

Research Article

Zero-Downtime Migration Strategies for Decomposing Monoliths into Microservices in the Healthcare Industry: A Multivocal Literature Review

Thomas Paul* 

Independent Researcher, International Scientific Society, San Francisco, United States of America

Abstract

Healthcare organizations are modernizing core platforms such as electronic health records (EHRs), order-entry, and billing systems, but decomposing long-lived monoliths into microservices introduces a critical challenge: how to migrate without interrupting mission-critical care operations. Regulated healthcare environments impose strict requirements for availability, patient safety, privacy of protected health information (PHI), auditability, and compliance (e.g., HIPAA and GDPR), which makes conventional migration approaches that rely on downtime or temporary service degradation unsuitable. This study reports a multivocal literature review that synthesizes evidence from peer-reviewed research and high-relevance practitioner sources on strategies that enable continuous service during migration in healthcare settings. Across the reviewed studies, recurring patterns include event-driven integration, Change Data Capture (CDC), coordinated dual-write, backward-compatible schema evolution, progressive traffic shifting (canary and blue-green), and resilience controls such as circuit breakers, idempotent consumers, and controlled failover. These patterns are complemented by observability, governance, and security controls (encryption, access control, and immutable audit logs) that preserve compliance during transitional states. A total of 87 records were retrieved from IEEE Xplore, ACM Digital Library, and SpringerLink (Scopus and Web of Science returned zero records for the search string). After screening and full-text assessment using predefined criteria, a focused set of studies was selected for detailed synthesis. The findings provide practical guidance for planning and executing live migrations in regulated, data-intensive healthcare systems and highlight areas where additional empirical validation is needed.

Keywords

Microservices, Zero-Downtime Migration, Healthcare IT, EHR, Change Data Capture, Event-Driven Architecture, Transactional Consistency, Compliance Architecture

1. Introduction

Healthcare delivery depends on highly available information systems, including EHRs, laboratory and imaging pipelines, patient portals, and revenue-cycle platforms. Many of these systems evolved as monoliths over decades, which

complicates change delivery, scaling, and fault isolation. As healthcare organizations adopt digital pathways (telehealth, remote monitoring, clinical decision support, and interoperability initiatives), the architectural limitations of monoliths

*Correspondence: Thomas Paul (mail2thomaspaul@gmail.com)

Received: 19 January 2026; **Accepted:** 31 January 2026; **Published:** 9 February 2026



Copyright: © The Author(s), 2026. Published by Science Publishing Group. This is an **Open Access** article, distributed under the terms of the Creative Commons Attribution 4.0 License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

increasingly constrain responsiveness and resilience [1, 2].

Microservices are frequently adopted to improve modularity, fault isolation, and independent delivery. In healthcare, these benefits must be realized without violating privacy and safety constraints and while sustaining operational continuity across clinical workflows. Unlike general enterprise domains, healthcare migration must protect PHI, preserve traceability for audits, and avoid disruptions that could affect patient care [3, 4]. Interoperability mandates also shape migration design: healthcare organizations increasingly need standards-based, API-enabled exchange of clinical data and must avoid practices that block authorized access, which affects service boundaries, logging, and audit trails [11, 13, 16]. Migration governance commonly aligns with recognized security baselines, including HIPAA Security Rule implementation guidance and information security management system controls, and API security guidance is used to reduce exposure during incremental rollout and traffic shifting [14, 15, 17].

In this paper, “zero-downtime” refers to migration activities that introduce no planned service interruption perceptible to end users, achieved through redundancy, traffic shifting, and rollback-capable releases; “near-zero downtime” refers to brief, controlled interruptions (typically seconds to minutes) that are bounded by safety procedures and do not disrupt patient-safety-critical workflows.

Zero-downtime decomposition therefore requires incremental migration strategies that limit blast radius and keep legacy and modern components consistent during transition. Prior work highlights that interface-level strangling alone is insufficient when internal coupling and shared data dependencies dominate; instead, migration must address data synchronization, transactional semantics, and operational controls end-to-end [5]. The present study consolidates available evidence on strategies that enable continuous availability during migration and identifies recurring technical and operational patterns that are particularly relevant in regulated healthcare environments.

2. Materials and Methods

This study investigates architectural strategies, technical mechanisms, and operational considerations that enable zero-downtime migration from monolithic systems to microservices in the healthcare industry. Because healthcare systems are safety-critical, highly regulated, and expected to remain continuously available, the review protocol emphasizes rigor and replicability. The study follows a multivocal literature review (MLR) structure, combining a systematic review of peer-reviewed literature (white literature) with a structured review of high-relevance practitioner and technical reports (grey literature) [6]. Reporting and synthesis were structured to align with PRISMA 2020 guidance where applicable to systematic reviews [18].

The methodology comprises four phases: (1) definition of research objectives and questions (Section 2.1), (2) search

strategy and study selection (Section 2.2), (3) data extraction and synthesis (Section 2.3), and (4) transparency and replicability measures (Section 2.4).

2.1. Research Objectives and Questions

The objective is to consolidate evidence on migration patterns that preserve continuous availability and compliance while decomposing healthcare monoliths into microservices. While microservice migration is widely discussed, fewer studies focus specifically on zero- or near-zero downtime under healthcare-specific constraints such as PHI protection, auditability, and clinical safety requirements.

The study addresses the following research questions (RQs):

RQ1: What architectural and migration strategies are used to achieve zero-downtime decomposition of healthcare monoliths into microservices?

RQ2: What technical mechanisms are employed to preserve data integrity, transactional consistency, and latency requirements during zero-downtime migration?

RQ3: What operational and governance practices support continuous availability, security, and regulatory compliance throughout the migration process?

2.2. Search Strategy and Study Selection

The search strategy was designed to identify peer-reviewed studies and authoritative technical documents addressing zero-downtime migration, monolith-to-microservices decomposition, and regulated healthcare system constraints. The search protocol defined (i) search strings, (ii) bibliographic sources, (iii) inclusion and exclusion criteria, and (iv) a staged screening workflow (de-duplication, title/abstract screening, full-text review, and snowballing).

2.2.1. Search String

The search string was derived from the research goals and iteratively refined to capture alternative spellings and closely related terms. Three concept groups were operationalized: (A) microservices, (B) monolith/legacy healthcare systems, and (C) migration with continuous-availability constraints. A healthcare domain qualifier (D) was included to focus results on clinical and health information contexts.

The final query combined these groups using Boolean operators:

```
(microservice* OR micro-service* OR "micro service*")
AND
(monolith* OR "legacy system" OR "monolithic architecture" OR EHR OR EMR OR "health information system") AND
(migrat* OR moderni* OR refactor* OR rearchitect* OR decompos* OR transform*)
AND
("zero downtime" OR "near-zero downtime" OR "continuous availability" OR "live migration" OR "traffic shifting" OR "blue-green" OR canary OR "no downtime")
```

AND

(healthcare OR hospital OR clinical OR "life sciences")

The query was adapted to each library's advanced search syntax and applied to titles, abstracts, and keywords where supported. The search window targeted 2015–2025 to reflect contemporary cloud-native and microservice practices.

To ensure healthcare-specific relevance, the query included domain qualifiers (e.g., healthcare, hospital, clinical) and health-information-system terms. This scoping reduces the risk of over-generalizing from other regulated sectors but may exclude potentially transferable evidence from adjacent regulated domains (e.g., finance or aviation). This trade-off is acknowledged in the limitations and motivates future cross-domain comparative reviews.

2.2.2. White Literature

Sources: Peer-reviewed literature was searched in IEEE Xplore (<https://ieeexplore.ieee.org/>) ACM Digital Library (<https://dl.acm.org/>), and SpringerLink (<https://link.springer.com/>) (all accessed on 28 December 2025). Scopus (<https://www.scopus.com/>) and Web of Science (<https://www.webofscience.com/>) were additionally consulted as citation indexes to validate coverage; however, no relevant records were retrieved from these platforms under the defined query and screening criteria.

Initial search results for white literature are summarized in Table 1.

Table 1. Initial literature search by library.

Library (White Literature)	Records
Scopus	0
IEEE Xplore	45
ACM Digital Library	10
SpringerLink	32
Web of Science	0
Total	87
Non-duplicates	83

Inclusion and Exclusion Criteria: Studies were screened to ensure relevance and technical rigor. Inclusion required (i) explicit focus on monolith-to-microservices migration or decomposition, (ii) continuous availability constraints (zero/near-zero downtime, live cutover, progressive rollout), and (iii) sufficient technical detail (architecture, mechanisms, evaluation, or reproducible procedures). Studies were excluded if they were duplicated, not in English, purely conceptual without technical grounding, lacked verifiable provenance, or were unrelated to healthcare or comparable regulated/mission-critical contexts.

2.2.3. Search and Selection Process

The screening workflow followed a staged process. First, retrieved records were consolidated into a single working set and de-duplicated using title, author list, venue, and year. When a preprint and a publisher version of the same work were both present, the publisher version was retained. Second, an applicability test was performed on a small subset of records to validate the clarity and suitability of the inclusion/exclusion criteria. Third, titles and abstracts were screened; records were removed when they did not address healthcare migration, did not involve continuous availability constraints, or lacked technical specificity. Fourth, full-text review was conducted on shortlisted studies; records were excluded when migration was only mentioned superficially, when continuity mechanisms were not described, or when the work could not be accessed or verified. Finally, backward and forward snowballing was applied to the retained studies to identify missing relevant work through reference lists and citation links.

Based on the above process, a focused subset of studies was selected for data extraction and synthesis.

2.2.4. Grey Literature

Sources: The same search string was executed in Google Scholar (<https://scholar.google.com/>) and arXiv (<https://arxiv.org/>) (both accessed on 28 December 2025) to identify high-relevance technical reports, practitioner articles, and preprints. Broad forum threads and general web search results were not used as primary sources to reduce low-verifiability evidence.

Grey literature items were screened using the same relevance criteria, while relaxing the peer-review requirement. Additional quality filters were applied: identifiable authorship and provenance (e.g., institution, venue, or traceable technical report origin) and substantial technical depth (architecture, mechanisms, evaluation metrics, or reproducible methods).

2.3. Data Extraction and Synthesis

For each retained study, data elements relevant to the RQs were extracted into a structured form. Extracted data included (i) migration patterns and architectural approaches (e.g., strangler-style routing, hybrid coexistence), (ii) data synchronization mechanisms (e.g., CDC, dual-write, outbox), (iii) transaction and consistency strategies (e.g., 2PC, sagas with isolation mitigations), (iv) deployment and traffic management techniques (e.g., canary, blue-green, feature flags), (v) resilience and observability mechanisms (e.g., circuit breakers, idempotency, distributed tracing), and (vi) healthcare-specific compliance controls (e.g., PHI encryption, audit logs, access control, interoperability standards such as HL7/FHIR).

Data synthesis used narrative thematic analysis. Initial categories were derived deductively from the RQs and refined inductively as recurring patterns emerged across studies.

2.4. Transparency and Replicability

To support replicability, the study documented the search strings, source selection, screening criteria, and extraction categories. Where possible, persistent identifiers and bibliographic metadata were retained to enable re-execution of the search and verification of included evidence.

3. Data Collection

3.1. Data Sources

The data collection process was designed to capture verifiable, technically detailed evidence on zero-downtime monolith-to-microservices migration in healthcare. Peer-reviewed studies were sourced from IEEE Xplore, ACM Digital Library, and SpringerLink. Practitioner-facing technical reports and preprints were considered as complementary grey literature when authorship and provenance were identifiable.

3.2. Data Extraction Process

For each included study, data were extracted using a structured form aligned to the research questions. Extracted items included system context (e.g., EHR, imaging, claims), availability constraints, migration patterns (e.g., Strangler Fig), data synchronization methods (e.g., CDC, dual-write), transactional coordination approaches, deployment and traffic-management mechanisms, and compliance controls (privacy, auditability, access control).

3.3. Data Categorization

Extracted data were categorized iteratively. Initial categories were derived deductively from the research questions; additional categories were added inductively when recurring patterns emerged. The final categories emphasized (i) data and migration integrity mechanisms and (ii) operational resilience and compliance enablement.

3.4. Data Validation

To improve reliability, extracted evidence was cross-checked against the original source text and only explicitly supported claims were retained. When multiple versions of the same work were encountered (e.g., preprint and publisher copy), the publisher version was preferred.

4. Results

Following the search and screening process, four healthcare-focused sources were retained for synthesis, comprising two peer-reviewed studies and two grey-literature sources. Despite the small evidence base, the sources con-

verged on a consistent message: zero-downtime migration in healthcare is primarily constrained by patient-safety workflows, protected health information (PHI) governance, and interoperability obligations, which together make data correctness and controlled cutover more critical than raw deployment speed.

4.1. Data and Migration Integrity Mechanisms

Across the selected studies, the first cluster of findings concerned how organizations preserve clinical continuity while decomposing monolithic EHR and ancillary systems. The dominant mechanisms address incremental functional replacement, real-time state synchronization, and the preservation of transactional semantics for safety-critical operations.

4.1.1. Incremental Decomposition with Clinical Interface Continuity

The included sources favored incremental replacement rather than big-bang rewrites. The typical approach routes external clinical requests through a stable gateway or façade while progressively moving capabilities (e.g., scheduling, results viewing, billing) into services. This aligns with the Strangler Fig pattern discussed in software migration research, but the healthcare-specific emphasis is on keeping clinician workflows stable: interfaces and identifiers remain backward-compatible while implementations change behind the boundary [5].

4.1.2. Change Data Capture, Dual-Write Coordination, and Reconciliation

For data-layer migration, the reviewed sources consistently positioned Change Data Capture (CDC) as the baseline technique for maintaining near-real-time synchronization between the legacy database and service-owned stores. Where dual writes were required, the studies stressed coordination and reconciliation to prevent silent divergence. Practical tactics include write-ahead logging pipelines, outbox-style publication of domain events, and periodic integrity checks over patient and encounter records to detect drift before expanding traffic to the new services [7, 8].

4.1.3. Transactional Consistency for Safety-Critical Workflows

A recurring healthcare finding is that several workflows cannot tolerate inconsistent reads (e.g., medication orders, allergy checks, care-plan updates). The sources therefore describe stronger coordination than “eventual consistency” for these paths, including distributed commit where necessary and carefully scoped sagas where compensation is clinically meaningful. When sagas are used, the literature recommends additional safeguards (idempotency, deduplication, and explicit state machines) to avoid repeated side effects in the

presence of retries or network partitions [9, 10].

4.2. Operational Resilience and Compliance Enablement

The second cluster of findings addressed how organizations execute migration under continuous-service constraints while maintaining auditability and privacy controls. Compared with general enterprise migrations, healthcare sources emphasized operational guardrails that limit blast radius and generate reviewable evidence of correctness.

4.2.1. Progressive Traffic Shifting and Safe Cutover

The studies describe progressive traffic shifting (canary and shadow modes) as the practical mechanism for approaching zero downtime. Shadowing compares responses between monolith and services without affecting clinicians, while canarying limits risk to defined clinics, facilities, or user cohorts. Safe cutover criteria include functional parity checks, performance thresholds at peak clinical load, and rollback procedures that restore consistent state rather than merely reverting traffic [2, 5].

4.2.2. Observability and Audit-Ready Telemetry for PHI Workflows

To support regulatory obligations, the sources highlight observability as a compliance capability: end-to-end tracing, immutable audit logs, and structured event histories provide evidence for incident review and access investigations. Event-driven designs (including event sourcing in selected contexts) are used to retain a tamper-evident record of state

changes, which can simplify audit preparation and post-incident reconstruction under HIPAA-style accountability expectations [3, 11].

4.2.3. Resilience Controls and Privacy Safeguards During Transitional States

Finally, the reviewed studies emphasize that transitional architectures increase failure modes (more hops, more data copies) and therefore require layered resilience: circuit breakers, bulkheads, controlled failover, and automated rollback tied to clinical KPIs. In parallel, PHI protections must remain consistent across old and new components through encryption in transit/at rest, least-privilege identities, and policy enforcement at gateways or service meshes. This coupling of resilience and privacy controls is a distinctive healthcare requirement because outages and data exposures both directly affect patient trust and safety [3, 4, 12].

4.3. Healthcare-Specific Adaptations and Constraints

Although many migration patterns are broadly applicable across regulated industries, healthcare systems require additional safeguards due to patient-safety-critical workflows, protected health information, and interoperability requirements. Table 2 summarizes how general zero-/near-zero-downtime migration patterns map to healthcare-specific constraints such as privacy regulation (e.g., HIPAA and GDPR), clinical auditability, and HL7 FHIR-based integration [3, 4, 11].

Table 2. Mapping general migration patterns to healthcare-specific requirements.

General pattern / mechanism	Healthcare-specific requirements	Typical adaptation in practice
Incremental decomposition (Strangler Fig, API facade)	No interruption to clinical workflows; safe rollback; coexistence with legacy EHR interfaces	Introduce FHIR-compliant façade or adapter layer; route traffic progressively by endpoint and cohort; maintain backward-compatible contracts
Change data capture and coordinated dual write	Protected health information integrity; traceable change history; low-latency read-after-write for care delivery	CDC pipelines with end-to-end encryption; dual write with reconciliation jobs; immutable audit logging for patient record updates
Event-driven integration and event sourcing	Auditability for medico-legal traceability; reproducible clinical state changes; provenance for reporting	Use immutable event logs for clinical events; attach metadata for consent and source system; replay for recovery and audit
State-aware deployment and traffic shifting (blue-green/canary plus clinical safety gates)	Prevent partial execution of multi-step clinical transactions; minimize risk during peak clinical hours	Gate releases on clinical safety checks; drain in-flight sessions; gradual rollout with rapid rollback and incident playbooks
Service mesh and policy enforcement	Segmentation of protected health information; zero-trust communication; observability for compliance monitoring	Mutual transport authentication, fine-grained authorization, and audit metrics; enforce data access policies at service boundaries
Disaster recovery and	Consistent restoration of patient record references	Coordinated recovery workflows; integrity checks for identifiers and references; prioritize restoration

General pattern / mechanism	Healthcare-specific requirements	Typical adaptation in practice
cross-service reconciliation	across services; bounded recovery objectives	of patient-safety-critical services

5. Discussion

This section interprets the synthesized evidence in relation to the research questions and highlights implications for healthcare organizations pursuing zero- or near-zero-downtime modernization. Because the final evidence base comprised only four sources (two peer-reviewed studies and two grey literature items), the conclusions are intended as design propositions and plausible hypotheses rather than definitive, generalizable claims; stronger causal statements require additional empirical validation in diverse healthcare settings.

5.1. Interconnected Findings Across Research Questions

The results for RQ1 and RQ2 together suggest that healthcare migration success depends on treating data integrity and operational continuity as inseparable concerns. Incremental decomposition patterns reduce architectural risk, but they only remain safe when paired with continuous synchronization and reconciliation mechanisms that preserve the semantics of PHI and clinical events [7, 8]. This link explains why CDC, dual-write coordination, and strengthened transaction designs emerged as central themes: they provide the correctness substrate required to allow progressive traffic shifting without exposing clinicians or patients to inconsistent state.

The results for RQ2 and RQ3 further indicate that operational controls are not optional add-ons. In healthcare, resilience mechanisms (circuit breakers, controlled failover) are directly connected to safety and compliance obligations, because downtime or partial failure can delay care and create audit gaps [3]. Observability and immutable audit logging therefore act as bridging mechanisms between architectural migration and regulatory expectations, reinforcing the view that compliance must be embedded into system behavior rather than checked retrospectively [3, 8].

5.2. Implications for Healthcare Providers and Vendors

For healthcare providers, the findings imply that modernization programs should prioritize data-plane engineering early. Migration plans that focus primarily on service decomposition without a parallel strategy for CDC, reconciliation, and schema evolution are likely to encounter safety and compliance blockers. Providers should invest in explicit

clinical invariants and validation checks (for example, medication state transitions, encounter timelines) as part of cutover criteria, and they should design rollback as a first-class capability.

For EHR vendors and platform teams, the results suggest that interoperability boundaries must be treated as stable contracts throughout migration. Maintaining HL7 and FHIR compatibility while services evolve reduces downstream integration risk and enables staged adoption by partner systems [11]. Vendors should also expect that transitional states increase the number of data replicas and access paths, which elevates privacy risk. Security controls must therefore be systematically propagated into new services using least-privilege identities, consistent authorization policies, and PHI-aware logging practices [3, 4]. These constraints are reinforced by interoperability and transparency obligations that increasingly favor standards-based APIs and discourage information blocking, making FHIR-aligned service interfaces and traceable access logs a first-class migration concern [13, 16].

5.3. Limitations of Current Evidence and Future Investigation

The evidence base remains limited in two ways. First, only four sources met the inclusion criteria, which reduces confidence in prevalence estimates and increases the risk that some patterns are underrepresented. As a result, the synthesis should be interpreted as a set of candidate mechanisms that appear repeatedly in the included sources, not as an exhaustive or universally validated migration playbook. Second, many healthcare modernization accounts are documented as practitioner guidance or high-level case narratives rather than rigorously evaluated studies, which constrains the strength of claims about effectiveness, cost, and operational risk. These limitations were addressed by framing findings as design propositions, and by highlighting where additional empirical evidence is required.

The search strategy was intentionally scoped to healthcare and to selected bibliographic databases and vetted grey sources. This improves traceability and reduces low-verifiability evidence, but it may exclude transferable lessons from other highly regulated domains (e.g., finance, aviation, or public-sector systems) where zero-downtime constraints are also common. Accordingly, cross-domain generalization is discussed cautiously, and future multivocal reviews could broaden the scope to compare healthcare-specific constraints against patterns reported in adjacent regulated industries.

Future work should therefore emphasize empirical validation under realistic healthcare workloads and integration topologies. Promising directions include controlled evaluations of reconciliation strategies for large EHR datasets, comparative studies of workflow coordination approaches for safety-critical transactions, and disaster-recovery experiments that quantify recovery-point consistency across polyglot persistence. In addition, research on compliance-as-code and automated audit evidence generation could help organizations reduce the operational burden of proving control effectiveness during continuous migration [3, 12]. Empirical studies should also evaluate security and governance controls used during phased cutovers, including HIPAA Security Rule-aligned safeguards and information security management practices, and assess API-level risk mitigation during gateway and service-mesh adoption [14, 15, 17].

6. Conclusions

This study synthesized peer-reviewed and practitioner evidence on zero- and near-zero-downtime migration strategies for legacy healthcare systems transitioning from monolithic architectures to cloud-native microservices. Across the included sources, the most consistently reported mechanisms were incremental decomposition combined with progressive traffic shifting, supported by data-migration foundations such as change data capture, coordinated dual write, and reconciliation workflows. The synthesis further indicates that healthcare deployments require governance and observability practices that treat privacy, auditability, and interoperability as first-class constraints during transitional states, rather than as post-migration checks. Because the final dataset comprised only four sources, the results are best interpreted as practice-oriented design propositions that can guide planning and risk assessment, while motivating rigorous empirical evaluation in broader healthcare contexts.

Beyond architecture and data management, the review highlights that sustained availability during migration depends on operational resilience controls, including circuit breakers, controlled failover, state-aware release management, and incident-ready rollback procedures. For healthcare organizations, these controls are particularly important because even brief service interruptions can affect clinical coordination and downstream reporting obligations. Overall, the review contributes a structured synthesis of candidate patterns and healthcare-specific adaptations that can inform modernization roadmaps for electronic health records and related clinical platforms, and it identifies open research needs in validation metrics, safety-oriented testing, and recovery consistency across distributed patient-data services.

Abbreviations

CDC Change Data Capture

EHR Electronic Health Record
 FHIR Fast Healthcare Interoperability Resources
 HL7 Health Level Seven
 HIPAA Health Insurance Portability and Accountability Act
 MTTR Mean Time to Recovery

Author Contributions

Thomas Paul is the sole author. The author read and approved the final manuscript.

Data Availability Statement

The data is available from the corresponding author upon reasonable request.

Conflicts of Interest

The author declares no conflicts of interest.

Appendix

Appendix I: Search String and Database-Specific Syntax

The following base search string was adapted to each library's advanced search syntax and applied to the title, abstract, and keywords fields where supported:

```
(microservice* OR micro-service* OR "micro service*")
AND (monolith* OR "legacy system" OR "monolithic architecture")
AND (migration OR refactor* OR rearchitect*
OR decomposition OR transformation)
AND (healthcare OR "health information systems" OR "regulated industries")
```

Minor syntax adjustments were made per database (e.g., phrase quoting and field selectors). The searches were executed on 10 November 2025 in IEEE Xplore, ACM Digital Library, and SpringerLink; complementary queries were also run in Scopus and Web of Science to validate coverage.

Appendix II: Screening Summary for Study Selection

The initial peer-reviewed search retrieved 87 records (IEEE Xplore: 45; ACM Digital Library: 10; SpringerLink: 32; Scopus: 0; Web of Science: 0). After removing duplicates, 79 unique peer-reviewed records remained. Titles and abstracts were screened against the inclusion/exclusion criteria, followed by full-text review of shortlisted papers. Ultimately, 4 studies (peer-reviewed and grey combined) were retained for detailed analysis.

During screening, common exclusion reasons included: (i) general cloud migration without explicit zero-/near-zero

downtime constraints, (ii) microservices adoption discussion without migration mechanisms, (iii) healthcare context without technical architecture detail, and (iv) insufficient methodological detail for comparison.

Appendix III: Data Extraction Fields and Coding Categories

For each included study, the following fields were extracted and coded:

- 1) System context (healthcare domain, workload criticality, regulatory constraints)
- 2) Migration strategy (incremental decomposition, interface strangling, hybrid coexistence period)
- 3) Data migration method (change data capture, dual-write coordination, backfill approach)
- 4) Consistency approach (distributed commit, workflow coordination, compensation strategy)
- 5) Release strategy (progressive traffic shifting, rollback strategy, deployment orchestration)
- 6) Resilience and observability (circuit breakers, controlled failover, monitoring and alerting)
- 7) Compliance mechanisms (auditability, immutable logs, access control, governance controls)
- 8) Reported outcomes (availability impact, performance impact, operational complexity)

References

- [1] Newman, S. Building Microservices: Designing Fine-Grained Systems. O'Reilly Media; 2015.
- [2] Fowler, M., Lewis, J. Microservices. Available from: <https://martinfowler.com/articles/microservices.html> (accessed 10 November 2025).
- [3] Publishing body. Health Insurance Portability and Accountability Act (HIPAA), Pub. L. 104–191. United States; 1996. Available from: <https://www.govinfo.gov/content/pkg/PLAW-104publ191/pdf/PLAW-104publ191.pdf> (accessed 10 November 2025).
- [4] Publishing body. Regulation (EU) 2016/679 (General Data Protection Regulation—GDPR). European Union; 2016. Available from: <https://eur-lex.europa.eu/eli/reg/2016/679/oj> (accessed 10 November 2025).
- [5] Richardson, C. Microservices Patterns. Manning Publications; 2018.
- [6] Garousi, V., Felderer, M., Mäntylä M. V. Guidelines for including grey literature and conducting multivocal literature reviews in software engineering. Information and Software Technology. 2019, 106, 101–121. <https://doi.org/10.1016/j.infsof.2018.09.006>
- [7] Kansara, M. Cloud Migration Strategies and Challenges in Highly Regulated and Data-Intensive Industries: A Technical Perspective. International Journal of Applied Machine Learning and Computational Intelligence. 2021. Available from: https://www.researchgate.net/publication/389254166_Cloud_Migration_Strategies_and_Challenges_in_Highly_Regulated_and_Data-Intensive_Industries_A_Technical_Perspective (accessed 10 November 2025).
- [8] Kagga, S. R. Migrating Legacy Healthcare Systems to Cloud-Native Microservices with AI: Best Practices and Pitfalls. International Journal of Applied Mathematics. 2025, 38(2s). Available from: <https://ijamjournal.org/ijam/publication/index.php/ijam/article/view/123/119>
- [9] Garcia-Molina, H., Salem, K. Sagas. ACM SIGMOD Record. 1987, 16(3), 249–259. <https://doi.org/10.1145/38714.38742>
- [10] Bernstein, P. A., Newcomer, E. Principles of Transaction Processing. 2nd ed. Morgan Kaufmann; 2009.
- [11] HL7 International. FHIR Release 4 (R4). 2019. Available from: <https://www.hl7.org/fhir/R4/> (accessed 10 November 2025).
- [12] Fowler, M. Event Sourcing. Available from: <https://martinfowler.com/eaDev/EventSourcing.html>
- [13] Office of the National Coordinator for Health Information Technology (ONC). 21st Century Cures Act: Interoperability, Information Blocking, and the ONC Health IT Certification Program. Federal Register. 2020. Available from: <https://www.federalregister.gov/documents/2020/05/01/2020-07419/21st-century-cures-act-interoperability-information-blocking-and-the-onc-health-it-certification>
- [14] National Institute of Standards and Technology. Implementing the Health Insurance Portability and Accountability Act (HIPAA) Security Rule: A Cybersecurity Resource Guide. NIST Special Publication 800-66 Revision 2. 2024. <https://doi.org/10.6028/NIST.SP.800-66r2>
- [15] International Organization for Standardization. ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection—Information security management systems—Requirements. Geneva, Switzerland: ISO; 2022.
- [16] HL7 International. FHIR Release 5 (R5). 2023. Available from: <https://hl7.org/fhir/R5/>
- [17] OWASP Foundation. OWASP API Security Top 10 – 2023. 2023. Available from: <https://owasp.org/www-project-api-security/>
- [18] Page, M. J., McKenzie, J. E., Bossuyt, P. M., et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. BMJ. 2021, 372, n71. <https://doi.org/10.1136/bmj.n71>

Biography



Thomas Paul is an enterprise software engineer and researcher focused on modernizing legacy systems in regulated domains. His work centers on zero-downtime migration, microservices decomposition, and cloud-native architecture for mission-critical environments, with particular emphasis on data integrity, resilience engineering, and regulatory compliance. He investigates and documents practical strategies for safely transitioning monolithic healthcare and financial platforms to distributed architectures using techniques such as progressive traffic shifting, change data capture, event-driven integration, and automated rollback. His research interests also include reliability patterns, observability, distributed transaction management, and operational governance for high-assurance systems.

Research Field

Thomas Paul: Zero downtime system migration, Microservices architecture modernization, Healthcare information systems engineering, Data consistency and integrity, Distributed transaction management, Event driven system design, Cloud native deployment strategies, Reliability and resilience engineering, Security and compliance by design, Observability and operational.