

Event Driven Change Data Capture Architectures for High-Volume Enterprise Data

Srujana Parepalli
Senior Data Engineer, USA

Abstract

By May 2016, enterprises operating high-volume transactional platforms were under growing pressure to integrate operational data across systems with lower latency and greater reliability than traditional batch-oriented approaches could provide. Core databases supporting orders, payments, customer records, and inventory were increasingly required to feed analytics platforms, search indexes, downstream services, and reporting systems in near real time. Existing extract, transform, and load pipelines, typically executed on hourly or nightly schedules, introduced unacceptable delays and operational risk as data volumes and integration complexity increased. At the same time, direct querying of production databases for integration purposes imposed additional load and created tight coupling between systems, threatening transactional stability. Change Data Capture emerged as a pragmatic technique for addressing these challenges by enabling systems to observe committed data changes directly from transactional sources. Rather than periodically extracting full datasets, CDC focused on incrementally capturing inserts, updates, and deletes as they occurred. This approach significantly reduced data movement overhead and improved freshness while preserving the performance characteristics of source systems. By mid 2016, CDC was increasingly implemented using database transaction logs, allowing changes to be captured after commit with strong ordering guarantees and minimal intrusion into application logic. As enterprises adopted asynchronous messaging and event-driven integration models, CDC evolved beyond a replication mechanism into a foundational architectural pattern. Captured data changes were increasingly represented as immutable events and published to a messaging infrastructure that decoupled producers from consumers. This enabled multiple downstream systems to consume the same change stream independently, each applying its own processing logic without coordinating directly with the source database. Event-driven CDC is aligned with broader distributed systems principles that favor loose coupling, scalability, and resilience through asynchronous communication. High-volume integration scenarios placed particular emphasis on scalability and fault tolerance. Organizations processing large transaction volumes require CDC pipelines capable of sustaining continuous throughput while tolerating consumer slowdowns, infrastructure failures, and transient network issues. Event-driven CDC architectures addressed these requirements through durable message storage, partitioned processing, and replay capabilities. However, these benefits came with trade-offs related to eventual consistency, operational complexity, and the need for disciplined monitoring of replication lag and processing health. This paper examines event-driven Change Data Capture patterns as they were understood and applied by May 2016, focusing on architectural approaches suitable for high-volume data integration. It analyzes capture mechanisms, messaging integration models, consistency semantics, and operational considerations relevant to enterprise environments. Rather than presenting CDC as a universal solution, the paper situates it as a critical integration pattern whose effectiveness depends on careful design, governance, and alignment with transactional system constraints.

Keywords: Event-driven architecture, change data capture, high volume data integration, log-based replication, distributed messaging systems, transactional databases, asynchronous data pipelines, data

consistency, fault-tolerant data integration, enterprise data platforms. These keywords reflect the architectural and operational themes associated with CDC based integration as practiced by May 2016, emphasizing scalable event propagation, minimal source system impact, and reliable downstream data consumption.

DOI: 10.21590/ijtmh.2.02.2

1 Introduction

By May 2016, enterprise data integration requirements had expanded well beyond the periodic batch synchronization models that dominated earlier generations of data architecture. Organizations increasingly relied on large-scale transactional systems to support customer-facing applications, operational analytics, and distributed service-oriented platforms. These systems generated continuous streams of data changes that needed to be propagated reliably and with low latency to multiple downstream consumers. As a result, the timeliness and accuracy of data movement became a critical factor in enabling responsive business processes and data-driven decision-making. Traditional data integration techniques, including scheduled extracts and bulk data transfers, struggled to meet these demands in high-volume environments. Batch-oriented approaches introduced inherent delays between data creation and availability, often measured in hours or days. In addition to latency, such approaches placed a significant load on source systems during extraction windows, increasing the risk of performance degradation and operational instability. As transactional workloads grew in size and complexity, these limitations became increasingly difficult to mitigate through incremental optimization alone.

Change Data Capture emerged as a foundational technique for addressing these challenges by focusing on the propagation of incremental changes rather than full dataset replication. By capturing inserts, updates, and deletes at the time they occurred, CDC enabled downstream systems to remain closely synchronized with source databases while minimizing redundant data movement. In enterprise environments, CDC also provides a more precise and auditable representation of data evolution, supporting both operational and analytical use cases. The introduction of event-driven architectural principles further transformed how CDC pipelines were designed and operated. Instead of tightly coupling change extraction with downstream consumption, event-driven CDC architectures treated each data change as an independent event that could be published, transported, and processed asynchronously. Messaging systems acted as intermediaries that decoupled producers from consumers, allowing multiple downstream systems to subscribe to change streams according to their own processing requirements and capacity constraints.

By May 2016, this combination of CDC and event-driven messaging represented a significant shift in enterprise data integration strategy. However, not yet standardized or universally adopted, event-driven CDC patterns offered a promising approach for scaling data movement, improving reliability, and supporting increasingly complex integration topologies. This paper explores these patterns in detail, examining the architectural motivations, design choices, and operational considerations that shaped early implementations in high-volume enterprise environments.

2 Limitations of Traditional Batch-Based Data Integration

Prior to the adoption of CDC-driven integration approaches, most enterprise data movement relied on batch-based mechanisms that operated on fixed schedules and processed large volumes of data in discrete execution windows. These mechanisms typically involve querying source systems for changed records using timestamps or control tables, followed by bulk extraction and downstream loading. While effective for relatively static datasets and low change volumes, this approach became increasingly problematic as transaction rates and data volumes grew. The rigidity of batch schedules constrained how quickly data could be made available and limited the ability of organizations to respond to real-time operational needs. One of the most significant limitations of batch-based integration was inherent latency. Data changes occurring shortly after a batch window closed would not be propagated until the next scheduled execution, creating delays that could span several hours or longer. For use cases such as operational reporting, fraud detection, or system synchronization, this delay undermined the value of downstream data. Attempts to reduce latency by increasing batch frequency often resulted in overlapping jobs, increased system load, and diminishing returns as extraction overhead accumulated.

Batch extraction processes also imposed substantial stress on source systems, particularly transactional databases optimized for concurrent read and write operations. Large scans of tables or indexes during extraction windows competed with production workloads for input output bandwidth, memory, and CPU resources. In high-volume environments, this contention increased the risk of query slowdowns, lock escalation, and, in extreme cases, service outages. These operational risks made aggressive batch scheduling untenable for mission-critical systems. Another limitation involved correctness and consistency guarantees. Batch-based approaches often relied on approximate mechanisms such as last updated timestamps or surrogate keys to identify changed records. These techniques were susceptible to edge cases, including clock skew, out-of-order updates, and multi-row transactional changes that spanned batch boundaries. As a result, downstream systems could observe partial or inconsistent views of source data, requiring additional reconciliation logic and manual intervention to restore correctness.

Dimension	Batch-Based Integration	Event-Driven Integration	CDC
Data movement model	Periodic bulk extraction	Continuous incremental changes	
Latency	Hours to days	Seconds to minutes	
Source system impact	High during batch windows	Minimal via log-based capture	
Failure recovery	Re run entire batch	Replay from event log	
Scalability	Limited by extraction windows	Scales with partitions and consumers	
Consistency handling	Approximate	Transaction aware	

Operational complexity further compounded these issues as integration pipelines grew in scale and scope. Batch workflows frequently involved long chains of dependent jobs with implicit assumptions about data availability and execution order. Failures in upstream jobs could cascade across multiple pipelines, delaying data delivery and complicating recovery procedures. Restarting failed jobs often required reprocessing large volumes of data, increasing both recovery time and operational overhead. Collectively, these limitations highlighted the structural mismatch between batch based integration models and the emerging needs of high volume enterprise systems. As organizations sought to support near real time data propagation, continuous availability, and scalable consumption patterns, it became clear that incremental change driven approaches were required. This realization set the stage for broader adoption of Change Data Capture techniques and for the integration of CDC with event driven architectural patterns.

3 Change Data Capture Techniques in Enterprise Systems as of 2016

By May 2016, Change Data Capture had evolved into a set of well understood techniques for identifying and extracting incremental data changes from transactional systems. These techniques varied in their implementation complexity, performance characteristics, and impact on source systems. Enterprises selected CDC mechanisms based on database capabilities, workload characteristics, and operational constraints, often balancing simplicity against the need for low latency and high fidelity change capture. Trigger based CDC represented one of the earliest approaches to capturing data changes. In this model, database triggers were defined on tables of interest to record inserts, updates, and deletes into auxiliary tables or message queues. While trigger based CDC provided immediate visibility into data changes and strong consistency guarantees, it also introduced additional overhead on transactional operations. In high volume environments, trigger execution could significantly impact transaction latency and throughput, making this approach unsuitable for systems with strict performance requirements.

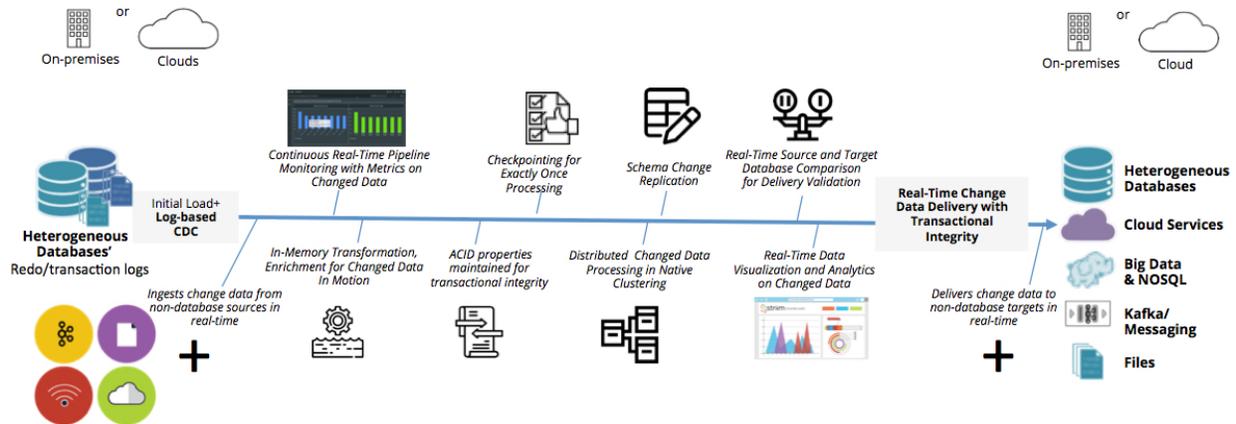
Polling based CDC techniques relied on periodically querying source tables to identify changed records using timestamp columns, version numbers, or change flags. This approach was relatively simple to implement and required minimal database specific features. However, polling based CDC suffered from the same latency and correctness challenges as batch extraction, particularly in environments with high update rates. Additionally, frequent polling increased load on source systems and required careful tuning to avoid contention and missed updates. Log based CDC emerged as the preferred approach for high volume enterprise systems due to its ability to capture changes directly from database transaction logs. By reading and interpreting redo or write ahead logs, log based CDC could observe committed changes without interfering with transactional execution paths. This technique preserved the order and atomicity of database transactions, enabling downstream systems to reconstruct data changes accurately. As of 2016, commercial tools and emerging open source projects supported log based CDC for major relational database platforms.

Despite its advantages, log-based CDC introduced its own set of complexities. Implementations required deep understanding of database internals, log formats, and recovery semantics. Managing log positions, handling schema changes, and ensuring compatibility across database versions demanded careful operational oversight.

CDC Technique	Capture Method	Advantages	Limitations
Trigger based CDC	Database triggers	Immediate visibility	High transactional overhead
Polling based CDC	Timestamps or flags	Simple to implement	Latency and correctness gaps
Log-based CDC	Transaction logs	Low overhead, ordered	Operational complexity

Nevertheless, for enterprises operating at scale, these complexities were often outweighed by the benefits of low-latency change capture and minimal source system impact. Across these techniques, a common theme was the increasing emphasis on incremental, transaction-aware data movement. CDC shifted the focus of integration from periodic snapshots to continuous change streams, laying the groundwork for more responsive and scalable architectures. When combined with asynchronous messaging systems, CDC techniques enabled event-driven patterns that further decoupled data producers from consumers, a topic explored in subsequent sections.

Log-based Change Data Capture



4 Event-Driven Architecture Principles Applied to CDC

Event-driven architecture provided a conceptual framework for extending Change Data Capture beyond simple replication into a scalable and flexible integration mechanism. By May 2016, event-driven principles were increasingly applied to data integration use cases, particularly in environments where multiple downstream systems needed timely access to data changes. In this context, CDC events were treated as immutable records of state change that could be published, stored, and consumed independently of the systems that produced them. A core principle of event-driven CDC architectures was decoupling between change producers and consumers. Source databases and CDC extractors were responsible solely for emitting change events, without knowledge of how or when those events would be processed downstream. Messaging systems acted as durable intermediaries that absorbed fluctuations in event volume and isolated producers from consumer failures. This decoupling improved system resilience and allowed integration pipelines to scale as new consumers were added.

Event ordering and transactional consistency were critical concerns when applying event-driven patterns to CDC. Since database changes occurred within transactions, CDC events needed to preserve commit order to maintain correctness in downstream systems. Event-driven CDC architectures addressed this requirement by partitioning event streams based on primary keys or table identifiers, ensuring that related changes were delivered in sequence. Downstream consumers could then apply changes deterministically, reconstructing consistent data states. Another important principle involved designing for replay and recovery. In contrast to point-to-point replication, event-driven CDC pipelines retained change events for a defined period, allowing consumers to reprocess data in the event of failures or schema changes. This capability was particularly valuable in high-volume environments, where reextracting data from source systems was costly and disruptive. Retained event logs served as a durable source of truth for downstream state reconstruction.

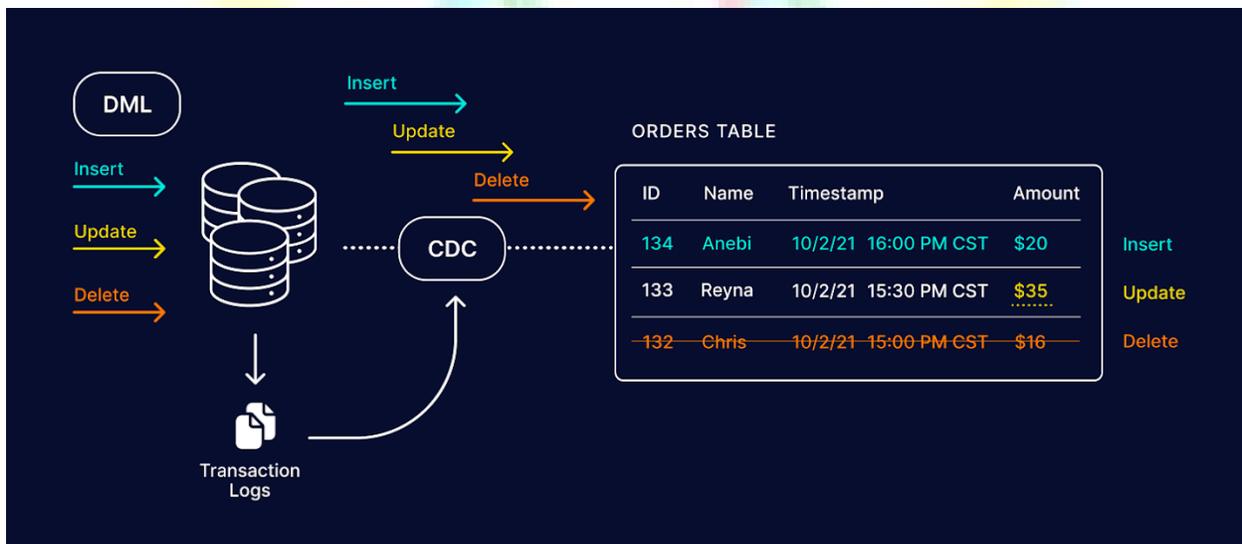
Backpressure handling also shaped early event-driven CDC designs. Messaging systems provided buffering mechanisms that allowed consumers to process events at their own pace without overwhelming downstream resources. This capability was essential for integrating heterogeneous systems with varying performance characteristics. By absorbing bursts of change activity, event-driven CDC pipelines reduced the likelihood of cascading failures across the integration landscape. Through the application of these principles, event-driven CDC architectures enabled more flexible and resilient data integration patterns than traditional replication approaches. Although tooling and operational practices were still evolving in May 2016, the conceptual alignment between CDC and event driven design laid a strong foundation for scalable high volume data integration. The following sections examine how these principles were implemented in practice through messaging platforms and consumer design patterns.

5 Messaging Platforms as the Backbone of Event Driven CDC Pipelines

Messaging platforms played a central role in operationalizing event driven CDC architectures by providing durable, scalable transport for change events. By May 2016, enterprises increasingly relied on distributed messaging systems to decouple CDC producers from downstream consumers and to manage the high throughput associated with transactional change streams.

These platforms offered persistent storage, configurable retention policies, and partitioned delivery models that aligned well with the requirements of CDC based integration. Apache Kafka emerged as a prominent messaging backbone for CDC pipelines due to its append only log design and ability to handle large volumes of events with predictable performance. Kafka topics were commonly used to represent streams of database changes, with partitions enabling horizontal scaling and ordered delivery within defined boundaries. This model allowed CDC events to be consumed by multiple independent applications, including data warehouses, search indexes, and operational caches, without imposing additional load on source systems.

Durability and retention were particularly important characteristics in CDC use cases. Messaging platforms retained events for a configurable period, enabling consumers to recover from failures or to bootstrap new applications by replaying historical changes. This capability reduced reliance on ad hoc backfills and simplified recovery procedures in complex integration environments. Retained event logs also provided a form of operational audit trail, supporting traceability and debugging of data propagation issues. Partitioning strategies were carefully designed to balance ordering guarantees with parallelism. Events were often keyed by primary key or logical entity identifier to ensure that changes affecting the same record were delivered in order. At the same time, distributing events across multiple partitions enabled consumers to process change streams concurrently, improving throughput and reducing end to end latency. Selecting appropriate partitioning schemes was a critical design decision that influenced both correctness and performance.



Messaging platforms also served as a buffer between variable change rates and downstream processing capacity. During peak transactional activity, message queues absorbed bursts of events without overwhelming consumers. This buffering capability supported graceful degradation and allowed downstream systems to scale independently based on their own resource constraints. In high volume environments, such elasticity was essential for maintaining

overall system stability. By providing durable storage, ordered delivery, and scalable consumption models, messaging platforms formed the backbone of event driven CDC pipelines in early enterprise implementations. Their integration with CDC extractors and downstream consumers enabled a shift away from tightly coupled replication mechanisms toward more flexible and resilient data integration architectures. The next section examines how consumers applied CDC events to maintain consistent downstream state.

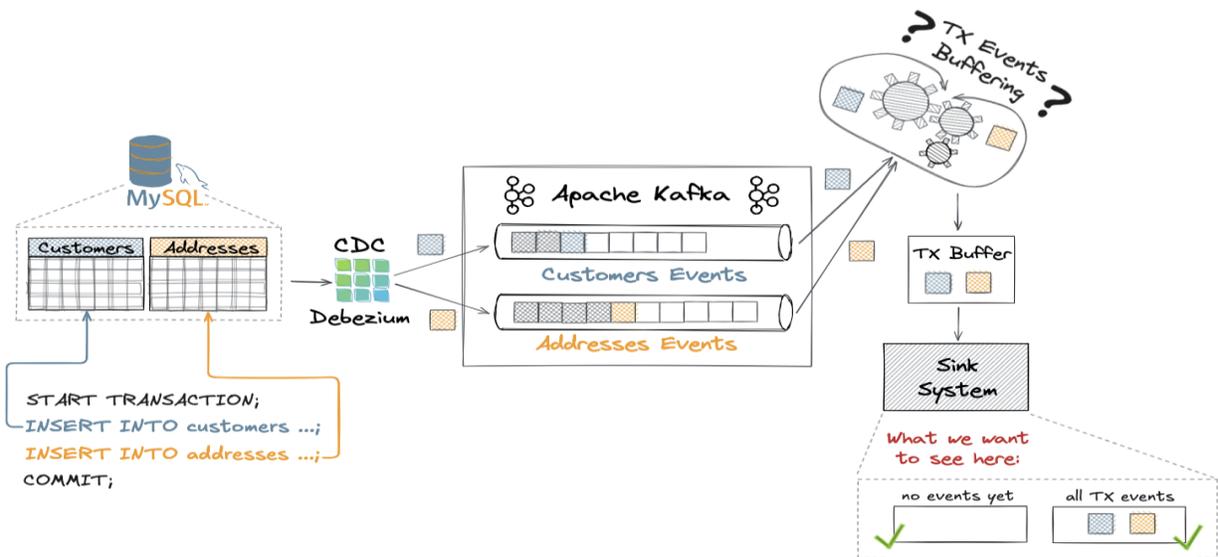
6 CDC Event Modeling and Downstream Consumer Patterns

Effective event driven CDC architectures depended not only on reliable change extraction and transport, but also on careful modeling of CDC events and robust consumer design. By May 2016, enterprises recognized that the structure and semantics of change events directly influenced downstream correctness, extensibility, and operational stability. CDC events were commonly modeled as representations of row level changes, including metadata such as operation type, primary key values, and commit timestamps, allowing consumers to interpret and apply changes deterministically. Event payload design balanced completeness against payload size and processing overhead. Some implementations emitted full row images for each change to simplify downstream consumption and reduce dependency on historical state. Others favored minimal change representations that included only modified fields, reducing event size but requiring consumers to maintain local state. The choice between these approaches reflected trade offs between network utilization, consumer complexity, and recovery behavior. In high volume environments, consistency and replay simplicity often outweighed the benefits of smaller payloads.

Downstream consumers applied CDC events to a variety of target systems, each with distinct consistency and performance requirements. Data warehouses and analytical stores typically consumed events in micro batch or near real time modes, applying changes incrementally to maintain fresh analytical views. Search platforms and caching layers favored low latency consumption to support responsive query workloads. Consumers were designed to handle idempotent processing, ensuring that reprocessed events did not result in duplicate or inconsistent data. Handling schema evolution represented a persistent challenge for CDC consumers. As source database schemas changed over time, event formats needed to evolve without disrupting existing consumers. Early solutions involved versioned event schemas and backward compatible field additions, allowing consumers to adapt incrementally. Although tooling support for schema management was limited in 2016, disciplined schema design and explicit versioning practices helped mitigate integration risks.

Error handling and isolation were essential aspects of consumer design. Consumers were typically built to process events independently, isolating failures to specific pipelines rather than impacting the entire integration landscape. Failed events could be redirected to quarantine queues or logged for manual inspection, allowing the main processing flow to continue. This approach improved overall system resilience and reduced the operational impact of isolated data issues. Through thoughtful event modeling and consumer patterns, enterprises were able to translate raw CDC streams into consistent and actionable downstream state. These designs reinforced the benefits of event driven CDC architectures by enabling flexible consumption,

reliable recovery, and scalable processing. The next section examines operational governance and reliability considerations that shaped real world deployments of these patterns.



7 Governance, Reliability, and Operational Considerations

As event driven CDC pipelines became integral to enterprise data integration, governance and operational reliability emerged as first class concerns. By May 2016, organizations recognized that propagating high volume change events across distributed systems introduced new risks related to data quality, security, and operational control. Unlike traditional point to point replication mechanisms, event driven CDC architectures spanned multiple infrastructure layers, requiring governance to be enforced consistently across extraction, transport, and consumption stages.

Data quality assurance in CDC pipelines focused on validating the integrity and completeness of change events. Since downstream systems relied on incremental updates, missing or malformed events could lead to persistent inconsistencies that were difficult to detect and correct. Enterprises implemented validation checks at multiple points in the pipeline, including verification of event structure, primary key presence, and operation semantics. Monitoring metrics such as event lag, consumer offsets, and processing error rates provided early indicators of data quality issues. Security considerations were particularly important in environments where CDC events contained sensitive transactional data. Access controls were applied at both the messaging layer and consumer endpoints to restrict who could read or process change streams. Encryption of data in transit and, where supported, at rest within messaging platforms helped mitigate the risk of unauthorized access. In early implementations, security controls were often enforced through infrastructure level policies and network segmentation rather than fine grained application level authorization.

Operational reliability depended on designing CDC pipelines to tolerate partial failures without compromising data consistency. Messaging platforms provided durable storage and replay capabilities that allowed consumers to recover from outages by resuming consumption from

known offsets. CDC extractors were designed to track source log positions persistently, ensuring that no committed changes were lost or duplicated during restarts. These mechanisms supported at least once delivery semantics, with idempotent consumer logic used to achieve effective exactly once outcomes. Monitoring and observability played a critical role in managing high volume CDC pipelines. Operational teams tracked metrics related to event throughput, lag between source commit and downstream application, and consumer health. Alerting thresholds were established to detect abnormal conditions such as stalled consumers or rapidly growing backlogs. In the absence of fully integrated tooling, enterprises often relied on custom dashboards and log analysis to maintain visibility into pipeline behavior.

Concern	Design Approach	Risk Mitigated
Ordering	Key based partitioning	Out of order updates
Recovery	Offset tracking and replay	Data loss
Throughput	Buffered messaging	Consumer overload
Security	Topic level access control	Unauthorized access
Observability	Lag and offset metrics	Silent failures

Through disciplined governance practices and robust operational controls, enterprises were able to deploy event driven CDC architectures with confidence despite their inherent complexity. These considerations underscored that CDC was not merely a data extraction technique, but a core integration capability requiring the same level of rigor as transactional application systems. The following section describes the methodology used to analyze and evaluate these architectural patterns within the constraints of the May 2016 technology landscape.

8 Methodology

This paper adopts a qualitative architectural analysis methodology grounded in enterprise data integration practices and documented technology adoption patterns as of May 2016. Rather than relying on controlled experiments or benchmark driven evaluations, the methodology emphasizes synthesis of industry experience, architectural documentation, and peer reviewed technical literature related to Change Data Capture and event driven systems. This approach is appropriate for examining architectural patterns that were emerging in practice but had not yet converged into standardized platforms or frameworks. Primary sources for the analysis include technical papers, vendor documentation, and case descriptions of large scale CDC implementations in enterprise environments. These sources provide insight into how organizations designed and operated CDC pipelines under real world constraints, including high transaction volumes, heterogeneous systems, and strict operational requirements. Particular attention is given to materials that describe log based CDC mechanisms and the use of distributed messaging platforms to transport change events.

The methodology involves decomposing event driven CDC architectures into functional components, including change extraction, event modeling, transport, consumption, and operational governance. Each component is examined in terms of its responsibilities, interaction patterns, and failure modes. By analyzing these components individually and in combination, the paper identifies recurring design patterns and architectural principles that characterize effective CDC based integration solutions. Comparative analysis is also employed to contrast event driven CDC approaches with traditional batch based and point to point replication models. This comparison focuses on dimensions such as latency, scalability, fault tolerance, and operational complexity. The objective is not to quantify performance differences, but to highlight structural trade offs that influence system behavior in high volume environments.

To maintain historical accuracy, the scope of analysis is explicitly constrained to technologies, tools, and organizational practices that were realistically available by May 2016. Concepts and capabilities that gained prominence in later years are deliberately excluded to avoid retrospective interpretation. This constraint ensures that the findings reflect the decision making context faced by enterprises during the period under study. Finally, the methodology prioritizes operational feasibility and architectural clarity as evaluation criteria. Patterns are assessed based on their ability to support sustained throughput, preserve transactional consistency, and integrate with existing enterprise governance frameworks. By grounding the analysis in practical considerations, the methodology aims to produce insights that are relevant and actionable within the technological landscape of May 2016.

9 Findings and Observations

The analysis indicates that event driven Change Data Capture architectures provided a substantial improvement in both scalability and timeliness compared to traditional batch based integration approaches. Enterprises that adopted CDC driven pipelines were able to propagate data changes with significantly lower latency, often reducing delays from hours to minutes or seconds depending on workload characteristics. This improvement enabled downstream systems to operate on fresher data, supporting operational reporting and near real time analytics without imposing additional load on transactional source systems. A key finding is that log based CDC emerged as the most viable technique for high volume environments due to its minimal impact on database performance and its ability to preserve transactional ordering. By capturing changes directly from database logs, enterprises avoided intrusive query patterns and reduced contention on production systems. This approach also ensured that downstream consumers received a faithful representation of committed transactions, simplifying consistency management and reducing the need for complex reconciliation logic.

The study also finds that messaging platforms fundamentally changed the integration topology by enabling multiple independent consumers to subscribe to the same change streams. Rather than maintaining separate replication pipelines for each downstream system, enterprises centralized change propagation through shared event streams. This fan out model improved operational efficiency and reduced duplication of integration logic, while allowing individual consumers to evolve independently based on their specific processing and availability requirements. Another important observation concerns the role of replayability in operational resilience. Retained CDC event logs allowed downstream systems to recover from failures,

rebuild state, or onboard new consumers without re extracting data from source systems. This capability reduced recovery time and improved confidence in pipeline correctness. Enterprises that were explicitly designed for replay and idempotent consumption experienced fewer prolonged outages and less operational disruption.

The findings further suggest that effective CDC implementations required close alignment between data engineering and operational teams. Decisions related to event schema design, partitioning strategy, and retention policies had direct implications for both correctness and performance. Organizations that treated CDC pipelines as shared infrastructure, rather than isolated integration scripts, were more successful in managing growth in data volume and consumer diversity. Finally, the analysis highlights that early event driven CDC architectures were most successful when applied selectively to high value integration paths. Enterprises often prioritized systems with high transaction volumes or strong latency requirements, while continuing to use batch integration for less critical workloads. This pragmatic adoption strategy allowed organizations to realize immediate benefits from CDC while gradually building operational expertise and confidence in event driven integration models.

10 Challenges and Limitations

Despite the clear advantages of event driven Change Data Capture architectures, enterprises adopting these patterns by May 2016 faced several technical and organizational challenges. One significant limitation involved the operational complexity of CDC infrastructure. Log based CDC systems required careful management of database log positions, connector health, and compatibility with database versions and configurations. Failures in CDC extractors could lead to gaps or duplication in change streams if not handled correctly, increasing the operational burden on data engineering teams. Another challenge stemmed from the immaturity of tooling and standards for CDC event schemas. In early implementations, event formats were often custom defined and tightly coupled to specific consumers. This coupling made schema evolution difficult, particularly as source databases changed over time. Without robust schema governance mechanisms, changes to table structures risked breaking downstream consumers or introducing silent data inconsistencies. Enterprises mitigated this risk through conservative schema design and explicit versioning, but these practices required discipline and coordination.

Performance variability also presented a limitation in high volume environments. Although messaging platforms could absorb bursts of change activity, downstream consumers were sometimes unable to keep pace during sustained peaks. This resulted in growing backlogs and increased end to end latency. Capacity planning for CDC pipelines proved challenging due to fluctuating transaction rates and uneven consumer workloads. Enterprises often relied on empirical tuning and over provisioning to maintain acceptable performance. Data consistency across multiple consumers posed additional challenges. While CDC preserved transaction order within individual streams, coordinating consistent state across heterogeneous downstream systems was complex. Consumers processed events at different rates and applied different transformation logic, leading to temporary divergence in data state. For use cases requiring

strong cross system consistency, additional coordination mechanisms or reconciliation processes were sometimes necessary.

Security and compliance requirements further constrained CDC adoption in regulated environments. Exposing detailed change events raised concerns about sensitive data leakage and access control enforcement. Implementing fine grained security policies across messaging platforms and consumers was difficult with the tooling available in 2016. As a result, some organizations limited CDC usage to internal systems or applied masking and filtering logic at the point of consumption. Finally, organizational readiness influenced the effectiveness of event driven CDC initiatives. The shift from batch integration to continuous event processing required new skills, operational mindsets, and collaboration models. Teams accustomed to scheduled data movement had to adapt to always on pipelines and real time monitoring. Resistance to change and uncertainty about long term architectural direction led some enterprises to adopt CDC cautiously, limiting its scope and delaying broader integration.

11 Conclusion

By May 2016, event driven Change Data Capture patterns had emerged as a powerful architectural approach for addressing the limitations of traditional batch based data integration in high volume enterprise environments. Increasing demands for low latency data propagation, scalable integration, and reduced impact on transactional systems drove organizations to adopt incremental change driven models supported by distributed messaging platforms. These patterns represented a meaningful evolution in how enterprises designed and operated data integration pipelines. This paper examined the foundational techniques and architectural principles underlying event driven CDC as practiced during this period. Through analysis of CDC extraction methods, messaging backbones, event modeling strategies, and consumer design patterns, the paper highlighted how enterprises achieved greater decoupling, resilience, and scalability in data movement. The findings emphasize that log based CDC combined with durable event transport provided a reliable mechanism for propagating transactional changes across diverse downstream systems.

The analysis also demonstrates that early CDC architectures were shaped by practical constraints related to tooling maturity, operational complexity, and governance requirements. Enterprises adopted event driven CDC selectively, focusing on high value integration paths while continuing to rely on batch processing for less critical workloads. Success depended on disciplined schema management, careful capacity planning, and close collaboration between data engineering and operational teams. In conclusion, event driven Change Data Capture patterns in May 2016 represented an important transitional stage in enterprise data integration architecture. While challenges and limitations remained, the principles established during this period laid the groundwork for more advanced real time data platforms in subsequent years. Understanding

these early patterns provides valuable context for the evolution of scalable, reliable, and event oriented data integration systems.

12 References

1. Leonardo Neumeyer, Bruce Robbins, Anish Nair, Anand Kesari (2010). S4: Distributed Stream Computing Platform. 2010 IEEE International Conference on Data Mining Workshops (ICDMW), 170-177. <https://doi.org/10.1109/ICDMW.2010.172>
2. Rajendra K. Gupta, Jayant R. Haritsa, Krithi Ramamritham (1996). Commit Processing in Distributed Real-Time Database Systems. Proceedings of the 17th IEEE Real-Time Systems Symposium, 289-298. <https://doi.org/10.1109/REAL.1996.563719>
3. Rajendra K. Gupta, Jayant R. Haritsa, Krithi Ramamritham (2000). The PROMPT Real-Time Commit Protocol. IEEE Transactions on Parallel and Distributed Systems, 11(2), 160-181. <https://doi.org/10.1109/71.841752>
4. Flaviu Cristian, Christof Fetzer (1999). The Timed Asynchronous Distributed System Model. IEEE Transactions on Parallel and Distributed Systems, 10(6), 642-657. <https://doi.org/10.1109/71.774912>
5. Einar Broch Johnsen, Olaf Owe (2004). An Asynchronous Communication Model for Distributed Concurrent Objects. Proceedings of the 2nd International Conference on Software Engineering and Formal Methods (SEFM 2004), 188-197. <https://doi.org/10.1109/SEFM.2004.1347520>
6. Peter Bailis, Ali Ghodsi (2013). Eventual Consistency Today: Limitations, Extensions, and Beyond. Communications of the ACM, 56(5), 55-63. <https://doi.org/10.1145/2460276.2462076>
7. Sebastian Burckhardt (2014). Principles of Eventual Consistency. Foundations and Trends in Programming Languages, 1(1-2), 1-150. <https://doi.org/10.1561/25000000011>
8. Tushar Deepak Chandra, Vassos Hadzilacos, Sam Toueg (1996). The Weakest Failure Detector for Solving Consensus. Journal of the ACM, 43(4), 685-722. <https://doi.org/10.1145/234533.234549>
9. Cynthia Dwork, Nancy Lynch, Larry Stockmeyer (1988). Consensus in the Presence of Partial Synchrony. Journal of the ACM, 35(2), 288-323. <https://doi.org/10.1145/42282.42283>
10. Seth Gilbert, Nancy Lynch (2002). Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-Tolerant Web Services. ACM SIGACT News, 33(2), 51-59. <https://doi.org/10.1145/564585.564601>
11. Maurice P. Herlihy, Jeannette M. Wing (1990). Linearizability: A Correctness Condition for Concurrent Objects. ACM Transactions on Programming Languages and Systems, 12(3), 463-492. <https://doi.org/10.1145/78969.78972>

12. Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, David G. Andersen (2011). Don't Settle for Eventual: Scalable Causal Consistency for Wide-Area Storage with COPS. SOSP '11: Proceedings of the 23rd ACM Symposium on Operating Systems Principles, 401-416. <https://doi.org/10.1145/2043556.2043593>
13. C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz (1992). ARIES: A Transaction Recovery Method Supporting Fine-Granularity Locking and Partial Rollbacks Using Write-Ahead Logging. ACM Transactions on Database Systems, 17(1), 94-162. <https://doi.org/10.1145/128765.128770>
14. Marc Shapiro, Nuno Preguiça, Carlos Baquero, Marek Zawirski (2011). Conflict-Free Replicated Data Types. SSS 2011: Proceedings of the 13th International Conference on Stabilization, Safety, and Security of Distributed Systems, LNCS vol. 6976, 386-400. https://doi.org/10.1007/978-3-642-24550-3_29
15. Douglas B. Terry, Marvin M. Theimer, Karin Petersen, Alan J. Demers, Mike J. Spreitzer, Carl H. Hauser (1995). Managing Update Conflicts in Bayou, a Weakly Connected Replicated Storage System. SOSP '95: Proceedings of the 15th ACM Symposium on Operating Systems Principles, 172-182. <https://doi.org/10.1145/224056.224070>

