primary business data. Because both the business data update and the insertion into the outbox table occur within the same ACID transaction, they are guaranteed to either both succeed or both fail. A separate process, often referred to as a message relay or publisher, then reliably reads the records from the outbox table and publishes them to a message broker like Apache Kafka.

The relay mechanism can be implemented using either a polling approach or log-based CDC. A polling-based relay periodically checks the outbox table for unprocessed rows, which is simple but can introduce latency and additional database load. A CDC-based relay, such as Debezium, monitors the database's transaction log for changes to the outbox table, enabling sub-second latency from the time a write occurs to when the event is published. Once the event is successfully published to the message broker, the relay marks the outbox record as processed or deletes it to prevent re-publication.

The Transactional Outbox pattern typically ensures at-least-once publication, meaning that in certain failure scenarios, an event might be published more than once. To maintain system-wide consistency, the consumers of these events must be idempotent, meaning they can handle duplicate messages without unintended side effects. This is often achieved by including unique event IDs in each message and having consumers track the IDs of the events they have already processed.

This pattern is particularly foundational for systems where transactional integrity is non-negotiable, such as financial platforms and inventory management systems. For example, when a user initiates a payment, the core system updates the user's balance and records an event in the outbox table atomically. The relay then publishes the event to downstream services for fraud detection and accounting, ensuring that the system's overall state remains synchronized even in the event of partial failures.

### **STRATEGIC MODERNIZATION: THE STRANGLER FIG PATTERN**

In the case of large legacy systems, mission-critical legacy systems often carry a prohibitively high "Big Bang" replacement risk. The Strangler Fig pattern offers a less risky and more measured way to modernize systems through the gradual replacement of specific functionality with new functionality via applications and services. The pattern is named after a plant that grows around a host tree, eventually replacing it. Similarly, a legacy system will continue to function as a new system gets transmitted, minimizing business disruption.

To technically implement the Strangler Fig pattern, we introduce a facade or routing layer (like an API gateway/proxy) between the client applications and the legacy system. At first, all requests are routed to legacy application through facade. As microservices are created to replace functionality, and part of the facade is reconfigured to forward the call to the new service. The teams thus solve complexity in increments. The process of incremental approach allows the system to be functional and stable.

Choosing which features to "strangle" first is a critical choice. Organizations frequently begin with modules that are impactful, have clear boundaries, and are simpler to develop and deploy. For instance, authentication, reporting, or a certain feature module. The project's momentum is built through these early wins that would allow team members to learn and adapt their strategy based on actual feedback. As

additional features are transferred to the new system, the legacy application's functionality reduces to nothing until it can be retired.

The Strangler Fig pattern provides essential advantages in risk management and resilience. (13 words) Investigating the two approaches (Strangler Fig v/s big bang) one will observe that strangler fig consistently has lower operational risk and higher resilience in the long run. Organizations that use the Strangler Fig pattern see 45% fewer production incidents while migrating and achieve a successful migration rate of 70%. Moreover, Strangler Fig migration average cost (\$850K) is less than the average cost of Big Bang migration (\$1.32M). In an incremental model, costs are amortized over a longer period. However, the Strangler Fig pattern is not suitable for all environments. It requires the ability to intercept requests to the back-end system, and it may not be appropriate for highly entangled monoliths where it is difficult to isolate individual functions. Additionally, there is a risk of ending up in a "perpetual hybrid" mode if the migration is not clearly planned and executed, potentially leading to increased complexity and inconsistent user experiences.

### **IMPACT ANALYSIS: OPERATIONAL AGILITY AND DEVELOPER EXPERIENCE**

Switching from monolithic ERP systems to event-driven microservices is significantly improving organizational agility and the efficiency of productive development teams. Decomposing complex applications into smaller, independently deployable units can help organizations significantly reduce the time to market for new features. Research shows that organizations leveraging microservices enjoy up to 75% faster time to market, since teams release their updates without needing to coordinate on the whole code.

Microservices architecture makes developers more productive because they have greater autonomy. Decoupling offers a "One microservice per Team" policy; i.e. every team specializes in a microservice and handles it throughout its lifecycle. This independence is accompanied by the theatre to choose the best technology stack for each service. This polyglot programming has been seen to boost developer productivity by 28% in early adopting organizations. You can reduce the time spent on coordinating on different parts of the system by up to 65%.

The ability of a system to recover from incidents has improved considerably. "Microservices architectures follow 'design for failure' guidance, implementing circuit breakers and bulkheads to limit the impact of cascading failures." It's 24x faster to recover from incidents if the organization is aligned with service boundaries, while a 'you build it, you run it' mindset reduces the mean time to detect (MTTD) of issues by 63%. Only 89% are reduced in cascading failures with only the implementation of circuit breakers in the production environment.

Event-driven microservices scale well in busy e-commerce situations, providing key benefits. When a service-oriented application is developed using an event-based approach, it helps in communicating between processes, reducing protocol overhead, and fast message delivery.

Research has demonstrated that properly decoupled systems can sustain low response times despite an enormous increase in transaction volume, allowing institutions to conduct business at tremendous scale, effectively processing billions of transactions each day.

Despite these benefits, the move to microservices also introduces new challenges related to operational complexity and data consistency. Managing multiple services, databases, and APIs requires robust service orchestration and comprehensive observability tooling. Research shows that while a large majority of microservices implementations initially struggle with data consistency, this rate drops significantly once proper monitoring and distributed tracing tools are implemented.

## CASE STUDIES AND EMPIRICAL EVIDENCE

The theoretical benefits of event streaming and decoupling are supported by a range of empirical studies and real-world implementations. In the financial services sector, for instance, major institutions have utilized event-driven architectures to modernize their transaction processing. Citi successfully scaled its commercial cards API platform using event-driven microservices, enabling the processing of millions of monthly requests while maintaining operational excellence. The implementation of these systems allowed for independent scaling during demand spikes and improved fraud detection capabilities by 75%.

In the e-commerce domain, Target's implementation of event-driven systems across its stores and digital platform serves as a notable example of real-time inventory management. By using events to synchronize data between local stores and the central system, the organization was able to provide immediate stock checks and personalized customer assistance, significantly accelerating fulfillment times. During peak demand periods, e-commerce platforms with mature event-driven implementations have been shown to consistently maintain 99.8% availability while handling up to 8.5x their normal transaction volumes.

The manufacturing sector has also seen successful digital factory initiatives that link production orders with factory systems and ERPs through event-driven architecture. The use of IoT sensor data processed through event streams allows for predictive maintenance, where potential equipment failures are identified before they occur, preventing costly downtime. These initiatives demonstrate how event streaming can be used to integrate legacy operational technologies with modern cloud platforms.

A comparative investigation of 471 enterprise migration projects provides strong empirical evidence for the efficacy of the Strangler Fig pattern. The study found that the Strangler Fig approach, while having a longer project duration on average (26 months compared to 18 months for Big Bang), demonstrated significantly lower operating risk and much higher resilience. Big Bang migrations were characterized by a high failure rate (12.10%) and "deplorable survivability," often leading to total project failure. The findings highlight the value of an incremental "pay-as-you-go" model that delivers value throughout the migration process.

Furthermore, experimental data from e-commerce system designs shows the performance advantages of asynchronous communication. In a comparative study, an event-driven system achieved up to 495 requests per second while maintaining low response times, whereas a synchronous RPC-based system peaked at only 5.1 requests per second with significant increases in latency. These results provide a compelling argument for the adoption of decoupled, event-driven models in large-scale, transaction-intensive environments.

## CONCLUSION

The process of using event streaming to decouple legacy ERP systems is a very effective enterprise modernisation strategy. According to this study, monolithic architectures have previously performed excellently. However, now they cannot support fast and scalable digital commerce as they used to. According to experts, the growing constraints posed by outdated legacy systems can deprive an organization of value. The specific configuration of NACs across a wide range of devices causes them to be inflexible, cost-centric and bottlenecks in performance.

The backbone of event streaming with Apache Kafka is the bridge that allows reliability of the legacy and agile of today to exist together. Moving towards event-driven architectures can offer services the temporal and spatial decoupling required for independent evolution by organizations. Methods like Change Data Capture (CDC) and the Outbox pattern make it easy to extract and propagate data nearly in real-time with no transaction loss or dirty reads, while not risking a high-impact invasive change to the legacy core.

The Strangler Fig pattern is a winning strategy for migration that allows lower risk phased replacement rather than “Big Bang” type replacements. Studies show that this incremental approach diminishes the likelihood of project failure and consistently produces business value throughout the transformation process. The shift brings multiple advantages including a faster time to market, higher deployment frequency, and improved system resilience.

Nonetheless, there are difficulties in the way of the decoupled enterprise. The transition to eventual consistency models and handling a distributed system introduces several complications that need sound architectural thinking and observability tools. Organizational alignment of service boundaries (based on Conway’s Law) also helps maximise the benefits of microservices and event-driven design.

In conclusion, leveraging event streaming to decouple legacy ERP systems is a strategic imperative for organizations seeking to navigate the complexities of the modern digital landscape. By adopting these architectural patterns and migration strategies, enterprises can unlock the full potential of their data, empower their development teams, and build the resilient, scalable systems necessary for the future of digital commerce.

## REFERENCES:

1. SAP Business Blog. (2024). Monolithic ERP systems in medium-sized businesses: challenges, solutions and risk management. [sap-b1-blog.com](http://sap-b1-blog.com).
2. Smith, K. (2023). Key Challenges in Migrating from Legacy Systems to Modern Platforms. Medium.
3. Emerline. (2024). ERP System Modernization: Solving the Problems That Hold You Back. [emerline.com](http://emerline.com).
4. Singleclitic. (2024). Challenges of Integrating Legacy Systems with Modern Apps. [singleclitic.com](http://singleclitic.com).
5. ResearchGate. (2024). Challenges and Solutions in Legacy System Modernization for Cloud Readiness. [publication/387750421](https://www.researchgate.net/publication/387750421).
6. RST Software. (2024). Event-driven architecture foundations for legacy system transformation. [rst.software](http://rst.software).
7. SCITEPRESS. (2025). Event-driven architecture: simplified migration and microservice communication. [scitepress.org](http://scitepress.org).
8. Solace. (2023). Event-driven architecture patterns: ECST, CQRS, and CDC. [solace.com](http://solace.com).
9. Squer. (2024). Data synchronization in action: leveraging change data capture for system modernization. [squer.io](http://squer.io).
10. IJACT. (2023). Tracing the methods and approaches used by CDC. Volume 1, Issue 2.
11. Confluent. (2024). How change data capture works: patterns, solutions, implementation. [confluent.io](http://confluent.io).
12. Red Hat. (2023). What is change data capture? [redhat.com](http://redhat.com).
13. IJFMR. (2022). Building Resilient Microservices with Apache Kafka. Volume 3.
14. DZone. (2024). Kafka event-driven ecommerce migration case study. [dzone.com](http://dzone.com).
15. AIM Journals. (2024). Apache Kafka as the backbone for MSA: benchmarking and governance. [aimjournals.com](http://aimjournals.com).
16. Platform Engineers. (2024). Loose coupling in microservices: the role of Apache Kafka. [medium.com](http://medium.com).
17. Microsoft Azure Architecture Center. (2025). Strangler Fig pattern for incremental modernization. [learn.microsoft.com](http://learn.microsoft.com).
18. Three North. (2024). The economics of the Strangler Fig pattern. [threenorth.io](http://threenorth.io).
19. Step Software. (2024). Modernizing legacy systems with the Strangler Fig pattern. [stepsoftware.com](http://stepsoftware.com).

20. Future Processing. (2025). How the Strangler Fig Pattern supports legacy system replacement. [future-processing.com](http://future-processing.com).
21. International Journal of Emerging Research in Engineering and Technology. (2022). Modernizing Legacy ERP Systems with AI and Machine Learning. 3(4).
22. IJARCAST. (2024). AI-First Banking conceptual model for legacy ERP integration. Volume 2.
23. ResearchGate. (2024). The Role of Data Governance in Strengthening ERP and MDM Collaboration. [publication/395166584](https://www.researchgate.net/publication/395166584).
24. Rapydo. (2024). SQL databases as the backbone of an event-driven system: CDC and Outbox patterns. [rapydo.io](http://rapydo.io).
25. GeekyAnts. (2024). Transactional Outbox pattern for real-world mobile applications. [geekyants.com](http://geekyants.com).
26. ResearchGate. (2025). Outbox Pattern for Reliable Distributed Systems. [publication/398270614](https://www.researchgate.net/publication/398270614).
27. Carr, J. (2026). The Transactional Outbox Pattern: Reliable Event Publishing. [james-carr.org](http://james-carr.org).
28. Journal of International Crisis and Risk Communication Research. (2026). De-Risking SAP S/4HANA Migrations: A Strangler Fig Pattern. 71–81.
29. SAR Council. (2025). Integrating AI with established architectural patterns like the Strangler Pattern. SJECS-109.
30. IJCSS. (2025). Comparative investigation of migration approaches: Big Bang vs Strangler Fig. Volume 19.
31. OARJET. (2025). Core Transformation Techniques in Microservices Transformation. OARJET-2025-0053.
32. KDI. (2024). Asynchronous vs synchronous models in e-commerce: RabbitMQ vs REST. [jurnal.kdi.or.id](http://jurnal.kdi.or.id).
33. WJAETS. (2025). Event-Driven Architecture in modern distributed systems. WJAETS-2025-0402.
34. ResearchGate. (2023). Event-Driven Architecture to Improve Performance and Scalability. [publication/368431218](https://www.researchgate.net/publication/368431218).
35. Zenodo. (2025). Microservices and event-driven architecture in e-commerce systems. [record/17299559](https://zenodo.org/record/17299559).
36. Fried, U. (2024). The reality of eventual consistency in distributed commerce. [ufried.com](http://ufried.com).
37. All Multidisciplinary Journal. (2025). Benefits of cloud-native and distributed architectures. MGE-2025-2-018.
38. ER Publications. (2024). Hybrid cloud strategies for modernizing global e-commerce. [erpublications.com](http://erpublications.com).
39. Parra, J. (2024). Justifying the use of eventual consistency patterns in modern software architecture. Medium.
40. IJSAT. (2025). Interface design strategies for distributed systems: modularity and loose coupling. [ijsat.org](http://ijsat.org).
41. DiVA Portal. (2025). Transitioning a legacy monolithic system to microservices: a case study. [diva2:1986594](https://diva2.org/1986594).
42. EA Journals. (2025). Microservices Transformation: principles, methodologies, and outcomes. [ejournals.org](http://ejournals.org).
43. The Journal of Scientific and Engineering Research. (2022). Architecting Resilient Cloud-Native APIs: Fault Recovery. 9(3).
44. WJAETS. (2025). The Evolution of Banking Architecture from Monoliths to Microservices. WJAETS-2025-1144.
45. Carr, J. (2026). Implementation Guides for the Transactional Outbox. [james-carr.org](http://james-carr.org).
46. ResearchGate. (2024). Toward Organizational Decoupling in Microservices Through Key Developer Allocation. [publication/388494976](https://www.researchgate.net/publication/388494976).

47. SSRN. (2024). The impact of microservices decoupling on developer productivity and system integration. 5010352.
48. RSIS International. (2024). Microservices Architecture: a software engineering perspective. [rsisinternational.org](http://rsisinternational.org).
49. Medium. (2024). The Rise of Microservices Architecture: Transforming Software Development. [tech-stack-insights](https://tech-stack-insights.com).
50. WJARR. (2025). Event-driven data synchronization between Workday and downstream applications. [record/1116](https://record/1116).
51. WJAETS. (2025). AI-powered data migration from legacy systems to cloud platforms. WJAETS-2025-0183.
52. IJACT. (2023). Technical Process and Benefits of Log-based CDC. Volume 1, Issue 2.
53. IJCSS. (2025). Quantitative Data and Comparative Analysis of Migration Strategies. Volume 19.
54. ResearchGate. (2025). Theoretical discussion on eventual consistency and transactional integrity. [publication/398270614](https://publication/398270614).
55. EA Journals. (2025). Key findings on developer productivity and organizational agility. [eajournals.org](http://eajournals.org).