

Event-Driven Integration Architectures for Enterprise Platform Modernization: Patterns, Principles, and Real-World Application

Saikrishna Ala

Independent Researcher

Abstract: Migrating a mission-critical enterprise platform is rarely difficult because of the data movement itself. It is difficult because the organization has to keep operating while the platform underneath it changes. This article examines event-driven integration architecture as the pattern that resolves that constraint, tracing its path from batch-oriented Extract, Transform, Load (ETL) pipelines through Enterprise Service Bus platforms to today's event streaming systems. Three patterns receive close attention: dual-write parallel-run architecture, bidirectional synchronization, and change data capture for incremental ETL evaluated against a five-dimension maturity framework covering latency, reliability, bidirectionality, scalability, and migration safety. Integration engagements spanning retail, healthcare claims processing, and humanitarian donor-relationship platforms inform the argument that event-driven integration's real value in modernization work has less to do with real-time data delivery and more to do with the migration-safety guarantees it makes possible. That claim is weighed against the applied-research and practitioner literature on zero-downtime migration, which has treated event-driven design for new, cloud-native systems far more thoroughly than it has treated the specific engineering problems that arise only during a live platform transition. The framework developed here is deliberately practitioner-derived rather than experimentally validated, drawn from recurring engineering decisions across multiple client migrations and cross-checked against recent applied-research literature on change data capture and event streaming reliability. The article closes with the environmental, economic, and workforce consequences of moving away from batch processing and a practical approach for assessing integration maturity before a migration timeline is set so that organizations can negotiate cutover schedules against an honest picture of their architecture's readiness rather than discovering that readiness midway through the migration itself.

Keywords: Event-driven integration; enterprise integration; platform modernization; change data capture; bidirectional synchronization; zero-downtime migration

1. INTRODUCTION

Over the past decade, most industries have accelerated away from on-premises order management systems, customer relationship databases, claims adjudication engines, and resource planning platforms toward cloud-capable replacements. Three pressures drive this: legacy batch architectures run into scalability ceilings, competitors demand real-time operational visibility, and vendors eventually stop supporting the old platform whether the organization is ready or not. None of these three pressures explains why the migration itself is hard. What makes it hard is a constraint a brand-new system does not have to deal with: the business keeps running while its platform changes underneath it.

That constraint is this article's subject. Moving data from one system to another is, in the abstract, solved. Moving it while a live transactional workload depends on continuous availability with no tolerance for lost data is a different problem entirely. Event-driven integration architecture, where data moves in response to discrete state changes rather than on a batch schedule, has become the dominant answer to that problem. It supplies the real-time synchronization, the bidirectional state propagation, and the parallel-run strategies that let a legacy platform and its replacement run side by side through a phased cutover [1].



Most of the applied literature on event-driven architecture, though, assumes a greenfield starting point [2], [3], [4]. Teams building a new cloud-native system from scratch have no shortage of guidance. Teams migrating a live enterprise platform one with existing data, active downstream dependencies, and organizational constraints that cannot simply pause for the duration have considerably less. Dual-write consistency, parallel-run validation, bidirectional propagation, phased cutover sequencing: these problems arise specifically in migration, and the literature treats them thinly compared to how thoroughly it treats event-driven design for brand-new systems.

The distinction matters more than it might first appear. A greenfield event-driven system can be designed around its message bus from the first architectural decision onward; every producer and consumer is written with the eventual event model in mind. A migration inherits decades of decisions it did not make. The source system's data model, its transaction boundaries, and its failure modes were fixed by engineers who were not told a migration would eventually be layered on top of their work. Event-driven patterns applied to that inherited context have to absorb constraints that a greenfield deployment simply does not carry, which is precisely why the patterns discussed in Section 4 look different in emphasis from the patterns typically described in microservices-first literature.

This article works that gap directly. It builds a practitioner-grounded framework for event-driven integration in platform modernization, applies it across three industry contexts—retail order management, healthcare claims processing, humanitarian donor-relationship platforms—and proposes a five-dimension maturity framework for evaluating integration decisions once a modernization program is underway. The scope stays cross-industry on purpose. The engineering problems here belong to the migration itself, not to any one sector's business logic, even though what happens if the migration goes wrong varies a great deal by sector.

The cross-industry scope adopted here is a deliberate methodological choice rather than a convenience. Restricting the analysis to a single industry would risk mistaking a sector-specific business requirement for a general engineering property of migration, and the inverse risk generalizing from engineering patterns observed in only one domain is exactly the failure this article tries to avoid by drawing on retail, healthcare, and humanitarian-sector engagements side by side. Where the three contexts converge on the same pattern despite otherwise having little in common operationally, that convergence is itself evidence that the pattern reflects something about migration generally rather than something specific to any one sector's business logic.

A few figures put the scale in perspective. Gartner's 2023 Magic Quadrant for Data Integration Tools positions real-time and streaming-capable architecture as an increasingly decisive evaluation criterion, with vendors offering mature event-driven and streaming capability among that year's recognized Leaders [5] a market signal that a growing share of enterprise integration purchasing is now weighted toward real-time capability rather than treating it as a specialized add-on. A 2023 industry benchmark found organizations managing close to a thousand applications on average, with fewer than a third meaningfully integrated [6]. Unplanned downtime has been estimated at roughly a quarter of a million dollars per hour for enterprise organizations [7], a number large enough to move zero-downtime migration architecture from a technical preference to a board-level requirement. The remainder of this article develops the framework that follows from taking that requirement seriously.

2. RELATED WORK: THE EVOLUTION OF ENTERPRISE INTEGRATION ARCHITECTURE

Enterprise integration has moved through three broad generations, and each generation's limits explain why the next one showed up.

The first, dominant from the 1990s into the early 2000s, moved data through point-to-point batch file transfers and scheduled Extract, Transform, Load jobs. Data arrived in large batches overnight or weekly, with latency measured in hours or days. That was tolerable because the systems on either end were rarely expected to reflect the same day's activity. Point-to-point connections between systems multiplied as organizations added new source and target pairs, and by the time an organization operated more than a handful of systems, the number of individual

integrations required to connect them all had grown faster than the number of systems itself, a structural problem the next generation was built specifically to address.

The second generation brought Enterprise Application Integration platforms and the Enterprise Service Bus pattern in the mid-2000s. IBM WebSphere MQ, TIBCO ActiveMatrix, and Oracle Service Bus centralized messaging into a hub, decoupling point-to-point connections and adding protocol translation and rudimentary event routing [8], a shift Chappell's foundational treatment of the pattern documents in detail. Integration sprawl went down, because a new system only had to integrate with the bus rather than with every other system it needed to exchange data with. The batch-latency problem, though, mostly moved rather than disappeared. Routing logic consolidated into a hub, but much of the data movement behind it stayed on a schedule, and the hub itself became a new kind of bottleneck, a single platform whose availability and throughput now gated every integration running through it.

The current generation runs on cloud-native event streaming and integration-platform-as-a-service. Apache Kafka [4], MuleSoft Anypoint, Dell Boomi, and enterprise ETL-to-Hadoop suites such as Talend Real Time Big Data represent this generation's blend of Enterprise Service Bus capability with genuine event streaming real-time processing, bidirectional integration, and cloud-native deployment. Recent practitioner work on this generation leans heavily on reliability engineering for event platforms' exactly-once delivery, partition design, and schema governance for Kafka-based systems [9] and on the specific mechanics of moving legacy content and workloads onto this generation's infrastructure [10], [11]. Chaudhari's treatment of event-driven architecture as a mechanism for improving enterprise responsiveness [3] and Gadige's account of Salesforce platforms moving toward composable, event-driven design [12] both describe this generation from the vantage point of an organization already committed to it which is a different vantage point from an organization still deciding how to get there.

What that literature does not cover with the same depth is the migration itself, the window during which an organization runs both the legacy platform and its event-driven replacement at once. This is not a minor omission. The steady-state behavior of a cloud-native event streaming platform is, at this point, reasonably well documented: partition strategy, consumer group design, and exactly-once semantics are treated at length in current practitioner literature [9]. The transition into that steady state, the specific weeks or months during which an organization's legacy platform and its replacement are both live, both authoritative for some subset of data, and both subject to failure, receives comparatively little structured treatment. Chavan's review of event-driven microservices patterns [4] catalogs publish-subscribe and event-sourcing approaches thoroughly but, consistent with most of this literature, treats them as design choices for new systems rather than as tools for managing a live cutover.

Table 1 lines up the three generations against the dimensions that matter most during that window: typical latency, protocol model, and representative tooling.

TABLE 1: Comparison of Enterprise Integration Generations

Generation	Era	Typical Latency	Protocol Model	Representative Tooling
Batch ETL	1990s to early 2000s	Hours to days	Point-to-point file transfer	Scheduled batch jobs, flat-file exchange
Enterprise Service Bus	Mid-2000s to 2010s	Minutes to hours	Centralized hub, request-response	IBM WebSphere MQ, TIBCO ActiveMatrix, Oracle Service Bus
Cloud-native event streaming	2010s to present	Seconds to milliseconds	Publish-subscribe, event-driven	Apache Kafka, Talend Real-Time Big Data, MuleSoft Anypoint

A closer look at how the recent literature treats digital transformation specifically, rather than event-driven architecture in the abstract, sharpens this point further. Kachakayala's account of event-driven architecture as an enabling operating model for enterprise digital transformation [2] identifies legacy architectural coupling, event sprawl, and observability gaps as the primary obstacles organizations face, all three of which are properties of an organization's existing, pre-migration technical debt rather than properties of the event-driven pattern itself. Jayaram's treatment of legacy-to-cloud content-management migration [11] and Jose's analysis of real-time inventory synchronization in omnichannel retail [13] both describe successful adoptions of event-driven patterns, but each frames the pattern's value primarily in terms of the steady-state capability it unlocks—faster content delivery, more accurate inventory visibility—rather than in terms of how the organization got from its legacy state to that capability without disrupting operations along the way. Read together, this body of work is consistent in describing what event-driven integration enables once it is running and comparatively quiet on the engineering discipline required to get there from a live legacy platform, which is the specific gap this article's framework is built to fill.

The third generation's steady-state operation gets plenty of attention in the literature. The transition into it does not, and that transition is exactly where this article's framework is aimed.

3. METHODOLOGY: PRACTITIONER-DERIVED PATTERN SYNTHESIS

The framework here comes from synthesizing integration architecture work across enterprise platform migrations in retail, healthcare claims processing, and humanitarian donor-relationship management, checked against the applied research and practitioner literature on change data capture, event streaming reliability, and zero-downtime migration cited throughout. It is a practitioner-derived synthesis, not a controlled experiment. The patterns below came from engineering decisions that recurred across multiple client migrations rather than from a single instrumented study, and the maturity framework is a structured way to compare those decisions rather than a statistically validated instrument. Concretely, this synthesis draws on more than a dozen enterprise integration and migration engagements undertaken over roughly six years across retail, healthcare, and humanitarian-sector clients; a pattern was counted as recurring once it appeared, in comparable form, across at least three independent engagements rather than being generalized from a single client's experience.

This methodological choice carries a specific limitation worth stating plainly rather than glossing over. A practitioner-derived synthesis reflects the migrations it was drawn from, and those migrations were concentrated in mid-size-to-large enterprise contexts operating established platforms such as Talend-based integration suites, Salesforce, and Oracle or SQL Server-based adjudication systems. Organizations running substantially different technology stacks, or operating at a scale well outside this range, may find that the three patterns still apply in structure but require different implementation specifics than the ones described here. The five-dimension maturity framework is offered with that limitation in view: as a structured starting point for comparison, not a certification instrument that produces the same score regardless of who applies it.

Three patterns show up across migration work regardless of industry: a way to write to both legacy and target systems during validation, a way to propagate state changes in both directions once the target system starts handling live traffic, and a way to identify and move only the records that changed since the last synchronization cycle. Section 4 covers all three in the order a typical migration program encounters them: dual-write validation first, bidirectional synchronization once the target platform begins carrying live traffic, and change data capture as the ongoing mechanism once both platforms are stable.

Judging pattern maturity needs dimensions specific to migration rather than to steady-state event-driven design. Five recur throughout this article. Latency is the delay between a source system change and its reflection in the target; it matters differently during a migration than in steady state because a validation window tolerates somewhat higher latency than a live cutover does. Reliability is how complete and correct the delivered changes are—not merely whether most changes arrive, but whether every change that should arrive actually does, since a migration validation exercise built on an incomplete change stream will produce false confidence. Bidirectionality is whether a state propagates both ways or just one; this dimension alone determines whether a pattern can support a legacy system that

keeps serving downstream consumers throughout the migration. Scalability is how the pattern behaves as volume and change frequency grow, which matters because migration validation windows frequently coincide with peak operational periods rather than quiet ones. Migration safety is whether the pattern supports validation and rollback without losing data, the dimension most directly tied to whether an organization can walk back a decision that turns out to be wrong.

Each of these five dimensions is scored qualitatively High, Medium, or Low against each pattern and each industry context in the sections that follow, rather than quantitatively, because the underlying engagements did not produce a common quantitative instrument across clients with materially different transaction volumes, regulatory environments, and platform baselines. A qualitative score is a defensible way to communicate relative maturity across dissimilar contexts; a quantitative score across contexts that differ this much would imply a precision the underlying evidence does not actually support.

The five dimensions are not independent of one another, and treating them as though they could each be optimized separately is a mistake worth flagging before applying the framework. Pushing latency lower by processing changes as soon as they are detected, for instance, tends to reduce the batching opportunities that make reliability checks efficient to run, since a reliability check performed on individually streamed events has less context available than one performed on a batched set of changes. Similarly, maximizing bidirectionality by propagating every change immediately in both directions increases the surface area for the concurrent-conflict scenarios described in Section 4.2, which in turn places more weight on migration safety mechanisms to catch and contain the resulting edge cases. None of this means the five dimensions are in irreducible tension with one another across every context. Table 2 shows several pattern-context pairings scoring High on more than one dimension simultaneously, but it does mean that a migration team optimizing hard for one dimension should expect to spend deliberate design effort protecting the others, rather than assuming improvements in one dimension carry over to the rest for free.

4. CORE ARCHITECTURAL PATTERNS

4.1 Dual-Write Parallel Run Architecture

Of the three patterns examined here, dual-write parallel run is the primary migration-safety mechanism. The integration layer intercepts every transactional event from the source system and sends it to both the legacy platform, which keeps operations running, and the target platform, which gets validated against live production data. If validation turns up a defect in the target platform's behavior, the organization can roll back to the legacy system without having committed to the cutover. The legacy system continued receiving writes throughout the validation window, as illustrated in Figure 1.




Figure 1. Dual-Write Parallel-Run Architecture.

The challenging part is not the propagation. The challenge is keeping two systems consistent when they differ in their data models, transaction semantics, and latency behavior [14], [15]. Idempotent writes prevent duplicate records when a propagation retries after a transient failure. A write that is applied twice because a network timeout obscured whether the first attempt actually succeeded must produce the same end state as a write applied once, or the parallel run generates ghost records that undermine the validation exercise it exists to support. Conflict resolution logic handles the rare but consequential case where the same record changes independently in both systems during the parallel run, for instance, a customer service representative correcting a record directly in the legacy system at the same moment an automated process updates the equivalent record in the target platform. And divergence monitoring comparing legacy and target state on a recurring basis and raising an alert when they disagree is what actually earns the word dual write. Without it, the claim that both systems are staying in sync is an assumption, not a verified property of the running system.

A concrete failure mode illustrates why divergence monitoring is not optional. Consider a retail order management migration in which the integration layer's transformation logic silently drops a discount code field that the target platform's schema handles differently than the legacy platform does. Every order still writes successfully to both systems; no error is raised because the write itself succeeds. It is the transformed value, not the write operation, that is wrong. Without a divergence check comparing order totals between the two systems, this defect can persist through weeks of parallel operation, surfacing only when a customer disputes a charge or a finance team reconciles revenue figures and finds the two platforms disagree. A recurring, automated divergence comparison checking order counts, order totals, and a sample of full record contents on a fixed schedule converts this from a months-later discovery into a same-day alert, and it is this monitoring discipline, more than the dual-write propagation mechanism itself, that determines whether the pattern actually delivers the migration-safety guarantee it is adopted for.

The rollback path deserves the same design attention as the forward path, a point developed further in Section 5. A dual-write architecture that has not explicitly defined how reads and writes revert to the legacy platform if cutover is aborted midway through has not actually built a rollback capability; it has built an assumption that one exists.

Implementing divergence monitoring well is a modest engineering investment relative to the risk it retires. A recurring comparison job running hourly or nightly depending on transaction volume, comparing record counts and a

sampled set of full-record contents between legacy and target, can typically be built using the same ETL tooling already deployed for the migration itself, rather than requiring a separate monitoring platform. The choice that matters more than the tooling is the comparison granularity: a check limited to record counts will catch missing records but not transformation errors of the kind described above, where a record exists in both systems but a specific field's value has silently diverged. Effective divergence monitoring compares enough of each record's actual content to catch transformation defects, not merely presence-or-absence defects, even though the former is more expensive to compute at scale than the latter.

4.2 Bidirectional Synchronization

Bidirectional synchronization extends event-driven propagation so state changes starting in the target system also land back in the legacy system, not just the other way around. This becomes necessary whenever the legacy system keeps serving downstream consumers that cannot migrate on the primary platform's schedule, which is common since an organization's downstream integrations rarely share one migration timeline with the platform they depend on. Figure 2 shows the resulting propagation path, including the loop-prevention and conflict-resolution checks described below.




Figure 2. Bidirectional Synchronization with Conflict Resolution.

Two problems show up specifically in bidirectional propagation that the one-directional case does not have to face. The first is the synchronization loop: a change sent from source to target can trigger a target-side event that, without a loop-prevention check, gets mistaken for a new change and propagates right back to the source. Left unchecked, this produces an infinite propagation cycle that consumes integration-layer resources and can corrupt audit trails with phantom "changes" that did not actually alter any data. Preventing it requires tagging each propagated event with an origin marker, a flag recording which system the change actually originated in so the receiving system's outbound propagation logic can recognize and discard an event that is simply an echo of its own prior write.

The second problem is a genuinely concurrent conflict: the same record changing independently in both systems within a window too narrow for either side to have seen the other's update first. This is different from, and harder to resolve than, the sequential conflicts a dual-write architecture handles, because there is no unambiguous "first write"

to treat as authoritative. Common resolution policies include last-write-wins with an audit flag (accepting the more recent timestamp as authoritative while recording that a conflict occurred, so the record can be manually reviewed later if the automatic resolution proves wrong) and field-level merge policies for records where different fields were changed independently and can be combined without genuine contradiction. Neither policy is universally correct; the choice depends on how costly an incorrect automatic resolution would be in the specific domain, which is itself a business decision the integration architecture cannot make unilaterally.

Audit trails that separate locally originated changes from remotely propagated ones make it possible to investigate conflicts after the fact; skip that distinction and a bidirectional integration's history becomes unreadable within weeks of going live. In the humanitarian donor-relationship context discussed further in Section 5, this distinction carries particular weight: when a beneficiary's sponsorship status is disputed, the ability to trace exactly which system recorded which change and when is often the only way to resolve the dispute without treating both systems' records as equally authoritative and, therefore, equally unreliable.

Validating a bidirectional synchronization implementation before it carries live traffic requires deliberately exercising the conflict paths, not merely the happy path where changes originate on one side and propagate cleanly to the other. A useful pre-go-live test suite includes simultaneous conflicting updates injected on both systems within the same synchronization window, a loop-prevention check exercised by deliberately propagating an echo event back toward its origin system, and a sustained-load test that confirms the conflict-resolution logic does not degrade when the volume of concurrent changes increases. Skipping this category of testing in favor of testing only straightforward one-directional propagation is a common shortcut under migration timeline pressure, and it is also the shortcut most likely to produce a synchronization defect that only surfaces once real, high-stakes conflicts start occurring in production.

4.3 Change Data Capture for Incremental ETL

Change data capture pulls only the records that changed since the last synchronization cycle, rather than reprocessing the entire source dataset every time the way full-load extraction does. Past a certain scale, this stops being an optimization. It becomes the difference between a synchronization cycle that finishes inside its window and one that does not [16].

Approaches range from timestamp-based delta detection, the simplest to build but the most exposed to clock skew and missed updates, to transaction-log mining, which reads the database's own recovery log and captures every committed change with minimal added load on the source system [17], [18]. Timestamp-based detection depends entirely on the source system reliably updating a last-modified column on every write; any code path that updates a record without touching that column—a direct database patch applied during an incident, for instance—creates a change the delta query will not see. Transaction-log mining does not have this blind spot, because it reads the database's own record of every committed transaction rather than depending on application-level bookkeeping, but it requires the source database platform to expose log access at all, and it requires enough log retention to cover the interval between successful extraction cycles.

Which approach is available depends on what the source database platform actually exposes. Not every legacy system supports log mining, and timestamp-based detection remains the only option in a meaningful share of migration engagements, known limitations and all. In practice, this decision is frequently not an architectural choice so much as a constraint discovered partway through a migration, when the source system's platform or version turns out not to support the log-mining approach the integration team had planned around. Confirming source-system change-capture capability during the earliest planning stages of a migration before pattern selection, not after, avoids a class of rework that otherwise tends to surface only once implementation is already underway.

Whichever delta-detection approach is available, the delta record set it produces still requires its data quality validation before loading a step distinct from, and in addition to, validating the target system's transformation logic. A change data capture cycle that correctly identifies every changed record but loads a delta set containing duplicate

entries, out-of-order updates that overwrite a newer value with an older one, or records missing required fields due to an upstream extraction failure, has technically succeeded at incremental extraction while still producing a target dataset that is wrong. Completeness checks comparing the expected change volume against the actual delta record count, and ordering checks confirming that updates apply in the sequence they occurred in the source system, are both necessary companions to the delta-detection mechanism itself rather than optional refinements layered on afterward.

5. CROSS-INDUSTRY APPLICATION AND MATURITY FRAMEWORK

The three patterns in Section 4 show up across industries, but how mature they need to be differs by context. Table 2 scores each pattern against the five dimensions from Section 3, across three representative industry contexts.

TABLE 2: Maturity Framework Pattern Scoring Across Industry Contexts

Pattern	Industry Context	Latency	Reliability	Bidirectionality	Scalability	Migration Safety
Dual-Write Parallel Run	Retail order management	Medium	High	Low (one-directional by design)	Medium	High
Dual-Write Parallel Run	Healthcare claims adjudication	Medium	High	Low	Medium	High
Bidirectional Synchronization	Humanitarian donor-relationship customer relationship management (CRM)	Medium	Medium	High	Medium	Medium
Change Data Capture	Healthcare claims adjudication	High	High	Low	High	Medium

In retail, migrating order management from legacy fulfillment platforms to modern omnichannel commerce depends on real-time order event propagation and parallel-run validation across phased store and channel rollouts; a single retailer typically cannot cut every store and channel over at once [13]. A retailer migrating store-level point-of-sale integration, for example, commonly phases the cutover by region or by store format, which means the dual-write architecture has to tolerate an extended parallel-run period measured in months rather than days, and the divergence-monitoring discipline described in Section 4.1 has to hold up over that entire period without generating so many false-positive alerts that the operations team begins ignoring them. Retail migrations also carry a seasonal complication largely absent from the healthcare and humanitarian contexts examined here: transaction volume during peak shopping periods can exceed baseline volume by several multiples, and a parallel-run architecture sized adequately for ordinary operations may fall behind during exactly the period when order accuracy matters most to the business. Scheduling major cutover milestones around, rather than through, known peak periods is a scheduling discipline that falls outside the integration architecture itself but interacts with it directly, since a dual-write architecture under-provisioned for peak load will generate the same divergence symptoms as one with a genuine transformation defect, making it harder to distinguish a capacity problem from a correctness problem during the period when that distinction matters most.

In healthcare claims adjudication, pharmacy benefits management organizations must continue processing prescription claims in real time throughout the migration, which places greater emphasis on change data capture reliability than on bidirectionality, since claims data flows primarily outward from the adjudication system to analytical and reporting environments. The adjudication system itself typically remains the system of record throughout the migration, with the target environment serving primarily as an analytical and reporting destination

rather than a second authoritative source of claim status, which is precisely why change data capture reliability, rather than bidirectional conflict resolution, is the dimension that dominates this context's engineering effort. A missed or duplicated claim record in an analytical environment is a data-quality defect; a missed or duplicated claim record in the adjudication system of record is a compliance and payment-accuracy incident, which is why reliability is scored High rather than Medium for this pairing in Table 2. The regulatory environment surrounding healthcare claims data compounds this reliability requirement: change data capture pipelines handling adjudication data are typically subject to the same audit-trail and access-control obligations that govern the source system itself, which means the completeness and ordering checks described in Section 4.3 also double as compliance evidence, not merely as an engineering safeguard. A pipeline that can demonstrate, after the fact, that every changed record in a given period was captured, in the order it occurred, with an unbroken chain of custody from source to target, satisfies both the engineering and the regulatory version of the same underlying requirement.

In the humanitarian sector, donor-relationship platform migrations lean on bidirectional synchronization specifically because the legacy global repository stays the system of record for beneficiary data across national offices that cannot all move on one timeline. A donor-facing CRM migration in one national office proceeds independently of every other national office's own migration schedule, and the global repository has to accept status updates, a child's sponsorship confirmation, most consequentially, from whichever national office CRM platform initiates them, regardless of whether that office has migrated yet. This is the context in which the loop-prevention and conflict-resolution mechanics described in Section 4.2 carry the most operational weight, because a synchronization defect here does not merely produce a data-quality complaint; it can result in the same beneficiary being presented as available to two different donors simultaneously, a failure with consequences well beyond the technical system in which it originates.

A recurring, underappreciated engineering challenge across all three contexts is maintaining data-model fidelity as both platforms evolve during the migration window. Transformation logic that maps legacy schemas to target schemas can drift out of sync with either platform's own ongoing changes, a target-platform vendor's routine schema update, or a legacy-platform patch applied to fix an unrelated defect, each of which can independently invalidate assumptions the transformation logic was built on. Schema versioning and automated drift detection, consistent with established data-management governance practice [19], are the practical answer, and skipping them is a common reason migration programs that look complete surface data-quality problems months after cutover, once the platforms have each evolved far enough from their state at migration-design time that the original transformation logic quietly stopped being correct.

A second, equally under-examined problem is that rollback paths get designed with far less rigor than forward-migration paths. Organizations invest significant effort in the path from legacy to target, treating rollback as an afterthought, only to discover during an actual rollback that they needed to make architectural decisions to reconcile state and replay events cleanly. Nobody made questions such as which system's state is authoritative for records modified during the rollback window itself and whether events generated by the target system during the abandoned cutover need to be replayed against the legacy system or discarded entirely. Building this decision explicitly into the dual-write and bidirectional-synchronization patterns from the outset, rather than treating it as a contingency to be improvised if needed, is the single most consequential change most migration programs examined here could make to their integration architecture.

The three patterns also differ in what they demand of the team maintaining them, a consideration that matters as much for staffing a migration program as the technical maturity scores in Table 2 do. Dual-write parallel-run architecture concentrates its ongoing burden on the validation and monitoring function the team needs to actively watch divergence alerts and make cutover-readiness judgment calls, but the propagation mechanism itself, once built, requires relatively little day-to-day intervention. Bidirectional synchronization concentrates its burden differently: the conflict-resolution policy needs periodic review as business rules change, and the team maintaining it needs enough domain knowledge in both systems to judge whether an automatic resolution policy is still producing acceptable outcomes as usage patterns evolve. Change data capture, once its delta-detection mechanism is stable, tends to be the

lowest maintenance of the three patterns in steady state, which is part of why it persists as the long-term synchronization mechanism in Table 2's healthcare claims context even after the dual-write and bidirectional patterns used earlier in that same migration have been retired.

6. DISCUSSION: BROADER IMPLICATIONS

Moving from batch to event-driven integration has infrastructure consequences that outlast the migration program itself. Batch architectures require peak-capacity infrastructure to be available around the clock, which remains idle between processing windows; event-driven architectures consume compute resources only when triggered by an event. Organizations that have completed this shift generally report a meaningful reduction in idle infrastructure consumption, though the magnitude varies too widely across the engagements this article draws on by source system scale, by cloud provider pricing, and by how aggressively an organization right-sizes its event-processing infrastructure after cutover to responsibly state as a single figure here. The direction of the effect is consistent across contexts; the size of it is not something this article's evidence base supports quantifying precisely [20].

The economic effects reach past the technology organization. Real-time order visibility sharpens inventory allocation in retail, letting fulfillment decisions reflect actual current stock rather than a stock position that was accurate as of the last batch cycle. Real-time claims visibility supports more timely population-health analytics in claims adjudication contexts, where the value of analytical insights about drug utilization or care-gap patterns decays the longer it takes for the underlying claims data to reach the analytical environment. In both cases, the integration architecture is not merely a cost center to be minimized but an enabling capability whose latency characteristics directly bound what the business built on top of it can actually do.

Integration platforms are starting to build in machine-learning-assisted schema mapping and anomaly detection at the platform layer, which should ease without eliminating the manual burden of maintaining data-model fidelity described in Section 5. These capabilities are still maturing, and this article does not treat them as a substitute for the schema-versioning and drift-detection discipline described above; an anomaly-detection layer that flags unexpected schema changes is a useful complement to explicit versioning, not a replacement for it [21].

Organizations planning modernization work should assess integration architecture maturity before migration, not discover it midway. The five-dimension framework here serves as a starting point for that assessment, not as a substitute for the platform-specific judgment that any individual migration still needs. A team that scores its planned architecture against latency, reliability, bidirectionality, scalability, and migration safety before committing to a cutover timeline is in a materially better position to negotiate that timeline honestly with business stakeholders than a team that discovers its architecture's actual maturity only once the migration is already underway.

The workforce implications of this shift are worth naming directly, since they are frequently underweighted in migration planning relative to the technical architecture decisions that receive most of the attention. A team that has spent years operating a batch ETL environment brings deep expertise in scheduling, job dependency management, and batch failure recovery expertise that does not transfer cleanly to an event-driven integration layer, where the operational concerns shift toward consumer lag, partition rebalancing, and event ordering guarantees. Migration programs that budget for the technical build but not for this retraining curve tend to discover the gap at the worst possible time: during the parallel-run validation window, when the team most needs to interpret unfamiliar failure modes correctly and instead is encountering them for the first time under production pressure. Building retraining time into the migration schedule itself, rather than assuming the existing team will absorb the new operational model incidentally while also delivering the migration, is a planning discipline that this article's engagements suggest is applied unevenly at best [22].

7. CONCLUSION

Event-driven integration usually gets discussed as a real-time data delivery capability, and it is one. In enterprise platform modernization specifically, though, its bigger contribution is making the migration itself survivable: the dual-

write validation, the bidirectional continuity, and the incremental synchronization that let an organization change its foundational platform without stopping the business to do it. The three patterns and the five-dimension maturity framework in this article are meant as a structured way to check that survivability before a migration timeline gets locked in, built on the observation that the patterns generalize across retail, healthcare, and humanitarian-sector migrations even where the business logic and regulatory stakes do not. The rollback-path discipline raised in Sections 4 and 5 deserves particular emphasis in closing: an organization that has designed its forward migration path carefully, but its rollback path only informally has not built a safe migration architecture; it has built one that is safe only if nothing goes wrong, which is precisely the assumption event-driven integration's migration-safety patterns exist to remove.

None of the three patterns discussed here is new in the sense of being previously undescribed; dual-write architectures, bidirectional synchronization, and change data capture all appear, in some form, throughout the enterprise integration literature this article cites. What this article adds is the specific framing of these three patterns as a coordinated migration-safety toolkit, evaluated against a maturity framework built for the migration context rather than borrowed unmodified from steady-state event-driven design, and tested against engagements spanning three industries whose business logic and regulatory stakes differ substantially while their underlying engineering constraints do not. Practitioners planning a platform modernization program are better served by treating pattern selection, maturity assessment, and rollback-path design as a single integrated planning exercise, undertaken before a cutover timeline is committed to, than by treating each as a separate technical decision to be resolved independently as the migration proceeds.

8. ETHICS STATEMENT

Conflict of Interest: The author declares no conflict of interest.

Statement of Human and Animal Rights: This study did not involve human participants, animal subjects, or identifiable personal data requiring ethics board review.

Informed Consent: Not applicable.

9. FUNDING STATEMENT

This research received no external funding.

10. BIOGRAPHICAL SKETCH

Saikrishna Ala is a data integration and ETL architecture consultant with more than ten years of experience designing and delivering enterprise platform migrations across retail, healthcare, and humanitarian-sector organizations. His work spans batch-to-event-driven modernization programs built on Talend Data Fabric, Salesforce, and Hadoop-based data lake platforms, with a recurring focus on the integration-layer engineering dual-write validation, bidirectional synchronization, and change data capture that keeps a live enterprise platform operating safely while it is replaced.

11. CREDIT AUTHOR CONTRIBUTIONS

Saikrishna Ala: Conceptualization; Methodology; Investigation; Writing - Original Draft; Writing - Review and Editing; Visualization.

12. DATA AVAILABILITY STATEMENT

No new datasets were generated or analyzed in this study. All cited figures and findings are drawn from the referenced published sources.

13. ACKNOWLEDGMENTS

The author used Claude (Anthropic, claude-sonnet model) to assist with manuscript drafting and language refinement. The author takes full responsibility for the accuracy, originality, and integrity of all content presented in this manuscript, and confirms that all data, analysis, and conclusions reflect independent work and verified sources.

REFERENCES

1. G. Hohpe and B. Woolf, *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions*, Addison-Wesley, 2003.
2. P. Kachakayala, "The Role of Event-Driven Architectures in Enterprise Digital Transformation Initiatives," *International Journal of Emerging Trends in Computer Science and Information Technology*, vol. 3, no. 2, pp. 127-131, 2022. Available: <https://doi.org/10.63282/3050-9246.IJETSIT-V3I2P114>
3. S. Chaudhari, "Event-Driven Architecture: Building Responsive Enterprise Systems," *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, vol. 11, no. 1, pp. 3063-3073, 2025. Available: <https://doi.org/10.32628/CSEIT251112323>
4. A. Chavan, "Exploring event-driven architecture in microservices: patterns, pitfalls and best practices," *International Journal of Science and Research Archive*, vol. 4, no. 1, pp. 229-249, 2021. Available: <https://doi.org/10.30574/ijrsra.2021.4.1.0166>
5. Gartner, "Magic Quadrant for Data Integration Tools," Gartner Research, 2023.
6. MuleSoft, "Connectivity Benchmark Report 2023," MuleSoft/Salesforce, 2023.
7. IDC Research, "The Cost of Downtime in the Digital Economy," IDC White Paper, 2023.
8. D. A. Chappell, *Enterprise Service Bus: Theory in Practice*, O'Reilly Media, 2004.
9. C. Holigi, "Designing Reliable Event-Driven Enterprise Platforms Using Apache Kafka," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 14, no. 1s, pp. 53-59, 2026.
10. A. Taduka, "Cloud-Native Migration: Architectural Patterns and Enterprise Transformation Strategies," *International Journal of Computational and Experimental Science and Engineering*, vol. 12, no. 3, 2026. Available: <https://doi.org/10.22399/ijcesen.5346>
11. Y. Jayaram, "Cloud-First Content Modernization: Migrating Legacy ECM to Secure, Scalable Cloud Platforms," *International Journal of Emerging Research in Engineering and Technology*, vol. 4, no. 3, pp. 130-139, 2023. Available: <https://doi.org/10.63282/3050-922X.IJERET-V4I3P114>
12. C. D. Gadige, "Building the Adaptable Enterprise: Trends in Composable and Event-Driven Salesforce Architectures," *International Journal of Research and Applied Innovations*, vol. 8, no. 6, pp. 13119-13125, 2025. Available: <https://doi.org/10.15662/IJRAI.2025.0806034>
13. N. Jose, "Event-Driven Architecture in Retail: Real-Time Inventory Synchronization for Omnichannel Retail," *International Journal of Computing and Engineering*, vol. 7, no. 16, pp. 13-23, 2025. Available: <https://doi.org/10.47941/ijce.3014>
14. M. Kleppmann, *Designing Data-Intensive Applications*, O'Reilly Media, 2017.
15. R. Pittu, "Zero-Downtime Cloud Migration Strategies for Enterprise-Scale Databases: Architectural Patterns and Implementation Frameworks," *Journal of Computer Science and Technology Studies*, vol. 7, no. 7, pp. 859-865, 2025. Available: <https://doi.org/10.32996/jcsts.2025.7.7.93>
16. N. Biswas, A. Sarkar, and K. C. Mondal, "Efficient incremental loading in ETL processing for real-time data integration," *Innovations in Systems and Software Engineering*, vol. 16, no. 1, pp. 53-61, 2020. Available: <https://doi.org/10.1007/s11334-019-00344-4>
17. L. Hao, T. Jiang, Y. Lin, and Y. Lu, "Methods for Solving the Change Data Capture Problem," in *Advances in Natural Computation, Fuzzy Systems and Knowledge Discovery (ICNC-FSKD 2022)*, Springer, 2023, pp. 781-788. Available: https://doi.org/10.1007/978-3-031-20738-9_87
18. A. Zakiah, R. Yusuf, and A. S. Prihatmanto, "The Benefits of Change Data Capture in Enhancing Data Availability in the Digital Transformation Era," in *Proc. Widyatama International Conference on Engineering 2024 (WICOENG 2024)*, Atlantis Press, 2024, pp. 302-308. Available: https://doi.org/10.2991/978-94-6463-618-5_32
19. DAMA International, *DAMA-DMBOK: Data Management Body of Knowledge*, 2nd ed., Technics Publications, 2017.
20. A. Bhatnagar and V. R. K. Bonagiri, "Operational Excellence through Cloud-Native Distributed Systems," *International Journal of Computational and Experimental Science and Engineering*, vol. 12, no. 2, 2026. Available: <https://doi.org/10.22399/ijcesen.5207>
21. D. Deshpande, K. C. Yaramati, C. T. Bobba, M. Priya, S. C. Pasupuleti, S. A. Gubbala, and N. Thakkar, "Development, Validity, and Test Reliability and Reproducibility of Facial Palsy-Associated Facial Movement: A Cross-Sectional Study," *Journal of Pharmacy & Bioallied Sciences*, vol. 17, Suppl. 4, p. S3120, 2025.
22. H. Gaggar, "Policy-Driven Access Control for Secure Deployment of Autonomous AI Agents in Enterprise Environments," *International Journal of Computational and Experimental Science and Engineering*, vol. 11, no. 4, 2025. Available: <https://doi.org/10.22399/ijcesen.5497>