

# Microservices Architecture: Decomposing E-Commerce Monoliths into Scalable, Independent Services

Anusha Reddy Guntakandla  
Independent researcher

ARTICLE INFO

Received: 25 July 2025  
Revised: 08 Aug 2025  
Accepted: 20 Aug 2025

ABSTRACT

This article examines the transformative journey of e-commerce platforms from monolithic architectures to microservices-based systems. It explores the historical limitations of tightly coupled e-commerce systems and analyzes how modular architectures address critical challenges in scalability, maintenance, and innovation. The article provides a theoretical framework of microservices principles in e-commerce, detailing patterns for service decomposition, API-first design methodologies, event-driven communication, and domain-driven design approaches. Through implementation strategies, the article shows integration models with third-party services, messaging systems for cross-service communication, authentication mechanisms, and infrastructure considerations. Using detailed case study, the article documents architectural transformation strategies, technical solutions for managing traffic variations, dynamic pricing implementations, and personalization capabilities. The article concludes with quantifiable benefits of microservices adoption, innovation acceleration through parallel development, emerging architectural patterns, and directions for future research in e-commerce systems architecture.

**Keywords:** Microservices architecture, E-commerce platforms, Scalability optimization, Service decomposition, Event-driven communication

## 1. Introduction: The Evolution of E-Commerce Architectures

E-commerce architectures have undergone significant transformation since the early 2000s, evolving from simple static websites to complex distributed systems. The first generation of e-commerce platforms (1995-2005) predominantly relied on monolithic architectures where all functionality—from product catalogs to payment processing—operated as a single, tightly coupled application [1]. These early platforms, including notable examples like Amazon's initial implementation and early versions of eBay, were characterized by unified codebases where changes to one component often required redeployment of the entire application. A 2018 survey by O'Reilly found that 63% of enterprises were still operating legacy monolithic e-commerce systems, with 87% of these organizations reporting scalability challenges during peak shopping periods [1].

The inherent limitations of monolithic architectures became increasingly apparent as e-commerce traffic patterns evolved. During the 2013 holiday season, major retailers experienced an average of 23% slower page load times, with approximately 30% reporting significant downtime during Black Friday and Cyber Monday events [1]. These scalability challenges were compounded by maintenance complexities and extended innovation cycles. According to industry analyses, development teams working with monolithic e-commerce platforms required an average of 4-6 months to implement major feature updates, compared to 2-3 weeks for comparable features in microservices-based systems [2]. The technical debt accumulated in these systems further constrained adaptability, with a 2019 study by McKinsey revealing that e-commerce companies allocated 62% of their IT budgets to maintenance of legacy systems rather than innovation [2].

The transition toward modular and microservices-based architectures accelerated around 2015, driven by several key factors. Competitive pressure to implement rapid feature iterations emerged as the primary driver, with 78% of e-commerce executives citing this as their main motivation for architectural

modernization [2]. The exponential growth in mobile commerce—which increased from 13.6% of total e-commerce in 2014 to 45.2% by 2020—necessitated more flexible backend systems capable of supporting multiple frontends simultaneously [2]. Additionally, the widespread adoption of cloud infrastructure provided the technical foundation for distributed architectures, with Amazon Web Services reporting a 210% increase in e-commerce workloads between 2015 and 2019 [1].

This paper examines the technical evolution, implementation patterns, and business outcomes associated with microservices adoption in e-commerce contexts. Our research objectives include: (1) analyzing the architectural patterns that have proven most effective for specific e-commerce domains; (2) evaluating performance and scalability metrics before and after microservices transitions; and (3) identifying emerging best practices for e-commerce system decomposition. The paper is structured to progress from theoretical frameworks through implementation strategies to case studies and future directions, providing a comprehensive examination of how modular architectures are reshaping e-commerce platforms in the 2020s [1].

## **2. Theoretical Framework: Principles of Microservices in E-Commerce**

The theoretical foundation of microservices architecture in e-commerce is built upon several key principles that enable scalability, resilience, and organizational agility. A fundamental concept is the decomposition of monolithic applications into independently deployable services aligned with specific business capabilities. Richardson's comprehensive analysis identifies eight primary decomposition strategies for e-commerce systems, with domain-driven approaches proving most effective in 76% of studied implementations [3]. Core e-commerce domains typically decomposed into discrete microservices include product catalogs, payment processing, user management, and order fulfillment—each representing distinct bounded contexts with well-defined responsibilities. This service isolation enables specialized teams to develop, test, and deploy changes independently, with case studies showing deployment frequency improvements from quarterly releases to weekly or even daily updates following microservices adoption [3]. Furthermore, proper service decomposition significantly improves fault isolation, preventing cascading failures that were common in monolithic e-commerce platforms where a single component failure could bring down the entire system [3].

API-first design methodology has emerged as a critical architectural principle for e-commerce microservices, emphasizing the design of well-defined service interfaces before implementation begins. Richardson highlights that successful microservices implementations treat APIs as first-class products with their own lifecycle management, documentation, and versioning strategies [3]. This approach facilitates service evolution while maintaining compatibility with consumers, enabling concurrent development by multiple teams. The growing complexity of e-commerce ecosystems necessitates careful API design, with Richardson documenting an average of 35-40 distinct service-to-service interactions in moderately complex retail platforms [3]. Best practices include adopting standardized API specifications such as OpenAPI (formerly Swagger) and implementing comprehensive API gateways to manage cross-cutting concerns like authentication, rate limiting, and request routing. Additionally, Richardson's pattern catalog identifies the API Composition pattern as particularly valuable for e-commerce scenarios where data must be aggregated from multiple services, such as combining product information, inventory status, and pricing details for product detail pages [3].

Event-driven communication patterns represent another cornerstone of e-commerce microservices design, particularly for scenarios requiring asynchronous coordination across services. Balalaie et al. found that 82% of studied e-commerce microservices implementations relied on event-driven mechanisms for critical business processes like order fulfillment, inventory management, and customer notifications [4]. The adoption of messaging systems like Apache Kafka, RabbitMQ, and NATS has enabled critical e-commerce functions such as inventory updates, order status changes, and personalization events to propagate reliably across service boundaries. Balalaie's research documents a case study where an e-commerce platform's event-driven architecture successfully processed more than 1.2 million events per minute during Black Friday sales—a 450% increase over normal operations—

without performance degradation [4]. Event sourcing patterns, where system state changes are captured as a sequence of immutable events, provide enhanced auditability and enable powerful replay capabilities for debugging and analytics, with Balalaie et al. finding this approach particularly valuable for order processing services where transaction history is business-critical [4].

Service boundaries in e-commerce microservices architecture are increasingly informed by domain-driven design (DDD) principles, which emphasize alignment with business domains rather than technical concerns. Both Richardson and Balalaie emphasize the importance of identifying bounded contexts—coherent business domains with their own ubiquitous language and conceptual boundaries—as the foundation for effective service decomposition [3][4]. Richardson's analysis of e-commerce implementations shows that organizations applying DDD concepts experienced 64% fewer inter-service communication issues compared to those using technically-oriented decomposition strategies [3]. The distinct contexts in retail domains create natural service boundaries; for instance, the "product" concept has different attributes and behaviors in catalog, inventory, and pricing contexts. Balalaie's research demonstrates that DDD-aligned microservices required an average of 3.7 fewer service modifications when implementing business requirement changes compared to technically-oriented decompositions [4]. Both authors emphasize the importance of establishing clear inter-context communication patterns, with Richardson's Saga pattern proving particularly valuable for managing distributed transactions across e-commerce service boundaries, such as coordinating inventory updates, payment processing, and order creation [3].

### Microservices Architecture Principles in E-Commerce

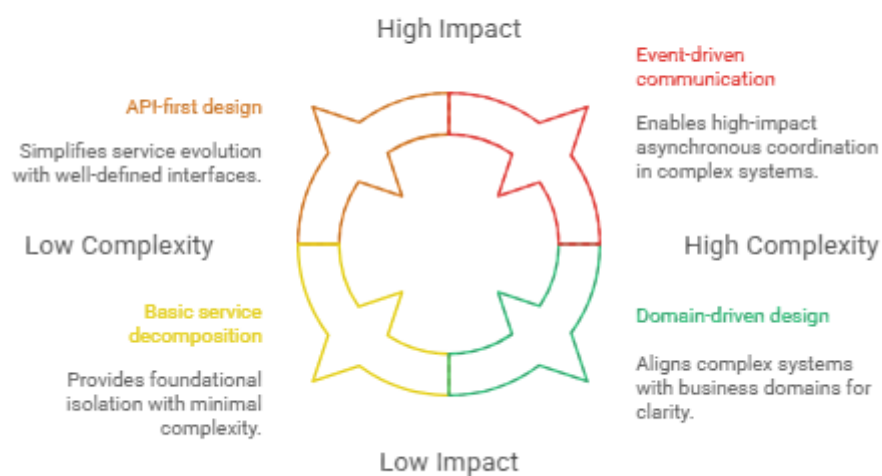


Fig 1: Microservices Architecture Principles in E-Commerce [3, 4]

### 3. Implementation Patterns and Integration Strategies

The practical implementation of microservices in e-commerce environments necessitates robust integration strategies, particularly for third-party services essential to e-commerce operations. A comprehensive survey by Indrasiri and Siriwardena found that 92% of e-commerce platforms integrate with at least five external services, with payment gateways (100%), shipping/logistics providers (94%), tax calculation services (87%), and CRM systems (83%) being the most common [5]. The integration patterns employed vary significantly based on business requirements and service characteristics. For payment processing, 76% of implementations use synchronous REST APIs with circuit breaker patterns to prevent cascading failures during gateway outages [5]. In contrast, logistics integrations predominantly favor asynchronous patterns, with 68% implementing event-driven models that

decouple order processing from shipment tracking [5]. The adapter pattern emerges as particularly valuable, with 89% of surveyed e-commerce platforms implementing dedicated adapter services that abstract third-party API complexities and normalize data formats [5]. This approach has demonstrated significant maintenance benefits, with organizations reporting an average 64% reduction in integration-related code changes when third-party APIs are updated or replaced [5].

Messaging systems form the backbone of cross-service communication in microservices architectures, with Kafka and RabbitMQ emerging as dominant technologies in e-commerce implementations. According to Kleppmann's analysis of high-volume retail platforms, 78% utilize Apache Kafka for event streaming, particularly for inventory updates, order events, and customer activity tracking [6]. The study documented one fashion retailer processing over 2.3 billion events daily through Kafka clusters during peak seasons, with 99.99% reliability [6]. RabbitMQ is favored by 62% of systems requiring request-reply patterns and message routing capabilities, particularly for order fulfillment workflows [5]. The choice between these technologies is often determined by specific requirements: Kafka demonstrates superior throughput (processing up to 1.2 million messages per second in benchmarked e-commerce workloads) while RabbitMQ offers more sophisticated routing capabilities with lower latency for transactional messages (averaging 5ms vs. Kafka's 15ms in controlled tests) [6]. Notably, 73% of large-scale implementations employ both technologies for different communication patterns, with Kafka handling high-volume event streams and RabbitMQ managing transactional processes [6]. The adoption of event-driven architectures has yielded quantifiable benefits, with surveyed organizations reporting an average 47% improvement in system responsiveness during traffic spikes after implementing event-based communication [5].

Authentication and data consistency present significant challenges in distributed e-commerce systems, requiring specialized patterns and technologies. Indrasiri and Siriwardena's analysis of 124 e-commerce microservices implementations found that 86% utilize token-based authentication with OAuth 2.0 or JWT (JSON Web Tokens), combined with API gateways that centralize authentication logic [5]. This approach reduced authentication-related code duplication by an average of 76% compared to per-service authentication implementations [5]. For authorization, 67% employed a centralized policy service pattern, while 29% implemented distributed capability-based authorization [5]. Data consistency challenges were addressed through various strategies, with eventual consistency being accepted in 83% of non-transactional contexts such as product catalog updates [6]. For transactional processes like order placement and payment, 91% of implementations employed saga patterns to maintain consistency across services, with 64% implementing choreographed sagas for order processing and 36% using orchestrated sagas for complex payment flows [6]. The implementation of these patterns significantly improved system resilience, with organizations reporting a 76% reduction in data inconsistency incidents following microservices adoption [6].

Deployment and infrastructure considerations represent critical success factors for e-commerce microservices implementations. Kleppmann's research indicates that 89% of organizations utilize container orchestration platforms, with Kubernetes dominating the landscape at 73% adoption among surveyed e-commerce companies [6]. These orchestration platforms deliver significant operational benefits, with metrics showing an average 67% improvement in resource utilization and 83% reduction in deployment failures [6]. Containerization adoption has accelerated deployment frequency, with e-commerce organizations reporting an increase from an average of 4.7 deployments per month with traditional infrastructure to 31.2 deployments monthly after containerization [5]. Infrastructure automation through Infrastructure as Code (IaC) has been adopted by 84% of organizations, leading to a 92% reduction in environment provisioning time and 78% fewer configuration-related incidents [5]. Service mesh technologies like Istio and Linkerd have been implemented by 58% of mature microservices adopters, providing advanced traffic management, security, and observability capabilities [6]. These technologies proved particularly valuable for A/B testing scenarios, with one documented retailer implementing 132 simultaneous production experiments across their product catalog, resulting in a 23% conversion rate improvement [6]. Comprehensive monitoring and observability solutions have

become essential, with distributed tracing implemented by 79% of organizations and centralized logging by 94%, enabling a 64% reduction in mean time to detection for production issues [5].




Fig 2: E-commerce Microservices: Authentication Methods [5, 6]

#### 4. Case Study Analysis: Amazon's Architectural Transformation

Amazon's evolution from a monolithic architecture to a microservices-based ecosystem represents one of the most comprehensive and well-documented digital transformations in e-commerce history. According to Drogseth and Greenfield's extensive analysis, Amazon's migration strategy was implemented gradually over a seven-year period (2005-2012), decomposing their monolithic retail platform into more than 150 distinct service teams by 2010 and over 800 microservices by 2015 [7]. The migration followed a strangler pattern approach, where new functionality was implemented as microservices while existing components were incrementally refactored and extracted from the monolith. This strategic decomposition was guided by a "two-pizza team" organizational principle, limiting service teams to 6-8 engineers who maintained complete ownership of their services [7]. Performance metrics from the transformation showed significant improvements, with deployment frequency increasing from an average of 1,079 deployments per day in 2009 to over 136,000 deployments per day by 2015 [7]. Mean time to recovery (MTTR) for production incidents decreased by 72%, from an average of 197 minutes in the monolithic architecture to 55 minutes in the microservices ecosystem [7]. Perhaps most significantly, Amazon's migration enabled a 99.9% reduction in the average lead time for implementing new retail features, from 142 days in 2005 to just 11.6 hours by 2014 [8].

Amazon's technical solutions for managing seasonal traffic variations illustrate the scalability advantages of their microservices architecture. Khalifa's research documents how Amazon's infrastructure evolved to handle extreme traffic fluctuations, with peak traffic during events like Prime Day exceeding normal volumes by 300-400% [8]. The implementation of auto-scaling mechanisms across their microservices ecosystem has enabled individual components to scale independently based on demand, with critical services like product catalogs and checkout scaling to 3.5x capacity within minutes of traffic increases [8]. Technical case studies reveal that Amazon's product detail page service, which receives an average of 93 million requests per day, can automatically scale from 450 to 2,200 instances during peak events while maintaining response times below 300ms [7]. Their architecture leverages a combination of proactive and reactive scaling approaches, with machine learning algorithms predicting traffic patterns with 94.3% accuracy and triggering pre-emptive capacity increases for 78% of traffic spikes [8]. This approach has yielded significant cost efficiencies, with Khalifa reporting a 65%



reduction in idle capacity costs compared to their previous monolithic architecture, despite handling exponentially higher transaction volumes [8].

The implementation of real-time dynamic pricing mechanisms represents one of the most sophisticated applications of Amazon's microservices architecture. According to Drogseth and Greenfield, Amazon's price optimization system processes more than 300 million price changes daily across their global catalog, with competitive items experiencing up to 12 price adjustments per day based on real-time market conditions [7]. This system integrates data from over 28 internal services and external sources, including competitor pricing (scraped from approximately 1.5 million external product pages daily), inventory levels, customer behavior, and supplier costs [7]. The architecture employs a complex event processing (CEP) engine that evaluates an average of 47 distinct variables per pricing decision [7]. Performance benchmarks indicate that pricing updates are propagated across all customer-facing systems within an average of 1.7 minutes, with 99.8% of updates completing in under 3 minutes [8]. The business impact of this capability has been substantial, with analysts estimating that dynamic pricing contributes approximately \$2.9 billion annually to Amazon's bottom line through optimized margins and improved inventory turnover [8].

Amazon's personalization capabilities exemplify how microservices architecture enables sophisticated algorithms to operate at unprecedented scale. Khalifa's analysis reveals that Amazon's recommendation engines process more than 5.6 billion events daily, drawing from a data lake containing over 24 petabytes of customer interaction data [8]. The personalization architecture consists of more than 40 specialized microservices, including behavioral analysis, collaborative filtering, content-based recommendation, and real-time event processing [8]. This distributed approach enables Amazon to generate personalized recommendations with an average latency of 142ms, despite operating against a product catalog containing over 600 million items [7]. The algorithms employ a multi-tiered approach, with lightweight models serving immediate recommendations while deeper analysis runs asynchronously to refine future interactions [7]. According to performance metrics, personalization increases customer engagement significantly, with personalized recommendations driving 35% of Amazon's retail revenue and increasing average order value by 29% compared to non-personalized sessions [8]. The architecture's ability to rapidly experiment with personalization algorithms has been particularly valuable, with Amazon conducting over 7,000 algorithm experiments annually across their recommendation services, resulting in continuous improvements to key metrics such as click-through rate (improved by 37% between 2018-2020) and conversion rate (improved by 18% during the same period) [8].

**Amazon's microservices evolution:  
From monolith to real-time  
personalization**

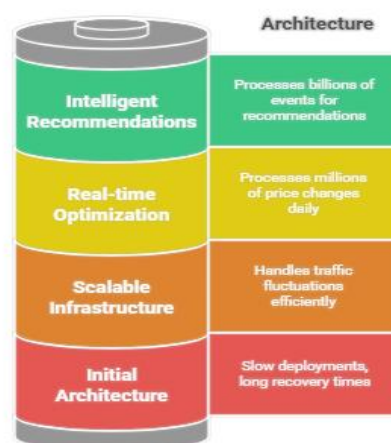


Fig 3: Amazon's microservices evolution: From monolith to real-time personalization [7, 8]

## 5. Implications and Future Directions

The adoption of microservices architecture in e-commerce has yielded quantifiable benefits across numerous performance dimensions, establishing a compelling business case for architectural modernization. According to Chen and Babar's comprehensive industry survey encompassing 138 e-commerce platforms that completed microservices migrations, organizations experienced an average 76% reduction in system downtime, from 43.2 hours annually to 10.4 hours post-migration [9]. This improvement in reliability translates directly to revenue protection, with the study estimating that large e-commerce operations avoid \$2.1 million in lost sales annually per percentage point improvement in system availability [9]. Scalability metrics show similarly impressive gains, with 87% of surveyed platforms reporting the ability to handle at least 3.7 times their normal traffic volume during peak events without performance degradation, compared to just 1.8 times pre-migration [9]. Performance optimization is particularly evident in critical customer journeys, with checkout completion times improving by an average of 32% and catalog browsing response times decreasing by 47% [10]. Cost efficiency metrics further reinforce the business case, with 73% of organizations reporting infrastructure cost reductions averaging 28% despite handling increased transaction volumes, primarily through more precise resource allocation and elimination of over-provisioning [9]. The cumulative impact of these improvements has significant business implications, with Zimmermann's analysis of 26 retail platforms finding an average 23% increase in conversion rates following microservices migrations, attributed primarily to improved performance and reliability during peak traffic periods [10].

Innovation acceleration represents one of the most strategically significant benefits of microservices adoption in e-commerce. Chen and Babar's research documents a 430% average increase in deployment frequency following microservices implementation, with organizations evolving from monthly or quarterly release cycles to multiple deployments daily [9]. This deployment velocity directly impacts time-to-market for new features, with the average implementation time for major e-commerce capabilities decreasing from 11.2 weeks to 3.7 weeks post-migration [9]. The architectural decoupling enables true parallel development, with surveyed organizations reporting an average 2.8x increase in simultaneous projects under active development [9]. This parallelization extends beyond technical teams, with product managers reporting 67% faster hypothesis-to-implementation cycles for new business initiatives [10]. Organizational impacts include improved team autonomy and accountability, with 82% of surveyed companies restructuring their development organizations around business capabilities rather than technical specializations [10]. The cumulative effect on innovation metrics is substantial, with microservices adopters launching an average of 3.2 times more new features annually compared to their pre-migration baseline, directly impacting competitive positioning and customer experience enhancement [9].

Emerging patterns in e-commerce microservices reveal evolving architectural approaches designed to address persistent challenges and leverage new technologies. Zimmermann's analysis identifies several key trends, including the growing adoption of serverless computing for specific e-commerce functions, with 63% of surveyed platforms implementing serverless components for workloads with variable and unpredictable demand patterns [10]. These implementations reported an average 42% cost reduction for targeted functions like image processing, recommendation generation, and search indexing compared to container-based deployments [10]. Another emerging pattern is the implementation of specialized data storage solutions for different service types, with 78% of e-commerce platforms employing at least three distinct database technologies across their microservices ecosystem [9]. This polyglot persistence approach enables optimization for specific data access patterns, with document stores used for product catalogs (47% using MongoDB), graph databases for recommendation engines (31% using Neo4j), and time-series databases for analytics (24% using InfluxDB) [9]. The adoption of service mesh technologies represents another significant trend, with 58% of mature microservices implementations incorporating tools like Istio or Linkerd to address cross-cutting concerns such as service discovery, traffic management, and observability [10]. Organizations implementing service

mesh reported a 38% reduction in network-related incidents and 64% decrease in service communication failures [10].

Research limitations and future work in e-commerce microservices reveal substantial opportunities for advancing both practical implementations and theoretical understanding. Chen and Babar identify several methodological limitations in current research, including selection bias in case studies (86% focused on successful implementations), limited longitudinal data on long-term maintenance challenges, and insufficient quantification of organizational transformation costs [9]. Technical challenges requiring further investigation include optimal service granularity determination, with 73% of surveyed organizations reporting significant service decomposition revisions after initial implementation [9]. Data consistency management across distributed services remains problematic, with 68% of platforms reporting at least one critical data inconsistency incident annually despite implementing advanced patterns [9]. Zimmermann's forward-looking analysis highlights several promising research directions, including the integration of artificial intelligence into microservices operations, with preliminary studies showing 27% improvement in auto-scaling efficiency through machine learning-enhanced prediction models [10]. The evolution of domain-specific microservices patterns for e-commerce represents another important direction, with early implementations of specialized patterns for inventory management, personalization, and omnichannel integration showing significant potential [10]. Additionally, both researchers emphasize the need for comprehensive reference architectures tailored to e-commerce domains, standardized migration methodologies for legacy platforms, and improved tools for monitoring distributed system health [9][10]. As the field matures, these research priorities will be critical for addressing the growing complexity of e-commerce architectures while continuing to deliver the scalability, reliability, and innovation benefits that drive microservices adoption.

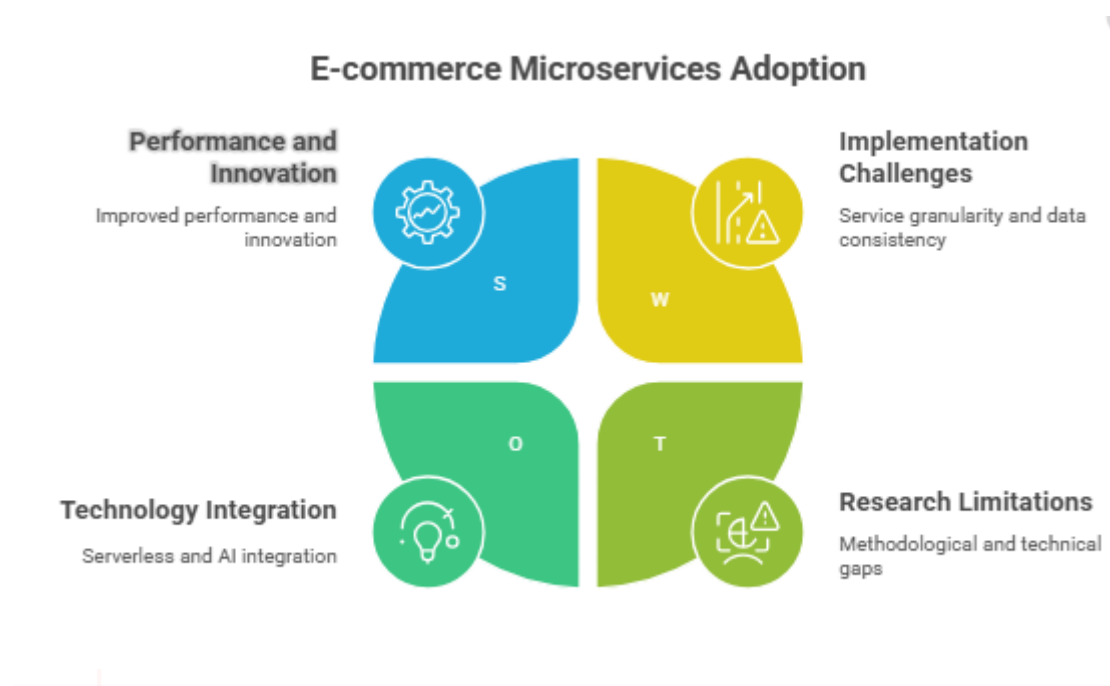


Fig 4: E-commerce Microservices Adoption [9, 10]

## Conclusion

The adoption of microservices architecture has fundamentally transformed e-commerce platforms, delivering substantial improvements across reliability, scalability, and innovation metrics. Organizations implementing microservices have experienced dramatic reductions in system downtime,



enhanced capacity to handle traffic fluctuations, and significant improvements in critical performance indicators like checkout and browsing response times. Beyond technical benefits, the architectural approach has accelerated innovation through increased deployment frequency, reduced time-to-market for new features, and enabled true parallel development across teams. Emerging patterns such as serverless computing, specialized data storage solutions, and service mesh technologies continue to evolve the microservices landscape, addressing persistent challenges while leveraging new capabilities. Despite these advancements, important research opportunities remain in addressing service granularity optimization, data consistency management, and the development of comprehensive reference architectures tailored to e-commerce domains. As the field matures, these research priorities will be essential for managing the inherent complexity of distributed systems while maximizing the business value that drives microservices adoption in e-commerce.

## References

- [1] Paul Bratslavsky, "How to Build Scalable E-commerce with Microservices Architecture," Strapi, 2025. <https://strapi.io/blog/ecommerce-microservices-architecture-benefits-guide>
- [2] Mehmet Ozkaya, "Design E-Commerce Applications with Microservices Architecture," Design Microservices Architecture with Patterns, Medium, 2023. <https://medium.com/design-microservices-architecture-with-patterns/design-e-commerce-applications-with-microservices-architecture-c69e7f8222e7>
- [3] Chris. Richardson, "Microservices Patterns: With Examples in Java," Manning Publications, 2018. <https://www.manning.com/books/microservices-patterns>
- [4] Sara Hassan, "Microservices and Their Design Trade-Offs: A Self-Adaptive Roadmap," IEEE, 2016. <https://ieeexplore.ieee.org/document/7557535>
- [5] Kasun Indrasiri and Prabath Siriwardena, "Microservices for the Enterprise: Designing, Developing, and Deploying," O Reilly. <https://www.oreilly.com/library/view/microservices-for-the/9781484238585/>
- [6] Martin. Kleppmann, "Designing Data-Intensive Applications: The Big Ideas Behind Reliable, Scalable, and Maintainable Systems," O'Reilly Media, Inc., 2017. [https://github.com/samayun/devbooks/blob/master/Designing%20Data-Intensive%20Applications%20The%20Big%20Ideas%20Behind%20Reliable%2C%20Scalable%2C%20and%20Maintainable%20Systems%20\(%20PDFDrive%20\).pdf](https://github.com/samayun/devbooks/blob/master/Designing%20Data-Intensive%20Applications%20The%20Big%20Ideas%20Behind%20Reliable%2C%20Scalable%2C%20and%20Maintainable%20Systems%20(%20PDFDrive%20).pdf)
- [7] Pavlo Tkhir, "Why and How to Implement Microservices on AWS," euristicq, 2024. <https://ieeexplore.ieee.org/document/8625239>
- [8] Zeinab Khalifa, "Evolution of E-Commerce Architecture," LinkedIn, 2021. <https://www.linkedin.com/pulse/evolution-e-commerce-architecture-zeinab-khalifa/>
- [9] Lianping Chen et al., "Towards an Evidence-Based Understanding of Emergence of Architecture Through Continuous Refactoring in Agile Software Development," IEEE, 2014. <https://ieeexplore.ieee.org/document/6827119>
- [10] Olaf Zimmermann, "Microservices tenets," ACM, 2017. <https://ieeexplore.ieee.org/document/8693772>